

Robust Safety-Critical Control: A Lyapunov and Barrier Approach

Thesis by
Andrew James Taylor

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy, Control and Dynamical Systems

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2023
Defended May 22nd, 2023

© 2023

Andrew James Taylor
ORCID: 0000-0002-5990-590X

All rights reserved

ACKNOWLEDGEMENTS

When I show my father my travel pictures, he often responds that “your pictures need people in them”. He will say that thirty years from now, the picture of that beautiful landscape I took will not matter very much to me, but the ones capturing memories of the people I love will be ones I cherish forever.

I have come to believe that a similar sentiment is true about graduate school. In thirty years, I likely will not remember every detail of every proof or every small mathematical technicality that has seemed so important at one point or another. But rather, I will remember the people that were there with me sharing in the discovery process. Those there with me checking proofs, running last-minute experiments on robots, and pushing to get our ideas down on paper for an impending deadline. I will also remember those people who were there when research was not “clicking” and gave me the support to just keep pushing. Without them, this thesis would not have been possible.

My thesis advisor, Professor Aaron Ames, has played an integral role in my development as a control theorist, researcher, and engineer. Under his guidance, I have developed in my capacity to identify interesting problems, propose insightful solutions, and perhaps most importantly, communicate *why* the problems and solutions matter. My writing skills have improved tremendously because of his teaching, and for this, I am very grateful. Professor Yisong Yue has provided invaluable advice to me throughout my PhD. I came in to my PhD knowing nearly nothing about the field of machine learning, and through his mentoring, have come to understand many of the important challenges that must be solved in building truly autonomous systems. Professors Richard Murray and Joel Burdick have also provided me with much wisdom on the relationship between control theory research and control systems practice. This has been fundamental in forming my perspectives on which problems are worth solving. Professor James Forbes, my undergraduate advisor, introduced me to control systems with his course on attitude dynamics. He was instrumental in my decision to pursue graduate studies at Caltech, and has continually encouraged my research efforts.

There have been several colleagues of mine who have been a fundamental part of my research experience. I met Victor Dorobantu at the very first wel-

come event for our program (something I often tease him for forgetting). Throughout the entirety of our PhDs, we have spent countless hours working on homework assignments, whiteboarding ideas, pair-programming, and arguing over notation as we tried to write our papers as precisely as possible. His attention to detail and mathematical rigor has vastly impacted my abilities as a researcher, and I am so glad to call him a colleague and friend. I was exceptionally fortunate to have had the opportunity to serve as a research mentor for Noel Csomay-Shanklin and Ryan Cosner. Their boundless enthusiasm and constant excitement for solving the next big problem in robotics has been a great source of inspiration for me throughout my PhD, and for their friendships I am deeply grateful.

There are many additional colleagues I have had the chance to collaborate with throughout my PhD, including Andrew Singletary, Ruben Grandia, Anil Alan, Tamás Molnar, Pio Ong, Jason Choi, Kim Wabersich, Preston Culbertson, Maegan Tucker, Amy (Kejun) Li, Wyatt Ubellacker, Min Dai, Ugo Rosolia, Hoang Le, Chaozhe He, Manuel Galliker, Farbod Farshidian, Mohamed Maghenem, and Meera Krishnamoorthy, as well as Professors Sarah Dean, Gábor Orosz, Marco Hutter, Benjamin Recht, Paulo Tabuada, Jorge Cortés, Melanie Zeilinger, Claire Tomlin, Koushil Sreenath, Katherine Bouman, and Ricardo Sanfelice. Every one of them has in some way contributed in building my perspective on research. Likewise, there are other colleagues who have provided me with fruitful discussions on various research and engineering problems, including Riley Murray, Apurva Badithela, Fernando Castenada Garcia-Rozas, Bike Zhang, Thomas Youmans, James Croughan, and Mrdjan Janković. I am also grateful for the other past and present members of the AMBER Lab, including Jenna Reher, Rachel Gehlhar, Thomas Gurriet, Ivan Jimenez Rodriguez, Albert Li, Gilbert Bahati, Sergio Esteban, and Adrian Ghansah.

My ability to succeed in research has depended on my happiness when not doing research. There have been many people who have helped me achieve that during my time at Caltech. From our Dungeons & Dragons group to our trips out in the wilderness, my time spent with my friends Bijan Mazaheri, Aidan Fenwick, Kristjan Eldjarn, Ellande Tang, Prithvi Akella, and Rebecca Gallivan has brought much joy. Another huge facet of my life outside of research has been music. My friends Connor McMahan and Christian Leefmans in the Caltech Classical Guitar Ensemble and all of the individuals I have played with

at St. Philip the Apostle have helped recharge my research efforts through our shared musical experiences. Some of my favorite memories from graduate school were with James Holper, Varun Wadia, Cam Gray, David Needell, and Janis Hesse in our band Sunshine & The Deep Jams. I loved every night we spent practicing, every post-practice stop for a beer, and every show we played. When I had the *brilliant* idea to leap from the Gene Pool and break my leg during a show, they were incredibly good friends supporting my recovery. Beyond our music, they have been some of my best friends through my hardest moments in graduate school. My guitar teacher, Matthew Elgart has taught me the importance of having *an* idea, of nothing being done unconsidered. This is perhaps one of the most important lessons I have learned with respect to my research, as it now underlies every word I write. For everything he has taught me, about music, and life, I am so immensely grateful.

Throughout my PhD, I have been carried by the overwhelming support and love of my family. Roya Saadat and Farhad Razavian have been my California parents, opening their home and lives to me in the heart of the pandemic and supporting me throughout many difficult times during my studies. My Nannie and Dzia, Terenia and Andrew Zawisza, have given so much to build a life for their grandchildren, including me. They have nurtured my curiosity and always expressed their belief in my ability to achieve anything I set out to. My grandparents, Glen and Christine Taylor, lived a life that has been deeply inspirational for me. In the face of many challenges, they created a loving environment for their children and grandchildren that has shaped who I have grown into. My aunts, Suzette Dalrymple and Lisa Byerley, have been there throughout my entire life, and have never failed to support me whenever I have needed it. Jason Eickmann has been an uncle, friend, and mentor to me, and has never hesitated to support me in any way possible. He taught me to be honest with myself, which has been an invaluable guiding principle whenever I make difficult decisions.

The unconditional love and limitless support of my parents, James and Rachelle Taylor, has made all of my achievements possible. They have spent countless hours encouraging me when I have been frustrated, shared in the joy of my successes, and offered me profound wisdom on how to find happiness in my life. I can not fathom navigating many of the situations I have faced without the guidance of my father's experience. My mother's endless compassion has

always brought me back to the fundamental truth that “it will be okay”. The completion of my education is as much their achievement as it is mine. My sister, Isabel Taylor, is the best sibling that a brother could ask for. She has never faltered in showing me support, never failed to brighten my mood, and never allowed me to take myself too seriously. I am so glad that I have got to share so many life experiences with her, and I am excited to share so many more. I love all three of you.

To my partner, Noosha Razavian, you and I met right before the start of the pandemic, and had to choose to likely never see each other again, or go for it. We chose the latter, and it has been one of the greatest decisions of my life. You, more than anybody, have been in the trenches with me through the thick and thin of graduate school. You have always recognized and celebrated my highest points, and shown me the deepest of compassion when I have been at my lowest points. You have demonstrated immense patience when I had to commit myself fully to research, reminded me the virtue of modesty, and rallied the people who care about me when I have needed it the most. We have shared so many exhilarating adventures together, and I am thrilled for the ones that lie ahead of us. I love you.

Above all, I thank God. You have been with me at every step of this journey, and when I was not strong enough to walk it, you carried me. Each individual in this acknowledgement is a blessing you have bestowed upon me, for which I am eternally grateful. Without you, I am nothing.

ABSTRACT

Accompanying the technological advances of the past decade has been the promise for widespread growth of autonomous systems into nearly all domains of human society, including manufacturing, transportation, and healthcare. At the same time, there have been several tragic failures that reveal potential risks with the expansion of autonomous systems into everyday life, and indicate that it is vital for safety to be accounted for in the design of control systems.

This thesis seeks to develop a theory of robust safety-critical control for autonomous systems. This theory will be built upon the foundational tools of Control Lyapunov Functions (CLFs) and Control Barrier Functions (CBFs), which provide a powerful paradigm for the design of model-based safety-critical controllers. The dependence of CLF and CBF-based controllers on a system model makes them susceptible to modeling inaccuracies, potentially resulting in unsafe behavior when deploying these controllers on real-world systems.

In this thesis I present methods for resolving four classes of model inaccuracies referred to as model error, disturbances, measurement error, and input sampling, which are commonly faced challenges when designing controllers for robotic systems. The proposed methods are unified by their shared use of CLFs and CBFs to produce controllers possessing rigorous and robust safety guarantees that can be demonstrated in simulation or experimentally. A hallmark of these methods is a focus on enabling control synthesis through convex optimization, which ensures that controllers can be efficiently computed on real-world robotic hardware platforms.

In addressing model error, I consider both data-driven learning approaches and adaptive control approaches. I present three episodic learning frameworks that iteratively augment existing CLF and CBF-based controllers specified via convex optimization problems to improve the stability and safety properties of a system, which I demonstrate in simulation and experimentally. I also establish a relationship between the degradation of stability and safety properties with the magnitude of residual learning error through the perspective of Input-to-State Stability (ISS) and Input-to-State Safety (ISSf). Lastly, I develop an adaptive safety-critical control framework for systems with parametric model error through the notion of adaptive CBFs.

In addressing disturbances, I resolve challenges in balancing performance and robustness with ISSf-based controllers through the notion of Tunable Input-to-State Safety (TISSf), which permits prioritizing robustness to disturbances only when safety requirements are close to being violated. I demonstrate the capabilities of TISSf-based control design experimentally on an autonomous semi-trailer truck system that is subject to input disturbances due to complex unmodeled actuator dynamics. Lastly, I develop a framework for achieving ISSf-like finite-time safety guarantees for discrete-time systems subject to stochastic disturbances through the use of CBFs and convex optimization.

In addressing measurement error, I develop the notion of Measurement-Robust CBFs (MR-CBFs), which permit control synthesis through convex optimization in the presence of imperfect measurements. I demonstrate the capability of MR-CBFs on an experimental Segway system using a vision-based measurement system, validating the tractability of using controllers specified through increasingly complex classes of convex optimization problems on real-world systems. Lastly, I present an application of Preference-Based Learning (PBL) in tuning the robustness parameters of a CBF-based controller, demonstrating the first use of PBL with CBFs and providing a tool for tuning the safety and performance of the robust controllers proposed in this thesis.

In addressing input sampling, I consider both sampled-data and event-triggered paradigms for modeling input sampling. I provide a method for synthesizing CLF-based controllers for sampled-data systems by integrating feedback linearization with approximate discrete-time models, leading to a significant improvement over continuous-time CLF-based controllers implemented with input sampling. I then develop a framework for achieving safety of sampled-data systems through approximate discrete-time models through the notion of practical safety and Sampled-Data CBFs (SD-CBFs), which I demonstrate with convex-optimization based controllers in simulation. Lastly, I develop a method for event-triggered safety-critical control that uses ISSf to achieve safety while satisfying the requirement of a minimum interevent time.

Collectively, these contributions constitute a significant advance in the theory of robust safety-critical control by establishing a framework, unified by the use of CLFs and CBFs in conjunction with convex optimization, that addresses a wide class of challenges faced in the design of safety-critical control systems.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, “Robust safety under stochastic uncertainty with discrete-time control barrier functions,” *accepted in Robotics: Science and Sys. (RSS) XIX*, 2023. [Online]. Available: <https://arxiv.org/abs/2302.07469>,
A. J. Taylor participated in the conception of the project, theoretical analysis, and writing of the article. The contents are presented in Chapter 4.
- [2] A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, “Control barrier functions and input-to-state safety with application to automated vehicles,” *conditionally accepted as a Full Paper in IEEE Trans. on Control Sys. Tech.*, 2022. [Online]. Available: <https://arxiv.org/abs/2206.03568>,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 4.
- [3] A. Alan, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, “Safe controller synthesis with tunable input-to-state safe control barrier functions,” *IEEE Control Sys. Let.*, vol. 6, pp. 908–913, 2022. DOI: 10.1109/LCSYS.2021.3087443,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article. The contents are presented in Chapter 4.
- [4] R. K. Cosner, M. Tucker, A. J. Taylor, *et al.*, “Safety-aware preference-based learning for safety-critical control,” in *Proc. 4th Learning for Dynamics and Control (L4DC)*, vol. 168, Stanford, CA, USA, 2022, pp. 1020–1033. [Online]. Available: <https://proceedings.mlr.press/v168/cosner22a.html>,
A. J. Taylor participated in the conception of the project, theoretical analysis, and writing of the article. The contents are presented in Chapter 5.
- [5] A. J. Taylor, V. D. Dorobantu, R. K. Cosner, Y. Yue, and A. D. Ames, “Safety of sampled-data systems with control barrier functions via approximate discrete time models,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 7127–7134. DOI: 10.1109/CDC51059.2022.9993226,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 6.

- [6] A. J. Taylor, V. D. Dorobantu, Y. Yue, P. Tabuada, and A. D. Ames, “Sampled-data stabilization with control lyapunov functions via quadratically constrained quadratic programs,” *IEEE Control Sys. Let.*, vol. 6, pp. 680–685, 2022. DOI: 10.1109/LCSYS.2021.3085172,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 6.
- [7] A. J. Taylor, P. Ong, T. G. Molnar, and A. D. Ames, “Safe backstepping with control barrier functions,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 5775–5782. DOI: 10.1109/CDC51059.2022.9992763,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article.
- [8] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, “Measurement-robust control barrier functions: Certainty in safety with uncertainty in state,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Prague, Czech Republic, 2021, pp. 6286–6291. DOI: 10.1109/IROS51168.2021.9636584,
A. J. Taylor participated in the conception of the project, algorithm design, and writing of the article. The contents are presented in Chapter 5.
- [9] N. Csomay-Shanklin, R. K. Cosner, M. Dai, A. J. Taylor, and A. D. Ames, “Episodic learning for safe bipedal locomotion with control barrier functions and projection-to-state safety,” in *Proc. 3rd Learning for Dynamics and Control (L4DC)*, vol. 144, Zürich, Switzerland, 2021, pp. 1041–1053. [Online]. Available: <http://proceedings.mlr.press/v144/csomay-shanklin21a/csomay-shanklin21a.pdf>,
A. J. Taylor participated in the conception of the project, algorithm design, and writing of the article. The contents are presented in Chapter 3.
- [10] A. J. Taylor, V. D. Dorobantu, S. Dean, B. Recht, Y. Yue, and A. D. Ames, “Towards robust data-driven control synthesis for nonlinear systems with actuation uncertainty,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6469–6476. DOI: 10.1109/CDC45484.2021.9683511,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 3.
- [11] A. J. Taylor, P. Ong, J. Cortés, and A. D. Ames, “Safety-critical event triggered control via input-to-state safe barrier functions,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 749–754, 2021. DOI: 10.1109/LCSYS.2020.3005101,
A. J. Taylor participated in the conception of the project, algorithm

design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 6.

- [12] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “A control barrier perspective on episodic learning via projection-to-state safety,” *IEEE Control Sys. Lett.*, vol. 5, no. 3, pp. 1019–1024, 2021. DOI: 10.1109/LCSYS.2020.3009082,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation and experimental code implementation, conduction of experiments, and writing of the article. The contents are presented in Chapter 3.
- [13] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, “Guaranteeing safety of learned perception modules via measurement-robust control barrier functions,” in *Proc. Conf. on Robotics Learning (CoRL)*, Cambridge, MA, USA, 2020. [Online]. Available: <https://proceedings.mlr.press/v155/dean21a/dean21a.pdf>,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article. The contents are presented in Chapter 5.
- [14] A. J. Taylor and A. D. Ames, “Adaptive safety with control barrier functions,” in *Proc. IEEE American Control Conf. (ACC)*, Denver, CO, USA, 2020, pp. 1399–1405. DOI: 10.23919/ACC45564.2020.9147463,
A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 3.
- [15] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “Learning for safety-critical control with control barrier functions,” in *Proc. 2nd Learning for Dynamics and Control (L4DC)*, vol. 120, Berkeley, CA, USA, 2020, pp. 708–717. [Online]. Available: <https://proceedings.mlr.press/v120/taylor20a.html>,
A. J. Taylor participated in the conception of the project, algorithm design, simulation and experimental code implementation, conduction of experiments, and writing of the article. The contents are presented in Chapter 3.
- [16] A. J. Taylor, V. D. Dorobantu, M. Krishnamoorthy, H. M. Le, Y. Yue, and A. D. Ames, “A control lyapunov perspective on episodic learning via projection to state stability,” in *Proc. IEEE 58th Conf. on Decision and Control (CDC)*, Nice, France, 2019, pp. 1448–1455. DOI: 10.1109/CDC40024.2019.9029226,
A. J. Taylor participated in the conception of the project, theoretical analysis, simulation code implementation, and writing of the article. The contents are presented in Chapter 3.

- [17] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, “Episodic learning with control lyapunov functions for uncertain robotic systems,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp. 6878–6884. DOI: 10.1109/IROS40897.2019.8967820,
- A. J. Taylor participated in the conception of the project, algorithm design, simulation code implementation, and writing of the article. The contents are presented in Chapter 3.

TABLE OF CONTENTS

| | |
|---|-------|
| Acknowledgements | iii |
| Abstract | vii |
| Published Content and Contributions | ix |
| Table of Contents | xii |
| List of Illustrations | xv |
| List of Tables | xxvi |
| Nomenclature | xxvii |
| Chapter I: Introduction | 1 |
| 1.1 General Mathematical Outlook | 12 |
| Chapter II: Background | 30 |
| 2.1 Nonlinear Dynamics | 30 |
| 2.2 Stability & Safety | 31 |
| 2.3 Control Lyapunov Functions | 33 |
| 2.4 Feedback Linearization | 39 |
| 2.5 Control Barrier Functions | 46 |
| 2.6 Nonlinear Dynamics with Disturbances | 51 |
| 2.7 Input-to-State Stability & Input-to-State Safety | 53 |
| 2.8 Input-to-State Stable Control Lyapunov Functions | 55 |
| 2.9 Input-to-State Safe Control Barrier Functions | 60 |
| Chapter III: Learning & Adaptive Control for Nonlinear Systems | 64 |
| 3.1 Related Work | 65 |
| 3.2 Model Errors | 70 |
| 3.3 Learning with Control Lyapunov Functions | 72 |
| 3.4 Projection-to-State Stability | 83 |
| 3.5 Learning with Control Barrier Functions | 97 |
| 3.6 Projection-to-State Safety | 105 |
| 3.7 Projected Disturbance Learning with Control Barrier Functions | 112 |
| 3.8 Data-Driven Nonlinear Control | 125 |
| 3.9 Adaptive Safety via Control Barrier Functions | 144 |
| 3.10 Conclusion and Future Work | 163 |
| Chapter IV: Disturbance-Robust Safety-Critical Control | 165 |
| 4.1 Related Work | 167 |
| 4.2 Limitations of Input-to-State Safety | 171 |
| 4.3 Tunable Input-to-State Safety | 174 |
| 4.4 Tunable Input-to-State Safety on Automated Truck | 188 |
| 4.5 Safety with Stochastic Disturbances | 202 |
| 4.6 Conclusion | 221 |
| Chapter V: Measurement-Robust Safety-Critical Control | 224 |
| 5.1 Related Work | 226 |

| | |
|--|-----|
| 5.2 Measurement-Robust Control Barrier Functions | 228 |
| 5.3 Integration with Backup-Set Methods | 240 |
| 5.4 Preference-Based Learning for Robust Controller Tuning | 249 |
| 5.5 Conclusion | 262 |
| Chapter VI: Sampled-Data & Event-Triggered Control | 264 |
| 6.1 Related Work | 266 |
| 6.2 Sampling Paradigms | 271 |
| 6.3 Sampled-Data Stabilization | 279 |
| 6.4 Sampled-Data Safety | 297 |
| 6.5 Event-Triggered Safety | 316 |
| 6.6 Conclusion | 329 |
| Chapter VII: Conclusion | 333 |
| 7.1 Infeasible Optimization Problems | 336 |
| 7.2 Future Work | 341 |
| Bibliography | 343 |
| Appendix A: Unmatched Tunable Input-to-State Safety Proof | 376 |

LIST OF ILLUSTRATIONS

| <i>Number</i> | <i>Page</i> |
|--|-------------|
| 1.1 Ideal control system diagram. | 2 |
| 1.2 Practical control system diagram featuring model error, disturbances, measurement modules, and input sampling. | 3 |
| 1.3 Spectrum of robustness properties of results presented in this thesis. | 16 |
| 3.1 (Left) Nominal model-based (CLF-QP) controller fails to stabilize the system to the trajectory. (Right) Improvement in angle tracking of system with learning-informed (CLF-QP) controller over nominal PD controller. (Bottom) Corresponding visualizations of state data. Video of this simulation is found at https://youtu.be/cB5MY_8vCrQ | 81 |
| 3.2 Augmented controllers consistently improve stabilization to trajectory across episodes. Ten instances of the algorithm are executed with the shaded region formed from minimum and maximum angles for each time step within an episode. The corresponding average angle trajectories are also displayed. | 82 |
| 3.3 Comparison of tracking performance for PD controller and final learning-informed (CLF-QP) controller in inverted pendulum example. The final learning-informed (CLF-QP) controller tracks the desired angle trajectory significantly more accurately. | 95 |

- 3.4 Comparison of worst-case residual learning errors (3.80) with nominal model-based (CLF-QP) controller (Left) and (3.81) with final learning-informed (CLF-QP) controller (Right). Closed-loop trajectories using the two controllers are displayed with dashed lines. The learning-informed (CLF-QP) controller keeps the system in regions with lower residual learning error, while the system leaves the region around the training data under the nominal model-based (CLF-QP) controller. The heat maps were generated by sampling states randomly about training data points and evaluating the upper bound for each sampled state. The results were then discretized for ease of visualization. Each bin is colored by the maximum disturbance observed in the bin. 96
- 3.5 Simulation results using the nominal model-based (CBF-QP) controller (green) and learning-informed (CBF-QP) controller (3.88) (blue). Episodic data appears as red traces. (Left) Segway simulation CAD model. (Middle) The phase portrait showing the evolution of the state of the system leaves the safe set (black ellipse) with the nominal model-based (CBF-QP) controller, while it remains within using the learning-informed (CBF-QP) controller. (Right) The value of h drops below zero (twice) using the nominal model-based (CBF-QP) controller, while it remains above zero using the learning-informed (CBF-QP) controller. Video can be found at https://youtu.be/tD8zH4GF_8U. 103
- 3.6 Experimental results using the nominal model-based (CBF-QP) controller (green) and learning-informed (CBF-QP) controller (3.88) (blue). (Left) Custom robotic Ninebot Segway platform. (Middle) The phase portrait showing the evolution of the state of the system leaves the safe set (black strip) with the nominal model-based (CBF-QP) controller, while it remains within using the learning-informed (CBF-QP) controller. (Right) The value of h drops below zero (at $t = 3$) using the nominal model-based (CBF-QP) controller, while it remains above zero using the learning-informed (CBF-QP) controller. Video can be found at https://youtu.be/YCesvD_Ae00 104

- 3.7 Simulation results with Segway platform demonstrating improvement in PSSf behavior. (Left) Absolute value of the projected disturbance along the trajectory without learning models ((3.115), red) and with learning models ((3.116), blue), with learning reducing the worse case projected disturbance ($\bar{\delta}/k$). (Right) The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound is raised with learning (the blue dashed line lies above the red dashed line). 111
- 3.8 Experimental results with Segway platform demonstrating improvement in PSSf behavior. (Left) Absolute value of the projected disturbance along the trajectory without learning models ((3.115), red) and with learning models ((3.116), blue), with learning reducing the worse case projected disturbance ($\bar{\delta}/k$). (Right) The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound is raised with learning (the blue dashed line lies above the red dashed line). 112
- 3.9 Schematic for challenges in learning CBF time derivatives with structured models. 114
- 3.10 Schematic for introducing additional data for resolving challenges in learning CBF time derivatives with structured models. . . . 115
- 3.11 (Left) Schematic diagram of the AMBER-3M robot with position coordinates. (Center) Schematic of the foot placement in the stepping-stone problem. The boundaries of virtual stepping-stones are captured via the blue and orange vertical lines. (Right) Virtual stepping stone width as function of the phase variable $\tau(\mathbf{q})$. 121

- 3.12 Simulation (S) and Hardware (H) data where model mismatch causes violations. (Far Left) Simulation where the barrier functions h_1 (solid blue) and h_2 (solid orange) are enforced via a nominal model-based (CBF-QP) controller. A conservative controller using a worst-case upper bound on the projected disturbance is also shown (dashed blue, dashed orange), which results in more conservative behavior over many steps. (Mid Left) After three episodes of learning using the (SS-QP) controller in simulation, the maximum barrier violation decreases from 2.0 to 0.3 [cm]. (Mid Right) Hardware where the barrier functions h_1 (blue) and h_2 (orange) are enforced via a nominal model-based (CBF-QP) controller. (Far Right) After two episodes of learning on hardware, the maximum barrier violation decreases from 9.2 to 1.9 [cm] using the (SS-QP) controller. 123
- 3.13 Gait tiles for Episode 2 of learning showing the AMBER-3M robot safely traversing a set of stepping stones. Notice the change in step width and added lean of the torso induced by the barrier functions. Video can be found at <https://vimeo.com/481809664124>
- 3.14 Schematic of model error set arising from multiple data points. 127
- 3.15 Projected view of model error set used in proving feasibility of the (DR-CCF-OP). The green region corresponds to $\Pi(\mathcal{U}_c(\mathbf{x}))$ while the orange region highlights the set $\bar{\mathcal{U}}_c(\mathbf{x})$ lying above the hyperplane \mathcal{H} . The line \mathcal{H}_β represents the strictly separating hyperplane with normal \mathbf{s} and offset β , while \mathcal{H}_0 is the hyperplane shifted to pass through the origin. 139
- 3.16 The value of the certificate function $\mathcal{C}(\mathbf{x})$ (Top) and input u (Bottom) over time for the stability (Top Pair) and safety (Bottom Pair) experiments. I compare the behavior of a true model-based (CCF-QP) controller (oracle, black square), a nominal-model based (CCF-QP) controller (baseline, red diamond), and the proposed (DR-CCF-SOCP) controller using dense input data (R-dense, blue triangle) and sparse input data (R-Sparse, green circle). For stability, the desired behavior for $\mathcal{C}(\mathbf{x})$ is to converge towards 0, while for safety it is to remain non-positive. 143
- 3.17 Schematic of aCBF counterexample. 157

| | | |
|------|--|-----|
| 3.18 | Simulation trajectory of aCBF counterexample. The system state leaves the prescribed safe set as it converges to the limit cycle. | 159 |
| 3.19 | Comparison of different adaptive and non-adaptive control methodologies. (Top, Left) An aCBF controller is able to enforce safety of the nominal controller k_{nom} . (Top, Right) An aCLF-aCBF controller is able to track the desired velocity with zero steady state error (Top, Right). (Bottom) Both aCBF controllers are able to keep the vehicle within the safe region for all time. . . . | 161 |
| 4.1 | Simple example demonstrating limits of ISSf in achieving performance and meaningful theoretical safety guarantees. (a) Safe controller with no disturbances. (b) Impact of harmonic disturbance signal without modifications to controller design. (c) ISSf based controller designs for different values of the parameter ϵ_0 . Large values of ϵ_0 lead to weak theoretical guarantees, while small values lead to conservative behavior. | 173 |
| 4.2 | Examples of the function ϵ that satisfy the condition in (4.14). | 177 |
| 4.3 | Simple example using TISSf to achieve both meaningful theoretical safety guarantees and non-conservative behavior. | 180 |
| 4.4 | Simulation results for the inverted pendulum system. The gold ellipse is the set \mathcal{C} as defined in (4.37). The black line is the set where $L_{\mathbf{g}}h(\mathbf{x}) = 0$ as defined in (4.38). The dashed blue line is the trajectory of the system evolving under k_{nom} as defined in (4.43), which leaves the set \mathcal{C} . The green and purple regions indicate where the controller k_{nom} meets and fails to meet the CBF condition, respectively. The dashed red line is the trajectory of the system evolving under k_{CBF} which remains inside the set \mathcal{C} . | 183 |
| 4.5 | Disturbance signal for the inverted pendulum system example as defined in (4.45). | 185 |
| 4.6 | (Right) Simulation results for the inverted pendulum system with disturbances. The gold ellipse is the safe set \mathcal{C} defined in (4.37). The blue line is the trajectory of the system evolving under k_{CBF} defined in (CBF-QP), which is not robust to disturbances and leaves the safe set. The black, green, and dashed red lines are trajectories of the system evolving under k_{TISSf} defined in (TISSf-CBF-QP) with different values for ϵ_0 and λ | 186 |

| | | |
|------|--|-----|
| 4.7 | (Left) Curves corresponding to the value of h^* solving (4.49) across the (ϵ_0, λ) parameter space for the inverted pendulum example. (Right) The boundary of the set $\mathcal{C}_{d,T}$ rendered forward invariant for different choices of the parameters ϵ_0 and λ for the inverted pendulum example. Note that the $\mathcal{C}_{d,T}$ contains the set \mathcal{C} for each parameter choice. | 187 |
| 4.8 | Experimental configuration for heavy-duty semi-trailer truck problem. (Top) Controller design without robustifying element yields safety violation. (Bottom) Robust safety-critical controller ensures semi-trailer truck brakes early and aggressively enough to maintain safe distance. | 189 |
| 4.9 | A connected automated truck following a connected vehicle. . . | 190 |
| 4.10 | (Left) The value of the function ρ defined in (4.55), which defines the minimum safe following distance as a function of the leader's velocity v_L and truck velocity v . (Right) The value of $L_f h(\mathbf{x}) + \alpha(h(\mathbf{x}))$ as defined in (4.58) when $L_g h(\mathbf{x}) = 0$ for various distances. As the function is strictly positive over the domain of interest, h is a CBF on \mathcal{C} for (4.52). | 192 |
| 4.11 | (Top) Range policy V defined in (4.60). (Bottom) Speed policy W defined in (4.61). | 192 |
| 4.12 | Example profile for acceleration a_L of lead vehicle used in simulation. | 193 |
| 4.13 | (Left) Velocity v_L of lead vehicle (black) and velocity of the truck using the nominal controller (4.59) (blue) and (CBF-QP) controller (red). (Right) Following distance D and value of CBF h using the nominal controller (4.59) and (CBF-QP) controller. . | 194 |
| 4.14 | (a) Vehicles used in experiments. (b) Image from the dashboard of the truck during an experimental run. (c,d,e) Final configurations of separate experiments. See https://youtu.be/9dJtC1TCBbA for a video. | 196 |
| 4.15 | Mean value (lines) and standard deviations (fills) of the distance D and the CBF h when using the nominal controller (4.59) (blue) and the (CBF-QP) controller (red) in the truck experiment. The experiments with these controllers are highly consistent. . . . | 197 |

| | | |
|------|--|-----|
| 4.16 | (Left) Discrepancy between acceleration commanded by safety-critical controller (black) and actual acceleration of the automated truck (red). (Right) Disturbance signal in input seen by the truck used to define the model (4.62). | 198 |
| 4.17 | Following distance D and value of h using the nominal controller (4.59) (blue), (CBF-QP) controller (red), and (TISSf-CBF-QP) controller (green) in the disturbed simulation. | 199 |
| 4.18 | Mean value (lines) and standard deviations (fills) of the following distance D and h using the nominal controller (4.59) (blue), (CBF-QP) controller (red), and (TISSf-CBF-QP) controller (green) in experiment. | 199 |
| 4.19 | (Left) Parameter values for ϵ_0 and λ used in the truck experiments, with contours showing theoretical values of h^* . Green markers denote parameter sets which achieve the original safety goal ($h(\varphi_d(t)) \geq 0$), while red markers denote parameter sets for which the original safety goal is violated. (Right) Theoretical values of h^* and the shift in steady-state tracking distance \tilde{D}_{ss} , for the parameter sets used in the truck experiments. The blue markers denote parameter sets with $\lambda > 0$, while the black markers denote parameter sets with $\lambda = 0$ | 200 |
| 4.20 | Experimental results for parameter pairs (B), (C), and (D) in Table 4.3. Case (B) is conservative as indicated by the large steady-state tracking distance. Cases (C) and (D) display nearly identical behavior, but case (C) possesses a stronger theoretical guarantee. | 202 |
| 4.21 | The dashed lines represent the theoretical probability bounds for the system as in Theorem 34. The solid lines represent the Monte Carlo (MC) estimated P_u across 500 experiments. . . . | 218 |
| 4.22 | (Top Left) System diagram of the inverted pendulum. (Top Right) 500 one second long example trajectories starting at $\mathbf{x}_0 = 0$. (Bottom Left) Monte Carlo estimates of P_u for $\gamma = 0$ using 500 one second long trials for each initial conditions represented by a black dot. (Bottom Right) The (conservative) theoretical bounds on P_u from Theorem 34. | 219 |

- 4.23 Safety of a simulated quadrupedal robot locomoting on a narrow path for a variety of controllers. (Top Left) The safe region that the quadruped is allowed to traverse. (Bottom Left) A system diagram depicting the reduced-order model states of the quadruped $\begin{bmatrix} x & y & \theta \end{bmatrix}^\top$. (Top Right) 50 trajectories for 3 controllers: one without any knowledge of safety (\mathbf{k}_{nom}), one with the standard (DT-CBF-OP) controller, and finally the proposed method which accounts for stochasticity (JED). (Bottom Right) Plots of $h(\mathbf{x})$, a scalar value representing safety. The system is safe (i.e., in the green safe region) if $h(\mathbf{x}) \geq 0$ 220
- 5.1 Visualization of the point-wise sets $\hat{\mathcal{X}}(\mathbf{x})$ and $\mathcal{X}(\mathbf{y}, \hat{\mathbf{x}})$. The set $\hat{\mathcal{X}}(\mathbf{x})$ consists of all possible estimates that may correspond to the state \mathbf{x} , and the set $\mathcal{X}(\mathbf{y}, \hat{\mathbf{x}})$ consists of all the possible states for a given measurement-state estimate pair $(\mathbf{y}, \hat{\mathbf{x}})$ 230
- 5.2 (Left) The Segway model used in simulation from the perspective of the fixed virtual camera used to estimate its state. (Right) Simulation results for worst-case measurement model uncertainty of $\epsilon = 0.2$ subtracted from the true pitch angle θ_y when measured. A state trajectory generated using the the (CBF-QP) controller (red) and the (MR-CBF-SOCP) controller (blue) are shown as projections onto their pitch angle and pitch rate components. The set \mathcal{C} is plotted in green. Given the same initial condition, the (MR-CBF-SOCP) controller ensured safety of the trajectory whereas the (CBF-QP) controller did not. 236
- 5.3 Simulation results for a measurement model of $\hat{x} = x - 0.4$ [m] and constant desired velocity of 1 [m/s]. (Left) Trajectories generated using the (BS-QP) controller. Solid line represents the true state, dashed line shows the estimated state, and green region indicates the safe set \mathcal{C} . The true trajectory fails to be safe and exits the safe set at $t = 3$ [s]. (Right) Trajectories generated using the (MR-BS-SOCP) controller. An additional robustness region is plotted in blue to indicate the set of true states which the control input renders safe. Both the true and measured trajectories are safe demonstrating the robustness of the (MR-BS-SOCP) controller when compared to the (BS-QP) controller. 246

- 5.4 Experimental results using SLAM from the onboard Intel RealSense T265 and constant desired velocity of 1 m/s. The notation and color schemes are the same as in Figure (5.3). (Left) Trajectories generated using the (BS-QP) controller. The true trajectory exits the safe set at $t = 6.7$ [s]. The measurement error is plotted in blue. (Right) Trajectories generated using the (MR-BS-SOCP) controller. Both the true and measured trajectories are safe demonstrating the robustness of the (MR-BS-SOCP) controller compared to the (BS-QP) controller. 248
- 5.5 Images from the experiment using the (MR-BS-SOCP) controller. The Segway is piloted towards a wall of yellow boxes and the controller ensures that it remains safe, i.e. that it does not crash into the boxes. (Top) Time lapse of the Segway trajectory. (Bottom) Camera images taken from the perspective of the Segway throughout the experiment. The images are displayed in chronological order from left to right. A video can be found at <https://vimeo.com/520247516>. 248
- 5.6 An overview of the Safety-Aware Preference-Based Learning design paradigm. Safety-Aware LineCoSpar is used to generate actions which are rolled out in experiments as parameters of the CBF-based safety filter to obtain user preferences and safety ordinal labels which are then used to update the user's estimated utility and generate new actions. 250
- 5.7 A comparison of Safety-Aware LineCoSpar and standard LineCoSpar on a synthetic utility function (drawn from the Gaussian prior) averaged over 50 runs with standard error shown by the shaded region. The safety-aware criteria reduces the number of sampled unsafe actions with a minimal effect on the prediction error, defined as $|\hat{\mathbf{a}}_i^* - \mathbf{a}^*|$ with $\hat{\mathbf{a}}_i^* \triangleq \operatorname{argmax}_{\mathbf{a}} \hat{r}_{S_i}$ and $\mathbf{a}^* \triangleq \operatorname{argmax}_{\mathbf{a}} r(\mathbf{a})$. 254

| | | |
|------|--|-----|
| 5.8 | Actions sampled during simulation in 30 iterations with 3 new actions in each iteration. The preferred action, $\hat{\mathbf{a}}_{30} = (3, 0.6, 0.5, 0.015)$, is shown in black and white. A conservative action, $\mathbf{a} = (2, 0.5, 0.0651, 0.485)$, is indicated by the black circle, where a and b were determined by estimating the Lipschitz coefficients present in the proof of Theorem 40. The conservative action fails to progress whereas Safety-Aware LineCoSpar provides an action which successfully navigates between obstacles. | 260 |
| 5.9 | Seven additional iterations of 3 actions executed indoors. The preferred action, $\hat{\mathbf{a}}_{37}^* = (4, 0.6, 0.4, 0)$, successfully traverses between the obstacles. | 261 |
| 5.10 | The preferred action, $\hat{\mathbf{a}}_{40}^* = (5, 0.1, 0.4, 0.02)$, after simulation, indoor experiments, and 3 additional iterations of 3 actions in an outdoor environment is shown alongside views from the onboard camera. Video can be found at https://youtu.be/QEuwRDTG7TE | 262 |
| 6.1 | Results using sample-and-hold implementation of (CLF-QP) controller on inverted pendulum at various sample rates. | 272 |
| 6.2 | Double Inverted Pendulum Schematic | 273 |
| 6.3 | Results using sample-and-hold implementation of (CLF-QP) controller on double inverted pendulum at various sample rates. . . | 274 |
| 6.4 | Planar Quadrotor Schematic | 275 |
| 6.5 | Results using sample-and-hold implementation of (CLF-QP) controller on planar quadrotor at various sample rates. | 276 |
| 6.6 | With inputs applied via a zero-order hold, the (CLF-QP) controller does not stabilize the system (6.84) (Top), while the (CLF-QCQP) does (Bottom). Simulation code listed at https://bit.ly/CLF-QCQP | 296 |
| 6.7 | Simulation results for the inverted pendulum system using the (CLF-QCQP) controller. | 297 |
| 6.8 | Simulation results for the double inverted pendulum system using the (CLF-QCQP) controller. | 298 |
| 6.9 | Simulation results for the planar quadrotor system using the (CLF-QCQP) controller. | 299 |

| | | |
|------|---|-----|
| 6.10 | Visualizing the properties (6.106)-(6.109) of SD-BF candidates. The dark green region represents the lower bound $T\alpha(h_T(\mathbf{x}_4))$ and $h_T(\mathbf{x}_6)$ cannot be in the red region due to the Lipschitz bound. | 307 |
| 6.11 | A visual representation of the main sets and three cases discussed in the proof of Theorem 44. | 309 |
| 6.12 | The single (left) and double (right) inverted pendulums and their sets \mathcal{C} . The inverted pendulum sets \mathcal{C} are, from left to right, a configuration ellipsoid, a halfspace, and a Lyapunov sublevel set. The far right image shows the set \mathcal{C} for the double inverted pendulum where $(\theta_1, \theta_2) \in \mathbb{R}^2$ are constrained to an ellipse and no constraint is placed on $(\dot{\theta}_1, \dot{\theta}_2) \in \mathbb{R}^2$ (represented by vertical fibers). | 315 |
| 6.13 | The maximum distance from the safe set \mathcal{C} (lower is better) achieved during trials vs. the sampling frequency. The simulations and animations be found at https://bit.ly/CBF-OP and https://vimeo.com/690803272 . (Top) The inverted pendulum for 3 different safe sets. (Bottom) The double inverted pendulum. | 316 |
| 6.14 | Simulation results for the system (6.167) using the trigger law (6.169). Even as the boundary of the safe set is approached ($h(\varphi_{\text{ET}}(t)) \rightarrow 0$), the growth rate of the error does not diminish, leading to arbitrarily small interevent times. | 325 |
| 6.15 | Simulation results demonstrating safety achieved with an event-triggered controller. (Top, Left) State trajectories for both triggers remain within the safe set for the length of the simulation. (Top, Right) The value of the ISSf-BF h remains above zero for the length of the simulation, corresponding to the system remaining safe. (Right) The interevent times of the two trigger laws. The interevent times of the trigger law (6.169) decreases towards 0 as predicted by Lemma 11 while the trigger law (6.189) satisfies the theoretical bound. | 329 |

LIST OF TABLES

| <i>Number</i> | <i>Page</i> |
|---|-------------|
| 1.1 General stability and safety properties. | 13 |
| 1.2 General unstable and unsafe properties. | 15 |
| 1.3 General partially robust stability and safety properties. | 16 |
| 4.1 Parameters sets used in (4.47) and (4.49) for the inverted pendulum example. | 186 |
| 4.2 Parameter values used in controller design. | 193 |
| 4.3 Sets of parameter values used for the exponential function (4.63) in the automated truck experiments with theoretical safety guarantee h^* , minimum experimental value of the TISSf-CBF h_{\min} , and shift in the steady-state tracking distance \tilde{D}_{ss} in (4.65). | 201 |
| 5.1 The safety-aware hyperparameters, and action space bounds (minimum and maximum) with discretizations Δ | 259 |

NOMENCLATURE

- aBF.** Adaptive Barrier Function.
- aCBF.** Adaptive Control Barrier Function.
- ACC.** Adaptive Cruise Control.
- aCLF.** Adaptive Control Lyapunov Function.
- aLF.** Adaptive Lyapunov Function.
- BF.** Barrier Function.
- BO.** Bayesian Optimization.
- BS.** Backup Set.
- CAN.** Controller Area Network.
- CBF.** Control Barrier Function.
- CCC.** Connected Cruise Controller.
- CCF.** Control Certificate Function.
- CCP.** Continuous Control Property.
- CLF.** Control Lyapunov Function.
- CTLE.** Continuous Time Lyapunov Equation.
- Dagger.** Dataset Aggregation.
- DT.** Discrete-Time.
- DT-BF.** Discrete-Time Barrier Function.
- DT-CBF.** Discrete-Time Control Barrier Function.
- ECU.** Engine Control Unit.
- ERM.** Empirical Risk Minimization.
- GPR.** Gaussian Process Regression.
- GPS.** Global Positioning System.
- i.i.d.** Independently and Identically Distributed.
- ID.** Inverse Dynamics.

IMU. Inertial Measurement Unit.

ISS. Input-to-State Stability.

ISS-CCP. Input-to-State Stable Continuous Control Property.

ISS-CLF. Input-to-State Stable Control Lyapunov Function.

ISS-LF. Input-to-State Stable Lyapunov Function.

ISS-SCP. Input-to-State Stable Small Control Property.

ISSf. Input-to-State Safety.

ISSf-BF. Input-to-State Safe Barrier Function.

ISSf-CBF. Input-to-State Safe Control Barrier Function.

LF. Lyapunov Function.

LQR. Linear Quadratic Regulator.

MAP. Maximum A Posteriori.

MIET. Minimum Interevent Time.

MPC. Model Predictive Control.

MR-BF. Measurement-Robust Barrier Function.

MR-CBF. Measurement-Robust Control Barrier Function.

MRISS. Measurement-Robust Implicit Safe Set.

NMPC. Nonlinear Model Predictive Control.

OBU. Onboard Unit.

PBL. Preference-Based Learning.

PD. Proportional-Derivative.

PSS. Projection-to-State Stability.

PSSf. Projection-to-State Safety.

QCQP. Quadratically Constrained Quadratic Program.

QP. Quadratic Program.

ReLU. Rectified Linear Unit.

RL. Reinforcement Learning.

ROS. Robot Operating System.

RTOS. Real-Time Operating Systems.

SCP. Small Control Property.

SGD. Stochastic Gradient Descent.

SLAM. Simultaneous Localization And Mapping.

SOCP. Second-Order Cone Program.

SOS. Sums-of-Squares.

TISSf. Tunable Input-to-State Safety.

TISSf-BF. Tunable Input-to-State Safe Barrier Function.

TISSf-CBF. Tunable Input-to-State State Safe Control Barrier Function.

V2X. Vehicle-to-Everything.

VIO. Visual Inertial Odometry.

Chapter 1

INTRODUCTION

The goal of this thesis is to make progress towards a complete theory of robust safety-critical control for autonomous systems. The technological developments of the past decade have set the stage for the rapid proliferation of autonomous systems that will fundamentally change the landscape of human activity, including labor, transportation, and healthcare. It is estimated that 87% of hours spent on production activities performed by workers in manufacturing roles can be automated [1], while it is predicted that by 2030, 12% of new passenger cars sold will feature L3+ autonomous driving technologies, with that number increasing to 37% by 2035 [2]. At the same time, there have been a number of high profile incidents of failures in autonomous driving systems [3], [4], including some that have produced fatal accidents for both passengers in the vehicle [5] and others outside the vehicle [6], [7]. These tragic failures reveal the potential risks with the expansion of autonomous systems into everyday life, and indicate that it is more important than ever for safety to be accounted for in the design of control systems.

I will build this theory upon Control Lyapunov Functions (CLFs) [8], [9] and Control Barrier Functions (CBFs) [10], [11], which are a foundation of model-based safety-critical control theory, but which can break down when deployed on real-world systems. The contributions of this thesis will extend the capabilities of CLFs and CBFs to provide a coherent approach for synthesizing robust safety-critical controllers that explicitly address the gap between system models and real-world systems. This direction is motivated by the fact that CLFs and CBFs have been shown to be a promising tool for the constructive synthesis of safety-critical controllers through a wealth of results demonstrating their capabilities in experimental applications including mobile robots [12]–[14], robotic swarms [15], autonomous aerial vehicles [16], robotic arms [17], [18], robotic manipulators [19], brachiating robots [20], automotive systems [21], quadrupedal robots [22], [23], and bipedal robots [24], [25].

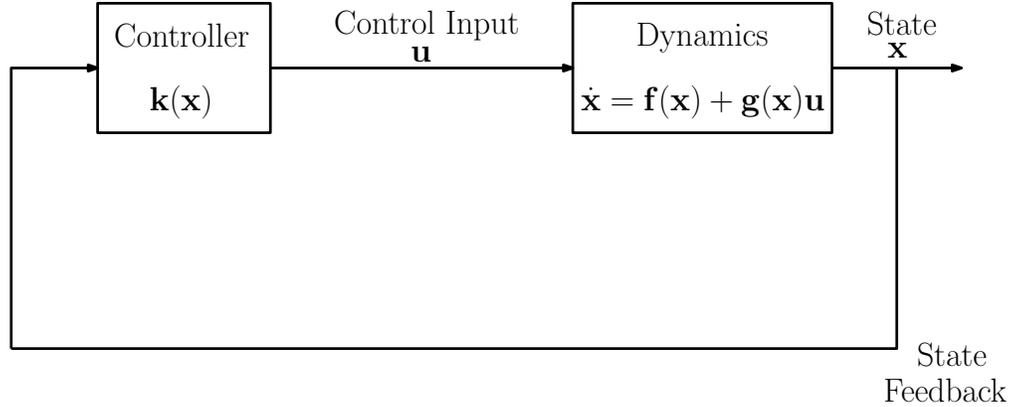


Figure 1.1. Ideal control system diagram.

The starting point of control design with CLFs and CBFs is often described by the ideal control system seen in Figure 1.1. The system that is to be controlled is captured by a nonlinear differential equation:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (1.1)$$

with state $\mathbf{x} \in \mathbb{R}^n$, input $\mathbf{u} \in \mathbb{R}^m$, and dynamics functions $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$. The goal of control design is to produce a state-feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that yields a closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}), \quad (1.2)$$

displaying desired stability and safety properties.

In practice, real systems are subject to a number of discrepancies from this ideal model, and are more accurately described by the control system seen in Figure 1.2. This diagram highlights four challenges faced when designing safety-critical controllers for real-world systems that I will address in this thesis. The first challenge, highlighted in red, is referred to as *model error*, and refers to the fact that access to the true dynamics in (1.1) is not available when designing a controller. Rather, a nominal model for the dynamics is utilized:

$$\hat{\dot{\mathbf{x}}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}), \quad (1.3)$$

where the functions $\hat{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are not necessarily equal to their true counterparts \mathbf{f} and \mathbf{g} . The second challenge, highlighted in green, is referred to as *disturbances*, and refers to both internal perturbations in the

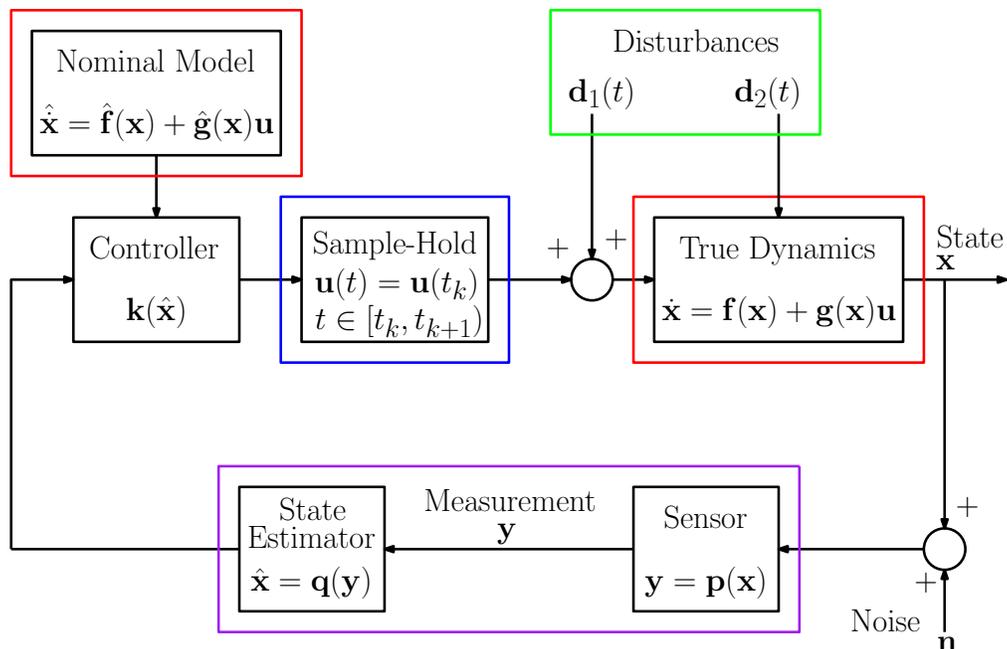


Figure 1.2. Practical control system diagram featuring model error, disturbances, measurement modules, and input sampling.

actuators of a system and external perturbations that are generally difficult to model, leading to a nonlinear system with disturbances:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})(\mathbf{k}(\mathbf{x}) + \mathbf{d}_1(t)) + \mathbf{d}_2(t), \quad (1.4)$$

where $\mathbf{d}_1 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ and $\mathbf{d}_2 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ capture disturbance signals. The third challenge, highlighted in purple, is referred to as *measurement error*, and refers to the fact the state \mathbf{x} is almost never perfectly measured and provided to a controller. Rather, an estimate of the state $\hat{\mathbf{x}}$ is formed through potentially imperfect measurement models and state estimators and used in a controller:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\hat{\mathbf{x}}). \quad (1.5)$$

Notably, the system dynamics \mathbf{f} and \mathbf{g} depend on the true state of the system, and thus if the estimate is imperfect, such that $\hat{\mathbf{x}} \neq \mathbf{x}$, the controller may not be using an accurate representation of the system dynamics. The last challenge, highlighted in blue, is referred to as *input sampling*, and refers to the fact that for cyber-physical systems such as robots, the underlying dynamics of the system evolve according to a differential equation that is continuous in time, while controllers are implemented on computing resources that are inherently discrete-time in nature, leading to a mixture of continuous and discrete-time

dynamics where inputs are held constant between sampling times t_k and t_{k+1} :

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t_k)), \quad t \in [t_k, t_{k+1}). \quad (1.6)$$

As I will show in the respective chapters of this thesis, neglecting these challenges in control design can lead to a failure to obtain the stability and safety properties that make the paradigm of CLF and CBF-based control so powerful. Thus, addressing these challenges will make substantial progress in developing a theory of robust safety-critical control.

A key focus in my development of this theory will be on obtaining robust safety-critical controllers that can be specified as *convex optimization* problems, which is a critical feature of the CLF and CBF-based control paradigm [11], [26], [27]. The importance of this focus is best captured by the insightful remark that “the great watershed in optimization isn’t between linearity and nonlinearity, but convexity and nonconvexity” [28]. Convex optimization [29], [30] provides a theoretically rigorous framework for minimizing a cost function subject to a collection of problem-specific constraints, both of which must possess the important property of convexity. By satisfying these requirements, a convex optimization problem enjoys two significant benefits. First, there exists strong statements about the ability to solve the problem numerically to an arbitrary degree of accuracy with a complexity that is polynomial in the problem dimensionality [31], [32]. This is a vital property for safety-critical control systems as decisions must often be made both accurately and quickly. Second, there exists significant software infrastructure and commercial solvers including tools such as CVXPY [33], YALMIP [34], ECOS [35], MOSEK [36], and Gurobi [37], enabling not only rapid-prototyping of control systems, but robust real-time implementations needed for real-world safety-critical systems. These benefits suggest that extending the CLF and CBF-based control paradigm to achieve robust theoretical guarantees and improved performance while preserving convexity of the control synthesis process is a valuable endeavor.

Summary of Contributions

The contributions of this thesis are a collection of methods for addressing the challenges of model error, disturbances, measurement error, and input sampling in safety-critical control. Utilizing the design primitives of CLFs and CBFs in conjunction with convex optimization, these methods constitute a framework for robust safety-critical control that efficiently addresses a wide set of challenges seen with real-world systems.

Below I provide a statement of related works pertaining to each contribution (identified by a bullet point). Note that a detailed discussion of related work on each challenge is presented at the beginning of each chapter.

Chapter III: Learning & Adaptive Control for Nonlinear Systems

There has been a recent wealth of results addressing model error in safety-critical control through data-driven learning techniques in conjunction with CLFs [17], [38]–[41] and CBFs [16], [41]–[45]. I will focus on three directions that have been unexplored in this work. The first is studying what properties can ensure that CLFs and CBFs constructively synthesized for a nominal model (such as with feedback linearization) can still be utilized in well-posed control design for the true system, allowing learning to be used to augment a nominal model-based controller that partially realizes a desired stability or safety task encoded by a CLF or CBF. The second is utilizing the fact that only model error that impacts the time derivative of a CLF or a CBF is relevant for achieving stability or safety, respectively, and that learning model error in the full-system dynamics is not required to achieve these objectives. The third is analyzing the impacts of residual learning error in order to formalize a relationship between learning accuracy and stability and safety.

- My first contribution is a collection of three episodic learning methods that capture model error directly as it impacts the evolution of CLFs and CBFs. I consider structural assumptions on the true system that enable transferring CLFs and CBFs from a nominal model to the true system. This allows me to utilize structured learning models that augment CLF and CBF-based controllers specified via convex optimization problems while preserving convexity and yielding significant improvements in stability and safety (Sections 3.3, 3.5, and 3.7). My second contribution is a characterization of the impact of residual learning errors on stability and safety built through the lens of Input-to-State Stability (ISS) and Input-to-State Safety (ISSf) (Sections 3.4 and 3.6). I demonstrate the episodic learning methods and analyze the effects of residual learning error experimentally on a Segway and bipedal robotic system (Sections 3.5, 3.6, and 3.7).

Establishing the preceding contributions reveals that the frequent inability to densely sample the control input space at a given state leads to potentially

large worst-case residual learning errors associated with input directions not seen in the data set. Conversely, residual learning error associated with input directions seen in the data set tends to be significantly smaller, motivating studying how to capitalize on these input directions during control synthesis. The only work studying this question with CLFs and CBFs developed in parallel to my contributions and used Gaussian Process Regression (GPR) with affine kernels and CLFs [46].

- My contribution is a data-driven control method that utilizes CLFs and CBFs with robust convex optimization to ensure the stability and safety of a system with only a partial characterization of model errors from data. I integrate a characterization of possible model errors that are permitted by a data set with CLF and CBF conditions to yield a controller specified by a convex second-order cone program (SOCP) that I prove feasibility results for and demonstrate in simulation (Section 3.8).

A key development in stabilizing adaptive control for nonlinear systems in the presence of parametric model error was the concept of an adaptive CLF introduced in [47] and further developed in [20], [48], [49]. Absent from this work has been a focus on adaptively enforcing the forward invariance of a particular set in the state space in the presence of parametric model error.

- My contribution is a method for safety-critical control of systems with parametric model error. Inspired by [47], I achieve this by proposing adaptive Control Barrier Functions as a tool for constructively synthesizing adaptive controllers that render a particular set in the state space forward invariant. I highlight challenges in adaptively ensuring forward invariance that are not seen in adaptive stabilization through a counterexample, and demonstrate the method in simulation (Section 3.9).

Chapter IV: Disturbance-Robust Safety Critical Control

The notion of CBFs that are robust to disturbances have primarily focused on incorporating bounds on the worst-case disturbance in the controller to enforce forward invariance of a particular set [50]–[55], which can often lead to unnecessarily conservative behavior if the bounds are not tight. Building upon the notion of Input-to-State Stability (ISS) [56], [57], an alternative notion of robustness that describes the expansion of a forward invariant set in

the presence of disturbances was captured by Input-to-State Safety (ISSf) and ISSf-CBFs in [58]. A challenge with robust control design through ISSf-CBFs is that it can still be difficult to balance meaningful safety guarantees (a small expansion of the forward invariant set) with conservativeness (restricting the system to deep within a forward invariant set), as I explore in Section 4.2.

- My contribution is an extension of ISSf-CBFs in the form of Tunable Input-to-State Safe CBFs (TISSf-CBFs), which permit robust yet performant control synthesis in the presence of disturbances (Section 4.3). I provide a thorough analysis and design of a TISSf-CBF-based controller that is experimentally demonstrated on a semi-trailer truck, the first demonstration of a CBF-based controller on such a system (Section 4.4).

Early work studied the impact of stochastic disturbances on stability guarantees through the perspective of Lyapunov functions [59], which was extended to the setting of barrier functions and set invariance in [60], [61]. The advent of CBFs in [11], [62] has led to a wide set of results addressing the safety of systems with stochastic disturbances, both in continuous-time [43], [63]–[71] and discrete-time systems [69], [72]. While the work in [69] considers using CBFs to provide finite-time safety guarantees for discrete-time systems with stochastic disturbances using tools from [59], it did not address how to synthesize control policies given a CBF for a system with stochastic disturbances.

- My contribution is an approach for synthesizing control policies given a CBF for a discrete-time system with stochastic disturbances. Inspired by both [69] and [59], I propose a method for establishing ISSf-like finite-time safety guarantees for discrete-time systems subject to stochastic disturbances. I explore how convexity properties of a CBF can be used to permit convex optimization-based controllers that provide such safety guarantees, which are demonstrated in simulation (Section 4.5).

Chapter V: Measurement-Robust Safety Critical Control

The work in [50] was the first to consider the impact of measurement errors on CBF-based controllers by incorporating the covariance of an unscented Kalman filter into the standard CBF inequality as an additive robustness term. Subsequently, additional work has considered measurement noise using stochastic differential equations [66], required a restricted sub-tangentiality

condition on the boundary of a set that is to be kept forward invariant [73], or studied measurement error using interval arithmetic and incremental stability [74]. Absent from these works was a consideration of the interaction between measurement errors and convexity as it impacts the control synthesis process.

- My contribution is the definition of Measurement-Robust CBFs (MR-CBFs), which permit safe control synthesis through convex optimization in the presence of measurement errors (Section 5.2), and the experimental demonstration of an MR-CBF-based controller specified via a convex SOCP on a Segway platform using a vision-based measurement system (Section 5.3). This experimental demonstration validates that the controllers proposed throughout this thesis that are specified via increasingly complex classes of convex optimization problems can be solved efficiently enough to be deployed on real-world hardware platforms.

Preference-Based Learning (PBL) provides an approach for searching complex parameter spaces via subjective feedback, without an explicitly defined reward function [75]. In the context of control systems, there has recently been significant effort in developing PBL approaches for tuning existing tools built from control theory [76]–[80]. Notably, the use of PBL to tune parameters of CBF controllers to balance the tradeoff between conservativeness in safety and performance has yet to be considered.

- I will present an application of PBL to tune robustness parameters of a CBF-based controller to achieve both safe and performant behavior on an experimental quadrupedal robotic system. My particular contributions are the conceptualization of the idea to use PBL to tune robustness parameters in a CBF-based controller, and an analysis of safety guarantees using reduced-order models in the presence of disturbances and measurement error (Section 5.4). This work highlights a tool for tuning the robust safety-critical control methods presented in this thesis.

Chapter VI: Sampled-Data & Event-Triggered Control

A key development in the design of stabilizing controllers for sampled-data systems came with the framework using approximate discrete-time models proposed in [81], [82], which led to a wide range of results focused on synthesizing controllers for sampled-data systems using approximate discrete-time

models [83]–[88]. This work has not addressed the design of controllers for sampled-data systems with approximate discrete-time dynamics using CLFs and convex optimization, which I highlight the importance of by demonstrating the fragility of continuous-time CLF-based controllers specified through convex optimization to input sampling in Section 6.2.

- My contribution is an extension of the work in [81] that uses approximate discrete-time models in conjunction with CLFs to synthesize stabilizing convex optimization-based controllers that significantly outperforms continuous-time CLF-based controllers. I address how a CLF synthesized for a system’s continuous-time dynamics through feedback linearization possesses the necessary properties of a CLF for a set of approximate discrete-time dynamics, and show how stability of zero-dynamics is preserved with input sampling, providing a full practical stability result for a partially feedback linearizable system (Section 6.3).

There has been a number of recent works exploring safety-critical control for sampled-data systems from the perspective of CBFs, [19], [73], [74], [89]–[93] much of which has taken an emulation approach, where an additive correction term is factored in to the standard continuous-time CBF inequality. This correction term often relies on Lipschitz constants of the system dynamics and bounds on the control inputs, which can be difficult to produce, and over-approximations will often lead to conservative behavior that can only be resolved with exceptionally high sample rates [90]. None of this work has considered using approximate discrete-time models along the lines of the work in [81], which often works exceptionally well even at low sample rates [94].

- My contribution is a method for studying safety of sampled-data systems using approximate discrete-time models. I propose a notion of practical safety and Sampled-Data CBFs that can be utilized with approximate discrete-time models to enable safety-critical control synthesis through convex optimization (Section 6.4), which I demonstrate in simulation.

Event-triggered control is a paradigm for effectively using actuation resources when it is costly to switch the inputs to the system [95]. The work developed in [96] utilized Input-to-State Stable Lyapunov functions (ISS-LFs) to capitalize on inherent robustness in a continuous-time system to produce a

stabilizing event-triggered control implementation while maintaining a minimum interevent time (MIET), which ensures that events (and input switching) do not occur arbitrarily close together [97]. Notably, event-triggered work using CBFs has primarily focused on how they can be used to improve stability [89], [98], rather than trying to enforce the forward invariance of a specific set.

- My contribution is a method for safety-critical control of event-triggered systems. Drawing inspiration from [96], I utilize the concept of ISSf barrier functions to provide guarantees of both safety and a minimum interevent time, which I show in simulation (Section 6.5).

Outline of Thesis

Chapter II: Background presents a review of background material on core tools from nonlinear control theory that will be utilized throughout this thesis, including nonlinear dynamics, stability, safety, CLFs, feedback linearization, CBFs, nonlinear dynamics with disturbances, ISS and ISSf, and their corresponding variants of CLFs and CBFs. While this section is not a contribution of the thesis, it provides pedagogical examples highlighting some of the finer points of these concepts.

Chapter III: Learning & Adaptive Control for Nonlinear Systems presents my first set of contributions on the topic of addressing model error as it impacts safety-critical control. I begin by exploring related work addressing model error in the context of CLF and CBF-based controllers as well safety-critical adaptive control. Next, I present a description of general and parametric model error. In Section 3.3 I present an episodic learning method for mitigating model error in the time derivative of a CLF, and establish theoretical results on the robustness of stability to residual learning error in Section 3.4. I extend this episodic learning method to the setting of CBFs in Section 3.5 with an experimental demonstration on a Segway platform, and provide theoretical safety guarantees in the presence of residual learning error in Section 3.6. In Section 3.7 I study some of the challenges with the preceding learning method for high-dimensional systems, and propose an alternative episodic learning method that overcomes these challenges and is demonstrated experimentally on a bipedal robotic platform. Building off of these previous results, in Section 3.8 I present a data-driven control method that meets stability or safety requirements in the presence of a partial characterization of

model error from data using robust convex optimization. Lastly, in Section 3.9, I provide a method for safety-critical control in the presence of parametric model error through the notion of adaptive CBFs (aCBFs). I conclude with thoughts on directions for future work.

Chapter IV: Disturbance-Robust Safety-Critical Control presents my second set of contributions on the topic of addressing robustness to disturbances in safety-critical control. I begin by exploring related work on disturbance-robust safety critical control and safety-critical control in the presence of stochastic disturbances. In Section 4.2 I present some of the limitations of ISSf and ISSf-CBFs as posed in [58] regarding robustness and performance. Drawing inspiration from this, in Section 4.3 I present an extension of ISSf in the form of Tunable Input-to-State Safety (TISSf), which prioritizes robustness to disturbances near the boundary of a set that is to be kept forward invariant as opposed to equally throughout the set. After demonstrating TISSf on a simple inverted pendulum example, in Section 4.4 I present a detailed deployment of TISSf-based controllers on an experimental semi-trailer truck, once again showing that TISSf can achieve safety in the presence of disturbances without significantly conceding performance. Lastly, in Section 4.5 I present a method for safety-critical control of discrete-time systems in the presence of stochastic disturbances. I conclude with directions for future work.

Chapter V: Measurement-Robust Safety-Critical Control presents my third set of contributions on the topic of addressing robustness to measurement errors in safety-critical control. I begin by exploring related work in the areas of measurement-robust safety critical control and the tuning of controllers using PBL. In Section 5.2 I present a definition of Measurement-Robust Control Barrier Functions (MR-CBFs) that enable robust control synthesis through convex optimization for systems with measurement errors. Following this, in Section 5.3 I present the integration of MR-CBFs with backup set CBF methods and an experimental demonstration of a MR-CBF controller specified via a SOCP on an experimental Segway system. Lastly, in Section 5.4 I present results on using PBL to tune disturbance and measurement robustness parameters in a CBF-based controller to achieve both safety and performance.

Chapter VI: Sampled-Data & Event-Triggered Control presents my fourth set of contributions on the topic of addressing robustness to sampling in safety-critical control. I will begin exploring related work in the areas

of sampled-data control through approximate discrete-time models, sampled-data control and feedback linearization, and event-triggered control. In Section 6.2 I highlight the fragility of continuous-time CLF-based controllers specified through convex optimization to sampling with simulation results and define the sampled-data and event-triggered sampling paradigms. Following this, in Section 6.3 I present the design of stabilizing controllers using approximate discrete-time models and CLFs through convex optimization, with a focus on how feedback linearizability of a continuous-time system can produce a CLF for an approximate discrete-time model. I revisit the simulation settings in Section 6.2, showing that the proposed control approach greatly outperforms the continuous-time approach. In Section 6.4 I present a framework for safety-critical sampled-data control using approximate discrete-time models and CBFs. I develop notions of practical safety, Sampled-Data barrier functions, and Sampled-Data CBFs, which I use to produce controllers through convex optimization that I demonstrate in simulation. Lastly, in Section 6.5, I present a framework for event-triggered safety-critical control. I review stabilizing event-triggered control as presented in [96], and show how naïvely transferring the methods of this paper to the safety-critical setting fail to produce a minimum interevent time with a counterexample. I then propose modifications that yield both safety and a minimum interevent time, which I demonstrate in simulation. I conclude with directions for future work.

Chapter VII: Conclusion provides my final conclusions, as well as a discussion on future work that extends beyond the results presented in the thesis.

1.1 General Mathematical Outlook

This section briefly presents the mathematical methodology of robust control that will be utilized throughout this thesis. There are two primary goals of this presentation. The first is to highlight the two fundamental perspectives on robustness that underlie the collection of results in this thesis. It is my hope that readers will discover that many challenges in robust safety-critical control are connected by the types of outcomes that can be produced through robust control synthesis. The second is to preface the technical contributions of each section in an effort to prepare a reader with an idea of the destination that mathematical developments are building towards. This math in this section is presented at an informal level, and should be interpreted to serve merely as a collection of guide posts to the more rigorous theoretical descriptions

provided in the subsequent chapters. A reader well-versed in nonlinear systems, Control Lyapunov Functions (CLFs), Control Barrier Functions (CBFs), Input-to-State Stability (ISS), and Input-to-State Safety (ISSf) will find the details of this section establish a conceptual understanding of what types of technical contributions will be made in this thesis, while readers newer to these concepts will be best served by reading the background material in Chapter 2 before returning to this section.

Nonlinear Dynamics

Consider a nonlinear control-affine system:

$$\hat{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (1.7)$$

with state $\mathbf{x} \in \mathbb{R}^n$, input $\mathbf{u} \in \mathbb{R}^m$, and functions $\hat{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$. Given a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we may define a closed-loop system:

$$\hat{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}). \quad (1.8)$$

Assume that for any initial condition in $\mathbf{x}_0 \in \mathbb{R}^n$, there exists a unique continuously differentiable solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ solving this closed-loop system.

Stability & Safety

Stability and safety properties for these types of system will be encoded through a scalar function $\mathbf{C} : \mathbb{R}^n \rightarrow \mathbb{R}$ that is continuously differentiable on \mathbb{R}^n . More specifically, stability and safety properties will require that:

| | |
|-----------|--|
| Stability | $\lim_{t \rightarrow \infty} \mathbf{C}(\varphi(t)) = 0$ |
| Safety | $\mathbf{C}(\varphi(t)) \leq 0, \forall t \in \mathbb{R}_{\geq 0}$ |

Table 1.1. General stability and safety properties.

Achieving this goal through control will rely on the fact that the time derivative of \mathbf{C} can be expressed as:

$$\hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) = \frac{\partial \mathbf{C}}{\partial \mathbf{x}}(\mathbf{x}) \left(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} \right), \quad (1.9)$$

noting that the input \mathbf{u} appears. Consequently, we will seek to design a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that satisfies:

$$\hat{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -\alpha \mathbf{C}(\mathbf{x}), \quad (1.10)$$

for some $\alpha \in \mathbb{R}_{>0}$, as such a controller will lead to the stability and safety properties defined above. The ability to find a controller satisfying this inequality will be the special property that defines the function \mathcal{C} as a CLF or CBF in the settings of stability and safety, respectively. An optimization-based controller structure that will serve as a starting point for more complex robust control formulations that enforce this derivative inequality will be given by:

$$\mathbf{k}(\mathbf{x}) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \quad (1.11)$$

$$\text{s.t. } \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) \leq -\alpha \mathcal{C}(\mathbf{x}), \quad (1.12)$$

where $\mathbf{k}_{\text{nom}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is some nominal controller. Importantly, this controller is specified by a convex quadratic program, ensuring that it can be efficiently implemented on real-world hardware platforms.

Modeling Inaccuracies

The modeling inaccuracies of model error, disturbances, measurement error, and input sampling will lead to a deviation from the nominal model in (1.7) in various ways. Instead, let us consider a nonlinear control-affine system:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{u} + \mathbf{\Delta}(\mathbf{x}, \mathbf{u}, t), \quad (1.13)$$

with a function $\mathbf{\Delta} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ capturing a general perturbation to the nominal model in (1.7). Note that this general perturbation not only can depend on the state of the system, but it may depend on the input and time. Given a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, denote the closed-loop system:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}) + \mathbf{\Delta}(\mathbf{x}, \mathbf{k}(\mathbf{x}), t), \quad (1.14)$$

and assume that for any initial condition $\mathbf{x}_0 \in \mathbb{R}^n$, there exists a unique continuously differentiable solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ to this closed-loop system. The time derivative of \mathcal{C} is given by:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}, t) = \underbrace{\frac{\partial \mathcal{C}}{\partial \mathbf{x}}(\mathbf{x}) \left(\widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{u} \right)}_{\widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u})} + \underbrace{\frac{\partial \mathcal{C}}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{\Delta}(\mathbf{x}, \mathbf{u}, t)}_{\delta(\mathbf{x}, \mathbf{u}, t)}. \quad (1.15)$$

Observe that the perturbation $\mathbf{\Delta}$ affects the time derivative of \mathcal{C} in the form of δ and will play a role in the ability to achieve stability and safety properties.

No Robustness

The idea of robustness can be seen as living on a spectrum. At one end, we can forgo any effort in ensuring our control design is robust to Δ by designing a controller that simply satisfies (1.10). Consequently, we have that:

$$\dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x}), t) \leq -\alpha\mathbf{C}(\mathbf{x}) + \delta(\mathbf{x}, \mathbf{k}(\mathbf{x}), t). \quad (1.16)$$

The presence of the term $\delta(\mathbf{x}, \mathbf{k}(\mathbf{x}), t)$ without further assumptions will typically lead to a failure in rigorous stability and safety guarantees, such that one may have an unstable or unsafe system:

| | |
|----------|---|
| Unstable | $\lim_{t \rightarrow \infty} \mathbf{C}(\varphi(t)) = \infty$ |
| Unsafe | $\exists t \in \mathbb{R}_{\geq 0}$ s.t. $\mathbf{C}(\varphi(t)) > 0$ |

Table 1.2. General unstable and unsafe properties.

Absolute Robustness

At the opposite end of the spectrum is an *absolute* form of robustness, in which we will seek to design a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that:

$$\dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x}), t) \leq -\alpha\mathbf{C}(\mathbf{x}), \quad (1.17)$$

and therefore we will recover the desired stability and safety properties by effectively eliminating the affect of δ on the time derivative of \mathbf{C} . This perspective on robustness is akin to binary notion of robustness in linear systems theory [99], in which a system is either stable, or unstable. In this case, the nonlinear system is either stable (safe), or unstable (unsafe). As we will see in this thesis, this perspective on robustness typically leads to more complex controllers that account for the value of $\delta(\mathbf{x}, \mathbf{k}(\mathbf{x}), t)$ when specifying \mathbf{k} .

Partial Robustness

In between these two ends of the robustness spectrum is an approach that seeks to manage the degradation of stability or safety properties. In particular, control design is done to produce controllers that achieve:

$$\dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x}), t) \leq -\alpha\mathbf{C}(\mathbf{x}) + \delta(\mathbf{x}, \mathbf{k}(\mathbf{x}), t) \leq -\alpha\mathbf{C}(\mathbf{x}) + \bar{\delta}, \quad (1.18)$$

for some $\bar{\delta} \in \mathbb{R}_{\geq 0}$. By achieving such a bound, one may establish the following partially robust stability and safety properties for a function $R : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$:

| | |
|----------------------------|--|
| Partially Robust Stability | $\exists t^* \in \mathbb{R}_{\geq 0}$ s.t. $ \mathbf{C}(\boldsymbol{\varphi}(t)) \leq R(\bar{\delta}) \forall t \geq t^*$ |
| Partially Robust Safety | $\mathbf{C}(\boldsymbol{\varphi}(t)) \leq R(\bar{\delta}) \forall t \in \mathbb{R}_{\geq 0}$ |

Table 1.3. General partially robust stability and safety properties.

An important property of the function R is that $\lim_{\bar{\delta} \rightarrow 0^+} R(\bar{\delta}) = 0$, such that for small perturbation bounds $\bar{\delta}$, the degradation in stability and safety properties is small. This concept of robustness was first proposed in [56] in the form of Input-to-State Stability (ISS) and later as Input-to-State Safety (ISSf) in [58]. Instead of trying to completely eliminate $\delta(\mathbf{x}, \mathbf{k}(\mathbf{x}), t)$, one tries to manage it to some value ($\bar{\delta}$) that is tolerable in terms of its impact on stability and safety properties. While I will explicitly use the technical definitions of ISS and ISSf throughout this thesis when addressing disturbances (for which they were originally posed), I will also make use of this idea of relaxing stability and safety guarantees in a amount proportional to a perturbation at a conceptual level throughout this thesis for settings outside of disturbances (such as residual learning error and input sampling effects).

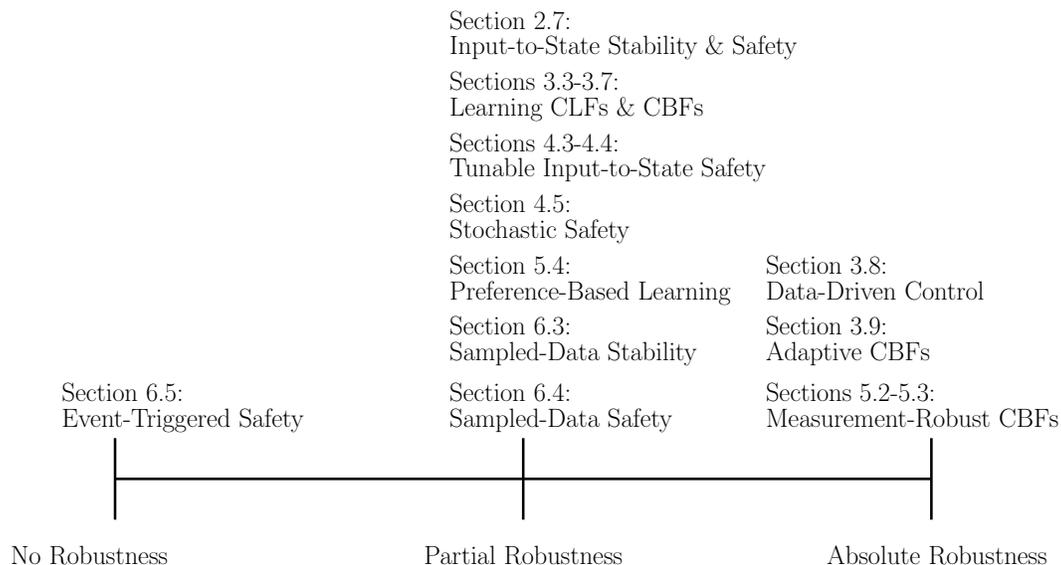


Figure 1.3. Spectrum of robustness properties of results presented in this thesis.

Chapter 3: Model Error

In Chapter 3 I will explore the problem of *model error*. In Sections 3.3-3.8 I will consider a general form of model error yielding the perturbation:

$$\Delta(\mathbf{x}, \mathbf{u}) = \underbrace{\mathbf{f}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x})}_{\tilde{\mathbf{f}}(\mathbf{x})} + \underbrace{(\mathbf{g}(\mathbf{x}) - \hat{\mathbf{g}}(\mathbf{x}))}_{\tilde{\mathbf{g}}(\mathbf{x})} \mathbf{u}, \quad (1.19)$$

with unknown functions $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ which capture the *true* dynamics of the system. These unknown functions lead to the model error terms $\tilde{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\tilde{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$. Importantly, the time derivative of the function \mathcal{C} will be given by:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}) = \hat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \underbrace{\frac{\partial \mathcal{C}}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x})}_{b(\mathbf{x})} + \underbrace{\frac{\partial \mathcal{C}}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{g}}(\mathbf{x}) \mathbf{u}}_{\mathbf{a}(\mathbf{x})^\top}. \quad (1.20)$$

As previously noted, neglecting the function δ can lead to unstable and unsafe behavior.

Sections 3.3-3.6: Learning CLFs and CBFs

The focus in Sections 3.3 (CLFs and stability) and 3.5 (CBFs and safety) will be attempting to learn the functions b and \mathbf{a} using data-driven estimators, such that an estimator for δ can be constructed as follows:

$$\hat{\delta}(\mathbf{x}, \mathbf{u}) = \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u} \approx b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u}, \quad (1.21)$$

with $\hat{\delta} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$, $\hat{b} : \mathbb{R}^n \rightarrow \mathbb{R}$, and $\hat{\mathbf{a}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The control-affine form of this estimator is explicitly enforced in its construction to ensure that it can be integrated into a convex optimization-based controller:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ \text{s.t. } &\hat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u} \leq -\alpha \mathcal{C}(\mathbf{x}). \end{aligned} \quad (1.22)$$

Challenges that arise due to non-independently and identically distributed training data will motivate constructing the estimator $\hat{\delta}$ in an episodic manner in which data is collected, learning models are trained, a new controller is synthesized using the learning models, and the new controller is deployed to produce additional data that is aggregated with previous data. The learning-informed controllers will enforce that:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -\alpha \mathcal{C}(\mathbf{x}) + \underbrace{b(\mathbf{x}) - \hat{b}(\mathbf{x})}_{\tilde{b}(\mathbf{x})} + \underbrace{(\mathbf{a}(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x}))^\top \mathbf{k}(\mathbf{x})}_{\tilde{\mathbf{a}}(\mathbf{x})^\top}. \quad (1.23)$$

Notably, *residual learning error* leads to the function δ' appearing, which will still degrade the properties of stability and safety. In Sections 3.4 (Projection-to-State Stability) and 3.6 (Projection-to-State Safety), I will establish theoretical stability and safety properties of the form:

Proposition 1 (Informal). *For a sufficiently large (and problem dependent) set $D \subseteq \mathbb{R}^n$, we have that:*

$$|\delta'(\mathbf{x}, \mathbf{k}(\mathbf{x}))| \leq \bar{\delta} \quad \forall \mathbf{x} \in D \implies \begin{cases} \text{Section 3.4} & \exists t^* \in \mathbb{R}_{\geq 0} \text{ s.t.} \\ & |\mathcal{C}(\boldsymbol{\varphi}(t))| \leq R(\bar{\delta}) \quad \forall t \geq t^* \\ \text{Section 3.6} & \mathcal{C}(\boldsymbol{\varphi}(t)) \leq R(\bar{\delta}) \quad \forall t \in \mathbb{R}_{\geq 0}. \end{cases},$$

This informal result captures the fact that when using CLF and CBF-based controllers, residual learning error present after identifying the functions b and \mathbf{a} with corresponding estimators leads to degradation in stability and safety properties that is proportional to the magnitude of residual learning error. Importantly, small residual learning errors do not lead to catastrophic failures in stability and safety. This is aligned with the notion of partial robustness described above.

Section 3.7: Projected Disturbance Learning

In Section 3.7 I will explore some of the challenges related to data richness that arise using the control-affine learning models in Sections 3.3-3.6. In particular, given a data point $(\mathbf{x}_i, \mathbf{u}_i, \dot{\mathbf{c}}_i) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}$, there are infinitely many values of the estimators $\hat{b}(\mathbf{x}_i)$ and $\hat{\mathbf{a}}(\mathbf{x}_i)$ that satisfy:

$$\dot{\mathbf{c}}_i = \hat{\mathbf{C}}(\mathbf{x}_i, \mathbf{u}_i) + \hat{b}(\mathbf{x}_i) + \hat{\mathbf{a}}(\mathbf{x}_i)^\top \mathbf{u}_i, \quad (1.24)$$

as the problem can be viewed as an underdetermined linear system. In the absence of a variety of inputs in the data set (which is typical of high-dimensional systems), this challenge can lead to poor estimators and a large residual learning error bounds $\bar{\delta}$ that significantly degrade stability and safety properties. Instead, given a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, a purely state-based estimator can be constructed:

$$\hat{\delta}(\mathbf{x}) \approx b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}). \quad (1.25)$$

This estimator is built to capture the difference between $\dot{\mathbf{c}}$ and $\hat{\mathbf{C}}$ using the particular controller \mathbf{k} . Importantly, it does not have the same underdetermined problem that appeared with the control-affine estimator and its accuracy can be evaluated when it is built. The estimator can subsequently be integrated into an optimization-based controller as follows:

$$\begin{aligned} \mathbf{k}'(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ \text{s.t.} & \hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \hat{\delta}(\mathbf{x}) \leq -\alpha \mathcal{C}(\mathbf{x}). \end{aligned} \quad (1.26)$$

A consequence of this state based approach is that we then have that:

$$\dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}'(\mathbf{x})) \leq -\alpha\mathbf{C}(\mathbf{x}) - \widehat{\delta}(\mathbf{x}) + \underbrace{\delta(\mathbf{x}, \mathbf{k}'(\mathbf{x}))}_{\delta'(\mathbf{x})}, \quad (1.27)$$

such that a new perturbation $\delta' : \mathbb{R}^n \rightarrow \mathbb{R}$ appears due to using the new controller \mathbf{k}' , while the estimator $\widehat{\delta}$ was built based on information from using the previous controller \mathbf{k} . Despite this difference, the structure of the optimization-based controller \mathbf{k}' will enable the following result:

Proposition 2 (Informal). *If the estimator $\widehat{\delta}$ is an accurate estimator of the function $\delta(\cdot, \mathbf{k}(\cdot))$, then:*

$$\|\mathbf{a}(\mathbf{x})\| \leq \left\| \frac{\partial \mathbf{C}}{\partial \mathbf{x}}(\mathbf{x}) \widehat{\mathbf{g}}(\mathbf{x}) \right\| \text{ and } \widehat{\delta}(\mathbf{x}) \leq 0 \implies \dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}'(\mathbf{x})) \leq \dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x})). \quad (1.28)$$

The key part of this result is the first part of the precedent, which states that the model error affecting the input to the system is small compared to how the nominal model says input will affect the system. If this true, and the estimator $\widehat{\delta}$ built using the previous controller \mathbf{k} indicates the system is not meeting the necessary inequality on $\dot{\mathbf{C}}$ at the state \mathbf{x} (the second term in the precedent), the new controller \mathbf{k}' will improve upon the stability and safety of the system at the state \mathbf{x} .

Section 3.8: Data-Driven Control

In Section 3.8 I will explore an alternative approach to resolving the issues with control-affine estimators by using robust convex optimization. I will construct a convex set $\mathcal{U}(\mathbf{x}) \subset \mathbb{R}^{m+1}$ of all possible values of the model errors $(b(\mathbf{x}), \mathbf{a}(\mathbf{x}))$ that are permissible with a data set:

$$(b(\mathbf{x}), \mathbf{a}(\mathbf{x})) \in \mathcal{U}(\mathbf{x}). \quad (1.29)$$

These sets can be integrated into an optimization-based controller as:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ \text{s.t. } \widehat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + q + \mathbf{p}^\top \mathbf{u} \leq -\alpha\mathbf{C}(\mathbf{x}), \quad \forall (q, \mathbf{p}) \in \mathcal{U}(\mathbf{x}). \end{aligned} \quad (1.30)$$

This controller has an infinite number of constraints (one for each $(q, \mathbf{p}) \in \mathcal{U}(\mathbf{x})$). Addressing this leads to the following result:

Proposition 3 (Informal). *The controller \mathbf{k} can be represented as a finite-dimensional convex second-order cone program (SOCP) and under Lipschitz continuity assumptions of $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$, leads to:*

$$\dot{\mathbf{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -\alpha \mathbf{C}(\mathbf{x}). \quad (1.31)$$

This result state that is possible to tractably implement the robust controller, and by doing so, ensure that stability and safety requirements are satisfied for the true values of the model error ($b(\mathbf{x}), \mathbf{a}(\mathbf{x})$). I note that this aligns with the form of absolute robustness described above, where the original stability or safety requirement is still met in the presence of a perturbation.

Section 3.9: Adaptive CBFs

In Section 3.9 I consider a special form of parametric model error of the form:

$$\Delta(\mathbf{x}) = \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^*, \quad (1.32)$$

where $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times p}$ is a known collection of basis functions and $\boldsymbol{\theta}^*$ is an unknown vector of parameters. I note that this setting does not consider model error impacting how input affects the system. Motivated by the work studying adaptive CLFs in [47], I will consider an adaptive control approach for ensuring safety. In particular, I will explore controllers that maintain an estimate of the parameters, $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$, and update them according to an adaptation law:

$$\dot{\hat{\boldsymbol{\theta}}} = \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}), \quad (1.33)$$

where $\boldsymbol{\tau} : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^p$ is the adaptation law. This will result in the definition of adaptive Control Barrier Functions, which lead to optimization-based controllers of the form:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ \text{s.t. } &\hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \frac{\partial \mathbf{C}}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x})\hat{\boldsymbol{\theta}} \leq 0, \end{aligned} \quad (1.34)$$

that utilize the parameter estimate $\hat{\boldsymbol{\theta}}$ and a (stronger in this case) bound on the time derivative of \mathbf{C} . This type of adaptation law and controller lead to the following result:

Proposition 4 (Informal). *For any initial parameter estimate $\hat{\boldsymbol{\theta}}_0 \in \mathbb{R}^p$, there exists an update law (with potentially high adaptive gains) such that the controller \mathbf{k} enforces:*

$$\mathbf{C}(\boldsymbol{\varphi}(t)) \leq 0, \quad \forall t \in \mathbb{R}_{\geq 0}. \quad (1.35)$$

This result states that it is possible to produce an adaptation law and controller that ensures a system remains safe in the presence of potentially large parametric model errors. Furthermore, this is an instance of absolute robustness, where a safety requirement is met despite the perturbation Δ .

Chapter 4: Disturbances

In Chapter 4 I will explore the problem of *disturbances*. In Sections 4.2-4.4 I will consider a perturbation of the form:

$$\Delta(\mathbf{x}, t) = \mathbf{g}(\mathbf{x})\mathbf{d}_1(t) + \mathbf{d}_2(t), \quad (1.36)$$

with *matched* disturbance signal $\mathbf{d}_1 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ and *unmatched* disturbance signal $\mathbf{d}_2 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$. The difference between matched and unmatched disturbances will matter in transferring CLFs and CBFs for systems without disturbances to systems with disturbances. Both types of disturbance signal can capture a very large class of unstructured perturbations to a system, and are generally assumed not to have any known underlying structure that might be identifiable with data. The time derivative of the function \mathcal{C} is given by:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}, t) = \hat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \overbrace{\frac{\partial \mathcal{C}}{\partial \mathbf{x}}(\mathbf{x}) (\mathbf{g}(\mathbf{x})\mathbf{d}_1(t) + \mathbf{d}_2(t))}^{\delta(\mathbf{x}, t)}. \quad (1.37)$$

Once again, ignoring δ can lead to unstable and unsafe behavior.

Section 4.2 Limits of Input-to-State Safety

In Section 4.2 I will explore some of the limitations of the partial robustness notion of Input-to-State Safety (ISSf). In particular, ISSf-based design focuses on producing a controller $\mathbf{k} : \mathbb{R}^n \times \mathbb{R}_{>0} \rightarrow \mathbb{R}^m$ such that:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x}, \epsilon), t) \leq -\alpha \mathcal{C}(\mathbf{x}) + \bar{\delta}(\epsilon), \quad (1.38)$$

where $\epsilon \in \mathbb{R}_{>0}$ is a robustness parameter that determines the size of the perturbation $\bar{\delta}(\epsilon)$ related to disturbances in the time derivative of $\dot{\mathcal{C}}$. As ISSf is a notion of partial robustness, this leads to the outcome that:

$$\mathcal{C}(\varphi(t)) \leq R(\bar{\delta}(\epsilon)), \quad \forall t \in \mathbb{R}_{\geq 0}. \quad (1.39)$$

In this way, it is possible to change the degradation of safety guarantees through the parameter ϵ .

In practice, this design approach struggles to balance performance and theoretical safety guarantees. In particular, for small values of ϵ , one achieves a strong theoretical statement about safety, such that $R(\bar{\delta}(\epsilon)) \approx 0$, but the closed-loop system displays extremely conservative behavior, such that $\mathcal{C}(\varphi(t)) \ll 0$ for all $t \in \mathbb{R}_{\geq 0}$. For large values of ϵ , one does not have that $\mathcal{C}(\varphi(t)) \ll 0$, such that performance is recovered, but instead vacuous safety guarantees occur as the value of $R(\bar{\delta}(\epsilon))$ can be quite large.

Sections 4.3 and 4.4: Tunable Input-to-State Safety

In Sections 4.3 and 4.4 the limits of ISSf will be approached by viewing the robustness parameter ϵ as a function of the value of $\mathcal{C}(\mathbf{x})$. This will allow prioritizing robustness when the system is close to unsafe ($\epsilon(\mathcal{C}(\mathbf{x}))$ is small when $\mathcal{C}(\mathbf{x}) \approx 0$) and relaxing robustness when the system is very safe ($\epsilon(\mathcal{C}(\mathbf{x}))$ is large when $\mathcal{C}(\mathbf{x}) \ll 0$). In particular, a bound of the form:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{k}(\mathbf{x}, \epsilon(\mathcal{C}(\mathbf{x}))), t) \leq -\alpha\mathcal{C}(\mathbf{x}) + \bar{\delta}(\epsilon(\mathcal{C}(\mathbf{x}))), \quad (1.40)$$

will be achieved. Theoretical safety properties with this modification will be encoded by *Tunable* Input-to-State Safety, leading to a result of the form:

Proposition 5 (Informal). *For a monotonically decreasing function $\epsilon : \mathbb{R} \rightarrow \mathbb{R}_{>0}$, the controller $\mathbf{k} : \mathbb{R}^n \times \mathbb{R}_{>0} \rightarrow \mathbb{R}^m$ using a parameter function $\epsilon : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ leads to a safety result of the form:*

$$\mathcal{C}(\varphi(t)) \leq R(\bar{\delta}(\epsilon(\mathcal{C}(\varphi(t)))))) \approx 0, \quad (1.41)$$

without leading to $\mathcal{C}(\varphi(t)) \ll 0$ for all $t \in \mathbb{R}_{\geq 0}$.

This result states that Tunable Input-to-State Safety allows for tuning controllers to achieve strong theoretical safety properties (the right-hand side of the preceding equation is approximately 0) while not requiring that the system display excessive degrees of conservativeness in maintaining safety. This makes Tunable Input-to-State Safety appealing for designing controllers that enforce partial robustness as it also allows for strong performance.

Section 4.5: Stochastic Disturbances

In Section 4.5 I will consider safety of discrete-time systems of the form:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k, \quad k \in \mathbb{Z}_{\geq 0}, \quad (1.42)$$

where $\mathbf{F} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and \mathbf{d}_k is a *stochastic* disturbance signal that is distributed according to some distribution \mathcal{D} , such that $\mathbf{d}_k \sim \mathcal{D}$. Given a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, one may consider a closed-loop system:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)) + \mathbf{d}_k, \quad k \in \mathbb{Z}_{\geq 0}, \quad (1.43)$$

which is assumed to have a solution $\varphi : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}^n$ for any initial condition $\mathbf{x}_0 \in \mathbb{R}^n$ and sufficiently well-behaved sequence of random disturbances. The potential for unbounded disturbance signals (such as disturbances drawn from a Gaussian distribution) will motivate the need to consider a notion of *finite-time* safety, in which an inequality on $\mathbf{C}(\varphi(k))$ will only need to be met for some finite values of k with some probability. Pursuing this will lead to a result of the following form:

Proposition 6 (Informal). *If a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ leads to a closed-loop system satisfying:*

$$\mathbb{E}[\mathbf{C}(\mathbf{x}_{k+1})] - \mathbf{C}(\mathbf{x}_k) \leq -\alpha \mathbf{C}(\mathbf{x}_k) + \bar{\delta}(\mathcal{D}), \quad k \in \mathbb{Z}_{\geq 0}, \quad (1.44)$$

then we have that for any $K \in \mathbb{Z}_{\geq 0}$:

$$\mathbb{P}[\mathbf{C}(\varphi(k)) \leq R(\bar{\delta}(\mathcal{D})) \text{ for all } k \leq K] \geq 1 - \epsilon(\mathcal{D}, K). \quad (1.45)$$

This result states that if a bound on the expected increment of the value of \mathbf{C} can be achieved, the probability of meeting a given safety requirement ($R(\bar{\delta}(\mathcal{D}))$) for a finite number (K) of time-steps can be lower-bounded. This is a form of partial robustness, as it allows the value of \mathbf{C} to increase above 0, and only holds over a finite-time horizon. This work will further explore how such a controller \mathbf{k} can be designed using convex optimization, and what additional modifications must occur to enforce the desired expectation condition.

Chapter 5: Measurement Error

In Chapter 5 I will explore the problem of *measurement error*. In Sections 5.2-5.4 I will consider a perturbation of the form:

$$\Delta(\mathbf{x}, \mathbf{u}) = \mathbf{f}(\mathbf{x} + \mathbf{e}(\mathbf{x})) - \mathbf{f}(\mathbf{x}) + (\mathbf{g}(\mathbf{x} + \mathbf{e}(\mathbf{x})) - \mathbf{g}(\mathbf{x}))\mathbf{u}, \quad (1.46)$$

where $\mathbf{e} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an unknown function capturing error in the estimate of the state used by a controller. This perturbation can be seen as replacing the

dynamics of the system with the dynamics of the system at the state $\mathbf{x} + \mathbf{e}(\mathbf{x})$, but the controller only has access to the state \mathbf{x} .

Sections 5.2 and 5.3: Measurement-Robust CBFs

The focus in Sections 5.2 and 5.3 is on designing safety-critical controllers in the presence of measurement error. To accomplish this, this setting will be based on access to a state-based measurement $\mathbf{y} = \mathbf{p}(\mathbf{x}) \in \mathbb{R}^p$. This measurement may be low-dimensional and directly include parts of the state, or it may be high-dimensional, such as an image, and need to be converted into a low-dimensional representation of the system state. Consequently, it will be assumed that given a measurement \mathbf{y} , an estimate of the state \mathbf{x} can be produced (though the state dynamics still use $\mathbf{x} + \mathbf{e}(\mathbf{x})$). It will be assumed that accompanying the measurement \mathbf{y} is an error set $\mathcal{E}(\mathbf{y})$, such that it is known that:

$$\mathbf{e}(\mathbf{x}) \in \mathcal{E}(\mathbf{y}). \quad (1.47)$$

This assumption states that given a measurement \mathbf{y} , we can form an estimate of the state \mathbf{x} and know that the error $\mathbf{e}(\mathbf{x})$ is restricted to the set $\mathcal{E}(\mathbf{y})$. It will additionally be assumed that the error set satisfies the following property:

$$\sup_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} \|\mathbf{e}\| \leq \epsilon(\mathbf{y}), \quad (1.48)$$

for a known function $\epsilon : \mathbb{R}^p \rightarrow \mathbb{R}_{\geq 0}$, such that it is known that $\|\mathbf{e}(\mathbf{x})\| \leq \epsilon(\mathbf{y})$. This construction will lead to the following result:

Proposition 7 (Informal). *There exists minimum values of the parameters $a, b \in \mathbb{R}_{\geq 0}$ such that for all larger values, the convex second-order cone program-based controller:*

$$\begin{aligned} \mathbf{k}(\mathbf{x}, \mathbf{y}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \|\mathbf{u} - \mathbf{k}_{nom}(\mathbf{x})\|^2 \\ \text{s.t. } & \hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + (b + a\|\mathbf{u}\|)\epsilon(\mathbf{y}) \leq -\alpha\mathbf{C}(\mathbf{x}), \end{aligned} \quad (1.49)$$

enforces:

$$\mathbf{C}(\varphi(t)) \leq 0, \quad t \in \mathbb{R}_{\geq 0}. \quad (1.50)$$

This results states that by incorporating the bound $\epsilon(\mathbf{y})$ on the measurement error into the controller design, it is possible to meet the safety requirement without needing to access $\mathbf{e}(\mathbf{x})$. This is a form of absolute robustness, as the

original safety requirement is met despite the perturbation. Furthermore, this controller is parameterized (in terms of a and b), and there is minimal values of these parameters beyond which robust guarantees are achieved.

Section 5.4: Preference-Based Learning

The disturbance and measurement-robust controllers explored in Chapters 4 and 5 are often defined using a collection of robustness parameters (ϵ for disturbance robustness and a, b for measurement robustness). As multiple forms of robustness are built into a controller, it often leads to excessively conservative behavior as the system is robust to more than it needs to be at any given point in time. Tuning the robustness parameters is often a tedious task, and control designers often have a limited understanding of how the parameters interact to balance performance and robustness. Motivated by this challenge, Section 5.4 explores how feedback on performance and robust safety from a human control designer can be used to automate the process of tuning control gains using *preference-based learning*. In particular, I will consider controllers of the form:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ &\text{s.t. } \widehat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \phi(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}) \leq -\alpha \mathbf{C}(\mathbf{x}), \end{aligned} \quad (1.51)$$

with a robustness function $\phi : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p \rightarrow \mathbb{R}_{\geq 0}$ capturing multiple forms of robustness (disturbance and measurement error) and robustness parameters $\boldsymbol{\theta} \in \mathbb{R}^p$. The robustness function ϕ will be constructed using the formulations for disturbance and measurement robustness to ensure this controller is defined by a convex optimization problem. This controller construction will lead to the following observation:

Proposition 8 (Informal). *Utilizing human feedback and preferences about closed-loop performance and safety, preference-based learning can be successfully used to iteratively tune the parameters $\boldsymbol{\theta}$ to achieve safety of the form:*

$$\mathbf{C}(\boldsymbol{\varphi}(t)) \leq R(\bar{\delta}(\boldsymbol{\theta})) \approx 0, \quad (1.52)$$

without requiring that $\mathbf{C}(\boldsymbol{\varphi}(t)) \ll 0$.

This result states that it is possible to automate the process of tuning robustness parameters using feedback about a control designer’s qualitative observations about performance and safety properties to achieve strong safety

results without requiring that the system display excessive degrees of conservativeness. This result, while more algorithmic and experimental in nature, is promising as it may be deployed with many of the robust control approaches in this thesis which may be parameterized.

Chapter 6: Input Sampling

In Chapter 6 I will explore the problem of *input sampling*. In Sections 6.3-6.5 I will consider a perturbation of the form:

$$\Delta(\mathbf{x}, \mathbf{u}, t) = -\mathbf{g}(\mathbf{x})\mathbf{u} + \mathbf{g}(\mathbf{x})\mathbf{u}(t_k), \quad \forall t \in [t_k, t_{k+1}), \quad (1.53)$$

where $t_k, t_{k+1} \in \mathbb{R}_{\geq 0}$ are sequential *sample times* with $t_k < t_{k+1}$, and $\mathbf{u}(t_k)$ is the input applied to the system at time t_k . I note that this type of perturbation requires more care in modeling in this continuous-time formulation as multiple points in time must be considered, but the approach taken in this thesis will actually not begin its investigation by considering input-sampling as a perturbation to a continuous-time system.

Sections 6.3 and 6.4: Approximate Discrete-Time Model Design

Instead of viewing input sampling from a continuous-time perspective, I will view it from a discrete-time perspective. Given a sample period $T \in \mathbb{R}_{> 0}$ separating sample times ($t_{k+1} - t_k = T$), the evolution of a sample-and-hold system can be described by:

$$\mathbf{x}_{k+1} = \mathbf{F}_T^e(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{x}_k + \int_{t_k}^{t_{k+1}} \mathbf{f}(\varphi(\tau)) + \mathbf{g}(\varphi(\tau))\mathbf{u}_k \, d\tau, \quad (1.54)$$

where \mathbf{u}_k is the input taken at sample time t_k , and φ is the continuous-time solution of the system under the sample-and-hold input. The function \mathbf{F}_T^e is known as the *exact* discrete-time map, and is generally difficult to compute due to requiring an integral of general nonlinear functions. Motivated by the work in [81], I will consider design using an approximation $\mathbf{F}_T^a : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$:

$$\mathbf{F}_T^a(\mathbf{x}_k, \mathbf{u}_k) \approx \mathbf{F}_T^e(\mathbf{x}_k, \mathbf{u}_k). \quad (1.55)$$

In order to achieve theoretical stability and safety properties, the approximation will need to satisfy a *one-step consistency* property defined as:

$$\|\mathbf{F}_T^e(\mathbf{x}_k, \mathbf{u}_k) - \mathbf{F}_T^a(\mathbf{x}_k, \mathbf{u}_k)\| \leq T\rho(T), \quad (1.56)$$

for sufficiently small sample periods T , where $\rho : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is monotonically increasing and satisfies $\rho(0) = 0$. This property dictates the accuracy required

by the approximation, stating that it must more accurately represent the exact discrete-time map for smaller sample periods. I will demonstrate that standard approximation schemes such as Euler approximations and Runge-Kutta approximations satisfy this property, and thus can be used in control design while preserving stability and safety properties. In particular, I will develop results of the form:

Proposition 9 (Informal). *If the approximate discrete-time map \mathbf{F}_T^a is one-step consistent with the exact discrete-time map \mathbf{F}_T^e , then for controllers of the form:*

$$\begin{aligned} \mathbf{k}_T(\mathbf{x}_k) &= \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} \|\mathbf{u} - \mathbf{k}_{nom}(\mathbf{x}_k)\|^2 \\ \text{s.t. } \mathbf{C}(\mathbf{F}_T^a(\mathbf{x}_k, \mathbf{u})) - \mathbf{C}(\mathbf{x}_k) &\leq -T\alpha\mathbf{C}(\mathbf{x}_k), \end{aligned} \quad (1.57)$$

we have stability properties of the form:

$$\lim_{k \rightarrow \infty} |\mathbf{C}(\varphi(t_k))| \leq R(T), \quad (1.58)$$

and safety properties of the form:

$$\mathbf{C}(\varphi(t_k)) \leq R(T), \quad (1.59)$$

for all $k \in \mathbb{Z}_{\geq 0}$.

There are a number of interesting observations about this result. First observe that the controller considers a finite difference of the value of \mathbf{C} across one sample period using an approximate discrete-time map, and the inequality with the difference scales with the sample period. This differs from the continuous-time formulations used in the previous chapters. Next, the stability and safety properties are of a partial robustness form, as the violation of stability and safety scales with the sample period, and can be made arbitrarily small with sufficiently small sample periods. Lastly, the stability and safety guarantees hold at the sample-times, and do not make statements about inter-sample behavior (this can be done with additional analysis of this framework if so desired). The unique variations in these controllers and properties arise because rather than trying to view input sampling as a perturbation in continuous-time dynamics, a completely different paradigm of design is taken using approximations of discrete-time systems. Lastly, a key focus in this work will be in

ensuring that the controller in the preceding proposition is specified by a convex optimization problem, and I will observe that different classes of convex optimization problems arise due to the class of approximation and geometry of stability and safety constraints, which is not seen in continuous-time.

Section 6.5: Event-Triggered Safety

In Section 6.5 I will explore a different paradigm of input-sampling known as event-triggered control. In this setting, input-sampling is actually introduced to a continuous-time system (it is assumed the controller could be implemented continuously in time) in an effort to conserve actuation resources. In particular, this problem setting is concerned with minimizing costs that happen due to *switching* the input to the system.

The goal in event-triggered control will be to determine a *trigger-law* that dictates how sampling times are specified:

$$t_{k+1} = \min\{t \geq t_k \mid \rho(\|\mathbf{x} - \mathbf{x}(t_k)\|) = 0\}, \quad (1.60)$$

where $\rho : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is defined to ensure that:

$$\dot{\mathcal{C}}(\mathbf{x}(t), \mathbf{u}_k) \leq -c\alpha\mathcal{C}(\mathbf{x}(t)), \quad (1.61)$$

where $\mathbf{x}(t)$ is the continuous-time evolution of the state at a time $t \geq t_k$ using the constant input \mathbf{u}_k , and $c \in (0, 1)$. This states that a continuous-time controller implementation that meets the original inequality on $\dot{\mathcal{C}}$ is implemented with a sample-and-hold implementation that is updated to meet a relaxed inequality. In this sense, inherent robustness in the continuous-time controller is being traded for more efficient use of actuation resources.

The burden of proof in this setting is to ensure the existence of a minimum interevent time (MIET), such that the trigger law leads to the existence of some $\tau \in \mathbb{R}_{>0}$ such that $t_{k+1} \geq t_k + \tau$. This is not only important for ensuring there is some efficiency benefit to an event-triggered implementation, but it also is important for eliminating undesirable Zeno-like behavior, where infinitely many events occur in a finite amount of time. Exploring this challenge will lead to the following result:

Proposition 10 (Informal). *There exists a function $\rho : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ such that the corresponding trigger-law leads to a minimum interevent time and a safety guarantee of the form:*

$$\mathcal{C}(\varphi(t)) \leq \epsilon, \quad (1.62)$$

where $\epsilon \in \mathbb{R}_{>0}$ may be taken arbitrarily small at the expense of a smaller minimum interevent time.

This result states that it is possible to use event-triggered implementations of CBF-based controllers and satisfy a partial robustness property that can be tuned to be an arbitrarily small degradation of the original safety property. The cost for making this degradation smaller is a potentially smaller minimum interevent time, which may degrade resource efficiency. The key challenge that is addressed in this result is regulating the behavior of the system when it is on the boundary between safe and unsafe states ($\mathcal{C}(\mathbf{x}) = 0$, where motion along the boundary can lead to a failure to have an MIET).

BACKGROUND

In this chapter I will present various background materials in nonlinear control systems that underlie my contributions presented later in this thesis. I will begin by reviewing nonlinear dynamics and the properties of stability and safety that are often an objective of control design. I next review Control Lyapunov Functions (CLFs) as a tool for stabilizing control synthesis, followed by feedback linearization as a constructive tool for synthesizing both stabilizing controllers and CLFs. I then proceed to review Control Barrier Functions (CBFs) as a tool for safe control synthesis. Following this, I review nonlinear dynamics with disturbances, the robustness properties of Input-to-State Stability (ISS) and Input-to-State Safety (ISSf), and corresponding variants of CLFs and CBFs for achieving these properties.

2.1 Nonlinear Dynamics

Let $E \subseteq \mathbb{R}^n$ be an open set and consider a nonlinear control-affine system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (2.1)$$

with state $\mathbf{x} \in E$, input $\mathbf{u} \in \mathbb{R}^m$, and functions $\mathbf{f} : E \rightarrow \mathbb{R}^n$ and $\mathbf{g} : E \rightarrow \mathbb{R}^{n \times m}$ assumed to be locally Lipschitz continuous on E . This system is referred to as an *open-loop* system, as the input \mathbf{u} is unspecified. Let $\mathbf{k} : E \rightarrow \mathbb{R}^m$ be a controller that is locally Lipschitz continuous on E . This controller yields the following *closed-loop* system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}). \quad (2.2)$$

As the functions \mathbf{f} , \mathbf{g} , and \mathbf{k} are locally Lipschitz continuous on E , for any initial condition $\mathbf{x}_0 \in E$ there exists a maximal time interval $I(\mathbf{x}_0) = [0, t_{\max}(\mathbf{x}_0))$ and a unique continuously differentiable solution $\varphi : I(\mathbf{x}_0) \rightarrow E$ satisfying:

$$\dot{\varphi}(t) = \mathbf{f}(\varphi(t)) + \mathbf{g}(\varphi(t))\mathbf{k}(\varphi(t)), \quad (2.3)$$

$$\varphi(0) = \mathbf{x}_0, \quad (2.4)$$

for all $t \in I(\mathbf{x}_0)$ [100], [101].

A specific set of nonlinear dynamics that I will consider in various contributions in this thesis are *robotic system dynamics*. The state of a robotic system will be described by a set of n_c *generalized coordinates* $\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^{n_c}$, where \mathcal{Q} is referred to as the *configuration space*, and a set of n_c *generalized coordinate rates* $\dot{\mathbf{q}} \in \mathbb{R}^{n_c}$, yielding the state $\mathbf{x} \triangleq \begin{bmatrix} \mathbf{q}^\top & \dot{\mathbf{q}}^\top \end{bmatrix}^\top \in \mathcal{Q} \times \mathbb{R}^{n_c} \triangleq \mathcal{X} \subseteq \mathbb{R}^n$. The dynamics of a robotic system are governed by the *Euler-Lagrange* equations:

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) = \mathbf{B}(\mathbf{q})\boldsymbol{\tau}, \quad (2.5)$$

with positive definite inertia matrix $\mathbf{D} : \mathcal{Q} \rightarrow \mathbb{S}_{>0}^{n_c}$, Coriolis matrix $\mathbf{C} : \mathcal{X} \rightarrow \mathbb{R}^{n_c}$, gravitational matrix $\mathbf{G} : \mathcal{Q} \rightarrow \mathbb{R}^{n_c}$, actuation matrix $\mathbf{B} : \mathcal{Q} \rightarrow \mathbb{R}^{n_c \times m}$, and input torques $\boldsymbol{\tau} \in \mathbb{R}^m$. The functions \mathbf{D} , \mathbf{G} , and \mathbf{B} are assumed to be locally Lipschitz continuous on \mathcal{Q} and the function \mathbf{C} is assumed to be locally Lipschitz continuous on \mathcal{X} . These dynamics can be written in the form of (2.1) as follows:

$$\underbrace{\frac{d}{dt} \begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} \dot{\mathbf{q}} \\ -\mathbf{D}(\mathbf{q})^{-1}(\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q})) \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} \mathbf{0}_{n_c} \\ \mathbf{D}(\mathbf{q})^{-1}\mathbf{B}(\mathbf{q}) \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} \underbrace{\boldsymbol{\tau}}_{\mathbf{u}}, \quad (2.6)$$

noting that the functions \mathbf{f} and \mathbf{g} are locally Lipschitz continuous on \mathcal{X} as the matrix inverse is a locally Lipschitz continuous function [102, Theorem 9.8]. Lastly, a robotic system is said to be *fully-actuated* if $m = n_c$ and the matrix $\mathbf{B}(\mathbf{q})$ is full-rank for each $\mathbf{q} \in \mathcal{Q}$.

2.2 Stability & Safety

The concept of stability has long been a focus in studying dynamic systems [103], and is often an important objective in control design. Stability properties of the closed-loop system (2.2) are defined with respect to an equilibrium point, which is captured in the following definition:

Definition 1 (*Equilibrium Point*). A point $\mathbf{x}_e \in E$ is said to be an *unforced equilibrium point* of the open-loop system (2.1) if:

$$\mathbf{f}(\mathbf{x}_e) = \mathbf{0}_n. \quad (2.7)$$

A point $\mathbf{x}_e \in E$ is said to be a *forced equilibrium point* of the open-loop system (2.1) if there exists a $\mathbf{u}_e \in \mathbb{R}^m$ such that:

$$\mathbf{f}(\mathbf{x}_e) + \mathbf{g}(\mathbf{x}_e)\mathbf{u}_e = \mathbf{0}_n. \quad (2.8)$$

A point $\mathbf{x}_e \in E$ is said to be an *equilibrium point* of the closed-loop system (2.2) if:

$$\mathbf{f}(\mathbf{x}_e) + \mathbf{g}(\mathbf{x}_e)\mathbf{k}(\mathbf{x}_e) = \mathbf{0}_n. \quad (2.9)$$

This enables the following definition of exponential stability [100]:

Definition 2 (*Exponential Stability*). Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2). The closed-loop system (2.2) is said to be *locally exponentially stable with respect to \mathbf{x}_e* if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ and $M, \lambda \in \mathbb{R}_{>0}$ such that:

$$\mathbf{x}_0 \in D \implies \|\varphi(t) - \mathbf{x}_e\| \leq M\|\mathbf{x}_0 - \mathbf{x}_e\|e^{-\lambda t} \quad (2.10)$$

for all $t \in \mathbb{R}_{\geq 0}$.

Note the requirement for local exponential stability that the unique continuously differentiable solution φ exists for all time. If it is possible to find an open set $D' \subseteq E$ with $\mathbf{x}_e \in D'$ such that:

$$\mathbf{x}_0 \in D' \implies \|\varphi(t) - \mathbf{x}_e\| \leq M\|\mathbf{x}_0 - \mathbf{x}_e\|e^{-\lambda t}. \quad (2.11)$$

for all $t \in I(\mathbf{x}_0)$, then an open set $D \subseteq D'$ with $\mathbf{x}_e \in D$ and a compact set K with $D \subset K \subset D'$ can be found such that for any initial condition $\mathbf{x}_0 \in D$, $\varphi(t) \in K$ for all $t \in I(\mathbf{x}_0)$, and thus $I(\mathbf{x}_0) = \mathbb{R}_{\geq 0}$. [101, Chapter 2.4, Corollary 2]. This is a useful requirement in the definition of stability as it establishes a well-posedness when discussing the evolution of the system (2.2) for all time.

The notion of safety I will consider for nonlinear dynamic systems is codified through the concept of forward invariance [104]:

Definition 3 (*Forward Invariance & Safety*). A set $\mathcal{C} \subseteq E$ is said to be *forward invariant* for the closed-loop system (2.2) if for any initial condition $\mathbf{x}_0 \in \mathcal{C}$, we have that $\varphi(t) \in \mathcal{C}$ for all $t \in I(\mathbf{x}_0)$. If a set \mathcal{C} is forward for the closed-loop system (2.2), then the closed-loop system (2.2) is said to be *safe with respect to \mathcal{C}* .

If a closed-loop system (2.2) is safe with respect to a compact set \mathcal{C} , then $I(\mathbf{x}_0) = \mathbb{R}_{\geq 0}$ for any initial condition $\mathbf{x}_0 \in \mathcal{C}$ [101, Chapter 2.4, Corollary 2], and thus $\varphi(t) \in \mathcal{C}$ for all $t \in \mathbb{R}_{\geq 0}$.

2.3 Control Lyapunov Functions

Lyapunov theory is a fundamental tool for both certifying stability properties of nonlinear systems and constructively synthesizing stabilizing controllers. I will begin by reviewing exponential Lyapunov functions and their connection with exponential stability [100]:

Definition 4 (*Exponential Lyapunov Function*). Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2). A function $V : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E is said to be a *local exponential Lyapunov function* for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ and $k_1, k_2, k_3, a \in \mathbb{R}_{>0}$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.12)$$

$$\dot{V}(\mathbf{x}) \triangleq \underbrace{\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})\mathbf{f}(\mathbf{x})}_{L_{\mathbf{f}}V(\mathbf{x})} + \underbrace{\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})\mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x})}_{L_{\mathbf{g}}V(\mathbf{x})} \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.13)$$

for all $\mathbf{x} \in D$.

Exponential Lyapunov functions enable the certification of exponential stability, as codified in the following theorem [100, Theorem 4.10]:

Theorem 1. *Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2). If there exists a local exponential Lyapunov function $V : E \rightarrow \mathbb{R}_{\geq 0}$ for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e , then the closed-loop system (2.2) is locally exponentially stable with respect to \mathbf{x}_e .*

A converse result to this preceding theorem establishes the existence of a local exponential Lyapunov function for a system that is locally exponentially stable under mild regularity assumptions [100, Theorem 4.14]:

Theorem 2. *Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2), and suppose that the functions \mathbf{f} , \mathbf{g} , and \mathbf{k} are continuously differentiable on E . If the closed-loop system (2.2) is locally exponentially stable with respect to \mathbf{x}_e , then there exists a local exponential Lyapunov function $V : E \rightarrow \mathbb{R}_{\geq 0}$ for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e with $a = 2$ and corresponding open set $D \subseteq E$. Furthermore, there exists a $k_4 \in \mathbb{R}_{>0}$ such that:*

$$\left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\| \leq k_4 \|\mathbf{x} - \mathbf{x}_e\|, \quad (2.14)$$

for all $\mathbf{x} \in D$.

A result arising from the existence of a local exponential Lyapunov function pertains to the existence and uniqueness of solutions when a controller is not locally Lipschitz continuous at the equilibrium point \mathbf{x}_e [105, Theorem 1.2]:

Theorem 3. *Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2), and suppose that the function $\mathbf{k} : E \rightarrow \mathbb{R}^m$ is continuous on E , but only locally Lipschitz continuous on $E \setminus \{\mathbf{x}_e\}$. If there exists a local exponential Lyapunov function $V : E \rightarrow \mathbb{R}_{\geq 0}$ for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e , then for the initial condition $\mathbf{x}_0 = \mathbf{x}_e$, there exists a unique continuously differentiable solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow E$ to the closed-loop system (2.2) given by $\varphi(t) = \mathbf{x}_e$ for all $t \in \mathbb{R}_{\geq 0}$.*

This result states that if a controller is not locally Lipschitz continuous at the equilibrium point \mathbf{x}_e , but a local exponential Lyapunov function exists for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e , then an issue with non-uniqueness of solutions for the closed-loop system (2.2) does not arise from the initial condition \mathbf{x}_e . This precludes the pedagogical closed-loop system:

$$\dot{x} = f(x) = x^{\frac{2}{3}}, \quad (2.15)$$

which is not locally Lipschitz continuous at the equilibrium point $x_e = 0$, and has two solutions from the initial condition $x_0 = 0$ given by $\varphi_1(t) = 0$ and $\varphi_2(t) = \left(\frac{t}{3}\right)^3$ for all $t \in \mathbb{R}_{\geq 0}$. It is straightforward to see this system is not locally exponentially stable (it is unstable), and thus we could not find a local exponential Lyapunov function V that would certify that $\varphi_1(t) = 0$ for all $t \in \mathbb{R}_{\geq 0}$ is the unique solution from $x_0 = 0$. This further shows how stability and Lyapunov functions play a role in ensuring the existence and uniqueness of solutions to the differential equations describing the system dynamics.

A powerful tool for jointly synthesizing both an exponentially stabilizing controller and a corresponding local exponential Lyapunov function as a certificate of that exponential stability is the notion of a local exponential Control Lyapunov Function (CLF). CLFs were first introduced in [8] and were subsequently used to constructively synthesize stabilizing controllers in [9] and [106]. Their use in synthesizing exponentially stabilizing controllers through convex optimization was first considered in [26], [107]. I will explore how CLFs can be synthesized using feedback linearization in Section 2.4, but other recent work has synthesized CLFs using proportional-derivative (PD) control [108].

Definition 5 (*Control Lyapunov Function (CLF)*). Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (2.1). A function $V : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E is said to be a *local exponential Control Lyapunov Function (CLF)* for the open-loop system (2.1) and equilibrium point \mathbf{x}_e if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$, and $k_1, k_2, k_3, a \in \mathbb{R}_{>0}$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.16)$$

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}(\mathbf{x}', \mathbf{u}) \triangleq L_{\mathbf{f}}V(\mathbf{x}') + L_{\mathbf{g}}V(\mathbf{x}')\mathbf{u} < -k_3 \|\mathbf{x}' - \mathbf{x}_e\|^a, \quad (2.17)$$

for all $\mathbf{x} \in D$ and all $\mathbf{x}' \in D \setminus \{\mathbf{x}_e\}$.

There are a few notable observations to make about the definition of a local exponential CLF. First, I note that a local exponential CLF is constructed with respect to the open-loop system (2.1), rather than the closed-loop system (2.2). Because of this, local exponential CLFs serve as a statement about how a dynamic system *could* be stabilized. Defining stability properties still requires specifying a controller \mathbf{k} and producing a closed-loop system. This distinction will be important in Chapter 3 when understanding the well-posedness of addressing model error using local exponential CLFs.

Second, I note the strict inequality used in (2.17). This strict inequality is consistent with the original definition of CLFs in [8, Theorem 4.1], and the more recent definition in [51, Definition 1]. The importance of this strict inequality is two-fold. First, the infimum in (2.17) may not be attained, and thus if the inequality was non-strict, it may be impossible to produce a controller that meets the exponential Lyapunov condition in (2.13) (for a particular value of k_3). Second, the strictness of this inequality is important when establishing local Lipschitz continuity of controllers synthesized via local exponential CLFs, as seen in [51, Theorem 1]. Furthermore, this strict inequality contrasts with the non-strict inequality appearing in (2.13). Satisfying the non-strict inequality is sufficient for ensuring stability of a closed-loop system, while the importance of the strict inequality is in synthesizing controllers to go from an open-loop system to a closed-loop system that satisfies the non-strict inequality. Lastly, the positive definiteness of a local exponential CLF as enforced by (2.16) requires that the gradient of the local exponential CLF with respect to the state vanishes at the equilibrium point \mathbf{x}_e , implying that $\dot{V}(\mathbf{x}_e, \mathbf{u}) = 0$ for all $\mathbf{u} \in \mathbb{R}^m$. Thus, the strict inequality in (2.13) cannot be enforced at this

point, and thus it is omitted from consideration. This omission is acceptable because any controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ satisfies $\dot{V}(\mathbf{x}_e, \mathbf{k}(\mathbf{x}_e)) = 0$, meeting the condition (2.13) for stability at the point \mathbf{x}_e .

Third, I note that the infimum condition in (2.17) can be equivalently rewritten as the following implication:

$$L_{\mathbf{g}}V(\mathbf{x}') = \mathbf{0}_m \implies L_{\mathbf{f}}V(\mathbf{x}') < -k_3\|\mathbf{x}' - \mathbf{x}_e\|^a, \quad (2.18)$$

for all $\mathbf{x}' \in D \setminus \{\mathbf{x}_e\}$. This alternative perspective provides a geometrical interpretation of the local exponential CLF condition (2.17). In the set of states which the time derivative of the exponential CLF can not be manipulated through the choice of control input (\mathbf{x}' satisfies the precedent), the unactuated dynamics of the system (given by \mathbf{f}) must display the necessary stability properties (the antecedent). This interpretation will be useful later in Section 3.8 for proving feasibility results for optimization-based controllers synthesized via local exponential CLFs in the presence of model error.

The utility of local exponential CLFs is in their characterization of exponentially stabilizing inputs and subsequent role as local exponential Lyapunov functions. Given a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e , define the following pointwise set:

$$K_{\text{CLF}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \dot{V}(\mathbf{x}, \mathbf{u}) \leq -k_3\|\mathbf{x} - \mathbf{x}_e\|^a \right\}. \quad (2.19)$$

This set leads to the following result [62], [109]:

Theorem 4. *Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (2.1), and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$. Then, the set $K_{\text{CLF}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in D$, and for any controller¹ $\mathbf{k} : D \rightarrow \mathbb{R}^m$ that is continuous on D , locally Lipschitz continuous on $D \setminus \{\mathbf{x}_e\}$, renders \mathbf{x}_e an equilibrium point of the closed-loop system (2.2), and satisfies $\mathbf{k}(\mathbf{x}) \in K_{\text{CLF}}(\mathbf{x})$ for all $\mathbf{x} \in D$, the function V is a local exponential Lyapunov function for the closed-loop system (2.2) and equilibrium point \mathbf{x}_e .*

¹The feedback linearization and optimization-based controllers I will work with in this thesis will often not be naturally defined outside of an open set $D \subset E$, nor often can they be extended beyond D while meeting regularity properties on all of E . When working with these controllers it is natural to restrict the notion of the solutions to the closed-loop system (2.2) to initial conditions $\mathbf{x}_0 \in D$ and solutions $\varphi : I(\mathbf{x}_0) \rightarrow D$. Note that D is not the set typically referred to as the “region of attraction” but rather contains the region of attraction.

This result characterizes requirements on the controller \mathbf{k} for which the exponential stability of (2.2) with respect to the equilibrium point \mathbf{x}_e can be certified via the local exponential CLF V serving as a local exponential Lyapunov function. Producing such a controller meeting the required regularity properties was first considered in [9], wherein the notion of the small control property was proposed for unforced equilibrium points, and generalized in [106, Definition 4.6] to forced equilibrium points in the form of the continuous control property:

Definition 6 (*Small Control Property (SCP)*). Let $\mathbf{x}_e \in E$ be an unforced equilibrium point of the open-loop system (2.1), and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e . The local exponential CLF V is said to satisfy the *small control property (SCP)* if for any $\varepsilon \in \mathbb{R}_{>0}$, there exists an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ such that for any $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, there exists a $\mathbf{u} \in \mathbb{R}^m$ satisfying $\|\mathbf{u}\| < \varepsilon$ and:

$$\dot{V}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \quad (2.20)$$

Definition 7 (*Continuous Control Property (CCP)*). Let $\mathbf{x}_e \in E$ be a forced equilibrium point of the open-loop system (2.1) for equilibrium input \mathbf{u}_e , and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e . The local exponential CLF V is said to satisfy the *continuous control property (CCP)* if for any $\varepsilon \in \mathbb{R}_{>0}$, there exists an open set $D \subseteq \mathbb{R}^n$ with $\mathbf{x}_e \in D$ such that for any $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, there exists a $\mathbf{u} \in \mathbb{R}^m$ satisfying $\|\mathbf{u} - \mathbf{u}_e\| < \varepsilon$ and:

$$\dot{V}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \quad (2.21)$$

These properties provide a statement about the behavior of a local exponential CLF V as it approaches the equilibrium point \mathbf{x}_e . To highlight the importance of the SCP in achieving the necessary regularity properties on a local exponential CLF-based controller, consider the following simple example:

Example 1. Consider the simple scalar open-loop system:

$$\dot{x} = x + x^2 u. \quad (2.22)$$

The origin of this system is an unforced equilibrium point ($x_e = 0$), and the quadratic function $V(x) = \frac{1}{2}x^2$ is a local exponential CLF for the open-loop

system (2.22) and equilibrium point x_e , as we have that:

$$\inf_{u \in \mathbb{R}} \dot{V}(x, u) = \inf_{u \in \mathbb{R}} x^2 + x^3 u = -\infty \quad (2.23)$$

for all $x \in \mathbb{R} \setminus \{0\}$. The exponential CLF V does not satisfy the SCP. To see this, consider an $x \in \mathbb{R}_{>0}$, and observe that the SCP condition (2.20) requires that there exists a $u \in \mathbb{R}^m$ satisfying:

$$x^2 + x^3 u < -k_3 x^a \iff u < -k^3 x^{a-3} - x^{-1}. \quad (2.24)$$

As x approaches 0 from the right, the input u must become unbounded in the negative direction, and thus can not be made to satisfy $|u| < \varepsilon$ for arbitrarily small $\varepsilon \in \mathbb{R}_{>0}$ in some open set containing 0. Intuitively, we observe that in the absence of input, the system (2.22) is an unstable linear system. Furthermore, its input is multiplied by a quadratic term of the state. Thus, as x approaches the origin, the magnitude of the input must become unbounded to compensate for the faster decay of the quadratic term compared to the unstable linear term.

As I will discuss when exploring feedback linearization in the next section, establishing that a local exponential CLF meets the small or continuous control property is often most easily done by constructing the local exponential CLF through a controller \mathbf{k} that is continuous at \mathbf{x}_e and for which \mathbf{x}_e is an equilibrium point of the closed-loop system (2.2).

Let $\mathbf{x}_e \in E$ be an unforced ($\mathbf{u}_e = \mathbf{0}_m$) or forced equilibrium point with equilibrium input \mathbf{u}_e for the open-loop system (2.1), and let V be a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$. Consider the following optimization-based controller $\mathbf{k}_{\text{CLF}} : D \rightarrow \mathbb{R}^m$, first introduced in [26], which is synthesized directly using the local exponential CLF V :

$$\begin{aligned} \mathbf{k}_{\text{CLF}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 && (\text{CLF-QP}) \\ \text{s.t. } \dot{V}(\mathbf{x}, \mathbf{u}) &\leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \end{aligned}$$

I make the following observations regarding this controller. First, for $\mathbf{x} = \mathbf{x}_e$, the constraint of the optimization problem is trivially satisfied for any value of \mathbf{u} , and thus the input that minimizes the cost function is $\mathbf{u} = \mathbf{u}_e$, which renders \mathbf{x}_e an equilibrium point of the closed-loop system (2.2). Second, the cost function of the optimization problem defining the constraint is convex

with respect to the input, and the function \dot{V} is affine in the input, making this optimization problem a *convex quadratic program* (QP). The minimizing solutions of these problems are well understood (there is a unique solution to this particular optimization problem as the cost is *strictly* convex with respect to the input), and they may be efficiently solved numerically [29], [110], permitting their use on real-world systems such as bipedal walking platforms [24], [107]. This optimization-based control framework will serve as a cornerstone in the contributions of this thesis. Lastly, by construction, the resulting closed-loop system (2.2) satisfies the local exponential Lyapunov function condition (2.13), such that the closed-loop system is locally exponentially stable by design. This is codified in the following result [51, Theorem 1]:

Theorem 5. *Let $\mathbf{x}_e \in E$ be an unforced (forced) equilibrium point of the open-loop system (2.1) with equilibrium input \mathbf{u}_e , let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$, let the function $\frac{\partial V}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ be locally Lipschitz continuous on D , and assume that V satisfies the small (continuous) control property. Then the controller $\mathbf{k}_{CLF} : D \rightarrow \mathbb{R}^m$ is continuous on D , locally Lipschitz continuous on $D \setminus \{\mathbf{x}_e\}$, renders \mathbf{x}_e an equilibrium point of the closed-loop system (2.2), and satisfies $\mathbf{k}_{CLF}(\mathbf{x}) \in K_{CLF}(\mathbf{x})$ for all $\mathbf{x} \in D$.*

This result states that the CLF-QP controller meets the conditions of Theorem 4, and thus locally exponentially stabilizes the closed-loop system (2.2) with respect to the equilibrium point \mathbf{x}_e .

2.4 Feedback Linearization

Feedback linearization is a geometric approach for synthesizing stabilizing controllers for a special class of nonlinear systems [111], [112]. A feedback linearizing controller eliminates the nonlinear components in a subset of the system dynamics, and replaces them with stable linear dynamics. While eliminating nonlinear dynamics through control may result in less efficient controllers [109], be fragile to model uncertainty [113], and induce unstable internal dynamics [114], feedback linearization is an important constructive tool in nonlinear control design as it enables the synthesis of local exponential CLFs with which efficient and robust control can be performed. My review of feedback linearization will focus on only the elements necessary to support the contributions of this thesis, and thus will begin from an advanced starting point rather than

the usual starting point of defining a *relative degree*, but I note that a thorough introduction to the material can be found in [111]. Following this definition, I will provide a concrete example in the form of a fully-actuated robotic system.

I now provide the following definition of feedback linearizability [115]:

Definition 8 (*Feedback Linearizability*). Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (2.1) with equilibrium input \mathbf{u}_e . The open-loop system (2.1) is said to be *locally feedback linearizable* with respect to the equilibrium point \mathbf{x}_e if there exists an open set $D \subseteq E$ with $\mathbf{x}_e \in D$, a function $\Phi : D \rightarrow \mathbb{R}^n$ that is a diffeomorphism² between D and $\Phi(D)$ satisfying $\Phi(\mathbf{x}_e) = \mathbf{0}_n$, constants $\gamma, k \in \mathbb{Z}_{>0}$ satisfying $\gamma \leq n$ and $k \leq m$, a controller $\mathbf{k}_{\text{fbl}} : D \times \mathbb{R}^k \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on $D \times \mathbb{R}^k$ and satisfies $\mathbf{k}_{\text{fbl}}(\mathbf{x}_e, \mathbf{0}_k) = \mathbf{u}_e$, a controllable pair $(\mathbf{A}, \mathbf{B}) \in \mathbb{R}^{\gamma \times \gamma} \times \mathbb{R}^{\gamma \times k}$, and functions $\mathbf{f}_\eta : \Phi(D) \rightarrow \mathbb{R}^\gamma$, $\mathbf{g}_\eta : \Phi(D) \rightarrow \mathbb{R}^{\gamma \times m}$ and $\boldsymbol{\omega} : \Phi(D) \rightarrow \mathbb{R}^{n-\gamma}$ that are continuously differentiable³ on $\Phi(D)$ such that:

1. The functions $\mathbf{f}_\xi : \Phi(D) \rightarrow \mathbb{R}^n$ and $\mathbf{g}_\xi : \Phi(D) \rightarrow \mathbb{R}^{n \times m}$ defined as:

$$\mathbf{f}_\xi(\boldsymbol{\xi}) \triangleq \frac{\partial \Phi}{\partial \mathbf{x}}(\Phi^{-1}(\boldsymbol{\xi}))\mathbf{f}(\Phi^{-1}(\boldsymbol{\xi})), \quad \mathbf{g}_\xi(\boldsymbol{\xi}) \triangleq \frac{\partial \Phi}{\partial \mathbf{x}}(\Phi^{-1}(\boldsymbol{\xi}))\mathbf{g}(\Phi^{-1}(\boldsymbol{\xi})), \quad (2.25)$$

for all $\boldsymbol{\xi} \in \Phi(D)$ satisfy:

$$\mathbf{f}_\xi(\boldsymbol{\xi}) = \begin{bmatrix} \mathbf{f}_\eta(\boldsymbol{\xi}) \\ \boldsymbol{\omega}(\boldsymbol{\xi}) \end{bmatrix}, \quad \mathbf{g}_\xi(\boldsymbol{\xi}) = \begin{bmatrix} \mathbf{g}_\eta(\boldsymbol{\xi}) \\ \mathbf{0}_{n-\gamma} \end{bmatrix}, \quad (2.26)$$

for all $\boldsymbol{\xi} \in \Phi(D)$.

2. We have that:

$$\mathbf{f}_\eta(\boldsymbol{\xi}) + \mathbf{g}_\eta(\boldsymbol{\xi})\mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \boldsymbol{\nu}) = \mathbf{A}\boldsymbol{\eta} + \mathbf{B}\boldsymbol{\nu}, \quad (2.27)$$

for all $\boldsymbol{\xi} \in \Phi(D)$ and $\boldsymbol{\nu} \in \mathbb{R}^k$, where $\boldsymbol{\eta} \in \mathbb{R}^\gamma$ and $\mathbf{z} \in \mathbb{R}^{n-\gamma}$ satisfy $(\boldsymbol{\eta}, \mathbf{z}) = \boldsymbol{\xi}$.

² $\Phi : D \rightarrow \mathbb{R}^n$ is a diffeomorphism between D and $\Phi(D)$ if it is smooth on D , and there exists an inverse function $\Phi^{-1} : \Phi(D) \rightarrow D$ that is smooth on $\Phi(D)$. Note that this requires that $\frac{\partial \Phi}{\partial \mathbf{x}}(\mathbf{x})$ is full rank for all $\mathbf{x} \in D$.

³This assumption on \mathbf{f}_η , \mathbf{g}_η , and $\boldsymbol{\omega}$ is typically stronger than necessary. It is made here only to reason about stability using linearizations of $\boldsymbol{\omega}$. I note that if \mathbf{f} and \mathbf{g} are continuously differentiable on E , then the smoothness of Φ implies these properties.

If $\gamma = n$, the open-loop system (2.1) is said to be locally *full-state feedback linearizable* with respect to the equilibrium point \mathbf{x}_e , and the function $\boldsymbol{\omega}$ does not appear in (2.26). The corresponding *normal-form* open-loop system is:

$$\dot{\boldsymbol{\xi}} = \begin{bmatrix} \dot{\boldsymbol{\eta}} \\ \dot{\mathbf{z}} \end{bmatrix} \triangleq \begin{bmatrix} \mathbf{f}_\eta(\boldsymbol{\xi}) \\ \boldsymbol{\omega}(\boldsymbol{\xi}) \end{bmatrix} + \begin{bmatrix} \mathbf{g}_\eta(\boldsymbol{\xi}) \\ \mathbf{0}_{n-\gamma} \end{bmatrix} \mathbf{u} = \mathbf{f}_\xi(\boldsymbol{\xi}) + \mathbf{g}_\xi(\boldsymbol{\xi})\mathbf{u}. \quad (2.28)$$

The term $\boldsymbol{\nu}$ is referred to as an *auxiliary input* to the controller \mathbf{k}_{fbl} . The requirement that $k \leq m$ means that the number of *outputs* being feedback linearized is less than or equal to the number of control inputs. The controllable pair (\mathbf{A}, \mathbf{B}) are the linear dynamics achieved through feedback linearization, and for which standard linear control design techniques such as pole-placement, linear quadratic regulation (LQR), or robust control designs [99] can be used. The constant γ captures the notion of a *vector relative degree* for the system, while the function $\boldsymbol{\omega}$ captures the *internal dynamics* that arise due to creating a linear relationship between control inputs and outputs through feedback when $\gamma < n$. I will refer to $\boldsymbol{\xi}$ as the *normal state*, $\boldsymbol{\eta}$ as the *output coordinates*, and \mathbf{z} as the *zero coordinates*, and note I consider a specialized setting of feedback linearization in which the dynamics of the zero coordinates do not depend on the input \mathbf{u} , which occurs frequently in robotics [116].

If $\boldsymbol{\nu}$ is specified by an auxiliary controller $\mathbf{k}_{\text{aux}} : \Phi(D) \rightarrow \mathbb{R}^k$ that is locally Lipschitz continuous on $\Phi(D)$ and satisfies $\mathbf{k}_{\text{aux}}(\mathbf{0}_n) = \mathbf{0}_k$, then the definitions in (2.25) imply that $\mathbf{0}_n$ is an equilibrium point of the closed-loop system:

$$\dot{\boldsymbol{\xi}} = \mathbf{f}_\xi(\boldsymbol{\xi}) + \mathbf{g}_\xi(\boldsymbol{\xi})\mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \mathbf{k}_{\text{aux}}(\boldsymbol{\xi})). \quad (2.29)$$

Defining $\mathbf{k}'_{\text{fbl}} : D \rightarrow \mathbb{R}^m$ as:

$$\mathbf{k}'_{\text{fbl}}(\mathbf{x}) = \mathbf{k}_{\text{fbl}}(\mathbf{x}, \mathbf{k}_{\text{aux}}(\Phi(\mathbf{x}))), \quad (2.30)$$

for all $\mathbf{x} \in D$, observe that \mathbf{k}'_{fbl} is locally Lipschitz continuous on D by the smoothness of Φ and:

$$\mathbf{f}(\mathbf{x}_e) + \mathbf{g}(\mathbf{x}_e)\mathbf{k}'_{\text{fbl}}(\mathbf{x}_e) = \mathbf{0}_n, \quad (2.31)$$

such that \mathbf{x}_e is an equilibrium point of the closed-loop system (2.2) using the controller \mathbf{k}'_{fbl} .

More generally, (2.25) and the fact that $\frac{\partial \Phi}{\partial \mathbf{x}}(\mathbf{x}_e)$ is full rank implies that \mathbf{x}_e is an unforced (forced) equilibrium point of the open-loop system (2.1) if and only

if the point $\mathbf{0}_n$ is an unforced (forced) equilibrium point of the normal-form open-loop system (2.28), noting that this requires $\boldsymbol{\omega}(\mathbf{0}_n) = \mathbf{0}_{n-\gamma}$ whether the equilibrium point is unforced or forced. Given a controller $\mathbf{k}_\xi : \Phi(D) \rightarrow \mathbb{R}^m$, define a normal-form closed-loop system:

$$\dot{\boldsymbol{\xi}} = \mathbf{f}_\xi(\boldsymbol{\xi}) + \mathbf{g}_\xi(\boldsymbol{\xi})\mathbf{k}_\xi(\boldsymbol{\xi}). \quad (2.32)$$

Similarly, $\mathbf{0}_n$ is an equilibrium point of the normal-form closed-loop system if and only if \mathbf{x}_e is an equilibrium point of the closed-loop system (2.2) using the controller $\mathbf{k} : D \rightarrow \mathbb{R}^m$ defined as $\mathbf{k}(\mathbf{x}) = \mathbf{k}_\xi(\Phi(\mathbf{x}))$. The diffeomorphic relationship between the closed-loop system (2.2) and normal-form closed-loop system (2.32) allows for the following results regarding the perseverance of local exponential stability⁴:

Theorem 6. *Let the open-loop system (2.1) be locally feedback linearizable with respect to an equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$ with $\mathbf{x}_e \in D$ and function $\Phi : D \rightarrow \mathbb{R}^n$ that is a diffeomorphism between D and $\Phi(D)$ and satisfies $\Phi(\mathbf{x}_e) = \mathbf{0}_n$. Given a controller $\mathbf{k}_\xi : \Phi(D) \rightarrow \mathbb{R}^m$ that renders $\mathbf{0}_n$ an equilibrium point of the normal-form closed-loop system (2.32), the normal-form closed-loop system is locally exponentially stable with respect to the equilibrium point $\mathbf{0}_n$ if and only if the closed-loop system (2.2) with the controller $\mathbf{k} : D \rightarrow \mathbb{R}^m$ defined as $\mathbf{k}(\mathbf{x}) = \mathbf{k}_\xi(\Phi(\mathbf{x}))$ is locally exponentially stable with respect to the equilibrium point \mathbf{x}_e .*

Considering the normal-form open-loop system (2.28), let $\mathbf{K} \in \mathbb{R}^{k \times \gamma}$ be a gain matrix such that the matrix $\mathbf{A}_{\text{cl}} \triangleq \mathbf{A} - \mathbf{BK}$ is Hurwitz (all of its eigenvalues have a negative real part). Such a gain matrix is guaranteed to exist as the pair (\mathbf{A}, \mathbf{B}) is controllable [117]. Using the auxiliary controller $\mathbf{k}_{\text{aux}}(\boldsymbol{\xi}) = -\mathbf{K}\boldsymbol{\eta}$ implies the output coordinates of the closed-loop system (2.29) evolve as the following output closed-loop system:

$$\dot{\boldsymbol{\eta}} = \mathbf{f}_\eta(\boldsymbol{\xi}) + \mathbf{g}_\eta(\boldsymbol{\xi})\mathbf{k}_{\text{fb}}(\Phi^{-1}(\boldsymbol{\xi}), -\mathbf{K}\boldsymbol{\eta}) = \mathbf{A}_{\text{cl}}\boldsymbol{\eta}. \quad (2.33)$$

Momentarily neglecting the dynamics of the zero coordinates, this output closed-loop system is locally exponentially stable with respect to the unique equilibrium point $\mathbf{0}_\gamma$. In this way, a subset of the system dynamics can be

⁴I prove a similar result to this in Section 6.3, for which the proof steps can be followed to establish this theorem.

made into a stable linear system through an appropriate choice of feedback. To achieve more efficient or robust control, this linearizing feedback can serve as a starting point for synthesizing a CLF for the output open-loop system:

$$\dot{\boldsymbol{\eta}} = \mathbf{f}_{\boldsymbol{\eta}}(\boldsymbol{\xi}) + \mathbf{g}_{\boldsymbol{\eta}}(\boldsymbol{\xi})\mathbf{u}, \quad (2.34)$$

and equilibrium point $\mathbf{0}_{\gamma}$. A core result in linear systems theory is that for any positive definite matrix $\mathbf{Q} \in \mathbb{S}_{>0}^{\gamma}$, there exists a unique positive definite matrix $\mathbf{P} \in \mathbb{S}_{>0}^{\gamma}$ satisfying the *continuous-time Lyapunov equation* (CTLE):

$$\mathbf{A}_{\text{cl}}^{\top} \mathbf{P} + \mathbf{P} \mathbf{A}_{\text{cl}} = -\mathbf{Q}. \quad (2.35)$$

For a given $\mathbf{Q} \in \mathbb{S}_{>0}^{\gamma}$ with a corresponding solution $\mathbf{P} \in \mathbb{S}_{>0}^{\gamma}$ to the CTLE, the function $V : \mathbb{R}^{\gamma} \rightarrow \mathbb{R}_{\geq 0}$ defined as:

$$V(\boldsymbol{\eta}) = \boldsymbol{\eta}^{\top} \mathbf{P} \boldsymbol{\eta}, \quad (2.36)$$

satisfies:

$$\lambda_{\min}(\mathbf{P}) \|\boldsymbol{\eta}\|^2 \leq V(\boldsymbol{\eta}) \leq \lambda_{\max}(\mathbf{P}) \|\boldsymbol{\eta}\|^2, \quad (2.37)$$

$$\dot{V}(\boldsymbol{\xi}, \mathbf{k}_{\text{fbl}}(\boldsymbol{\Phi}^{-1}(\boldsymbol{\xi}), -\mathbf{K}\boldsymbol{\eta})) \leq -\lambda_{\min}(\mathbf{Q}) \|\boldsymbol{\eta}\|^2, \quad (2.38)$$

$$\left\| \frac{\partial V}{\partial \boldsymbol{\eta}}(\boldsymbol{\eta}) \right\| \leq \lambda_{\max}(\mathbf{P}) \|\boldsymbol{\eta}\|, \quad (2.39)$$

for all $\boldsymbol{\xi} \in \boldsymbol{\Phi}(D)$, where (2.39) is an instance of the bound in (2.14). Furthermore, note that $\frac{\partial V}{\partial \boldsymbol{\eta}}$ is locally Lipschitz continuous on its domain. Consequently, V serves as a local exponential Lyapunov function for the output closed-loop system (2.33) and the equilibrium point $\mathbf{0}_{\gamma}$.

As in [118], to ensure that \dot{V} satisfies the strict inequality necessary for a local exponential CLF, consider a positive definite matrix $\mathbf{Q}_2 \in \mathbb{S}_{>0}^{\gamma}$ satisfying $\mathbf{Q}_2 \prec \mathbf{Q}$, yielding:

$$\dot{V}(\boldsymbol{\xi}, \mathbf{k}_{\text{fbl}}(\boldsymbol{\Phi}^{-1}(\boldsymbol{\xi}), -\mathbf{K}\boldsymbol{\eta})) < -\lambda_{\min}(\mathbf{Q}_2) \|\boldsymbol{\eta}\|^2, \quad (2.40)$$

for all $\boldsymbol{\xi} \in \boldsymbol{\Phi}(D) \setminus (\{\mathbf{0}_{\gamma}\} \times \mathbb{R}^{n-\gamma})$. Consequently, we have that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}(\boldsymbol{\xi}, \mathbf{u}) < -\lambda_{\min}(\mathbf{Q}_2) \|\boldsymbol{\eta}\|^2, \quad (2.41)$$

for all $\boldsymbol{\xi} \in \boldsymbol{\Phi}(D) \setminus (\{\mathbf{0}_{\gamma}\} \times \mathbb{R}^{n-\gamma})$, and thus V is a CLF for the output open-loop dynamics (2.34) and equilibrium point $\mathbf{0}_{\gamma}$. We can define the non-empty pointwise set:

$$K_{\text{CLF}}(\boldsymbol{\xi}) \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid \dot{V}(\boldsymbol{\xi}, \mathbf{u}) \leq -\lambda_{\min}(\mathbf{Q}_2) \|\boldsymbol{\eta}\|^2\}, \quad (2.42)$$

for all $\boldsymbol{\xi} \in \Phi(D)$, noting that $\mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), -\mathbf{K}\boldsymbol{\eta}) \in K_{\text{CLF}}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in \Phi(D)$, and synthesize an optimization-based controller $\mathbf{k}_{\text{CLF}} : \Phi(D) \rightarrow \mathbb{R}^m$ as follows:

$$\begin{aligned} \mathbf{k}_{\text{CLF}}(\boldsymbol{\xi}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \mathbf{0}_k)\|^2 && \text{(Output-CLF-QP)} \\ \text{s.t. } \dot{V}(\boldsymbol{\xi}, \mathbf{u}) &\leq -\lambda_{\min}(\mathbf{Q}_2)\|\boldsymbol{\eta}\|^2. \end{aligned}$$

Using the techniques in [51] with the local Lipschitz continuity of \mathbf{k}_{fbl} on $D \times \mathbb{R}^k$, this controller can be shown to be locally Lipschitz continuous on $\Phi(D)$ (including $\Phi(D) \cap (\{\mathbf{0}_\gamma\} \times \mathbb{R}^{n-\gamma})$). While the techniques in [51] cover the case when $\boldsymbol{\xi} \in \Phi(D) \setminus (\{\mathbf{0}_\gamma\} \times \mathbb{R}^{n-\gamma})$, I will briefly sketch why this is true for $\boldsymbol{\xi} \in \Phi(D) \cap (\{\mathbf{0}_\gamma\} \times \mathbb{R}^{n-\gamma})$. Note that for such $\boldsymbol{\xi}$, we have that $\mathbf{k}_{\text{CLF}}(\boldsymbol{\xi}) = \mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \mathbf{0}_k)$. Let $\boldsymbol{\xi}' = (\boldsymbol{\eta}', \mathbf{z}') \in \Phi(D)$. Observe that $\mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}'), -\mathbf{K}\boldsymbol{\eta}') \in K_{\text{CLF}}(\boldsymbol{\xi}')$ implies:

$$\|\mathbf{k}_{\text{CLF}}(\boldsymbol{\xi}') - \mathbf{k}_{\text{CLF}}(\boldsymbol{\xi})\| = \|\mathbf{k}_{\text{CLF}}(\boldsymbol{\xi}') - \mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \mathbf{0}_k)\|, \quad (2.43)$$

$$\leq \|\mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}'), -\mathbf{K}\boldsymbol{\eta}') - \mathbf{k}_{\text{fbl}}(\Phi^{-1}(\boldsymbol{\xi}), \mathbf{0}_k)\|, \quad (2.44)$$

$$\leq L_{\mathbf{k}_{\text{fbl}}} \left\| \begin{bmatrix} \Phi^{-1}(\boldsymbol{\xi}') - \Phi^{-1}(\boldsymbol{\xi}) \\ -\mathbf{K}\boldsymbol{\eta}' \end{bmatrix} \right\|, \quad (2.45)$$

$$\leq L_{\mathbf{k}_{\text{fbl}}} L_{\Phi^{-1}} \|\boldsymbol{\xi}' - \boldsymbol{\xi}\| + L_{\mathbf{k}_{\text{fbl}}} \|\mathbf{K}\| \|\boldsymbol{\eta}'\|, \quad (2.46)$$

$$\leq L_{\mathbf{k}_{\text{fbl}}} (\|\mathbf{K}\| + L_{\Phi^{-1}}) \|\boldsymbol{\xi}' - \boldsymbol{\xi}\|, \quad (2.47)$$

where $L_{\mathbf{k}_{\text{fbl}}}, L_{\Phi^{-1}} \in \mathbb{R}_{\geq 0}$ are local Lipschitz constants that hold by restricting $\boldsymbol{\xi}'$ to a sufficiently small neighborhood of $\boldsymbol{\xi}$, and the last inequality used the fact that $\|\boldsymbol{\eta}'\| \leq \|\boldsymbol{\xi}' - \boldsymbol{\xi}\|$. Because Φ is smooth, the controller $\mathbf{k}'_{\text{CLF}} : D \rightarrow \mathbb{R}^m$ defined as $\mathbf{k}'_{\text{CLF}}(\mathbf{x}) = \mathbf{k}_{\text{CLF}}(\Phi(\mathbf{x}))$ is locally Lipschitz continuous on D .

While the output closed-loop dynamics may be locally exponentially stable with respect to $\mathbf{0}_\gamma$, local exponential stability of the zero coordinate dynamics is still necessary for the full normal-form closed-loop system (2.2) to be locally exponentially stable with respect to $\mathbf{0}_n$. This is important to consider because if the zero coordinates do not display stability properties, they can lead to solutions $\boldsymbol{\varphi}$ that exit the domain D where the controllers \mathbf{k}_{fbl} and \mathbf{k}_{CLF} are well-defined (the interval of existence of the solution would be finite). To this end, define the *zero dynamics* system as:

$$\dot{\mathbf{z}} = \boldsymbol{\omega}((\mathbf{0}_\gamma, \mathbf{z})), \quad (2.48)$$

which correspond to the evolution of the zero coordinates when the output $\boldsymbol{\eta} = \mathbf{0}_\gamma$. Note that $\mathbf{0}_{n-\gamma}$ is an equilibrium point for this system as \mathbf{x}_e is assumed

to be an unforced or forced equilibrium point of the open-loop system (2.1), implying $\boldsymbol{\omega}(\mathbf{0}_n) = \mathbf{0}_{n-\gamma}$, yielding the following result [111, Proposition 4.4.2]:

Theorem 7. *If the zero dynamics system (2.48) is locally exponentially stable with respect to $\mathbf{0}_{n-\gamma}$, then for any controller $\mathbf{k}_\xi : \Phi(D) \rightarrow \mathbb{R}^m$ that is continuous on $\Phi(D)$, locally Lipschitz continuous on $\Phi(D) \setminus \{\mathbf{0}_n\}$, renders $\mathbf{0}_n$ an equilibrium point of the normal-form closed-loop system (2.32), and satisfies $\mathbf{k}_\xi(\boldsymbol{\xi}) \in K_{CLF}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in \Phi(D)$, the normal-form closed-loop system (2.32) is locally exponentially stable with respect to the equilibrium point $\mathbf{0}_n$.*

Combining Theorems 6 and 7, local exponential stability of the zero dynamics in conjunction with a locally exponentially stabilizing controller \mathbf{k}_ξ yields that the closed-loop system (2.2) using $\mathbf{k}(\mathbf{x}) = \mathbf{k}_\xi(\Phi(\mathbf{x}))$ is locally exponentially stable with respect to the equilibrium point \mathbf{x}_e .

To provide a concrete example of feedback linearization, consider a fully-actuated robotic system governed by (2.5) with a forced equilibrium point $\mathbf{x}_e = (\mathbf{q}_e, \mathbf{0}_{n_c}) \in \mathcal{X}$ and corresponding equilibrium torque:

$$\boldsymbol{\tau}_e = \mathbf{B}(\mathbf{q}_e)^{-1}(\mathbf{C}(\mathbf{q}_e, \mathbf{0}_{n_c}) + \mathbf{G}(\mathbf{q}_e)). \quad (2.49)$$

Let us define the function $\Phi : \mathcal{X} \rightarrow \mathbb{R}^n$ as:

$$\Phi(\mathbf{x}) = \begin{bmatrix} \mathbf{q} - \mathbf{q}_e \\ \dot{\mathbf{q}} \end{bmatrix} = \boldsymbol{\xi}. \quad (2.50)$$

It is easy to see that Φ is smooth on \mathcal{X} and has an inverse $\Phi^{-1} : \Phi(\mathcal{X}) \rightarrow \mathcal{X}$ that is smooth on $\Phi(\mathcal{X})$. Consider a controller $\mathbf{k}_{\text{fbl}} : \mathcal{X} \times \mathbb{R}^{n_c} \rightarrow \mathbb{R}^m$ given by:

$$\mathbf{k}_{\text{fbl}}(\mathbf{x}, \boldsymbol{\nu}) = \mathbf{B}(\mathbf{q})^{-1} \mathbf{D}(\mathbf{q}) (\mathbf{D}(\mathbf{q})^{-1} (\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q})) + \boldsymbol{\nu}), \quad (2.51)$$

which is locally Lipschitz continuous on $\mathcal{X} \times \mathbb{R}^{n_c}$, and observe that $\mathbf{k}_{\text{fbl}}(\mathbf{x}_e, \mathbf{0}_{n_c}) = \boldsymbol{\tau}_e$. This controller yields:

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{0}_{n_c \times n_c} & \mathbf{I}_{n_c} \\ \mathbf{0}_{n_c \times n_c} & \mathbf{0}_{n_c \times n_c} \end{bmatrix}}_{\mathbf{A}} \mathbf{x} + \underbrace{\begin{bmatrix} \mathbf{0}_{n_c \times n_c} \\ \mathbf{I}_{n_c} \end{bmatrix}}_{\mathbf{B}} \boldsymbol{\nu}, \quad (2.52)$$

where $(\mathbf{A}, \mathbf{B}) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n_c}$ are a controllable pair. Noting that $\mathbf{A}\mathbf{x}_e = \mathbf{0}_n$ and denoting $\boldsymbol{\eta} = \mathbf{x} - \mathbf{x}_e$, these dynamics can be rewritten as:

$$\dot{\boldsymbol{\eta}} = \mathbf{A}\boldsymbol{\eta} + \mathbf{B}\boldsymbol{\nu}, \quad (2.53)$$

mirroring the form in (2.27) with functions $\mathbf{f}_\eta(\boldsymbol{\xi}) = \mathbf{f}(\Phi^{-1}(\boldsymbol{\xi}))$ and $\mathbf{g}_\eta(\boldsymbol{\xi}) = \mathbf{g}(\Phi^{-1}(\boldsymbol{\xi}))$ with \mathbf{f} and \mathbf{g} defined in (2.6), and no function $\boldsymbol{\omega}$, such that $\gamma = n$ and $k = n_c$. Consequently, a fully-actuated robotic system is locally full-state feedback linearizable with respect to the equilibrium point \mathbf{x}_e with $D = \mathcal{X}$. A detail that will appear in Chapter 3 when considering the impact of model error on CLF-based controllers is the model dependence of the feedback linearizing controller (2.51).

2.5 Control Barrier Functions

Historically, barrier methods were first developed in the context of constrained optimization [119], wherein constraint satisfaction could be achieved through increasingly large penalties on constraint violation. The idea to use barrier functions (BFs) in the context of nonlinear dynamical systems was first proposed in [120] with the goal of certifying the forward invariance of a set for a closed-loop system, similar to the way in which local exponential Lyapunov functions serve as a certificate of local exponential stability. This idea was further developed in [10], yielding the first definition of Control Barrier Functions (CBFs) as a tool for simultaneously synthesizing a safety-critical controller and a BF for the corresponding closed-loop system. The controller in this work was based on a structured design developed with CLFs for stabilization in [9]. A consequence of this structured design was that the controller’s only goal was safety, which, while of vital importance, is often not the only control objective that needs to be achieved (such as tracking performance).

An important pair of changes to the formulation of CBFs that increased their potential to provide safety-critical yet performant control were proposed in [11]. The first change was incorporating an extended class- \mathcal{K} function (defined below) into the CBF time derivative condition required for safety. This change allowed the system state to approach the boundary of the forward invariant set as long as it displayed a safe degree of “braking”, reducing the conservative nature of the original definition of CBFs. The second change was realizing that the CBF time derivative was affine in the control input, and thus could be directly incorporated as a constraint in a convex optimization problem, as was done with CLFs in Section 2.3. This resulted in a way to minimally “filter” a controller designed for performance such that it meets safety requirements. Subsequently, the notion of *zeroing* CBFs were introduced in [62], which reframed the preceding definition of CBFs in a way that was well-

defined outside of a set being kept safe [121, Remark 3]. I note that in this thesis my definition of CBFs will correspond to zeroing CBFs.

Before defining BF and CBFs, I recall the following useful definitions:

Definition 9 (*Class \mathcal{K} Function* [122]). Let $a \in \mathbb{R}_{>0}$. A function $\alpha : [0, a) \rightarrow \mathbb{R}_{\geq 0}$ that is continuous on $[0, a)$ is said to be *class- \mathcal{K}* ($\alpha \in \mathcal{K}$) if $\alpha(0) = 0$ and $\alpha(r_1) < \alpha(r_2)$ for all $r_1, r_2 \in [0, a)$ satisfying $r_1 < r_2$. The function α is said to be *class- \mathcal{K}_∞* ($\alpha \in \mathcal{K}_\infty$) if $a = \infty$ and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$.

Definition 10 (*Extended Class \mathcal{K} Function* [121]). Let $a, b \in \mathbb{R}_{>0}$. A function $\alpha : (-b, a) \rightarrow \mathbb{R}$ that is continuous on $(-b, a)$ is said to be *extended class- \mathcal{K}* ($\alpha \in \mathcal{K}^e$) if $\alpha(0) = 0$ and $\alpha(r_1) < \alpha(r_2)$ for all $r_1, r_2 \in (-b, a)$ satisfying $r_1 < r_2$. The function α is said to be *extended class- \mathcal{K}_∞* ($\alpha \in \mathcal{K}_\infty^e$) if $a = \infty$, $b = \infty$, $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, and $\lim_{r \rightarrow -\infty} \alpha(r) = -\infty$.

A barrier function is used to describe the safety of the closed-loop system (2.2) with respect to a set $\mathcal{C} \subset E$ specified as the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E :

$$\mathcal{C} \triangleq \{\mathbf{x} \in E \mid h(\mathbf{x}) \geq 0\}, \quad (2.54)$$

$$\text{Int}(\mathcal{C}) \triangleq \{\mathbf{x} \in E \mid h(\mathbf{x}) > 0\}, \quad (2.55)$$

$$\partial\mathcal{C} \triangleq \{\mathbf{x} \in E \mid h(\mathbf{x}) = 0\}. \quad (2.56)$$

Given this construction, barrier functions are defined as in [62]:

Definition 11 (*Barrier Function (BF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is a *barrier function (BF)* for the closed-loop system (2.2) on \mathcal{C} if there exists an $\alpha \in \mathcal{K}^e$ such that:

$$\dot{h}(\mathbf{x}) \triangleq \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{f}(\mathbf{x})}_{L_{\mathbf{f}}h(\mathbf{x})} + \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x})}_{L_{\mathbf{g}}h(\mathbf{x})} \geq -\alpha(h(\mathbf{x})), \quad (2.57)$$

for all $\mathbf{x} \in E$.

Note that the function $\alpha \in \mathcal{K}^e$ in this definition must be defined on the interval:

$$\left(\inf_{\mathbf{x} \in E} h(\mathbf{x}), \sup_{\mathbf{x} \in E} h(\mathbf{x}) \right) \subseteq \mathbb{R}. \quad (2.58)$$

The following result establishes how the existence of a barrier function for the closed-loop system (2.2) on \mathcal{C} certifies the safety of the closed-loop system (2.2) with respect to \mathcal{C} :

Theorem 8. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is a BF for the closed-loop system (2.2) on \mathcal{C} , then the closed-loop system (2.2) is safe with respect to the set \mathcal{C} .*

This result was first established in [62, Proposition 1] only requiring that the BF inequality (2.57) be satisfied for $\mathbf{x} \in \mathcal{C}$ by using Nagumo's theorem [104], [123] under the requirement that the function h has 0 as a *regular value*:

$$h(\mathbf{x}) = 0 \implies \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n, \quad (2.59)$$

for all $\mathbf{x} \in E$. Subsequently, an alternative proof of this result was provided in [124, Theorem 1] that does not require h have 0 as regular value, but rather uses the fact that $\alpha \in \mathcal{K}^e$ and the barrier function inequality (2.57) is satisfied for $\mathbf{x} \in E \setminus \mathcal{C}$. The definition of CBFs that I next consider will implicitly require that h has 0 as a regular value.

CBFs provide a tool for synthesizing both a safety-critical controller and a BF that certifies the corresponding closed-loop system (2.2) is safe with respect to the set \mathcal{C} :

Definition 12 (*Control Barrier Function (CBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is a *Control Barrier Function (CBF)* for the open-loop system (2.1) on \mathcal{C} if there exists an $\alpha \in \mathcal{K}^e$ such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) \triangleq \sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} > -\alpha(h(\mathbf{x})), \quad (2.60)$$

for all $\mathbf{x} \in E$

As with CLFs, the strict inequality in (2.60), which appears in [51], is important not only for ensuring a controller \mathbf{k} satisfying (2.57) can exist, but also that a controller meeting suitable regularity properties can be found. Note that the strictness also requires h to have 0 as a regular value. Similarly, an equivalent requirement to (2.60) is given by:

$$L_{\mathbf{g}}h(\mathbf{x}) = \mathbf{0}_m \implies L_{\mathbf{f}}h(\mathbf{x}) > -\alpha(h(\mathbf{x})), \quad (2.61)$$

for all $\mathbf{x} \in E$.

Given a CBF h for the open-loop system (2.1) on \mathcal{C} and a corresponding $\alpha \in \mathcal{K}^e$, define the point-wise set:

$$K_{\text{CBF}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \dot{h}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x})) \right\}. \quad (2.62)$$

This yields the following result [62]:

Theorem 9. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is a CBF for the open-loop system (2.1) on \mathcal{C} , then the set $K_{\text{CBF}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in E$, and for any controller \mathbf{k} that is locally Lipschitz continuous on E with $\mathbf{k}(\mathbf{x}) \in K_{\text{CBF}}(\mathbf{x})$ for all $\mathbf{x} \in E$, the function h is a BF for the closed-loop system (2.2) on \mathcal{C} .*

This result states that for a given $\alpha \in \mathcal{K}^e$, a CBF describes the set of inputs that a controller can take such that the function h serves as a BF for the closed-loop system (2.2) on \mathcal{C} . To see how CBFs can minimally modify a performant controller while ensuring safety, consider a *nominal* controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E , and define the optimization-based controller as was first presented in [27]:

$$\begin{aligned} \mathbf{k}_{\text{CBF}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && \text{(CBF-QP)} \\ \text{s.t. } &\dot{h}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x})). \end{aligned}$$

As with CLFs, because \dot{h} is affine with respect to the input \mathbf{u} , this controller is defined by a convex QP which can be efficiently solved. Moreover, if at a given state $\mathbf{x} \in E$ the nominal controller \mathbf{k}_{nom} satisfies $\dot{h}(\mathbf{x}, \mathbf{k}_{\text{nom}}(\mathbf{x})) \geq -\alpha(h(\mathbf{x}))$, then $\mathbf{k}_{\text{CBF}}(\mathbf{x}) = \mathbf{k}_{\text{nom}}(\mathbf{x})$. Thus the nominal controller can be designed to meet some performance criteria, and the controller (CBF-QP) can be used to modify the nominal controller only when it is unsafe, as specified by the BF requirement (2.57). The following result describes the properties of this controller [51, Theorem 1]:

Theorem 10. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E , let the function $\frac{\partial h}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ be locally Lipschitz continuous on E , let the nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ be locally Lipschitz continuous on E , and let h be a CBF for the open-loop system (2.1) on \mathcal{C} with α locally Lipschitz continuous on its domain. Then the*

controller $\mathbf{k}_{CBF} : E \rightarrow \mathbb{R}^m$ is locally Lipschitz continuous on E and satisfies $\mathbf{k}_{CBF}(\mathbf{x}) \in K_{CBF}(\mathbf{x})$ for all $\mathbf{x} \in E$.

Suppose that beyond ensuring the closed-loop system (2.2) is safe with respect to the set \mathcal{C} , an objective of control design is to locally exponentially stabilize the (2.2) system with respect to an unforced equilibrium point or forced equilibrium point $\mathbf{x}_e \in \text{Int}(\mathcal{C})$ with equilibrium input \mathbf{u}_e . Furthermore, suppose that there exists a CBF h for the open-loop system (2.1) on \mathcal{C} with corresponding $\alpha \in \mathcal{K}^e$, and a local exponential CLF V for the open-loop system (2.1) and equilibrium point \mathbf{x}_e with a corresponding open set $D \subseteq E$ that satisfies the small or continuous control property. A controller $\mathbf{k} : D \rightarrow \mathbb{R}^m$ that incorporates both the CLF and CBF can be specified as follows [62]:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \underset{(\mathbf{u}, \delta) \in \mathbb{R}^{m+1}}{\text{argmin}} \quad & \|\mathbf{u} - \mathbf{u}_e\|^2 + c\delta^2 && \text{(CLF-CBF-QP)} \\ \text{s.t.} \quad & \dot{V}(\mathbf{x}, \mathbf{u}) \leq -k_3\|\mathbf{x} - \mathbf{x}_e\|^a + \delta, \\ & \dot{h}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x})), && (2.63) \end{aligned}$$

for $c \in \mathbb{R}_{>0}$. This controller ensures the closed-loop system (2.2) is safe with respect to \mathcal{C} , and attempts to locally exponentially stabilize the closed-loop system (2.2) with respect to the equilibrium point \mathbf{x}_e . Because it may not be possible to simultaneously satisfy both the CLF and CBF inequality constraints, the CLF constraint is relaxed with the slack variable δ which is highly penalized by a large value of c in the cost function. While this controller does not endow the system with a formal stability guarantee, it can often effectively stabilize a system while ensuring safety. Furthermore, alternative relaxation schemes can be considered that provide local exponential stability results [51, QP Problem $-\gamma m$ Version]. The continuity on D and local Lipschitz continuity on $D \setminus \{\mathbf{x}_e\}$ of this controller can be established along the lines of [51, Theorem 1] under assumptions on the small or continuous control property and the local Lipschitz continuity of $\frac{\partial V}{\partial \mathbf{x}}$ and $\frac{\partial h}{\partial \mathbf{x}}$.

I will now make a useful observation regarding the relationship between local exponential CLFs and CBFs. First, as noted in [62, Proposition 2], outside of the set \mathcal{C} , a barrier function $h : E \rightarrow \mathbb{R}$ for the closed-loop system (2.2) can be used to construct a Lyapunov-like function for the closed-loop system (2.2) and the set \mathcal{C} , such that the closed-loop system (2.2) displays asymptotic stability properties with respect to the set \mathcal{C} . Considering the other direction, suppose

that $V : E \rightarrow \mathbb{R}_{\geq 0}$ is a local exponential CLF for the open-loop system (2.1) and an unforced or forced equilibrium point $\mathbf{x}_e \in E$ with corresponding open set D . In this development, I will restrict the open set E in the definition of forward invariance and CBFs to be $E = D$. The lower bound in (2.12) implies the existence of a $\bar{c} \in \mathbb{R}_{>0}$ such that for any $c \in [0, \bar{c})$, the set \mathcal{C} defined as:

$$\mathcal{C} \triangleq \{\mathbf{x} \in E \mid V(\mathbf{x}) \leq c\}, \quad (2.64)$$

satisfies $\mathcal{C} \subset D$. Next, define function the $h : D \rightarrow \mathbb{R}$ as follows:

$$h(\mathbf{x}) = c - V(\mathbf{x}), \quad (2.65)$$

noting that \mathcal{C} is the 0-superlevel set of h . We then have that:

$$\begin{aligned} \sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) &= \sup_{\mathbf{u} \in \mathbb{R}^m} -\dot{V}(\mathbf{x}, \mathbf{u}), & (2.66) \\ &= - \inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}(\mathbf{x}, \mathbf{u}), \\ &> k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \\ &\geq \frac{k_3}{k_2} V(\mathbf{x}), \\ &> \frac{k_3}{k_2} (V(\mathbf{x}) - c) = -\frac{k_3}{k_2} h(\mathbf{x}), & (2.67) \end{aligned}$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$. We also have that $\dot{V}(\mathbf{x}_e, \mathbf{u}) = \dot{h}(\mathbf{x}_e, \mathbf{u}) = 0$ for all $\mathbf{u} \in \mathbb{R}^m$, and $h(\mathbf{x}_e) = c > 0$, and consequently:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}_e, \mathbf{u}) > -\frac{k_3}{k_2} h(\mathbf{x}_e). \quad (2.68)$$

Thus we have that h is a CBF for the open-loop system (2.1) on the set \mathcal{C} . This is a useful observation as constructive tools such as feedback linearization can be used to produce local exponential CLFs, and from these CBFs can be constructed. The key step in this process occurs in (2.67), where the inequality is relaxed by the offset $-\frac{k_3}{k_2}c$. This term allows a system using a CBF to approach the boundary of the safe set, rather than explicitly converge to the point \mathbf{x}_e as is achieved with an exponential CLF and exponential stability.

2.6 Nonlinear Dynamics with Disturbances

The preceding developments of CLFs, feedback linearization, and CBFs were presented in the context of the ideal control system seen in Figure 1.1. For many real-world systems, control design is often complicated by challenges

such as unmodeled actuator dynamics or external perturbations. These challenges can be studied by considering nonlinear systems with *disturbances*.

I will consider two classes of systems with disturbances. The first class is the setting of *unmatched* disturbances described by the open-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} + \mathbf{d}(t), \quad (2.69)$$

where the function $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ is piecewise continuous⁵ on $\mathbb{R}_{\geq 0}$ and reflects a disturbance signal. A controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz on E yields the closed loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}) + \mathbf{d}(t). \quad (2.70)$$

As before, for any initial condition $\mathbf{x}_0 \in E$ and piecewise continuous disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$, there exists a maximal time interval $I_d(\mathbf{x}_0, \mathbf{d}) \triangleq [0, t_{\max}(\mathbf{x}_0, \mathbf{d}))$ and a unique piecewise continuously differentiable⁶ solution $\varphi_d : I_d(\mathbf{x}_0, \mathbf{d}) \rightarrow E$ to the closed-loop system (2.70) satisfying:

$$\dot{\varphi}_d(t) = \mathbf{f}(\varphi_d(t)) + \mathbf{g}(\varphi_d(t))\mathbf{k}(\varphi_d(t)) + \mathbf{d}(t), \quad (2.71)$$

$$\varphi_d(0) = \mathbf{x}_0, \quad (2.72)$$

for almost all $t \in I_d(\mathbf{x}_0, \mathbf{d})$ [100]. Note that the interval of existence and uniqueness $I_d(\mathbf{x}_0, \mathbf{d})$ depends not only on the initial condition \mathbf{x}_0 as it did in the absence of disturbances, but now it also depends on the particular disturbance signal the system experiences. The second class I will consider is the setting of *matched* disturbances, described by the open-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})(\mathbf{u} + \mathbf{d}(t)), \quad (2.73)$$

where the function $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ is piecewise continuous on $\mathbb{R}_{\geq 0}$ and reflects a disturbance in the input signal to the system. I will later show that this setting permits specialized robust control designs to be developed using the previous formulations of CLF and CBFs. A controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz on E yields the closed loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})(\mathbf{k}(\mathbf{x}) + \mathbf{d}(t)). \quad (2.74)$$

⁵This definition is taken as in [100], where for every compact interval $[a, b] \subset \mathbb{R}_{\geq 0}$, there exists a finite partition of the interval $[a, b]$, $\mathcal{T} = \{t_0, t_1, \dots, t_N\}$, such that \mathbf{d} is continuous on the domain (t_{i-1}, t_i) and the one-sided limits $\lim_{\delta \rightarrow 0^+} \mathbf{d}(t_{i-1} + \delta)$ and $\lim_{\delta \rightarrow 0^-} \mathbf{d}(t_i + \delta)$ exist for $i = 1, \dots, N$.

⁶Piecewise continuous differentiability is defined such that the function φ_d is continuous on $I_d(\mathbf{x}_0, \mathbf{d})$ with a derivative that is piecewise continuous on $I_d(\mathbf{x}_0, \mathbf{d})$.

As before, for any initial condition $\mathbf{x}_0 \in E$ and piecewise continuous disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$, there exists a maximal time interval $I_d(\mathbf{x}_0, \mathbf{d}) \triangleq [0, t_{\max}(\mathbf{x}_0, \mathbf{d}))$ and a unique piecewise continuously differentiable solution $\varphi_d : I_d(\mathbf{x}_0, \mathbf{d}) \rightarrow E$ to the closed-loop system (2.74) satisfying:

$$\dot{\varphi}_d(t) = \mathbf{f}(\varphi_d(t)) + \mathbf{g}(\varphi_d(t))(\mathbf{k}(\varphi_d(t)) + \mathbf{d}(t)), \quad (2.75)$$

$$\varphi_d(0) = \mathbf{x}_0, \quad (2.76)$$

for almost all $t \in I_d(\mathbf{x}_0, \mathbf{d})$ [100].

2.7 Input-to-State Stability & Input-to-State Safety

When studying nonlinear systems in the presence of disturbances, there are two perspectives of robust stability and safety I will consider. The first approach is to ensure that exponential stability or safety are achieved no matter what the disturbance signal may be, as is done with CBFs in [51]. In the setting of safety, this may induce significant conservativeness as the system is robust to worst-case disturbance signals that it may not experience. Achieving exponential stability of an equilibrium point in the presence of disturbances is often impossible, simply because it is impossible to choose one control input that will enforce \mathbf{x}_e is an equilibrium point for all possible disturbance signals without extremely restrictive assumptions [109, Section 3.4.1].

This particular challenge with stability in the presence of disturbances led to the proposal of *Input-to-State Stability* (ISS) [56], [57], [125], and later in the setting of safety, *Input-to-State Safety* (ISSf) [58]. These definitions consider the *degradation* of stability and safety in the presence of disturbances. In particular, systems that display ISS or ISSf properties gracefully degrade in the presence of disturbances, such that the magnitude of degradation is correlated with the magnitude of the disturbance that is experienced. This can provide for a powerful design paradigm when working with real-world systems in which disturbances are present [21], [126]. In this thesis I will explore several ways in which this perspective on robust control synthesis can be powerful for synthesizing performant yet robust controllers, including settings involving residual learning error, disturbances, and input-sampling effects.

I will now mathematically define these robustness properties. I first note these definitions are the same for both the unmatched and matched disturbance settings, with the distinction between those settings arising when considering

control synthesis in the next section. First, I will denote the following:

$$\|\mathbf{d}\|_\infty = \operatorname{ess\,sup}_{t \in \mathbb{R}_{\geq 0}} \|\mathbf{d}(t)\|, \quad (2.77)$$

where the essential supremum $\operatorname{ess\,sup}$ allows us to neglect the values of the piecewise continuous disturbance signal \mathbf{d} at the times at which it is discontinuous when assessing a bound on the disturbance. This permits the following definition of Input-to-State Stability:

Definition 13 (*Input-to-State Stability (ISS)*). Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2). The closed-loop system with unmatched disturbances (2.70) or matched disturbances (2.74) is said to be *locally exponentially Input-to-State Stable (ISS) with respect to \mathbf{x}_e* if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$, constants $M, \lambda, \bar{d} \in \mathbb{R}_{>0}$, and $\gamma \in \mathcal{K}$ such that for all $d \in [0, \bar{d}]$:

$$\mathbf{x}_0 \in D \implies \|\varphi_d(t) - \mathbf{x}_e\| \leq M\|\mathbf{x}_0 - \mathbf{x}_e\|e^{-\lambda t} + \gamma(d), \quad (2.78)$$

for all $t \in \mathbb{R}_{\geq 0}$ and disturbance signals $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ (\mathbb{R}^m for matched disturbances) that are piecewise continuous on $\mathbb{R}_{\geq 0}$ and satisfy $\|\mathbf{d}\|_\infty \leq d$.

This definition modifies the definition of local exponential stability by incorporating a class- \mathcal{K} function of the worst-case magnitude disturbance experienced by the system into the bound on the solution's deviation from the equilibrium point. If $d = 0$, such that the disturbance signal satisfies $\mathbf{d}(t) = \mathbf{0}_n$ ($\mathbf{0}_m$ for matched disturbances) for all $t \in \mathbb{R}_{\geq 0}$, we recover that the closed-loop system (2.2) without disturbances must be locally exponentially stable with respect to the equilibrium point \mathbf{x}_e . In the presence of disturbances, the system must converge exponentially to a ball of radius $\gamma(d)$ centered at \mathbf{x}_e , and remain within that ball. ISS is inherently a statement about the behavior of the closed-loop systems (2.70) or (2.74) across the spectrum of disturbances which are essentially bounded by the maximum disturbance magnitude \bar{d} . If the disturbance the system experiences is small, then the ball the system converges to is small, and likewise, if it is large (but satisfies the bound), the ball it will converge to is bounded. Producing a concrete statement about the size of the ball a closed-loop system with disturbances will actually remain in during its evolution requires producing an estimate of the worst-case possible disturbance signal the system will experience.

Before defining Input-to-State Safety, I present the following definition:

Definition 14 (*Disturbed Forward Invariance & Safety*). Let $d \in \mathbb{R}_{\geq 0}$. A set $\mathcal{C} \subset E$ is said to be *forward invariant up to d* for the closed-loop system with unmatched disturbances (2.70) or closed-loop system with matched disturbances (2.74) if for any initial condition $\mathbf{x}_0 \in \mathcal{C}$ and disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ (\mathbb{R}^m for matched disturbances) that is piecewise continuous on $\mathbb{R}_{\geq 0}$ and satisfies $\|\mathbf{d}\|_\infty \leq d$, we have $\varphi_d(t) \in \mathcal{C}$ for all $t \in I_d(\mathbf{x}_0, \mathbf{d})$. In this case, we call the system (2.70) or (2.74) *safe up to d* with respect to the set \mathcal{C} .

The first perspective on robustness to disturbances, where a system must satisfy safety properties with respect to the set \mathcal{C} for any disturbance it may experience, is described by the notion of the system being safe up to d with respect to \mathcal{C} . Input-to-State Safety will relax this definition as follows:

Definition 15 (*Input-to-State Safety (ISSf)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The closed-loop system with unmatched disturbances (2.70) or closed-loop system with matched disturbances (2.74) is said to be *Input-to-State Safe (ISSf)* with respect to the set \mathcal{C} if there exist $\bar{d} \in \mathbb{R}_{> 0}$ and $\gamma \in \mathcal{K}$ such that for all $d \in [0, \bar{d}]$, the set $\mathcal{C}_d \subset E$ defined as:

$$\mathcal{C}_d = \{\mathbf{x} \in E \mid h(\mathbf{x}) + \gamma(d) \geq 0\}, \quad (2.79)$$

is forward invariant up to d for the closed-loop system (2.70) or (2.74). If the system (2.70) or (2.74) is ISSf with respect to the set \mathcal{C} , the set \mathcal{C} is referred to as an *Input-to-State Safe set (ISSf set)*.

This definition states that the set that is kept forward invariant scales with the magnitude of the disturbance up to some maximum disturbance, and that in the absence of disturbances, safety of the original set \mathcal{C} is retained. Note that ISSf is explicitly constructed with a set that is the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E , rather than just a general set $\mathcal{C} \subset E$.

2.8 Input-to-State Stable Control Lyapunov Functions

Establishing ISS will be achieved through corresponding Lyapunov functions that accommodate the disturbance signal. Producing these certificate functions will be achieved through corresponding modifications to CLFs, where

the importance of unmatched and matched disturbances will become apparent. I now define Input-to-State Stable Lyapunov functions (ISS-LFs) as in [125, Definition 2.2]:

Definition 16 (*Input-to-State Stable Lyapunov Function (ISS-LF)*). Let \mathbf{x}_e be an equilibrium point of the closed-loop system (2.2). A function $V : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E is said to be a *local exponential Input-to-State Lyapunov function (ISS-LF)* for the closed-loop system with unmatched disturbances (2.70) and equilibrium point \mathbf{x}_e if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$, constants $k_1, k_2, k_3, a, \bar{d} \in \mathbb{R}_{>0}$, and $\rho \in \mathcal{K}$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.80)$$

$$\dot{V}(\mathbf{x}, \mathbf{d}) \triangleq L_f V(\mathbf{x}) + L_g V(\mathbf{x}) \mathbf{k}(\mathbf{x}) + \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{d} \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.81)$$

for all $\mathbf{x} \in D$ and $\mathbf{d} \in \mathbb{R}^n$ such that $\|\mathbf{d}\| \leq \bar{d}$ and $\|\mathbf{x} - \mathbf{x}_e\| \geq \rho(\|\mathbf{d}\|)$. The function V is said to be a ISS-LF for the closed-loop system with matched disturbances (2.70) and equilibrium point \mathbf{x}_e if (2.81) is replaced with:

$$\dot{V}(\mathbf{x}, \mathbf{d}) \triangleq L_f V(\mathbf{x}) + L_g V(\mathbf{x}) \mathbf{k}(\mathbf{x}) + L_g V(\mathbf{x}) \mathbf{d} \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (2.82)$$

for all $\mathbf{x} \in D$ and $\mathbf{d} \in \mathbb{R}^m$ such that $\|\mathbf{d}\| \leq \bar{d}$ and $\|\mathbf{x} - \mathbf{x}_e\| \geq \rho(\|\mathbf{d}\|)$.

This definition of ISS-LFs requires that the Lyapunov decay conditions (2.81) or (2.82) be met only for disturbances that are smaller (with the class- \mathcal{K} function ρ) than the distance of the state from the equilibrium point \mathbf{x}_e . A useful alternative understanding of the requirement in (2.81) and (2.82) is captured in the following result [125, Remark 2.4]:

Theorem 11. *Let \mathbf{x}_e be an equilibrium point of the closed-loop system (2.2). A function $V : E \rightarrow \mathbb{R}_{\geq 0}$ satisfying (2.80) for some $k_1, k_2, a \in \mathbb{R}_{>0}$ is a local exponential ISS-LF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) with corresponding open set $D \subseteq E$ if and only if there exist $k'_3, \bar{d} \in \mathbb{R}_{>0}$ and $\sigma \in \mathcal{K}$ such that:*

$$\dot{V}(\mathbf{x}, \mathbf{d}) \leq -k'_3 \|\mathbf{x} - \mathbf{x}_e\|^a + \sigma(\|\mathbf{d}\|), \quad (2.83)$$

for all $\mathbf{x} \in D$ and $\mathbf{d} \in \mathbb{R}^n$ ($\mathbf{d} \in \mathbb{R}^m$) such that $\|\mathbf{d}\| \leq \bar{d}$.

Note that this alternative condition uses the quantity $k'_3 \in \mathbb{R}_{>0}$ which is potentially different than the k_3 in (2.81) and (2.82). ISS-LFs serve as certificate

of local exponential ISS of a closed-loop system with respect to an equilibrium point as follows [125, Theorem 1]:

Theorem 12. *Let $\mathbf{x}_e \in E$ be an equilibrium point of the closed-loop system (2.2). If there exists a local exponential ISS-LF $V : E \rightarrow \mathbb{R}_{\geq 0}$ for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) and equilibrium point \mathbf{x}_e , then the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) is locally exponentially ISS with respect to \mathbf{x}_e .*

Extending the notion of an ISS-LF to an ISS Control Lyapunov Function (ISS-CLF) requires more care than the extension of Lyapunov functions to CLFs. One potential way for defining an ISS-CLF in the setting of matched disturbances would be to require that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}(\mathbf{x}, \mathbf{u}, \mathbf{d}) \triangleq \inf_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})(\mathbf{u} + \mathbf{d}) < -k_3\|\mathbf{x} - \mathbf{x}_e\|^a + \sigma(\|\mathbf{d}\|), \quad (2.84)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$ and $\mathbf{d} \in \mathbb{R}^n$ ($\mathbf{d} \in \mathbb{R}^m$) such that $\|\mathbf{d}\| \leq \bar{d}$. The issue with such a definition is that it is unclear how control inputs should be chosen, given the appearance of \mathbf{d} , which is not a decision variable, but rather an unknown perturbation. This issue of defining ISS-CLFs in a way that permits controller design has been considered in [127, Theorem 2] and [106, Section 6]. In this thesis, I will consider the following definition of an ISS-CLF:

Definition 17 (*Input-to-State Stable Control Lyapunov Function (ISS-CLF)*).

Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (2.1). A function $V : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E is said to be a *local exponential Input-to-State Stable Control Lyapunov Function (ISS-CLF)* for the open-loop system with unmatched disturbances (2.69) if there exist an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ and $k_1, k_2, k_3, \epsilon \in \mathbb{R}_{>0}$ such that:

$$k_1\|\mathbf{x} - \mathbf{x}_e\|^2 \leq V(\mathbf{x}) \leq k_2\|\mathbf{x} - \mathbf{x}_e\|^2, \quad (2.85)$$

$$\inf_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}V(\mathbf{x}') + L_{\mathbf{g}}V(\mathbf{x}')\mathbf{u} + \frac{1}{\epsilon} \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}') \right\|^2 < -k_3\|\mathbf{x}' - \mathbf{x}_e\|^2, \quad (2.86)$$

for all $\mathbf{x} \in D$ and $\mathbf{x}' \in D \setminus \{\mathbf{x}_e\}$. The function V is said to be a local exponential ISS-CLF for the open-loop system with matched disturbances (2.73) and equilibrium point \mathbf{x}_e if (2.86) is replaced with:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}V(\mathbf{x}') + L_{\mathbf{g}}V(\mathbf{x}')\mathbf{u} + \frac{1}{\epsilon} \|L_{\mathbf{g}}V(\mathbf{x}')\|^2 < -k_3\|\mathbf{x}' - \mathbf{x}_e\|^2, \quad (2.87)$$

for all $\mathbf{x}' \in D \setminus \{\mathbf{x}_e\}$.

First observe that I have specialized this definition to use a value of $a = 2$. Second, observe that the ISS-CLF condition for the matched disturbance setting (2.87) can be rewritten as the alternative CLF condition (2.18). This implies that if V is a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e , it is a local exponential ISS-CLF for the open-loop system with matched disturbances (2.73) and equilibrium point \mathbf{x}_e . This is a very useful property given the ability to synthesize CLFs through feedback linearization demonstrated in Section 2.4, and is a significant difference between the unmatched and matched disturbance settings.

Given a local exponential ISS-CLF V for the open-loop system with unmatched disturbances (2.69) and equilibrium point \mathbf{x}_e , define the following pointwise set:

$$K_{\text{ISS}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{u} \\ + \frac{1}{\epsilon} \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2 \end{array} \right\}, \quad (2.88)$$

and given a local exponential ISS-CLF V for the open-loop system with matched disturbances (2.73) and equilibrium point \mathbf{x}_e , define the following pointwise set:

$$K_{\text{ISS}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{u} \\ + \frac{1}{\epsilon} \|L_{\mathbf{g}}V(\mathbf{x})\|^2 \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2 \end{array} \right\}. \quad (2.89)$$

This set leads to the following result⁷:

Theorem 13. *Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (2.1), and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential ISS-CLF for the open-loop system with unmatched disturbances (2.69) (matched disturbances (2.73)) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$. Then, the set $K_{\text{ISS}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in D$, and for any controller $\mathbf{k} : D \rightarrow \mathbb{R}^m$ that is continuous on D , locally Lipschitz continuous on $D \setminus \{\mathbf{x}_e\}$, renders \mathbf{x}_e an equilibrium point of the closed-loop system (2.2), and satisfies $\mathbf{k}(\mathbf{x}) \in K_{\text{ISS}}(\mathbf{x})$ for all $\mathbf{x} \in D$, the function V is a local exponential ISS-LF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) and equilibrium point \mathbf{x}_e with $k'_3 = k_3$ and $\sigma \in \mathcal{K}$ defined as $\sigma(r) = \frac{\epsilon r^2}{4}$.*

⁷This result can be constructed along the proof steps used in Section 4.3.

This theorem not only states that an ISS-CLF V characterizes the set of inputs that allow V to serve as an ISS-LF, but it also provides a characterization of the constant k'_3 and class- \mathcal{K} function σ in the alternative definition of an ISS-LF given in (2.83). Note that the impact of the disturbance can be made smaller through small values of ϵ , but this requires that the controller provide more aggressive stabilizing control inputs. This is an analog to the fundamental balance between mitigating disturbances and high-gain control commonly studied in linear systems [128]. Building controllers using ISS-CLFs meeting the conditions of Theorem 13 will require the following variants of the small and continuous control properties:

Definition 18 (*ISS Small Control Property (ISS-SCP)*). Let $\mathbf{x}_e \in E$ be an unforced equilibrium point of the open-loop system (2.1), and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential ISS-CLF for the open-loop system with unmatched disturbances (2.69) (matched disturbances (2.73)) and equilibrium point \mathbf{x}_e with constant $\epsilon \in \mathbb{R}_{>0}$. The local exponential ISS-CLF V is said to satisfy the *ISS small control property (ISS-SCP)* if for any $\varepsilon \in \mathbb{R}_{>0}$, there exists an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ such that for any $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, there exists a $\mathbf{u} \in \mathbb{R}^m$ satisfying $\|\mathbf{u}\| < \varepsilon$ such that:

$$L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \quad (2.90)$$

for the unmatched disturbance setting, or such that:

$$L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \|L_{\mathbf{g}}V(\mathbf{x})\|^2 < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \quad (2.91)$$

for the matched disturbance setting.

Definition 19 (*ISS Continuous Control Property (ISS-CCP)*). Let $\mathbf{x}_e \in E$ be a forced equilibrium point of the open-loop system (2.1) for equilibrium input \mathbf{u}_e , and let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential ISS-CLF for the open-loop system with unmatched disturbances (2.69) (matched disturbances (2.73)) and equilibrium point \mathbf{x}_e with constant $\epsilon \in \mathbb{R}_{>0}$. The local exponential ISS-CLF V is said to satisfy the *ISS continuous control property (ISS-CCP)* if for any $\varepsilon \in \mathbb{R}_{>0}$, there exists an open set $D \subseteq E$ with $\mathbf{x}_e \in D$ such that for any $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, there exists a $\mathbf{u} \in \mathbb{R}^m$ satisfying $\|\mathbf{u} - \mathbf{u}_e\| < \varepsilon$ such that:

$$L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \quad (2.92)$$

for the unmatched disturbance setting, or such that:

$$L_f V(\mathbf{x}) + L_g V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \|L_g V(\mathbf{x})\|^2 < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \quad (2.93)$$

for the matched disturbance setting.

I note that in the matched disturbance setting, the relationship between a local exponential CLF satisfying the small or continuous property and a local ISS-CLF satisfying the small or continuous control property is non-trivial. Rather, it is necessary, but not sufficient, for the CLF to satisfy the small or continuous control property for the ISS-CLF to satisfy the ISS small or continuous control property. Lastly, I will define the convex optimization-based controllers:

$$\begin{aligned} \mathbf{k}_{ISS}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 && \text{(ISS-CLF-QP)} \\ \text{s.t. } & L_f V(\mathbf{x}) + L_g V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \end{aligned}$$

and:

$$\begin{aligned} \mathbf{k}_{ISS}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 && \text{(ISS-CLF-QP)} \\ \text{s.t. } & L_f V(\mathbf{x}) + L_g V(\mathbf{x})\mathbf{u} + \frac{1}{\epsilon} \|L_g V(\mathbf{x})\|^2 \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \end{aligned}$$

for the unmatched and matched disturbance settings, respectively. This leads to the following result, which can be proven similarly to [51, Theorem 1]:

Theorem 14. *Let $\mathbf{x}_e \in E$ be an unforced (forced) equilibrium point of the open-loop system (2.1) with equilibrium input \mathbf{u}_e , let $V : E \rightarrow \mathbb{R}_{\geq 0}$ be a local exponential ISS-CLF for the open-loop system with unmatched disturbances (2.69) or matched disturbances (2.73) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$, let the function $\frac{\partial V}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ be locally Lipschitz continuous on D , and assume that V satisfies the ISS small (continuous) control property. Then the controller $\mathbf{k}_{ISS} : D \rightarrow \mathbb{R}^m$ is continuous on D , locally Lipschitz continuous on $D \setminus \{\mathbf{x}_e\}$, renders \mathbf{x}_e an equilibrium point of the closed-loop system (2.2), and satisfies $\mathbf{k}_{ISS}(\mathbf{x}) \in K_{ISS}(\mathbf{x})$ for all $\mathbf{x} \in D$.*

2.9 Input-to-State Safe Control Barrier Functions

Establishing ISSf will be achieved through corresponding barrier functions that accommodate the disturbance signal. Producing these certificate functions will

be achieved through corresponding modifications to CBFs, where the importance of unmatched and matched disturbances will become apparent. I now define Input-to-State Safe barrier functions (ISSf-BFs) as in [58, Definition 4]:

Definition 20 (*Input-to-State Safe Barrier Function (ISSf-BF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is said to be an *Input-to-State Safe barrier function (ISSf-BF)* for the closed-loop system with unmatched disturbances (2.70) on \mathcal{C} if there exist $\alpha \in \mathcal{K}^e$, $\iota \in \mathcal{K}$, and $\bar{d} \in \mathbb{R}_{>0}$ such that:

$$\dot{h}(\mathbf{x}, \mathbf{d}) \triangleq L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x}) + \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{d} \geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{d}\|), \quad (2.94)$$

for all $\mathbf{x} \in E$ and $\mathbf{d} \in \mathbb{R}^n$ such that $\|\mathbf{d}\| \leq \bar{d}$. The function h is said to be an ISSf-BF for the closed-loop system with matched disturbances (2.74) on \mathcal{C} if (2.94) is replaced with:

$$\dot{h}(\mathbf{x}, \mathbf{d}) \triangleq L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{d} \geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{d}\|), \quad (2.95)$$

for all $\mathbf{x} \in E$ and $\mathbf{d} \in \mathbb{R}^m$ such that $\|\mathbf{d}\| \leq \bar{d}$.

ISSf-BFs serve as a certificate of ISSf of a closed-loop system with respect to a set \mathcal{C} as follows [58, Theorem 1]:

Theorem 15. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is an ISSf-BF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) on \mathcal{C} , then the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) is ISSf with respect to the set \mathcal{C} .*

In the same way that extending ISS-LFs to ISS-CLFs required considering the difference between the unmatched and matched disturbances, the notion of an ISSf-CBF will depend on what disturbance setting is being considered:

Definition 21 (*Input-to-State Safe Control Barrier Function (ISSf-CBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is said to be an *Input-to-State Safe Control Barrier Function (ISSf-CBF)* for the open-loop system with unmatched disturbances (2.69) on \mathcal{C} if there exist $\alpha \in \mathcal{K}^e$ and $\epsilon \in \mathbb{R}_{>0}$ such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon} \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 > -\alpha(h(\mathbf{x})), \quad (2.96)$$

for all $\mathbf{x} \in E$. The function h is said to be an ISSf-CBF for the open-loop system with matched disturbances (2.73) on \mathcal{C} if (2.96) is replaced with:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 > -\alpha(h(\mathbf{x})), \quad (2.97)$$

for all $\mathbf{x} \in E$.

As with ISS-CLFs, the condition for the matched disturbance setting (2.97) can be rewritten as the alternative CBF condition (2.61). This implies that if h is a CBF for the open-loop system (2.1) on \mathcal{C} , it is an ISSf-CBF for the open-loop system with matched disturbances (2.73) on \mathcal{C} .

Given an ISSf-CBF h for the open-loop system with unmatched disturbances (2.69) on \mathcal{C} , define the following pointwise set:

$$K_{\text{ISSf}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} \\ -\frac{1}{\epsilon} \|\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\|^2 \geq -\alpha(h(\mathbf{x})) \end{array} \right\}, \quad (2.98)$$

and given an ISS-CBF h for the open-loop system with matched disturbances (2.73) on \mathcal{C} , define the following pointwise set:

$$K_{\text{ISSf}}(\mathbf{x}) = \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} \\ -\frac{1}{\epsilon} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 \geq -\alpha(h(\mathbf{x})) \end{array} \right\}. \quad (2.99)$$

This set leads to the following result:

Theorem 16. *Let $h : E \rightarrow \mathbb{R}$ be an ISSf-CBF for the open-loop system with unmatched disturbances (2.69) (matched disturbances (2.73)) on \mathcal{C} . Then, the set $K_{\text{ISSf}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in E$, and for any controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E and satisfies $\mathbf{k}(\mathbf{x}) \in K_{\text{ISSf}}(\mathbf{x})$ for all $\mathbf{x} \in E$, the function h is an ISSf-BF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) on \mathcal{C} with $\iota \in \mathcal{K}$ defined as $\iota(r) = \frac{\epsilon r^2}{4}$.*

This theorem not only states that an ISSf-CBF h characterizes the set of inputs that allow h to serve as an ISS-BF, but it also provides a characterization of the class- \mathcal{K} function ι in the definition of an ISSf-BF. Again, note that the impact of the disturbance can be made smaller through small values of ϵ , but this requires that the controller provide inputs that meet a stronger safety criterion. This will be a trade-off that is studied in Chapter 4. Lastly, given a

nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E , I will define the following convex optimization-based controllers:

$$\begin{aligned} \mathbf{k}_{\text{ISSf}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && (\text{ISSf-CBF-QP}) \\ \text{s.t. } & L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon} \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 \geq -\alpha(h(\mathbf{x})), \end{aligned}$$

and:

$$\begin{aligned} \mathbf{k}_{\text{ISSf}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && (\text{ISSf-CBF-QP}) \\ \text{s.t. } & L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 \geq -\alpha(h(\mathbf{x})), \end{aligned}$$

for the unmatched and matched disturbance settings, respectively. This leads to the following result, which can be proven similarly to [51, Theorem 1]:

Theorem 17. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E , the function $\frac{\partial h}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ be locally Lipschitz continuous on E , the nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ be locally Lipschitz continuous on E , and h be an ISSf-CBF for the open-loop system with unmatched disturbances (2.69) or matched disturbances (2.73) on \mathcal{C} . Then the controller $\mathbf{k}_{\text{ISSf}} : E \rightarrow \mathbb{R}^m$ is locally Lipschitz continuous on E and satisfies $\mathbf{k}_{\text{ISSf}}(\mathbf{x}) \in K_{\text{ISSf}}(\mathbf{x})$ for all $\mathbf{x} \in E$.*

LEARNING & ADAPTIVE CONTROL FOR NONLINEAR SYSTEMS

The construction of controllers for nonlinear systems through Control Lyapunov Functions (CLFs), feedback linearization, and Control Barrier Functions (CBFs) is inherently a model-based design paradigm. In particular, the functions \mathbf{f} and \mathbf{g} describing the open-loop system (2.1) appear in the inequality constraints in the optimization-based controllers (CLF-QP), (CBF-QP), (CLF-CBF-QP), (ISS-CLF-QP), and (ISSf-CBF-QP), and similarly appear in the feedback linearizing controller (2.51) synthesized for robotic systems. Consequently, if there is error in a nominal model of the system (which will be mathematically described in Section 3.2), these controllers will not necessarily lead their respective true closed-loop systems to possess corresponding Lyapunov and barrier functions that certify local exponential stability and safety.

Robustness to model error is one of the core challenges in control theory that the use of feedback is known to solve [99], [128]. A limitation of a purely robust approach is that it requires a characterization of the possible model errors, and a control design that meets stability or safety requirements for all possible model errors. Importantly, the resulting controller will be robust to model errors that are not actually present between the nominal model of the system and the true dynamics of the system. Thus, a coarse characterization of possible model error can lead to unnecessary conservativeness which significantly degrades the performance of a controller

Alternatively, one may try to *identify* the actual model error between a nominal model of the system and true system, and design a controller using a corrected version of the nominal model. Classically, control theory has explored this question through the tools of *system identification* [129] and *adaptive control* [130]. System identification seeks to reduce model error by fitting models that more accurately capture the functions \mathbf{f} and \mathbf{g} using data, while adaptive control seeks to identify parametric model errors online while modifying parameters of a controller to ensure that stability and safety properties are

met as the model errors are identified. Recently, the advent of *machine learning* techniques for identifying complex models from data has led to significant development of data-driven techniques for control design beyond the classical tools of system identification [131]–[133], including both learning model errors and constructing updated model-based controllers [134]–[137], as well as direct model-free controller design [138]–[142].

My work presented in this chapter focuses on data-driven techniques for reducing model errors in a way that is amenable to the convex optimization-based control framework using CLFs and CBFs discussed in Chapter 2, as well as an adaptive control approach for ensuring safety in the presence of parametric model error. In Section 3.1 I discuss related work in the area of integrating learning techniques with CLFs and CBFs, as well as some of the historical and recent work in adaptive control for safety-critical systems. In Section 3.2 I will describe the mathematical structure of model errors that I consider. In Sections 3.3–3.6 I will present algorithms for learning model errors as they impact the time derivative of a CLF or CBF with simulation and experimental demonstrations, as well as provide a mathematical framework for understanding the impact of residual learning error on stability and safety guarantees through the lens of ISS and ISSf. In Section 3.7 I will consider some structural challenges that occur with the learning approaches in Sections 3.3 and 3.5 for high-dimensional systems, and explore an alternative learning framework using projected disturbances that can overcome these challenges. This work culminates in Section 3.8, wherein I will draw upon the lessons learned in the previous sections to propose an alternative approach for data-driven control using CLFs and CBFs that is constructed through robust convex optimization. This leads to a controller specified via a second-order cone program (SOCP), which is a convex optimization problem of a greater complexity than the QP-based controllers considered in Chapter 2 and Sections 3.3–3.7. Lastly, in Section 3.9, I will propose a framework for achieving safety of systems with parametric model error through the notion of adaptive CBFs. Key contributions of this work are described at the beginning of each respective section.

3.1 Related Work

Learning with Control Lyapunov Functions

The certification of stability provided by Lyapunov functions is an appealing tool for learning and control methods that seek to ensure a controller built

using learning is stabilizing. Early work in the topic of learning for dynamic systems used Lyapunov functions as a tool to ensure the stable learning of dynamics and regions of attraction of nonlinear systems [143], [144]. Similarly, Lyapunov functions were paired with model-based reinforcement learning (RL) to synthesize optimal controllers that provide stability guarantees [145], [146]. Further work has looked how existing models and controllers can be refined online while being kept stable using Lyapunov functions [147], and how Lyapunov functions can be learned from counterexamples [148]. A comprehensive survey on methods for learning Lyapunov functions can be found in [45].

The combination of learning with CLFs in an effort to synthesize stabilizing controllers, rather than just certify stability, was first considered in [17] to ensure stability of robotic reaching motions. The work in [38] proposed an uncertainty-based CLF using Gaussian Process Regression (GPR) that allowed the system to be driven into regions of the state space where there was high confidence in the learning model’s ability to predict model error in the system dynamics. Other work has looked at how CLFs can be learned from counterexamples [39], or how CLFs can be used to in the cost function of an RL problem to reduce sample-complexity [40].

My work in [149], which is presented in Section 3.3, was the first to look at learning model error as it directly impacts the time derivative of a CLF. This approach was motivated by the observation that only model errors appearing in the time derivative of a CLF impacted the local exponential stability of the system, and a complete characterization of model errors in the full state dynamics was unnecessary to certify stability. To ensure the well-posedness of this approach, I develop an understanding of structural assumptions that can be made about the true system that allow transferring a CLF for a nominal model to the true system. This enables enforcing a specific structure upon the learning models that allow them to be integrated into the optimization-based control framework described in Section 2 without destroying convexity. An analysis of the impact of residual learning error on local exponential stability guarantees was developed in my work in [150], which is presented in Section 3.4. Subsequent approaches have built upon the idea of learning model error directly as it impacts a CLF time derivative, such as in the context of RL [151]. Other work used affine GPR models to learn model errors in a CLF time derivative and produce a robust optimization-based controller [46] based on

the optimization-based control framework described in Section 2. My work in [152], discussed in Section 3.8, also takes a robust optimization-based approach for synthesizing a stabilizing controller directly using data, but takes a different perspective on quantifying model error. A comprehensive survey of CLFs being integrated with learning methods can be found in [41].

Learning with Control Barrier Functions

A significant amount of recent work has looked at utilizing learning with CBFs for dynamic systems. An indicator of this is the several recent survey papers focused on the two topics, including one focused on learning CBFs [45], one focused on learning applications using CBFs [41], one looking at safe learning in robotics [133], and my own work surveying connections between data-driven safety-filter techniques developed through Hamilton-Jacobi Reachability, CBFs, and predictive control [153]. One key reason for this magnitude of development compared to learning with CLFs is the permissive behaviors that can be achieved when enforcing safety of systems with CBFs. In particular, rather than stabilizing a system to a particular point or a trajectory, CBFs can allow a system to explore a safe region as it acquires data, builds learning models, and discovers optimal behaviors.

The first, and perhaps most obvious, way in which CBFs and learning can be integrated is by learning model errors that impact CBF-based controllers. Some approaches seek to do this task online, with the system required to remain safe during the process of collecting data and building learning models. This challenge was first explored in [16], where the dynamics of a quadrotor were safely learned using GPR by enforcing CBF constraints during the data collection process, while further work has looked at this task of safely learning dynamics online, either through Bayesian learning approaches [43], [154], [155] or online disturbance estimation [156]. Other approaches have taken an offline approach in which the system collects data, learning models are built offline, and the models are integrated into subsequent controller designs. The work in [42] looks at learning nonlinear dynamics using GPR and integrating the learning models into a CBF-based controller. Similar approaches using GPR and CBFs have been explored for learning in the context of decentralized collision avoidance [157] and elastic joint robots [158]. Additional approaches use Koopman operators to learn open-loop dynamics to be used in CBF-based controller synthesis [159], or learn the trajectories of closed-loop systems under

back-up set CBF-based controllers to reduce online computations [160].

Inspired by my work on learning model error as it impacts the time derivatives of CLFs presented in Sections 3.3 and 3.4, my work in [161], discussed in Section 3.5, looked at how model error could be learned directly as it impacts the time derivative of a CBF. This was an offline learning approach, in which data is episodically collected, integrated into structured learning models that preserve convexity of downstream control synthesis, and iteratively improve the safety of a system. Based on this approach, I developed a rigorous analysis of the impact of residual learning error on the safety of the system through the lens of ISSf in my work in [162], discussed in Section 3.6. This idea of learning model error only as it impacts a CBF time derivative was extended to the RL setting in [151], and used in online learning with extreme learning machines [163]. This idea was further explored in my work in [25], discussed in Section 3.7, which looked at learning unstructured model errors in a CBF time derivative when working with high-dimensional systems.

One direction of research that extended on the idea of learning model uncertainty directly as it impacts the time derivative of a CBF was through the use of robust optimization. This work began in the context of CLFs in [46], but my work in [152] done in parallel provided an alternative perspective that was applied to both CLFs and CBFs. The work in [46] was later extended to the setting of CBFs with a focus on verifying the feasibility of the robust optimization-based controllers in [164], and was subsequently utilized in an online learning setting in [165]. I note that the robust control approach developed in my work [152] has been specialized for certain cases of model uncertainty that allow for more efficient robust optimization problems in [166] and [167].

Another area that has seen rapid development in integrating learning with CBFs is safe RL. First steps in this area occurred in [44], [168], where CBFs were used to restrict policy updates such that the RL algorithm would only return safe policies. Work beyond this has sought to improve optimality and relax requirements on knowledge of a conservative but safe controller when starting the RL algorithm [169], [170]. Other work has looked at using approximate dynamic programming methods for safely solving the optimal control problem posed by RL online [171], multi-agent RL using multiple CBFs [172], and safe inverse RL [173]. Though not explicitly RL, the recent work in [174] looks at how online imitation learning can be safely performed using CBFs.

Instead of learning model errors or controllers that optimize reward functions, another avenue of recent work has explored how CBFs can be learned for nonlinear control systems. The first work considering this challenge was [175], which looked to learn CBFs for motion planning of robotic manipulators from human demonstrations. Following this, a framework for building CBFs from data was proposed in [176], and extended to hybrid nonlinear systems in [177]. Other developments have looked at the construction of CBFs online for robotic systems operating in unknown environments using distance measurements [178], using RL to synthesize CBFs that capture tolerable levels of failure risk [179], integrating learned CBFs into Nonlinear Model Predictive Control (NMPC) [180], and iteratively learning linear CBFs to approximate a complex set that is to be kept forward invariant [181].

Learning and CBFs have been integrated in several other creative research directions. One direction has looked at the tuning of CBFs and CBF-based controllers, such as using RL to tune CBF parameters to enable efficient collision avoidance of robotic manipulators [182], using RL to ensure input constraints are satisfied [183], or my own work in [23], discussed in Section 5.4, which uses Preference-Based Learning (PBL) to tune robustness parameters of a robust CBF controller in an effort to meet performance requirements. Another direction has been in ensuring safety of controllers using learned vision modules, such as my own work [184], or work on safe self-supervised online learning using stereo vision [185]. Other ideas have looked at how CBFs can be used to define cost functions for learning problems that produce bipedal walking gaits [186], how CBFs can be used to quantify the safety of learning-based autonomous driving (without controlling the system itself) [187], or how CBF safety filters can be cheaply approximated using learning models [188].

Adaptive Safety

My work in adaptive safety through CBFs was motivated by the existing notion of adaptive CLFs [47] that were used to stabilize systems in the presence of parametric model error. Adaptive CLFs were extended into a modular framework focused on ISS [189], considered in an inverse optimal control problem [48], and integrated with sums-of-squares (SOS) programming [49]. Recently, they have been used in an experimental setting on a brachiating robot [20]. Absent from these works was a focus on adaptively enforcing the safety of a particular set in the state space, with the set being kept invariant using adap-

tive CLFs lying in the composite state and parameter estimate space. My work in [190] sought to translate much of the framework for adaptive CLFs into the setting of adaptive safety, yielding a definition of adaptive CBFs. This work explored how enforcing safety of a particular set in the state space in the presence of parametric model error is more difficult than stabilizing a system in the presence of parametric model error. A key extension of this work integrated previously collected data by the system into the adaptation algorithm [191], similarly to concurrent learning [192], allowing for a reduction in conservativeness with respect to safety as a system evolves. Further extensions have looked at fixed time adaptation laws [193], hybrid adaptation laws [194], multiple CBFs [195], higher-order CBFs [196], utilizing the certainty-equivalence principle to simplify requirements for adaptive safety [197], dynamic regressor extension and mixing [198], modular adaptation law design [199], and approaches for permitting parametric model error in the actuation matrix \mathbf{g} [200].

3.2 Model Errors

Recall the open-loop system (2.1). In many practical settings, a control designer may not have perfect knowledge of the functions \mathbf{f} and \mathbf{g} . This lack of knowledge may arise due to unmodeled effects such as friction, or due to errors in system parameters such as inertias. Instead, control design is done with a nominal model that estimates the true open-loop system (2.1):

$$\hat{\dot{\mathbf{x}}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (3.1)$$

where the functions $\hat{\mathbf{f}} : E \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : E \rightarrow \mathbb{R}^{n \times m}$ are assumed to be locally Lipschitz continuous on E . I note that this construction implicitly assumes that the true system can be described by the state vector \mathbf{x} , and that the true open-loop system (2.1) is control-affine and time-invariant. I further note that this permits error in how the actuation matrix \mathbf{g} is modeled. Many existing approaches assume certainty in how control inputs enter the system [38], [201], [202]. These structural assumptions will ensure that the learning problem in Sections 3.3 and 3.5 is well-posed, while they are necessary for the stability and safety guarantees achieved by the data-driven control approach in Section 3.8. Adding and subtracting the nominal model open-loop system (3.1) from the true open-loop system (2.1) yields the following description of the true

open-loop system:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x}) + \underbrace{\mathbf{f}(\mathbf{x}) - \widehat{\mathbf{f}}(\mathbf{x})}_{\widetilde{\mathbf{f}}(\mathbf{x})} + \underbrace{(\mathbf{g}(\mathbf{x}) - \widehat{\mathbf{g}}(\mathbf{x}))}_{\widetilde{\mathbf{g}}(\mathbf{x})} \mathbf{u}, \quad (3.2)$$

where $\widetilde{\mathbf{f}} : E \rightarrow \mathbb{R}^n$ and $\widetilde{\mathbf{g}} : E \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous on E (by the assumptions on $\mathbf{f}, \mathbf{g}, \widehat{\mathbf{f}}$, and $\widehat{\mathbf{g}}$) and capture the error between the nominal model open-loop system (3.1) and the true open-loop system (2.1).

Example 2. Consider an inverted pendulum with input attenuation near the upright configuration described by the true open-loop system:

$$\dot{\mathbf{x}} = \begin{bmatrix} x_2 \\ \frac{g}{\ell} \sin(x_1) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1-0.75 \exp(-x_1^2)}{m\ell^2} \end{bmatrix} u, \quad (3.3)$$

with state $\mathbf{x} \in \mathbb{R}^2$, input $u \in \mathbb{R}$, and gravitational acceleration, length, and mass parameters $g, \ell, m \in \mathbb{R}_{>0}$, respectively. I will consider a nominal model of the true open-loop system given by:

$$\widehat{\dot{\mathbf{x}}} = \begin{bmatrix} x_2 \\ \frac{\widehat{g}}{\widehat{\ell}} \sin(x_1) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{\widehat{m}\widehat{\ell}^2} \end{bmatrix} u, \quad (3.4)$$

where $\widehat{g}, \widehat{\ell}, \widehat{m} \in \mathbb{R}_{>0}$ are estimates of the gravitational acceleration, length, and mass parameters. The corresponding functions $\widetilde{\mathbf{f}}$ and $\widetilde{\mathbf{g}}$ are given by:

$$\widetilde{\mathbf{f}}(\mathbf{x}) = \begin{bmatrix} 0 \\ \left(\frac{g}{\ell} - \frac{\widehat{g}}{\widehat{\ell}}\right) \sin(x_1) \end{bmatrix}, \quad \widetilde{\mathbf{g}}(\mathbf{x}) = \begin{bmatrix} 0 \\ \frac{1-0.75 \exp(-x_1^2)}{m\ell^2} - \frac{1}{\widehat{m}\widehat{\ell}^2} \end{bmatrix}. \quad (3.5)$$

A special case of model error that is typically considered in adaptive control is *parametric* model error in the drift dynamics, such that the true open-loop system (2.1) can be written as:

$$\dot{\mathbf{x}} = \underbrace{\widehat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^*}_{\mathbf{f}(\mathbf{x})} + \underbrace{\widehat{\mathbf{g}}(\mathbf{x})}_{\mathbf{g}(\mathbf{x})} \mathbf{u}, \quad (3.6)$$

where the function $\mathbf{F} : E \rightarrow \mathbb{R}^n \times \mathbb{R}^{n \times p}$ is locally Lipschitz continuous on E and $\boldsymbol{\theta}^* \in \mathbb{R}^p$. The function \mathbf{F} captures a collection of basis functions that are known to parameterize the model error, and $\boldsymbol{\theta}^*$ is the vector of true parameters that describe the system. Note that this construction does not permit error in the actuation matrix ($\widetilde{\mathbf{g}}(\mathbf{x}) = \mathbf{0}_{n \times m}$ for all $\mathbf{x} \in E$). Recalling the open-loop robotic system (2.5), this can be a limiting setting as the mass and inertia terms which appear in the matrix \mathbf{D} are often estimated imperfectly, leading to error between \mathbf{g} and $\widehat{\mathbf{g}}$.

Example 3. Considering the inverted pendulum with input attenuation in Example 2, suppose that $\hat{m} = m$, $\hat{\ell} = \ell$, and the attenuation term $-0.75 \exp(-x_1^2)$ is absent. Then the true open-loop system is of the form in (3.6) with:

$$\mathbf{F}(\mathbf{x}) = \begin{bmatrix} 0 \\ \sin(x_1) \end{bmatrix}, \quad (3.7)$$

and $\theta^* = \frac{g-\hat{g}}{\ell}$.

3.3 Learning with Control Lyapunov Functions

In this section I will present an episodic learning approach for mitigating model errors as they directly impact the time derivative of a local exponential CLF. The motivation for learning model error in this way is the idea that not all of the model error needs to be identified to achieve local exponential stability, but rather, only the model uncertainty that impacts the time derivative of a corresponding local exponential CLF. This learning approach will utilize a structured learning model that can be directly integrated into the convex optimization-based (CLF-QP) controller to improve the stability of a system, without destroying convexity.

One challenge in developing learning-based methods for controller improvement is how best to collect training data that accurately reflects the desired operating environment and control goals. In particular, exhaustive data collection typically scales exponentially with dimensionality of the joint state and control input space, and so should be avoided. But first pre-collecting data upfront can lead to poor performance as downstream control behavior may enter states that are not present in the pre-collected training data. I will leverage the episodic learning idea of Dataset Aggregation (DAgger) [203] to address these challenges in a data-efficient manner, and yield iteratively refined controllers.

The contributions of this section are as follows:

- An episodic learning algorithm that identifies model errors as they directly impact the time derivative of a local exponential CLF using a structured learning model that permits use of the convex optimization-based control framework described in Section 2.3.
- A characterization of the well-posedness of learning and control problems using CLFs by justifying the transfer of a local exponential CLF for a

nominal model open-loop system to the true open-loop system using assumptions on structural properties of the true open-loop system.

The text for this section is adapted from:

A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, “Episodic learning with control lyapunov functions for uncertain robotic systems,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp.6878-6884.

A. J. Taylor participated in the conception of the project, algorithm design, simulation code implementation, and writing of the article.

Uncertainty in Control Lyapunov Functions

Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point with equilibrium input $\hat{\mathbf{u}}_e \in \mathbb{R}^m$ for the nominal model open-loop system (3.1). Consider a local exponential CLF $V : E \rightarrow \mathbb{R}_{\geq 0}$ for the nominal model open-loop system (3.1) and equilibrium point $\mathbf{x}_e \in E$, with corresponding open set $D \subseteq E$ with $\mathbf{x}_e \in D$, and $k_3, a \in \mathbb{R}_{>0}$, such that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \hat{V}(\mathbf{x}, \mathbf{u}) \triangleq \inf_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.8)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$. The time derivative of the function V uses the true open-loop system (2.1), and thus using (3.2), we have that:

$$\dot{V}(\mathbf{x}, \mathbf{u}) = \hat{V}(\mathbf{x}, \mathbf{u}) + \underbrace{\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})\tilde{\mathbf{f}}(\mathbf{x})}_{b(\mathbf{x})} + \underbrace{\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x})\tilde{\mathbf{g}}(\mathbf{x})}_{\mathbf{a}(\mathbf{x})^\top} \mathbf{u}. \quad (3.9)$$

Consequently, choosing an input $\mathbf{u} \in \mathbb{R}^m$ such that the inequality in (3.8) is satisfied is insufficient to conclude that the time derivative of the local exponential CLF V will meet the required inequality to guarantee local exponential stability because of the functions $b : E \rightarrow \mathbb{R}$ and $\mathbf{a} : E \rightarrow \mathbb{R}^m$. Note that the functions b and \mathbf{a} capture the model errors in $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ only as they impact the evolution of the local exponential CLF V , and consequently, local exponential stability. This motivates considering a learning-based approach to identify the functions b and \mathbf{a} such that their effects on the CLF time derivative can be accounted for in the control design process, and local exponential stability can be achieved. Note that by explicitly learning the functions b and \mathbf{a} , instead

of learning their combined effect in a single function $\delta(\mathbf{x}, \mathbf{u}) = b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u}$ that takes both the state and control input as inputs, highly accurate nonlinear learning models can be used without destroying the control-affine nature of \dot{V} that permits control synthesis through convex optimization. To ensure the well-posedness of this approach, I make the following assumptions:

Assumption 1. The point $\mathbf{x}_e \in E$ is an unforced or forced equilibrium point of true open-loop system (2.1) with known equilibrium input \mathbf{u}_e .

Assumption 2. The function V is a CLF for the true open-loop system (2.1) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$ with $\mathbf{x}_e \in D$, and constants $k_3, a \in \mathbb{R}_{>0}$.

The first assumption states that trying to stabilize the true open-loop system (2.1) to the point \mathbf{x}_e is well-posed, as \mathbf{x}_e can be rendered an equilibrium point of the true open-loop system (2.1). It also states that the equilibrium input \mathbf{u}_e needed to render \mathbf{x}_e an equilibrium point of the true open-loop system (2.1) is known (it may be different than the one for the nominal model, i.e., $\hat{\mathbf{u}}_e \neq \mathbf{u}_e$). The second assumption states that V is a local exponential CLF for the true open-loop system (2.1) and equilibrium point \mathbf{x}_e with the same open set D and constants $k_3, a \in \mathbb{R}_{>0}$ used for the nominal model open-loop system (3.1). It is important to note that this is an assumption on how the true open-loop system (2.1) *can* be controlled. In particular, this assumption states that for a given $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, one knows that there *exists* an input $\mathbf{u} \in \mathbb{R}^m$ such that:

$$\dot{V}(\mathbf{x}, \mathbf{u}) = \hat{V}(\mathbf{x}, \mathbf{u}) + b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u} < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.10)$$

but not necessarily *which* value of \mathbf{u} satisfies the inequality. If one knew the functions b and \mathbf{a} perfectly, the value of \mathbf{u} satisfying this inequality could be found. Thus, if one could learn these functions from data, the control synthesis problem through CLFs would be well-posed. To better understand this assumption consider the following example of a fully-actuated robotic system.

Example 4. Consider the following nominal model for a fully-actuated robotic system:

$$\hat{\mathbf{D}}(\mathbf{q})\ddot{\mathbf{q}} + \hat{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}}) + \hat{\mathbf{G}}(\mathbf{q}) = \hat{\mathbf{B}}(\mathbf{q})\boldsymbol{\tau}, \quad (3.11)$$

where $\hat{\mathbf{D}} : \mathcal{Q} \rightarrow \mathbb{S}_{>0}^{n_c}$, $\hat{\mathbf{C}} : \mathcal{X} \rightarrow \mathbb{R}^{n_c}$, $\hat{\mathbf{G}} : \mathcal{Q} \rightarrow \mathbb{R}^{n_c}$, and $\hat{\mathbf{B}} : \mathcal{Q} \rightarrow \mathbb{R}^m$ are locally Lipschitz continuous functions on their respective domains. These functions

are nominal models of the corresponding functions \mathbf{D} , \mathbf{C} , \mathbf{G} , and \mathbf{B} in the true open-loop robotic system (2.5). Let $(\mathbf{q}_e, \mathbf{0}) \in \mathcal{X}$ be a forced equilibrium point of the nominal model open-loop robotic system (3.11) with equilibrium input:

$$\hat{\boldsymbol{\tau}}_e = \hat{\mathbf{B}}(\mathbf{q}_e)^{-1}(\hat{\mathbf{C}}(\mathbf{q}_e, \mathbf{0}_{n_c}) + \hat{\mathbf{G}}(\mathbf{q}_e)). \quad (3.12)$$

Let us suppose that the true open-loop robotic system is fully-actuated, such that $\mathbf{B}(\mathbf{q})$ is invertible for each $\mathbf{q} \in \mathcal{Q}$. Then $(\mathbf{q}_e, \mathbf{0})$ is a forced equilibrium point of the true open-loop robotic system (2.5) with equilibrium input:

$$\boldsymbol{\tau}_e = \mathbf{B}(\mathbf{q}_e)^{-1}(\mathbf{C}(\mathbf{q}_e, \mathbf{0}_{n_c}) + \mathbf{G}(\mathbf{q}_e)). \quad (3.13)$$

Let $\boldsymbol{\eta} = [(\mathbf{q} - \mathbf{q}_e)^\top \quad \dot{\mathbf{q}}^\top]^\top$. A feedback linearizing controller $\hat{\mathbf{k}}_{\text{fbl}} : \mathcal{X} \times \mathbb{R}^{n_c} \rightarrow \mathbb{R}^m$ using the nominal model open-loop robotic system (3.11) can be designed:

$$\hat{\mathbf{k}}_{\text{fbl}}(\mathbf{x}, -\mathbf{K}\boldsymbol{\eta}) = \hat{\mathbf{B}}(\mathbf{q})^{-1}\hat{\mathbf{D}}(\mathbf{q})(\hat{\mathbf{D}}(\mathbf{q})^{-1}(\hat{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}}) + \hat{\mathbf{G}}(\mathbf{q})) - \mathbf{K}\boldsymbol{\eta}), \quad (3.14)$$

with $\mathbf{K} \in \mathbb{R}^{n_c \times n}$. For the nominal model open-loop robotic system (3.11), this yields the linear closed-loop output dynamics:

$$\dot{\boldsymbol{\eta}} = \mathbf{A}_{\text{cl}}\boldsymbol{\eta}, \quad (3.15)$$

which are locally exponentially stable for an appropriate choice of \mathbf{K} . As in Section 2.4, for a given $\mathbf{Q} \in \mathbb{S}_{\zeta_0}^n$, the CTLE can be solved to produce a matrix $\mathbf{P} \in \mathbb{S}_{\zeta_0}^n$ that defines a local exponential CLF V for the nominal model open-loop robotic system (3.11) and equilibrium point $(\mathbf{q}_e, \mathbf{0})$. Considering the true open-loop robotic system (2.5), a feedback linearizing controller $\mathbf{k}_{\text{fbl}} : \mathcal{X} \times \mathbb{R}^{n_c} \rightarrow \mathbb{R}^m$ can be specified as:

$$\mathbf{k}_{\text{fbl}}(\mathbf{x}, -\mathbf{K}\boldsymbol{\eta}) = \mathbf{B}(\mathbf{q})^{-1}\mathbf{D}(\mathbf{q})(\mathbf{D}(\mathbf{q})^{-1}(\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q})) - \mathbf{K}\boldsymbol{\eta}). \quad (3.16)$$

Similarly, for the true open-loop robotic system (2.5), this controller leads to the same closed-loop linear output dynamics in (3.15), for which the CTLE produces the same matrix $\mathbf{P} \in \mathbb{S}_{\zeta_0}^n$ that defined the local exponential CLF V . While the controller (3.16) can not be evaluated as the functions \mathbf{D} , \mathbf{C} , \mathbf{G} , and \mathbf{B} are not precisely known, the input it would return exists, and thus it is sufficient to conclude that V is a local exponential CLF for the true open-loop robotic system (2.5) and equilibrium point $(\mathbf{q}_e, \mathbf{0})$. Thus, the assumption of V being a local exponential CLF for the true open-loop robotic system (2.5) is justified under other structural assumptions such as the robotic system being fully-actuated. This ensures that successfully learning the functions b and \mathbf{a} will lead to a well-posed control problem.

Data-Driven Episodic Learning Framework

Given that the formulation in (3.2) and (3.9) defines a general class of model error (rather than parametric), it is natural to consider a data-driven method to estimate the unknown functions b and \mathbf{a} on the domain D . To motivate a learning-based framework, first consider a simple approach of learning b and \mathbf{a} via supervised regression [204]: the true closed-loop system (2.2) is operated using a given controller $\mathbf{k} : D \rightarrow \mathbb{R}^m$, data points are gathered along the system's evolution, and functions that approximate b and \mathbf{a} are produced via supervised learning.

Concretely, let $\mathbf{x}_0 \in D$ be an initial state. An experiment is defined as the evolution of the true closed-loop system over a finite time interval from the initial condition \mathbf{x}_0 using a sample-and-hold implementation of the controller \mathbf{k} . A resulting discrete-time state history is obtained, which is then transformed with the local exponential CLF V and finally differentiated numerically to estimate \dot{V} throughout the experiment. This yields a data set composed of input-output pairs:

$$\mathfrak{D} = \{((\mathbf{x}_i, \mathbf{u}_i), \dot{V}_i)\}_{i=1}^N \subseteq (D \times \mathbb{R}^m) \times \mathbb{R}. \quad (3.17)$$

Consider a class \mathcal{H}_b of nonlinear functions mapping \mathbb{R}^n to \mathbb{R} and a class \mathcal{H}_a of nonlinear functions mapping \mathbb{R}^n to \mathbb{R}^m . For a given $\hat{b} \in \mathcal{H}_b$ and $\hat{\mathbf{a}} \in \mathcal{H}_a$, define the structured estimator \hat{W} as:

$$\hat{W}(\mathbf{x}, \mathbf{u}) = \hat{V}(\mathbf{x}, \mathbf{u}) + \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}, \quad (3.18)$$

and let \mathcal{H} be the class of all such estimators mapping $\mathbb{R}^n \times \mathbb{R}^m$ to \mathbb{R} . Observe that the structure of this estimator ensures it is affine in the control input, and thus can be directly integrated into the convex optimization-based control framework presented in Section 2.3. Defining a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, the supervised regression task is then to find a function in \mathcal{H} via empirical risk minimization (ERM):

$$\inf_{\substack{\hat{\mathbf{a}} \in \mathcal{H}_a \\ \hat{b} \in \mathcal{H}_b}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{W}(\mathbf{x}_i, \mathbf{u}_i), \dot{V}_i). \quad (3.19)$$

This experiment protocol can be executed either in simulation or directly on hardware. While being simple to implement, supervised learning critically assumes independently and identically distributed (i.i.d) training data. Each

experiment violates this assumption, as the regression target of each data point is coupled with the input data of the previous time step. As a consequence, standard supervised learning with sequential, non-i.i.d data collection often leads to error cascades [142].

Episodic learning refers to learning procedures that iteratively alternate between executing an intermediate controller (also known as a roll-out in reinforcement learning [132]) to collect data from that roll-out, and designing a new controller using the newly collected data. My approach integrates learning b and \mathbf{a} with improving the performance and stability of the controller in such an iterative fashion. First, assume we are given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$, which may not locally exponentially stabilize the resulting true closed-loop system (2.2) to the equilibrium point \mathbf{x}_e , but satisfies $\mathbf{k}_{\text{nom}}(\mathbf{x}_e) = \mathbf{u}_e$. With an estimator $\widehat{W} \in \mathcal{H}$ defined in (3.18), specify a learning-informed (CLF-QP) controller as:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & (3.20) \\ \text{s.t. } \widehat{W}(\mathbf{x}, \mathbf{u}) &\leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \end{aligned}$$

This learning-informed (CLF-QP) controller finds the closest input to $\mathbf{k}_{\text{nom}}(\mathbf{x})$ such that the local exponential CLF inequality is met using the estimator \widehat{W} for \dot{V} . However, error in the estimator \widehat{W} degrades the local exponential stability of the true closed-loop system (2.2) with respect to the equilibrium point \mathbf{x}_e .

In an effort to reduce the remaining error in the estimator \widehat{W} , an experiment can be run using the learning-informed (CLF-QP) controller to produce data by which to obtain better estimates of b and \mathbf{a} . To overcome the limitations in needing i.i.d data for conventional supervised learning, I leverage reduction techniques: a sequential prediction problem is reduced to a sequence of supervised learning problems over multiple episodes [203], [205]. In particular, in each episode, an experiment generates data using a different controller. The data set is aggregated and a new ERM problem is solved after each episode. This episodic learning implementation is inspired by the Data Aggregation (Dagger) algorithm [203], with some key differences:

- Dagger is a model-free policy learning algorithm, which trains a policy directly in each episode using optimal control input oracles. As such

an oracle is not available in my problem setting, my algorithm defines a controller indirectly via a local exponential CLF as in (3.20) to explicitly drive the system towards stable behavior.

- The ERM problem is underdetermined, i.e., different approximations $(\hat{b}, \hat{\mathbf{a}})$ may achieve similar loss for a given data set while failing to accurately model b and \mathbf{a} . This issue will be explored more in Sections 3.7 and 3.8. This potentially introduces error in estimating \dot{V} for control inputs not reflected in the training data, and necessitates the use of exploratory control action to constrain the estimators \hat{b} and $\hat{\mathbf{a}}$. Such exploration can be achieved by randomly perturbing the controller used in an experiment at each time step. This need for exploration is an analog to the notion of persistent excitation for adaptive systems [206].

Algorithm 1 specifies a method of computing a sequence of local exponential CLF derivative estimates \widehat{W}_k and augmented¹ controllers \mathbf{k}_k . During the k^{th} episode, the $k - 1^{\text{th}}$ augmented controller is used in an experiment, data is collected and aggregated with the data set \mathfrak{D} and an estimate of the local exponential CLF derivative \widehat{W}_k is formed. After this, a corresponding learning-informed (CLF-QP) controller as in (3.20) is produced, scaled by a heuristically chosen factor $w_k \in [0, 1]$ reflecting trust in the estimate and added to the nominal controller $\mathbf{k}_0 = \mathbf{k}_{\text{nom}}$ for use in the subsequent experiment as an augmented controller \mathbf{k}_k . The trust coefficients form a monotonically non-decreasing sequence. Importantly, this experiment need not take place in simulation; the same procedure may be executed directly on hardware. At a high level, this episodic approach makes progress by gathering more data in relevant regions of the state space, such as states close to \mathbf{x}_e . This extends the generalizability of the estimator in its use by subsequent controllers.

There are a few practical concerns with this algorithm that should be observed. First, the system will not necessarily be locally exponentially stable during the episodic learning process. In general, one will often begin the learning process with a controller that is unstable, or displays desirable behavior but possess no guarantees of local exponential stability (such as a Lyapunov function). This algorithm instead tolerates instability during the episodic learning process to

¹I refer to these as *augmented* controllers if $w_k \neq 1$, as the learning-informed (CLF-QP) controller in (3.20) is not fully used, but blended with the nominal controller \mathbf{k}_0 .

Algorithm 1 Dataset Aggregation for Control Lyapunov Functions (Da-CLyF)

Require: Local exponential CLF V , derivative estimate $\widehat{W}_0 = \widehat{V}$, model classes \mathcal{H}_a and \mathcal{H}_b , loss function \mathcal{L} , nominal controller $\mathbf{k}_0 = \mathbf{k}_{\text{nom}}$, number of experiments T , sequence of trust coefficients $0 \leq w_1 \leq \dots \leq w_T \leq 1$.

```

 $\mathcal{D} = \emptyset$  ▷ Initialize data set
for  $k = 1, \dots, T$  do
   $\mathbf{x}_0 \leftarrow \text{sample } D$  ▷ Get initial condition
   $\mathcal{D}_k \leftarrow \text{experiment}(\mathbf{x}_0, \mathbf{k}_{k-1})$  ▷ Run experiment
   $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_k$  ▷ Aggregate data set
   $\widehat{\mathbf{a}}_k, \widehat{b}_k \leftarrow \text{ERM}(\mathcal{H}_a, \mathcal{H}_b, \mathcal{L}, \mathcal{D}, \widehat{W}_0)$  ▷ Fit estimators
   $\widehat{W}_k \leftarrow \widehat{W}_0 + \widehat{\mathbf{a}}_k^\top \mathbf{u} + \widehat{b}_k$  ▷ Update derivative estimator
   $\mathbf{k}_k \leftarrow (1 - w_k) \cdot \mathbf{k}_0 + w_k \cdot \text{augment}(\mathbf{k}_0, \widehat{W}_k)$  ▷ Update controller
end for
return  $\widehat{W}_T, \mathbf{k}_T$ 

```

ensure that the system performs adequate exploration in the data acquisition process and return a controller at the end of the episodic learning for which meaningful stability properties can be established, as explored in Section 3.4.

Second, there can be issues that arise with the constraint in the learning-informed (CLF-QP) controller (3.20). In particular, while we know that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \widehat{V}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.21)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, and we assume that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.22)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, such that the local exponential CLF constraint can be met for the nominal model open-loop system (3.1) and the true open-loop system (2.1), we do not know that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \widehat{W}_k(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.23)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$ for $k = 1, \dots, T$. In particular, the learned models \widehat{b} and $\widehat{\mathbf{a}}$ can lead to this condition being violated. This means that the learning-informed (CLF-QP) controller will not necessarily be feasible for all $\mathbf{x} \in D$. This issue is noticeable at the equilibrium point \mathbf{x}_e , as we must have:

$$\widehat{b}_k(\mathbf{x}_e) + \widehat{\mathbf{a}}_k(\mathbf{x}_e)^\top \mathbf{u}_e \leq 0, \quad (3.24)$$

or the equilibrium input will not satisfy $\widehat{W}_k(\mathbf{x}_e, \mathbf{u}_e) \leq 0$ (as $\widehat{V}(\mathbf{x}_e, \mathbf{u}_e) = 0$), and thus the controller will not use the equilibrium input at the equilibrium point. In practice, the relaxation technique used in the (CLF-CBF-QP) controller is utilized, such that the problem always remains feasible at the expense of rigorous stability guarantees. This highlights a particular challenge in learning and control, which is the analytic nature of the stability and safety guarantees that come from CLFs and CBFs, and the numerical nature of learned models. In reality, one will likely never be able to learn the functions b and \mathbf{a} perfectly. Instead, stability and safety guarantees will need to be built keeping in mind that residual learning error arising from numerical learning methods will be present in learning models, as explored in Sections 3.4 and 3.6.

Simulation Results

I apply this episodic learning algorithm to a planar Segway platform seen in Figure 3.1. In particular, I consider a 4-dimensional planar Segway model based on the simulation model in [27]. The system states consist of horizontal position $x \in \mathbb{R}$, horizontal velocity $\dot{x} \in \mathbb{R}$, pitch angle $\theta \in \mathbb{R}$, and pitch angle rate $\dot{\theta} \in \mathbb{R}$. Control is specified as a single voltage input supplied to both wheel motors. The parameters of the model (including mass, inertias, and motor parameters but excluding gravity) are randomly modified by up to 10% of their nominal values and then held fixed for the simulations. Note that even though this is parametric model error, it introduces error between $\widehat{\mathbf{g}}$ and \mathbf{g} , which is typically not considered by adaptive control approaches.

I seek to stabilize the pitch angle and pitch angle rate to a trajectory generated using the GPOPS-II Optimal Control Software [207] for the nominal model². The nominal controller $k_0 = k_{\text{nom}}$ is a linear proportional-derivative (PD) controller on angle and angle rate error. A set of $T = 20$ experiments are conducted with trust values w_k varying from 0.01 to 0.99 in a sigmoid fashion. The exploratory control is drawn uniformly at random between -20% and 20% of the norm of the control input specified by the augmented controller k_k for

²More precisely, the definition of local exponential stability in Definition 2 can easily be extended to use a desired trajectory $\mathbf{x}_d : \mathbb{R}_{\geq 0} \rightarrow E$ that satisfies the differential equation (2.1) for some input signal $\mathbf{u}_d : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ instead of an equilibrium point \mathbf{x}_e . Furthermore, the Segway system is underactuated, and thus the CLF considered is the one for the outputs (the pitch angle tracking error and pitch angle rate tracking error) (2.36) produced using feedback linearization. The goal is thus to drive the outputs to $\mathbf{0}_2$, but the stability of the zero dynamics (and full system stability) are not considered.

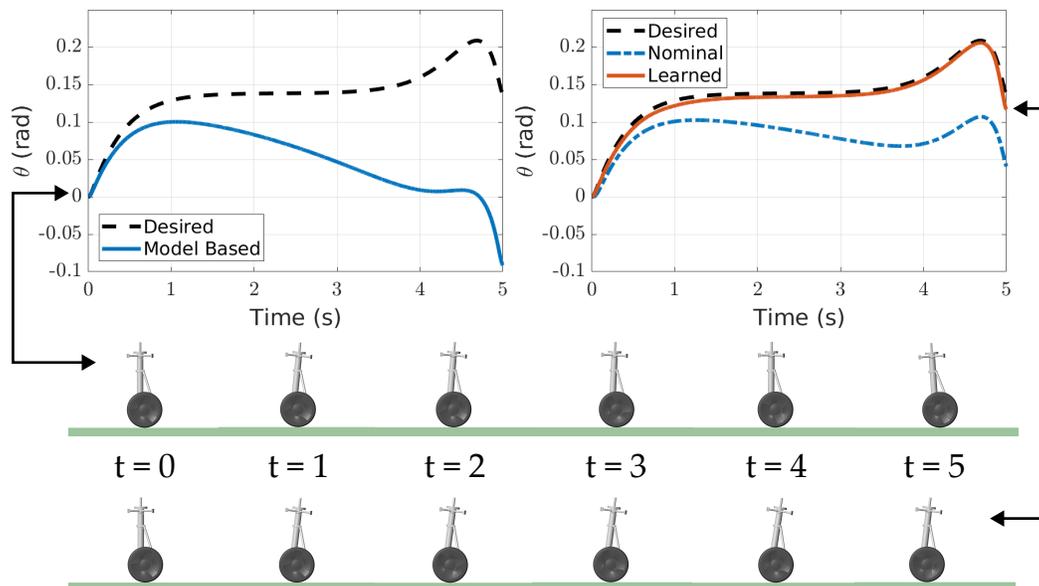


Figure 3.1. (Left) Nominal model-based (CLF-QP) controller fails to stabilize the system to the trajectory. (Right) Improvement in angle tracking of system with learning-informed (CLF-QP) controller over nominal PD controller. (Bottom) Corresponding visualizations of state data. Video of this simulation is found at https://youtu.be/cB5MY_8vCrQ.

the first ten episodes. The percentages decay linearly to zero in the remaining ten episodes. The model classes selected are sets of two-layer neural networks with Rectified Linear Unit (ReLU) nonlinearities with a hidden layer width of 2000 nodes that are implemented in Keras [208].

Failure of the nominal model-based (CLF-QP) controller is seen in the left panel of Figure 3.1. The nominal PD controller $k_0 = k_{\text{nom}}$ and the final learning-informed (CLF-QP) controller k_{20} produced using $w_{20} = 1$ can be seen in the right panel of Figure 3.1. Corresponding visualizations of the Segway states are displayed in the bottom panel of Figure 3.1. The final learning-informed (CLF-QP) controller exhibits a notable improvement over the nominal model-based (CLF-QP) controller and nominal PD controller in its ability to track the desired trajectory. To verify the robustness of the learning algorithm, the 20 experiment process was conducted ten times. After each experiment the augmented controller k_k was tested without exploratory perturbations. For the last three experiments and a test of the final learning-informed (CLF-QP) controller k_{20} , the minimum, mean, and maximum angles across all ten instances are displayed for each time step in Figure 3.2. The mean trajectory consistently improves in these later episodes as the trust fac-

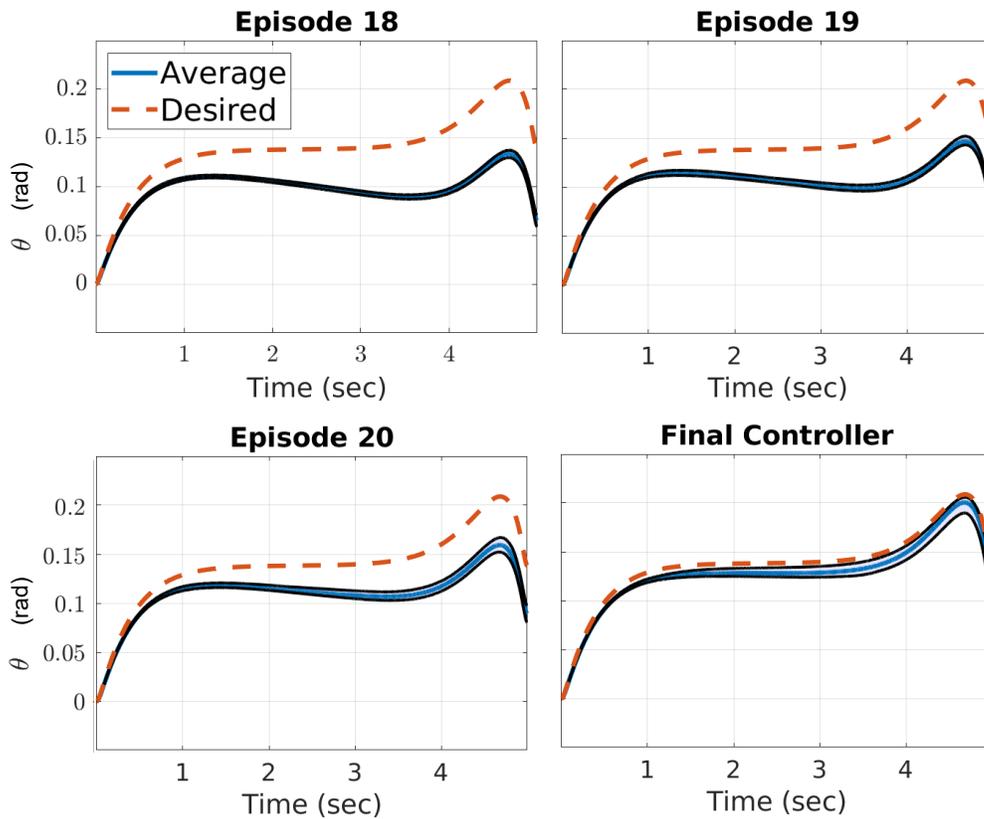


Figure 3.2. Augmented controllers consistently improve stabilization to trajectory across episodes. Ten instances of the algorithm are executed with the shaded region formed from minimum and maximum angles for each time step within an episode. The corresponding average angle trajectories are also displayed.

tor increases. The variation increases but remains small, indicating that the learning problem is robust to randomness in the initialization of the neural networks, in the network training algorithm, and in the noise added during the experiments. The performance of the controller in the earlier episodes displayed negligible variation from the nominal PD controller $k_0 = k_{\text{nom}}$ due to small trust factors.

Conclusion

In this section I have presented an episodic learning framework for learning model errors between a nominal model open-loop system and the true open-loop system as they directly impact the time derivative of a local exponential CLF V . I provide a discussion on assumptions that ensure a local exponential CLF V for the nominal model open-loop system is a local exponential CLF for the true open-loop system, and provide justification of these assumptions

through an example in the context of a fully actuated robotic system. I then describe the episodic learning framework, and address practical concerns that arise when integrating the analytic tools of local exponential CLFs with the numerical tools of learning models. Lastly I provide a demonstration of the efficacy of the algorithm in simulation on a planar Segway system.

3.4 Projection-to-State Stability

In this section I will present an analysis of the effects of residual learning error on the local exponential stability of the true closed-loop system (2.2) using a learning-informed (CLF-QP) controller. Building off of the episodic learning algorithm in Section 3.3, I will consider learning model errors as they impact the time derivative of a local exponential CLF. I will first demonstrate that the inherent robustness properties of local exponential CLF-based controllers yield a general robustness to residual learning errors in this learning and control setting. This will be achieved by looking at the evolution of a local exponential CLF as a disturbed dynamic system through a modification of ISS known as Projection-to-State Stability (PSS). Second, I will consider the possible residual learning errors that can occur given a data set \mathcal{D} and estimator \widehat{W} as in Section 3.3, and what this will mean for the behavior of the true closed-loop system (2.2) using an imperfect learning-informed controller.

The contributions of this section are as follows:

- A notion of robust stability in the form of *Projection-to-State Stability (PSS)*, which captures the impact of disturbances on stability only as they impact the time derivative of a CLF, rather than the full-order dynamics, as is done by ISS discussed in Section 2.7.
- A data-driven characterization of possible residual learning errors when learning model errors as they impact the time derivative of a local exponential CLF. The convex representations that arise from this characterization are the foundation of the work in Section 3.8.

The text for this section is adapted from:

A. J. Taylor, V. D. Dorobantu, M. Krishnamoorthy, H. M. Le, Y. Yue, and A. D. Ames, “A control lyapunov perspective on episodic learning

via projection to state stability,” in *Proc. IEEE 58th Conf. on Decision and Control (CDC)*, Nice, France, 2019, pp.1448-1455.

A. J. Taylor participated in the conception of the project, theoretical analysis, simulation code implementation, and writing of the article.

Projection-to-State Stability and Residual Learning Error

Let $\mathbf{x}_e \in E$ be an equilibrium point of the nominal model open-loop system (3.1) and the true open-loop system (2.1) with known equilibrium inputs $\hat{\mathbf{u}}_e$ and \mathbf{u}_e , respectively. Suppose that $V : E \rightarrow \mathbb{R}_{\geq 0}$ is a local exponential CLF for the nominal model open-loop system (3.1) and true open-loop system (2.1) and equilibrium point \mathbf{x}_e with corresponding open set $D \subseteq E$ with $\mathbf{x}_e \in D$. This amounts to the requirements of Assumptions 1 and 2 made in Section 3.3. Furthermore, consider the nominal model-based (CLF-QP) controller:

$$\begin{aligned} \hat{\mathbf{k}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 \\ &\text{s.t. } \hat{V}(\mathbf{x}, \mathbf{u}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \end{aligned} \quad (3.25)$$

This controller corresponds to the (CLF-QP) controller using the nominal model open-loop system (3.1) for the time derivative of the local exponential CLF V . The true closed-loop system using the controller $\hat{\mathbf{k}}$ may be written as:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\hat{\mathbf{k}}(\mathbf{x}) = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\hat{\mathbf{k}}(\mathbf{x}) + \tilde{\mathbf{f}}(\mathbf{x}) + \tilde{\mathbf{g}}(\mathbf{x})\hat{\mathbf{k}}(\mathbf{x}). \quad (3.26)$$

For a given initial condition $\mathbf{x}_0 \in D$, denote the solution to the true closed-loop system (3.26) as $\varphi : I(\mathbf{x}_0) \rightarrow D$. Alternatively, this system can be understood to evolve as a nominal model closed-loop system with unmatched disturbances:

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\hat{\mathbf{k}}(\mathbf{x}) + \mathbf{d}(t), \quad (3.27)$$

where for any initial condition $\mathbf{x}_0 \in D$ the disturbance signal is defined as:

$$\mathbf{d}(t) = \tilde{\mathbf{f}}(\varphi(t)) + \tilde{\mathbf{g}}(\varphi(t))\hat{\mathbf{k}}(\varphi(t)). \quad (3.28)$$

Viewing model error as an unmatched disturbance in the nominal model closed-loop system motivates studying it through the lens of Input-to-State Stability (ISS). If the nominal model closed-loop system with unmatched disturbances (3.27) is ISS with respect to the equilibrium point \mathbf{x}_e , then small model errors will lead to small disturbances and a small degradation in the

local exponential stability that would be demonstrated by the nominal model closed-loop system without disturbances:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\widehat{\mathbf{k}}(\mathbf{x}). \quad (3.29)$$

In the context of learning the time derivative of a local exponential CLF as in Section 3.3, the disturbance signal is not as readily understood in the state dynamics. In particular, recall that:

$$\dot{V}(\mathbf{x}, \mathbf{u}) = \widehat{V}(\mathbf{x}, \mathbf{u}) + b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u}. \quad (3.30)$$

Let $\widehat{b} \in \mathcal{H}_b$ and $\widehat{\mathbf{a}} \in \mathcal{H}_a$ compose the estimator $\widehat{W} \in \mathcal{H}$ as:

$$\widehat{W}(\mathbf{x}, \mathbf{u}) = \widehat{V}(\mathbf{x}, \mathbf{u}) + \widehat{b}(\mathbf{x}) + \widehat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}, \quad (3.31)$$

and let us assume that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \widehat{W}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.32)$$

for all $\mathbf{x} \in D \setminus \{\mathbf{x}_e\}$, and³ that $\widehat{W}(\mathbf{x}_e, \mathbf{u}_e) \leq 0$. Adding and subtracting the estimators \widehat{b} and $\widehat{\mathbf{a}}$ from \dot{V} yields:

$$\dot{V}(\mathbf{x}, \mathbf{u}) = \widehat{W}(\mathbf{x}, \mathbf{u}) + \underbrace{b(\mathbf{x}) - \widehat{b}(\mathbf{x})}_{\widetilde{b}(\mathbf{x})} + \underbrace{(\mathbf{a}(\mathbf{x}) - \widehat{\mathbf{a}}(\mathbf{x}))^\top}_{\widetilde{\mathbf{a}}(\mathbf{x})^\top} \mathbf{u}. \quad (3.33)$$

The functions $\widetilde{b} : E \rightarrow \mathbb{R}$ and $\widetilde{\mathbf{a}} : E \rightarrow \mathbb{R}^m$ capture *residual learning error* left after trying to learn the functions b and \mathbf{a} with the estimators \widehat{b} and $\widehat{\mathbf{a}}$. As these error terms appear in the time derivative of the local exponential CLF V , they will consequently have an impact on the stability properties of the system. Using the learning-informed (CLF-QP) controller:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 \\ &\text{s.t. } \widehat{W}(\mathbf{x}, \mathbf{u}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \end{aligned} \quad (3.34)$$

³This assumption will lead the learning-informed controller in (3.34) to satisfy $\mathbf{k}(\mathbf{x}_e) = \mathbf{u}_e$, such that \mathbf{x}_e is an equilibrium point of the true closed-loop system (2.2). This assumption can often practically be enforced at the expense of additional residual learning error near the equilibrium point \mathbf{x}_e by providing “artificial” training data specifying that $\widehat{W}(\mathbf{x}_e, \mathbf{u}_e) = -\epsilon$ for some $\epsilon \in \mathbb{R}_{>0}$. A consequence of continuous learning models is that this artificial modification will lead to $\widehat{W}(\mathbf{x}, \mathbf{u}_e) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a$ for all \mathbf{x} in a neighborhood of the equilibrium point \mathbf{x}_e , such that $\mathbf{k}(\mathbf{x}) = \mathbf{u}_e$, even if $\dot{V}(\mathbf{x}, \mathbf{u}_e) > -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a$ for some values of \mathbf{x} in that neighborhood. Thus, local exponential stability will be compromised. As we are interested in studying the effects of residual learning errors as disturbances in a dynamic system, we will only really be concerned with convergence to a neighborhood of the origin as captured by ISS, making the loss of local exponential stability tolerable.

yields:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) = \widehat{W}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \widetilde{b}(\mathbf{x}) + \widetilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}). \quad (3.35)$$

For any initial condition $\mathbf{x}_0 \in D$, let $\varphi : I(\mathbf{x}_0) \rightarrow D$ denote the solution to the true closed-loop system (2.2) using the learning-informed (CLF-QP) controller (3.34). The evolution of (3.35) can be viewed as a closed-loop system with unmatched disturbances:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) = \widehat{W}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \delta(t), \quad (3.36)$$

where for any initial condition $\mathbf{x}_0 \in D$ the function $\delta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ defined as:

$$\delta(t) = \widetilde{b}(\varphi(t)) + \widetilde{\mathbf{a}}(\varphi(t))^\top \mathbf{k}(\varphi(t)). \quad (3.37)$$

is a *projected disturbance*. As with model error in the state dynamics, this motivates studying residual learning error through the lens of ISS. Observe that the learning-informed controller enforces:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a + |\delta(t)|, \quad (3.38)$$

$$\leq -\frac{k_3}{k_2} V(\mathbf{x}) + |\delta(t)|. \quad (3.39)$$

Consequently, we have that the identity map $\iota : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ defined as $\iota(r) = r$ for all $r \in \mathbb{R}_{\geq 0}$ satisfies:

$$|V(\mathbf{x})| \leq \iota(V(\mathbf{x})) \leq |V(\mathbf{x})|, \quad (3.40)$$

$$\frac{\partial \iota}{\partial r}(V(\mathbf{x})) \dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -\frac{k_3}{k_2} V(\mathbf{x}) + |\delta(t)| = -\frac{k_3}{k_2} |V(\mathbf{x})| + |\delta(t)|, \quad (3.41)$$

and thus is a local exponential ISS-LF for the closed-loop system with unmatched disturbances (3.36) with respect to 0. Thus, there exist an open set $D' \subseteq D$ with $\mathbf{x}_e \in D'$, constants $M, \lambda, \bar{d} \in \mathbb{R}_{> 0}$, and a function $\gamma \in \mathcal{K}$ such that for any initial condition $\mathbf{x}_0 \in D'$ and $d \in [0, \bar{d}]$, if $\|\delta\|_\infty \leq d$, then the solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow D$ to the true closed-loop system (2.2) satisfies:

$$V(\varphi(t)) \leq MV(\mathbf{x}_0)e^{-\lambda t} + \gamma(d), \quad (3.42)$$

for all $t \in \mathbb{R}_{\geq 0}$. Using the bounds on V in (2.16), we have that:

$$\|\varphi(t) - \mathbf{x}_e\|^a \leq \frac{Mk_2}{k_1} \|\mathbf{x}_0 - \mathbf{x}_e\|^a e^{-\lambda t} + \frac{1}{k_1} \gamma(d). \quad (3.43)$$

Using the weak class- \mathcal{K} triangle inequality [56, Equation 12], we have that:

$$\|\varphi(t) - \mathbf{x}_e\| \leq \sqrt[a]{\frac{2Mk_2}{k_1}} \|\mathbf{x}_0 - \mathbf{x}_e\| e^{-\frac{\lambda}{a}t} + \left(\frac{2}{k_1}\gamma(d)\right)^{\frac{1}{a}}, \quad (3.44)$$

$$\triangleq M' \|\mathbf{x}_0 - \mathbf{x}_e\| e^{-\lambda't} + \gamma'(d), \quad (3.45)$$

where $M' = \sqrt[a]{\frac{2Mk_2}{k_1}}$, $\lambda' = \frac{\lambda}{a}$, and $\gamma'(r) = \left(\frac{2}{k_1}\gamma(r)\right)^{\frac{1}{a}}$. This bound appears like a statement of ISS for the true closed-loop system (2.2) using the learning-informed controller (3.34), but it is defined now in terms of a bound on the projected disturbance signal δ rather than a disturbance signal \mathbf{d} in the state dynamics. Consequently, the true closed-loop system (2.2) with the learning-informed controller (3.34) is said to be *Projection-to-State Stable (PSS)* with respect to the equilibrium point \mathbf{x}_e and projection V . The utility of PSS is that it considers the impacts of model error only as it impacts the time derivative of the local exponential CLF. By doing so, it provides a mathematical framework for understanding the impact of introducing learning models for the time derivative of a local exponential CLF, as it is not obvious how to represent residual learning error in learning models for a CLF time derivative as a disturbance in the full state dynamics (2.2), making an analysis by standard ISS difficult. Furthermore, it may potentially produce less conservative statements on the degradation of local exponential stability by neglecting model error that does not affect stability.

Quantifying Residual Learning Error

In the preceding construction, I use a local exponential ISS-LF to establish the existence of *some* open set $D' \subseteq D$ with $\mathbf{x}_e \in D'$, constants $M, \lambda, \bar{d} \in \mathbb{R}_{>0}$, and function $\gamma \in \mathcal{K}$ such that if $\|\delta\|_\infty \leq \bar{d}$, then conclusions can be drawn about the PSS properties of the true closed-loop system (2.2). Recalling that given a data set \mathfrak{D} and corresponding estimators \hat{b} and $\hat{\mathbf{a}}$, the system experiences the particular projected disturbance signal in (3.37), this in and of itself is not a particularly useful result if we do not actually know that $\|\delta\|_\infty \leq \bar{d}$ for all initial conditions $\mathbf{x}_0 \in D'$. Consequently, an important question is whether or not we can find an open set $D' \subseteq D$ with $\mathbf{x}_e \in D'$, constants $M, \lambda, \bar{d} \in \mathbb{R}_{>0}$, and a function $\gamma \in \mathcal{K}$ satisfying the relationship in (3.42) for all $t \in \mathbb{R}_{\geq 0}$ and $d \in [0, \bar{d}]$, and for which we know that the particular disturbance signal in (3.37) at least satisfies $\|\delta\|_\infty \leq \bar{d}$ for all initial conditions $\mathbf{x}_0 \in D'$. This task can be reduced to finding just an open set $D' \subseteq D$ with $\mathbf{x}_e \in D'$ and a constant

$\bar{d} \in \mathbb{R}_{>0}$ such that for any initial condition $\mathbf{x}_0 \in D'$, the solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow D$ to the true closed-loop system (2.2) is defined for all time and the projected disturbance defined in (3.37) satisfies $|\delta(t)| \leq \bar{d}$ for all $t \in \mathbb{R}_{\geq 0}$. With these, the constants M, λ and function γ can be established using the standard proof showing ISS through an ISS-LF [56, Theorem 1].

While these concerns seem technical, the concrete concern in the context of learning is if the residual learning error near the equilibrium point \mathbf{x}_e is large, it may cause the system to evolve into states not present in the data set \mathfrak{D} , leading to even more residual learning error and a larger projected disturbance, leading to further deviation from \mathbf{x}_e and repetition of this cycle. This motivates using a data-driven approach to quantify the projected disturbance signal with the goal of finding the desired open set D' and constant $\bar{d} \in \mathbb{R}_{>0}$, for which we know that for all initial conditions $\mathbf{x}_0 \in D'$, the solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow D$ to the true closed-loop system (2.2) is defined for all time and the projected disturbance defined in (3.37) satisfies $|\delta(t)| \leq \bar{d}$ for all $t \in \mathbb{R}_{\geq 0}$. To this end, I will first describe the notion of a set-valued *uncertainty function*, and show how such a function can be used to produce a set D' and constant \bar{d} . I will then provide a data-driven construction of such an uncertainty function.

Let $\mathcal{P}(\mathbb{R} \times \mathbb{R}^m)$ denote the powerset of $\mathbb{R} \times \mathbb{R}^m$, and define the following:

Definition 22 (*Uncertainty Function*). A function $\Delta : D \rightarrow \mathcal{P}(\mathbb{R} \times \mathbb{R}^m)$ is an *uncertainty function* if $(\tilde{b}(\mathbf{x}), \tilde{\mathbf{a}}(\mathbf{x})) \in \Delta(\mathbf{x})$ for all $\mathbf{x} \in D$.

It is straightforward to see that given an uncertainty function Δ , we have that:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq \widehat{W}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} (b + \mathbf{a}^\top \mathbf{k}(\mathbf{x})), \quad (3.46)$$

for all $\mathbf{x} \in D$. To see how this can be used, let us denote the c -sublevel set of the local exponential CLF V as:

$$\Omega(c) = \{\mathbf{x} \in E \mid V(\mathbf{x}) \leq c\}. \quad (3.47)$$

For a given $c \in \mathbb{R}_{>0}$, I will denote the closure of $\Omega(c)$ in \mathbb{R}^n by $\bar{\Omega}(c)$. For simplicity, I will assume that there exists some $\bar{c} \in \mathbb{R}_{>0}$ such that $\bar{\Omega}(c) \subset E$ for all $c \in [0, \bar{c})$. This is a mild technical assumption that can generally be shown to hold with large values of \bar{c} for commonly used local exponential CLFs using the lower bound in (2.12). This leads to the following result:

Theorem 18. *If a value $c \in (0, \bar{c})$ satisfies:*

$$\sup_{\mathbf{x} \in \Omega(c)} \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}) \leq \frac{k_3}{k_2} c, \quad (3.48)$$

and $\Omega(c) \subset D$, then for the set $D' = \text{Int}(\Omega(c))$ and constant $\bar{d} = \frac{k_3}{k_2} c$, we have that for any initial condition $\mathbf{x}_0 \in D'$, the solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow D$ to the true closed-loop system (2.2) exists for all time and the projected disturbance in (3.37) satisfies $|\delta(t)| \leq \bar{d}$ for all $t \in \mathbb{R}_{\geq 0}$.

This theorem states that if the worst-case residual learning error permitted by an uncertainty set Δ satisfies a bound over a sublevel set of the local exponential CLF V , and the sublevel set is contained in the domain D where the time derivative of the local exponential CLF V can be bounded, we can produce the desired open set D' and constant \bar{d} .

Proof. Note that for any $c \in (0, \bar{c})$, the continuity of V yields that $\Omega(c)$ is closed ($\Omega(c) = \bar{\Omega}(c)$), and thus $\partial\Omega(c) \subset \Omega(c)$. Additionally, the lower bound in (2.12) ensures that the set $\Omega(c)$ is bounded, such that $\Omega(c)$ is compact.

Consider an $\mathbf{x} \in \partial\Omega(c)$, noting that $\mathbf{x} \in \Omega(c)$ and thus $\mathbf{x} \in D$. First, observe that $V(\mathbf{x}) = c$. To see this, assume that $V(\mathbf{x}) < c$ ($\mathbf{x} \notin \Omega(c)$ if $V(\mathbf{x}) > c$). By the continuity of V on E , there exists an open set U with $\mathbf{x} \in U$ such that $V(\mathbf{y}) < c$ for all $\mathbf{y} \in U$, and thus $U \subset \Omega(c)$. Thus $\mathbf{x} \in \text{Int}(\Omega(c))$, and thus $\mathbf{x} \notin \partial\Omega(c)$, yielding a contradiction. Next observe that $\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$. To see this, observe that $\mathbf{x} \neq \mathbf{x}_e$, as $V(\mathbf{x}) = c > 0$. As we have that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \hat{V}(\mathbf{x}, \mathbf{u}) < -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a < 0, \quad (3.49)$$

for all $\mathbf{x} \in D$, we must have that $\frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$, otherwise the left-hand side would be 0. Now, we have that:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a + \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}), \quad (3.50)$$

$$\leq -\frac{k_3}{k_2} V(\mathbf{x}) + \sup_{\mathbf{y} \in \Omega(c)} \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{y})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{y}), \quad (3.51)$$

$$\leq -\frac{k_3}{k_2} c + \frac{k_3}{k_2} c = 0. \quad (3.52)$$

Thus $\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq 0$ for any $\mathbf{x} \in \partial\Omega(c)$, and by Nagumo's theorem [104], [123], the set $\Omega(c)$ is forward invariant for the true closed-loop system (2.2).

Consequently, by the compactness by $\Omega(c)$, for any initial condition $\mathbf{x}_0 \in D' \triangleq \text{Int}(\Omega(c))$, we have that the solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow \Omega(c)$ for the true closed-loop system (2.2) exists for all time ($I(\mathbf{x}_0) = \mathbb{R}_{\geq 0}$) [101, Chapter 2.4, Corollary 2]. Furthermore, from any initial condition $\mathbf{x}_0 \in D'$ we have that the actual projected disturbance encountered by the system satisfies:

$$|\delta(t)| \leq \frac{k_3}{k_2} c \triangleq \bar{d}, \quad (3.53)$$

for all $t \in \mathbb{R}_{\geq 0}$, such that $\|\delta\|_\infty \leq \bar{d}$. \square

This theorem highlights a critical challenge when working with local exponential stability and data-driven methods. As the value of c increases, the bound on the worst-case residual learning error is allowed to grow. This corresponds to the fact that the inequality enforced in the learning-informed controller (3.34) drives the system to the equilibrium point \mathbf{x}_e more aggressively the further the state of the system is from \mathbf{x}_e , overpowering the projected disturbance. At the same time, the worst-case residual learning error at additional states (as $\Omega(c_1) \subset \Omega(c_2)$ for $c_1 < c_2$) must be bounded. Furthermore, the value of c must be small enough that $\Omega(c) \subset D$, where the learning-informed controller (3.34) is feasible. This creates a “racing” behavior, where the learning error needs to grow slowly enough with c such that a small enough sublevel set $\Omega(c)$ meeting the inequality in the theorem can be found inside D . While it may be possible to change the learning-informed controller \mathbf{k} to ensure that the conditions of this theorem are met, one straightforward approach is to improve the uncertainty function Δ , as captured in the following theorem:

Theorem 19. *Consider uncertainty functions $\Delta, \Delta' : D \rightarrow \mathcal{P}(\mathbb{R} \times \mathbb{R}^m)$, and let $\Omega(c) \subset D$. If $\Delta'(\mathbf{x}) \subseteq \Delta(\mathbf{x})$ for all $\mathbf{x} \in \Omega(c)$, then:*

$$\sup_{\mathbf{x} \in \Omega(c)} \sup_{(b, \mathbf{a}) \in \Delta'(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}) \leq \sup_{\mathbf{x} \in \Omega(c)} \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}). \quad (3.54)$$

Proof. The proof is straightforward as for any $\mathbf{x} \in \Omega(c)$:

$$\sup_{(b, \mathbf{a}) \in \Delta'(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}) \leq \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}). \quad (3.55)$$

\square

This theorem states that if the possible uncertainties can be reduced (such that $\Delta'(\mathbf{x}) \subseteq \Delta(\mathbf{x})$ for all $\mathbf{x} \in \Omega(c)$), the bound on the worst-case projected

disturbance can be lowered, making it easier to satisfy the inequality in (3.48) and thus achieve the forward invariance that yields a valid D' and \bar{d} .

I will now explore how an uncertainty function Δ can be built from a data set \mathfrak{D} of the form in (3.17), restated here:

$$\mathfrak{D} = \{((\mathbf{x}_i, \mathbf{u}_i), \dot{V}_i)\}_{i=1}^N \subseteq (D \times \mathbb{R}^m) \times \mathbb{R}. \quad (3.56)$$

For notational convenience denote:

$$\mathfrak{D}_0 = \{(\mathbf{x}_i, \mathbf{u}_i) \in D \times \mathbb{R}^m \mid ((\mathbf{x}_i, \mathbf{u}_i), \dot{V}_i) \in \mathfrak{D}\}. \quad (3.57)$$

This leads to the following result:

Theorem 20. *Suppose⁴ that the functions $\tilde{\mathbf{f}}, \tilde{\mathbf{g}}$ are bounded on D by known constants $M_{\tilde{\mathbf{f}}}$ and $M_{\tilde{\mathbf{g}}}$, respectively, and are Lipschitz continuous on D with known Lipschitz constants $L_{\tilde{\mathbf{f}}}$ and $L_{\tilde{\mathbf{g}}}$, respectively. Then an uncertainty function Δ can be constructed as:*

$$\Delta(\mathbf{x}) = \{(b, \mathbf{a}) \in \mathbb{R} \times \mathbb{R}^m \mid \pm (b + \mathbf{a}^\top \mathbf{u}_i) \leq \epsilon(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i) \text{ for all } (\mathbf{x}_i, \mathbf{u}_i) \in \mathfrak{D}_0\}, \quad (3.58)$$

for all $\mathbf{x} \in D$, where $\epsilon : D \times D \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$ is continuous on its domain.

Observe that for each $\mathbf{x} \in D$, the set $\Delta(\mathbf{x})$ is a closed polyhedron, and bounded given sufficiently diverse control inputs in the data set (as will be made rigorous in Section 3.8). In this case, $\Delta(\mathbf{x})$ is a compact, convex set, and the supremum in (3.46) can be computed via a convex linear program (LP) that yields a finite optimal value.

Proof. Define the observed learning error as:

$$\ell(\mathbf{x}_i, \mathbf{u}_i) = \left| \dot{V}_i - \widehat{W}(\mathbf{x}_i, \mathbf{u}_i) \right|, \quad (3.59)$$

which can be computed from the data set \mathfrak{D} and the estimator \widehat{W} . Observe that we can rewrite this observed learning error as:

$$\ell(\mathbf{x}_i, \mathbf{u}_i) = \left| \tilde{b}(\mathbf{x}_i) + \tilde{\mathbf{a}}(\mathbf{x}_i)^\top \mathbf{u}_i \right|. \quad (3.60)$$

⁴Recall that D is a set associated with a *local* exponential CLF, and thus these bounds and Lipschitz properties generally only need to hold over a small neighborhood of the equilibrium point \mathbf{x}_e .

Given a test point $\mathbf{x} \in D$, we have that:

$$\ell(\mathbf{x}_i, \mathbf{u}_i) = \left| \tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}_i + \tilde{b}(\mathbf{x}_i) - \tilde{b}(\mathbf{x}) + (\tilde{\mathbf{a}}(\mathbf{x}_i) - \tilde{\mathbf{a}}(\mathbf{x}))^\top \mathbf{u}_i \right|, \quad (3.61)$$

$$\geq \left| \tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}_i \right| - \left| \tilde{b}(\mathbf{x}_i) - \tilde{b}(\mathbf{x}) \right| - \|\tilde{\mathbf{a}}(\mathbf{x}_i) - \tilde{\mathbf{a}}(\mathbf{x})\| \|\mathbf{u}_i\|, \quad (3.62)$$

following from the reverse triangle inequality, triangle inequality, and Cauchy-Schwarz inequality. This can be rearranged as:

$$\left| \tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}_i \right| \leq \ell(\mathbf{x}_i, \mathbf{u}_i) + \left| \tilde{b}(\mathbf{x}_i) - \tilde{b}(\mathbf{x}) \right| + \|\tilde{\mathbf{a}}(\mathbf{x}_i) - \tilde{\mathbf{a}}(\mathbf{x})\| \|\mathbf{u}_i\|. \quad (3.63)$$

We have that:

$$\left| \tilde{b}(\mathbf{x}_i) - \tilde{b}(\mathbf{x}) \right| = \left| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) \tilde{\mathbf{f}}(\mathbf{x}_i) - \hat{b}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}) + \hat{b}(\mathbf{x}) \right|, \quad (3.64)$$

$$\leq \left| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) \tilde{\mathbf{f}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}) \right| + \left| \hat{b}(\mathbf{x}) - \hat{b}(\mathbf{x}_i) \right|. \quad (3.65)$$

We can then perform the following manipulation:

$$\left| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) \tilde{\mathbf{f}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}) \right| = \left| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) \tilde{\mathbf{f}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}_i) + \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \tilde{\mathbf{f}}(\mathbf{x}) \right|, \quad (3.66)$$

$$\leq \left\| \tilde{\mathbf{f}}(\mathbf{x}_i) \right\| \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\| + \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\| \left\| \tilde{\mathbf{f}}(\mathbf{x}_i) - \tilde{\mathbf{f}}(\mathbf{x}) \right\|, \quad (3.67)$$

$$\leq M_{\tilde{\mathbf{f}}} \epsilon_\infty(\mathbf{x}, \mathbf{x}_i) + L_{\tilde{\mathbf{f}}} \epsilon_L(\mathbf{x}, \mathbf{x}_i), \quad (3.68)$$

where:

$$\epsilon_\infty(\mathbf{x}, \mathbf{x}_i) = \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}_i) - \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\|, \quad (3.69)$$

$$\epsilon_L(\mathbf{x}, \mathbf{x}_i) = \left\| \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) \right\| \|\mathbf{x} - \mathbf{x}_i\|. \quad (3.70)$$

This yields:

$$\left| \tilde{b}(\mathbf{x}_i) - \tilde{b}(\mathbf{x}) \right| \leq M_{\tilde{\mathbf{f}}} \epsilon_\infty(\mathbf{x}, \mathbf{x}_i) + L_{\tilde{\mathbf{f}}} \epsilon_L(\mathbf{x}, \mathbf{x}_i) + \left| \hat{b}(\mathbf{x}) - \hat{b}(\mathbf{x}_i) \right|. \quad (3.71)$$

A similar decomposition can be performed to establish:

$$\|\tilde{\mathbf{a}}(\mathbf{x}_i) - \tilde{\mathbf{a}}(\mathbf{x})\| \leq M_{\tilde{\mathbf{g}}} \epsilon_\infty(\mathbf{x}, \mathbf{x}_i) \|\mathbf{u}_i\| + L_{\tilde{\mathbf{g}}} \epsilon_L(\mathbf{x}, \mathbf{x}_i) \|\mathbf{u}_i\| + \|\hat{\mathbf{a}}(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x}_i)\| \|\mathbf{u}_i\|. \quad (3.72)$$

Denoting:

$$\epsilon_{\mathcal{H}}(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i) = \left| \widehat{b}(\mathbf{x}) - \widehat{b}(\mathbf{x}_i) \right| + \|\widehat{\mathbf{a}}(\mathbf{x}) - \widehat{\mathbf{a}}(\mathbf{x}_i)\| \|\mathbf{u}_i\|, \quad (3.73)$$

we have that:

$$\begin{aligned} \left| \widetilde{b}(\mathbf{x}) + \widetilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}_i \right| &\leq \ell(\mathbf{x}_i, \mathbf{u}_i) + (M_{\widetilde{f}} + M_{\widetilde{g}} \|\mathbf{u}_i\|) \epsilon_{\infty}(\mathbf{x}, \mathbf{x}_i) \\ &\quad + (L_{\widetilde{f}} + L_{\widetilde{g}} \|\mathbf{u}_i\|) \epsilon_L(\mathbf{x}, \mathbf{x}_i) + \epsilon_{\mathcal{H}}(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i), \end{aligned} \quad (3.74)$$

$$\triangleq \epsilon(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i). \quad (3.75)$$

Note that if the estimators \widehat{b} and $\widehat{\mathbf{a}}$ are continuous functions on D , the function ϵ is continuous on its domain. Breaking the absolute value on the left-hand side in to two cases, we have:

$$\pm \left(\widetilde{b}(\mathbf{x}) + \widetilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}_i \right) \leq \epsilon(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i). \quad (3.76)$$

Consequently, defining the uncertainty function Δ as in (3.58), we have that $(\widetilde{b}(\mathbf{x}), \widetilde{\mathbf{a}}(\mathbf{x})) \in \Delta(\mathbf{x})$ as required. \square

Observe that given a test point $\mathbf{x} \in D$, the function ϵ can be computed from the data set \mathfrak{D} , estimator \widehat{W} , and constants $M_{\widetilde{f}}$, $M_{\widetilde{g}}$, $L_{\widetilde{f}}$, and $L_{\widetilde{g}}$, and thus it is possible to evaluate the supremum in (3.46) at the test point by solving a linear program. Furthermore, we have that:

$$\epsilon(\mathbf{x}_i, \mathbf{x}_i, \mathbf{u}_i) = \ell(\mathbf{x}_i, \mathbf{u}_i). \quad (3.77)$$

This means that close to a data point, the bound on the uncertainty is only determined by the observed learning error $\ell(\mathbf{x}_i, \mathbf{u}_i)$. Thus, the uncertainty set can be kept small by sufficiently accurate learning models and sufficiently dense data covering the set D .

Lastly, let us consider an estimator \widehat{W} trained on a data set \mathfrak{D} . Suppose that additional data is produced, and added to the data set \mathfrak{D} to produce a data set \mathfrak{D}' with $\mathfrak{D} \subset \mathfrak{D}'$, but the estimator \widehat{W} is not retrained. For uncertainty functions $\Delta, \Delta' : D \rightarrow \mathcal{P}(\mathbb{R} \times \mathbb{R}^m)$ constructed using the data sets \mathfrak{D} and \mathfrak{D}' , respectively, we have that $\Delta'(\mathbf{x}) \subseteq \Delta(\mathbf{x})$ for all $\mathbf{x} \in D$. This is because the function ϵ is the same between the two constructions, such that for a data point $(\mathbf{x}_i, \mathbf{u}_i) \in \mathfrak{D}_0$, at a test point $\mathbf{x} \in D$ we have that $\epsilon_i(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i)$ returns the same value between the two constructions. But, by adding data points in

going from \mathfrak{D} to \mathfrak{D}' , we have added additional constraints in the construction of Δ' compared to Δ , thus reducing possible residual learning errors. Thus, continuing to add data allows the improvement of residual learning error with a fixed estimator. If the estimator \widehat{W} is retrained with the new data, this may not necessarily be the case, as the observed residual learning error $\ell(\mathbf{x}_i, \mathbf{u}_i)$ and the function $\epsilon_{\mathcal{H}}(\mathbf{x}, \mathbf{x}_i, \mathbf{u}_i)$ can increase for data points $(\mathbf{x}_i, \mathbf{u}_i) \in \mathfrak{D}_0$. This must be balanced with the benefits of achieving low observed residual learning error on the data points appearing in $\mathfrak{D}' \setminus \mathfrak{D}$, which may be in regions of interest.

Simulation Results

In this section I apply the episodic learning DaCLyF algorithm in Algorithm 1 to an inverted pendulum subject to parametric model error. The nominal model open-loop system is given by:

$$\widehat{\dot{\mathbf{x}}} = \begin{bmatrix} x_2 \\ \widehat{\frac{g}{\ell}} \sin(x_1) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{\widehat{m\ell^2}} \end{bmatrix} u, \quad (3.78)$$

while the true open-loop system is given by:

$$\dot{\mathbf{x}} = \begin{bmatrix} x_2 \\ \frac{g}{\ell} \sin(x_1) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{m\ell^2} \end{bmatrix} u. \quad (3.79)$$

The true mass m and length ℓ are perturbations of the nominal model mass \widehat{m} and length $\widehat{\ell}$ by up to 30% of the nominal model values. The nominal controller k_{nom} is a linear PD controller set to track angle and angle rate trajectories. The estimators \widehat{b} and $\widehat{\mathbf{a}}$ are chosen from the class of two layer neural networks with 200 hidden units and ReLU nonlinearities. The trust factors are chosen in a sigmoid fashion. Naive exploratory control is introduced as in Section 3.3, with perturbations chosen uniformly at random between -50% and 50% of the value of the augmented controller k_k . A comparison of the nominal model-based (CLF-QP) controller and final learning-informed (CLF-QP) controller demonstrating improved tracking performance is shown in Figure 3.3.

After performing the episodic learning process, a final data set \mathfrak{D} and estimator \widehat{W} are produced, and an uncertainty function Δ can be built. I evaluate the following quantities across the state space:

$$\sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \widehat{\mathbf{k}}(\mathbf{x}), \quad (3.80)$$

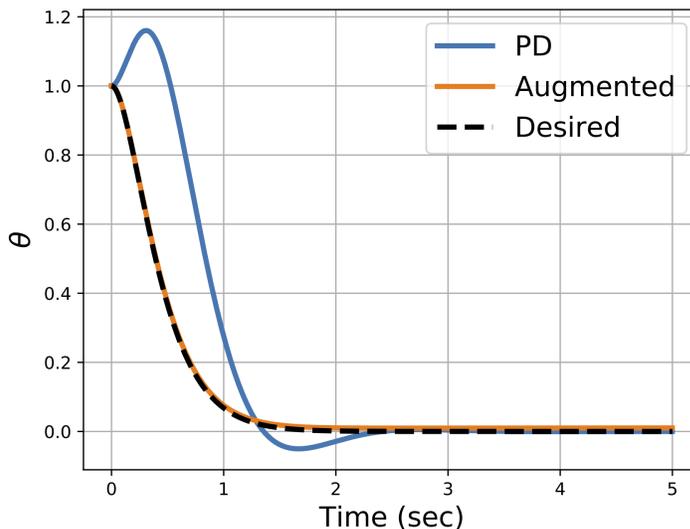


Figure 3.3. Comparison of tracking performance for PD controller and final learning-informed (CLF-QP) controller in inverted pendulum example. The final learning-informed (CLF-QP) controller tracks the desired angle trajectory significantly more accurately.

and:

$$\sup_{(b,\mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}). \quad (3.81)$$

The first quantity corresponds to the worst-case residual learning error when using the nominal model-based (CLF-QP) controller in (3.25), while the second quantity is the worst-case residual learning error using the learning-informed (CLF-QP) controller in (3.34). The closed-loop trajectories using these controllers from the same initial conditions can be seen in Figure 3.4, on the left and right, respectively. We see that the nominal model-based (CLF-QP) controller does not stabilize the system to the desired trajectory and fails to bring it to the origin, while the learning-informed (CLF-QP) controller does stabilize the system and brings it to the origin. The value of the computed supremums is indicated by the color-maps. Observe that there is some mild change between the values for the nominal model-based (CLF-QP) controller and learning-informed (CLF-QP) controller (in particular, in the top right of the plot), but that the changes throughout the state space are fairly small. This indicates that although the learning-informed (CLF-QP) controller is able to stabilize the system, it has not greatly reduced the worst-case residual learning error possible by the uncertainty function Δ . This suggests that the system is not actually experiencing the worst-case residual learning error,

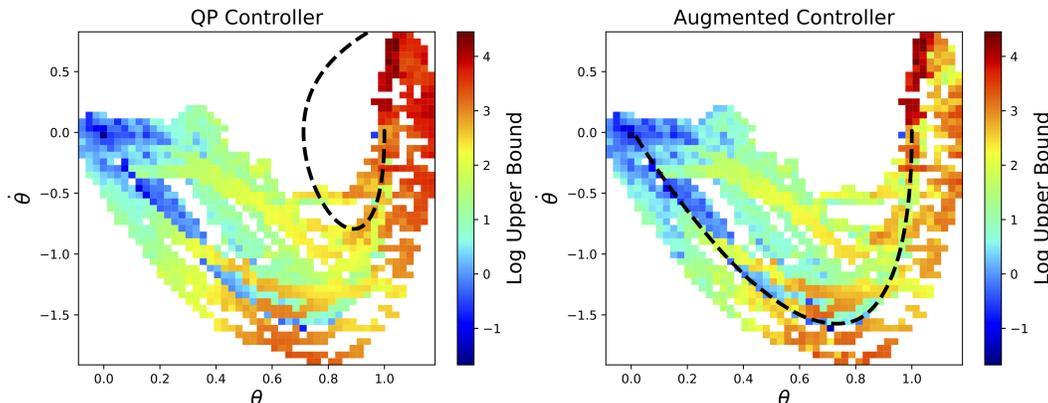


Figure 3.4. Comparison of worst-case residual learning errors (3.80) with nominal model-based (CLF-QP) controller (Left) and (3.81) with final learning-informed (CLF-QP) controller (Right). Closed-loop trajectories using the two controllers are displayed with dashed lines. The learning-informed (CLF-QP) controller keeps the system in regions with lower residual learning error, while the system leaves the region around the training data under the nominal model-based (CLF-QP) controller. The heat maps were generated by sampling states randomly about training data points and evaluating the upper bound for each sampled state. The results were then discretized for ease of visualization. Each bin is colored by the maximum disturbance observed in the bin.

and that reducing the worst-case residual learning error is not necessary for successful stabilization if a controller does not “excite” the worst-case residual learning error. This observation is the basis for the work in Section 3.8.

Conclusion

In this section I have presented an approach for analyzing the impacts of residual learning errors when learning model errors between a nominal model open-loop system and the true open-loop system as they directly impact the time derivative of a local exponential CLF V . I propose a novel notion of robust stability based on ISS in the form of Projection-to-State Stability (PSS), which only considers the impacts of disturbances on local exponential stability as they affect the time derivative of a local exponential CLF. I then provide an analysis of closed-loop system behavior given a particular data set and learning models through a data-driven construction of uncertainty functions. Lastly I perform this data-driven analysis in simulation on an inverted pendulum, revealing that stabilization can be achieved without needing to explicitly reduce worst-case residual learning errors, which will motivate the work in Section 3.8.

3.5 Learning with Control Barrier Functions

The inherent similarities between local exponential CLFs and CBFs suggest that the episodic learning algorithm developed in Section 3.3 could naturally be adapted to address model error directly as it impacts the time derivative of a CBF. In this section I will adapt the DaCLyF episodic learning algorithm in Algorithm 1 to the setting of CBFs, producing a corresponding algorithm named DaCBarF. I will explore how model uncertainty impacts the time derivative of a CBF, and how an ERM problem can be solved sequentially to overcome challenges that arise when deploying standard supervised learning techniques on non-i.i.d data sets. While these adaptations are fairly straightforward and provided for completeness, the main contribution of this section is the demonstration of the episodic learning algorithm on a Segway hardware platform. This result establishes the possibility of identifying error between a nominal model and a real-world system using real data, and constructing learning models that can reflect this model error while remaining computationally efficient enough to be evaluated in a closed-loop controller for a robotic system.

The contributions of this section are as follows:

- An adaption of the episodic learning algorithm presented in Section 3.3 for identifying model errors as they directly impact the time derivative of a CBF, including using structured learning models to enable using the convex optimization-based control framework described in Section 2.5.
- An experimental demonstration of the proposed episodic learning algorithm on a Segway hardware platform.

The text for this section is adapted from:

A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “Learning for safety-critical control with control barrier functions,” in *Proc. 2nd Learning for Dynamics and Control (L4DC)*, vol. 120, Berkeley, CA, USA, 2020, pp.708-717.

A. J. Taylor participated in the conception of the project, algorithm design, simulation and experimental code implementation, conduction of experiments, and writing of the article.

Uncertainty in Control Barrier Functions

Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . Suppose that h is a CBF for the nominal model open-loop system (3.1) on \mathcal{C} with corresponding function $\alpha \in \mathcal{K}^e$, such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \widehat{h}(\mathbf{x}, \mathbf{u}) \triangleq \sup_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{u}) > -\alpha(h(\mathbf{x})), \quad (3.82)$$

for all $\mathbf{x} \in E$. The time derivative of the function h uses the true open-loop system (2.1), and thus using (3.2), we have that:

$$\dot{h}(\mathbf{x}, \mathbf{u}) = \widehat{h}(\mathbf{x}, \mathbf{u}) + \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\widetilde{\mathbf{f}}(\mathbf{x})}_{b(\mathbf{x})} + \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\widetilde{\mathbf{g}}(\mathbf{x})}_{\mathbf{a}(\mathbf{x})^\top} \mathbf{u}. \quad (3.83)$$

As with CLFs, simply choosing an input \mathbf{u} such that the inequality in (3.82) is satisfied is insufficient to conclude that the time derivative of the CBF h will meet the required inequality to guarantee safety because of the functions $b : E \rightarrow \mathbb{R}$ and $\mathbf{a} : E \rightarrow \mathbb{R}^m$. Note the functions b and \mathbf{a} capture the model errors in $\widetilde{\mathbf{f}}$ and $\widetilde{\mathbf{g}}$ only as they impact the evolution of the CBF h , and consequently, safety. Once again, I will consider a learning-based approach to identify the functions b and \mathbf{a} , such that their effects on the CBF time derivative can be accounted for in the control design process, and safety can be achieved. To ensure the well-posedness of such an approach, I make the following assumption:

Assumption 3. The function h is a CBF for the true open-loop system (2.1) on \mathcal{C} with corresponding function $\alpha \in \mathcal{K}^e$.

This assumption is similar to Assumption 2 and states that h is a CBF for the true open-loop system (2.1) on \mathcal{C} with the same function $\alpha \in \mathcal{K}^e$ used for the nominal model open-loop system (3.1). It is important to note that once again, this is an assumption on how the true open-loop system (2.1) *can* be controlled. In particular, this assumption states that for a given $\mathbf{x} \in E$, one knows that there *exists* an input $\mathbf{u} \in \mathbb{R}^m$ such that:

$$\dot{h}(\mathbf{x}, \mathbf{u}) = \widehat{h}(\mathbf{x}, \mathbf{u}) + b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u} > -\alpha(h(\mathbf{x})), \quad (3.84)$$

but not necessarily *which* value of \mathbf{u} satisfies the inequality. If one knew the functions b and \mathbf{a} perfectly, the value of \mathbf{u} satisfying this inequality could be found. Thus, if one could learn these functions from data, the control synthesis

problem through CBFs would be well-posed. Recalling the construction of a CBF from a local exponential CLF in Section 2.5 and the construction of a local exponential CLF for both a fully-actuated robotic system nominal model open-loop system and true open-loop system in Section 3.3, assumptions on the structural properties of the true-open loop system can be sufficient to ensure this assumption is met.

Data-Driven Episodic Learning Framework

As with CLFs, I consider a learning-based framework that learns the functions b and \mathbf{a} via supervised regression [204]. In particular, the true closed-loop system (2.2) is operated using a given controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$, data points are gathered along the system's evolution, and functions that approximate b and \mathbf{a} are produced via supervised learning.

Concretely, let $\mathbf{x}_0 \in E$ be an initial state. An experiment is defined as the evolution of the true closed-loop system over a finite time interval from the initial condition \mathbf{x}_0 using a sample-hold implementation of the controller \mathbf{k} . A resulting discrete-time state history is obtained, which is then transformed with the CBF h and finally differentiated numerically to estimate \dot{h} throughout the experiment. This yields a data set comprised of input-output pairs:

$$\mathfrak{D} = \{((\mathbf{x}_i, \mathbf{u}_i), \dot{h}_i)\}_{i=1}^N \subseteq (E \times \mathbb{R}^m) \times \mathbb{R}. \quad (3.85)$$

Consider a class \mathcal{H}_b of nonlinear functions mapping \mathbb{R}^n to \mathbb{R} and a class \mathcal{H}_a of nonlinear functions mapping \mathbb{R}^n to \mathbb{R}^m . For a given $\hat{b} \in \mathcal{H}_b$ and $\hat{\mathbf{a}} \in \mathcal{H}_a$, define \hat{S} as:

$$\hat{S}(\mathbf{x}, \mathbf{u}) = \hat{h}(\mathbf{x}, \mathbf{u}) + \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}, \quad (3.86)$$

and let \mathcal{H} be the class of all such estimators mapping $\mathbb{R}^n \times \mathbb{R}^m$ to \mathbb{R} . Defining a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, the supervised regression task is then to find a function in \mathcal{H} via empirical risk minimization (ERM):

$$\inf_{\substack{\hat{\mathbf{a}} \in \mathcal{H}_a \\ \hat{b} \in \mathcal{H}_b}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{S}(\mathbf{x}_i, \mathbf{u}_i), \dot{h}_i). \quad (3.87)$$

This process can be executed either in simulation or directly on hardware.

Once again, standard supervised learning with sequential, non-i.i.d data collection often leads to error cascades [142], motivating the use of an episodic learning framework. My approach integrates learning b and \mathbf{a} with improving

the safety of the controller in such an iterative fashion. First, assume we are given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$, which may not render the set \mathcal{C} safe for the resulting true closed-loop system (2.2). With an estimator $\hat{S} \in \mathcal{H}$ defined in (3.18), specify a learning-informed (CBF-QP) controller as:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ \text{s.t. } & \hat{S}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x})). \end{aligned} \quad (3.88)$$

This learning-informed (CBF-QP) controller finds the closest input to $\mathbf{k}_{\text{nom}}(\mathbf{x})$ such that the CBF inequality is met using the estimator \hat{S} for \dot{h} . However, error in the estimator \hat{S} degrades the safety of the true closed-loop system (2.2) with respect to the set \mathcal{C} .

In an effort to reduce the remaining error in the estimator \hat{S} , an experiment can be run using the learning-informed (CBF-QP) controller to produce data by which to obtain better estimates of b and \mathbf{a} . In particular, in each episode, an experiment generates data using a different controller. The data set is aggregated and a new ERM problem is solved after each episode, once again inspired by the Data Aggregation (DAgger) algorithm [203]. The DaCBarF algorithm in Algorithm 2 specifies a method of computing a sequence of CBF derivative estimates \hat{S}_k and augmented controllers \mathbf{k}_k . During the k^{th} episode, the learning-informed (CBF-QP) controller (3.88) associated with the estimate of the CBF time derivative \hat{S}_k is scaled by a heuristically chosen factor $w_k \in [0, 1]$ reflecting trust in the estimate and added to the nominal model-based (CBF-QP) controller $\mathbf{k}_0 = \mathbf{k}_{\text{CBF}}$ for use in the subsequent experiment as an augmented controller \mathbf{k}_k . The trust coefficients form a monotonically non-decreasing sequence. Importantly, this experiment need not take place in simulation; the same procedure may be executed directly on hardware. At a high level, this episodic approach makes progress by gathering more data in relevant regions of the state space, such as states inside \mathcal{C} . This extends the generalizability of the estimator in its use by subsequent controllers.

There are a few practical concerns with this algorithm that should be observed. First, the system will not necessarily be safe with respect to the set \mathcal{C} during the episodic learning process. In general, one will often begin the learning process with a controller that is unsafe, or displays desirable behavior but possess no guarantees of safety (such as a barrier function). This algorithm

Algorithm 2 Dataset Aggregation for Control Barrier Functions (DaCBarF)

Require: CBF h , derivative estimate $\hat{S}_0 = \hat{h}$, model classes \mathcal{H}_a and \mathcal{H}_b , loss function \mathcal{L} , nominal model-based (CBF-QP) controller $\mathbf{k}_0 = \mathbf{k}_{\text{CBF}}$, number of experiments T , sequence of trust coefficients $0 \leq w_1 \leq \dots \leq w_T \leq 1$

```

 $\mathcal{D} = \emptyset$  ▷ Initialize data set
for  $k = 1, \dots, T$  do
   $\mathbf{x}_0 \leftarrow \text{sample } E$  ▷ Get initial condition
   $\mathcal{D}_k \leftarrow \text{experiment}(\mathbf{x}_0, \mathbf{k}_{k-1})$  ▷ Run experiment
   $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_k$  ▷ Aggregate data set
   $\hat{\mathbf{a}}_k, \hat{b}_k \leftarrow \text{ERM}(\mathcal{H}_a, \mathcal{H}_b, \mathcal{L}, \mathcal{D}, \hat{S}_0)$  ▷ Fit estimators
   $\hat{S}_k \leftarrow \hat{S}_0 + \hat{\mathbf{a}}_k^\top \mathbf{u} + \hat{b}_k$  ▷ Update derivative estimator
   $\mathbf{k}_k \leftarrow (1 - w_k) \cdot \mathbf{k}_0 + w_k \cdot \text{augment}(\mathbf{k}_0, \hat{S}_k)$  ▷ Update controller
end for
return  $\hat{S}_T, \mathbf{k}_T$ 

```

instead tolerates unsafe behavior during the episodic learning process to ensure that the system performs adequate exploration in the data acquisition process and return a controller at the end of the episodic learning for which meaningful safety properties can be established, as explored in Section 3.6.

Second, there can be issues that arise with the constraint in the learning-informed controller (3.88). In particular, while we begin knowing that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \hat{h}(\mathbf{x}, \mathbf{u}) > -\alpha(h(\mathbf{x})), \quad (3.89)$$

for all $\mathbf{x} \in E$, and we assume that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) > -\alpha(h(\mathbf{x})), \quad (3.90)$$

for all $\mathbf{x} \in E$, such that the CBF constraint is met for the nominal model open-loop system (3.1) and the true open-loop system (2.1), we do not know:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \hat{S}_k(\mathbf{x}, \mathbf{u}) > -\alpha(h(\mathbf{x})), \quad (3.91)$$

for all $\mathbf{x} \in E$ for $k = 1, \dots, T$. In particular, it is possible that the learned models \hat{b}_k and $\hat{\mathbf{a}}_k$ can lead to this condition being violated. This means that the learning-informed controller will not necessarily be feasible for all $\mathbf{x} \in E$. In practice, the relaxation technique used in (CLF-CBF-QP) is utilized, such that the problem always remains feasible at the expense of rigorous safety

guarantees. In Section 3.6 safety guarantees will be built keeping in mind that some amount of residual learning error arising from numerical learning methods will be present in learning models.

Simulation & Experimental Results

I apply the proposed episodic learning algorithm to the Segway platform considered in Section 3.3, which can be seen in Figures 3.5 and 3.6. The nominal controller $k_{\text{nom}} : E \rightarrow \mathbb{R}$ is a hand-tuned proportional-derivative (PD) controller which locally exponentially stabilizes the system to an equilibrium configuration. A particular requirement for safety of the Segway system can be encoded as limitations on how far and how fast the Segway is allowed to tip from this equilibrium configuration. This appears mathematically as a CBF on the pitch angle and pitch angle rate:

$$h(\theta, \dot{\theta}) = \frac{1}{2} \left(\theta_{\max}^2 - (\theta - \theta_e)^2 - c\dot{\theta}^2 \right), \quad (3.92)$$

where $\theta \in \mathbb{R}$ is the pitch angle, $\dot{\theta} \in \mathbb{R}$ is the pitch angle rate, $\theta_{\max} \in \mathbb{R}_{>0}$ is a limit on how far the pitch angle can deviate from the equilibrium pitch angle, $\theta_e \in \mathbb{R}$, and the coefficient $c \in \mathbb{R}_{>0}$ scales the importance of the pitch angle rate in defining safety. This construction of h is performed such that it is a CBF for the nominal model open-loop system (3.1), and thus the (CBF-QP) controller can be used to ensure the safety of the system (in the absence of model errors).

To evaluate the impact of model errors on the ability of the nominal model-based (CBF-QP) controller to ensure the safety of the system, I perturb the mass and electrical parameters of the Segway model by up to 15%, but withhold these artificial model errors from the nominal model-based (CBF-QP) controller. Figure 3.5 shows that the nominal model-based (CBF-QP) controller is unable to achieve safety, with the state of the system leaving the 0-superlevel set of h (left), and the value of h falling below 0 (right).

I deploy the episodic learning framework on the Segway system to overcome this model error and improve safety. In particular, I conduct a sequence of ten experiments with varying initial conditions. Data from each experiment is aggregated into a data set \mathfrak{D} that is used to train learning models \hat{b} and \hat{a} and construct the estimator \hat{S} . The learning models \hat{b} and \hat{a} are chosen to be neural networks with one hidden layer of 200 nodes and ReLU nonlinearities.

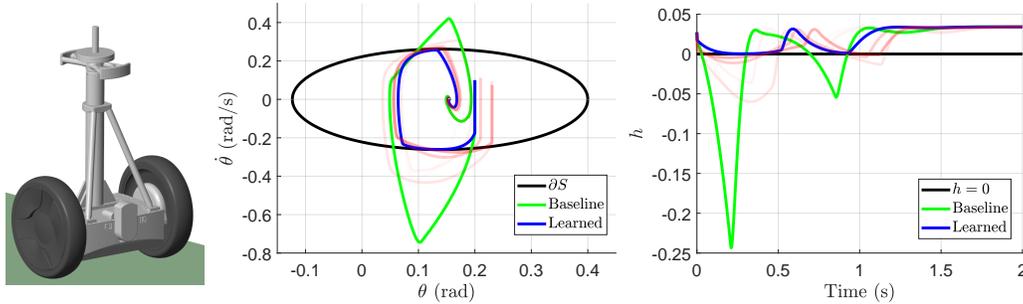


Figure 3.5. Simulation results using the nominal model-based (CBF-QP) controller (green) and learning-informed (CBF-QP) controller (3.88) (blue). Episodic data appears as red traces. (Left) Segway simulation CAD model. (Middle) The phase portrait showing the evolution of the state of the system leaves the safe set (black ellipse) with the nominal model-based (CBF-QP) controller, while it remains within using the learning-informed (CBF-QP) controller. (Right) The value of h drops below zero (twice) using the nominal model-based (CBF-QP) controller, while it remains above zero using the learning-informed (CBF-QP) controller. Video can be found at https://youtu.be/tD8zH4GF_8U.

These networks are implemented in Keras [208] and trained to minimize mean absolute error using stochastic gradient descent (SGD). After each episode the learning-informed (CBF-QP) controller (3.88) is blended with the nominal model-based (CBF-QP) using trust factor weights that grow linearly across the episodes. Figure 3.5 shows the behavior of the system when the learning-informed (CBF-QP) controller (3.88) is deployed after episodic learning. The state of the system is kept within the 0-superlevel set of h (left), with the value of h being kept above 0 (right). This indicates that the proposed learning approach is able to sufficiently identify model errors as they impact safety such that the learning-informed (CBF-QP) controller is able to keep the system safe.

To further demonstrate the ability of this learning approach, I deploy it on the Segway hardware platform seen in Figure 3.6. The nominal controller k_{nom} was designed to track a desired horizontal velocity profile while satisfying a safety constraint on the pitch angle rate specified by the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ defined as:

$$h(\theta, \dot{\theta}) = \frac{1}{2}(1 - c\dot{\theta}^2). \quad (3.93)$$

The desired velocity profile is chosen to yield a violation of this safety constraint when using the nominal model-based (CBF-QP) controller, as seen in Figure 3.6. A series of three episodes were conducted tracking this velocity profile without introducing learning models into the controller, after which

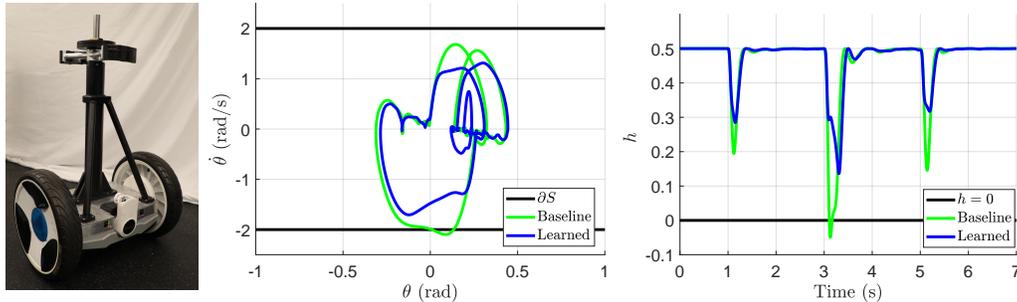


Figure 3.6. Experimental results using the nominal model-based (CBF-QP) controller (green) and learning-informed (CBF-QP) controller (3.88) (blue). (Left) Custom robotic Ninebot Segway platform. (Middle) The phase portrait showing the evolution of the state of the system leaves the safe set (black strip) with the nominal model-based (CBF-QP) controller, while it remains within using the learning-informed (CBF-QP) controller. (Right) The value of h drops below zero (at $t = 3$) using the nominal model-based (CBF-QP) controller, while it remains above zero using the learning-informed (CBF-QP) controller. Video can be found at https://youtu.be/YCesvD_Ae00

estimators $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ were implemented as neural networks with two hidden layers of 50 nodes and ReLU nonlinearities, trained on this collected data, and incorporated into the learning-informed controller (3.20). Note that this implementation deviates from the DaCBarF algorithm by not updating the controller between episodes, but performs an initial set of experiments to form a data set and learning just using this data. In this experimental configuration, natural variations in the initial condition of the system and behaviors of the hardware platform led to enough variation of data to support successful learning for this safety task. The closed-loop behavior of the system can be seen in Figure 3.6. The state of the system is kept within the 0-superlevel set of h (left), with the value of h being kept above 0 (right). This indicates that the proposed learning approach is able to sufficiently identify model errors as they impact safety, the assumption of h being a CBF for the true-open loop system is justified, and the learning-informed (CBF-QP) controller is able to keep the system safe.

Conclusion

In this section I have presented an adaptation of the episodic learning framework in Section 3.3 for learning model errors between a nominal model open-loop system and the true open-loop system as they directly impact the time derivative of a CBF h . I formalize how model errors appear in the time derivative of a CBF and adapt the episodic learning algorithm for the setting of CBFs,

yielding the DaCBarf algorithm. I demonstrate the algorithm in simulation, and as the main contribution of this section, experimentally on a Segway hardware platform.

3.6 Projection-to-State Safety

In this section I will present an analysis of the effects of residual learning error on the safety of the true closed-loop system (2.2) using a learning-informed (CBF-QP) controller. This analysis will mirror the one presented in Section 3.4 that considered residual learning error as a disturbance in the time derivative of a local exponential CLF and culminated in the idea of Projection-to-State Stability (PSS). In the setting of safety, I will build off of the episodic learning algorithm in Section 3.5 and consider learning model errors as they impact the time derivative of a CBF. This will lead to a characterization of the inherent robustness of CBF-based controllers to residual learning errors in a modification of ISSf known as Projection-to-State Safety (PSSf). Given a data set and corresponding learning models, this will lead to a statement on the superlevel set of a CBF that can be kept forward invariant in the presence of residual learning errors. While the extensions in going from PSS to PSSf are relatively straightforward due to the similarities between CLFs and CBFs, a key contribution of this work will be the demonstration of PSSf behavior on a Segway hardware platform.

The contributions of this section are as follows:

- A notion of robust safety in the form of *Projection-to-State Safety (PSSf)*, which adapts the analysis via CLFs and PSS in Section 3.4 to CBFs, thereby capturing the impact of disturbances on safety only as they impact the time derivative of a CBF, rather than the full-order dynamics, as is done by ISSf discussed in Section 2.7.
- An experimental demonstration of PSSf behavior with respect to residual learning errors on a Segway hardware platform, including a demonstration of how learning models can improve PSSf safety guarantees.

The text for this section is adapted from:

A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “A control barrier perspective on episodic learning via projection-to-state safety,” in *IEEE*

Control Sys. Let., vol. 5, no. 3, pp.1019-1024, 2021.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation and experimental code implementation, conduction of experiments, and writing of the article.

Projection-to-State Safety

In the context of learning model errors as they appear in the time derivative of a CBF as in Section 3.5, it can be difficult to formalize residual learning errors as a disturbance signal in the full state dynamics. Rather, as in Section 3.4, I will look at residual learning error as it appears directly in the time derivative of the CBF. In particular, recall that:

$$\dot{h}(\mathbf{x}, \mathbf{u}) = \hat{h}(\mathbf{x}, \mathbf{u}) + b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{u}. \quad (3.94)$$

Let $\hat{b} \in \mathcal{H}_b$ and $\hat{\mathbf{a}} \in \mathcal{H}_a$ compose the estimator $\hat{S} \in \mathcal{H}$ as:

$$\hat{S}(\mathbf{x}, \mathbf{u}) = \hat{h}(\mathbf{x}, \mathbf{u}) + \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}, \quad (3.95)$$

and let us assume that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \hat{S}(\mathbf{x}, \mathbf{u}) > -\alpha(h(\mathbf{x})), \quad (3.96)$$

for all $\mathbf{x} \in E$. Adding and subtracting the estimators \hat{b} and $\hat{\mathbf{a}}$ from \dot{h} yields:

$$\dot{h}(\mathbf{x}, \mathbf{u}) = \hat{S}(\mathbf{x}, \mathbf{u}) + \underbrace{b(\mathbf{x}) - \hat{b}(\mathbf{x})}_{\tilde{b}(\mathbf{x})} + \underbrace{(\mathbf{a}(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x}))^\top}_{\tilde{\mathbf{a}}(\mathbf{x})^\top} \mathbf{u}. \quad (3.97)$$

Once again, the functions $\tilde{b} : E \rightarrow \mathbb{R}$ and $\tilde{\mathbf{a}} : E \rightarrow \mathbb{R}^m$ capture residual learning error left after trying to learn the functions b and \mathbf{a} with the estimators \hat{b} and $\hat{\mathbf{a}}$. As these error terms appear in the time derivative of the CBF h , they will consequently have an impact on the safety properties of the system. Letting $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ be a nominal controller, using the learning-informed (CBF-QP) controller:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \\ &\text{s.t. } \hat{S}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x})), \end{aligned} \quad (3.98)$$

yields:

$$\dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) = \hat{S}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}). \quad (3.99)$$

For any initial condition $\mathbf{x}_0 \in E$, let $\boldsymbol{\varphi} : I(\mathbf{x}_0) \rightarrow E$ denote the solution to the true closed-loop system (2.2) using the learning-informed (CBF-QP) controller (3.98). We can then view the evolution of (3.99) as a closed-loop system with unmatched disturbances:

$$\dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) = \widehat{S}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \delta(t), \quad (3.100)$$

where for any initial condition $\mathbf{x}_0 \in E$, the function $\delta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ defined as:

$$\delta(t) = \widetilde{b}(\boldsymbol{\varphi}(t)) + \widetilde{\mathbf{a}}(\boldsymbol{\varphi}(t))^\top \mathbf{k}(\boldsymbol{\varphi}(t)). \quad (3.101)$$

is a projected disturbance. This perspective motivates studying residual learning error through the lens of ISSf. Observe that the learning-informed (CBF-QP) controller enforces:

$$\dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \geq -\alpha(h(\mathbf{x})) - |\delta(t)|. \quad (3.102)$$

Considering the identity map $\iota : \mathbb{R} \rightarrow \mathbb{R}$ defined as $\iota(r) = r$ for all $r \in \mathbb{R}$, note that the 0-superlevel set of ι is $\mathbb{R}_{\geq 0}$. Furthermore, the identity map satisfies:

$$\frac{\partial \iota}{\partial r}(h(\mathbf{x})) \dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \geq -\alpha(h(\mathbf{x})) - |\delta(t)| = -\alpha(\iota(h(\mathbf{x}))) - |\delta(t)|, \quad (3.103)$$

and thus is a ISSf-BF for the closed-loop system (3.100) on $\mathbb{R}_{\geq 0}$. Thus, noting the relationship between the 0-superlevel set of h and 0-superlevel set of the identity map ι , there exists a constant $\bar{d} \in \mathbb{R}_{> 0}$ and a function $\gamma \in \mathcal{K}$ such that for all $d \in [0, \bar{d}]$, the set:

$$\mathcal{C}_d = \{\mathbf{x} \in E \mid \iota(h(\mathbf{x})) + \gamma(d) \geq 0\}, \quad (3.104)$$

is forward invariant for the closed-loop system (3.100) if the projected disturbance satisfies⁵ $\|\delta\|_\infty \leq d$ for all initial conditions $\mathbf{x}_0 \in \mathcal{C}_d$.

This appears like a statement of ISSf for the true closed-loop system (2.2) using the learning-informed controller (3.98), but it is defined now in terms of a bound on the projected disturbance signal δ rather than a disturbance signal \mathbf{d} in the state dynamics. Consequently, the true closed-loop system (2.2) with the learning-informed controller (3.98) is said to be *Projection-to-State Safe*

⁵In the setting of forward invariance and safety, we only care about the solution $\boldsymbol{\varphi}$ being contained in a set for the interval on which it exists. Consequently, the particular projected disturbance signal δ defined in (3.101) may not be defined on all of $\mathbb{R}_{\geq 0}$. In this case, the $\|\cdot\|_\infty$ norm is understood to be over the interval of existence of the solution.

(*PSSf*) with respect to the set \mathcal{C} and projection h . The utility of *PSSf* is that it considers the impacts of model error only as it impacts the time derivative of the CBF h . By doing so, it provides a mathematical framework for understanding the impact of introducing learning models for the time derivative of a CBF, as it is not obvious how to represent residual learning error in learning models for a CBF time derivative as a disturbance in the full state dynamics (2.2), making an analysis by standard *ISSf* difficult. Furthermore, it may potentially produce less conservative statements on the degradation of safety by neglecting model error that does not affect safety.

Ensuring Projection-to-State Safety with Residual Learning Error

As we saw with *PSS* in Section 3.4, the preceding analysis through *PSSf* produced some value of \bar{d} which is not particularly useful if we do not know that the particular projected disturbance signal in (3.101) at least satisfies $\|\delta\|_\infty \leq \bar{d}$. We may not simply increase the value of \bar{d} to achieve this, as we must then consider initial conditions and subsequent projected disturbances from larger sets \mathcal{C}_d , which can lead to the need for larger bounds on the projected disturbance, and a repetition of this cycle. This is particularly problematic noting that the data-driven construction of an uncertainty function in Section 3.4 revealed that proximity to training data is important for constructing meaningful bounds on the projected disturbance signal in (3.101). This means we need data further and further outside of the set \mathcal{C} if we are to establish a viable bound \bar{d} , which may not be tractable to acquire safely.

I will use uncertainty functions (Definition 22) to establish conditions yielding a function $\gamma \in \mathcal{K}$ and a value of \bar{d} such that the projected disturbance is bounded from any initial condition \mathbf{x}_0 in the set $\mathcal{C}_{\bar{d}}$ defined using the function γ as in (3.104). While I will not establish this in the proof, it is straightforward to show using the same proof technique that this function γ can be used to evaluate the sets that are kept forward invariant if the projected disturbance is bounded by a value of $d \in [0, \bar{d}]$, capturing the behavior of *PSSf* if the actual projected disturbance the system experiences is bounded by a lower value than the worst-case bound \bar{d} . Denote the $-c$ -superlevel set of the CBF h as:

$$\mathcal{C}(c) = \{\mathbf{x} \in E \mid h(\mathbf{x}) \geq -c\}, \quad (3.105)$$

noting that $\mathcal{C}(0) = \mathcal{C}$. I denote the closure of $\mathcal{C}(c)$ in \mathbb{R}^n by $\bar{\mathcal{C}}(c)$, and for simplicity, assume that there exists some $\bar{c} \in \mathbb{R}_{>0}$ such that $\bar{\mathcal{C}}(c) \subset E$ for all

$c \in [0, \bar{c})$. This is generally a weak assumption in many practical cases. This leads to the following theorem:

Theorem 21. *Let $\Delta : E \rightarrow \mathcal{P}(\mathbb{R} \times \mathbb{R}^m)$ be a uncertainty function and suppose there exists a $c \in (0, \bar{c})$ such that:*

$$\sup_{\mathbf{x} \in \mathcal{C}(c)} \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}) \leq -\alpha(-c). \quad (3.106)$$

Then for any initial condition $\mathbf{x}_0 \in \mathcal{C}(c)$, we have that the projected disturbance in (3.101) satisfies $|\delta(t)| \leq \bar{d} = -\alpha(-c)$ for all $t \in I(\mathbf{x}_0)$, and the function $\gamma \in \mathcal{K}$ appearing in (3.104) is defined as:

$$\gamma(r) = -\alpha^{-1}(-r). \quad (3.107)$$

Proof. First observe that by the continuity of h , the set $\mathcal{C}(c)$ is closed, such that $\bar{\mathcal{C}}(c) = \mathcal{C}(c)$. Consider a point $\mathbf{x} \in \partial\mathcal{C}(c)$, noting that this implies $\mathbf{x} \in E$ by our assumption that $\bar{\mathcal{C}}(c) \subset E$. We can show by continuity of h that $\mathbf{x} \in \partial\mathcal{C}(c)$ implies that $h(\mathbf{x}) = -c$. As h is a CBF on \mathcal{C} for the nominal model (3.1), we have:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}) > -\alpha(h(\mathbf{x})) = -\alpha(-c) > 0, \quad (3.108)$$

as $\alpha \in \mathcal{K}^e$. This implies that $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$, and thus $-c$ is a regular value of the function h . Now, we have that:

$$\dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \geq -\alpha(h(\mathbf{x})) - \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{x})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{x}), \quad (3.109)$$

$$\geq -\alpha(h(\mathbf{x})) - \sup_{\mathbf{y} \in \mathcal{C}(c)} \sup_{(b, \mathbf{a}) \in \Delta(\mathbf{y})} b + \mathbf{a}^\top \mathbf{k}(\mathbf{y}), \quad (3.110)$$

$$\geq -\alpha(-c) + \alpha(-c) = 0. \quad (3.111)$$

Thus $\dot{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \geq 0$ for any $\mathbf{x} \in \partial\mathcal{C}(c)$, and by Nagumo's theorem [104], [123], we may conclude that the set $\mathcal{C}(c)$ is forward invariant for the true closed-loop system (2.2). Consequently, for any initial condition $\mathbf{x}_0 \in \mathcal{C}(c)$, we have:

$$|\delta(t)| \leq -\alpha(-c) \triangleq \bar{d}, \quad (3.112)$$

for all $t \in I(\mathbf{x}_0)$, such that $\|\delta\|_\infty \leq \bar{d}$. Lastly, because the set $\mathcal{C}(c)$ is forward invariant for the disturbance bound \bar{d} , we wish for $\mathcal{C}_{\bar{d}}$ to be equal to $\mathcal{C}(c)$. Correspondingly, we have that:

$$\gamma(-\alpha(-c)) = c, \quad (3.113)$$

from which we can manipulate to establish that:

$$\gamma(r) = -\alpha^{-1}(-r) \quad (3.114)$$

for any $r \in \mathbb{R}_{\geq 0}$. \square

While not repeated here, similar results to Theorem 19 and Theorem 20 regarding the refinement of uncertainty sets yielding improved bounds and the construction of an uncertainty function capturing residual learning errors when learning a CBF time derivative through data can easily be constructed.

Simulation & Experimental Results

To demonstrate the ability of learning to improve safety guarantees via PSSf, I deploy the episodic learning framework with CBFs presented in Section 3.5 on the Segway system described in Sections 3.3 and 3.5 in simulation⁶ and experimentally. In both simulation and experimentally I run a sequence of episodes to train the estimators \hat{b} and $\hat{\mathbf{a}}$, represented by 4-layer neural networks with ReLU nonlinearities. In each episode the Segway was set to track a desired trajectory in the pitch angle space without violating a barrier function on a portion of its state, using an augmented (CBF-QP) controller. After the sequence of episodes, the Segway was ran once more with a learning-informed (CBF-QP) controller, and the projected disturbance δ as defined in (3.101) was computed. This computation is performed by looking at the difference between the measured value of \dot{h} (numerically approximated) and an estimated value of \dot{h} at sample times t_i . The first estimate I use only uses the nominal-model based estimate of \dot{h} , given by \hat{h} , with no learning models, to produce a projected disturbance signal:

$$\delta(t_i) = \dot{h}(\boldsymbol{\varphi}(t_i), \mathbf{k}(\boldsymbol{\varphi}(t_i))) - \hat{h}(\boldsymbol{\varphi}(t_i), \mathbf{k}(\boldsymbol{\varphi}(t_i))). \quad (3.115)$$

The second estimate I consider uses the estimator \hat{S} produced after the episodic learning process to produce a projected disturbance signal:

$$\delta_l(t_i) = \dot{h}(\boldsymbol{\varphi}(t_i), \mathbf{k}(\boldsymbol{\varphi}(t_i))) - \hat{S}(\boldsymbol{\varphi}(t_i), \mathbf{k}(\boldsymbol{\varphi}(t_i))). \quad (3.116)$$

The worst case disturbances $\bar{\delta} = \|\delta\|_\infty$ and $\bar{\delta}_l = \|\delta_l\|_\infty$ were found, and a lower bound on h for that trajectory was determined using the function $\alpha \in \mathcal{K}^e$ defined as $\alpha(r) = kr$ for $k \in \mathbb{R}_{>0}$, to produce $\gamma(r) = \frac{r}{k}$ as in (3.104).

⁶The simulation code (including an expression for the full system dynamics) can be found at https://github.com/DrewSingletary/cyberpod_sim_ros

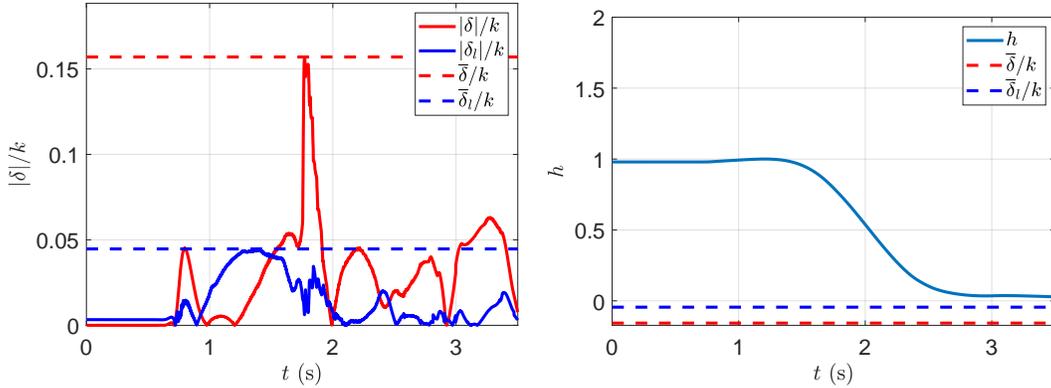


Figure 3.7. Simulation results with Segway platform demonstrating improvement in PSSf behavior. (Left) Absolute value of the projected disturbance along the trajectory without learning models ((3.115), red) and with learning models ((3.116), blue), with learning reducing the worst case projected disturbance ($\bar{\delta}/k$). (Right) The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound is raised with learning (the blue dashed line lies above the red dashed line).

In simulation, the Segway was given a bound on its position in space, constraining it to a one meter distance from its starting location. A CBF on the entire system state that enforces this constraint was generated through the backup controller method [13] (reviewed in Section 5.3). The simulation was done in a Robot Operating System (ROS) based C++ environment similar to [209]. The simulation environment accurately simulates the physical system by adding input delay, sensor noise, and state estimation. The code is identical to that on the physical hardware for the state estimator, controller, and CBF, apart from the ROS functionality. Experimentally, a CBF was specified to limit the pitch angle and pitch angle rate of the Segway to an ellipse about the Segway’s equilibrium state. The desired pitch angle trajectory leads the Segway to tip quickly, leaving the safety set in the absence of the CBF and safety-critical controller.

In both cases, we see that introducing learning estimators into the computation of the projected disturbance decreases the worst case disturbance ($\bar{\delta} > \bar{\delta}_l$). This leads to a greater lower bound on h , and thus a stronger guarantee on the PSSf behavior of the system. I note that the conservative nature of the lower bounds on h arise from the fact that the worst case disturbance $\bar{\delta}$ along the trajectory is used, while the system experiences a much lower projected disturbance throughout most of its evolution. If the worst case disturbance

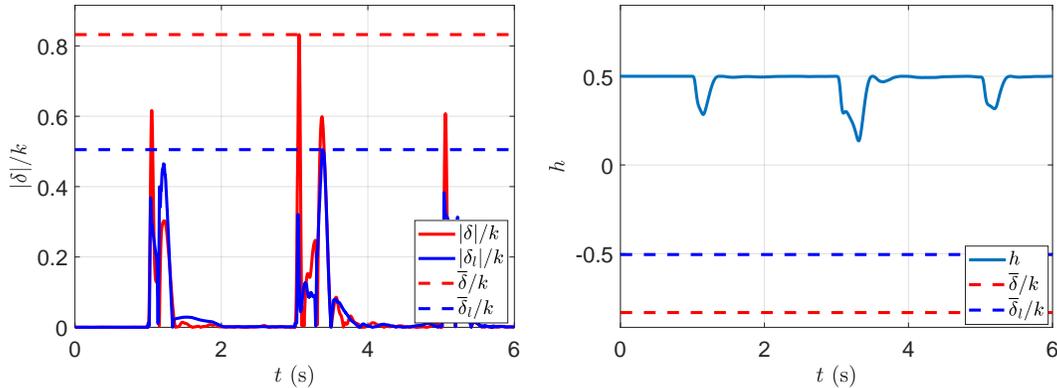


Figure 3.8. Experimental results with Segway platform demonstrating improvement in PSSf behavior. (Left) Absolute value of the projected disturbance along the trajectory without learning models ((3.115), red) and with learning models ((3.116), blue), with learning reducing the worst case projected disturbance ($\bar{\delta}/k$). (Right) The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound is raised with learning (the blue dashed line lies above the red dashed line).

can be reduced (by data-aware control synthesis), stronger guarantees on safety can be made that reflect this behavior.

Conclusion

In this section I have presented an approach for analyzing the impacts of residual learning errors when learning model errors between a nominal model open-loop system and the true open-loop system as they directly impact the time derivative of a CBF h . I propose a notion of robust safety based on ISSf in the form of Projection-to-State Safety (PSSf), which only considers the impacts of disturbances on safety as they affect the time derivative of a CBF. I demonstrate the safety guarantees ensured by PSSf in both simulation and experimentally, with a focus on the fact that learning model errors directly as they impact the time derivative of a CBF can reduce the projected disturbance a system experiences and therefore lead to stronger conclusions on safety.

3.7 Projected Disturbance Learning with Control Barrier Functions

In this section I will present an episodic learning approach for learning a projected disturbance directly as it impacts a CBF time derivative. In contrast to the learning approach taken in Section 3.5, I will not utilize structured learning models that attempt to learn the functions b and \mathbf{a} individually in an effort to

preserve convexity of downstream control synthesis. Rather, I will learn the sum of these functions under a given feedback controller. I will first motivate the reasoning behind this learning approach by considering some challenges that arise from insufficiently rich data when trying to learn using structured models on high-dimensional systems. I will then describe the episodic learning algorithm for directly learning projected disturbances as they impact the time derivative of a CBF. I will conclude with an application of the learning algorithm to a bipedal walking system, both in simulation and experimentally.

The contributions of this section are as follows:

- An analysis of challenges in using the structured learning models considered in Sections 3.3 and 3.5, and an episodic learning algorithm for learning projected disturbances that circumnavigates these challenges.
- An implementation of the proposed algorithm on a bipedal walking platform, both in simulation and experimentally.

The text for this section is adapted from:

N. Csomay-Shanklin, R. K. Cosner, M. Dai, A. J. Taylor and A. D. Ames, “Episodic learning for safe bipedal locomotion with control barrier functions and projection-to-state safety,” in *Proc. 3rd Learning for Dynamics and Control (L4DC)*, vol. 144, Zürich, Switzerland, 2021, pp.1041-1053.

A. J. Taylor participated in the conception of the project, algorithm design, and writing of the article.

Challenges in Learning with Control Barrier Functions

To understand some of the challenges that can arise when trying to learn structured models with insufficiently rich data, I will consider a simple single-input example. Consider a data point $((\mathbf{x}_i, u_i), \dot{h}_i)$ as in (3.85), and suppose $u_i \neq 0$. The model error we are trying to reduce can be expressed as:

$$\tilde{h}_i = \dot{h}_i - \hat{h}(\mathbf{x}_i, u_i) = b(\mathbf{x}_i) + a(\mathbf{x}_i)u_i. \quad (3.117)$$

This can be represented in an affine form as:

$$\tilde{h}_i = \begin{bmatrix} u_i & 1 \end{bmatrix} \begin{bmatrix} a(\mathbf{x}_i) \\ b(\mathbf{x}_i) \end{bmatrix}. \quad (3.118)$$

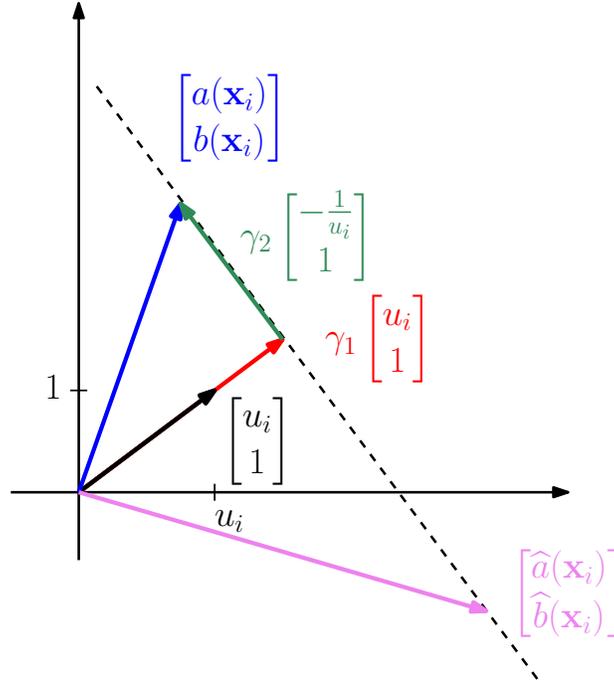


Figure 3.9. Schematic for challenges in learning CBF time derivatives with structured models.

The vectors $(u_i, 1)$ and $(a(\mathbf{x}_i), b(\mathbf{x}_i))$ are visualized in Figure 3.9 (black and blue, respectively). We can decompose the vector $(a(\mathbf{x}_i), b(\mathbf{x}_i))$ as follows:

$$\begin{bmatrix} a(\mathbf{x}_i) \\ b(\mathbf{x}_i) \end{bmatrix} = \gamma_1 \begin{bmatrix} u_i \\ 1 \end{bmatrix} + \gamma_2 \begin{bmatrix} -\frac{1}{u_i} \\ 1 \end{bmatrix}, \quad (3.119)$$

for some $\gamma_1, \gamma_2 \in \mathbb{R}$, noting we have used an orthogonal basis. This decomposition can be seen in Figure 3.9 (red and green vectors). Noting this decomposition, we have that:

$$\tilde{h}_i = \gamma_1 \left\| \begin{bmatrix} u_i \\ 1 \end{bmatrix} \right\|^2. \quad (3.120)$$

This indicates that in our data point we have information about the constant γ_1 , but we do not have information about γ_2 . Suppose that we build accurate learning models $\hat{b} : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\hat{a} : \mathbb{R}^n \rightarrow \mathbb{R}$ such that:

$$\tilde{h}_i \approx \hat{b}(\mathbf{x}_i) + \hat{a}(\mathbf{x}_i)u_i. \quad (3.121)$$

This can be represented in affine form as:

$$\tilde{h}_i \approx \begin{bmatrix} u_i & 1 \end{bmatrix} \begin{bmatrix} \hat{a}(\mathbf{x}_i) \\ \hat{b}(\mathbf{x}_i) \end{bmatrix}. \quad (3.122)$$

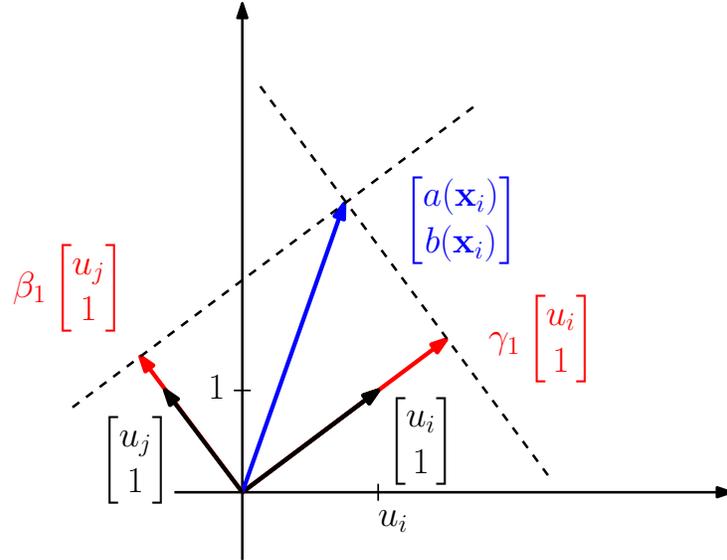


Figure 3.10. Schematic for introducing additional data for resolving challenges in learning CBF time derivatives with structured models.

We can perform a similar decomposition on the estimators:

$$\begin{bmatrix} \widehat{a}(\mathbf{x}_i) \\ \widehat{b}(\mathbf{x}_i) \end{bmatrix} = \widehat{\gamma}_1 \begin{bmatrix} u_i \\ 1 \end{bmatrix} + \widehat{\gamma}_2 \begin{bmatrix} -\frac{1}{u_i} \\ 1 \end{bmatrix}, \quad (3.123)$$

for some $\widehat{\gamma}_1, \widehat{\gamma}_2 \in \mathbb{R}$. What this reveals is that an accurate learning model only needs to satisfy $\widehat{\gamma}_1 \approx \gamma_1$, and there are no requirements on $\widehat{\gamma}_2 \approx \gamma_2$. This can be seen in Figure 3.9, where the estimator (pink) satisfies $\widehat{\gamma}_1 = \gamma_1$ (lies on the dotted black line), but is dramatically different than the true value of the error terms (the blue vector). Consequently, inserting this estimator into the optimization-based controllers as in Section 3.3 and 3.5 can lead to poor performance and failure to satisfy stability and safety requirements.

Reducing this problem can be achieved through richness in the inputs composing the data sets. Suppose that we have another data point at the same state, but with a different input $u_j \neq 0$, and that the learning models are trained accurately on this data to yield:

$$\widetilde{h}_j \approx \widehat{b}(\mathbf{x}_i) + \widehat{a}(\mathbf{x}_i)u_j. \quad (3.124)$$

This additional data point is visualized in Figure 3.10. Considering the two

data points together, we have that:

$$\begin{bmatrix} \tilde{h}_i \\ \tilde{h}_j \end{bmatrix} \approx \underbrace{\begin{bmatrix} u_i & 1 \\ u_j & 1 \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \hat{a}(\mathbf{x}_i) \\ \hat{b}(\mathbf{x}_i) \end{bmatrix}. \quad (3.125)$$

Accurately identifying the values of $b(\mathbf{x}_i)$ and $a(\mathbf{x}_i)$ with the corresponding estimators can thus be seen as solving a least-squares problem locally at the state \mathbf{x}_i . While learning is often done numerically through gradient-based updates rather than taking a (psuedo)inverse of the matrix \mathbf{A} , the numerical stability in solving the learning problem accurately is effectively determined by the condition number of the matrix \mathbf{A} . If the input $u_j = -u_i$, then the $\text{cond}(\mathbf{A}) = 1$. In practice, taking dramatically different inputs at the same state may be unsafe or damaging to a system, making it difficult to acquire data that leads to a well conditioned matrix \mathbf{A} . If instead $u_j = u_i + \epsilon$ for some small non-zero value $\epsilon \in \mathbb{R}$, we have that:

$$\lim_{\epsilon \rightarrow 0} \text{cond}(\mathbf{A}) = \infty. \quad (3.126)$$

Thus small perturbations to a nominal input do not significantly improve the numerical stability of the estimation problem at this state.

This highlights an important challenge in learning with structured models that ensure an affine dependence on the control input. Data collection that does not damage a system is often limited to small perturbations around baseline inputs which are already stabilizing or safe, but this does not provide strong numerical foundations for the downstream learning problem. Moreover, this challenge is exacerbated for high-dimensional systems with multiple inputs. The same geometric decomposition can be performed in high dimensions, revealing the need for $m + 1$ data points with varying inputs at the same state to make sure the learning problem is well conditioned. Additionally, for high-dimensional systems, such as bipedal walking platforms, there are strong requirements on what the input must be at a given state to achieve walking, prohibiting large exploratory perturbations to inputs. In Section 3.8 I will explore how this challenge can be approached through robust convex optimization. In the remainder of this section, I will consider learning projected disturbances without using structured learning models in an effort to avoid this problem.

Learning Projected Disturbances

To motivate learning projected disturbances with unstructured models, I will first indicate some challenges that can arise, and what assumptions justify taking such a learning approach. First, note that for a given controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E , in the absence of learning models \hat{b} and $\hat{\mathbf{a}}$, the value of the projected disturbance the system experiences at any given state is defined in (3.101) and given by:

$$\delta(\mathbf{x}) = b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}). \quad (3.127)$$

Importantly, this projected disturbance depends on the particular controller \mathbf{k} . If another controller $\mathbf{k}' : E \rightarrow \mathbb{R}^m$ was designed to account for this projected disturbance, then the value of the new projected disturbance would be:

$$\delta'(\mathbf{x}) = b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{k}'(\mathbf{x}). \quad (3.128)$$

This introduces a type of feedback loop in which modifying the controller based on an estimate of the projected disturbance leads to a new projected disturbance for which the estimator may not be accurate.

To understand why one may still consider learning the projected disturbance directly, we will understand the impacts of introducing such a learning model into the optimization-based controllers in Section 3.5. In particular, the nominal model-based (CBF-QP) controller can be expressed as [21]:

$$\mathbf{k}_{\text{CBF}}(\mathbf{x}) = \mathbf{k}_{\text{nom}}(\mathbf{x}) + \lambda(\mathbf{x}) \underbrace{\left(\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \hat{\mathbf{g}}(\mathbf{x}) \right)^\top}_{L_{\hat{\mathbf{g}}}h(\mathbf{x})^\top}, \quad (3.129)$$

for a function $\lambda : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ that is locally Lipschitz continuous on E . Suppose that we have a function $\hat{\delta} : E \rightarrow \mathbb{R}$ that is locally Lipschitz continuous on E , and we define a (CBF-QP) controller $\mathbf{k}'_{\text{CBF}} : E \rightarrow \mathbb{R}^m$ as:

$$\begin{aligned} \mathbf{k}'_{\text{CBF}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && (\hat{\delta}\text{-CBF-QP}) \\ \text{s.t. } &\hat{h}(\mathbf{x}, \mathbf{u}) + \hat{\delta}(\mathbf{x}) \geq -\alpha(h(\mathbf{x})). \end{aligned}$$

This controller can be expressed as:

$$\mathbf{k}'_{\text{CBF}}(\mathbf{x}) = \mathbf{k}_{\text{nom}}(\mathbf{x}) + \lambda'(\mathbf{x}) L_{\hat{\mathbf{g}}}h(\mathbf{x})^\top, \quad (3.130)$$

for a function $\lambda' : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ that is locally Lipschitz continuous on E . It can be shown that:

$$\widehat{\delta}(\mathbf{x}) \geq 0 \implies \lambda'(\mathbf{x}) \leq \lambda(\mathbf{x}), \quad (3.131)$$

$$\widehat{\delta}(\mathbf{x}) \leq 0 \implies \lambda'(\mathbf{x}) \geq \lambda(\mathbf{x}). \quad (3.132)$$

Denoting $L_{\widehat{\mathbf{f}}}h(\mathbf{x}) = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\widehat{\mathbf{f}}(\mathbf{x})$, the time derivative of the CBF h using both of these controllers is given by:

$$\begin{aligned} \dot{h}(\mathbf{x}, \mathbf{k}_{\text{CBF}}(\mathbf{x})) &= L_{\widehat{\mathbf{f}}}h(\mathbf{x}) + b(\mathbf{x}) + (L_{\widehat{\mathbf{g}}}h(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top)\mathbf{k}_{\text{nom}}(\mathbf{x}) \\ &\quad + \lambda(\mathbf{x})L_{\widehat{\mathbf{g}}}h(\mathbf{x})L_{\widehat{\mathbf{g}}}h(\mathbf{x})^\top + \lambda(\mathbf{x})\mathbf{a}(\mathbf{x})^\top L_{\widehat{\mathbf{g}}}h(\mathbf{x}), \end{aligned} \quad (3.133)$$

and:

$$\begin{aligned} \dot{h}(\mathbf{x}, \mathbf{k}'_{\text{CBF}}(\mathbf{x})) &= L_{\widehat{\mathbf{f}}}h(\mathbf{x}) + b(\mathbf{x}) + (L_{\widehat{\mathbf{g}}}h(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top)\mathbf{k}_{\text{nom}}(\mathbf{x}) \\ &\quad + \lambda'(\mathbf{x})L_{\widehat{\mathbf{g}}}h(\mathbf{x})L_{\widehat{\mathbf{g}}}h(\mathbf{x})^\top + \lambda'(\mathbf{x})\mathbf{a}(\mathbf{x})^\top L_{\widehat{\mathbf{g}}}h(\mathbf{x})^\top. \end{aligned} \quad (3.134)$$

Let us suppose that:

$$\|L_{\widehat{\mathbf{g}}}h(\mathbf{x})\| \geq \|\mathbf{a}(\mathbf{x})\|, \quad (3.135)$$

for all $\mathbf{x} \in E$. This assumption states that the impact of model error in how the input affects the time derivative of h is less significant than the impact of the input on the time derivative of h captured by our nominal model. Consequently, we have that:

$$L_{\widehat{\mathbf{g}}}h(\mathbf{x})L_{\widehat{\mathbf{g}}}h(\mathbf{x})^\top \geq \mathbf{a}(\mathbf{x})^\top L_{\widehat{\mathbf{g}}}h(\mathbf{x})^\top. \quad (3.136)$$

Now let us suppose that for a particular state $\mathbf{x} \in E$, we have $\widehat{\delta}(\mathbf{x}) \leq 0$, such that $\lambda'(\mathbf{x}) \geq \lambda(\mathbf{x})$. If the function $\widehat{\delta}$ is an (accurate) estimator for the projected disturbance (3.127) using the controller \mathbf{k}_{CBF} , this indicates that:

$$\dot{h}(\mathbf{x}, \mathbf{k}_{\text{CBF}}(\mathbf{x})) \leq \widehat{\dot{h}}(\mathbf{x}, \mathbf{k}_{\text{CBF}}(\mathbf{x})), \quad (3.137)$$

such that the system may not satisfy the safety requirement:

$$\dot{h}(\mathbf{x}, \mathbf{k}_{\text{CBF}}(\mathbf{x})) \geq -\alpha(h(\mathbf{x})). \quad (3.138)$$

Because $\lambda'(\mathbf{x}) \geq \lambda(\mathbf{x})$, we have that:

$$\dot{h}(\mathbf{x}, \mathbf{k}'_{\text{CBF}}(\mathbf{x})) \geq \dot{h}(\mathbf{x}, \mathbf{k}_{\text{CBF}}(\mathbf{x})), \quad (3.139)$$

such that the introduction of $\widehat{\delta}$ to the controller is doing more to ensure the safety requirement is met. In this way by learning the projected disturbance

using one controller, we can construct another controller that improves upon the previous controller with respect to safety. As this modifies the projected disturbance, this process can be iterated on to identify the new projected disturbance and improve the safety of the controller further.

More concretely, consider the following setting: in an experiment, the system is allowed to evolve forward in time from a particular initial condition $\mathbf{x}_0 \in E$ under a given controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E . During this experiment, state data $\mathbf{x}_i \in E$ is collected at discrete sample-times $t_i \in \mathbb{R}_{\geq 0}$, providing a discrete-time history of the CBF, $h_i \triangleq h(\mathbf{x}_i)$. This time history is smoothed and numerically differentiated to compute an approximate time history of the true value of the time derivative of the CBF, $\dot{h}_i \triangleq \dot{h}(\mathbf{x}_i, \mathbf{k}(\mathbf{x}_i))$. This yields a collection of input-output pairs:

$$\mathfrak{D} = \left\{ ((\mathbf{x}_i, \mathbf{k}(\mathbf{x}_i)), \dot{h}_i) \right\}_{i=1}^N \subseteq (E \times \mathbb{R}^m) \times \mathbb{R}. \quad (3.140)$$

Given a nonlinear function class $\mathcal{H} : \mathbb{R}^n \rightarrow \mathbb{R}$ and a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, a learning problem can be specified as finding a function $\hat{\delta} \in \mathcal{H}$ to estimate δ via empirical risk minimization (ERM):

$$\inf_{\hat{\delta} \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^N \mathcal{L} \left(\hat{h}(\mathbf{x}_i, \mathbf{k}(\mathbf{x}_i)) + \hat{\delta}(\mathbf{x}_i), \dot{h}_i \right). \quad (\text{ERM})$$

A controller can then be synthesized as in ($\hat{\delta}$ -CBF-QP). I directly build upon the episodic learning framework from Sections 3.3 and 3.5 as outlined in Algorithm 3. In each episode, the algorithm runs the current controller to collect data, learns a new estimator $\hat{\delta}$ using the newly collected data, and synthesizes a new controller. A notable difference is that data is not aggregated during this process. This is because data from a previous episode corresponds to a projected disturbance using a different controller, and thus would conflict with data collected using a modified controller.

Bipedal System Application

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in conference proceedings, and edited by Andrew Taylor for their appearance in this thesis.

We will now consider learning projected disturbances in the setting of bipedal locomotion on stepping-stones. We will briefly introduce the theory of bipedal

Algorithm 3 Projected Disturbance Learning (PDL)

Require: CBF h , CBF derivative estimate \hat{h} , model class \mathcal{H} , loss function \mathcal{L} , nominal model-based (CBF-QP) controller $\mathbf{k}_0 = \mathbf{k}_{\text{CBF}}$, number of episodes T

```

for  $k = 1, \dots, T$  do
   $\mathbf{x}_0 \leftarrow \text{sample}E$  ▷ Sample initial condition
   $\mathcal{D} \leftarrow \text{experiment}(\mathbf{x}_0, \mathbf{k}_{k-1})$  ▷ Run experiment
   $\hat{\delta}_k \leftarrow \text{ERM}(\mathcal{H}, \mathcal{L}, \mathcal{D}, \hat{h})$  ▷ Fit estimator
   $\mathbf{k}_k \leftarrow \text{augment}(\mathbf{k}_0, \hat{\delta}_k)$  ▷ Update controller
end for
return  $\hat{\delta}_T, \mathbf{k}_T$ 

```

locomotion and then describe the barrier function formulations which allow achieving safe bipedal locomotion across stepping-stones. A deeper exploration of this material may be found in [210].

The bipedal robotic system we consider is the AMBER-3M robotic platform seen in Figure 3.11, modeled as an underactuated, planar five-link robot with point feet [211] and physical parameters reported in [212, Table 1]. The configuration coordinates $\mathbf{q} \in \mathcal{Q} \subset \mathbb{R}^5$ are given by $\mathbf{q} = [q_{sf}, q_{sk}, q_{sh}, q_{nsh}, q_{nsk}]^\top$, with stance foot angle q_{sf} , stance knee angle q_{sk} , stance hip angle q_{sh} , non-stance hip angle q_{nsh} and non-stance knee angle q_{nsk} . The full system state is given by $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{Q} \times \mathbb{R}^5$. The continuous-time equations of motion are given by the Euler-Lagrange equations as in (2.5). Note that this is the pinned model of the robot dynamics; for the unpinned model, refer to [213]. For AMBER-3M, the number of inputs is one fewer than the degrees of freedom, meaning the system has one degree of underactuation.

Taking $p^v : \mathcal{Q} \rightarrow \mathbb{R}$ to represent the vertical position (height) of the swing foot, the admissible states are given by the domain $\mathcal{D} = \{(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{Q} \times \mathbb{R}^5 \mid p^v(\mathbf{q}) \geq 0\}$. The switching surface on which the impact events occur, also known as the guard, is defined by:

$$\mathcal{S} = \{(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{Q} \times \mathbb{R}^5 \mid p^v(\mathbf{q}) = 0, \dot{p}^v(\mathbf{q}, \dot{\mathbf{q}}) < 0\} \subset \mathcal{D}. \quad (3.141)$$

Assuming that an impact at happens at time $t \in \mathbb{R}_{\geq 0}$, the impact dynamics⁷

⁷I do not discuss the notion of solutions for hybrid systems in this thesis, but an introduction to the topic can be found in [214]. The relevant detail for this example is that the corresponding solution φ will be discontinuous at impact times.

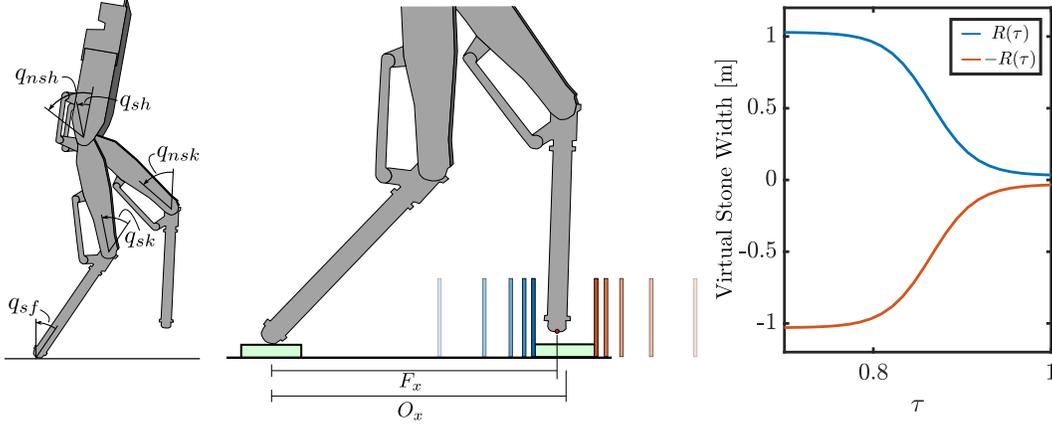


Figure 3.11. (Left) Schematic diagram of the AMBER-3M robot with position coordinates. (Center) Schematic of the foot placement in the stepping-stone problem. The boundaries of virtual stepping-stones are captured via the blue and orange vertical lines. (Right) Virtual stepping stone width as function of the phase variable $\tau(\mathbf{q})$.

[215] are defined by a reset map $\Delta : \mathcal{S} \rightarrow \mathcal{D}$ relating pre-impact states $\mathbf{x}^- \triangleq \lim_{\tau \rightarrow t^-} \varphi(\tau)$ and post-impact states $\mathbf{x}^+ = \lim_{\tau \rightarrow t^+} \varphi(\tau)$ via $\mathbf{x}^+ = \Delta(\mathbf{x}^-)$. Combining these concepts yields the following hybrid control system:

$$\mathcal{HC} = \begin{cases} \dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} & \mathbf{x}^- \in \mathcal{D} \setminus \mathcal{S}, \\ \mathbf{x}^+ = \Delta(\mathbf{x}^-) & \mathbf{x}^- \in \mathcal{S}. \end{cases} \quad (3.142)$$

Control of bipedal systems uses a phasing variable, $\tau : \mathcal{Q} \rightarrow [0, 1]$, given by:

$$\tau(\mathbf{q}) = \frac{\delta_{\text{hip}}^+(\mathbf{q}) - \delta_{\text{hip}}^+}{\delta_{\text{hip}}^- - \delta_{\text{hip}}^+}, \quad (3.143)$$

where $\delta_{\text{hip}} : \mathcal{Q} \rightarrow \mathbb{R}$ defined as $\delta_{\text{hip}}(\mathbf{q}) = [-l_t - l_f, -l_f, 0, 0, 0]\mathbf{q}$ is the linearized hip position with l_t and l_f the length of the robots tibia and femur, respectively. The constants δ_{hip}^+ and δ_{hip}^- are the linearized hip positions at the beginning and the end of a step, ensuring that $\tau(\mathbf{q})$ increases monotonically in time within a step. Desired trajectories resulting in walking gaits for the robot can be synthesized via a hybrid zero dynamics framework [107], [210].

We are now equipped to define the relative degree 2 ([112]) outputs $\mathbf{y} : \mathcal{Q} \times \mathbb{R}^p \rightarrow \mathbb{R}^4$ as the difference between the actual output $\mathbf{y}_a : \mathcal{Q} \rightarrow \mathbb{R}^4$ and the desired output trajectory $\mathbf{y}_d : \mathbb{R} \times \mathbb{R}^p \rightarrow \mathbb{R}^4$:

$$\mathbf{y}(\mathbf{q}, \boldsymbol{\alpha}) \triangleq \mathbf{y}_a(\mathbf{q}) - \mathbf{y}_d(\tau(\mathbf{q}), \boldsymbol{\alpha}), \quad (3.144)$$

with $\boldsymbol{\alpha} \in \mathbb{R}^p$ being the coefficients of a Bézier polynomial coming from the trajectory generation step. The actual output is given by the actuated coordinates:

$$\mathbf{y}_a(\mathbf{q}) = \begin{bmatrix} \mathbf{0}_{4 \times 1} & \mathbf{I}_{4 \times 4} \end{bmatrix} \mathbf{q}. \quad (3.145)$$

The nominal controller for this system is then given by the proportional-derivative controller $\mathbf{k}_{\text{nom}}(\mathbf{x}) = -\mathbf{K}_P \mathbf{y}(\mathbf{q}) - \mathbf{K}_D \dot{\mathbf{y}}(\mathbf{q}, \dot{\mathbf{q}})$ with proportional gain $\mathbf{K}_P \in \mathbb{S}_{>0}^4$ and derivative gain $\mathbf{K}_D \in \mathbb{S}_{>0}^4$.

The stepping-stone problem uses *virtual* stepping-stones, which shrink over the course of a step to confine foot placement to a targeted stone [22]. The CBFs used to specify these foot position constraints are given by:

$$h_1(\mathbf{q}) = R(\tau(\mathbf{q})) - (O_x - F_x(\mathbf{q})), \quad (3.146)$$

$$h_2(\mathbf{q}) = R(\tau(\mathbf{q})) + (O_x - F_x(\mathbf{q})), \quad (3.147)$$

where $F_x(\mathbf{q})$ is the horizontal position of the swing foot, $O_x > 0$ is the horizontal position of the center of stepping-stone, and the virtual stone width is given by the function $R : \mathbb{R} \rightarrow \mathbb{R}$ defined as:

$$R(\tau(\mathbf{q})) = \frac{ar - 1}{1 + ar(e^{-\lambda(\tau(\mathbf{q})-1)} - 1)} + 1 + r, \quad (3.148)$$

where $\lambda \in \mathbb{R}_{>0}$ determines the decay rate of the barrier function, $(1+a)r$ is half of the targeted stone width, and $1+r$ defines the half the width of the virtual stepping-stone when $\tau = 0$. These functions are visualized in Figure 3.11. The safety constraints can be interpreted as keeping the swing foot horizontal position in an interval centered at the middle of the stepping-stone, where the interval shrinks as τ increases. As this formulation of CBFs is position-based and therefore relative degree two, we employ the exponential CBF (ECBF) extension technique [216] to both CBFs to attain the relative degree 1 CBFs: $h_{e,i}(\mathbf{x}) \triangleq L_{\hat{\mathbf{f}}} h_i(\mathbf{x}) + \alpha_e h_i(\mathbf{q})$ with $\alpha_e \in \mathbb{R}_{>0}$. The final Stepping Stone QP (SS-QP) controller combines the robustifying term of the ($\hat{\delta}$ -CBF-QP) controller with the stepping-stone ECBF extensions of (3.146) and (3.147):

$$\mathbf{k}(\mathbf{x}) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \quad (\text{SS-QP})$$

$$\text{s.t. } L_{\hat{\mathbf{f}}}^2 h_1(\mathbf{x}) + L_{\hat{\mathbf{g}}} L_{\hat{\mathbf{f}}} h_1(\mathbf{x}) \mathbf{u} + \alpha_e L_{\hat{\mathbf{f}}} h_1(\mathbf{x}) + \hat{\delta}_1(\mathbf{x}) \geq -\alpha(h_{e,1}(\mathbf{x})),$$

$$L_{\hat{\mathbf{f}}}^2 h_2(\mathbf{x}) + L_{\hat{\mathbf{g}}} L_{\hat{\mathbf{f}}} h_2(\mathbf{x}) \mathbf{u} + \alpha_e L_{\hat{\mathbf{f}}} h_2(\mathbf{x}) + \hat{\delta}_2(\mathbf{x}) \geq -\alpha(h_{e,2}(\mathbf{x})).$$

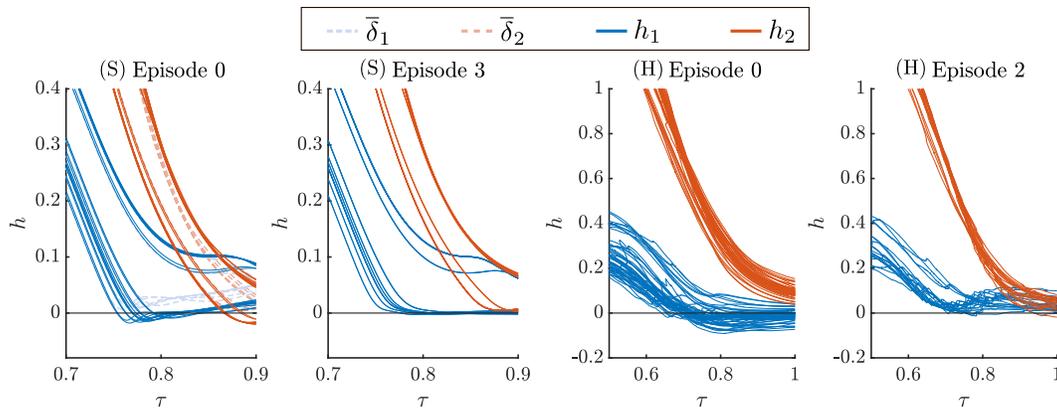


Figure 3.12. Simulation (S) and Hardware (H) data where model mismatch causes violations. (Far Left) Simulation where the barrier functions h_1 (solid blue) and h_2 (solid orange) are enforced via a nominal model-based (CBF-QP) controller. A conservative controller using a worst-case upper bound on the projected disturbance is also shown (dashed blue, dashed orange), which results in more conservative behavior over many steps. (Mid Left) After three episodes of learning using the (SS-QP) controller in simulation, the maximum barrier violation decreases from 2.0 to 0.3 [cm]. (Mid Right) Hardware where the barrier functions h_1 (blue) and h_2 (orange) are enforced via a nominal model-based (CBF-QP) controller. (Far Right) After two episodes of learning on hardware, the maximum barrier violation decreases from 9.2 to 1.9 [cm] using the (SS-QP) controller.

Simulation & Experimental Results

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in conference proceedings, and edited by Andrew Taylor for their appearance in this thesis.

We now apply the episodic learning algorithm in Algorithm 3, to the AMBER-3M platform in both simulation with artificial model error and on hardware with the model error inherent to real-world systems. In each instance the estimator $\hat{\delta}$ was implemented as a neural network with two hidden layers of 50 hidden units using the ReLU activation function. The network was trained minimizing mean absolute error using mini-batch gradient descent. Mean absolute error was chosen over other loss functions for its robustness to outliers. The SS-QP controller was deployed in the RaiSim [217] simulation environment and on the AMBER-3M hardware platform.

The controllers and learning algorithm were first validated in simulation. Model error was introduced by increasing the inertia of all limbs on the nominal model by a factor of ten while maintaining constant mass to produce the true system. Due to the underactuated nature of the robot and the relationship between step length and zero dynamics stability, not every set of stepping stones

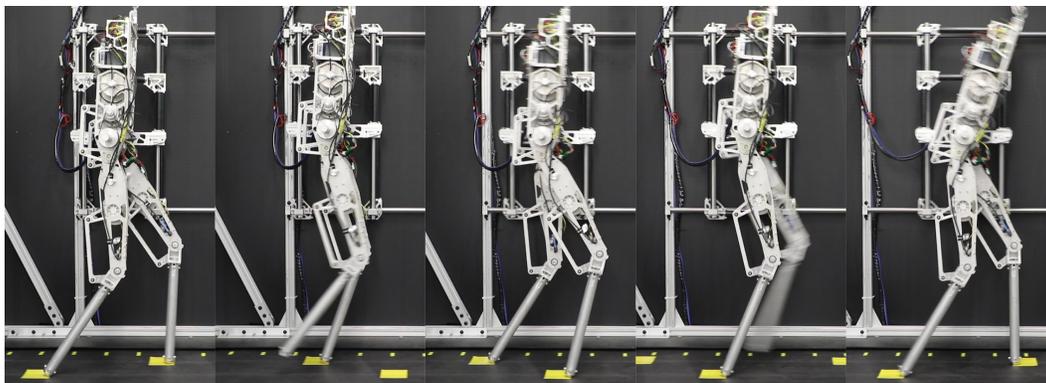


Figure 3.13. Gait tiles for Episode 2 of learning showing the AMBER-3M robot safely traversing a set of stepping stones. Notice the change in step width and added lean of the torso induced by the barrier functions. Video can be found at <https://vimeo.com/481809664>

is navigable, even if safety is perfectly enforced with respect to the CBFs. Therefore, a feasible stepping stone configuration was first generated for the robot to traverse with stones of 4 [cm] in width. Without knowledge of the modified model ($\hat{\delta}_1(\mathbf{x}) = \hat{\delta}_2(\mathbf{x}) = 0$ for all $\mathbf{x} \in E$), the controller did not satisfy the CBF constraints, resulting in a maximum violation at foot placement of 2.0 [cm], causing the robot to miss the stepping stone and fall over. Three episodes of the PDL algorithm were run, after which the maximum violation was reduced to be 0.3 [cm], only 15% of the original violation. Additionally, a controller using a worst-case global upper bound on the projected disturbance was implemented, which ensured safety but resulted in extremely conservative behavior, resulting in poor qualitative walking, i.e. harsh foot strikes and an over-bending torso. A comparison of the barrier functions h_1 and h_2 over the steps with these controllers can be seen in Figure 3.12.

The same nominal model for the robot was used in the hardware experiments as in simulation, with model uncertainty presenting itself as significant friction in the joints, as well as imperfect mass and inertia measurements. The PDL algorithm was implemented on the AMBER-3M robot across a sequence of two episodes. The controllers ran on an off-board i7-6700HQ CPU at 2.6 [GHz] with 16 GB RAM, which computed desired torques and communicated to ELMO motor drivers on the 137 [cm] tall, 22 [kg] robot. The motor driver communication ran at 2 [kHz], and the (SS-QP) ran at 1 [kHz]. The stepping stone configuration was specified to the controller with stones of 8 [cm] in width. The nominal model-based (CBF-QP) controller resulted in a maximum

violation of the barriers of 9.2 [cm] due to model error. After running the PDL algorithm for two episodes, the maximum violation of the barriers was 1.9 [cm], only 21% of the original violation, as depicted in Figure 3.12. Gait tiles for this improved traversal of the stepping stones are shown in Figure 3.13.

Conclusion

In this section I have presented an approach for learning projected disturbances in the time derivative of a CBF with unstructured learning models. I begin by analyzing some of the challenges that can occur with learning projected disturbances with structured learning models that preserve a control-affine relationship. I then motivate why directly learning a projected disturbance produced using a particular controller and modifying the controller can lead to improved safety, and produce an episodic learning algorithm. I then present the problem of bipedal walking on stepping-stones, and present results from deploying the proposed learning algorithm in simulation and experimentally on the AMBER-3M bipedal platform.

3.8 Data-Driven Nonlinear Control

In this section I will present a data-driven control framework. This framework directly incorporates data into a convex optimization problem that ensures that a CLF or CBF condition is met for all possible model errors permitted by the data. This approach is fundamentally different than the supervised learning approaches in Sections 3.3, 3.5, and 3.7. Rather than building learning models offline and deploying them in a controller online, which typically leads to residual learning errors (that may be quite large as discussed in Section 3.7) that impact stability or safety (as studied in Sections 3.4 and 3.6), this approach actively chooses control inputs that work for all possible model errors allowed by a data set. A key result of this work will be the ability to synthesize stabilizing or safe controllers without requiring the degree of variety in inputs in a data set that was needed to ensure a well-conditioned learning problem as discussed in Section 3.7. Rather, the use of robust convex optimization will allow relying only on inputs that have been seen in the data set.

I will begin by considering the data-driven characterization of uncertainty in Section 3.4, and establish some intuition behind the approach with geometric visualizations. I will then introduce Control Certificate Functions (CCFs), which is a generalization that captures both CLFs and CBFs and permits

a single unifying framework. Next I will describe how model errors will be characterized, (it will differ slightly than in Section 3.4) and how this leads to a robust controller specified by a second-order cone program (SOCP), which is a convex optimization problem of a more complex class than the quadratic programs considered in Chapter 2 and Sections 3.3-3.7. After this, I will address questions regarding the boundedness of the possible model errors given properties of a data set, and provide a thorough proof of the feasibility of the robust controller by generalizing the conditions required by CLFs and CBFs to accommodate a set of model errors. Lastly, I will present simulation results utilizing the robust controller, demonstrating that performance can be achieved with limited variation of the inputs in the data set.

The contributions of this section are as follows:

- A data-driven control framework that is robust to model errors permitted by a data set through the use of robust convex optimization.
- A characterization of feasibility of a robust controller by generalizing the conditions required of CLFs and CBFs to include a set of model errors.

The text for this section is adapted from:

A. J. Taylor, V. D. Dorobantu, S. Dean, B. Recht, Y. Yue, and A. D. Ames, “Towards robust data-driven control synthesis for nonlinear systems with actuation uncertainty,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, pp.6469-6476, 2021.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Data-Driven Intuition

I will begin by developing some intuition behind how partial characterizations of model error from data can still permit controllers that meet stability or safety requirements. Consider the true open-loop system (2.1), and the nominal model open-loop system (3.1), and suppose there is a single input ($m = 1$). Suppose that \mathbf{x}_e is an unforced or forced equilibrium point of both the true open-loop system (2.1) and the nominal model open-loop system (3.1), and

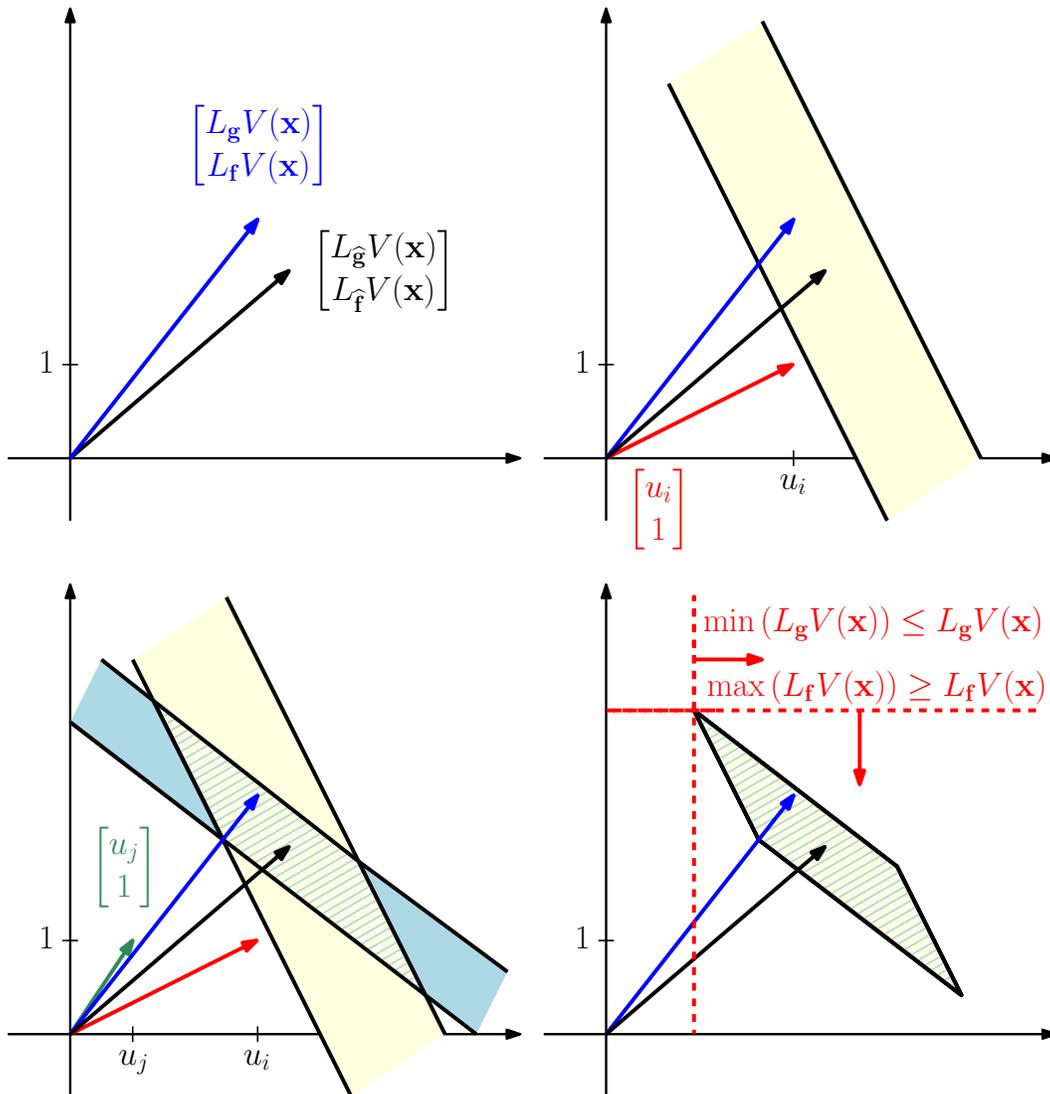


Figure 3.14. Schematic of model error set arising from multiple data points.

the function $V : E \rightarrow \mathbb{R}^n$ is a local exponential CLF for both the true-open system (2.1) and nominal model (3.1) and equilibrium point \mathbf{x}_e . This amounts to the requirements of Assumptions 1 and 2 in Section 3.3. Recall that:

$$\dot{V}(\mathbf{x}, u) = L_f V(\mathbf{x}) + L_g V(\mathbf{x})u, \quad (3.149)$$

$$\hat{V}(\mathbf{x}, u) = L_{\hat{f}} V(\mathbf{x}) + L_{\hat{g}} V(\mathbf{x})u. \quad (3.150)$$

For a given $\mathbf{x} \in D$, the vectors $(L_g V(\mathbf{x}), L_f V(\mathbf{x}))$ and $(L_{\hat{g}} V(\mathbf{x}), L_{\hat{f}} V(\mathbf{x}))$ can be visualized in Figure 3.14 (Top, Left). Model error between the true open-loop system (2.1) and the nominal model (3.1) leads to a mismatch between the two vectors.

Consider the data-driven model error characterization in Section 3.4 in the absence of estimators \widehat{b} and \widehat{a} (such that $\widehat{W}(\mathbf{x}, u) = \widehat{V}(\mathbf{x}, u)$ for all $\mathbf{x} \in D$ and $u \in \mathbb{R}$ and $\epsilon_{\mathcal{H}}$ does not appear). Given a data point $((\mathbf{x}_i, u_i), V_i)$, we have that:

$$|b(\mathbf{x}) + a(\mathbf{x})u_i| \leq \epsilon(\mathbf{x}, \mathbf{x}_i, u_i). \quad (3.151)$$

This data point tells us that the vector $(L_{\mathbf{g}}V(\mathbf{x}), L_{\mathbf{f}}V(\mathbf{x}))$ lies in a strip that runs perpendicular to the vector $(u_i, 1)$, is centered about the vector $(L_{\widehat{\mathbf{g}}}V(\mathbf{x}), L_{\widehat{\mathbf{f}}}V(\mathbf{x}))$, and has a width determined by $\epsilon(\mathbf{x}, \mathbf{x}_i, u_i)$, as visualized in Figure 3.14 (Top, Right). If we consider another data point $((\mathbf{x}_j, u_j), V_j)$ with $u_j \neq u_i$, we similarly have that:

$$|b(\mathbf{x}) + a(\mathbf{x})u_j| \leq \epsilon(\mathbf{x}, \mathbf{x}_j, u_j). \quad (3.152)$$

This again defines another strip that the vector $(L_{\mathbf{g}}V(\mathbf{x}), L_{\mathbf{f}}V(\mathbf{x}))$ must lie in, as seen in Figure 3.14 (Bottom, Left). Thus the vector $(L_{\mathbf{g}}V(\mathbf{x}), L_{\mathbf{f}}V(\mathbf{x}))$ must lie in the intersection of the two strips, which is a compact polytope as visualized in Figure 3.14 (Bottom, Right). In this way, multiple data points (even without dramatically different inputs) can lead to bounded sets of possible model errors, as will be explored later in this section. Intuitively, the greater the difference between u_i and u_j (such as if $u_j = -u_i$), the smaller the set of possible model errors will be. This is a manifestation of the well-conditioning of the learning problem explored in Section 3.7. Still, even a small amount of variation is sufficient to introduce a bounded set of possible model errors, even if it would not lead to a well conditioned learning problem.

Observe that for all values of $L_{\mathbf{f}}V(\mathbf{x})$ and $L_{\mathbf{g}}V(\mathbf{x})$ permitted by the set of possible model errors, we have that:

$$\max(L_{\mathbf{f}}V(\mathbf{x})) \geq L_{\mathbf{f}}V(\mathbf{x}), \quad (3.153)$$

$$0 < \min(L_{\mathbf{g}}V(\mathbf{x})) \leq L_{\mathbf{g}}V(\mathbf{x}), \quad (3.154)$$

as seen in Figure 3.14 (Bottom, Right). Consequently, if a negative input $u \in \mathbb{R}_{<0}$ is chosen to satisfy:

$$\max(L_{\mathbf{f}}V(\mathbf{x})) + \min(L_{\mathbf{g}}V(\mathbf{x}))u \leq -k_3\|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.155)$$

then it satisfies:

$$L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})u \leq -k_3\|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.156)$$

and thus the LF condition needed for local exponential stability is satisfied for the true closed-loop system. In this way, inputs can be chosen based on worst-case values in the possible set of model errors permitted by a data set to ensure the desired stability or safety condition is met. Adding more data (with more variety in the inputs) serves to reduce these worst-case model errors and would yield a potentially less conservative (or more optimal) choice of u , but importantly, is not *needed* for feasibility, just optimality. The choice of inputs that are robust to the worst-case model errors will be rigorously codified via a convex optimization problem later in this section.

Control Certificate Functions and Model Error

To develop a control framework that can flexibly accommodate both CLFs and CBFs, I will use the notion of a *Control Certificate Function (CCF)*, which was first explored in [218], [219]:

Definition 23 (*Control Certificate Function (CCF)*). A function $\mathbf{c} : E \rightarrow \mathbb{R}$ that is continuously differentiable on E is said to be a *Control Certificate Function (CCF)* for the open-loop system (2.1) on the set $\mathcal{D} \subseteq E$ if there exists a function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ such that:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{\mathbf{c}}(\mathbf{x}, \mathbf{u}) \triangleq \inf_{\mathbf{u} \in \mathbb{R}^m} \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{f}(\mathbf{x})}_{L_{\mathbf{f}}\mathbf{c}(\mathbf{x})} + \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{g}(\mathbf{x}) \mathbf{u}}_{L_{\mathbf{g}}\mathbf{c}(\mathbf{x})} < -\alpha(\mathbf{c}(\mathbf{x})), \quad (3.157)$$

for all $\mathbf{x} \in \mathcal{D}$.

It is straightforward to see that if \mathbf{x}_e is an unforced or forced equilibrium point of the open-loop system (2.1) and $V : E \rightarrow \mathbb{R}_{\geq 0}$ is a local exponential CLF for the open-loop system (2.1) and equilibrium point \mathbf{x}_e with a corresponding open set D , then V is a CCF for the open-loop system (2.1) on the set $\mathcal{D} = D \setminus \{\mathbf{x}_e\}$ with a function $\alpha(r) = -\frac{k_3}{k_2}r$. Showing that a CBF $h : E \rightarrow \mathbb{R}$ for the open-loop system (2.1) on a set \mathcal{C} defined as the 0-*super*level set of h is a bit more difficult. It is easier if \mathcal{C} is defined as the 0-*sub*level set, such that $\mathbf{x} \in \mathcal{C} \iff h(\mathbf{x}) \leq 0$, in which case the CBF condition is replaced by:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) < -\alpha(h(\mathbf{x})), \quad (3.158)$$

for some $\alpha \in \mathcal{K}^e$ and all $\mathbf{x} \in E$. It is straightforward to transfer all of the results relating BFs and CBFs to safety to this sign convention, and using

CCFs is done simply to unify CLFs and CBFs under one mathematical tool. In particular, h is a CCF for the open-loop system (2.1) on the set $\mathcal{D} = E$. As was done with both CLFs and CBFs, given a CCF \mathcal{C} for the open-loop system (2.1) on the set $\mathcal{D} \subseteq E$ and a function α , define the point-wise set:

$$K_{\text{CCF}}(\mathbf{x}) \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid \dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}) \leq -\alpha(\mathcal{C}(\mathbf{x}))\}. \quad (3.159)$$

Any nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous E can be modified to take values in the set $K_{\text{CCF}}(\mathbf{x})$ via following controller:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && (\text{CCF-QP}) \\ \text{s.t. } &\nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \leq -\alpha(\mathcal{C}(\mathbf{x})). \end{aligned}$$

Considering a nominal model as in (3.1) and the dynamics decomposition in (3.2), we can see that the time derivative of a CCF for the open-loop system (2.1) on the set \mathcal{D} can be written as:

$$\dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}) = \overbrace{\underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \hat{\mathbf{f}}(\mathbf{x})}_{L_{\hat{\mathbf{f}}}\mathcal{C}(\mathbf{x})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \hat{\mathbf{g}}(\mathbf{x}) \mathbf{u}}_{L_{\hat{\mathbf{g}}}\mathcal{C}(\mathbf{x})}}^{\hat{\mathcal{C}}(\mathbf{x}, \mathbf{u})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \tilde{\mathbf{f}}(\mathbf{x})}_{L_{\tilde{\mathbf{f}}}\mathcal{C}(\mathbf{x})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \tilde{\mathbf{g}}(\mathbf{x}) \mathbf{u}}_{L_{\tilde{\mathbf{g}}}\mathcal{C}(\mathbf{x})}. \quad (3.160)$$

Consequently, even if \mathcal{C} is a CCF for the nominal model (3.1) on the set \mathcal{D} , it is not necessarily a CCF for the open-loop system (2.1) on \mathcal{D} . To ensure that the data-driven control problem is well-posed, I will make the following assumption as was seen in Sections 3.3 and 3.5:

Assumption 4. If the function \mathcal{C} is a CCF for the nominal model (3.1) on the set \mathcal{D} , then it is a CCF for the open-loop system (2.1) on the set \mathcal{D} .

Lastly, I will also make the following assumption:

Assumption 5. The functions $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ are Lipschitz continuous on E with known Lipschitz constants $L_{\tilde{\mathbf{f}}}$ and $L_{\tilde{\mathbf{g}}}$.

Observe that this preceding assumption is about Lipschitz constants on the error terms $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$. This is an important detail because it suggests that developing a better nominal model (reducing the error terms, and (likely) their Lipschitz constants) will make the subsequent constructions less conservative, and thus work invested in developing a good model-based controller before employing data-driven techniques will require the data-driven techniques to accomplish less.

Data-Driven Robust Control Synthesis

Consider a dataset consisting of N tuples of states, inputs, and corresponding state time derivatives, $\mathfrak{D} = \{(\mathbf{x}_i, \mathbf{u}_i, \dot{\mathbf{x}}_i)\}_{i=1}^N$, with $\mathbf{x}_i \in E$, $\mathbf{u}_i \in \mathbb{R}^m$, and $\dot{\mathbf{x}}_i \in \mathbb{R}^n$ for $i = 1, \dots, N$. I now show how data allows us to reduce model error by constraining the possible values of the functions $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ directly, without a parametric estimator. Considering the true open-loop system (2.1) evaluated at a state and input pair $(\mathbf{x}_i, \mathbf{u}_i)$ in the dataset yields:

$$\tilde{\mathbf{F}}_i \triangleq \dot{\mathbf{x}}_i - (\hat{\mathbf{f}}(\mathbf{x}_i) + \hat{\mathbf{g}}(\mathbf{x}_i)\mathbf{u}_i) = \tilde{\mathbf{f}}(\mathbf{x}_i) + \tilde{\mathbf{g}}(\mathbf{x}_i)\mathbf{u}_i, \quad (3.161)$$

where $\tilde{\mathbf{F}}_i \in \mathbb{R}^n$ can be interpreted as the error between the true state time derivative and the nominal model (3.1) evaluated at the state and input pair $(\mathbf{x}_i, \mathbf{u}_i)$. Considering a state $\mathbf{x} \in E$ (not necessarily present in the dataset \mathfrak{D}), the second equality in (3.161) implies:

$$\tilde{\mathbf{f}}(\mathbf{x}) + \tilde{\mathbf{g}}(\mathbf{x})\mathbf{u}_i - \tilde{\mathbf{F}}_i = \tilde{\mathbf{f}}(\mathbf{x}) - \tilde{\mathbf{f}}(\mathbf{x}_i) + (\tilde{\mathbf{g}}(\mathbf{x}) - \tilde{\mathbf{g}}(\mathbf{x}_i))\mathbf{u}_i. \quad (3.162)$$

This expression provides a relationship between the possible values of the unmodeled dynamics $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ at the state \mathbf{x} and the values of the unmodeled dynamics at the data point \mathbf{x}_i . Using the fact the functions $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ are Lipschitz continuous on E yields the following bound:

$$\left\| \tilde{\mathbf{f}}(\mathbf{x}) + \tilde{\mathbf{g}}(\mathbf{x})\mathbf{u}_i - \tilde{\mathbf{F}}_i \right\| = \left\| \tilde{\mathbf{f}}(\mathbf{x}) - \tilde{\mathbf{f}}(\mathbf{x}_i) + (\tilde{\mathbf{g}}(\mathbf{x}) - \tilde{\mathbf{g}}(\mathbf{x}_i))\mathbf{u}_i \right\|, \quad (3.163)$$

$$\leq (L_{\tilde{\mathbf{f}}} + L_{\tilde{\mathbf{g}}}\|\mathbf{u}_i\|) \|\mathbf{x} - \mathbf{x}_i\| \triangleq \epsilon_i(\mathbf{x}), \quad (3.164)$$

where $\epsilon_i : E \rightarrow \mathbb{R}_{\geq 0}$ for $i = 1, \dots, N$. The bound grows with the magnitude of the Lipschitz constants $L_{\tilde{\mathbf{f}}}$ and $L_{\tilde{\mathbf{g}}}$ and distance of the state \mathbf{x} from the data point \mathbf{x}_i . The values of $L_{\tilde{\mathbf{f}}}$ and $L_{\tilde{\mathbf{g}}}$ are not explicitly data dependent, and thus the bound can be improved for a given dataset by reducing the possible model error through improved modeling. Given this construction, I define the point-wise error set:

$$\mathcal{U}_i(\mathbf{x}) \triangleq \left\{ (\mathbf{A}, \mathbf{b}) \in \mathbb{R}^{n \times m} \times \mathbb{R}^n \mid \left\| \mathbf{b} + \mathbf{A}\mathbf{u}_i - \tilde{\mathbf{F}}_i \right\| \leq \epsilon_i(\mathbf{x}) \right\} \subset \mathbb{R}^{n \times m} \times \mathbb{R}^n, \quad (3.165)$$

noting that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}_i(\mathbf{x})$ and $\mathcal{U}_i(\mathbf{x})$ is closed and convex. Performing this construction over the entire dataset \mathfrak{D} yields a point-wise error set:

$$\mathcal{U}(\mathbf{x}) \triangleq \bigcap_{i=1}^N \mathcal{U}_i(\mathbf{x}) \subset \mathbb{R}^{n \times m} \times \mathbb{R}^n, \quad (3.166)$$

noting that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}(\mathbf{x})$ and $\mathcal{U}(\mathbf{x})$ is closed and convex. Therefore, $\mathcal{U}(\mathbf{x})$ consists of all possible model errors that are consistent with the observed data. This allows us to pose the following robust control problem:

Definition 24 (*Data Robust CCF Optimization Problem*).

$$\begin{aligned} \mathbf{k}_{\text{ROB}}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & (\text{DR-CCF-OP}) \\ \text{s.t. } & \hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathbf{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) \leq -\alpha(\mathbf{C}(\mathbf{x})) \\ & \text{for all } (\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x}). \end{aligned}$$

By construction we have that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}(\mathbf{x})$, implying that $\mathbf{k}_{\text{ROB}}(\mathbf{x}) \in K_{\text{CCF}}(\mathbf{x})$ when the problem is feasible. I next present one the main results of this section by using robust optimization [220] to yield a convex problem for synthesizing such a robust controller.

Theorem 22. *Let $\mathbf{C} : E \rightarrow \mathbb{R}$ be a CCF with on the set $\mathcal{D} \subset E$ with corresponding function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$. The robust controller (DR-CCF-OP) is equivalently expressed as:*

$$\begin{aligned} \mathbf{k}_{\text{ROB}}(\mathbf{x}) = \operatorname{argmin}_{\substack{\mathbf{u} \in \mathbb{R}^m \\ \boldsymbol{\lambda}_i \in \mathbb{R}^n}} & \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & (\text{DR-CCF-SOCP}) \\ \text{s.t. } & \hat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \tilde{\mathbf{F}}_i - \|\boldsymbol{\lambda}_i\| \epsilon_i(\mathbf{x}) \right) \leq -\alpha(\mathbf{C}(\mathbf{x})), \\ & \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top = -\nabla \mathbf{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathbf{C}(\mathbf{x}). \end{aligned}$$

Observe that this optimization problem is not a quadratic program, but rather, a *second-order cone program* (SOCP) [29]. This is due to the fact that the terms $\|\boldsymbol{\lambda}_i\|$ appear in an inequality constraint. SOCPs are a more complex type of convex optimization problem (every QP can be written as an SOCP), but may still be efficiently solved through various numerical solvers such as ECOS [35] and MOSEK [36]. An important observation that I will also show in other contributions of this thesis synthesizing robust controllers for various modeling challenges through convex optimization is that robustifying an optimization problem of a particular class (a QP) will generally make it more complex, such as an SOCP. While I do not explore the local Lipschitz continuity properties of the controller above, I will explore some local Lipschitz continuity results for

SOCP-based controllers in Section 5.2. A meaningful area of future work in robust convex optimization-based control through CLFs and CBFs will seek to study the regularity properties of controllers specified through more complex optimization problems than QPs (which have been well covered in [51], [221], [222]) using concepts from variational analysis [223], [224]. Another interesting feature of this form of the controller is the dependence of the control input \mathbf{u} on the inputs \mathbf{u}_i seen in the data set. In particular, the last two constraints establish a relationship between \mathbf{u} and \mathbf{u}_i where \mathbf{u} must be in some sense a linear combination of inputs \mathbf{u}_i in the data set. This relationship is best captured by the colloquial description of *do what has worked before, and avoid what you don't know*. Lastly, the size of this optimization problem scales with the numbers of data points N , as each data point adds a decision variable λ_i . Solving this optimization problem in a computationally efficient manner may require data segmentation for higher-dimensional problems [225], but a detailed consideration is outside the scope of this thesis, which is focused on theoretical foundations.

Proof. An input $\mathbf{u} \in \mathbb{R}^m$ is feasible if the optimal value of the optimization problem:

$$\begin{aligned} & \sup_{\substack{\mathbf{A} \in \mathbb{R}^{n \times m} \\ \mathbf{b} \in \mathbb{R}^n}} \widehat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathbf{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) & (3.167) \\ & \text{s.t. } \|\mathbf{b} + \mathbf{A}\mathbf{u}_i - \widetilde{\mathbf{F}}_i\| \leq \epsilon_i(\mathbf{x}) \text{ for all } i \in 1, \dots, N, \end{aligned}$$

is less than or equal to $-\alpha(\mathbf{C}(\mathbf{x}))$. Each inequality constraint can be rewritten as set membership in a second-order cone. Thus we can rewrite this optimization problem as:

$$\begin{aligned} & \sup_{\substack{\mathbf{A} \in \mathbb{R}^{n \times m} \\ \mathbf{b} \in \mathbb{R}^m}} \widehat{\mathbf{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathbf{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) & (3.168) \\ & \text{s.t. } \left(\mathbf{b} + \mathbf{A}\mathbf{u}_i - \widetilde{\mathbf{F}}_i, \epsilon_i(\mathbf{x}) \right) \in \mathcal{Q}_n \text{ for all } i \in 1, \dots, N, \end{aligned}$$

where $\mathcal{Q}_n \subset \mathbb{R}^{n+1}$ denotes the Lorentz (second-order) cone:

$$\mathcal{Q}_n = \{(\mathbf{y}, t) \in \mathbb{R}^n \times \mathbb{R} : \|\mathbf{y}\| \leq t\}. \quad (3.169)$$

The associated Lagrangian for this optimization problem is given by:

$$\begin{aligned} & \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) + \sum_{i=1}^N \begin{bmatrix} \boldsymbol{\lambda}_i \\ \nu_i \end{bmatrix}^\top \begin{bmatrix} \mathbf{b} + \mathbf{A}\mathbf{u}_i - \widetilde{\mathbf{F}}_i \\ \epsilon_i(\mathbf{x}) \end{bmatrix} \\ &= \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathcal{C}(\mathbf{x})^\top \mathbf{b} + \langle \nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \mathbf{A} \rangle \\ & \quad + \sum_{i=1}^N \left(\langle \boldsymbol{\lambda}_i \mathbf{u}_i^\top, \mathbf{A} \rangle + \boldsymbol{\lambda}_i^\top \mathbf{b} - \boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i + \nu_i \epsilon_i(\mathbf{x}) \right), \end{aligned} \quad (3.170)$$

$$\begin{aligned} &= \left\langle \nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top + \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top, \mathbf{A} \right\rangle + \left(\nabla \mathcal{C}(\mathbf{x}) + \sum_{i=1}^N \boldsymbol{\lambda}_i \right)^\top \mathbf{b} \\ & \quad + \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \nu_i \epsilon_i(\mathbf{x}) \right), \end{aligned} \quad (3.171)$$

where $(\boldsymbol{\lambda}_i, \nu_i) \in \mathcal{Q}_n$ are Lagrange multipliers corresponding to the constraints imposed by the data point $(\mathbf{x}_i, \mathbf{u}_i)$ and $\langle \cdot, \cdot \rangle$ denotes the trace inner product. The dual optimization problem is therefore:

$$\begin{aligned} & \inf_{\substack{\boldsymbol{\lambda}_i \in \mathbb{R}^n \\ \nu_i \in \mathbb{R}}} \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \nu_i \epsilon_i(\mathbf{x}) \right) \\ & \text{s.t.} \quad \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top = -\nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathcal{C}(\mathbf{x}), \\ & \quad \|\boldsymbol{\lambda}_i\| \leq \nu_i \text{ for all } i \in 1, \dots, N. \end{aligned} \quad (3.172)$$

For any $(\mathbf{x}_i, \mathbf{u}_i)$, choosing Lagrange multipliers such that $\nu_i > \|\boldsymbol{\lambda}_i\|$ increases the corresponding value of the problem compared to choosing $\nu_i = \|\boldsymbol{\lambda}_i\|$. I therefore simplify the problem as:

$$\begin{aligned} & \inf_{\boldsymbol{\lambda}_i \in \mathbb{R}^n} \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \|\boldsymbol{\lambda}_i\| \epsilon_i(\mathbf{x}) \right) \\ & \text{s.t.} \quad \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top = -\nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathcal{C}(\mathbf{x}). \end{aligned} \quad (3.173)$$

This optimization problem can be substituted in for the robust constraint in (DR-CCF-OP) to yield the optimization problem (DR-CCF-SOCP). \square

Controller Feasibility

I now provide an analysis of the feasibility of the preceding controller. The feasibility of the controller (DR-CCF-SOCP) at a given state \mathbf{x} is determined by the structure of the model error set $\mathcal{U}(\mathbf{x})$ defined in (3.166). The following

lemma provides a condition on the inputs in the dataset that implies $\mathcal{U}(\mathbf{x})$ is bounded for all $\mathbf{x} \in E$:

Lemma 1. *Consider a dataset \mathcal{D} with N data points satisfying $N \geq m + 1$. If there exists a set of data points $\{(\mathbf{x}_i, \mathbf{u}_i, \tilde{\mathbf{x}}_i)\}_{i=1}^{m+1} \subseteq \mathcal{D}$ such that the set of vectors:*

$$\mathcal{M} \triangleq \left\{ \begin{bmatrix} \mathbf{u}_i \\ 1 \end{bmatrix} \right\}_{i=1}^{m+1}, \quad (3.174)$$

are linearly independent, then the model error set $\mathcal{U}(\mathbf{x})$ is bounded (and thus compact) for any $\mathbf{x} \in E$.

Proof. Consider the set \mathcal{M} , and define:

$$\mathcal{M}_{\mathbf{u}} \triangleq \left\{ \mathbf{u}_i \in \mathbb{R}^m \mid \begin{bmatrix} \mathbf{u}_i \\ 1 \end{bmatrix} \in \mathcal{M} \right\}. \quad (3.175)$$

For arbitrary $\mathbf{x} \in E$, let $(\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x})$. By definition of the model error sets $\mathcal{U}_i(\mathbf{x})$ and $\mathcal{U}(\mathbf{x})$ given in (3.165) and (3.166), respectively, we have that:

$$\left\| \mathbf{A}\mathbf{u}_i + \mathbf{b} - \tilde{\mathbf{F}}_i \right\| \leq \epsilon_i(\mathbf{x}), \quad (3.176)$$

for $\mathbf{u}_i \in \mathcal{M}_{\mathbf{u}}$. Defining $\mathbf{v}_i \triangleq \mathbf{A}\mathbf{u}_i + \mathbf{b} - \tilde{\mathbf{F}}_i$, we have that:

$$\sum_{i=1}^{m+1} \left\| \mathbf{A}\mathbf{u}_i + \mathbf{b} - \tilde{\mathbf{F}}_i \right\|^2 = \sum_{i=1}^{m+1} \sum_{j=1}^n |V_{ji}|^2 = \|\mathbf{V}\|_F^2, \quad (3.177)$$

where $\|\cdot\|_F$ is the Frobenius norm and $\mathbf{V} = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_{m+1}]$. Noting that $\|\mathbf{V}\|_F = \|\mathbf{V}^\top\|_F$, factoring and using the fact $\|\mathbf{P}\| \leq \|\mathbf{P}\|_F$ for any $\mathbf{P} \in \mathbb{R}^{m+1 \times n}$ in conjunction with (3.176) yields:

$$\left\| \begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} - \begin{bmatrix} \tilde{\mathbf{F}}_1^\top \\ \vdots \\ \tilde{\mathbf{F}}_{m+1}^\top \end{bmatrix} \right\|^2 \leq \sum_{i=1}^{m+1} \epsilon_i(\mathbf{x})^2. \quad (3.178)$$

Taking the square root of both sides and employing the reverse triangle inequality we arrive at:

$$\left\| \underbrace{\begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix}}_{\mathbf{U}} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\| \leq \sqrt{\sum_{i=1}^{m+1} \epsilon_i(\mathbf{x})^2} + \left\| \begin{bmatrix} \tilde{\mathbf{F}}_1^\top \\ \vdots \\ \tilde{\mathbf{F}}_{m+1}^\top \end{bmatrix} \right\|. \quad (3.179)$$

Noting that \mathbf{U} has full rank by the linear independence of the vectors in the set \mathcal{M} , we have that:

$$\sigma_{m+1}(\mathbf{U}) \left\| \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\| \leq \left\| \begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\|, \quad (3.180)$$

where $\sigma_{m+1}(\mathbf{U}) > 0$ is the smallest singular value of \mathbf{U} . Combining the two previous inequalities shows that the model error set $\mathcal{U}(\mathbf{x})$ is bounded. \square

This result shows that variety in input directions is sufficient to assert boundedness of the model error set as a uniform property over the entire set E . As seen in the proof, the bound on this set may be very large if the values of $\epsilon_i(\mathbf{x})$ are large (as is the case when \mathbf{x} is far away from the data points \mathbf{x}_i associated with the inputs used to construct the bound). Alternatively, the model error set will be small for a point \mathbf{x} if there is local variety in input directions within the training dataset. To understand how the size of $\mathcal{U}(\mathbf{x})$ impacts feasibility of the optimization problem, I define the following set:

$$\begin{aligned} \tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}) \triangleq \left\{ (\mathbf{a}, b) \in \mathbb{R}^m \times \mathbb{R} \mid \exists (\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x}) \text{ s.t.} \right. \\ \left. \mathbf{a} = (\nabla \mathcal{C}(\mathbf{x})^\top \mathbf{A})^\top, b = \nabla \mathcal{C}(\mathbf{x})^\top \mathbf{b} \right\}. \end{aligned} \quad (3.181)$$

The set $\tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ can be interpreted as the projection of the dynamics model error set $\mathcal{U}(\mathbf{x})$ along the gradient of the CCF \mathcal{C} , creating an $m + 1$ dimensional set. Additionally define the set:

$$\mathcal{U}_{\mathcal{C}}(\mathbf{x}) \triangleq \left\{ (L_{\hat{\mathbf{g}}}\mathcal{C}(\mathbf{x})^\top, L_{\hat{\mathbf{f}}}\mathcal{C}(\mathbf{x})) \right\} \oplus \tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}), \quad (3.182)$$

where \oplus denotes a Minkowski sum. The set $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ is the recentering of $\tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ around the nominal model terms $(L_{\hat{\mathbf{g}}}\mathcal{C}(\mathbf{x})^\top, L_{\hat{\mathbf{f}}}\mathcal{C}(\mathbf{x})) \in \mathbb{R}^{m+1}$, such that it captures the possible values of $(L_{\mathbf{g}}\mathcal{C}(\mathbf{x})^\top, L_{\mathbf{f}}\mathcal{C}(\mathbf{x}))$. As multiplication by $\nabla \mathcal{C}(\mathbf{x})^\top$ is a linear transformation, $\tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ and $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ are convex, and if $\mathcal{U}(\mathbf{x})$ is bounded (and therefore compact), then $\tilde{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ and $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ are compact. I now present the second main result of this section in the form of a necessary and sufficient condition for feasibility of (DR-CCF-OP):

Theorem 23. *For a state $\mathbf{x} \in E$, let the sets $\mathcal{U}(\mathbf{x})$ and $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ be defined as in (3.166) and (3.182), respectively. Define the ray $\mathcal{R} \subset \mathbb{R}^{m+1}$ as $\mathcal{R} =$*

$\{\mathbf{0}_m\} \times (-\alpha(\mathbf{C}(\mathbf{x})), \infty)$. Assuming that $\mathcal{U}(\mathbf{x})$ is bounded, the data-driven robust controller (DR-CCF-SOCP) is feasible if and only if:

$$\mathcal{U}_c(\mathbf{x}) \cap \mathcal{R} = \emptyset. \quad (3.183)$$

Intuitively, the ray \mathcal{R} represents scenarios in which the CCF time derivative can not be modified by actuation, and the CCF condition is not satisfied without actuation. By Assumption 4, the tuple $(L_{\mathbf{g}}\mathbf{C}(\mathbf{x})^\top, L_{\mathbf{f}}\mathbf{C}(\mathbf{x}))$ is not contained in \mathcal{R} , but the possible model errors permitted by data need not necessarily reflect this. If one possible model error in the set $\mathcal{U}_c(\mathbf{x})$ is contained in \mathcal{R} , it is impossible to meet the CCF condition for that model error. Observe that the error set not intersecting this ray is a direct analog of the alternative CLF and CBF condition in (2.18) and (2.61), respectively. As those conditions enable proving feasibility of the resulting optimization-based controller, the analog in this setting will also permit proving feasibility.

Proof. The proof proceeds from the original structure of the robust control problem (DR-CCF-OP). The constraint on the input \mathbf{u}^* specified by this controller is given by:

$$L_{\mathbf{f}}\mathbf{C}(\mathbf{x}) + \nabla\mathbf{C}(\mathbf{x})^\top \mathbf{b} + (L_{\mathbf{g}}\mathbf{C}(\mathbf{x}) + \nabla\mathbf{C}(\mathbf{x})^\top \mathbf{A})\mathbf{u}^* \leq -\alpha(\mathbf{C}(\mathbf{x})), \quad (3.184)$$

for all $(\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x})$. Given the definitions of $\tilde{\mathcal{U}}_c(\mathbf{x})$ and $\mathcal{U}_c(\mathbf{x})$ in (3.181) and (3.182), this constraint can be expressed as:

$$q + \mathbf{p}^\top \mathbf{u}^* \leq -\alpha(\mathbf{C}(\mathbf{x})), \quad (3.185)$$

for all $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$.

Necessity

Assume that $\mathcal{U}_c(\mathbf{x}) \cap \mathcal{R} \neq \emptyset$. This implies that there exists a $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$ such that $q > -\alpha(\mathbf{C}(\mathbf{x}))$ and $\mathbf{p} = \mathbf{0}_m$. Thus for any input $\mathbf{u} \in \mathbb{R}^m$, we have that:

$$q + \mathbf{p}^\top \mathbf{u} = q > -\alpha(\mathbf{C}(\mathbf{x})), \quad (3.186)$$

violating the constraint in (3.185). Thus the optimization problem is infeasible.

Sufficiency

Begin by defining the hyperplane $\mathcal{H} \subset \mathbb{R}^{m+1}$ with unit normal $\mathbf{n} = \begin{bmatrix} \mathbf{0}_m^\top & 1 \end{bmatrix}^\top$ and offset $-\alpha(\mathcal{C}(\mathbf{x}))$ and define the set:

$$\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}) = \left\{ (\mathbf{p}, q) \in \mathcal{U}_{\mathcal{C}}(\mathbf{x}) \mid \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle \geq -\alpha(\mathcal{C}(\mathbf{x})) \right\}, \quad (3.187)$$

which corresponds to the tuples in the set $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ that do not meet the CCF condition (3.185) strictly under no input ($\mathbf{u}^* = \mathbf{0}_m$). Note that if $\mathbf{u}^* = \mathbf{0}_m$ satisfies (3.185) for all $(\mathbf{p}, q) \in \bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$, then $\mathbf{u}^* = \mathbf{0}_m$ satisfies (3.185) for all $(\mathbf{p}, q) \in \mathcal{U}_{\mathcal{C}}(\mathbf{x})$. Also note that the set $\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ is the closed subset of the compact and convex set $\mathcal{U}_{\mathcal{C}}(\mathbf{x})$ that is also contained in a (convex) half-space defined by the hyperplane \mathcal{H} , meaning $\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ is compact and convex (as it is the intersection of convex sets). Therefore, we can define:

$$q^{**} = \max_{(\mathbf{p}, q) \in \bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})} q = \max_{(\mathbf{p}, q) \in \bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})} \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle, \quad (3.188)$$

which exists as the function $\left\langle \begin{bmatrix} \mathbf{0}_m^\top & 1 \end{bmatrix}^\top, \cdot \right\rangle : \bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}) \rightarrow \mathbb{R}$ is continuous on a compact domain. We consider the case that $q^{**} > -\alpha(\mathcal{C}(\mathbf{x}))$, as otherwise $\mathbf{u}^* = \mathbf{0}_m$ satisfies (3.185).

Define the projection function $\mathbf{\Pi} : \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}^m \times \mathbb{R}$ such that for $(\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R}$ we have:

$$\mathbf{\Pi}((\mathbf{p}, q)) = \mathbf{p}. \quad (3.189)$$

As $\mathbf{\Pi}$ is a linear transform, the image of the compact and convex set $\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})$ under $\mathbf{\Pi}$, denoted as $\mathbf{\Pi}(\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}))$, is compact and convex and by assumption satisfies:

$$\mathbf{\Pi}(\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x})) \cap \{\mathbf{0}_m\} = \emptyset. \quad (3.190)$$

We may use the strict separating hyperplane theorem [29] to separate the set $\mathbf{\Pi}(\bar{\mathcal{U}}_{\mathcal{C}}(\mathbf{x}))$ from $\{\mathbf{0}_m\}$ with the hyperplane \mathcal{H}_β with unit normal $\mathbf{s} \in S^{m-1}$ and offset $\beta \in \mathbb{R}_{>0}$. This hyperplane can also be shifted to pass through the origin, given by \mathcal{H}_0 (with unit normal $\mathbf{s} \in S^{m-1}$ and offset 0). This results in the configuration seen in Figure 3.15. These hyperplanes can be extended back into the ambient space $\mathbb{R}^m \times \mathbb{R}$ by defining hyperplanes $\mathcal{H}'_\beta \subset \mathbb{R}^m \times \mathbb{R}$ and $\mathcal{H}'_0 \subset \mathbb{R}^m \times \mathbb{R}$ with normal vector $\mathbf{s}' = \begin{bmatrix} \mathbf{s}^\top & 0 \end{bmatrix}^\top$ and respective offsets β and 0.

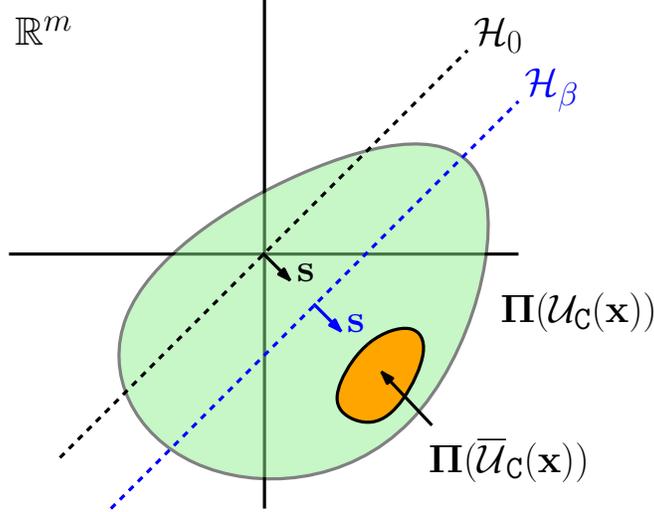


Figure 3.15. Projected view of model error set used in proving feasibility of the (DR-CCF-OP). The green region corresponds to $\Pi(\mathcal{U}_c(\mathbf{x}))$ while the orange region highlights the set $\bar{\mathcal{U}}_c(\mathbf{x})$ lying above the hyperplane \mathcal{H} . The line \mathcal{H}_β represents the strictly separating hyperplane with normal \mathbf{s} and offset β , while \mathcal{H}_0 is the hyperplane shifted to pass through the origin.

These hyperplanes serve to separate the vertical axis $\mathbf{0}_m \times \mathbb{R}$ from the cylinder defined by $\Pi(\bar{\mathcal{U}}_c(\mathbf{x})) \times \mathbb{R}$. Define the two following open halfspaces:

$$\mathcal{H}_0^+ \triangleq \left\{ (\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R} \mid \left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{p}^\top \mathbf{s} > 0 \right\}, \quad (3.191)$$

$$\mathcal{H}_0^- \triangleq \left\{ (\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R} \mid \left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{p}^\top \mathbf{s} < 0 \right\}. \quad (3.192)$$

We will now find a feasible input that lies anti-parallel to \mathbf{s} , i.e., we will consider inputs of the form $\mathbf{u}^* = -\gamma \mathbf{s}$ for $\gamma \in \mathbb{R}_{\geq 0}$ throughout the rest of this proof. Finding a feasible input satisfying constraint (3.185) amounts to finding a γ^* such that $q - \gamma^* \mathbf{p}^\top \mathbf{s} \leq -\alpha(\mathbf{C}(\mathbf{x}))$ for all $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$. By definition we have that for any $(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})$:

$$\left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{s}^\top \mathbf{p} \geq \beta. \quad (3.193)$$

Let $\gamma_0 \in \mathbb{R}_{>0}$ be defined as:

$$\gamma_0 = \frac{q^{**} + \alpha(\mathbf{C}(\mathbf{x}))}{\beta} > 0. \quad (3.194)$$

This value of γ_0 implies that for any $(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})$ we have:

$$q - \gamma_0 \mathbf{p}^\top \mathbf{s} \leq q^{**} - \gamma_0 \beta \leq -\alpha(\mathbf{C}(\mathbf{x})). \quad (3.195)$$

Define the set:

$$\mathcal{V}_c(\mathbf{x}, \gamma) = \{(\mathbf{p}, q - \gamma \mathbf{p}^\top \mathbf{s}) \in \mathbb{R}^m \times \mathbb{R} \mid (\mathbf{p}, q) \in \overline{\mathcal{U}}_c(\mathbf{x})\}, \quad (3.196)$$

and the function:

$$\psi(\gamma) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} q, \quad (3.197)$$

that satisfies $\psi(0) = q^{**}$ and $\psi(\gamma_0) \leq -\alpha(\mathcal{C}(\mathbf{x}))$. I later show that ψ is continuous, but note the intermediate value theorem implies the existence of a $\gamma^* \in (0, \gamma_0]$ with:

$$\psi(\gamma^*) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma^*)} q = -\alpha(\mathcal{C}(\mathbf{x})). \quad (3.198)$$

Fixing $\mathbf{u}^* = -\gamma^* \mathbf{s}$, we have that condition (3.185) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap \overline{\mathcal{U}}_c(\mathbf{x})$. Likewise, consider a point $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap (\mathcal{U}_c(\mathbf{x}) \setminus \overline{\mathcal{U}}_c(\mathbf{x}))$, for which the condition (3.185) is satisfied with no input as $q < -\alpha(\mathcal{C}(\mathbf{x}))$, and note that $\mathbf{p}^\top \mathbf{s} > 0$. We have that:

$$q + \mathbf{p}^\top \mathbf{u}^* = q - \gamma^* \mathbf{p}^\top \mathbf{s} < -\alpha(\mathcal{C}(\mathbf{x})). \quad (3.199)$$

Combining these results, we have that condition (3.185) is satisfied for all $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap \mathcal{U}_c(\mathbf{x})$. Additionally, for any $(\mathbf{p}, q) \in \mathcal{H}'_0 \cap \mathcal{U}_c(\mathbf{x})$, $\mathbf{p}^\top \mathbf{s} = 0$, as \mathbf{s}' is normal to \mathcal{H}'_0 . As strict separation implies that $q < -\alpha(\mathcal{C}(\mathbf{x}))$ for any $(\mathbf{p}, q) \in \mathcal{H}'_0 \cap \mathcal{U}_c(\mathbf{x})$, we have $q - \gamma^* \mathbf{p}^\top \mathbf{s} < -\alpha(\mathcal{C}(\mathbf{x}))$ and condition (3.185) is satisfied.

Lastly, we must consider points $(\mathbf{p}, q) \in \mathcal{H}_0^- \cap \mathcal{U}_c(\mathbf{x})$. To this end, define the set:

$$\mathcal{W}_c(\mathbf{x}, \gamma^*) = \{(\mathbf{p}, q - \gamma^* \mathbf{p}^\top \mathbf{s}) \in \mathbb{R}^m \times \mathbb{R} \mid (\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})\}. \quad (3.200)$$

This set can be interpreted as an invertible linear transformation of the set $\mathcal{U}_c(\mathbf{x})$ as we have:

$$\begin{bmatrix} \mathbf{p}' \\ q' \end{bmatrix} = \begin{bmatrix} \mathbf{p} \\ q - \gamma^* \mathbf{p}^\top \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{m \times m} & \mathbf{0}_m \\ -\gamma^* \mathbf{s}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix}, \quad (3.201)$$

implying that $\mathcal{W}_c(\mathbf{x}, \gamma^*)$ is compact and convex. Similarly, define the set:

$$\overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \left\{ (\mathbf{p}, q) \in \mathcal{W}_c(\mathbf{x}, \gamma^*) \mid \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle \geq -\alpha(\mathcal{C}(\mathbf{x})) \right\}. \quad (3.202)$$

The set $\overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$ is convex and contains points in $\mathcal{W}_c(\mathbf{x}, \gamma^*)$ that are in or above the hyperplane \mathcal{H} , or points that meet with equality or violate (3.185),

respectively. By the specific choice of γ^* , there exists at least one point $\mathbf{v}^+ \in \mathcal{H}_0^+ \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. Furthermore, by the preceding strict separation, we have that $\mathcal{H}_0^+ \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$. Now assume for contradiction that $\mathcal{H}_0^- \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*) \neq \emptyset$, or that there exists $\mathbf{v}^- \in \mathcal{H}_0^- \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. As the set $\overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$ is convex and $\langle \mathbf{s}', \mathbf{v}^+ \rangle > 0$ and $\langle \mathbf{s}', \mathbf{v}^- \rangle < 0$, there is a $\lambda^* \in (0, 1)$ satisfying:

$$\langle \mathbf{s}', (1 - \lambda^*)\mathbf{v}^+ + \lambda^*\mathbf{v}^- \rangle = 0, \quad (3.203)$$

implying $(1 - \lambda^*)\mathbf{v}^+ + \lambda^*\mathbf{v}^- \in \mathcal{H}_0' \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. This contradicts the fact $\mathcal{H}_0' \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$, implying that $\mathcal{H}_0^- \cap \overline{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$, or that condition (3.185) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{H}_0^- \cap \mathcal{U}_c(\mathbf{x})$. Combining all previous results, condition (3.185) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$ for the input $\mathbf{u}^* = -\gamma^*\mathbf{s}$, ensuring feasibility.

I now show that ψ is continuous. For $\gamma, \gamma' > 0$ and $(\mathbf{p}, q) \in \overline{\mathcal{U}}_c(\mathbf{x})$, we have:

$$\min_{(\mathbf{p}', q') \in \mathcal{V}_c(\mathbf{x}, \gamma')} \left\| \begin{bmatrix} \mathbf{p} \\ q - \gamma \mathbf{p}^\top \mathbf{s} \end{bmatrix} - \begin{bmatrix} \mathbf{p}' \\ q' \end{bmatrix} \right\| \quad (3.204)$$

$$\leq \left\| \begin{bmatrix} \mathbf{p} \\ q - \gamma \mathbf{p}^\top \mathbf{s} \end{bmatrix} - \begin{bmatrix} \mathbf{p} \\ q - \gamma' \mathbf{p}^\top \mathbf{s} \end{bmatrix} \right\| \leq \mathbf{p}^\top \mathbf{s} |\gamma - \gamma'|, \quad (3.205)$$

and similarly, for $(\mathbf{p}', q') \in \overline{\mathcal{U}}_c(\mathbf{x})$, we have:

$$\min_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} \left\| \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} - \begin{bmatrix} \mathbf{p}' \\ q' - \gamma' (\mathbf{p}')^\top \mathbf{s} \end{bmatrix} \right\| \leq (\mathbf{p}')^\top \mathbf{s} |\gamma - \gamma'|. \quad (3.206)$$

Therefore, the Hausdorff distance between $\mathcal{V}_c(\mathbf{x}, \gamma)$ and $\mathcal{V}_c(\mathbf{x}, \gamma')$ is bounded:

$$d_H(\mathcal{V}_c(\mathbf{x}, \gamma), \mathcal{V}_c(\mathbf{x}, \gamma')) \leq \left(\max_{(\mathbf{p}, q) \in \overline{\mathcal{U}}_c(\mathbf{x})} \mathbf{p}^\top \mathbf{s} \right) |\gamma - \gamma'|,$$

implying $\mathcal{V}_c(\mathbf{x}, \gamma)$ is Lipschitz continuous (with respect to the Hausdorff metric) as a function of γ . The support function of a nonempty, compact, and convex set $\mathcal{A} \subset \mathbb{R}^m \times \mathbb{R}$ given by $h_{\mathcal{A}} : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ is defined as:

$$h_{\mathcal{A}}(\mathbf{v}) = \max_{(\mathbf{p}, q) \in \mathcal{A}} \left\langle \mathbf{v}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle. \quad (3.207)$$

Recalling $\mathbf{n} = \begin{bmatrix} \mathbf{0}_m^\top & 1 \end{bmatrix}^\top$, the Hausdorff distance between two nonempty, compact, and convex sets $\mathcal{A}, \mathcal{B} \subset \mathbb{R}^m \times \mathbb{R}$ satisfies:

$$d_H(\mathcal{A}, \mathcal{B}) = \max_{\mathbf{v} \in S^m} |h_{\mathcal{A}}(\mathbf{v}) - h_{\mathcal{B}}(\mathbf{v})| \geq |h_{\mathcal{A}}(\mathbf{n}) - h_{\mathcal{B}}(\mathbf{n})|.$$

Noting that ψ can be expressed in terms of the support function as a composition, with:

$$\psi(\gamma) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} q = h_{\mathcal{V}_c(\mathbf{x}, \gamma)}(\mathbf{n}),$$

implies ψ is a continuous function. \square

Simulation Results

To demonstrate the proposed controller, I simulate the inverted pendulum with input attenuation in (3.3) with gravitational acceleration $g = 10$, length $\ell = 0.7$, and mass $m = 0.7$. The nominal model for the system is given by (3.4) with gravitational acceleration estimate $\hat{g} = 10$, length estimate $\hat{\ell} = 0.63$, and mass estimate $\hat{m} = 0.63$. Consider functions $V : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$ and $h : \mathbb{R}^2 \rightarrow \mathbb{R}$, given by $V(\mathbf{x}) = \mathbf{x}^\top \mathbf{P} \mathbf{x}$ and $h(\mathbf{x}) = \mathbf{x}^\top \mathbf{P} \mathbf{x} - c$ with:

$$\mathbf{P} = \begin{bmatrix} \sqrt{3} & 1 \\ 1 & \sqrt{3} \end{bmatrix}, \quad (3.208)$$

and a constant $c = 0.2$. Noting that both the nominal model open-loop system (3.4) and the true open-loop system (3.3) are feedback linearizable, V and h satisfy the CCF condition (3.157) for the comparison function $\alpha(r) = \sigma r / \lambda_{\max}(\mathbf{P})$ with $\sigma \in (0, 1)$, for both the nominal model open-loop system (3.4) and the true open-loop system (3.3) (implying Assumption 4 is met). In particular, V may serve as a CLF, and h as a CBF.

I explore both data-driven stability and safety with this system. I compare the robust data-driven controller with an oracle controller, which is a true model-based (CCF-QP) controller, and a baseline controller, which is a nominal model-based (CCF-QP) controller. In each setting, the system model underestimates the pendulum mass and length and assumes that the input gain is independent of the state. As a result, the Lipschitz constants of the errors can be bounded by $L_{\tilde{\mathbf{f}}} = g|\ell - \hat{\ell}|/(\ell\hat{\ell})$ and $L_{\tilde{\mathbf{g}}} = 0.75\sqrt{2} \exp(-\frac{1}{2})/(m\ell^2)$.

For the stability experiment, I generate data sets by gridding the state and input spaces. I consider x_1 in the interval $[0, 1]$ and x_2 in the interval $[-0.25, 0.25]$, with grid sizes $\epsilon_\theta = \epsilon_{\dot{\theta}} = \frac{1}{40}$. I generate a *sparse* data set by considering $u \in \{-5, -1\}$ and a *dense* data set by considering $u \in \{-5, -3, -1, 1, 3, 5\}$. I set the controller \mathbf{k}_{nom} to be a feedback linearizing controller designed using the nominal model and linear gains $k_p = 1/2$ and $k_d = \sqrt{3}/2$. The system is simulated from the initial condition $\mathbf{x}_0 = [0.8, 0.1]^\top$ for 10 [s], with control

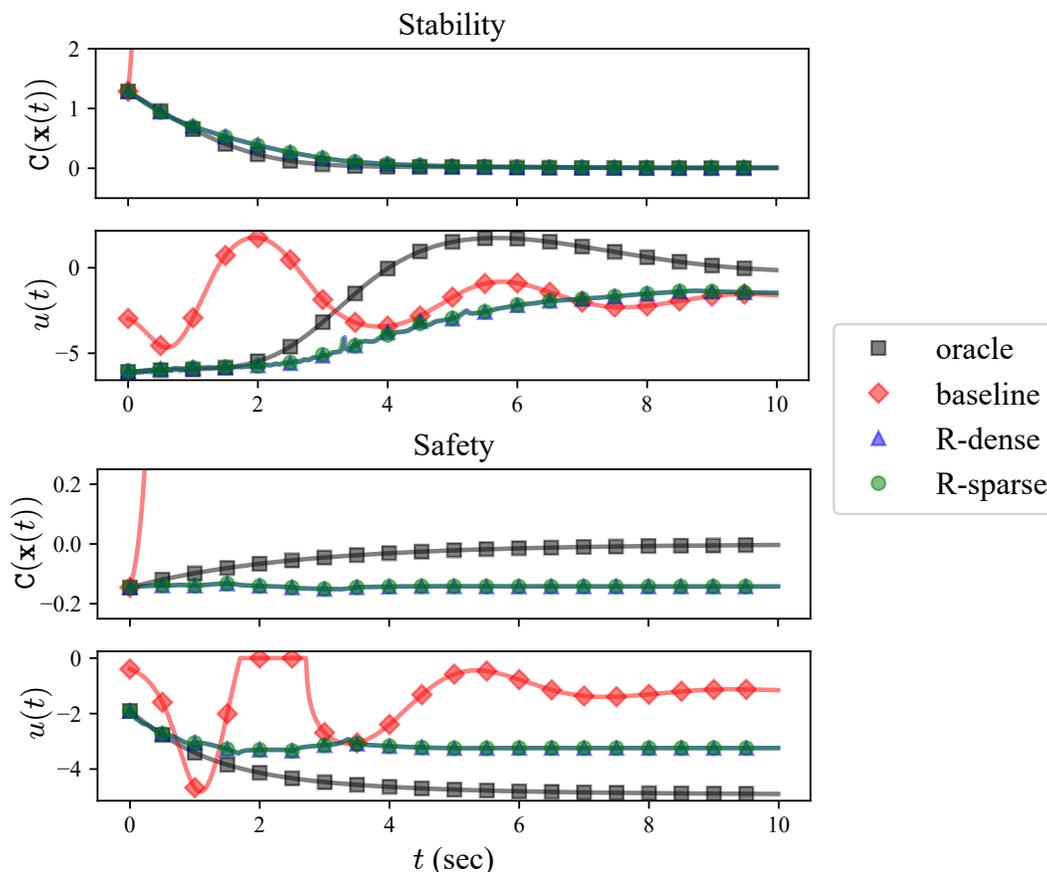


Figure 3.16. The value of the certificate function $C(\mathbf{x})$ (Top) and input u (Bottom) over time for the stability (Top Pair) and safety (Bottom Pair) experiments. I compare the behavior of a true model-based (CCF-QP) controller (oracle, black square), a nominal-model based (CCF-QP) controller (baseline, red diamond), and the proposed (DR-CCF-SOCP) controller using dense input data (R-dense, blue triangle) and sparse input data (R-Sparse, green circle). For stability, the desired behavior for $C(\mathbf{x})$ is to converge towards 0, while for safety it is to remain non-positive.

inputs specified at 100 [Hz]. For the safety experiment, I consider a similar pair of sparse and dense data sets with x_1 in the interval $[0, 0.25]$ and x_2 in the same interval as the stability experiment, the same grid sizes, and the same sets of control inputs. The system is simulated with no nominal controller ($\mathbf{k}_{\text{nom}} = 0$ for all $\mathbf{x} \in E$) from the initial condition $\mathbf{x}_0 = [0.1, 0.1]^\top$ for the same amount of time and the same control input frequency.

The results of the simulations may be seen in Figure 3.16. In both experiments, we see that the baseline controller fails to achieve the specified objective, while the oracle controller succeeds. Furthermore, we see that the data-driven controllers perform nearly identically for both the sparse and dense input data sets. This similarity indicates that greater variety in input directions and cov-

erage of the input space by the data set are not needed to achieve satisfactory closed-loop behavior.

Conclusion

In this section I have presented an approach for data-driven robust control through convex optimization. I first provide a geometric understanding of model error as it impacts the time derivative of a CLF, and indicate how data can be used to bound potential model errors in a way that permits robust control. I then consider Control Certificate Functions as a generalization that captures both CLFs and CBFs. This leads to a robust optimization problem with an infinite number of constraints which can be shown to be equivalent to a finite dimensional second-order cone program through a duality argument. I then provide a statement about the boundedness of model error sets given sufficiently rich input data, and use this to provide a proof on the feasibility of the robust controller. I demonstrate the proposed control approach in simulation on the inverted pendulum with input attenuation in Section 3.2.

3.9 Adaptive Safety via Control Barrier Functions

In this section I will present work on the problem of adaptive safety-critical control through CBFs. I will first review the work in [47] which describes the adaptive control setting considered in this work, and defines adaptive CLFs as a tool for stabilization in the presence of parametric model error. Building off of the ideas in this work, I will propose a novel notion of an *adaptive* barrier function (aBF) and Control Barrier Function (aCBF) that must meet an inequality condition over both states and possible parameter values. I will then show how an aBF yields safety of a closed-loop system and how an aCBF can be used to synthesize a controller that yields an aBF for a closed-loop system. I will then analyze some of the conservative conditions that define aBFs and aCBFs through a rigorous counterexample, demonstrating how a naïve approach to defining aCBFs can lead to a failure of rigorous safety guarantees. Lastly, I demonstrate the tool of aCBFs in simulation on an adaptive cruise control (ACC) example.

The contributions of this section are as follows:

- The definition of adaptive Control Barrier Functions for ensuring the safety of a nonlinear control system with parametric model error.

- A detailed analysis of the strong requirements needed for forward invariance of a set in the presence of parametric model error through a theoretically rigorous counterexample.

The text for this section is adapted from:

A. J. Taylor and A. D. Ames, “Adaptive safety with control barrier functions,” in *Proc. IEEE American Control Conf. (ACC)*, Denver, CO, USA, 2020, pp.1399-1405.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Adaptive Systems

Recall the formulation of parametric model error described in Section 3.2 yielding the open-loop system (3.6) which is restated here:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (3.209)$$

where $\boldsymbol{\theta}^* \in \mathbb{R}^p$ is a vector of unknown parameters and $\mathbf{F} : E \rightarrow \mathbb{R}^{n \times p}$ is a collection of known basis functions and is assumed to be locally Lipschitz continuous on E . The challenge of designing explicit controllers that are robust to a wide range of unknown parameters suggests considering a larger class of controllers to ensure the safety of (3.209). In particular, controllers that *update* an estimate of the unknown parameters. These are called adaptive controllers, and take the form:

$$\mathbf{u} = \mathbf{k}(\mathbf{x}, \widehat{\boldsymbol{\theta}}), \quad (3.210)$$

$$\dot{\widehat{\boldsymbol{\theta}}} = \mathbf{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \widehat{\boldsymbol{\theta}}), \quad (3.211)$$

where $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$ represents an *estimate* of the unknown parameters $\boldsymbol{\theta}^*$ maintained by the controller, $\mathbf{k} : E \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ is locally Lipschitz continuous on $E \times \mathbb{R}^p$, $\mathbf{\Gamma} \in \mathbb{S}_{>0}^p$ is a matrix adaptive gain, and $\boldsymbol{\tau} : E \times \mathbb{R}^p \rightarrow \mathbb{R}^p$ is a adaptation law that is locally Lipschitz continuous on $E \times \mathbb{R}^p$. Introducing this parameter update results in a composite closed-loop system:

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\widehat{\boldsymbol{\theta}}} \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \\ \mathbf{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \end{bmatrix}. \quad (3.212)$$

I refer to the state portion of this system as the state closed-loop system, and the parameter portion as the parameter closed-loop system. The composite system evolves in $E \times \mathbb{R}^p$, such that for any initial condition $(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) \in E \times \mathbb{R}^p$, there exists a maximal interval $I(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) = [0, t_{\max}(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0))$ and a unique continuously differentiable solution $(\boldsymbol{\varphi}_x, \boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}) : I(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) \rightarrow E \times \mathbb{R}^p$ satisfying:

$$\begin{bmatrix} \dot{\boldsymbol{\varphi}}_x(t) \\ \dot{\boldsymbol{\varphi}}_{\widehat{\boldsymbol{\theta}}}(t) \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{f}}(\boldsymbol{\varphi}_x(t)) + \mathbf{F}(\boldsymbol{\varphi}_x(t))\boldsymbol{\theta}^* + \widehat{\mathbf{g}}(\boldsymbol{\varphi}_x(t))\mathbf{k}(\boldsymbol{\varphi}_x(t), \boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}(t)) \\ \boldsymbol{\Gamma}\boldsymbol{\tau}(\boldsymbol{\varphi}_x(t), \boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}(t)) \end{bmatrix}, \quad (3.213)$$

for all $t \in I(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0)$. For simplicity, in this work I will assume that the closed-loop system (3.212) is *forward complete*, such that for any initial condition $(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) \in E \times \mathbb{R}^p$, the interval $I(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) = \mathbb{R}_{\geq 0}$.

Adaptive Control Lyapunov Functions

Stabilization of a system with parametric model errors has long been a goal of adaptive control, though the work in [47] was the first to consider achieving it through CLFs. It is useful to review this work as it will guide the formulation taken for achieving safety in the presence of parametric model errors through the use of adaptive CBFs. In this setting we are interested in driving the system to a forced or unforced equilibrium $\mathbf{x}_e \in E$ of the open-loop system (3.209). Because we are considering potentially unbounded parameter errors, we will restrict our attention to global stability results. To ensure the well-posedness of the stabilization problem, I will assume that $\mathbf{F}(\mathbf{x}_e) = \mathbf{0}_n$, such that \mathbf{x}_e only needs to be an unforced or forced equilibrium point of the nominal model open-loop system:

$$\dot{\mathbf{x}} = \widehat{\mathbf{f}}(\mathbf{x}) + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{u}. \quad (3.214)$$

Next, I will relax the assumption that the controllers I consider be locally Lipschitz continuous on $E \times \mathbb{R}^p$, but rather require that they only be locally Lipschitz continuous on $(E \setminus \{\mathbf{x}_e\}) \times \mathbb{R}^p$. Furthermore, the controllers will need to render \mathbf{x}_e an equilibrium point of the state portion of the closed-loop system (3.212), such that:

$$\widehat{\mathbf{f}}(\mathbf{x}_e) + \widehat{\mathbf{g}}(\mathbf{x}_e)\mathbf{k}(\mathbf{x}_e, \widehat{\boldsymbol{\theta}}) = \mathbf{0}_n, \quad (3.215)$$

for all $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$. This leads to the following notion of adaptive stability:

Definition 25 (*Adaptive Stability*). Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (3.209). The composite closed-loop

system (3.212) is said to be *adaptively stable* with respect to \mathbf{x}_e if for any initial condition $(\mathbf{x}_0, \widehat{\boldsymbol{\theta}}_0) \in E \times \mathbb{R}^p$, we have that the solution $(\boldsymbol{\varphi}_{\mathbf{x}}, \boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}) : \mathbb{R}_{\geq 0} \rightarrow E \times \mathbb{R}^p$ is bounded on $\mathbb{R}_{\geq 0}$ and satisfies $\lim_{t \rightarrow \infty} \boldsymbol{\varphi}_{\mathbf{x}}(t) = \mathbf{x}_e$.

Note that the requirements for adaptive stability are rather weak in the sense that $\boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}(t)$ is not required to converge to $\boldsymbol{\theta}^*$. We will see, in fact, that convergence of $\boldsymbol{\varphi}_{\widehat{\boldsymbol{\theta}}}(t)$ to $\boldsymbol{\theta}^*$ is not necessary for $\boldsymbol{\varphi}_{\mathbf{x}}(t)$ to converge to \mathbf{x}_e . Establishing that a system is adaptively stable with respect to an equilibrium point \mathbf{x}_e can be achieved through an adaptive Lyapunov Function (aLF):

Definition 26 (*Adaptive Lyapunov Function (aLF)*). Let $\mathbf{x}_e \in E$ be an equilibrium point of the state closed-loop system of (3.212). A function $V_a : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E with $\frac{\partial V_a}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ locally Lipschitz continuous on E is said to be an *adaptive Lyapunov function (aLF)* for the state closed-loop system in (3.212) and equilibrium point \mathbf{x}_e if there exists $k_1, k_2, k_3, a \in \mathbb{R}_{>0}$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V_a(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.216)$$

$$\dot{V}_a(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \triangleq \frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\widehat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\widehat{\boldsymbol{\theta}} + \widehat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \right) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.217)$$

for all $\mathbf{x} \in E$ and $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$.

The importance of $\frac{\partial V_a}{\partial \mathbf{x}}$ being locally Lipschitz continuous on E will be important for ensuring the update law $\boldsymbol{\tau}$ is locally Lipschitz continuous on $E \times \mathbb{R}^p$. An aLF can be used to certify the adaptive stability of the closed-loop system (3.212), as captured in the following theorem:

Theorem 24. *Let $\mathbf{x}_e \in E$ be an equilibrium point of the state closed-loop system in (3.212). If there exists an adaptive Lyapunov function $V_a : E \rightarrow \mathbb{R}_{\geq 0}$ for the state closed-loop system in (3.212) and equilibrium point \mathbf{x}_e , then for any $\boldsymbol{\Gamma} \in \mathbb{S}_{>0}^p$, the update law:*

$$\boldsymbol{\tau}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) = \left(\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{F}(\mathbf{x}) \right)^\top, \quad (3.218)$$

renders the closed-loop system (3.212) adaptively stable with respect to \mathbf{x}_e .

It is useful to review this proof, as similar steps will be taken when considering adaptive safety through adaptive barrier functions:

Proof. Denote the parameter error:

$$\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}, \quad (3.219)$$

to be the difference between the actual and estimated parameters. Consider now the candidate composite Lyapunov function:

$$V(\mathbf{x}, \hat{\boldsymbol{\theta}}) = V_a(\mathbf{x}) + \frac{1}{2} \tilde{\boldsymbol{\theta}}^T \boldsymbol{\Gamma}^{-1} \tilde{\boldsymbol{\theta}}, \quad (3.220)$$

noting that it satisfies:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a + \frac{1}{2\lambda_{\max}(\boldsymbol{\Gamma})} \|\tilde{\boldsymbol{\theta}}\|^2 \leq V(\mathbf{x}, \hat{\boldsymbol{\theta}}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a + \frac{1}{2\lambda_{\min}(\boldsymbol{\Gamma})} \|\tilde{\boldsymbol{\theta}}\|^2. \quad (3.221)$$

Computing its derivative we can obtain:

$$\dot{V}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \dot{V}_a(\mathbf{x}, \boldsymbol{\theta}^*) - \tilde{\boldsymbol{\theta}}^T \boldsymbol{\Gamma}^{-1} \dot{\tilde{\boldsymbol{\theta}}}, \quad (3.222)$$

$$\leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a + \tilde{\boldsymbol{\theta}}^T \left(\left(\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{F}(\mathbf{x}) \right)^T - \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right). \quad (3.223)$$

It is now easy to see that using the update law:

$$\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \left(\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{F}(\mathbf{x}) \right)^T, \quad (3.224)$$

implies:

$$\dot{V}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.225)$$

from which we conclude that the solution $(\boldsymbol{\varphi}_{\mathbf{x}}, \boldsymbol{\varphi}_{\hat{\boldsymbol{\theta}}})$ is bounded on $\mathbb{R}_{\geq 0}$. It now follows from LaSalle's invariance principle [100, Theorem 4.4] that $\boldsymbol{\varphi}_{\mathbf{x}}$ converges to the largest invariant subset of the collection of points $\mathbf{x} \in E$ satisfying $k_3 \|\mathbf{x} - \mathbf{x}_e\|^a = 0$, which is $\mathbf{x} = \mathbf{x}_e$, such that $\lim_{t \rightarrow \infty} \boldsymbol{\varphi}_{\mathbf{x}}(t) = \mathbf{x}_e$. \square

The proof utilizes LaSalle's invariance principle, which has two effects. First, it does not allow us to draw conclusions about the convergence of $\boldsymbol{\varphi}_{\hat{\boldsymbol{\theta}}}$ to the true parameter $\boldsymbol{\theta}^*$. This is generally achieved by requiring conditions such as persistence of excitation [206], but this will not be considered in this thesis. Second, it only provides a weak bound on the transient behavior of the state of the system. In particular, the time derivative of V can be made negative by decreasing the parameter error, while increasing the deviation of $\boldsymbol{\varphi}_{\mathbf{x}}$ from the equilibrium point \mathbf{x}_e . It is this behavior that will have to be carefully

considered when considering forward invariance of a set in the state space when developing adaptive CBFs.

The notion of adaptive LFs naturally generalizes to adaptive CLFs, as captured in the following definition:

Definition 27 (*Adaptive Control Lyapunov Function (aCLF)*). Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (3.209). A function $V_a : E \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on E is said to be an *adaptive Control Lyapunov Function (aCLF)* for the open-loop system (3.209) and equilibrium point \mathbf{x}_e if there exists $k_1, k_2, k_3, a \in \mathbb{R}_{>0}$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V_a(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (3.226)$$

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \dot{V}_a(\mathbf{x}', \hat{\boldsymbol{\theta}}, \mathbf{u}) \triangleq \inf_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}') \left(\hat{\mathbf{f}}(\mathbf{x}') + \mathbf{F}(\mathbf{x}') \hat{\boldsymbol{\theta}} + \hat{\mathbf{g}}(\mathbf{x}') \mathbf{u} \right) < -k_3 \|\mathbf{x}' - \mathbf{x}_e\|^a, \quad (3.227)$$

for all $\mathbf{x} \in E$, $\mathbf{x}' \in E \setminus \{\mathbf{x}_e\}$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$.

Given an aCLF for the open-loop system (3.209) and equilibrium point $\mathbf{x}_e \in E$, we can define the following pointwise set:

$$K_{\text{aCLF}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \{\mathbf{u} \in \mathbb{R}^m \mid \dot{V}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a\}. \quad (3.228)$$

This leads to the following theorem typically seen with CLFs:

Theorem 25. *Let $\mathbf{x}_e \in E$ be an unforced or forced equilibrium point of the open-loop system (3.209) and let $V_a : E \rightarrow \mathbb{R}_{\geq 0}$ be an aCLF for the open-loop system (3.209) and equilibrium point \mathbf{x}_e . Then, the set $K_{\text{aCLF}}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ is non-empty for all $\mathbf{x} \in E$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$, and for any controller $\mathbf{k} : E \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ that is continuous on $E \times \mathbb{R}^p$, locally Lipschitz continuous on $(E \setminus \{\mathbf{x}_e\}) \times \mathbb{R}^p$, renders \mathbf{x}_e an equilibrium point of the state closed-loop system in (3.212), and satisfies $\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \in K_{\text{aCLF}}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ for all $\mathbf{x} \in E$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$, the function V_a is an aLF for the state closed-loop system in (3.212) and equilibrium point \mathbf{x}_e .*

Naturally, this motivates synthesizing a controller through convex optimization as follows:

$$\begin{aligned} \mathbf{k}_{\text{aCLF}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{u}_e\|^2 && (\text{aCLF-QP}) \\ \text{s.t. } \dot{V}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) &\leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a. \end{aligned}$$

While not provided in the thesis, continuity and local Lipschitz continuity results for controllers built through CLFs and optimization in Section 2.3 can be extended to this controller under appropriate small or continuous control property assumptions on V_a that hold for all parameter values $\hat{\theta}$ in some bounded set (noting that the proof of adaptive stability through aLFs yields that $\varphi_{\hat{\theta}}(t)$ is bounded).

Adaptive Control Barrier Functions

The strong parallels between CLFs and CBFs motivate the consideration of achieving safety for systems with parametric model error through the use of adaptive control in conjunction with CBFs. In spite of these parallels, one challenge that is hidden in the proof relating an aLF and adaptive stability is the transient behavior of the solution $\varphi_{\mathbf{x}}$. In particular, one does not have that the sublevel sets of the aLF V_a are forward invariant, but rather that the sublevel sets of the composite LF V are forward invariant. Because of this, the system state may leave sublevel sets of V_a while V decreases because the error between the parameter estimates and true parameters is sufficiently decreasing. This will be an issue if it is necessary to keep the state of the system within a particular set during the evolution of the system.

In an effort to translate the work on aCLFs to the setting of safety and CBFs and overcome the aforementioned challenges, I will propose a novel notion of adaptive safety:

Definition 28 (*Adaptive Safety*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h_a : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The composite closed-loop system (3.212) is said to be *adaptively safe* with respect to \mathcal{C} if there exists a $\delta \in \mathbb{R}_{>0}$ such that the set:

$$\mathcal{C}_\delta \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid h_a(\mathbf{x}) \geq \delta\}, \quad (3.229)$$

is non-empty and for any initial condition $(\mathbf{x}_0, \hat{\theta}_0) \in \mathcal{C}_\delta \times \Theta$ where $\Theta \subset \mathbb{R}^p$ is some collection of possible initial parameter estimates, we have that $\varphi_{\mathbf{x}}(t) \in \mathcal{C}$ for all $t \in \mathbb{R}_{\geq 0}$.

There are a few notable observations about this definition. First, it is posed for a set \mathcal{C} that can be represented as the 0-superlevel set of a function h_a that is continuously differentiable on E , rather than any general set in E . This

structure allows us to construct a contraction of \mathcal{C} by shifting the value of the superlevel set up to produce \mathcal{C}_δ . The state initial conditions are restricted to this set \mathcal{C}_δ to provide a buffer between when the system begins evolving (when adaptation begins) and when the system can encounter the boundary of the set \mathcal{C} . We will see when proving that an adaptive barrier function renders a system adaptively safe with respect to the set \mathcal{C} , the size of this buffer determines requirements on the adaption rate (the magnitude of the adaptive gain Γ). Second, we consider a restricted initial parameter set Θ , rather than all of \mathbb{R}^p . This initial parameter set typically contains the true parameters θ^* , its size will be related to possible initial parameter errors, and it will similarly yield requirements on the adaptive gain Γ .

Given this definition, I now provide a novel definition of adaptive barrier functions (aBFs):

Definition 29 (*Adaptive Barrier Function (aBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h_a : E \rightarrow \mathbb{R}$ that is continuously differentiable on E , with $\frac{\partial h_a}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ locally Lipschitz continuous on E , $\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ when $h_a(\mathbf{x}) = 0$, and suppose there exists an $\mathbf{x} \in E$ such that $h_a(\mathbf{x}) > 0$. The function h is an *adaptive barrier function (aBF)* for the state closed-loop system in (3.212) on \mathcal{C} if:

$$\dot{h}_a(\mathbf{x}, \hat{\theta}) \triangleq \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\hat{\theta} + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}, \hat{\theta}) \right) \geq 0, \quad (3.230)$$

for all $\mathbf{x} \in E$ and $\hat{\theta} \in \mathbb{R}^p$.

Observe that the typical extended class- \mathcal{K} function that lower bounds the time derivative has been replaced with 0. It is natural to wonder if this could be replaced with $-\alpha(h_a(\mathbf{x}))$ for some $\alpha \in \mathcal{K}^e$. I will show that this weakening of the lower bound will not lead to adaptive safety in the next subsection through a rigorous counterexample.

The following theorem relates adaptive safety and aBFs:

Theorem 26. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h_a : E \rightarrow \mathbb{R}$. If h_a is an aBF for the state closed-loop system in (3.212) on \mathcal{C} , then for any $c \in \mathbb{R}_{>0}$, there exists an adaptive gain $\Gamma \in \mathbb{S}_{>0}^p$ such that the update law:*

$$\tau(\mathbf{x}, \hat{\theta}) = - \left(\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x}) \right)^\top, \quad (3.231)$$

renders the closed-loop system (3.212) adaptively safe with respect to \mathcal{C} with:

$$\Theta = \left\{ \widehat{\boldsymbol{\theta}}_0 \in \mathbb{R}^p \mid \|\widehat{\boldsymbol{\theta}}_0 - \boldsymbol{\theta}^*\| \leq c \right\}. \quad (3.232)$$

This theorem states that for any maximum initial parameter error c , there exists an adaptive gain $\boldsymbol{\Gamma}$ and update law $\boldsymbol{\tau}$ that renders the composite closed-loop system (3.212) adaptively safe with respect to \mathcal{C} . In contrast to adaptive stabilization, the same adaptive gain may not work for all possible initial parameter errors, as shown in the following proof:

Proof. By assumption, there exists an $\mathbf{x} \in \mathcal{C}$ such that $h_a(\mathbf{x}) \triangleq \delta > 0$. Let $\boldsymbol{\Gamma} \in \mathbb{S}_{>0}^p$ satisfy:

$$\lambda_{\min}(\boldsymbol{\Gamma}) \geq \frac{c^2}{2\delta}. \quad (3.233)$$

Note that this inequality requires the adaptive gain to grow with larger possible initial parameter errors, as well as grow as the buffer between the set \mathcal{C} and \mathcal{C}_δ becomes smaller. Notice that for a fixed adaptive gain, the size of the buffer can be smaller for smaller initial parameter errors. Consider a continuously differentiable function $h : E \times \mathbb{R}^p \rightarrow \mathbb{R}$ defined as:

$$h(\mathbf{x}, \widehat{\boldsymbol{\theta}}) = h_a(\mathbf{x}) - \frac{1}{2}(\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}})^\top \boldsymbol{\Gamma}^{-1}(\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}). \quad (3.234)$$

Observe that rather than add the quadratic function of the parameter error as was done with aCLFs, I have subtracted it. To see that 0 is a regular value of h , consider its gradient:

$$\begin{bmatrix} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \\ \frac{\partial h}{\partial \widehat{\boldsymbol{\theta}}}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \end{bmatrix} = \begin{bmatrix} \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \\ (\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}})^\top \boldsymbol{\Gamma}^{-1} \end{bmatrix}. \quad (3.235)$$

Suppose that $h(\mathbf{x}, \widehat{\boldsymbol{\theta}}) = 0$. If $h_a(\mathbf{x}) = 0$, then $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) = \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$. If $h_a(\mathbf{x}) \neq 0$, then $\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}} \neq \mathbf{0}_p$, such that $\frac{\partial h}{\partial \widehat{\boldsymbol{\theta}}}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) = (\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}})^\top \boldsymbol{\Gamma}^{-1} \neq \mathbf{0}_p$ as $\boldsymbol{\Gamma}^{-1}$ is full rank. Thus 0 is a regular value of h . Next, note that by construction:

$$(\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}})^\top \boldsymbol{\Gamma}^{-1}(\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}) \leq \lambda_{\max}(\boldsymbol{\Gamma}^{-1}) \|\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}\|^2, \quad (3.236)$$

$$= \frac{1}{\lambda_{\min}(\boldsymbol{\Gamma})} \|\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}\|^2. \quad (3.237)$$

Thus for any initial condition $(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0) \in \mathcal{C}_\delta \times \Theta$, we have that:

$$h(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0) = h_a(\mathbf{x}_0) - \frac{1}{2}(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0)^\top \boldsymbol{\Gamma}^{-1}(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0), \quad (3.238)$$

$$\geq \delta - \frac{1}{2\lambda_{\min}(\boldsymbol{\Gamma})} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0\|^2, \quad (3.239)$$

$$\geq \delta - \frac{2\delta}{2c^2} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0\|^2 \geq 0. \quad (3.240)$$

Thus, the adaptive gain $\boldsymbol{\Gamma}$ has been chosen large enough to ensure that the value of the composite barrier function h is positive for all initial conditions $(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0) \in \mathcal{C}_\delta \times \Theta$. The time derivative of h is given by:

$$\begin{aligned} \dot{h}(\mathbf{x}, \hat{\boldsymbol{\theta}}) &= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \hat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right) + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\Gamma}^{-1} \dot{\hat{\boldsymbol{\theta}}}, \\ &= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \hat{\mathbf{g}}(\mathbf{x})\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right) + \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x})\hat{\boldsymbol{\theta}} \\ &\quad - \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x})\tilde{\boldsymbol{\theta}} + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}), \\ &= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\hat{\boldsymbol{\theta}} + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right) \\ &\quad + \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x})\tilde{\boldsymbol{\theta}} + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}), \\ &\geq \tilde{\boldsymbol{\theta}}^\top \left(\left(\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x})\mathbf{F}(\mathbf{x}) \right)^\top + \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right), \\ &\geq 0, \end{aligned}$$

where the first inequality follows as h_a is an aBF for the state closed-loop system in (3.212) on \mathcal{C} and the second inequality follows by the choice of update law $\boldsymbol{\tau}$. Thus by Nagumo's theorem [104], [123], we have that $h(\boldsymbol{\varphi}_x(t), \boldsymbol{\varphi}_{\hat{\boldsymbol{\theta}}}(t)) \geq 0$ for all $t \in \mathbb{R}_{\geq 0}$. Lastly, we may conclude that:

$$h_a(\boldsymbol{\varphi}_x(t)) \geq \frac{1}{2}(\boldsymbol{\theta}^* - \boldsymbol{\varphi}_{\hat{\boldsymbol{\theta}}}(t))^\top \boldsymbol{\Gamma}^{-1}(\boldsymbol{\theta}^* - \boldsymbol{\varphi}_{\hat{\boldsymbol{\theta}}}(t)), \quad (3.241)$$

$$\geq 0, \quad (3.242)$$

such that $\boldsymbol{\varphi}_x(t) \in \mathcal{C}$ for all $t \in \mathbb{R}_{\geq 0}$ as desired. \square

The proof reveals that all superlevel sets of h are forward invariant, which is a potentially conservative requirement, but also a non-trivial one to circumnavigate. As h can not be computed without knowing the true parameters $\boldsymbol{\theta}^*$, it is not possible to set $\dot{h}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \geq -\alpha(h(\mathbf{x}, \hat{\boldsymbol{\theta}}))$ for some extended class- \mathcal{K} function α as is typical with CBFs. Furthermore, we have that $h_a(\mathbf{x}) \geq h(\mathbf{x}, \hat{\boldsymbol{\theta}})$ for all

$\mathbf{x} \in E$ and $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$, implying that $-\alpha(h_a(\mathbf{x})) \leq -\alpha(h(\mathbf{x}, \widehat{\boldsymbol{\theta}}))$ for all $\mathbf{x} \in E$, $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$, and $\alpha \in \mathcal{K}^e$. Thus setting $\dot{h}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \geq -\alpha(h_a(\mathbf{x}))$ does not yield the desired lower bound on $\dot{h}(\mathbf{x}, \widehat{\boldsymbol{\theta}})$. One may note that setting $\dot{h}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \geq -\alpha(h_a(\mathbf{x}))$ leads to $\dot{h}(\mathbf{x}, \widehat{\boldsymbol{\theta}}) \geq 0$ when $h_a(\mathbf{x}) = 0$, or when the state is on the boundary of \mathcal{C} . This fact is concurrent with the common forward invariance proof technique utilizing Nagumo's theorem [104], [123]. Despite this, it is in fact possible to construct simple examples (in \mathbb{R}^2) such that the state must leave the set \mathcal{C} defined by h_a for any choice of differentiable $\alpha \in \mathcal{K}^e$ and $\boldsymbol{\Gamma} \in \mathbb{S}_{>0}^p$, as shown in the following subsection. The reason these challenges are avoided with adaptive CLFs is because the proof concludes by using LaSalle's invariance principle, and neglecting the transient behavior of $\boldsymbol{\varphi}_{\mathbf{x}}$. It is precisely the transient behavior which describes whether or not the state remains within a particular set during the course of its evolution.

Extending the notion of adaptive BFs to adaptive CBFs requires a small amount of care. In particular, we must be careful when translating the non-strict inequality in (3.244) to a strict inequality as is typical when defining a CBF variant. If there exists a state $\mathbf{x} \in \text{Int}(\mathcal{C})$ such that $\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) = \mathbf{0}_n$, then it will be impossible to satisfy the strict inequality at that state through the choice of input. This eliminates considering any set \mathcal{C} that is compact, as any continuously differentiable function h_a must⁸ have a critical point in the interior of \mathcal{C} . Fortunately, at any such state $\mathbf{x} \in \text{Int}(\mathcal{C})$, $\dot{h}_a(\mathbf{x}, \widehat{\boldsymbol{\theta}}, \mathbf{u}) = 0$ for all $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$ and $\mathbf{u} \in \mathbb{R}^m$, so the aBF property is still satisfied at these states as required for adaptive safety even if the strict inequality can't be satisfied there. Note that such states must not appear on the boundary of \mathcal{C} so that Nagumo's theorem may be used. To this end, I will denote the following set:

$$\mathcal{A}_{\mathcal{C}} = \left\{ \mathbf{x} \in E \mid \mathbf{x} \in \text{Int}(\mathcal{C}) \text{ and } \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) = \mathbf{0}_n \right\}, \quad (3.243)$$

and propose the following definition:

Definition 30 (*Adaptive Control Barrier Function (aCBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h_a : E \rightarrow \mathbb{R}$ that is continuously differentiable on E with $\frac{\partial h_a}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ locally Lipschitz continuous on E , $\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ when $h_a(\mathbf{x}) = 0$, and suppose there exists an $\mathbf{x} \in E$ such that $h_a(\mathbf{x}) > 0$. The function h_a is said to be an *adaptive Control Barrier Function (aCBF)* for the

⁸I do not offer a proof of this in this thesis, but it is relatively easy to show through first-order optimality conditions.

open-loop system (3.209) on \mathcal{C} if:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) \triangleq \sup_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) \left(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\hat{\boldsymbol{\theta}} + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} \right) > 0, \quad (3.244)$$

for all $\mathbf{x} \in E \setminus \mathcal{A}_{\mathcal{C}}$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$.

Given an aCBF for the open-loop system (3.209) on \mathcal{C} , we can define the following pointwise set:

$$K_{\text{aCBF}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \{\mathbf{u} \in \mathbb{R}^m \mid \dot{h}(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) \geq 0\}. \quad (3.245)$$

This leads to the following theorem:

Theorem 27. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h_a : E \rightarrow \mathbb{R}$ and let h_a be an adaptive CBF for the open-loop system (3.209) on \mathcal{C} . Then, the set $K_{\text{aCBF}}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ is non-empty for all $\mathbf{x} \in E$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$, and for any controller $\mathbf{k} : E \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on $E \times \mathbb{R}^p$, and satisfies $\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \in K_{\text{aCBF}}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ for all $\mathbf{x} \in E$ and $\hat{\boldsymbol{\theta}} \in \mathbb{R}^p$, the function h_a is an aBF for the state closed-loop system in (3.212) on \mathcal{C} .*

Letting $\mathbf{k}_{\text{nom}} : E \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ be locally Lipschitz continuous on $E \times \mathbb{R}^p$, we may consider the following controller:

$$\begin{aligned} \mathbf{k}_{\text{aCBF}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\|^2 && \text{(aCBF-QP)} \\ &\text{s.t. } \dot{h}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) \geq 0. \end{aligned}$$

This controller can be shown to be locally Lipschitz continuous on $E \times \mathbb{R}^p$ under a variant of the continuous control property that requires for each $(\mathbf{x}, \hat{\boldsymbol{\theta}}) \in \mathcal{A}_{\mathcal{C}} \times \mathbb{R}^p$, there exists an open set $D_{(\mathbf{x}, \hat{\boldsymbol{\theta}})} \subseteq E \times \mathbb{R}^p$ and a constant $L_{(\mathbf{x}, \hat{\boldsymbol{\theta}})} \in \mathbb{R}_{\geq 0}$ such that for any $(\mathbf{y}, \hat{\boldsymbol{\psi}}) \in D_{(\mathbf{x}, \hat{\boldsymbol{\theta}})}$, there exists an input $\mathbf{u} \in \mathbb{R}^m$ such that:

$$\|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\| \leq L_{(\mathbf{x}, \hat{\boldsymbol{\theta}})} \|\mathbf{y} - \mathbf{x}\| + L_{(\mathbf{x}, \hat{\boldsymbol{\theta}})} \|\hat{\boldsymbol{\psi}} - \hat{\boldsymbol{\theta}}\|, \quad (3.246)$$

and:

$$\dot{h}_a(\mathbf{y}, \hat{\boldsymbol{\psi}}, \mathbf{u}) \geq 0. \quad (3.247)$$

I note that it is important that \mathbf{k}_{aCBF} be locally Lipschitz continuous on $E \times \mathbb{R}^p$, and not just locally Lipschitz continuous on $(E \setminus \mathcal{A}_{\mathcal{C}}) \times \mathbb{R}^p$ and continuous on $E \times \mathbb{R}^p$. An important requirement of Nagumo's theorem [104], [123] used in the

proof of adaptive safety are that solutions are unique. Theorem 3 establishes the existence of a unique solution to the closed-loop system (2.2) from the initial condition $\mathbf{x}_e \in E$ at which the controller \mathbf{k} is only continuous (but not locally Lipschitz continuous) by requiring that \mathbf{x}_e be an equilibrium point of the closed-loop system (2.2) and the existence a local exponential Lyapunov function. If \mathbf{k}_{aCBF} was only locally Lipschitz continuous on $(E \setminus \mathcal{A}) \times \mathbb{R}^p$, using Theorem 3 would require that each point $\mathbf{x}_a \in \mathcal{A}_C$ be an equilibrium point of the state closed-loop system in (3.212). Furthermore, if $\mathbf{x} \in \mathcal{A}_C$ is a local minimum of the function h_a we would require an open set $D_{(\mathbf{x}, \hat{\boldsymbol{\theta}})}$ such that $\dot{h}_a(\mathbf{y}, \hat{\boldsymbol{\psi}}) \leq 0$ for all $(\mathbf{y}, \hat{\boldsymbol{\psi}}) \in D_{(\mathbf{x}, \hat{\boldsymbol{\theta}})}$. Given the requirements of an aBF, this would require that $\dot{h}_a(\mathbf{y}, \hat{\boldsymbol{\psi}}) = 0$ for all $(\mathbf{y}, \hat{\boldsymbol{\psi}}) \in D_{(\mathbf{x}, \hat{\boldsymbol{\theta}})}$. We can avoid inducing these unwanted stable equilibrium points throughout the set \mathcal{C} by requiring the Lipschitz-like continuous control property noted above.

Counterexample

I now analyze the aCBF condition to verify that, in fact, it is not overly conservative. In particular, changing the aCBF condition $\dot{h}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) \geq 0$ to $\dot{h}_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) \geq -\alpha(h_a(\mathbf{x}))$ does not necessarily lead to adaptive safety. Consider the simple open-loop system given by:

$$\dot{x} = \theta^* + u, \quad (3.248)$$

with $\theta^* \in \mathbb{R}$ unknown and the function $h_a : \mathbb{R} \rightarrow \mathbb{R}$ defined as:

$$h_a(x) = 1 - x^2, \quad (3.249)$$

yielding the set $\mathcal{C} = \{x \in \mathbb{R} \mid x^2 \leq 1\}$. It is easy to see that h_a is an aCBF for the open-loop system (3.248) on \mathcal{C} . Let $\delta \in (0, 1)$, $c \in \mathbb{R}_{\geq 0}$, and consider a composite barrier function given by:

$$h(x, \hat{\boldsymbol{\theta}}) = h_a(x) - \frac{1}{2}\gamma^{-1}\tilde{\boldsymbol{\theta}}^2, \quad (3.250)$$

with $\gamma \in \mathbb{R}_{>0}$ satisfying:

$$\gamma \geq \frac{c^2}{2\delta}. \quad (3.251)$$

Define the following sets, which are seen in Figure 3.17:

$$U = \{(x, \tilde{\boldsymbol{\theta}}) \in \mathbb{R}^2 \mid x \in \mathcal{C}\}, \quad (3.252)$$

$$H_0 = \{(x, \tilde{\boldsymbol{\theta}}) \in \mathbb{R}^2 \mid h(x, \tilde{\boldsymbol{\theta}}) \geq 0\}. \quad (3.253)$$

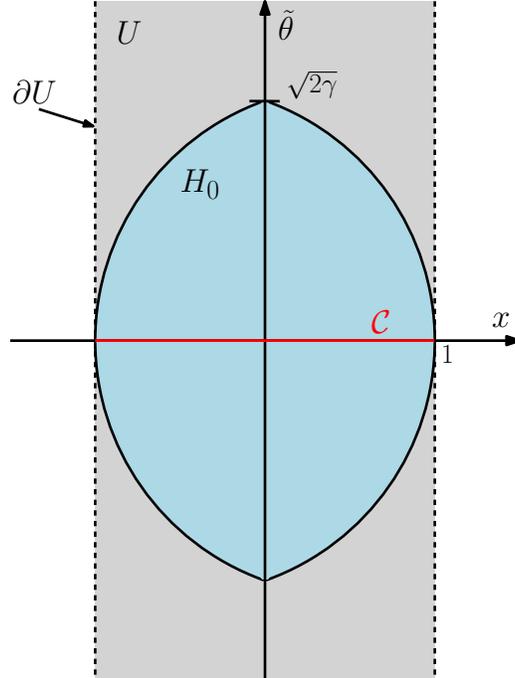


Figure 3.17. Schematic of a CBF counterexample.

Note that the set U extends infinitely along the $\tilde{\theta}$ -axis, and completely contains H_0 . Furthermore, $H_0 \cap \partial U = \{(-1, 0), (1, 0)\}$. The time derivative of the composite barrier function h is given by:

$$\dot{h}(x, \hat{\theta}, u) = -2x(\hat{\theta} + u) + \tilde{\theta}(-2x + \tau(x, \hat{\theta})), \quad (3.254)$$

for $\hat{\theta} = \gamma\tau(x, \hat{\theta})$. Choose the update law $\tau(x, \hat{\theta}) = 2x$ and controller $k(x, \hat{\theta}) = -\hat{\theta} + \frac{1}{2}x\alpha(h_a(x))$, with continuously differentiable extended class- \mathcal{K} function α . Note that both τ and k are locally Lipschitz continuous on $E \times \mathbb{R}^p$. These choices yield:

$$\dot{h}(x, \hat{\theta}) = -x^2\alpha(h_a(x)) \geq -\alpha(h_a(x)), \quad (3.255)$$

as when $\alpha(h_a(x)) \geq 0$, $x^2 \leq 1$, and when $\alpha(h_a(x)) \leq 0$, $x^2 \geq 1$. Noting the construction of U , we have the implication that $(x, \tilde{\theta}) \in U \implies \dot{h}(x, \hat{\theta}) \leq 0$. Instead of considering the composite dynamics in (3.212), we will instead consider an equivalent composite closed-loop system built using the parameter error as a state:

$$\begin{bmatrix} \dot{x} \\ \dot{\tilde{\theta}} \end{bmatrix} = \begin{bmatrix} \tilde{\theta} + \frac{1}{2}x\alpha(h_a(x)) \\ -2\gamma x \end{bmatrix} = \begin{bmatrix} \tilde{\theta} - F(x) \\ -g(x) \end{bmatrix}, \quad (3.256)$$

which has an unstable equilibrium point at the origin. This system is an example of a Liénard system (like the Van der Pol oscillator) as in [101, Section

3.8], with $F(x) = -\frac{1}{2}x\alpha(h_a(x))$ and $g(x) = 2\gamma x$. For systems of this form, the following theorem [101, Section 3.8, Theorem 1], attributed to Liénard, provides the existence of a unique, stable limit cycle:

Theorem 28. *Under the assumption that F, g are continuously differentiable on \mathbb{R} , F and g are odd functions of x , $xg(x) > 0$ for $x \neq 0$, $F(0) = 0$, $F'(0) < 0$, F has a single positive zero at $x = a$, and F increases monotonically to infinity for $x \geq a$ as $x \rightarrow \infty$, it follows that the Liénard system (3.256) has exactly one limit cycle and it is stable.*

As α is continuously differentiable in addition to being extended \mathcal{K}_∞ , the assumptions of this theorem are met by the functions given in (3.256). Note that $a = 1$ in this given example. Thus we can conclude that the system (3.256) has a stable periodic orbit, which is denoted $\Phi \subset \mathbb{R}^2$. Denote the open set in \mathbb{R}^2 enclosed by the limit cycle as $\text{int}(\Phi)$. Additionally, the proof of this theorem as in [101, Section 3.8] implies the following corollary:

Corollary 1. *The stable limit cycle Φ is symmetric about the origin and passes through a point, denoted as $P_2 = (x_2, \tilde{\theta}_2)$, such that $x_2 > a$.*

Given that $a = 1$, this corollary reveals that the stable limit cycle leaves the set U , for which the state is considered safe. Additionally, as the limit cycle is symmetric about the origin, and the origin is an unstable equilibrium, the origin is contained in $\text{int}(\Phi)$.

This corollary also implies that $H_0 \subset (\Phi \cup \text{int}(\Phi))$. To see this, note that as the limit cycle encircles the origin, it must reenter the set U after leaving the point P_2 . At any point $\mathbf{v} = (v_1, v_2) \in U$ that the limit cycle enters, we must have $h(v_1, v_2) \leq 0$, given the two points in $H_0 \cap \partial U$. Once the limit cycle enters U , we have $\dot{h}(x, \hat{\theta}) \leq 0$ until the limit cycle leaves U as previously noted. Thus, $h(x, \hat{\theta}) \leq 0$ along the portion of the limit cycle contained in U , implying $H_0 \subset (\Phi \cup \text{int}(\Phi))$. To complete the proof, we will employ the following definition [100, p. 127] and lemma [100, Lemma 4.1]:

Definition 31. The positive limit set L^+ of a solution $(\varphi_x, \varphi_{\tilde{\theta}})$ is defined as all points $p \in \mathbb{R}^2$ such that there is a sequence $\{t_n\}$ with $t_n \rightarrow \infty$ as $n \rightarrow \infty$, and $(\varphi_x(t_n), \varphi_{\tilde{\theta}}(t_n)) \rightarrow p$ as $n \rightarrow \infty$.

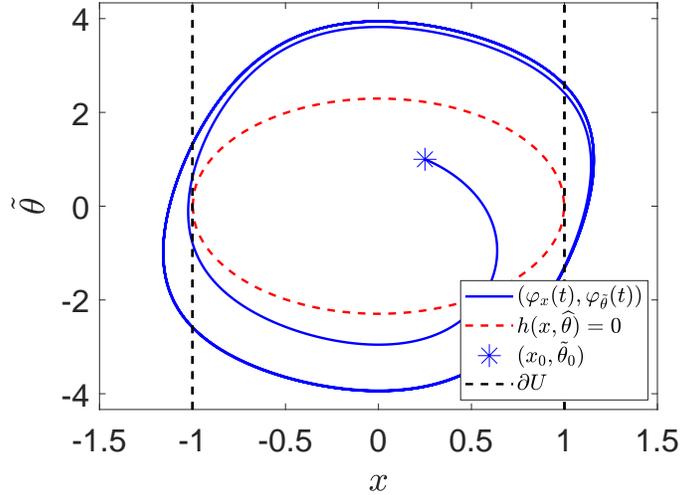


Figure 3.18. Simulation trajectory of a CBF counterexample. The system state leaves the prescribed safe set as it converges to the limit cycle.

Lemma 2. *If a solution $(\varphi_x, \varphi_{\hat{\theta}})$ of (3.256) is bounded on $\mathbb{R}_{\geq 0}$, then its positive limit set L^+ is a nonempty, compact, invariant set, and $(\varphi_x, \varphi_{\hat{\theta}})$ approaches L^+ as $t \rightarrow \infty$.*

Note that the unstable equilibrium point is not contained within the positive limit set L^+ of a solution $(\varphi_x, \varphi_{\hat{\theta}})$ starting from any initial condition $(x_0, \hat{\theta}_0) \in H_0 \setminus \{\mathbf{0}_2\}$. As the 0-superlevel set of h , and thus all possible initial conditions given our bound on $\tilde{\theta}_0$, are contained inside the limit cycle, all solutions to (3.256) are bounded (by the limit cycle). Furthermore, for the solution $(\varphi_x, \varphi_{\hat{\theta}})$ starting from any initial condition $(x_0, \hat{\theta}_0) \in H_0 \setminus \{\mathbf{0}_2\}$, we have that $L^+ = \Phi$, and thus all solutions starting in the 0-superlevel of h except for the origin approach Φ . As the point $P_2 \in \Phi$, and $P_2 \notin U$, we see that any solution starting in the 0-superlevel set of h leaves the desired state safe set \mathcal{C} . Hence, the relaxation does not achieve safety of the state, as seen in Figure 3.18.

Simulation Results

To demonstrate how an aCBF can be used to render a system adaptively safe, I consider the problem of adaptive cruise control (ACC) as posed in [11]. The

dynamics of the system⁹ are given by:

$$\frac{d}{dt} \begin{bmatrix} v \\ D \end{bmatrix} = \begin{bmatrix} 0 \\ v_0 - v \end{bmatrix} - \frac{1}{m} \begin{bmatrix} 1 & v & v^2 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} f_0 \\ f_1 \\ f_2 \end{bmatrix} + \begin{bmatrix} \frac{1}{m} \\ 0 \end{bmatrix} u, \quad (3.257)$$

with v the velocity of the vehicle, D the distance between the vehicle and a leading vehicle traveling at a fixed velocity v_0 , m the vehicle's mass, and f_0 , f_1 , and f_2 unknown parameters associated with rolling frictional force. Define the state as $\mathbf{x} = (v, D)$. In this problem, I seek to drive the velocity to a desired velocity, v_d , while simultaneously ensuring the distance between the vehicles satisfies a safety constraint given by:

$$D \geq 1.8v. \quad (3.258)$$

The parameters f_0 , f_1 , and f_2 are often determined empirically, and if they are not accurate, the desired velocity may not be accurately tracked. Furthermore, if the parameters do not exactly match the true parameters, it may not be possible to certify that the system will satisfy the safety constraint. The control objective of tracking a desired velocity can be achieved with a hand-designed controller k_{nom} , or encoded using a CLF, and the safety constraint can be encoded using a CBF. Additional constraints on the maximum acceleration and deceleration can be enforced to maintain passenger comfort. To handle error in the parameters, I utilize the tool of aCBFs to maintain and update estimates of these parameters. An aCBF that yields desirable results is defined as the following continuously differentiable function:

$$h_a(v, D) = \begin{cases} \beta^2 & \text{if } D - 1.8v \geq \beta, \\ \beta^2 - (D - 1.8v - \beta)^2 & \text{if } D - 1.8v < \beta, \end{cases} \quad (3.259)$$

for $\beta > 0$. This particular construction of h_a is constant away from the safety boundary and diminishes to 0 (quadratically to preserve continuous differentiability) as the boundary is approached. In practice, this is to handle the fact that superlevel sets of the composite safety function h are forward invariant. In regions where h_a is constant, $\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}) = \mathbf{0}_n$, and thus the update law in (3.231), is $\mathbf{0}_p$, thus making $\dot{h}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = 0$.

⁹This system will violate the assumption that $\mathbf{F}(\mathbf{x}_e) = \mathbf{0}_2$ for the point $\mathbf{x}_e \in \mathbb{R}^2$ that I try to stabilize the system to. In general, this will compromise stability guarantees, which are not well-posed. Still, for this example, I am able to achieve interesting behavior showing the benefit of using adaptive elements.

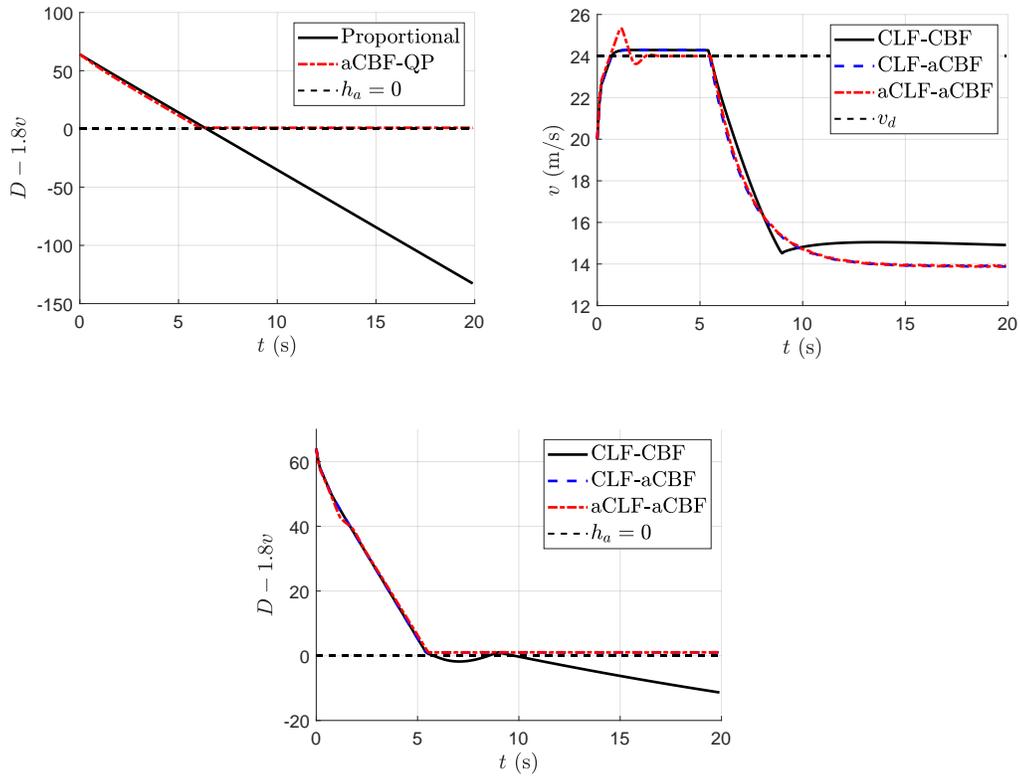


Figure 3.19. Comparison of different adaptive and non-adaptive control methodologies. (Top, Left) An aCBF controller is able to enforce safety of the nominal controller k_{nom} . (Top, Right) An aCLF-aCBF controller is able to track the desired velocity with zero steady state error (Top, Right). (Bottom) Both aCBF controllers are able to keep the vehicle within the safe region for all time.

aCBF-QP Controller: A simple proportional controller k_{nom} on tracking error $v - v_d$ can be implemented and achieve good tracking performance, but is not necessarily safe. A CBF alone would not ensure the safety of this controller with model parametric model error, but by using the proportional controller k_{nom} in a (aCBF-QP) controller, safety can be achieved.

aCLF-aCBF-QP Controller: Additionally, we can unify an aCLF and an aCBF in a quadratic program based controller to receive the benefits of optimal and adaptive tracking while remaining safe. Separate estimates of the parameters are maintained for the aCLF and the aCBF, as the form of the update laws in (3.218) and (3.231) may not be simultaneously satisfiable for only one estimate of the parameters. The CLF in [11] on the velocity tracking error $v - v_d$, given by $V_a = (v - v_d)^2$, also satisfies the aCLF condition (3.227). Letting $\hat{\theta}$ and $\hat{\psi}$ be parameter estimates associated with the aCLF and aCBF,

respectively, I formulate a QP-based controller:

$$\begin{aligned}
 k(\mathbf{x}, \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\psi}}) = \operatorname{argmin}_{u \in \mathbb{R}^m} & J(u) + c_V(\mathbf{x})\delta_V + c_p(\mathbf{x})\delta_p \\
 \text{s.t. } & \dot{V}_a(\mathbf{x}, \widehat{\boldsymbol{\theta}}, u) \leq -k_3(v - v_d)^2 + \delta_V, \\
 & \dot{h}_a(\mathbf{x}, \widehat{\boldsymbol{\psi}}, u) \geq 0, \\
 & u \leq u_{max} + \delta_p, \\
 & u \geq -u_{max} - \delta_p, \\
 & \delta_V, \delta_p \geq 0,
 \end{aligned}$$

with parameter updates for $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\psi}}$ as in (3.224) and (3.231), respectively. The function J is a cost function on the input, and the variables δ_V and δ_p are relaxations to the optimization problem to ensure its feasibility, while safety is ensured. The functions c_V and c_p are Lipschitz continuous and are used to ensure the optimization problem is well-conditioned.

With initial parameter estimates $[\widehat{f}_0 \ \widehat{f}_1 \ \widehat{f}_2] = 10 [f_0^* \ f_1^* \ f_2^*]$ (less friction than modeled), the results of this controller appear in Figure 3.19. We see that the proportional controller fails to keep the vehicle safe, but filtering it with the aCBF-QP keeps it safe (with $D \geq 1.8v$ for all time) even with parametric model error. A CLF-CBF controller with no adaptive elements fails to either track the desired velocity (with steady state error) or keep the vehicle safe. The CLF-aCBF controller keeps the vehicle safe but has steady state tracking error, while an aCLF-aCBF controller accurately tracks the desired velocity with no steady state error, and keeps the vehicle safe.

Conclusion

In this section I have presented a novel definition of adaptive Control Barrier Functions as a tool for safety-critical control synthesis in the presence of parametric model error. I begin by reviewing adaptive CLFs as formulated in [47] as a starting point for developing adaptive CBFs. I then propose a notion of adaptive safety and a definition of adaptive BFs that ensure adaptive safety. Subsequently, I introduce a notion of adaptive CBFs and explain some of the challenges that must be considered when ensuring the controllers producing using adaptive CBFs satisfy regularity properties. I then show that the conservative nature of adaptive CBFs is not unnecessary through a rigorous counter example. I lastly highlight the usefulness of adaptive elements in simulation with an adaptive cruise control system.

3.10 Conclusion and Future Work

In this chapter I have presented a collection of results for addressing general model error and parametric model error. In Section 3.3 I presented an episodic learning framework for learning general model errors as they impact the time derivative of a CLF and local exponential stability. In Section 3.4 I explored how residual learning error could be understood as a disturbance in a closed-loop system, and how robustness to these disturbances can be codified with Projection-to-State Stability. I additionally explore how possible model errors can be restricted through the use of a data set. In Section 3.5 I extend the preceding episodic learning algorithm to the setting of safety by considering how model error can be learned directly as it impacts the time derivative of a CBF, with an experimental demonstration. In Section 3.6 I rigorously codify safety guarantees that can be achieved in the presence of residual learning error through the notion of Projection-to-State Safety, and demonstrate the ability to improve safety guarantees through learning in simulation and experimentally. Building off of these previous results, in Section 3.7 I consider some of the challenges that arise when learning with structured models that preserve convexity of downstream control synthesis. To resolve these challenges, I present an episodic learning framework that uses an unstructured learning model, and demonstrate the ability to achieve safe walking on stepping-stones with a bipedal platform both in simulation and experimentally. This work culminates in Section 3.8 in which I propose a data-driven control framework that ensures robustness to all model errors permitted by a data set. I propose a convex optimization-based controller specified via a second-order cone program, and develop a rigorous theoretical understanding of this controller’s feasibility using an analog of the alternative CLF and CBF condition. Lastly, in Section 3.9 I propose a notion of adaptive CBFs for achieving safety in the presence of parametric model error. I further highlight how the conservative nature of the proposed definition is necessary through a rigorous counterexample.

I believe there are two major directions of future work worthwhile pursuing based on the contributions of this chapter. The first direction pertains to the data-driven control framework in Section 3.8. In the context of data-driven control, there are challenges with implementing the proposed controller with large data sets. This is due to the fact that the size of the optimization problem scales with the size of the data set. A similar challenge is seen in non-parametric learning methods, such as GPR, where evaluating a model scales

in complexity with the size of the data set [226]. One potential approach for resolving this issue is developing heuristics for selecting relevant data close to a test point while achieving necessary input variance, similar to the notion of k -nearest neighbors with diversity [227]. Another approach would be to build outer-approximations of the model error sets offline that can be quickly utilized in online control synthesis. There is perhaps great potential in this direction because the model error sets are defined by a collection of convex constraints, and there exist efficient ways for approximating such sets [30, Chapter 11]. Another worthwhile direction is extending this framework to the sampled-data setting in Chapter 6, as well as capturing robustness to disturbances and measurement error discussed in Chapters 4 and 5, respectively.

The other direction of future work looks at further developing the notion of adaptive CBFs established in Section 3.9, with some work already being considered in this direction [191], [196], [197], [199], [200]. The main objective of much of this work is in reducing the conservative nature of the adaptive CBFs presented in this thesis through the use of data. In particular, these approaches seek to find ways to introduce the extended class- \mathcal{K} function back into the aCBF formulation, at the expense of potentially more conservativeness until adequate data can be collected. Building upon this work with methods for collecting relevant data necessary for adaptation is one future direction. Another direction of work is to address the challenge of parametric model error in the function \mathbf{g} , which has not been considered until only very recently in [200]. This opens up the classes of systems that adaptive systems can be deployed on, such as robotic systems that have error in their mass and inertia parameters.

DISTURBANCE-ROBUST SAFETY-CRITICAL CONTROL

In this chapter I will present results from my work on developing methods for safety-critical control that are both performant and robust to disturbances. Disturbances, either unmatched or matched, can serve as a sort of catch-all for various types of challenges in nonlinear control design. In fact, model errors, measurement errors, and sampling effects, as presented in Chapters 3, 5, and 6, respectively, can all be cast as disturbances in a nonlinear system. Often times though, this perspective can obscure important structural properties of these different types of challenges that can be used to synthesize controllers that are both robust and performant. Instead, disturbances are best for capturing challenges for which there is no underlying structure that can easily be understood. For robotic systems, this may include things such as perturbing external forces created by the environment (objects, people, or other robots bumping into the robot), or complex actuator dynamics that are abstracted away when assuming that the inputs to the system are torques (for systems with electric motors, the inputs are typically voltage or current commands which are passed to a motor controller that attempts to produce a certain motor torque). In the former case, there may be no reasonable model that could be constructed to capture the myriad of potential unexpected interactions, even from data, while in the latter case, constructing a first-principles model may require significant work, and collecting data to build a data-driven model may require extensive additional sensing capabilities. Moreover, if a robot is subjected to multiple challenges, such as both external forces and actuator dynamics (which they likely will be), the modeling challenges becomes increasingly difficult. Instead, one may simply capture the collective effects as a disturbance that the control design must be robust to. Thus, disturbance robustness is a general form of robustness to the many different challenges that can be captured as disturbances (even if they can be captured using more structure, as in the other chapters of this thesis).

Instead of trying to ascribe structure to the system's environment or produce models for complex underlying structures, designs that consider disturbances generally require the control designer to produce a bound on the potential

disturbance signal. This in and of itself can be a difficult problem, and producing a weak bound can have significant consequences for the performance of a system. Intuitively, for larger potential disturbances, a larger degree of robustness is required. If the system never actually experiences such a large disturbance, then the system is potentially conceding a significant amount of performance for robustness to something it never encounters. Moreover, if it were to encounter this disturbance, robustness in the controller would be most important near the boundary of a set that is to be kept forward invariant, rather than deep inside the set, as the disturbance must take time to make the system unsafe if it is deep inside the safe set, and feedback control can be used to respond during this time.

My work in this chapter focuses on balancing robustness to disturbances with performance in safety-critical control through a modification to Input-to-State Safety, as well as a characterization of safety guarantees in the presence of stochastic disturbance signals. In Section 4.1 I discuss related work in the setting of ISSf, disturbance robust control, and stochastic safety-critical control. In Section 4.2 I will discuss some of the limitations of ISSf and ISSf-CBFs as presented in Section 2.7 and 2.9. In Section 4.3 I will resolve some of these limitations through the proposal of a novel modification of ISSf known as *Tunable Input-to-State Safety (TISSf)* that allows the prioritization of robustness to disturbances near the boundary of a set that is to be kept forward invariant, rather than over the entire set. I will then highlight the ability for TISSf to balance performance with robustness through an inverted pendulum example. In Section 4.4 I demonstrate the efficacy of TISSf with an application on an autonomous semi-trailer truck. I first present a model, safety specification captured via a CBF, and nominal controller design for this system. I show that the simplified model for this system neglects important dynamics in the actuation of an experimental system, leading to violations of safety when using the CBF-QP controller presented in Section 2.5. Using the tool of TISSf, I design a robust yet performant controller that successfully keeps the experimental autonomous semi-trailer truck system safe. Lastly, in Section 4.5, I will consider CBFs for discrete-time systems subjected to stochastic disturbances. I develop a notion of probabilistic safety related to the idea of ISSf, and show how CBFs satisfying convexity properties can be used to meet the necessary requirements to produce probabilistic safety guarantees. Key contributions of this work are described at the beginning of each respective section.

4.1 Related Work

Disturbance Robust Safety-Critical Control

The inherent robustness properties of CBFs were first observed in [221]. In particular, this work observed that if the CBF property was satisfied on the set E , including outside of the set \mathcal{C} , then asymptotic stability of the set \mathcal{C} could be established under a CBF-based controller. This asymptotic stability property was then used to establish an ISS result with respect to the set \mathcal{C} for a system experiencing unmatched disturbances [221, Proposition 5]. This idea was codified through the concept of Input-to-State Safety (ISSf) in [58] for systems with matched disturbances. This work introduced the notion of ISS-BFs and ISSf-CBFs as tool for certifying and constructively achieving ISSf. Importantly, this work established a paradigm for safety-critical control design that focuses on controlling the enlargement of the set \mathcal{C} that is kept forward invariant, rather than trying to explicitly enforce the forward invariance of the set \mathcal{C} itself. It is this paradigm that will be utilized in my work. My contributions in this chapter focus on improving the balance of performance and robustness to disturbances that is achieved using ISSf-CBFs. In particular, my work in [228], presented in Sections 4.2 and 4.3, modifies the definitions of ISSf and ISSf-CBFs from those in [58] to permit controllers that achieve rigorous safety guarantees without greatly degrading performance. This is accomplished by relaxing the robustness terms in an ISSf-CBF deep in the interior of the set \mathcal{C} , and only requiring robustness to be enforced near the boundary of \mathcal{C} . My work in [21] demonstrates the effectiveness of this approach on the challenging real-world problem of keeping an automated semi-trailer truck at a safe following distance behind a lead vehicle, including an experimental demonstration.

In parallel to the development of ISSf, other notions of robust CBFs that ensure the forward invariance of the set \mathcal{C} in the presence of disturbances were formulated in [50]–[55], [229]–[231]. These approaches require knowledge of a bound on the disturbance that is utilized directly in the controller. Consequently, if the bound is weak, the system is more robust than necessary and displays conservative behavior, as was seen in the results of [50]. The work in [52] models dynamics of disturbances through the use of differential inclusions in which the disturbance is seen as taking values in some set (rather than just satisfying a bound on its magnitude). The work in [55] considers sector-bounded uncertainties affecting the input of the system (such as saturation) as a matched disturbance signal that takes a worst-case value based on

the input to the system, while the work in [232] considers unmodeled actuator dynamics leading to a matched disturbance. An interesting direction that could be paired with much of this work is the synthesis of CBFs that can be rendered robust to disturbances through control by utilizing a modified form of Hamilton-Jacobi reachability [229]. Other work that could be paired with many of these approaches considers using a disturbance observer to reduce the worst-case disturbance the system needs to be robust to [230], [231].

Stochastic Safety-Critical Control

Instead of considering disturbances simply as piecewise continuous signals, it is often interesting to ascribe more structure through the form of a statement on the probability distribution of the disturbance value, leading to a stochastic disturbance signal. This perspective allows one to weight the frequency that a certain disturbance value occurs more highly than others. This contrasts much of the work studying disturbances through ISSf and other robustness approaches mentioned above, which typically assumes that the disturbance can (and will) always take the worst-case value when establishing safety guarantees. Because it is still possible for a stochastic disturbance signal to take these worst-case values (perhaps with low-probability), it is important that the resulting safety guarantees be designed to account for not only this situation, but make use of the fact that a stochastic disturbance is frequently not taking worst-case values.

Results using CBFs for safety-critical control in the setting of stochastic disturbances can generally be divided along two axes, the first being the use of continuous versus discrete-time dynamics, and the second being chance-constraints versus finite-time safety guarantees. Results in continuous-time consider stochastic differential equations, with the standard closed-loop system (2.2) perturbed by a Brownian motion [233]. This assumption typically implies that the stochastic disturbances experienced by the system at two different points in time are correlated based on how far those points in time are spaced apart, imparting significant structure that can be utilized in proving safety guarantees. Early work on this challenge developed results using Lyapunov functions [59], but work building off of this in the context of barrier functions focused on autonomous continuous-time systems (without control inputs) [60], and subsequently, synthesizing stochastic barrier certificates through Sums-of-Squares (SOS) programming [61]. The advent of CBFs in

[62] led to the rapid development of CBFs designed to enable control synthesis for stochastic differential equations [43], [63]–[71]. Results in discrete-time also stem from the early work [59] which additionally studied safety guarantees using Lyapunov functions for discrete-time systems with stochastic disturbances. Typically, stochastic disturbances in this setting are not assumed to be correlated across discrete-time instances in contrast to continuous-time. The work in [61] also considered verifying safety for discrete-time systems with stochastic disturbances using SOS with the tools in [59]. Developed shortly after continuous-time CBFs, discrete-time CBFs [234] provided an analog for studying safety-critical control synthesis for discrete-time systems with stochastic disturbances [69], [72]. My work in [235], discussed in Section 4.5, also takes a discrete-time CBF-based approach for studying safety-critical control synthesis, and utilizes the framework in [59] to establish rigorous safety guarantees. One direction that has remained fairly unexplored is the mixture of continuous and discrete-time stochastic systems. In particular, control systems are often based on discrete-time measurements with a sample-and-hold control implementation, as I will discuss in Chapter 6. The discrete-time measurements are often susceptible to noise, suggesting a discrete-time approach for modeling stochastic disturbances. Underlying the sample-and-hold inputs to the system is the continuous-time evolution of the system dynamics (with the input held constant). The work I present in Section 4.5 is a first step in this direction, with future work beyond this thesis seeking to unify it with the developments in Chapter 6.

The other major axis in stochastic safety-critical control pertains to chance-constraints versus finite-time safety guarantees. Chance-constraints typically require that a CBF-based inequality constraint be satisfied at any given time instance with high probability as a means for prescribing safety [43], [70], [72]. One challenge with this approach is that one instantaneous failure of the CBF constraint in closed-loop does not necessarily imply the system is unsafe, but rather that the change in the value of the CBF is larger than desired. Instead, the failure of safety typically arises from the compounding of several successive failures to satisfy the CBF inequality.

Understanding this compounding behavior has been the focus of stochastic safety-critical control seeking finite-time safety guarantees. Much of this work focuses on bounding the probability that the state of the system leaves the set

\mathcal{C} over some finite time horizon. I note that this differs from the typical notion of safety, which considers whether or not the state remains in the set \mathcal{C} over the entire interval on which a solution exists (potentially for all time). Considering a finite time horizon allows for considering potentially unbounded disturbances (such as a discrete-time disturbance signal with a Gaussian distribution), for which considering an infinite horizon leads to a system being unsafe with probability 1 [68, Footnote 7]. This pursuit of finite-time safety guarantees using Lyapunov functions can be seen in the early work [59], which established safety guarantees using Martingale theory and has served as a source of motivation for many subsequent works. Examples of such work include [60], [61], which were the first to use barrier functions to certify finite-time safety properties in the presence of stochastic disturbances. Subsequently, the introduction of CBFs in [62] led to the rapid development of stochastic variants for achieving finite-time safety guarantees by synthesizing CBFs and polynomial controllers for continuous and discrete-time systems through SOS [63], [69], considering online optimization-based control for continuous-time systems [67], considering continuous-time hybrid systems [65], and proposing definitions for continuous-time stochastic CBFs that achieve less conservative guarantees [71].

My work in [235], discussed in Section 4.5, also develops off of the work in [59] in conjunction with the work in [69] to develop online optimization-based CBF-based controllers for discrete-time systems that provide finite-time safety guarantees. The focus of this work is on the properties that need to be satisfied by a CBF and the dynamic system to permit robust control synthesis through convex optimization. Additionally, this work takes an approach akin to ISSf by providing probabilistic safety guarantees based on enlargements of the set \mathcal{C} , rather than only \mathcal{C} . An alternative perspective that forgoes Martingale-based analysis but focuses on finite-time safety guarantees is proposed in [68]. This approach focuses on the continuous-time evolution of the entire distribution of the value of a CBF, which can be captured by the solution to a convection-diffusion partial differential equation. Alternatively, the tools developed in this work were considered again in [64] to propose a method for achieving long-horizon safety guarantees using the forward invariance of level sets of the safety probability. Lastly, the work in [66] considers reciprocal CBFs (as defined in [11]) and continuous-time dynamics, allowing for strong claims on infinite horizon safety with probability 1.

4.2 Limitations of Input-to-State Safety

Note that the contents of this section were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

In this section we will explore some of the limitations of ISSf and ISSf-CBFs as proposed in [58] and presented in Section 2.9.

The text for this section is adapted from:

A. Alan, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, "Safe controller synthesis with tunable input-to-state safe control barrier functions", *IEEE Control Sys. Let.*, vol. 6, pp. 908-913, 2022.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article.

This will be done through a simple illustrative example.

Consider a open-loop system:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} -x_2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u, \quad (4.1)$$

with state $\mathbf{x} = [x_1 \ x_2]^\top \in \mathbb{R}^2$ and input $u \in \mathbb{R}$. A controller $k : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is locally Lipschitz continuous on \mathbb{R}^2 can be specified as follows:

$$k(\mathbf{x}) = x_1 - 2x_2 - 1, \quad (4.2)$$

for which the resulting closed-loop system (2.2) has a locally exponentially stable equilibrium point $\mathbf{x}_e = [1 \ 0]^\top$. Next consider a function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is continuously differentiable on \mathbb{R}^2 that is defined as:

$$h(\mathbf{x}) = x_1 - x_2, \quad (4.3)$$

yielding the set \mathcal{C} defined as:

$$\mathcal{C} = \{ \mathbf{x} \in \mathbb{R}^2 \mid x_1 - x_2 \geq 0 \}. \quad (4.4)$$

The time derivative of h is given by:

$$\dot{h}(\mathbf{x}) = \underbrace{-x_1 + x_2}_{-h(\mathbf{x})} + 1 > -h(\mathbf{x}). \quad (4.5)$$

Thus for $\alpha \in \mathcal{K}^e$ defined as $\alpha(r) = r$, we have that:

$$\dot{h}(\mathbf{x}) \geq -\alpha(h(\mathbf{x})), \quad (4.6)$$

and thus h is a BF for the closed-loop system (2.2) on \mathcal{C} . Simulation results for the closed-loop system can be seen in Figure 4.1(a), where all of the trajectories initiated from different initial conditions $\mathbf{x}_0 \in \mathcal{C}$ approach the locally exponentially stable equilibrium point \mathbf{x}_e while remaining inside \mathcal{C} .

Now introduce a matched disturbance to the system (4.1), yielding:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} -x_2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u + \begin{bmatrix} 0 \\ 1 \end{bmatrix} d(t), \quad (4.7)$$

where $d : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is piecewise continuous on $\mathbb{R}_{\geq 0}$. Using a harmonic disturbance signal $d(t) = \delta \sin t$ with $\delta = 3$, we see the evolution of the closed-loop system with matched disturbances (2.74) from the previous initial conditions $\mathbf{x}_0 \in \mathcal{C}$ in Figure 4.1(b). Observe that the disturbance makes the state trajectories leave the set \mathcal{C} periodically. Note that $L_{\mathbf{g}}h(\mathbf{x}) = -1$ for all $\mathbf{x} \in \mathbb{R}^2$. Inspired by the definition of ISSf-CBFs in Section 2.9, consider a controller $k' : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is locally Lipschitz continuous on \mathbb{R}^2 that is defined as:

$$k'(\mathbf{x}) = k(\mathbf{x}) + \frac{L_{\mathbf{g}}h(\mathbf{x})}{\epsilon_0}, \quad (4.8)$$

for $\epsilon_0 \in \mathbb{R}_{>0}$, noting that for the corresponding closed-loop system with matched disturbances (2.74), we have that:

$$\dot{h}(\mathbf{x}, t) \geq -\alpha(h(\mathbf{x})) + \frac{|L_{\mathbf{g}}h(\mathbf{x})|^2}{\epsilon_0} - L_{\mathbf{g}}h(\mathbf{x})d(t), \quad (4.9)$$

for all $\mathbf{x} \in \mathbb{R}^2$ and $t \in \mathbb{R}_{\geq 0}$. Completing the square, we have that:

$$\dot{h}(\mathbf{x}, t) \geq -\alpha(h(\mathbf{x})) - \frac{\epsilon_0}{4}|d(t)|^2, \quad (4.10)$$

for all $\mathbf{x} \in \mathbb{R}^2$ and $t \in \mathbb{R}_{\geq 0}$. Consequently, the closed-loop system with matched disturbances (2.74) is ISSf with respect to the set \mathcal{C} with $\gamma \in \mathcal{K}$ defined as $\gamma(r) = -\alpha^{-1}(-\frac{\epsilon_0}{4}r^2)$. Observe that we can make the enlargement of the set \mathcal{C} that is kept forward invariant smaller with smaller values of ϵ_0 at the expense of the controller k' specifying larger inputs. Using the harmonic disturbance with $\delta = 3$, the set \mathcal{C}_δ (the set \mathcal{C}_d defined in (2.79) with $d = \delta$), we can see the evolution of the closed-loop system with matched disturbances (2.74) using the

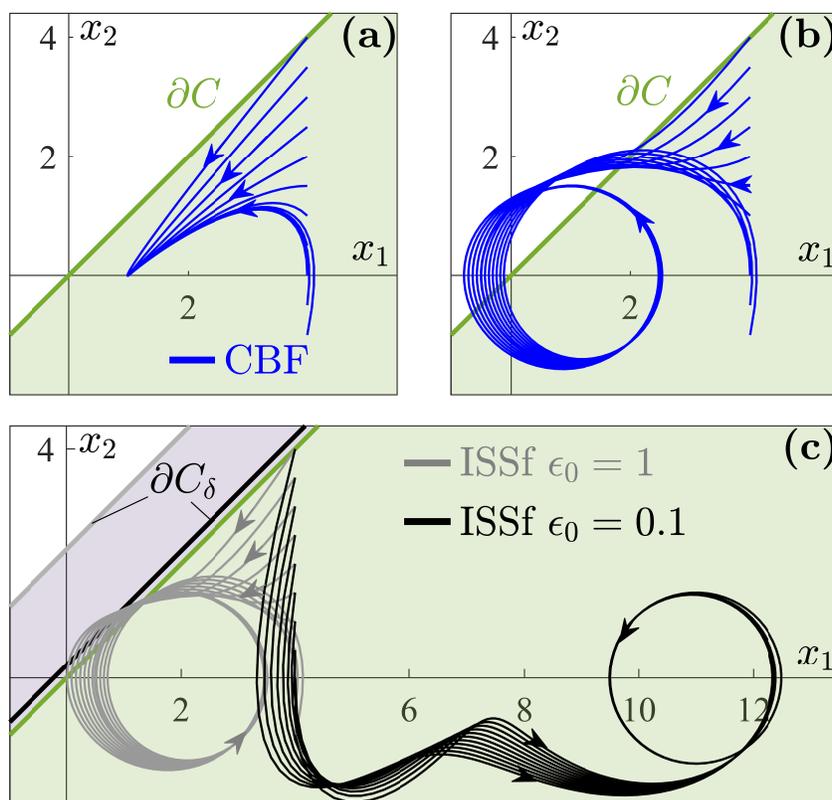


Figure 4.1. Simple example demonstrating limits of ISSf in achieving performance and meaningful theoretical safety guarantees. (a) Safe controller with no disturbances. (b) Impact of harmonic disturbance signal without modifications to controller design. (c) ISSf based controller designs for different values of the parameter ϵ_0 . Large values of ϵ_0 lead to weak theoretical guarantees, while small values lead to conservative behavior.

controller k' with different values of ϵ_0 in Figure 4.1(c). Observe that for both values of ϵ_0 , the system remains in the corresponding set \mathcal{C}_δ as required by ISSf. For the value of $\epsilon_0 = 1$, the boundary of the set \mathcal{C}_δ is relatively far away from the boundary of the set \mathcal{C} compared to how close the solutions remain to the set \mathcal{C} . In this case, a small violation of the original safety requirement is occurring, but the ISSf safety guarantee is fairly weak and does not accurately describe the closed-loop behavior. The weakness of this guarantee can make it useless as a guiding principle when designing robust controllers. In contrast, for $\epsilon_0 = 0.1$, the boundary of \mathcal{C}_δ is very close to the boundary of \mathcal{C} , providing a rigorous theoretical statement on small violations of safety. Unfortunately, with this strong guarantee comes undesirable closed-loop behavior, as the system evolves deep into the safe set, potentially degrading performance if the system is to remain near the equilibrium point \mathbf{x}_e . In fact, this introduces an even wider gap between the theoretical guarantee on the closed-loop behavior and the actual

realized closed-loop behavior, suggesting that once again, this is a vacuous guiding principle in producing a robust control design. The developments in the next sections in this chapter will seek to balance this tradeoff between strong theoretical guarantees and performant closed-loop behavior by replacing the parameter ϵ_0 with a function that prioritizes safety only at the boundary and outside of the set \mathcal{C} .

4.3 Tunable Input-to-State Safety

In this section I will present work on balancing performance with robustness to disturbances through a modification of ISSf in the form of *Tunable Input-to-State Safety (TISSf)*. I will first present the idea of TISSf, and establish what theoretical safety guarantees are achieved using it. Next I will consider its application on the simple example presented in Section 4.2. Following this, I will present a detailed inverted pendulum example that shows how the parameters of a TISSf-based controller can be used to balance performance and theoretical safety guarantees.

The contributions of this section are as follows:

- The definition of Tunable Input-to-State Safety (TISSf), which modifies the original definition of ISSf in [58] to permit performant behavior with meaningful theoretical safety guarantees.
- A detailed application of TISSf-based control design to an inverted pendulum with a focus on how the parameters of a TISSf controller determine performance and theoretical safety guarantees.

The text for this section is adapted from:

A. Alan, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, "Safe controller synthesis with tunable input-to-state safe control barrier functions", *IEEE Control Sys. Let.*, vol. 6, pp. 908-913, 2022.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article.

A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, "Control barrier functions and input-to-state safety with application to automated

vehicles", *conditionally accepted as a Full Paper in IEEE Trans. on Control Sys. Tech.*, 2022.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Tunable Input-to-State Safety

I will begin by defining the notion of Tunable Input-to-State Safety, which will lead to corresponding variants of ISSf-BFs and ISSf-CBFs.

Definition 32 (*Tunable Input-to-State Safety (TISSf)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The closed-loop system with unmatched disturbances (2.70) or closed-loop system with matched disturbances (2.74) is said to be *Tunable Input-to-State Safe (TISSf)* with respect to the set \mathcal{C} if there exists a function $\gamma_T : \mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with $\gamma_T(a, \cdot) \in \mathcal{K}_{\infty}^e$ for all $a \in \mathbb{R}_{\geq 0}$ and $\gamma_T(\cdot, b)$ continuously differentiable on \mathbb{R} for all $b \in \mathbb{R}_{\geq 0}$ such that defining the function $h_T : E \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ as:

$$h_T(\mathbf{x}, d) = h(\mathbf{x}) + \gamma_T(h(\mathbf{x}), d), \quad (4.11)$$

for all $d \in \mathbb{R}_{\geq 0}$ the set $\mathcal{C}_{d,T} \subset E$ defined as:

$$\mathcal{C}_{d,T} = \{\mathbf{x} \in E \mid h_T(\mathbf{x}, d) \geq 0\}, \quad (4.12)$$

is forward invariant up to d for the closed-loop system (2.70) or (2.74). If the system (2.70) or (2.74) is TISSf with respect to the set \mathcal{C} , the set \mathcal{C} is referred to as a *Tunable Input-to-State Safe set (TISSf set)*.

TISSf simply modifies the original definition of ISSf presented in Section 2.7 by allowing the function γ_T to not only depend on the disturbance signal, but to also depend on the value of the function h itself. While the utility of this modification is not immediately clear, it will be more easily understood when considering control synthesis. Another small modification is dropping the maximum disturbance magnitude \bar{d} , and just requiring the property to hold for disturbances of arbitrarily large magnitudes. This is a simplifying assumption, and is not necessary if one keeps track of the appropriate domains for various functions throughout subsequent proofs. The objective of control synthesis will be producing a type of ISSf-BF, as captured in the following:

Definition 33 (*Tunable Input-to-State Safe Barrier Function (TISSf-BF)*).

Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E with $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ when $h(\mathbf{x}) = 0$. The function h is said to be a *Tunable Input-to-State Safe barrier function (TISSf-BF)* for the closed-loop system with unmatched disturbances (2.70) on \mathcal{C} if there exist functions $\alpha \in \mathcal{K}_\infty^e$ with $\alpha^{-1} \in \mathcal{K}_\infty^e$ continuously differentiable on \mathbb{R} and $\epsilon : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ that is continuously differentiable on \mathbb{R} such that:

$$L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x}) - \frac{1}{\epsilon(h(\mathbf{x}))} \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 \geq -\alpha(h(\mathbf{x})), \quad (4.13)$$

for all $\mathbf{x} \in E$, and:

$$\frac{\partial \epsilon}{\partial r}(r) \geq 0, \quad (4.14)$$

for all $r \in \mathbb{R}$. The function h is said to be a TISSf-BF for the closed-loop system with matched disturbances (2.74) on \mathcal{C} if (4.13) is replaced with:

$$L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x}) - \frac{1}{\epsilon(h(\mathbf{x}))} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 \geq -\alpha(h(\mathbf{x})), \quad (4.15)$$

for all $\mathbf{x} \in E$.

I note this definition looks closer to the definition of ISSf-CBFs than ISSf-BFs in Section 2.9. This is done as it will allow explicitly highlighting the importance of the function ϵ , which has replaced the constant parameter ϵ_0 that appears in the definition of ISSf-CBFs and was shown in Section 4.2 to heavily dictate performance and theoretical safety guarantees in the presence of disturbances. Note that the function ϵ must always take positive values, thereby requiring that the controller \mathbf{k} always meet a stricter requirement than the standard CBF-QP condition (2.60). It also must increase (or not decrease) with its argument. This suggests that when $h(\mathbf{x})$ is larger, or the system state is deep inside the set \mathcal{C} , the modification to the standard CBF inequality (2.60) is reduced, as $\epsilon(h(\mathbf{x}))$ is larger. As $h(\mathbf{x})$ gets smaller, such as when the state of the system approaches the boundary of \mathcal{C} or moves further away from \mathcal{C} , the value of ϵ must get smaller, and the importance of the robustifying term increases. This reflects the intuitive idea that robustness is more important near the boundary of \mathcal{C} , and less important when the system is sufficiently far inside \mathcal{C} . Examples of such functions can be seen in Figure 4.2

TISSf-BFs serve as a certificate of TISSf of a closed-loop system with respect to a set \mathcal{C} as follows:

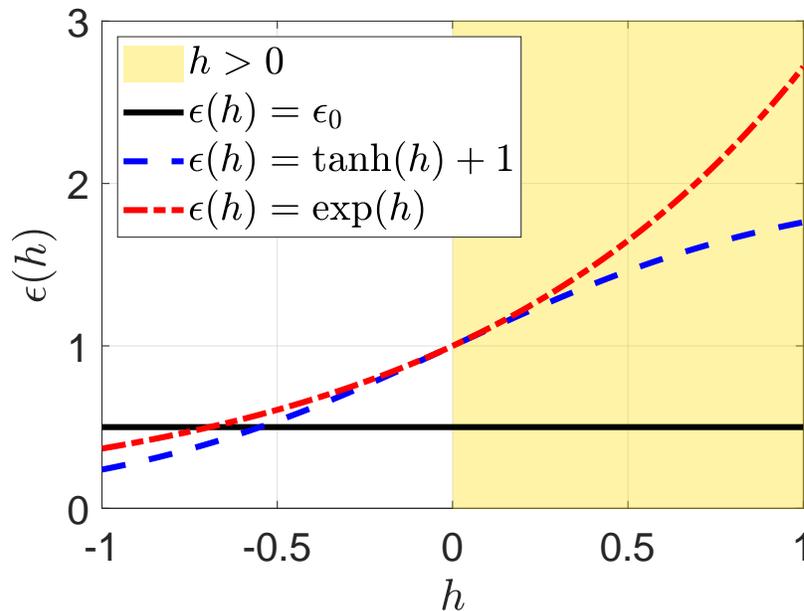


Figure 4.2. Examples of the function ϵ that satisfy the condition in (4.14).

Theorem 29. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is a TISSf-BF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) on \mathcal{C} , then the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) is TISSf with respect to the set \mathcal{C} with a function $\gamma_T : \mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ defined as:*

$$\gamma_T(r, d) \triangleq -\alpha^{-1} \left(-\frac{\epsilon(r)d^2}{4} \right). \quad (4.16)$$

I will prove this theorem for the matched disturbance setting here, but a similar proof for the unmatched disturbance case is provided in Appendix A.

Proof. Observe that by assumption α^{-1} and ϵ are continuously differentiable on \mathbb{R} , such that the function $\gamma_T(\cdot, b)$ is continuously differentiable for any $b \in \mathbb{R}_{\geq 0}$. Similarly, as $\epsilon(r) > 0$ for all $r \in \mathbb{R}$ and $\alpha^{-1} \in \mathcal{K}_\infty^e$, the function $\gamma_T(a, \cdot) \in \mathcal{K}_\infty^e$ for any $a \in \mathbb{R}$ as required. Our next goal is to verify that for all $d \in \mathbb{R}_{\geq 0}$, the set $\mathcal{C}_{d,T}$ defined in (4.12) is forward invariant up to d . Consider a value of $d \in \mathbb{R}_{\geq 0}$ and consider a disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ that is piecewise continuous on $\mathbb{R}_{\geq 0}$ and satisfies $\|\mathbf{d}\|_\infty \leq d$. As h is a TISSf-BF for

the closed-loop system with matched disturbances (2.74), we have that:

$$\dot{h}(\mathbf{x}, t) \triangleq L_f h(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{k}(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{d}(t), \quad (4.17)$$

$$\geq -\alpha(h(\mathbf{x})) + \frac{1}{\epsilon(h(\mathbf{x}))} \|L_g h(\mathbf{x})\|^2 + L_g h(\mathbf{x}) \mathbf{d}(t). \quad (4.18)$$

Noting that:

$$L_g h(\mathbf{x}) \mathbf{d}(t) \geq -\|L_g h(\mathbf{x})\| \|\mathbf{d}\|_\infty \geq -\|L_g h(\mathbf{x})\| d, \quad (4.19)$$

for all $\mathbf{x} \in E$ and $t \in \mathbb{R}_{\geq 0}$ and $\epsilon(h(\mathbf{x})) > 0$ for all $\mathbf{x} \in E$, adding and subtracting $\frac{\epsilon(h(\mathbf{x}))d^2}{4}$, and completing the squares yields:

$$\dot{h}(\mathbf{x}, t) \geq -\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4}. \quad (4.20)$$

Next, observe that:

$$\frac{\partial \gamma_T}{\partial r}(r, d) = \frac{d^2}{4} \frac{\partial \alpha^{-1}}{\partial r} \left(-\frac{\epsilon(r)d^2}{4} \right) \frac{\partial \epsilon}{\partial r}(r) \geq 0, \quad (4.21)$$

for all $r \in \mathbb{R}$. Taking the time derivative of the function h_T defined in (4.11) yields:

$$\dot{h}_T(\mathbf{x}, d, t) = \left(1 + \frac{\partial \gamma_T}{\partial r}(h(\mathbf{x}), d) \right) \dot{h}(\mathbf{x}, t). \quad (4.22)$$

noting that d is a fixed constant in this proof. By (4.21), we have that:

$$1 + \frac{\partial \gamma_T}{\partial r}(h(\mathbf{x}), d) > 0, \quad (4.23)$$

for all $\mathbf{x} \in E$. Substituting (4.20) into (4.22), we obtain:

$$\dot{h}_T(\mathbf{x}, d, t) \geq \left(1 + \frac{\partial \gamma_T}{\partial r}(h(\mathbf{x}), d) \right) \left(-\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4} \right). \quad (4.24)$$

Next, consider a state $\mathbf{x} \in \partial \mathcal{C}_{d,T}$, such that $h_T(\mathbf{x}, d) = 0$, for which (4.11) and (4.16) imply:

$$-\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4} = 0, \quad (4.25)$$

yielding:

$$\dot{h}_T(\mathbf{x}, d, t) \geq 0, \quad (4.26)$$

for all $t \in \mathbb{R}_{\geq 0}$. In the case that $d = 0$, such that $h_T(\mathbf{x}, 0) = h(\mathbf{x})$, we have that $h_T(\mathbf{x}, 0) = 0$ implies $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$. Considering $d > 0$, because $\gamma(a, \cdot) \in \mathcal{K}_\infty^e$ for

all $a \in \mathbb{R}$, when $h_T(\mathbf{x}, d) = 0$, we have $-\alpha(h(\mathbf{x})) > 0$. Thus, the inequality in (4.15) requires that $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ for $\mathbf{x} \in \partial \mathcal{C}_{d,T}$. Finally, we have:

$$\frac{\partial h_T}{\partial \mathbf{x}}(\mathbf{x}, d) = \underbrace{\left(1 + \frac{\partial \gamma_T}{\partial r}(h(\mathbf{x}), d)\right)}_{>0} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n, \quad (4.27)$$

using (4.23). Therefore, Nagumo's theorem [104], [123] implies the set $\mathcal{C}_{d,T}$ is forward invariant as $h_T(\mathbf{x}, d) = 0$ implies $\dot{h}_T(\mathbf{x}, d, t) \geq 0$ for all $t \in \mathbb{R}_{\geq 0}$, and $\frac{\partial h_T}{\partial \mathbf{x}}(\mathbf{x}, d) \neq \mathbf{0}_n$. \square

Let us now consider a TISSf-based controller for the simple example we considered in Section 4.2. For the open-loop system with matched disturbances in (4.7) with the particular harmonic disturbance signal $d(t) = \delta \sin t$ with $\delta = 3$, I pick the following differentiable function:

$$\epsilon(h(\mathbf{x})) \triangleq \epsilon_0 e^{\lambda h(\mathbf{x})}, \quad (4.28)$$

with constants $\epsilon_0, \lambda \in \mathbb{R}_{>0}$. Considering a controller $k' : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is locally Lipschitz continuous on \mathbb{R}^2 that is defined as:

$$k'(x) = k(\mathbf{x}) + \frac{L_{\mathbf{g}} h(\mathbf{x})}{\epsilon(h(\mathbf{x}))} = x_1 - 2x_2 - 1 - \frac{1}{\epsilon_0 e^{\lambda(x_1 - x_2)}}, \quad (4.29)$$

the function h as defined in (4.3) is a TISSf-BF with $\alpha(r) = r$ for corresponding closed-loop system on \mathcal{C} . Thus, the set:

$$\mathcal{C}_{\delta,T} = \left\{ \mathbf{x} \in \mathbb{R}^2 \mid x_1 - x_2 + \frac{\epsilon_0 e^{\lambda(x_1 - x_2)} \delta^2}{4} \geq 0 \right\}, \quad (4.30)$$

is forward invariant up to δ (which the system experiences). It is noted that $\lambda = 0$ recovers the ISSf-based controller with $\epsilon(h(\mathbf{x})) = \epsilon_0$ for all $\mathbf{x} \in \mathbb{R}^2$, whereas a larger λ pulls $\partial \mathcal{C}_{\delta,T}$ closer to $\partial \mathcal{C}$, and decreases the effect of the corresponding term in the controller (4.29) for $h(\mathbf{x}) > 0$. The set $\mathcal{C}_{\delta,T}$ is depicted in Figure 4.3(d) with $\epsilon_0 = e^{-2}$, $\lambda = 2$ chosen such that $\partial \mathcal{C}_{\delta,T}$ corresponds to $\partial \mathcal{C}_{\delta}$ using $\epsilon_0 = 0.1$. In this sense, it is possible to achieve the same theoretical safety guarantee using both TISSf and ISSf with small values of ϵ_0 . Additionally, solution trajectories from the same set of initial conditions are also shown. Observe that all solution trajectories stay within the set $\mathcal{C}_{\delta,T}$ as required by TISSf. Moreover, the system is allowed to evolve close to the boundary of the

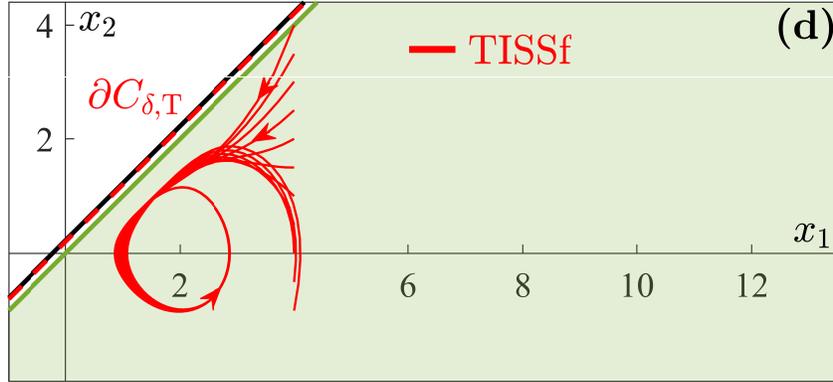


Figure 4.3. Simple example using TISSf to achieve both meaningful theoretical safety guarantees and non-conservative behavior.

set $\mathcal{C}_{\delta,T}$ and is not forced deep into the set \mathcal{C} . Thus, both meaningful theoretical guarantees and non-conservative behavior can be achieved using TISSf.

Given a value of $d \in \mathbb{R}_{\geq 0}$, the boundary of the set $\mathcal{C}_{d,T}$ that is rendered forward invariant is defined as the 0-level set of the TISSf-BF h . The value of h on this level set, denoted as $h^* \leq 0$, can be found by solving the equation:

$$h^* - \underbrace{\alpha^{-1} \left(-\frac{\epsilon(h^*)d^2}{4} \right)}_{\gamma_T(h^*,d)} = 0. \quad (4.31)$$

By definition, $\gamma_T(h^*,d)$ must be strictly positive for $d > 0$, implying that $h^* < 0$ in the presence of disturbances. Moreover, as d increases, h^* must get more negative, implying that the boundary of $\mathcal{C}_{d,T}$ falls farther from the boundary of \mathcal{C} . Control over this degradation in safety can be achieved by modifying the function ϵ to yield different values of h^* as specified in (4.31). From a practical perspective, $h(\varphi_d(t)) \geq h^*$ can be seen as a design objective with h^* chosen to satisfy a minimum tolerable safety requirement in the presence of disturbances. With this mindset, a control designer must appropriately choose ϵ to ensure h^* meets these minimum tolerable safety requirements.

It is straightforward to extend the idea of an TISSf-BF to a TISSf-CBF:

Definition 34 (*Tunable Input-to-State Safe Control Barrier Function (TISSf-CBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E with $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ when $h(\mathbf{x}) = 0$. The function h is said to be a *Tunable Input-to-State Safe Control Barrier Function*

(TISSf-CBF) for the open-loop system with unmatched disturbances (2.69) on \mathcal{C} if there exist functions $\alpha \in \mathcal{K}_\infty^e$ with $\alpha^{-1} \in \mathcal{K}_\infty^e$ continuously differentiable on \mathbb{R} and $\epsilon : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ that is continuously differentiable on \mathbb{R} such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon(h(\mathbf{x}))} \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 > -\alpha(h(\mathbf{x})), \quad (4.32)$$

for all $\mathbf{x} \in E$, and:

$$\frac{\partial \epsilon}{\partial r}(r) \geq 0, \quad (4.33)$$

for all $r \in \mathbb{R}$. The function h is said to be a TISSf-CBF for the open-loop system with matched disturbances (2.73) on \mathcal{C} if (4.32) is replaced with:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon(h(\mathbf{x}))} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 > -\alpha(h(\mathbf{x})), \quad (4.34)$$

for all $\mathbf{x} \in E$.

As with the notion of ISSf-CBFs presented in Section 2.9, in the setting of matched disturbances, if h is a CBF for the open-loop system (2.1) with a function $\alpha \in \mathcal{K}_\infty^e$ that has an inverse $\alpha^{-1} \in \mathcal{K}_\infty^e$ that is continuously differentiable on \mathbb{R} , then for any $\epsilon : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ that is continuously differentiable on \mathbb{R} and satisfies (4.33), the function h is a TISSf-CBF for the open-loop system with matched disturbances on \mathcal{C} . This makes it relatively easy to robustify standard CBF-based controllers to matched disturbances if necessary. As typical, given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E , an optimization-based controller can be specified as follows:

$$\begin{aligned} \mathbf{k}_{\text{TISSf}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 && \text{(TISSf-CBF-QP)} \\ \text{s.t. } & L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} - \frac{1}{\epsilon(h(\mathbf{x}))} \|L_{\mathbf{g}}h(\mathbf{x})\|^2 \geq -\alpha(h(\mathbf{x})), \end{aligned}$$

with a similar specification possible for the unmatched disturbance setting. The optimization problem defining this controller remains a quadratic program, as there is no modification to how the constraint depends on the decision variable \mathbf{u} from the standard CBF-QP controller. As with the CBF-QP controller, under assumptions of local Lipschitz continuity of the function $\frac{\partial h}{\partial \mathbf{x}} : E \rightarrow \mathbb{R}^n$ and α this controller can be shown to be locally Lipschitz continuous on E .

Inverted Pendulum Example

Consider a control system for an inverted pendulum described by the model:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} = \underbrace{\begin{bmatrix} \dot{\theta} \\ \frac{g}{l} \sin \theta \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} 0 \\ \frac{1}{ml^2} \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} u, \quad (4.35)$$

with pendulum angle $\theta \in \mathbb{R}$ and angular velocity $\dot{\theta} \in \mathbb{R}$ defining the state $\mathbf{x} = [\theta, \dot{\theta}]^\top$, and parameters given by the mass $m \in \mathbb{R}_{>0}$, length $l \in \mathbb{R}_{>0}$, and gravitational acceleration constant $g \in \mathbb{R}_{>0}$. In this example I will use the parameter values $m = 2$ [kg], $l = 1$ [m] and $g = 10$ [m/s²]. The single input $u \in \mathbb{R}$ is a torque applied on the pendulum. A set \mathcal{C} that we wish to keep safe for the inverted pendulum that restricts the angular position and velocity is given by the 0-superlevel set of a function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined as:

$$h(\theta, \dot{\theta}) = 1 - \frac{\theta^2}{a^2} - \frac{\dot{\theta}^2}{b^2} - \frac{\theta\dot{\theta}}{ab}, \quad (4.36)$$

with parameters $a, b \in \mathbb{R}_{>0}$. In this example I will use the parameter values $a = 0.25$ [rad] and $b = 0.5$ [rad/s]. The resulting set:

$$\mathcal{C} = \left\{ \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} \in \mathbb{R}^2 \mid 1 - \frac{\theta^2}{a^2} - \frac{\dot{\theta}^2}{b^2} - \frac{\theta\dot{\theta}}{ab} \geq 0 \right\}, \quad (4.37)$$

is an ellipse as depicted in Figure 4.4 by a gold region. The function h given as in (4.36) is a CBF for the inverted pendulum system (4.35) on \mathcal{C} . To see this, consider a function $\alpha \in \mathcal{K}_\infty^e$ defined as $\alpha(r) = \alpha_c r$ with $\alpha_c > 0$ satisfying $\alpha_c \leq b/a$. In this example I use the parameter value $\alpha_c = 0.2$ [1/s]. The CBF condition defined in (2.61) yields that:

$$L_{\mathbf{g}}h(\theta_0, \dot{\theta}_0) = 0 \implies \dot{\theta}_0 = -\frac{b}{2a}\theta_0. \quad (4.38)$$

This equation defines a line as depicted in Figure 4.4 by a solid black line. We have that on this line:

$$L_{\mathbf{f}}h(\theta_0, \dot{\theta}_0) + \alpha(h(\theta_0, \dot{\theta}_0)) = \alpha_c + \frac{3}{4a^2} \left(\frac{b}{a} - \alpha_c \right) \theta_0^2 > 0, \quad (4.39)$$

such that the condition (2.61) is met for our choice of α_c .

I note that if we consider a set $\tilde{\mathcal{C}}$ defined as the 0-superlevel set of a function $\tilde{h} : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by:

$$\tilde{h}(\theta, \dot{\theta}) = 1 - \frac{\theta^2}{a^2} - \frac{\dot{\theta}^2}{b^2}, \quad (4.40)$$

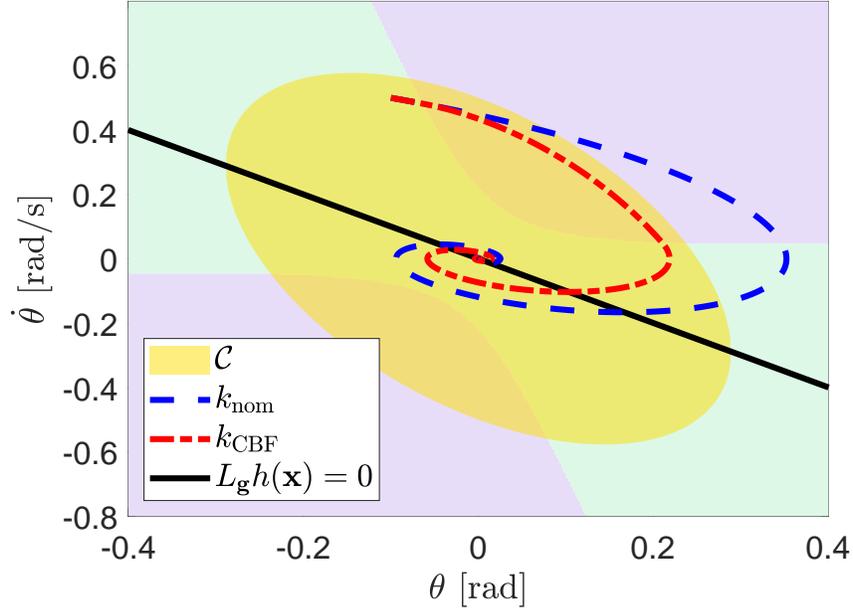


Figure 4.4. Simulation results for the inverted pendulum system. The gold ellipse is the set \mathcal{C} as defined in (4.37). The black line is the set where $L_{\mathbf{g}}h(\mathbf{x}) = 0$ as defined in (4.38). The dashed blue line is the trajectory of the system evolving under k_{nom} as defined in (4.43), which leaves the set \mathcal{C} . The green and purple regions indicate where the controller k_{nom} meets and fails to meet the CBF condition, respectively. The dashed red line is the trajectory of the system evolving under k_{CBF} which remains inside the set \mathcal{C} .

which does not include the term $\theta\dot{\theta}/ab$, then the function \tilde{h} is not a CBF for the system (4.35) on $\tilde{\mathcal{C}}$. To see this, note that:

$$L_{\mathbf{g}}\tilde{h}(\theta_0, \dot{\theta}_0) = 0 \implies \dot{\theta}_0 = 0. \quad (4.41)$$

In turn, we have that for any $\alpha \in \mathcal{K}_{\infty}^e$:

$$L_{\mathbf{f}}\tilde{h}(\theta_0, \dot{\theta}_0) + \alpha(\tilde{h}(\theta_0, \dot{\theta}_0)) = \alpha(1 - \theta_0^2/a^2). \quad (4.42)$$

The condition (2.61) is not satisfied for $|\theta_0| \geq a$, including $|\theta_0| = a$ which is in $\partial\tilde{\mathcal{C}}$. Thus, it is important to choose the safe set and design the CBF to be compatible with the system dynamics, eliminating points where the CBF condition is not met.

Consider a controller $k_{\text{nom}} : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is locally Lipschitz continuous on \mathbb{R}^2 that locally exponentially stabilizes the pendulum to an upright position, given by the feedback linearization controller:

$$k_{\text{nom}}(\theta, \dot{\theta}) = ml^2 \left(-\frac{g}{l} \sin \theta - K_p \theta - K_d \dot{\theta} \right), \quad (4.43)$$

with controller gains $K_p, K_d \in \mathbb{R}_{>0}$. This controller yields the closed-loop system:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -K_p & -K_d \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix}, \quad (4.44)$$

such that the equilibrium point $\mathbf{x}_e = [0, 0]^\top$ is locally exponentially stable. In this example I will use the parameter values $K_p = 0.6 [1/s^2]$ and $K_d = 0.6 [1/s]$. I simulate the system from the initial condition $\mathbf{x}_0 = [-0.1, 0.5]^\top \in \mathcal{C}$. This trajectory is depicted in Figure 4.4 by a dashed blue curve. Although the controller k_{nom} stabilizes the system to the upright position, in doing so it causes the state of the system to leave the set \mathcal{C} .

I next deploy the (CBF-QP) controller $k_{\text{CBF}} : \mathbb{R}^2 \rightarrow \mathbb{R}$ for the inverted pendulum system using the nominal controller k_{nom} defined in (4.43). I simulate the system from the initial condition $\mathbf{x}_0 = [-0.1, 0.5]^\top \in \mathcal{C}$. This trajectory is depicted in Figure 4.4 by a dashed red curve. Observe that the (CBF-QP) controller k_{CBF} ensures that the system state remains within the set \mathcal{C} .

To see how this controller is impacted by disturbances, I will consider a matched disturbance signal specified as:

$$d(t) = M(1 - s(t - 5) - s(t - 10) + s(t - 15)), \quad (4.45)$$

where $M \in \mathbb{R}_{\geq 0}$ and $s : \mathbb{R} \rightarrow \mathbb{R}$ is the *Heaviside function*:

$$s(\tau) = \begin{cases} 0 & \text{if } \tau < 0, \\ 1 & \text{if } \tau \geq 0. \end{cases} \quad (4.46)$$

With this disturbance we have $\|d\|_\infty = M$. In this example I use the parameter value $M = 0.75 [\text{N}\cdot\text{m}]$ and the corresponding disturbance signal is depicted in Figure 4.5. I deploy the (CBF-QP) controller with the nominal controller k_{nom} as in (4.43) to the inverted pendulum system without considering the disturbance d defined in (4.45). We see in Figure 4.6 that this controller fails to keep the system in the set \mathcal{C} , and deviates from it significantly.

I will next consider a design through TISSf. In particular, I will use the (TISSf-CBF-QP) controller k_{TISSf} with a function ϵ defined as:

$$\epsilon(r) = \epsilon_0 e^{\lambda r}, \quad (4.47)$$

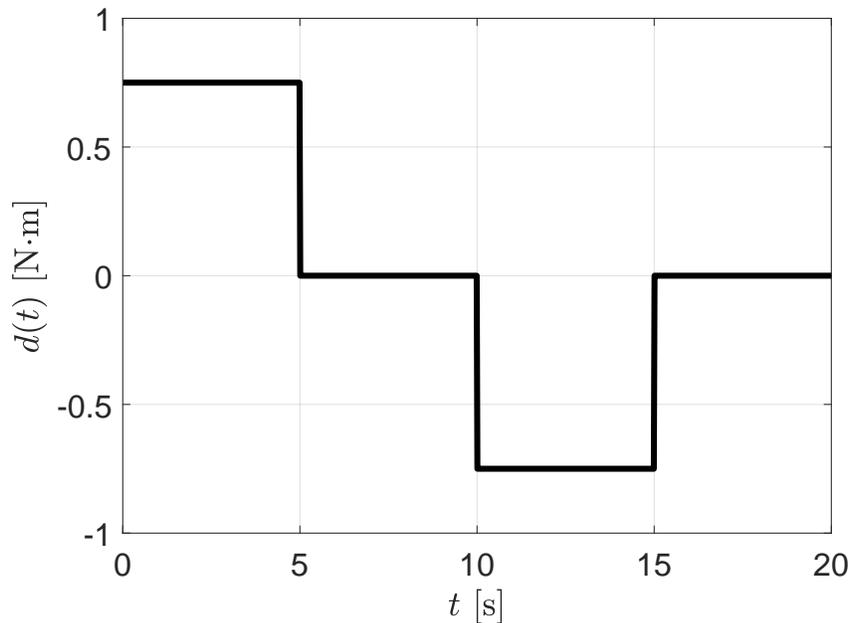


Figure 4.5. Disturbance signal for the inverted pendulum system example as defined in (4.45).

with parameters $\epsilon_0 \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}_{\geq 0}$. Using this controller will render h a TISSf-BF for the corresponding closed-loop system, and thus given the choice of α , Theorem 29 yields the sets $\mathcal{C}_{d,T}$ defined in (4.12) are forward invariant up to d with the function γ_T defined as:

$$\gamma_T(h(\theta, \dot{\theta}), d) = \frac{\epsilon(h(\theta, \dot{\theta}))d^2}{4\alpha_c}. \quad (4.48)$$

As the disturbance signal is bounded by M , the forward invariant set we will be interested in studying is determined by considering $d = M$. Thus, I use $d = 0.75$ [N·m]. Next, we have that (4.31) reduces to:

$$h^* + \frac{\epsilon_0 e^{\lambda h^*} d^2}{4\alpha_c} = 0. \quad (4.49)$$

Once ϵ_0 and λ are specified, (4.49) can be solved for h^* to find the value of h that corresponds to the boundary $\partial\mathcal{C}_{d,T}$. The left panel of Figure 4.7 shows the value of h^* for the different choices of ϵ_0 and λ specified in Table 4.1. The boundary $\partial\mathcal{C}_{d,T}$ corresponding to each set of parameters is shown in the right panel of Figure 4.7. The black and red parameter sets return (approximately) the same value of h^* , and thus the produce the same boundary $\partial\mathcal{C}_{d,T}$. In contrast, the green parameter set yields a larger set $\mathcal{C}_{d,T}$ as indicated by the smaller value of h^* in Table 4.1.

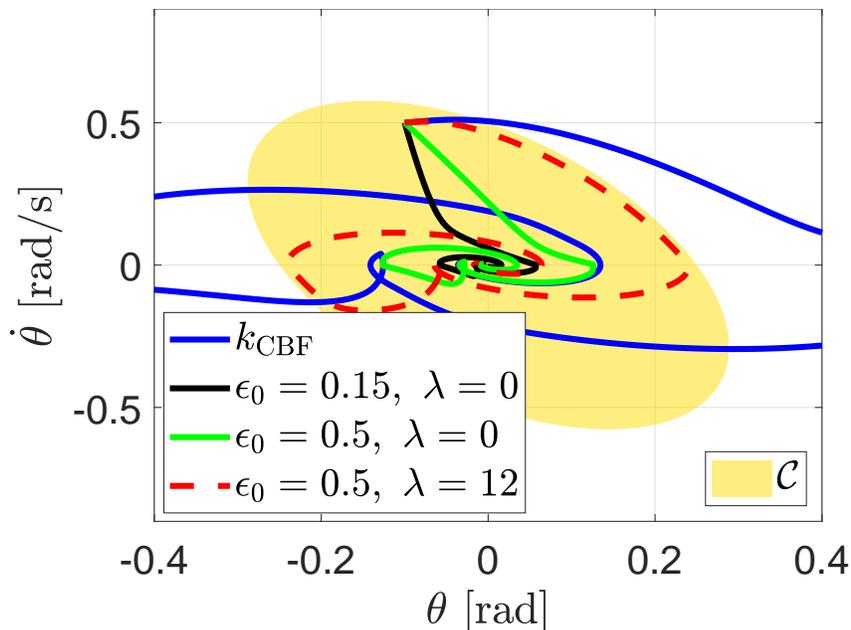


Figure 4.6. (Right) Simulation results for the inverted pendulum system with disturbances. The gold ellipse is the safe set \mathcal{C} defined in (4.37). The blue line is the trajectory of the system evolving under k_{CBF} defined in (CBF-QP), which is not robust to disturbances and leaves the safe set. The black, green, and dashed red lines are trajectories of the system evolving under k_{TISSf} defined in (TISSf-CBF-QP) with different values for ϵ_0 and λ .

| Color | Black | Red | Green |
|---|---------|---------|---------|
| ϵ_0 [$\frac{1}{\text{N}^2\text{m}^2\text{s}}$] | 0.15 | 0.5 | 0.5 |
| λ | 0 | 12 | 0 |
| h^* | -0.1055 | -0.1026 | -0.3516 |

Table 4.1. Parameters sets used in (4.47) and (4.49) for the inverted pendulum example.

I next deploy the (TISSf-CBF-QP) controller k_{TISSf} with the nominal controller k_{nom} as in (4.43) in simulation with the exponential function given in (4.47) utilizing the parameter pairs specified in Table 4.1. I first observe in Figure 4.6 that with the black parameter set, which corresponds to a standard ISSf design, the system is very conservative and the trajectory of the system rapidly converges to the equilibrium. Increasing the value of ϵ_0 returns the green parameter set, which is also a standard ISSf design. We see in the right panel of Figure 4.7 that a larger value of ϵ_0 leads to a larger set $\mathcal{C}_{d,T}$ as specified by the lesser value of h^* in Table 4.1, while in Figure 4.6 we see that a larger value of ϵ_0 allows the system to evolve less conservatively. This highlights that larger values of ϵ_0 generally lead to less conservative closed-loop behavior and an expansion of the set possessing theoretical safety guarantees.

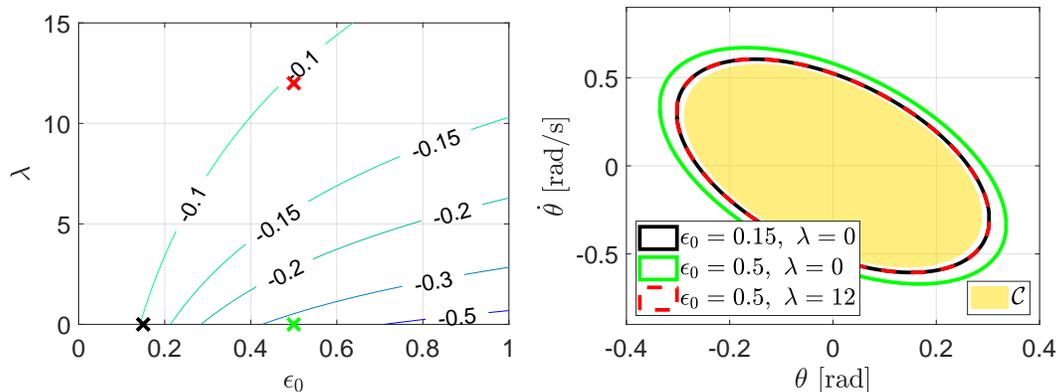


Figure 4.7. (Left) Curves corresponding to the value of h^* solving (4.49) across the (ϵ_0, λ) parameter space for the inverted pendulum example. (Right) The boundary of the set $\mathcal{C}_{d,T}$ rendered forward invariant for different choices of the parameters ϵ_0 and λ for the inverted pendulum example. Note that the $\mathcal{C}_{d,T}$ contains the set \mathcal{C} for each parameter choice.

Comparing the red and black parameter sets, we can see in the left panel of Figure 4.7 that the red parameters lie on the same h^* -level-set curve as the black parameters, and thus return the same expanded safe set $\mathcal{C}_{d,T}$ as seen in the right panel of Figure 4.7. In contrast, we see in the right panel of Figure 4.7 that the closed-loop system is much less conservative and approaches the boundary of the safe set. This behavior indicates that ϵ_0 and λ can be jointly tuned to return the same theoretical guarantees as a small value of ϵ_0 without inducing conservativeness. The red parameter set has the same value of ϵ_0 as the green parameter set, allowing us to observe directly the impact of λ . We see in the right panel of Figure 4.7 that the red parameters yield a smaller set $\mathcal{C}_{d,T}$, implying stronger theoretical safety guarantees, while in Figure 4.6 we see that the closed-loop behavior of the red parameter set is much less conservative. Thus, the introduction of λ not only allows one to improve the theoretical safety guarantees for a fixed value of ϵ_0 , but also allows one to reduce the conservativeness of the system. Lastly, we see in Figure 4.6 that for all three parameter sets, the controller keeps the trajectories within \mathcal{C} , and thus, within $\mathcal{C}_{d,T}$, that is, it guarantees $h(\varphi_d(t)) \geq h^*$.

Conclusion

In this section I have presented an approach for robust safety-critical control synthesis in the presence of disturbances. This approach builds off the intuition that robustness to disturbances is most important when near the boundary of the set \mathcal{C} to propose a novel modification of ISSf in the form

of Tunable Input-to-State Safety (TISSf). The definition of TISSf allows for static parameters in ISSf-CBF-based controllers to be replaced by functions that weight robustness to disturbances according to the value of a CBF h . I then show how this allows for meaningful theoretical safety guarantees while permitting a significant reduction in conservativeness of a closed-loop system on an inverted pendulum.

4.4 Tunable Input-to-State Safety on Automated Truck

In this section I will present work deploying TISSf-based controllers on an automated semi-trailer truck. I will begin by describing a simplified model for the system dynamics, followed by codifying a safety specification of maintaining a safe following distance behind a lead vehicle in the form of a CBF. Next, I will design a nominal controller for this system which tries to track a desired velocity while maintaining ride comfort. I show in simulation that this nominal controller does not satisfy safety constraints, while the controller (CBF-QP) does meet safety constraints. In deploying the controller (CBF-QP) on the experimental system as seen in Figure 4.8, I find it fails to keep the system safe due to actuator dynamics that can be captured by a matched disturbance. I then deploy the TISSf-based control framework from Section 4.3 using different parameters, demonstrating the ability to achieve meaningful theoretical safety guarantees with desirable performance in both simulation and experimentally.

The contributions of this section are as follows:

- A design and evaluation of TISSf-based controllers for a heavy-duty semi-trailer truck system, both in simulation and experimentally.

The text for this section is adapted from:

A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, “Control barrier functions and input-to-state safety with application to automated vehicles,” *conditionally accepted as a Full Paper in IEEE Trans. on Control Sys. Tech.*, 2022.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.



Figure 4.8. Experimental configuration for heavy-duty semi-trailer truck problem. (Top) Controller design without robustifying element yields safety violation. (Bottom) Robust safety-critical controller ensures semi-trailer truck brakes early and aggressively enough to maintain safe distance.

System Model, Safety Specification, & Control Design

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis. Andrew Taylor’s particular contributions in this subsection are the verification of the CBF from equation (4.57) to (4.58).

We will begin by designing a safety-critical longitudinal controller for a connected automated semi-trailer truck. We first introduce the physical system and define a safe set via a CBF. Following this, we will present a nominal performance-based controller, and synthesize a safety-critical controller that modifies this nominal controller in a minimally invasive way while ensuring safety. We note that these constructions do not yet explicitly consider disturbances in the system model.

In this work we consider a rear-axle-driven truck without a trailer. Assuming the truck’s tires roll without slipping and the truck travels on a flat road with no headwind, the longitudinal dynamics of the truck are described by the following model:

$$\dot{v} = \frac{T}{m_{\text{eff}}R} - \frac{k_d v^2 + mg\sigma}{m_{\text{eff}}}. \quad (4.50)$$

Here the state is given by the truck’s longitudinal speed $v \in \mathbb{R}$, the input is the torque applied on the rear axle $T \in \mathbb{R}$, and the parameters in the model are the mass of the truck $m \in \mathbb{R}_{>0}$, the mass moment of inertia of the rotating elements $I \in \mathbb{R}_{>0}$, the tire radius $R \in \mathbb{R}_{>0}$, the effective mass $m_{\text{eff}} = m + \frac{I}{R^2}$, the air drag constant $k_d \in \mathbb{R}_{>0}$, gravitational acceleration $g \in \mathbb{R}_{>0}$, and rolling resistance coefficient $\sigma \in \mathbb{R}_{>0}$. Note that the second term in (4.50) is dissipative

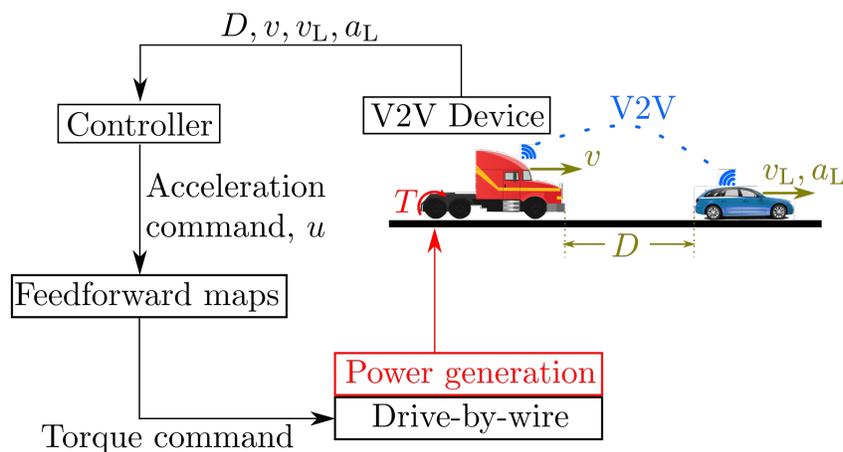


Figure 4.9. A connected automated truck following a connected vehicle.

in nature, and slows down the vehicle when it has a positive velocity. This term may be accounted for in the control design via feedback linearization, or may be ignored as its omission introduces a factor of conservativeness to the controller in terms of safety. The torque input to the system is computed from a desired longitudinal acceleration command $u \in \mathbb{R}$ via feed-forward maps. This torque input command is provided by a drive-by-wire system to the braking systems that produce the actual torque T ; see Figure 4.9. With these feed-forward maps in mind, we simplify the longitudinal dynamics model to:

$$\dot{v} = u. \quad (4.51)$$

Now let us consider the scenario when the truck follows a connected vehicle as depicted in Figure 4.9. Using the truck model in (4.51), the dynamics of this connected system are given by:

$$\begin{aligned} \dot{D} &= v_L - v, \\ \dot{v} &= u, \\ \dot{v}_L &= a_L, \end{aligned} \quad (4.52)$$

where $v_L, a_L \in \mathbb{R}$ are the speed and acceleration of the lead vehicle, respectively, and $D \in \mathbb{R}$ denotes the bumper-to-bumper headway distance between the truck and the lead vehicle, yielding the state $\mathbf{x} = [D, v, v_L]^\top \in \mathbb{R}^3$. The truck and lead vehicle are outfitted with vehicle-to-everything (V2X) communication systems, permitting the truck to receive motion information from

the lead vehicle such as its Global Positioning System (GPS) position which yields the distance D , its speed v_L , and its acceleration a_L . We assume that the leader's behavior satisfies:

$$a_L \in [-\underline{a}_L, \bar{a}_L], \quad v_L \in [0, \bar{v}_L], \quad (4.53)$$

where $\underline{a}_L, \bar{a}_L, \bar{v}_L \in \mathbb{R}_{>0}$ reflect a city-driving scenario; see Table 4.2 below.

The safety task for the truck is to maintain a safe distance behind the leader. This task motivates a CBF $h : \mathbb{R}^3 \rightarrow \mathbb{R}$ of the form:

$$h(D, v, v_L) = D - \rho(v, v_L), \quad (4.54)$$

where the headway function $\rho : \mathbb{R}^2 \rightarrow \mathbb{R}$ describes the minimum safe distance between the vehicles given their current velocities, v and v_L . Motivated by [62] and [236], we define the headway function as:

$$\rho(v, v_L) = c_0 + c_1 v + c_2 v_L + c_3 v^2 + c_4 v v_L + c_5 v_L^2, \quad (4.55)$$

with parameters $c_i \in \mathbb{R}$ for $i = 0, \dots, 5$; see Table 4.2. The value of the function ρ is seen in Figure 4.10. The corresponding set \mathcal{C} defined by h is given by:

$$\mathcal{C} = \left\{ \begin{bmatrix} D \\ v \\ v_L \end{bmatrix} \in \mathbb{R}^3 \mid D \geq \rho(v, v_L) \right\}. \quad (4.56)$$

To verify that the function h is a CBF for (4.52), observe that:

$$L_{\mathbf{g}} h(D_0, v_0, v_{L,0}) = 0 \implies c_1 + 2c_3 v_0 + c_4 v_{L,0} = 0, \quad (4.57)$$

which describes a line in (v, v_L) space where the CBF condition (2.61) must be met for all $D \in \mathbb{R}$. We consider $\alpha(r) = \alpha_c r$ with $\alpha_c \in \mathbb{R}_{>0}$, yielding:

$$\begin{aligned} L_{\mathbf{f}} h(D, v_0, v_{L,0}) + \alpha(h(D, v_0, v_{L,0})) &= v_{L,0} - v_0 - a_L(c_2 + c_4 v_0 + 2c_5 v_{L,0}) \\ &\quad + \alpha_c(D - \rho(v_0, v_{L,0})). \end{aligned} \quad (4.58)$$

Since checking the alternative CBF condition (2.61) analytically may be cumbersome using (4.58), we graphically evaluate it over a range of D and $v_{L,0}$ (and v_0 defined implicitly through (4.57)) that are states of interest in this problem, while taking the worst case value of a_L making (4.58) as negative

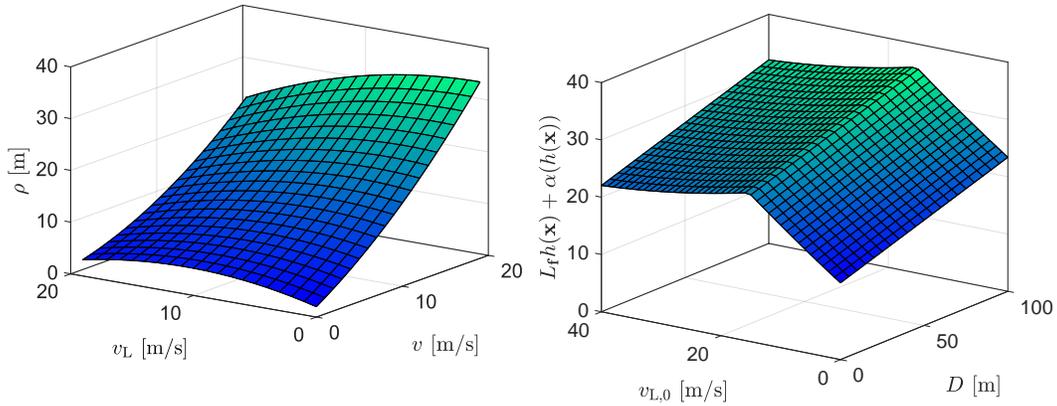


Figure 4.10. (Left) The value of the function ρ defined in (4.55), which defines the minimum safe following distance as a function of the leader's velocity v_L and truck velocity v . (Right) The value of $L_f h(\mathbf{x}) + \alpha(h(\mathbf{x}))$ as defined in (4.58) when $L_g h(\mathbf{x}) = 0$ for various distances. As the function is strictly positive over the domain of interest, h is a CBF on \mathcal{C} for (4.52).

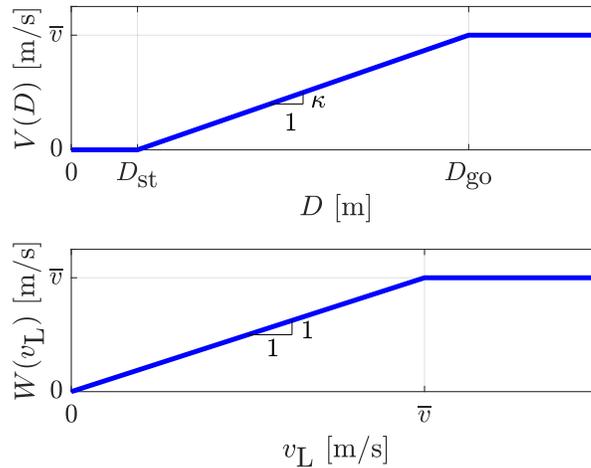


Figure 4.11. (Top) Range policy V defined in (4.60). (Bottom) Speed policy W defined in (4.61).

as possible; see Figure 4.10. For $\alpha_c = 0.1$ [1/s], the value of (4.58) is strictly positive, ensuring the CBF condition (2.61) is met.

Beyond the task of safety, we wish for the controller to have other desirable properties such as plant stability and string stability in the presence of communication delay [237], or optimal performance regarding energy efficiency and passenger comfort [238]. To accomplish this, we first design a nominal controller that prioritizes performance. In particular, we design a connected cruise controller (CCC) for the truck that utilizes information about the lead vehicle

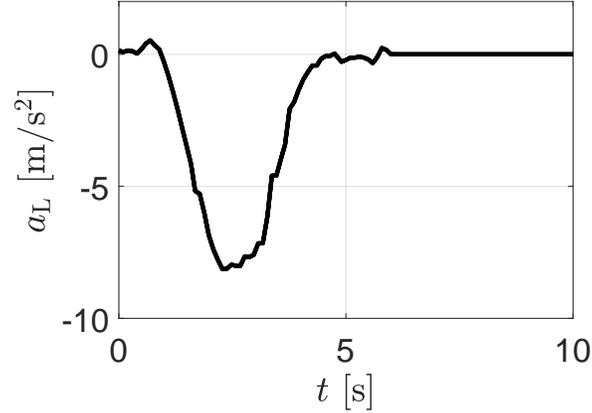


Figure 4.12. Example profile for acceleration a_L of lead vehicle used in simulation.

available through V2X connectivity. We propose the controller structure:

$$k_{\text{nom}}(D, v, v_L) = A(V(D) - v) + B(W(v_L) - v), \quad (4.59)$$

with parameters $A, B \in \mathbb{R}_{>0}$ and functions $V : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and $W : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. The first term in (4.59) specifies a distance-based speed error with:

$$V(D) = \max \{0, \min\{\kappa(D - D_{\text{st}}), \bar{v}\}\}, \quad (4.60)$$

depicted in the top panel of Figure 4.11, producing a desired speed based on the distance D . Here $D_{\text{st}} \in \mathbb{R}_{>0}$ is the desired stopping distance, $1/\kappa > 0$ is the desired time headway, and $D_{\text{go}} = \bar{v}/\kappa + D_{\text{st}}$. The second term in (4.59) specifies the error related to the relative speed with the speed policy:

$$W(v_L) = \min\{v_L, \bar{v}\}, \quad (4.61)$$

depicted in the bottom panel of Figure 4.11, which bounds the speed error if the lead vehicle violates $v_L \leq \bar{v}$.

| | | |
|---|--------------------------------------|----------------------------------|
| $\bar{a}_L = 5 \text{ [m/s}^2\text{]}$ | $c_0 = 2 \text{ [m]}$ | $\kappa = 0.8 \text{ [1/s]}$ |
| $\underline{a}_L = 10 \text{ [m/s}^2\text{]}$ | $c_1 = 1.1 \text{ [s]}$ | $\alpha_c = 0.1 \text{ [1/s]}$ |
| $\bar{v}_L = 20 \text{ [m/s]}$ | $c_2 = 0.6 \text{ [s]}$ | $D_{\text{st}} = 5 \text{ [m]}$ |
| $\delta = 4.5 \text{ [m/s}^2\text{]}$ | $c_3 = 0.03 \text{ [s}^2\text{/m]}$ | $D_{\text{go}} = 30 \text{ [m]}$ |
| $\epsilon_0 = 0.5 \text{ [s}^3\text{/m]}$ | $c_4 = -0.03 \text{ [s}^2\text{/m]}$ | $A = 0.4 \text{ [1/s]}$ |
| $\lambda = 0.4 \text{ [1/m]}$ | $c_5 = -0.03 \text{ [s}^2\text{/m]}$ | $B = 0.5 \text{ [1/s]}$ |

Table 4.2. Parameter values used in controller design.

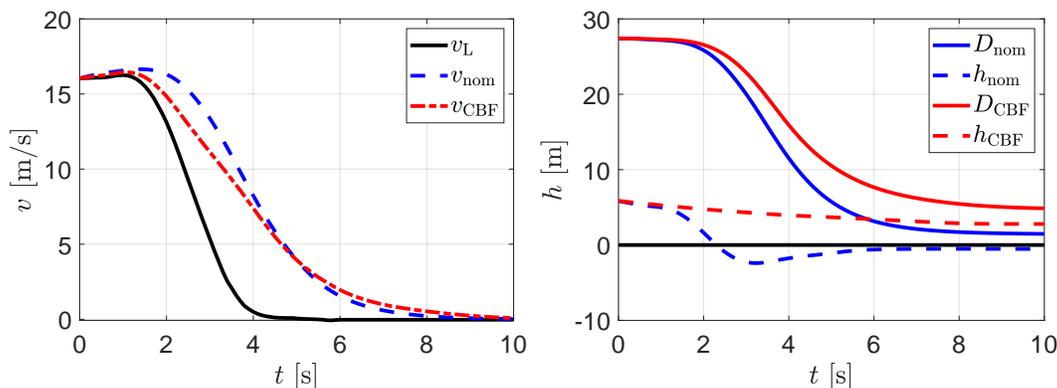


Figure 4.13. (Left) Velocity v_L of lead vehicle (black) and velocity of the truck using the nominal controller (4.59) (blue) and (CBF-QP) controller (red). (Right) Following distance D and value of CBF h using the nominal controller (4.59) and (CBF-QP) controller.

We simulate both the nominal controller k_{nom} and (CBF-QP) controller from the initial condition $\mathbf{x}_0 = [27.4, 16, 16]^\top \in \mathcal{C}$. We use parameter values as specified in Table 4.2. The acceleration a_L of the lead vehicle is given by a time profile reflecting a hard braking event, as seen in Figure 4.12. The velocity of the truck converges to zero and a crash does not occur for both controllers, but only the (CBF-QP) controller ensures the truck maintains a safe distance (indicated by $h_{\text{CBF}}(\varphi(t)) \geq 0$) as seen in Figure 4.13. Observe that the nominal controller brakes less aggressively than the (CBF-QP) controller, and thus does not react quickly enough to avoid violating the safe following distance requirement.

Robust Control Design & Experimental Results

We now provide a description of the automated semi-trailer truck experimental configuration and present results using the nominal controller (4.59) and (CBF-QP) controller. Subsequently, we will deploy TISSf-based controllers and demonstrate their advantages experimentally.

Note that the contents of the following paragraph were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

The automated truck used in the experiments is an International ProStar+ Class-8 truck developed by the Navistar [239]; see Figure 4.14(a). Both the automated truck and the lead vehicle are equipped with a V2X Onboard Unit (OBU) developed by Commsignia [240]. These units are equipped with an

accelerometer, gyroscope, magnetometer, and GPS unit. Furthermore, these OBUs support peer-to-peer communication such that the automated truck may receive position, velocity, and acceleration data from the lead vehicle through V2X antennas shown in Figure 4.14(a). The automated truck is additionally equipped with a Mobile Real-Time Targeting Machine developed by Speedgoat [241], which interfaces with the V2X OBU and the truck’s Engine Controller Unit (ECU) through a Controller Area Network (CAN) bus. The Speedgoat runs the controller for the system given a stream of values for D , v , v_L , and a_L coming from the V2X OBUs. It computes a desired acceleration input and converts it to a corresponding torque value through a feed-forward map. A drive-by-wire system controls the engine and the brake torques accordingly. The steering of the truck is done by a human driver in the experiments.

In an effort to evaluate the repeatability of the experiments, it is necessary to eliminate variation in the lead vehicle’s behavior, which is being driven by a human. To achieve this, we use a recorded time profile of position, velocity, and acceleration of the lead vehicle while it performs a hard braking event. This profile for a_L and v_L is seen in Figure 4.12 and the left panel of Figure 4.13, respectively, and was used to produce the simulation results. During the experiments, this data is played back to the truck controller as a *perceived lead vehicle*, which the truck uses in its control computations. Experiments also include a physical lead vehicle for visualization purposes which travels on the other lane for safety reasons; see Figure 4.14(b)-(e) and notice the “collision” in panel (c). We note that this lead vehicle closely follows the previously recorded time profiles, but does not provide any data to the truck during the actual experiments. Importantly, the quantitative analysis of safety is performed by evaluating the CBF using the truck’s measured state with the recorded time profiles of the perceived lead vehicle.

The nominal controller and (CBF-QP) controller are deployed on the automated truck with results as seen in Figure 4.15. Observe that not only does the nominal controller consistently fail to meet the safety requirements imposed by the CBF h , but the (CBF-QP) controller also consistently fails to meet the safety requirements. The top row in Figure 4.8 illustrates an experimental run with the nominal controller. To understand why the (CBF-QP) controller fails, we examine the discrepancy between the commanded acceler-

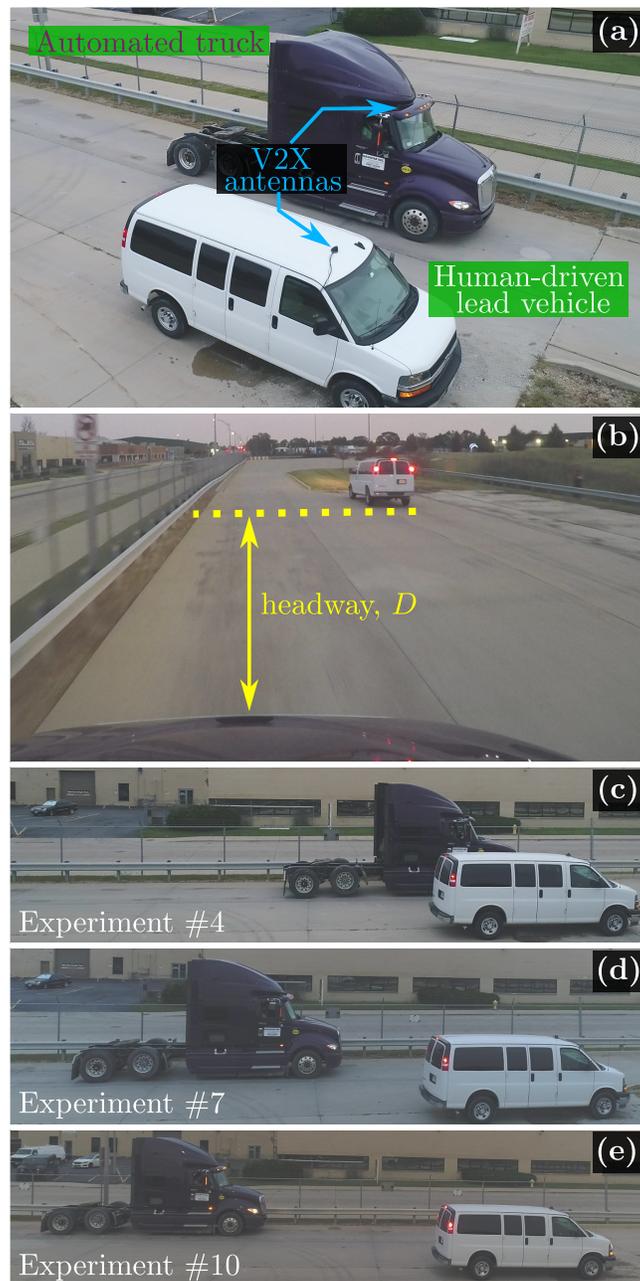


Figure 4.14. (a) Vehicles used in experiments. (b) Image from the dashboard of the truck during an experimental run. (c,d,e) Final configurations of separate experiments. See <https://youtu.be/9dJtC1TCBbA> for a video.

ation and actual acceleration of the automated truck, as seen in the left panel of Figure 4.16. One may observe a delay between the commanded acceleration and the achieved acceleration. This delay in acceleration is due to the fact that the brakes of the truck are a complex nonlinear dynamical system that has been imperfectly abstracted away by the feed-forward maps that allow the

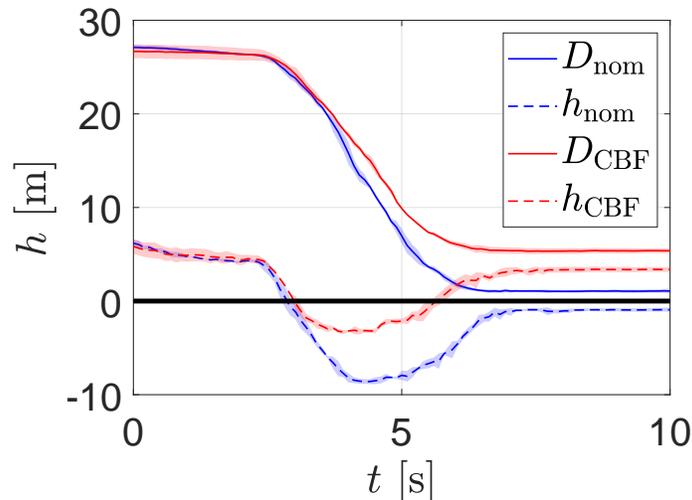


Figure 4.15. Mean value (lines) and standard deviations (fills) of the distance D and the CBF h when using the nominal controller (4.59) (blue) and the (CBF-QP) controller (red) in the truck experiment. The experiments with these controllers are highly consistent.

simplified model in (4.51). Rather than working with the braking dynamic system and improving the feed-forward maps, we describe the discrepancy in commanded and actual acceleration as a disturbance in the simplified model:

$$\dot{v} = u + d(t), \quad (4.62)$$

where $d: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ reflects the difference between commanded acceleration and actual acceleration. As the disturbance d is caused by the complicated interactions of the drive-by-wire system and brake dynamics, it may be difficult to use model-based techniques to construct a meaningful bound $\delta \in \mathbb{R}_{\geq 0}$ for the worst-case disturbance magnitude. Instead, we estimate the worst-case disturbance empirically by comparing the actual acceleration $\dot{v}(t)$ to the commanded acceleration $u(t)$. Observe in the right panel of Figure 4.16 that the largest difference in the commanded and actual acceleration is around $4 \text{ [m/s}^2\text{]}$. Thus, we study the degradation of safety taking a slightly larger value $\delta = 4.5 \text{ [m/s}^2\text{]}$.

To overcome this disturbance and improve the safe behavior of the truck, we deploy the tool of TISSf-CBFs described in Section 4.3. As h satisfies the CBF condition (2.61), and our disturbance it is matched, it also satisfies the TISSf-CBF condition (4.34), where we take:

$$\epsilon(r) = \epsilon_0 e^{\lambda r}, \quad (4.63)$$

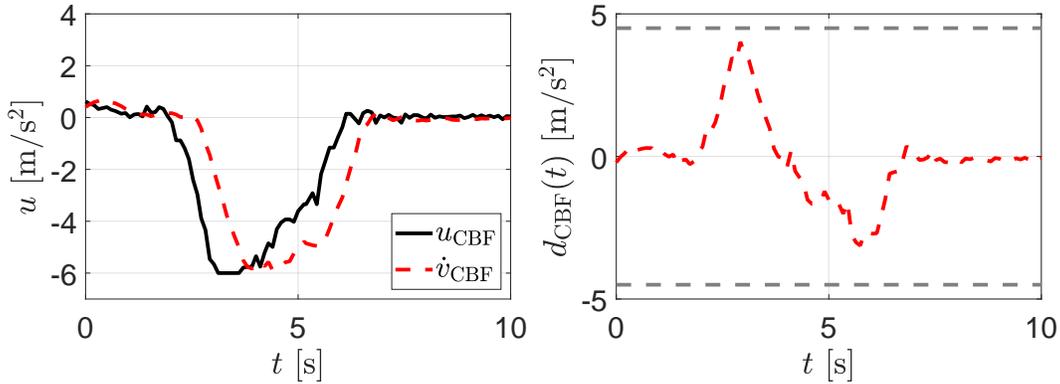


Figure 4.16. (Left) Discrepancy between acceleration commanded by safety-critical controller (black) and actual acceleration of the automated truck (red). (Right) Disturbance signal in input seen by the truck used to define the model (4.62).

with $\epsilon_0 \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}_{\geq 0}$. The parameter λ introduces a measure of flexibility by allowing one to require a greater degree of robustness when the truck is close to the leading vehicle, and less robustness when the distance is greater. Given (4.63) and the particular value of δ we consider, the smallest set that is forward invariant up to δ is given by:

$$\mathcal{C}_\delta = \left\{ \begin{bmatrix} D \\ v \\ v_L \end{bmatrix} \in \mathbb{R}^3 \mid h(D, v, v_L) \geq -\frac{\epsilon_0 e^{\lambda h(D, v, v_L)} \delta^2}{4\alpha_c} \right\}. \quad (4.64)$$

As in the inverted pendulum example in Section 4.3, the set \mathcal{C}_δ being forward invariant up to δ implies that $h(\varphi_d(t)) \geq h^*$, where h^* is the value of the TISSf-CBF h on the boundary of \mathcal{C}_δ , which can be calculated by solving (4.31). The value of h^* for different choices of ϵ_0 and λ can be seen in Table 4.3 below.

We simulate the nominal controller, (CBF-QP) controller, and (TISSf-CBF-QP) controller from the initial condition $\mathbf{x}(0) = [27.4, 16, 16]^\top \in \mathcal{C}$ while disturbing the input using the signal shown in the right panel of Figure 4.16. We use parameter values as specified in Table 4.2. Observe in Figure 4.17 that introducing the disturbance signal into our simulation allows the recreation of the failures of the nominal controller and (CBF-QP) controller that were seen experimentally in Figure 4.15. Furthermore, observe that the (TISSf-CBF-QP) controller maintains the safety of the system even in the presence of the disturbance.

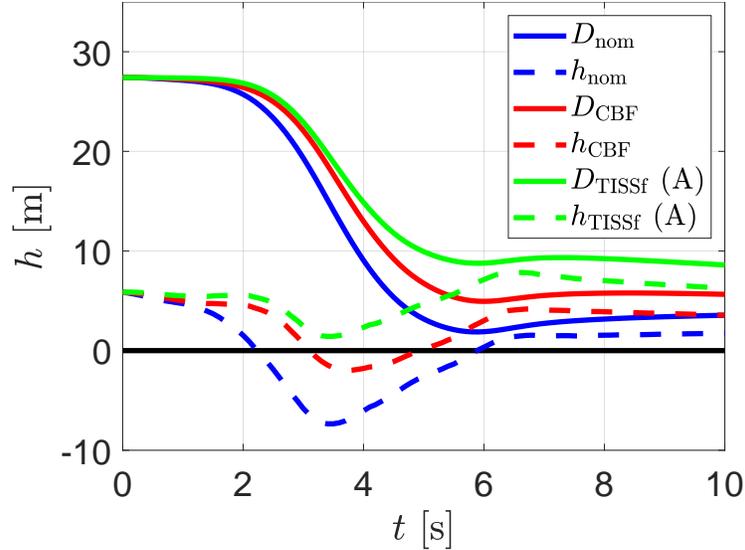


Figure 4.17. Following distance D and value of h using the nominal controller (4.59) (blue), (CBF-QP) controller (red), and (TISSf-CBF-QP) controller (green) in the disturbed simulation.

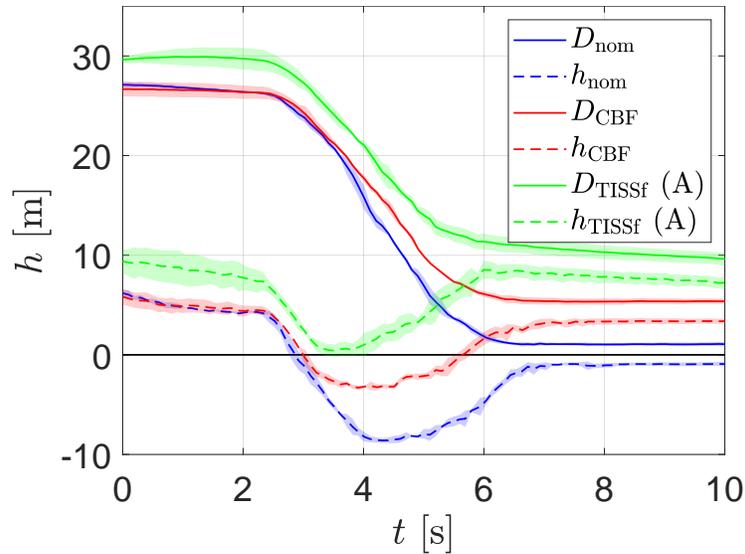


Figure 4.18. Mean value (lines) and standard deviations (fills) of the following distance D and h using the nominal controller (4.59) (blue), (CBF-QP) controller (red), and (TISSf-CBF-QP) controller (green) in experiment.

We next deploy the TISSf-based control design on the experimental platform. Sets of three experimental runs were conducted using each parameter pair ϵ_0 and λ shown in Table 4.3. The experimental results using the parameter set $\epsilon_0 = 0.5$ [s³/m] and $\lambda = 0.4$ [1/m] (labeled as parameter pair (A)) can be seen in Figure 4.17 and are visualized in the bottom row of Figure 4.8. With these parameters the system is safe, as the value of h does not drop below

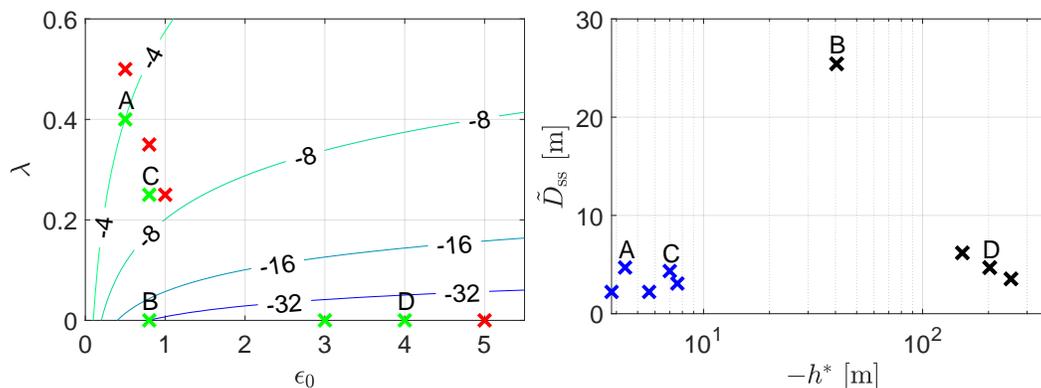


Figure 4.19. (Left) Parameter values for ϵ_0 and λ used in the truck experiments, with contours showing theoretical values of h^* . Green markers denote parameter sets which achieve the original safety goal ($h(\varphi_d(t)) \geq 0$), while red markers denote parameter sets for which the original safety goal is violated. (Right) Theoretical values of h^* and the shift in steady-state tracking distance \tilde{D}_{ss} , for the parameter sets used in the truck experiments. The blue markers denote parameter sets with $\lambda > 0$, while the black markers denote parameter sets with $\lambda = 0$.

0. Although the TISSf-based controller has a larger standard deviation across the experimental runs, it consistently satisfies the original safety requirement.

To evaluate how the system behavior depends on the values of the parameters ϵ_0 and λ , we first consider whether the original safety requirement is met, i.e., whether or not the value of h remains positive. While the TISSf-based controller does not provide a theoretical guarantee that h will remain non-negative (it guarantees that $h(\varphi_d(t)) \geq h^*$), for certain values of ϵ_0 and λ the original safety requirement is still met. The minimum value h_{\min} of the TISSf-CBF observed during the experimental runs is shown in Table 4.3. In the left panel of Figure 4.19, green and red markers indicate parameter values for which the original safety requirement is and is not met, respectively. Safety is achieved using the ISSf-CBF formulation presented in Section 2.9 (with $\lambda = 0$) for small values of ϵ_0 , but is also achieved using small values of λ .

We remark that when changing the controller from k_{CBF} to k_{TISSf} , the equilibrium of the system is shifted as can be noticed when comparing the runs in Figure 4.18. We characterize this by the shift in the steady-state tracking distance error:

$$\tilde{D}_{ss} \triangleq D_{ss}^{\text{exp}} - D^*. \quad (4.65)$$

Here D_{ss}^{exp} is the steady-state distance captured in experiments when the

| Label | ϵ_0 [s ³ /m] | λ [1/m] | h^* [m] | h_{\min} [m] | \tilde{D}_{ss} [m] |
|-------|-------------------------------------|--------------------|--------------|-------------------|-------------------------|
| (B) | 0.8 | 0 | -40.50 | 22.09 | 25.43 |
| | 3 | 0 | -151.88 | 2.99 | 6.19 |
| (D) | 4 | 0 | -202.50 | 1.02 | 4.69 |
| | 5 | 0 | -253.13 | -0.45 | 3.54 |
| (A) | 0.5 | 0.4 | -4.38 | 0.35 | 4.70 |
| | 0.5 | 0.5 | -3.80 | -1.27 | 2.22 |
| (C) | 0.8 | 0.25 | -7.01 | 0.78 | 4.34 |
| | 0.8 | 0.35 | -5.64 | -1.03 | 2.22 |
| | 1.0 | 0.25 | -7.59 | -0.86 | 3.07 |

Table 4.3. Sets of parameter values used for the exponential function (4.63) in the automated truck experiments with theoretical safety guarantee h^* , minimum experimental value of the TISSf-CBF h_{\min} , and shift in the steady-state tracking distance \tilde{D}_{ss} in (4.65).

leader is moving with the steady-state speed $v^* \in (0, \bar{v})$ before braking. The term $D^* = V^{-1}(v^*)$ captures the desired steady-state distance given by the inverse of the range policy (4.60). In the experiments $v^* = 16$ [m/s], yielding $D^* = 25$ [m]. The values of \tilde{D}_{ss} corresponding to different parameter pairs are given in Table 4.3. In the right panel of Figure 4.19 we visualize the theoretical values of h^* and the experimental values of \tilde{D}_{ss} for different parameter sets. The black markers indicate parameter sets with $\lambda = 0$, while the blue markers show parameter sets with $\lambda > 0$. With $\lambda = 0$, the theoretical guarantees are nearly meaningless (observe the large negative values of h^*), and improving them requires dramatically increasing \tilde{D}_{ss} . In contrast, the parameter sets with $\lambda > 0$ allow us to obtain significantly (an order of magnitude) stronger theoretical guarantees without greatly increasing \tilde{D}_{ss} , thereby also achieving good performance. In the Figure 4.20, we give experimental results of three other parameter pairs labeled as (B), (C), and (D) in Table 4.3. The poor performance of case (B) is indicated by the large value of \tilde{D}_{ss} . The results for cases (C) and (D) nearly overlap, but the introduction of λ yield a strong theoretical guarantee for case (C), which is missing for case (D).

Conclusion

In conclusion, in this section I have developed a robust yet performant safety-critical controller for a connected automated semi-trailer truck system. I begin by presenting a simple model for this system, a safety specification via a CBF, a nominal controller, and simulation and experimental results using a (CBF-QP)

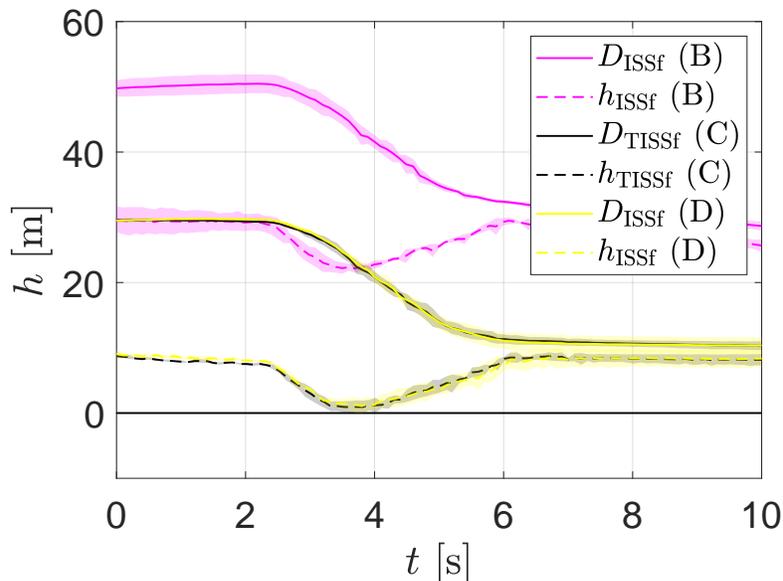


Figure 4.20. Experimental results for parameter pairs (B), (C), and (D) in Table 4.3. Case (B) is conservative as indicated by the large steady-state tracking distance. Cases (C) and (D) display nearly identical behavior, but case (C) possesses a stronger theoretical guarantee.

controller. Unmodeled actuator dynamics lead to a matched disturbance in the experimental system, resulting in the (CBF-QP) controller failing to keep the experimental system safe. I develop a collection of TISSf-based controllers and explore how the parameters of these controllers can be tuned to achieve both meaningful theoretical safety guarantees and strong performance. These conclusions are then realized on the experimental system.

4.5 Safety with Stochastic Disturbances

In this section I will present work on safety-critical control in the presence of stochastic disturbances. This work will study conditions on a discrete-time system and a corresponding CBF for which it is possible to synthesize controllers that return finite-time ISSf-like safety guarantees using convex optimization. I will begin by reviewing discrete-time dynamics and CBFs for discrete-time systems. I will then briefly review martingales, Ville’s Inequality, and Jensen’s Inequality, which are tools for analyzing stochastic systems. Next, I will propose an ISSf-like notion of finite-time safety for discrete-time systems with stochastic disturbances, and establish how a discrete-time BF meeting an expectation condition on its change across time steps certifies this safety property. Following this, I will explore how a discrete-time CBF satisfying certain convexity properties can be used to synthesize a safety-critical

controller that yields a discrete-time BF. Lastly, I show the proposed method in simulation on several examples, including a quadrupedal robot.

The contributions of this section are as follows:

- An ISSf-like characterization of finite-time safety for discrete-time systems experiencing stochastic disturbances.
- An exploration of how convexity properties of a discrete-time CBF must be utilized with a stochastic disturbance signal to ensure the convexity of downstream control synthesis.

The text for this section is adapted from:

R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, “Robust safety under stochastic uncertainty with discrete-time control barrier functions,” *accepted in Robotics: Science and Sys. (RSS) XIX*, 2023.

A. J. Taylor participated in the conception of the project, theoretical analysis, and writing of the article.

Discrete Dynamics

In this section I provide a review of safety-critical control systems for discrete-time nonlinear systems via CBFs. Consider a discrete-time (DT) open-loop system with dynamics given by:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k), \quad \forall k \in \mathbb{N}, \quad (4.66)$$

with state $\mathbf{x}_k \in E$, input $\mathbf{u}_k \in \mathbb{R}^m$, and dynamics $\mathbf{F} : E \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ that are continuous on $E \times \mathbb{R}^m$. A controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is continuous on E yields the DT closed-loop system:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)), \quad \forall k \in \mathbb{Z}_{\geq 0}. \quad (4.67)$$

I assume for any $\mathbf{x}_0 \in E$, there exists a set $I(\mathbf{x}_0) = \{0, 1, \dots, k_{\max}(\mathbf{x}_0)\} \subseteq \mathbb{N}$ and a unique solution $\varphi : I(\mathbf{x}_0) \rightarrow E$ satisfying:

$$\varphi(k+1) = \mathbf{F}(\varphi(k), \mathbf{k}(\varphi(k))), \quad \forall k \in I(\mathbf{x}_0), \quad (4.68)$$

$$\varphi(0) = \mathbf{x}_0. \quad (4.69)$$

The notion of safety for systems of this form is similarly formalized using the concept of forward invariance [104]:

Definition 35 (*Discrete-Time Forward Invariance & Safety*). A set $\mathcal{C} \subset E$ is *forward invariant* for the closed-loop system (4.67) if $\mathbf{x}_0 \in \mathcal{C}$ implies that $\boldsymbol{\varphi}(k) \in \mathcal{C}$ for all $k \in I(\mathbf{x}_0)$. In this case, the closed-loop system (4.67) is said to be *safe* with respect to the set \mathcal{C} .

Discrete-time barrier functions (DT-BFs) are a tool for guaranteeing the safety of discrete-time systems. Consider a set \mathcal{C} expressed as the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuous on E . I refer to such a function h as a DT-BF¹ if it satisfies the following properties [234]:

Definition 36 (*Discrete-Time Barrier Function (DT-BF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuous on E . The function h is a *discrete-time barrier function (DT-BF)* for the closed-loop system (4.67) on \mathcal{C} if there exists an $\alpha \in [0, 1]$ such that:

$$h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x}))) \geq \alpha h(\mathbf{x}). \quad (4.70)$$

for all $\mathbf{x} \in E$,

DT-BFs serve as a certificate of forward invariance as captured in the following theorem [234]:

Theorem 30. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuous on E . If h is a DT-BF for the closed-loop system (4.67) on \mathcal{C} , then the closed-loop system (4.67) is safe with respect to the set \mathcal{C} .*

Intuitively, the value of $h(\boldsymbol{\varphi}(k))$ can only decay as fast as the geometric sequence $\alpha^k h(\mathbf{x}_0)$, which is lower-bounded by 0, thus ensuring the safety (i.e., forward invariance) of \mathcal{C} .

Discrete-time Control Barrier Functions (DT-CBFs) [234] provide a tool for constructively synthesizing controllers that yield closed-loop systems that possess a DT-BF:

Definition 37 (*Discrete-Time Control Barrier Function (DT-CBF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuous on E .

¹The state constraint $\mathbf{x}_k \in \mathcal{C}$ for all $k \in I(\mathbf{x}_0)$, when expressed as $h(\mathbf{x}_k) \geq 0$ for all $k \in I(\mathbf{x}_0)$, is the special case of a DT-BF with $\alpha = 0$.

The function h is a *discrete-time Control Barrier Function (DT-CBF)* for the open-loop system (4.66) on \mathcal{C} if there exists an $\alpha \in [0, 1]$ such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} h(\mathbf{F}(\mathbf{x}, \mathbf{u})) > \alpha h(\mathbf{x}), \quad (4.71)$$

for all $\mathbf{x} \in E$.

Given a CBF h for (4.66) and a corresponding $\alpha \in [0, 1]$, I define the point-wise set of control values:

$$K_{\text{CBF}}(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^m \mid h(\mathbf{F}(\mathbf{x}, \mathbf{u})) \geq \alpha h(\mathbf{x})\}. \quad (4.72)$$

This yields the following result [242]:

Theorem 31. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuous on E . If h is a DT-CBF for the open-loop system (4.66) on \mathcal{C} , then the set $K_{\text{CBF}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in E$, and for any controller \mathbf{k} that is continuous on E with $\mathbf{k}(\mathbf{x}) \in K_{\text{CBF}}(\mathbf{x})$ for all $\mathbf{x} \in E$, the function h is a DT-BF for the closed-loop system (4.67) on \mathcal{C} .*

Given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is continuous on E and a DT-CBF h for the open-loop system (4.66) on \mathcal{C} , a controller can be specified via the following optimization problem:

$$\begin{aligned} \mathbf{k}_{\text{CBF}}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & (\text{DT-CBF-OP}) \\ \text{s.t. } & h(\mathbf{F}(\mathbf{x}, \mathbf{u})) \geq \alpha h(\mathbf{x}), \end{aligned}$$

noting that $\mathbf{k}_{\text{CBF}}(\mathbf{x}) \in K_{\text{CBF}}(\mathbf{x})$ for all $\mathbf{x} \in E$. I note that unlike the affine inequality constraint that arises in the continuous-time (CBF-QP), the DT-CBF inequality constraint (4.71) is not necessarily convex with respect to the input, preventing it from being integrated into a convex optimization-based controller. To solve this issue, it is often assumed that the function $h \circ \mathbf{F} : E \times \mathbb{R}^m \rightarrow \mathbb{R}$ is concave with respect to its second argument [234], [243], [244]. This assumption will be studied later in Section 6.4.

Martingales & Jensen's Inequality

In this section I review tools from probability theory that will allow utilizing information about the distribution of a stochastic disturbance signal in constructing a notion of stochastic safety and corresponding safety-critical controllers. I choose to provide this background material at the level necessary

to understand subsequent constructions of stochastic safety and safety-critical controllers, but refer to [245] for a precise measure-theoretic presentation of the following concepts. The key tool underlying the construction of a notion of stochastic safety is a nonnegative supermartingale, a specific type of expectation-governed random process:

Definition 38 (*Martingales*). Let \mathbf{x}_k , $k \in \mathbb{Z}_{\geq 0}$, be a sequence of random variables that take values in \mathbb{R}^n , $W : \mathbb{R}^n \times \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}$, and suppose that $\mathbb{E}[|W(\mathbf{x}_k, k)|] < \infty$ for all $k \in \mathbb{Z}_{\geq 0}$. The process $W_k \triangleq W(\mathbf{x}_k, k)$ is a supermartingale if:

$$\mathbb{E}[W_{k+1} \mid \mathbf{x}_{0:k}] \leq W_k \text{ almost surely for all } k \in \mathbb{Z}_{\geq 0}, \quad (4.73)$$

where $\mathbf{x}_{0:k}$ indicates the random variables $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k\}$. If, additionally, $W_k \geq 0$ for all $k \in \mathbb{Z}_{\geq 0}$, W_k is a nonnegative supermartingale. If the process is non-decreasing in expectation, the process W_k is a submartingale. If the inequality (4.73) holds with equality, the process W_k is a martingale.

An important result from martingale theory that I will use to develop probabilistic safety guarantees is *Ville's inequality*, which allows bounding the probability that a nonnegative supermartingale will rise above a certain value [246]:

Theorem 32. *If W_k is a nonnegative supermartingale, then for all $\lambda \in \mathbb{R}_{>0}$:*

$$\mathbb{P} \left\{ \sup_{k \in \mathbb{Z}_{\geq 0}} W_k > \lambda \right\} \leq \frac{\mathbb{E}[W_0]}{\lambda}. \quad (4.74)$$

Intuitively, Ville's inequality can be compared with Markov's inequality [245] for nonnegative random variables; since the process W_k is non-increasing in expectation, Ville's inequality allows for controlling the probability the process instead moves upward above λ . Lastly, as I will show when synthesizing safety-critical controllers in the presence of stochastic disturbances, conditions will need to be enforced on the expectation of a DT-CBF. In doing so, I will need to relate the expectation of the DT-CBF $h(\mathbf{x}_{k+1})$ to the expectation of the state \mathbf{x}_{k+1} . This will be achieved using Jensen's inequality [247]:

Theorem 33. *Consider a function $h : E \rightarrow \mathbb{R}$ that is continuous on E and a random variable \mathbf{x} that takes values in E with $\mathbb{E}[||\mathbf{x}||] < \infty$. We have that:*

$$\begin{cases} \text{if } h \text{ is convex,} & \text{then } \mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]), \\ \text{if } h \text{ is concave,} & \text{then } \mathbb{E}[h(\mathbf{x})] \leq h(\mathbb{E}[\mathbf{x}]). \end{cases} \quad (4.75)$$

Safety of Stochastic Systems

In this section I provide one of the main results of this work in the form of a bound on the probability that a system with stochastic disturbances will exit a given superlevel set of a DT-BF over a finite time horizon. Consider the following the DT open-loop system with additive disturbances:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k, \quad \forall k \in \mathbb{Z}_{\geq 0}, \quad (4.76)$$

with $\mathbf{d}_k \in \mathbb{R}^n$. A controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is continuous on E yields the following closed-loop system:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)) + \mathbf{d}_k, \quad \forall k \in \mathbb{Z}_{\geq 0}. \quad (4.77)$$

I assume that for any initial condition $\mathbf{x}_0 \in E$ and sequence of disturbances $\mathbf{D} = \{\mathbf{d}_k\}$ for $k \in \mathbb{Z}_{\geq 0}$, there exists a set $I_d(\mathbf{x}_0, \mathbf{D}) = \{0, 1, \dots, k_{\max}(\mathbf{x}_0, \mathbf{D})\} \subseteq \mathbb{Z}_{\geq 0}$ and a unique solution $\varphi_d : I(\mathbf{x}_0, \mathbf{D}) \rightarrow E$ satisfying:

$$\varphi_d(k+1) = \mathbf{F}(\varphi_d(k), \mathbf{k}(\varphi_d(k))) + \mathbf{d}_k, \quad \forall k \in I_d(\mathbf{x}_0, \mathbf{D}), \quad (4.78)$$

$$\varphi_d(0) = \mathbf{x}_0. \quad (4.79)$$

I assume that $\mathbf{x}_0 \in E$ is known and consider sequences of disturbances $\mathbf{D} = \{\mathbf{d}_k\}$ for $k \in \mathbb{Z}_{\geq 0}$ that are independently and identically distributed (with distribution \mathcal{D}) random variables² with (potentially unbounded) support on \mathbb{R}^n , generating the random process $\mathbf{x}_{1:k} \triangleq \{\varphi_d(1), \varphi_d(1), \dots, \varphi_d(k)\}$ (this is done for notational convenience). To study the safety of this system, I will use the following definition:

Definition 39 (*K-Step Exit Probability*). Let the function $h : E \rightarrow \mathbb{R}$ be continuous on E . For any $K \in \mathbb{Z}_{\geq 0}$, $\gamma \in \mathbb{R}_{\geq 0}$, and initial condition $\mathbf{x}_0 \in E$, the *K-step exit probability* of the closed-loop system (4.77) is given by:

$$P_u(K, \gamma, \mathbf{x}_0) = \mathbb{P} \left\{ \min_{k \in \{0, \dots, K\}} h(\mathbf{x}_k) < -\gamma \right\}. \quad (4.80)$$

This definition describes the probability that the state of the closed-loop system will leave the $-\gamma$ superlevel set of h within K steps. This is inherently a finite-time safety guarantee, as it only considers finite values of K . The

²This implies the dynamics define a Markov process, i.e. $\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)) + \mathbf{d}_k) \mid \mathbf{x}_{0:k}] = \mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)) + \mathbf{d}_k) \mid \mathbf{x}_k]$, since the state \mathbf{x}_{k+1} at time $k+1$ only depends on the state \mathbf{x}_k and disturbance \mathbf{d}_k at time k .

finite-time aspect of K -step exit probabilities is critical since systems exposed to unbounded disturbances will exit a bounded set with probability $P_u = 1$ over an infinite horizon [61], [68]. Intuitively, this is because a sufficiently large sample will eventually be drawn from the tail of the distribution that forces the system out in a single step. Additionally, observe that this definition is directly related ISSf as it reasons about all superlevel sets of the function h , rather than just the 0-superlevel set. For the remainder of this work, I will omit the dependence of P_u on K , γ , and \mathbf{x}_0 for notational simplicity.

Given this definition, I now provide one of the main results of this work by relating DT-BFs to K -step exit probabilities. I note that this result is a reframing of the stochastic invariance theorem in [59], [63]. This reframing features three key components. First, it develops the results using the standard formulation of DT-BFs covered above. Second, it yields a probability bound not only for \mathcal{C} (defined as the 0-superlevel set of h , such that $\gamma = 0$), but for all non-positive superlevel sets of h ($\gamma \geq 0$), providing a stochastic variant of ISSf. Third, I present a complete proof of the result, with the goal of illuminating how to leverage tools from martingale theory to reason about the safety of discrete-time stochastic systems.

Theorem 34. *Let the function $h : E \rightarrow \mathbb{R}$ be continuous and upper-bounded on E with upper bound $M \in \mathbb{R}_{>0}$. Suppose there exists an $\alpha \in (0, 1)$ and³ $\delta \leq M(1 - \alpha)$ such that the closed-loop system (4.77) satisfies:*

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbf{d}) \mid \mathbf{x}] \geq \alpha h(\mathbf{x}) + \delta, \quad (4.81)$$

for all $\mathbf{x} \in E$, with $\mathbf{d} \sim \mathcal{D}$. For any $K \in \mathbb{N}$ and $\gamma \in \mathbb{R}_{\geq 0}$, if $\delta < -\gamma(1 - \alpha)$, we have that:

$$P_u \leq \left(\frac{M - h(\mathbf{x}_0)}{M + \gamma} \right) \alpha^K + \frac{M(1 - \alpha) - \delta}{M + \gamma} \sum_{i=1}^K \alpha^{i-1}. \quad (4.82)$$

Alternatively if $\delta \geq -\gamma(1 - \alpha)$, then:

$$P_u \leq 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma} \left(\frac{M\alpha + \gamma + \delta}{M + \gamma} \right)^K. \quad (4.83)$$

³The original presentation of Theorem 34 in [59] considers variable δ_k for $k \in \{0, \dots, K\}$, which are known *a priori*. In most practical applications, one assumes a lower bound that holds for all δ_k , motivating the use of a constant δ . Moreover, the use of a constant δ significantly clarifies the proof.

The upper bound $\delta \leq M(1 - \alpha)$ is relatively non-restrictive, as not only is δ typically negative, but it must hold such that, in expectation, $h(\mathbf{x}_{k+1})$ cannot rise above the upper bound M on h . The switching condition between (4.82) and (4.83) of $\delta = \gamma(1 - \alpha)$ corresponds to whether, in expectation, the one-step evolution of the system remains in the set $\mathcal{C}_\gamma = \{\mathbf{x} \in E \mid h(\mathbf{x}) \geq -\gamma\}$ when it begins on the boundary of \mathcal{C}_γ . To make the following argument clear at a high level, I begin with a short proof sketch before proceeding in detail.

Proof sketch: The key tool in proving Theorem 34 is Ville's inequality (4.74). Since $h(\mathbf{x}_k)$, in general, is not a super- or submartingale, I will first construct a nonnegative supermartingale, $W_k \triangleq W(\mathbf{x}_k, k)$, by scaling and shifting $h(\mathbf{x}_k)$. I can then apply Ville's inequality (4.74) to bound the probability of W_k going above any $\lambda > 0$. Next I find a particular value of λ , denoted λ^* , such that:

$$\max_{k \in \{0, \dots, K\}} W_k \leq \lambda^* \implies \min_{k \in \{0, \dots, K\}} h(\mathbf{x}_k) \geq -\gamma. \quad (4.84)$$

Intuitively, this means that any sequence W_k that remains below λ^* ensures that the corresponding sequence $h(\mathbf{x}_k)$ remains (safe) above $-\gamma$. This allows bounding the K -step exit probability P_u of the original process $h(\mathbf{x}_k)$ with the probability that W_k will rise above λ^* :

$$P_u \leq \mathbb{P} \left\{ \max_{k \in \{0, \dots, K\}} W_k > \lambda^* \right\} \leq \frac{\mathbb{E}[W_0]}{\lambda^*} = \frac{W_0}{\lambda^*}, \quad (4.85)$$

where the last equality will follow as it is assumed \mathbf{x}_0 is known *a priori*. Particular choices of W and λ^* will yield the bounds stated in the theorem, completing the proof.

Proof. **Proof: Constructing a Nonnegative Supermartingale**

I will begin by constructing a nonnegative supermartingale, allowing the use of Ville's inequality. To construct this supermartingale, first note that by rearranging terms in the inequality in (4.81), the process $M - h(\mathbf{x}_k)$ resembles a supermartingale:

$$\begin{aligned} \mathbb{E}[M - h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k] &\leq \alpha(M - h(\mathbf{x}_k)) + M(1 - \alpha) - \delta, \\ &\triangleq \alpha(M - h(\mathbf{x}_k)) + \varphi, \end{aligned} \quad (4.86)$$

but with a scaling α and additive term $\varphi \triangleq M(1 - \alpha) - \delta$ that makes $\mathbb{E}[M - h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k] \not\leq M - h(\mathbf{x}_k)$ in general. To remove the effects of α

and φ , consider the function $W : \mathbb{R}^n \times \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}$ defined as:

$$W(\mathbf{x}_k, k) \triangleq \underbrace{(M - h(\mathbf{x}_k))\theta^k}_{\text{negate and scale}} - \underbrace{\varphi \sum_{i=1}^k \theta^i}_{\text{cancel } \varphi} + \underbrace{\varphi \sum_{i=1}^K \theta^i}_{\text{ensure } W \geq 0}, \quad (4.87)$$

where $\theta \in [1, \infty)$ will be used to cancel the effect of α , but is left as a free variable that I will later use to tighten the bound on P_u . Denoting $W_k \triangleq W(\mathbf{x}_k, k)$, I now verify W_k is a nonnegative supermartingale. I first show that $W_k \geq 0$ for all $k \in \{0, \dots, K\}$. Combining the two sums in (4.87) yields:

$$W_k = (M - h(\mathbf{x}_k))\theta^k + \varphi \sum_{i=k+1}^K \theta^i, \quad (4.88)$$

which is nonnegative as $h(\mathbf{x}) \leq M$ for all $\mathbf{x} \in E$, $\theta \geq 1$, and $\varphi \geq 0$ since $\delta \leq M(1-\alpha)$ by assumption. I now show that W_k satisfies the supermartingale inequality (4.73):

$$\mathbb{E}[W_{k+1} \mid \mathbf{x}_{0:k}] = \mathbb{E}[W_{k+1} \mid \mathbf{x}_k], \quad (4.89)$$

$$= (M - \mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k])\theta^{k+1} + \varphi \sum_{i=k+2}^K \theta^i, \quad (4.90)$$

$$\leq (M - \alpha h(\mathbf{x}_k) - \delta)\theta^{k+1} + \varphi \sum_{i=k+2}^K \theta^i, \quad (4.91)$$

$$\begin{aligned} &= \alpha\theta(M - h(\mathbf{x}_k))\theta^k + \theta^{k+1} \underbrace{((1-\alpha)M - \delta)}_{=\varphi} + \varphi \sum_{i=k+2}^K \theta^i, \\ &= \underbrace{\alpha\theta}_{\text{req. } \leq 1} (M - h(\mathbf{x}_k))\theta^k + \varphi \sum_{i=k+1}^K \theta^i \leq W_k, \end{aligned} \quad (4.92)$$

where (4.89) is due to the Markovian nature of system (4.77), (4.90) comes from using (4.88) to write W_{k+1} , (4.91) follows from (4.81), and (4.92) follows from the preceding line using the definition of φ and assuming the further requirement that $\theta \leq \frac{1}{\alpha}$. Thus, I have shown that W_k is a nonnegative supermartingale.

Proof: Bounding the Exit Probability via Ville's inequality

Since W_k is a nonnegative supermartingale, Ville's inequality can be applied to establish:

$$\mathbb{P} \left\{ \max_{k \in \{0, \dots, K\}} W_k > \lambda \right\} \leq \frac{\mathbb{E}[W_0]}{\lambda} = \frac{W_0}{\lambda}. \quad (4.93)$$

for all $\lambda \in \mathbb{R}_{>0}$. To relate this bound to the K -step exit probability P_u , I seek a value of λ , denoted λ^* , such that:

$$\max_{k \in \{0, \dots, K\}} W_k \leq \lambda^* \implies \min_{k \in \{0, \dots, K\}} h(\mathbf{x}_k) \geq -\gamma. \quad (4.94)$$

In short, I will choose a value of λ^* such that all trajectories of W_k that remain below λ^* must also have $h_k \geq -\gamma$. To this end, I use the geometric series identity⁴ $\sum_{i=1}^k \theta^{i-1} = \frac{1-\theta^k}{1-\theta}$ to rewrite W_k as:

$$W_k = (M - h(\mathbf{x}_k))\theta^k + \varphi\theta \frac{\theta^K - \theta^k}{\theta - 1}. \quad (4.95)$$

Define:

$$\lambda_k = \left(\gamma + M - \frac{\varphi\theta}{\theta - 1} \right) \theta^k + \frac{\varphi\theta}{\theta - 1} \theta^K > 0, \quad (4.96)$$

which, intuitively, applies the same time-varying scaling and shift to a constant, $-\gamma$, that was applied to $h(\mathbf{x}_k)$ to yield W_k (4.95). I then choose:

$$\lambda^* \triangleq \min_{k \in \{0, \dots, K\}} \lambda_k. \quad (4.97)$$

Since I assume $\max_{k \in \{0, \dots, K\}} W_k \leq \lambda^*$, we can write, for all $k \in \{0, \dots, K\}$:

$$0 \geq W_k - \lambda^* \geq W_k - \lambda_k = (-\gamma - h_k)\theta^k. \quad (4.98)$$

Since $\theta > 1$, this implies that $-\gamma - h_k \leq 0$ for all $k \in \{0, \dots, K\}$, and thus $\min_{k \in \{0, \dots, K\}} h(\mathbf{x}_k) \geq -\gamma$, as needed.

Proof: Choosing θ to Minimize the Ville's Inequality

Since the supermartingale W_k includes a free parameter $\theta \in (1, \frac{1}{\alpha}]$, I will choose a value of θ in this interval which provide the tightest bound on P_u .

Case 1: Consider the first case where $\delta < -\gamma(1 - \alpha)$, implying $\varphi > (M + \gamma)(1 - \alpha)$. In this case $\frac{1}{\alpha} < \frac{M+\gamma}{M+\gamma-\varphi}$ and thus all of the allowable choices of $\theta \in (1, \frac{1}{\alpha})$ are such that $\theta < \frac{M+\gamma}{M+\gamma-\varphi}$. Denoting k^* such that $\lambda^* = \lambda_{k^*}$, we have that:

$$\lambda^* = \underbrace{\left(\gamma + M - \frac{\varphi\theta}{\theta - 1} \right)}_{\leq 0} \theta^{k^*} + \frac{\varphi\theta}{\theta - 1} \theta^K. \quad (4.99)$$

⁴At $\theta = 1$, the fraction $\frac{1-\theta^k}{1-\theta}$ is not well-defined. However, the proof can be carried out using the summation notation. In this case $\lambda^* = M + \gamma$, and (4.93) yields $P_u \leq 1 - \frac{h(\mathbf{x}_0) + \gamma - \varphi K}{M + \gamma}$.

Thus, we know $\min_{k \in \{0, \dots, K\}} \lambda_k$ occurs at $k^* = K$ and so:

$$P_u \leq \frac{W_0}{\lambda^*} = \frac{M - h(\mathbf{x}_0) + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}{(M + \gamma)\theta^K}. \quad (4.100)$$

Since this bound is a decreasing function of θ (as shown in Lemma 3 following this proof), I choose the largest allowable value $\theta^* = \frac{1}{\alpha}$ to achieve the bound:

$$P_u \leq \frac{W_0}{\lambda^*} = \frac{M - h(\mathbf{x}_0) + \frac{\varphi}{1-\alpha} (\alpha^{-K} - 1)}{(M + \gamma)\alpha^{-K}}, \quad (4.101)$$

$$= \left(\frac{M - h(\mathbf{x}_0)}{M + \gamma} \right) \alpha^K + \frac{M(1 - \alpha) - \delta}{M + \gamma} \sum_{i=1}^K \alpha^{i-1}, \quad (4.102)$$

where I again use the geometric series identity.

Case 2: Now consider the second case where $\delta \geq -\gamma(1 - \alpha)$, so $\varphi \leq (M + \gamma)(1 - \alpha)$, which implies that the set $[\frac{M+\gamma}{M+\gamma-\varphi}, \frac{1}{\alpha}]$ is nonempty. Choosing a value of θ in this set ensures that:

$$\lambda^* = \underbrace{\left(\gamma + M - \frac{\varphi\theta}{\theta-1} \right)}_{\geq 0} \theta^{k^*} + \frac{\varphi\theta}{\theta-1} \theta^K. \quad (4.103)$$

Thus $\min_{k \in \{0, \dots, K\}} \lambda_k$ occurs at $k^* = 0$ and:

$$P_u \leq \frac{W_0}{\lambda} = \frac{(M - h(\mathbf{x}_0)) + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}{(M + \gamma) + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}, \quad (4.104)$$

$$= 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}. \quad (4.105)$$

Since this bound is increasing in θ (as shown in Lemma 4 following this proof), I choose $\theta^* = \frac{M+\gamma}{M+\gamma-\varphi}$ to achieve the bound:

$$P_u \leq 1 - \left(\frac{h(\mathbf{x}_0) + \gamma}{M + \gamma} \right) \left(\frac{M\alpha + \gamma + \delta}{M + \gamma} \right)^K. \quad (4.106)$$

If, alternatively, I choose $\theta \in \left(1, \frac{M+\gamma}{M+\gamma-\varphi} \right]$, then the inequality in (4.99) holds, $k^* = K$, and the bound is decreasing in θ as in Case 1. Evaluating this bound for the minimizing value $\theta^* = \frac{M+\gamma}{M+\gamma-\varphi}$ again yields:

$$P_u \leq \frac{M - h(\mathbf{x}_0) + (M + \gamma)(\theta^K - 1)}{(M + \gamma)\theta^K}, \quad (4.107)$$

$$= 1 - \left(\frac{h(\mathbf{x}_0) + \gamma}{M + \gamma} \right) \left(\frac{M\alpha + \gamma + \delta}{M + \gamma} \right)^K. \quad (4.108)$$

□

I now state and prove the lemmas used in the preceding proof to prove optimality of the bound in Theorem 34 Cases 1 and 2. These lemmas were originally stated without proof in [59].

Lemma 3. For $M \in \mathbb{R}_{>0}$, $\gamma, \varphi \in \mathbb{R}_{\geq 0}$, $h(\mathbf{x}_0) \in [-\gamma, M]$, and $K \in \mathbb{Z}_{\geq 1}$, the function $\Psi_1 : (1, \infty) \rightarrow \mathbb{R}$ defined as:

$$\Psi_1(\theta) = \frac{M - h(\mathbf{x}_0) + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}{(M + \gamma)\theta^K}, \quad (4.109)$$

is monotonically decreasing.

Proof. The geometric series identity yields:

$$\Psi_1(\theta) = \frac{M - h(\mathbf{x}_0)}{M + \gamma} \theta^{-K} + \frac{\varphi}{(M + \gamma)} \sum_{i=1}^K \theta^{i-K}, \quad (4.110)$$

$$\begin{aligned} \frac{d\Psi_1}{d\theta} &= -\frac{M - h(\mathbf{x}_0)}{M + \gamma} K \theta^{-K-1} - \varphi \sum_{i=1}^K \frac{(K - i) \theta^{i-K-1}}{M + \gamma}, \\ &\leq 0, \end{aligned} \quad (4.111)$$

for all $\theta \in (1, \infty)$. □

Lemma 4. For $M \in \mathbb{R}_{>0}$, $\gamma, \varphi \in \mathbb{R}_{\geq 0}$, $h(\mathbf{x}_0) \in [-\gamma, M]$, and $K \in \mathbb{Z}_{\geq 1}$, the function $\Psi_2 : (1, \infty) \rightarrow \mathbb{R}$ defined as:

$$\Psi_2(\theta) = 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \frac{\varphi\theta}{\theta-1} (\theta^K - 1)}, \quad (4.112)$$

is monotonically increasing.

Proof. The geometric series identity yields:

$$\Psi_2(\theta) = 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \varphi \sum_{i=1}^K \theta^i}, \quad (4.113)$$

$$\frac{d\Psi_2}{d\theta} = \frac{(h(\mathbf{x}_0) + \gamma) \left(\varphi \sum_{i=1}^K i \theta^{i-1} \right)}{\left(M + \gamma + \varphi \sum_{i=1}^K \theta^i \right)^2}, \quad (4.114)$$

$$\geq 0, \quad (4.115)$$

for all $\theta \in (1, \infty)$. □

Practical Considerations in Enforcing Safety of Stochastic Systems

Theorem 34 allows reasoning about the finite-time safety of systems governed by DT-BFs. To utilize the results of this theorem in a control setting, I aim to use DT-CBFs to develop controllers which enforce the expectation condition:

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k) \mid \mathbf{x}_k] \geq \alpha h(\mathbf{x}_k). \quad (4.116)$$

Like the DT-CBF-OP controller, I seek to enforce this constraint using an optimization-based controller that enforces safety while achieving pointwise minimal deviation from a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ in the form of an Expectation-based DT-CBF (ED) controller:

$$\begin{aligned} \mathbf{k}_{\text{ED}}(\mathbf{x}_k) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k)\|^2 & (\text{ED}) \\ \text{s.t. } & \mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}) + \mathbf{d}_k) \mid \mathbf{x}_k] \geq \alpha h(\mathbf{x}_k). \end{aligned}$$

The expectation in (ED) adds complexity that is not generally considered in the application of deterministic DT-CBFs. Typically, CBF-based controllers solve “certainty-equivalent” optimization programs, like this Certainty-Equivalent DT-CBF (CED) controller, that replace the expected barrier value $\mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k]$ with the barrier at the expected next state, $h(\mathbb{E}[\mathbf{x}_{k+1} \mid \mathbf{x}_k])$:

$$\begin{aligned} \mathbf{k}_{\text{CED}}(\mathbf{x}_k) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k)\|^2 & (\text{CED}) \\ \text{s.t. } & h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}) + \mathbb{E}[\mathbf{d}_k]) \geq \alpha h(\mathbf{x}_k), \end{aligned}$$

where $\mathbb{E}[\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) \mid \mathbf{x}_k] = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k)$ and $\mathbb{E}[\mathbf{d}_k \mid \mathbf{x}_k] = \mathbb{E}[\mathbf{d}_k]$. This constraint is often easier to evaluate than (4.116) since it allows control actions to be selected with respect to the expected disturbance $\mathbb{E}[\mathbf{d}_k]$ without needing to model the disturbance distribution \mathcal{D} . If the disturbance is zero-mean, then this form of the constraint is implicitly enforced by DT-CBF controllers such as those presented in [234], [244]. However, when replacing ED with CED it is important to consider the effect of Jensen’s inequality in Theorem 33.

If the “certainty-equivalent” constraint in CED is strictly concave⁵, then we can apply the results of Theorem 34 directly since Jensen’s inequality tightens the constraint and ensures satisfaction of the expectation condition (4.81). Unfortunately, using such a controller is a non-convex optimization program which

⁵The constraint $h(\mathbf{x}_k + \mathbf{u}) \geq \alpha h(\mathbf{x}_k)$ is concave in \mathbf{u} when h is convex and it is convex in \mathbf{u} when h is concave.

can be impractical to solve. If, instead, the constraint is convex, then CED is a convex program, but does not necessarily enforce the expectation condition (4.81) in Theorem (34) due to the gap introduced by Jensen's inequality. In order to apply the results of Theorem 34 to controllers of the form (CED) with convex constraints, we must first provide a bound on the gap introduced by Jensen's Inequality. In particular, for any concave function $h : E \rightarrow \mathbb{R}$ and random variable $\mathbf{d} \sim \mathcal{D}$, I seek to determine a value $\psi \in \mathbb{R}_{\geq 0}$ such that, for all $\mathbf{x} \in E$ and $\mathbf{u} \in \mathbb{R}^m$:

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{d}) \mid \mathbf{x}] \geq h(\mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbb{E}[\mathbf{d}]) - \psi, \quad (4.117)$$

thus quantifying the gap introduced by Jensen's inequality. A large body of work has studied methods for finding the smallest possible ψ that satisfies (4.117). Here I adapt a result in [248] to achieve a relatively loose, but straightforward bound:

Lemma 5. *Consider a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ that is twice-continuously differentiable on \mathbb{R}^n , concave, and satisfies $\sup_{\mathbf{x} \in \mathbb{R}^n} \|\nabla^2 h(\mathbf{x})\|_2 \leq \lambda_{\max}$ for some $\lambda_{\max} \in \mathbb{R}_{\geq 0}$, and a random variable \mathbf{x} that takes values in \mathbb{R}^n with $\mathbb{E}[\|\mathbf{x}\|] < \infty$ and $\|\text{cov}(\mathbf{x})\| < \infty$. Then we have that:*

$$\mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]) - \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{x})). \quad (4.118)$$

Proof. Consider the function $\eta : \mathbb{R}^n \rightarrow \mathbb{R}$ which is twice-continuously differentiable on E and convex that is defined as $\eta = -h$. The Intermediate Value Theorem [249, Theorem 4.5.1] implies that for all $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$, there exists an $\omega \in [0, 1]$ such that:

$$\eta(\mathbf{z}) = \eta(\mathbf{y}) + \nabla \eta(\mathbf{y})^\top \mathbf{e} + \frac{1}{2} \mathbf{e}^\top \nabla^2 \eta(\mathbf{c}) \mathbf{e}, \quad (4.119)$$

where $\mathbf{e} \triangleq \mathbf{z} - \mathbf{y}$, $\mathbf{c} \triangleq \omega \mathbf{z} + (1 - \omega) \mathbf{y}$, and $\nabla^2 \eta(\mathbf{c})$ is the Hessian of η evaluated at \mathbf{c} . We then have that:

$$\eta(\mathbf{z}) = \eta(\mathbf{y}) + \nabla \eta(\mathbf{y})^\top \mathbf{e} + \frac{1}{2} \text{tr}(\nabla^2 \eta(\mathbf{c}) \mathbf{e} \mathbf{e}^\top), \quad (4.120)$$

$$\leq \eta(\mathbf{y}) + \nabla \eta(\mathbf{y})^\top \mathbf{e} + \frac{1}{2} \|\nabla^2 \eta(\mathbf{c})\|_2 \text{tr}(\mathbf{e} \mathbf{e}^\top), \quad (4.121)$$

$$\leq \eta(\mathbf{y}) + \nabla \eta(\mathbf{y})^\top \mathbf{e} + \frac{\lambda_{\max}}{2} \text{tr}(\mathbf{e} \mathbf{e}^\top), \quad (4.122)$$

where the first inequality is a property of the trace operator for positive semi-definite matrices [250] (and $\nabla^2 \eta(\mathbf{c})$ is positive semi-definite as η is convex),

and the second inequality follows by our definition of λ_{\max} . Let \mathbf{x} be a random variable taking values in \mathbb{R}^n with probability density function $p : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$, and let $\boldsymbol{\mu} \triangleq \mathbb{E}[\mathbf{x}]$. We then have that:

$$\mathbb{E}[\eta(\mathbf{x})] - \eta(\mathbb{E}[\mathbf{x}]) = \int_{\mathbb{R}^n} (\eta(\mathbf{x}) - \eta(\boldsymbol{\mu}))p(\mathbf{x})d\mathbf{x}, \quad (4.123)$$

$$\leq \int_{\mathbb{R}^n} \nabla\eta(\boldsymbol{\mu})^\top \mathbf{e} + \frac{\lambda_{\max}}{2} \text{tr}(\mathbf{e}\mathbf{e}^\top) p(\mathbf{x})d\mathbf{x}, \quad (4.124)$$

$$= \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{x})), \quad (4.125)$$

where $\mathbf{e} = \mathbf{x} - \boldsymbol{\mu}$. Replacing η with $-h$ yields:

$$\mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]) - \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{x})). \quad (4.126)$$

□

I first note that this theorem requires considering \mathbb{R}^n instead of an open set $E \subset \mathbb{R}^n$. This is important for ensuring the the first term in the integral in (4.124) evaluates to 0. This could likely be relaxed with additional assumptions on the nature of the distribution \mathcal{D} . Next, I note that although this value of $\psi = \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{x}))$ is easy to interpret, tighter bounds exist which have less restrictive assumptions than a globally bounded Hessian [247]. Lastly, I note that sampling-based methods could be used to approximately satisfy the constraint (4.117) by estimating ψ empirically.

Next, I present a controller which combines the mean-based control of the ‘‘certainty equivalent’’ (CED) while also accounting for Jensen’s inequality. This Jensen-Enhanced DT-CBF (JED) controller includes an additional control parameter $c_J \geq 0$ to account for Jensen’s inequality:

$$\begin{aligned} \mathbf{k}_{\text{JED}}(\mathbf{x}_k) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k)\|^2 & (\text{JED}) \\ \text{s.t. } & h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbb{E}[\mathbf{d}_k]) - c_J \geq \alpha h(\mathbf{x}_k). \end{aligned}$$

Given this controller and a method for bounding ψ , we can now apply Theorem 34 while accounting for (or analyzing) the effects of Jensen’s inequality on the (JED) controller:

Theorem 35. *Consider the closed-loop system (4.77) and let the function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be twice-continuously differentiable on \mathbb{R}^n , concave, and satisfy $\sup_{\mathbf{x} \in \mathbb{R}^n} h(\mathbf{x}) \leq M$ for $M \in \mathbb{R}_{>0}$ and $\sup_{\mathbf{x} \in \mathbb{R}^n} \|\nabla^2 h(\mathbf{x})\|_2 \leq \lambda_{\max}$ for $\lambda_{\max} \in$*

$\mathbb{R}_{\geq 0}$. Suppose there exists an $\alpha \in (0, 1)$ and a $c_J \in [0, \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d})) + M(1 - \alpha)]$ such that:

$$h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbb{E}[\mathbf{d}]) - c_J \geq \alpha h(\mathbf{x}), \quad (4.127)$$

for all $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{d} \sim \mathcal{D}$. Then we have that:

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbf{d}) \mid \mathbf{x}] \geq \alpha h(\mathbf{x}) + \delta, \quad (4.128)$$

for all $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{d} \sim \mathcal{D}$ and $\delta = c_J - \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d}))$.

Proof. Given $\mathbf{x} \in \mathbb{R}^n$, Lemma 5 ensures that:

$$0 \leq h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbb{E}[\mathbf{d}]) - c_J - \alpha h(\mathbf{x}), \quad (4.129)$$

$$\leq \mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbf{d}) \mid \mathbf{x}] + \psi - c_J - \alpha h(\mathbf{x}), \quad (4.130)$$

where $\psi = \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d}))$. Letting $\delta = c_J - \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d}))$ yields the desired result. \square

Simulation Results

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in conference proceedings, and edited by Andrew Taylor for their appearance in this thesis.

We now consider a variety of simulation examples that highlight the key features of the proposed approach.

Linear 1D System

Here we analyze the finite-time safety probability bounds by considering the case of unbounded i.i.d. disturbances $d_k \sim \mathcal{N}(0, 1)$ (normal distribution with zero mean and variance of one) for the system:

$$x_{k+1} = x_k + 2 + u_k + \sigma d_k, \quad (4.131)$$

with a function $h : \mathbb{R} \rightarrow \mathbb{R}$ defined as $h = 1 - x^2$ yielding the set $\mathcal{C} = \{x \mid 1 - x^2 \geq 0\}$. Observe that h is upper bounded by 1 on \mathbb{R}^n and as a constant Hessian of value -2 . The Jensen gap for this system and DT-CBF is bounded by $\psi = \sigma^2$. For simulation, we employ the JED controller with $c_J = \sigma^2$, $\alpha = 1 - \sigma^2$, and nominal controller $\mathbf{k}_{\text{nom}}(\mathbf{x}_k) = 0$. Figure 4.21 shows the results of 500 one second long trials run with a variety of $\sigma \in [0, 0.2]$ and also displays how the bound on P_u decreases as γ increases.

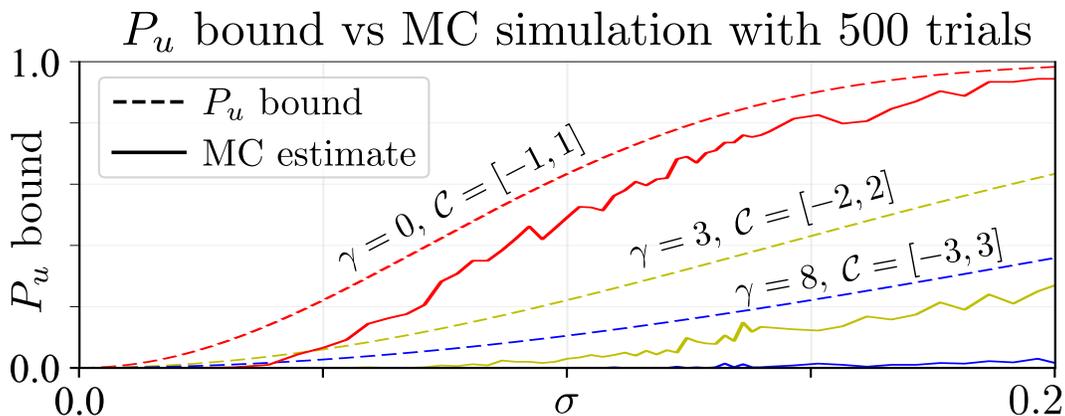


Figure 4.21. The dashed lines represent the theoretical probability bounds for the system as in Theorem 34. The solid lines represent the Monte Carlo (MC) estimated P_u across 500 experiments.

Simple Pendulum

Next we consider an inverted pendulum about its upright equilibrium point with the DT dynamics:

$$\begin{bmatrix} \theta_{k+1} \\ \dot{\theta}_{k+1} \end{bmatrix} = \begin{bmatrix} \theta_k + \Delta t \dot{\theta}_k \\ \dot{\theta}_k + \Delta t \sin(\theta_k) \end{bmatrix} + \begin{bmatrix} 0 \\ \Delta t \end{bmatrix} \mathbf{u} + \mathbf{d}_k, \quad (4.132)$$

with time step $\Delta_t = 0.01$, i.i.d disturbances⁶ $\mathbf{d}_k \sim \mathcal{N}(\mathbf{0}_2, \text{Diag}([0.005^2, 0.025^2]))$, and set \mathcal{C} defined as:

$$\mathcal{C} = \left\{ \mathbf{x} \in \mathbb{R}^n \mid \underbrace{1 - \frac{6^2}{\pi^2} \mathbf{x}^\top \begin{bmatrix} 1 & 3^{-\frac{1}{2}} \\ 3^{-\frac{1}{2}} & 1 \end{bmatrix} \mathbf{x}}_{h_{\text{pend}}(\mathbf{x})} \geq 0 \right\}, \quad (4.133)$$

which is constructed using the CTLE as in [115] and for which $|\theta| \leq \pi/6$ for all $\mathbf{x} \in \mathcal{C}$. Figure 4.22 shows the results of 500 one second long trials for each $\mathbf{x}_0 \in \mathcal{C}$ using the JED controller with parameters $\alpha = 1 - \psi$, $c_J = \psi$, where $\psi = \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d}_k))$. This figure highlights the influence of \mathbf{x}_0 and shows how the bound on P_u increases as $h(\mathbf{x}_0)$ decreases.

Quadruped

Finally, we consider controlling a simulated quadrupedal robot locomoting along a narrow path. The simulation is based on a Unitree A1 robot as shown

⁶Diag: $\mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ generates a square diagonal matrix with its argument along the main diagonal.

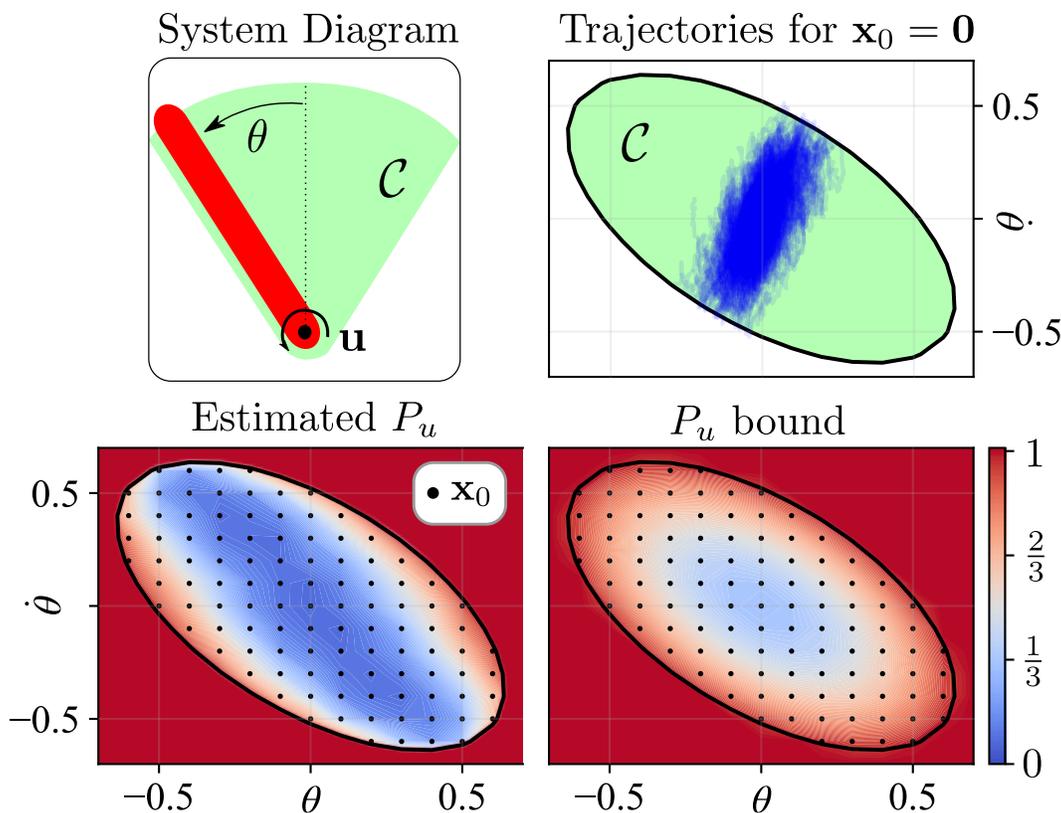


Figure 4.22. (Top Left) System diagram of the inverted pendulum. (Top Right) 500 one second long example trajectories starting at $\mathbf{x}_0 = \mathbf{0}$. (Bottom Left) Monte Carlo estimates of P_u for $\gamma = 0$ using 500 one second long trials for each initial conditions represented by a black dot. (Bottom Right) The (conservative) theoretical bounds on P_u from Theorem 34.

in Figure (4.23) which has 18 degrees of freedom and 12 actuators. An inverse dynamics-based (ID) QP controller designed using concepts in [251] and implemented at 1 [kHz] is used to track stable walking gaits with variable planar velocities and angle rate using the motion primitive framework presented in [252]. The full 18 degree of freedom quadruped dynamics are simulated at 1 [kHz], but the methodology in [253] is used by considering the following simplified discrete-time single-integrator system for DT-CBF-based control:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_k^x \\ v_k^y \\ \theta_k \end{bmatrix} + \mathbf{d}_k, \quad (4.134)$$

where $\mathbf{x}_k = [x, y, \theta]^\top$. In order to represent the error caused by uncertain terrain, zero mean Gaussian disturbances are added to the quadruped's (x, y) body position and velocity with variances of 2.25×10^6 and 0.01 respectively. This random noise along with the dynamics-mismatch between the full-order

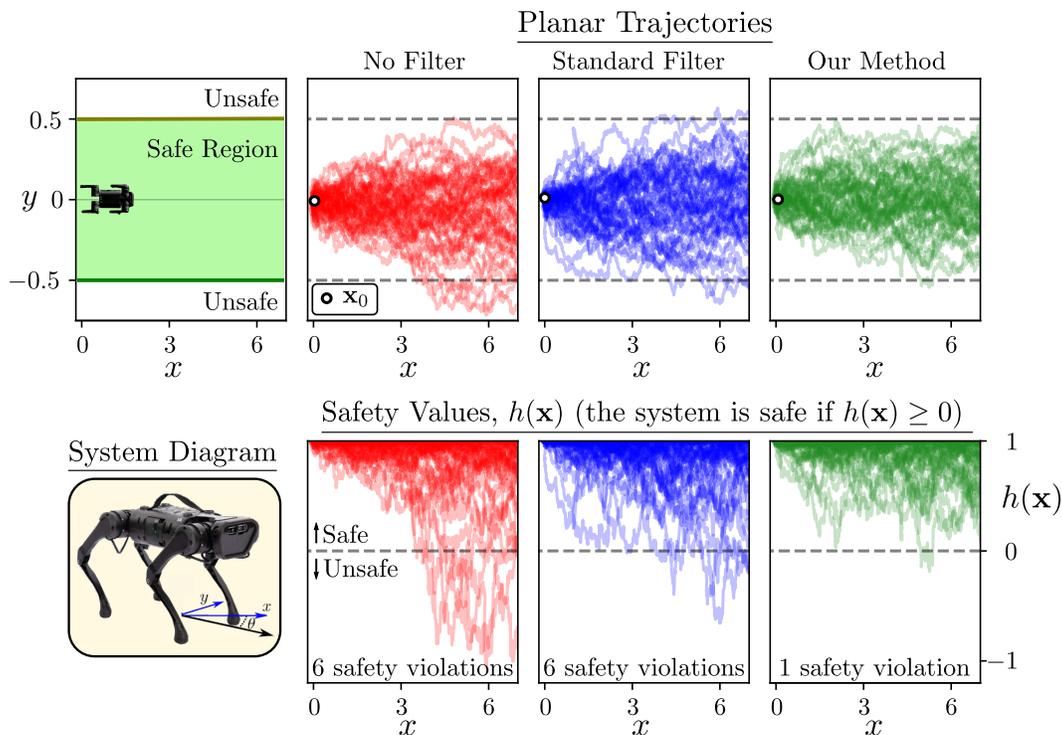


Figure 4.23. Safety of a simulated quadrupedal robot locomoting on a narrow path for a variety of controllers. (Top Left) The safe region that the quadruped is allowed to traverse. (Bottom Left) A system diagram depicting the reduced-order model states of the quadruped $[x \ y \ \theta]^\top$. (Top Right) 50 trajectories for 3 controllers: one without any knowledge of safety (\mathbf{k}_{nom}), one with the standard (DT-CBF-OP) controller, and finally the proposed method which accounts for stochasticity (JED). (Bottom Right) Plots of $h(\mathbf{x})$, a scalar value representing safety. The system is safe (i.e., in the green safe region) if $h(\mathbf{x}) \geq 0$.

quadrupedal dynamics and (4.134) is modeled as an i.i.d. random process \mathbf{d}_k .

The quadruped is commanded to stand and then traverse a 7 [m] path that is 1 [m] meter wide, with the set $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid 0.5^2 - y^2 \geq 0\}$. For this simulation, three controllers are compared: a simple nominal controller $\mathbf{k}_{\text{nom}}(\mathbf{x}) = [0.2, 0, -\theta]^\top$ with no understanding of safety, the DT-CBF-OP controller with $\alpha = 0.99$, and the proposed JED controller with $\alpha = 0.99$ and $c_J = \psi$ using the mean and covariance estimates, $\mathbb{E}[\mathbf{d}_k] \approx [-0.0132, -0.0034, -0.0002]^\top$ and $\text{tr}(\text{cov}(\mathbf{d}_k)) \approx \psi = 0.000548$, which were generated using 15 minutes of walking data controlled by \mathbf{k}_{nom} . The results of 50 trials for each controller can be seen in Figure 4.23. As expected, \mathbf{k}_{nom} generated the largest safety violations and JED the smallest and fewest safety violations.

Conclusion

In conclusion, in this section I have developed a framework for safety-critical control of discrete-time systems with stochastic disturbances using CBFs and convex optimization. I begin by reviewing discrete-time nonlinear dynamics and corresponding variants of BFs and CBFs for such systems, followed by a review of tools from probability theory that are useful for studying systems with stochastic disturbances. I then consider nonlinear discrete-time systems with stochastic disturbances and propose a notion of finite-time safety akin to ISSf. I show how discrete-time BFs certify this finite-time safety property, and then show how convexity properties of discrete-time CBFs can be used to synthesize controllers that yield discrete-time BFs through convex optimization. I conclude with a set of simulations on various systems.

4.6 Conclusion

In this chapter I have presented a collection of results on designing controllers that are performant while providing robustness to disturbances. In Section 4.2 I looked at some of the limitations of controllers designed through ISSf-CBFs as originally posed in [58] and presented in Section 2.9. This investigation focuses on balancing meaningful theoretical guarantees with conservative behavior. Drawing motivation from this analysis, in Section 4.3 I present a novel modification of ISSf in the form of Tunable Input-to-State Safety (TISSf) that permits controllers that provide not only meaningful theoretical guarantees, but also strong closed-loop performance. I then explore the capabilities of TISSf in the context of a simple inverted pendulum. In Section 4.4, I consider the significantly more challenging problem of designing a safety-critical controller for a connected automated semi-trailer truck. A simplified model for this system neglects important actuator dynamics in an experimental system that lead to the failure of standard CBF-based controllers to render the system safe. Capturing the effect of these dynamics as a matched disturbance, I design a host of controllers utilizing ISSf-CBFs and TISSf-CBFs, and demonstrate the ability of TISSf to produce improvements in theoretical guarantees and closed-loop performance over ISSf-CBF based controllers, including on the experimental system. In Section 4.5, I consider safety-critical control for discrete-time systems subject to stochastic disturbances. I review tools for discrete-time safety critical control and probabilistic tools for studying stochastic disturbances. I then propose a notion of finite-time safety that

draws inspiration from the set-expansion paradigm posed by ISSf, and show how discrete-time barrier functions certify this safety property. I conclude with studying how convexity properties of a discrete-time CBF play an important role in permitting the design of controllers that are robust to stochastic disturbances through convex optimization, and demonstrate such controllers on a collection of simulation environments.

I think there are a number of meaningful directions of future work stemming from these results. In the context of ISSf and TISSf, TISSf makes significant strides in achieving theoretical guarantees that more accurately describe the closed-loop behavior of a system operating under a controller that is robust to disturbances. Still, I believe that there still remains the potential to more closely match theoretical safety guarantees with the closed-loop behavior that is actually observed, thereby ensuring that designs focused on pursuing quantitative theoretical guarantees don't produce excessively conservative controllers. One particular observation is that both ISSf and TISSf rely on the ∞ -norm of the disturbance. This contrasts the fact that disturbance signals impact the evolution of the state by being integrated over a period of time. If the disturbance is large but present for a small period of time, it will yield a large ∞ -norm and have potentially little impact on the evolution of the system (this is why the essential supremum is used in the first place!). This is not a new observation in terms of studying the impact of disturbances on systems, which has been well characterized for linear systems [99] and codified for stability through the notion of \mathcal{L} -stability [100, Chapter 5]. I believe that developing notions of ISSf that are based on \mathcal{L}^p -norms may permit stronger theoretical guarantees for a system exposed to a specific class of disturbances. An another idea that would likely share connections with this line of thinking is that of *integral Input-to-State Stability* [254].

In the context of safety-critical control in the presence of stochastic disturbances, I also think there are a number of future work directions that would directly improve the results in Section 4.5. First, there is likely significant room for improving the quality of the bounds achieved through this analysis through the intelligent choice of CBFs, or synthesis of CBFs that optimize these bounds (similar to the work in [61]). Other approaches for improving the bound could look at making stronger claims on the higher-order moments of the disturbance distributions and making use of stronger tools than Ville's

inequality. A potentially better avenue would consider the fact that the nominal controller \mathbf{k}_{nom} is rarely attempting to make a system unsafe, such that the required CBF inequality constraint is rarely active. I hypothesize this may require studying the state distribution, similar to the line of work in [68]. Lastly, I think there is a natural opportunity to consider the impact of noisy measurements on safety-critical control through this framework. While the work in the next chapter makes some initial steps into measurement-robust safety-critical control, it takes a deterministic model of measurement error, and thus could be expanded to capture the fact that most measurement errors are better captured through a stochastic model.

MEASUREMENT-ROBUST SAFETY-CRITICAL CONTROL

In this chapter I will present results from my work developing methods for safety-critical control that are robust to measurement errors. Feedback control is predicated on the ability to observe some portion (or all) of the system state, which is then used as a state estimate in subsequent control computations. Forming a state estimate is generally done through a combination of direct measurements of states and observers such as Kalman filters which produce estimates of states that are not directly measured through an internal model of the system dynamics [255]. Direct measurements are almost always subject to some amount of noise, while observers can be subjected to errors in an internal model, and may need to use a simplified internal model to permit online computation (such as the linearization step that happens in an extended Kalman filter [255]). The recent development of vision-based measurement systems for robotic systems offers the promise of operating in highly unstructured environments, but also comes with its own challenges. In particular, vision-based measurement systems, and especially those built using learning methods, must convert a high-dimensional (and noisy) measurement (such as an image) down into a low-dimensional estimate of the system state (such as position, orientation, and velocities), which is an often imperfect process. Thus, the resulting state estimate is often not exactly equal to the true system state, introducing a mismatch between the state being used in the theory-driven control computations and the true system state. Consequently, this can not only lead to a failure to achieve theoretical guarantees, but poor closed-loop system behavior. It is worth noting that this is an issue that will face the practical deployment of all of the methods developed in Chapters 2, 3, 4, and 6 of this thesis. Thus, a thorough characterization of robust methods for safety-critical control synthesis should consider the impacts of errors between a state estimate and the true system state.

At the same time, building systems that are increasingly robust to different challenges comes with its own challenge of over-conservative closed-loop behavior. As studied in Chapter 4, if a controller is built to provide rigorous theoretical safety guarantees in the presence of potentially large disturbances,

and worst-case disturbances are not experienced, it can lead to unnecessarily poor performance. If a system must be robust to both worst-case disturbances, and worst-case measurement errors, it is reasonable to expect that this will only lead to greater degrees of conservativeness. In practice, the designer of a control system is often left performing the tedious task of tuning various robustness parameters to try and balance the various types of robustness until they get an acceptable level of closed-loop performance. This can be a time consuming task, as often the designer of the control system doesn't have a complete understanding of how the various parameters interact to dictate overall robustness and closed-loop performance. One thing the control designer is often capable of easily providing feedback on is whether or not they think that the system is too conservative, or not safe enough. It is this observation that will motivate the contents of Section 5.4.

My work in this chapter focuses on developing CBF-based control designs that are robust to measurement errors, as well as an algorithm for tuning robustness parameters of a CBF-based controller through *preference-based learning (PBL)*. In Section 5.1 I discuss related work in the topics of safety-critical control techniques that are robust to measurement errors as well as the use of preference-based learning for tuning controller parameters. In Section 5.2 I will present the notion of a Measurement-Robust Control Barrier Function (MR-CBF). I begin by discussing how measurement error appears when trying to certify safety of a system. I will then present MR-CBFs and show how they can be used to synthesize controllers that are robust to measurement errors using convex optimization, resulting in a controller specified as a second-order cone program (SOCP). Following this, I show how this controller can be converted to a standard form SOCP and establish some local Lipschitz properties of a relaxed form of the controller. In Section 5.3 I will present the integration of MR-CBFs with backup set CBFs as formulated in [256]. I review backup set CBFs, present their integration with MR-CBFs, and demonstrate the ability to keep a Segway system using vision-based measurements safe in the presence of measurement errors in both simulation and experimentally. Lastly, in Section 5.4, I will present work on tuning robustness parameters in a CBF-based controller through the use of PBL. I will review PBL, frame the tuning of robustness parameters as a PBL problem, and demonstrate the proposed approach in simulation and experimentally. Key contributions of this work are described at the beginning of each respective section.

5.1 Related Work

Measurement-Robust and Vision-Based Control Barrier Functions

The consideration of measurement errors and their impact on CBF-based controllers has been a topic of recent interest. The work in [50] was the first to consider the impact of measurement errors by incorporating the covariance of an unscented Kalman filter into the CBF inequality as an additive term. In [66] measurements are modeled using stochastic differential equations, and a CBF is robustified against potential estimation error, while in [73] a bounded set of measurement errors is considered by requiring that a controller satisfy a restricted sub-tangentiality condition on the boundary of a set \mathcal{C} that is to be kept forward invariant. The work in [74], [93] utilizes interval arithmetic to consider all possible evolutions of the value of a CBF in the presence of measurement errors, and chooses control inputs to require that all possible evolutions satisfy a safety constraint. In [257], measurement errors are treated as errors in the specification of safety via an uncertain CBF, and made subsequently robust. The work in [258] forms barriers based on locally sensed constraints drawn from online measurements. Additional work considering observers for forming state estimates has used the notion of ISSf explored in Section 2.7 to quantify the impact of observation errors on safety [259], used function approximation techniques to capture estimation error and synthesize robust controllers [260], and provided a framework for understanding safety with general observers that achieve ISS or bounded error properties [261].

Vision-based control in robotics has also seen significant development [262]–[264] (these reflect only a tiny fraction of the extensive body of recent work). A collection of this work has focused primarily on the problem of achieving theoretical stability or safety guarantees in the presence of imperfect perception systems. The work in [265] considers linear systems with imperfect perception-based measurements and achieves robust stability guarantees through the perspective of System Level Synthesis (SLS), while the work in [266] considers certainty-equivalent perception-based control for linear systems. The work in [267], [268] focuses on discrete-time system with perception-based measurements that are subject to adversarial attacks, and provides rigorous theoretical guarantees utilizing Lyapunov-based arguments. The use of CBFs with vision-based systems (not including my work and its extensions, discussed below) has considered producing barrier functions for LiDAR-based vision systems subject to sensor faults and attacks [269], or considered the use of Neural Radiance

Fields (NeRFs) in conjunction with CBFs [270].

My work in [184], presented in Section 5.2 proposes the notion of a Measurement-Robust Control Barrier Function (MR-CBF). This definition considers how error in a state estimate propagates through nonlinear control-affine dynamics, revealing specific structure on how measurement error reacts with the input to the system when determining the time derivative of a MR-CBF. This leads to a convex optimization-based controller specified as a second-order cone program (SOCP), which I evaluate the local Lipschitz continuity of. I demonstrate the proposed approach in simulation on a planar Segway system. In Section 5.3, which features my work in [23], I explore how MR-CBFs can be unified with backup set CBFs [256], leading to the first experimental demonstration of MR-CBFs. Extensions of this work have looked at performing self-supervised learning using stereo vision and MR-CBFs [185] and end-to-end imitation learning with MR-CBFs [174], as well as ways in which perception systems can be improved to ensure feasibility of CBF-based inequalities [271].

Preference-Based Learning

Preference-based learning (PBL) provides an approach for searching complex parameter spaces via subjective feedback, without an explicitly defined reward function [75]. The main advantage of online PBL is its ability to interactively infer a user’s latent utility function using only subjective feedback such as pairwise preferences and ordinal labels [272], [273]. It has seen application on a wide-range of problems including recommender systems [274], search engines [275], trajectory planning [276]–[278], and spinal cord stimulation [279]. A particular area of development relevant for control systems is safe PBL, which seeks to ensure the recommendations to users do not violate some prespecified safety criterion [135], [279], [280]. In the context of control systems, there has recently been significant effort in developing PBL approaches for tuning existing tools built from control theory. One such example is the selection of walking gaits preferred by a patient that is utilizing a lower-body exoskeleton, which has been explored extensively in [76]–[78]. These ideas were further extended to autonomous bipedal systems to tune walking gait parameters for the AMBER-3M planar biped explored in Section 3.7 [79], or used to tune the parameters of a (CLF-QP) controller for the AMBER-3M planar biped and Cassie 3D biped [80]. A toolbox and comprehensive survey of the use of PBL for robotic systems can be found in [281].

My work in [23], presented in Section 5.4, draws inspiration from the ideas in [80] to tune parameters of a convex optimization-based controller. While the work in [80] focused specifically on learning the parameters of a (CLF-QP) focused on stabilizing a bipedal system, I will focus on tuning the parameters of a CBF-based controller that is robust to both disturbances and measurement errors. In particular, this controller will unify the reduced-order model tools for CBFs developed in [253] with ISSf-CBFs presented in Section 2.9 and MR-CBFs developed in Section 5.2 to design a robust safety-critical controller for a quadrupedal robotic system. This controller is initially extremely conservative, with almost negligible performance. PBL is deployed to adjust the parameters of this controller until a balance can be struck between robustness and performance in both simulation and experimentally.

5.2 Measurement-Robust Control Barrier Functions

In this section I will present work on developing Measurement-Robust Control Barrier Functions (MR-CBFs) for safety-critical control in the presence of measurement errors. I begin by describing the setting of measurement error that I consider. Following this, I present novel definitions of Measurement-Robust Barrier Functions (MR-BFs), show how an MR-BF leads to safety in the presence of measurement errors. I then extend the notion of an MR-BF to an MR-CBF, and demonstrate its application on a Segway system in simulation. Lastly, I show how the convex optimization problem defining an MR-CBF-based controller can be cast as an SOCP in standard form, and then establish local Lipschitz properties of the controller under a mild relaxation.

The contributions of this section are as follows:

- The definition of Measurement-Robust Control Barrier Functions (MR-CBFs) as a tool for achieving safety-critical control in the presence of measurement errors through convex optimization.
- A detailed analysis of the regularity properties of the resulting SOCP-based controllers synthesized using MR-CBFs.

The text for this section is adapted from:

S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, “Guaranteeing safety of learned perception modules via measurement-robust

control barrier functions,” in *Proc. Conf. on Robotics Learning (CoRL)*, Cambridge, MA, USA, 2020.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, and writing of the article.

Measurement Error

In many applications, the state \mathbf{x} is not directly available to the controller, but rather a corresponding measurement $\mathbf{y} \in \mathbb{R}^k$ which satisfies the relationship:

$$\mathbf{y} = \mathbf{p}(\mathbf{x}), \quad (5.1)$$

for a function $\mathbf{p} : E \rightarrow \mathbb{R}^k$ that is assumed to be locally Lipschitz continuous on E . In this work I assume the relationship between the measurement and the true state is deterministic, such that the evolution of the system can be described by an ordinary differential equation rather than a stochastic differential equation. I further assume that there exists a function $\mathbf{q} : \mathbf{p}(E) \rightarrow E$ that is locally Lipschitz continuous on $\mathbf{p}(E)$ and for any $\mathbf{x} \in E$ with a corresponding measurement $\mathbf{y} \in \mathbf{p}(E)$, satisfies $\mathbf{q}(\mathbf{y}) = \mathbf{x}$. This assumption implies that the state \mathbf{x} could be uniquely determined from its corresponding measurement \mathbf{y} if the function \mathbf{q} was known, but that is often not the case in many modern control applications (such as when using vision-based measurement systems).

While the function \mathbf{p} is often determined by the physical attributes of a system and its sensors, an estimate of the function \mathbf{q} , given by the function $\widehat{\mathbf{q}} : \mathbf{p}(E) \rightarrow E$ which is assumed to be locally Lipschitz continuous on $\mathbf{p}(E)$, is constructed such that given an $\mathbf{x} \in E$ with a corresponding measurement $\mathbf{y} \in \mathbf{p}(E)$, an estimate of the state, $\widehat{\mathbf{x}} \in E$, can be determined as $\widehat{\mathbf{x}} = \widehat{\mathbf{q}}(\mathbf{y})$. For notational simplicity, I define the measurement-estimate function $\widehat{\mathbf{v}} : E \rightarrow \mathbf{p}(E) \times E$ such that $\widehat{\mathbf{v}}(\mathbf{x}) = (\mathbf{p}(\mathbf{x}), \widehat{\mathbf{q}}(\mathbf{p}(\mathbf{x})))$ and $\widehat{\mathbf{v}}(E)$ as the image of the set E under the measurement-estimate function.

The function $\widehat{\mathbf{q}}$ is constructed either via system and measurement models, or from data using learning methods, and thus its accuracy in estimating \mathbf{q} degrades with imperfections in sensor fabrication and integration, or imperfections in learning models and training data. Thus, given a state $\mathbf{x} \in E$ and a corresponding measurement $\mathbf{y} \in \mathbf{p}(E)$, I assume that the corresponding state estimate $\widehat{\mathbf{x}}$ is related to the true state by:

$$\widehat{\mathbf{x}} = \widehat{\mathbf{q}}(\mathbf{y}) \triangleq \mathbf{x} + \mathbf{e}(\mathbf{y}), \quad (5.2)$$

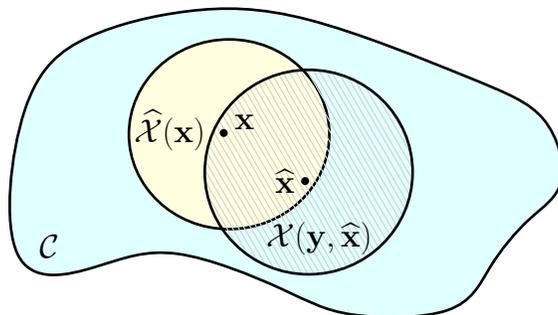


Figure 5.1. Visualization of the point-wise sets $\hat{\mathcal{X}}(\mathbf{x})$ and $\mathcal{X}(\mathbf{y}, \hat{\mathbf{x}})$. The set $\hat{\mathcal{X}}(\mathbf{x})$ consists of all possible estimates that may correspond to the state \mathbf{x} , and the set $\mathcal{X}(\mathbf{y}, \hat{\mathbf{x}})$ consists of all the possible states for a given measurement-state estimate pair $(\mathbf{y}, \hat{\mathbf{x}})$.

for an unknown function $\mathbf{e} : \mathbf{p}(E) \rightarrow \mathbb{R}^n$ that is defined implicitly via $\hat{\mathbf{q}}$. In practice, the function \mathbf{e} can often be characterized via upper bounds on measurement model errors or via data-driven arguments for learning models. In particular, I assume that while $\mathbf{e}(\mathbf{y})$ is not known for a particular measurement $\mathbf{y} \in \mathbf{p}(E)$, it is known that $\mathbf{e}(\mathbf{y}) \in \mathcal{E}(\mathbf{y})$ for a measurement-dependent, compact pointwise set $\mathcal{E}(\mathbf{y}) \subset \mathbb{R}^n$. This leads to the definition of the following two pointwise sets:

$$\hat{\mathcal{X}}(\mathbf{x}) \triangleq \{\hat{\mathbf{x}} \in \mathbb{R}^n \mid \exists \mathbf{e} \in \mathcal{E}(\mathbf{p}(\mathbf{x})) \text{ s.t. } \hat{\mathbf{x}} = \mathbf{x} + \mathbf{e}\}, \quad (5.3)$$

$$\mathcal{X}(\mathbf{y}, \hat{\mathbf{x}}) \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid \exists \mathbf{e} \in \mathcal{E}(\mathbf{y}) \text{ s.t. } \hat{\mathbf{x}} = \mathbf{x} + \mathbf{e}\}. \quad (5.4)$$

The first of these two pointwise sets can be interpreted as all possible state estimates $\hat{\mathbf{x}}$ corresponding to a particular state \mathbf{x} , restricted by the possible error dictated by the set $\mathcal{E}(\mathbf{p}(\mathbf{x}))$. The second pointwise set consists of all potential states that may yield a measurement-state estimate pair. This set is directly computable given a measurement and state estimate, and will play an important role in control synthesis. These sets are visualized in Figure 5.1.

Measurement-Robust Control Barrier Functions

Consider a controller $\mathbf{k} : \mathbf{p}(E) \times E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on $\mathbf{p}(E) \times E$ and depends on a measurement \mathbf{y} and state estimate $\hat{\mathbf{x}}$, rather than explicitly on the state \mathbf{x} . This leads to the closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{y}, \hat{\mathbf{x}}). \quad (5.5)$$

I assume that for any initial condition $\mathbf{x}_0 \in E$, there exists an interval $I(\mathbf{x}_0) = [0, t_{\max}(\mathbf{x}_0)) \subseteq \mathbb{R}_{\geq 0}$ and a unique continuously differentiable solution $\varphi :$

$I(\mathbf{x}_0) \rightarrow E$ satisfying:

$$\dot{\boldsymbol{\varphi}}(t) = \mathbf{f}(\boldsymbol{\varphi}(t)) + \mathbf{g}(\boldsymbol{\varphi}(t))\mathbf{k}(\mathbf{p}(\boldsymbol{\varphi}(t)), \widehat{\mathbf{q}}(\mathbf{p}(\boldsymbol{\varphi}(t))))), \quad t \in I(\mathbf{x}_0), \quad (5.6)$$

$$\boldsymbol{\varphi}(0) = \mathbf{x}_0. \quad (5.7)$$

Consider a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E and its corresponding 0-superlevel set \mathcal{C} . Observe that satisfying the BF condition in (2.57) would require that:

$$\dot{h}(\mathbf{x}) = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) \geq -\alpha(h(\mathbf{x})), \quad (5.8)$$

for each $\mathbf{x} \in E$ with corresponding measurement-state estimate pair $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$ for some $\alpha \in \mathcal{K}^e$. Importantly, while the controller \mathbf{k} only has access to the measurement-state estimate pair, the safety constraint depends on the true state in several ways, including the value of h and $\frac{\partial h}{\partial \mathbf{x}}$ as well as the dynamics \mathbf{f} and \mathbf{g} . Thus, I will propose an alternative condition which depends on only a measurement \mathbf{y} , a state estimate $\widehat{\mathbf{x}}$, and the set $\mathcal{X}(\mathbf{y}, \widehat{\mathbf{x}})$. Observe that to ensure safety with a BF, it is sufficient for the following condition to hold for all $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$:

$$\inf_{\mathbf{x} \in \mathcal{X}(\mathbf{y}, \widehat{\mathbf{x}})} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) \geq -\alpha(h(\mathbf{x})). \quad (5.9)$$

This condition implies that the controller \mathbf{k} renders the system safe for all possible states corresponding to a given measurement-state estimate pair. Verifying that this condition holds can be difficult for an arbitrary CBF, and it is not easily (or possibly) enforced in a convex-optimization based controller. To resolve these problems, I introduce the following definition:

Definition 40 (*Measurement-Robust Barrier Function (MR-BF)*). Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is a *Measurement-Robust barrier function (MR-BF)* for the closed-loop system (5.5) on \mathcal{C} with *parameter function* $(b, a) : \widehat{\mathbf{v}}(E) \rightarrow \mathbb{R}_{\geq 0}^2$ that is locally Lipschitz continuous on $\widehat{\mathbf{v}}(E)$ if there exists $\alpha \in \mathcal{K}^e$ such that:

$$\frac{\partial h}{\partial \mathbf{x}}(\widehat{\mathbf{x}}) (\mathbf{f}(\widehat{\mathbf{x}}) + \mathbf{g}(\widehat{\mathbf{x}})\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - (b(\mathbf{y}, \widehat{\mathbf{x}}) + a(\mathbf{y}, \widehat{\mathbf{x}})\|\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})\|) \geq -\alpha(h(\widehat{\mathbf{x}})). \quad (5.10)$$

for all $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$.

The definition of a MR-BF introduces the non-positive term $-(b(\mathbf{y}, \hat{\mathbf{x}}) + a(\mathbf{y}, \hat{\mathbf{x}})\|\mathbf{k}(\mathbf{y}, \hat{\mathbf{x}})\|)$ to the BF condition (2.57), requiring that a stronger degree of safety be enforced compared to the typical BF. Furthermore, the norm of the input appears in this term, indicating that for large values of $a(\mathbf{y}, \hat{\mathbf{x}})$, large inputs will often violate this constraint. The advantage of this is that knowledge of the true state \mathbf{x} can be removed from the BF condition and replaced with only knowledge of the measurement \mathbf{y} and state estimate $\hat{\mathbf{x}}$ which are available to the controller \mathbf{k} . This enables the following safety result:

Theorem 36. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . Assume¹ the functions $L_{\mathbf{f}}h$, $L_{\mathbf{g}}h$, and $\alpha \circ h : E \rightarrow \mathbb{R}$ are Lipschitz continuous on E with Lipschitz constants $L_{L_{\mathbf{f}}h}$, $L_{L_{\mathbf{g}}h}$, and $L_{\alpha \circ h}$, respectively. Further let a function $\epsilon : \mathbf{p}(E) \rightarrow \mathbb{R}_{\geq 0}$ that is locally Lipschitz continuous on $\mathbf{p}(E)$ satisfy:*

$$\max_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} \|\mathbf{e}\| \leq \epsilon(\mathbf{y}) \quad (5.11)$$

for all $\mathbf{y} \in \mathbf{p}(E)$. If h is a MR-BF for (5.5) on \mathcal{C} with parameter function $(\epsilon(\mathbf{y})(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h}), \epsilon(\mathbf{y})L_{L_{\mathbf{g}}h})$, then the system (5.5) is safe with respect to \mathcal{C} .

Before proving this theorem, note that if h is an MR-BF for the closed-loop system (5.5) on \mathcal{C} with a parameter function $(b, a) : \hat{\mathbf{v}}(E) \rightarrow \mathbb{R}_{\geq 0}^2$ such that $b(\mathbf{y}, \hat{\mathbf{x}}) \geq \epsilon(\mathbf{y})(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h})$ and $a(\mathbf{y}, \hat{\mathbf{x}}) \geq \epsilon(\mathbf{y})L_{L_{\mathbf{g}}h}$ for all measurement-state estimate pairs $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$, then h is an MR-BF for the closed-loop system (5.5) on \mathcal{C} with the parameter function $(\epsilon(\mathbf{y})(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h}), \epsilon(\mathbf{y})L_{L_{\mathbf{g}}h})$. Thus, even if tight Lipschitz constants of these functions are not known, over-approximations can be used and theoretical safety guarantees will be retained (at the expense of potential conservativeness).

Proof. Define the function $c : E \times \mathbb{R}^m \rightarrow \mathbb{R}$ as

$$c(\mathbf{x}, \mathbf{u}) = L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} + \alpha(h(\mathbf{x})). \quad (5.12)$$

This proof will follow from the proof of Theorem 8 by showing that for any $\mathbf{x} \in E$, with corresponding measurement-state estimate pair $(\mathbf{y}, \hat{\mathbf{x}}) = \hat{\mathbf{v}}(\mathbf{x})$:

$$c(\mathbf{x}, \mathbf{k}(\mathbf{y}, \hat{\mathbf{x}})) \geq 0. \quad (5.13)$$

¹The following results developed using MR-CBFs will focus on maintaining forward invariance of the set \mathcal{C} , rather than an enlargement of \mathcal{C} as was considered in Chapter 4. Due to this, the set that this assumption needs to hold over can be any arbitrarily small open set containing \mathcal{C} , to which E is understood to be restricted to. This assumption is particularly reasonable if the set \mathcal{C} is compact.

To show that (5.13) is true, consider a measurement-state estimate pair $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$. A sufficient condition for (5.13) to hold is given by:

$$\inf_{\mathbf{x} \in \mathcal{X}(\mathbf{y}, \widehat{\mathbf{x}})} c(\mathbf{x}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) \geq 0. \quad (5.14)$$

Recalling that $\widehat{\mathbf{x}} = \mathbf{x} + \mathbf{e}(\mathbf{x})$, we have:

$$\inf_{\mathbf{x} \in \mathcal{X}(\mathbf{y}, \widehat{\mathbf{x}})} c(\mathbf{x}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) = \inf_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} c(\widehat{\mathbf{x}} - \mathbf{e}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})), \quad (5.15)$$

$$= c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) + \inf_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} c(\widehat{\mathbf{x}} - \mathbf{e}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})), \quad (5.16)$$

$$\geq c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - \sup_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} |c(\widehat{\mathbf{x}} - \mathbf{e}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}}))|. \quad (5.17)$$

The assumption on Lipschitz continuity of $L_{\mathbf{f}}h$, $L_{\mathbf{g}}h$, and $\alpha \circ h$ on E enables the following bound:

$$|c(\mathbf{x}', \mathbf{u}) - c(\mathbf{x}, \mathbf{u})| \leq |L_{\mathbf{f}}h(\mathbf{x}') - L_{\mathbf{f}}h(\mathbf{x})| + |L_{\mathbf{g}}h(\mathbf{x}')\mathbf{u} - L_{\mathbf{g}}h(\mathbf{x})\mathbf{u}| + |\alpha(h(\mathbf{x}')) - \alpha(h(\mathbf{x}))|, \quad (5.18)$$

$$\leq L_{L_{\mathbf{f}}h} \|\mathbf{x}' - \mathbf{x}\| + \|L_{\mathbf{g}}h(\mathbf{x}') - L_{\mathbf{g}}h(\mathbf{x})\| \|\mathbf{u}\| + L_{\alpha \circ h} \|\mathbf{x}' - \mathbf{x}\|, \quad (5.19)$$

$$\leq (L_{L_{\mathbf{f}}h} + L_{L_{\mathbf{g}}h} \|\mathbf{u}\| + L_{\alpha \circ h}) \|\mathbf{x}' - \mathbf{x}\|. \quad (5.20)$$

Therefore, using the definition of $\epsilon(\mathbf{y})$ we have:

$$\begin{aligned} \sup_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} |c(\widehat{\mathbf{x}} - \mathbf{e}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}}))| \\ \leq \sup_{\mathbf{e} \in \mathcal{E}(\mathbf{y})} (L_{L_{\mathbf{f}}h} + L_{L_{\mathbf{g}}h} \|\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})\| + L_{\alpha \circ h}) \|\mathbf{e}\|, \end{aligned} \quad (5.21)$$

$$\leq (L_{L_{\mathbf{f}}h} + L_{L_{\mathbf{g}}h} \|\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})\| + L_{\alpha \circ h}) \epsilon(\mathbf{y}). \quad (5.22)$$

Thus:

$$\inf_{\mathbf{x} \in \mathcal{X}(\mathbf{y}, \widehat{\mathbf{x}})} c(\mathbf{x}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) \geq c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - (L_{L_{\mathbf{f}}h} + L_{L_{\mathbf{g}}h} \|\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})\| + L_{\alpha \circ h}) \epsilon(\mathbf{y}). \quad (5.23)$$

By the MR-BF condition and the design of \mathbf{k} we have that:

$$c(\widehat{\mathbf{x}}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})) - \epsilon(\mathbf{y})(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h}) - \epsilon(\mathbf{y})L_{L_{\mathbf{g}}h} \|\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})\| \geq 0, \quad (5.24)$$

implying the condition (5.13). \square

As with BFs and CBFs, the notion of an MR-BF extends to an MR-CBF:

Definition 41 (*Measurement-Robust Control Barrier Function (MR-CBF)*).

Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . The function h is a *Measurement-Robust Control Barrier Function (MR-CBF)* for the open-loop system (2.1) on \mathcal{C} with parameter function $(b, a) : \widehat{\mathbf{v}}(E) \rightarrow \mathbb{R}_{\geq 0}^2$ that is locally Lipschitz continuous on $\widehat{\mathbf{v}}(E)$ if there exists $\alpha \in \mathcal{K}^e$ such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \frac{\partial h}{\partial \mathbf{x}}(\widehat{\mathbf{x}})(\mathbf{f}(\widehat{\mathbf{x}}) + \mathbf{g}(\widehat{\mathbf{x}})\mathbf{u} - (b(\mathbf{y}, \widehat{\mathbf{x}}) + a(\mathbf{y}, \widehat{\mathbf{x}})\|\mathbf{u}\|)) > -\alpha(h(\widehat{\mathbf{x}})). \quad (5.25)$$

for all $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$.

This condition is equivalently stated as:

$$\|L_{\mathbf{g}}h(\widehat{\mathbf{x}})\| \leq a(\mathbf{y}, \widehat{\mathbf{x}}) \implies L_{\mathbf{f}}h(\widehat{\mathbf{x}}) > -\alpha(h(\widehat{\mathbf{x}})) + b(\mathbf{y}, \widehat{\mathbf{x}}). \quad (5.26)$$

In contrast to the implication in (2.61), we must not only consider where $L_{\mathbf{g}}h(\widehat{\mathbf{x}}) = \mathbf{0}_m$, but also where it is smaller than some threshold. Intuitively, this captures the fact that if we are unsure of how the input will impact the time derivative of the MR-CBF (we don't know the direction of $L_{\mathbf{g}}h(\widehat{\mathbf{x}})$ sufficiently well), then large inputs can be highly detrimental, and we can't take them to achieve safety. As $\epsilon(\mathbf{y})$ becomes smaller, the level of robustness required by an MR-CBF approaches that of a regular CBF for the same set \mathcal{C} , and recovers the original CBF condition with no measurement error.

Given a MR-CBF h for (2.1) on \mathcal{C} with parameter function (b, a) and a corresponding $\alpha \in \mathcal{K}^e$, define the pointwise set:

$$K_{\text{MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}}) \triangleq \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} L_{\mathbf{f}}h(\widehat{\mathbf{x}}) + L_{\mathbf{g}}h(\widehat{\mathbf{x}})\mathbf{u} \\ -(b(\mathbf{y}) + a(\mathbf{y})\|\mathbf{u}\|) \geq -\alpha(h(\widehat{\mathbf{x}})) \end{array} \right\}, \quad (5.27)$$

for $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$. Given this construction, we have the following result:

Theorem 37. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is an MR-CBF for the open-loop system (2.1) on \mathcal{C} with parameter function $(b, a) : \widehat{\mathbf{v}}(E) \rightarrow \mathbb{R}_{\geq 0}^2$, then the set $K_{\text{MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}})$ is non-empty for all $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$, and for any controller \mathbf{k} that is locally Lipschitz continuous on $\widehat{\mathbf{v}}(E)$ with $\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}}) \in K_{\text{MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}})$ for all $(\mathbf{y}, \widehat{\mathbf{x}}) \in \widehat{\mathbf{v}}(E)$, the function h is a MR-BF for the closed-loop system (5.5) on \mathcal{C} with parameter function (b, a) .*

One advantage of this approach for resolving the impact of measurement model error on safety is that the constraint in (5.27) remains convex. Given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E , this constraint can then be directly integrated into an optimization based controller as follows:

$$\begin{aligned} \mathbf{k}_{\text{MR-CBF}}(\mathbf{y}, \hat{\mathbf{x}}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\hat{\mathbf{x}})\|^2 && \text{(MR-CBF-SOCP)} \\ \text{s.t. } & L_{\mathbf{f}}h(\hat{\mathbf{x}}) - b(\mathbf{y}, \hat{\mathbf{x}}) + L_{\mathbf{g}}h(\hat{\mathbf{x}})\mathbf{u} - a(\mathbf{y}, \hat{\mathbf{x}})\|\mathbf{u}\| \geq -\alpha(h(\hat{\mathbf{x}})). \end{aligned}$$

This optimization problem is a second-order cone program (SOCP), with an explicit conversion to standard form provided below. As this constraint is non-smooth, existing methods for computing closed-form solutions via Lagrangian duality and assessing local Lipschitz continuity [51] are not applicable. Future work will consider methods from variational analysis to study the Lipschitz continuity of solutions to this problem [223]. In practice, a slack variable, δ , is often added to ensure constraint feasibility. This relaxation is penalized in the cost with a large coefficient $c \in \mathbb{R}_{>0}$:

$$\begin{aligned} \mathbf{k}_{\text{R-MR-CBF}}(\mathbf{y}, \hat{\mathbf{x}}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\hat{\mathbf{x}})\|^2 + c\delta^2 && \text{(R-MR-CBF-SOCP)} \\ \text{s.t. } & L_{\mathbf{f}}h(\hat{\mathbf{x}}) - b(\mathbf{y}, \hat{\mathbf{x}}) + L_{\mathbf{g}}h(\hat{\mathbf{x}})\mathbf{u} - a(\mathbf{y}, \hat{\mathbf{x}})\|\mathbf{u}\| \geq -\alpha(h(\hat{\mathbf{x}})) - \delta. \end{aligned}$$

While this relaxed controller does not necessarily enforce the desired safety constraint, if δ remains small the impact on safety can be understood through the notion of Projection-to-State Safety (PSSf) discussed in Section 3.6. Furthermore, this relaxation ensures that the resulting controller is locally Lipschitz continuous on $\hat{\mathbf{v}}(E)$, and with the assumptions on \mathbf{p} and $\hat{\mathbf{q}}$, is locally Lipschitz continuous with respect to the true state \mathbf{x} on E , as shown below.

Simulation

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

We now present simulation results using MR-CBFs on a simulated planar Segway platform. The Segway can be seen in the left panel of Figure 5.2, and uses the model in [27] as was done in Sections 3.3, 3.5, and 3.6. The nominal controller \mathbf{k}_{nom} is a simple PD controller as was used in Section 3.5. The

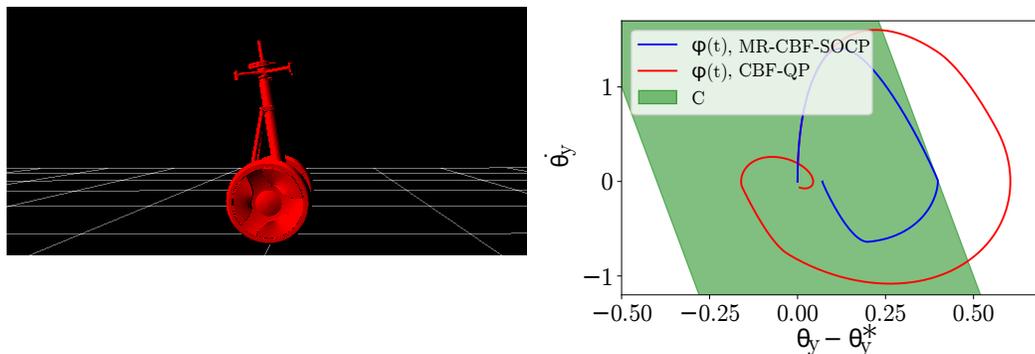


Figure 5.2. (Left) The Segway model used in simulation from the perspective of the fixed virtual camera used to estimate its state. (Right) Simulation results for worst-case measurement model uncertainty of $\epsilon = 0.2$ subtracted from the true pitch angle θ_y when measured. A state trajectory generated using the (CBF-QP) controller (red) and the (MR-CBF-SOCP) controller (blue) are shown as projections onto their pitch angle and pitch rate components. The set \mathcal{C} is plotted in green. Given the same initial condition, the (MR-CBF-SOCP) controller ensured safety of the trajectory whereas the (CBF-QP) controller did not.

simulation² was written in a ROS based environment using a mixture of C++ and Python to mimic the code structure of the existing hardware platform [27]. The set \mathcal{C} was defined as $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^4 \mid h_1(\mathbf{x}) \geq 0, \text{ and } h_2(\mathbf{x}) \geq 0\}$ with:

$$h_1(\mathbf{x}) = -\dot{\theta} + \alpha_e(c - \theta + \theta^*) \quad h_2(\mathbf{x}) = \dot{\theta} + \alpha_e(c + \theta - \theta^*), \quad (5.28)$$

where $\theta \in \mathbb{R}$ is the pitch angle, $\dot{\theta} \in \mathbb{R}$ is the pitch angle rate, θ^* is the pitch angle at equilibrium, and $c, \alpha_e \in \mathbb{R}_{>0}$. The MR-CBF constraint in (5.25) was applied simultaneously to both functions in (5.28) and was then implemented using the ECOS SOCP solver [35]. The parameter function for each MR-CBF was given by parameter function $(\epsilon(\mathbf{y})(L_{L_f}h + L_{\alpha_o}h), \epsilon(\mathbf{y})L_{L_g}h)$, with the Lipschitz constants in this constraint being estimated by sampling $L_f h$, $L_g h$ and $\alpha \circ h$ on a set of gridded values around the system's equilibrium point by taking the maximum of the slopes between any two adjacent grid points. For measurement error, we assumed that direct measurements of the pitch angle θ were offset by a constant factor of $\epsilon > 0$, such that $\hat{\theta} = \theta - \epsilon$. Implementing the (MR-CBF-SOCP) controller with $\epsilon(\mathbf{y}) = \epsilon$ for all $\mathbf{y} \in \mathbf{p}(E)$ ensures safety for this measurement error. As a baseline comparison, a (CBF-QP) controller was implemented and applied using both functions in (5.28). The effect of this type of worst-case measurement model error in the Segway system with the (CBF-QP) and (MR-CBF-SOCP) controllers can be seen in Figure 5.2.

²Simulation code at https://github.com/rkcosner/cyberpod_sim_ros.git.

Conversion to Second-Order Cone Program Standard Form

I now show how the optimization problem specified by (MR-CBF-SOCP) can be written in standard SOCP form. Recalling the (MR-CBF-SOCP) controller, first, the constant term can be removed from the cost such that it becomes: $\frac{1}{2}\|\mathbf{u}\|^2 - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})^\top \mathbf{u}$. Additionally, the constraint can be written in terms of a second order cone:

$$\mathbf{q} = \begin{bmatrix} L_{\mathbf{g}}h(\widehat{\mathbf{x}}) \\ a(\mathbf{y}, \widehat{\mathbf{x}})\mathbf{I}_m \end{bmatrix} \mathbf{u} + \begin{bmatrix} \alpha(h(\widehat{\mathbf{x}})) + L_{\mathbf{f}}h(\widehat{\mathbf{x}}) - b(\mathbf{y}, \widehat{\mathbf{x}}) \\ \mathbf{0}_m \end{bmatrix}, \quad \mathbf{q} \in \mathcal{Q}_m, \quad (5.29)$$

where³ $\mathcal{Q}_m \triangleq \{(q_0, \mathbf{q}_1) \in \mathbb{R} \times \mathbb{R}^m \mid q_0 \geq \|\mathbf{q}_1\|\}$ is the second-order cone in \mathbb{R}^{m+1} . Furthermore, by adding the decision variable $t \in \mathbb{R}$, the quadratic cost function can be converted to an equivalent linear cost, $t - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})^\top \mathbf{u}$, with a rotated second order cone constraint, $\|\mathbf{u}\|^2 \leq 2t$, that can be converted to a standard second order cone constraint via a rotation matrix $\mathbf{R} \in \mathbb{R}^{(m+2) \times (m+2)}$ [30]. Thus the additional constraint can be written as:

$$\mathbf{r} = \mathbf{R} \begin{bmatrix} t & 1 & \mathbf{u}^\top \end{bmatrix}^\top, \quad \mathbf{r} \in \mathcal{Q}_{m+1}. \quad (5.30)$$

Combining the two second order cone constraints, the problem is rewritten as:

$$\begin{aligned} \mathbf{k}_{\text{MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}}) &= \underset{(\mathbf{u}, t, \mathbf{s}) \in \mathbb{R}^{3m+4}}{\text{argmin}} \begin{bmatrix} 1 & 0 & -\mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})^\top \end{bmatrix} \begin{bmatrix} t & 1 & \mathbf{u}^\top \end{bmatrix}^\top \\ &\text{s.t. } \mathbf{G}(\mathbf{y}, \widehat{\mathbf{x}}) \begin{bmatrix} t \\ 1 \\ \mathbf{u} \end{bmatrix} + \mathbf{s} = \mathbf{h}(\mathbf{y}, \widehat{\mathbf{x}}), \\ &\mathbf{s} \in \mathcal{Q}_{m+1} \times \mathcal{Q}_m, \end{aligned} \quad (5.31)$$

where:

$$\mathbf{G}(\mathbf{y}, \widehat{\mathbf{x}}) = - \begin{bmatrix} \mathbf{R}_1 & 0 & \mathbf{R}_{3:m+2} \\ 0 & 0 & L_{\mathbf{g}}h(\widehat{\mathbf{x}}) \\ \mathbf{0}_m & \mathbf{0}_m & a(\mathbf{y})\mathbf{I}_m \end{bmatrix}, \quad (5.32)$$

$$\mathbf{h}(\mathbf{y}, \widehat{\mathbf{x}}) = \begin{bmatrix} \mathbf{R}_2 \\ \alpha(h(\widehat{\mathbf{x}})) + L_{\mathbf{f}}h(\widehat{\mathbf{x}}) - b(\mathbf{y}, \widehat{\mathbf{x}}) \\ \mathbf{0}_m \end{bmatrix}. \quad (5.33)$$

Here \mathbf{R}_i represents the i^{th} column of \mathbf{R} , and $\mathbf{R}_{3:m+2}$ represents the columns 3 through $m+2$. This equivalent formulation of (MR-CBF-SOCP) is in standard SOCP form as in [35].

³This is a slight abuse of notation from Section 3.8 for the sake of notational convenience.

Lipschitz Continuity of Second-Order Cone Program Controller

I now show that the controller (R-MR-CBF-SOCP) is locally Lipschitz continuous on E in terms of the true state \mathbf{x} . Recall this controller is given by:

$$\begin{aligned} \mathbf{k}_{\text{R-MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})\|^2 + c\delta^2 && \text{(R-MR-CBF-SOCP)} \\ \text{s.t. } & L_{\mathbf{f}}h(\widehat{\mathbf{x}}) - b(\mathbf{y}, \widehat{\mathbf{x}}) + L_{\mathbf{g}}h(\widehat{\mathbf{x}})\mathbf{u} - a(\mathbf{y}, \widehat{\mathbf{x}})\|\mathbf{u}\| \geq -\alpha(h(\widehat{\mathbf{x}})) - \delta, \end{aligned}$$

where $\mathbf{y} = \mathbf{p}(\mathbf{x})$ and $\widehat{\mathbf{x}} = \widehat{\mathbf{q}}(\mathbf{p}(\mathbf{x}))$, with \mathbf{p} and $\widehat{\mathbf{q}}$ locally Lipschitz continuous on E . Thus if we prove that $\mathbf{k}_{\text{R-MR-CBF}}$ is locally Lipschitz continuous on $\widehat{\mathbf{v}}(E)$ with respects to its arguments, the composition of $\mathbf{k}_{\text{R-MR-CBF}}$ with the measurement functions will be locally Lipschitz continuous on E . Note this optimization problem is always feasible, even if h is not an MR-CBF, due to the relaxation δ , and there is a unique minimizer as the cost is strictly convex.

The constraint in this problem can be restated as a conic constraint:

$$\begin{aligned} \mathbf{k}_{\text{R-MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}}) &= \underset{(\mathbf{u}, \delta) \in \mathbb{R}^m \times \mathbb{R}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})\|^2 + c\delta^2 && \text{(R-MR-CBF-SOCP)} \\ \text{s.t. } & \begin{bmatrix} L_{\mathbf{g}}h(\widehat{\mathbf{x}})\mathbf{u} + \delta + \alpha(h(\widehat{\mathbf{x}})) + L_{\mathbf{f}}h(\widehat{\mathbf{x}}) - b(\mathbf{y}, \widehat{\mathbf{x}}) \\ a(\mathbf{y}, \widehat{\mathbf{x}})\mathbf{u} \end{bmatrix} \in \mathcal{Q}_m. \end{aligned}$$

To prove local Lipschitz continuity of this controller, I will make use of Theorem 4 in [282], stated below. To draw parallels with the notation of this work, I define the following:

$$\mathbf{w} \triangleq (\mathbf{u}, \delta) \in \mathbb{R}^{m+1}, \quad (5.34)$$

$$C_{(\mathbf{y}, \widehat{\mathbf{x}})}(\mathbf{w}) \triangleq \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})\|^2 + c\delta^2, \quad (5.35)$$

$$G_{(\mathbf{y}, \widehat{\mathbf{x}})}(\mathbf{w}) \triangleq \begin{bmatrix} L_{\mathbf{g}}h(\widehat{\mathbf{x}})\mathbf{u} + \delta + \alpha(h(\widehat{\mathbf{x}})) + L_{\mathbf{f}}h(\widehat{\mathbf{x}}) - b(\mathbf{y}, \widehat{\mathbf{x}}) \\ a(\mathbf{y}, \widehat{\mathbf{x}})\mathbf{u} \end{bmatrix}, \quad (5.36)$$

The optimization problem can then be written as:

$$\begin{aligned} \mathbf{k}_{\text{R-MR-CBF}}(\mathbf{y}, \widehat{\mathbf{x}}) &= \underset{\mathbf{w} \in \mathbb{R}^{m+1}}{\operatorname{argmin}} C_{(\mathbf{y}, \widehat{\mathbf{x}})}(\mathbf{w}) && (5.37) \\ \text{s.t. } & G_{(\mathbf{y}, \widehat{\mathbf{x}})}(\mathbf{w}) \in \mathcal{Q}_m. \end{aligned}$$

I note that $G_{(\mathbf{y}, \widehat{\mathbf{x}})} : \mathbb{R}^{m+1} \rightarrow \mathbb{R}^{m+1}$, with \mathbb{R}^{m+1} a Banach space, and that $\mathcal{Q}_m \subset \mathbb{R}^{m+1}$ is a closed, convex cone with its vertex at the origin. Furthermore, $C_{(\mathbf{y}, \widehat{\mathbf{x}})}$ and $G_{(\mathbf{y}, \widehat{\mathbf{x}})}$ are twice differentiable with respect to \mathbf{w} , and these derivatives are locally Lipschitz continuous with respect to $(\mathbf{y}, \widehat{\mathbf{x}})$ by local Lipschitz continuity

of \mathbf{k}_{nom} and (b, a) on their domains and Lipschitz continuity of $L_f h$, $L_g h$, $\alpha \circ h$ on E (an assumption that appears in Theorem 36).

Let $H_{(\mathbf{y}, \hat{\mathbf{x}})}$ denote the Lagrangian:

$$H_{(\mathbf{y}, \hat{\mathbf{x}})}(\mathbf{w}, \boldsymbol{\lambda}) = C_{(\mathbf{y}, \hat{\mathbf{x}})}(\mathbf{w}) - \langle \boldsymbol{\lambda}, G_{(\mathbf{y}, \hat{\mathbf{x}})}(\mathbf{w}) \rangle, \quad (5.38)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^{m+1}$. The first-order necessary conditions associated with a solution $(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*, \boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}$ to (5.37) can be expressed as:

$$\frac{\partial H_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*, \boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}) = \mathbf{0}_{m+1}, \quad (5.39)$$

$$\langle \boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*, G_{(\mathbf{y}, \hat{\mathbf{x}})}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}) \rangle = 0, \quad (5.40)$$

$$G_{(\mathbf{y}, \hat{\mathbf{x}})}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}) \in \mathcal{Q}_m, \quad (5.41)$$

where $\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^* \in \mathbb{R}^{m+1}$ and $\boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^* \in \mathcal{Q}_m$. We also have that the following coercivity condition is met:

$$\left\langle \frac{\partial H_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*, \boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)})(\mathbf{v}_2 - \mathbf{v}_1), \mathbf{v}_2 - \mathbf{v}_1 \right\rangle \geq \|\mathbf{v}_2 - \mathbf{v}_1\|^2, \quad (5.42)$$

for any $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^{m+1}$ as:

$$\frac{\partial H_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*, \boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}) = \mathbf{I}_{m+1}, \quad (5.43)$$

for all measurement-state estimate pairs $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$. Lastly, I note that:

$$\frac{\partial G_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)}) = \begin{bmatrix} L_g h(\hat{\mathbf{x}}) & 1 \\ a(\mathbf{y}, \hat{\mathbf{x}})\mathbf{I}_m & \mathbf{0}_m \end{bmatrix}. \quad (5.44)$$

Under the assumption that $a(\mathbf{y}, \hat{\mathbf{x}}) \neq 0$ for all $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$ (or that every measurement has some amount of worst-case error), we have the map $\frac{\partial G_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)})$ is surjective from \mathbb{R}^{m+1} to \mathbb{R}^{m+1} . Thus we have that all of the conditions of the following theorem are met [282]:

Theorem 38. *If $\frac{\partial G_{(\mathbf{y}, \hat{\mathbf{x}})}}{\partial \mathbf{w}}(\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*)})$ is surjective and the coercivity condition (5.42) holds, then there exists $s \in \mathbb{R}_{>0}$ such that (5.37) has a strict local minimizer $\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*$ for each $(\mathbf{y}', \hat{\mathbf{x}}') \in B_s((\mathbf{y}, \hat{\mathbf{x}}))$, and both $\mathbf{w}_{(\mathbf{y}, \hat{\mathbf{x}}}^*$, and the associated (unique) multiplier $\boldsymbol{\lambda}_{(\mathbf{y}, \hat{\mathbf{x}}}^*$ satisfying the first-order necessary condition (5.39), are Lipschitz continuous functions of $(\mathbf{y}', \hat{\mathbf{x}}') \in B_s((\mathbf{y}, \hat{\mathbf{x}}))$.*

As $(\mathbf{y}, \hat{\mathbf{x}})$ were arbitrary, and are locally Lipschitz continuous functions of the true state \mathbf{x} , we have that $\mathbf{k}_{\text{R-MR-CBF}}$ is locally Lipschitz with respect to \mathbf{x} .

Conclusion

In this section I have presented an approach for safety-critical control in the presence of measurement errors through the notion of Measurement-Robust Control Barrier Functions (MR-CBFs). I begin by introducing the notion of measurement error that I consider, and propose definitions of MR-BFs and MR-CBFs that allow controllers to be synthesized that provide rigorous safety guarantees in the presence of these measurement errors. I demonstrate the proposed approach in simulation on a planar Segway system that has an imperfect measurement of its pitch angle. Lastly, I show how the MR-CBF-based controller can be specified in standard SOCP form, and prove that a relaxed version of the controller is a locally Lipschitz continuous function of the underlying state of the system.

5.3 Integration with Backup-Set Methods

In this section I will present work on unifying MR-CBFs with backup set CBF methods to produce performant yet robust safe behavior in the presence of measurement errors. I will begin by briefly reviewing backup set CBF methods, followed by integrating them with MR-CBFs. Following this, I will provide demonstrations of the proposed approach in both simulation and experimentally on a planar Segway system. Experimentally, the Segway system will utilize a vision-based measurement system for describing objects in its environment, providing the not only the first experimental demonstration of MR-CBFs, but also the first experimental demonstration of MR-CBFs using vision-based measurement systems.

The contributions of this section are as follows:

- The integration of backup set CBF methods with Measurement-Robust CBFs as a way of achieving both performant and safe behavior in the presence of measurement errors.
- An experimental demonstration of MR-CBFs on a Segway platform using an onboard vision system.

The text for this section is adapted from:

R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, “Measurement-robust control barrier func-

tions: certainty in safety with uncertainty in state,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Prague, Czech Republic, 2021, pp. 6286-6291.

A. J. Taylor participated in the conception of the project, algorithm design, and writing of the article.

Backup-Set Control Barrier Functions

I now briefly review the topic of backup set CBFs methods, which a more detailed description of can be found in [256]. Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function h that is continuously differentiable on E . To guarantee that a safe control input exists, one needs to ensure that the function h satisfies the CBF condition (2.60) for the open-loop system (2.1). For a given set \mathcal{C} , fulfilling this requirement can be nontrivial and potentially impossible (h may not be a CBF on \mathcal{C} for the open-loop system (2.1)). To this end, I will focus on finding a set $\mathcal{C}_I \subseteq \mathcal{C}$ which is control invariant:

Definition 42 (*Control Invariance*). A set $\mathcal{C}_I \subseteq E$ is *control invariant* for the open-loop system (2.1) if there exists a controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E such that the set \mathcal{C}_I is forward invariant for the closed-loop system (2.2).

A control invariant set \mathcal{C}_I is a set that can be kept forward invariant, and thus makes logical sense as a set we would wish to find a CBF for (we will never find one for a set that is not control invariant)⁴. If a control invariant set \mathcal{C}_I can be found such that $\mathcal{C}_I \subseteq \mathcal{C}$, then for initial conditions in \mathcal{C}_I , it is possible to keep the system state within the set \mathcal{C} (as they can be kept in \mathcal{C}_I). Thus, while \mathcal{C} itself will not necessarily be forward invariant (considering initial conditions $\mathbf{x}_0 \in \mathcal{C} \setminus \mathcal{C}_I$), it will be possible to satisfy $\varphi(t) \in \mathcal{C}$ for all $t \in I(\mathbf{x}_0)$ if $\mathbf{x}_0 \in \mathcal{C}_I$.

While directly computing control invariant sets remains challenging in general, we may define one implicitly via a backup set [256]. Suppose that we desire to keep the system state within the set \mathcal{C} , but \mathcal{C} is not necessarily control invariant for the open-loop system (2.1). Additionally, suppose there exists a set $\mathcal{C}_B \subset \mathcal{C}$, defined as the 0-superlevel set of a function $h_B : E \rightarrow \mathbb{R}$ that is continuously

⁴Throughout this thesis, when taking a CBF h on a set \mathcal{C} for the open-loop system (2.1), I am implicitly assuming that \mathcal{C} is control invariant. In this way, backup-set CBF methods can offer a strong starting point for getting to a CBF that can later be robustified, as is done in this section.

differentiable on E and which is known *a priori* to be control invariant for the open-loop system (2.1) and can be rendered forward invariant for all time ($I(\mathbf{x}_0) = \mathbb{R}_{\geq 0}$ for all $\mathbf{x}_0 \in \mathcal{C}_B$) by a known *backup controller* $\mathbf{k}_B : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E . I refer to \mathcal{C}_B as the *backup set*. For simple backup controllers (such as linear controllers designed for the linearization of a system) it is possible to find analytical expressions for local regions of attraction to serve as backup sets. Alternatively, numerical tools such as SOS may be used to synthesize control invariant sets [283]. I extend the backup set to a larger control invariant set $\mathcal{C}_I \subset E$, satisfying $\mathcal{C}_B \subseteq \mathcal{C}_I \subseteq \mathcal{C}$, by considering the *backup trajectory* over a finite and fixed time $T \in \mathbb{R}_{>0}$ as follows. Assume that for any $\mathbf{x} \in E$ there exists a unique solution $\varphi_B : [0, T] \rightarrow E$ satisfying:

$$\dot{\varphi}_B(\tau) = \mathbf{f}(\varphi(\tau)) + \mathbf{g}(\varphi(\tau))\mathbf{k}_B(\varphi(\tau)), \quad \tau \in [0, T], \quad (5.45)$$

$$\varphi_B(0) = \mathbf{x}. \quad (5.46)$$

The solution φ_B may be interpreted as the evolution of the system over the interval $[0, T]$ from a state, \mathbf{x} , under the backup controller \mathbf{k}_B . I denote $\phi_\tau^B(\mathbf{x}) \triangleq \varphi_B(\tau)$ for the initial condition \mathbf{x} and $\tau \in [0, T]$. Using this notation, I define the set $\mathcal{C}_I \subseteq \mathcal{C}$ as:

$$\mathcal{C}_I = \{ \mathbf{x} \in \mathcal{C} \mid h(\phi_\tau^B(\mathbf{x})) \geq 0 \ \forall \tau \in [0, T] \text{ and } h_B(\phi_T^B(\mathbf{x})) \geq 0 \}. \quad (5.47)$$

The first inequality implies safety under the backup policy ($\phi_\tau^B(\mathbf{x}) \in \mathcal{C}$ for all $\tau \in [0, T]$), and the second inequality implies the backup trajectory reaches \mathcal{C}_B by time T ($\phi_T^B(\mathbf{x}) \in \mathcal{C}_B$). The set \mathcal{C}_I is thus control invariant as there exists at least one controller that is locally Lipschitz continuous on E , \mathbf{k}_B , which renders it forward invariant⁵. While \mathcal{C}_I is not necessarily the largest control invariant subset of \mathcal{C} due to its implicit dependence on \mathcal{C}_B (see *viability kernel*, [284]), backup-sets provide a computationally tractable method for finding an under-approximation of the largest control invariant set.

⁵This is a non-obvious fact. Note that for any state $\mathbf{x} \in \mathcal{C}_B$, we have $\phi_\tau^B(\mathbf{x}) \in \mathcal{C}_B$ for all $\tau \in [0, T]$ by our assumption that \mathbf{k}_B renders \mathcal{C}_B forward invariant for all time. As $\mathcal{C}_B \subset \mathcal{C}$, we have that both inequalities in (5.47) are satisfied, and thus $\mathcal{C}_B \subset \mathcal{C}_I$. Additionally, the first inequality defining \mathcal{C}_I requires that $\mathbf{x} \in \mathcal{C}_I \implies \mathbf{x} \in \mathcal{C}$, such that $\mathcal{C}_I \subseteq \mathcal{C}$. Now suppose that $\mathbf{x} \in \mathcal{C}_I$, and consider an $\mathbf{x}' \in \mathcal{C}$ such that $\mathbf{x}' = \phi_{\tau'}^B(\mathbf{x})$ for some $\tau' \in [0, T]$. If \mathcal{C}_I is forward invariant using the backup controller \mathbf{k}_B , then we must have $\mathbf{x}' \in \mathcal{C}_I$. To see this, observe that $\phi_\tau^B(\mathbf{x}') \in \mathcal{C}$ for all $\tau \in [0, T - \tau']$ and $\phi_{T-\tau'}^B(\mathbf{x}') \in \mathcal{C}_B$ because $\mathbf{x} \in \mathcal{C}_I$. Moreover, by the forward invariance of \mathcal{C}_B using the backup controller \mathbf{k}_B , we have that $\phi_\tau^B(\mathbf{x}') \in \mathcal{C}_B \subset \mathcal{C}$ for all $\tau \in [T - \tau', T]$, and thus $\mathbf{x}' \in \mathcal{C}_I$.

For notational simplicity, define the continuously differentiable functions $\bar{h}_\tau : E \rightarrow \mathbb{R}$ and $\bar{h}_B : E \rightarrow \mathbb{R}$ as:

$$\bar{h}_\tau(\mathbf{x}) \triangleq h(\phi_\tau^B(\mathbf{x})), \quad \bar{h}_B(\mathbf{x}) \triangleq h_B(\phi_T^B(\mathbf{x})). \quad (5.48)$$

Given these definitions, the CBF condition (2.60) can then be specified as:⁶:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{\bar{h}}_\tau(\mathbf{x}, \mathbf{u}) \triangleq \sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}\bar{h}_\tau(\mathbf{x}) + L_{\mathbf{g}}\bar{h}_\tau(\mathbf{x})\mathbf{u} \geq -\alpha(\bar{h}_\tau(\mathbf{x})), \quad \forall \tau \in [0, T], \quad (5.49)$$

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{\bar{h}}_B(\mathbf{x}, \mathbf{u}) \triangleq \sup_{\mathbf{u} \in \mathbb{R}^m} L_{\mathbf{f}}\bar{h}_B(\mathbf{x}) + L_{\mathbf{g}}\bar{h}_B(\mathbf{x})\mathbf{u} \geq -\alpha(\bar{h}_B(\mathbf{x})), \quad (5.50)$$

for⁷ $\alpha \in \mathcal{K}^e$. Any controller $\mathbf{k} : E \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on E that satisfies both of the constraints (5.49) and (5.50) for all $\mathbf{x} \in E$ (importantly, not just the backup controller \mathbf{k}_B) will render the closed-loop system (2.2) safe with respect to \mathcal{C}_I [13, p. 6]. I note that enforcing the constraint (5.49) is not necessarily tractable as it must hold for all $\tau \in [0, T]$. To resolve this, it can be reduced to a finite collection of more conservative constraints through constraint tightening. Given a nominal controller $\mathbf{k}_{\text{nom}} : E \rightarrow \mathbb{R}^m$ which is locally Lipschitz continuous on E , a controller which implements the finite number of tightened constraints, and thus renders (2.2) safe with respect to \mathcal{C}_I , is given by the following quadratic program-based controller:

$$\begin{aligned} \mathbf{k}_{BS}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} & \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & (\text{BS-QP}) \\ \text{s.t. } & L_{\mathbf{f}}\bar{h}_{\tau_j}(\mathbf{x}) + L_{\mathbf{g}}\bar{h}_{\tau_j}(\mathbf{x})\mathbf{u} \geq -\alpha(\bar{h}_{\tau_j}(\mathbf{x}) - \mu), \\ & L_{\mathbf{f}}\bar{h}_B(\mathbf{x}) + L_{\mathbf{g}}\bar{h}_B(\mathbf{x})\mathbf{u} \geq -\alpha(\bar{h}_B(\mathbf{x})), \end{aligned}$$

for all $\tau_j \in \{0, \Delta_t, \dots, T\}$, where $\Delta_t \in \mathbb{R}_{>0}$ is a time-step such that $T/\Delta_t \in \mathbb{Z}_{\geq 0}$ and $\mu \in \mathbb{R}_{>0}$ satisfies:

$$\mu \geq \frac{\Delta_t}{2} L_h \sup_{\mathbf{x} \in E} \|\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}_B(\mathbf{x})\|, \quad (5.51)$$

⁶Taking the time derivative of the function \bar{h}_τ and \bar{h}_B requires evaluating $\frac{\partial \phi_\tau^B}{\partial \mathbf{x}}(\mathbf{x})$, which is a first-order approximation of the change in the value of $\varphi_B(\tau)$ if the initial condition \mathbf{x} is changed. This evaluation requires integrating a *sensitivity matrix* as explained in [256, Section B.2]. Effort to make this computation easier through learning was proposed in [160].

⁷Note that α need not be the same between both constraints. Further, observe that just because \mathbf{k}_B keeps \mathcal{C}_I forward invariant, it will not necessarily satisfy these barrier constraints for a fixed α . In practice, this is usually overcome by parameterizing the function α with a linear constant ($\alpha(r) = kr$ for $k \in \mathbb{R}_{>0}$) and allowing the constant k to be a decision variable in the resulting optimization-based controller [256, Section 2.C].

with $L_h \in \mathbb{R}_{>0}$ a Lipschitz constant⁸ for h on E [256, Theorem 1].

Backup Set & Measurement-Robust Control Barrier Functions

In this section I will unify the backup set CBF method with MR-CBFs. I will assume that the functions $L_f \bar{h}_{\tau_j}$, $L_g \bar{h}_{\tau_j}$, $L_f \bar{h}_B$, $L_g \bar{h}_B$, \bar{h}_{τ_j} , \bar{h}_B , and α are Lipschitz continuous on E with known Lipschitz constants $L_{L_f \bar{h}_{\tau_j}}$, $L_{L_g \bar{h}_{\tau_j}}$, $L_{L_f \bar{h}_B}$, $L_{L_g \bar{h}_B}$, $L_{\bar{h}_{\tau_j}}$, $L_{\bar{h}_B}$, $L_\alpha \in \mathbb{R}_{\geq 0}$. Using the MR-CBF condition (5.25) the finite set of constraints in the BS-QP become:

$$L_f \bar{h}_{\tau_j}(\hat{\mathbf{x}}) + L_g \bar{h}_{\tau_j}(\hat{\mathbf{x}})\mathbf{u} - (b_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}}) + a_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}})\|\mathbf{u}\|) \geq -\alpha(\bar{h}_{\tau_j}(\hat{\mathbf{x}}) - \mu), \quad (5.52)$$

$$L_f \bar{h}_B(\hat{\mathbf{x}}) + L_g \bar{h}_B(\hat{\mathbf{x}})\mathbf{u} - (b_B(\mathbf{y}, \hat{\mathbf{x}}) + a_B(\mathbf{y}, \hat{\mathbf{x}})\|\mathbf{u}\|) \geq -\alpha(\bar{h}_B(\hat{\mathbf{x}})), \quad (5.53)$$

with parameter functions:

$$b_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}}) = (L_{L_f \bar{h}_{\tau_j}} + L_\alpha L_{\bar{h}_{\tau_j}})\epsilon(\mathbf{y}), \quad (5.54)$$

$$a_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}}) = L_{L_g \bar{h}_{\tau_j}}\epsilon(\mathbf{y}), \quad (5.55)$$

$$b_B(\mathbf{y}, \hat{\mathbf{x}}) = (L_{L_f \bar{h}_B} + L_\alpha L_{\bar{h}_B})\epsilon(\mathbf{y}), \quad (5.56)$$

$$a_B(\mathbf{y}, \hat{\mathbf{x}}) = L_{L_g \bar{h}_B}\epsilon(\mathbf{y}), \quad (5.57)$$

for all $\tau_j \in \{0, \Delta_t, \dots, T\}$, with $\epsilon : \mathbf{p}(E) \rightarrow \mathbb{R}_{\geq 0}$ satisfying the relationship in (5.11) for the pointwise error set $\mathcal{E}(\mathbf{y})$. These constraints can be used to define the following set:

$$K_{\text{BS-MR-CBF}}(\mathbf{y}, \hat{\mathbf{x}}) = \{\mathbf{u} \in \mathbb{R}^m \mid \mathbf{u} \text{ satisfies (5.52) and (5.53)}\}. \quad (5.58)$$

These constructions enable the following definition:

Definition 43 (*Measurement-Robust Implicit Safe Set*). The set $\mathcal{C}_I \subseteq \mathcal{C} \subseteq E$ defined as in (5.47) is a *Measurement-Robust Implicit Safe Set (MRISS)* for the open-loop system (2.1), error function $\epsilon : \mathbf{p}(E) \rightarrow \mathbb{R}_{\geq 0}$, and parameter functions $a_0, b_0, \dots, a_T, b_T, a_B, b_B : \hat{\mathbf{v}}(E) \rightarrow \mathbb{R}_{\geq 0}$ if:

- the constant $\mu \in \mathbb{R}_{\geq 0}$ satisfies (5.51),

⁸As in Section 5.2, we are interested in keeping the system state inside \mathcal{C} , and thus this constant and norm bound only needs to hold over some open set E that contains \mathcal{C} , thus allowing us to restrict our attention to a potentially small set E containing \mathcal{C} , especially if \mathcal{C} is compact.

- The set $K_{\text{BS-MR-CBF}}(\mathbf{y}, \hat{\mathbf{x}})$ is non-empty for each measurement-state estimate pair $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$.

Next, using this definition, I show that the safety of such sets can be made robust to measurement model error.

Theorem 39. *Given a MRISS \mathcal{C}_I for the open-loop system (2.1), error function ϵ , and parameter functions $a_0, b_0, \dots, a_T, b_T, a_B, b_B$, if a controller $\mathbf{k} : \hat{\mathbf{v}}(E) \rightarrow \mathbb{R}^m$ is locally Lipschitz continuous on $\hat{\mathbf{v}}(E)$ and satisfies $\mathbf{k}(\mathbf{x}) \in K_{\text{BS-MR-CBF}}(\mathbf{y}, \hat{\mathbf{x}})$ for all $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$, then the closed-loop system (5.5) is safe with respect to \mathcal{C}_I .*

The proof of this follows the proof of Theorem 36 in Section 5.2 and is thus not provided. This result allows an alternative to the BS-QP controller which adds the measurement-robustness of MR-CBFs. The constraints (5.52) and (5.53) can be directly integrated into a Measurement-Robust Backup Set Second-Order Cone Program controller MR-BS-SOCP as:

$$\begin{aligned} \mathbf{k}(\mathbf{y}, \hat{\mathbf{x}}) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \quad & \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\hat{\mathbf{x}})\|^2 & (\text{MR-BS-SOCP}) \\ \text{s.t.} \quad & L_{\mathbf{f}} \bar{h}_{\tau_j}(\hat{\mathbf{x}}) + L_{\mathbf{g}} \bar{h}_{\tau_j}(\hat{\mathbf{x}}) \mathbf{u} \\ & - (b_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}}) + a_{\tau_j}(\mathbf{y}, \hat{\mathbf{x}}) \|\mathbf{u}\|) \geq -\alpha(\bar{h}_{\tau_j}(\hat{\mathbf{x}}) - \mu), \\ & L_{\mathbf{f}} \bar{h}_B(\hat{\mathbf{x}}) + L_{\mathbf{g}} \bar{h}_B(\hat{\mathbf{x}}) \mathbf{u} \\ & - (b_B(\mathbf{y}, \hat{\mathbf{x}}) + a_B(\mathbf{y}, \hat{\mathbf{x}}) \|\mathbf{u}\|) \geq -\alpha(\bar{h}_B(\hat{\mathbf{x}})), \end{aligned}$$

for all $\tau_j \in \{0, \Delta_t, \dots, T\}$. Since this controller is a second-order cone program (SOCP), there exist a variety of solvers capable of implementing it including ECOS [35] and MOSEK [36]. Notably, the conservative nature of the method scales with the bound on the measurement-model error $\epsilon(\mathbf{y})$ and the MR-BS-SOCP reduces to the BS-QP when $\epsilon(\mathbf{y}) = 0$ for all $\mathbf{y} \in \mathbf{p}(E)$. I remark that the feasibility of MR-BS-SOCP for all $(\mathbf{y}, \hat{\mathbf{x}}) \in \hat{\mathbf{v}}(E)$ can be ensured by adding a slack variable to the optimization problem, the impact of which can be understood via the concept of PSSf in Section 3.6.

Simulation & Experimental Results

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

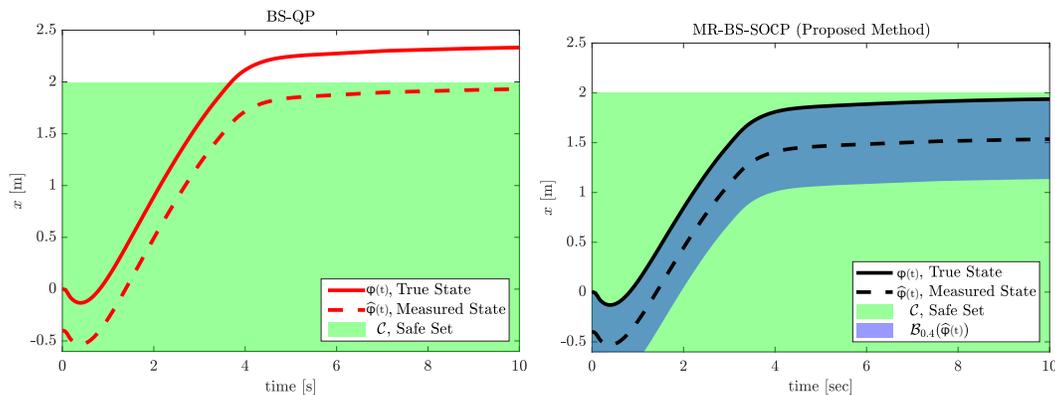


Figure 5.3. Simulation results for a measurement model of $\hat{x} = x - 0.4$ [m] and constant desired velocity of 1 [m/s]. (Left) Trajectories generated using the (BS-QP) controller. Solid line represents the true state, dashed line shows the estimated state, and green region indicates the safe set \mathcal{C} . The true trajectory fails to be safe and exits the safe set at $t = 3$ [s]. (Right) Trajectories generated using the (MR-BS-SOCP) controller. An additional robustness region is plotted in blue to indicate the set of true states which the control input renders safe. Both the true and measured trajectories are safe demonstrating the robustness of the (MR-BS-SOCP) controller when compared to the (BS-QP) controller.

We now demonstrate the efficacy of the proposed (MR-BS-SOCP) controller on Segway platform from Section 5.2 in both simulation and experimentally. The backup set method for generating control invariant sets is particularly relevant for this system due to its non-minimum phase dynamics.

The set \mathcal{C} was chosen to be the set of states with position less than 2 [m] from the origin, i.e. $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^4 \mid x \leq 2\}$ and $h(\mathbf{x}) = 2 - x$. The backup controller was an LQR controller on the linearized system dynamics and the backup set was an estimate of the region of attraction of the LQR controller to the upright equilibrium state, given by a quadratic Lyapunov function. This set is then translated to match the current position of the Segway, while not allowing it to exceed the set boundary. The functions \bar{h}_τ , $\tau \in [0, T]$ were converted into four CBFs \bar{h}_{τ_j} . Lastly, the Lipschitz constants for \bar{h}_{τ_j} were found explicitly by inspection of the Segway dynamics and the Lipschitz constants for \bar{h}_B were estimated by sampling the state space in simulation and taking the largest numerical gradient.

The (MR-BS-SOCP) controller was first validated in simulation in a ROS-based environment⁹. Measurement error was achieved by adding a constant error of -0.4 [m] to the true state, such that $\mathbf{e}(\mathbf{y}) = -0.4$ [m] for all $\mathbf{y} \in \mathbf{p}(E)$.

⁹Simulation code at github.com/rkcosner/mrcbf_IROS21.git.

The simple test scenario involved driving the Segway forward with a constant desired velocity of 1 [m/s]. As seen in Figure 5.3, the MR-BS-SOCP controller provided robustness to this error. Importantly, without measurement-robustness, the system would be unsafe due to error in the state estimate.

The MR-BS-SOCP controller was then implemented on hardware. State estimates for \dot{x} , θ , and $\dot{\theta}$ were found using wheel incremental encoders and a VectorNav VN-100 Inertial Measurement Unit (IMU). The position estimate for x was obtained from an Intel RealSense T265 onboard camera running proprietary Visual Inertial Odometry (VIO) based SLAM. Onboard computation was performed by a Jetson TX2 which computes control inputs and relays them to the low-level motor controllers. The TX2 concurrently runs Linux with ROS, enabling external communication and logging, and the ERIKA3 real-time operating system, which enables real-time low-level communication and computation of the control input.

As the $(\dot{x}, \theta, \dot{\theta})$ state estimates provided by the encoders and IMU are highly accurate, the focus is on making the system robust to measurement error in its vision-based position estimate \hat{x} . An OptiTrack motion capture system was used in laboratory experiments to provide x estimates which are considered true. These closely matched the encoder position estimates for short trials, so the encoder x estimates were considered true in the outdoor experiments. This data was used to determine the error bound $\epsilon(\mathbf{y})$ that appears in the MR-CBF constraint when using the onboard camera. The function $\epsilon(\mathbf{y}) = 0.4$ for all $\mathbf{y} \in \mathbf{p}(E)$ was chosen as an upper bound on the measurement error. The (MR-BS-SOCP) controller was implemented at the embedded level in the ERIKA3 operating system using the ECOS SOCP solver [35]. The nominal controller \mathbf{k}_{nom} was a PD controller tracking user velocity inputs. The backup trajectory $\phi_{\tau}^B(\hat{\mathbf{x}})$ and its partial derivatives were approximated via Euler integration using a time step of $\Delta t = 5$ [ms] and the time used to expand the backup set \mathcal{C}_B to \mathcal{C}_I was $T = 1$ [s]. The MR-BS-SOCP controller ran at 250 [Hz] with 5 decision variables, 4 linear constraints, and 6 second order cone constraints and saturated at ± 20 [N·m].

To demonstrate the method, a simple scenario is executed on the Segway in which it is driven forward at a desired velocity of 1 [m/s]. This scenario is performed with both the BS-QP and MR-BS-SOCP controllers. The results of

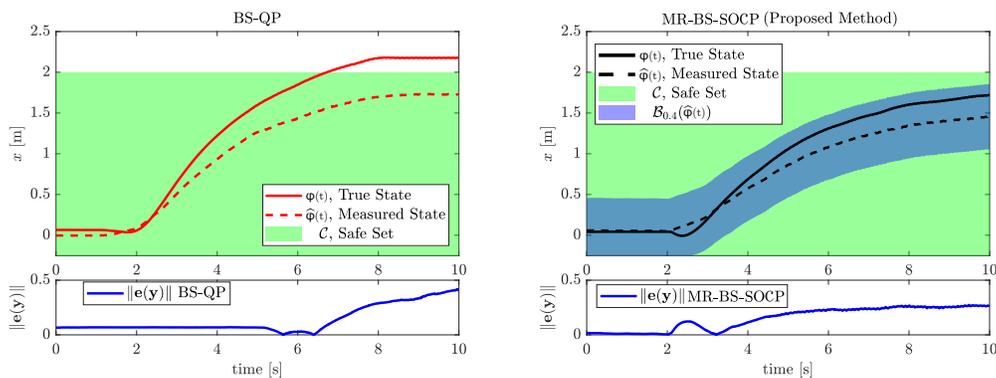


Figure 5.4. Experimental results using SLAM from the onboard Intel RealSense T265 and constant desired velocity of 1 m/s. The notation and color schemes are the same as in Figure (5.3). (Left) Trajectories generated using the (BS-QP) controller. The true trajectory exits the safe set at $t = 6.7$ [s]. The measurement error is plotted in blue. (Right) Trajectories generated using the (MR-BS-SOCP) controller. Both the true and measured trajectories are safe demonstrating the robustness of the (MR-BS-SOCP) controller compared to the (BS-QP) controller.



Figure 5.5. Images from the experiment using the (MR-BS-SOCP) controller. The Segway is piloted towards a wall of yellow boxes and the controller ensures that it remains safe, i.e. that it does not crash into the boxes. (Top) Time lapse of the Segway trajectory. (Bottom) Camera images taken from the perspective of the Segway throughout the experiment. The images are displayed in chronological order from left to right. A video can be found at <https://vimeo.com/520247516>.

these experiments can be found in Figure 5.4, and images from the experiment can be seen in Figure 5.5. With the BS-QP controller the estimated state $\hat{\mathbf{x}}$ remains safe, but the true state \mathbf{x} becomes unsafe whereas with the MR-BS-SOCP controller both the estimated and the true state are kept safe. This highlights the importance of providing robustness against measurement error, as achieved by Theorem 39.

Conclusion

In conclusion, in this section I have presented work on integrating backup set CBF methods with MR-CBFs. I begin by reviewing backup set methods. Following this, I show how the modifications to CBF constraints that occur when forming MR-CBFs can be translated to CBFs synthesized using backup set methods. Following this, I demonstrate the proposed combination in simulation and experimentally on a Segway system using a vision system to estimate its position relative to obstacles.

5.4 Preference-Based Learning for Robust Controller Tuning

In this section I present work on utilizing preference-based learning (PBL) to tune the robustness parameters of a CBF-based controller to achieve both robust safety and performant behavior. I begin by reviewing PBL, including an algorithm developed in this work (not my explicit contribution) for performing PBL in a safety-aware way. Next, I consider the setting of designing safety-critical controllers with reduced-order models and measurement error, producing a controller that unifies the concepts of ISSf-CBFs and MR-CBFs to provide a wide range of robustness properties. As worst-case robustness parameters for this controller can yield excessively conservative behavior, PBL is used to adjust the parameters according to a user's feedback on performance and safety. Following this, I provide demonstrations of the proposed approach in both simulation and experimentally on a quadrupedal system using a vision-based measurement system for obstacle detection.

The contributions of this section are as follows:

- The utilization preference-based learning (PBL) as a tool for tuning the robustness parameters of a CBF-based controller to achieve both robust safety and performant behavior.
- A rigorous unification of reduced-order CBFs methods with ISSf and MR-CBFs to produce a convex optimization-based controller that provides robustness to a wide-range of challenges addressed in this thesis.

The text for this section is adapted from:

R. K. Cosner, M. Tucker, A. J. Taylor, *et al.*, "Safety-aware preference based learning for safety-critical control," in *Proc. 4th Learning for Dy-*

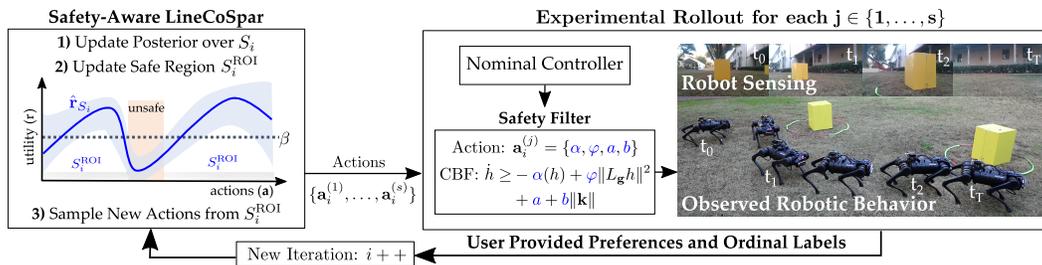


Figure 5.6. An overview of the Safety-Aware Preference-Based Learning design paradigm. Safety-Aware LineCoSpar is used to generate actions which are rolled out in experiments as parameters of the CBF-based safety filter to obtain user preferences and safety ordinal labels which are then used to update the user’s estimated utility and generate new actions.

namics and Control (L4DC), vol. 168, Stanford, CA, USA, 2022, pp. 1020-1033.

A. J. Taylor participated in the conception of the project, theoretical analysis, and writing of the article.

Safety-Aware Preference-Based Learning

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

I now present Safety-Aware LineCoSpar, as outlined in Algorithm 4. I note that this algorithm is not explicitly a contribution of mine (Andrew Taylor) in this work, but introducing it will provide background information on PBL that will support the later developments in the context of tuning robust safety-critical controllers. This is a modification of the LineCoSpar algorithm [77], which iteratively selects actions to query user for subjective feedback and updates its belief of the user’s underlying utility function via Bayesian inference.

Let \mathbf{a} denote an action, such as a collection of l parameters used in a controller, that takes values in a finite search space $A \subset \mathbb{R}^l$. We assume that each action $\mathbf{a} \in A$ has an unknown utility to the user, defined by a function $r : A \rightarrow \mathbb{R}$. These utilities are given by $\mathbf{r}_A = [r(\mathbf{a}_1), \dots, r(\mathbf{a}_{|A|})]^\top \in \mathbb{R}^{|A|}$. In each iteration, $s \in \mathbb{Z}_{>0}$ actions are sampled from A and executed. Then, the user is queried for two forms of feedback: pairwise preferences and ordinal labels, describing *performance* and *safety*, respectively, which is collected into a dataset \mathcal{D} .

Since collecting an exhaustive dataset to estimate the unknown utility \mathbf{r}_A is ex-

pensive for non-trivial action spaces, Bayesian optimization (BO), a sampling efficient paradigm, is used to identify the optimizer. In BO, \mathbf{r}_A is modeled as a Gaussian process with prior $\mathcal{N}(\mathbf{0}_{|A|}, \Sigma^{\text{pr}})$, where each element of the covariance matrix $\Sigma^{\text{pr}} \in \mathbb{S}_{>\mathbf{0}}^{|A| \times |A|}$ is computed as $\Sigma_{ij}^{\text{pr}} = k(\mathbf{a}_i, \mathbf{a}_j)$ with a kernel function $k : A \times A \rightarrow \mathbb{R}$ and $\mathbf{a}_i \in A$ denoting the i^{th} action in A . The function k is selected to be the squared exponential kernel, yielding a prior given by the multivariate Gaussian:

$$\mathbb{P}(\mathbf{r}_A) = \frac{1}{(2\pi)^{|A|/2} |\Sigma^{\text{pr}}|^{1/2}} \exp\left(-\frac{1}{2} \mathbf{r}_A^\top (\Sigma^{\text{pr}})^{-1} \mathbf{r}_A\right). \quad (5.59)$$

Given a dataset \mathfrak{D} , the posterior is proportional to the likelihood and the prior by Bayes' theorem, i.e., $\mathbb{P}(\mathbf{r}_A \mid \mathfrak{D}) \propto \mathbb{P}(\mathfrak{D} \mid \mathbf{r}_A) \mathbb{P}(\mathbf{r}_A)$. Denote the maximum a posteriori (MAP) estimate of the posterior by $\hat{\mathbf{r}}_A \in \mathbb{R}^{|A|}$, which is defined as $\hat{\mathbf{r}}_A \triangleq \operatorname{argmax}_{\mathbf{r}_A \in \mathbb{R}^{|A|}} \mathbb{P}(\mathbf{r}_A \mid \mathfrak{D})$, noting that $\hat{\mathbf{r}}_A$ is equivalent to the minimizer of $\mathcal{S}(\mathbf{r}_A) = -\ln(\mathbb{P}(\mathfrak{D} \mid \mathbf{r}_A)) + \frac{1}{2} \mathbf{r}_A^\top (\Sigma^{\text{pr}})^{-1} \mathbf{r}_A$. As is common in BO, the posterior is modeled as a multivariate Gaussian centered at $\hat{\mathbf{r}}_A$ with the covariance $\Sigma_A \in \mathbb{S}_{>\mathbf{0}}^{|A| \times |A|}$ defined¹⁰ as $\Sigma_A = \left(\frac{\partial^2 \mathcal{S}}{\partial \mathbf{r}_A^2}(\hat{\mathbf{r}}_A)\right)^{-1}$ [75]. Additionally, tractability of calculating $\hat{\mathbf{r}}_A$ can be improved by reducing the action space A to a subset $S \subset A$, forming a partial characterization of the utilities denoted by $\mathbb{P}(\mathbf{r}_S \mid \mathfrak{D}) \approx \mathcal{N}(\hat{\mathbf{r}}_S, \Sigma_S)$, with $\mathbf{r}_S, \hat{\mathbf{r}}_S \in \mathbb{R}^{|S|}$.

A pairwise preference is defined as a relation between two actions $\mathbf{a}_1, \mathbf{a}_2 \in A$, where $\mathbf{a}_1 \succ \mathbf{a}_2$ if action \mathbf{a}_1 is preferred to \mathbf{a}_2 . Since user preferences are expected to be corrupted by noise, individual pairwise preferences are modeled via a likelihood function:

$$\mathbb{P}(\mathbf{a}_1 \succ \mathbf{a}_2 \mid r(\mathbf{a}_1), r(\mathbf{a}_2)) = g_p\left(\frac{r(\mathbf{a}_1) - r(\mathbf{a}_2)}{c_p}\right), \quad (5.60)$$

where $g_p : \mathbb{R} \rightarrow [0, 1]$ is any monotonically-increasing link function, and $c_p \in \mathbb{R}_{>0}$ accounts for preference noise. The function g_p is chosen to be the sigmoid function, i.e., $g_p(r) = 1/(1 + e^{-r})$. Assuming conditional independence, the likelihood function for a collection of $K \in \mathbb{Z}_{>0}$ preferences, \mathfrak{D}_p , can be modeled as the product of each individual preference likelihood:

$$\mathbb{P}(\mathfrak{D}_p \mid r(\mathbf{a}_{11}), r(\mathbf{a}_{12}), \dots, r(\mathbf{a}_{K2})) = \prod_{k=1}^K \mathbb{P}(\mathbf{a}_{k1} \succ \mathbf{a}_{k2} \mid r(\mathbf{a}_{k1}), r(\mathbf{a}_{k2})), \quad (5.61)$$

¹⁰This is known as the Laplace approximation of the distribution $\mathbb{P}(\mathbf{r}_A \mid \mathfrak{D})$, i.e., $\mathbb{P}(\mathbf{r}_A \mid \mathfrak{D}) \approx \mathcal{N}(\hat{\mathbf{r}}_A, \Sigma_A)$.

where $\mathbf{a}_{k1}, \mathbf{a}_{k2} \in A$ are the preferred and non-preferred actions, respectively, in the k^{th} preference.

The action space is partitioned into “unsafe” and “safe” actions by leveraging the ordinal nature of these definitions (i.e., unsafe actions are always considered worse than safe actions). A user provides this feedback as an ordinal label, which assigns an action to a discrete ordered category such as “bad” and “good” [285]. While ordinal labels can be generalized to any number of ordinal categories [78], only two categories are used to represent “unsafe” and “safe”. In this case, the action space is decomposed into two disjoint sets, $A = O_1 \cup O_2$, with $\mathbf{a} \in O_1$ if $r(\mathbf{a}) < \beta$ and $\mathbf{a} \in O_2$ if $r(\mathbf{a}) \geq \beta$, with the ordinal threshold $\beta \in \mathbb{R}$. As with preferences, it is assumed that ordinal label feedback is corrupted by noise and is modeled as:

$$\mathbb{P}(\mathbf{a} \in O_1 \mid r(\mathbf{a})) = g_o \left(\frac{\beta - r(\mathbf{a})}{c_o} \right), \quad (5.62)$$

$$\mathbb{P}(\mathbf{a} \in O_2 \mid r(\mathbf{a})) = 1 - g_o \left(\frac{\beta - r(\mathbf{a})}{c_o} \right), \quad (5.63)$$

where $g_o : \mathbb{R} \rightarrow [0, 1]$ is any monotonically-increasing link function and c_o quantifies the noise in the ordinal label feedback. Again, g_o is chosen to be the sigmoid function $g_o(r) = 1/(1 + e^{-r})$. Assuming conditional independence of ordinal label queries, the likelihood function for a collection of $M \in \mathbb{Z}_{>0}$ ordinal labels, \mathfrak{D}_o , can be modeled as the product of individual ordinal likelihoods:

$$\mathbb{P}(\mathfrak{D}_o \mid r(\mathbf{a}_1), \dots, r(\mathbf{a}_k)) = \prod_{k=1}^M \mathbb{P}(\mathbf{a}_k \in O_{o(k)} \mid r(\mathbf{a}_k)), \quad (5.64)$$

where $\mathbf{a}_k \in A$ refers to the action corresponding to the k^{th} ordinal label, $o(k) \in \{1, 2\}$. For the simulation and experimental results in this section, the hyperparameters c_p , c_o , β are determined in advance. Lastly, assuming conditional independence of the feedback mechanisms, the combined likelihood function is calculated as the product of the individual likelihoods, $\mathbb{P}(\mathfrak{D} \mid r) = \mathbb{P}(\mathfrak{D}_p \mid r)\mathbb{P}(\mathfrak{D}_o \mid r)$.

In the first iteration ($i = 1$), $s \in \mathbb{Z}_{>0}$ actions are sampled randomly from A , recorded as the set of visited actions $V_1 = \{\mathbf{a}_1^{(1)}, \dots, \mathbf{a}_1^{(s)}\}$, executed on the system, and the preferences and ordinal labels are collected into a dataset \mathfrak{D}_1 . In each subsequent iteration ($i > 1$), s new actions are sampled using Thompson sampling, which is shown to have desirable regret minimization

Algorithm 4 Safety-Aware LineCoSpar (SA-LineCoSpar)**Require:** s uniform random actions $V_1 \subset A$, corresponding feedback \mathfrak{D}_1 ,

```

for  $i = 2, \dots, N$  do
   $\mathbb{P}(\hat{\mathbf{r}}_A \mid \mathfrak{D}) \leftarrow V_{i-1}$  ▷ Update posterior
   $\hat{\mathbf{a}}_{i-1}^* \leftarrow \operatorname{argmax}_{\mathbf{a} \in V_{i-1}} \hat{\mathbf{r}}_{V_{i-1}}(\mathbf{a})$  ▷ Find optimal action
   $L_i \leftarrow \hat{\mathbf{a}}_{i-1}^*$  ▷ Sample random intersecting line
   $S_i = L_i \cup V_{i-1}$  ▷ Construct subspace
   $\mathbb{P}(\hat{\mathbf{r}}_{S_i} \mid \mathfrak{D}) \leftarrow S_i$  ▷ Update posterior over  $S_i$ 
   $S_i^{\text{ROI}} \leftarrow \mathbb{P}(\hat{\mathbf{r}}_{S_i} \mid \mathfrak{D})$  ▷ Determine ROI
  for  $j = 1, \dots, s$  do
     $r^{(j)} \sim \mathcal{N}(\hat{\mathbf{r}}_{S_i}, \Sigma_{S_i})$  ▷ Sample utility
     $\mathbf{a}_i^{(j)} \leftarrow \operatorname{argmax}_{\mathbf{a} \in S_i^{\text{ROI}}} r^{(j)}$  ▷ Sample action from utility sample
  end for
   $\mathfrak{D}_p, \mathfrak{D}_o \leftarrow \{\mathbf{a}_i^{(1)}, \dots, \mathbf{a}_i^{(s)}\}$  ▷ Collect preferences and ordinal labels
   $V_i \leftarrow V_{i-1} \cup \{\mathbf{a}_i^{(1)}, \dots, \mathbf{a}_i^{(s)}\}$  ▷ Aggregate Actions
   $\mathfrak{D}_i \leftarrow \mathfrak{D}_{i-1} \cup \mathfrak{D}_p \cup \mathfrak{D}_o$  ▷ Aggregate preferences and ordinal labels
end for

```

properties [286]. Ideally, Thompson sampling draws s samples from the posterior $\mathbb{P}(\mathbf{r}_A \mid \mathfrak{D}_{i-1})$, i.e. $\mathbf{r}^{(j)} \sim \mathbb{P}(\mathbf{r}_A \mid \mathfrak{D}_{i-1})$ for $j \in \{1, \dots, s\}$, and the action $\mathbf{a}_i^{(j)} \in A$ maximizing each $\mathbf{r}^{(j)}$ is selected to execute on the system. These sampled actions $\{\mathbf{a}_i^{(1)}, \dots, \mathbf{a}_i^{(s)}\}$ are concatenated with V_{i-1} to produce V_i , executed on the system, and the resulting preferences and ordinal labels are concatenated with \mathfrak{D}_{i-1} to produce \mathfrak{D}_i . However, since it is intractable to approximate $\mathcal{P}(\mathbf{r}_A \mid \mathfrak{D})$ for high-dimensional action spaces, a dimensionality-reduction technique introduced in [77] that instead updates the posterior over a subset $S_i \subset A$ is utilized. Motivated by [287], a subset of actions is constructed as $S_i = L_i \cup V_{i-1}$, where $L_i \subset A$ is the collection of $e \in \mathbb{Z}_{>0}$ actions in A closest to a randomly drawn line $\ell_i \subset \mathbb{R}^l$. This line is drawn to intersect with the believed best action, computed as $\hat{\mathbf{a}}_{i-1}^* = \operatorname{argmax}_{\mathbf{a} \in V_{i-1}} \hat{\mathbf{r}}_{V_{i-1}}(\mathbf{a})$ where $\hat{\mathbf{r}}_{V_{i-1}}$ is the MAP estimate of the posterior $\mathbb{P}(\mathbf{r}_{V_{i-1}} \mid \mathfrak{D}_i)$. See [77] for more details.

It is important to avoid unsafe actions during sequential decision making in certain applications, such as learning robotic controllers on hardware, where low-reward actions might lead to physical damage of the platform. Safe exploration algorithms [135], [279], [280] considered the setting where actions below a prespecified safety threshold are catastrophic and must be avoided at

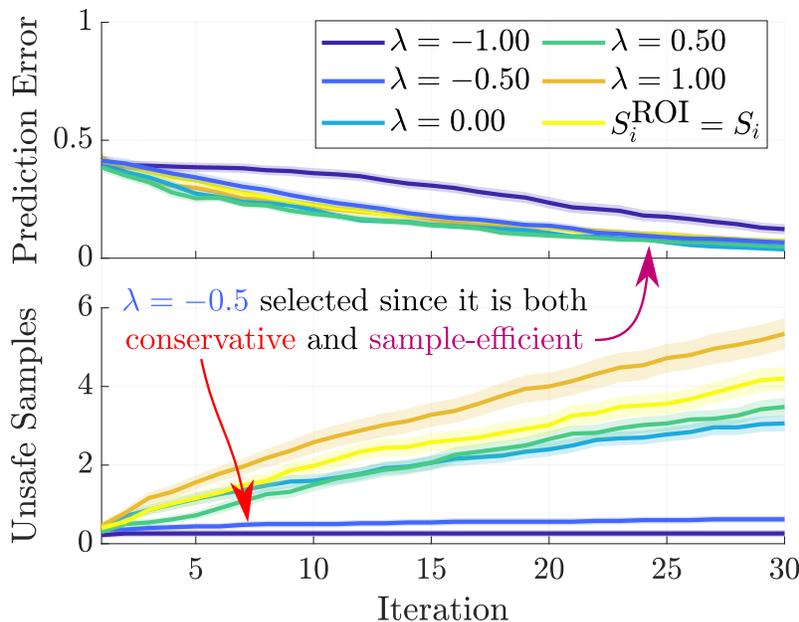


Figure 5.7. A comparison of Safety-Aware LineCoSpar and standard LineCoSpar on a synthetic utility function (drawn from the Gaussian prior) averaged over 50 runs with standard error shown by the shaded region. The safety-aware criteria reduces the number of sampled unsafe actions with a minimal effect on the prediction error, defined as $|\hat{\mathbf{a}}_i^* - \mathbf{a}^*|$ with $\hat{\mathbf{a}}_i^* \triangleq \operatorname{argmax}_{\mathbf{a}} \hat{\mathbf{r}}_{S_i}$ and $\mathbf{a}^* \triangleq \operatorname{argmax}_{\mathbf{a}} r(\mathbf{a})$.

all cost. As this work constructs controllers that account for safety, a more optimistic learning approach called *safety-aware* is adopted. In this case, actions labeled by a human as “unsafe” are not catastrophic but undesirable. Thus, the algorithm *avoids* these actions; whereas the safe exploration algorithms guarantee that no such actions are sampled which can be sometimes exceedingly conservative in settings like that considered in this work.

To achieve this safety-awareness, the approach introduced in [78] is leveraged, which uses ordinal labels to identify a *region of interest* (ROI) in A . In this work, the ROI is defined to be the actions labeled as “safe”. In each iteration i an estimate an ROI within the set S_i is formed as:

$$S_i^{\text{ROI}} = \{\mathbf{a} \in S_i \mid \hat{\mathbf{r}}_{S_i}(\mathbf{a}) + \lambda \sigma_{S_i}(\mathbf{a}) > \beta\}, \quad (5.65)$$

where $\hat{\mathbf{r}}_{S_i}(\mathbf{a})$ and $\sigma_{S_i}(\mathbf{a})$ are the posterior mean and standard deviation, respectively, evaluated at the action $\mathbf{a} \in S_i$. The variable $\lambda \in \mathbb{R}$ determines how conservative the algorithm would be in estimating the safety region, as illustrated in Figure 5.7. We see that lower values of λ result in fewer unsafe actions being sampled, with only a slight effect on sample-efficiency. The restriction to S_i^{ROI} is added to LineCoSpar by only considering actions in S_i^{ROI}

during Thompson sampling. This algorithm is referred to as Safety-Aware LineCoSpar (SA-LineCoSpar), and is outlined in Algorithm 4.

Preference-Based Learning for Robust Controller Tuning

I now propose a design paradigm that leverages SA-LineCoSpar to select parameters for a CBF-based controller that achieves performance and safety for a multi-layered control system. Drawing inspiration from [253], I note that many real-life engineering systems have high-dimensional state spaces and complex dynamics. Hence control systems are often designed as a set of interconnected subsystems, such as a low-dimensional subsystem that provides reference signals capturing safe behavior and a high-dimensional subsystem that tracks these reference signals. In particular, let $n_\xi, m_\xi \in \mathbb{Z}_{>0}$, let $E \triangleq E_x \times E_\xi \subset \mathbb{R}^n \times \mathbb{R}^{n_\xi}$ be an open set, and consider the following cascaded open-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\boldsymbol{\kappa}(\boldsymbol{\xi}), \quad (5.66)$$

$$\dot{\boldsymbol{\xi}} = \mathbf{f}_\xi(\mathbf{x}, \boldsymbol{\xi}) + \mathbf{g}_\xi(\mathbf{x}, \boldsymbol{\xi})\mathbf{u}, \quad (5.67)$$

with state $(\mathbf{x}, \boldsymbol{\xi}) \in E$, control input $\mathbf{u} \in \mathbb{R}^{m_\xi}$, and functions $\mathbf{f} : E_x \rightarrow \mathbb{R}^n$, $\mathbf{g} : E_x \rightarrow \mathbb{R}^{n \times m}$, $\boldsymbol{\kappa} : E_\xi \rightarrow \mathbb{R}^m$, $\mathbf{f}_\xi : E \rightarrow \mathbb{R}^{n_\xi}$, and $\mathbf{g}_\xi : E \rightarrow \mathbb{R}^{n_\xi \times m_\xi}$ assumed to be locally Lipschitz continuous on their domains. I note that the input \mathbf{u} from (2.1) was replaced by $\boldsymbol{\kappa}(\boldsymbol{\xi})$ in the top-level system (5.66). These dynamics may represent Euler-Lagrange systems such as robots, where \mathbf{x} reflects base position, $\boldsymbol{\xi}$ captures base velocities and joint positions and velocities, and the input \mathbf{u} reflects the torques applied to the joints. Control in this setting will be based on a measurement $\mathbf{y} = \mathbf{p}(\mathbf{x}, \boldsymbol{\xi}) \in \mathbf{p}(E)$ and state estimate $\hat{\mathbf{x}} = \hat{\mathbf{q}}(\mathbf{y}) \in E_x$ (I will address $\boldsymbol{\xi}$ shortly) defined using the measurement model setting in Section 5.2, noting that $\hat{\mathbf{v}}(\mathbf{x}) = (\mathbf{y}, \hat{\mathbf{x}})$.

Given this cascaded system, I will seek to establish the Input-to-State Safety of a set \mathcal{C} that can be described as the 0-superlevel set of a function $h : E_x \rightarrow \mathbb{R}$ that is continuously differentiable on E_x and only depends on the states \mathbf{x} and not the states $\boldsymbol{\xi}$. For example, in the context of a robotic system, this assumption is justified if safety is described as keeping the base position of the robot away from obstacles. Second, I assume there exists a controller $\boldsymbol{\pi} : \hat{\mathbf{v}}(E) \times E_\xi \times \mathbb{R}^m \rightarrow \mathbb{R}^{m_\xi}$ and $\mu \in \mathbb{R}_{\geq 0}$ such that for any continuous,

bounded signal $\mathbf{s} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$, the closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\boldsymbol{\kappa}(\boldsymbol{\xi}), \quad (5.68)$$

$$\dot{\boldsymbol{\xi}} = \mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}, \boldsymbol{\xi}) + \mathbf{g}_{\boldsymbol{\xi}}(\mathbf{x}, \boldsymbol{\xi})\boldsymbol{\pi}(\mathbf{y}, \widehat{\mathbf{x}}, \boldsymbol{\xi}, \mathbf{s}(t)), \quad (5.69)$$

satisfies the following implication:

$$\|\boldsymbol{\kappa}(\boldsymbol{\xi}_0) - \mathbf{s}(0)\| \leq \mu \implies \|\boldsymbol{\kappa}(\boldsymbol{\varphi}_{\boldsymbol{\xi}}(t)) - \mathbf{s}(t)\| \leq \mu, \quad t \in \mathbb{R}_{\geq 0}, \quad (5.70)$$

for all initial conditions $(\mathbf{x}_0, \boldsymbol{\xi}_0) \in E$, where $\boldsymbol{\varphi}_{\boldsymbol{\xi}}$ is the part of the corresponding solution $(\boldsymbol{\varphi}_{\mathbf{x}}, \boldsymbol{\varphi}_{\boldsymbol{\xi}})$ from the initial condition $(\mathbf{x}_0, \boldsymbol{\xi}_0)$. This assumption reflects that a separate controller may be designed¹¹ for the high-dimensional dynamics to track well-behaved reference signals synthesized via the low-dimensional model. In particular, if a continuous measurement-based controller $\mathbf{k} : \widehat{\mathbf{v}}(E) \rightarrow \mathbb{R}^m$ is designed for the low-dimensional system (2.1) and $\|\boldsymbol{\kappa}(\boldsymbol{\xi}_0) - \mathbf{k}(\widehat{\mathbf{v}}(\mathbf{x}_0))\| \leq \mu$, then we have that the controller $\boldsymbol{\pi}$ ensures $\|\boldsymbol{\kappa}(\boldsymbol{\varphi}_{\boldsymbol{\xi}}(t)) - \mathbf{k}(\widehat{\mathbf{v}}(\boldsymbol{\varphi}_{\mathbf{x}}(t)))\| \leq \mu$ for $t \in \mathbb{R}_{\geq 0}$. With this assumption in mind, we may explicitly define the signal \mathbf{s} and study the ISSf behavior of the closed-loop system with matched disturbances:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})(\mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}}) + \mathbf{d}(t)), \quad (5.71)$$

$$\dot{\boldsymbol{\xi}} = \mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}, \boldsymbol{\xi}) + \mathbf{g}_{\boldsymbol{\xi}}(\mathbf{x}, \boldsymbol{\xi})\boldsymbol{\pi}(\mathbf{y}, \widehat{\mathbf{x}}, \boldsymbol{\xi}, \mathbf{k}(\mathbf{y}, \widehat{\mathbf{x}})), \quad (5.72)$$

where similar to Sections 3.4 and 3.6, the disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ satisfies the relationship $\mathbf{d}(t) = \boldsymbol{\kappa}(\boldsymbol{\varphi}_{\boldsymbol{\xi}}(t)) - \mathbf{k}(\widehat{\mathbf{v}}(\boldsymbol{\varphi}_{\mathbf{x}}(t)))$ for any initial condition $(\mathbf{x}_0, \boldsymbol{\xi}_0) \in E$, and thus must satisfy $\|\mathbf{d}\|_{\infty} \leq \mu$.

I now combine the robustness properties of MR-CBFs and ISSf-CBFs to account for measurement uncertainty and the disturbance, \mathbf{d} , allowing us to make robust safety guarantees for the full closed-system (5.72). This is formalized in the following theorem:

Theorem 40. *Let $\mathcal{C} \subset E_{\mathbf{x}}$ be the 0-superlevel set of a function $h : E_{\mathbf{x}} \rightarrow \mathbb{R}$ that is continuously differentiable on $E_{\mathbf{x}}$. Suppose that the functions $L_{\mathbf{f}}h$, $L_{\mathbf{g}}h$, $\|L_{\mathbf{g}}h\|^2$, and $\alpha \circ h$ are Lipschitz continuous on their domains, and assume that*

¹¹The controller $\boldsymbol{\pi}$ is assumed to have perfect measurements of $\boldsymbol{\xi}$. Note that it yields bounded tracking error (by a particular value of μ , the implication is not for all values of μ) for $\boldsymbol{\varphi}_{\boldsymbol{\xi}}$. This bound tracking can be seen as an abstraction that captures the effect of explicit measurement error in $\widehat{\mathbf{x}}$ as well as implicit measurement error in the value of $\boldsymbol{\xi}$ being passed to $\boldsymbol{\pi}$.

$\|\widehat{\mathbf{x}} - \mathbf{x}\| \leq \epsilon$ for some $\epsilon \in \mathbb{R}_{\geq 0}$ and all $\mathbf{x} \in E_{\mathbf{x}}$. Then there exists $\underline{a}, \underline{b} \in \mathbb{R}_{\geq 0}$ such that if h satisfies:

$$\sup_{\boldsymbol{\nu} \in \mathbb{R}^m} L_{\mathbf{f}}h(\widehat{\mathbf{x}}) + L_{\mathbf{g}}h(\widehat{\mathbf{x}})\boldsymbol{\nu} - \phi\|L_{\mathbf{g}}h(\widehat{\mathbf{x}})\|^2 - b - a\|\boldsymbol{\nu}\| > -\alpha(h(\widehat{\mathbf{x}})), \quad (5.73)$$

for all $\widehat{\mathbf{x}} \in \widehat{\mathbf{q}}(\mathbf{p}(E))$ and some $a, b \in \mathbb{R}_{\geq 0}$ satisfying $a \geq \underline{a}$ and $b \geq \underline{b}$, then there exists a controller $\mathbf{k} : E_{\mathbf{x}} \rightarrow \mathbb{R}^m$ such that (5.72) is ISSf with respect to \mathcal{C} with $\gamma(d) = -\alpha^{-1}(-d^2/(4\phi))$.

Before proving this result, I note that given a nominal controller $\mathbf{k}_{\text{nom}} : E_{\mathbf{x}} \rightarrow \mathbb{R}^m$ that is locally Lipschitz continuous on $E_{\mathbf{x}}$, the constraint in (5.73) can be incorporated into a convex second-order cone program-based controller $\mathbf{k}_{\text{TR-SOCP}} : E_{\mathbf{x}} \rightarrow \mathbb{R}^m$ referred to as a Tunable Robustified Second-Order Cone Program (TR-SOCP) controller with tunable parameters α, ϕ, a , and b :

$$\begin{aligned} \mathbf{k}_{\text{TR-SOCP}}(\widehat{\mathbf{x}}) = \underset{\boldsymbol{\nu} \in \mathbb{R}^m}{\text{argmin}} \quad & \|\boldsymbol{\nu} - \mathbf{k}_{\text{nom}}(\widehat{\mathbf{x}})\|^2 && (\text{TR-SOCP}) \\ \text{s.t.} \quad & L_{\mathbf{f}}h(\widehat{\mathbf{x}}) + L_{\mathbf{g}}h(\widehat{\mathbf{x}})\boldsymbol{\nu} - \phi\|L_{\mathbf{g}}h(\widehat{\mathbf{x}})\|^2 \\ & - b - a\|\boldsymbol{\nu}\| \geq -\alpha h(\widehat{\mathbf{x}}). \end{aligned}$$

Here a linear extended class- \mathcal{K}_{∞} function with coefficient $\alpha \in \mathbb{R}_{>0}$ is used to set up the use of PBL, but is not necessary for the following proof.

Proof. I will show that the controller $\mathbf{k}_{\text{TR-SOCP}} : E_{\mathbf{x}} \rightarrow \mathbb{R}^m$ using values of a and b that are greater than a particular pair $(\underline{a}, \underline{b})$ implies satisfaction of the ISSf-BF constraint for matched disturbances (2.95). For this I choose:

$$\underline{b} = \epsilon(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h} + L_{\phi\|L_{\mathbf{g}}h\|^2}), \quad (5.74)$$

$$\underline{a} = \epsilon L_{L_{\mathbf{g}}h}, \quad (5.75)$$

where $L_{(\cdot)}$ indicates the Lipschitz coefficient of the subscripted function on its domain. Define the function $c : E_{\mathbf{x}} \times \mathbb{R}^m \rightarrow \mathbb{R}$ such that:

$$c(\mathbf{x}, \boldsymbol{\nu}) \triangleq L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\boldsymbol{\nu} - \phi\|L_{\mathbf{g}}h(\mathbf{x})\|^2 + \alpha(h(\mathbf{x})). \quad (5.76)$$

Using this definition we have for all $\boldsymbol{\nu} \in \mathbb{R}^m$ that:

$$c(\mathbf{x}, \boldsymbol{\nu}) = c(\widehat{\mathbf{x}}, \boldsymbol{\nu}) + c(\mathbf{x}, \boldsymbol{\nu}) - c(\widehat{\mathbf{x}}, \boldsymbol{\nu}), \quad (5.77)$$

$$\geq c(\widehat{\mathbf{x}}, \boldsymbol{\nu}) - \underbrace{\epsilon(L_{L_{\mathbf{f}}h} + L_{\alpha \circ h} + L_{\phi\|L_{\mathbf{g}}h\|^2})}_{\underline{b}} - \underbrace{\epsilon L_{L_{\mathbf{g}}h}}_{\underline{a}} \|\boldsymbol{\nu}\|, \quad (5.78)$$

$$\geq c(\widehat{\mathbf{x}}, \boldsymbol{\nu}) - b - a\|\boldsymbol{\nu}\|. \quad (5.79)$$

Above I added zero in (5.77) and used the Lipschitz coefficients and the worst-case measurement error ϵ to achieve the bound in (5.78). Thus we have that:

$$c(\mathbf{x}, \mathbf{k}_{\text{TR-SOCP}}(\hat{\mathbf{x}})) \geq c(\hat{\mathbf{x}}, \mathbf{k}_{\text{TR-SOCP}}(\hat{\mathbf{x}})) - b - a\|\mathbf{k}_{\text{TR-SOCP}}(\hat{\mathbf{x}})\| \geq 0, \quad (5.80)$$

and thus we may follow the proof of Theorem 29 to show that \mathcal{C} is ISSf with $\gamma(d) = -\alpha^{-1}(-d^2/(4\phi))$ where $\alpha^{-1} \in \mathcal{K}^e$. \square

The parameter selection process of the TR-SOCP controller is particularly important, since the parameters \underline{a} and \underline{b} guaranteed to exist by Theorem 40 are worst-case approximations of the uncertainty generated using Lipschitz constants. Such approximations often lead to undesired conservatism and may render the system incapable of performing its goal (as seen in Figure 5.8). Thus, as illustrated in Figure 5.6, SA-LineCoSpar is used to identify user-preferred parameters of the TR-SOCP controller. This relaxes the worst-case over-approximation to experimentally realize performant and safe behavior. This design paradigm relies on the tunable construction of the TR-SOCP controller, allowing the definition of actions for SA-LineCoSpar as $\mathbf{a} = (\alpha, \phi, a, b)$. I note the construction of the TR-SOCP controller assures that unsafe actions are not necessarily catastrophic, as any $\alpha, \phi, a, b > 0$ endows the system with a non-zero degree of robustness to disturbances and measurement error. This allows using a safety-aware approach where unsafe actions are considered undesirable as opposed to a more conservative safety-critical approach to learning where unsafe actions are considered catastrophic.

Simulation & Experimental Results

Note that the contents of this subsection were primarily written by collaborators, edited by Andrew Taylor for their appearance in journal proceedings, and edited by Andrew Taylor for their appearance in this thesis.

The proposed design paradigm is applied to a perception-based obstacle avoidance task with a Unitree A1 quadrupedal robot as seen in Figure 5.6 in simulation and on hardware for both indoor and outdoor environments. The action space A and hyperparameters of PBL are defined in Table 5.1. A unicycle

model is used as the low-order model (5.66):

$$\underbrace{\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \end{bmatrix}}_{\dot{\mathbf{x}}} = \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} \cos(\psi) & 0 \\ \sin(\psi) & 0 \\ 0 & 1 \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} \left(\underbrace{\begin{bmatrix} v \\ \omega \end{bmatrix}}_{\mathbf{v}} + \mathbf{d}(t) \right), \quad (5.81)$$

where (x, y) is the planar position of the robot, ψ is the yaw angle, v is planar velocity, and ω is the yaw rate. The nominal controller \mathbf{k}_{nom} is given by:

$$\mathbf{k}_{\text{nom}}(\mathbf{x}) = \begin{bmatrix} K_v d_g + C \\ -K_\omega \left(\sin(\psi) - \frac{(y_g - y)}{d_g} \right) \end{bmatrix}, \quad (5.82)$$

where (x_g, y_g) is the goal position of the robot, $d_g = \|(x_g - x, y_g - y)\|$ is the distance to the goal, and $K_v, K_\omega, C \in \mathbb{R}_{>0}$.

Obstacle avoidance is encoded via the 0-superlevel set of the function:

$$h(\mathbf{x}) = d_{\text{obs}} - r_{\text{obs}} - \zeta \cos(\psi - \theta), \quad (5.83)$$

where $(x_{\text{obs}}, y_{\text{obs}})$ is the location of an obstacle, $d_{\text{obs}} = \|(x_{\text{obs}} - x, y_{\text{obs}} - y)\|$ and $\theta = \arctan((y_{\text{obs}} - y)/(x_{\text{obs}} - x))$ are the distance and angle from the obstacle, r_{obs} is the sum of the radii of the obstacle and robot, and $\zeta > 0$ determines the effect of the heading angle on safety. The TR-SOCP controller with the nominal controller \mathbf{k}_{nom} from (5.82) is used to drive the system. In practice, infeasibility of the (TR-SOCP) controller was assigned an ordinal label of unsafe and the inputs were saturated such that $v \in [-0.2, 0.3]$ [m/s] and $\omega \in [-0.4, 0.4]$ [rad/s]. The velocity command \mathbf{v} is computed at 20 [Hz] and error introduced by this sampling scheme is captured by the tracking error $\mathbf{d}(t)$. Tracking of \mathbf{v} is performed by an inverse dynamics quadratic program

| Hyperparameter | Value | Name | Min. | Max. | Δ |
|----------------|-------|----------|------|------|----------|
| | | α | 0.5 | 5 | 0.5 |
| | | ϕ | 0 | 1 | 0.1 |
| | | b | 0 | 1 | 0.1 |
| | | a | 0 | 0.05 | 0.005 |

Table 5.1. The safety-aware hyperparameters, and action space bounds (minimum and maximum) with discretizations Δ .

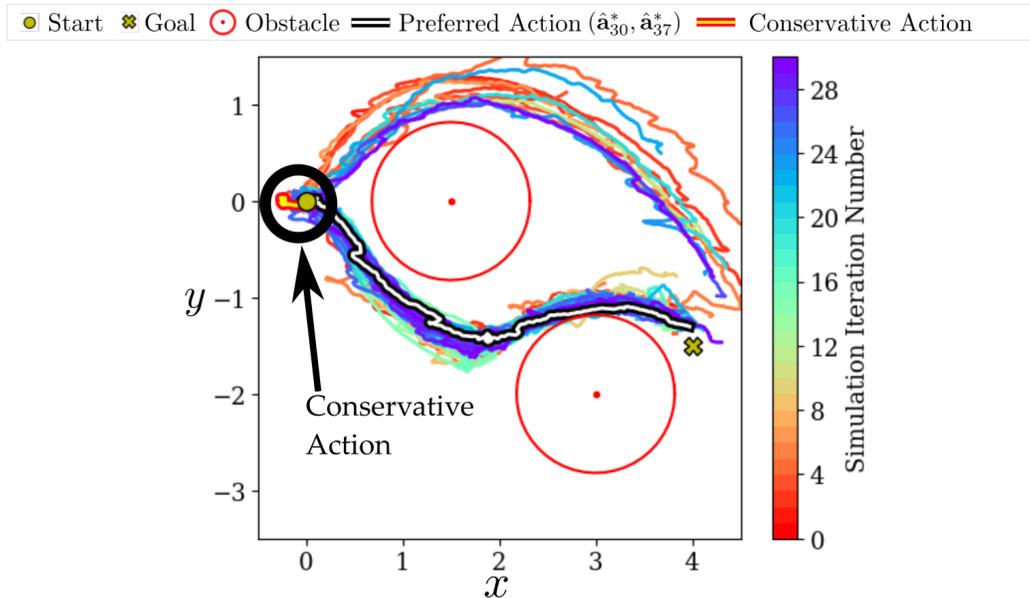


Figure 5.8. Actions sampled during simulation in 30 iterations with 3 new actions in each iteration. The preferred action, $\hat{\mathbf{a}}_{30} = (3, 0.6, 0.5, 0.015)$, is shown in black and white. A conservative action, $\mathbf{a} = (2, 0.5, 0.0651, 0.485)$, is indicated by the black circle, where a and b were determined by estimating the Lipschitz coefficients present in the proof of Theorem 40. The conservative action fails to progress whereas Safety-Aware LineCoSpar provides an action which successfully navigates between obstacles.

(ID-QP) walking controller designed using the concepts in [251], which realizes a stable walking gait for (5.67) at 1 [kHz].

The quadruped is simulated executing the proposed controller with parameters provided by SA-LineCoSpar, and the resulting trajectories and the position of the obstacles¹² are shown in Figure 5.8. Thirty iterations were ran, with 3 new actions sampled in each iteration ($s = 3$), to obtain user preferences and ordinal labels in between each set of actions. To simulate perception error, the measurements of the obstacles were shifted by -0.1 m in the y -direction. The parameters found with SA-LineCoSpar allow the robot to navigate between obstacles. For comparison, a conservative action is also shown, which is safe but fails to progress towards the goal. SA-LineCoSpar eliminates this conservatism with only minor safety violations and determines a parameter set which is both safe and performant.

After simulation, learning was performed with hardware experiments in a lab-

¹²The simulation uses two CBF constraints, one for each obstacle, and both are enforced in the (TR-SOCP) controller.

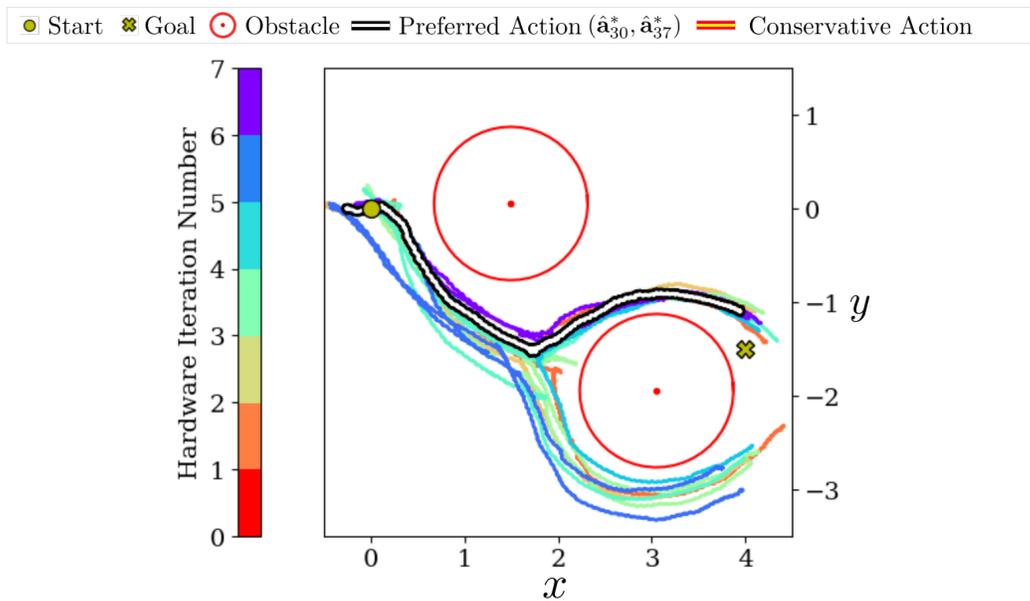


Figure 5.9. Seven additional iterations of 3 actions executed indoors. The preferred action, $\hat{\mathbf{a}}_{37}^* = (4, 0.6, 0.4, 0)$, successfully traverses between the obstacles.

oratory setting for 7 additional iterations until the user was satisfied with the experimental behavior. The robot and obstacle positions were estimated using Intel RealSense T265 and D415 cameras to perform SLAM and segmentation. Centroids of segmented clusters in the occupancy map were used as the measured obstacle positions. The true robot and obstacle positions were obtained for comparison using an OptiTrack motion capture system. The results of these experiments is seen in Figure 5.9. Afterwards, three iterations were conducted outdoors on grass until the user was satisfied with the experimental behavior. The resulting best trajectory can be seen in Figure 5.10. The preferred action was tested on other obstacle arrangements to confirm its generalizability.

Conclusion

In conclusion, in this section I have presented work on using PBL to tune the robustness parameters of CBF-based controllers. I begin by reviewing PBL, and present an algorithm (not my explicit contribution) for safety-aware PBL. Following this, I unify reduced-order CBF methods with ISSf-CBFs and MR-CBFs to produced a second-order cone program-based controller that provides a wide range of robustness properties. Following this, I demonstrate the use of PBL to tune the parameters of this controller for a quadrupedal system until safety and performance are adequately balanced.

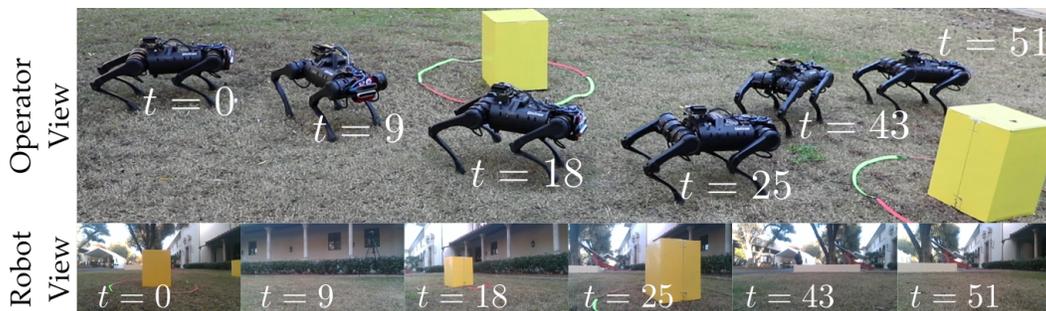


Figure 5.10. The preferred action, $\hat{\mathbf{a}}_{40}^* = (5, 0.1, 0.4, 0.02)$, after simulation, indoor experiments, and 3 additional iterations of 3 actions in an outdoor environment is shown alongside views from the onboard camera. Video can be found at <https://youtu.be/QEuwRDTG7TE>

5.5 Conclusion

In this chapter I have presented a collection of results on designing safety-critical controllers that are robust to measurement errors. In Section 5.2 I look at how measurement errors can arise in a control system and how they present challenges for standard CBF-based design techniques. This leads to the definition of Measurement-Robust Barrier Function (MR-BFs), which certify safety in the presence of measurement errors. These are then extended to Measurement-Robust Control Barrier Functions (MR-CBFs) which can be incorporated to controllers specified via convex second-order cone programs, one of which I explore local Lipschitz continuity properties of. Next, in Section 5.3 I unify MR-CBFs with backup set CBF methods. I begin by reviewing backup set CBF methods, and then show how they can be modified to incorporate MR-CBF constraints, providing a tool for specifying measurement-robust control invariant sets. This leads to the first demonstration of an MR-CBF-based controller on an experimental Segway system using a vision-based measurement system for object detection. Importantly, this result reveals that more complex convex optimization problems can be efficiently solved to the point they can be deployed in real-time robotic applications. Lastly, in Section 5.4, I consider using preference-based learning (PBL) for tuning the parameters of a robust CBF-based controller. I provide a review of PBL, and then consider the setting of safety-critical control with a reduced-order model and measurement errors. By unifying the tools of ISSf-CBFs and MR-CBFs, I provide a parameterized robust CBF-based controller that address the reduced-order and full-order model gap utilizing ISSf guarantees that hold in the presence of measurement errors. As this robust controller using worst-case parameter estimates can be excessively conservative, PBL is deployed to adjust parameters

according to user feedback on performance and safety, producing a safety-critical controller that is robust to a wide set of real-world challenges while remaining performant.

There are a number of potential directions for future work stemming from the contributions in this chapter. One direction would be exploring MR-CBFs in different settings of measurement error than the one considered in Section 5.2. In particular, it is often not possible to determine the entire state \mathbf{x} from a single measurement, but rather, a time history of measurements are used to produce a state estimate $\hat{\mathbf{x}}$ using a model-based observer, such as in [50], [261]. Thus, quantifying the error in an observer for a nonlinear system and integrating that with MR-CBFs may expand their capabilities to a wider range of systems. Furthermore, direct measurements are almost always corrupted by noise, and observers often provide probabilistic state estimates, suggesting that extending MR-CBF guarantees to the stochastic measurement setting would more accurately describe real-world systems. The work explored in Section 4.5 is a potential path forward for this task, but will need to address how probabilistic properties of state estimates are propagated through MR-CBF-based controllers while ideally preserving convexity.

Another direction would continue to explore how PBL can be further integrated into the tuning process of robust controllers. The significant body of work [76]–[80], [281] and the results in Section 5.4 suggests that PBL can successfully capture the underlying preferences of a control system designer to produce desirable closed-loop behavior. This raises the question of how can problems in control be set up parametrically so that PBL can be deployed using user feedback. One example may include parameterizing a reduced-order model for a full-order system and querying feedback from the control designer on the ability of the full-order system to track the behaviors produced by the reduced-order model (seeking to produce the constant μ in Section 5.4). A second example could include parameterizing the model error sets in Section 3.8 instead of constructing them from data, and querying control designer feedback on the performance and robustness of the controller, similar to the results in Section 5.4.

SAMPLED-DATA & EVENT-TRIGGERED CONTROL

In this chapter I will present results from my work developing methods for safety-critical control in the presence of input sampling. Very early feedback control systems, such as governors [288], employed mechanical means for implementing feedback loops, such that the implementation of a controller was intrinsically a continuous process that is accurately captured by the closed-loop model (2.2). The advent of modern compute resources has led control system design to dramatically depart from this type of implementation. Now, the majority of autonomous systems feature a host of computers sampling measurements from various analog and digital sensors, performing numerically-intensive control computations, and sending commands to actuators. Importantly, this process is *not instantaneous*. Rather, there is time required to sample sensors, time required to perform control computations, and time required to send an input command to actuators. During these collective times, the actuators typically continue to perform the previously commanded input, such that inputs are applied with a *sample-and-hold* implementation. This fundamentally differs¹ from the modeling paradigm that CLFs and CBFs utilize in providing rigorous guarantees on stability and safety, and thus, can compromise said guarantees (including the ones developed in Chapters 2-5 of this thesis). Thus, it is critical that a framework for robust safety-critical control address the effects of sampling and ensure that theoretically rigorous stability and safety guarantees can be made.

There are different paradigms for considering the effect of sampling on control systems, including sampled-data control [289], event-triggered control [95], [96], and time-delays in control [290]. Of these, my results will focus on sampled-data control and event-triggered control. The perhaps most natural setting for describing sample-and-hold input implementations is sampled-data control. In this setting, inputs are held constant across a fixed period of time,

¹I note that classical linear system theory focuses on the notion of *bandwidth* to describe the underlying frequency at which things are being performed in the control-loop. The same issues with limited bandwidth for linear systems are bound to occur for nonlinear systems as well, they are just not as easily mathematically quantified.

referred to as a *sample period*, and updated at sequential *sample times*. This is an accurate description for many cyber-physical systems, especially those running a real-time operating systems (RTOS), in which the control-loop is ran at a fixed *sample rate*. Within the sampled-data paradigm, approaches can generally be broken into three categories [289]. The first category focuses on *emulation*, or how a controller designed for a continuous-time system model can be modified in implementation to account for the fact it will be sampled. The second category, and the direction my work will take, is discrete-time design, which uses a discrete-time model of the sampled-data dynamics to design controllers. The last category, focuses on design that is both continuous and discrete-time in nature, but there have been few results in this direction given the complexity of the problem. In contrast, event-triggered control does not assume a fixed sample period, but assumes that the control system can monitor² for *events* while holding the control input constant, and when an event occurs, update the input. While this is perhaps not as natural in the setting of robotic systems, it provides a powerful tool for efficient utilization of resources when there is a cost to changing the input, such as context switching in compute tasks [291], or firing thrusters on a spacecraft [292]. The challenge in developing rigorous event-triggered control paradigms is to ensure that the *trigger-law* defining events leads to some minimum amount of time between events, more commonly known as a minimum interevent time (MIET).

My work in this chapter focuses on CLF and CBF-based control designs for sampled-data systems using *approximate* discrete-time models, as well as developing a method for safety-critical event-triggered control through ISSf-CBFs. In Section 6.1 I will review related work in the area of sampled-data control with CLFs and CBFs as well as safety-critical event-triggered control. Following this, in Section 6.2 I will provide some empirical studies of the effects of sampling on the (CLF-QP) controller for different systems, highlighting its fragility as motivation for my results. I will then provide a mathematical description of sampled-data control and event-triggered control. Next, in Section 6.3, I will present work on sampled-data stabilization using CLFs. I will first review some important structural results regarding feedback linearization and sampled-data control which will underlie the subsequent theory. Then, I will

²For most problems in which event-triggered control is of interest, events can be monitored significantly more quickly than they occur, such that event monitoring is abstracted to be in continuous-time. That being said, a complete addressing of the problem would consider the fact that event monitoring is also a process that involves sampling.

review approximate discrete-time models and practical stability as a goal when designing stabilizing controllers for sampled-data systems. This will lead to the main result of this work, which is a constructive approach for extending feedback linearizability of continuous-time dynamics to a set of approximate discrete-time dynamics that allows the synthesis of practically stabilizing controllers through convex optimization. I will conclude by revisiting the simulation results from Section 6.2 with this control paradigm and demonstrate a significant improvement in closed-loop stability properties.

In Section 6.4 I will develop a framework for safety-critical control of sampled-data systems with CBFs and approximate discrete-time models. I will begin by considering a more general class of approximate discrete-time models that allow working with CBFs that do not possess a relative degree of one. Next, I will define practical safety and Sampled-Data Barrier Functions (SD-BFs) and Sampled-Data Control Barrier Functions (SD-CBFs), and show how SD-BFs can be used to certify practical safety of a sampled-data system. I will conclude by exploring how convexity properties of a SD-CBF lead to a convex optimization-based controller that I demonstrate in simulation. Lastly, in Section 6.5, I will consider the problem of safety-critical event-triggered control using ISSf-CBFs. I will briefly review event-triggered control for stability as in [96] to motivate the approach for event-triggered control for safety. Then, I will consider how errors introduced by an event-triggered control scheme can be captured by ISSf-BFs when considering safety. I will consider a trigger-law similar to that proposed for stability in [96], and show that it fails to produce a MIET through a counterexample, highlighting the importance of how a system behaves near the boundary of a set that is to be kept forward invariant. Following this, I will modify the trigger-law slightly to achieve safety-critical event-triggered control while ensuring the existence of a MIET, which I demonstrate in simulation. Key contributions of this work are described at the beginning of each respective section.

6.1 Related Work

Sampled-Data Stabilization

A key overview and starting point for exploring sampled-data control for non-linear systems is the review in [289]. My review of the literature will focus in particular on discrete-time design with approximate models, its relationship with feedback linearizability, and methods that have used CLFs. A key

challenge in discrete-time design for sampled-data nonlinear systems is that it is difficult to produce a computationally tractable representation of the *exact* discrete-time dynamics of the sampled-data system. Finding an analytic expression requires solving a general nonlinear differential equation (known to be quite a difficult task), and numerically integrating the differential equations poses computational challenges in synthesis that have been prohibitive (until perhaps recently with modern compute resources, such as with the backup set CBF methods in Section 5.3). The idea to instead use approximations of the discrete-time sampled-data dynamics was proposed early in the pursuit of controllers for nonlinear sampled-data systems [293]–[295]. A key development in this direction was the sequence of work in [81], [82], [296], which built a framework which identified properties of both approximation schemes and controllers that ensure a property known as *practical stability*. This framework permitted simple approximations such as first-order Euler methods (which I will frequently use in this chapter), which greatly simplified the synthesis problem. This led to a wealth of work extending control techniques such as backstepping [83], [84], Lyapunov redesign [85], and finite-horizon optimization-based schemes (such as Model Predictive Control (MPC)) [86]–[88]. One of the observations of this work was that sampled-data designs with approximate discrete-time models often significantly outperform their counterparts designed with continuous-time models, and even successfully operate at low sample rates [94]. This will be something that I observe in Sections 6.2 and 6.3 with my proposed method. I also note that the notion of practical stability was extended to corresponding ideas of ISS [297] and integral ISS [298].

Given the constructive relationship between feedback linearization and CLFs established in Chapter 2, CLF-based controllers developed for sampled-data nonlinear systems should account for the interactions between feedback linearization and sampling. The relationship between feedback linearization and sampled-data control was a deeply studied topic in the early formulation of controllers for nonlinear sampled-data systems [293], [294], [299]–[302]. A key result from this work appeared in [293]. In particular, this work identified that feedback linearizability and sampling *do not commute*. In particular, there were three critical observations.

- Even if the continuous-time dynamics of a system are feedback linearizable, the exact discrete-time sampled dynamics are not necessarily feed-

back linearizable.

- The feedback linearizability of approximate discrete-time dynamics for feedback linearizable continuous-time dynamics can depend on the order of the approximation used.
- The feedback linearizability of approximate discrete-time dynamics for feedback linearizable continuous-time dynamics can depend on the coordinates in which the continuous-time dynamics are expressed before being approximately discretized.

The second and third point of this preceding list will be important in Section 6.3, as producing CLF-based controllers will rely on the order of the approximate discrete-time dynamics considered, as well as what coordinates the discretization is performed in. I note that zero-dynamics that arise due to sampling and higher-order approximations have been studied [303]–[305], but the effects of sampling on the stability of zero-dynamics for a continuous-time system that is feedback linearizable have yet to be studied (I will consider them in Section 6.3). Lastly, sampled-data control with CLFs has been explored in [306], [307], but it did not make use of approximate discrete-time models or convex optimization as I will do in my work.

My work in [115], presented in Section 6.3, will translate the pipeline from feedback linearization to the optimization-based (CLF-QP) controller in Chapter 2 to the sampled-data setting by integrating with the sampled-data stabilization framework through approximate discrete-time models proposed in [81], thereby yielding a constructive method for optimization-based stabilizing control for sampled-data systems. To do this, I will first review important tools from [81] regarding approximate discrete-time models and practical stability, and show that properties of these tools are preserved when considering coordinate transformations needed to preserve feedback linearizability between continuous-time and approximate discrete-time models. I will then show that a CLF synthesized for a feedback linearizable continuous-time system still satisfies the necessary properties to permit CLF-based control synthesis for the approximate discrete-time dynamics. Given a controller from this process, I will then show that if the zero-dynamics of the continuous-time feedback linearizable system are locally exponentially stable, that stability is preserved with sampling. Lastly, I will use the CLF to propose a novel optimization-

based controller specified as a convex quadratically constrained quadratic program (QCQP). I note that a convex QCQP is a more complex class of convex optimization problem than a QP, but is a sub-class of the class of second-order cone programs (SOCPs) seen throughout this thesis. Lastly, I demonstrate the significant improvement in closed-loop behavior of this controller over deploying a continuous-time (CLF-QP) controller with a sampled implementation on the problem settings considered in Section 6.2.

Safety-Critical Sampled-Data Control

The development of safety-critical controllers for sampled-data systems is a significantly more nascent field than the development of stabilizing controllers, in large part due only to the recent development of CBFs in [11], [62]. The first work looking at the effect of sampled-and-hold inputs through the perspective of barrier functions [308], which considered only the setting of a double integrator. Subsequently, a wide collection of work was developed for safety-critical control of sampled-data systems [19], [73], [74], [89]–[93]. Much of this work takes an emulation approach, where an additive correction term is factored in to the CBF inequality [19], [73], [90]–[93]. This correction term often relies on Lipschitz constants of the system dynamics and bounds on the control inputs, which can be difficult to produce, and over-approximations will often lead to conservative behavior that can only be resolved with exceptionally high sample rates, as studied in [90]. The work in [74] takes a computationally intensive approach to reduce conservatism by propagating sensitivity functions, which may be difficult for high-dimensional systems. Notably, none of this work considers using discrete-time designs with approximate discrete-time models.

My work in [309], presented in Section 6.4, is the first to consider safety-critical control using CBFs and discrete-time design with approximate discrete-time models for nonlinear sampled-data systems. In particular, it will draw inspiration from the framework in [81] and my preceding work in Section 6.3. I will begin by establishing that the class of higher-order Runge-Kutta approximation schemes satisfy the necessary consistency properties to provide rigorous safety-critical control guarantees. Then, by considering the notions of practical stability and corresponding Lyapunov functions in [81] and discrete-time CBFs in [234], I will produce a notion of practical safety and corresponding Sampled-Data barrier functions (SD-BFs) that certify practical safety. I will then extend this to a definition of a Sampled-Data Control Barrier Function

(SD-CBF). To ensure that SD-CBFs can be used in convex optimization-based control synthesis, I will establish a result relating convexity properties of an SD-CBF with the order of a Runge-Kutta approximate discrete-time model to ensure that the inequality constraint that leads to practical safety is convex with respect to the input, enabling its use in convex optimization-based controllers. Lastly, in simulation I will demonstrate the efficacy of this method in ensuring practical safety of nonlinear sampled-data systems, including the ability to operate at low sample rates without being conservative.

Event-Triggered Safety-Critical Control

Event-triggered control provides a framework that allows the prescription of, in a principled way, when certain resources (such as actuators, sensors, access to a communication network with neighboring agents, or even a human) should be utilized in order to guarantee the satisfaction of said control objectives efficiently. The first exploration into advantages of event-triggered control was the work in [310], [311], which was later expanded to a collection of results focused on stabilization [96], [312]–[314]. Following this, a number of frameworks and generalizations have been proposed to extend the capabilities and improve the performance of event-triggered control schemes [95], [97], [315], [316]. A key observation in the event-triggered control literature is the need to ensure the existence of a minimum interevent time (MIET) [97], [317], [318], thereby preventing the accumulation of events in “Zeno-like” like behavior. One notable investigation into output-based event-triggered control observed challenges in ensuring an MIET if the entire system state did not converge to an equilibrium point [319], which is a challenge that will translate into my work on safety-critical event-triggered control.

Early work in safety-critical event-triggered control considered using Barrier-Lyapunov functions, which unify both stability and safety into one function, and enforce forward invariance by requiring stability [320]. The first work that looked at using BFs and CBFs in the context of event-triggered control were focused on improving the stabilizing properties of a controller, rather than enforcing the forward invariance of a particular set in the state space [89], [98]. Taking inspiration from the approach for event-triggered stabilization using ISS in [96], my work in [321], presented in Section 6.5, was the first to explicitly explore safety-critical event-triggered control (without requiring stabilization) through ISSf. My work shows that naïvely extending the ap-

proach in [96] does not lead to a MIET using a counterexample that does not converge to an equilibrium point, similarly to the issues seen in [319]. I instead propose an alternative triggering scheme that ensures that a system remains safe while possessing an MIET, which I both prove and show in simulation. There have been follow-on to this work that has considered relaxing assumptions on boundedness of dynamics (made locally in my work, but a strong assumption to take globally) [322], safety-critical event-triggered control with model error [323], and greedy approaches to maximize event separation [324]. Similarly, there have been a number of applications of the ideas in this work to different settings, including spacecraft orbit adjustment [292], connected automated vehicles [325], and a Stefan partial differential equation system with actuator dynamics [326].

6.2 Sampling Paradigms

In this section I will present some motivating simulation results highlighting the fragility of the (CLF-QP) controller to sampling on a set of systems, and then provide a mathematical description of the sampled-data and event-triggered sampling schemes. I will first consider the inverted pendulum system:

$$\ddot{\theta} = \frac{g}{\ell} \sin(\theta) + \frac{1}{m\ell^2}u, \quad (6.1)$$

with angle θ , angle rate $\dot{\theta}$, gravity g , length ℓ , and mass m . The point $\mathbf{x}_e = \mathbf{0}_2$ is an unforced equilibrium point of this system, and the system is full-state feedback linearizable to this point. Thus, a CLF V can be synthesized as in Section 2.4. In this example, I will use a matrix $\mathbf{Q}_2 = 0.99\mathbf{Q}$, where \mathbf{Q} is used to find the matrix \mathbf{P} via the CTLE. I will deploy the (CLF-QP) controller on this system with a sample-and-hold implementation at a variety of different sample rates, as seen in Figure 6.1. Observe that at 10 [Hz], the controller is not stabilizing. At 20 [Hz], the controller does stabilize the system, but it displays undesirable properties where the control input periodically drops to 0 before reactivating. At 50 [Hz], this controller is stabilizing and the input profile does not display undesirable behavior. This aligns with the mentality of “*just run it fast enough and it will work*” that controllers designed with continuous-time models are often implemented with.

Now let us consider the more complicated system of a double inverted pendu-

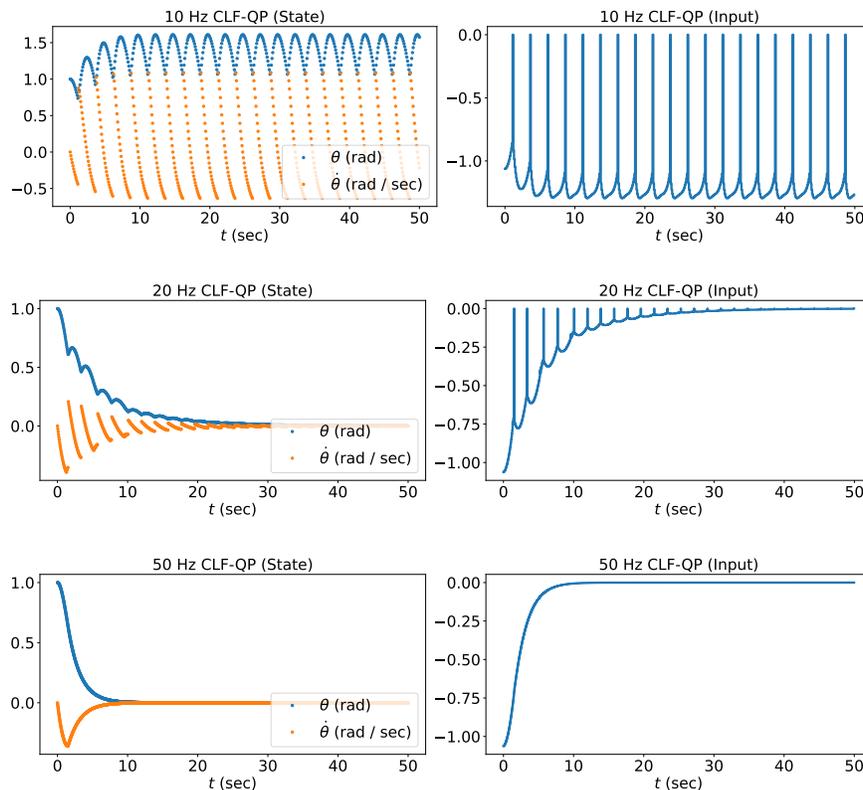


Figure 6.1. Results using sample-and-hold implementation of (CLF-QP) controller on inverted pendulum at various sample rates.

lum, seen in Figure 6.2, with the dynamics:

$$\begin{aligned}
 & \begin{bmatrix} (m_1 + m_2)\ell_1^2 + m_2\ell_2^2 + 2m_2\ell_1\ell_2 \cos(\theta_2) & m_2\ell_2^2 + m_2\ell_1\ell_2 \cos(\theta_2) \\ m_2\ell_2^2 + m_2\ell_1\ell_2 \cos(\theta_2) & m_2\ell_2^2 \end{bmatrix} \begin{bmatrix} \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{bmatrix} \\
 & + \begin{bmatrix} 0 & -m_2\ell_1\ell_2(2\dot{\theta}_1 + \dot{\theta}_2) \sin(\theta_2) \\ \frac{1}{2}m_2\ell_1\ell_2(2\dot{\theta}_1 + \dot{\theta}_2) \sin(\theta_2) & -\frac{1}{2}m_2\ell_1\ell_2\dot{\theta}_1 \sin(\theta_2) \end{bmatrix} \begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \end{bmatrix} \\
 & + g \begin{bmatrix} (m_1 + m_2)\ell_1 \sin(\theta_1) + m_2\ell_2 \sin(\theta_1 + \theta_2) \\ m_2\ell_2 \sin(\theta_1 + \theta_2) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}. \quad (6.2)
 \end{aligned}$$

This system has an unforced equilibrium point $\mathbf{x}_e = \mathbf{0}_4$ to which it is full-state feedback linearizable. This can be used to synthesize a CLF (using $\mathbf{Q}_2 = 0.99\mathbf{Q}$) and a corresponding (CLF-QP) controller. The results of deploying this controller on the system with a sample-and-hold implementation at various frequencies can be seen in Figure 6.3. Observe that this controller fails to stabilize the system at sample rates of 10 [Hz], 100 [Hz], and even at 1000 [Hz], which is likely beyond a reasonable sample rate that can be achieved on lightweight compute platforms. Moreover, the system is wholly unstable

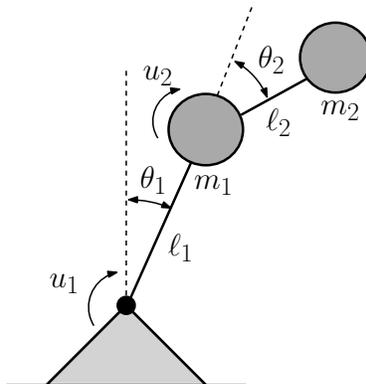


Figure 6.2. Double Inverted Pendulum Schematic

and displays very frequency high oscillations in the input that would not be possible with real actuators (such as motors that have inertia).

Lastly, I will consider a planar quadrotor system, as seen in Figure 6.4. This system has dynamics given by:

$$\ddot{x} = \frac{f_t}{m} \sin(\theta), \quad (6.3)$$

$$\ddot{z} = -g + \frac{f_t}{m} \cos(\theta), \quad (6.4)$$

$$\ddot{\theta} = \frac{1}{J} \tau, \quad (6.5)$$

with horizontal position x , vertical position z , horizontal velocity \dot{x} , vertical velocity \dot{z} , roll angle θ , roll angle rate $\dot{\theta}$, thrust force f_t , torque τ , mass m , inertia J , and gravity g . If this system is dynamically extended by treating f_t and \dot{f}_t as states and the inputs to the system as \dot{f}_t and τ , the system is full-state feedback linearizable [327] to the unforced equilibrium point $\mathbf{x}_e = [0 \ 0 \ z_e \ 0 \ 0 \ 0 \ mg \ 0]^\top$ for any $z_e \in \mathbb{R}_{>0}$. This permits the constructive synthesis of a CLF. Once again synthesizing a CLF with $\mathbf{Q}_2 = 0.99\mathbf{Q}$, the results of using the (CLF-QP) controller deployed with a sample-and-hold implementation can be see in Figure 6.5 Observe that at both 10 [Hz] and 100 [Hz] the controller fails to stabilize the system. At 500 [Hz] the system is stabilized, but the controller displays large, high frequency oscillations that would likely be undesirable (or impossible) on a real-world hardware platform.

These simulation results suggest that not only does a sample-and-hold implementation compromise theoretical stability guarantees, but the degree of the failures also indicate that the sample rates needed to justify the mentality of

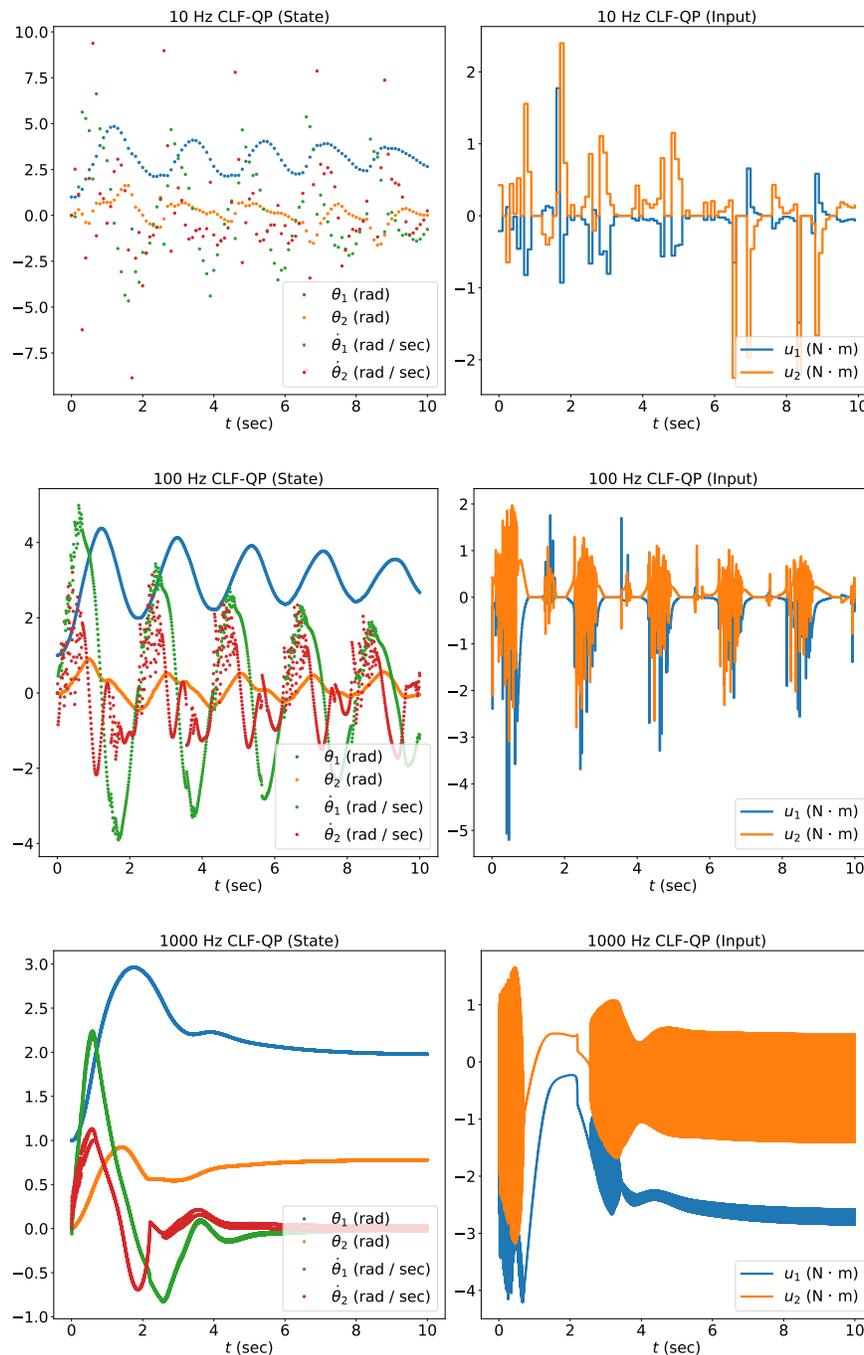


Figure 6.3. Results using sample-and-hold implementation of (CLF-QP) controller on double inverted pendulum at various sample rates.

“just run it fast enough and it will work” for these continuous-time CLF-based controllers may not be obtainable on real-world hardware platforms. In Section 6.3 I will revisit each of these examples using a controller using discrete-time approximations, and demonstrate the ability of said control approach to work at every one of the frequencies presented in these results (including 10 [Hz]).

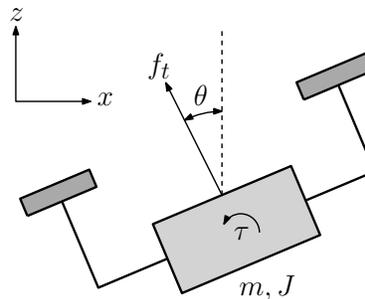


Figure 6.4. Planar Quadrotor Schematic

This supports the idea that addressing how a controller is implemented in its design is vital for achieving desirable behavior.

Sampled-Data Sampling

I now describe the sampled-data control paradigm for describing systems with input sampling. Consider the open-loop system (2.1), and consider an open subset $\mathcal{Z} \subseteq E \times \mathbb{R}^m$ and its projection onto the state space:

$$\mathcal{X} \triangleq \{\mathbf{x} \in E \mid \exists \mathbf{u} \in \mathbb{R}^m \text{ s.t. } (\mathbf{x}, \mathbf{u}) \in \mathcal{Z}\} \subseteq E. \quad (6.6)$$

I will assume that there exists a $T_{\max} \in \mathbb{R}_{>0}$ such that for every state-input pair $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$, there exists a unique continuously differentiable solution $\varphi : [0, T_{\max}] \rightarrow E$ satisfying:

$$\dot{\varphi}(t) = \mathbf{f}(\varphi(t)) + \mathbf{g}(\varphi(t))\mathbf{u}, \quad (6.7)$$

$$\varphi(0) = \mathbf{x}, \quad (6.8)$$

for all $t \in (0, T_{\max})$. Given a sample period $T \in (0, T_{\max}]$, a controller $\mathbf{k} : \mathcal{X} \rightarrow \mathbb{R}^m$ is *T-admissible* if for any $\mathbf{x} \in \mathcal{X}$, the state-input pair $(\mathbf{x}, \mathbf{k}(\mathbf{x}))$ satisfies $(\mathbf{x}, \mathbf{k}(\mathbf{x})) \in \mathcal{Z}$ and the corresponding solution φ satisfies $\varphi(t) \in \mathcal{X}$ for all $t \in [0, T]$. This requirement on *T-admissible* controllers will ensure that in the sampled-data context, the closed-loop system is forward complete and its evolution may be described by iterative solutions to (6.7) and (6.8). Though verifying *T-admissibility* of a controller may be intractable, assuming that a controller is *T-admissible* and renders the set \mathcal{X} invariant is relatively weak as \mathcal{X} is defined to ensure the continued existence of solutions rather than reflecting a task-specific set.

The preceding construction of solutions and admissible controllers describes the sampled-data control setting, in which inputs are applied to the system

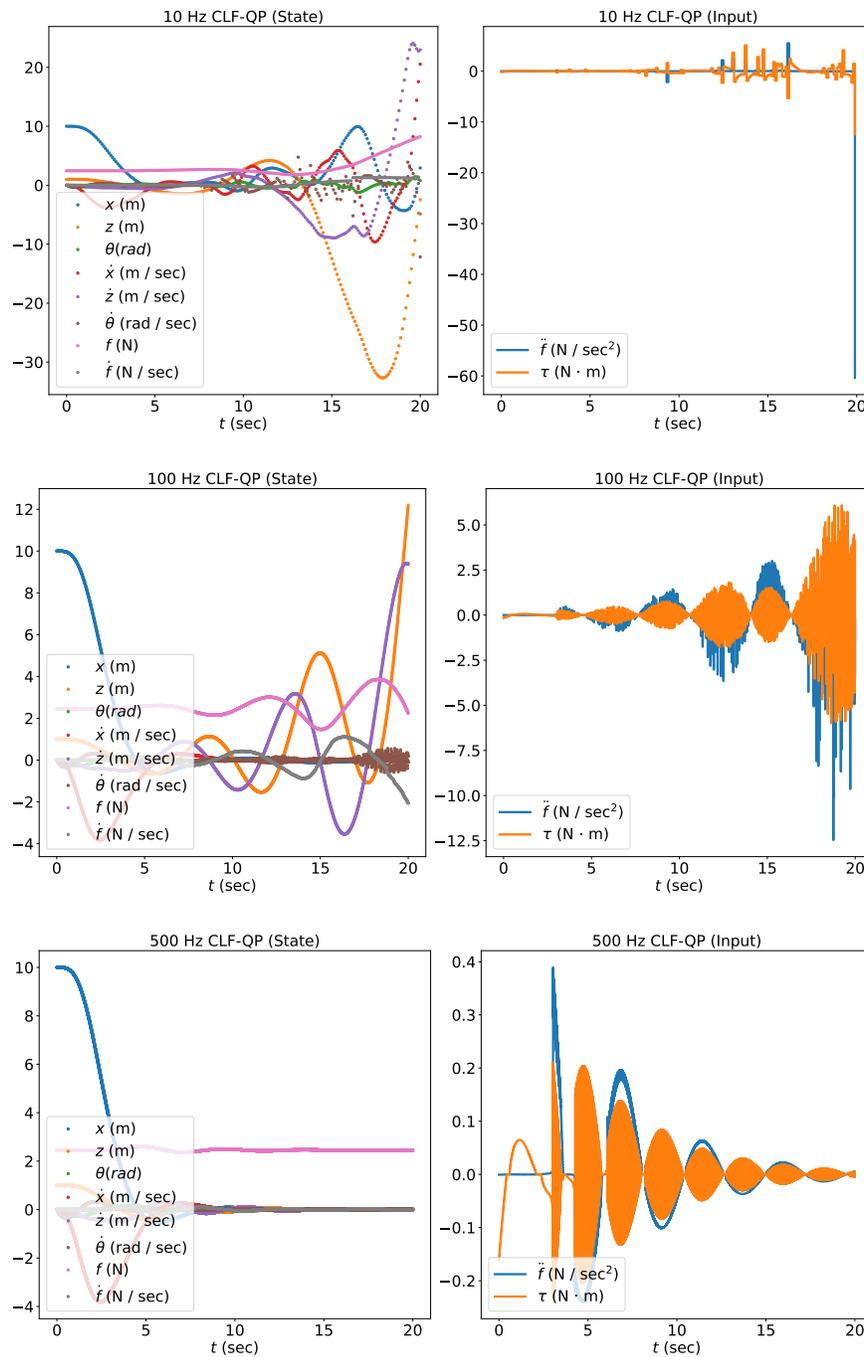


Figure 6.5. Results using sample-and-hold implementation of (CLF-QP) controller on planar quadrotor at various sample rates.

with a zero-order hold over a sample period. More precisely, the set of possible sample periods is given by $I = (0, T_{\max}]$. Given a sample period $T \in I$ and a T -admissible controller $\mathbf{k} : \mathcal{X} \rightarrow \mathbb{R}^m$, we have that for any initial condition $\mathbf{x}_0 \in \mathcal{X}$, there exists a unique solution $\varphi_{\text{SD}} : \mathbb{R}_{\geq 0} \rightarrow \mathcal{X}$ that is piecewise

continuously differentiable on $\mathbb{R}_{\geq 0}$ and satisfies:

$$\dot{\varphi}_{\text{SD}}(t) = \mathbf{f}(\varphi_{\text{SD}}(t)) + \mathbf{g}(\varphi_{\text{SD}}(t))\mathbf{k}(\varphi_{\text{SD}}(t_k)), \quad \forall t \in [t_k, t_{k+1}) \quad (6.9)$$

$$\varphi_{\text{SD}}(0) = \mathbf{x}_0, \quad (6.10)$$

with sample times satisfying $t_{k+1} - t_k = T$ for all $k \in \mathbb{Z}_{\geq 0}$. Given an $\mathbf{x}_0 \in \mathcal{X}$, I will denote the solution at sample times as $\mathbf{x}_k \triangleq \varphi_{\text{SD}}(t_k)$ for $k \in \mathbb{Z}_{\geq 0}$. An alternative way of representing the evolution of the system over a sample period is given by the *exact map* $\mathbf{F}_T^e : \mathcal{Z} \rightarrow \mathbb{R}^n$:

$$\mathbf{F}_T^e(\mathbf{x}, \mathbf{u}) = \mathbf{x} + \int_0^T [\mathbf{f}(\varphi(\tau)) + \mathbf{g}(\varphi(\tau))\mathbf{u}] \, d\tau, \quad (6.11)$$

for all state-input pairs $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$. For all $T \in I$ such that \mathbf{k}_T is T -admissible and any initial condition $\mathbf{x}_0 \in \mathcal{X}$, the solution φ_{SD} satisfies the recursion:

$$\mathbf{x}_{k+1} = \mathbf{F}_T^e(\mathbf{x}_k, \mathbf{k}_T(\mathbf{x}_k)) \in \mathcal{X}, \quad (6.12)$$

for all $k \in \mathbb{Z}_{\geq 0}$. Closed-form expressions for exact maps are rarely obtainable, suggesting the use of approximations in control synthesis, which I will consider in Sections 6.3 and 6.4. I call $\{\mathbf{k}_T : \mathcal{X} \rightarrow \mathbb{R}^m \mid T \in I\}$ a *family of admissible controllers* if there is a $T^* \in I$ such that for each $T \in (0, T^*)$, \mathbf{k}_T is T -admissible. This enables the following definition:

Definition 44 (*Exact Family*). I define the *exact family of maps* $\{\mathbf{F}_T^e \mid T \in I\}$, and for a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$, I define the *exact family of controller-map pairs* $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$.

Lastly, if $\mathbf{x}_e \in \mathcal{X}$ is an unforced or forced equilibrium point of the open-loop system (2.1) for an equilibrium input $\mathbf{u}_e \in \mathbb{R}^m$ such that $(\mathbf{x}_e, \mathbf{u}_e) \in \mathcal{Z}$, then \mathbf{x}_e is said³ to be an unforced or forced equilibrium point of the exact family of maps $\{\mathbf{F}_T^e \mid T \in I\}$, noting that:

$$\mathbf{F}_T^e(\mathbf{x}_e, \mathbf{u}_e) = \mathbf{x}_e. \quad (6.13)$$

³Throughout this work I will require that a point be an unforced or forced equilibrium point of an underlying continuous-time open-loop system to define it as an unforced or forced equilibrium point of a discrete-time family of maps. Notably, just satisfying a condition (6.13) will not define an equilibrium point for a discrete-time family of maps, as it does not say anything about equilibrium behavior in the inter-sample behavior. Rather, relationships like (6.13) are a consequence of the definition I take.

Given a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$, the point \mathbf{x}_e is said to be an equilibrium point of the exact family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ if there exists a $T^* \in I$ such that:

$$\mathbf{k}_T(\mathbf{x}_e) = \mathbf{u}_e, \quad (6.14)$$

for all $T \in (0, T^*)$.

Event-Triggered Sampling

I now describe the event-triggered control paradigm for systems with input sampling. I will assume that for every state-input pair⁴ $(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^n \times \mathbb{R}^m$, there exists a unique continuously differentiable solution $\varphi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ satisfying:

$$\dot{\varphi}(t) = \mathbf{f}(\varphi(t)) + \mathbf{g}(\varphi(t))\mathbf{u}, \quad (6.15)$$

$$\varphi(0) = \mathbf{x}, \quad (6.16)$$

for all $t \in \mathbb{R}_{\geq 0}$. The assumption that solutions exist for all time may appear strong, but note that the control input is fixed, and thus this is mainly a requirement that the functions \mathbf{f} and \mathbf{g} won't lead to finite escape times. Importantly, it will later permit arbitrarily large spacing between events, which is a desirable property in event-triggered control.

The implementation of a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is done by sampling the state at sample times t_k for $k \in \mathbb{Z}_{\geq 0}$ with $t_0 = 0$ and $t_k \leq t_{k+1}$, evaluating the controller on the corresponding state $\mathbf{x}(t_k)$, and holding the input constant until the next sample-time. This leads to an event-triggered closed-loop system:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t_k)) \quad \forall t \in [t_k, t_{k+1}), \quad (6.17)$$

where the argument t is used for \mathbf{x} to denote that it continues to evolve continuously in time while the input $\mathbf{k}(\mathbf{x}(t_k))$ is held constant. As opposed to the sampled-data sampling paradigm, the sample times at which the controller is updated are determined by a state-dependent execution rule or trigger-law:

$$t_{k+1} = \min\{t \geq t_k \mid \rho(\mathbf{x}(t), t_k) = 0\}, \quad (6.18)$$

for a function $\rho : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$. Given a trigger-law, for any initial condition $\mathbf{x}_0 \in \mathbb{R}^n$, there exists an interval $I(\mathbf{x}_0) \triangleq [0, t_{\max}(\mathbf{x}_0))$, sample times $t_k \in \mathbb{R}_{\geq 0}$

⁴I work with \mathbb{R}^n instead of an open set $E \subseteq \mathbb{R}^n$ with this topic for simplicity. This is not needed, but otherwise requires some additional labor maintaining function domains that is not particularly illuminating.

with $t_0 = 0$ and $t_k \leq t_{k+1} < t_{\max}(\mathbf{x}_0)$ for all $k \in \mathbb{Z}_{\geq 0}$, and a function $\varphi_{\text{ET}} : I(\mathbf{x}_0) \rightarrow \mathbb{R}^n$ that is piecewise continuously differentiable on $I(\mathbf{x}_0)$ and satisfies:

$$\dot{\varphi}_{\text{ET}}(t) = \mathbf{f}(\varphi_{\text{ET}}(t)) + \mathbf{g}(\varphi_{\text{ET}}(t))\mathbf{k}(\varphi_{\text{ET}}(t_k)), \quad \forall t \in [t_k, t_{k+1}) \quad (6.19)$$

$$\varphi_{\text{ET}}(0) = \mathbf{x}_0. \quad (6.20)$$

There are two interesting behaviors to observe regarding this paradigm. First, it maybe possible for t_{k+1} to be undefined ($t_{k+1} = t_{\max} = \infty$) for all k greater than some $k^* \in \mathbb{Z}_{\geq 0}$ if $\rho(\varphi_{\text{ET}}(t))$ never equals 0 for $t \in [t_k, \infty)$. This is actually a desirable probably in the event-triggered paradigm, as it means the control input to the system does not need to be changed (with switching assumed to have some undesirable cost). The other, more pressing behavior is if t_{k+1} can become arbitrarily close to t_k as $k \rightarrow \infty$, such that there is no minimum interevent time (MIET). This is described as Zeno behavior⁵, and is undesirable as it requires the controller essentially be implemented continuously in time, which not only defeats the purpose of an event-triggered implementation, but is likely not possible on real-world hardware platforms.

6.3 Sampled-Data Stabilization

In this section I will present work on stabilizing sampled-data control through feedback linearization and CLF-based controllers specified via convex optimization problems. I will begin by extending the sampled-data sampling paradigm in Section 6.2 to account for the coordinate transformations that occur with feedback linearization. Next I will consider approximate discrete-time models, with a focus on the Euler approximation. I review an important property for relating the accuracy of approximate discrete-time models known as one-step consistency, and show how regularity properties on the dynamics and controller lead to one-step consistency. Next, I will review the notion of practical stability and asymptotic stability by equi-Lipschitz Lyapunov functions, which allows for establishing stability results for sampled-data systems based on approximate discrete-time models. Following this, in the first main contribution of this work, I show that a feedback linearizable continuous-

⁵What I will show in the counterexample in Section 6.5 is Zeno-like behavior. Zeno behavior requires that the t_k converge to a finite $t_{\max}(\mathbf{x}_0)$ as $k \rightarrow \infty$. Instead, I will not require t_k converge to a limit, but I will show that for any $\epsilon \in \mathbb{R}_{>0}$, there exists some $k \in \mathbb{Z}_{\geq 0}$ such that $|t_{k+1} - t_k| < \epsilon$, which is equally undesirable. The distinction is best understood by the two series $\sum_{k=1}^{\infty} \frac{1}{k}$ and $\sum_{k=1}^{\infty} \frac{1}{k^2}$. The terms in each series converge to 0, but the first series does not converge (the behavior I achieve) while the second series converges (Zeno behavior).

time system with locally exponentially stable zero-dynamics can be rendered practically stable with a sample-data implementation of the continuous-time feedback linearizing controller. After this, I will present the second main contribution of this work by extending the results from feedback linearization to produce CLF-based controllers specified via convex optimization for sampled-data systems. I conclude by demonstrating the proposed control approach on the simulation environments in Section 6.2, demonstrating significant improvements over continuous-time counterparts.

The contributions of this section are as follows:

- Adaptation of the framework for stabilizing sampled-data control with approximate discrete-time models [81] to support the coordinate transformations and local stability results that are common with feedback linearization and zero-dynamics, leading to a full system stability result for a sampled-data system using feedback linearization.
- A method for synthesizing CLF-based controllers specified as convex optimization problems for sampled-data systems that greatly outperforms continuous-time counterparts in simulation.

The text for this section is adapted from:

A. J. Taylor, V. D. Dorobantu, Y. Yue, P. Tabuada, and A. D. Ames, “Sampled-data stabilization with control lyapunov functions via quadratically constrained quadratic programs,” *IEEE Control Sys. Let.*, vol.6, pp.680–685, 2022.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Feedback Linearization for Sampled-Data Control

I will begin by expanding the sampled-data sampling paradigm in Section 6.2 to account for the coordinate transformation that accompanies feedback linearization. First, define the following reachable set using the solution satisfying (6.7) and (6.8):

$$\mathcal{D} \triangleq \{ \mathbf{x} \in E \mid \exists (\mathbf{x}_0, \mathbf{u}_0) \in \mathcal{Z}, t \in [0, T_{\max}] \text{ s.t. } \mathbf{x} = \boldsymbol{\varphi}(t) \}, \quad (6.21)$$

noting that $\mathcal{X} \subseteq \mathcal{D}$. This set can be interpreted as the region of the state space which can be reached by the continuous-time solution φ within T_{\max} time from an initial condition $\mathbf{x}_0 \in \mathcal{X}$ with an input \mathbf{u}_0 that satisfies $(\mathbf{x}_0, \mathbf{u}_0) \in \mathcal{Z}$. Let $\mathbf{x}_e \in \mathcal{X}$ be an unforced or forced equilibrium point of the open-loop system (2.1) such that $(\mathbf{x}_e, \mathbf{u}_e) \in \mathcal{Z}$, and recalling the definition of feedback linearization in Section 2.4, assume that the open-loop system (2.1) is locally feedback linearizable with respect to the equilibrium point \mathbf{x}_e with an open set D that satisfies $\mathcal{D} \subseteq D$. This additional assumption that $\mathcal{D} \subseteq D$ is made simply so that the diffeomorphism Φ is well defined on all of the states that can be encountered when evolving from an initial condition $\mathbf{x}_0 \in \mathcal{X}$ with an input \mathbf{u}_0 satisfying $(\mathbf{x}_0, \mathbf{u}_0) \in \mathcal{Z}$, and may require considering a restricted set \mathcal{Z} that restricts \mathcal{X} to achieve $\mathcal{D} \subseteq D$.

As shown in [293], feedback linearizability of a continuous-time system does not guarantee feedback linearizability of the resulting discrete-time sampled-data system, even when using approximate discrete-time models. In particular, this property may be lost due to a change of coordinates. The preservation of this property motivates studying the evolution of the normal-form open-loop system in the sampled-data context. To this end, define the set:

$$\mathcal{Z}_\xi = \{(\xi, \mathbf{u}) \in \Phi(\mathcal{X}) \times \mathbb{R}^m \mid (\Phi^{-1}(\xi), \mathbf{u}) \in \mathcal{Z}\}, \quad (6.22)$$

noting that by assumption, for every state-input pair $(\xi_0, \mathbf{u}_0) \in \mathcal{Z}_\xi$, there exists a unique solution $\psi : [0, T_{\max}] \rightarrow \Phi(\mathcal{D})$ satisfying:

$$\dot{\psi}(t) = \mathbf{f}_\xi(\psi(t)) + \mathbf{g}_\xi(\psi(t))\mathbf{u}_0, \quad (6.23)$$

$$\psi(0) = \xi_0, \quad (6.24)$$

for all $t \in (0, T_{\max})$. For $T \in (0, T_{\max}]$, a controller $\mathbf{k} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$ is a T -admissible controller if the corresponding controller $\mathbf{k}' : \mathcal{X} \rightarrow \mathbb{R}^m$ given by $\mathbf{k}'(\mathbf{x}) = \mathbf{k}(\Phi(\mathbf{x}))$ for all $\mathbf{x} \in \mathcal{X}$ is T -admissible. A controller $\mathbf{k}_{\text{aux}} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^k$ is a T -admissible auxiliary controller if $\mathbf{k}_{\text{fbl}}^\xi : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$ given by:

$$\mathbf{k}_{\text{fbl}}^\xi(\xi) = \mathbf{k}_{\text{fbl}}(\Phi^{-1}(\xi), \mathbf{k}_{\text{aux}}(\xi)), \quad (6.25)$$

is a T -admissible controller.

Because we are working with two different coordinate systems (state coordinates and normal-form coordinates), the evolution of the system over a sample

period can be described by both the *exact state discrete map* $\mathbf{F}_T^{e,\mathbf{x}} : \mathcal{Z} \rightarrow \mathcal{D}$ and *exact normal discrete map* $\mathbf{F}_T^{e,\xi} : \mathcal{Z}_\xi \rightarrow \Phi(\mathcal{D})$:

$$\mathbf{F}_T^{e,\mathbf{x}}(\mathbf{x}_0, \mathbf{u}_0) = \mathbf{x}_0 + \int_0^T [\mathbf{f}(\varphi(\tau)) + \mathbf{g}(\varphi(\tau))\mathbf{u}_0] d\tau, \quad (6.26)$$

$$\mathbf{F}_T^{e,\xi}(\xi_0, \mathbf{u}_0) = \xi_0 + \int_0^T [\mathbf{f}_\xi(\psi(\tau)) + \mathbf{g}_\xi(\psi(\tau))\mathbf{u}_0] d\tau, \quad (6.27)$$

for all state-input pairs $(\mathbf{x}_0, \mathbf{u}_0) \in \mathcal{Z}$ and all normal state-input pairs $(\xi_0, \mathbf{u}_0) \in \mathcal{Z}_\xi$. The exact maps are related by:

$$\mathbf{F}_T^{e,\xi}(\xi_0, \mathbf{u}_0) = \Phi(\mathbf{F}_T^{e,\mathbf{x}}(\Phi^{-1}(\xi_0), \mathbf{u}_0)), \quad (6.28)$$

for all normal state-input pairs $(\xi_0, \mathbf{u}_0) \in \mathcal{Z}_\xi$. While an equivalence between the exact state discrete map and exact normal discrete map is achieved via the diffeomorphism Φ , it is useful to define both maps as the notion of stability I later consider for sampled-data systems is defined for a particular exact map. I call $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$ a *family of admissible controllers* if there is an $T^* \in I$ such that for each $T \in (0, T^*)$, \mathbf{k}_T is T -admissible. This enables the following definition:

Definition 45 (*Exact Families*). I define the *exact state family* of maps $\{\mathbf{F}_T^{e,\mathbf{x}} \mid T \in I\}$ and *exact normal family* of maps $\{\mathbf{F}_T^{e,\xi} \mid T \in I\}$, and for a family of admissible controllers $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$, I define the *exact state family* of controller-map pairs $\{(\mathbf{k}_T \circ \Phi, \mathbf{F}_T^{e,\mathbf{x}}) \mid T \in I\}$ and *exact normal family* of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\xi}) \mid T \in I\}$.

For all $T \in I$ such that \mathbf{k}_T is T -admissible, the recursion:

$$\xi_{k+1} = \mathbf{F}_T^{e,\xi}(\xi_k, \mathbf{k}_T(\xi_k)) \in \Phi(\mathcal{X}), \quad (6.29)$$

is well-defined for all $\xi_0 \in \Phi(\mathcal{X})$ and $k \in \mathbb{Z}_{\geq 0}$. As before, if $\mathbf{x}_e \in \mathcal{X}$ is an unforced or forced equilibrium point of the open-loop system (2.1) for an equilibrium input $\mathbf{u}_e \in \mathbb{R}^m$ such that $(\mathbf{x}_e, \mathbf{u}_e) \in \mathcal{Z}$, it is said to be an unforced or forced equilibrium point of the exact state family of maps $\{\mathbf{F}_T^{e,\mathbf{x}} \mid T \in I\}$. The relationship in (2.25) implies that:

$$\mathbf{f}_\xi(\mathbf{0}_n) + \mathbf{g}_\xi(\mathbf{0}_n)\mathbf{u}_e = \mathbf{0}_n, \quad (6.30)$$

such that \mathbf{x}_e is an unforced or forced equilibrium point of the exact state family of maps $\{\mathbf{F}_T^{e,\mathbf{x}} \mid T \in I\}$ if and only if $\mathbf{0}_n$ is an unforced or forced equilibrium

point of the exact normal family of maps $\{\mathbf{F}_T^{e,\xi} \mid T \in I\}$, implying that:

$$\mathbf{F}_T^{e,\xi}(\mathbf{0}_n, \mathbf{u}_e) = \mathbf{0}_n, \quad (6.31)$$

for all $T \in I$. Moreover, given a family of admissible controllers $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$, (2.25) also implies that the point \mathbf{x}_e is an equilibrium point of the exact state family of controller-map pairs $\{(\mathbf{k}_T \circ \Phi, \mathbf{F}_T^{e,\mathbf{x}}) \mid T \in I\}$, such that $\mathbf{k}_T(\Phi(\mathbf{x}_e)) = \mathbf{u}_e$ for all $T \in I$, if and only if the point $\mathbf{0}_n$ is an equilibrium point of the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\xi}) \mid T \in I\}$ as $\mathbf{k}_T(\mathbf{0}_n) = \mathbf{u}_e$ for all $T \in I$. This is an important point as it suggests that relating stability results for the exact normal family of controller-map pairs to the exact state family of controller-map pairs will be well-posed.

Approximate Discrete-Time Models

In practice, closed-form expressions for the exact maps are rarely obtainable, suggesting the use of approximations in the control synthesis process. While there are many approaches to approximating this map, I will use the following approximation of the exact normal map:

Definition 46 (*Euler Approximation Family*). For every sample period $T \in I$, define the map $\mathbf{F}_T^{a,\xi} : \mathcal{Z}_\xi \rightarrow \mathbb{R}^n$ as:

$$\mathbf{F}_T^{a,\xi}(\xi_0, \mathbf{u}_0) = \xi_0 + T(\mathbf{f}_\xi(\xi_0) + \mathbf{g}_\xi(\xi_0)\mathbf{u}_0), \quad (6.32)$$

for all $(\xi_0, \mathbf{u}_0) \in \mathcal{Z}_\xi$, yielding the *Euler approximation family* of maps $\{\mathbf{F}_T^{a,\xi} \mid T \in I\}$. For a family of admissible controllers $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$, I define the corresponding *Euler approximation family* of controller-map pairs is $\{(\mathbf{k}_T, \mathbf{F}_T^{a,\xi}) \mid T \in I\}$.

The motivation behind this particular approximation will be preserving the strict feedback nature of the normal form [84], which I will utilize later when considering control synthesis. Observe that by (6.30), $\mathbf{0}_n$ is an unforced or forced equilibrium point of the exact normal family of maps $\{\mathbf{F}_T^{e,\xi} \mid T \in I\}$ if and only if:

$$\mathbf{F}_T^{a,\xi}(\mathbf{0}_n, \mathbf{u}_e) = \mathbf{0}_n. \quad (6.33)$$

Also note that if a family of admissible controllers $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$ satisfies $\mathbf{k}_T(\mathbf{0}_n) = \mathbf{u}_e$ for all $T \in (0, T^*)$ for some $T^* \in I$, then:

$$\mathbf{F}_T^{a,\xi}(\mathbf{0}_n, \mathbf{k}_T(\mathbf{0}_n)) = \mathbf{0}_n. \quad (6.34)$$

This is an important point as it suggests that designing for stability with the Euler approximation family of controller-map pairs in an effort to stabilize the exact normal family of controller-map pairs will be well-posed.

For $T \in I$, we can also define $\mathbf{F}_T^{a,\eta} : \mathcal{Z}_\xi \rightarrow \mathbb{R}^\gamma$ and $\mathbf{F}_T^{a,z} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^{n-\gamma}$ such that for all $(\boldsymbol{\xi}, \mathbf{u}) = ((\boldsymbol{\eta}, \mathbf{z}), \mathbf{u}) \in \mathcal{Z}_\xi$:

$$\mathbf{F}_T^{a,\boldsymbol{\xi}}(\boldsymbol{\xi}, \mathbf{u}) = \begin{bmatrix} \mathbf{F}_T^{a,\eta}(\boldsymbol{\xi}, \mathbf{u}) \\ \mathbf{F}_T^{a,z}(\boldsymbol{\xi}) \end{bmatrix} = \begin{bmatrix} \boldsymbol{\eta} + T(\mathbf{f}_\eta(\boldsymbol{\xi}) + \mathbf{g}_\eta(\boldsymbol{\xi})\mathbf{u}) \\ \mathbf{z} + T\boldsymbol{\omega}(\boldsymbol{\xi}) \end{bmatrix}. \quad (6.35)$$

First, observe that using $\mathbf{k}_{\text{fb1}}^\xi$ yields:

$$\mathbf{F}_T^{a,\eta}(\boldsymbol{\xi}, \mathbf{k}_{\text{fb1}}^\xi(\boldsymbol{\xi})) = (\mathbf{I}_\gamma + T\mathbf{A})\boldsymbol{\eta} + T\mathbf{B}\mathbf{k}_{\text{aux}}(\boldsymbol{\xi}), \quad (6.36)$$

such that the Euler approximation of the output dynamics appear as a linear discrete-time system. Moreover, it is straightforward to show that (\mathbf{A}, \mathbf{B}) being a controllable pair implies that $(\mathbf{I}_\gamma + T\mathbf{A}, T\mathbf{B})$ is a discrete-time controllable pair for all $T \in I$, such that the Euler approximation of the output dynamics can be rendered a stable discrete-time system through an appropriate choice of \mathbf{k}_{aux} . I note that if the approximation was taken using the state coordinates instead of the normal coordinates, this structure would not necessarily appear. This is an instance of sampling and feedback linearization failing to commute, but it also suggests that if the coordinate transformation associated with feedback linearization is performed before sampling, it is possible to retain the desirable control properties of feedback linearizability. Second, observe that there may be a $T \in I$ such that the controller \mathbf{k}_T is T -admissible but the recursion $\boldsymbol{\xi}_{k+1} = \mathbf{F}_T^{a,\boldsymbol{\xi}}(\boldsymbol{\xi}_k, \mathbf{k}_T(\boldsymbol{\xi}_k))$ is not well-defined for all $\boldsymbol{\xi}_0 \in \Phi(\mathcal{X})$ and $k \in \mathbb{Z}_{\geq 0}$. This is due to this map enabling $\boldsymbol{\xi}_k \notin \Phi(\mathcal{X})$ for some $k > 0$. While my results do not need this recursion to be well-defined, this can be achieved by extending the domains of \mathbf{f}_ξ , \mathbf{g}_ξ , and \mathbf{k}_T to \mathbb{R}^n .

The following definition characterizes how accurately an approximate map captures the exact map:

Definition 47 (*One-Step Consistency*). A family of controller-map pairs, $\{(\mathbf{k}_T, \mathbf{F}_T) : T \in I\}$, is *one-step consistent* with the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\boldsymbol{\xi}}) \mid T \in I\}$ if, for each compact set $K \subseteq \Phi(\mathcal{X})$, there exist a function $\rho \in \mathcal{K}_\infty$ and $T^* \in I$ such that for all $\boldsymbol{\xi} \in K$ and $T \in (0, T^*)$, we have:

$$\|\mathbf{F}_T^{e,\boldsymbol{\xi}}(\boldsymbol{\xi}, \mathbf{k}_T(\boldsymbol{\xi})) - \mathbf{F}_T(\boldsymbol{\xi}, \mathbf{k}_T(\boldsymbol{\xi}))\| \leq T\rho(T). \quad (6.37)$$

The next lemma establishes that the Euler approximation family of controller-map pairs is one-step consistent with the exact normal family of controller-map pairs. I note that this result is based on [81, Lemma 1], which establishes the one-step consistency of the Euler approximation family of controller-map pairs with an exact family of controller-map pairs. The proof of my lemma will only provide conditions for which I show the sufficient conditions of [81, Lemma 1] are met. The conclusion of one-step consistency in my lemma can then be established by directly following the proof of [81, Lemma 1]. In this sense, this lemma should be understood as at most a minor contribution of my work, but important for understanding my later results. I also note that in Section 6.4 I will provide a complete proof of one-step consistency for a larger class of approximate discrete-time maps that includes the Euler approximation family which will follow the steps I have omitted here.

Lemma 6. *Suppose \mathbf{f}_ξ and \mathbf{g}_ξ are locally Lipschitz continuous on $\Phi(\mathcal{X})$. Consider a family of admissible controllers $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$ and suppose that for any compact set $K \subset \Phi(\mathcal{X})$ there exist $T^* \in I$ and a bound $M \in \mathbb{R}_{>0}$ such that for every sample time $T \in (0, T^*)$, the controller \mathbf{k}_T is bounded by M on K . Then the Euler approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a,\xi}) \mid T \in I\}$ is one-step consistent with the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\xi}) \mid T \in I\}$.*

I note that since \mathbf{f} , \mathbf{g} , Φ , and $\frac{\partial \Phi}{\partial \mathbf{x}}$ are assumed to be locally Lipschitz continuous on at least D and $\mathcal{X} \subseteq \mathcal{D} \subseteq D$, the first condition of Lemma 6 is met.

Proof. Consider a compact set $K \subset \Phi(\mathcal{X})$ and corresponding $T^* \in I$ and $M \in \mathbb{R}_{>0}$, and fix a sample period $T \in (0, T^*)$. By assumption, \mathbf{k}_T is bounded on K , and since \mathbf{f}_ξ and \mathbf{g}_ξ are continuous, \mathbf{f}_ξ and \mathbf{g}_ξ are also bounded on K , implying there exists a bound $M' \in \mathbb{R}_{>0}$ with:

$$\|\mathbf{f}_\xi(\xi_2) + \mathbf{g}_\xi(\xi_2)\mathbf{k}_T(\xi_1)\| \leq M', \quad (6.38)$$

for all normal states $\xi_1, \xi_2 \in K$. As \mathbf{f}_ξ and \mathbf{g}_ξ are locally Lipschitz continuous on the compact set K , it follows that \mathbf{f}_ξ and \mathbf{g}_ξ are Lipschitz continuous on

K . Therefore:

$$\begin{aligned} & \| \mathbf{f}_\xi(\xi_2) + \mathbf{g}_\xi(\xi_2)\mathbf{k}_T(\xi_1) - (\mathbf{f}_\xi(\xi_1) + \mathbf{g}_\xi(\xi_1)\mathbf{k}_T(\xi_1)) \| \\ & \leq \| \mathbf{f}_\xi(\xi_2) - \mathbf{f}_\xi(\xi_1) \| + \| \mathbf{g}_\xi(\xi_2) - \mathbf{g}_\xi(\xi_1) \| \| \mathbf{k}_T(\xi_1) \|, \end{aligned} \quad (6.39)$$

$$\leq (L_{\mathbf{f}_\xi} + L_{\mathbf{g}_\xi}M) \| \xi_2 - \xi_1 \|, \quad (6.40)$$

$$\triangleq \rho(\| \xi_2 - \xi_1 \|), \quad (6.41)$$

for all states $\xi_1, \xi_2 \in K$, where $L_{\mathbf{f}_\xi}, L_{\mathbf{g}_\xi} \in \mathbb{R}_{>0}$ are Lipschitz constants for \mathbf{f}_ξ and \mathbf{g}_ξ , respectively, on K , and $\rho \in \mathcal{K}_\infty$ satisfies $\rho(r) = (L_{\mathbf{f}_\xi} + L_{\mathbf{g}_\xi}M)r$ for all $r \in \mathbb{R}_{\geq 0}$. The remainder of the proof directly follows the proof of Lemma 1 in [81] by substituting (using the notation of [81]) $\mathcal{X} = N(K, \epsilon) \subset \Phi(\mathcal{X})$, with proper containment implied for some $\epsilon \in \mathbb{R}_{>0}$ as $\Phi(\mathcal{X})$ is open, and substituting $T_1^* = \min\{T^*, \epsilon/M'\}$. \square

Practical Stability

Before defining the notion of practical stability which will be the objective of stabilizing controller design for sampled-data systems, I recall the following definition [100]:

Definition 48 (*Class- \mathcal{KL} Function*). Let $a \in \mathbb{R}_{>0}$. A continuous function $\beta : [0, a) \times \mathbb{R}_{\geq 0}$ is said to be a *class- \mathcal{KL} function* ($\beta \in \mathcal{KL}$) if for all $s \in \mathbb{R}_{\geq 0}$ the function $\beta(\cdot, s) : [0, a) \rightarrow \mathbb{R}_{\geq 0}$ is a class- \mathcal{K} function, and for all $r \in [0, a)$ and $s_1, s_2 \in \mathbb{R}_{\geq 0}$ satisfying $s_1 < s_2$, the function $\beta(r, \cdot) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ satisfies $\beta(r, s_2) < \beta(r, s_1)$ and $\lim_{s \rightarrow \infty} \beta(r, s) = 0$. A class- \mathcal{KL} function is said to be a *class- \mathcal{KL}_∞ function* ($\beta \in \mathcal{KL}_\infty$) if $a = \infty$.

This allows the following definition of practical stability, defined for both the exact state discrete map and the exact normal discrete map:

Definition 49 (*Practical Stability*). Let ζ_e be an equilibrium point of a family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$, let $\beta \in \mathcal{KL}_\infty$, and let $N \subseteq \mathbb{R}^n$ be an open set with $\zeta_e \in N$. The family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ is *(β, N)-practically stable* with respect to ζ_e if for each $R \in \mathbb{R}_{>0}$, there exists a $T^* \in I$ such that for each sample period $T \in (0, T^*)$, initial state $\zeta_0 \in N$, and number of steps $k \in \mathbb{Z}_{\geq 0}$, the recursion $\zeta_{k+1} = \mathbf{F}_T(\zeta_k, \mathbf{k}_T(\zeta_k))$ is well-defined and:

$$\| \zeta_k - \zeta_e \| \leq \beta(\| \zeta_0 - \zeta_e \|, kT) + R. \quad (6.42)$$

This definition requires that the state of the system converge to a neighborhood of the equilibrium point ζ_e (standing in for both \mathbf{x}_e and $\mathbf{0}_n$) that can be made arbitrarily small (smaller values of R) by requiring potentially higher sample rates (a smaller value of T^*). The following lemma relates the practical stability of the exact normal family and the exact state family of controller-map pairs. Importantly, it justifies considering the sampled normal form dynamics, which can be feedback linearized, rather than the sampled state dynamics that may not be feedback linearizable. I note that this lemma is a contribution of my work.

Lemma 7. *Suppose that for any compact sets $K, K' \subset \mathbb{R}^n$, Φ and Φ^{-1} are Lipschitz continuous on $K' \cap \mathcal{X}$ and $K \cap \Phi(\mathcal{X})$, respectively. If the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\xi}) \mid T \in I\}$ is (β, N) -practically stable with respect to $\mathbf{0}_n$ for an open set $N \subseteq \Phi(\mathcal{X})$ with $\mathbf{0}_n \in N$, then there exist $\beta' \in \mathcal{KL}_\infty$ and a bounded open set $N' \subseteq \mathcal{X}$ with $\mathbf{x}_e \in N'$ such that the exact state family of controller-map pairs $\{(\mathbf{k}_T \circ \Phi, \mathbf{F}_T^{e,\mathbf{x}}) \mid T \in I\}$ is (β', N') -practically stable with respect to \mathbf{x}_e .*

Proof. Let $N' \subseteq \Phi^{-1}(N)$ be a bounded open set satisfying $\text{cl}(N') \subseteq \mathcal{X}$ and $\mathbf{x}_e \in N'$, noting that such an N' exists as \mathcal{X} is open and $\mathbf{x}_e \in \Phi^{-1}(N) \subseteq \mathcal{X}$. As $\text{cl}(N')$ is compact and Φ is a homeomorphism (because it is a diffeomorphism) between \mathcal{X} and $\Phi(\mathcal{X})$, we have $\Phi(\text{cl}(N')) \subseteq \Phi(\mathcal{X})$ is compact and satisfies $\Phi(\text{cl}(N')) = \text{cl}(\Phi(N'))$. Fix $\bar{R} \in \mathbb{R}_{>0}$ and define the corresponding radii $r, r' \in \mathbb{R}_{>0}$ as:

$$r = \max_{\xi_0 \in \text{cl}(\Phi(N'))} \beta(\|\xi_0\|, 0) + \bar{R}, \quad (6.43)$$

$$r' = \max_{\mathbf{x}_0 \in \text{cl}(N')} \|\mathbf{x}_0\|, \quad (6.44)$$

and let $B_r, B_{r'} \subset \mathbb{R}^n$ be closed norm-balls centered at of radius r centered at $\mathbf{0}_n$ and r' centered at \mathbf{x}_e , respectively. By assumption, Φ^{-1} and Φ are Lipschitz continuous on $\Phi(\mathcal{X}) \cap B_r$ and $\mathcal{X} \cap B_{r'}$ with Lipschitz constants $L_{\Phi^{-1}}, L_\Phi \in \mathbb{R}_{>0}$, respectively. Given an arbitrary $R' \in \mathbb{R}_{>0}$, pick $R < \min\{\bar{R}, R'/L_{\Phi^{-1}}\}$. Let $\mathbf{x}_0 \in N'$ and let $\xi_0 = \Phi(\mathbf{x}_0) \in \Phi(N') \subseteq N$. By the (β, N) -practical stability of the exact normal family of controller-map pairs, corresponding to R , there exists a $T^* \in I$ such that for any sample period $T \in (0, T^*)$, the recursion $\xi_{k+1} = \mathbf{F}_T^{e,\xi}(\xi_k, \mathbf{k}_T(\xi_k)) \in \Phi(\mathcal{X})$ satisfies $\|\xi_k\| \leq \beta(\|\xi_0\|, kT) + R$, and implies $\xi_k \in \Phi(\mathcal{X}) \cap B_r$ for all $k \in \mathbb{Z}_{\geq 0}$. Letting $\mathbf{x}_{k+1} = \mathbf{F}_T^{e,\mathbf{x}}(\mathbf{x}_k, \mathbf{k}_T(\Phi(\mathbf{x}_k)))$ for all

$k \in \mathbb{Z}_{\geq 0}$, note that $\mathbf{x}_k = \Phi^{-1}(\boldsymbol{\xi}_k)$. It follows that for all $k \in \mathbb{Z}_{\geq 0}$:

$$\|\mathbf{x}_k - \mathbf{x}_e\| = \|\Phi^{-1}(\boldsymbol{\xi}_k) - \Phi^{-1}(\mathbf{0}_n)\|, \quad (6.45)$$

$$\leq L_{\Phi^{-1}} \|\boldsymbol{\xi}_k - \mathbf{0}_n\|, \quad (6.46)$$

$$\leq L_{\Phi^{-1}} (\beta(\|\boldsymbol{\xi}_0\|, kT) + R), \quad (6.47)$$

$$= L_{\Phi^{-1}} (\beta(\|\Phi(\mathbf{x}_0) - \Phi(\mathbf{x}_e)\|, kT) + R), \quad (6.48)$$

$$\leq L_{\Phi^{-1}} (\beta(L_{\Phi} \|\mathbf{x}_0 - \mathbf{x}_e\|, kT) + R), \quad (6.49)$$

$$< \beta'(\|\mathbf{x}_0 - \mathbf{x}_e\|, kT) + R', \quad (6.50)$$

where $\beta'(r, s) = L_{\Phi^{-1}} \beta(L_{\Phi} r, s)$ for all $r, s \in \mathbb{R}_{\geq 0}$ is a class- \mathcal{KL}_{∞} function. Note that the first inequality results from the fact we have taken $R < \bar{R}$, such that $\boldsymbol{\xi}_k \in \Phi(\mathcal{X}) \cap B_r$ for all $k \in \mathbb{Z}_{\geq 0}$, and thus the Lipschitz bound can be used, while the second inequality follows from the way r' is defined and the fact $\mathbf{x}_0 \in N'$. Therefore, the exact state family of controller-map pairs is (β', N') -practically stable. \square

Certifying practical stability will come through the following class of Lyapunov functions defined in [81]:

Definition 50 (*Asymptotic Stability by Equi-Lipschitz Lyapunov Functions*).

A family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ is *asymptotically stable by equi-Lipschitz Lyapunov functions* if for some open set $N \subseteq \Phi(\mathcal{X})$ with $\mathbf{0}_n \in N$ and any compact set $K \subset N$, there exist $T^* \in I$, comparison functions $\alpha_1, \alpha_2 \in \mathcal{K}_{\infty}$ and $\alpha_3 \in \mathcal{K}$, a family $\{V_T : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0} \mid T \in (0, T^*)\}$, and a Lipschitz constant $M \in \mathbb{R}_{>0}$ such that:

$$\alpha_1(\|\boldsymbol{\xi}_1\|) \leq V_T(\boldsymbol{\xi}_1) \leq \alpha_2(\|\boldsymbol{\xi}_1\|), \quad (6.51)$$

$$V_T(\mathbf{F}_T(\boldsymbol{\xi}_2, \mathbf{k}_T(\boldsymbol{\xi}_2))) - V_T(\boldsymbol{\xi}_2) \leq -T\alpha_3(\|\boldsymbol{\xi}_2\|), \quad (6.52)$$

$$|V_T(\boldsymbol{\xi}_3) - V_T(\boldsymbol{\xi}_4)| \leq M\|\boldsymbol{\xi}_3 - \boldsymbol{\xi}_4\|, \quad (6.53)$$

for all $\boldsymbol{\xi}_1 \in \mathbb{R}^n$, normal states $\boldsymbol{\xi}_2 \in N$ and $\boldsymbol{\xi}_3, \boldsymbol{\xi}_4 \in K$, and sample times $T \in (0, T^*)$.

The first condition is the standard positive-definite condition seen with Lyapunov functions. The second condition is a discrete-time decrement condition on the value of the Lyapunov function that scales with the sample-rate T . Observe that (6.51) and (6.52) require that:

$$\mathbf{F}_T(\mathbf{0}_n, \mathbf{k}_T(\mathbf{0}_n)) = \mathbf{0}_n. \quad (6.54)$$

The last condition requires that the function V_T be Lipschitz continuous on any compact set $K \subseteq N$. Notably, the functions $\alpha_1, \alpha_2, \alpha_3$ and the constant M must hold *uniformly* for all $T \in (0, T^*)$, leading to the name *equi-Lipschitz*.

These functions link one-step consistency to practical stability in the next lemma. I note that this lemma is a modification of [81, Theorem 2], which is the main result of that work and provides a global practical stability result. My modification focuses on making the technically necessary adjustments that allow for local stability results, as I will be later studying the local exponential stability of zero-dynamics. A full proof of this lemma would require significant repetition of the steps in [81], and thus my “proof” should be seen as a list of instructions that can be followed to modify the proof steps of [81, Theorem 2] rather than a full proof. Again, in this sense, this lemma should be seen as at most a minor contribution of my work, but it is important for justifying my later results.

Lemma 8. *If the corresponding Euler approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{\alpha, \xi}) \mid T \in I\}$ is asymptotically stable by equi-Lipschitz Lyapunov functions, then there exist $\beta \in \mathcal{KL}_\infty$ and a bounded open set $U \subseteq \Phi(\mathcal{X})$ with $\mathbf{0}_n \in U$ such that the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e, \xi) \mid T \in I\}$ is (β, N) -practically stable for any open set $N \subseteq U$ with $\mathbf{0}_n \in N$.*

The importance of this lemma is that it states that controllers designed using the Euler approximation family of maps can endow the corresponding exact normal family of controller-map pairs with practical stability guarantees, and using Lemma 7, endow the exact state family of controller-map pairs with practical stability guarantees.

Proof. Consider the open set $N \subseteq \Phi(\mathcal{X})$ and the functions $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ and $\alpha_3 \in \mathcal{K}$ as specified by Definition 50. Let $K \subset N$ be a compact set with $\mathbf{0}_n \in \text{int}(K)$. By one-step consistency and asymptotic stability by equi-Lipschitz Lyapunov functions, there exist a $\rho \in \mathcal{K}_\infty$, a Lipschitz constant $M \in \mathbb{R}_{\geq 0}$, and an $T_0^* \in I$ such that for all $T \in (0, T_0^*)$, (6.37), (6.51), (6.52), and (6.53) hold for all $\xi, \xi_1, \dots, \xi_4 \in K$. There exists a radius $R \in \mathbb{R}_{> 0}$ such that the closed norm-ball of radius R centered at $\mathbf{0}_n$ is contained in K . I modify the claim of [81, Equation 37] for the local setting in this work as follows:

Claim 1. For any $d, D \in \mathbb{R}_{>0}$ with $D \leq \alpha_2^{-1}(\alpha_1(\frac{R}{2}))$ and $d \leq 2\alpha_2(R)$, there exists an $T^* \in (0, T_0^*)$ such that for every $\boldsymbol{\xi} \in \Phi(\mathcal{X})$ and $T \in (0, T^*)$, if $\|\boldsymbol{\xi}\| \leq D$ and $\max\{V_T(\mathbf{F}_T^{e,\boldsymbol{\xi}}(\boldsymbol{\xi}, \mathbf{k}_T(\boldsymbol{\xi}))), V_T(\boldsymbol{\xi})\} \geq d$, then:

$$V_T(\mathbf{F}_T^{e,\boldsymbol{\xi}}(\boldsymbol{\xi}, \mathbf{k}_T(\boldsymbol{\xi}))) - V_T(\boldsymbol{\xi}) \leq -\frac{T}{2}\alpha_3(\|\boldsymbol{\xi}\|). \quad (6.55)$$

In the language of [81], the restrictions on d and D imply $\Delta \leq R$, and the proof follows by replacing the sets \mathcal{X} and \mathcal{X}_1 with K , the constant M with the Lipschitz constant M given above, and the constants T_1^* and T_2^* with T_0^* . Letting $U \subset K$ be the open ball of radius $\alpha_2^{-1}(\alpha_1(\alpha_2^{-1}(\alpha_1(\frac{R}{2}))))$, the modified claim may be used to prove the existence of a $\beta \in \mathcal{KL}_\infty$ such that the exact normal family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{e,\boldsymbol{\xi}}) \mid T \in I\}$ is (β, N') -practically stable for any open set $N' \subseteq U$ containing the origin by following the proof of Theorem 2 in [81]. \square

Sampled-Data Stabilization

In now present my main result establishing feedback linearization as a method for practically stabilizing sampled-data nonlinear systems. The first result builds on [328] to make a claim on the stabilizability of the output dynamics:

Lemma 9. *Let the open-loop system (2.1) be feedback linearizable to an unforced or forced equilibrium point $\mathbf{x}_e \in \mathcal{X}$, consider $\mathbf{K} \in \mathbb{R}^{k \times \gamma}$ such that $\mathbf{A}_{\text{cl}} \triangleq \mathbf{A} - \mathbf{BK}$ is Hurwitz, and let $\mathbf{P}_\eta \in \mathbb{S}_{>0}^\gamma$ solve the CTLE:*

$$\mathbf{A}_{\text{cl}}^\top \mathbf{P}_\eta + \mathbf{P}_\eta \mathbf{A}_{\text{cl}} = -\mathbf{Q}_\eta, \quad (6.56)$$

for some $\mathbf{Q} \in \mathbb{S}_{>0}^\gamma$. Define the function $V_\eta : \mathbb{R}^\gamma \rightarrow \mathbb{R}_{\geq 0}$ as $V_\eta(\boldsymbol{\eta}) = \boldsymbol{\eta}^\top \mathbf{P}_\eta \boldsymbol{\eta}$ for all $\boldsymbol{\eta} \in \mathbb{R}^\gamma$. For any $\mathbf{Q}_\eta \in \mathbb{S}_{>0}^\gamma$ such that $\mathbf{Q}_\eta \prec \mathbf{Q}$ and $c \in (0, 1)$, there exists $T_\eta^* \in I$ such that for any $\boldsymbol{\eta}_0 \in \mathbb{R}^\gamma$, $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X}) \setminus \{\mathbf{0}_n\}$, and $T \in (0, T_\eta^*)$, there exists an input $\mathbf{u} \in \mathbb{R}^m$ such that:

$$\lambda_{\min}(\mathbf{P}_\eta) \|\boldsymbol{\eta}_0\|^2 \leq V_\eta(\boldsymbol{\eta}_0) \leq \lambda_{\max}(\mathbf{P}_\eta) \|\boldsymbol{\eta}_0\|^2, \quad (6.57)$$

$$V_\eta(\mathbf{F}_T^{a,\eta}(\boldsymbol{\xi}, \mathbf{u})) - V_\eta(\boldsymbol{\eta}) < -Tc\lambda_{\min}(\mathbf{Q}_\eta) \|\boldsymbol{\eta}\|^2. \quad (6.58)$$

This result states that if the continuous-time open-loop system (2.1) is feedback linearizable to an equilibrium point \mathbf{x}_e and we pick a stabilizing linear gain \mathbf{K} for the output dynamics, we can solve the CTLE for the output portion of the corresponding closed-loop normal-form dynamics to produce a type of

discrete-time CLF for the output portion of the Euler approximation family of maps. While I have yet to show that the equi-Lipschitz property is met (I will shortly), note that V_η satisfies the necessary uniformity properties in the sample period T with respect to α_1, α_2 and α_3 required by equi-Lipschitz Lyapunov functions.

Proof. The bounds in (6.57) follow from the definition of V_η . Define the auxiliary controller $\mathbf{k}_{\text{aux}}((\boldsymbol{\eta}, \mathbf{z})) = -\mathbf{K}\boldsymbol{\eta}$ for all $(\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X})$. For the controller $\mathbf{k}_{\text{fbl}}^\xi$ defined in (6.25), we have:

$$V_\eta(\mathbf{F}_T^{a,\eta}(\boldsymbol{\xi}, \mathbf{k}_{\text{fbl}}^\xi(\boldsymbol{\xi}))) - V_\eta(\boldsymbol{\eta}) = V_\eta((\mathbf{I}_\gamma + T\mathbf{A}_{\text{cl}})\boldsymbol{\eta}) - V_\eta(\boldsymbol{\eta}), \quad (6.59)$$

$$= T\boldsymbol{\eta}^\top (\mathbf{A}_{\text{cl}}^\top \mathbf{P}_\eta + \mathbf{P}_\eta \mathbf{A}_{\text{cl}} + T\mathbf{A}_{\text{cl}}^\top \mathbf{P}_\eta \mathbf{A}_{\text{cl}}) \boldsymbol{\eta} \quad (6.60)$$

$$< -T(\lambda_{\min}(\mathbf{Q}_\eta) - T\lambda_{\max}(\mathbf{A}_{\text{cl}}^\top \mathbf{P}_\eta \mathbf{A}_{\text{cl}})) \|\boldsymbol{\eta}\|^2, \quad (6.61)$$

for all $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X}) \setminus \{\mathbf{0}_n\}$ and $T \in I$. Picking $T_\eta^* \in I$ with:

$$T_\eta^* \leq (1 - c)\lambda_{\min}(\mathbf{Q}_\eta)/\lambda_{\max}(\mathbf{A}_{\text{cl}}^\top \mathbf{P}_\eta \mathbf{A}_{\text{cl}}), \quad (6.62)$$

implies that for all $T \in (0, T_\eta^*]$ and $\boldsymbol{\xi} \in \Phi(\mathcal{X}) \setminus \{\mathbf{0}_n\}$, the input $\mathbf{k}_{\text{fbl}}^\xi(\boldsymbol{\xi})$ satisfies (6.58). \square

I note that the preceding result did not include $\mathbf{0}_n$, as the strict inequality can not be satisfied there. Observe that using the controller $\mathbf{k}_{\text{fbl}}^\xi$ with \mathbf{k}_{aux} as defined in the proof, (2.33) implies that $\mathbf{F}_T^{a,\eta}(\mathbf{0}_n, \mathbf{k}_{\text{fbl}}^\xi(\mathbf{0}_n)) = \mathbf{0}_\gamma$. Furthermore, if \mathbf{x}_e is an unforced or forced equilibrium point of the open-loop system (2.1), the assumption that the dynamics of \mathbf{z} do not depend on the input implies that $\boldsymbol{\omega}(\mathbf{0}_n) = \mathbf{0}_{n-\gamma}$. Thus, with this controller, we have that $\mathbf{F}_T^{a,\xi}(\mathbf{0}_n, \mathbf{k}_{\text{fbl}}^\xi(\mathbf{0}_n)) = \mathbf{0}_n$.

The function V_η is said to be a discrete-time CLF for the output portion of the Euler approximation family of maps for $T \in (0, T_\eta^*]$. For each $T \in (0, T_\eta^*)$, define $\mathcal{U}_T : \Phi(\mathcal{X}) \rightarrow \mathcal{P}(\mathbb{R}^m)$ as:

$$\mathcal{U}_T(\boldsymbol{\xi}) = \left\{ \mathbf{u} \in \mathbb{R}^m \left| \begin{array}{l} (\boldsymbol{\xi}, \mathbf{u}) \in \mathcal{Z}_\xi \text{ and} \\ V_\eta(\mathbf{F}_T^{a,\eta}(\boldsymbol{\xi}, \mathbf{u})) - V_\eta(\boldsymbol{\eta}) \leq -Tc\lambda_{\min}(\mathbf{Q}_\eta)\|\boldsymbol{\eta}\|^2 \end{array} \right. \right\}, \quad (6.63)$$

for all $\boldsymbol{\xi} \in \Phi(\mathcal{X})$, noting that it if \mathbf{k}_{aux} as in the proof of Lemma 9 is T -admissible for all $T \in (0, T_\eta^*]$, then this set is non-empty for all $T \in (0, T_\eta^*]$

and $\boldsymbol{\xi} \in \Phi(\mathcal{X})$, as $\mathbf{k}_{\text{fb}}^{\boldsymbol{\xi}}(\boldsymbol{\xi}) \in \mathcal{U}_T(\boldsymbol{\xi})$. The next result connects these functions to continuous-time stability of the zero-dynamics, implying the conditions of Lemma 8 are met for controllers other than feedback linearizing controllers:

Theorem 41. *Let $V_{\boldsymbol{\eta}}$ and $T_{\boldsymbol{\eta}}^*$ be defined as in Lemma 9, and the zero-dynamics system governed by the differential equation:*

$$\dot{\mathbf{z}} = \boldsymbol{\omega}((\mathbf{0}_{\gamma}, \mathbf{z})), \quad (6.64)$$

for zero-coordinate signal \mathbf{z} is locally exponentially stable with respect to $\mathbf{0}_{n-\gamma}$. Let $\{\mathbf{k}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m \mid T \in I\}$ be a family of admissible controllers satisfying $\mathbf{k}_T(\boldsymbol{\xi}) \in \mathcal{U}_T(\boldsymbol{\xi})$ for all $T \in (0, T_{\boldsymbol{\eta}}^*]$ and $\boldsymbol{\xi} \in \Phi(\mathcal{X})$. Then the Euler approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a, \boldsymbol{\xi}}) \mid T \in I\}$ is asymptotically stable by equi-Lipschitz Lyapunov functions.

Proof. The local exponential stability of (6.64) implies that for any $\mathbf{Q}_{\mathbf{z}} \in \mathbb{S}_{>0}^n$ and $d \in (0, 1)$, there exist an open neighborhood of the origin $N \subseteq \mathbb{R}^{n-\gamma}$, a $T_{\mathbf{z}}^* \in I$, a $\mathbf{P}_{\mathbf{z}} \in \mathbb{S}_{>0}^n$, and a quadratic Lyapunov function $V_{\mathbf{z}} : \mathbb{R}^{n-\gamma} \rightarrow \mathbb{R}_{\geq 0}$ defined as $V_{\mathbf{z}}(\mathbf{z}) = \mathbf{z}^{\top} \mathbf{P}_{\mathbf{z}} \mathbf{z}$ for all $\mathbf{z} \in \mathbb{R}^{n-\gamma}$ and satisfying:

$$\lambda_{\min}(\mathbf{P}_{\mathbf{z}}) \|\mathbf{z}_0\|^2 \leq V_{\mathbf{z}}(\mathbf{z}_0) \leq \lambda_{\max}(\mathbf{P}_{\mathbf{z}}) \|\mathbf{z}_0\|^2, \quad (6.65)$$

$$V_{\mathbf{z}}(\mathbf{F}_T^{a, \mathbf{z}}((\mathbf{0}_{\gamma}, \mathbf{z}))) - V_{\mathbf{z}}(\mathbf{z}) \leq -Td\lambda_{\min}(\mathbf{Q}_{\mathbf{z}}) \|\mathbf{z}\|^2, \quad (6.66)$$

for all $\mathbf{z}_0 \in \mathbb{R}^{n-\gamma}$, $\mathbf{z} \in N$, and $T \in (0, T_{\mathbf{z}}^*)$. Construction of $V_{\mathbf{z}}$, N , and $T_{\mathbf{z}}^*$ first follows the proof steps of [100, Theorem 4.7] with the linearization of $\boldsymbol{\omega}(\mathbf{0}_{\gamma}, \cdot)$ at $\mathbf{0}_{n-\gamma}$ to produce a suitable neighborhood N and then follows the proof steps of Lemma 9 to produce $T_{\mathbf{z}}^*$. Let $\sigma \in \mathbb{R}_{>0}$ be a coefficient to be specified later. Let $V_{\boldsymbol{\eta}}$, c , and $T_{\boldsymbol{\eta}}^*$ be defined as in Lemma 4, and define the composite Lyapunov function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ as:

$$V(\boldsymbol{\xi}) = \sigma V_{\boldsymbol{\eta}}(\boldsymbol{\eta}) + V_{\mathbf{z}}(\mathbf{z}), \quad (6.67)$$

for all $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \mathbb{R}^n$. First, note that (6.57) and (6.65) imply:

$$\min \{ \sigma \lambda_{\min}(\mathbf{P}_{\boldsymbol{\eta}}), \lambda_{\min}(\mathbf{P}_{\mathbf{z}}) \} \|\boldsymbol{\xi}\|^2 \leq V(\boldsymbol{\xi}) \leq \underbrace{\max \{ \sigma \lambda_{\max}(\mathbf{P}_{\boldsymbol{\eta}}), \lambda_{\max}(\mathbf{P}_{\mathbf{z}}) \}}_{\triangleq \mu} \|\boldsymbol{\xi}\|^2, \quad (6.68)$$

for all $\boldsymbol{\xi} \in \mathbb{R}^n$, implying (6.51) is met. Second, note that:

$$\begin{aligned} \left\| \frac{\partial V}{\partial \boldsymbol{\xi}}(\boldsymbol{\xi}) \right\| &\leq 2(\sigma \lambda_{\max}(\mathbf{P}_{\boldsymbol{\eta}}) \|\boldsymbol{\eta}\| + \lambda_{\max}(\mathbf{P}_{\mathbf{z}}) \|\mathbf{z}\|), \\ &\leq 2(\mu \|\boldsymbol{\xi}\| + \mu \|\boldsymbol{\xi}\|) = 4\mu \|\boldsymbol{\xi}\|, \end{aligned} \quad (6.69)$$

for all $\xi = (\eta, \mathbf{z}) \in \Phi(\mathcal{X})$, implying that for any compact set $K \subset \Phi(\mathcal{X})$, (6.53) is met as we have:

$$|V(\xi_1) - V(\xi_2)| \leq 4\mu \left(\max_{\xi \in K} \|\xi\| \right) \|\xi_1 - \xi_2\|, \quad (6.70)$$

for all $\xi_1, \xi_2 \in K$. Third, define a bounded open set $N_\xi \subset \mathbb{R}^n$ with $\mathbf{0}_n \in N_\xi$ and closure $\text{cl}(N_\xi) \subset \Phi(\mathcal{X}) \cap (\mathbb{R}^\gamma \times N)$, let $L_\omega \in \mathbb{R}_{>0}$ be a Lipschitz constant of ω on N_ξ , and let $T_1^* = \min\{T_\eta^*, T_z^*\}$. Note that L_ω exists as ω is continuously differentiable on $\Phi(\mathcal{X}) \subseteq \Phi(D)$, and $\text{cl}(N_\xi)$ is a compact set contained in $\Phi(\mathcal{X})$. For all $\xi = (\eta, \mathbf{z}) \in N_\xi$ and $T \in (0, T_1^*)$, we have:

$$\begin{aligned} & V(\mathbf{F}_T^{a,\xi}(\xi, \mathbf{k}_T(\xi))) - V(\xi) \\ &= \sigma(V_\eta(\mathbf{F}_T^{a,\eta}(\xi, \mathbf{k}_T(\xi))) - V_\eta(\eta)) + V_z(\mathbf{F}_T^{a,z}(\xi)) - V_z(\mathbf{z}), \end{aligned} \quad (6.71)$$

$$\begin{aligned} & \leq -\sigma T c \lambda_{\min}(\mathbf{Q}_\eta) \|\eta\|^2 + V_z(\mathbf{F}_T^{a,z}(\mathbf{0}_\gamma, \mathbf{z})) - V_z(\mathbf{z}) \\ & \quad + V_z(\mathbf{F}_T^{a,z}((\eta, \mathbf{z}))) - V_z(\mathbf{F}_T^{a,z}(\mathbf{0}_\gamma, \mathbf{z})), \end{aligned} \quad (6.72)$$

$$\begin{aligned} & \leq -\sigma T c \lambda_{\min}(\mathbf{Q}_\eta) \|\eta\|^2 - T d \lambda_{\min}(\mathbf{Q}_z) \|\mathbf{z}\|^2 \\ & \quad + 2T \mathbf{z}^\top \mathbf{P}_z(\omega((\eta, \mathbf{z})) - \omega(\mathbf{0}_\gamma, \mathbf{z})) \\ & \quad + T^2(\omega((\eta, \mathbf{z}))^\top \mathbf{P}_z \omega((\eta, \mathbf{z})) - \omega(\mathbf{0}_\gamma, \mathbf{z})^\top \mathbf{P}_z \omega(\mathbf{0}_\gamma, \mathbf{z})), \end{aligned} \quad (6.73)$$

$$\begin{aligned} & \leq -\sigma T c \lambda_{\min}(\mathbf{Q}_\eta) \|\eta\|^2 - T d \lambda_{\min}(\mathbf{Q}_z) \|\mathbf{z}\|^2 \\ & \quad + 2T \lambda_{\max}(\mathbf{P}_z) L_\omega \|\eta\| \|\mathbf{z}\| + T^2 \lambda_{\max}(\mathbf{P}_z) L_\omega \|\xi\|^2, \end{aligned} \quad (6.74)$$

$$= -T \begin{bmatrix} \|\eta\| \\ \|\mathbf{z}\| \end{bmatrix}^\top \underbrace{\begin{bmatrix} \theta_\eta(\sigma, T) & -\theta_\times \\ -\theta_\times & \theta_z(T) \end{bmatrix}}_{\triangleq \Theta_\sigma(T)} \begin{bmatrix} \|\eta\| \\ \|\mathbf{z}\| \end{bmatrix}, \quad (6.75)$$

where $\theta_\eta(\sigma, T) = \sigma c \lambda_{\min}(\mathbf{Q}_\eta) - T \lambda_{\max}(\mathbf{P}_z) L_\omega$, $\theta_\times = \lambda_{\max}(\mathbf{P}_z) L_\omega$, and $\theta_z(T) = d \lambda_{\min}(\mathbf{Q}_z) - T \lambda_{\max}(\mathbf{P}_z) L_\omega$. Observe that in going from (6.73) to (6.74) I dropped the last term (as it is negative), and used the Lipschitz property of ω to produce the term $\|\xi\|^2$, which can be decoupled into the diagonal terms of the matrix $\Theta_\sigma(T)$. Pick $T_2^* \in (0, T_1^*)$ such that $T_2^* < d \lambda_{\min}(\mathbf{Q}_z) / \theta_\times$, implying $\theta_z(T) > 0$ for all $T \in (0, T_2^*)$, and note that $\theta_z(T_2) \geq \theta_z(T_1)$ for all $T_1, T_2 \in (0, T_2^*)$ such that $T_2 < T_1$. Choosing σ such that:

$$\sigma > (\theta_\times^2 / \theta_z(T_2^*) + T_2^* \theta_\times) / (c \lambda_{\min}(\mathbf{Q}_\eta)), \quad (6.76)$$

then implies that $\Theta_\sigma(T) \in \mathbb{S}_{>0}^n$ for all $T \in [0, T_2^*]$. The composition $\lambda_{\min} \circ \Theta_\sigma$ is continuous and $\mathbb{R}_{>0}$ -valued for all $T \in [0, T_2^*]$ as Θ_σ is an affine function,

and thus attains its minimum on the compact interval $[0, T_2^*]$, such that:

$$\alpha \triangleq \min_{T' \in [0, T_2^*]} \lambda_{\min}(\Theta_\sigma(T')) > 0. \quad (6.77)$$

Therefore:

$$V(\mathbf{F}_T^{a, \xi}(\xi, \mathbf{k}_T(\xi))) - V(\xi) \leq -T \lambda_{\min}(\Theta_\sigma(T)) \|\xi\|^2, \quad (6.78)$$

$$\leq -T \alpha \|\xi\|^2, \quad (6.79)$$

for all $\xi \in N_\xi$ and $T \in (0, T_2^*]$. \square

As the controller $\mathbf{k}_{\text{fbl}}^\xi$ with \mathbf{k}_{aux} defined in the proof of Lemma 9 does not explicitly depend on the sample-period T , if $(\xi, \mathbf{k}_{\text{fbl}}^\xi(\xi)) \in \mathcal{Z}_\xi$ for all $\xi \in \Phi(\mathcal{X})$ and $\mathbf{k}_{\text{fbl}}^\xi$ is T_0 -admissible for some $T_0 \in I$, then it is T -admissible for all $T \in (0, T_0)$ (as $(\xi, \mathbf{k}_{\text{fbl}}^\xi(\xi)) \in \mathcal{Z}_\xi$ for all $T \in (0, T_0)$). Consequently, this result implies that if such a T_0 exists, the exact normal family of controller-map pairs $\{(\mathbf{k}_{\text{fbl}}^\xi, \mathbf{F}_T^{e, \xi}) \mid T \in I\}$ is (β, N) -practically stable with respect to $\mathbf{0}_n$ for some $\beta \in \mathcal{KL}_\infty$ and open set $N \subseteq \Phi(\mathcal{X})$ with $\mathbf{0}_n \in N$. In this way, I have shown that a feedback linearizing controller for a continuous-time system will render the corresponding sampled-data system practically stable, without modification of the feedback linearizing controller.

Optimization-Based Control Synthesis

The existence of the feedback linearizing controller ensures the function V_η is also a local exponential CLF for the continuous-time output dynamics and equilibrium point $\mathbf{0}_\gamma$. For a sample period $T \in I$, continuous-time design yields the standard (CLF-QP) controller $\mathbf{k}_T^{\text{CLF}} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$ specified as:

$$\mathbf{k}_T^{\text{CLF}}(\xi) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u}\|^2 \quad (\text{CLF-QP})$$

$$\text{s.t. } L_{\mathbf{f}_\eta} V_\eta(\xi) + L_{\mathbf{g}_\eta} V_\eta(\xi) \mathbf{u} \leq -\lambda_{\min}(\mathbf{Q}_\eta) \|\eta\|^2,$$

for all $\xi = (\eta, \mathbf{z}) \in \Phi(\mathcal{X})$. As shown in Section 6.2, this controller often displays degradation in performance with sampling, motivating the specification of a sampled-data controller. For $T \in (0, T_\eta^*]$, using the Euler approximate model $\mathbf{F}_T^{a, \eta}$, consider a controller $\mathbf{k}_T^{\text{SD-CLF}} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$ specified by the following convex quadratically constrained quadratic program (QCQP):

$$\mathbf{k}_T^{\text{SD-CLF}}(\xi) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u}\|^2 \quad (\text{CLF-QCQP})$$

$$\text{s.t. } V_\eta(\mathbf{F}_T^{a, \eta}(\xi, \mathbf{u})) - V_\eta(\eta) \leq -T c \lambda_{\min}(\mathbf{Q}_\eta) \|\eta\|^2,$$

$$(6.80)$$

for all $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X})$. To see that this controller is a convex QCQP, we can expand the constraint to reveal the alternative form:

$$\begin{aligned} \mathbf{k}_T^{\text{SD-CLF}}(\boldsymbol{\xi}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \|\mathbf{u}\|^2 && \text{(CLF-QCQP)} \\ \text{s.t. } \mathbf{u}^\top \boldsymbol{\Lambda}_T(\boldsymbol{\xi}) \mathbf{u} + 2\boldsymbol{\lambda}_T(\boldsymbol{\xi})^\top \mathbf{u} + l_T(\boldsymbol{\xi}) &\leq 0, \end{aligned}$$

for all $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X})$, where $\boldsymbol{\Lambda}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{S}_{\geq 0}^m$, $\boldsymbol{\lambda}_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$, and $l_T : \Phi(\mathcal{X}) \rightarrow \mathbb{R}$ are defined with \mathbf{P}_η , \mathbf{Q}_η , and c from Lemma 9:

$$\boldsymbol{\Lambda}_T(\boldsymbol{\xi}) = T\mathbf{g}_\eta(\boldsymbol{\xi})^\top \mathbf{P}_\eta \mathbf{g}_\eta(\boldsymbol{\xi}), \quad (6.81)$$

$$\boldsymbol{\lambda}_T(\boldsymbol{\xi}) = \mathbf{g}_\eta(\boldsymbol{\xi})^\top \mathbf{P}_\eta (\boldsymbol{\eta} + T\mathbf{f}_\eta(\boldsymbol{\xi})), \quad (6.82)$$

$$l_T(\boldsymbol{\xi}) = \mathbf{f}_\eta(\boldsymbol{\xi})^\top \mathbf{P}_\eta (2\boldsymbol{\eta} + T\mathbf{f}_\eta(\boldsymbol{\xi})) + c\lambda_{\min}(\mathbf{Q}_\eta)\|\boldsymbol{\eta}\|^2, \quad (6.83)$$

for all $\boldsymbol{\xi} = (\boldsymbol{\eta}, \mathbf{z}) \in \Phi(\mathcal{X})$. For any $\boldsymbol{\xi} \in \Phi(\mathcal{X})$, the input $\mathbf{k}_{\text{fb1}}^\xi(\boldsymbol{\xi})$ is in the feasible set of the corresponding optimization problem, and as the feasible set is closed and the $\|\mathbf{k}_{\text{fb1}}^\xi(\boldsymbol{\xi})\|^2$ -sublevel set of the continuous objective function is compact, there exists a minimizer in this sublevel set. Since the objective function is strictly convex and the feasible set is convex, this minimizer is unique and can be found in polynomial time [35]. For each $T \in (T_\eta^*, T_{\max}]$, define $\mathbf{k}_T^{\text{SD-CLF}} : \Phi(\mathcal{X}) \rightarrow \mathbb{R}^m$ arbitrarily. If $\{\mathbf{k}_T^{\text{SD-CLF}} \mid T \in I\}$ is a family of admissible controllers, then the exact family $\{(\mathbf{k}_T^{\text{SD-CLF}}, \mathbf{F}_T^{e,\xi}) \mid T \in I\}$ is (β, N) -practically stable for some $\beta \in \mathcal{KL}_\infty$ and open set $N \subseteq \Phi(\mathcal{X})$ by Theorem 41. This follows as the feasibility of the feedback linearizing control input implies the family $\{(\mathbf{k}_T^{\text{SD-CLF}}, \mathbf{F}_T^{a,\xi}) \mid T \in I\}$ is asymptotically stable by the same equi-Lipschitz Lyapunov functions as the family $\{(\mathbf{k}_{\text{fb1}}^\xi, \mathbf{F}_T^{a,\xi}) \mid h \in I\}$.

Simulation Results

To illustrate the advantage of sampled-data design, I will first consider the following system with exponentially stable zero-dynamics:

$$\dot{\eta}_1 = \eta_2, \quad (6.84)$$

$$\dot{\eta}_2 = 10 \sin(\eta_1) + u, \quad (6.85)$$

$$\dot{z} = \eta_1^2 - z, \quad (6.86)$$

where (η_1, η_2) , z , and u denote the output, zero-coordinate, and control signal, respectively. For $\mathbf{K} = \begin{bmatrix} 1/2 & \sqrt{3}/2 \end{bmatrix}$, $\mathbf{Q}_\eta = \mathbf{I}_2$, $c = 0.5$, $T = 0.2$, and initial condition $(1, 0, 1)$, the (CLF-QP) fails to stabilize the system, while the (CLF-QCQP) stabilizes the system as seen in Figure 6.6.

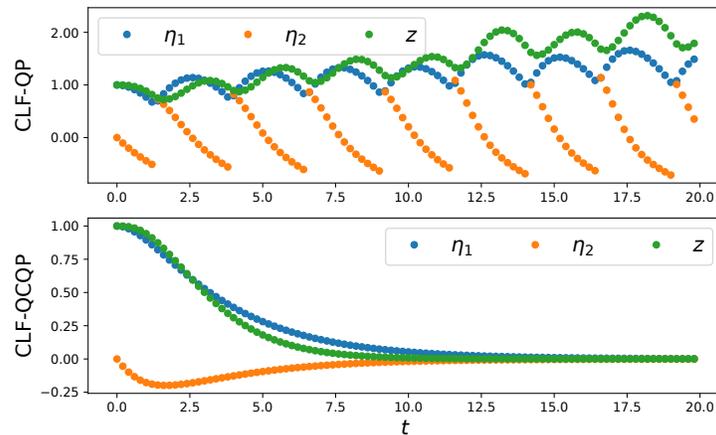


Figure 6.6. With inputs applied via a zero-order hold, the (CLF-QP) controller does not stabilize the system (6.84) (Top), while the (CLF-QCQP) does (Bottom). Simulation code listed at <https://bit.ly/CLF-QCQP>.

Next, I deploy the proposed (CLF-QCQP) controller on the inverted pendulum, double inverted pendulum, and planar quadrotor simulation environments considered in Section 6.2, with the results seen in Figures 6.7, 6.8, and 6.9, respectively. We see that the proposed controller stabilizes each system at each sample rate, including the exceptionally low sample rate of 10 [Hz]. Moreover, the input does not display high frequency oscillations for any system at any frequency. I do note the large input jump that occurs with the double inverted pendulum near the beginning of the simulation, which suggests the controller is quickly pushing the system into a region of the state space from which it continues to produce smooth inputs. These simulation results show that addressing sampling in the control design can lead to significantly better performance, and my results in this section highlight that there is a rigorous theoretical path towards doing so.

Conclusion

In this section I presented my work on stabilizing sampled-data control through feedback linearization and CLF-based controllers specified via convex optimization. I begin by reviewing the stabilizing sampled-data control framework with approximate discrete-time models in [81] and providing necessary adjustments to support local stability results and the coordinate transformations that accompany feedback linearization. I then demonstrate how a continuous-time

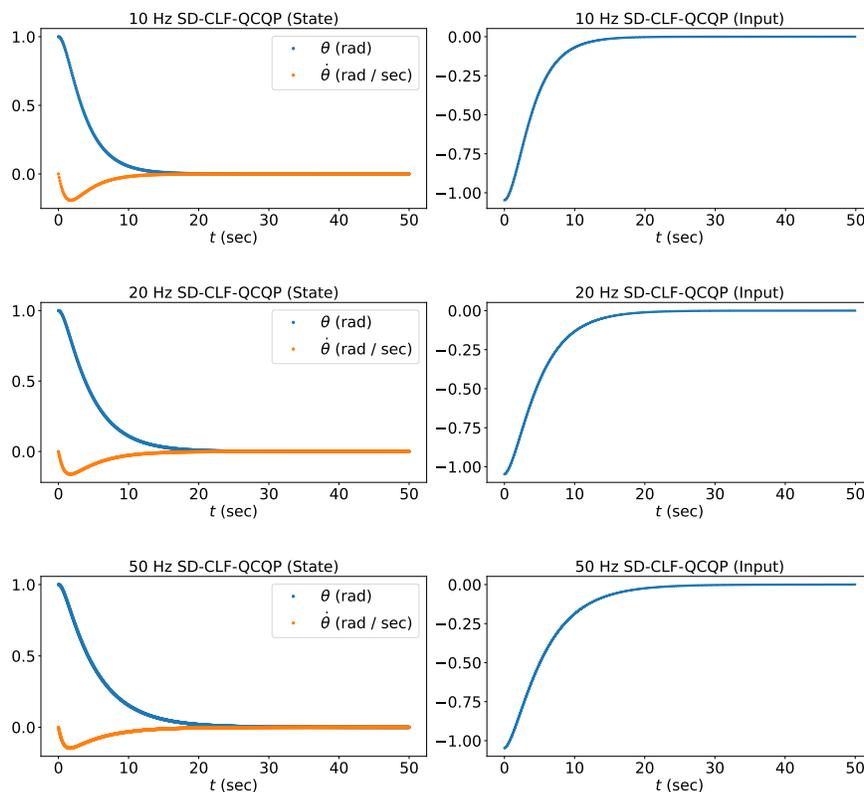


Figure 6.7. Simulation results for the inverted pendulum system using the (CLF-QCQP) controller.

system that is feedback linearizable with locally exponentially stable zero-dynamics can be rendered practically stable using the continuous-time feedback linearization controller without modification. I then use this result to develop CLF-based controllers specified via convex quadratically constrained quadratic programs (QCQP), a more complex class of convex optimization problems than the standard QPs. I conclude by demonstrating this control approach in simulation on the systems in Section 6.2, demonstrating a significant improvement over the standard continuous-time (CLF-QP) controller.

6.4 Sampled-Data Safety

In this section I will present work on sampled-data safety-critical control through approximate discrete-time models. Building off of the work in Section 6.3, I will propose a framework for achieving theoretically meaningful safety guarantees for sampled-data systems through controllers that are specified via convex optimization problems. I will begin by considering the class of Runge-Kutta discrete-time approximations, which include the Euler approximation used in Section 6.3, but also permit working with CBFs that have a rela-

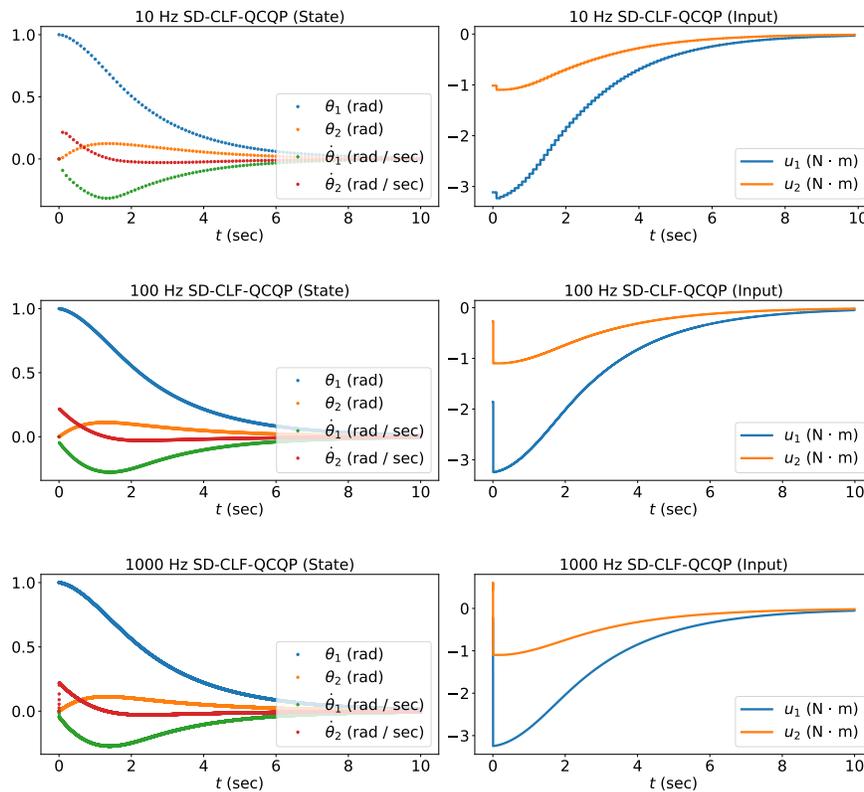


Figure 6.8. Simulation results for the double inverted pendulum system using the (CLF-QCQP) controller.

tive degree higher than one. I will show that these approximations satisfy a one-step consistency property, suggesting that designing controllers with these approximations can endow the exact family of controller-map pairs with safety guarantees. Next, I will present the notion of practical safety as a theoretical safety property for sampled-data systems. I subsequently define a notion of sampled-data barrier functions (SD-BFs) and sampled-data Control Barrier Functions (SD-CBFs) that satisfy properties uniformly in the sample rate. In the largest contribution of this work, I provide a rigorous proof that a family of SD-BFs for a Runge-Kutta approximation family of controller-map pairs leads to practical safety of the exact family of controller-map pairs. Then, I consider a common class of nonlinear systems (with integrator structure), and show that given a SD-CBF satisfying a convexity property, the order of a Runge-Kutta approximation can be chosen to ensure that the SD-CBF inequality constraint is convex, and thus can be incorporated into a convex optimization-based controller. I conclude by demonstrating the proposed control strategy in simulation.

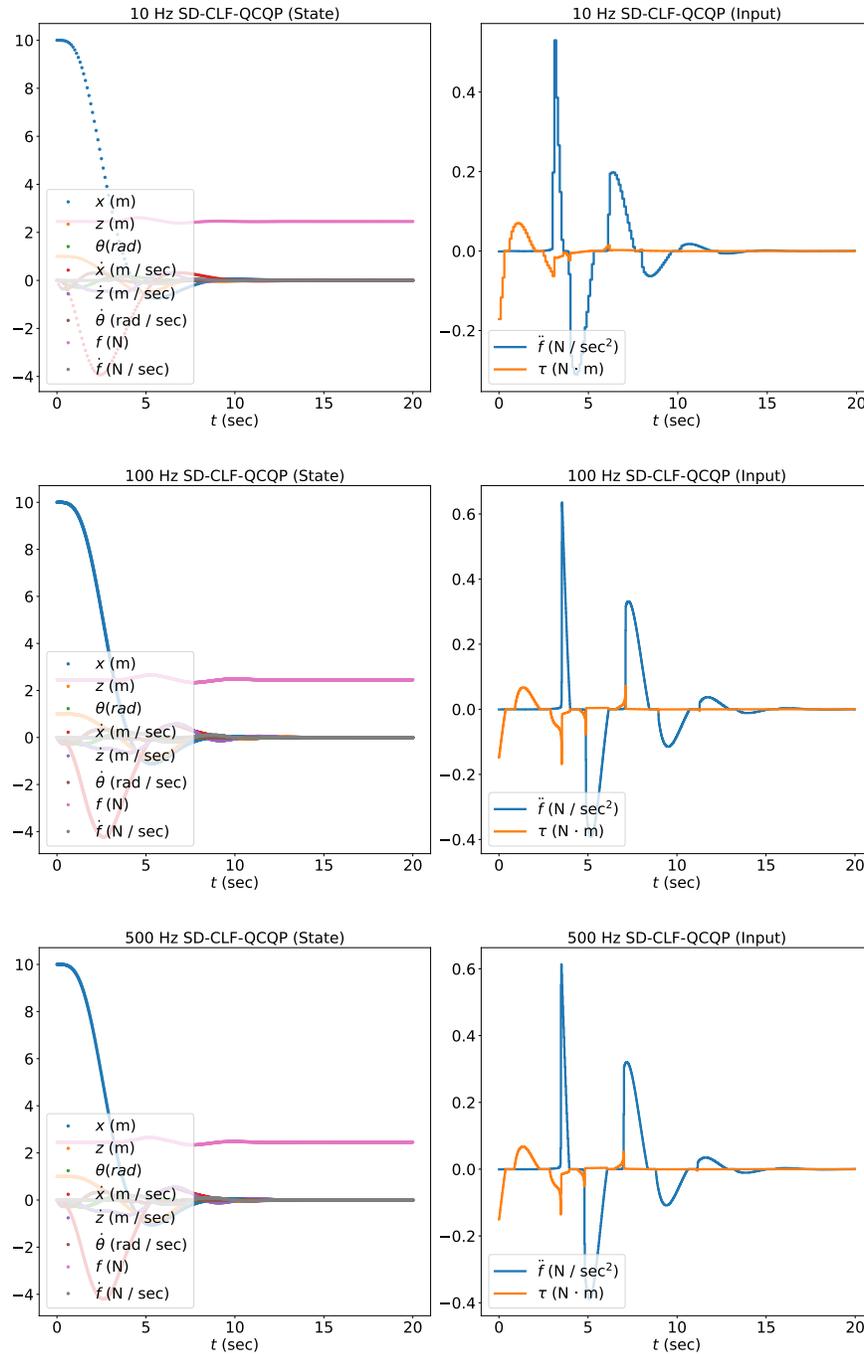


Figure 6.9. Simulation results for the planar quadrotor system using the (CLF-QCQP) controller.

The contributions of this section are as follows:

- A complete framework for safety-critical of sampled-data systems through approximate discrete-time models, including a characterization of one-step consistency of Runge-Kutta approximate discrete-time models, a

definition of practical safety, sampled-data barrier functions, and sampled-data Control Barrier Functions, and a rigorous statement relating control designs for approximate discrete-time models to practical safety of the exact family of controller-map pairs.

- A characterization of the relationship between system dynamics, SD-CBF convexity properties, and Runge-Kutta approximation order that yields controllers that can be specified via convex optimization problems, which I demonstrate in simulation.

The text for this section is adapted from:

A. J. Taylor, V. D. Dorobantu, R. K. Cosner, Y. Yue, and A. D. Ames, “Safety of sampled-data systems with control barrier functions via approximate discrete time models,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp.7127–7134.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Approximate Discrete-Time Models

I will begin by considering a larger class of discrete-time approximations that include the Euler approximation used in Section 6.3. The purpose for considering higher-order approximations is that CBFs are often not relative degree one (motivating the use of extended CBFs as was done in Section 3.7), where the CLFs considered in Section 6.3 are typically relative degree one. The importance of this is that the input \mathbf{u} may not show up in a first-order approximation of the evolution of a function with relative degree greater than one⁶. Later in this section I will show how these higher-order approximations can ensure the input appears when studying the evolution of a CBF, and provide simulation examples where higher-order approximations are necessary.

Definition 51 (*Runge-Kutta Approximation Family*). Let $p \in \mathbb{Z}_{>0}$. The *Runge-Kutta approximation family* of maps $\{\mathbf{F}_T^{a,p} \mid T \in I\}$ is defined such

⁶This is relatively easy to see for a double integrator with a CBF $h(\mathbf{x}) = x_1$. Using an Euler approximation returns $h(\mathbf{F}_T^{a,1}(\mathbf{x}, u)) = x_1 + Tx_2$, such that the input does not appear in the value h at the next time step.

that for every sample period $T \in I$, the function $\mathbf{F}_T^{a,p} : \mathcal{Z} \rightarrow \mathbb{R}^n$ is defined recursively as:

$$\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{u}) = \mathbf{x} + T \sum_{i=1}^p b_i (\mathbf{f}(\mathbf{z}_i) + \mathbf{g}(\mathbf{z}_i)\mathbf{u}), \quad (6.87)$$

$$\mathbf{z}_i = \mathbf{x} + T \sum_{j=1}^{i-1} a_{i,j} (\mathbf{f}(\mathbf{z}_j) + \mathbf{g}(\mathbf{z}_j)\mathbf{u}), \quad (6.88)$$

for all pairs $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$, with $\mathbf{z}_1 = \mathbf{x}$. Here, $b_1, \dots, b_p \in \mathbb{R}_{\geq 0}$ satisfy $\sum_{i=1}^p b_i = 1$ and $a_{i,j} \in \mathbb{R}$ for each $i \in \{1, \dots, p\}$ and $j \in \{1, \dots, i-1\}$. For a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$, I define the *Runge-Kutta approximation family* of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a,p}) \mid T \in I\}$.

I will once again make use of one-step consistency, which I define slightly differently (the introduction of the set A) from the definition in Section 6.3:

Definition 52 (*One-Step Consistency*). A family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ is *one-step consistent* with the exact family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ over a set $A \subseteq \mathcal{X}$ if there exist $\rho \in \mathcal{K}_\infty$ and $T^* \in I$ such that for all $\mathbf{x} \in A$ and $T \in (0, T^*)$, we have:

$$\|\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) - \mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))\| \leq T\rho(T). \quad (6.89)$$

Before establishing a relationship between a Runge-Kutta approximation family and one-step consistency, I state the following lemma I will use throughout this work:

Lemma 10. *For any compact set $K \subset \mathcal{X}$, there is an $\varepsilon \in \mathbb{R}_{>0}$ such that $K \oplus \overline{B}_\varepsilon \subset \mathcal{X}$ and $K \oplus \overline{B}_\varepsilon$ is compact, where \overline{B}_ε is the closed norm-ball of radius ε and \oplus is the Minkowski sum.*

Proof. As \mathcal{X} is open, for every $\mathbf{x} \in K$, there is a corresponding open ball centered at \mathbf{x} with radius $\delta_{\mathbf{x}} \in \mathbb{R}_{>0}$ that is contained in \mathcal{X} . Let $B_{\mathbf{x}} \subset \mathcal{X}$ be the open ball centered at \mathbf{x} of radius $\delta_{\mathbf{x}}/2$. Consider the collection $\{B_{\mathbf{x}} : \mathbf{x} \in K\}$; this is an open cover for the compact set K , so some finite collection $B_{\mathbf{x}_1}, \dots, B_{\mathbf{x}_N}$ for some $\mathbf{x}_1, \dots, \mathbf{x}_N \in K$, respectively, also covers K [249, Theorem 3.3.8]. Let $\delta = \min_i \delta_{\mathbf{x}_i}$, and consider any $\mathbf{z} \in K \oplus \overline{B}_{\delta/4}$. There is some $\mathbf{x} \in K$ such that $\|\mathbf{z} - \mathbf{x}\| \leq \delta/4$ and some $i \in \{1, \dots, N\}$ such that $\|\mathbf{x} - \mathbf{x}_i\| < \delta_{\mathbf{x}_i}/2$. Thus, $\|\mathbf{z} - \mathbf{x}_i\| < \delta/4 + \delta_{\mathbf{x}_i}/2 < \delta_{\mathbf{x}_i}$, so $\mathbf{z} \in \mathcal{X}$. As \mathbf{z} was

arbitrary, $K \oplus \overline{B}_{\delta/4} \subseteq \mathcal{X}$, so pick $\varepsilon \leq \delta/4$. The set $K \oplus \overline{B}_\varepsilon$ is compact as $K \times \overline{B}_\varepsilon$ is compact and $(\mathbf{x}, \mathbf{y}) \mapsto \mathbf{x} + \mathbf{y}$ is continuous. \square

I now provide the first result of this work showing how properties of the dynamics and a family of controllers can be used to establish one-step consistency of the Runge-Kutta approximation family with the exact family of controller-map pairs:

Theorem 42. *Let $K \subset \mathcal{X}$ be compact, consider a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$, and suppose there exists $T_1 \in I$ and a bound $M_K \in \mathbb{R}_{\geq 0}$ such that for every sample period $T \in (0, T_1)$, the controller \mathbf{k}_T is bounded in norm by M_K over K . Then for any $p \in \mathbb{Z}_{>0}$, the Runge-Kutta approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a,p}) \mid T \in I\}$ is one-step consistent with the family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ over K .*

Proof. Consider a compact set $K \subset \mathcal{X}$ and corresponding $T_1 \in I$ and $M_K \in \mathbb{R}_{>0}$, and fix a sample period $T \in (0, T_1)$. By Lemma 10, there exists an $\varepsilon \in \mathbb{R}_{>0}$ such that the compact set $N = K \oplus \overline{B}_\varepsilon$ satisfies $N \subset \mathcal{X}$. By assumption, \mathbf{k}_T is bounded on K , and \mathbf{f} and \mathbf{g} are bounded on N by continuity, implying there exists an $M \in \mathbb{R}_{>0}$ such that:

$$\|\mathbf{f}(\mathbf{z}) + \mathbf{g}(\mathbf{z})\mathbf{k}_T(\mathbf{y})\| \leq M, \quad (6.90)$$

for all $\mathbf{y} \in K$ and $\mathbf{z} \in N$. As \mathbf{f} and \mathbf{g} are locally Lipschitz continuous over \mathcal{X} , they are Lipschitz continuous over the compact set N . Therefore:

$$\|\mathbf{f}(\mathbf{z}) + \mathbf{g}(\mathbf{z})\mathbf{k}_T(\mathbf{y}) - (\mathbf{f}(\mathbf{y}) + \mathbf{g}(\mathbf{y})\mathbf{k}_T(\mathbf{y}))\| \quad (6.91)$$

$$\leq \|\mathbf{f}(\mathbf{z}) - \mathbf{f}(\mathbf{y})\| + \|\mathbf{g}(\mathbf{z}) - \mathbf{g}(\mathbf{y})\| \|\mathbf{k}_T(\mathbf{y})\|, \quad (6.92)$$

$$\leq (L_f + L_g M_K) \|\mathbf{z} - \mathbf{y}\|, \quad (6.93)$$

$$= \rho(\|\mathbf{z} - \mathbf{y}\|), \quad (6.94)$$

for all $\mathbf{y} \in K$ and $\mathbf{z} \in N$, with $L_f, L_g \in \mathbb{R}_{>0}$ Lipschitz constants of \mathbf{f} and \mathbf{g} on N , respectively, and $\rho \in \mathcal{K}_\infty$ satisfying $\rho(r) = (L_f + L_g M_K)r$ for all $r \in \mathbb{R}_{\geq 0}$.

Let $\mathbf{x} \in K$. Then:

$$\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) - \mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \quad (6.95)$$

$$= \int_0^T [\mathbf{f}(\boldsymbol{\varphi}(t)) + \mathbf{g}(\boldsymbol{\varphi}(t))\mathbf{k}_T(\mathbf{x})] dt - T \sum_{i=1}^p b_i (\mathbf{f}(\mathbf{z}_i) + \mathbf{g}(\mathbf{z}_i)\mathbf{k}_T(\mathbf{x})) \quad (6.96)$$

$$= \int_0^T [\mathbf{f}(\boldsymbol{\varphi}(t)) + \mathbf{g}(\boldsymbol{\varphi}(t))\mathbf{k}_T(\mathbf{x}) - (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}_T(\mathbf{x}))] dt \\ + T \sum_{i=1}^p b_i [\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}_T(\mathbf{x}) - (\mathbf{f}(\mathbf{z}_i) + \mathbf{g}(\mathbf{z}_i)\mathbf{k}_T(\mathbf{x}))], \quad (6.97)$$

where I use the fact $\sum_{i=1}^p b_i = 1$. To bound the first term in (6.97), let $T_2 \in (0, T_1)$ satisfy $T_2 < \varepsilon/M$. By continuity of the solution $\boldsymbol{\varphi}$, if $\boldsymbol{\varphi}(t_0) \notin N$ for any $t_0 \in I$, then there is a minimal $t^* \in (0, t_0)$ such that $\|\boldsymbol{\varphi}(t) - \mathbf{x}\| < \varepsilon$ ($\boldsymbol{\varphi}(t) \in N$) for all $t \in [0, t^*)$ and $\|\boldsymbol{\varphi}(t^*) - \mathbf{x}\| = \varepsilon$. We have:

$$\|\boldsymbol{\varphi}(t) - \mathbf{x}\| \leq \int_0^t \|\mathbf{f}(\boldsymbol{\varphi}(s)) + \mathbf{g}(\boldsymbol{\varphi}(s))\mathbf{k}_T(\mathbf{x})\| ds \leq Mt, \quad (6.98)$$

for all $t \in [0, t^*]$. Since $\varepsilon = \|\boldsymbol{\varphi}(t^*) - \mathbf{x}\| \leq Mt^*$, we know that $t^* \geq \varepsilon/M > T_2$. Thus if $T \in (0, T_2)$, then:

$$\|\boldsymbol{\varphi}(t) - \mathbf{x}\| \leq Mt \leq MT < MT_2 < \varepsilon, \quad (6.99)$$

for all $t \in [0, T]$, implying $\boldsymbol{\varphi}(t) \in N$ for all $t \in [0, T]$. To bound the second term in (6.97), I show by induction that if T is sufficiently small, then $\mathbf{z}_i \in N$ for all $i \in \{1, \dots, p\}$. First, since $\mathbf{z}_1 = \mathbf{x}$, we have $\mathbf{z}_1 \in N$. Next, for $i \in \{1, \dots, p\}$, suppose $\mathbf{z}_j \in N$ for all $j \in \{1, \dots, i-1\}$. Considering the definition of \mathbf{z}_i in (6.88) and the bound (6.90):

$$\|\mathbf{z}_i - \mathbf{x}\| \leq T \sum_{j=1}^{i-1} |a_{i,j}| \|\mathbf{f}(\mathbf{z}_j) + \mathbf{g}(\mathbf{z}_j)\mathbf{k}_T(\mathbf{x})\|, \quad (6.100)$$

$$\leq MT \sum_{j=1}^{i-1} |a_{i,j}| \leq MT(p-1) \max_{j,k} |a_{j,k}| \triangleq LT. \quad (6.101)$$

Let $T^* \in (0, T_2)$ satisfy $T^* < \varepsilon/L$. Then for $T \in (0, T^*)$, we have $\|\mathbf{z}_i - \mathbf{x}\| < \varepsilon$, or $\mathbf{z}_i \in N$. Since this choice of T^* does not depend on i , we can conclude by induction that if $T \in (0, T^*)$, then $\mathbf{z}_i \in N$ for all $i \in \{1, \dots, p\}$. I have shown that if $T \in (0, T^*)$, then $\boldsymbol{\varphi}(t) \in N$ for all $t \in [0, T]$, and $\mathbf{z}_i \in N$ for

$i \in \{1, \dots, p\}$. Thus, using the bound (6.94) in (6.97), we have that:

$$\|\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) - \mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))\|, \quad (6.102)$$

$$\leq \int_0^T \rho(\|\boldsymbol{\varphi}(t) - \mathbf{x}\|) dt + T \sum_{i=1}^p b_i \rho(\|\mathbf{z}_i - \mathbf{x}\|), \quad (6.103)$$

$$\leq T\rho(MT) + T \sum_{i=1}^p b_i \rho(LT), \quad (6.104)$$

$$\leq T\tilde{\rho}(T), \quad (6.105)$$

and $\tilde{\rho} \in \mathcal{K}_\infty$ satisfies $\tilde{\rho}(r) = \rho(Mr) + \rho(Lr)$ for all $r \in \mathbb{R}_{\geq 0}$. \square

Practical Safety & Sampled-Data Control Barrier Functions

I now develop a notion of *practical safety* for sampled-data systems, define *Sampled-Data Control Barrier Functions* (SD-CBFs) as a tool for safety-critical sampled-data control synthesis, and highlight how a function with 0 as a regular value satisfies a specific property required by SD-CBFs. I begin with the following definition relating the evolution of a sampled-data system and a set:

Definition 53 (*Discrete-Time Forward Invariance*). A set $\mathcal{C} \subseteq \mathcal{X}$ is *forward invariant* for a controller-map pair (\mathbf{k}, \mathbf{F}) if for every $\mathbf{x}_0 \in \mathcal{C}$ and number of steps $k \in \mathbb{Z}_{\geq 0}$, the recursion $\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k))$ is well-defined and satisfies $\mathbf{x}_k \in \mathcal{C}$.

This definition of forward invariance requires that the system state be contained in the set \mathcal{C} only at sample times, which is aligned with the notion of practical stability for sampled-data systems presented in [81] and used in Section 6.3. This differs from the standard definition of forward invariance used in the existing sampled-data safety literature, which additionally requires that the solution remain in the set \mathcal{C} between sample times, i.e. $\boldsymbol{\varphi}(t) \in \mathcal{C}$ for $t \in [t_k, t_{k+1}]$. As seen in this literature, requiring inter-sample safety typically requires selecting control actions that meet a robustified continuous-time CBF time derivative condition. This robust condition typically depends on parameters of the system that are difficult to estimate, and using over-approximations may produce very conservative behavior [90]. Reducing this conservativeness usually amounts to operating at exceedingly high sample rates, which may not be practical, and which may excite unmodeled features of the system dynamics. Moreover, in practice, inter-sample safety violations at high sample rates can be inconsequential (and may not even be detectable).

As we are focused on designing controllers using an approximate family of maps, the following definition will describe the safety properties of the exact family of controller-map pairs when design uses approximations:

Definition 54 (*Practical Safety*). The family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ is *practically safe* with respect to a set $\mathcal{C} \subseteq \mathcal{X}$ if for each $R \in \mathbb{R}_{>0}$, there exists a $T^* \in I$ such that for each sample period $T \in (0, T^*)$, there is a corresponding set $\mathcal{C}_T \subseteq \mathcal{X}$ that is forward invariant for the controller-map pair $(\mathbf{k}_T, \mathbf{F}_T)$ and satisfies $\mathcal{C} \subseteq \mathcal{C}_T \subseteq \mathcal{C} \oplus \overline{B}_R$.

This definition is posed to mirror that of practical stability for sampled-data systems proposed in [81] and used in Section 6.3. In particular, the burden of proof lies with small values of R . If $R' \geq R$ and \mathcal{C}_T is a forward invariant subset of $\mathcal{C} \oplus \overline{B}_R$, then it is automatically a forward invariant subset of $\mathcal{C} \oplus \overline{B}_{R'}$.

Before defining Sampled-Data Control Barrier Functions, for a non-empty set $\mathcal{C} \subseteq \mathcal{X}$, denote $d_{\mathcal{C}}(\mathbf{x}) = \inf_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{x}\|$ for all $\mathbf{x} \in \mathcal{X}$. I now define Sampled-Data Barrier Functions and Sampled-Data Control Barrier Functions:

Definition 55 (*Sampled-Data Barrier Function Candidate*). Consider a set $\mathcal{C} \subseteq \mathcal{X}$. A collection of functions $\{h_T : \mathcal{X} \rightarrow \mathbb{R} \mid T \in I\}$ is a *family of Sampled-Data Barrier Function Candidates* on \mathcal{C} if there exist $T^* \in I$, a function $\alpha \in \mathcal{K}^e$, a radius $\varepsilon \in \mathbb{R}_{>0}$, and a Lipschitz constant $M \in \mathbb{R}_{>0}$ such that:

$$h_T(\mathbf{x}_1) > 0, \quad h_T(\mathbf{x}_2) = 0, \quad h_T(\mathbf{x}_3) < 0, \quad (6.106)$$

$$T\alpha(h_T(\mathbf{x}_4)) \leq h_T(\mathbf{x}_4), \quad (6.107)$$

$$|h_T(\mathbf{x}_5) - h_T(\mathbf{x}_6)| \leq M\|\mathbf{x}_5 - \mathbf{x}_6\|, \quad (6.108)$$

for all states $\mathbf{x}_1 \in \text{Int}(\mathcal{C})$, $\mathbf{x}_2 \in \partial\mathcal{C}$, $\mathbf{x}_3 \in \mathcal{X} \setminus \mathcal{C}$, $\mathbf{x}_4 \in \mathcal{C}$, $\mathbf{x}_5, \mathbf{x}_6 \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_\varepsilon)$, and sample periods $T \in (0, T^*)$. Additionally, I require that for each $\eta \in \mathbb{R}_{>0}$ there exists a $\delta \in \mathbb{R}_{>0}$ such that⁷:

$$d_{\mathcal{C}}(\mathbf{x}) > \eta \implies h_T(\mathbf{x}) < -\delta, \quad (6.109)$$

for all $\mathbf{x} \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_\varepsilon)$ and $T \in (0, T^*)$.

Definition 56 (*Sampled-Data Control Barrier Functions (SD-CBFs)*). A family of Sampled-Data Barrier Function Candidates $\{h_T \mid T \in I\}$ is a *family of*

⁷See Theorem 43 for how this property relates to regular values.

Sampled-Data Control Barrier Functions (SD-CBFs) on \mathcal{C} for the family of maps $\{\mathbf{F}_T \mid T \in I\}$ if for each state $\mathbf{x} \in \mathcal{X}$ and sample time $T \in (0, T^*)$, there exists a corresponding input $\mathbf{u} \in \mathbb{R}^m$ such that $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$ and:

$$h_T(\mathbf{F}_h(\mathbf{x}, \mathbf{u})) - h_T(\mathbf{x}) > -T\alpha(h_T(\mathbf{x})). \quad (6.110)$$

Definition 57 (*Sampled-Data Barrier Functions (SD-BFs)*). Given a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$, a family of Sampled-Data Barrier Function Candidates $\{h_T \mid T \in I\}$ is a *family of Sampled-Data Barrier Functions (SD-BFs)* on \mathcal{C} for the family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ if:

$$h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{x}) \geq -T\alpha(h_T(\mathbf{x})), \quad (6.111)$$

for all states $\mathbf{x} \in \mathcal{X}$ and sample times $T \in (0, T^*)$.

I note that the conditions in (6.106) and (6.107) are standard conditions required by barrier functions for discrete-time systems [234]. The inequalities in (6.106) imply that for each $T \in (0, T^*)$, \mathcal{C} is the 0-superlevel set of h_T . The inequality in (6.107) places a requirement on the SD-BF decrement through (6.111) that implies that for each $T \in (0, T^*)$, \mathcal{C} is forward invariant for $(\mathbf{k}_h, \mathbf{F}_h)$. The condition in (6.108) requires the SD-BF to be Lipschitz continuous over a slightly larger set than \mathcal{C} with a Lipschitz constant that is uniform in the sample period, and will be used to relate exact and approximate families of controller-map pairs through one-step consistency. The implication in (6.109) resembles a *coercivity* condition, requiring the SD-BF value to decrease locally outside of the set \mathcal{C} in a way that is uniform in the sample period. This property will be critical for producing forward invariant sets contained in $\mathcal{C} \oplus \bar{B}_R$ for arbitrarily small values of R . The distinction between (6.110) and (6.111) is that the former condition states the possibility of safe control synthesis for an open-loop system, while the latter applies as a *certificate* for a closed-loop system. These properties are illustrated in Figure 6.10.

To more clearly understand the nature of the coercivity condition (6.109), the following result makes a connection with regular values that are common with continuous-time CBFs:

Theorem 43. *If $h : \mathcal{X} \rightarrow \mathbb{R}$ is twice continuously differentiable with a compact 0-superlevel set \mathcal{C} and 0 as a regular value, then there is an $\varepsilon \in \mathbb{R}_{>0}$ such that*

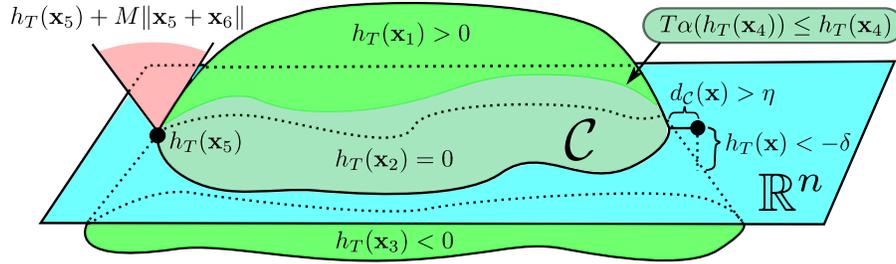


Figure 6.10. Visualizing the properties (6.106)-(6.109) of SD-BF candidates. The dark green region represents the lower bound $T\alpha(h_T(\mathbf{x}_4))$ and $h_T(\mathbf{x}_6)$ cannot be in the red region due to the Lipschitz bound.

each $\eta \in \mathbb{R}_{>0}$ corresponds to a $\delta \in \mathbb{R}_{>0}$ satisfying:

$$d_{\mathcal{C}}(\mathbf{x}) > \eta \implies h(\mathbf{x}) < -\delta, \quad (6.112)$$

for all states $\mathbf{x} \in \mathcal{X} \cap (\mathcal{C} \oplus \bar{B}_\varepsilon)$.

Proof. The boundary $\partial\mathcal{C}$ is compact as a closed subset of the compact set \mathcal{C} . Thus, $\sigma \triangleq \min_{\mathbf{x} \in \partial\mathcal{C}} \|\nabla h(\mathbf{x})\|$ is strictly positive since 0 is a regular value. By Lemma 10, there is an $\varepsilon' \in \mathbb{R}_{>0}$ with $\mathcal{C} \oplus \bar{B}_{\varepsilon'} \subset \mathcal{X}$ and $\mathcal{C} \oplus \bar{B}_{\varepsilon'}$ compact. Consider a state $\mathbf{x} \in \mathcal{C} \oplus \bar{B}_{\varepsilon'}$ with $\mathbf{x} \notin \mathcal{C}$. There exists a $\mathbf{y} \in \partial\mathcal{C}$ such $d_{\mathcal{C}}(\mathbf{x}) = \|\mathbf{y} - \mathbf{x}\| > 0$. Since h has 0 as a regular value, by [329, Proposition 1.1.9] we have that:

$$\nabla h(\mathbf{y}) = -\|\nabla h(\mathbf{y})\| \frac{\mathbf{x} - \mathbf{y}}{\|\mathbf{x} - \mathbf{y}\|}, \quad (6.113)$$

that is $\nabla h(\mathbf{y})$ is anti-parallel to $\mathbf{x} - \mathbf{y}$. As $\bar{B}_{\varepsilon'}$ is convex, $(1 - \lambda)\mathbf{y} + \lambda\mathbf{x} \in \mathcal{C} \oplus \bar{B}_{\varepsilon'}$ for all $\lambda \in [0, 1]$, and Lagrange's Remainder Formula [249, Theorem 6.6.3] implies that for some $\lambda^* \in [0, 1]$, $\boldsymbol{\xi} \triangleq (1 - \lambda^*)\mathbf{y} + \lambda^*\mathbf{x}$ satisfies:

$$h(\mathbf{x}) = h(\mathbf{y}) + (\mathbf{x} - \mathbf{y})^\top \nabla h(\mathbf{y}) + \frac{1}{2}(\mathbf{x} - \mathbf{y})^\top \nabla^2 h(\boldsymbol{\xi})(\mathbf{x} - \mathbf{y}), \quad (6.114)$$

$$= -\|\nabla h(\mathbf{y})\| \|\mathbf{x} - \mathbf{y}\| + \frac{1}{2}(\mathbf{x} - \mathbf{y})^\top \nabla^2 h(\boldsymbol{\xi})(\mathbf{x} - \mathbf{y}). \quad (6.115)$$

Since $\mathcal{C} \oplus \bar{B}_{\varepsilon'}$ is compact, there is an upper bound $\mu \in \mathbb{R}_{\geq 0}$ such that:

$$\max_{\mathbf{z} \in \mathcal{C} \oplus \bar{B}_{\varepsilon'}} \|\nabla^2 h(\mathbf{z})\| = \mu, \quad (6.116)$$

and thus we have:

$$h(\mathbf{x}) \leq -(\sigma - \frac{\mu}{2} \|\mathbf{x} - \mathbf{y}\|) \|\mathbf{x} - \mathbf{y}\|. \quad (6.117)$$

If $\|\mathbf{x} - \mathbf{y}\| \leq \sigma/\mu$, then:

$$h(\mathbf{x}) \leq -\frac{\sigma}{2}\|\mathbf{x} - \mathbf{y}\| = -\frac{\sigma}{2}d_{\mathcal{C}}(\mathbf{x}). \quad (6.118)$$

I pick $\varepsilon \in \mathbb{R}_{>0}$ such that $\varepsilon \leq \min\{\varepsilon', \sigma/\mu\}$, and for any $\eta \in \mathbb{R}_{>0}$, I pick $\delta \in \mathbb{R}_{>0}$ such that $\delta < \sigma\eta/2$. \square

Practical Safety via Sampled-Data Control Barrier Functions

In this section I present the main contribution of this work by establishing how a family of SD-BFs for an approximate family of controller-map pairs can be used to ensure the practical safety of the exact family of controller-map pairs via one-step consistency:

Theorem 44. *Consider a set $\mathcal{C} \subseteq \mathcal{X}$ and a family of admissible controllers $\{\mathbf{k}_T \mid T \in I\}$. Suppose that:*

1. *There exists a family of Sampled-Data Barrier Functions on \mathcal{C} for a family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$.*
2. *There exists an $\varepsilon' \in \mathbb{R}_{>0}$ such that the family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T) \mid T \in I\}$ is one-step consistent with the exact family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ over the set $\mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{\varepsilon'})$.*

Then the exact family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ is practically safe with respect to \mathcal{C} .

Proof. Let T_1^* , α , ε , and M be defined as in Definition 55. By assumption, there exists a $T_2^* \in I$ and $\rho \in \mathcal{K}_\infty$ such that (6.89) holds for all $\mathbf{x} \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{\varepsilon'})$ and $T \in (0, T_2^*)$. As the family of controllers is assumed admissible, there is an $T_3^* \in I$ such that \mathbf{k}_T is T -admissible for each $T \in (0, T_3^*)$.

Let $R \in \mathbb{R}_{>0}$, and pick $R' \in \mathbb{R}_{>0}$ such that $R' \leq \min\{\varepsilon, \varepsilon', R\}$. By (6.109), there exist $\delta, \Delta \in \mathbb{R}_{>0}$ such that:

$$d_{\mathcal{C}}(\mathbf{x}) > R'/2 \implies h_T(\mathbf{x}) < -\delta, \quad (6.119)$$

$$d_{\mathcal{C}}(\mathbf{x}) > \delta/(2M) \implies h_T(\mathbf{x}) < -\Delta, \quad (6.120)$$

for all $\mathbf{x} \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{\varepsilon})$ and $T \in (0, T_1^*)$. Fix $T \in I$ with $T < \min\{T_1^*, T_2^*, T_3^*\}$.

For any $c \in \mathbb{R}$, I denote the c -superlevel set of h_T as:

$$\Omega_{c,T} = \{\mathbf{x} \in \mathcal{X} \mid h_T(\mathbf{x}) \geq c\}. \quad (6.121)$$

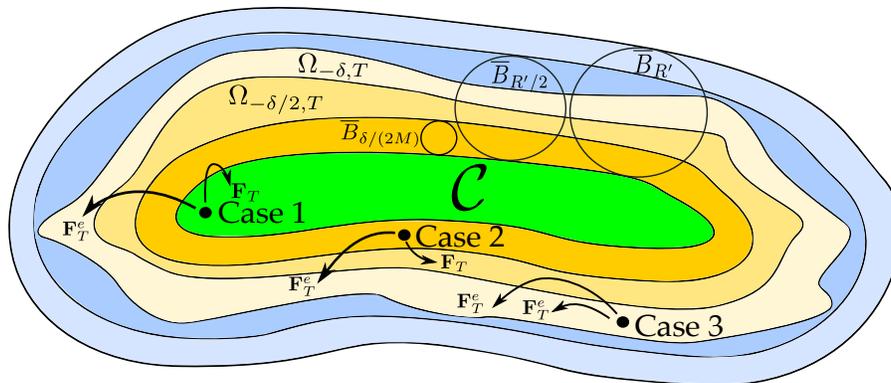


Figure 6.11. A visual representation of the main sets and three cases discussed in the proof of Theorem 44.

For any state $\mathbf{x} \in \Omega_{-\delta, T}$, we have $d_C(\mathbf{x}) \leq R'/2$, and thus $\mathcal{C} \subseteq \Omega_{-\delta, T} \subseteq \mathcal{X} \cap (\mathcal{C} \oplus \overline{B_{R'/2}}) \subseteq \mathcal{C} \oplus \overline{B_R}$.

I will prove that for small enough T , the set $\Omega_{-\delta, T}$ is forward invariant for the controller-map pair $(\mathbf{k}_T, \mathbf{F}_T^e)$. I denote three cases, as seen in Figure 6.11, considering a state $\mathbf{x} \in \mathcal{X}$ such that either:

1. $\mathbf{x} \in \mathcal{C}$,
2. $\mathbf{x} \in \Omega_{-\delta, T} \setminus \mathcal{C}$ and $d_C(\mathbf{x}) \leq \delta/(2M)$,
3. $\mathbf{x} \in \Omega_{-\delta, T} \setminus \mathcal{C}$ and $d_C(\mathbf{x}) > \delta/(2M)$.

Case 1: Suppose $\mathbf{x} \in \mathcal{C}$. From (6.111) and (6.107), we have:

$$h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{x}) \geq -T\alpha(h_T(\mathbf{x})) \geq -h_T(\mathbf{x}), \quad (6.122)$$

so $h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq 0$, or $\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \mathcal{C}$. By one-step consistency, we have:

$$\|\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) - \mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))\| \leq T\rho(T), \quad (6.123)$$

so if $T\rho(T) \leq \varepsilon$, then $\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B_\varepsilon})$. Thus:

$$|h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x})))| \leq MT\rho(T), \quad (6.124)$$

and if $MT\rho(T) \leq \delta$ as well, then:

$$h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - MT\rho(T) \geq -\delta, \quad (6.125)$$

giving us $\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \Omega_{-\delta, T}$. The analysis of this case gives us the requirement $T\rho(T) \leq \min\{\varepsilon, \delta/M\}$.

Before continuing to cases 2 and 3, I establish some additional properties. First, note that the superlevel sets have the containment property $\Omega_{-\delta/2, T} \subseteq \Omega_{-\delta, T}$. Next, for any $\eta \in \mathbb{R}_{>0}$ and any $\mathbf{x} \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_\varepsilon)$ with $\mathbf{x} \notin \mathcal{C}$, there is a state $\mathbf{y} \in \mathcal{C}$ such that $\|\mathbf{x} - \mathbf{y}\| < d_{\mathcal{C}}(\mathbf{x}) + \eta$. Therefore:

$$h_T(\mathbf{x}) \geq h_T(\mathbf{y}) - M\|\mathbf{x} - \mathbf{y}\| \geq -Md_{\mathcal{C}}(\mathbf{x}) - M\eta, \quad (6.126)$$

since $h_T(\mathbf{y}) \geq 0$. Since η can be chosen arbitrarily small, we have $h_T(\mathbf{x}) \geq -Md_{\mathcal{C}}(\mathbf{x})$. If $d_{\mathcal{C}}(\mathbf{x}) \leq \delta/(2M)$, then $h_T(\mathbf{x}) \geq -\delta/2$, so $\mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{\delta/(2M)}) \subseteq \Omega_{-\delta/2, T} \subseteq \Omega_{-\delta, T}$.

Next, consider $\mathbf{x} \in \Omega_{-\delta, T} \setminus \mathcal{C}$. Since $\mathbf{x} \notin \mathcal{C}$, meaning $h_T(\mathbf{x}) < 0$ and thus $\alpha(h_T(\mathbf{x})) < 0$, we have from (6.111) that:

$$h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq h_T(\mathbf{x}) - T\alpha(h_T(\mathbf{x})) > -\delta. \quad (6.127)$$

Thus $\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \Omega_{-\delta, T} \subseteq \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{R'/2})$ so we can apply one-step consistency to achieve:

$$\|\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) - \mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))\| \leq T\rho(T). \quad (6.128)$$

If $T\rho(T) \leq R'/2$, then $\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{R'})$, in which case the Lipschitz property of h_T yields the bound:

$$|h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x})))| \leq MT\rho(T). \quad (6.129)$$

Note that because $R'/2 < \varepsilon$, the requirement from Case 1 can be replaced by $T\rho(T) \leq \min\{R'/2, \delta/M\}$.

Case 2: Suppose $\mathbf{x} \in \Omega_{-\delta, T} \setminus \mathcal{C}$ and $d_{\mathcal{C}}(\mathbf{x}) \leq \delta/(2M)$. Since $\mathbf{x} \notin \mathcal{C}$ and $\mathcal{X} \cap (\mathcal{C} \oplus \overline{B}_{\delta/(2M)}) \subseteq \Omega_{-\delta/2, T}$, we have $-\delta/2 \leq h_T(\mathbf{x}) < 0$. Therefore:

$$h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq h_T(\mathbf{x}) - T\alpha(h_T(\mathbf{x})) \geq -\delta/2, \quad (6.130)$$

so $\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \Omega_{-\delta/2, T}$. By adding and subtracting $h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})))$ and using (6.129), we have:

$$h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq -MT\rho(T) - \delta/2, \quad (6.131)$$

when $T\rho(T) \leq R'/2$. If $MT\rho(T) \leq \delta/2$ as well, then $h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) \geq -\delta$, or $\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \Omega_{-\delta, T}$. Thus I update the requirements to be $T\rho(T) \leq \min\{R'/2, \delta/(2M)\}$.

Case 3: Suppose $\mathbf{x} \in \Omega_{-\delta, T} \setminus \mathcal{C}$ and $d_{\mathcal{C}}(\mathbf{x}) > \delta/(2M)$. From (6.120), we have:

$$h_T(\mathbf{F}_T(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{x}) > -T\alpha(-\Delta). \quad (6.132)$$

Adding and subtracting $h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})))$ and (6.129) yield:

$$h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) > h_T(\mathbf{x}) - MT\rho(T) - T\alpha(-\Delta), \quad (6.133)$$

$$= h_T(\mathbf{x}) - T(M\rho(T) + \alpha(-\Delta)), \quad (6.134)$$

when $T\rho(T) \leq R'/2$. If $M\rho(T) \leq -\alpha(-\Delta)$ as well, then $h_T(\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) > h_T(\mathbf{x}) \geq -\delta$, or $\mathbf{F}_T^e(\mathbf{x}, \mathbf{k}_T(\mathbf{x})) \in \Omega_{-\delta, T}$.

To conclude, if both:

1. $T < \min\{T_1^*, T_2^*, T_3^*, \rho^{-1}(-\alpha(-\Delta)/M)\}$,
2. $T\rho(T) \leq \min\{R'/2, \delta/(2M)\}$,

then the set $\mathcal{C}_T \triangleq \Omega_{-\delta, T} \subseteq \mathcal{C} \oplus \overline{B}_R$ is forward invariant for the exact controller-map pair $(\mathbf{k}_T, \mathbf{F}_T^e)$, and thus the exact family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^e) \mid T \in I\}$ is practically safe with respect to \mathcal{C} . \square

Optimization-Based Controller Synthesis

I will now explore convexity of the CBF decrement condition, and define an optimization-based controller via an SD-CBF for achieving practical safety. The following result establishes how for a system with a block integrator structure, a Runge-Kutta approximation family of maps of the appropriate order can preserve a convexity property of a family $\{h_T \mid T \in I\}$:

Theorem 45. Consider $\ell, \gamma, q \in \mathbb{Z}_{\geq 0}$ such that $n = \ell\gamma$ and $q \leq \gamma$. Suppose the system dynamics have the form:

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{0}_{\ell \times \ell} & \mathbf{I}_{\ell} & & \\ & \ddots & \ddots & \\ & & \mathbf{0}_{\ell \times \ell} & \mathbf{I}_{\ell} \\ & & & \mathbf{0}_{\ell \times \ell} \end{bmatrix}}_{\mathbf{A}} \mathbf{x} + \underbrace{\begin{bmatrix} \mathbf{0}_{\ell} \\ \vdots \\ \mathbf{0}_{\ell} \\ \mathbf{f}_{\gamma}(\mathbf{x}) + \mathbf{g}_{\gamma}(\mathbf{x})\mathbf{u} \end{bmatrix}}_{\mathbf{r}(\mathbf{x}, \mathbf{u})}, \quad (6.135)$$

where $\mathbf{f}_\gamma : E \rightarrow \mathbb{R}^\ell$ and $\mathbf{g}_\gamma : E \rightarrow \mathbb{R}^{\ell \times m}$. For each $T \in I$, consider a function $h_T : \mathcal{X} \rightarrow \mathbb{R}$, and suppose there exists a function $\tilde{h}_T : (\mathbb{R}^\ell)^q \rightarrow \mathbb{R}$ satisfying:

$$h_T(\mathbf{x}) = \tilde{h}_T(\boldsymbol{\zeta}_1, \dots, \boldsymbol{\zeta}_q), \quad (6.136)$$

for all $\mathbf{x} = (\boldsymbol{\zeta}_1, \dots, \boldsymbol{\zeta}_\gamma) \in \mathcal{X}$. If the function \tilde{h}_T is concave with respect to its last argument and $p = \gamma - q + 1$, then for $\alpha \in \mathcal{K}^e$, the function $\phi_T : \mathcal{Z} \rightarrow \mathbb{R}$ defined as:

$$\phi_T(\mathbf{x}, \mathbf{u}) = -h_T(\mathbf{F}_T^{\alpha,p}(\mathbf{x}, \mathbf{u})) + h_T(\mathbf{x}) - T\alpha(h_T(\mathbf{x})), \quad (6.137)$$

is convex in its second argument.

Proof. For all $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$, denote:

$$\mathbf{F}_T^{\alpha,p}(\mathbf{x}, \mathbf{u}) = ((\mathbf{F}_1)_T^{\alpha,p}(\mathbf{x}, \mathbf{u}), \dots, (\mathbf{F}_\gamma)_T^{\alpha,p}(\mathbf{x}, \mathbf{u})), \quad (6.138)$$

where $(\mathbf{F}_i)_T^{\alpha,p} : \mathcal{Z} \rightarrow \mathbb{R}^\ell$ for all $i \in \{1, \dots, \gamma\}$. For $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$, the block vector $\mathbf{r}(\mathbf{x}, \mathbf{u})$ can be nonzero only in the last (γ th) block. Noting the block chain-of-integrators structure of \mathbf{A} , for any degree $d \in \{0, \dots, \gamma - 1\}$, $\mathbf{A}^d \mathbf{r}(\mathbf{x}, \mathbf{u})$ can be nonzero only in the $(\gamma - d)$ th block, and for a degree d polynomial ρ_d , $\rho_d(\mathbf{A})\mathbf{r}(\mathbf{x}, \mathbf{u})$ can be nonzero only in the last $d + 1$ blocks (that is, blocks $\gamma - d$ through γ).

Consider a state-input pair $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$. We have:

$$\mathbf{F}_T^{\alpha,p}(\mathbf{x}, \mathbf{u}) = \mathbf{x} + T \sum_{i=1}^p b_i(\mathbf{A}\mathbf{z}_i + \mathbf{r}(\mathbf{z}_i, \mathbf{u})), \quad (6.139)$$

$$\mathbf{z}_i = \mathbf{x} + T \sum_{j=1}^{i-1} a_{i,j}(\mathbf{A}\mathbf{z}_j + \mathbf{r}(\mathbf{z}_j, \mathbf{u})), \quad (6.140)$$

with $\mathbf{z}_1 = \mathbf{x}$. By induction, for any $i \in \{1, \dots, p\}$, I show:

$$\mathbf{z}_i = \rho_{i,i-1}(\mathbf{A})\mathbf{x} + \sum_{j=1}^{i-1} \sigma_{i,i-j-1}(\mathbf{A})\mathbf{r}(\mathbf{z}_j, \mathbf{u}), \quad (6.141)$$

where $\rho_{i,i-1}$ is a degree $i - 1$ polynomial, and for $j \in \{1, \dots, i - 1\}$, $\sigma_{i,i-j-1}$ is a degree $i - j - 1$ polynomial. Indeed, $\mathbf{z}_1 = \mathbf{I}_n \cdot \mathbf{x}$, and assuming (6.141) holds

for $0, \dots, i-1$, substituting (6.141) into (6.140) yields the following:

$$\begin{aligned} \mathbf{z}_i &= \underbrace{\left(\mathbf{I} + T \sum_{j=1}^{i-1} a_{i,j} \overbrace{\mathbf{A} \rho_{j,j-1}(\mathbf{A})}^{\text{degree } j} \right)}_{\triangleq \rho_{i,i-1}(\mathbf{A})} \mathbf{x} + T \sum_{j=1}^{i-1} a_{i,j} \mathbf{r}(\mathbf{z}_j, \mathbf{u}) \\ &\quad + T \sum_{k=1}^{i-1} \sum_{j=1}^{k-1} a_{i,k} \mathbf{A} \sigma_{k,k-j-1}(\mathbf{A}) \mathbf{r}(\mathbf{z}_j, \mathbf{u}), \end{aligned} \quad (6.142)$$

which we may further manipulate to obtain:

$$\mathbf{z}_i - \rho_{i,i-1}(\mathbf{A}) \mathbf{x} = \sum_{j=1}^{i-1} T \underbrace{\left(a_{i,j} + \sum_{k=j+1}^{i-1} a_{i,k} \overbrace{\mathbf{A} \sigma_{k,k-j-1}(\mathbf{A})}^{\text{degree } k-j} \right)}_{\text{degree } i-j-1} \mathbf{r}(\mathbf{z}_j, \mathbf{u}), \quad (6.143)$$

$$\triangleq \sum_{j=1}^{i-1} \sigma_{i,i-j-1}(\mathbf{A}) \mathbf{r}(\mathbf{z}_j, \mathbf{u}), \quad (6.144)$$

establishing (6.141) holds for i . Substituting the expression (6.141) into (6.139) and following a similar sequence of steps, we find a degree p polynomial $\tilde{\rho}_p$, and for each $i \in \{1, \dots, p\}$, a degree $p-i$ polynomial $\tilde{\sigma}_{p-i}$ such that:

$$\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{u}) = \tilde{\rho}_p(\mathbf{A}) \mathbf{x} + \sum_{i=1}^p \tilde{\sigma}_{p-i}(\mathbf{A}) \mathbf{r}(\mathbf{z}_i, \mathbf{u}). \quad (6.145)$$

For $i \in \{1, \dots, p\}$, the term $\tilde{\sigma}_{p-i}(\mathbf{A}) \mathbf{r}(\mathbf{z}_i, \mathbf{u})$ can be nonzero only in blocks $\gamma - (p-i) = q + i - 1$ through γ . The highest-order polynomial multiplying the block vectors $\mathbf{r}(\mathbf{z}_1, \mathbf{u}), \dots, \mathbf{r}(\mathbf{z}_p, \mathbf{u})$ is $\tilde{\sigma}_{p-1} = \tilde{\sigma}_{\gamma-q}$. Therefore, the functions $(\mathbf{F}_1)_h^{a,p}, \dots, (\mathbf{F}_{q-1})_T^{a,p}$ are independent of their second argument (they depend only on state). Moreover, $(\mathbf{F}_q)_T^{a,p}(\mathbf{x}, \mathbf{u})$ depends on the block vector $\mathbf{r}(\mathbf{z}_1, \mathbf{u}) = \mathbf{r}(\mathbf{x}, \mathbf{u})$, which depends on \mathbf{u} affinely, and does not depend on the block vectors $\mathbf{r}(\mathbf{z}_2, \mathbf{u}), \dots, \mathbf{r}(\mathbf{z}_p, \mathbf{u})$, which may depend on \mathbf{u} nonlinearly. The composition $h_T \circ \mathbf{F}_T^{a,p} : \mathcal{Z} \rightarrow \mathbb{R}$ satisfies:

$$h_T(\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{u})) = \tilde{h}_T((\mathbf{F}_1)_T^{a,p}(\mathbf{x}, \mathbf{u}), \dots, (\mathbf{F}_q)_T^{a,p}(\mathbf{x}, \mathbf{u})), \quad (6.146)$$

for all $(\mathbf{x}, \mathbf{u}) \in \mathcal{Z}$. The composition of concave and affine functions is concave, so $h_T \circ \mathbf{F}_T^{a,p}$ is concave in its second argument, and ϕ_T in (6.137) is convex in its second argument. \square

The following result highlights how we may synthesize a family of controllers that achieve practical safety through optimization using only a Runge-Kutta approximation family of maps and not requiring knowledge of the exact family of maps:

Theorem 46. *Let $\{h_T \mid T \in I\}$ be a family of SD-CBFs on \mathcal{C} for a Runge-Kutta approximation family of map $\{\mathbf{F}_T^{a,p} \mid T \in I\}$ such that the set:*

$$\mathcal{F}_T(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^m \mid (\mathbf{x}, \mathbf{u}) \in \mathcal{Z} \text{ and } \phi_T(\mathbf{x}, \mathbf{u}) \leq 0\}, \quad (6.147)$$

is closed and convex for each $T \in I$ and $\mathbf{x} \in \mathcal{X}$. Consider a set of controllers $\{\mathbf{k}_T \mid T \in I\}$ satisfying:

$$\begin{aligned} \mathbf{k}_T(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 & \quad (\text{SD-CBF-OP}) \\ \text{s.t. } h_T(\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{u})) - h_T(\mathbf{x}) & \geq -T\alpha(h_T(\mathbf{x})), \end{aligned}$$

for each $\mathbf{x} \in \mathcal{X}$ and $T \in (0, T^)$, where $\mathbf{k}_{\text{nom}} : \mathcal{X} \rightarrow \mathbb{R}^m$ is a nominal controller. If $\{\mathbf{k}_T \mid T \in I\}$ is a family of admissible controllers, then $\{h_T \mid T \in I\}$ is a family of Sampled-Data Barrier Functions on \mathcal{C} for the Runge-Kutta approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a,p}) \mid T \in I\}$.*

Proof. Consider $T \in (0, T^*)$ and $\mathbf{x} \in \mathcal{X}$. As h_T is a SD-CBF on \mathcal{C} , there is a $\mathbf{u}' \in \mathbb{R}^m$ such that $(\mathbf{x}, \mathbf{u}') \in \mathcal{Z}$ and:

$$h_T(\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{u}')) - h_T(\mathbf{x}) \geq -T\alpha(h_T(\mathbf{x})), \quad (6.148)$$

implying that $\mathbf{u}' \in \mathcal{F}_T(\mathbf{x})$. Thus the optimization problem in (SD-CBF-OP) is feasible. Define the compact, convex set:

$$A = \{\mathbf{u} \in \mathbb{R}^m \mid \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \leq \|\mathbf{u}' - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2\}. \quad (6.149)$$

Note that $\mathbf{u}' \in A$. As $\mathcal{F}_T(\mathbf{x})$ is closed and convex, $A \cap \mathcal{F}_T(\mathbf{x})$ is compact, convex, and non-empty. As the cost is continuous and strictly convex with respect to \mathbf{u} , there is a unique minimizer $\mathbf{u}^* \in A \cap \mathcal{F}_T(\mathbf{x})$. We have $\|\mathbf{u}^* - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 \leq \|\mathbf{u}' - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2 < \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x})\|^2$ for all $\mathbf{u} \in \mathcal{F}_T(\mathbf{x}) \setminus A$, implying \mathbf{u}^* is the unique minimizer in $\mathcal{F}_T(\mathbf{x})$ (the argmin is well-defined). Thus:

$$h_T(\mathbf{F}_T^{a,p}(\mathbf{x}, \mathbf{k}_T(\mathbf{x}))) - h_T(\mathbf{x}) \geq -T\alpha(h_T(\mathbf{x})), \quad (6.150)$$

and as \mathbf{x} and T were arbitrary, we have that $\{h_T \mid T \in I\}$ is a family of SD-BFs on \mathcal{C} for the Runge-Kutta approximation family of controller-map pairs $\{(\mathbf{k}_T, \mathbf{F}_T^{a,p}) \mid T \in I\}$. \square

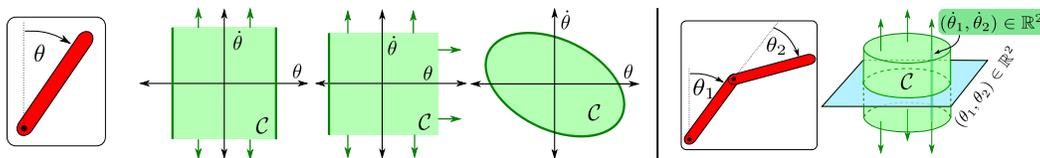


Figure 6.12. The single (left) and double (right) inverted pendulums and their sets \mathcal{C} . The inverted pendulum sets \mathcal{C} are, from left to right, a configuration ellipsoid, a halfspace, and a Lyapunov sublevel set. The far right image shows the set \mathcal{C} for the double inverted pendulum where $(\theta_1, \theta_2) \in \mathbb{R}^2$ are constrained to an ellipse and no constraint is placed on $(\dot{\theta}_1, \dot{\theta}_2) \in \mathbb{R}^2$ (represented by vertical fibers).

Simulation Results

I now evaluate the proposed (SD-CBF-OP) on the inverted pendulum and double inverted pendulum in Section 6.2. With state vector $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^n$, the dynamics of these systems can be expressed in the form (6.135), where $\ell = m$ and $\gamma = 2$. For the inverted pendulum, I use sets \mathcal{C} with the form of a Lyapunov sublevel ($h_T(\mathbf{x}) = 1 - \mathbf{x}^\top \mathbf{P} \mathbf{x}$ with $\mathbf{P} \in \mathbb{S}_{>0}^2$), a configuration ellipsoid ($h_T(\theta) = 1 - \theta^2$), and a halfspace ($h_T(\theta) = \theta + 0.1$). For the double inverted pendulum I enforce safety of a configuration ellipsoid ($h_T(\mathbf{q}) = 1 - \|\mathbf{q}\|^2$). These sets are visualized in Figure 6.12. I use Runge-Kutta approximations with $p = 1$ (forward Euler) for the Lyapunov sublevel set and $p = 2$ (midpoint rule) for the other settings. Controllers of the form (SD-CBF-OP) are employed with identity comparison functions; for the Lyapunov sublevel set, \mathbf{k}_{nom} is a feedback linearizing controller with auxiliary PD control (proportional gain 1, derivative gain 2), and for the other settings, \mathbf{k}_{nom} is a zero (constant) controller. With 11 sample periods spaced logarithmically (over $[0.05, 0.5]$ and $[0.01, 0.1]$ seconds for the single and double inverted pendulums, respectively) and initial conditions sampled from each set \mathcal{C} , the sampled-data closed-loop systems are simulated for 10 seconds. For the inverted pendulum, 500 initial states are sampled uniformly from the Lyapunov sublevel set, and 41×41 grids of initial states cover $[-1, 1] \times [-5, 5]$ for the configuration ellipsoid and $[-0.1, 1] \times [-5, 5]$ for the halfspace. For the double inverted pendulum, 500 initial states are uniformly sampled with configurations in the unit Euclidean ball in \mathbb{R}^2 and velocities in $[-1, 1]^2$. The worst-case distances from the safe sets are reported as a function of sample period in Figure 6.13. These distances decrease for sufficiently small sample periods, as predicted by the definition of practical safety.

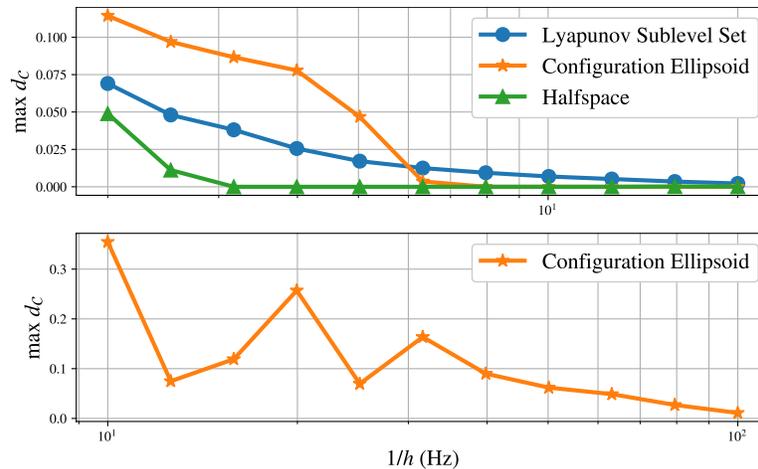


Figure 6.13. The maximum distance from the safe set \mathcal{C} (lower is better) achieved during trials vs. the sampling frequency. The simulations and animations be found at <https://bit.ly/CBF-OP> and <https://vimeo.com/690803272>. (Top) The inverted pendulum for 3 different safe sets. (Bottom) The double inverted pendulum.

Conclusion

In this section I have presented my work on sampled-data safety-critical control through approximate discrete-time models. I show how Runge-Kutta approximation models satisfy important consistency properties, and define the concepts of practical safety, sampled-data barrier functions, and sampled-data Control Barrier Functions. In the main contribution of this work, I show how a family of sampled-data barrier functions for a Runge-Kutta approximation family of controller-map pairs implies the exact family of controller-map pairs is practically safe, permitting theoretical guarantees for the exact discrete-time systems using controllers designed with approximations. I then show how for a common class of nonlinear systems, the convexity properties of a SD-CBF can be preserved with a Runge-Kutta approximation family of a specific order, allow safety-critical controllers using approximate discrete-time models to be specified as convex optimization problems. Lastly I demonstrate these controllers in a variety of simulation settings, highlighting that worst-case safety violations decrease with increasing sample rate as predicted by practical safety.

6.5 Event-Triggered Safety

In this section I will present work on safety-critical event-triggered control. I will begin by reviewing event-triggered control for stabilization using ISS-LFs as presented in [96]. Drawing inspiration from this, I will propose a direct

translation of the ideas in this paper to the safety-critical setting using ISSf-BFs. I will then demonstrate that this direct translation does not lead to a minimum interevent time (MIET) through a rigorous counterexample. I resolve the issues shown with this counter example through a concept known as the *strong event-triggered ISSf barrier property*, which requires a more stringent degree of safe behavior on the boundary of a set that is to be kept forward invariant. Lastly, I demonstrate that this trigger law ensures both safety and an MIET in simulation on the system used in the counterexample.

The contributions of this section are as follows:

- The first framework for safety-critical event-triggered control that takes inspiration from [96] by utilizing the tools of ISSf-BFs.
- A rigorous counterexample highlighting challenges not seen when performing event-triggered stabilization, and a proposed method for resolving these challenges.

The text for this section is adapted from:

A. J. Taylor, P. Ong, J. Cortés, and A. D. Ames, “Safety-critical event triggered control via input-to-state safe barrier functions,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 749–754, 2021.

A. J. Taylor participated in the conception of the project, algorithm design and theoretical analysis, simulation code implementation, and writing of the article.

Event-Triggered Stability

I will now review the approach for event-triggered stabilization as presented in [96], which will inspire how I develop my proposed approach for safety-critical event-triggered control. Let $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a controller, and recall the event-triggered closed-loop system:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t_k)) \quad \forall t \in [t_k, t_{k+1}). \quad (6.151)$$

I define the sampling error $\mathbf{e} : [t_k, t_{k+1}) \rightarrow \mathbb{R}^n$ as:

$$\mathbf{e}(t) = \mathbf{x}(t_k) - \mathbf{x}(t) \quad \forall t \in [t_k, t_{k+1}), \quad (6.152)$$

noting that this in conjunction with (6.151) implies:

$$\dot{\mathbf{x}}(t) = -\dot{\mathbf{e}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t_k)) \quad \forall t \in [t_k, t_{k+1}). \quad (6.153)$$

This can alternatively be written as:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)) \quad \forall t \in [t_k, t_{k+1}). \quad (6.154)$$

Let $\mathbf{x}_e \in \mathbb{R}^n$ be an equilibrium point of the continuous-time closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}). \quad (6.155)$$

I will define an ISS-LF in the event-triggered setting as follows:

Definition 58 (*Event-Triggered ISS-LF (ET-ISS-LF)*). Let $\mathbf{x}_e \in \mathbb{R}^n$ be an equilibrium point of the continuous-time closed-loop system (6.155). A function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ that is continuously differentiable on \mathbb{R}^n is said to be an *exponential Event-Triggered Input-to-Stable Lyapunov function (ET-ISS-LF)* for the continuous-time closed-loop system (6.155) and equilibrium point \mathbf{x}_e if there exist constants k_1, k_2, k_3, a , and $\sigma \in \mathcal{K}_\infty$ such that:

$$k_1 \|\mathbf{x} - \mathbf{x}_e\|^a \leq V(\mathbf{x}) \leq k_2 \|\mathbf{x} - \mathbf{x}_e\|^a, \quad (6.156)$$

$$\dot{V}(\mathbf{x}, \mathbf{e}) \triangleq L_{\mathbf{f}}V(\mathbf{x}) + L_{\mathbf{g}}V(\mathbf{x})\mathbf{k}(\mathbf{x} + \mathbf{e}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^a + \sigma(\|\mathbf{e}\|), \quad (6.157)$$

for all $\mathbf{x}, \mathbf{e} \in \mathbb{R}^n$.

Suppose that we have an ET-ISS-LF for an equilibrium point \mathbf{x}_e of the continuous-time closed-loop system (6.155). If the trigger-law is defined to enforce:

$$\sigma(\|\mathbf{e}(t)\|) \leq ck_3 \|\mathbf{x}(t) - \mathbf{x}_e\|^a, \quad (6.158)$$

for $c \in (0, 1)$, then the ET-ISS-LF condition (6.157) implies that:

$$L_{\mathbf{f}}V(\mathbf{x}(t)) + L_{\mathbf{g}}V(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)) \leq (c - 1)k_3 \|\mathbf{x}(t) - \mathbf{x}_e\|^a, \quad (6.159)$$

for all $t \in [t_k, t_{k+1})$ and $k \in \mathbb{Z}_{\geq 0}$, such that V is an exponential Lyapunov function for the the event-triggered closed-loop system (6.151) and equilibrium point \mathbf{x}_e and the event-triggered closed-loop system (6.151) is exponentially stable with respect to the equilibrium point \mathbf{x}_e . In this way, a controller that is designed to be robust to sampling error (such that (6.157) is met) can be used with an event-triggered implementation and recover the desired exponential

stability property at the expense of convergence rate ($(c-1)k_3 < k_3$ in (6.159)). The inequality in (6.158) can be enforced by defining the trigger-law as:

$$t_{k+1} = \min \{t \geq t_k \mid \sigma(\|\mathbf{e}(t)\|) = ck_3\|\mathbf{x}(t) - \mathbf{x}_e\|^a\}. \quad (6.160)$$

As is typical in event-triggered control formulations, it is critical to show that such a trigger-law does not lead to the control being updated at arbitrarily close time instances [315], or that the *interevent times* $\{t_{k+1} - t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ are lower bounded by a positive constant $\tau \in \mathbb{R}_{>0}$, referred to as the minimum interevent time (MIET). The results of [96] ensure that a MIET exists under the trigger-law (6.159) under the assumption that the functions \mathbf{f} , \mathbf{g} , and \mathbf{k} are Lipschitz continuous on any compact set $K \subset \mathbb{R}^n$ and σ is Lipschitz continuous on any compact set $K \subset \mathbb{R}_{\geq 0}$, noting that if these function are locally Lipschitz continuous on their domains, then they are Lipschitz continuous on any compact set contained in their domain.

Counterexample

I will now begin to set up an approach for safety-critical event-triggered control. In attempting to mirror the approach for stabilization, I will reveal a challenge that occurs with a naïve trigger-law near the boundary of a set \mathcal{C} that is to be kept safe. I will explore resolving this in the subsequent section.

I will define an ISSf-BF in the event-triggered context as follows:

Definition 59 (*Event-Triggered ISSf-BF (ET-ISSf-BF)*). Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ that is continuously differentiable on \mathbb{R}^n . The function h is said to be an *Event-Triggered Input-to-State Safe barrier function (ET-ISSf-BF)* for the continuous-time closed-loop system (6.155) on \mathcal{C} if there exist $\alpha \in \mathcal{K}_\infty^e$ and $\iota \in \mathcal{K}_\infty$ such that:

$$\dot{h}(\mathbf{x}, \mathbf{e}) \triangleq L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x} + \mathbf{e}) \geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{e}\|), \quad (6.161)$$

for all $\mathbf{x}, \mathbf{e} \in \mathbb{R}^n$.

Given the similarity of the ET-ISS-LF constraint (6.157) and the ET-ISSf-BF constraint (6.161), it is natural to propose a trigger-law that enforces:

$$\iota(\|\mathbf{e}(t)\|) \leq c\alpha(h(\mathbf{x}(t))), \quad (6.162)$$

for some $c \in \mathbb{R}_{>0}$, implying:

$$L_{\mathbf{f}}h(\mathbf{x}(t)) + L_{\mathbf{g}}h(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)) \geq -(1+c)\alpha(h(\mathbf{x}(t))), \quad (6.163)$$

for all $t \in [t_k, t_{k+1})$ and $k \in \mathbb{Z}_{\geq 0}$, such that h is a barrier function for the event-triggered closed-loop system (6.151) on \mathcal{C} , and thus the event-triggered closed-loop system is safe with respect to the set \mathcal{C} . This can be interpreted as conserving actuation resources at the expense of allowing the system to more quickly approach the boundary of the set \mathcal{C} from the inside. It is important to note that inside \mathcal{C} it is possible to satisfy (6.162) and thus enforce safety, but it is impossible to satisfy (6.162) outside \mathcal{C} as $\alpha(h(\mathbf{x}(t))) < 0$ if $\mathbf{x}(t) \notin \mathcal{C}$. This type of behavior does not arise in the context of event-triggered stabilization, where convergence is to a point. One way to solve this issue is to instead define the trigger-law to enforce:

$$\iota(\|\mathbf{e}(t)\|) \leq c|\alpha(h(\mathbf{x}(t)))|, \quad (6.164)$$

for $c \in (0, 1)$, which enforces (6.163) if $\mathbf{x}(t) \in \mathcal{C}$ and enforces:

$$L_{\mathbf{f}}h(\mathbf{x}(t)) + L_{\mathbf{g}}h(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)) \geq -(1 - c)\alpha(h(\mathbf{x}(t))), \quad (6.165)$$

if $\mathbf{x}(t) \notin \mathcal{C}$. In this formulation, the system is not only allowed to more quickly approach the boundary, but is also not required to converge to the set as quickly when outside of the set. This is a generalization of event-triggered stabilization to a set. Even with this solution, it is not guaranteed that this trigger law will have an MIET. Although ruling out Zeno behavior is not required to guarantee safety, unlike stabilization (unique solutions must exist for all time), it is important to have an MIET in term of implementation of the controller [97]. The key difference between stability and safety leading to the failure of an MIET to exist for a safe event-triggered controller lies in how the system dynamics must behave close to an equilibrium point compared to how they can behave close to the boundary of the set \mathcal{C} . In stabilization, the continuity of the dynamics requires that they vanish as the equilibrium is approached, leading to the error dynamics in (6.153) vanishing. In safety, the dynamics close to the boundary of the safe set need not vanish as the boundary is approached, such that the error dynamics in (6.153) need not vanish. I provide the following counterexample to illustrate how this difference can lead to a MIET failing to exist for the trigger-law design (6.164).

Consider the following system:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \underbrace{\begin{bmatrix} x_2 \\ -x_1 \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} u, \quad (6.166)$$

for which we wish to ensure the safety of the set \mathcal{C} , given by the 0-superlevel set of the continuously differentiable function $h(\mathbf{x}) = 1 - x_1^2 - x_2^2 = 1 - \|\mathbf{x}\|^2$. The time derivative of this function along solutions to (6.166) is given by $\dot{h}(\mathbf{x}, u) = -2(x_1^2 + x_2^2)u$, for which the controller $k(\mathbf{x}) = \frac{1}{2}(1 - x_1^2 - x_2^2) = \frac{1}{2}h(\mathbf{x})$ yields $\dot{h}(\mathbf{x}) = -(x_1^2 + x_2^2)h(\mathbf{x}) \geq -h(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^2$, which implies h is a BF for the resulting closed-loop system on the set \mathcal{C} .

In an event-triggered context, the closed-loop system is given by:

$$\dot{\mathbf{x}}(t) = \begin{bmatrix} k(\mathbf{x}(t_k)) & 1 \\ -1 & k(\mathbf{x}(t_k)) \end{bmatrix} \mathbf{x}(t) \quad \forall t \in [t_k, t_{k+1}). \quad (6.167)$$

This leads to the time derivative of h along solutions to (6.167) being given by:

$$\dot{h}(\mathbf{x}(t), \mathbf{e}(t)) = -\|\mathbf{x}(t)\|^2 h(\mathbf{x}(t_k)) = -\|\mathbf{x}(t)\|^2 h(\mathbf{x}(t) + \mathbf{e}(t)), \quad (6.168)$$

for $t \in [t_k, t_{k+1})$, where $\mathbf{e}(t) = \mathbf{x}(t_k) - \mathbf{x}(t)$. To see that h is in fact an ET-ISSf-BF, note that its time derivative can be bounded as follows:

$$\begin{aligned} \dot{h}(\mathbf{x}(t), \mathbf{e}(t)) &= -\|\mathbf{x}(t)\|^2 h(\mathbf{x}(t) + \mathbf{e}(t)), \\ &= -\|\mathbf{x}(t)\|^2 (1 - \|\mathbf{x}(t)\|^2 - 2\mathbf{x}(t)^\top \mathbf{e}(t) - \|\mathbf{e}(t)\|^2), \\ &\geq -\|\mathbf{x}(t)\|^2 (1 - \|\mathbf{x}(t)\|^2) - 2\|\mathbf{x}(t)\|^3 \|\mathbf{e}(t)\|, \\ &\geq -(1 - \|\mathbf{x}(t)\|^2) - 2\|\mathbf{x}(t)\|^3 \|\mathbf{e}(t)\|, \\ &\geq -h(\mathbf{x}(t)) - 2r^3 \|\mathbf{e}(t)\|, \end{aligned}$$

for all $r \in \mathbb{R}_{\geq 0}$ such that $r \geq 1$ and $\mathbf{x}(t) \in \mathbb{R}^2$ such that $\|\mathbf{x}(t)\| \leq r$. Note that for a fixed value of r , this restricts the claim to a neighborhood of the compact set \mathcal{C} . The value of r can be chosen arbitrarily large to produce an arbitrarily large neighborhood, but this detail is not critical for the subsequent results regarding the non-existence of an MIET. Given that h is an ET-ISSf-BF on some domain containing the unit circle (choose $r > 1$), the trigger-law enforcing (6.164) is given by:

$$t_{k+1} = \min\{t \geq t_k \mid 2r^3 \|\mathbf{e}(t)\| = c|h(\mathbf{x}(t))|\}, \quad (6.169)$$

with $c \in (0, 1)$. This will guarantee that \mathcal{C} is safe as:

$$\dot{h}(\mathbf{x}(t), \mathbf{e}(t)) \geq \begin{cases} -(1+c)h(\mathbf{x}(t)), & \|\mathbf{x}(t)\| \leq 1, \\ -(1-c)h(\mathbf{x}(t)), & 1 < \|\mathbf{x}(t)\| \leq r. \end{cases} \quad (6.170)$$

Despite the fact \mathbf{f} , \mathbf{g} , k , α ($\alpha(r) = r$), and ι ($\iota(r) = 2r^3$) are Lipschitz continuous on compacts, which was sufficient in the case of stabilization, the following result shows the trigger-law fails to provide an MIET:

Lemma 11. *The system (6.167) with the trigger-law defined as in (6.169) does not possess an MIET.*

Proof. To show that the interevent times $\{t_{k+1} - t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ are not lower bounded, I will proceed via contradiction. In particular, let us assume that there exists an MIET $\tau \in \mathbb{R}_{>0}$ such that $t_{k+1} - t_k \geq \tau$ for all $k \in \mathbb{Z}_{\geq 0}$. If the state $\mathbf{x}_k = \mathbf{x}(t_k)$ at event time t_k is used as an initial condition, the solution to the event-triggered closed-loop system (6.167) is:

$$\mathbf{x}(t) = \exp\left(\frac{h(\mathbf{x}_k)\Delta t_k}{2}\right) \begin{bmatrix} \cos(\Delta t_k) & \sin(\Delta t_k) \\ -\sin(\Delta t_k) & \cos(\Delta t_k) \end{bmatrix} \mathbf{x}_k, \quad (6.171)$$

$$= \mathbf{M}_k(\Delta t_k)\mathbf{x}_k, \quad (6.172)$$

for $t \in [t_k, t_{k+1})$ with $\Delta t_k = t - t_k$. Denote $\omega_k = h(\mathbf{x}_k)\Delta t_k$ and observe that the norm of the error is lower bounded by a function monotonically increasing in time:

$$\|\mathbf{e}(t)\| = \|(\mathbf{I}_2 - \mathbf{M}_k(\Delta t_k))\mathbf{x}_k\|, \quad (6.173)$$

$$= \sqrt{\left(\exp(\omega_k) - 2\exp\left(\frac{\omega_k}{2}\right)\cos(\Delta t_k) + 1\right)}\|\mathbf{x}_k\|, \quad (6.174)$$

$$\geq \sqrt{\left(\exp(\omega_k) - 2\exp\left(\frac{\omega_k}{2}\right) + 1\right)}\|\mathbf{x}_k\|, \quad (6.175)$$

$$= \sqrt{\left(\exp\left(\frac{\omega_k}{2}\right) - 1\right)^2}\|\mathbf{x}_k\|, \quad (6.176)$$

$$\geq \left|\exp\left(\frac{\omega_k}{2}\right) - 1\right|\|\mathbf{x}_k\|. \quad (6.177)$$

This lower bound on the error grows unbounded in time. This implies that no matter the state in $\text{Int}(\mathcal{C})$ that an event occurs, another event must occur at some time in the future (or the bound in (6.169) will be violated as h is upper bounded on \mathcal{C}). Thus, for all $T > 0$, there exists an event time $t_k > T$.

Next, I show that $\lim_{t \rightarrow \infty} h(\mathbf{x}(t)) = 0$. Note that:

$$h(\mathbf{x}(t)) = 1 - \|\mathbf{x}(t)\|^2 = 1 - \exp(h(\mathbf{x}_k)\Delta t_k)\|\mathbf{x}_k\|^2, \quad (6.178)$$

with time derivative:

$$\dot{h}(\mathbf{x}(t)) = -h(\mathbf{x}_k)\exp(h(\mathbf{x}_k)\Delta t_k)\|\mathbf{x}_k\|^2, \quad (6.179)$$

for $t \in [t_k, t_{k+1})$. Within the safe set we have that $\dot{h}(\mathbf{x}(t)) \leq 0$, such that $h(\mathbf{x}(t))$ is monotonically decreasing in time. The safety of \mathcal{C} implies $h(\mathbf{x}(t))$ is lower bounded by 0, and thus we can conclude that $\lim_{t \rightarrow \infty} h(\mathbf{x}(t))$ exists by the Monotone Convergence Theorem [249, Theorem 2.4.2] Assume that this limit is some value $0 < a < 1$. For any $\delta \in \mathbb{R}_{>0}$, there exists $T > 0$ such that for $t > T$, $h(\mathbf{x}(t)) < a + \delta$. Since there are an infinite number of events, we deduce there exists $t_k > T$ such that $h(\mathbf{x}(t_k)) < a + \delta$. As $h(\mathbf{x}(t))$ is monotonically decreasing, it also follows $h(\mathbf{x}(t)) \geq a$ for all t . This implies:

$$\dot{h}(\mathbf{x}(t)) \leq -a \exp(a\Delta t_k)(1 - (a + \delta)) \leq -a + a^2 + a\delta < 0. \quad (6.180)$$

for $t \geq t_k$ where δ can be chosen small enough to enforce the strict inequality with 0 as $a^2 < a$. Thus, between two events we have:

$$h(\mathbf{x}(t_{k+1})) \leq h(\mathbf{x}(t_k)) + \tau(-a + a^2 + a\delta), \quad (6.181)$$

where τ is the assumed MIET. Choosing $\delta < \tau(a - a^2)/(1 + \tau a)$ in conjunction with $h(\mathbf{x}(t_k)) < a + \delta$ implies $h(\mathbf{x}(t_{k+1})) < a$, contradicting the assumption that $a \neq 0$ (and maintaining the assumption on the existence of τ).

To complete the proof, note $\mathbf{e}(t_k) = \mathbf{0}_n$ and take the second-order (one-sided) Taylor expansion of $\|\mathbf{e}(t)\|^2$ at $t = t_k$:

$$\|\mathbf{e}(t)\|^2 = (\dot{\mathbf{e}}(t_k)^\top \dot{\mathbf{e}}(t_k)) (t - t_k)^2 + \mathcal{O}((t - t_k)^3) \quad (6.182)$$

$$= (1 + k(\mathbf{x}(t_k))) \|\mathbf{x}(t_k)\|^2 (t - t_k)^2 + \mathcal{O}((t - t_k)^3) \quad (6.183)$$

$$\geq \|\mathbf{x}(t_k)\|^2 (t - t_k)^2 - c_3 (t - t_k)^3, \quad (6.184)$$

with $c_3 \in \mathbb{R}_{>0}$. The first term in the inequality follows from $k(\mathbf{x}(t_k)) \geq 0$ for $\mathbf{x}(t_k) \in \mathcal{C}$. To understand the second term, note that for any $t \in [t_k, t_{k+1})$, Lagrange's Remainder Formula [249, Theorem 6.6.3] implies there exists a $t' \in (t_k, t)$ such that:

$$\|\mathbf{e}(t)\|^2 = (\dot{\mathbf{e}}(t_k)^\top \dot{\mathbf{e}}(t_k)) (t - t_k)^2 + \frac{1}{6} \frac{d^3}{dt^3} \|\mathbf{e}(t')\|^2 (t - t_k)^3. \quad (6.185)$$

As $\frac{d^3}{dt^3} \|\mathbf{e}(t)\|^2$ is a continuous function of $\mathbf{x}(t)$ and $\mathbf{x}(t_k)$ (can be shown by expanding it), both of which remain in the compact set \mathcal{C} , we must have that $\frac{d^3}{dt^3} \|\mathbf{e}(t)\|^2$ is bounded by a constant, denoted by $6c_3 \in \mathbb{R}_{\geq 0}$, for all $t \in \mathbb{R}_{\geq 0}$, including $t = t'$. This yields the desired bound.

At trigger time t_k , let $h(\mathbf{x}(t_k)) = \epsilon_k$ for $\epsilon_k \in \mathbb{R}_{>0}$. Note that for arbitrarily small $\epsilon \in \mathbb{R}_{>0}$, there exists a trigger time t_k such that $\epsilon_k < \epsilon$ due to the

existence of infinite triggers and convergence of $h(\mathbf{x}(t))$ to 0. This implies $\|\mathbf{x}(t_k)\|^2 = 1 - \epsilon_k$. Let $n \in \mathbb{Z}_{\geq 0}$ be such that $\frac{1}{c_3} < n\tau$ and define $t_k^* = t_k + \frac{1}{n} \left(\frac{1-\epsilon_k}{c_3} \right)$, noting $t_k^* < t_k + \tau$. It follows from the Taylor expansion that:

$$\|\mathbf{e}(t_k^*)\|^2 \geq \frac{(1 - \epsilon_k)^3}{c_3^2} \frac{n - 1}{n^3}. \quad (6.186)$$

As ϵ_k can be chosen arbitrarily small, I choose it such that:

$$(1 - \epsilon_k)^3 \geq \frac{c^2 n^3}{4r^6(n - 1)} \epsilon_k^2,$$

which indicates that:

$$2r^3 \|\mathbf{e}(t_k^*)\| \geq c|h(\mathbf{x}(t_k))| \geq c|h(\mathbf{x}(t_k^*))|,$$

as $h(\mathbf{x}(t))$ is monotonically decreasing. As $t_k^* < t_k + \tau$, this contradicts that τ is the MIET. Figure 6.14 shows the number of events as a function of time and distance from the barrier. The blue curves in Figure 6.15 correspond to the interevent times. \square

In the proof of Lemma 11, we see the fact that the dynamics of the event-triggered closed-loop system (6.167) are not required to vanish on the boundary of the set \mathcal{C} . This leads to the derivative of the measurement error, $\frac{d}{dt} \|\mathbf{e}(t_k)\|$, being uniformly lower bounded at event times t_k , which, together with the convergence of h to 0, as seen in Figure 6.14, leads to arbitrarily small interevent times. In particular, the dynamics evolve tangentially to the boundary of the set \mathcal{C} , leading to growing measurement error while moving arbitrarily close to the 0-level set of h . As the original controller may have additional objectives beyond safety (such as stabilization), it is desirable that the event-triggered implementation not completely eliminate tangential motion near the boundary that may be necessary to achieve the other objectives. To accommodate this, I will introduce a trigger-law that limits dynamic evolution tangential to the boundary of the safe set.

Event-Triggered Safety & Simulation Results

I now propose an alternative trigger-law that ensures an MIET exists. To resolve the issues in the preceding example, I introduce the following definition:

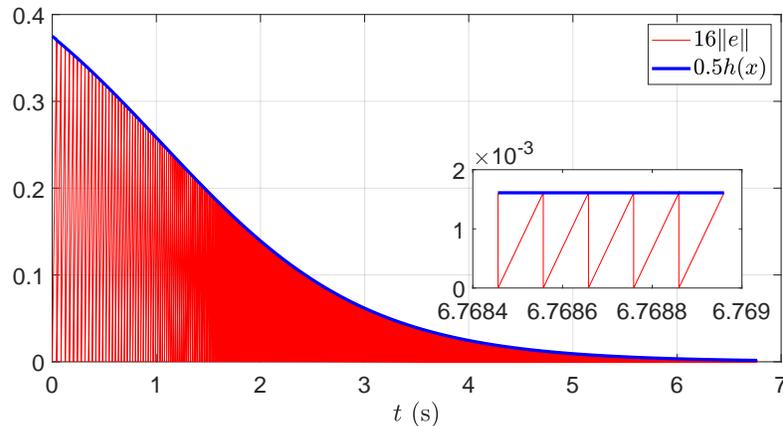


Figure 6.14. Simulation results for the system (6.167) using the trigger law (6.169). Even as the boundary of the safe set is approached ($h(\varphi_{\text{ET}}(t)) \rightarrow 0$), the growth rate of the error does not diminish, leading to arbitrarily small interevent times.

Definition 60 (*Strong ET-ISSf Barrier Property*). An ET-ISSf-BF h for the continuous-time closed-loop system (6.155) satisfies the *strong ET-ISSf barrier property* if there exists a $d \in \mathbb{R}_{>0}$ such that:

$$L_f h(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{k}(\mathbf{x} + \mathbf{e}) \geq -\alpha(h(\mathbf{x})) + d - \iota(\|\mathbf{e}\|), \quad (6.187)$$

for all $\mathbf{x}, \mathbf{e} \in \mathbb{R}^n$.

This property introduces a positive constant, d , into the ET-ISSf-BF condition (6.161). In the presence of zero measurement error, this enforces that the state dynamics must lie in the interior of the tangent cone [104] when on the boundary of the safe set \mathcal{C} . It also enforces that $\frac{d}{dt}|h(\mathbf{x}(t_k))|$ will be greater than a positive constant as we approach the boundary, similarly to $\frac{d}{dt}\|\mathbf{e}(t_k)\|$. I now show this property is sufficient to design a trigger-law that ensures safety with a MIET.

Theorem 47. Let h be an ET-ISSf-BF for the continuous-time closed-loop system (6.155) on a set $\mathcal{C} \subset \mathbb{R}^n$ defined as the 0-superlevel set of a function h that is continuously differentiable on \mathbb{R}^n , with corresponding functions $\alpha \in \mathcal{K}_\infty^e$ and $\iota \in \mathcal{K}_\infty$. Let $\beta \in \mathcal{K}_\infty^e$ and $c \in (0, 1]$. If the following assumptions hold:

1. h satisfies the strong ET-ISSf barrier property for a constant $d \in \mathbb{R}_{>0}$,
2. ι is Lipschitz continuous with Lipschitz constant L_ι ,
3. there exists $F \in \mathbb{R}_{>0}$, such that for all $\mathbf{x}, \mathbf{e} \in \mathbb{R}^n$:

$$\|\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x} + \mathbf{e})\| \leq F, \quad (6.188)$$

4. $\beta(r) \geq \alpha(r)$ for all $r \in \mathbb{R}$,

then the trigger-law:

$$t_{k+1} = \min \{t \geq t_k \mid \iota(\|\mathbf{e}(t)\|) = \beta(h(\mathbf{x}(t))) - \alpha(h(\mathbf{x}(t))) + cd\}, \quad (6.189)$$

deployed recursively enforces:

$$\dot{h}(\mathbf{x}(t), \mathbf{e}(t)) = L_{\mathbf{f}}h(\mathbf{x}(t)) + L_{\mathbf{g}}h(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)) \geq -\beta(h(\mathbf{x}(t))), \quad (6.190)$$

thus rendering the set \mathcal{C} safe. Furthermore, there exists a MIET given by:

$$t_{k+1} - t_k \geq \tau \triangleq \frac{cd}{L_{\iota}F}, \quad \forall k \in \mathbb{Z}_{\geq 0}. \quad (6.191)$$

Before proving the result, I make a few observations regarding its assumptions. Assumption 3 on the boundedness of the dynamics need not hold over the entire state space for safety, but can hold for $(\mathbf{x}, \mathbf{e}) \in \mathcal{C} \times \mathbb{R}^n$. Furthermore, if \mathcal{C} is compact, the trigger-law enforces the existence of a compact set $K \subset \mathbb{R}^n$ such that $\mathbf{e}(t) \in K$ for all $t \in \mathbb{R}_{\geq 0}$. Thus, the continuity of \mathbf{f} , \mathbf{g} , and \mathbf{k} on \mathbb{R}^n would imply the assumption is satisfied on $\mathcal{C} \times K$, which also would be sufficient for the result to hold. Assumption 4 ensures the right-hand side of the equality in the trigger will always be positive. The function β can be viewed as a tuning function which can raise interevent times (but not the MIET) at the expense of less “braking” within the safe set and convergence outside of the safe set. One choice is $\beta = \alpha$, in which case interevent times are lowered for more braking and faster convergence. Lastly, I note this trigger-law can be used in conjunction with (6.160) to jointly achieve event-triggered stabilization and safety if there exists an ET-ISS-LF for the continuous-time closed-loop system (6.155).

Proof. To see the set \mathcal{C} is rendered safe, observe that:

$$\dot{h}(\mathbf{x}(t), \mathbf{e}(t)) = L_{\mathbf{f}}h(\mathbf{x}(t)) + L_{\mathbf{g}}h(\mathbf{x}(t))\mathbf{k}(\mathbf{x}(t) + \mathbf{e}(t)), \quad (6.192)$$

$$\geq -\alpha(h(\mathbf{x}(t))) + d - \iota(\|\mathbf{e}(t)\|), \quad (6.193)$$

$$\geq -\alpha(h(\mathbf{x}(t))) + d - (\beta(h(\mathbf{x}(t))) - \alpha(h(\mathbf{x}(t))) + cd), \quad (6.194)$$

$$= -\beta(h(\mathbf{x}(t))) + (1 - c)d \quad (6.195)$$

$$\geq -\beta(h(\mathbf{x}(t))), \quad (6.196)$$

for all $t \in [t_k, t_{k+1})$ and $k \in \mathbb{Z}_{\geq 0}$, implying that h is a barrier function for the event-triggered closed-loop system (6.151) on \mathcal{C} , such that the event-triggered closed-loop system (6.151) is safe with respect to \mathcal{C} . To see the interevent times are lower bounded, observe that:

$$\|\mathbf{e}(t)\| = \left\| \mathbf{e}(t_k) + \int_{t_k}^t (-\mathbf{f}(\mathbf{x}(\tau)) - \mathbf{g}(\mathbf{x}(\tau))\mathbf{k}(\mathbf{x}(\tau) + \mathbf{e}(\tau)))d\tau \right\|, \quad (6.197)$$

$$= \left\| \int_{t_k}^t (-\mathbf{f}(\mathbf{x}(\tau)) - \mathbf{g}(\mathbf{x}(\tau))\mathbf{k}(\mathbf{x}(\tau) + \mathbf{e}(\tau)))d\tau \right\|, \quad (6.198)$$

$$\leq \int_{t_k}^t F d\tau. \quad (6.199)$$

This inequality together with the trigger law (6.189) yields:

$$t_{k+1} \geq \min \{t \geq t_k \mid L_\iota \|\mathbf{e}(t)\| = cd\}, \quad (6.200)$$

$$\geq \min \{t \geq t_k \mid L_\iota F(t - t_k) = cd\} \quad (6.201)$$

$$= \frac{cd}{L_\iota F} + t_k, \quad (6.202)$$

ensuring the desired result. \square

In the case that an ET-ISSf-BF h does not satisfy the strong barrier property, an auxiliary ET-ISSf-BF, h_b , satisfying the strong ISSf barrier property can be synthesized via h at the expense of keeping a larger set safe:

Theorem 48. *Let h be an ET-ISSf-BF for the continuous-time closed-loop system (6.155) on a set $\mathcal{C} \subset \mathbb{R}^n$ defined as the 0-superlevel set of a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ that is continuously differentiable on \mathbb{R}^n , with corresponding functions $\alpha \in \mathcal{K}_\infty^e$ and $\iota \in \mathcal{K}_\infty$. Then the function $h_b : \mathbb{R}^n \rightarrow \mathbb{R}$ defined as $h_b(\mathbf{x}) = h(\mathbf{x}) + b$, with $b \in \mathbb{R}_{>0}$, is an ET-ISSf-BF satisfying the strong ET-ISSf barrier property on the set $\mathcal{C}_b \subset \mathbb{R}^n$ defined as:*

$$\mathcal{C}_b \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid h_b(\mathbf{x}) \geq 0\}. \quad (6.203)$$

Proof. Observe that:

$$L_{\mathbf{f}}h_b(\mathbf{x}) + L_{\mathbf{g}}h_b(\mathbf{x})\mathbf{k}(\mathbf{x} + \mathbf{e}) = L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x} + \mathbf{e}), \quad (6.204)$$

$$\geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{e}\|), \quad (6.205)$$

$$\geq -\alpha(h_b(\mathbf{x}) - b) + \alpha(-b) - \alpha(-b) - \iota(\|\mathbf{e}\|), \quad (6.206)$$

$$= -\alpha_b(h_b(\mathbf{x})) + d_b - \iota(\|\mathbf{e}\|), \quad (6.207)$$

for all $\mathbf{x}, \mathbf{e} \in \mathbb{R}^n$, where $\alpha_b \in \mathcal{K}_\infty^e$ is defined as $\alpha_b(r) = \alpha(r - b) - \alpha(-b)$ and $d_b = -\alpha(-b) > 0$. \square

Thus an ET-ISSf-BF can be used with the trigger law (6.189) by enlarging the safe set an arbitrarily small amount, captured by the following result:

Corollary 2. *If h is an ET-ISSf-BF for (6.155) on the set \mathcal{C} satisfying Assumptions (2-4) of Theorem 47, then h_b is an ET-ISSf-BF for (6.155) on the set \mathcal{C}_b satisfying Assumptions (1-4) of Theorem 47 such that the corresponding trigger-law renders the event-triggered closed-loop system (6.151) safe with respect to \mathcal{C}_b , and yields an MIET.*

This is effectively an instance of Input-to-State Safety, in which case the original set \mathcal{C} defined via h becomes an ISSf safe set. We note that the larger the set is made (via a larger choice of b), the larger the MIET will be. This effectively highlights a trade-off that arises in the context of safety but not in stabilization: allowing motion near the boundary of a safe set requires additional relaxations to achieve the additional desirable property of an MIET. To verify the ability of this trigger to keep the system safe and have a MIET, I simulate the event-triggered closed-loop system (6.167) using both the naïve trigger law (6.169) and the corrected trigger-law (6.189). The results of these simulations can be seen in Figure 6.15. Observe that although both systems are kept safe, the trigger-law not using the strong ISSf barrier property has interevent times that approach 0.

Conclusion

In this section I have presented work on safety-critical event-triggered control. I begin by reviewing the formulation of event-triggered stabilization through ISS-LFs as originally presented in [96]. In attempting to translate these results to the safety-critical setting through the use of ISSf-BFs, I reveal a challenge that occurs due to motion of the system along the boundary of a set \mathcal{C} that is to be kept forward invariant. I rigorously show that this challenge can prevent the existence of an MIET through a counterexample. I conclude by proposing the notion of the strong ET-ISSf barrier property, which limits tangential motion on the boundary of the set \mathcal{C} and allows for a trigger-law that possesses an MIET, as I show in simulation.

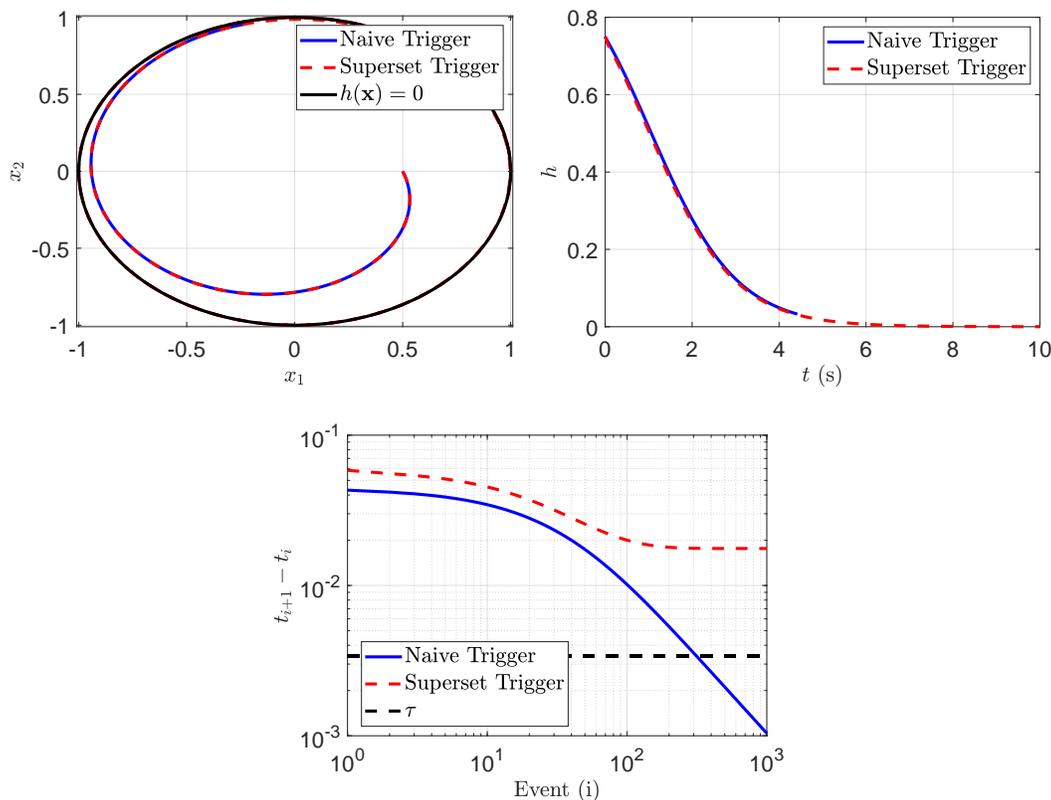


Figure 6.15. Simulation results demonstrating safety achieved with an event-triggered controller. (Top, Left) State trajectories for both triggers remain within the safe set for the length of the simulation. (Top, Right) The value of the ISSf-BF h remains above zero for the length of the simulation, corresponding to the system remaining safe. (Right) The interevent times of the two trigger laws. The interevent times of the trigger law (6.169) decreases towards 0 as predicted by Lemma 11 while the trigger law (6.189) satisfies the theoretical bound.

6.6 Conclusion

In this chapter I have presented a collection of results on designing safety-critical controllers that are robust to input sampling. In Section 6.2 I look at a few simulation examples that demonstrate the fragility of the (CLF-QP) controller to sampling, even at high sample rates. Following this, I provide mathematical descriptions of the sampled-data and event-triggered sampling paradigms. Next, in Section 6.3 I present my first set of contributions in the form of a framework for using feedback linearization and CLF-based controllers specified via convex optimization problems for sampled-data systems. I augment the framework of sampled-data control through approximate discrete-time models presented in [81] to support local stability results and the coordinate transformations that accompany feedback linearization. Given this setup,

I show that a continuous-time feedback linearizable system with locally exponentially stable zero dynamics can be practically stabilized at sufficiently high sample rates using a sample-and-hold implementation of the continuous-time feedback linearizing controller without modification. As was reviewed in Section 2.4, I then use feedback linearization to produce a discrete-time CLF which I can incorporate into a CLF-based controller specified as a convex quadratically constrained quadratic program (QCQP). This controller is shown to greatly outperform a sample-and-hold implementation of the continuous-time (CLF-QP) controller in the simulation environments in Section 6.2. In Section 6.4, I present my second set of contributions in the form of a framework for safety-critical sampled-data control through Control Barrier Functions and approximate discrete-time models. I consider a larger class of approximate discrete-time models that support working with CBFs that have a relative degree greater than one, and show how these models satisfy certain consistency properties. Following this, I present definitions of practical safety, sampled-data barrier functions, and sampled-data Control Barrier Functions, and establish how a family of sampled-data barrier functions for a family of approximate discrete-time models implies the exact discrete-time system is practically safe. Following this, I consider a large class of nonlinear systems for which I can show that convexity properties of a sampled-data CBF can be preserved through the use of a Runge-Kutta approximation of the appropriate order, leading to controller that can be specified as convex optimization problems. I demonstrate these controllers in simulation, showing that safety violations decrease with increased sample rate as described by practical safety. Lastly, in Section 6.5 I present my third set of contributions in the form of a framework for safety-critical event-triggered control through ISSf-BFs. I first review the framework for event-triggered stabilization through ISS-LFs presented in [96]. I then consider a naïve transfer of this framework to the safety-critical setting, and show that the importance requirement of a minimum interevent time fails to hold through a rigorous counterexample. Building off the challenges seen in this counterexample, I present a modification that leads to event-triggered controllers that ensure safety and a minimum interevent time, which I show in simulation.

There are several directions for future work building off the results presented in this chapter. The controllers that were produced through CLFs and CBFs with convex optimization in Sections 6.3 and 6.4, respectively, are more complex

forms of convex optimization problems than QPs. In particular, the CLF-based controllers are QCQPs, and the CBF-based controllers depend on the safe set specification, returning QPs for half-space CBFs, QCQPs for ellipsoidal CBFs, SOCPs for norm CBFs, and more general convex programs for CBFs that utilize other convex functions. While the sampled-data control paradigm doesn't require controllers be locally Lipschitz continuous, or even continuous (because they are not evaluated continuously in time in the differential equation, they are held constant), these types of regularity properties are often desirable in practically implementing these controllers. Thus, it would be worthwhile to characterize the regularity properties of these controllers similar to the way that regularity properties of the (CLF-QP) and (CBF-QP) controllers have been characterized in [51]. Doing so may not be able to rely on the sort of tools used in [51], and may require a more analysis-driven approach using tools from variational analysis [223], [224].

Another direction would consider the approximate discrete-time models used. My work uses the relatively simple Euler and Runge-Kutta approximations, in large part due to their wide existing use, but also because they can be shown to preserve the feedback linearization and convexity properties ideal for control synthesis. I first note that my theory did not require any thing about the coefficients $a_{i,j}$ in the Runge-Kutta approximation. Many methods (of the same-order) for numerical approximation fall under the umbrella of Runge-Kutta approximations with different values for these coefficients, and can be shown to demonstrate different desirable properties for specific problem structures. Thus, a comparative study, both theoretically and empirically, evaluating controller performance using different values of the coefficients associated with common numerical approximation schemes may be a worthwhile direction of exploration. Second, other approximation schemes outside of Runge-Kutta approximations could be explored, with the goal of obtaining higher approximation accuracy or conforming more closely to underlying structure in the dynamics. An important focus of this study would be on what sorts of approximations preserve convexity in a way that allows for efficient downstream control synthesis, as was considered in my work.

Lastly, I think there are interesting avenues to consider event-triggered paradigms. As noted in [95], the widespread adoption and use of event-triggered control schemes has been lacking. It is my belief that the adoption of such methods in

robotics has not happened because the cost for switching inputs to a robot at a torque level are fairly negligible, and it is generally best to run a controller as fast as possible on a given set of hardware. Instead, I think that the use of event-triggered ideas in robotics need to focus on tasks that are costly to perform, and for which it is desirable not to have to perform many times. One such example could be event-triggered learning, where new data is incorporated into learning models only after it has been deemed to contain enough information to warrant the expensive learning process. In fact, recent work has even begun to explore this question [330]. I believe it would be worthwhile to develop an understanding of the stability and safety of these types of dynamic processes through tools from event-triggered control as explored in this thesis.

CONCLUSION

Motivated by the tragic failures seen with the ongoing proliferation of autonomous systems in modern society, this thesis has focused on making progress towards a complete theory of robust safety-critical. Built upon the foundational tools of Control Lyapunov Functions (CLFs) and Control Barrier Functions (CBFs), my contributions in this thesis constitute a collection of methods for resolving the challenges of model error, disturbances, measurement error, and input sampling seen in real-world applications. In each chapter of this thesis, I show the consequences for safety in failing to address these challenges in control design through simulation or experimental demonstration. Equipped with an understanding of these failures, I develop methods that provide theoretically rigorous guarantees of robust safety. A hallmark of my contributions is the focus on developing controllers that can be specified via convex optimization problems, enabling their computationally efficient implementation on real-world platforms, as demonstrated in the collection of experimental results presented in this thesis.

In summary, the contributions of this thesis are as follows:

- **Learning Frameworks:** In Sections 3.3-3.8 I present a series of three episodic learning frameworks for resolving model error as it directly impacts stability or safety as encoded by CLFs and CBFs, respectively. In particular, these frameworks address the fact that only model error that affects the time derivative of a CLF or CBF needs to be characterized to achieve a control objective. The frameworks in Sections 3.3 and 3.5 take an episodic learning approach built on supervised learning with a focus on learning models that are affine in the control input, permitting their direct integration into existing convex optimization-based controllers. Accompanying these frameworks in Sections 3.4 and 3.6 are a rigorous characterization of the impacts of residual learning error on stability and safety through the notions of Projection-to-State Stability (PSS) and Projection-to-State Safety (PSSf). In Section 3.7 I explore some of the challenges with these frameworks on high-dimensional system, and utilize

a learning model that instead learns a projected disturbance associated with a particular controller, allowing iterative improvement in achieving safety objectives. Lastly, in Section 3.8, I draw upon the lessons collected in developing the preceding frameworks to present a data-driven control method through robust convex optimization. Importantly, this method capitalizes on the fact that it is not necessary to completely characterize model error to choose robust control inputs.

- **Adaptive Safety-Critical Control:** In Section 3.9 I present the first instance of an adaptive Control Barrier Function for safety-critical control in the presence of parametric model error. Building upon the tool of adaptive Control Lyapunov Functions in [47], I develop an understanding of the challenges in achieving forward invariance as opposed to stability with adaptive control techniques, and present a notion of aCBFs that resolve these challenges.
- **Tunable Input-to-State Safety:** In Section 4.2 I highlight some of the challenges in balancing performance and robustness to disturbances with controllers designed with Input-to-State Safe Control Barrier Functions as presented in [58]. In Sections 4.3 and 4.4 I greatly improve upon this balance of performance and robustness through the the notion of Tunable Input-to-State Safety, which allows prioritizing robustness to disturbances near the boundary of a set that is to be kept forward invariant, and relaxing robustness requirements deep within the set. This methodology is experimentally demonstrated on an autonomous semi-trailer truck, showcasing it on a challenging real-world problem.
- **Stochastic Safety-Critical Control:** In Section 4.5 I present a framework for safety-critical control of discrete-time systems with stochastic disturbances. This work focuses on using the inherent robustness properties of CBFs to produce ISSf-like finite-time safety guarantees. Additionally, it studies how the the interaction between convexity properties of CBFs and the distributions of stochastic disturbances must be accounted for in producing theoretically rigorous safety-critical controllers specified as convex optimization problems.
- **Measurement-Robust Control Barrier Functions:** In Sections 5.2 I explore how model error impacts safety-critical control, and develop

the notion of Measurement-Robust Control Barrier Functions as a tool for measurement-robust control synthesis through convex optimization. In Section 5.3 MR-CBFs are integrated with backup-set CBF methods to produce a measurement-robust controller specified as a second-order cone program (SOCP) that is demonstrated experimentally on a Segway platform. Importantly, this result indicates that the controllers specified via convex optimization problems that are more complex than QPs (such as QCQPs and SOCPs) in this thesis are efficient enough to be deployed in real-time on real-world robotic systems.

- **Preference Based Learning:** In Section 5.4 I explore how preference-based learning can be used to tune the robustness parameters of a CBF-based controller, which is the first instance of using preference-based learning with CBFs. While my particular contributions in this work focus on the control theoretic aspects of safety guarantees, it is an important result as it presents a methodology that may be deployed with many of the methods presented in this thesis to produce controllers that are both robust and performant.
- **Sampled-Data Control:** In Section 6.2 I highlight the fragility of the continuous-time (CLF-QP) controller to a sample-and-hold implementation for a number of systems in simulation. I resolve this issue in Section 6.3 by presenting a framework for stabilizing sampled-data control through approximate discrete-time models. Building upon the framework in [81], I show how continuous-time feedback linearizable systems can be practically stabilized through a sample-and-hold implementation of a feedback linearizing controller. Furthermore, I extend the pipeline from feedback linearization to CLFs to this setting, producing discrete-time CLFs which can be incorporated into convex-optimization based controllers specified as quadratically constrained quadratic programs (QCQPs). I demonstrate the significant improvement in performance of these controllers over their continuous-time counterparts with a sampled-data implementation in simulation. In Section 6.4 I develop a framework for safety-critical sampled-data control through CBFs. I propose notions of practical safety and sampled-data CBFs, and rigorously connect designs using approximate discrete-time models to the practical safety of the exact discrete-time sampled-data dynamics of a system. Lastly, I

show how convexity properties of a SD-CBF can be utilized with an appropriately chosen approximate discrete-time model to produce convex optimization-based controllers.

- **Event-Triggered Safety-Critical Control:** In Section 6.5 I provide a framework for event-triggered safety-critical control. Drawing inspiration from the work in [96] that studies stabilizing event-triggered control through ISS-LFs, my approach studies event-triggered safety through the perspective of ISSf-BFs. I show that that directly transferring the approach in [96] is sufficient to ensure safety, but will not guarantee a minimum interevent time. To resolve this, I propose the strong ISSf-BF property and show how it can be used to guarantee both safety and a minimum interevent time.

Together, these contributions constitute a significant advance in the theory of robust safety-critical control, and unified by the use of CLFs and CBFs, provide a cohesive approach for addressing the vital challenges faced by real-world systems.

7.1 Infeasible Optimization Problems

Throughout this thesis there has been a focus on developing controllers specified via convex optimization problems. In implementing these controllers, a relevant question that practitioners may face is what to do if the optimization problem defining these controllers returns that there is no feasible solution (or does not return a result at all). This is certainly an important question for developing the ability to trace faults in complex systems. I wish to note that in some sense, the ability to entirely eliminate the possibility of infeasibility while maintaining rigorous theoretical stability and safety guarantees is tied directly to whether the assumptions a control designer makes actually hold. If exceptional events can happen (as they are want to do in the real world) that violate these assumptions, there is always some potential for infeasibility or failure to achieve stability and safety. The following discussion seeks to highlight common reasons for infeasibility, and discuss practical methods that can be used to improve upon feasibility. A practitioner will hopefully find this helpful in identifying why their controller went infeasible, and provide a spectrum of increasingly involved options to consider in trying to improve upon their design.

Numerical Issues

First, the cost functions in these controllers tend to be fairly well-behaved, and issues with infeasibility (either certifiable infeasibility, or due to the solver not returning any result) are almost always due to the constraints. One annoying, but fairly mild, source of infeasibility is numerical conditioning, which typically leads to the solver not returning. If we recall the CLF condition:

$$L_f V(\mathbf{x}) + L_g V(\mathbf{x})\mathbf{u} \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2, \quad (7.1)$$

with $a = 2$, we note that the term $L_g V(\mathbf{x})$ multiplies the decision variable \mathbf{u} in the optimization problem. At the equilibrium point \mathbf{x}_e , we will often have that $L_f V(\mathbf{x})$ and $L_g V(\mathbf{x})$ will converge to 0 and $\mathbf{0}_m$, respectively, linearly with the value of $\|\mathbf{x} - \mathbf{x}_e\|$, while the right hand-side converges to 0 quadratically in $\|\mathbf{x} - \mathbf{x}_e\|$. In this case, the vanishing of the term multiplying the decision variable can be rectified by rescaling the problem by dividing by $\|\mathbf{x} - \mathbf{x}_e\|$, noting that the quadratic term will remain small. In fact, most off-the-shelf optimization tools will perform this sort of rescaling for a user. But such rescaling away from \mathbf{x}_e is not so straightforward. In particular, if we have a state $\mathbf{x}^* \neq \mathbf{x}_e$ such that $L_g V(\mathbf{x}^*) = \mathbf{0}_m$, then for a state \mathbf{x} close to \mathbf{x}^* , dividing each term of the constraint by $\|L_g V(\mathbf{x})\|$ will amplify the terms $L_f V(\mathbf{x})$ and $-k_3 \|\mathbf{x} - \mathbf{x}_e\|^2$ greatly, thereby still leading to a poorly scaled problem. Many off-the-shelf optimization tools will struggle with this problem, but it reflects important geometric properties of the interaction between the underlying system dynamics and CLF. Instead, observe that if we have a CLF, then in a neighborhood of \mathbf{x}^* , we have that:

$$L_f V(\mathbf{x}) \leq -k_3 \|\mathbf{x} - \mathbf{x}_e\|^2 + \epsilon, \quad (7.2)$$

for some small value of $\epsilon \approx 0$. This states that by taking no input (or using \mathbf{k}_{nom}), we are *almost* satisfying the inequality we need for local exponential stability. Thus, rather than solve the optimization problem, it may be prudent to set a tolerance on the value of $\|L_g V(\mathbf{x})\|$ under which no input (or \mathbf{k}_{nom}) is used. The price of this is a small degradation in the convergence rate of stability, or ISS-like behavior. The same sort of issue arises with CBFs when $L_g h(\mathbf{x}^*) = \mathbf{0}_m$, and can be similarly addressed by using \mathbf{k}_{nom} directly at the expense of ISSf behavior. This type of case-based solution is an effective tool for many of the CLF and CBF-based methods throughout this thesis, and generally leads to insignificant degradation in stability and safety guarantees.

Infeasible Constraints

Another source of infeasibility (that often produces certifiable infeasibility) is when a function V or h is not actually the purported type of CLF or CBF it is claimed to be when it is used as a constraint in the optimization problem. Why does this happen? One example is if the CLF or CBF is simply specified for the system without using a constructive technique like feedback linearization as shown in Chapter 2. In this case, the CLF or CBF is being used as a heuristic design primitive, but not necessarily as a rigorous control theoretic tool. As shown by many of the simulation and experimental results in this thesis, this is often an okay approach (though I will discuss alternatives if it is not). Another example is if the CLF or CBF has numerical models (like learning models in Chapter 3) integrated into the constraint. These numerical models often do not display exactly the right analytical behavior at the “pain points” (such as when $L_{\mathbf{g}}V(\mathbf{x})$ or $L_{\mathbf{g}}h(\mathbf{x})$ approach $\mathbf{0}_m$), and can lead to infeasibility in the constraint. Note that this sort of problem can occur even if the CLF or CBF is in fact a CLF or CBF for the true dynamics of the system, as discussed in Section 3.3. An additional example is when a CLF or CBF for a system (one that actually is a CLF or CBF) is modified with robustness terms. Examples include using ISS-CLFs or ISS-CBFs with unmatched disturbances (this is why matched disturbances are a nicer problem to work with), adding model error sets into the CLF or CBF constraint as in Section 3.8 (where I explicitly address feasibility), and MR-CBFs in Section 5.2. The robustness modifications in these formulations may lead to infeasibility, especially if large robustness parameters are used. This is not unintuitive, as requiring robustness to an increasingly large quantity of model inaccuracies will at some point become impossible in most control formulations. Lastly, and perhaps most importantly, even if a CLF or CBF meets the necessary infimum and supremum conditions used throughout this thesis, nearly all real-world systems face limitations to their inputs, such as input bounds. It is generally not possible to take any input in \mathbb{R}^m . In this case, if these input constraints are enforced in the optimization with the CLF or CBF constraint, infeasibility may occur. If instead input saturation is applied after solving the optimization problem, it will not only be unclear if theoretical stability or safety properties will be preserved, but it will also be unclear if a potential feasible solution satisfying input bounds did in fact exist.

Relaxations

What are our options in the face of these potential causes of infeasibility? There are three main directions, all of which I believe are currently areas of research interest. The first is simply to relax the optimization problems. Such a relaxation was employed in the (CLF-CBF-QP) controller, and used throughout many of the experimental results in this thesis. This approach can guarantee the “feasibility” of the optimization problem in the face of not having a valid CLF or CBF, poorly behaved numerical models, excessive robustness, and input constraints. Of course, it is not a silver bullet, and will compromise the theoretical stability and safety guarantees of the controller. If the relaxation is kept small (one is only a small distance away from being feasible) over a region of the interest in the state space, one could employ the ideas of PSS and PSSf explored in Sections 3.4 and 3.6, respectively, to state what the degradation of the guarantees will be. But there are perhaps more interesting and principled approaches for relaxation, which suggest the potential for further research into the topic. Examples include the formulation in [51, QP Problem $-\gamma m$ Version], which uses a particular structure of the relaxation term to ensure that local exponential stability guarantees are preserved for equilibrium points that remain inside the interior of a set \mathcal{C} that is to be kept forward invariant. An other idea is to add additional decision variables to the controller that allow the strength of stability or safety requirements to be modulated online to achieve feasibility. Examples would be making the parameter k_3 in a CLF constraint a decision variable that is constrained to remain above some small positive value, or parameterizing the function $\alpha \in \mathcal{K}^e$ in a CBF constraint to be a linear function with a coefficient that is a decision variable constrained above some small positive value, as was done in [256, Section 2.C]. This type of scheme relaxes the convergence rate of local exponential stability, or allows a system to more rapidly approach the boundary of a set that is to be kept forward invariant, but can improve the feasibility of the controller while still giving some form of stability and safety guarantees.

Constructing CLFs and CBFs

The second direction is the offline construction of CLFs and CBFs that satisfy the necessary requirements arising from system dynamics, safety specifications, robustness considerations, and input constraints. In this way, one can reduce the number of assumptions made on the satisfaction of these requirements by

explicitly addressing them. Examples would include tools such as sums-of-squares (SOS) programming which has been used for computing CLFs [331] and CBFs [332], [333], solving the Hamilton-Jacobi-Bellman (HJB) equations to produce control invariant sets that yield CLFs [334] and CBFs [229], [335], or constructing CLFs and CBFs through data-driven methods [176], [177]. A significant benefit of these approaches is the potential to utilize large computational capabilities to produce solutions meeting all of the above requirements simultaneously in a way that can be verified before deployment on a safety-critical system. Of course, if these requirements are insufficient for the actual control task at hand, infeasibility may still occur, and these offline computations may need to be performed again with adjusted requirements. One limit of this approach is the well-known “curse of dimensionality”, with many of these methods being limited to single-digit state and input dimensions and offline computations. For a high-dimensional system, a prudent course of action may be to reduce the system to a collection of smaller, interconnected subsystems, and deploy the computational techniques with additional requirements to compensate for any issues that arise due to using an interconnected system modeling perspective.

Multi-Rate Control

Drawing some degree of inspiration from this previous comment, the third direction is to offload some of the requirements that a CLF or CBF-based controller needs to satisfy to a higher-level in the autonomy stack. For instance, one may require that the equilibrium state or trajectory that a CLF controller is stabilizing to be designed (online) to ensure that the CLF controller can accurately track the trajectory while meeting robustness requirements and input constraints. This was the focus on my work in [336], which modified the constraints of a low-frequency model predictive controller (MPC) producing desired trajectories to ensure the resulting trajectory could be tracked by a (CLF-QP) controller. Other recent work has also begun to explore this sort of multi-rate control paradigm, where high-level and mid-level tasks effectively reduce the requirements that need to be satisfied by low-level CLF and CBF-based controllers [337]–[339]. A particular interesting observation in this direction occurred in my work in [22]. By fusing CBF constraints into both a low-level controller and a mid-level model predictive controller producing trajectories, it was possible to achieve consistently safer behavior than merely

using CBFs only at the low-level, using CBFs only at the mid-level, and using hard state constraints at the mid-level and CBFs at the low-level. This suggests that truly robust safety is not a problem that can, nor should, be achieved at one level in the autonomy stack. Instead, feasibility at all levels of the autonomy stack is best achieved by allocating some portion of safety requirements to each level in a principled manner. I believe that this will be a critical direction of future research over the next 20 years as efforts continue to transfer theoretically rigorous control systems into complex real world applications.

7.2 Future Work

At the conclusion of each chapter I have provided directions for future work related to achieving robustness to each respective challenge. One future direction that stems from the results of this thesis is the construction of a safety-critical control framework that is *simultaneously* robust to each of the presented challenges. In reality, real-world systems will face model error, disturbances, measurement error, and input sampling all at once. As each chapter of this thesis has shown, neglecting any one of these challenges can lead to significant violations of safety requirements. I believe that capturing the results of this thesis into a singular framework will require addressing two challenges:

- A theme of my approaches is taking the (CLF-QP) and (CBF-QP) controllers and making them robust by producing controllers specified by more complex convex optimization problems than QPs, and in particular QCQPs and SOCPs. This can be observed in the data-driven control framework in Section 3.8, the measurement-robust control in 5.2, and the sampled-data controllers in Sections 6.3 and 6.4. The question that must be addressed is if it is possible to integrate multiple forms of robustness without compromising the convexity of the control synthesis process. Throughout this thesis structural assumptions such as Lipschitz continuity or approximate models are used in a way that permits convex synthesis. Answering this question will rely on finding the *right* set of structural assumptions and approximations that ensure multiple forms of robustness while maintaining convexity. It is my belief that the most fruitful starting point for addressing this question will follow along the lines of the robust optimization used in the data-driven control

framework in Section 3.8 in conjunction with the discrete-time approximations used in Sections 6.3 and 6.4. In particular, if model error, disturbances, and measurement error can be cast as the model error sets used in Section 3.8, it may be possible to consider satisfying CLF and CBF requirements for all discrete-time approximations permitted by these error sets. It is likely this will yield more complex convex optimization problems, such as semidefinite programs. As the capabilities of computing resources become increasingly powerful, the potential to deploy these complex controllers on real-world systems may become possible.

- In Section 5.4 both disturbance robustness using ISSf and measurement robustness using MR-CBFs are integrated into the same controller. The purpose of that work was in addressing the fact that having both forms of robustness led to extremely conservative behavior that was not desirable. As more forms of robustness are brought into the same controller, it is likely that conservativeness will continue to increase, rendering such a controller pointless. Thus, how to tune these robust controllers to attain performance and robustness is an important question. Preference-based learning was shown to be a powerful tool for achieving this, but there may be opportunities to achieve this through other data-driven techniques. In particular, if data can be used to identify robustness that is not necessary (because the model error, disturbances, or measurement error being robustified against isn't actually present), then the characterization of error being used in this controller can be improved until a desirable level of performance is achieved.

In conclusion, I believe there many important questions remaining in building safety-critical controllers. I think that the most important focus of this work is continuing to integrate with advancing computational capabilities, and utilizing real experimental data to balance conservativeness and performance.

BIBLIOGRAPHY

- [1] M. Chui, K. George, J. Manyika, and M. Miremadi, “Human + machine: A new era of automation in manufacturing,” 2017. [Online]. Available: <https://www.mckinsey.com/capabilities/operations/our-insights/human-plus-machine-a-new-era-of-automation-in-manufacturing>.
- [2] J. Diechmann, E. Ebel, K. Heineke, R. Heuss, M. Kellner, and F. Steiner, “Autonomous driving’s future: Convenient and connected,” 2023. [Online]. Available: <https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/autonomous-drivings-future-convenient-and-connected#/>.
- [3] M. McFarland, “Tesla ‘full self-driving’ triggered an eight-car crash, a driver tells police,” 2022. [Online]. Available: <https://www.cnn.com/2022/12/21/business/tesla-fsd-8-car-crash/index.html>.
- [4] A. Adler, “Feds close probe into tusimple autonomous truck crash,” 2023. [Online]. Available: <https://www.freightwaves.com/news/feds-close-probe-into-tusimple-autonomous-truck-crash>.
- [5] A. J. Hawkins, “Tesla’s autopilot was engaged when model 3 crashed into truck, report states,” 2019. [Online]. Available: <https://www.theverge.com/2019/5/16/18627766/tesla-autopilot-fatal-crash-delray-florida-ntsb-model-3>.
- [6] L. Wamsley, “Backup driver of autonomous uber suv charged with negligent homicide in arizona,” 2022. [Online]. Available: <https://www.npr.org/2020/09/16/913530100/backup-driver-of-autonomous-uber-suv-charged-with-negligent-homicide-in-arizona>.
- [7] T. Krisher and S. Dazio, “Felony charges are 1st in a fatal crash involving autopilot,” 2022. [Online]. Available: <https://apnews.com/article/tesla-autopilot-fatal-crash-charges-91b4a0341e07244f3f03051b5c2462ae>.
- [8] Z. Artstein, “Stabilization with relaxed controls,” *Nonlinear Analysis: Theory, Methods & App.*, vol. 7, no. 11, pp. 1163–1173, 1983. DOI: 10.1016/0362-546X(83)90049-4.
- [9] E. D. Sontag, “A ‘universal’ construction of artstein’s theorem on non-linear stabilization,” *Sys. & Control Let.*, vol. 13, no. 2, pp. 117–123, 1989. DOI: 10.1016/0167-6911(89)90028-5.
- [10] P. Wieland and F. Allgöwer, “Constructive safety using control barrier functions,” *IFAC Proc. Vol.*, vol. 40, no. 12, pp. 462–467, 2007. DOI: 10.3182/20070822-3-ZA-2920.00076.

- [11] A. Ames, J. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs with application to adaptive cruise control,” in *Proc. IEEE 53rd Conf. on Decision and Control (CDC)*, Los Angeles, CA, USA, 2014, pp. 6271–6278. DOI: 10.1109/CDC.2014.7040372.
- [12] X. Xu, T. Waters, D. Pickem, *et al.*, “Realizing simultaneous lane keeping and adaptive speed regulation on accessible mobile robot testbeds,” in *Proc. 1st IEEE Conf. on Control Tech. and App. (CCTA)*, Kohala Coast, HI, USA, 2017, pp. 1769–1775. DOI: 10.1109/CCTA.2017.8062713.
- [13] T. Gurriet, M. Mote, A. D. Ames, and E. Feron, “An online approach to active set invariance,” in *Proc. IEEE 57th Conf. on Decision and Control (CDC)*, Miami Beach, FL, USA, 2018, pp. 3592–3599. DOI: 10.1109/CDC.2018.8619139.
- [14] R. Grandia, A. J. Taylor, A. Singletary, M. Hutter, and A. D. Ames, “Nonlinear model predictive control of robotic systems with control lyapunov functions,” in *Proc. Robotics: Science and Sys. XVI (RSS)*, Bend, OR, USA, 2020. DOI: 10.15607/RSS.2020.XVI.098.
- [15] L. Wang, A. D. Ames, and M. Egerstedt, “Safety barrier certificates for collisions-free multirobot systems,” *IEEE Trans. on Robotics*, vol. 33, no. 3, pp. 661–674, 2017. DOI: 10.1109/TR0.2017.2659727.
- [16] L. Wang, E. A. Theodorou, and M. Egerstedt, “Safe learning of quadrotor dynamics using barrier certificates,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Brisbane, QLD, Australia, 2018, pp. 2460–2465. DOI: 10.1109/ICRA.2018.8460471.
- [17] S. M. Khansari-Zadeh and A. Billard, “Learning control lyapunov function to ensure stability of dynamical system-based robot reaching motions,” *Robotics and Autonomous Sys.*, vol. 62, no. 6, 2014. DOI: 10.1016/j.robot.2014.03.001.
- [18] A. Singletary, W. Guffey, T. G. Molnar, R. Sinnet, and A. D. Ames, “Safety-critical manipulation for collision-free food preparation,” *IEEE Robotics and Automation Let.*, vol. 7, no. 4, pp. 10954–10961, 2022. DOI: 10.1109/LRA.2022.3192634.
- [19] W. S. Cortez, D. Oetomo, C. Manzie, and P. Choong, “Control barrier functions for mechanical systems: Theory and application to robotic grasping,” *IEEE Trans. on Control Sys. Tech.*, vol. 29, no. 2, pp. 530–545, 2019. DOI: 10.1109/TCST.2019.2952317.
- [20] S. Farzan, V. Azimi, A.-P. Hu, and J. Rogers, “Adaptive control of wire-borne underactuated brachiating robots using control lyapunov and barrier functions,” *IEEE Trans. on Control Sys. Tech.*, vol. 30, no. 6, pp. 2598–2614, 2022. DOI: 10.1109/TCST.2022.3160058.

- [21] A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, “Control barrier functions and input-to-state safety with application to automated vehicles,” *arXiv preprint arXiv:2206.03568*, 2022. [Online]. Available: <https://arxiv.org/abs/2206.03568>.
- [22] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, “Multi-layered safety for legged robots via control barrier functions and model predictive control,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Xi’an, China, 2021, pp. 8352–8358. DOI: 10.1109/ICRA48506.2021.9561510.
- [23] R. K. Cosner, M. Tucker, A. J. Taylor, *et al.*, “Safety-aware preference-based learning for safety-critical control,” in *Proc. 4th Learning for Dynamics and Control (L4DC)*, vol. 168, Stanford, CA, USA, 2022, pp. 1020–1033. [Online]. Available: <https://proceedings.mlr.press/v168/cosner22a.html>.
- [24] K. Galloway, K. Sreenath, A. D. Ames, and J. W. Grizzle, “Torque saturation in bipedal robotic walking through control lyapunov function-based quadratic programs,” *IEEE Access*, vol. 3, pp. 323–332, 2015. DOI: 10.1109/ACCESS.2015.2419630.
- [25] N. Csomay-Shanklin, R. K. Cosner, M. Dai, A. J. Taylor, and A. D. Ames, “Episodic learning for safe bipedal locomotion with control barrier functions and projection-to-state safety,” in *Proc. 3rd Learning for Dynamics and Control (L4DC)*, vol. 144, Zürich, Switzerland, 2021, pp. 1041–1053. [Online]. Available: <http://proceedings.mlr.press/v144/csomay-shanklin21a/csomay-shanklin21a.pdf>.
- [26] A. D. Ames and M. Powell, “Towards the unification of locomotion and manipulation through control lyapunov functions and quadratic programs,” in *Control of Cyber-Physical Sys.* Springer, 2013, pp. 219–240. DOI: 10.1007/978-3-319-01159-2_12.
- [27] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames, “Towards a framework for realizable safety critical control through active set invariance,” in *Proc. ACM/IEEE 9th Int. Conf. on Cyber-Physical Sys. (ICCPS)*, Porto, Portugal, 2018, pp. 98–106. DOI: 10.1109/ICCPS.2018.00018.
- [28] R. T. Rockafellar, “Lagrange multipliers and optimality,” *SIAM Review*, vol. 35, no. 2, pp. 183–238, 1993. DOI: 10.1137/103504.
- [29] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004, ISBN: 9780521833783.
- [30] G. C. Calafiore and L. El Ghaoui, *Optimization models*. Cambridge university press, 2014, ISBN: 978-1-107-05087-7.
- [31] Y. Nesterov and A. Nemirovskii, *Interior-point polynomial algorithms in convex programming*. SIAM, 1994. DOI: 10.1137/1.9781611970791.

- [32] Y. E. Nesterov and M. J. Todd, “Self-scaled barriers and interior-point methods for convex programming,” *INFORMS Mathematics of Operations research*, vol. 22, no. 1, pp. 1–42, 1997. DOI: 10.1287/moor.22.1.1.
- [33] S. Diamond and S. Boyd, “Cvxpy: A python-embedded modeling language for convex optimization,” *J. of Machine Learning Research*, vol. 17, no. 1, pp. 2909–2913, 2016. [Online]. Available: <https://www.jmlr.org/papers/volume17/15-408/15-408.pdf>.
- [34] J. Lofberg, “Yalmip: A toolbox for modeling and optimization in matlab,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Taipei, Taiwan, 2004, pp. 284–289. DOI: 10.1109/CACSD.2004.1393890.
- [35] A. Domahidi, E. Chu, and S. Boyd, “Ecos: An socp solver for embedded systems,” in *Proc. IEEE European Control Conf. (ECC)*, Zürich, Switzerland, 2013, pp. 3071–3076. DOI: 10.23919/ECC.2013.6669541.
- [36] M. ApS, *Introducing the mosek optimization suite 10.0.40*, 2019. [Online]. Available: <https://docs.mosek.com/latest/intro/index.html>.
- [37] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, 2023. [Online]. Available: <https://www.gurobi.com>.
- [38] J. Umlauft, L. Pöhler, and S. Hirche, “An uncertainty-based control lyapunov approach for control-affine systems modeled by gaussian process,” *IEEE Control Sys. Let.*, vol. 2, no. 3, pp. 483–488, 2018. DOI: 10.1109/LCSYS.2018.2841961.
- [39] H. Ravanbakhsh and S. Sankaranarayanan, “Learning lyapunov (potential) functions from counterexamples and demonstrations,” Cambridge, MA, USA, 2017. DOI: 10.15607/RSS.2017.XIII.049.
- [40] T. Westenbroek, F. Castaneda, A. Agrawal, S. Sastry, and K. Sreenath, “Lyapunov design for robust and efficient robotic reinforcement learning,” *arXiv preprint arXiv:2208.06721*, 2022. DOI: 10.48550/arXiv.2208.06721.
- [41] A. Anand, K. Seel, V. Gjørsum, A. Håkansson, H. Robinson, and A. Saad, “Safe learning for control using control lyapunov functions and control barrier functions: A review,” *Procedia Computer Science*, vol. 192, pp. 3987–3997, 2021. DOI: 10.1016/j.procs.2021.09.173.
- [42] P. Jagtap, G. J. Pappas, and M. Zamani, “Control barrier functions for unknown nonlinear systems using gaussian processes,” in *Proc. IEEE 59th Conf. on Decision and Control (CDC)*, Jeju, South Korea, 2020, pp. 3699–3704. DOI: 10.1109/CDC42340.2020.9303847.

- [43] M. J. Khojasteh, V. Dhiman, M. Franceschetti, and N. Atanasov, “Probabilistic safety constraints for learned high relative degree system dynamics,” in *Proc. 2nd Learning for Dynamics and Control (L4DC)*, vol. 120, Berkeley, CA, USA, 2020, pp. 781–792. [Online]. Available: <https://proceedings.mlr.press/v120/khojasteh20a.html>.
- [44] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, “End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks,” in *Proc. AAAI Conf. on Artificial Intelligence*, vol. 33, 2019, pp. 3387–3395. DOI: 10.1609/aaai.v33i01.33013387.
- [45] C. Dawson, S. Gao, and C. Fan, “Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control,” *IEEE Trans. on Robotics*, pp. 1–19, 2023. DOI: 10.1109/TR0.2022.3232542.
- [46] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, “Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects and dynamics,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 3683–3690. DOI: 10.23919/ACC50511.2021.9483420.
- [47] M. Krstić and P. V. Kokotović, “Control lyapunov functions for adaptive nonlinear stabilization,” *Sys. & Control Let.*, vol. 26, no. 1, pp. 17–23, 1995. DOI: 10.1016/0167-6911(94)00107-7.
- [48] Z. Li and M. Krstić, “Optimal design of adaptive tracking controllers for non-linear systems,” *Automatica*, vol. 33, no. 8, pp. 1459–1473, 1997. DOI: 10.1016/S0005-1098(97)00072-1.
- [49] J. Moore and R. Tedrake, “Adaptive control design for underactuated systems using sums-of-squares optimization,” in *Proc. IEEE American Control Conf. (ACC)*, Portland, OR, USA, 2014, pp. 721–728. DOI: 10.1109/ACC.2014.6859508.
- [50] R. Takano and M. Yamakita, “Robust constrained stabilization control using control lyapunov and control barrier function in the presence of measurement noises,” in *Proc. 2nd IEEE Conf. on Control Tech. and App. (CCTA)*, Copenhagen, Denmark, 2018, pp. 300–305. DOI: 10.1109/CCTA.2018.8511637.
- [51] M. Jankovic, “Robust control barrier functions for constrained stabilization of nonlinear systems,” *Automatica*, vol. 96, pp. 359–367, 2018. DOI: 10.1016/j.automatica.2018.07.004.
- [52] Y. Emam, P. Glotfelter, and M. Egerstedt, “Robust barrier functions for a fully autonomous, remotely accessible swarm-robotics testbed,” in *Proc. IEEE 58th Conf. on Decision and Control (CDC)*, Nice, France, 2019, pp. 3984–3990. DOI: 10.1109/CDC40024.2019.9029886.

- [53] K. Garg and D. Panagou, “Robust control barrier and control Lyapunov functions with fixed-time convergence guarantees,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 2292–2297. DOI: 10.23919/ACC50511.2021.9482751.
- [54] X. Tan, W. S. Cortez, and D. V. Dimarogonas, “High-order barrier functions: Robustness, safety, and performance-critical control,” *IEEE Trans. on Automatic Control*, vol. 67, no. 6, pp. 3021–3028, 2021. DOI: 10.1109/TAC.2021.3089639.
- [55] J. Buch, S.-C. Liao, and P. Seiler, “Robust control barrier functions with sector-bounded uncertainties,” *IEEE Control Sys. Let.*, vol. 6, pp. 1994–1999, 2021. DOI: 10.1109/LCSYS.2021.3136653.
- [56] E. D. Sontag, “Smooth stabilization implies coprime factorization,” *IEEE Trans. on Automatic Control*, vol. 34, no. 4, pp. 435–443, 1989. DOI: 10.1109/9.28018.
- [57] E. D. Sontag, “Input to state stability: Basic concepts and results,” in *Nonlinear and Optimal Control Theory*, Springer, 2008, pp. 163–220. DOI: 10.1007/978-3-540-77653-6_3.
- [58] S. Kolathaya and A. D. Ames, “Input-to-state safety with control barrier functions,” *IEEE Control Sys. Let.*, vol. 3, no. 1, pp. 108–113, 2018. DOI: 10.1109/LCSYS.2018.2853698.
- [59] H. J. Kushner, “Stochastic stability and control,” Brown Univ Providence RI, Tech. Rep., 1967.
- [60] S. Prajna, A. Jadbabaie, and G. J. Pappas, “Stochastic safety verification using barrier certificates,” in *Proc. IEEE 43rd Conf. on Decision and Control (CDC)*, Nassau, Bahamas, 2004, pp. 929–934. DOI: 10.1109/CDC.2004.1428804.
- [61] J. Steinhardt and R. Tedrake, “Finite-time regional verification of stochastic non-linear systems,” *The Int. J. of Robotics Research*, vol. 31, no. 7, pp. 901–923, 2012. DOI: 10.1177/0278364912444146.
- [62] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Trans. on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017. DOI: 10.1109/TAC.2016.2638961.
- [63] C. Santoyo, M. Dutreix, and S. Coogan, “Verification and control for finite-time safety of stochastic systems via barrier functions,” in *Proc. IEEE Conf. on Control Tech. and App. (CCTA)*, Hong Kong, China, 2019, pp. 712–717. DOI: 10.1109/CCTA.2019.8920407.

- [64] Z. Wang, H. Jing, C. Kurniawan, A. Chern, and Y. Nakahira, “Myopically verifiable probabilistic certificate for long-term safety,” in *Proc. IEEE American Control Conf. (ACC)*, Atlanta, GA, USA, 2022, pp. 4894–4900. DOI: 10.23919/ACC53348.2022.9867675.
- [65] A. Lavaei, S. Soudjani, and E. Frazzoli, “Safety barrier certificates for stochastic hybrid systems,” in *Proc. IEEE American Control Conf. (ACC)*, Atlanta, GA, USA, 2022, pp. 880–885. DOI: 10.23919/ACC53348.2022.9867754.
- [66] A. Clark, “Control barrier functions for complete and incomplete information stochastic systems,” in *Proc. IEEE American Control Conf. (ACC)*, Philadelphia, PA, USA, 2019, pp. 2928–2935. DOI: 10.23919/ACC.2019.8814901.
- [67] S. Yaghoubi, K. Majd, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, “Risk-bounded control using stochastic barrier functions,” *IEEE Control Sys. Let.*, vol. 5, no. 5, pp. 1831–1836, 2020. DOI: 10.1109/LCSYS.2020.3043287.
- [68] A. Chern, X. Wang, A. Iyer, and Y. Nakahira, “Safe control in the presence of stochastic uncertainties,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6640–6645. DOI: 10.1109/CDC45484.2021.9683542.
- [69] C. Santoyo, M. Dutreix, and S. Coogan, “A barrier function approach to finite-time stochastic system verification and control,” *Automatica*, vol. 125, p. 109 439, 2021. DOI: 10.1016/j.automatica.2020.109439.
- [70] C. Wang, M. Bahreinian, and R. Tron, “Chance constraint robust control with control barrier functions,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 2315–2322. DOI: 10.23919/ACC50511.2021.9482973.
- [71] C. Wang, Y. Meng, S. L. Smith, and J. Liu, “Safety-critical control of stochastic systems using stochastic control barrier functions,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 5924–5931. DOI: 10.1109/CDC45484.2021.9683349.
- [72] M. Ahmadi, X. Xiong, and A. D. Ames, “Risk-averse control via cvar barrier functions: Application to bipedal robot locomotion,” *IEEE Control Sys. Let.*, vol. 6, pp. 878–883, 2021. DOI: 10.1109/LCSYS.2021.3086854.
- [73] T. Gurriet, P. Nilsson, A. Singletary, and A. D. Ames, “Realizable set invariance conditions for cyber-physical systems,” in *Proc. IEEE American Control Conf. (ACC)*, Philadelphia, PA, USA, 2019, pp. 3642–3649. DOI: 10.23919/ACC.2019.8815332.

- [74] A. Singletary, Y. Chen, and A. D. Ames, “Control barrier functions for sampled-data systems with input delays,” in *Proc. IEEE 59th Conf. on Decision and Control (CDC)*, Jeju, South Korea, 2020, pp. 804–809. DOI: 10.1109/CDC42340.2020.9304281.
- [75] W. Chu and Z. Ghahramani, “Preference learning with gaussian processes,” in *Proc. Int. Conf. on Machine Learning (ICML)*, 2005, pp. 137–144. DOI: 10.1145/1102351.1102369.
- [76] M. Tucker, E. Novoseller, C. Kann, *et al.*, “Preference-based learning for exoskeleton gait optimization,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Paris, France, 2020, pp. 2351–2357. DOI: 10.1109/ICRA40945.2020.9196661.
- [77] M. Tucker, M. Cheng, E. Novoseller, *et al.*, “Human preference-based learning for high-dimensional optimization of exoskeleton walking gaits,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Las Vegas, NV, USA, 2020, pp. 3423–3430. DOI: 10.1109/IROS45743.2020.9341416.
- [78] K. Li, M. Tucker, E. Biyik, *et al.*, “Roial: Region of interest active learning for characterizing exoskeleton gait preference landscapes,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Xi’an, China, 2021, pp. 3212–3218. DOI: 10.1109/ICRA48506.2021.9560840.
- [79] M. Tucker, N. Csomay-Shanklin, W.-L. Ma, and A. D. Ames, “Preference-based learning for user-guided hzd gait generation on bipedal walking robots,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Xi’an, China, 2021, pp. 2804–2810. DOI: 10.1109/ICRA48506.2021.9561515.
- [80] N. Csomay-Shanklin, M. Tucker, M. Dai, J. Reher, and A. D. Ames, “Learning controller gains on bipedal walking robots via user preferences,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, 2022, pp. 10 405–10 411. DOI: 10.1109/ICRA46639.2022.9811541.
- [81] D. Nešić, A. R. Teel, and P. V. Kokotović, “Sufficient conditions for stabilization of sampled-data nonlinear systems via discrete-time approximations,” *Sys. & Control Let.*, vol. 38, no. 4-5, pp. 259–270, 1999. DOI: 10.1016/S0167-6911(99)00073-0.
- [82] D. Nešić and A. R. Teel, “A framework for stabilization of nonlinear sampled-data systems based on their approximate discrete-time models,” *IEEE Trans. on Automatic Control*, vol. 49, no. 7, pp. 1103–1122, 2004. DOI: 10.1109/TAC.2004.831175.
- [83] D. Nešić and A. R. Teel, “Backstepping on the euler approximate model for stabilization of sampled-data nonlinear systems,” in *Proc. IEEE 40th*

- Conf. on Decision and Control (CDC)*, vol. 2, Orlando, FL, USA, 2001, pp. 1737–1742. DOI: 10.1109/CDC.2001.981153.
- [84] D. Nešić and A. R. Teel, “Stabilization of sampled-data nonlinear systems via backstepping on their euler approximate model,” *Automatica*, vol. 42, no. 10, pp. 1801–1808, 2006. DOI: 10.1016/j.automatica.2006.05.015.
- [85] D. Nešić and L. Grüne, “Lyapunov-based continuous-time nonlinear controller redesign for sampled-data implementation,” *Automatica*, vol. 41, no. 7, pp. 1143–1156, 2005. DOI: 10.1016/j.automatica.2005.03.001.
- [86] L. Grüne and D. Nešić, “Optimization-based stabilization of sampled-data nonlinear systems via their approximate discrete-time models,” *SIAM J. on Control and Optimization*, vol. 42, no. 1, pp. 98–122, 2003. DOI: 10.1137/S036301290240258X.
- [87] L. Grüne, D. Nešić, and J. Pannek, “Model predictive sampled-data redesign for nonlinear systems,” in *Proc. IEEE 44th Conf. on Decision and Control (CDC)*, Seville, Spain, 2005, pp. 36–41. DOI: 10.1109/CDC.2005.1582127.
- [88] L. Grüne, D. Nešić, and J. Pannek, “Model predictive control for nonlinear sampled-data systems,” in *Assessment and future directions of nonlinear model predictive control*, Springer, 2007, pp. 105–113. DOI: 10.1007/978-3-540-72699-9_8.
- [89] G. Yang, C. Belta, and R. Tron, “Self-triggered control for safety critical systems using control barrier functions,” in *Proc. IEEE American Control Conf. (ACC)*, Philadelphia, PA, USA, 2019, pp. 4454–4459. DOI: 10.23919/ACC.2019.8814657.
- [90] J. Breeden, K. Garg, and D. Panagou, “Control barrier functions in sampled-data systems,” *IEEE Control Sys. Let.*, vol. 6, pp. 367–372, 2021. DOI: 10.1109/LCSYS.2021.3076127.
- [91] L. Niu, H. Zhang, and A. Clark, “Safety-critical control synthesis for unknown sampled-data systems via control barrier functions,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6806–6813. DOI: 10.1109/CDC45484.2021.9683019.
- [92] J. Usevitch and D. Panagou, “Adversarial resilience for sampled-data systems using control barrier function methods,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 758–763. DOI: 10.23919/ACC50511.2021.9482659.
- [93] Y. Zhang, S. Walters, and X. Xu, “Control barrier function meets interval analysis: Safety-critical control with measurement and actuation uncertainties,” in *Proc. IEEE American Control Conf. (ACC)*, Atlanta,

- GA, USA, 2022, pp. 3814–3819. DOI: 10.23919/ACC53348.2022.9867681.
- [94] D. S. Laila and D. Nešić, “Changing supply rates for input–output to state stable discrete-time nonlinear systems with applications,” *Automatica*, vol. 39, no. 5, pp. 821–835, 2003. DOI: 10.1016/S0005-1098(03)00055-4.
- [95] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, “An introduction to event-triggered and self-triggered control,” in *Proc. IEEE 51st Conf. on Decision and Control (CDC)*, Maui, HI, USA, 2012, pp. 3270–3285. DOI: 10.1109/CDC.2012.6425820.
- [96] P. Tabuada, “Event-triggered real-time scheduling of stabilizing control tasks,” *IEEE Trans. on Automatic Control*, vol. 52, no. 9, pp. 1680–1685, 2007. DOI: 10.1109/TAC.2007.904277.
- [97] D. P. Borgers and W. P. M. H. Heemels, “Event-separation properties of event-triggered control systems,” *IEEE Trans. on Automatic Control*, vol. 59, no. 10, pp. 2644–2656, 2014. DOI: 10.1109/TAC.2014.2325272.
- [98] P. Ong and J. Cortés, “Event-triggered control design with performance barrier,” in *Proc. IEEE 57th Conf. on Decision and Control (CDC)*, Miami Beach, FL, 2018, pp. 951–956. DOI: 10.1109/CDC.2018.8619524.
- [99] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum, *Feedback control theory*. Courier Corporation, 2013, ISBN: 9780486469331.
- [100] H. K. Khalil and J. W. Grizzle, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 2002, vol. 3, ISBN: 9780130673893.
- [101] L. Perko, *Differential equations and dynamical systems*. Springer Science & Business Media, 2013, vol. 7, ISBN: 9780387951164.
- [102] W. Rudin, *Principles of mathematical analysis*. McGraw-hill New York, 1964, vol. 3, ISBN: 9780070542358.
- [103] R. I. Leine, “The historical development of classical stability concepts: Lagrange, poisson and lyapunov stability,” *Nonlinear Dynamics*, vol. 59, no. 1, pp. 173–182, 2010. DOI: 10.1007/s11071-009-9530-z.
- [104] F. Blanchini and S. Miani, *Set-theoretic methods in control*. Springer, 2008, ISBN: 9780817632557.
- [105] T. Yoshizawa, *Stability theory by Liapunov’s second method*. Mathematical Society of Japan, 1966, vol. 9.
- [106] R. A. Freeman and P. V. Kokotović, “Inverse optimality in robust stabilization,” *SIAM J. on Control and Optimization*, vol. 34, no. 4, pp. 1365–1391, 1996. DOI: 10.1137/S0363012993258732.

- [107] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle, “Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics,” *IEEE Trans. on Automatic Control*, vol. 59, no. 4, pp. 876–891, 2014. DOI: 10.1109/TAC.2014.2299335.
- [108] S. Kolathaya and S. Veer, “Pd based robust quadratic programs for robotic systems,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp. 6834–6841. DOI: 10.1109/IROS40897.2019.8968214.
- [109] R. Freeman and P. V. Kokotovic, *Robust nonlinear control design: state-space and Lyapunov techniques*. Birkhauser Basel, 1996, ISBN: 9780817647582.
- [110] A. Nemirovski, “Interior point polynomial time methods in convex programming,” *Lecture notes*, 1996. [Online]. Available: https://www2.isye.gatech.edu/~nemirovs/Lect_IPM.pdf.
- [111] A. Isidori, *Nonlinear control systems*. Springer-Verlag London, 1995, vol. 3, ISBN: 9783540199168.
- [112] S. S. Sastry, *Nonlinear Systems: Analysis, Stability and Control*. Springer New York, NY, 1999, ISBN: 9780387985138.
- [113] J.-J. E. Slotine and K. J. Hedrick, “Robust input-output feedback linearization,” *Int. J. of Control*, vol. 57, no. 5, pp. 1133–1139, 1993. DOI: 10.1080/00207179308934435.
- [114] A. Isidori, “The zero dynamics of a nonlinear system: From the origin to the latest progresses of a long successful story,” *European J. of Control*, vol. 19, no. 5, pp. 369–378, 2013. DOI: 10.1016/j.ejcon.2013.05.014.
- [115] A. J. Taylor, V. D. Dorobantu, Y. Yue, P. Tabuada, and A. D. Ames, “Sampled-data stabilization with control Lyapunov functions via quadratically constrained quadratic programs,” *IEEE Control Sys. Lett.*, vol. 6, pp. 680–685, 2022. DOI: 10.1109/LCSYS.2021.3085172.
- [116] A. Hereid, C. M. Hubicki, E. A. Cousineau, J. W. Hurst, and A. D. Ames, “Hybrid zero dynamics based multiple shooting optimization with applications to robotic walking,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Seattle, WA, USA, 2015, pp. 5734–5740. DOI: 10.1109/ICRA.2015.7140002.
- [117] C.-T. Chen, *Linear System Theory and Design*. New York, NY: Oxford University Press, Inc., 1999, ISBN: 9780195117776.
- [118] M. Jankovic, “Control barrier functions for constrained control of linear systems with input delay,” in *Proc. IEEE American Control Conf. (ACC)*, Milwaukee, WI, USA, 2018, pp. 3316–3321. DOI: 10.23919/ACC.2018.8430747.

- [119] A. V. Fiacco and G. P. McCormick, *Nonlinear programming: sequential unconstrained minimization techniques*. SIAM, 1990. DOI: 10.1093/comjnl/12.3.207-a.
- [120] S. Prajna, “Barrier certificates for nonlinear model validation,” *Automatica*, vol. 42, no. 1, pp. 117–126, 2006. DOI: 10.1016/j.automatica.2005.08.007.
- [121] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *Proc. IEEE European Control Conf. (ECC)*, Naples, Italy, 2019, pp. 3420–3431. DOI: 10.23919/ECC.2019.8796030.
- [122] C. M. Kellett, “A compendium of comparison function results,” *Mathematics of Control, Signals, and Sys.*, vol. 26, no. 3, pp. 339–374, 2014. DOI: 10.1007/s00498-014-0128-8.
- [123] M. Nagumo, “Über die lage der integralkurven gewöhnlicher differentialgleichungen,” *Proc. of the Physico-Mathematical Society of Japan. 3rd Series*, vol. 24, pp. 551–559, 1942. DOI: 10.11429/ppmsj1919.24.0_551.
- [124] R. Konda, A. D. Ames, and S. Coogan, “Characterizing safety: Minimal control barrier functions from scalar comparison systems,” *IEEE Control Sys. Let.*, vol. 5, no. 2, pp. 523–528, 2020. DOI: 10.1109/LCSYS.2020.3003887.
- [125] E. D. Sontag and Y. Wang, “On characterizations of the input-to-state stability property,” *Sys. & Control Let.*, vol. 24, no. 5, pp. 351–359, 1995. DOI: 10.1016/0167-6911(94)00050-6.
- [126] S. Kolathaya, J. Reher, A. Hereid, and A. D. Ames, “Input to state stabilizing control lyapunov functions for robust bipedal robotic locomotion,” in *Proc. IEEE American Control Conf. (ACC)*, Milwaukee, WI, USA, 2018, pp. 2224–2230. DOI: 10.23919/ACC.2018.8430946.
- [127] E. D. Sontag and Y. Wang, “On characterizations of input-to-state stability with respect to compact sets,” in *Nonlinear Control Sys. Design*, Elsevier, 1995, pp. 203–208. DOI: 10.1016/B978-0-08-042371-5.50039-3.
- [128] K. J. Aström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*. Princeton University Press, 2010, ISBN: 9780691135762.
- [129] L. Ljung, *System identification*. Springer, 1998. DOI: 10.1007/978-1-4612-1768-8_11.
- [130] M. Krstić, I. Kanellakopoulos, and P. V. Kokotović, *Nonlinear and adaptive control design*. Wiley New York, 1995, ISBN: 978-0-471-12732-1.

- [131] S. Schaal and C. G. Atkeson, “Learning control in robotics,” *IEEE Robotics & Automation Mag.*, vol. 17, no. 2, pp. 20–29, 2010. DOI: 10.1109/MRA.2010.936957.
- [132] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The Int. J. of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013. DOI: 10.1177/0278364913495721.
- [133] L. Brunke, M. Greeff, A. W. Hall, *et al.*, “Safe learning in robotics: From learning-based control to safe reinforcement learning,” *Ann. Review of Control, Robotics, and Autonomous Sys.*, vol. 5, pp. 411–444, 2022. DOI: 10.1146/annurev-control-042920-020211.
- [134] T. Beckers, D. Kulić, and S. Hirche, “Stable gaussian process based tracking control of euler–lagrange systems,” *Automatica*, vol. 103, pp. 390–397, 2019. DOI: 10.1016/j.automatica.2019.01.023.
- [135] F. Berkenkamp, A. P. Schoellig, and A. Krause, “Safe controller optimization for quadrotors with gaussian processes,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Stockholm, Sweden, 2016, pp. 491–496. DOI: 10.1109/ICRA.2016.7487170.
- [136] F. Berkenkamp and A. P. Schoellig, “Safe and robust learning control with gaussian processes,” in *Proc. IEEE European Control Conf. (ECC)*, Linz, Austria, 2015, pp. 2496–2501. DOI: 10.1109/ECC.2015.7330913.
- [137] G. Shi, X. Shi, M. O’Connell, *et al.*, “Neural lander: Stable drone landing control using learned dynamics,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Montréal, QC, Canada, 2019, pp. 9784–9790. DOI: 10.1109/ICRA.2019.8794351.
- [138] T. P. Lillicrap, J. J. Hunt, A. Pritzel, *et al.*, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015. [Online]. Available: <https://arxiv.org/abs/1509.02971>.
- [139] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017. [Online]. Available: <https://arxiv.org/abs/1707.06347>.
- [140] Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel, “Benchmarking deep reinforcement learning for continuous control,” in *Proc. Int. Conf. on Machine Learning (ICML)*, New York City, NY, USA, 2016, pp. 1329–1338. [Online]. Available: <https://proceedings.mlr.press/v48/duan16.html>.
- [141] J. Tan, T. Zhang, E. Coumans, *et al.*, “Sim-to-real: Learning agile locomotion for quadruped robots,” in *Proc. Robotics: Science and Sys. XIV (RSS)*, Pittsburgh, PA, USA, 2018. DOI: 10.15607/RSS.2018.XIV.010.

- [142] H. M. Le, A. Kang, Y. Yue, and P. Carr, “Smooth imitation learning for online sequence prediction,” in *Proc. Int. Conf. on Machine Learning (ICML)*, 2016, pp. 680–688. [Online]. Available: <https://proceedings.mlr.press/v48/le16.html>.
- [143] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, “Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes,” in *Proc. IEEE 55th Conf. on Decision and Control (CDC)*, Las Vegas, NV, USA, 2016, pp. 4661–4666. DOI: 10.1109/CDC.2016.7798979.
- [144] S. M. Richards, F. Berkenkamp, and A. Krause, “The lyapunov neural network: Adaptive stability certification for safe learning of dynamic systems,” in *Proc. Conf. on Robotic Learning (CoRL)*, Zurich, Switzerland, 2018. [Online]. Available: <https://proceedings.mlr.press/v87/richards18a.html>.
- [145] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, “Safe model-based reinforcement learning with stability guarantees,” in *Proc. Advances in Neural Information Processing Sys. (NeurIPS)*, vol. 30, Long Beach, CA, USA, 2017, pp. 908–918. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/766ebcd59621e305170616ba3d3dac32-Abstract.html>.
- [146] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, “A lyapunov-based approach to safe reinforcement learning,” in *Proc. Advances in Neural Information Processing Sys. (NeurIPS)*, vol. 31, Montréal, QC, Canada, 2018, pp. 8092–8101. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/hash/4fe5149039b52765bde64beb9f674940-Abstract.html>.
- [147] M. Gallieri, S. S. M. Salehian, N. E. Toklu, *et al.*, “Safe interactive model-based learning,” *arXiv preprint arXiv:1911.06556*, 2019. DOI: 10.48550/arXiv.1911.06556.
- [148] Y.-C. Chang, N. Roohi, and S. Gao, “Neural lyapunov control,” *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/hash/2647c1dba23bc0e0f9cdf75339e120d2-Abstract.html>.
- [149] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, “Episodic learning with control lyapunov functions for uncertain robotic systems,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp. 6878–6884. DOI: 10.1109/IROS40897.2019.8967820.
- [150] A. J. Taylor, V. D. Dorobantu, M. Krishnamoorthy, H. M. Le, Y. Yue, and A. D. Ames, “A control lyapunov perspective on episodic learning via projection to state stability,” in *Proc. IEEE 58th Conf. on Decision*

- and Control (CDC)*, Nice, France, 2019, pp. 1448–1455. DOI: 10.1109/CDC40024.2019.9029226.
- [151] J. Choi, F. Castañeda, C. Tomlin, and K. Sreenath, “Reinforcement Learning for Safety-Critical Control under Model Uncertainty, using Control Lyapunov Functions and Control Barrier Functions,” in *Proc. Robotics: Science and Sys. XVI (RSS)*, Corvallis, OR, USA, 2020. DOI: 10.15607/RSS.2020.XVI.088.
- [152] A. J. Taylor, V. D. Dorobantu, S. Dean, B. Recht, Y. Yue, and A. D. Ames, “Towards robust data-driven control synthesis for nonlinear systems with actuation uncertainty,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6469–6476. DOI: 10.1109/CDC45484.2021.9683511.
- [153] K. P. Wabersich, A. J. Taylor, J. J. Choi, *et al.*, “Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems,” 2023. [Online]. Available: https://andrewjamestaylor.github.io/assets/paper_materials/wabersich2022data_driven_safety_filters_hamilton_jacobi_reachability_control_barrier_functions_and_predictive_methods_for_uncertain_systems/preprint.pdf.
- [154] D. D. Fan, J. Nguyen, R. Thakker, N. Alatur, A.-a. Agha-mohammadi, and E. A. Theodorou, “Bayesian learning-based adaptive control for safety critical systems,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Paris, France, 2020, pp. 4093–4099. DOI: 10.1109/ICRA40945.2020.9196709.
- [155] V. Dhiman, M. J. Khojasteh, M. Franceschetti, and N. Atanasov, “Control barriers in bayesian learning of system dynamics,” *IEEE Trans. on Automatic Control*, vol. 68, pp. 214–229, 1 2023. DOI: 10.1109/TAC.2021.3137059.
- [156] Y. Emam, P. Glotfelter, S. Wilson, G. Notomista, and M. Egerstedt, “Data-driven robust barrier functions for safe, long-term operation,” *IEEE Trans. on Robotics*, vol. 38, no. 3, pp. 1671–1685, 2021. DOI: 10.1109/TR0.2021.3118965.
- [157] Y. Hu, J. Fu, and G. Wen, “Decentralized robust collision avoidance for cooperative multirobot systems: A gaussian process-based control barrier function approach,” *IEEE Trans. on Control of Network Sys.*, pp. 1–12, 2022. DOI: 10.1109/TCNS.2022.3203928.
- [158] A. Lederer, A. Begzadić, N. Das, and S. Hirche, “Safe learning-based control of elastic joint robots via control barrier functions,” *arXiv preprint arXiv:2212.00478*, 2022. DOI: 10.48550/arXiv.2212.00478.

- [159] V. Zinage and E. Bakolas, “Neural koopman control barrier functions for safety-critical control of unknown nonlinear systems,” *arXiv preprint arXiv:2209.07685*, 2022. DOI: 10.48550/arXiv.2209.07685.
- [160] C. Folkestad, Y. Chen, A. D. Ames, and J. W. Burdick, “Data-driven safety-critical control: Synthesizing control barrier functions with koopman operators,” *IEEE Control Sys. Let.*, vol. 5, no. 6, pp. 2012–2017, 2020. DOI: 10.1109/LCSYS.2020.3046159.
- [161] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “Learning for safety-critical control with control barrier functions,” in *Proc. 2nd Learning for Dynamics and Control (L4DC)*, vol. 120, Berekeley, CA, USA, 2020, pp. 708–717. [Online]. Available: <https://proceedings.mlr.press/v120/taylor20a.html>.
- [162] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “A control barrier perspective on episodic learning via projection-to-state safety,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 1019–1024, 2021. DOI: 10.1109/LCSYS.2020.3009082.
- [163] E. Munoz, D. Kalaria, Q. Lin, and J. M. Dolan, “Online adaptive compensation for model uncertainty using extreme learning machine-based control barrier functions,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Kyoto, Japan, 2022, pp. 10 959–10 966. DOI: 10.1109/IROS47612.2022.9981680.
- [164] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, “Pointwise feasibility of gaussian process-based safety-critical control under model uncertainty,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6762–6769. DOI: 10.1109/CDC45484.2021.9683743.
- [165] F. Castañeda, J. J. Choi, W. Jung, B. Zhang, C. J. Tomlin, and K. Sreenath, “Probabilistic safe online learning with control barrier functions,” *arXiv preprint arXiv:2208.10733*, 2022. DOI: 10.48550/arXiv.2208.10733.
- [166] M. H. Cohen, C. Belta, and R. Tron, “Robust control barrier functions for nonlinear control systems with uncertainty: A duality-based approach,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 174–179. DOI: 10.1109/CDC51059.2022.9992667.
- [167] Z. Jin, M. Khajenejad, and S. Z. Yong, “Robust data-driven control barrier functions for unknown continuous control affine systems,” *IEEE Control Sys. Let.*, vol. 7, pp. 1309–1314, 2023. DOI: 10.1109/LCSYS.2023.3235958.

- [168] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt, “Barrier-certified adaptive reinforcement learning with applications to brushbot navigation,” *IEEE Trans. on Robotics*, vol. 35, no. 5, pp. 1186–1205, 2019. DOI: 10.1109/TR0.2019.2920206.
- [169] S. Liu, L. Liu, and Z. Yu, “Safe reinforcement learning for affine nonlinear systems with state constraints and input saturation using control barrier functions,” *Neurocomputing*, vol. 518, pp. 562–576, 2023. DOI: 10.1016/j.neucom.2022.11.006.
- [170] Q. Zhao, Y. Zhang, and X. Li, “Safe reinforcement learning for dynamical systems using barrier certificates,” *Connection Science*, vol. 34, no. 1, pp. 2822–2844, 2022. DOI: 10.1080/09540091.2022.2151567.
- [171] M. H. Cohen and C. Belta, “Approximate optimal control for safety-critical systems with control barrier functions,” in *Proc. IEEE 59th Conf. on Decision and Control (CDC)*, 2020, pp. 2062–2067. DOI: 10.1109/CDC42340.2020.9303896.
- [172] Z. Cai, H. Cao, W. Lu, L. Zhang, and H. Xiong, “Safe multi-agent reinforcement learning through decentralized multiple control barrier functions,” *arXiv preprint arXiv:2103.12553*, 2021. DOI: 10.48550/arXiv.2103.12553.
- [173] Y. Yang, L. Chen, and M. Gombolay, “Safe inverse reinforcement learning via control barrier function,” *arXiv preprint arXiv:2212.02753*, 2022. DOI: 10.48550/arXiv.2212.02753.
- [174] R. K. Cosner, Y. Yue, and A. D. Ames, “End-to-end imitation learning with safety guarantees using control barrier functions,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 5316–5322. DOI: 10.1109/CDC51059.2022.9993193.
- [175] M. Saveriano and D. Lee, “Learning barrier functions for constrained motion planning with dynamical systems,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp. 112–119. DOI: 10.1109/IROS40897.2019.8967981.
- [176] A. Robey, H. Hu, L. Lindemann, *et al.*, “Learning control barrier functions from expert demonstrations,” in *Proc. IEEE 59th Conf. on Decision and Control (CDC)*, Jeju, South Korea, 2020, pp. 3717–3724. DOI: 10.1109/CDC42340.2020.9303785.
- [177] L. Lindemann, H. Hu, A. Robey, *et al.*, “Learning hybrid control barrier functions from data,” in *Proc. Conf. on Robot Learning (CoRL)*, London, UK, 2021, pp. 1351–1370. [Online]. Available: <https://proceedings.mlr.press/v155/lindemann21a.html>.

- [178] K. Long, C. Qian, J. Cortés, and N. Atanasov, “Learning barrier functions with memory for robust safe navigation,” *IEEE Robotics and Automation Let.*, vol. 6, no. 3, pp. 4931–4938, 2021. DOI: 10.1109/LRA.2021.3070250.
- [179] E. Scukins and P. Ögren, “Using reinforcement learning to create control barrier functions for explicit risk mitigation in adversarial environments,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, 2021. DOI: 10.1109/ICRA48506.2021.9561853.
- [180] A. S. Lafmejani, S. Berman, and G. Fainekos, “Nmpc-lbf: Nonlinear mpc with learned barrier function for decentralized safe navigation of multiple robots in unknown environments,” *arXiv preprint arXiv:2208.07974*, 2022.
- [181] Z. Liang and J. K. C. Lo, “An iterative method to learn a linear control barrier function,” *arXiv preprint arXiv:2211.09854*, 2022. DOI: 10.48550/arXiv.2211.09854.
- [182] S. McIlvanna, N. N. Minh, Y. Sun, M. Van, and W. Naeem, “Reinforcement learning-enhanced control barrier functions for robot manipulators,” *arXiv preprint arXiv:2211.11391*, 2022. DOI: 10.48550/arXiv.2211.11391.
- [183] H. Parwana and D. Panagou, “Foresee: Model-based reinforcement learning using unscented transform with application to tuning of control barrier functions,” *arXiv preprint arXiv:2209.12644*, 2022. DOI: 10.48550/arXiv.2209.12644.
- [184] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, “Guaranteeing safety of learned perception modules via measurement-robust control barrier functions,” in *Proc. Conf. on Robotics Learning (CoRL)*, Cambridge, MA, USA, 2020. [Online]. Available: <https://proceedings.mlr.press/v155/dean21a/dean21a.pdf>.
- [185] R. K. Cosner, I. D. J. Rodriguez, T. G. Molnar, *et al.*, “Self-supervised online learning for safety-critical control using stereo vision,” in *Proc. Int. Conf. on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, 2022, pp. 11487–11493. DOI: 10.1109/ICRA46639.2022.9812183.
- [186] I. D. J. Rodriguez, N. Csomay-Shanklin, Y. Yue, and A. D. Ames, “Neural gaits: Learning bipedal locomotion via control barrier functions and zero dynamics policies,” in *Proc. 4th Learning for Dynamics and Control (L4DC)*, vol. 168, Stanford, CA, USA, 2022, pp. 1060–1072. [Online]. Available: <https://proceedings.mlr.press/v168/rodriguez22a/rodriguez22a.pdf>.

- [187] Z. Qin, T.-W. Weng, and S. Gao, “Quantifying safety of learning-based self-driving control using almost-barrier functions,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Kyoto, Japan, 2022, pp. 12 903–12 910. DOI: 10.1109/IROS47612.2022.9982058.
- [188] A. Didier, R. C. Jacobs, J. Sieber, K. P. Wabersich, and M. N. Zeilinger, “Approximate predictive control barrier functions using neural networks: A computationally cheap and permissive safety filter,” *arXiv preprint arXiv:2211.15104*, 2022. DOI: 10.48550/arXiv.2211.15104.
- [189] M. Krstić and P. V. Kokotović, “Modular approach to adaptive nonlinear stabilization,” *Automatica*, vol. 32, no. 4, pp. 625–629, 1996. DOI: 10.1016/0005-1098(95)00179-4.
- [190] A. J. Taylor and A. D. Ames, “Adaptive safety with control barrier functions,” in *Proc. IEEE American Control Conf. (ACC)*, Denver, CO, USA, 2020, pp. 1399–1405. DOI: 10.23919/ACC45564.2020.9147463.
- [191] B. T. Lopez, J. E. Slotine, and J. P. How, “Robust adaptive control barrier functions: An adaptive and data-driven approach to safety,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 1031–1036, 2021. DOI: 10.1109/LCSYS.2020.3005923.
- [192] G. Chowdhary and E. Johnson, “Concurrent learning for convergence in adaptive control without persistency of excitation,” in *Proc. IEEE 49th Conf. on Decision and Control (CDC)*, Atlanta, GA, USA, 2010, pp. 3674–3679. DOI: 10.1109/CDC.2010.5717148.
- [193] M. Black, E. Arabi, and D. Panagou, “A fixed-time stable adaptation law for safety-critical control under parametric uncertainty,” in *Proc. IEEE European Control Conf. (ECC)*, Delft, Netherlands, 2021, pp. 1328–1333. DOI: 10.23919/ECC54610.2021.9655080.
- [194] M. Maghenem, A. J. Taylor, A. D. Ames, and R. G. Sanfelice, “Adaptive safety using control barrier functions and hybrid adaptation,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 2418–2423. DOI: 10.23919/ACC50511.2021.9482871.
- [195] A. Isaly, O. S. Patil, R. G. Sanfelice, and W. E. Dixon, “Adaptive safety with multiple barrier functions using integral concurrent learning,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 3719–3724. DOI: 10.23919/ACC50511.2021.9483208.
- [196] M. H. Cohen and C. Belta, “High order robust adaptive control barrier functions and exponentially stabilizing adaptive control lyapunov functions,” in *Proc. IEEE American Control Conf. (ACC)*, Atlanta, GA, USA, 2022, pp. 2233–2238. DOI: 10.23919/ACC53348.2022.9867633.
- [197] B. T. Lopez and J.-J. Slotine, “Unmatched control barrier functions: Certainty equivalence adaptive safety,” *arXiv preprint arXiv:2207.13873*, 2022. DOI: 10.48550/arXiv.2207.13873.

- [198] S. Wang, B. Lyu, S. Wen, K. Shi, S. Zhu, and T. Huang, “Robust adaptive safety-critical control for unknown systems with finite-time elementwise parameter estimation,” *IEEE Trans. on Sys., Man, and Cybernetics: Sys.*, vol. 53, no. 3, pp. 1607–1617, 2023. DOI: 10.1109/TSMC.2022.3203176.
- [199] M. Cohen and C. Belta, “Modular adaptive safety-critical control,” *arXiv preprint arXiv:2303.04241*, 2023. DOI: 10.48550/arXiv.2303.04241.
- [200] Y. Wang and X. Xu, “Adaptive safety-critical control for a class of non-linear systems with parametric uncertainties: A control barrier function approach,” *arXiv preprint arXiv:2302.08601*, 2023. DOI: 10.48550/arXiv.2302.08601.
- [201] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A general safety framework for learning-based control in uncertain robotic systems,” *IEEE Trans. on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2018. DOI: 10.1109/TAC.2018.2876389.
- [202] L. Zheng, R. Yang, J. Pan, H. Cheng, and H. Hu, “Learning-based safety-stability-driven control for safety-critical systems under model uncertainties,” in *Proc. IEEE Int. Conf. on Wireless Communications and Signal Processing (WCSP)*, Nanjing, China, 2020, pp. 1112–1118. DOI: 10.1109/WCSP49889.2020.9299833.
- [203] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *Proc. 14th Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, Fort Lauderdale, FL, USA, 2011, pp. 627–635. [Online]. Available: <https://proceedings.mlr.press/v15/ross11a.html>.
- [204] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk, *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006, vol. 1. DOI: 10.1007/b97848.
- [205] D. Ernst, P. Geurts, and L. Wehenkel, “Tree-based batch mode reinforcement learning,” *J. of Machine Learning Research*, vol. 6, no. Apr, pp. 503–556, 2005. [Online]. Available: <https://www.jmlr.org/papers/volume6/ernst05a/ernst05a.pdf>.
- [206] K. S. Narendra and A. M. Annaswamy, “Persistent excitation in adaptive systems,” *Int. J. of Control*, vol. 45, no. 1, pp. 127–160, 1987. DOI: 10.1080/00207178708933715.
- [207] M. A. Patterson and A. V. Rao, “Gpops-ii: A matlab software for solving multiple-phase optimal control problems using hp-adaptive gaussian quadrature collocation methods and sparse nonlinear programming,” *ACM Trans. on Mathematical Software (TOMS)*, vol. 41, no. 1, pp. 1–37, 2014. DOI: 10.1145/2558904.

- [208] F. Chollet *et al.*, *Keras*, <https://keras.io>, 2015.
- [209] A. Singletary, P. Nilsson, T. Gurriet, and A. D. Ames, “Online active safety for robotic manipulators,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Macau, China, 2019, pp. 173–178. DOI: 10.1109/IROS40897.2019.8968231.
- [210] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek, “Hybrid zero dynamics of planar biped walkers,” *IEEE Trans. on Automatic Control*, vol. 48, no. 1, pp. 42–56, 2003. DOI: 10.1109/TAC.2002.806653.
- [211] E. Ambrose, W.-L. Ma, C. Hubicki, and A. D. Ames, “Toward benchmarking locomotion economy across design configurations on the modular robot: Amber-3m,” in *Proc. 1st IEEE Conf. on Control Tech. and App. (CCTA)*, Kohala Coast, HI, USA, 2017, pp. 1270–1276. DOI: 10.1109/CCTA.2017.8062633.
- [212] W.-L. Ma, H.-H. Zhao, S. Kolathaya, and A. D. Ames, “Human-inspired walking via unified pd and impedance control,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Hong Kong, China, 2014, pp. 5088–5094. DOI: 10.1109/ICRA.2014.6907605.
- [213] A. Hereid, C. M. Hubicki, E. A. Cousineau, and A. D. Ames, “Dynamic humanoid locomotion: A scalable formulation for hzd gait optimization,” *IEEE Trans. on Robotics*, vol. 34, no. 2, pp. 370–387, 2018. DOI: 10.1109/TR0.2017.2783371.
- [214] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid dynamical systems*. Princeton University Press, 2012, ISBN: 9780691153896.
- [215] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. H. Choi, and B. Morris, *Feedback Control of Dynamic Bipedal Robot Locomotion*. Taylor & Francis/CRC Press, 2007, ISBN: 1-42005-372-8.
- [216] Q. Nguyen and K. Sreenath, “Exponential control barrier functions for enforcing high relative-degree safety-critical constraints,” in *Proc. IEEE American Control Conf. (ACC)*, Boston, MA, USA, 2016, pp. 322–328. DOI: 10.1109/ACC.2016.7524935.
- [217] J. Hwangbo, J. Lee, and M. Hutter, “Per-contact iteration method for solving contact dynamics,” *IEEE Robotics and Automation Let.*, vol. 3, no. 2, pp. 895–902, 2018. DOI: 10.1109/LRA.2018.2792536.
- [218] R. Dimitrova and R. Majumdar, “Deductive control synthesis for alternating-time logics,” in *Proc. IEEE 14th Int. Conf. on Embedded Software (EMSOFT)*, Uttar Pradesh, India, 2014, pp. 1–10. DOI: 10.1145/2656045.2656054.

- [219] N. Boffi, S. Tu, N. Matni, J.-J. Slotine, and V. Sindhvani, “Learning stability certificates from data,” in *Proc. Conf. on Robot Learning (CoRL)*, Cambridge, MA, USA, 2020, pp. 1341–1350. [Online]. Available: <https://proceedings.mlr.press/v155/boffi21a.html>.
- [220] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton University Press, 2009, ISBN: 978-0-691-14368-2.
- [221] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, “Robustness of control barrier functions for safety critical control,” *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015. DOI: 10.1016/j.ifacol.2015.11.152.
- [222] B. J. Morris, M. J. Powell, and A. D. Ames, “Continuity and smoothness properties of nonlinear optimization-based feedback controllers,” in *Proc. IEEE 54th Conf. on Decision and Control (CDC)*, Osaka, Japan, 2015, pp. 151–158. DOI: 10.1109/CDC.2015.7402101.
- [223] B. S. Mordukhovich, J. V. Outrata, and M. E. Sarabi, “Full stability of locally optimal solutions in second-order cone programs,” *SIAM J. on Optimization*, vol. 24, no. 4, pp. 1581–1613, 2014. DOI: 10.1137/130928637.
- [224] B. S. Mordukhovich, *Variational analysis and applications*. Springer, 2018, vol. 30. DOI: 10.1007/978-3-319-92775-6.
- [225] A. Lederer, A. J. O. Conejo, K. Maier, W. Xiao, and S. Hirche, “Real-time regression with dividing local gaussian processes,” *arXiv preprint arXiv:2006.09446*, 2020. DOI: 10.48550/arXiv.2006.09446.
- [226] C. E. Rasmussen, C. K. I. Williams, *et al.*, *Gaussian processes for machine learning*. Springer, 2006, vol. 1, pp. 63–71, ISBN: 3-540-23122-6.
- [227] A. Jain, P. Sarda, and J. R. Haritsa, “Providing diversity in k-nearest neighbor query results,” in *Advances in Knowledge Discovery and Data Mining: 8th Pacific-Asia Conf. (PAKDD)*, Springer, Sydney, Australia, 2004, pp. 404–413. DOI: 10.1007/978-3-540-24775-3_49.
- [228] A. Alan, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, “Safe controller synthesis with tunable input-to-state safe control barrier functions,” *IEEE Control Sys. Let.*, vol. 6, pp. 908–913, 2022. DOI: 10.1109/LCSYS.2021.3087443.
- [229] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, “Robust control barrier-value functions for safety-critical control,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6814–6821. DOI: 10.1109/CDC45484.2021.9683085.
- [230] E. Daş and R. M. Murray, “Robust safe control synthesis with disturbance observer-based control barrier functions,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 5566–5573. DOI: 10.1109/CDC51059.2022.9993032.

- [231] A. Alan, T. G. Molnar, E. Daş, A. D. Ames, and G. Orosz, “Disturbance observers for robust safety-critical control with control barrier functions,” *IEEE Control Sys. Let.*, vol. 7, pp. 1123–1128, 2022. DOI: 10.1109/LCSYS.2022.3232059.
- [232] P. Seiler, M. Jankovic, and E. Hellstrom, “Control barrier functions with unmodeled input dynamics using integral quadratic constraints,” *IEEE Control Sys. Let.*, vol. 6, pp. 1664–1669, 2021. DOI: 10.1109/LCSYS.2021.3130782.
- [233] L. C. Evans, *An introduction to stochastic differential equations*. American Mathematical Soc., 2012, vol. 82, ISBN: 978-1-4704-1054-4.
- [234] A. Agrawal and K. Sreenath, “Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation,” in *Proc. Robotics: Science and Sys. XIII (RSS)*, Cambridge, MA, USA, vol. 13, 2017. DOI: 10.15607/RSS.2017.XIII.073.
- [235] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, “Robust safety under stochastic uncertainty with discrete-time control barrier functions,” *arXiv preprint arXiv:2302.07469*, 2023. DOI: 10.48550/arXiv.2302.07469.
- [236] C. R. He and G. Orosz, “Safety guaranteed connected cruise control,” in *Proc. IEEE 21st Int. Conf. on Intelligent Transportation Sys. (ITSC)*, Maui, Hawaii, USA, 2018, pp. 549–554. DOI: 10.1109/ITSC.2018.8569979.
- [237] L. Zhang and G. Orosz, “Motif-based design for connected vehicle systems in presence of heterogeneous connectivity structures and time delays,” *IEEE Trans. on Intelligent Transportation Sys.*, vol. 17, no. 6, pp. 1638–1651, 2016. DOI: 10.1109/TITS.2015.2509782.
- [238] J. I. Ge, S. S. Avedisov, C. R. He, W. B. Qin, M. Sadeghpour, and G. Orosz, “Experimental validation of connected automated vehicle design among human-driven vehicles,” *Transportation Research Part C: Emerging Tech.*, vol. 91, pp. 335–352, 2018. DOI: 10.1016/j.trc.2018.04.005.
- [239] International-Trucks, *Prostar+ truck*. [Online]. Available: <https://www.internationaltrucks.com/trucks/prostar>.
- [240] Commsignia, *V2x onboard unit*. [Online]. Available: <https://www.commsignia.com/products/obu/>.
- [241] Speedgoat, *Mobile real-time target machine*. [Online]. Available: <https://www.speedgoat.com/products-services/real-time-target-machines/mobile>.

- [242] D. R. Agrawal, H. Parwana, R. K. Cosner, U. Rosolia, A. D. Ames, and D. Panagou, “A constructive method for designing safe multirate controllers for differentially-flat systems,” *IEEE Control Systems Let.*, vol. 6, pp. 2138–2143, 2021. DOI: 10.1109/LCSYS.2021.3136465.
- [243] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, “Safe policy synthesis in multi-agent pomdps via discrete-time barrier functions,” in *Proc. IEEE 58th Conf. on Decision and Control (CDC)*, Nice, France, 2019, pp. 4797–4803. DOI: 10.1109/CDC40024.2019.9030241.
- [244] J. Zeng, B. Zhang, and K. Sreenath, “Safety-critical model predictive control with discrete-time control barrier function,” in *Proc. IEEE American Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 3882–3889. DOI: 10.23919/ACC50511.2021.9483029.
- [245] G. Grimmett and D. Stirzaker, *Probability and random processes*. Oxford university press, 2020, ISBN: 978-0-19-884760-1.
- [246] J. Ville, “Etude critique de la notion de collectif,” 1939. [Online]. Available: http://www.numdam.org/item/THESE_1939__218__1_0.pdf.
- [247] J. G. Liao and A. Berg, “Sharpening jensen’s inequality,” *The American Statistician*, vol. 73, no. 3, pp. 278–281, 2019. DOI: 10.1080/00031305.2017.1419145.
- [248] R. A. Becker, “The variance drain and jensen’s inequality,” 2012. [Online]. Available: <https://caepri.indiana.edu/RePEc/inu/caeprp/CAEPR2012-004.pdf>.
- [249] S. Abbott, *Understanding analysis*. Springer, 2001, vol. 2, ISBN: 978-1-4939-2711-1.
- [250] K. Shebrawi and H. Albadawi, “Trace inequalities for matrices,” *Bul. of the Australian Mathematical Society*, vol. 87, no. 1, pp. 139–148, 2013. DOI: 10.1017/S0004972712000627.
- [251] J. Buchli, M. Kalakrishnan, M. Mistry, P. Pastor, and S. Schaal, “Compliant quadruped locomotion over rough terrain,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, St. Louis, MO, USA, 2009, pp. 814–820. DOI: 10.1109/IROS.2009.5354681.
- [252] W. Ubellacker, N. Csomay-Shanklin, T. G. Molnar, and A. D. Ames, “Verifying safe transitions between dynamic motion primitives on legged robots,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Prague, Czech Republic, 2021, pp. 8477–8484. DOI: 10.1109/IROS51168.2021.9636537.
- [253] T. G. Molnar, R. K. Cosner, A. W. Singletary, W. Ubellacker, and A. D. Ames, “Model-free safety-critical control for robotic systems,” *IEEE Robotics and Automation Let.*, vol. 7, no. 2, pp. 944–951, 2021. DOI: 10.1109/LRA.2021.3135569.

- [254] D. Angeli, E. D. Sontag, and Y. Wang, “A characterization of integral input-to-state stability,” *IEEE Trans. on Automatic Control*, vol. 45, no. 6, pp. 1082–1097, 2000. DOI: 10.1109/9.863594.
- [255] J. L. Crassidis and J. L. Junkins, *Optimal estimation of dynamic systems*. CRC press, 2011, ISBN: 978-1-4398-3985-0.
- [256] T. Gurriet, M. Mote, A. Singletary, P. Nilsson, E. Feron, and A. D. Ames, “A scalable safety critical control framework for nonlinear systems,” *IEEE Access*, vol. 8, pp. 187 249–187 275, 2020. DOI: 10.1109/ACCESS.2020.3025248.
- [257] K. Long, V. Dhiman, M. Leok, J. Cortés, and N. Atanasov, “Safe control synthesis with uncertain dynamics and constraints,” *IEEE Robotics and Automation Let.*, vol. 7, no. 3, pp. 7295–7302, 2022. DOI: 10.1109/LRA.2022.3182544.
- [258] Z. Li and N. Atanasov, “Governor-parameterized barrier function for safe output tracking with locally sensed constraints,” *Automatica*, vol. 152, p. 110 996, 2023. DOI: 10.1016/j.automatica.2023.110996.
- [259] I. Tezuka and H. Nakamura, “Time-varying obstacle avoidance by using high-gain observer and input-to-state constraint safe control barrier function,” *IFAC-PapersOnLine*, vol. 53, no. 5, pp. 391–396, 2020. DOI: 10.1016/j.ifacol.2021.04.192.
- [260] Y. Wang and X. Xu, “Observer-based control barrier functions for safety critical systems,” in *Proc. IEEE American Control Conf. (ACC)*, Atlanta, GA, USA, 2022, pp. 709–714. DOI: 10.23919/ACC53348.2022.9867262.
- [261] D. R. Agrawal and D. Panagou, “Safe and robust observer-controller synthesis using control barrier functions,” *IEEE Control Sys. Let.*, vol. 7, pp. 127–132, 2022. DOI: 10.1109/LCSYS.2022.3185142.
- [262] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, “End-to-end driving via conditional imitation learning,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Brisbane, QLD, Australia, 2018, pp. 4693–4700. DOI: 10.1109/ICRA.2018.8460487.
- [263] A. Lambert, A. Shaban, A. Raj, Z. Liu, and B. Boots, “Deep forward and inverse perceptual models for tracking and prediction,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, Brisbane, QLD, Australia, 2018, pp. 675–682. DOI: 10.1109/ICRA.2018.8461050.
- [264] S. Tang, V. Wüest, and V. Kumar, “Aggressive flight with suspended payloads using vision-based control,” *IEEE Robotics and Automation Let.*, vol. 3, no. 2, pp. 1152–1159, 2018. DOI: 10.1109/LRA.2018.2793305.

- [265] S. Dean, N. Matni, B. Recht, and V. Ye, “Robust guarantees for perception-based control,” in *Proc. 2nd Learning for Dynamics and Control (L4DC)*, Berkeley, CA, USA, 2020, pp. 350–360. [Online]. Available: <http://proceedings.mlr.press/v120/dean20a.html>.
- [266] S. Dean and B. Recht, “Certainty equivalent perception-based control,” in *Proc. 3rd Learning for Dynamics and Control (L4DC)*, vol. 144, Zürich, Switzerland, 2021, pp. 399–411. [Online]. Available: <http://proceedings.mlr.press/v144/dean21a.html>.
- [267] A. Khazraei, H. Pfister, and M. Pajic, “Resiliency of perception-based controllers against attacks,” in *Proc. 4th Learning for Dynamics and Control (L4DC)*, PMLR, Stanford, CA, USA, 2022, pp. 713–725. [Online]. Available: <https://proceedings.mlr.press/v168/khazraei22a.html>.
- [268] A. Khazraei and M. Pfister H.and Pajic, “Attacks on perception-based control systems: Modeling and fundamental limits,” *arXiv preprint arXiv:2206.07150*, 2022. DOI: 10.48550/arXiv.2206.07150.
- [269] H. Zhang, S. Cheng, L. Niu, and A. Clark, “Barrier certificate based safe control for lidar-based systems under sensor faults and attacks,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 2256–2263. DOI: 10.1109/CDC51059.2022.9992432.
- [270] M. Tong, C. Dawson, and C. Fan, “Enforcing safety for vision-based controllers via control barrier functions and neural radiance fields,” *arXiv preprint arXiv:2209.12266*, 2022. DOI: 10.48550/arXiv.2209.12266.
- [271] D. Sun, N. Musavi, G. Dullerud, S. Shakkottai, and S. Mitra, “Learning certifiably robust controllers using fragile perception,” *arXiv preprint arXiv:2209.11328*, 2022. DOI: 10.48550/arXiv.2209.11328.
- [272] Y. Yue, J. Broder, R. Kleinberg, and T. Joachims, “The k-armed dueling bandits problem,” *J. of Computer and Sys. Sciences*, vol. 78, no. 5, pp. 1538–1556, 2012. DOI: 10.1016/j.jcss.2011.12.028.
- [273] P. Shivaswamy and T. Joachims, “Online structured prediction via coactive learning,” *arXiv preprint arXiv:1205.4213*, 2012. DOI: 10.48550/arXiv.1205.4213.
- [274] M. De Gemmis, L. Iaquinta, P. Lops, C. Musto, F. Narducci, and G. Semeraro, “Preference learning in recommender systems,” *Preference Learning*, vol. 41, pp. 41–55, 2009. [Online]. Available: <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=e91613a8d5a48b698e1a4f3f5bb73c78fc22348e#page=45>.

- [275] K. Raman, T. Joachims, P. Shivaswamy, and T. Schnabel, “Stable coactive learning via perturbation,” in *Proc. Int. Conf. on Machine Learning (ICML)*, PMLR, 2013, pp. 837–845. [Online]. Available: <https://proceedings.mlr.press/v28/raman13.html>.
- [276] A. Jain, S. Sharma, T. Joachims, and A. Saxena, “Learning preferences for manipulation tasks from online coactive feedback,” *The Int. J. of Robotics Research*, vol. 34, no. 10, pp. 1296–1313, 2015. DOI: 10.1177/0278364915581193.
- [277] D. Sadigh, A. Dragan, S. Sastry, and S. Seshia, “Active preference-based learning of reward functions,” in *Proc. of Robotics: Science and Sys. XIII (RSS)*, Cambridge, MA, USA, 2017. DOI: 10.15607/RSS.2017.XIII.053.
- [278] E. Biyik, N. Huynh, M. Kochenderfer, and D. Sadigh, “Active Preference-Based Gaussian Process Regression for Reward Learning,” in *Proc. of Robotics: Science and Sys. XVI (RSS)*, Corvallis, OR, USA, 2020. DOI: 10.15607/RSS.2020.XVI.041.
- [279] Y. Sui, V. Zhuang, J. W. Burdick, and Y. Yue, “Stagewise safe bayesian optimization with gaussian processes,” in *Proc. Int. Conf. on Machine Learning (ICML)*, 2018, pp. 4781–4789. [Online]. Available: <https://proceedings.mlr.press/v80/sui18a.html>.
- [280] Y. Sui, A. Gotovos, J. Burdick, and A. Krause, “Safe exploration for optimization with gaussian processes,” in *Proc. Int. Conf. on Machine Learning (ICML)*, 2015, pp. 997–1005. [Online]. Available: <https://proceedings.mlr.press/v37/sui15.html>.
- [281] M. Tucker, K. Li, Y. Yue, and A. D. Ames, “Polar: Preference optimization and learning algorithms for robotics,” *arXiv preprint arXiv:2208.04404*, 2022. DOI: 10.48550/arXiv.2208.04404.
- [282] A. L. Dontchev and W. W. Hager, “Lipschitzian stability in nonlinear control and optimization,” *SIAM J. on Control and Optimization*, vol. 31, no. 3, pp. 569–603, 1993. DOI: 10.1137/0331026.
- [283] W. Tan and A. Packard, “Stability region analysis using sum of squares programming,” in *Proc. IEEE American Control Conf. (ACC)*, Minneapolis, MN, USA, 2006, pp. 2297–2302. DOI: 10.1109/ACC.2006.1656562.
- [284] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability theory: new directions*. Springer Science & Business Media, 2011, ISBN: 978-3-642-16683-9.
- [285] W. Chu, Z. Ghahramani, and C. K. I. Williams, “Gaussian processes for ordinal regression.” *J. of machine learning research*, vol. 6, no. 7, pp. 1019–1041, 2005. [Online]. Available: <https://www.jmlr.org/papers/volume6/chu05a/chu05a.pdf>.

- [286] O. Chapelle and L. Li, “An empirical evaluation of thompson sampling,” *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 24, 2011. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2011/hash/e53a0a2978c28872a4505bdb51db06dc-Abstract.html.
- [287] J. Kirschner, M. Mutny, N. Hiller, R. Ischebeck, and A. Krause, “Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces,” in *Proc. Int. Conf. on Machine Learning (ICML)*, 2019, pp. 3429–3438. [Online]. Available: <http://proceedings.mlr.press/v97/kirschner19a.html>.
- [288] J. C. Maxwell, “I. on governors,” *Proc. of the Royal Society of London*, no. 16, pp. 270–283, 1868. DOI: 10.1098/rspl.1867.0055.
- [289] S. Monaco and D. Normand-Cyrot, “Advanced tools for nonlinear sampled-data systems’ analysis and control,” in *Proc. IEEE European Control Conf. (ECC)*, 2007, pp. 1155–1158. DOI: 10.23919/ECC.2007.7069061.
- [290] M. Di Ferdinando and P. Pepe, “Robustification of sample-and-hold stabilizers for control-affine time-delay systems,” *Automatica*, vol. 83, pp. 141–154, 2017. DOI: 10.1016/j.automatica.2017.06.029.
- [291] C. Li, C. Ding, and K. Shen, “Quantifying the cost of context switch,” in *Proc. Work. on Experimental Computer Science (ExpCS)*, San Diego, CA, USA, 2007, pp. 1–4. DOI: 10.1145/1281700.1281702.
- [292] P. Ong, G. Bahati, and A. D. Ames, “Stability and safety through event-triggered intermittent control with application to spacecraft orbit stabilization,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 453–460. DOI: 10.1109/CDC51059.2022.9992757.
- [293] J. W. Grizzle and P. V. Kokotovic, “Feedback linearization of sampled-data systems,” *IEEE Trans. on Automatic Control*, vol. 33, no. 9, pp. 857–859, 1988. DOI: 10.1109/9.1316.
- [294] S. Monaco, D. Normand-Cyrot, and S. Stornelli, “On the linearizing feedback in nonlinear sampled data control schemes,” in *Proc. IEEE 25th Conf. on Decision and Control (CDC)*, Athens, Greece, 1986, pp. 2056–2060. DOI: 10.1109/CDC.1986.267399.
- [295] I. M. Mareels, H. Penfold, and R. J. Evans, “Controlling nonlinear time-varying systems via euler approximations,” *Automatica*, vol. 28, no. 4, pp. 681–696, 1992.
- [296] D. Nešić, A. Teel, and P. Kokotovic, “On the design of a controller based on the discrete-time approximation of the nonlinear plant model,” in *IEEE American Control Conf. (ACC)*, vol. 5, San Diego, CA, USA, 1999, pp. 3474–3478. DOI: 10.1109/ACC.1999.782411.

- [297] D. Nešić and D. S. Laila, “A note on input-to-state stabilization for nonlinear sampled-data systems,” *IEEE Trans. on Automatic Control*, vol. 47, no. 7, pp. 1153–1158, 2002. DOI: 10.1109/TAC.2002.800663.
- [298] D. Nešić and D. Angeli, “Integral versions of iss for sampled-data nonlinear systems via their approximate discrete-time models,” *IEEE Trans. on Automatic Control*, vol. 47, no. 12, pp. 2033–2037, 2002. DOI: 10.1109/TAC.2002.805681.
- [299] J. W. Grizzle, “Feedback linearization of discrete-time systems,” in *Analysis and Optimization of Systems*, A. Bensoussan and J. L. Lions, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 1986, pp. 273–281, ISBN: 978-3-540-39856-1.
- [300] B. Jakubczyk and E. D. Sontag, “The effect of sampling on feedback linearization,” in *Proc. IEEE 26th Conf. on Decision and Control (CDC)*, vol. 26, Los Angeles, CA, USA, 1987, pp. 1374–1379. DOI: 10.1109/CDC.1987.272636.
- [301] A. Arapostathis, B. Jakubczyk, H.-G. Lee, S. I. Marcus, and E. D. Sontag, “The effect of sampling on linear equivalence and feedback linearization,” *Sys. & Control Let.*, vol. 13, no. 5, pp. 373–381, 1989. DOI: 10.1016/0167-6911(89)90103-5.
- [302] J.-P. Barbot, S. Monaco, and D. Normand-Cyrot, “A sampled normal form for feedback linearization,” *Mathematics of Control, Signals and Sys.*, vol. 9, no. 2, pp. 162–188, 1996. DOI: 10.1007/BF01211752.
- [303] S. Monaco and D. Normand-Cyrot, “Zero dynamics of sampled nonlinear systems,” *Sys. & Control Let.*, vol. 11, no. 3, pp. 229–234, 1988. DOI: 10.1016/0167-6911(88)90063-1.
- [304] J. I. Yuz and G. C. Goodwin, “On sampled-data models for nonlinear systems,” *IEEE Trans. on Automatic Control*, vol. 50, no. 10, pp. 1477–1489, 2005. DOI: 10.1109/TAC.2005.856640.
- [305] M. Nishi, M. Ishitobi, and S. Kunimatsu, “Nonlinear sampled-data models and zero dynamics,” in *IEEE Int. Conf. on Networking, Sensing and Control (ICNSC)*, Okayama, Japan, 2009, pp. 373–378. DOI: 10.1109/ICNSC.2009.4919304.
- [306] P. Pepe, “Robustification of nonlinear stabilizers in the sample-and-hold sense,” *J. of The Franklin Institute*, vol. 352, no. 10, pp. 4107–4128, 2015. DOI: 10.1016/j.jfranklin.2015.06.007.
- [307] M. Di Ferdinando, P. Pepe, and S. Di Gennaro, “A new approach to the design of sampled-data dynamic output feedback stabilizers,” *IEEE Trans. on Automatic Control*, vol. 67, no. 2, pp. 1038–1045, 2021. DOI: 10.1109/TAC.2021.3062345.

- [308] A. Ghaffari, I. Abel, D. Ricketts, S. Lerner, and M. Krstić, “Safety verification using barrier certificates with application to double integrator with input saturation and zero-order hold,” in *Proc. IEEE American Control Conf. (ACC)*, Milwaukee, WI, USA, 2018, pp. 4664–4669. DOI: 10.23919/ACC.2018.8430929.
- [309] A. J. Taylor, V. D. Dorobantu, R. K. Cosner, Y. Yue, and A. D. Ames, “Safety of sampled-data systems with control barrier functions via approximate discrete time models,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 7127–7134. DOI: 10.1109/CDC51059.2022.9993226.
- [310] K. J. Åström and B. Bernhardsson, “Comparison of periodic and event based sampling for first-order stochastic systems,” *IFAC Proc. Vol.*, vol. 32, no. 2, pp. 5006–5011, 1999. DOI: 10.1016/S1474-6670(17)56852-4.
- [311] K.-E. Årzén, “A simple event-based pid controller,” *IFAC Proc. Vol.*, vol. 32, no. 2, pp. 8687–8692, 1999. DOI: 10.1016/S1474-6670(17)57482-0.
- [312] J. K. Yook, D. M. Tilbury, and N. R. Soparkar, “Trading computation for bandwidth: Reducing communication in distributed control systems using state estimators,” *IEEE Trans. on Control Sys. Tech.*, vol. 10, no. 4, pp. 503–518, 2002. DOI: 10.1109/TCST.2002.1014671.
- [313] W. P. M. H. Heemels, J. H. Sandee, and P. P. J. Van Den Bosch, “Analysis of event-driven controllers for linear systems,” *Int. J. of Control*, vol. 81, no. 4, pp. 571–590, 2008. DOI: 10.1080/00207170701506919.
- [314] T. Henningson, E. Johannesson, and A. Cervin, “Sporadic event-based control of first-order linear stochastic systems,” *Automatica*, vol. 44, no. 11, pp. 2890–2895, 2008. DOI: 10.1016/j.automatica.2008.03.026.
- [315] R. Postoyan, P. Tabuada, D. Nešić, and A. Anta, “A framework for the event-triggered stabilization of nonlinear systems,” *IEEE Trans. on Automatic Control*, vol. 60, no. 4, pp. 982–996, 2014. DOI: 10.1109/TAC.2014.2363603.
- [316] C. Nowzari, E. Garcia, and J. Cortés, “Event-triggered communication and control of networked systems for multi-agent consensus,” *Automatica*, vol. 105, pp. 1–27, 2019. DOI: 10.1016/j.automatica.2019.03.009.
- [317] J. Berneburg and C. Nowzari, “Distributed dynamic event-triggered coordination with a designable minimum inter-event time,” in *Proc. IEEE American Control Conf. (ACC)*, Philadelphia, PA, USA, 2019, pp. 1424–1429. DOI: 10.23919/ACC.2019.8815227.

- [318] R. Postoyan, R. Sanfelice, and W. P. M. H. Heemels, “Inter-event times analysis for planar linear event-triggered controlled systems,” in *Proc. IEEE 58th Conf. on Decision and Control (CDC)*, Nice, France, 2019, pp. 1662–1667. DOI: 10.1109/CDC40024.2019.9028888.
- [319] M. C. F. Donkers and W. P. M. H. Heemels, “Output-based event-triggered control with guaranteed \mathcal{L}_∞ -gain and improved event-triggering,” in *Proc. IEEE 49th Conf. on Decision and Control (CDC)*, Atlanta, GA, USA, 2010, pp. 3246–3251. DOI: 10.1109/CDC.2010.5718032.
- [320] K. P. Tee, S. S. Ge, and E. H. Tay, “Barrier lyapunov functions for the control of output-constrained nonlinear systems,” *Automatica*, vol. 45, no. 4, pp. 918–927, 2009. DOI: 10.1016/j.automatica.2008.11.017.
- [321] A. J. Taylor, P. Ong, J. Cortés, and A. D. Ames, “Safety-critical event triggered control via input-to-state safe barrier functions,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 749–754, 2021. DOI: 10.1109/LCSYS.2020.3005101.
- [322] L. Long and J. Wang, “Safety-critical dynamic event-triggered control of nonlinear systems,” *Sys. & Control Let.*, vol. 162, pp. 1–9, 2022. DOI: 10.1016/j.sysconle.2022.105176.
- [323] W. Xiao, C. Belta, and C. G. Cassandras, “Event-triggered safety-critical control for systems with unknown dynamics,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 540–545. DOI: 10.1109/CDC45484.2021.9682792.
- [324] M. Kishida, “Greedy synthesis of event-and self-triggered controls with control lyapunov-barrier function,” *arXiv preprint arXiv:2302.12435*, 2023. DOI: 10.48550/arXiv.2302.12435.
- [325] E. Sabouni, C. G. Cassandras, W. Xiao, and N. Meskin, “Optimal control of connected automated vehicles with event/self-triggered control barrier functions,” *arXiv preprint arXiv:2209.13053*, 2022. DOI: 10.48550/arXiv.2209.13053.
- [326] S. Koga, C. Demir, and M. Krstic, “Event-triggered safe stabilizing boundary control for the stefan pde system with actuator dynamics,” *arXiv preprint arXiv:2210.01454*, 2022. DOI: 10.48550/arXiv.2210.01454.
- [327] S. A. Al-Hiddabi, “Quadrotor control using feedback linearization with dynamic extension,” in *IEEE 6th Int. Sym. on Mechatronics and its App. (ISMA)*, Sharjah, United Arab Emirates, 2009, pp. 1–3. DOI: 10.1109/ISMA.2009.5164788.
- [328] P. Tabuada and L. Fraile, “Data-driven stabilization of siso feedback linearizable systems,” *arXiv preprint arXiv:2003.14240*, 2020. DOI: 10.48550/arXiv.2003.14240.

- [329] F. H. Clarke, Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski, *Nonsmooth analysis and control theory*. Springer Science & Business Media, 2008, vol. 178, ISBN: 0-387-98336-8.
- [330] F. Solowjow and S. Trimpe, “Event-triggered learning,” *Automatica*, vol. 117, pp. 1–10, 2020. DOI: 10.1016/j.automatica.2020.109009.
- [331] W. Tan and A. Packard, “Searching for control lyapunov functions using sums of squares programming,” *sibi*, vol. 1, no. 1, [Online]. Available: https://bcc.berkeley.edu/downloads/papers/weehong_3.pdf.
- [332] L. Wang, D. Han, and M. Egerstedt, “Permissive barrier certificates for safe stabilization using sum-of-squares,” in *Proc. IEEE American Control Conf. (ACC)*, Milwaukee, WI, USA, 2018, pp. 585–590. DOI: 10.23919/ACC.2018.8431617.
- [333] A. Clark, “Verification and synthesis of control barrier functions,” in *Proc. IEEE 60th Conf. on Decision and Control (CDC)*, Austin, TX, USA, 2021, pp. 6105–6112. DOI: 10.1109/CDC45484.2021.9683520.
- [334] Z. Gong, M. Zhao, T. Bewley, and S. Herbert, “Constructing control lyapunov-value functions using hamilton-jacobi reachability analysis,” *IEEE Control Sys. Let.*, vol. 7, pp. 925–930, 2022. DOI: 10.1109/LCSYS.2022.3228728.
- [335] S. Tonkens and S. Herbert, “Refining control barrier functions through hamilton-jacobi reachability,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Sys. (IROS)*, Kyoto, Japan, 2022, pp. 13 355–13 362. DOI: 10.1109/IROS47612.2022.9982203.
- [336] N. Csomay-Shanklin, A. J. Taylor, U. Rosolia, and A. D. Ames, “Multi-rate planning and control of uncertain nonlinear systems: Model predictive control and control lyapunov functions,” in *Proc. IEEE 61st Conf. on Decision and Control (CDC)*, Cancún, Mexico, 2022, pp. 3732–3739. DOI: 10.1109/CDC51059.2022.9992902.
- [337] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin, “Fastrack: A modular framework for fast and guaranteed safe motion planning,” in *Proc. IEEE 56th Conf. on Decision and Control (CDC)*, Melbourne, VIC, Australia, 2017, pp. 1517–1522. DOI: 10.1109/CDC.2017.8263867.
- [338] U. Rosolia and A. D. Ames, “Multi-rate control design leveraging control barrier functions and model predictive control policies,” *IEEE Control Sys. Let.*, vol. 5, no. 3, pp. 1007–1012, 2021. DOI: 10.1109/LCSYS.2020.3008326.

- [339] U. Rosolia, A. Singletary, and A. D. Ames, “Unified multirate control: From low-level actuation to high-level planning,” *IEEE Trans. on Automatic Control*, vol. 67, no. 12, pp. 6627–6640, 2022. DOI: 10.1109/TAC.2022.3184664.

UNMATCHED TUNABLE INPUT-TO-STATE SAFETY PROOF

In this appendix I will provide a proof of the unmatched disturbance case of Theorem 29 for completeness. Recall Theorem 29 as it appears in Section 4.3.

Theorem 29. *Let $\mathcal{C} \subset E$ be the 0-superlevel set of a function $h : E \rightarrow \mathbb{R}$ that is continuously differentiable on E . If h is a TISSf-BF for the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) on \mathcal{C} , then the closed-loop system with unmatched disturbances (2.70) (matched disturbances (2.74)) is TISSf with respect to the set \mathcal{C} with a function $\gamma_T : \mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ defined as:*

$$\gamma_T(r, d) \triangleq -\alpha^{-1} \left(-\frac{\epsilon(r)d^2}{4} \right). \quad (\text{A.1})$$

Proof. Observe that by assumption α^{-1} and ϵ are continuously differentiable on \mathbb{R} , such that the function $\gamma_T(\cdot, b)$ is continuously differentiable for any $b \in \mathbb{R}_{\geq 0}$. Similarly, as $\epsilon(r) > 0$ for all $r \in \mathbb{R}$ and $\alpha^{-1} \in \mathcal{K}_\infty^e$, the function $\gamma_T(a, \cdot) \in \mathcal{K}_\infty^e$ for any $a \in \mathbb{R}$ as required. Our next goal is to verify that for all $d \in \mathbb{R}_{\geq 0}$, the set $\mathcal{C}_{d,T}$ defined in (4.12) is forward invariant up to d . Consider a value of $d \in \mathbb{R}_{\geq 0}$ and consider a disturbance signal $\mathbf{d} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ that is piecewise continuous on $\mathbb{R}_{\geq 0}$ and satisfies $\|\mathbf{d}\|_\infty \leq d$. As h is a TISSf-BF for the closed-loop system with unmatched disturbances (2.74), we have that:

$$\dot{h}(\mathbf{x}, t) \triangleq L_f h(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{k}(\mathbf{x}) + \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{d}(t), \quad (\text{A.2})$$

$$\geq -\alpha(h(\mathbf{x})) + \frac{1}{\epsilon(h(\mathbf{x}))} \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\|^2 + \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{d}(t). \quad (\text{A.3})$$

Noting that:

$$\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \mathbf{d}(t) \geq - \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\| \|\mathbf{d}\|_\infty \geq - \left\| \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \right\| d, \quad (\text{A.4})$$

for all $\mathbf{x} \in E$ and $t \in \mathbb{R}_{\geq 0}$, and $\epsilon(h(\mathbf{x})) > 0$ for all $\mathbf{x} \in E$, adding and subtracting $\frac{\epsilon(h(\mathbf{x}))d^2}{4}$, and completing the squares yields:

$$\dot{h}(\mathbf{x}, t) \geq -\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4}. \quad (\text{A.5})$$

Next, observe that:

$$\frac{\partial \gamma_{\mathbb{T}}}{\partial r}(r, d) = \frac{d^2}{4} \frac{\partial \alpha^{-1}}{\partial r} \left(-\frac{\epsilon(r)d^2}{4} \right) \frac{\partial \epsilon}{\partial r}(r) \geq 0, \quad (\text{A.6})$$

for all $r \in \mathbb{R}$. Taking the time derivative of the function $h_{\mathbb{T}}$ defined in (4.11) yields:

$$\dot{h}_{\mathbb{T}}(\mathbf{x}, d, t) = \left(1 + \frac{\partial \gamma_{\mathbb{T}}}{\partial r}(h(\mathbf{x}), d) \right) \dot{h}(\mathbf{x}, t). \quad (\text{A.7})$$

noting that d is a fixed constant in this proof. By (A.6), we have that:

$$1 + \frac{\partial \gamma_{\mathbb{T}}}{\partial r}(h(\mathbf{x}), d) > 0, \quad (\text{A.8})$$

for all $\mathbf{x} \in E$. Substituting (A.5) into (A.7), we obtain:

$$\dot{h}_{\mathbb{T}}(\mathbf{x}, d, t) \geq \left(1 + \frac{\partial \gamma_{\mathbb{T}}}{\partial r}(h(\mathbf{x}), d) \right) \left(-\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4} \right). \quad (\text{A.9})$$

Next, consider a state $\mathbf{x} \in \partial \mathcal{C}_{d, \mathbb{T}}$, such that $h_{\mathbb{T}}(\mathbf{x}, d) = 0$, for which (4.11) and (A.1) imply:

$$-\alpha(h(\mathbf{x})) - \frac{\epsilon(h(\mathbf{x}))d^2}{4} = 0, \quad (\text{A.10})$$

yielding:

$$\dot{h}_{\mathbb{T}}(\mathbf{x}, d, t) \geq 0, \quad (\text{A.11})$$

for all $t \in \mathbb{R}_{\geq 0}$. In the case that $d = 0$, such that $h_{\mathbb{T}}(\mathbf{x}, 0) = h(\mathbf{x})$, we have that $h_{\mathbb{T}}(\mathbf{x}, 0) = 0$ implies $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$. Considering $d > 0$, because $\gamma(a, \cdot) \in \mathcal{K}_{\infty}^e$ for all $a \in \mathbb{R}$, when $h_{\mathbb{T}}(\mathbf{x}, d) = 0$, we have $-\alpha(h(\mathbf{x})) > 0$. Thus, the inequality in (4.13) requires that $\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n$ for $\mathbf{x} \in \partial \mathcal{C}_{d, \mathbb{T}}$. Finally, we have:

$$\frac{\partial h_{\mathbb{T}}}{\partial \mathbf{x}}(\mathbf{x}, d) = \underbrace{\left(1 + \frac{\partial \gamma_{\mathbb{T}}}{\partial r}(h(\mathbf{x}), d) \right)}_{>0} \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq \mathbf{0}_n, \quad (\text{A.12})$$

using (A.8). Therefore, Nagumo's theorem [104], [123] implies the set $\mathcal{C}_{d, \mathbb{T}}$ is forward invariant as $h_{\mathbb{T}}(\mathbf{x}, d) = 0$ implies $\dot{h}_{\mathbb{T}}(\mathbf{x}, d, t) \geq 0$ for all $t \in \mathbb{R}_{\geq 0}$, and $\frac{\partial h_{\mathbb{T}}}{\partial \mathbf{x}}(\mathbf{x}, d) \neq \mathbf{0}_n$. \square