Multicellular Control

Thesis by
Anish Anandsai Sarma

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2023
Defended June 17, 2022

Anish Anandsai Sarma
ORCID: 0000-0003-1261-0589

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support of many, many people. I have been lucky to have two remarkable advisors in John Doyle and Lea Goentoro. John took me on as his advisee and encouraged me to pursue an idea that, at the time, I could barely articulate. Lea welcomed me into her lab a few years later, encouraged me to think both creatively and rigorously, and trusted me to learn my way around as an experimentalist. I have learned so much about math and biology, especially a kind of optimism in the face of complexity, from John and Lea.

Many accomplished scientists, clinicians, mathematicians, and engineers have taught me and challenged me in the past six years, including my rotation advisors, committee members, and expert collaborators: David Brandman, Marie Csete, Michael Elowitz, Peter Lee, and Nikolai Matni. I owe special thanks to Ellen Rothenberg and Doris Tsao, who believed in this project early on.

From my time before Caltech, Leigh Hochberg, John Simeral, and Wilson Truccolo taught me how to build things in the world that help people directly. The participants in the BrainGate clinical trial, the handful of individuals who dedicated the last years of their lives to the development of science and technology from which they would never personally benefit, are never far from my mind.

Among my fellow graduate students and research staff there are too many friends and intellectual sparring partners to name, and I will likely leave many out. In the Doyle Group, Lisa Li and Fangzhou Xiao have been my frequent collaborators and have brought out the best in my work. Carmen Amo Alonso, James Anderson, Natalie Bernat, Jiexin Chen, Lauren Conger, Dimitar Ho, Mandy Huo, Olle Kjellqvist, Quanying Liu, Yorie Nakahira, Noah Olsman, Shih-Hao Tseng, and Jing Yu have been friends, critics, and co-creators of our group's wonderful, insular vernacular. In the Goentoro Lab, I owe gratitude to the people who taught me how to handle a pipette and a jellyfish, especially Michael Abrams, Ty Basinger, and Chris Frick; brainstorming sessions between experiments with Mengsha Gong, Martin Heithe, Aki Ohdera, and Fayth Tan, and later with Zevin Condiotte, Yutian Li, and Changhua Yu, were some of the most energetic, thrilling, and often comic moments of the past six years. My cohort-mates in the Computation and Neural Systems option, Jon Kenny, Matt Rosenberg, Tony Zhang, and Jeremy Bernstein, have been there from the very beginning, and we all pushed each other to pursue work we believed in,

wherever it led. I'm grateful to Gautam Goel, Juba Ziani, and Mason McGill, for always making time for wide-ranging conversations, informal jam sessions, and whatever else the moment called for.

My parents, Radha and Nagaraj Sarma, have been my constant supporters. They have read every scrap of science associated with this thesis and could teach a class on it. My sister-in-law, Rebecca Jacobs, became an integral part of the family in the years since I started at Caltech, and I have relied on her insight. My brother, Aartik Sarma, has been my longtime collaborator, my trusted friend, and my constant sounding board through every technical and personal challenge. Lastly, I write this in gratitude to my dear grandparents, whether here with us or with me in other ways.

ABSTRACT

Robust control theory was developed in the late twentieth century as a mathematical framework to enable the principled incorporation of uncertainty into engineering design in applications like aerospace. However, engineered technologies that interface with living systems in applications like medicine and ecology must accommodate uncertainties and unmodeled dynamics far beyond what robust control theory has historically achieved. This thesis develops a robust control foundation for overcoming large-scale uncertainty and designing interfaces with living systems, through formal theory and three case studies: neural control of movement, immune control of viruses, and homeostatic control of neoplasia in the moon jellyfish.

The central argument of this thesis is that these three systems, along with many others, have two key properties that enable new approaches to the uncertainty intrinsic to their study: they are themselves control systems, and they are multicellular systems. These properties motivate new work in control theory, blending recent results in localized and distributed control with older results from robust and modern control. The resulting theory framework answers domain-specific questions, guides the design of new experiments and technologies, and enables a conceptual synthesis. By leveraging the fact that these are *multicellular control* systems, we are able to make progress in theory, basic science, and engineering.

After an introduction in Chapter 1, I provide a high-level overview of the theory framework in Chapter 2, which is later elaborated in technical detail in Chapter 5 and Chapter 6. In each of three domain-specific investigations, I consider a gap between a classical theoretical paradigm and a more complicated recent body of evidence.

In Chapter 3, I consider the problem of localization of function in the neural control of movement. If, as is typically believed, different parts of the brain are specialized for different functions (for instance, vision and movement), why do detailed studies consistently find internal feedback signals between putatively specialized brain areas (such as movement-related signals in vision-related areas)? Why is neuronal activity in some brain areas explained primarily by external stimuli, whereas activity in other areas is explained primarily by autonomous internal dynamics? I show that internal feedback and internal dynamics do not contradict the old notion of localization of function, but in fact are necessary to achieve it.

In Chapter 4, I consider the evolution of pathogens in the context of viral-immune interactions. The theory that pathogens face a tradeoff between virulence and transmission been complicated by key examples, including SARS-CoV-2. I show that the classical virulence-transmission theory can be strengthened by incorporating models of immune variation between hosts into models of virulence and transmission, allowing us to identify potential future risks and potential interventions.

Chapter 7 provides an experimental test of the overall approach. I consider the regeneration-tradeoff model of cancer formation. I combine mathematical modeling with wet-lab experiments in the moon jellyfish *Aurelia coerulea*, an animal in which neoplasia has never been observed and in which a regenerative state can be induced or suppressed in the laboratory. I first use a theory-driven experiment to inhibit cancer resistance mechanisms in the moon jellyfish and thus induce neoplasia. I then propose a new model of cancer formation, the *flux* model, which makes the surprising prediction that the regenerative state in the jellyfish will be protective against neoplasia. I verify this prediction experimentally, inducing and then reversing a neoplasia-susceptible state.

PUBLISHED CONTENT AND CONTRIBUTIONS

Published journal papers and refereed conference papers:

[1] Michael J. Abrams*, Fayth Hui Tan*, Yutian Li*, Ty Basinger, Martin L. Heithe, Anish A. Sarma, Iris T. Lee, Zevin J. Condiotte, Misha Raffiee, John O. Dabiri, David A. Gold, and Lea Goentoro. A conserved strategy for inducing appendage regeneration in moon jellyfish, drosophila, and mice. *eLife*. doi: 10.7554/eLife.65092.
A.A.S. performed single-fly tracking experiments and imaging, and proposed and performed imaging analysis and statistical analysis.

[2] Nikolai Matni and Anish A. Sarma. Robust performance guarantees for system level synthesis. *American Control Conference*, 2020. doi: 10.23919/ACC45564.2020.9147436.
A.A.S. provided simulations and contributed to the text.

[3] Anish A. Sarma and John C. Doyle. Flexibility and cost-sensitivity in a quantized control loop. *American Control Conference*, 2019. doi: 10.23919/ACC.2019.8814926.
A.A.S. proposed the paper, proved the main theorem, and wrote the paper.

[4] Anish A. Sarma, Jing Shuang Li, Josefin Stenberg, Gwyneth Card, Elizabeth S. Heckscher, Narayan Kasthuri, Terrence Sejnowski, and John C. Doyle. Internal feedback in biological control: architectures and examples. *American Control Conference*. doi: arXiv:2110.05029.
A.A.S. proposed and wrote the paper.

[5] Josefin Stenberg, Jing Shuang Li, Anish A. Sarma, and John C. Doyle. Internal feedback in biological control: Diversity, delays, and standard theory. *American Control Conference*. doi: arXiv:2109.11752.
A.A.S. co-supervised the development and writing of the paper.

Some language and figures in Chapter 3 were contributed by Jing-Shuang (Lisa) Li, and in Chapter 6 by Nikolai Matni. The following chapters include IEEE copyrighted material, reprinted with permission: Chapter 5 (©2019), Chapter 6 (©2020), and Chapter 3 (©2021).

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

Chapter 1

INTRODUCTION

1.1 Goals of the thesis

This thesis distills a body of work in control theory under several guises: applied mathematics, electrical engineering, biomedical engineering, neuroscience, immunology, and developmental biology.[1] It includes theory, experiment, and engineering. This breadth of domains and techniques serves two goals, which happily coincide.

The first goal is the scientist and mathematician's goal of unification. This thesis has been animated by the belief that there exists, just out of view, a unified subject that meaningfully spans (but does not subsume) the above domains and others. Like any engineering subject, this new subject is at once theoretical, experimental, and applied; unlike classical engineering subjects, it aims to discover new phenomena in the natural world. Informally, this subject is the study of how flexible, coordinated action is achieved by simple, constrained components through a special composition we call *layering*, and what happens when components are lost, added, or modified. These issues often arise in the context of multicellular biological systems: how neurons work together to achieve sophisticated movements, how immune cells identify dangerous molecular patterns, or how developmental processes form and maintain patterned tissues like limbs.

Control theory is a natural home for these questions, providing a conceptual framework to study interconnected systems at scale, although we will travel some distance from what one would find in an undergraduate control textbook.[2]

The second goal is the engineer's goal of building a more humane world. This thesis has been animated by optimism about past and future triumphs over the material realities of the human condition: electrification, wastewater management, vaccines,

[1] Developmental biology is a broad umbrella, colloquially understood to cover embryonic and childhood development but also often used to include stem cell biology, regeneration biology, and some aspects of cancer biology. These are not identical systems, although there is considerable overlap in relevant pathways and processes, and I use *development* for brevity.

[2] Many people have noticed this subject and its connections to control. This type of observation was the foundation of the cybernetics movement in the 1940s and 1950s, the complex systems movement of the 1980s and 1990s, and more recent efforts in networks, active inference, and related theories.

antibiotics, agriculture, air filters, parks, libraries. Many contemporary engineering challenges, from medicine to ecology to infrastructure, are challenges not only of designing technologies but of interfacing with extant engineered or living systems. Layering control systems, including the layering of new systems, experiments, or decision-making processes onto legacy or evolved systems, is a general problem across many fields, with broad implications for the built and natural world.

In practice, skilled engineers have always built on top of legacy systems, and have often accounted for living systems (at the very least, human users) in the loop. Engineering skill often requires the practitioner to leverage the intuition from exact theory while accommodating real-world unknowns. This is where engineering, as any skilled practitioner will tell you, is both art and science. Skilled experimental biologists and clinicians will say similar. Nevertheless, a foundation in formal engineering theory is essential to high-performance and robust engineering design, and is increasingly useful to experimental design, even when intuition fills in gaps. Because my work has included engineering and experiment as well as theory, I am familiar not only with the limitations of existing theoretical assumptions but also with the limitations of existing theory-based intuition.[3] It was really this intuition that was most important to fortify with this work, but to do so, it was necessary to extend the formal theory as well.

The crucial technical engine of this work is **convex set descriptions** of all possible control systems that solve particular control problems – the profound fact, well known in control theory, that all solutions of particular types of problems necessarily look alike at the appropriate level of abstraction. Convex description allows us to access computational and theoretical techniques in optimization and machine learning, and therefore computerized analysis and design of large-scale systems. This thesis improves the state of the art of convex set descriptions of control systems, enabling new types of designs and faster algorithms for previous controller designs.

Although the theorems and concepts apply broadly, most of this thesis will consider the (still broad) class of multicellular biological systems. This level of description intersects with neuroscience, immunology, and developmental biology, and contrasts with a long history in biology and medicine that revolves around identifying single molecular mechanisms of disease and single drug targets – an approach that increasingly shows its limitations in responses to infectious disease, neurodegener-

[3] As I argue in Chapter 3, scientists often reason implicitly from control-like diagrams and structures that miss key details; updating the diagrammatic style by updating the math the diagrams represent is a major goal of this work.

ation, autoimmunity, and cancer. A multicellular control approach builds on work in systems biology that has moved towards multi-molecule descriptions, as well as work in systems neuroscience that has moved towards multicellular descriptions.

Three new conceptual and technical challenges are encountered in the leap to multicellularity: **locality**, **uncertainty**, and **self-renewal**. Each of these control challenges will be addressed in a case study that considers a particular biological system, a particular problem within that domain, and a particular control-theoretic technique. In the interest of maintaining focus, I will limit references to other biological systems within each case. However, each control challenge and each control technique is applicable across each of the relevant domains.

- **Locality (Chapter 3):** when moving signals around a system is slow, inaccurate, or expensive, it becomes necessary for control actions to be taken without complete information about the global state of the system. We study this problem in the **neural control of movement** by studying **implementations** of decentralized controllers.

- **Uncertainty (Chapter 4):** large-scale systems can have millions or billions of interacting components with dynamical parameters that can never be fully observed or predicted, yet we still want to design interpretable models, experiments, and reliable interventions. We study this problem in **viral-immune interactions** by studying **performance limits of controllers**.

- **Self-renewal (Chapter 7):** cells in biology are continually created from other cells, unlike physical components in electronic or mechanical control systems, which creates new constraints on achievable control – but also creates new opportunities for intervention that conventional engineering might overlook. We study this problem in a **cancer-resistant jellyfish** by developing **theory-driven experiments**.

Between the discussions of uncertainty and self-renewal, I will develop in greater detail new theorems in the convex set description of robust and localized control systems. These theorems are essential to the scholarly and technical contribution of this thesis and are presented in a denser technical style consistent with conventions in control theory and maybe necessary to study control systems in greater generality. By contrast, the mathematical exposition in each of the domain-oriented chapters is intended to be self-contained.

The theoretical work in total identifies and addresses a conceptual and technical gap that recurs across several fields. This is surprising, but perhaps not so strange when we consider that statistics, differential equation modeling, and machine learning have become core technologies across science and engineering. Control theory as we develop it here is likewise broadly applicable – a claim that I will justify by applying control theory broadly. Rather than hermetically developing a control theory with claims to generality, we will aim to answer old questions and pose new questions in other domains. These applications necessitate new theorems, but the theorems themselves are not the core contribution of this thesis. No single control theorem accounts for the subtleties of domain details, and every historical application of control theory, from aerospace to robotics, has included a mix of theoretical rigor and practitioner intuition. It is this way of thinking about multicellular biology that I hope this thesis provides: that the control theorist sees a new breadth of theoretical challenges and applications in multicellular systems, whether the systems we study here or other systems, and that the biologist sees new ways to use control theory, whether using the new tools we develop here or classical tools applied in new ways. There is an old joke from the poet Archilochus, by way of Isaiah Berlin, that applies here: a fox knows many things, but a hedgehog knows one big thing. This thesis looks like a fox, but it is a hedgehog.

1.2 The diagrammatic and mathematical language of control theory

Control theorists describe interconnected systems using a diagrammatic language of blocks and arrows. While control-theoretic block diagrams are rarely ambiguous within the established conventions of the field, they encompass subtleties that might differ from cosmetically similar diagrammatic languages used in other fields. Therefore, I will present here a brief introduction to the diagrammatic language of control theory as we use it, emphasizing three key ideas: **implementation**, **uncertainty**, and **delay**.

In general, a block in a control diagram describes a *transfer function*. The transfer function G between an input u and an output y is an operation on u, a vector-valued signal through time, that produces y, a vector-valued signal through time.

We can write this mathematically as:

$$y = G(u) \tag{1.1}$$

Typically, unless otherwise noted, G is *causal*, meaning that at any given time, the

Figure 1.1: A transfer function from u to y.

present value of y depends only on past values of u and y; the future does not affect the past.

Because G can depend on past values of both u and y, we can describe the *internal states* x of G as a dynamical system with a differential equation with respect to time t. We are almost always interested in derivatives with respect to time, so we will typically define $\dot{x} = \frac{dx}{dt}$.

$$\begin{aligned} \dot{x} &= f(x,t) + g(x,t)u \\ y &= h(x) \end{aligned} \tag{1.2}$$

We have now assigned to the simple block in Figure 1.2 two meanings, expressed as Equation 1.1 and Equation 1.2. These meanings are compatible, but not equivalent. Equation 1.1 expresses the *input-output* behavior of the system, while Equation 1.2 gives us one *realization* of the input-output behavior. A given realization corresponds to exactly one input-output, but a given input-output can have multiple realizations (infinitely many, though some are improbable). We make a further distinction in this work between realization and *implementation*, which is the distinction between a differential equation model of a system and a more complete characterization of the physical system that the differential equation model describes.[4] For example, the elementary mathematical operation of addition in Equation 1.2 might be implemented as the flows of water in and out of a tank. It might be suitable for some problems to treat addition and subtraction of water as positively and negatively signed values of the same term u. To build the system, or to diagnose problems in the system, it might be necessary to treat addition and subtraction separately. This

[4]These distinctions are reminiscent of the Marr's levels in neuroscience. Marr separated the computational, algorithmic, and implementational levels of analysis, which in control-theoretic language would be the input-output, realization, and implementation, respectively. To the reader familiar with Marr's levels, this work can be understood as creating a foundation by which to constrain the implementational level with the computational level and vice versa. In control problems, the algorithmic/realization level is sometimes a necessary technical intermediate, but it is not conceptually distinct from implementation.

distinction between realization and implementation is nonstandard even in control theory.

Input-output descriptions are a framework for organizing facts that we already know and highlighting facts that we do not know; they are not models or hypotheses in the typical scientific sense, nor designs in the typical engineering sense. However, in the particular capacity to characterize *all possible* models, hypotheses, and designs that result from given assumptions and data, input-output descriptions far exceed implementational descriptions, and facilitate the (intuitive or systematic) generation of testable implementational descriptions. Throughout this thesis, we will use input-output and implementational descriptions in tandem, for instance by characterizing one part of a larger system in implementational detail while subsuming other parts of the system into input-output blocks.

Four special types of blocks merit particular attention. One is a *controller*. Controllers are transfer functions, with at least one corresponding implementation, that are designed by the scientist or engineer to test a hypothesis or achieve an engineering goal. Controllers are contrasted with *plants*, blocks that represent extant transfer functions in the natural or technological world. The distinction between controller and plant, between what is designed and what is extant, is somewhat arbitrary, depending on the question being asked. A third special type of block that can be interconnected with other blocks is a *delay* block. While these look the same as any other block, they simply pass a signal forward untransformed after a time delay. Lastly, a special type of block that can be interconnected with other blocks is an *uncertainty* block. Rather than representing specific transfer functions, uncertainty blocks represent *bounded sets* of functions; a block diagram with an uncertainty block in it (often expressed as a Δ) should be understood to represent several possible functions, rather than just one. This set-based approach is important when we want to understand whether a model of a system is any good: if our decisions or conclusions about the system are narrowly dependent on the particular parametric assumptions of a single model, we call the model (or system) *fragile*. If our decisions or conclusions do not depend on particular parametric assumptions, we call the model (or system) *robust*. In general, even systems that are robust to some assumptions are fragile to others.[5]

[5]The reader may wonder here if we have gone beyond technical mathematical claims into a realm of epistemological claims. I admit that we have, but that this is we are no different than any use of mathematical tools in statistics, differential equation modeling, or machine learning in science or engineering. These tools are compatible with everything described here.

A diagrammatic style that includes uncertainty and delays, spanning input-output and implementation, is a central concern of this thesis, both technically and conceptually. We develop new tools to analyze systems described by such diagrams and apply this diagrammatic style to multicellular systems. The diagrammatic descriptions of multicellular systems in subsequent chapters are more than cartoons; they are arguments, associated with particular mathematical structures.

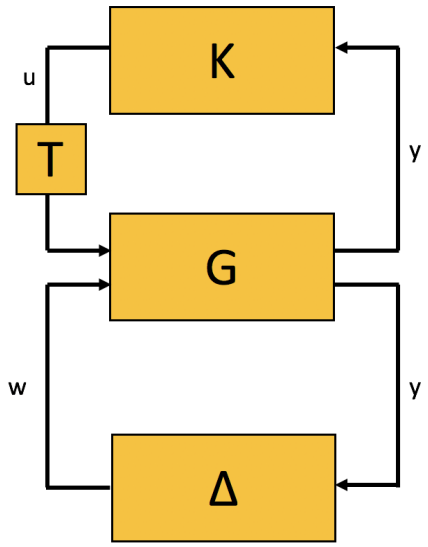

Figure 1.2: Special blocks in a controller-plant interconnection: a controller (K), a delay block (T) and an uncertainty block (Δ).

Chapter 2

THE THEORY AT A HIGH LEVEL

Suppose we observe a dynamical system that can be described approximately as:

$$x(t+1) = A(t)x(t) + B(t)u(t) + w(t) \tag{2.1}$$

x will typically be a small but important part of a large system, like the position of a limb during movement, the viral load in the lung during respiratory infection, or the number of pre-cancerous cells in the body during oncogenesis. We are interested in two kinds of problems: *discovery* problems, where we want to understand how this variable is controlled by extant biological control systems that generate actions u in response to perturbations w, and *design* problems, where we want to control the variable ourselves by providing actions u in response to perturbations w. We may know A and B exactly, or we may know that A and B belong to some uncertainty class corresponding to a set of nonnegative diagonal matrices $\mathcal{D}$. The core systems-theoretic result of this thesis, enabling the domain-specific work, is a theorem stating that the set of systems Φ that can generate ***either* the observed behavior *or* the desired behavior** can be characterized exactly by the following affine conditions:

$$\begin{gathered}\begin{bmatrix} zI - A & -B \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = D \\ \Phi_x, \Phi_u \in \frac{1}{z}\mathcal{RH}_\infty \text{ (stable and strictly causal)}, \; D \in \mathcal{D}\end{gathered} \tag{2.2}$$

Moreover, as long as such a Φ exists, at least one controller implementation is guaranteed to exist and preserve structural constraints on Φ. These conditions are explained in detail in Chapter 6. Variations are also given for different versions of this problem, such as the problems where the control system has partial observations or internally lossy communication. These new theorems belong to a family of *good systems theorems*, mathematical results that characterize all solutions to a class of systems-theoretic problems. Such results are at the heart of control theory [3–6]. The conditions 2.2 generalize the seminal results in [7], and recapitulate them exactly when $D = I$. The precise statement of the theorems is somewhat less important than the way they are used: while this is not the first work to connect design and discovery in this way, we provide a framework that makes this connection actionable in facing the major uncertainties that arise in multicellular systems.

In the following chapter, we consider the first of the three case studies in this thesis: the neural control of movement. This theoretical work is motivated by a question in neuroscience. The approach we take in this chapter is to consider the simplest possible tasks for which complex control systems are required, and the answers turn out to be broadly applicable in multicellular systems.

Chapter 3

LAYERED ARCHITECTURES AND INTERNAL FEEDBACK IN THE NEURAL CONTROL OF MOVEMENT

3.1 Introduction

Feedback control is an essential strategy for both engineered and biological systems to achieve reliable movement in unpredictable environments [8]. Optimal and robust control theory, which provide a general mathematical foundation to study feedback systems, have been used successfully to explain behavior-level observations in the sensorimotor domain by modeling the sensorimotor system as a single control loop from sensing to action and then feedback (through the environment) to sensing [9–11]. In such models, a sensorimotor system senses the environment, communicates signals from sensors through the body, computes actions, and then acts on the environment, creating a single unidirectional loop.

The assumptions of single-loop feedback are also implicit in interpretations of sensory and motor computations in neuronal populations. Consider the familiar canonical model of localized function in the primate visuomotor cortical pathway: a visual signal is captured on the retina, then travels to the lateral geniculate nucleus (LGN), then to the primary visual cortex (V1), progressing through successive transformations until it reaches the primary motor cortex (M1), the spinal cord, and ultimately the muscles. While intuitive, this model neglects several well-known features of sensorimotor processing that have made it increasingly untenable.

Internal feedback, signals that do not flow from sensing towards action, have an unclear role in visuomotor processing. Note that we can further subdivide internal feedback into *local*, *counterdirectional*, and *interareal* internal feedback, representing the anatomy of the feedback signal. For example, internal dynamics within a particular cortical area would be considered local, while signal flows from action to sensing would be considered counterdirectional. Interareal is a catch-all for internal feedback which is not local but not obviously counterdirectional.

Additional supporting concepts are *layering*, visuomotor processing that involves multiple complete sensing-to-action control loops; *diversity*, cellular and molecular components of the control loop that exhibit heterogeneity in properties like speed, accuracy, flexibility, and energetic cost; and *locality*, neuron-to-neuron computation

and communication that are spatially constrained, such that achieving fast and accurate computation and communication over large distances is difficult or impossible.

Internal feedback, especially counterdirectional internal feedback, is the main subject of this paper. One reason that the single-loop model has endured is that it offers a set of tools, including off-the-shelf tools from control theory, by which to reason about subsystems in isolation. However, these subsystems are not isolated, and internal feedback is ubiquitous. The eye is itself a site of computation and control: the eye moves around, sensing different parts of the visual scene, and retinal cells adapt and identify features of interest [12, 13]. Neurons project from LGN to V1 to V2 but also in comparable numbers from V2 to V1 and even greater numbers from V1 to LGN [14, 15], in addition to interneurons from V1 to V1 and projections from motor areas to visual areas – and these varied pathways include substantial heterogeneity in morphology, myelination, and synapse kinetics [16–20]. The cortical visual processing stream diverges into a ventral stream typically associated with object identification and a dorsal stream typically associated with object movement, with cross-talk between the two. In M1, a single-loop model is complicated still further; neuronal firing is better-explained by past firing than by present stimulus or movements, suggesting autonomous dynamics [21], and signals related to movements in the whole body are found in areas typically associated with particular parts of the body, like the hand [22, 23]. Although not typically studied together, we argue that all of these can be considered cases of internal feedback.

Internal feedback has been studied in neuroscience under several theoretical frameworks, notably predictive coding. Here, we build on the foundation of recent work in control theory [1, 24–29] and show that internal feedback is a natural consequence of component diversity in the presence of locality constraints. Our contribution is to present idealized systems for which we can prove mathematically that internal feedback is not merely plausible, but in fact necessary for function, and to explain why necessity arises under biological constraints. Along the way, we also synthesize several observations that have not previously been treated in a predictive coding framework. Internal feedback has at least three functions in our analysis: **estimation**, **localization of function**, and **attention**. Together, these functions facilitate fast control without loss of accuracy or flexibility. Lastly, we observe that this type of internal feedback is ubiquitous in biological signaling and multicellular coordination and likely results from similar constraints to those observed in the nervous system.

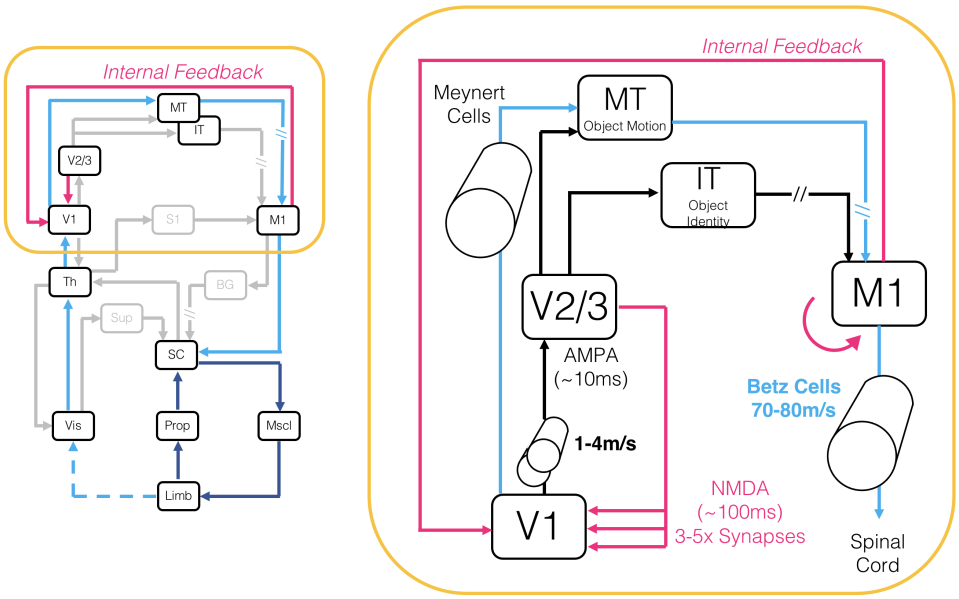

Figure 3.1: A partial schematic of sensorimotor control with layers and internal feedback. Sensorimotor control involves complex coordinated action across many parts of the nervous system. In order to study this system, it is often necessary to isolate single boxes or loops; here, we take steps towards understanding the function of layered loops and internal feedback. Note that even with all of these additions, we are simplifying away yet more diversity and complexity, and this diagram is schematic, intended to highlight particular features of the overall system. From bottom to top, more *layers* are added to the control architecture. Internal feedback pathways are indicated in pink, while a particular fast forward path for object tracking is indicated in light blue. Gray lines are connections we will discuss qualitatively, but which we do not address with quantitative theory. Dashed lines and broken lines indicate connections for which we discuss signal flows, but not necessarily direct neuronal projections. Abbreviations from bottom: Vis = vision, Prop = proprioception (muscle spindles), Mscl = muscles, SC = spinal cord, Sup = superior colliculus, Th = thalamus, BG = basal ganglia, V1 = primary visual cortex, S1 = somatosensory cortex, M1+ = primary motor cortex and additional motor areas, V2/3 = secondary and tertiary visual cortex, IT = inferotemporal cortex, MT = mediotemporal cortex.

The principles we describe are general, but we will ground the theory in familiar parts of the whole: differences in population responses between M1 and V1, differences in kinetics between AMPA and NMDA receptors, coordination between vision and proprioception, roles of giant pyramidal cells in visuomotor control, and localization of function in cortex.

3.2 Results

The Task

To develop a theory of layering and internal feedback, we analyze theoretical bounds on task performance for models of a simple, well-studied, and ethologically relevant task: reaching for and grasping a moving object (which we will simply call the *tracking* task). The complete task requires identification of the object amidst a visual scene, prediction of the object's movement, generation of bimanual limb movement, and execution of limb movement. The tracking task has a firm foundation in existing data and experimental approaches, making the theoretical arguments in this chapter directly testable in future experimental work.

Each of the subsequent models of this task will conform to the following simplifying assumptions, although we will not need every assumption for every model:

- In a visual scene, backgrounds change more slowly than moving objects.
- We treat objects as adding to the background scene rather than occluding it.
- Task success is defined as successful continuous pursuit, not one-time contact between limb and object.

These assumptions are not essential to our conclusions, but they allow is to study the essence of layering and internal feedback by showing that these can arise with small modifications to simple single-loop systems. As an added advantage, these simplifications allow us to use the familiar-in-neuroscience language of linear dynamical systems, so that our computations of best- and worst-case system performance are transparent.

Static Control

In the simplest case of the tracking problem, we consider tracking the endpoint of a limb on a plane, where the variable to be controlled is the distance between the limb endpoint and the object, or the *tracking error*, and where the system that controls the limb can perfectly sense the position of limb and object at every instant. We describe the time-evolving dynamics of the tracking error x, the action on the limb u, and the action of the object w in terms of A, a matrix which gives the intrinsic dynamics of x (e.g. the mechanical coupling between the two dimensions of limb movement). The actions u are feedback control actions on the tracking error, computed by an

arbitrary function $\boldsymbol{K}$ that has access to all past and present states $x(1:t)$.

$$\begin{aligned} x(t+1) &= Ax(t) + u(t) + w(t) \\ u(t) &= \boldsymbol{K}(x(1:t)) \end{aligned} \tag{3.1}$$

The optimal solution to this problem is straightforward, but interesting as a starting point: the optimal $\boldsymbol{K}(x(1:t))$ is exactly $-Ax(t)$. The optimal control action does not depend on any past states (the controller is *static*), and access to past states does not improve performance. There is no need for, and no advantage to, layering or internal feedback. This remains true for a large but special class of problems, not just problems where $A = -K$. This is mathematically remarkable, but it can be misleading about real-world systems; by relaxing some of the assumptions in this problem in subsequent sections, we will quickly encounter layering and internal feedback.

Estimation

Simple modifications to the control problem described above are sufficient to create a situation for which we can still compute a static K, but we can only implement it with internal feedback. One such modification is a delay in sensing. To rewrite a delayed-sensing problem as one that we can solve optimally using the standard control methods for the linear quadratic regulator (LQR), we introduce a virtual state $x^{(2)}$, the tracking error in the previous timestep, and write the tracking error itself as $x^{(1)}$. We can now treat this as a multivariate problem:

$$\begin{aligned} \hat{A} = \begin{bmatrix} A & 0 \\ I & 0 \end{bmatrix}, B = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}, C = \begin{bmatrix} 0 & I \end{bmatrix} \\ \begin{bmatrix} x^{(1)}(t+1) \\ x^{(2)}(t+1) \end{bmatrix} = \hat{A} \begin{bmatrix} x^{(1)}(t) \\ x^{(2)}(t) \end{bmatrix} + Bu(t) + w(t) \\ u(t) = KC \begin{bmatrix} x^{(1)}(t) \\ x^{(2)}(t) \end{bmatrix} \end{aligned} \tag{3.2}$$

In this new system, the controller senses x partly rather than fully, through the matrix C, so the controller does not directly sense the true state $x^{(1)}$. However, the controller can freely take actions that affect both the true state $x^{(1)}$ and its internal representation of the virtual delayed state $x^{(2)}$. This gives us our first opportunity to study internal feedback. Now the optimal solution can be computed as be a static control matrix K; in fact, K has the simple analytical form $\begin{bmatrix} K_1; & K_2 \end{bmatrix} = \begin{bmatrix} -A^2; & -A \end{bmatrix}$. We can directly

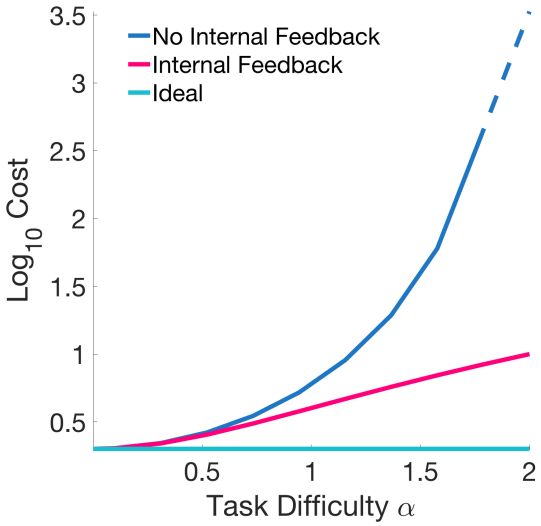

Figure 3.2: Delays in simple and otherwise unconstrained sensing and action necessitate internal feedback for good performance. For simulations, we consider the scalar problem of tracking a moving target over a line, varying the task difficulty (dynamics $A = \alpha$). Smaller costs indicate better tracking. As α approaches 2, the task becomes infeasible without internal feedback (broken line).

compare the performance of this controller (with internal feedback) to the optimal true static controller (Figure 3.2). In this case, the optimal static controller has the analytical form $\begin{bmatrix} K_1; & K_2 \end{bmatrix} = \begin{bmatrix} -A^2/4; & 0 \end{bmatrix}$. We can also characterize conditions under which internal feedback is necessary, not just for performance but for any solution to exist: a feasible static controller only exists when the spectral radius of A is less than 2.

Delays in sensing and delays in action are both relevant to sensorimotor control. Extending our model further, we can account for this in terms of a controller K and an estimator L. (For linear systems with Gaussian noise, the optimal Bayesian estimator L is the Kalman Filter.)

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + w(t) \\ y(t) &= Cx(t) \end{aligned} \tag{3.3}$$

The controller for this system inherently contains internal feedback irrespective of delays being present; these are represented by *IFP-Sense-2*, *IFP-State*, and *IFP-Act-2* in Figure 3.3. IFP-State estimates state evolution in the absence of noise and actuation, IFP-Act-2 accounts for controller action, and IFP-Sense-2 predicts incoming sensory signals based on the internal estimated state.

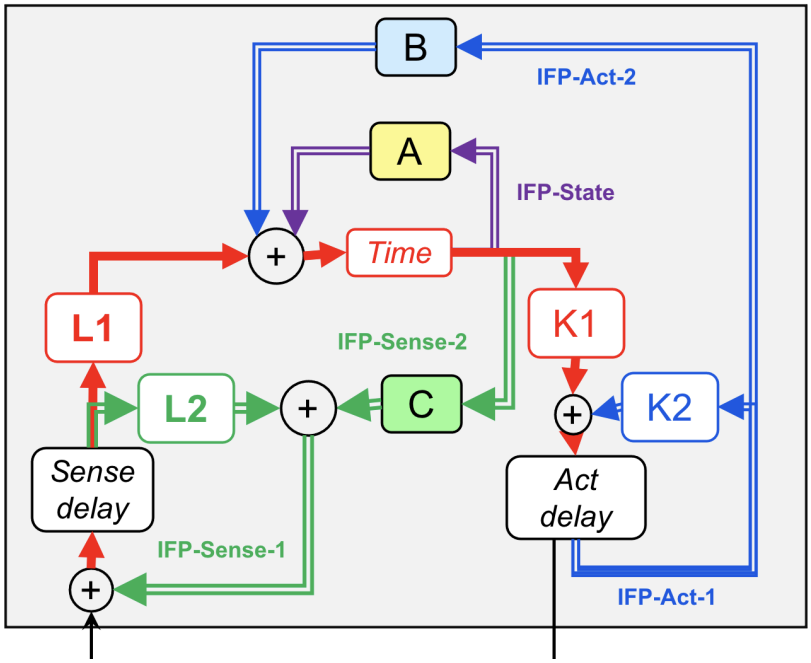

Figure 3.3: Internal feedback in a controller with sensing and actuation delays. A, B, and C represent the state, actuation, and sensing matrices of the physical plant; K_1, K_2, L_1, L_3 are submatrices of the optimal controller and observer gains. Delayed sensing and actuation induce two independent internal feedback pathways: IFP-Sense-1, and IFP-Act-1, which can be studied separately in the full control and state feedback subproblems. The remaining three internal feedback pathways (IFP-Sense-2, IFP-State, IFP-Act-2) are part of the delayed-sensing, delayed-action controller.

Let the augmented state be $\hat{x} = \begin{bmatrix} x^\top & x_a^\top & x_s^\top \end{bmatrix}^\top$, where x_a represents delayed actuation states and x_s represents delayed sensing states.

The system matrices for the augmented system are written in block matrix form:

$$\hat{A} = \begin{bmatrix} A & B & 0 \\ 0 & 0 & 0 \\ C & 0 & 0 \end{bmatrix}, \hat{B} = \begin{bmatrix} 0 & 0 \\ I & 0 \\ 0 & I \end{bmatrix}, \hat{C} = \begin{bmatrix} 0 & 0 & I \end{bmatrix} \tag{3.4}$$

We can then obtain the optimal estimator $L = \begin{bmatrix} L_1^\top & L_2^\top & L_3^\top \end{bmatrix}^\top$ and the optimal controller gain $K = \begin{bmatrix} K_1 & K_2 & K_3 \end{bmatrix}$. Because of the block-matrix structure of the system matrices, L_2 and K_3 are optimally zero; we then overload notation by reassigning L_3 to L_2. This allows us to simplify the resulting estimator and controller into the one shown in Fig. 3.3.

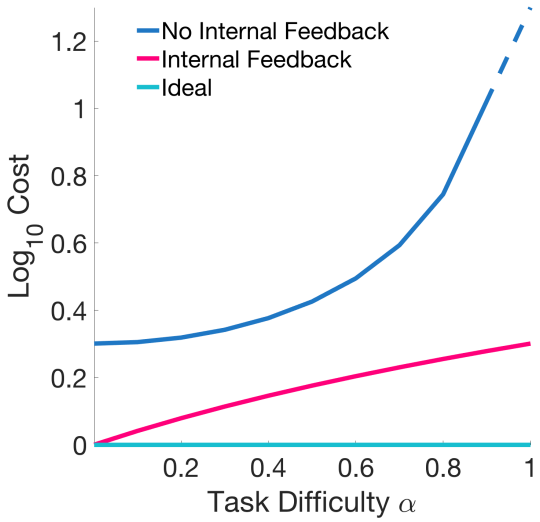

Figure 3.4: Localization of function within motor-related cortex: although different parts of the cortex control different parts of the body, these parts of the body are inherently mechanically coupled. As a result, internal feedback is useful and in some cases necessary to maintain localization of function. Centralized control performs best, but with internal feedback, control is still achievable. Without internal feedback, performance is substantially worse and can be infeasible (broken line).

$$\begin{aligned} \delta(t+1) &= Cx(t) - C\bar{x}(t) - L_2\delta(t) \\ \bar{x}(t+1) &= A\bar{x}(t) + B_2u_a(t) + L_1\delta(t) \\ u_a(t) &= K_1\bar{x}(t) + K_2x_a(t) \end{aligned} \tag{3.5}$$

Here, δ is the delayed difference between the estimated sensor output and true sensor output, discounted by the observer term $L_2\delta(t)$. The resulting controller, shown in Fig. 3.3, contains two sources of internal feedback due to delay: *IFP-Sense-1* and *IFP-Act-1*. The remaining internal feedback is inherent to the Kalman Filter.

Locality

We next consider locality constraints, which can include localization of function between sensory and motor areas and localization of function between one motor area and another. Here we focus on localization of function between one motor area and another.

Locality constraints have a physical interpretation, related to how quickly and accurately signals can be conveyed between subsystems. They have a less formal interpretation in the sense of coordination: who needs to know what, and when, for

the whole system to perform? Localization would mean that some subset of commands u_1 can be computed by sensing some subset of states x_1 alone, and another subset of commands u_2 can be computed by sensing another subset of states x_2 alone. For simplicity, we assume that there are only two subsets of actions and states, and that all commands and all states belong to one subset or the other. This is, again, a problem whose performance is substantially improved with internal feedback in the form of delayed cross-talk between the localized sub-controllers.

We start from a two-dimensional task without sensing or actuation delays:

$$x(t+1) = Ax(t) + Kx(t) \tag{3.6}$$

The optimal solution (minimizing $\|x\|_2$) is by inspection $K = -A$, which is static and has no internal feedback. Motor localization means that we can construct:

$$K = \begin{bmatrix} K_1 & 0 \\ 0 & K_2 \end{bmatrix} \tag{3.7}$$

If x_1 and x_2 are not coupled (if A is block-diagonal), then the optimal K is block-diagonal. However, in reality, all body movements are mechanically coupled, something which the motor system can conceal through effective localization.

Any controller can be implemented in a variety of ways. In order to study internal feedback, we assume particular non-unique controller implementations – which, when optimized, attain the optimal performance over all localized controllers. The internal feedback implementation we propose here is similar to the internal feedback we used in the case of sensing delays. We create additional virtual states, x_1' and x_2', which can be understood as what about x_1 can be communicated to K_2 without direct sensing and what about x_2 can be communicated to K_1 without direct sensing, respectively. The essence of localization is that u_1 never sees x_2 directly, in a timely manner; however, with internal feedback, u_2 can freely tell u_1 what it knows about x_2 at some delay. For simplicity, we will assume here that this is a unit delay.

$$\begin{bmatrix} x_1(t+1) \\ x_2'(t+1) \\ x_1'(t+1) \\ x_2(t+1) \end{bmatrix} = \begin{bmatrix} A_{11} & 0 & 0 & A_{12} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ A_{21} & 0 & 0 & A_{22} \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2'(t) \\ x_1'(t) \\ x_2(t) \end{bmatrix} + \begin{bmatrix} * & * & * & 0 \\ * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2'(t) \\ x_1'(t) \\ x_2(t) \end{bmatrix} + \begin{bmatrix} w_1(t) \\ 0 \\ 0 \\ w_2(t) \end{bmatrix} \tag{3.8}$$

The original localization structure of the problem is preserved because of the zeros in the top right and bottom left corners of the K matrix. We next determine

what the internal feedback signals must convey to achieve optimal performance. We choose this case because the optimal solution is relatively transparent. This kind of problem can be solved with System Level Synthesis (SLS) methods, which substantially generalize the LQR methods used in previous sections and can produce complex optimal controllers, but the technical details are not crucial to this special case, which produces a simple optimal K. We inspect the optimal K (which can be realized in multiple cost-equivalent ways, all with the relevant internal feedback):

$$K = \begin{bmatrix} -A_{11} & 0 & -A_{12} & 0 \\ 0 & 0 & -A_{12} & A_{12} \\ A_{21} & -A_{21} & 0 & 0 \\ 0 & -A_{21} & 0 & -A_{22} \end{bmatrix} \tag{3.9}$$

We can develop intuition for this implementation by following an impulse w through time.

$$x(1) = \begin{bmatrix} w_1 \\ 0 \\ 0 \\ w_2 \end{bmatrix} \rightarrow x(2) = \begin{bmatrix} A_{12}w_2 \\ A_{12}w_2 \\ A_{21}w_1 \\ A_{21}w_1 \end{bmatrix} \rightarrow x(3) = 0 \tag{3.10}$$

Here the internal states carry the predicted values (after control action) of the unsensed states. Without this internal feedback, task performance can be arbitrarily bad, or the task may be infeasible (Figure 3.4).

With internal feedback, task performance stays near the centralized optimal. Localization of motor function (specialization of parts of motor cortex to particular parts of the body) in fact requires ongoing information-sharing between local subcontrollers (signals related to the whole body appearing in, but not dominating, parts of motor cortex specialized to particular parts of the body). For the controller without internal feedback, we restrict to controllers that are linear; for the controller with internal feedback, this restriction is not needed, as linear solutions are globally optimal.

Attention

We have shown that in the presence of internal delays, corrections for self-movement are necessary for performance. We now want to extend this reasoning to a more physiologically grounded model, while also considering other predictable aspects of the visual scene beyond self-movement. Up to this point, we have made the assumption that the controller can directly sense the position of the object (perhaps

at some delay). In a real sensorimotor system, this assumption is not sensible; the position of the object must be computed from the scene, not directly sensed. This is a more complex information-processing task: it takes more bits to describe the whole visual scene than to describe the position of an object in the same scene.

However, moving objects, once identified as such, can be discriminated from a visual scene. A meaningful separation can be made between scene-related tasks (object identification) and error-related tasks (object movement). This separation mirrors the separation between bumps and trails in the mountain-biking task studied in [27], allowing us to build on the theorems in that earlier work. The main difference is that instead of separating into two control loops, we reuse the same sensor for each loop, necessitating internal feedback. We consider $\|x\|_\infty$ for this problem rather than $\|x\|_2$. The two norms have similar behaviors in this simple case. We also drop down from a two-dimensional problem (tracking on a plane) to a one-dimensional problem (tracking on a line).

Formally, assume that the absolute position of the object y is in some interval $[-H, H]$, while the change in position of the object w is in the interval $[-L, L]$ with $L << H$.

Consider a controller which has two pathways out of the sensor and a chance to see and store the whole visual scene prior to the appearance of the object. This controller can compute the absolute position of the object once, on a delay, and subsequently consider only the relative position, accurately updating the stored absolute position. In order for this to work in practice, the stored absolute position must be available near the sensor. In some architectures, this might be implemented fully at the sensor, which could reduce delays. In other architectures, more like the true tracking problem, the stored absolute position depends on some contextual knowledge (such as knowledge of the boundaries of the object) that makes longer processing necessary.

To model this with the tools in [27], we assume that the errors in measurement scale with the size of the interval according to some quantitative relationship. As the sensed interval $[-\beta, \beta]$ increases in size, the accuracy decreases, but on the other hand can be improved by time. We assume the quantizer is static and memoryless with a uniform partition; with no control cost, these assumptions can recover the optimal cost over all quantizers [30].

We formalize the speed-accuracy tradeoff:

$$R = \lambda T \tag{3.11}$$

And the worst-case errors for any given R:

$$\frac{\beta}{2^{R-1}} \tag{3.12}$$

Because the high-layer controller is neither perfect nor instant, we cannot simply take $\beta = L$. A simple modification wherein the background visual scene evolves slowly, rather than not evolving at all, adds more complexity. Consider the evolution of the visual background scene v:

$$v(t+1) = v(t) + w_v(t) \tag{3.13}$$

where we assume w_v is restricted to the interval $[-q, q]$. So if it takes T_H steps to sense v_t, the estimate provided has errors from the interval qT_H due to drift, and errors from the interval $H/(2^{\lambda T})$.

In this more realistic scenario, we can do no better than $\beta = L + qT_H + H/(2^{\lambda T_H})$, which in turn creates the overall state cost:

$$LT_L + \frac{(L + qT_H + H/(2^{\lambda T_H - 1}))}{2^{\lambda T_L - 1}} \tag{3.14}$$

Performance generally improves as T_L becomes small, while performance generally improves as T_H increases but can start to degrade if T_H becomes too large (Figure 3.5). The counterfactual system without internal feedback is one for which costs are $T_H + \frac{H}{2^{\lambda T_H}}$.

The key point is that fast, flexible, and accurate sensing can be achieved in a visual scene with predictable elements through internal feedback under reasonable assumptions about the dynamics of the environment.

3.3 Discussion

Interpretation

In the technical results, we considered a collection of simple models to isolate and illustrate different aspects of layering, diversity, locality, and internal feedback in sensorimotor control. We have not proposed an end-to-end parametric model of sensorimotor processing. While such a model may be an eventual goal of the field, we take a step in that direction by extracting mathematical principles in minimal models to explain features of sensorimotor control that have up until now been

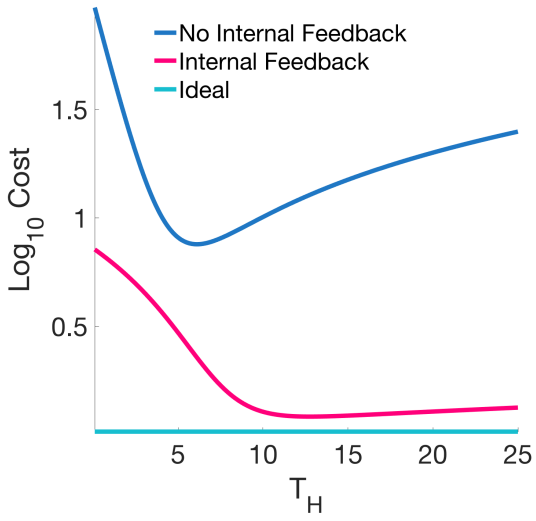

Figure 3.5: Feedback within the sensing hierarchy can compensate for speed-accuracy tradeoffs in estimation for suitable tasks, particularly when there are aspects of the visual scene (like the visual background or the object identity) that change slowly. This allows the creation of a fast forward pathway compensated by a slow internal feedback way. Costs depend on the speed and accuracy of both the slow and the fast paths. The fast path needs to be as fast as possible, while the slow path can take more time once the fast path is established. The layering of the fast and the slow paths (red line) achieves better performance than the unlayered system without internal feedback and subject to the same speed-accuracy tradeoff (blue line). Ideal performance, where the fast path can directly sense the moving object, is also shown. The layered system achieves performance close to the ideal.

cryptic. In this section we will use a narrative, non-technical style to synthesize evidence for readers from a range of backgrounds. We will use evidence from primate, mouse, and feline visuomotor systems. Mice, as small prey animals, have quite different visual systems from primates in several respects, so we incorporate evidence from mice cautiously and only where necessary in the context of the larger explanatory framework.

Across all of the models we have considered, diversity and speed-accuracy tradeoffs are intrinsic to the underlying components in the system as a result of evolutionary or biophysical tradeoffs [31]. Speed-accuracy tradeoffs can result from several mechanisms, such as a tradeoff between smoothing and timeliness at the synapse and a tradeoff between number of neurons and axonal diameter in long-range projections.

Layering and diversity together enable sweet spots, wherein diverse components are multiplexed in a task-specific way that approximates the performance a single-

loop system made up of ideal components. In order to achieve fast, accurate, and flexible control, the fastest components are reused flexibly in a fast forward loop, while internal feedback compensates for accuracy by filtering out slow-changing, predictable, or task-irrelevant stimuli, such that the fewest possible bits need to travel along the fastest possible neurons. From an evolutionary perspective, once a system can achieve fast responses, it becomes possible to add successive layers of more accurate and more flexible components. This is an essential aspect of layering: flexible reuse of existing parts to solve new problems.

Primary visual cortex (V1) and primary motor cortex (M1) have provided many examples in recent years of internal feedback that challenge the canonical view of sensorimotor control [17, 21]. In the canonical view (whose shortcomings are by now well-understood), visuomotor processing consists of a series of successive transformations from stimulus to response, with each cortical area along the way tuned to some aspect of stimulus space [32]. However, as these systems have been characterized in more detail, the empirical evidence has accumulated that complicates this model. V1 does include static representations of stimuli, but it also includes counterdirectional motor-based internal feedback and task-or attention-related modulatory internal feedback [16, 17, 20]. M1, by contrast, is dominated by its own past activity rather than static representations [21]. In the context of the estimation problem we considered in Figure 3.3, motor cortex is dominated neither by motor representations nor by pattern generation, but by predictions of the consequences of self-action through local internal feedback, which need to be sent throughout the body. By the same principle, the localization of function within motor cortex that we considered in Figure 3.4 would explain why the conventional view of homuncular organization becomes fuzzier, with, for instance, body-related signals found in putatively hand-related parts of motor cortex, as well as contralateral hand signals [22, 23, 33]. As with motor signals in visual cortex, these signals do not overturn the preceding view that function is localized. These additional signals, in fact, are *necessary* for localization of function.

Internal feedback is also observed between visual cortical areas. The neurons between V1 and V2 are a well-characterized case study. Roughly the same number of neurons, of similar conduction speed, project from V1 to V2 as from V2 to V1 [14, 15, 19]. However, these neurons are quite different in morphological and molecular characteristics: V1-V2 neurons primarily use the AMPA receptor, while while V2-V1 neurons primarily use the NMDA receptor [34]. While AMPA

and NMDA use the same neurotransmitter, glutamate, AMPA-mediated currents are cleared from the synapse in under 10ms, while NMDA-mediated currents can linger in the synapse for 100ms or longer [35]. This distinction is believed to be relevant for learning, and it is also relevant for top-down signaling to shape and direct perception during action. If we take the conventional view that V1 detects edges and V2 detects more involved elements of objects like contours, then V2 can facilitate edge detection by providing signals that inform V1 about where edges ought to be. Because the visual space cannot be sampled losslessly, these signals help target the resolution of V1. Accordingly, experiments that pharmacologically knock out NMDA show a loss of figure-ground discrimination, that is, a loss in capacity to contextually interpret the visual scene [34].

We have argued that the essential role of internal feedback and layering is to enable the fast, accurate, and flexible accomplishment of simple sensorimotor tasks in the presence of speed-accuracy tradeoffs. If this is true, we should see that fast-conducting, large-axon myelinated neurons are engaged in rapid responses to changes in the visual scene. Indeed in cortex, the largest and most striking neurons are the large pyramidal cells, also called Meynert Cells in visual cortex and Betz Cells in motor cortex, which transmit rapid moving-object changes in visual scenes or rapid responses to perturbations in planned movements, respectively [36–38]. Meynert Cells project from V1 to MT (an object motion area in the dorsal stream) but no equivalently large cells project from MT to V1, nor from V1 to IT (an object identity area in the ventral stream). In the context of our theory, it is necessary for predatory animals, in particular, to track object locations, which can change quickly. However, object identities are relatively slow-changing, and do not require a fast pathway. Because reaching and tracking tasks are widely studied in motor neuroscience, this theory therefore suggests additional experimental directions in reaching tasks that involve rapidly and unpredictably moving objects against predictable but not static backgrounds.

Relation to other theories

Optimal and robust control analyses have historically offered only limited insight into the *implementation* of control systems, which is a minor issue in most existing control theory and applications due to fast and accurate digital electronics. Spiking neurons are fast relative to other biological signaling mechanisms, but slow relative to electronics. From the perspective of control theory, biology must work with slow and inaccurate implementations of control, communications, and computing

components, so the associated internal architectures become critical. A single-loop model treats the entire nervous system as a single computing and control system, and supposes that this system can carry out sensing, decision-making, and action with communication speeds and computational power beyond what is available to the real nervous system. For any task with suitably defined performance criteria, a single-loop control system without component tradeoffs will perform at least as well as a layered control system with component tradeoffs. Thus single-loop models are entirely appropriate to analyze sensorimotor behavior in well-defined tasks. However, these models do not generate useful models of the *implementations* of sensorimotor control (e.g. the neural circuits and nervous system architecture).

The unclear role of internal feedback has been studied from different theoretical angles throughout neuroscience: computation through dynamics [39], recurrent networks [40], Bayesian inference [41, 42], predictive processing [43, 44], and many others. These frameworks have arisen in parallel with the development of methods for increasingly high-throughput and high-resolution measurements of biological systems, which support the idea that internal feedback and internal dynamics are essential for control – as well as for learning and computation.

As past work has argued, internal feedback can be interpreted as predictive; in recent neuroscience data, a major source of internal feedback is the known effects of past and current actions [17, 45]. These internal feedback signals carry signals about how actions propagate through the body and its environment, as well as about planned future actions, including how communication delays limit such plans and actions. These sources of internal feedback encompass more than Bayesian prediction, which is a special case of a general phenomenon. If the Bayesian framework is to be used, it must be used in consideration of task performance as well as sensory evidence, because costs shape optimal estimates [30].

These existing theoretical frameworks capture something useful about biological control and internal feedback. The missing piece we provide to tie them together is *speed*, which grounds these theories in physiological details and evolutionary tradeoffs.

In the previous chapter, we considered the nervous system as a multicellular control system. Because one of the overarching motivations of this thesis is to design interfaces with natural systems, I note here that a compatible conceptual approach has been used to improve performance and robustness of brain-computer interfaces for use by people with tetraplegia, including in clinical trial efforts to which I contributed [46–48].

In the next chapter, we move from the nervous system to the immune system, where many similar problems arise: layers, diversity, locality, and internal feedback are all features of the immune system as well as the nervous system. We have treated the nervous system as a multicellular decision-making and coordination systems, and we will treat the immune system as another such system with quite different components, tasks, and biophysics. Some of the techniques we used to model speed-accuracy tradeoffs in the nervous system will inform our view of the immune system.

Chapter 4

LIMITS ON IMMUNE ROBUSTNESS DRIVE VARIATION IN PATHOGEN VIRULENCE AND TRANSMISSION

4.1 Introduction

The risk that a pathogen poses to a population is often summarized in terms of two properties of infection: virulence (the severity of disease caused by the pathogen, often quantified as the case fatality rate) and transmission (the rate at which the pathogen spreads from host to host, often quantified as the basic reproductive number). In general, pathogens that cause both high virulence and high transmission create the greatest burden at the population level. The evolutionary relationship between pathogen virulence and transmission has been studied theoretically and empirically for decades [49–55]. Virulence and transmission depend on interactions between pathogen and host that can be studied using evolutionary game theory [50], with pathogen and host each facing evolutionary tradeoffs that can be enacted through specific host-pathogen interactions and host immune responses. At the extremes, these tradeoffs can be characterized simply if the replication rate of the pathogen is coupled to disease severity: a pathogen that does not replicate would not be very virulent, but would not produce enough copies to transmit; a pathogen that replicates very quickly would be so virulent as to cause severe disease and death in the host, and would therefore not transmit; because only pathogens that transmit are evolutionarily successful, the optimal replication rate and the optimal virulence for the pathogen must be somewhere between these two extremes.

Many studies have added theoretical and empirical nuance to this view of pathogen virulence and transmission. However, two major sources of additional complexity, inter-host variation and dynamics in the host immune response, are difficult to study with conventional mathematical models. Recently, the SARS-CoV-2 pandemic has shown that it is essential to understand inter-host variation and dynamics in order to predict pandemic spread [55–57].

What are the relationships between varying virulence and transmission at the individual host level? A theoretical answer to these questions should meet several criteria: It should be based on a quantitative understanding of the within-host dynamics of infection, address specific biological aspects of infection for a given virus, and

address strategies in the general evolutionary game between pathogens and hosts. The answer, then, depends on a framework to generate models of host-pathogen interactions, rather than a single highly parameterized model intended to address every possible interaction. In this chapter, we develop techniques to generate rigorous mathematical analyses of infection kinetics, pathogen strategies, and immune responses, despite the complexity of the immune response and despite the possibility of unknown dynamics in the immune response and in pathogen strategies. We introduce and adapt techniques from robust control theory to generate interpretable models and analyses for problems that cannot be conventionally modeled.

The organizing idea of our analysis is that many biological systems have evolved to be robust: survival-critical behavior is maintained despite internal and external perturbations and wide variation in kinetic parameters [58–60]. A healthy antiviral immune response is usually robust in this sense, repeatedly and reliably clearing a wide variety of respiratory viruses. Variation in virus virulence and transmission should be understood in the context of this typical robustness. Robust control theory is a mathematical framework that has been used to analyze robustness and variation in complex feedback systems in both engineered and biological settings [8, 59–64]. The robust control approach facilitates rigorous modeling of biology across scales despite the possibility that there might be important dynamics in either the virus or the immune response that are not explicitly described in the model equations. In essence, by analyzing evolved biological robustness to real-world uncertainties, we can make rigorous model- and theory-based conclusions that are robust to epistemic uncertainties about the underlying biology. A key insight of robust control theory in the engineered setting is that certain engineering problems are intrinsically and quantifiably more difficult than others, similar to the familiar notions of a problem's hardness in computational complexity theory. Here, we show that only a small number of host-viral interactions are needed to make individual- and population-level control of a theoretical virus intrinsically and quantifiably much harder than control of a closely related virus. We quantify the difficulty of the problem in terms of the best possible virulence and transmission achievable by any of a large set of immune responses and related host behaviors and public health policies.

In the main technical study of this chapter, we restrict our consideration to acute respiratory viruses. To introduce ideas and highlight the diversity of viral and host strategies we will here make qualitative comparisons between other well-known viruses. Low-virulence, high-transmission viruses (blue in Figure 4.1A-B) elicit

mild immune responses and symptoms, requiring minimal response at the behavioral and policy levels. Examples include rhinoviruses, mild influenza strains, and some coronaviruses that cause common colds. High-virulence, low-transmission viruses (red in Figure 4.1A-B) elicit severe immune responses but can be contained through diverse host behaviors and public health policies. Rabies, Ebola, and severe influenza could all be considered examples in this category, but with fundamentally different biology and behavioral and policy responses. Rabies is not transmissible between people except through unusual circumstances like organ transplantation [65], Ebola is transmitted through contact with body fluids from symptomatic patients [66], and influenza may transmit through droplets and direct contact presymptomatically [67, 68]. Therefore, regardless of population responses, rabies transmission will remain in the high-virulence, low-transmission category. Ebola and severe influenza may instead appear as high-virulence, high-transmission viruses (purple in Figure 4.1A-B) if population responses are ineffective.

Complex variation in pathogen virulence and transmission can be interpreted once the robust control strategies that create predictability in other viruses are established. Potent viral interferon suppression, layered over a sparsely but variably expressed host receptor, can drive substantial variation (Figure 4.1C). While these aspects of viral and host biology are well-known, the role of receptor sparsity as a major driver of the population-level risk would have been difficult to identify without a mathematical framework.

4.2 Results

A robust control approach to pathogen-immune dynamics

We assume the fundamental evolutionary pressures on pathogen and host to be familiar: pathogen evolution favors high transmission (and vice-versa), while host evolution favors low virulence. Cold-like viruses are examples of possible evolutionarily stable strategies (similar to the blue viruses in Figure 4.1); however, relatively high virulence and high transmission (purple) are achieved by some pathogens, including SARS-CoV-2, and these are typically the pathogens of greatest concern. These pathogens achieve high virulence and high transmission despite counter-efforts by the immune system and additional interrelated anti-pathogen measures in the host population.

The next step in our approach is to identify the best possible immune response strategy to a class of pathogens, corresponding to acute respiratory coronaviruses

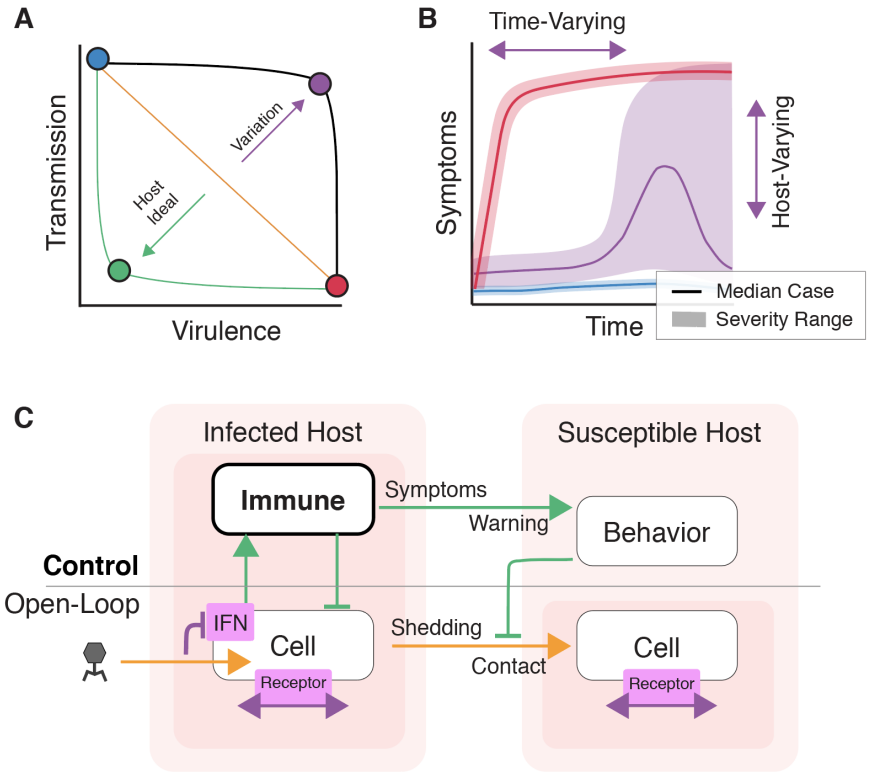

Figure 4.1: A control theory framework to analyze viral virulence and transmission. Qualitative cartoons show three levels of description explored quantitatively in this chapter: (A) is a population-level summary of two key properties of a pathogen, virulence and transmission, with variations in severity and transmission between individuals and across time in individual infections summarized into single points in the two-dimensional space. The diagram shows virulence and transmission combinations for four qualitative virus types (circles) as well as potential relationships between virulence and transmission (lines). Green viruses are the best case for the host population: They neither cause severe symptoms nor spread easily, so few individuals in the population are affected. Blue viruses spread easily but cause minimal harm to individuals and therefore minimal burden to the population. Red viruses cause severe symptoms and harm to individuals, but do not spread easily, burdening few affected individuals but not the whole population. Purple viruses with high virulence and high transmission are the worst case for the host population, causing many severe infections and a substantial burden to the population. Crucially, a virus's position in this space can depend both on its own biology and on host biological, technological, and behavioral interventions. (B) Symptom dynamics highlight some, but not all, additional complexity in the relationship between virulence and transmission. Solid lines represent median cases, while lighter colored clouds represent the range of severities. (C) We use robust control theory to derive rigorous lower bounds on the severity of virulence and transmission given specific virus-host interactions, despite significant epistemic uncertainties in how host responses are implemented.

that are identical except in a single lumped parameter that describes the rate of replication. The virulence and transmission associated with a given theoretical virus provide a quantification of the difficulty of control of that virus. Once we identify an immune response which achieves low virulence and low transmission against this class, we then add a small and well-characterized set of host and viral factors that constrain the set of possible immune strategies, such that even the best-case immune response leads to both high virulence and high transmission; that means that the virus is intrinsically more difficult to control. We will identify sufficient factors that drive variation in virulence and transmission.

While the eventual outcome of this analysis will use the well-known language of dynamical systems modeling and simulation, the models we produce are distinct in two ways. First, the immune models are selected rigorously out of the set of possible immune models that satisfy the explicit constraining assumptions, so that no assumption-satisfying immune models exist that achieve lower virulence and transmission than the models shown in simulation: when the simulated immune model yields high virulence and transmission, this is a strong result that describes the intrinsic difficulty of controlling these host-viral interactions. Second, the immune models have been constructed by a procedure that makes modeling a complex and intricate system like the immune system tractable despite the many unknowns in parameters, components, and interactions. We are therefore able to generate a simple model with a transparent relationship to more complex models.

Robust control theory can be used to reformulate and analyze evolutionary games, including the interaction between virus and host that we consider here [4, 69]. We use robust control theory to uncover the host and viral factors that lead to variation in virulence and transmission. Informally, we compute the best-case immune response, consolidating unmodeled immune dynamics into a control function K. The best-case K minimizes virulence and implicitly suppresses transmission. The variation in virulence and transmission computed for this best-case K can therefore be interpreted as follows: given particular viral replication dynamics and constraints on the model, any alternate immune system model we could have used would lead to greater variation in virulence and transmission and a greater population burden. Because this is true for any alternate immune system model, it is reasonable to conclude that the biological immune system corresponding to a high-variation K will also yield variation in virulence and transmission. The host-viral interactions captured in the constraints therefore capture the intrinsic difficulty of controlling the pathogen.

Robust control can be used to extend systems biology differential equation modeling. In a robust control formulation, we separate the system dynamics into viral loads v and immune actions u, with new virus entering the system as δ.

$$\begin{aligned} \dot{v} &= A_\Delta v + B_\Delta u + \delta \\ u &= K(v) \end{aligned} \tag{4.1}$$

A_Δ and B_Δ are sets of time-varying matrices describing uncertain linearized dynamics. In the case of the antiviral response, the uncertainty represents the wide range of viral kinetic parameters to which the immune system must be robust. K is a function that maps viral loads and kinetics to immune responses. For a linear A and B, the best-case K will also be linear. The linearizing assumptions will typically allow the immune system more capacity to suppress virulence and transmission than we expect in biological reality (for example, nonlinearities resulting from limits on the total number of white blood cells the immune system can produce to fight a virus could increase virulence and transmission). We note an exception to this rule when we introduce the linearized viral dynamics below. We leverage theorems guaranteeing that the best-case K can always be computed from a closely related convex set, so that the best-case K computed from the set will be the best out of all realizable functions. We implement kinetic details as constraints on the set of realizable control functions, and in this way identify specific aspects of host and viral biology (constraints) for which even a best-case K yields virulence and transmission variation. This best-case K bounds any immune system model that we could have used, allowing us to pose rigorous questions about host-viral interactions without a detailed model of immune dynamics.

The open-loop problem

We first consider the open-loop dynamics of viral replication, or equivalently $K = 0$. The open-loop dynamics directly inform both the achievable closed-loop responses and the consequences of control failures like delayed immune responses (Figure 4.3A).

We model viral infection in the individual host as a three-step process: cell entry, replication in the cell, and release of virus from the cell after an eclipse period T_e [70, 71]. We use the bracketed notation $\{T_e\}$ to represent a delayed discrete event, while unbracketed constants represent the standard continuous rates. This open-loop model is derived from standard mass-action viral kinetics, and its simplicity is consistent with the robust control approach, because complex unmodeled dynamics

can be consolidated into the controller K as described above.

$$\begin{gathered} V = \text{extracellular virus} \\ C_s = \text{susceptible cells} \\ C_I = \text{infected cells} \\ \text{Virus enters cells: } V + C_s \xrightarrow{r} C_I \\ \text{Virus exits cells: } C_I \xrightarrow{\{T_e\}} \beta V \\ \text{Virus degrades extracellularly: } V \xrightarrow{k} \oslash \end{gathered} \tag{4.2}$$

We discretize time with $n-1$ time-steps per T_e to derive an n-by-n matrix A in terms of α, the number of productively infected cells that result from a single infected cell. For example, with $T_e = 12h$ and a time-step of $6h$, the virus spends 2 time-steps within an infected cell and is then released into the extracellular space, so the total states are extracellular virions V, intracellular virions that entered within the last 6h C_{I1}, and intracellular virions that entered between $6h$ and $12h$ ago, C_{I2}. Taken together, with rates rescaled appropriately, A is a 3×3 matrix. Virion degradation and cell death due to the immune response are not included in the open-loop parameters. The dynamics are linearized around the origin. In linearizing, we assume that C_s remains constant in a given host through time. Because the uptake of host receptor by infected cells and the clearance of infected cells by the immune system will reduce C_s, viral replication in the full nonlinear system will slow at sufficiently high viral loads. This effect is not captured by the linearization. We show the linearized open-loop dynamics for the 3×3 case, but alternate discretizations would lead to alternate dimensions in A ($12h$ for a 2×2 matrix, $3h$ for a 5×5 matrix, and so on). The matrix entry A_{21} is 1, rather than k, because of the way the C_I states are scaled, as we explain below.

$$\begin{gathered} \begin{bmatrix} V \\ C_{I1} \\ C_{I2} \end{bmatrix}_{t+1} = A \begin{bmatrix} V \\ C_{I1} \\ C_{I2} \end{bmatrix}_t \\ A = \begin{bmatrix} 1-k & 0 & k\alpha \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \\ \alpha = \frac{r\beta}{k}[C_s] \end{gathered} \tag{4.3}$$

Here, α scales linearly with $[C_s]$. Each of the parameters k, r, and β can vary widely between viruses [72, 73]. Both r and k can also vary between strains of a virus. A

fraction ϕ of infected cells will constitutively turn over in a single eclipse period, which may also affect inter-host variation in α (higher turnover leading to lower α). To highlight inter-host variation, we assume that the virus has identical k, r, and β across hosts. We do not explicitly consider variation in ϕ, but the consequences of variation in ϕ will be inverse to consequences of variation in $[C_s]$.

A small susceptible cell percentage (SCP) enables large relative variation in α. In the open-loop dynamics, small absolute changes in $[C_s]$, consequent to changes in SCP, lead to large changes in viral load (Figure 4.3B-C).

Derivation of open-loop dynamics

We linearize and discretize the open-loop reaction-and-delay network. To simplify notation, we consider a discretization of exactly T_e, leading to a 2×2 matrix, but the steps followed here can be extend to any discretization that divides T_e evenly. We obtain a two-state system, with v as the extracellular viral load and y as the number of infected cells and. For each infected cell, β virions are produced.

$$\begin{aligned} v[t+1] &= v[t] + \beta(1-k_c)y[t] - rsv[t] - kv[t] \\ y[t+1] &= rsv[t] \end{aligned} \tag{4.4}$$

We control the system with two actuators: intracellular interferon-stimulated pathways (ISPs) and myeloid cells in the extracellular space. Myeloid cells (κ_1) remove virus from the compartment. ISPs (κ_2) reduce the number of virions produced per infected cell.

$$\begin{aligned} v[t+1] &= v[t] - \kappa_1 \\ v[t] + \beta(1-k_c)(1-\kappa_2)y[t] - rsv[t] - kv[t] \\ y[t+1] &= rsv[t] \end{aligned} \tag{4.5}$$

We rewrite the system equations so that viral replication and ISP control are represented within the cell rather than at the moment of exit.

$$\begin{aligned} v[t+1] &= v[t] - \kappa_1 v[t] + y[t] - rsv[t] - kv[t] \\ y[t+1] &= rs\beta(1-k_c)(1-\kappa_2)v[t] \end{aligned} \tag{4.6}$$

We express these dynamics in matrix form:

$$\begin{aligned} A &= \begin{bmatrix} 1-rs-k & 1 \\ rs\beta(1-k_c) & 0 \end{bmatrix} \\ B &= \begin{bmatrix} 1 & 0 \\ 0 & rs\beta(1-k_c) \end{bmatrix} \end{aligned} \tag{4.7}$$

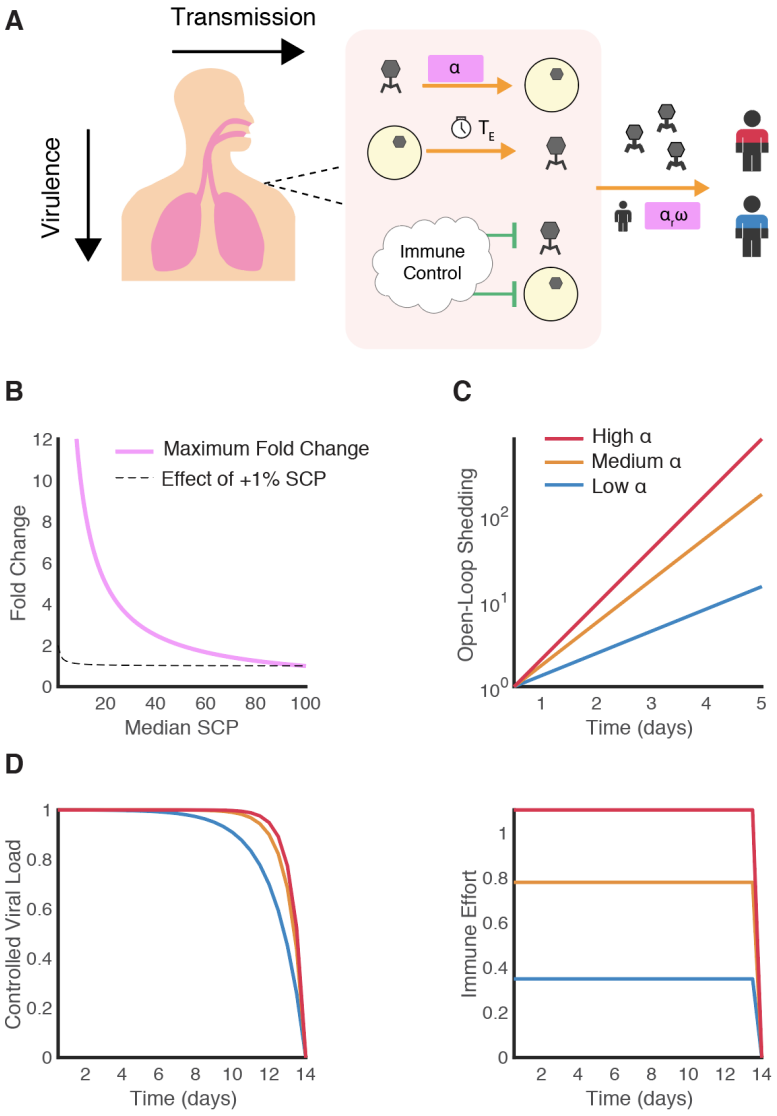

Figure 4.2: Host dynamics shape virulence and transmission. (A) Within each host, well-characterized kinetics govern viral replication. Immune responses are represented with a cloud to indicate that any additional dynamical elements in the sensing and coordination of viral removal are analyzed as part of a control system, for which we can compute best-case bounds with control theory (bottom two reactions). Transmission from a single host can lead to presymptomatic cases that go on to be severe (red-banded individuals), cases that are fully asymptomatic (blue-banded individuals), or any severity in between. (B) A low median susceptible cell percentage (SCP) in the host enables large relative variation in SCP and thus in α. (C) Open-loop variation in viral shedding varies dramatically on relevant time-scales, amplifying variations in SCP (note log units for shedding). (D) Ideal and rapid extracellular immune control can create similar, low-variation viral load trajectories between the same three hosts: high, medium, and low α (left). These similar simulated trajectories require differing immune effort (right). The underlying open-loop dynamics thus directly shape virulence.

From these dynamics, we define:

$$\alpha = \frac{r\beta(1-k_c)}{(r+k)}s \tag{4.8}$$

Because $\beta >> \alpha$ for typical values (e.g. 1000 vs 10), we approximate:

$$\alpha = \frac{r\beta(1-k_c)}{k}s = \frac{r\beta}{k\phi}s \tag{4.9}$$

Then substitute:

$$A = \begin{bmatrix} 1-k(\frac{\alpha\phi}{\beta}+1) & 1 \\ k\alpha & 0 \end{bmatrix} \tag{4.10}$$

And approximate:

$$\begin{aligned} A &= \begin{bmatrix} 1-k & 1 \\ k\alpha & 0 \end{bmatrix} \\ B &= \begin{bmatrix} 1 & 0 \\ 0 & k\alpha \end{bmatrix} \end{aligned} \tag{4.11}$$

At this point, we have matrices that can be used to compute the desired optimal state-feedback controllers. For a more compact and intuitive representation with $B = I$, we rescale y. We introduce a temporary variable y_2:

$$\begin{aligned} y_2 &= y/k\alpha \\ y[t+1] &= k\alpha v[t] + k\alpha u_2[t] \\ y_2[t+1] &= v[t] + u_2[t] \\ v[t+1] &= v[t] + k\alpha y_2[t] - rsv[t] - kv[t] + u_1[t] \end{aligned} \tag{4.12}$$

Then, replacing the original second state variable y with the new variable y_2, we obtain the new state-space dynamics:

$$\begin{aligned} A &= \begin{bmatrix} 1-k & k\alpha \\ 1 & 0 \end{bmatrix} \\ B &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \end{aligned} \tag{4.13}$$

To confirm that these dynamics retain the essential replication dynamics in α, we consider how many productively infected cells result from a single infected cell according to these dynamics. A single productive infection releases $k\alpha$ pre-infective virions. Cumulative new infected cells y^* are produced as follows:

$$\begin{aligned} y^*[1] &= k\alpha \\ y^*[2] &= k\alpha + (1-k)k\alpha \\ y^*[3] &= k\alpha + (1-k)k\alpha + (1-k)^2k\alpha \\ y^*[4] &= k\alpha(1 + (1-k) + (1-k)^2 \ldots) \end{aligned} \tag{4.14}$$

The sum on the right-hand side is a convergent geometric series:

$$k\alpha(1 + (1-k) + (1-k)^2 \ldots) = \frac{k\alpha}{(1-(1-k))} = \frac{k\alpha}{k} = \alpha \tag{4.15}$$

Therefore, y^* converges to α, as expected.

For our numerical computations, values of k and α are primarily important as relative quantities, so we select values in ranges reported for other influenza and coronaviruses: $T_e = 12h$, α varies 4-fold from 5 to 20, and $k = 0.16$ corresponding to a half-life of $2d$.

The controller is computed in the matrix 1-norm, weighted such that minimizing extracellular immune activity is the objective, i.e. $J_x = 0$, $Ju = \begin{bmatrix} 1 & 0; 0 & 0 \end{bmatrix}$. Interferon suppression and incubation period are treated as the same $5d$ period, during which $B[t] = 0$. The immune controller, better than reality by design, is instantaneously resynthesized for any of a set of matrices A or B.

After virulences are computed (as the weighted norm of the closed-loop map returned by controller synthesis), the warning w is scaled between 0 and 1, with $p = 5$ giving a suitable dynamic range from mild to severe. Similarly, the transmission parameter ω is scaled so that the interferon-suppressing virus achieves $R_{CL} = 2$. Closed-loop maps are constrained to be elementwise nonnegative to ensure realizability in a chemical reaction network.

The closed-loop problem

We sequentially introduce immune and viral complexity in the closed-loop problem. The open-loop problem suggests a natural distinction between intracellular and extracellular immune effectors. The key distinction is not the location of the effector but the site of action. Intracellular effectors include interferon-stimulated genes and immune cells that kill infected cells. Extracellular effectors include neutrophils, macrophages, and other immune cells that can trap or degrade viral particles found in the extracellular space. Like other taxonomies of immune responses (such as the distinction between cellular and humoral responses), this is an imperfect approximation that nevertheless captures a meaningful dimension of immune responses.

We start with innate extracellular immune effectors, corresponding to myeloid cells like neutrophils and macrophages. Higher α requires greater immune effort to achieve a comparable effect on viral load (Figure 4.3D). A higher SCP allows virions to leave the extracellular space more quickly. In order to clear the virus, extracellular effectors must then remove the virus more quickly, which can be achieved with more

myeloid cells in the tissue. Though notable, this variation in immune responses can be considered to fall within the minimal range needed to achieve robustness, because the virus is quickly contained to a tissue compartment and then removed. In a given tissue compartment, the minimal myeloid cell count depends on the viral replication rate rather than the viral load. However, if the virus spreads to other tissue compartments, more immune cells are required to achieve viral clearance.

To assess virulence from the simulated viral and immune dynamics, we make some simplifying assumptions. Respiratory viral infections can present with a diverse array of symptoms [55]. The correspondence of these diverse symptoms to a one-dimensional summary quantity is not straightforward. Rather than symptom heterogeneity, however, we are interested simply in variations in severity between hosts and across time. We therefore focus on infection of the respiratory tract. Patients with respiratory viral infections present with symtpoms ranging from mild nasal congestion to acute respiratory distress syndrome [74]. Often in respiratory viral infections, innate immune activation drives symptom severity [75], so we operationalize symptom severity as the modeled magnitude of innate immune activation in the respiratory tract. When considering a single infected host, we take the case virulence to mean the peak innate immune activation over the entirety of the acute infection. The analysis produces a range of virulence for different hosts, and we assume that at the more severe extreme of this range, fatality rates will increase and the heterogeneous symptoms and organ dysfunction associated with severe infection. However, we do not propose a direct mapping between our operationalized virulence and clinical outcomes for a specific virus.

Between the variation in open-loop shedding and the variation in immune responses, we have the building blocks to understand virulence and transmission variation more realistically. We define transmission R_{CL}, where $w(t)$ is a warning signal and $(t) = exp(-pw(t))$. Initially, we take $w(t)$ to be a scaled norm of the immune response, so that symptoms promote avoidance and isolation. We use a highly simplified model of avoidance and isolation, emphasizing the consequences of biological variation in the case where behavioral responses are ideal; our goal here is not to understand the epidemiology of transmission, but the effects of immune-viral interactions on transmission.

$$R_{CL} = \alpha_r \omega \int_0^T v(t)(t)dt \tag{4.16}$$

Note that transmission depends both on the infected host and the susceptible host;

a susceptible host with a higher SCP, expressed here as $\alpha_r \omega$, would more likely be infected.

We extend this simple transmission model to address a virus with a long presymptomatic period followed by uniformly severe infection (Figure 4.3A-B). Small variation in α leads to large variation in simulated viral loads. Despite these varied viral loads, advance warning and isolation measures can contain viral spread. However, fully asymptomatic cases (particularly in the absence of widespread and rapid testing) make advance warning more difficult. To illustrate this, we can consider a simple case involving individuals A, B, and C. Suppose A infects five people, including B, while presymptomatic. A goes on to develop severe symptoms. B infects C and no one else. If B goes on to develop severe symptoms (B was presymptomatic), then in an idealized case B can warn C, C can isolate while presymptomatic, and the chain of transmission can be stopped. By contrast, suppose B infects C and no one else, but does not develop symptoms (B was fully asymptomatic), and C goes on to infect five people. The advance warning for C, traced through B, depends on whether B is eventually symptomatic or whether B can be rapidly tested. Crucially, fully asymptomatic cases need not be as contagious as presymptomatic severe cases to prevent advance warning in the absence of testing. At low rates of fully asymptomatic cases, advance warning and isolation without testing can remain effective (Figure 4.3C).

Many respiratory viruses suppress interferon responses. We consider the case where interferon suppression delays the immune response for the entire incubation period; this assumption can be relaxed in more detailed models. Virulence when immune responses are interferon-mediated (intracellular) vary less with α than when immune responses are mediated by neutrophils and macrophages (extracellular), because intracellular responses act on rates of viral synthesis rather than on virions. However, virulence as a result of viral suppression of interferon, inducing delays and increasing the demand for extracellular responses, varies more than either intracellular responses or extracellular responses alone (Figure 4.3D-E). Interferon suppression thus drives atypical variation in virulence when the underlying open-loop dynamics are already different between hosts. Early intervention, even in the presymptomatic stage of infection, can potentially reduce the eventual symptom burden in what would otherwise be severe cases.

Taking these control layers together, we consider virulence and transmission as α varies. We consider a fixed five-day incubation and interferon-suppression period;

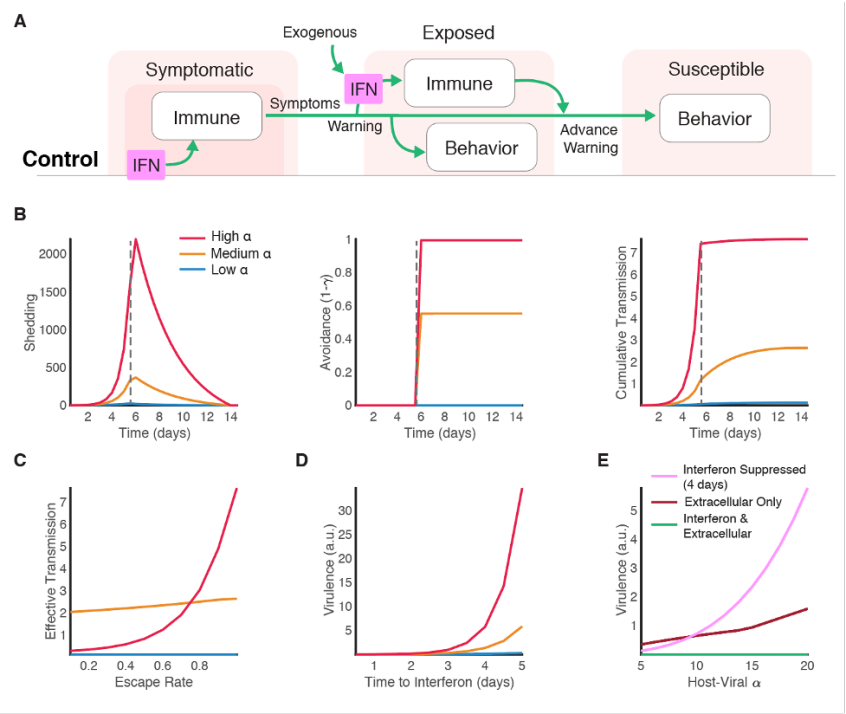

Figure 4.3: Layered control of virulence and transmission. (A) Several biological and behavioral responses are involved in the layered control of virulence and transmission. The challenge addressed in this cartoon is a delayed immune response, which both exacerbates virulence and allows for pre-symptomatic spread. (B) Interferon suppression of the immune response by the virus allows an extended period of viral replication and shedding during which avoidance behaviors are not prompted (without advance warning). Following symptom onset, shown here after five days, high α is associated with a high symptom burden, making high α individuals easiest to avoid (middle panel). Medium α individuals cumulatively transmit about as much over the presymptomatic period as the symptomatic period; however, the cumulative effect of the much higher viral loads in high α individuals in the pre-symptomatic period remains a dominant factor in transmission, while transmission by low α individuals remains low (right panel). (C) Advance warning can reduce the effective transmission rate of presymptomatic individuals, but asymptomatic cases facilitate escape by preventing effective information sharing. The escape rate describes the fraction of cases for a given level of α that cannot be localized by an advance warning (for instance, the fraction of cases directly or closely downstream of a fully asymptomatic transmission in the absence of testing). As the escape rate increases, the effective transmission from presymptomatic high α individuals increases sharply. (D) Timing is a crucial determinant of host and possibly therapeutic interferon efficacy. (E) Virulence mediated by extracellular immune responses (such as neutrophils and macrophages) varies more with α than intracellular immune responses (such as interferon and apoptosis).

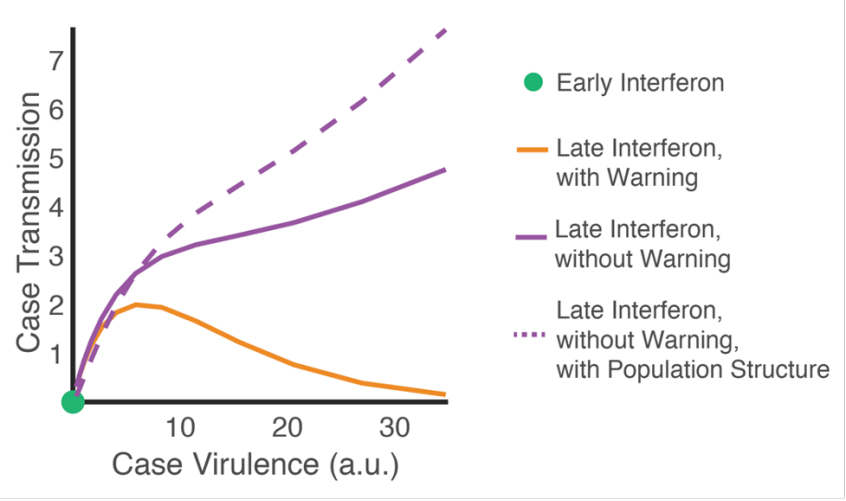

Figure 4.4: Tradeoffs in viral virulence and transmission depend on host control strategies. The relationship between virulence and transmission in individual hosts depends on control strategies such as interferon responses and isolation of contagious individuals. An optimal host immune control response, including timely interferon signaling, quickly blocks viral replication and neither serious symptoms nor substantial transmission occur. With interferon suppression, transmission peaks at low virulence. With interferon suppression and host variation in α, however, transmission is higher and peaks at higher virulence. This effect is amplified when high-α individuals interact primarily with other high-α individuals, leading to both high presymptomatic shedding and high susceptibility. These effects in combination enable the most contagious individuals in the most susceptible subpopulations to transmit at much higher rates (facilitating superspreading).

the additional effects of varied incubation period and strength of interferon suppression can be extrapolated. In this comparison, a rapid interferon-based immune response suppresses both virulence and transmission. A virulence-transmission tradeoff is maintained by the immune response in the presence of interferon suppression. With interferon suppression and host variation, presymptomatic severe high-α cases take a dominant role in spreading the pathogen (Figure 4.4). If the distribution of SCP through a population is structured – if high-α individuals interact more often with high-α individuals than low-α individuals and vice versa, which is likely if α depends on characteristics like age and smoking status – then variation in transmission will be further exaggerated.

4.3 Discussion

We have shown that variation in α and interferon suppression can drive variation in viral virulence and transmission, enabling a severe population burden. Additional

variations in underlying host health factors, innate immune dynamics, and aerosol exhalation may also account for inter-host variation in virulence and transmission [76–79]. We have described cases where interferon responses are identically affected across hosts, emphasizing interactions between virus and host. However, severe cases of respiratory viral infection are also associated with genetic variation in interferon-mediated immunity and with antagonism of the interferon response by the host immune system, suggesting an important role for additional variation in host immune responses in explaining variation in clinical outcomes [79–81]. Previous studies that have considered the evolutionary relationship between pathogen virulence and transmission studies can be revisited in light of our new results. A canonical example of these earlier studies showed, under assumptions coupling replication to virulence and virulence to the duration of infection, that a virulence and transmission tradeoff exists at very high virulence, but some intermediate virulence can be optimal [50]. We derive a similar tradeoff for a scenario based on entirely different specifics: interferon production and symptoms that enable advance warning. Our results are consistent with the spirit of this previous work, in which the authors argued that "depending on the specifics of this linkage, the coevolutionary course can be toward essentially zero virulence, or to very high virulence, or to some intermediate grade." Our mathematical approach allows for more detailed modeling of the within-host dynamics, between-host variation, and the immune response, generalizing the clever coupling of virulence and infection duration of the previous work. While evolutionary arguments emphasize the optimal strategy of pathogens over many generations, a robust control view allows us to consider how the immune response constrains feasible viral strategies even in a single infection as well as through an evolutionary trajectory.

The previous two chapters have been theoretical case studies of multicellular systems motivated by particular domain questions. In order to maintain focus on the domain questions, we have used simple models and special cases in a less formal mathematical style than is typical in control theory, which has limited our discussion of more general problems. In Chapter 5 and Chapter 6, we move into a more formal and abstract register in order to prove theorems about control systems in general that have some of the properties of multicellular control systems. Chapter 5 is more explicitly oriented towards illustrating some principles in biology – the optimality of logarithmic sensing and biased sensing schemes – while Chapter 6 steps away from the biology motivation in order to push the limits of the types of systems that we can describe with these techniques. After these chapters, we will return to biology in our final case study in Chapter 7.

Chapter 5

OPTIMALLY BIASED AND LOG-SCALED SENSING IN CONTROL LOOPS WITH INFORMATION LOSS

5.1 Introduction

Layered control is a type of composition of control systems, common in engineered and biological systems, in which the guarantees provided by one subset of controllers make it possible for another subset of controllers to achieve some other (qualitatively different) objective [82]. In settings where computation and communication are slow or costly, the time needed to achieve accurate state estimates become a major constraint on system performance; as a result, fast, low-bandwidth control loops must co-exist with slow, high-bandwidth control loops. In one recently proposed architecture, a fast loop is responsible for coarse disturbance rejection, while a slow loop is responsible for planning and reference tracking. The two controllers can operate independently, and control action on the plant is straightforwardly additive [83, 84]. The information loss in the fast layer is treated as quantization of the sensed state or the control action.

In real biological control architectures, however, low-layer controllers often receive control signals from higher layers. Low-layer neural and biomolecular systems dynamically re-allocate resources on multiple time-scales: rapidly via inhibitory and excitatory synaptic currents or inactivating and activating biomolecules, or more slowly via the turnover of synaptic and signal-transduction molecules. In

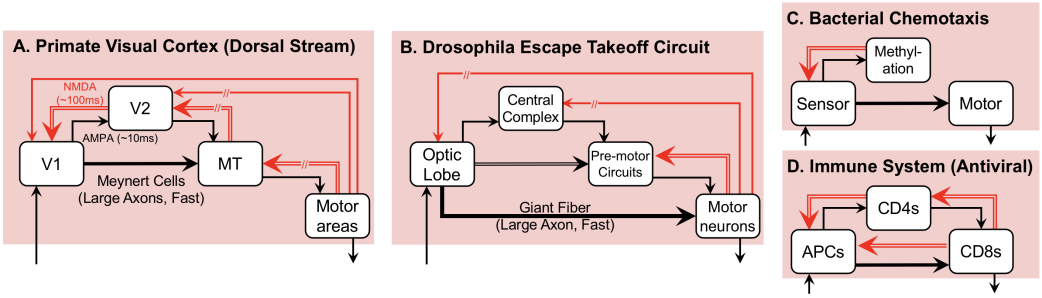

Figure 5.1: Feedback to sensors is a ubiquitous phenomenon in biology: (A) primate vision, (B) fly escape behaviors, (C) chemotaxis in a single bacterial cell and (D) immune responses to viruses.

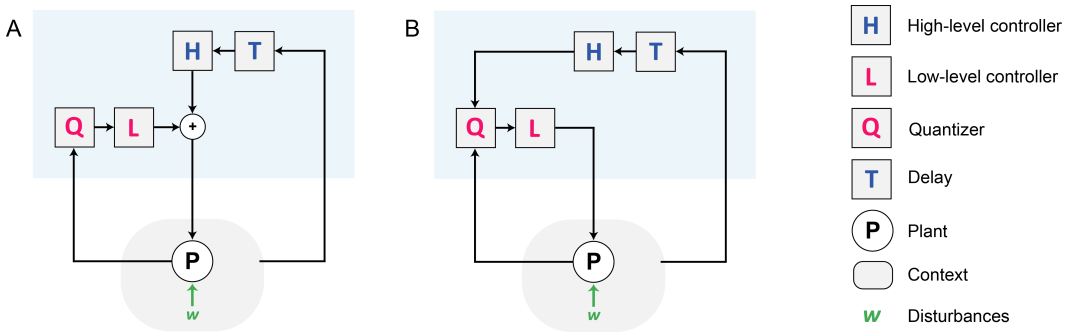

Figure 5.2: Hierarchical control schemes: (A) fully decoupled high-layer and low-layer loops; (B) scheme for incorporating high-layer context information into a low-layer quantized loop.

some cases, such as reflex suppression by conscious override or immune regulation by T helper cells, the high-layer controllers are allowed no direct actuation of the plant, and instead act by modulating parameters of the lower layer of control. In this chapter, we treat such high-layer signals as contextual re-optimization of the low-layer quantizer.

What are the theoretical advantages and limitations of a top-down modulation architecture, given a quantized low-layer control loop? Adaptive quantizers have been shown to be useful in exactly stabilizing plants [85, 86], but these adaptive strategies do not incorporate top-down information. We propose that an advantage of top-down quantizer modulation is that high-layer controllers can instruct low-layer controllers about context, allowing low-layer controllers to be repurposed flexibly across cost specifications. Because computation, communication, and data capacity are expensive in biological settings, biological sensing systems are particularly interesting: a low-layer biological controller's estimate of the state of the world will be lossy due to quantization, stochasticity, or both. We show here that given this constraint, what is sensed is a consequence of what is anticipated.

Low-information, time-sensitive, and context-sensitive control problems are common in biological systems at all scales. Individual cells under stress decide whether or not to self-destruct to protect their neighbors; stem cell populations skew towards different cell fates depending on demands such as wound healing; the nervous system decides whether to flee an approaching threat. Each of these control systems can be modulated by top-down mechanisms, but remain subject to the approximation losses inherent to the components.

Because the consequences of approximation in biology are asymmetric (prey ani-

mals are inconvenienced by false-positive detection of predators, but greatly more so by false negatives), we investigate here how the optimal approximation depends on context. While contextual flexibility is widely observed in biology and neuroscience, rigorous theory relating this flexibility to its component-level implementation remains elusive. This in turn leads to confusion about fundamental phenomena like plasticity and major diseases categories like autoimmunity and neuropsychiatric disorders. A rich tradition in control theory studies the inevitable tradeoffs resulting from given system and component constraints [3]. These tradeoffs are likely to underlie both the remarkable capabilities of biological systems and their evident limitations.

With simplifying assumptions, but crucially preserving the limitations on accuracy that are an inescapable feature of biological sensors and actuators, we take steps towards a rigorous characterization of processes by which high-bandwidth control can tune low-bandwidth control. We study a scalar quantized control loop which receives top-down instructions that set its control strategy. In this chapter, we show that the optimal quantization scheme depends on dynamics and costs, using an approach that recovers both logarithmic and uniform quantization schemes that have been reported previously [83, 87]. We interpret these results in relation to theories of biological learning, plausible mechanisms by which quantized sensors can be tuned to match context, and inevitable tradeoffs consequent to approximation.

5.2 Results

Consider a scalar discrete-time system with quantization on the sensed state $x_t \in \mathbb{R}$, control action $u_t \in \mathbb{R}$, and bounded disturbances $w_t \in W \subset \mathbb{R}$.

$$\begin{aligned} x_{t+1} &= ax_t + u_t + w_t \\ u_t &= k\mathbf{Q}_R(x_t) \end{aligned} \tag{5.1}$$

We seek to design the quantization intervals of the fixed R-bit quantizer $\mathbf{Q}_R$, given a fixed gain k, that minimizes costs. We assume that k stabilizes the infinite-accuracy case. For some $L \in \mathbb{R}_{++}$, the interval $[-L, L]$ will be partitioned into $2^{R-1} + 1$ sub-intervals at boundaries indicated by δ_j with $j \in [1 \ldots 2^{R-1}]$.

Formally, we propose that the cost due to quantization can be measured as the *control approximation loss*, the excess control cost paid that would have been saved with perfect accuracy, and the *state approximation loss*, the excess state cost paid.

We allow any cost function that takes the form $J(x, u)$ as a convex composition of $J_x(x_t)$ and $J_u(u_t)$ where J_x and J_u are nonnegative $\mathbb{R} \to \mathbb{R}$ convex costs on state and control, respectively. If we write the quantized system as:

$$x_{t+1} = ax_t + kx_t + k\eta(x) + w_t \tag{5.2}$$

with $\eta(x)$ giving the error due to approximation for a given true value of x, then the perfect-accuracy system is simply:

$$x_{t+1} = ax_t + kx_t + w_t \tag{5.3}$$

The loss due to quantization can then be given by the cost spent minus the ideal cost spent:

$$\begin{aligned} \psi_x(t) &= J_x(ax_t + kx_t + k\eta(x) + w_t) \\ &\quad -J_x(ax_t + kx_t + w_t) \\ \psi_u(t) &= J_u(kx_t + k\eta(x)) - J_u(kx_t) \end{aligned} \tag{5.4}$$

Suppose we design the partitions in $\mathbf{Q}_R$ to minimize the worst-case loss due to quantization, $\|\psi\|_\infty = \|\sup\{\psi_x(t), \psi_u(t)\}\|_\infty$.

Theorem 1. The partition of $[-L, L]$ minimizing $\|\psi\|_\infty$ will be a partition for which the worst-case loss due to quantization over each segment is equal.

Proof. Suppose a partition giving equal worst-case approximation loss over each segment is known and given by $\{\bar{\delta}_j\}$. For the sake of contradiction, suppose the partition given by $\{\bar{\delta}_j\}$ has a strictly lower worst-case approximation cost than $\{\delta_j^*\}$:

$$\sup_j \psi(\bar{\delta}_j) < \sup_j \psi(\delta_j^*) \tag{5.5}$$

For this to be true, each interval would need to reduce its approximation costs over a function J. The worst-case approximation cost over a given interval is given by:

$$\begin{aligned} \|\psi_j\|_\infty = \sup_{w \in W,\ \eta(x)} &\{J(ax_t, kx_t, k\eta, w_t) \\ &-J(ax_t, kx_t, 0, w_t)\} \end{aligned} \tag{5.6}$$

In order to reduce the worst-case approximation cost $\|\psi_j\|_\infty$ across a given interval, it is necessary to decrease $\|\delta_j - \delta_{j+1}\|$. This then necessitates increasing $\|\delta_k - \delta_{k+1}\|$ for some other interval that will include at least $[\delta_k^*, \delta_{k+1}^*)$, so its worst-case

approximation cost will at best equal the worst-case approximation cost associated with $\{\delta^*\}$:

$$\sup_j \psi(\bar{\delta}_j) \geq \sup_j \psi(\delta_j^*) \tag{5.7}$$

which gives the contradiction. We must also show that a partition with all interval costs equal is achievable; this follows from convexity of J. Note that a given partition is not unique in the case where J is constant. However, if J is convex and anywhere non-constant over $[-L, L]$, the bound becomes strict and the partition is guaranteed to be unique.

Special cases

The optimal partition problem $\inf_{\{\delta\}} \|\psi_j\|_\infty$ is not in general convex. However, closed-form or even numerical statements of the optimal partition are not necessary to study interesting features of the optimal quantizer across different cost functions. Though motivated by a two-layer control system, we will treat the high-layer controller as static for each of the three cost functions analyzed, then briefly address some implications of delay in the high-layer controller.

Quadratic case. We first consider the case of quadratic costs on state x with no cost on control: that is, $J_x = x^2$ and $J_u = 0$. While this problem formulation is not identical in some technical details to that posed in [87], we seek to verify that our formulation recaptures the key claim: namely, that optimal quantization intervals should expand geometrically with distance from the origin. Let $\bar{q}_j$ be the approximation loss on an interval with boundaries δ_j and δ_{j+1}. The approximation loss related to this interval for a given x is:

$$\bar{q} = (ax + kx + k\eta + w)^2 - (ax + kx + w)^2 \tag{5.8}$$

Set $w = 0$ to simplify expressions. The case with non-zero w follows the same steps and arrives at a comparable result. We solve the expression for $\hat{\delta_j}$ by expanding and applying the quadratic formula.

$$\eta = \frac{-2k(a+k)x \pm \sqrt{(2k(a+k)x)^2 + 4k^2\bar{q}}}{2k^2} \tag{5.9}$$

Because the minimax-optimal solution for δ will require that each worst-case $\bar{q}$ is equal, we can analyze the solution with $\bar{q}$ fixed. Observe that k and a are also fixed. Because $\bar{q}$ is positive, the radical is real for all x. The key question is therefore how

$\|\eta\|_\infty$ changes with x. It is apparent from (5.9) that this rate of change is sublinear near the origin and $O(x)$ with greater distance. If for a sequence δ the interval width expands linearly with δ, then δ itself grows geometrically. Notably, such geometric quantization is observed in muscle recruitment [84], reflecting quantizer tuning on an evolutionary timescale.

A notable exception to this rule is the case where $a = -k$, in which instabilities are canceled perfectly and instantaneously. In this case, the expression (5.9) simplifies considerably to $\eta = \frac{\sqrt{\bar{q}}}{|a|}$, and optimal intervals will be uniform and independent of x, as in [83, 84]. Intuitively, the control action in this case will always result in an interval-bounded over- or under-shoot of the origin, so the cost of errors will be the same regardless of the starting point.

Piecewise-quadratic case. The finite-bit, fixed-width formulation enforces controllers that never fully converge to a stable fixed point [88]. However, for many applications, an invariant set around the set point is adequate. While in the robust linear case an invariant set guarantee also guarantees stability at the origin, the nonlinearity resulting from quantization precludes this equivalence. Therefore, a piecewise cost function, with costs zero and flat over a region around the origin and increasing after, is of interest. At the other extreme, finite-bit quantization necessitates saturation, so any state outside the bounds of the quantizer $[-L, L]$ will be unstabilizable for an unstable plant. Motivated by [87] and problems in nonlinear control [89], we consider $J_u = 0$ and a piecewise state cost function composed of a no-cost safe region and a quadratic-cost dangerous region.

$$J_x = \begin{cases} 0 & \|x\| \leq \alpha \\ (\|x\| - \alpha)^2 & \alpha < \|x\| \leq L \end{cases} \tag{5.10}$$

In this case, the minimax-optimal sensing scheme ignores the safety region up until the point where the next time-step threatens danger; sensing is most precise at the transition from safety to danger, losing resolution further from safety.

Asymmetric case. We next consider a case with asymmetric costs over the domain $[-L, L]$:

$$\begin{aligned} J_x &= \frac{1}{(L+1) - x} \\ J_u &= \frac{|\alpha|}{(L+1) - |k|x} \end{aligned} \tag{5.11}$$

This cost asymmetry corresponds to an asymmetry in many realistic scenarios. For instance, as previous studies have suggested [90], it is far more important when

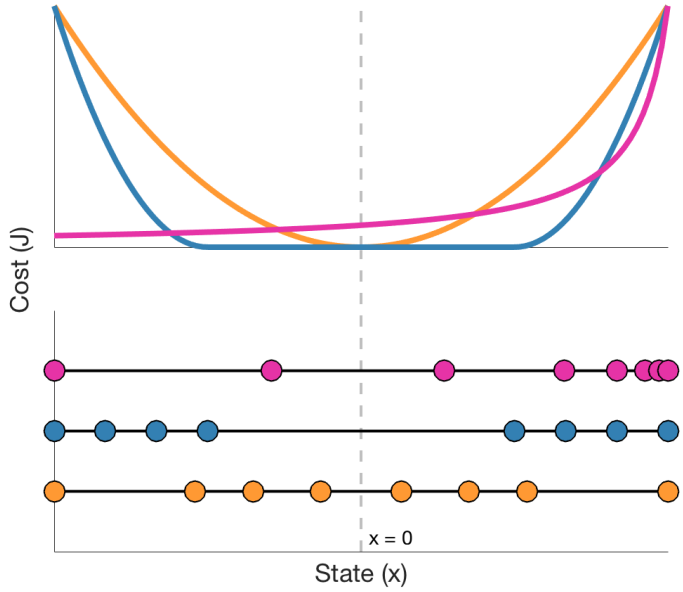

Figure 5.3: Optimal quantizer design depends on cost functions. Where is accurate sensing most valuable? From top to bottom, a quadratic case, a piecewise quadratic case with flat costs near the origin, and an asymmetric case. Quantizer partitions shown are numerical approximations with parameters $a = 1.25$, $L = 10$, and $k = -1$.

driving near a cliff to stay away from the hazard than to stay centered in the lane. When detecting dangerous predators or pathogens, false positives can be costly, but false negatives are potentially fatal.

In this case, minimax-optimal interval spacing will skew the limited number of sensors towards L. Moreover, for a given interval, the $\mathbf{Q}_R(x)$ that minimizes approximation cost will always be greater than x, because there is no state penalty for overshooting; the penalty instead comes from overpaying control costs. Unlike in the quadratic case, sensor resolution will increase closer to L. While a monotonic cost function does allow the system to blow up in the safe direction, a more realistic interpretation follows if we take the system to be open-loop stable, or allow the state to saturate in the state direction. In both the piecewise-quadratic and the asymmetric case, a desirable nonlinear behavior is achieved by a fixed linear control law via manipulation of the quantizer.

Interestingly, the quantized representation of the plant state that the controller uses (i.e. the representation that would be available to another observer in the system,

or to an experimenter) is systematically biased, given the cost functions (5.11), towards higher costs. As α approaches zero, the worst-case approximation cost also approaches zero, and the intervals collapse onto each other: it becomes optimal to always assume the hazard is imminent.

5.3 Discussion

Context-switching interpretation. In the above analyses, we assume that contextual changes occur on time scales irrelevant to control. When this is not true – i.e. when contextual changes occur faster than the high-layer control can compensate – the low-layer quanztizer can be severely suboptimal. For instance, if the asymmetric-cost case discussed here suddenly became the quadratic-cost case, then the low-layer controller optimized for the asymmetry would continue to tend systematically away from the origin; this uncorrected quantizer would be costly and potentially catastrophic.

Component-level interpretation. The control strategy described can be plausibly implemented on electronic or biological components. The electronic case is simplest and builds intuition: we treat the quantization step as analog sensed information passed through comparators before being transmitted to the controller. The comparators' reference signals can then be modulated by top-down analog signals based on goals. A comparator implementation extends naturally to both neural and biomolecular settings. In the neural case, a sensing neuron's threshold-to-fire can be regulated by inhibitory or excitatory synaptic inputs from top-down projections. In the biomolecular case, sequestration, inactivation, or degradation can enact comparison-like functionality. A particularly clear molecular implementation is given by mass-action Hill kinetics: re-allocation of the sensitive, linear range of the sigmoid is achievable by shifts of the curve, for instance via allosteric interactions [91].

Chapter 6

ROBUST PERFORMANCE WITH STRUCTURED UNCERTAINTIES IN SYSTEM LEVEL SYNTHESIS

6.1 Introduction

Robust control problems consider the gap between controller performance for idealized models of a system and controller performance in real-world settings. In the worst case, the gap between model and reality can create catastrophic instabilities in engineered systems [92, 93]. The robust controller *synthesis* problem is to synthesize a controller that guarantees robust stability in spite of uncertainties in the system model. The robust *performance* problem is to synthesize a controller that guarantees robust stability and to further compute and guarantee the worst-case bound on the performance cost under bounded disturbances and uncertainties in the dynamics. While robust control in the centralized linear setting is well-studied [6, 94, 95], the generalization of older optimal and robust control results to decentralized systems is a topic of more recent research [96–104].

In this chapter, we leverage the System Level Synthesis (SLS) framework [98, 105]. SLS reformulates robust and optimal control problems as an optimization over the achievable closed loop behavior, or *system responses*, of a linear-time-invariant (LTI) dynamical system, and in particular shows that it is necessary and sufficient to constrain these system responses to lie in an affine subspace defined by the dynamics. This parameterization has been successfully exploited in the context of the distributed optimal control of finite-dimensional LTI systems to scale controller synthesis and implementation techniques to systems of arbitrary size under practically realistic assumptions on the underlying system [106–108].

We connect the SLS framework to $\mathcal{L}_1$ robust synthesis techniques [6, 94] and derive necessary and sufficient conditions for robust performance in terms of *affine* constraints on the system response variables in the settings of both unstructured and structured uncertainty. These necessary and sufficient conditions are equally applicable when additional delay, sparsity, and locality constraints are imposed on the system responses and controller implementation. To the best of our knowledge, these are the first such necessary and sufficient conditions for robust performance that

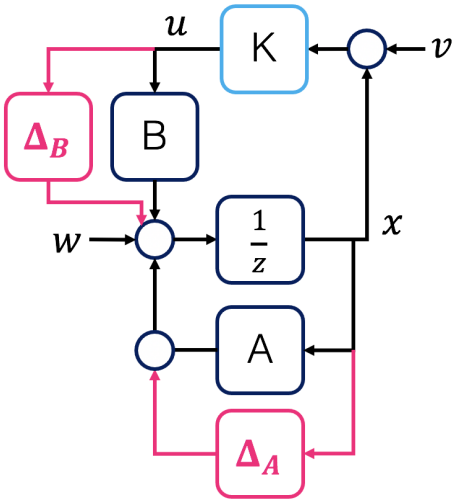

Figure 6.1: Block diagram of parametric uncertainty in a closed loop state feedback problem. The plant is shown in dark blue, and the controller in light blue, which represent the nominal problem. The uncertain loops are shown in pink.

can be verified at scale, and that produce a distributed controller with a corresponding scalable and distributed realization, making these results and their corresponding computational tools naturally applicable to large-scale uncertain systems. While the applications to decentralized control are the most interesting, the synthesis procedure also applies to the synthesis of centralized robust controllers.

System Level Synthesis

We present brief background on System Level Synthesis and $\mathcal{L}_1$ robust control. More detailed treatments of SLS can be found in [7, 109, 110] and more detailed treatments of $\mathcal{L}_1$ robustness in [6, 111, 112]. An initial result towards SLS robust performance can be found in [25].

In discrete-time, lightly modifying the tutorial notation in [109], we consider the following linear block-matrix dynamics:

$$\begin{bmatrix} x[t+1] \\ \bar{y}[t] \\ y[t] \end{bmatrix} = \left[\begin{array}{c|cc} A & B_w & B \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{array}\right] \begin{bmatrix} x[t] \\ w[t] \\ u[t] \end{bmatrix} \tag{6.1}$$

Here x are states, w disturbances, and u control inputs provided by an external controller. $\bar{y}$ and y are the measured and regulated output, respectively. Each of these variables can be a vector with compatible dimensions in the transfer functions

($x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$). The partitioned matrix dynamics in (6.1) can be written more compactly as a partitioned transfer function, taking the z-domain representation and consequent forward-shift operator z.

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \\ \mathbf{P}_{21} & \mathbf{P}_{22} \end{bmatrix} \tag{6.2}$$

$$\mathbf{P}_{ij} = C_i(zI - A)^{-1}B_j + D_{ij} \tag{6.3}$$

For an appropriate norm, an optimization problem can then be posed for the controller $\mathbf{K}$. (In subsequent usage in this chapter, we will always consider the $\|\cdot\|_{\ell_\infty\to\ell_\infty}$ norm unless otherwise specified.)

$$\begin{aligned} \underset{\mathbf{K}}{\text{minimize}} \quad & \|\mathbf{P}_{11} + \mathbf{P}_{12}\mathbf{K}(I - \mathbf{P}_{22})\mathbf{K})^{-1}\mathbf{P}_{21}\| \\ \text{subject to} \quad & \mathbf{K} \text{ internally stabilizes } \mathbf{P} \end{aligned} \tag{6.4}$$

This problem is non-convex in $\mathbf{K}$. In the Youla parameterization, the optimization of the controller is restated as optimization of a particular transfer function. The Youla problem is convex in this transfer function and an optimal $\mathbf{K}$ can be computed in terms of co-prime factors of $\mathbf{P}_{22}$. Both the computational and theoretical power of this parameterization are foundational.

SLS generalizes the Youla approach, parameterizing the same space of controllers but more easily accomodating sparsity, locality, and delay constraints [7]. In the state feedback case, we let $B_w = I$ in the dynamics (6.1), and restate synthesis in terms of the closed-loop maps $\Phi_x : w \to x$ and $\Phi_u : w \to u$ and the plant dynamics:

$$\begin{aligned} \underset{\mathbf{\Phi}}{\text{minimize}} \quad & \| \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix} \| \\ \text{subject to} \quad & \begin{bmatrix} zI - A & -B \end{bmatrix} \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix} = I \\ & \mathbf{\Phi}_x, \mathbf{\Phi}_u \in \frac{1}{z}\mathcal{RH}_\infty \end{aligned} \tag{6.5}$$

Here $\mathcal{RH}_\infty$ refers to the set of real-rational proper transfer functions, i.e. those are causal and implementable. $\frac{1}{z}\mathcal{RH}_\infty$ refers to the strictly proper set of the same. While these closed loop maps are in general infinite-dimensional, we will impose a finite-impulse response approximation on the optimization problem to ensure stability and computability. Other approximation bases can also be used. We show that this

problem can be posed as a linear program by reformulation in the $\ell_\infty \to \ell_\infty$ norm, and can therefore be solved quickly with typical computational methods. For large-scale problems, (6.5) has a partial separability structure amenable to computational speedups through iterative methods [113–115]. Because the controller associated with SLS is composed only of the closed-loop maps, any decentralization constraints on the closed-loop are also imposed on the controller (Fig. 6.2). This structure-preserving property will hold for the extensions and reformulations we describe below.

$\mathcal{L}_1$ Robust Control

We use results from $\mathcal{L}_1$ robust control as developed in [6, 111, 112] and related papers. The $\mathcal{L}_1$ setting uses the $\ell_\infty \to \ell_\infty$ induced norm on multivariate transfer functions, making it compatible with a finite-impulse approximation in SLS. Once the transfer function is configured appropriately as a matrix, the $\ell_\infty \to \ell_\infty$ norm can be computed conveniently as the maximum row sum of that matrix.

The key result we will use throughout this chapter concerns linear fractional transformations (LFTs). Specifically, if we want to ascertain that a particular coupling of a map $\mathbf{\Phi}$ and uncertainty $\mathbf{\Delta}$ (which is bounded in the $\ell_\infty \to \ell_\infty$ norm and structured in ways we describe below). Then the following structure which arises from the feedback interconnection:

$$M_{11} + M_{12}\Delta(I - M_{22}\Delta)^{-1}M_{12} \tag{6.6}$$

is stable for all admissible Δ if and only if the following augmented system is robustly stable.

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \tag{6.7}$$

Robust stability is certified by the existence of a diagonal matrix with positive entries D (with additional constraints that depend on Δ, discussed below) such that:

$$\|D^{-1}\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} D\|_{\ell_\infty \to \ell_\infty} < 1 \tag{6.8}$$

This linear fractional transformation, combined with a finite-impulse approximation, allows feedback-interconnected systems with uncertain Δs to be analyzed through simple-to-compute matrix operations. The $\ell_\infty \to \ell_\infty$ norm is a natural starting-point for robust performance in SLS because it can be computed row-wise, which gives

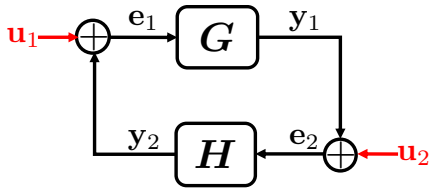

Figure 6.2: A feedback interconnection between systems **G** and **H**.

it favorable properties in large-scale systems [25, 114]. It is also a natural norm in which to consider challenges like quantization, which are encountered in large-scale systems. The $\ell_\infty \to \ell_\infty$ stability guarantee is exact, meaning it is the least conservative possible guarantee for systems with nonlinear or time-varying Δs with arbitrarily many blocks. However, the $\ell_\infty \to \ell_\infty$ norm is somewhat conservative with respect to disturbances when compared to $\mathcal{H}_\infty$ and especially μ synthesis [113, 116].

We note that computing a suitable M and D pair is typically an iterative process [111]. Although the two steps can be posed as a linear programming step and a spectral radius computation, respectively, the two variables cannot be optimized jointly. The spectral radius computation may make localized design challenging [113], and the iteration can get caught in local minima, although these local minima have in practice been handled with randomization.

Notation

One usage of nonstandard notation will facilitate the exposition. $[Q\mathbf{\Phi}]_{ij}$ will refer to the ij-th block in the partition of $Q\mathbf{\Phi}$. $[Q\mathbf{\Phi}]_{:j}$ will refer to the j-th column including all rows and $[Q\mathbf{\Phi}]_{i:}$ will refer to the i-th row including all columns.

6.2 Results

Robust Operator System Level Synthesis

Theorem 6.1[1]. Let $\boldsymbol{A} \in \mathcal{L}_{\mathrm{TV}}^{n,n}$ and $\boldsymbol{B} \in \mathcal{L}_{\mathrm{TV}}^{n,p}$, and suppose that $\left\{\hat{\mathbf{\Phi}}_x, \hat{\mathbf{\Phi}}_u\right\}$ satisfy

$$\begin{bmatrix} I - \boldsymbol{S}_+\boldsymbol{A} & -\boldsymbol{S}_+\boldsymbol{B} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{\Phi}}_x \\ \hat{\mathbf{\Phi}}_u \end{bmatrix} = \boldsymbol{S}_+(I - \mathbf{\Delta}), \tag{6.9}$$
$$\mathbf{\Phi}_x, \mathbf{\Phi}_u \text{ strictly causal, linear, and } \ell_\infty\text{-stable},$$

for $\mathbf{\Delta}$ a strictly causal linear operator from $\ell_{\infty,e}^n \to \ell_{\infty,e}^n$. Then the controller implementation defined in terms of the operators $\left\{\hat{\mathbf{\Phi}}_x, \hat{\mathbf{\Phi}}_u\right\}$ is well posed and achieves the

[1] This theorem and its proof were developed by Nikolai Matni and are reprinted from [25].

following response

$$\begin{bmatrix} \mathbf{x} \\ \boldsymbol{u} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} (I - \boldsymbol{\Delta})^{-1} \mathbf{w}. \tag{6.10}$$

Further, this interconnection is ℓ_∞-stable if and only if $(I - \boldsymbol{\Delta})^{-1}$ is ℓ_∞-stable.

Proof. As $\boldsymbol{\Delta}$ is strictly causal by assumption, $I_{\boldsymbol{\Delta}} := (I - \boldsymbol{\Delta})^{-1}$ exists as a map from $\ell^n_{\infty,e} \to \ell^n_{\infty,e}$. Going through a similar argument as that in the proof of Proposition 6.1, we observe that

$$\begin{bmatrix} \mathbf{x} \\ \boldsymbol{u} \\ \hat{\boldsymbol{w}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Phi}_x I_{\boldsymbol{\Delta}} & \boldsymbol{\Phi}_x I_{\boldsymbol{\Delta}}(-\boldsymbol{A}) & \boldsymbol{\Phi}_x I_{\boldsymbol{\Delta}} \boldsymbol{B} \\ \boldsymbol{\Phi}_u I_{\boldsymbol{\Delta}} & \boldsymbol{\Phi}_u I_{\boldsymbol{\Delta}}(-\boldsymbol{A}) & I + \boldsymbol{\Phi}_u I_{\boldsymbol{\Delta}} \boldsymbol{B} \\ \boldsymbol{S}_+ I_{\boldsymbol{\Delta}} & I_{\boldsymbol{\Delta}}(I - \boldsymbol{S}_+ \boldsymbol{A}) & I_{\boldsymbol{\Delta}} \boldsymbol{S}_+ \boldsymbol{B} \end{bmatrix} \begin{bmatrix} \boldsymbol{w} \\ \boldsymbol{\delta}_y \\ \boldsymbol{\delta}_u \end{bmatrix}. \tag{6.11}$$

Thus we see that the desired map (6.10) from $\mathbf{w} \to (\mathbf{x}, \boldsymbol{u})$ is achieved. Further, as $\boldsymbol{\Phi}_x$, $\boldsymbol{\Phi}_u$, $\boldsymbol{A}$, $\boldsymbol{B}$ are all ℓ_∞-stable by assumption, it follows that the ℓ_∞-stability of the map from $(\boldsymbol{w}, \boldsymbol{\delta}_y, \boldsymbol{\delta}_u) \to (\mathbf{x}, \boldsymbol{u}, \hat{\boldsymbol{w}})$ is determined by the ℓ_∞-stability of $I_{\boldsymbol{\Delta}}$, from which the result follows.

We now use the tools developed in the previous section to identify necessary and sufficient conditions for the robust stability and robust performance of a system subject to bounded perturbations in its $\boldsymbol{A}$ and $\boldsymbol{B}$ operators. In particular consider the system

$$\boldsymbol{x} = \boldsymbol{S}_+(\boldsymbol{A}_0 + \boldsymbol{\Delta}_{\boldsymbol{A}})\mathbf{x} + \boldsymbol{S}_+(\boldsymbol{B}_0 + \boldsymbol{\Delta}_{\boldsymbol{B}})\boldsymbol{u} + \boldsymbol{S}_+ \mathbf{w}, \tag{6.12}$$

where $\boldsymbol{A}_0 = \text{blkdiag}(\hat{A}, \hat{A}, \dots)$ and $\boldsymbol{B}_0 = \text{blkdiag}(\hat{B}, \hat{B}, \dots)$ are memoryless LTI operators defining a nominal LTI system $x_{t+1} = \hat{A}x_t + \hat{B}u_t + w_t$, and $\boldsymbol{\Delta}_{\boldsymbol{A}}$ and $\boldsymbol{\Delta}_{\boldsymbol{B}}$ are ℓ_∞-stable and satisfy

$$\|[\boldsymbol{\Delta}_{\boldsymbol{A}}, \boldsymbol{\Delta}_{\boldsymbol{B}}]\|_{\ell\infty \to \ell_\infty} \leq \varepsilon. \tag{6.13}$$

We first identify necessary and sufficient conditions for robust stability, and then build upon those to formulate a robust performance problem. Define $\hat{\boldsymbol{\Delta}}$:

$$\hat{\boldsymbol{\Delta}} := \boldsymbol{S}_+ \begin{bmatrix} \boldsymbol{\Delta}_{\boldsymbol{A}} & \boldsymbol{\Delta}_{\boldsymbol{B}} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix}, \tag{6.14}$$

is a strictly causal ℓ_∞-stable operator, we conclude that the controller implementation defined in terms of the LTI operators $\{\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u\}$ achieves the following closed loop behavior when applied to the uncertain dynamics (6.12):

$$\begin{bmatrix} \mathbf{x} \\ \boldsymbol{u} \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} (I - \hat{\boldsymbol{\Delta}})^{-1}. \tag{6.15}$$

This control law is internally stabilizing if and only if $(I - \hat{\boldsymbol{\Delta}})^{-1}$ is ℓ_∞-stable.

Defining the controlled output signal as

$$z = \boldsymbol{C}\mathbf{x} + \boldsymbol{D}\boldsymbol{u}, \tag{6.16}$$

for $\boldsymbol{C} = \text{blkdiag}(C, C, \dots)$ and $\boldsymbol{D} = \text{blkdiag}(D, D, \dots)$ user specified cost matrices,[2] and consider the goal of minimizing the $\ell_\infty \to \ell_\infty$ induced gain from $\boldsymbol{w} \to \boldsymbol{z}$ of the uncertain system (6.12). We can then pose the robust performance problem for a specified performance level $\gamma \geq 0$ as finding LTI operators $\{\hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u\}$ that satisfy

$$\begin{aligned}
&\begin{bmatrix} C & D \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} \left(I - \boldsymbol{S}_+ \begin{bmatrix} \boldsymbol{\Delta_A} & \boldsymbol{\Delta_B} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} \right)^{-1} \leq \gamma \\
&\begin{bmatrix} zI - \hat{A} - \hat{B} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} = I, \ \hat{\boldsymbol{\Phi}}_x, \hat{\boldsymbol{\Phi}}_u \in \frac{1}{z}\mathcal{RH}_\infty, \\
&\left(I - \boldsymbol{S}_+ \begin{bmatrix} \boldsymbol{\Delta_A} & \boldsymbol{\Delta_B} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix} \right)^{-1} \text{ is } \ell_\infty\text{-stable}
\end{aligned} \tag{6.17}$$

for all $(\boldsymbol{\Delta_A}, \boldsymbol{\Delta_B})$ satisfying bound (6.13), where we have combined equations (6.16) and (6.15) to derive the robust performance bound condition.

To lighten notation going forward, we let

$$\boldsymbol{Q} := \begin{bmatrix} \boldsymbol{C} & \boldsymbol{D} \end{bmatrix}, \ \hat{\boldsymbol{\Phi}} := \begin{bmatrix} \hat{\boldsymbol{\Phi}}_x \\ \hat{\boldsymbol{\Phi}}_u \end{bmatrix}, \ \boldsymbol{\Delta} := \frac{1}{\varepsilon}\boldsymbol{S}_+ \begin{bmatrix} \boldsymbol{\Delta_A} & \boldsymbol{\Delta_B} \end{bmatrix}, \ Z_{AB} := \begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix}. \tag{6.18}$$

With this notation, the robust performance problem (6.17) is equivalent to finding an LTI operator $\hat{\boldsymbol{\Phi}}$ satisfying

$$\begin{aligned}
&\frac{1}{\gamma}\boldsymbol{Q}\hat{\boldsymbol{\Phi}} + \frac{1}{\gamma}\boldsymbol{Q}\hat{\boldsymbol{\Phi}}\boldsymbol{\Delta}(I - (\varepsilon\hat{\boldsymbol{\Phi}})\boldsymbol{\Delta})^{-1}(\varepsilon\hat{\boldsymbol{\Phi}}) \leq 1 \\
&Z_{AB}\hat{\boldsymbol{\Phi}} = I, \ \hat{\boldsymbol{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty, \ (I - \hat{\boldsymbol{\Phi}}\boldsymbol{\Delta})^{-1} \text{ is } \ell_\infty\text{-stable}
\end{aligned} \tag{6.19}$$

for all $\boldsymbol{\Delta}$ satisfying $\boldsymbol{\Delta} \leq 1$, where we have used that $(I - \Delta\hat{\boldsymbol{\Phi}})^{-1} = I + \Delta(I - \hat{\boldsymbol{\Phi}}\Delta)^{-1}\hat{\boldsymbol{\Phi}}$ and that $(I - GH)^{-1}$ is ℓ_∞-stable if and only if $(I - HG)^{-1}$ is ℓ_∞-stable (see Proposition 1, [6]) to recast the expression (6.17) in a form that matches the linear-fractional-transform (LFT) structure studied in [6, 94].

[2] For simplicity, we assume $\boldsymbol{C}$ and $\boldsymbol{D}$ to be memoryless and LTI, however our results are equally applicable when the controlled output is defined in terms of LTI filters $\boldsymbol{C}(z)$ and $\boldsymbol{D}(z)$.

State feedback with unstructured uncertainty

The necessary and sufficient conditions for robust stability of the resulting two-block problem can be derived as a special case of Theorem 6.3 of [6]. The particular case of an augmented LTI system $\boldsymbol{M}$ satisfying $\boldsymbol{M}_{11} = \boldsymbol{M}_{12}$ and $\boldsymbol{M}_{21} = \boldsymbol{M}_{22}$, as is the case for our problem, is addressed in Ch 8.3 of [6], where a similarly structured augmented system arises in the context of bounding output sensitivity in the presence of output perturbations. The necessary and sufficient conditions specified in Theorem 6.3 of [6] reduce to the following *convex* constraints on the system response $\hat{\boldsymbol{\Phi}}$

$$\begin{aligned} &Z_{AB}\hat{\boldsymbol{\Phi}} = I,\ \hat{\boldsymbol{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty, \\ &\boldsymbol{Q}\hat{\boldsymbol{\Phi}} + \gamma\varepsilon\hat{\boldsymbol{\Phi}} < \gamma. \end{aligned} \tag{6.20}$$

Although the constraints (6.20) are in general infinite-dimensional due to the transfer matrix $\hat{\boldsymbol{\Phi}}$, principled finite-dimensional approximations, some of which enjoy provable sub-optimality guarantees, are available [105, 117–119]. Further, for the $\mathcal{L}_1$ problem considered here, the resulting optimization problem can be posed as a linear program, thus enjoying favorable computational complexity properties. It then follows that by bisecting on γ, e.g., by using golden search, we can find a performance level γ, and corresponding system responses and controller, satisfying $\gamma \leq \gamma_\star + \epsilon$ in $O\log_2(1/\epsilon)$ iterations, for $\gamma_\star$ the smallest γ such that the set defined by (6.20) is non-empty.

Convex unstructured state feedback

Consider again the partitioned system:

$$\begin{bmatrix} \frac{1}{\gamma}\boldsymbol{Q}\boldsymbol{\Phi} & \frac{1}{\gamma}\boldsymbol{Q}\boldsymbol{\Phi} \\ \varepsilon\boldsymbol{\Phi} & \varepsilon\boldsymbol{\Phi} \end{bmatrix} \tag{6.21}$$

Above, following [25], we used the special structure of the system to use simplified conditions given by [6]. Now we consider the case with more blocks, specifically blocks that affect different subsystems of the plant (e.g. block diagonal parametric uncertainty on $\boldsymbol{\Delta_A}$), to represent phenomena like demand surges in a power grid or injuries in an organism). For a system $\boldsymbol{\Phi}$ with structured perturbations, [6] gives two equivalent conditions for robust stability.

1. Take a block-partitioned system $\boldsymbol{G}$ with subsystems $\boldsymbol{G}_{ij}$ and construct a matrix M such that $m_{ij} = \|\boldsymbol{G}_{ij}\|_{\ell_\infty \to \ell_\infty}$. $\boldsymbol{G}$ is robustly stable if and only if $\rho(M) < 1$.

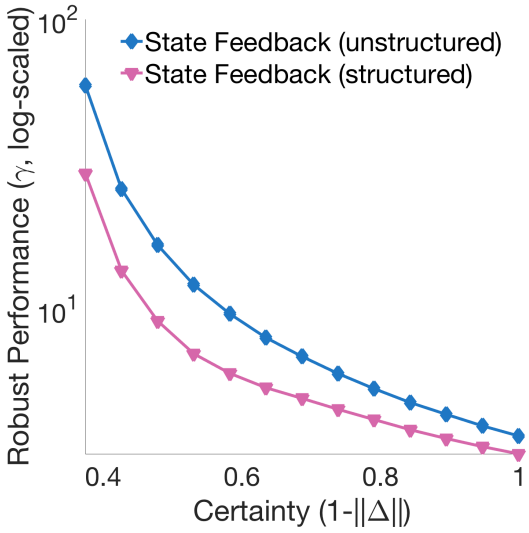

Figure 6.3: Robust performance on a 6-state ring ($\alpha = 1.5$) with structured disturbances and state feedback, with every other state actuated, computed with a linear program. For large Δ (towards the left), both problems eventually become infeasible, but the structured problem outperforms the unstructured until then.

2. Define a positive real diagonal matrix D that matches the partition structure of $\boldsymbol{G}$ and $\boldsymbol{\Delta}$. $\boldsymbol{G}$ is robustly stable if and only if $\inf_D \|D^{-1}\boldsymbol{G}D\|_{\ell_\infty \to \ell_\infty} < 1$.

We apply the second condition to our system (6.21). We seek d_1, d_2 satisfying:

$$\| \begin{bmatrix} \frac{1}{\gamma}Q\boldsymbol{\Phi} & \frac{d_2}{d_1}\frac{1}{\gamma}Q\boldsymbol{\Phi} \\ \frac{d_1}{d_2}\varepsilon\boldsymbol{\Phi} & \varepsilon\boldsymbol{\Phi} \end{bmatrix} \|_{\ell_\infty \to \ell_\infty} < 1 \tag{6.22}$$

The ratio $\frac{d_2}{d_1}$ appears in two places, so we introduce $\alpha = \frac{d_2}{d_1}$.

$$\| \begin{bmatrix} \frac{1}{\gamma}Q\boldsymbol{\Phi} & \alpha\frac{1}{\gamma}Q\boldsymbol{\Phi} \\ \frac{1}{\alpha}\varepsilon\boldsymbol{\Phi} & \varepsilon\boldsymbol{\Phi} \end{bmatrix} \|_{\ell_\infty \to \ell_\infty} < 1 \tag{6.23}$$

Because the norm is row-wise, we can separate this into two conditions:

$$\begin{aligned} \| \begin{bmatrix} Q\boldsymbol{\Phi} & \alpha Q\boldsymbol{\Phi} \end{bmatrix} \|_{\ell_\infty \to \ell_\infty} &< \gamma \\ \| \begin{bmatrix} \varepsilon\boldsymbol{\Phi} & \alpha\varepsilon\boldsymbol{\Phi} \end{bmatrix} \|_{\ell_\infty \to \ell_\infty} &< \alpha \end{aligned} \tag{6.24}$$

And because α is positive, we can factor:

$$\begin{aligned} \|(1+\alpha)Q\boldsymbol{\Phi}\|_{\ell_\infty \to \ell_\infty} &< \gamma \\ \|(1+\alpha)\varepsilon\boldsymbol{\Phi}\|_{\ell_\infty \to \ell_\infty} &< \alpha \end{aligned} \tag{6.25}$$

Because we are optimizing over γ, the first constraint can be ignored; we simply optimize the left-hand side and then choose γ. To present the results more clearly, we can further substitute $t = 1 + \alpha$, then substitute $\bar{\mathbf{\Phi}} = \mathbf{\Phi} t$.

Theorem 6.2. The following optimization problem is equivalent to (6.20)):

$$\begin{aligned} \underset{\bar{\mathbf{\Phi}}, t}{\text{minimize}} \quad & \|Q\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} \\ \text{subject to:} \quad & \|\varepsilon\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} < t - 1 \\ & Z_{AB}\bar{\mathbf{\Phi}} = tI \\ & \bar{\mathbf{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty \end{aligned} \tag{6.26}$$

Proof. The objective $\|Q\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty}$ is equal to the optimal γ^*, and we can recover the optimal map $\mathbf{\Phi} = \bar{\mathbf{\Phi}}/t$.

This problem is faster to solve, and can be verified computationally to give the same results. Instead of a bisection on γ requiring several consecutive linear programs with m variables, we solve one linear program with m variables (adding t but dropping γ). The problem also has some interesting properties: (a) it is convex; (b) it is structure-preserving for $\mathbf{\Phi}$; (c) it is almost identical to the nominal problem.

While the later steps of this derivation took advantage of some conveniences in the unstructured problem, we do not depend on the assumption that $\mathbf{\Delta}$ is unstructured. We will now generalize to structure.

State feedback with structured uncertainty

Assume that some n-block structure exists in $\mathbf{\Delta}$, with corresponding block partition in $\mathbf{\Phi}$. $\mathbf{\Delta_P}$ will remain unstructured and indexed as the first block, creating $n + 1$ blocks in the augmented robust stability problem.

We then consider a robust stability test on the following partitioned system. For clarity of exposition and without loss of generality, we take $n = 2$. In addition, we observe that ε here acts as a scalar weighting function on the input of Δ into Φ, and these weights can be considered more generally with diagonal Σ right-multiplying Φ.

$$\begin{bmatrix} \frac{1}{\gamma}Q\mathbf{\Phi} & [\frac{1}{\gamma}Q\mathbf{\Phi}]_{:1} & [\frac{1}{\gamma}Q\mathbf{\Phi}]_{:2} \\ [\mathbf{\Phi}\Sigma]_{1:} & [\mathbf{\Phi}\Sigma]_{11} & [\mathbf{\Phi}\Sigma]_{12} \\ [\mathbf{\Phi}\Sigma]_{2:} & [\mathbf{\Phi}\Sigma]_{21} & [\mathbf{\Phi}\Sigma]_{22} \end{bmatrix} \tag{6.27}$$

As in the optimization problem (6.36), we treat the first row (the fictitious perturbation row for robust performance) separately by fixing $d_1 = 1$ without loss of generality.

We define $\bar{\mathbf{\Phi}}D$, overloading D as follows. When right-multiplying $\mathbf{\Phi}$, $D = blkdiag(I, d_2 I, d_3 I)$ with each identity matrix sized appropriately. We will use the same notation with left-multiplication (e.g. by D^{-1}), but if $\mathbf{\Phi}$ is not square, the sizing of the respective identity matrices will be different. This overloading of D makes it easier to work with non-square $\mathbf{\Phi}$.

Theorem 6.3. We can solve the following convex robust performance problem for 2-block up to n-block structured uncertainty:

$$\begin{gathered}\underset{\bar{\mathbf{\Phi}},D}{\text{minimize}} \quad \|Q\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} \\ \text{subject to:} \\ \sum_j |\mathbf{\Phi}_{ij}| < d_i - 1 \quad \forall i \in 2, \ldots, n+1 \\ Z_{AB}\bar{\mathbf{\Phi}} = D \\ \bar{\mathbf{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty\end{gathered} \tag{6.28}$$

Proof. As before, the optimal cost of the above problem is exactly the robust performance level γ^* and the optimal $\mathbf{\Phi}$ can be recovered as $\bar{\mathbf{\Phi}}D^{-1}$.

Output feedback with unstructured uncertainty

We next consider the robust output feedback case with equations:

$$\begin{aligned} zx &= (\boldsymbol{A_0} + \mathbf{\Delta_A})x + (\boldsymbol{B_0} + \mathbf{\Delta_B})u \\ y &= (\boldsymbol{C_0} + \mathbf{\Delta_C})x + (\boldsymbol{D_0} + \mathbf{\Delta_D})u \end{aligned} \tag{6.29}$$

We define our nominal closed-loop maps as before, now with disturbances w and measurement noise v.

$$\begin{bmatrix} x \\ u \end{bmatrix} = \begin{bmatrix} \Phi_{xw} & \Phi_{xv} \\ \Phi_{uw} & \Phi_{uv} \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix} \tag{6.30}$$

Then in the robust case:

$$\begin{aligned} x &= \Phi_{xw}(w + \Delta_A x + \Delta_B u) + \Phi_{xv}(v + \Delta_C x + \Delta_D u) \\ u &= \Phi_{uw}(w + \Delta_A x + \Delta_B u) + \Phi_{uv}(v + \Delta_C x + \Delta_D u) \end{aligned} \tag{6.31}$$

Rearranging terms, we can proceed as we did in state feedback, optimizing γ such that:

$$\begin{bmatrix} \frac{1}{\gamma}Q\mathbf{\Phi} & \frac{1}{\gamma}Q\mathbf{\Phi} \\ \mathbf{\Phi}\Sigma & \mathbf{\Phi}\Sigma \end{bmatrix} \tag{6.32}$$

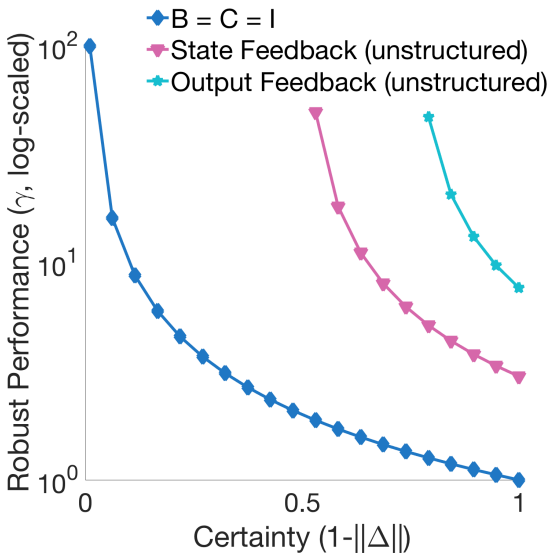

Figure 6.4: Robust performance for a simple two-state plant which is intrinsically fragile in the output feedback case with unstructured uncertainty. As $\|\mathbf{\Delta}\|$ increases towards the left, the output feedback problem quickly becomes infeasible.

is robustly stable and satisfies the output feedback constraints:

$$\begin{aligned} Z_{AB}\mathbf{\Phi} &= \begin{bmatrix} I & 0 \end{bmatrix} \\ \mathbf{\Phi} Z_{AC} &= \begin{bmatrix} I \\ 0 \end{bmatrix} \end{aligned} \tag{6.33}$$

where $Z_{AB} = [zI - A, \quad -B]$ as before and $Z_{AC} = [zI - A; -C]$. The robust performance criterion here is similar in structure to the criterion in the state feedback problem. However, the substitution used in state feedback, $\bar{\mathbf{\Phi}} = \mathbf{\Phi} D$, is helpful for the input feasibility constraint. but unhelpful for the output feasibility constraint; while $Z_{AB}\mathbf{\Phi} D = D$, it is not clear what $\mathbf{\Phi} D Z_{AC}$ is.

In the case with unstructured Δ, this issue is resolved easily; we fix the Δ block corresponding to the performance row to one, then set the D corresponding to the robust stability block to tI. The scalar t commutes to give us $\mathbf{\Phi} t Z_{AC} = \mathbf{\Phi} Z_{AC} t$ and we can proceed with the same substitutions that we used in state feedback.

Prototype Problems

The interplay between robust control and the local, sparse, and delayed constraints in the SLS framework leads to open problems and new application areas. Motivated by problems in cyberphysical systems, networked systems, and biological systems, we consider some prototype problems for which the techniques we have developed

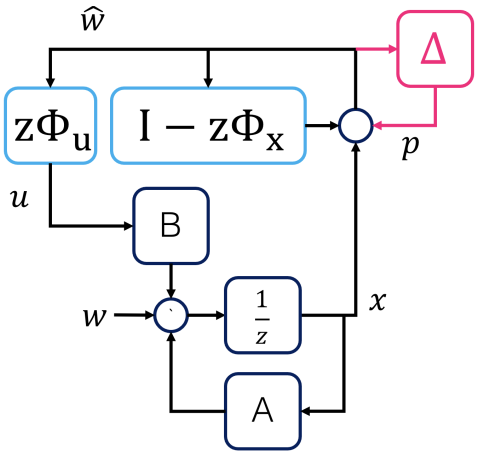

Figure 6.5: Block diagram of parametric uncertainty in a closed loop state feedback problem. The plant is shown in dark blue, and the controller in light blue, which represent the nominal problem. The uncertain loops are shown in pink.

in the preceding section can be used. In each case, we sketch out a type of problem for which a a class of instantiations can directly be solved with the techniques we have developed.

Uncertainty tolerances and speed-accuracy tradeoffs in communications. In the nominal (that is, not robust) SLS problem, it is possible to constrain the information structure of a large-scale controller in terms of sparsity, locality, and delay. For example, Sub-Controller A might not be able to send information about its state and disturbances to Sub-Controller B, or might only be able to send information after some time interval. Yet this existing theory for nominal systems assumes that communications, when sent, are perfect, when in practice internal communications in a large-scale system are likely to be imperfect. We are therefore motivated to solve a design problem where uncertainty enters the controller internally in communications (see Fig. (6.6). Such a problem depends on the controller structure, so we assume a controller structure. Let us therefore take the typical SLS state-feedback controller with predictive internal feedback pathways [1]. That is, instead of $\hat{w}$ we have $(I+\mathbf{\Delta})\hat{w}$, modeling a centralized controller with internal uncertainties in communication in the loop from sensors to actuators.

Using similar tools as we used in the state feedback problem with parametric uncertainty, we now consider a $\mathbf{\Delta}$ that enters at $\hat{w}$. Conveniently, this disturbance enters the system exactly where v enters the system in output feedback, now on a unit delay and multiplicative amplification. As in the state feedback problem, we define Σ to

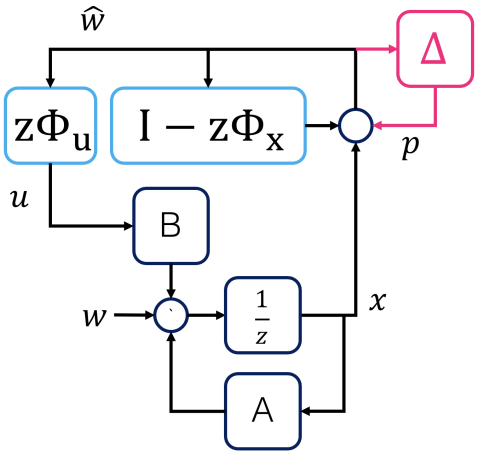

Figure 6.6: Block diagram of uncertainty entering communication lines in a state feedback SLS problem where communication between actuators and sensors is restricted (e.g. in [1]).

be the weighting matrix through which uncertainties enter the system.

We can therefore derive the following augmented system, which parallels the state feedback problem:

$$\begin{bmatrix} \frac{1}{\gamma}Q\mathbf{\Phi} - \frac{1}{z}\begin{bmatrix} I \\ 0 \end{bmatrix} & \frac{1}{\gamma}\frac{1}{z}Q\mathbf{\Phi}((zI - A)\Sigma \\ I & \Sigma \end{bmatrix} \tag{6.34}$$

When $\mathbf{\Phi}$ is a finite-impulse response convolution, it is straightforward to specify Σ as a finite impulse response filter instead of a static weighting matrix. This then allows us to study a realistic scenario where the immediate responses to a disturbance are uncertain (for instance based on low-resolution but widely available sensors) but later in the course of the FIR response to the disturbance, the controller has access to more certain information about the past disturbance (for instance based on high-resolution but selectively available sensors).

Trading off performance and certainty. Consider again the unstructured state feedback problem:

$$\begin{aligned} \underset{\bar{\mathbf{\Phi}},t}{\text{minimize}} \quad & \|Q\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} \\ \text{subject to:} \quad & \|\varepsilon\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} < t - 1 \\ & Z_{AB}\bar{\mathbf{\Phi}} = tI \\ & \bar{\mathbf{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty \end{aligned} \tag{6.35}$$

We can rearrange this problem to indirectly optimize the level of uncertainty tolerated by the closed loop:

$$\begin{aligned} \underset{\bar{\mathbf{\Phi}},t,\gamma,\theta}{\text{minimize}} \quad & \gamma + \theta \\ \text{subject to:} \quad & \|Q\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} < \gamma \\ & |\bar{\mathbf{\Phi}}\|_{\ell_\infty \to \ell_\infty} < \theta \\ & Z_{AB}\bar{\mathbf{\Phi}} = tI \\ & \bar{\mathbf{\Phi}} \in \frac{1}{z}\mathcal{RH}_\infty \end{aligned} \tag{6.36}$$

Here, one way to make θ small is to allow ε from the original problem to become large, which implicitly penalizes γ. This is easier to see using the augmented system conditions from (6.21):

$$\sum \left| \begin{bmatrix} Q\mathbf{\Phi} & \alpha Q\mathbf{\Phi} \\ \frac{1}{\alpha}\mathbf{\Phi} & \mathbf{\Phi} \end{bmatrix} \right| < \begin{bmatrix} \gamma \\ \theta \end{bmatrix} \tag{6.37}$$

We can improve γ by decreasing α, but this increases $\frac{1}{\alpha}$, which increases θ, which decreases ε. While tradeoffs between performance and robustness are an old subject, this formulation is relatively transparent, and may be useful in certain problem settings: Rather than designing for a particular ε, we might consider it a favorable property of the system to allow ε to be as large as possible. Alternatively, it might be favorable in an online setting like model-predictive control, or an adaptive control setting where richer information can be acquired at cost, to spend more on information in some contexts or to tolerate more risk in other contexts. We can place p in the weighted cost function and minimize it to further explore these kinds of tradeoffs.

Numerical experiments

In numerical experiments, we set control costs (costs on Φ_{uw} and Φ_{uv}) to be zero, $D_{22} = 0$ (no feedforward from u to y), and time horizon $T = 10$.

Chain system and ring system

We consider the chain system where A is a scaled tridiagonal matrix, and the related ring system, e.g. for $n = 4$:

$$A_{chain} = a \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \quad A_{ring} = a \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix}$$

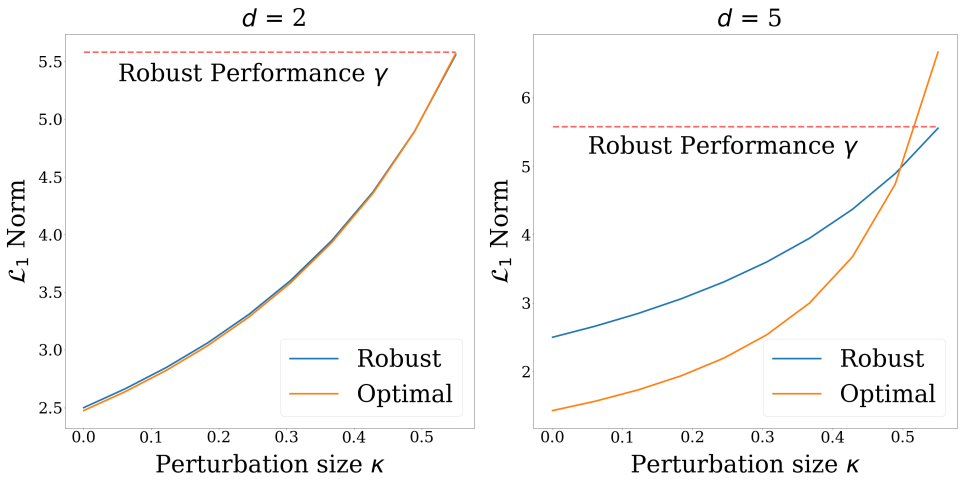

Figure 6.7: Performance of a robust controller and a nominally optimal controller for for two decentralized chains, with disturbances $\mathbf{\Delta}_A = \text{blkdiag}(\kappa I, \kappa I, \dots)$. As κ increases to $\epsilon = 0.55$, performance of the robust controller meets the robust performance bound γ (6.20).

Example plant fragile in output feedback

A simple system to consider for unstructured output feedback is the following:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \tag{6.38}$$

We allow only the first state to be actuated in both the state feedback and output feedback case, and only the second state sensed in the output feedback case (Fig. 6.4).

Localization on a chain

We consider a scaled doubly-stochastic chain system described by the following dynamics:

$$\begin{aligned} x^1_{t+1} &= \rho\left[(1-\alpha)x^1_t + \alpha x^2_t\right] + u^1_t \\ x^i_{t+1} &= \rho\left[\alpha x^{i-1}_t + (1-2\alpha)x^i_t + \alpha x^{i+1}_t\right] + u^i_t, \\ &\qquad \text{for } i = 2, \dots, N-1, \\ x^N_{t+1} &= \rho\left[\alpha x^{N-1}_t + (1-\alpha)x^N_t\right] + u^N_t \end{aligned} \tag{6.39}$$

where the $x^i_t, u^i_t \in \mathbb{R}$ are the scalar state and inputs, respectively, of the subsystems, and we set the number of scalar subsystems $N = 50$, the scaling factor $\rho = 0.5$, and the coupling constant $\alpha = 0.49$.

We solve the robust performance performance problem (6.20) under both centralized and localized distributed constraints with a norm bound on the uncertainty of $\epsilon =$

0.55, and cost matrices $C^\top = [I_N, 0^\top]^\top$ & $D^\top = [0^\top, 5I_N]$. For both settings, we impose an FIR horizon of $T = 10$ when solving the robust performance problem (6.20). Additionally, we enforce that the corresponding system responses satisfy d-locality constraints – intuitively, these constraints ensure that in closed loop, the disturbance striking node i only affects nodes j satisfying $|j - i| \leq d$.[3]

We determine that the optimal robust performance level is $\gamma = 5.57$ for both centralized and distributed controllers, where for the distributed localized controller we set the locality diameter to $d = 2$. That there is no gap between centralized and distributed is not surprising because: (i) we impose no communication delay constraints, and (ii) $\mathcal{L}_1$ optimal control leads to deadbeat optimal closed loop responses, which will consequently also be (approximately) localized in space as well. We note that the nominal $\mathcal{L}_1$ norms of the closed loop systems for the centralized and distributed localized controllers are both 2.5. Comparing these to the norms achieved by the optimal $\mathcal{L}_1$ controllers (i.e., those computed by minimizing the performance cost with $\epsilon = 0$) of 1.43 and 2.47, respectively, we see that while there is an appreciable degradation in nominal performance in the centralized setting, there is nearly no degradation in the localized distributed setting. We conjecture that this is due to the sparsity of the augmented plant $\boldsymbol{M}$ defined by the system response $\hat{\boldsymbol{\Phi}}$, which constrains both robust and nominal systems to behave similarly.

To empirically test this conjecture, we examine the evolution of the closed loop norm of the nominal and robust controllers to perturbations of the form $\boldsymbol{\Delta_A} = \text{blkdiag}(\kappa I, \kappa I, \dots)$ and $\boldsymbol{\Delta_B} = 0$, for $\kappa \in [0, \epsilon]$, for varying locality parameters $d \in \{2, 5, 10\}$, where $d = 10$ corresponds to the centralized setting. The results are displayed in Fig 6.7. We show only the results for $d = 2$ and $d = 5$, as the result for $d = 5$ and $d = 10$ are indistinguishable – as can be observed, in the tightly localized setting of $d = 2$, the degrees of freedom are limited such that robust and nominal control behave similarly; in contrast, when $d = 5$, the robust controller enjoys improved performance for larger values of κ, at the expense of degraded performance at lower values. Our approach therefore allows for a principled exploration of tradeoffs between synthesis/implementation complexity (as measured by d), nominal performance, and robust performance for large-scale distributed systems. For state feedback, we consider $C = I$ and B as a diagonal matrix with ones on every other entry starting with the first (i.e., $b_1 = 1, b_3 = 1$). For

[3] In the interest of clarity, we do not enforce communication delay constraints, but note that both communication delay and locality constraints can be enforced through suitable sparsity constraints on the system response variables: see [105] for details.

output feedback, we use the same B as in state feedback and define C as a diagonal matrix with ones on every other entry starting with the first (i.e., $c_2 = 1$, $c_4 = 1$).

6.3 Discussion

We generalized the SLS parameterization of LTI ℓ_∞-stabilizing controllers, as well as a robust counterpart, to systems described by bounded and causal linear operators. This extension, along with a simple algebraic transformation, allowed us to leverage tools from $\mathcal{L}_1$ robust control to derive necessary and sufficient conditions for the robust performance of an uncertain system in terms of convex constraints on the system response. These conditions remain necessary and sufficient when additional structure, such as that induced by delay, sparsity, and locality constraints, are imposed on the system response. Further, for $\mathcal{L}_1$ optimal control, the robust performance criteria satisfy the partial-separability properties needed to apply the distributed synthesis techniques described in [108], making our results applicable to large-scale systems.

The convexity of the solution space of robustly performing closed-loop maps is an interesting object in its own right. It has been previously observed that robustness is an essential feature not only of engineered systems but of natural systems [120].

In this work, we have in several places exploited special features of $\mathcal{L}_1$ robustness and the $\ell_\infty \to \ell_\infty$ norm. The substitutions and rearrangements we use are a consequence of positive, constant D-scaling matrices. The use of row-sums also allows partial separability of the synthesis problem. The connections between row and column separation and their interactions in estimation and control separation, large-scale synthesis, and iterative algorithms (including ADMM [113–115], and D-K iteration and similar methods [111, 121]) are an interesting topic of study that could facilitate more powerful synthesis algorithms or characterize the limitations of possible algorithms in these problem settings. The extent to which our convex sets are convenient artifacts of particular problem settings, as opposed to fundamental characteristics of robust control systems, remains unclear. It might be desirable to have a row- and column- separable induced norm, a topic discussed in [113, 116]. It is our hope that the theory and computational approaches developed here will generalize to many norms that may be appealing for large-scale networked systems.

In the final technical chapter of this thesis, we move from our most abstract theory to our most concrete case study. Chapter 7 is a theory-informed experimental study of resistance to neoplasia and cancer in the moon jellyfish *Aurelia coerulea*. This study draws on many of the themes and technical results of the previous chapters, studying decision-making by cells and cell populations in cancer detection in the moon jellyfish. While there are broader conceptual connections here, we will present this material in a way that emphasizes the experimental results.

Chapter 7

CELL FLUX PROOFREADS NEOPLASIA IN MOON JELLYFISH

7.1 Introduction

Tissue maintenance in multicellular organisms often requires precisely regulated cell proliferation and cell death. Dysregulation of processes is associated with hyperplasia, atrophy, and cancer [122]. Organisms in which cancer is rarely observed in nature or in the laboratory, such as cnidarians [123–125], suggest the intriguing possibility of novel anti-cancer and tissue-homeostatic mechanisms. In this chapter, we characterize a novel mechanism of resistance to neoplasia in the cnidarian moon jellyfish *Aurelia coerulea*.

The moon jellyfish has a distinct eight-lobed anatomy in its ephyra life stage (the life stage which occurs after metamorphosis from a polyp but before the development of the familiar medusa, Figure 7.1A). Recent experiments have shown that ephyrae do not typically regenerate these lobes after amputation, and instead rearrange the lobes to quickly re-symmetrize without cell proliferation [2, 126]. However, when fed with the amino acid L-leucine, ephyrae can be induced to proliferate regenerate lobes [2]. This inducible regenerative and proliferative state made the moon jellyfish an attractive model in which to study the relationship between proliferation, regeneration, and cancer. The moon jellyfish also has other features relevant to our experimental study: it is a basal metazoan, meaning that it may provide clues about the evolutionary origins of cancer; it has no known dedicated stem cells, instead retaining plasticity and self-renewal in most or all cells [127]; and it has a limited cellular immune response, so we can experimentally modulate proliferation dynamics without modulating immune cell population size [127, 128].

We use mathematical modeling and experiments to study resistance to neoplasia in the moon jellyfish. We first provide evidence that low rates of cancer in the moon jellyfish result from active resistance mechanisms, and establish an experimental model of inducible neoplasia in moon jellyfish ephyrae. We then propose a new model, the *flux* model, which makes the surprising prediction that proliferation in the moon jellyfish is protective against neoplasia, and conversely, suppressing proliferation in the ephyra should increase neoplastic susceptibility. We verify this prediction experimentally, first by inducing neoplasia, and then by rescuing the

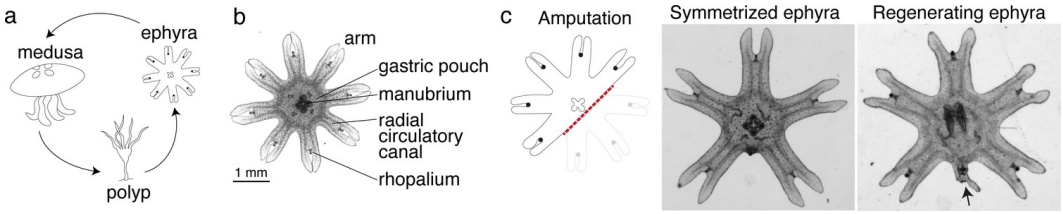

Figure 7.1: The moon jellyfish *Aurelia coerulea* has a three-stage life cycle. In the ephyra life stage, which we study here, the jellyfish has an eight-lobe morphology. When these lobes are amputated, the ephyra does not typically regenerate lobes. However, when treated with L-leucine, or L-leucine and insulin, the ephyra can regenerate lobes. Figure reprinted, under Creative Commons License, from [2].

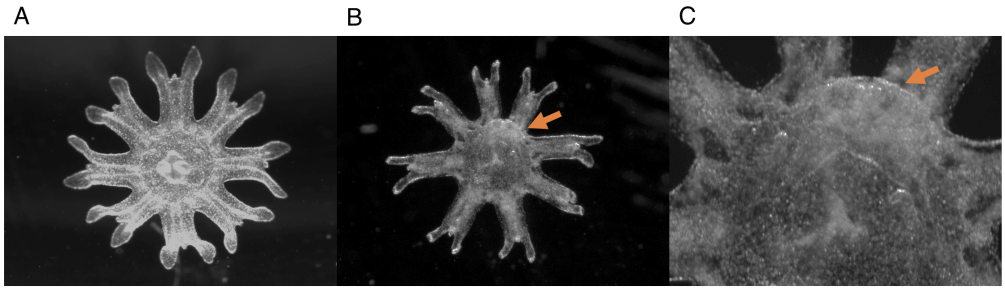

Figure 7.2: Representative images of jellyfish responses to carcinogen. (A) A typical healthy morphology, in this case despite treatment with the carcinogen 3-methylcholanthrene. (B) A representative neoplasia on the muscle band of an ephyra induced by co-treatment with 3-methylcholanthrene and an mTOR inhibitor (see text). (C) A zoomed in image ofthe same jellyfish as in (B).

neoplasia-resistant phenotype.

7.2 Results

Lifestyle vs resistance

We start by modeling the problem mathematically using a simple cell-level differential equation model. We use some concepts and simple techniques from control theory, but use them in ways that are familiar in systems biology. Because the model is simple, and because we are interested in modeling changes in qualitative features of organismal physiology and morphology, we will not need precise estimates of the model parameters; typically we will be interested in whether parameters are high or low relative to a physiological baseline.

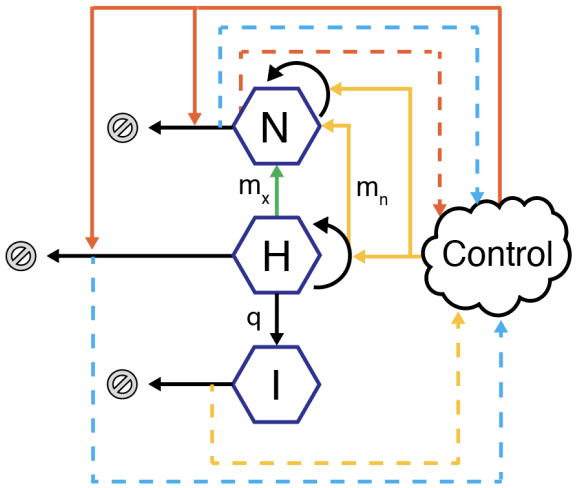

Figure 7.3: We consider a control-theoretic model derived from a base model. The base model considers three types of cells: injured cells (I), healthy (H), and neoplastic (N), each with a division and death rate. In the control-theoretic *flux* model, injury and cancer clearance are both coupled to cell division, so cancer clearance is facilitated by high cell division; more global cell proliferation predicts less cancer.

We model the problem as follows:

$$\begin{aligned} \dot{x}_h &= (1 - m_n)b_h - m_x x_h - q x_h - d_h \\ \dot{x}_m &= a_m x_m + m_n b_h + m_x x_h - d_m \\ \dot{x}_i &= q x_h - d_i \end{aligned} \tag{7.1}$$

x_h is the number of healthy cells per volume, x_m is the number of neoplastic cells per volume and x_i is the number of injured or damaged cells per volume. a_m is the rate of cell division of neoplastic cells, m_n is the rate at which healthy cells endogenously mutate during division, m_x is the rate at which healthy cells exogenously mutate, and q is the rate of injury or endogenous wear and tear on healthy cells. Lastly, four special parameters will be the focus of our analysis. b_h represents births of healthy cells and d_h, d_m, and d_i represent death of their respective cell types. These are the parameters that the organism, through internal processes, can modulate – the homeostatic parameters that govern cancer control.

We can make some immediate simplifications to the model. Because d_i is not coupled to any other parameters in the model, we can assume that organismal control clears x_i. We can therefore immediately simplify d_i and x_i from the model. Next, we define $x_h = y_h + h$ for some equilibrium healthy cell concentration h. y_h

will have the same derivative as x_h, but can be interpreted as the deviation from equilibrium; with this substitution, the equilibrium is zero in both coordinates.

$$\begin{aligned} \dot{y}_h &= (1-m_n)b_h - m_x y_h - m_x h - q y_h - q h - d_h \\ \dot{x}_m &= a_m x_m + m_n b_h + m_x(y_h + h) - d_m \end{aligned} \tag{7.2}$$

To temporarily cancel the h-dependent terms, we introduce u_h and u_m and set $b_h = u_h + (m_x+q)h/(1-m_n)$, $d_m = u_m + m_x h + m_n(m_x+q)h/(1-m_n)$, and $d_h = 0$. We are left with a linear control system in a standard form:

$$\begin{bmatrix} \dot{y}_h \\ \dot{x}_m \end{bmatrix} = \begin{bmatrix} -m_x - q & 0 \\ m_x & a_m \end{bmatrix} \begin{bmatrix} y_h \\ x_m \end{bmatrix} + \begin{bmatrix} 1-m_n & 0 \\ m_n & -1 \end{bmatrix} \begin{bmatrix} u_h \\ u_m \end{bmatrix} \tag{7.3}$$

with the further restriction that:

$$\begin{aligned} u &= K \begin{bmatrix} y_h \\ x_m \end{bmatrix} \\ K &= \begin{bmatrix} a_h & 0 \\ 0 & k_m \end{bmatrix} \end{aligned} \tag{7.4}$$

The control system is stable when $a_m < k_m$ and $a_h < (m_x + q)/(1 - m_n)$. Re-expressing these conditions in terms of the birth and death rates, we can write:

$$\begin{aligned} b_h &< \frac{m_x + q}{1 - m_n} x_h \\ d_m &> a_m x_m + \frac{m_x + m_n q}{1 - m_n} h \end{aligned} \tag{7.5}$$

To achieve a reliable no-cancer state, x_m must stay near zero throughout the time interval considered. This can happen under one of two conditions:

- Lifestyle: Cancerous mutations are so rare that cancer is never observed. m_x is small **and** $m_n b_h$ is small.
- Resistance: Cancerous mutations are detected and removed quickly enough that they are never observed, as described in (7.5).

If low rates of cancer in the moon jellyfish are primarily due to lifestyle, then increasing the rate of exogenous mutations might lead to neoplasia. A lack of neoplasia under carcinogen treatment would weigh against the lifestyle model, but not rule it out, because we would be left to wonder whether the dose and duration

of carcinogen were sufficient to cause malignant transformations. On the other hand, the resistance model makes a strong testable prediction. If d_m is sufficiently small and m_x is sufficiently large, then the resistance condition cannot be achieved. To test these predictions, we therefore designed experiments to modulate m_x and d_m. Exogenous mutations m_x could be induced by the chemical carcinogen 3-methylcholanthrene (MCA) [129]. d_m could be inhibited by the broad caspase inhibitor zVAD-FMK.

In experiments, co-treatment with MCA and zVAD-FMK induced neoplasia (Figure 7.4A). Neither MCA nor zVAD-FMK alone induced neoplasia. To our knowledge, these are the first neoplasia reported in the moon jellyfish or any scyphozoan, in natural observation or by laboratory induction. Neoplasia were consistently observed within the first day of treatment.

To further confirm that the neoplasia phenomenon did not depend on unmodeled effects of MCA or zVAD-FMK, we also tested co-treatment with MCA and a Bax-inhibiting peptide (MilliporeSigma 80603-646). Bax can initiate apoptosis downstream of signals from the tumor-suppressor gene p53 [130], making it a suitable and more specific alternative target for inhibiting apoptosis. We also tested an alternate carcinogen, ultraviolet light (Figure 7.4B). Neoplasia were again observed under these conditions after three days of treatment. MCA alone did not induce neoplasia after one day, three days, or seven days.

Proliferation vs flux

We next sought to perturb the homeostatic mechanisms in a different way, testing the relationship between injury and neoplasia. We injured ephyrae with the amputation protocol used in [2]. When injured ephyrae were co-treated with MCA, they were susceptible to neoplasia (Figure 7.6). Given that injury and cancer have traditionally been linked, one possible interpretation of this data might be that injury drives proliferation and growth, which create suitable conditions for neoplasia. However, some observations were not entirely consistent with this interpretation: neoplasia in the injured animals did not consistently appear at the injury site, suggesting that the neoplasia was something other than a failed attempt to regrow the lost lobes; (2) the MCA-exposed ephyrae were not smaller in size than the control ephyrae, suggesting that apoptosis alone did not account for the neoplasia resistance; and in previous experiments in which ephyrae were injured, cell proliferation did not increase, suggesting that proliferation did not account for the neoplasia phenotype

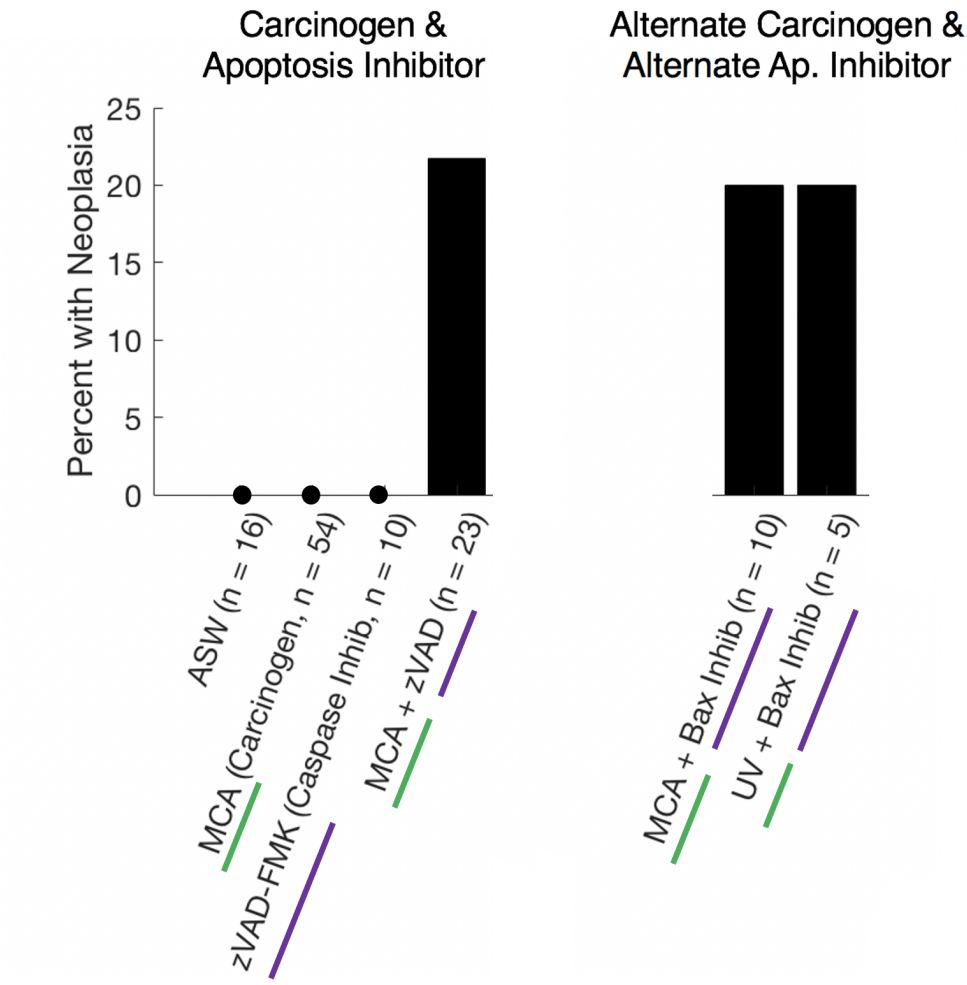

Figure 7.4: In artificial salt water and in exposure to 3-methylcholanthrene (MCA, 300-500 μM), neoplasia in ephyrae is not observed. By using the broad caspase inhibitor zVAD-FMK (100 μM), we inhibit one mechansism by which pre-neoplastic cells might be removed. Co-treatment with MCA and zVAD-FMK does cause neoplasia, while zVAD-FMK alone does not. These results are not specific to this carcinogen or this inhibitor; ultraviolet light (UV) and zVAD-FMK also causes neoplasia, as do MCA and a Bax-inhibiting peptide.

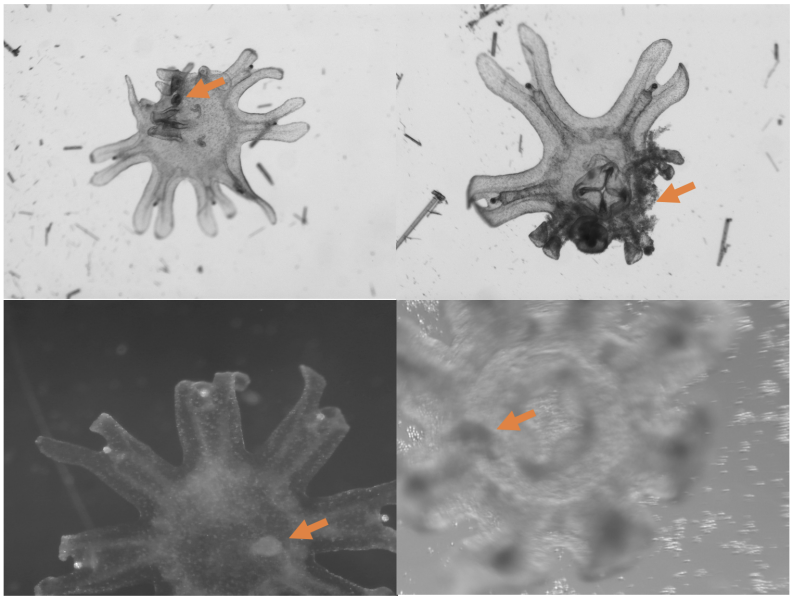

Figure 7.5: Additional examples of neoplasia in ephyrae, induced through a variety of conditions. Clockwise from top left: MCA and zVAD-FMK, MCA and zVAD-FMK, MCA and hydroxyurea, MCA and sapanisertib.

[2, 126].

We considered an alternate model of resistance to cancer which we call the *flux* model. Before we describe the mathematical argument, we develop the intuition. A fundamental challenge in active resistance to cancer is detecting neoplastic or pre-neoplastic cells, which may be difficult to distinguish from typical healthy cells. As a result, any decision process leading to clearance of a neoplastic cell will have a false positive rate and a context-dependent sensitivity to clearance. When cells are rapidly produced, they may also be easily cleared. This would then predict higher sensitivity to apoptosis in high proliferation conditions, perhaps through mechanical coupling as in [122, 131], although the model does not require that particular mechanism.

We can formalize this intuition by substituting $d_h = pd_m$ for p as the normalized false positive rate of clearance. In the linear control system, $a_h < (m_x+q+pk_m)/(1-m_n)$ and $a_m < k_m$. h is a concentration variable with respect to volume, so we can simplify further by setting $h = 1$ and implicitly rescaling the remaining parameters. With the approximation that near equilibrium x_m and y_h are small, we can write

each bound in terms of intrinsic parameters of the system.

$$\begin{aligned} d_h &> p\frac{m_x + m_n q}{1 - m_n} \\ b_h &> \frac{m_x + q + p a_m}{1 - m_n} \end{aligned} \tag{7.6}$$

Thus by coupling d_m and d_h, we have made visible not only an upper bound on b_h, which we had before, but a lower bound as well. Both b_h and d_h (that is, the *flux* in healthy cells), increase with m_x, q, p, and m_n.

The surprising model prediction is that proliferative conditions are cancer-protective We designed experiments to test this model prediction.

Testing the flux model

The most straightforward prediction of the flux model was that the tolerance of the moon jellyfish ephyrae to exogenous carcinogens depended on an active increase in cell flux. We tested this prediction using EdU staining, which allowed us to image newly synthesized DNA (and therefore newly divided cells) in a two-day period of incubation with MCA, compared to a two-day period in artificial salt water. As the model predicted, cell proliferation accelerated under MCA exposure (Figure 7.6A-B).

We reasoned that sensitivity to apoptosis could be modulated indirectly by modulating the rate of proliferation. Less proliferative jellyfish would have lower sensitivity to apoptosis, and therefore, higher susceptibility to neoplasia. We tested this with the antiproliferative chemotherapeutic hydroxyurea and with the mTOR inhibitor sapanisertib. Neither treatment caused neoplasia alone, but in each case, co-treatment with MCA did lead to neoplasia at a low but repeatable rate (Figure 7.6C).

Based on these findings, we revisited the earlier question of susceptibility to neoplasia during injury and recovery from injury. In the simplified model, injury increases flux – when body size and body plan are homeostatically maintained and neoplasia is resisted. In the ephyra, injury can lead to one of two outcomes: symmetrization, a body-plan reorganization that can proceed even without cell proliferation, or regeneration, which requires cell proliferation [2, 126]. Therefore, we reasoned that a transient *decrease* in cell flux after injury would explain the evidence.

Regeneration in moon jellyfish ephyra can be induced or suppressed through manipulation of organismal states. In particular, supplementation of L-leucine and insulin,

or just L-leucine, drives a regenerative state after injury. We used this inducibility to test three conditions: MCA with leucine, MCA with injury, and MCA with injury and leucine (recreating known regenerative conditions). As noted above, MCA with injury did make ephyrae susceptible to neoplasia. In the flux model, we would expect leucine alone to have no effect, and leucine after injury to *rescue* ephyrae from susceptibility to neoplasia.

The experimental findings were consistent with the flux model: leucine alone did not increase susceptibility to neoplasia, injury increased susceptibility to neoplasia, and leucine after injury restored resistance to neoplasia (Figure 7.6D). In totality, several experiments using different measurements and perturbations support the flux model.

7.3 Discussion

We used mathematical modeling, control theory, and experiments to investigate a curious phenomenon: the apparent absence of neoplasia in the moon jellyfish *Aurelia coerulea*. We introduced a *flux* model of resistance to cancer and neoplasia, which predicted that a high-flux state should be protective against neoplasia, while a low-flux state should be susceptible to neoplasia. We verified these predictions experimentally.

The flux model draws inspiration from, and may help synthesize, loosely related observations from several different model systems. Coupling between proliferation and cell death has been characterized through various mechanisms in other model organisms. Contact inhibition, a mechanism by which epithelial cells are sensitized to commit apoptosis when they are crowded (for instance, when they are dividing too quickly) is hypersensitive in long-lived naked mole rats. [131]. Embryonic stem cells are more sensitive to apoptosis after p53 signaling than differentiated cells [132]. Cell proliferation in the fly midgut is regulated by coupling between apoptotic enterocytes and stem cells, ensuring that new cells are produced when apoptosis drives demand [122]. Lastly, hypersensitive responses in the plant immune system can lead to the sacrifice of an entire branch to prevent spread of an infection; compared to vertebrates, plants lack motile immune cells and vital organs, suggesting an overall strategy that necessitates more immune false positives [133].

Taken together, these observations suggest interrelationships between cell flux, localization of immune responses, regeneration, and cancer. In our modeling, we have formalized some of these relationships in a relatively simple system. In a broader

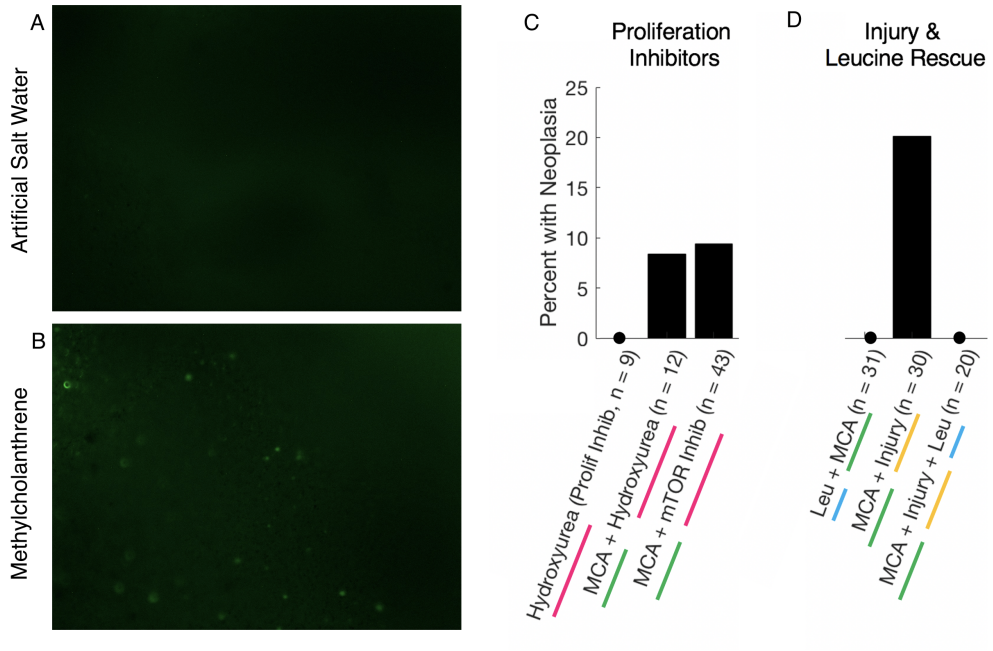

Figure 7.6: After initial mathematical modeling and experiments showed that caspase and Bax inhibition could make an ephyra susceptible to neoplasia, we pursued further modeling and experiments to test the flux models. (A,B) We used EdU staining to compare the rate of DNA synthesis, as a proxy for cell proliferation, between artificial salt water (ASW) and carcinogen (methylcholanthrene, MCA) in the muscle band of two ephyrae. As predicted by the flux model, cell cycle increased substantially with MCA. (C) A further prediction of the flux model is that decreasing global cell proliferation will increase neoplastic susceptilibity. We tested this with two anti-proliferative experiments: directly with the chemotherapeutic hydroxyurea, and indirectly with the mTOR inhibitor sapanisertib. (D) In previous experiments, leucine induced in ephyrae a regenerative and proliferative state. Leucine alone did not increase neoplastic susceptibility. However, injury did – again raising the question of whether a proliferation-associated process might create susceptibility. Instead, using leucine to recreate the post-injury regenerative conditions restore resistance to neoplasia.

evolutionary view, these elements may be considered part of a system of tissue homeostasis that allows a more unified understanding [134].

Limitations

With our experiments and mathematical model, we hope to identify general principles about cancer formation. The conclusions of this chapter leave many questions unanswered for future study. Our model considers cancer resistance, but we have not characterized in our experiments the distinction between neoplasia and cancer. It seems likely that an unpatterned neoplasia brought on by carcinogenic treatment is cancer in some or all cases, but we have not rigorously characterized, for example, in what percentage of cases the neoplasia is malignant.

The neoplasia phenotype and the experimental design allowed us to avoid modeling longer time-horizon scenarios in which pre-cancerous cells accumulate. We also avoided modeling body size and distinctions between stem cells and differentiated cells. These elements were not necessary for short timescale experiments in ephyrae, but would be necessary for a richer cross-species comparison.

Interpretation

Cancer has often been viewed as a subversion of self-repair mechanisms in wound healing and regeneration [135]. Early versions of this idea were proposed by Virchow in the 1850s [136], and a modern view began with the discovery of stem cells in the 1960s [137]. Because cancer is fundamentally a disease of cell division, it may be seen as an inevitable cost to regenerative tissue [138, 139]. In blood, for example, hematopoietic stem cells are more susceptible to cancer than differentiated cells [140]. Across tissue types in humans, lifelong cell divisions are associated with neoplastic susceptibility, with tissues like colon high in both division and susceptibility and tissues like bone low in both division and susceptibility [141]. Proposed mechanisms for a relationship between regeneration and cancer vary, and include the accumulated mutational burden from cell divisions [141], unrestrained stemness [142], and misuse of immune cells and pathways needed for wound healing [143].

Yet a relationship between regeneration and cancer, regardless of mechanism, is far from universal in the animal kingdom. Cancer rates and lifespan do not show the relationship that would be predicted by a simple mutation-accumulation argument, for example, with short-lived mammals like mice having similar rates of cancer to

long-lived mammals like humans [144]. Several mammals, such as naked mole rats, microbats, elephants and gray whales achieve long life without concomitantly high rates of cancer – either in comparison to similarly sized mammals or in comparison to a rate of cancer that would theoretically scale with body mass and lifespan [144, 145]. Intriguingly, highly regenerative non-mammalian organisms such as hydra, planaria, and axolotls often have low rates of cancer [146]. It remains unclear how these organisms balance extracellular signaling to regulate pre-neoplastic cells versus flux-like mechanisms to ultimately clear pre-neoplastic cells.

The flux model gives conditions under which cancer resistance and tissue homeostasis are jointly achieved, but it also implies failure modes. A high-flux tissue will resist exogenous mutations and wear and tear more easily, even with ongoing endogenous mutations. However, if flux declines, the high-flux system may be *more* vulnerable to cancer than the low-flux system – which may be why the neoplasia that formed in our experiments formed so quickly, typically in one to three days. Therefore, the overall slowdown of cell division in aging may be an important factor in oncogenesis.

Jellyfish are basal metazoans. If some evolutionary ancestor of humans had both a high cancer resistance and a high regenerative capacity, why did vertebrates – especially mammals and birds – lose those dual capacities? An intriguing possibility is that cell flux itself is a key variable: with the evolution of adaptive immunity, more targeted cancer defenses are possible, decreasing p in the flux model and decreasing the baseline necessary flux. This could then enable the development of tissues and evolved lifestyles that rely on long-lived cells – tissues like bones, hearts, and brains. Characterizing interrelationships between cell flux and more complex body plans will be an exciting direction in future work, both to clarify the evolutionary origins of cancer and to inform treatments that would seek to modulate flux rather than proliferation or apoptosis separately.

Chapter 8

CONCLUSION

In Chapter 1, I argued that nervous systems, immune systems, and developmental systems were connected by a broader theoretical framework based in special properties of multicellular control. Now that we have considered both theory and case studies, I will conclude with some informal observations.

- **Multicellular control builds on and enhances existing domain-specific models.** Each of the three domain-specific problems we considered has a history of mathematical modeling associated with it: predictive coding models in neuroscience [43], viral-immune kinetic models and host-parasite evolutionary models in immunology [50, 70, 71], and detailed statistical models of cancer risk [138, 144]. Once we recognize that each of the models analyzed in this preceding work either models a control system or includes a control system as a major unknown, we can use a multicellular control framework for both a technical and a conceptual enhancement of these existing models. Technically, we gain access to new methods for dealing with delays, uncertainty, and self-renewal in control systems. Conceptually, we gain the ability to transfer ideas and mechanisms between different systems, including systems we have not studied in detail. We take substantial inspiration, for example, from previous work on errors and fluctuations in molecular synthesis in single cells, even though our models primarily work on the multicellular scale [63, 147].

- **Even simple evolutionary problems are hard to solve, and solutions are rare.** In all three of the domains we considered in this work, we analyzed a simple task, in the sense that the task could be described by a dynamical system in relatively few variables. By using control theory to analyze all possible solutions to these simple tasks under more realistic constraints than had been used before, we were able to generate meaningful models of complexity, and in some cases directly testable predictions about the underlying biology.

- **A lot of complexity is coordination.** In Chapter 3, we considered a number of ways in which communication structure was shaped by task specifications

and locality constraints. When we defined locality as slow, inaccurate, or expensive communication between subsystems in a larger system, we found that the communication structure quickly became much more complex than the task itself. In order to separate a simple problem among different subsystems, it is necessary for each subsystem to tell at least some of the others what it is doing – that way, the control action of one subsystem would not be mistaken for additional disturbances in the regulated variables. This reasoning should be applicable to other systems as well; at least some of the signal complexity of cytokines, chemokines, and surface receptors in the immune system might be explained by the principle that the immune system is sensitive to stress and change, and immune responses themselves are sources of stress and change, so localized immune responses require a well-defined and complex communication structure.

- **Complex control systems often evolve on top of simple control systems.** When modeling the human immune system's response to viruses and cancer (in Chapter 4 and in additional work that is not part of this thesis), I often wondered why we need T cells and B cells to do a job that jellyfish accomplish more successfully with few if any dedicated immune cells. The mathematical model we use in Chapter 7 offers some clues. The challenge faced by the jellyfish in that model was how to discriminate between pre-cancerous cells and healthy cells, and one possible solution to that problem is to aggressively clear healthy cells along with pre-cancerous cells. But if you can target cells more carefully, you can build different kinds of bodies. Without spatially localized immune responses, it seems unlikely to me that humans would ever have evolved bodies and life histories that depend on protecting some parts of the body more than others – it seems unlikely that we would evolve, for example, brains or pregnancy. On the other hand, while our immune system can do many things the jellyfish's immune system cannot do, the jellyfish's capacity to resist cancer seems to exceed ours. What complex control often offers in practice is a more flexible set of systems that can be built that can exploit more detailed assumptions about the world.

- **The fact that simple interventions work is surprising, and suggests something profound.** One recurring theme across different fields in biology is that simple interventions can be surprisingly effective at changing whole-system behaviors, even when we do not understand many aspects of the underlying

mechanisms. Examples include implantable brain stimulators for epilepsy, immunotherapy for cancer, and nutrient supplementation for regeneration. If you told me, in my capacity as a control theorist, that you would present me with a system with almost no model, almost no measurements, and only one point of action, and you wanted me to exert a reliable effect on the system's behavior, I would probably tell you that the problem was impossible. The fact that it is not impossible in biology tells us something important; it tells us that biological systems belong to some very special class of systems. I have argued in this thesis that what makes these systems special is that they are multicellular control systems. I suspect that experiments and interventions that yield big effects tend to be those that are not fighting the extant control system – they are either working with it, or breaking it in just the right way. But what is the right way? The theorems and case studies in this thesis only scratch the surface of this question.

- **Theory informs, but does not replace, design and discovery.** We want math to help us understand biology more rigorously, but we do not always know what rigor should look like. The right statistical test? A thousand-parameter differential equation model? An opaque neural network with empirical predictive power? I would like to conclude this thesis with an appeal to mathematical *intuition*, not as the opponent of mathematical rigor but as its companion. The flux model that informed the experimental work in Chapter 7 would not have looked the way it did if not for the immune model in Chapter 4, the sensorimotor model in Chapter 3, and the theoretical foundation in Chapter 5 and Chapter 6. Yet the flux model itself was ultimately quite simple, and maybe the same experiments could have been done without any math at all. I have many stories like this, stories where the magic of math in biology is near-invisible. There is some knowledge to be won not just in reading the math but doing the math, not just analyzing the data but collecting the data, that cannot be conveyed in papers and textbooks. This is not the easy answer, but this is how control theory has always worked in engineering design: theory guides us to the boundary of what we can understand with certainty about the real world with rigorous mathematics, but equally importantly, it gives us intuition for the uncertain world just beyond that boundary. Most of multicellular biology remains in that uncertain world, and I hope this thesis helps us see it a little more clearly.

Bibliography

[1] Anish A. Sarma, Jing Shuang Li, Josefin Stenberg, Gwyneth Card, Elizabeth S. Heckscher, Narayan Kasthuri, Terrence Sejnowski, and John C. Doyle. Internal feedback in biological control: Architectures and examples. *American Control Conference*.

[2] Michael J. Abrams, Fayth Hui Tan, Yutian Li, Ty Basinger, Martin L Heithe, Anish A. Sarma, Iris T. Lee, Zevin J. Condiotte, Misha Raffiee, John O. Dabiri, David A Gold, and Lea Goentoro. A conserved strategy for inducing appendage regeneration in moon jellyfish, drosophila, and mice. *eLife*, 10, Dec 2021.

[3] Bode HW. *Network and Feedback Amplifier Design*. Van Nostrand, 1945.

[4] J.C. Doyle, K. Glover, P.P. Khargonekar, and B.A. Francis. State-space solutions to standard h/sub 2/ and h/sub infinity / control problems. *IEEE Transactions on Automatic Control*, 34(8):831–847, 1989.

[5] Dante C. Youla, Hamid A. Jabr, and Joseph J. Bongiorno. Modern wiener-hopf design of optimal controllers–part ii: The multivariable case. *IEEE Transactions on Automatic Control*, 21(3):319–338, Jun 1976.

[6] Mustafa Hani Khammash. *Stability and performance robustness of discrete-time systems with structured uncertainty*. PhD thesis, Rice University, 1990.

[7] Yuh-Shyang Wang, Nikolai Matni, and John C. Doyle. A system-level approach to controller synthesis. *IEEE Transactions on Automatic Control*, 64(10):4079–4093, Oct 2019.

[8] John Doyle, Bruce Francis, and Allen Tannenbaum. Feedback Control Theory. In *Feedback Control Theory*. Macmillan, 1992.

[9] Daniel M. Wolpert, Zhoubin Ghahramani, and Michael I. Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880–1882, Sep 1995.

[10] Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, 2004.

[11] David W. Franklin and Daniel M. Wolpert. Computational mechanisms of sensorimotor control. *Neuron*, 72(3):425–442, Nov 2011.

[12] Li Zhaoping. *Understanding Vision*. Oxford University Press, May 2014.

[13] Tim Gollisch and Markus Meister. Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron*, 65(2):150–164, Jan 2010.

[14] Alessandra Angelucci and Jean Bullier. Reaching beyond the classical receptive field of V1 neurons: Horizontal or feedback axons? *Journal of Physiology Paris*, 97(2-3):141–154, 2003.

[15] Y. El-Shamayleh, R. D. Kumbhani, N. T. Dhruv, and J. A. Movshon. Visual response properties of v1 neurons projecting to v2 in macaque. *Journal of Neuroscience*, 33(42):16594–16605, Oct 2013.

[16] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364, 2019.

[17] Simon Musall, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, 22(10):1677–1686, 2019.

[18] Daniel J. Felleman and David C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, pages 1–47, 1991.

[19] Edward M. Callaway. Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Networks*, 17(5-6):625–632, 2004.

[20] Lars Muckli and Lucy S. Petro. Network interactions: non-geniculate input to V1. *Current Opinion in Neurobiology*, 23(2):195–201, 2013.

[21] Mark M. Churchland, John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, Jun 2012.

[22] Francis R. Willett, Darrel R. Deo, Donald T. Avansino, Paymon Rezaii, Leigh R. Hochberg, Jaimie M. Henderson, and Krishna V. Shenoy. Hand knob area of premotor cortex represents the whole body in a compositional way, Apr 2020.

[23] Sergey D. Stavisky, Francis R. Willett, Guy H. Wilson, Brian A. Murphy, Paymon Rezaii, Donald T. Avansino, William D. Memberg, Jonathan P. Miller, Robert F. Kirsch, Leigh R. Hochberg, A. Bolu Ajiboye, Shaul Druckmann, Krishna V. Shenoy, and Jaimie M. Henderson. Neural ensemble dynamics in dorsal motor cortex during speech in people with paralysis, Dec 2019.

[24] James Anderson, John C. Doyle, Steven H. Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, 2019.

[25] Nikolai Matni and Anish A. Sarma. Robust performance guarantees for system level synthesis. *American Control Conference*, 2020.

[26] Yorie Nakahira, Quanying Liu, Terrence J. Sejnowski, and John C. Doyle. Diversity-enabled sweet spots in layered architectures and speed-accuracy trade-offs in sensorimotor control. 2019.

[27] Yorie Nakahira, Quanying Liu, Terrence J. Sejnowski, and John C. Doyle. Diversity-enabled sweet spots in layered architectures and speed–accuracy trade-offs in sensorimotor control. *Proceedings of the National Academy of Sciences of the United States of America*, 118(22):1–11, 2021.

[28] Anish A. Sarma, Jing Shuang Li, Josefin Stenberg, Gwyneth Card, Elizabeth S. Heckscher, Narayanan Kasthuri, Terrence Sejnowski, and John C. Doyle. Internal feedback in biological control: Architectures and examples. *American Control Conference*.

[29] Jing Shuang Li. Internal feedback in biological control: Locality and system level synthesis. To appear in IEEE American Control Conference 2022.

[30] Anish A. Sarma and John C. Doyle. Flexibility and cost-sensitivity in a quantized control loop. *American Control Conference*, 2019.

[31] Peter Sterling and Simon B. Laughlin. *Principles of neural design*. MIT Press, 2015.

[32] David H. Hubel and Thorsten N. Wiesel. Receptive fields of single neurones in the cat's striate cortex, Oct 1959.

[33] Katherine C. Ames and Mark M. Churchland. Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *eLife*, 8, Oct 2019.

[34] Matthew W. Self, Roxana N. Kooijmans, Hans Supèr, Victor A. Lamme, and Pieter R. Roelfsema. Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *Proceedings of the National Academy of Sciences of the United States of America*, 109(27):11031–11036, 2012.

[35] David Attwell and Alasdair Gibb. Neuroenergetics and the kinetic design of excitatory synapses. *Nature Reviews Neuroscience*, 6(11):841–849, 2005.

[36] Margaret S. Livingstone. Mechanisms of direction selectivity in macaque v1. *Neuron*, 20(3):509–526, Mar 1998.

[37] Victoria Chan-Palay, Sanford L. Palay, and Susan M. Billings-Gagliardi. Meynert cells in the primate visual cortex. *Journal of Neurocytology*, 3(5):631–658, Nov 1974.

[38] Eberhard E. Fetz. Functional organization of motor and sensory cortex: symmetries and parallels. In W.E. Gall G.M. Edelman and W.M. Cowan, editors, *Dynamic Aspects Of Neocortical Function*, pages 453–474. John Wiley, 1984.

[39] Saurabh Vyas, Matthew D. Golub, David Sussillo, and Krishna V. Shenoy. Computation through neural population dynamics. *Annual Review of Neuroscience*, 43(1):249–275, Jul 2020.

[40] Kohitij Kar, Jonas Kubilius, Kailyn Schmidt, Elias B. Issa, and James J. DiCarlo. Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nature Neuroscience*, 22(6):974–983, Apr 2019.

[41] Andre M. Bastos, W. Martin Usrey, Rick A. Adams, George R. Mangun, Pascal Fries, and Karl J. Friston. Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711, Nov 2012.

[42] E. Libby, T. J. Perkins, and P. S. Swain. Noisy information processing through transcriptional regulation. *Proceedings of the National Academy of Sciences*, 104(17):7151–7156, Apr 2007.

[43] Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, Jan 1999.

[44] Georg B. Keller and Thomas D. Mrsic-Flogel. Predictive processing: A canonical cortical computation, Oct 2018.

[45] Marcus Leinweber, Daniel R. Ward, Jan M. Sobczak, Alexander Attinger, and Georg B. Keller. A sensorimotor circuit in mouse cortex for visual flow predictions, Sep 2017.

[46] Vikash Gilja, Chethan Pandarinath, Christine H. Blabe, Paul Nuyujukian, John D. Simeral, Anish A. Sarma, Brittany L. Sorice, Janos A. Perge, Beata Jarosiewicz, Leigh R. Hochberg, Krishna V. Shenoy, and Jaimie M. Henderson. Clinical translation of a high-performance neural prosthesis. *Nat Med*, 21:1142–1145,.

[47] Beata Jarosiewicz, Anish A. Sarma, Daniel Bacher, Nicolas Y. Masse, John D. Simeral, Brittany L. Sorice, Erin M. Oakley, Christine H.. Blabe, Chethan Pandarinath, Vikash Gilja, Sydney S. Cash, Emad N. Eskandar, Gerhard Friehs, Jaimie M. Henderson, Krishna V. Shenoy, John P. Donoghue, and Leigh R. Hochberg. Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Sci Transl Med*, 7:313 179,.

[48] Tomislav Milekovic, Anish A. Sarma, Daniel Bacher, John D. Simeral, Jad Saab, Chethan Pandarinath, Brittany L. Sorice, Christine Blabe, Erin M. Oakley, Kathryn R. Tringale, Emad Eskandar, Sydney S. Cash, Jaimie M. Henderson, Krishna V. Shenoy, John P. Donoghue, and Leigh R. Hochberg. Stable long-term bci-enabled communication in als and locked-in syndrome using lfp signals. *Journal of Neurophysiology*, 120(1):343–360, Jul 2018.

[49] S. Alizon, A. Hurford, N. Mideo, and M. Van Baalen. Virulence evolution and the trade-off hypothesis: history, current state of affairs and the future. *Journal of Evolutionary Biology*, 22(2):245–259, December 2008. Publisher: Wiley.

[50] R. M. Anderson and R. M. May. Coevolution of hosts and parasites. *Parasitology*, 85(2):411–426, October 1982. Publisher: Cambridge University Press (CUP).

[51] James J. Bull and Adam S. Lauring. Theory and Empiricism in Virulence Evolution. *PLoS Pathogens*, 10(10):e1004387, October 2014. Publisher: Public Library of Science (PLoS).

[52] Oyungerel Byambasuren, Magnolia Cardona, Katy Bell, Justin Clark, Mary-Louise McLaws, and Paul Glasziou. Estimating the extent of asymptomatic COVID-19 and its potential for community transmission: Systematic review and meta-analysis. *Official Journal of the Association of Medical Microbiology and Infectious Disease Canada*, 5(4):223–234, December 2020. Publisher: University of Toronto Press Inc. (UTPress).

[53] Michael A. Johansson, Talia M. Quandelacy, Sarah Kada, Pragati Venkata Prasad, Molly Steele, John T. Brooks, Rachel B. Slayton, Matthew Biggerstaff, and Jay C. Butler. SARS-CoV-2 Transmission From People Without COVID-19 Symptoms. *JAMA Network Open*, 4(1):e2035057, January 2021. Publisher: American Medical Association (AMA).

[54] Angela L. Rasmussen and Saskia V. Popescu. SARS-CoV-2 transmission without symptoms. *Science*, 371(6535):1206–1207, March 2021. Publisher: American Association for the Advancement of Science (AAAS).

[55] W. Joost Wiersinga, Andrew Rhodes, Allen C. Cheng, Sharon J. Peacock, and Hallie C. Prescott. Pathophysiology, Transmission, Diagnosis, and Treatment of Coronavirus Disease 2019 (COVID-19). *JAMA*, 324(8):782, August 2020. Publisher: American Medical Association (AMA).

[56] Wei-jie Guan, Zheng-yi Ni, Yu Hu, Wen-hua Liang, Chun-quan Ou, Jian-xing He, Lei Liu, Hong Shan, Chun-liang Lei, David S.C. Hui, and et al. Clinical Characteristics of Coronavirus Disease 2019 in China. *New England Journal of Medicine*, February 2020. Publisher: Massachusetts Medical Society.

[57] Eskild Petersen, Marion Koopmans, Unyeong Go, Davidson H Hamer, Nicola Petrosillo, Francesco Castelli, Merete Storgaard, Sulien Al Khalili, and Lone Simonsen. Comparing SARS-CoV-2 with SARS-CoV and influenza pandemics. *The Lancet Infectious Diseases*, 20(9):e238–e244, September 2020. Publisher: Elsevier BV.

[58] N. Barkai and S. Leibler. Robustness in simple biochemical networks. *Nature*, 387(6636):913–917, June 1997. Publisher: Springer Science and Business Media LLC.

[59] T.-M. Yi, Y. Huang, M. I. Simon, and J. Doyle. Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences*, 97(9):4649–4653, April 2000. Publisher: Proceedings of the National Academy of Sciences.

[60] Jörg Stelling, Uwe Sauer, Zoltan Szallasi, III Doyle, Francis J., and John Doyle. Robustness of Cellular Functions. *Cell*, 118(6):675–685, September 2004. Publisher: Elsevier BV.

[61] F. A. Chandra, G. Buzi, and J. C. Doyle. Glycolytic Oscillations and Limits on Robust Efficiency. *Science*, 333(6039):187–192, July 2011. Publisher: American Association for the Advancement of Science (AAAS).

[62] H. El-Samad, H. Kurata, J. C. Doyle, C. A. Gross, and M. Khammash. Surviving heat shock: Control strategies for robustness and performance. *Proceedings of the National Academy of Sciences*, 102(8):2736–2741, January 2005. Publisher: Proceedings of the National Academy of Sciences.

[63] Ioannis Lestas, Glenn Vinnicombe, and Johan Paulsson. Fundamental limits on the suppression of molecular fluctuations. *Nature*, 467(7312):174–178, September 2010. Publisher: Springer Science and Business Media LLC.

[64] N. Li, J. Cruz, C. S. Chien, S. Sojoudi, B. Recht, D. Stone, M. Csete, D. Bahmiller, and J. C. Doyle. Robust efficiency and actuator saturation explain healthy heart rate control and variability. *Proceedings of the National Academy of Sciences*, 111(33):E3476–E3485, August 2014. Publisher: Proceedings of the National Academy of Sciences.

[65] Arjun Srinivasan, Elizabeth C. Burton, Matthew J. Kuehnert, Charles Rupprecht, William L. Sutker, Thomas G. Ksiazek, Christopher D. Paddock, Jeannette Guarner, Wun-Ju Shieh, Cynthia Goldsmith, and et al. Transmission of Rabies Virus from an Organ Donor to Four Transplant Recipients. *New England Journal of Medicine*, 352(11):1103–1111, March 2005. Publisher: Massachusetts Medical Society.

[66] Pauline Vetter, II Fischer, William A., Manuel Schibler, Michael Jacobs, Daniel G. Bausch, and Laurent Kaiser. Ebola Virus Shedding and Transmission: Review of Current Evidence. *Journal of Infectious Diseases*, 214(suppl 3):S177–S184, July 2016. Publisher: Oxford University Press (OUP).

[67] C. Fraser, S. Riley, R. M. Anderson, and N. M. Ferguson. Factors that make an infectious disease outbreak controllable. *Proceedings of the National Academy of Sciences*, 101(16):6146–6151, April 2004. Publisher: Proceedings of the National Academy of Sciences.

[68] Ben Killingley and Jonathan Nguyen-Van-Tam. Routes of influenza transmission. *Influenza and Other Respiratory Viruses*, 7:42–51, August 2013. Publisher: Wiley.

[69] Jason R. Marden and Jeff S. Shamma. Game theory and control. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):105–134, May 2018.

[70] Prasith Baccam, Catherine Beauchemin, Catherine A. Macken, Frederick G. Hayden, and Alan S. Perelson. Kinetics of Influenza A Virus Infection in Humans. *Journal of Virology*, 80(15):7590–7599, August 2006. Publisher: American Society for Microbiology.

[71] M. A. Nowak, S. Bonhoeffer, A. M. Hill, R. Boehme, H. C. Thomas, and H. McDade. Viral dynamics in hepatitis B virus infection. *Proceedings of the National Academy of Sciences*, 93(9):4398–4402, April 1996. Publisher: Proceedings of the National Academy of Sciences.

[72] Yinon M Bar-On, Avi Flamholz, Rob Phillips, and Ron Milo. SARS-CoV-2 (COVID-19) by the numbers. *eLife*, 9, March 2020. Publisher: eLife Sciences Publications, Ltd.

[73] Rob Phillips and Ron Milo. *Cell biology by the numbers*. Garland Science, 2015.

[74] Michael A. Matthay, Rachel L. Zemans, Guy A. Zimmerman, Yaseen M. Arabi, Jeremy R. Beitler, Alain Mercat, Margaret Herridge, Adrienne G. Randolph, and Carolyn S. Calfee. Acute respiratory distress syndrome. *Nature Reviews Disease Primers*, 5(1), March 2019. Publisher: Springer Science and Business Media LLC.

[75] Zhuo Zhou, Lili Ren, Li Zhang, Jiaxin Zhong, Yan Xiao, Zhilong Jia, Li Guo, Jing Yang, Chun Wang, Shuai Jiang, and et al. Overly Exuberant Innate Immune Response to SARS-CoV-2 Infection. *SSRN Electronic Journal*, 2020. Publisher: Elsevier BV.

[76] David A. Edwards, Dennis Ausiello, Jonathan Salzman, Tom Devlin, Robert Langer, Brandon J. Beddingfield, Alyssa C. Fears, Lara A. Doyle-Meyers, Rachel K. Redmann, Stephanie Z. Killeen, and et al. Exhaled aerosol increases with COVID-19 infection, age, and obesity. *Proceedings of the National Academy of Sciences*, 118(8):e2021830118, February 2021. Publisher: Proceedings of the National Academy of Sciences.

[77] Aartik Sarma, Stephanie A. Christenson, Eran Mick, Catherine DeVoe, Thomas Deiss, Angela Oliveira Pisco, Rajani Ghale, Alejandra Jauregui, Ashley Byrne, Farzad Moazed, and et al. *COVID-19 ARDS is characterized by a dysregulated host response that differs from cytokine storm and may be modified by dexamethasone*. Cold Spring Harbor Laboratory, January 2021.

[78] Joachim L. Schultze and Anna C. Aschenbrenner. COVID-19 and the human innate immune system. *Cell*, 184(7):1671–1692, April 2021. Publisher: Elsevier BV.

[79] Qian Zhang, Paul Bastard, Zhiyong Liu, Jérémie Le Pen, Marcela Moncada-Velez, Jie Chen, Masato Ogishi, Ira K. D. Sabli, Stephanie Hodeib, Cecilia Korol, and et al. Inborn errors of type I IFN immunity in patients with

life-threatening COVID-19. *Science*, 370(6515):eabd4570, September 2020. Publisher: American Association for the Advancement of Science (AAAS).

[80] Paul Bastard, Lindsey B. Rosen, Qian Zhang, Eleftherios Michailidis, Hans-Heinrich Hoffmann, Yu Zhang, Karim Dorgham, Quentin Philippot, Jérémie Rosain, Vivien Béziat, and et al. Autoantibodies against type I IFNs in patients with life-threatening COVID-19. *Science*, 370(6515):eabd4585, September 2020. Publisher: American Association for the Advancement of Science (AAAS).

[81] Alexis J. Combes, Tristan Courau, Nicholas F. Kuhn, Kenneth H. Hu, Arja Ray, William S. Chen, Nayvin W. Chew, Simon J. Cleary, Divyashree Kushnoor, and et al. Global absence and targeting of protective immune states in severe COVID-19. *Nature*, 591(7848):124–130, January 2021. Publisher: Springer Science and Business Media LLC.

[82] Matni N and Doyle JC. A theory of dynamics, control and optimization in layered architectures. *American Control Conference*, 2015.

[83] Yorie Nakahira, Nikolai Matni, and John C. Doyle. Hard limits on robust control over delayed and quantized communication channels with applications to sensorimotor control. *IEEE Conference on Decision and Control*, 2015.

[84] Yorie Nakahira, Quanying Liu, Natalie Bernat, Terrence Sejnowski, and John C. Doyle. Theoretical foundations for layered architectures and speed-accuracy tradeoffs in sensorimotor control. *American Control Conference*, 2019.

[85] Brockett RW and Liberzon D. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, 2000.

[86] Victoria Kostina, Yuval Peres, Gireeja Ranade, and Mark Sellke. Exact minimum number of bits to stabilize a linear system. *Arxiv*, 2018.

[87] Nicola Elia and Sanjoy Mitter. Stabilization of linear systems with limited information. *IEEE Transactions on Automatic Control*, 2001.

[88] Nair G and Evans RJ. Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM Journal on Control and Optimization*, 2004.

[89] Ames AD et al. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 2017.

[90] Sanger TD. Risk-aware control. *Neural Computation*, 2014.

[91] Olsman N and Goentoro L. Allosteric proteins as logarithmic sensors. *PNAS*, 2016.

[92] Gunter Stein. Respect the unstable. *IEEE Control Systems*, 23(4).

[93] John C. Doyle. Guaranteed margins for lqg regulators, Aug 1978.

[94] Munther A Dahleh and Ignacio J Diaz-Bobillo. *Control of uncertain systems: a linear programming approach*. Prentice-Hall, Inc., 1994.

[95] K. Zhou, J. C. Doyle, and K. Glover. *Robust and optimal control*. Prentice Hall New Jersey, 1996.

[96] Michael Rotkowitz and Sanjay Lall. A characterization of convex problems in decentralized control. *Automatic Control, IEEE Transactions on*, 51(2):274–286, 2006.

[97] A. Mahajan, N.C. Martins, M.C. Rotkowitz, and S. Yuksel. Information structures in optimal decentralized control. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 1291–1306, 2012.

[98] Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. A system level approach to controller synthesis. *IEEE Transactions on Automatic Control*, 2019.

[99] Yang Zheng, Luca Furieri, Antonis Papachristodoulou, Na Li, and Maryam Kamgarpour. On the equivalence of youla, system-level and input-output parameterizations. *arXiv preprint arXiv:1907.06256*, 2019.

[100] C. Langbort, R. S. Chandra, and R. D'Andrea. Distributed control design for systems interconnected over an arbitrary graph. *IEEE Transactions on Automatic Control*, 49(9):1502–1519, Sep. 2004.

[101] Nikolai Matni. Distributed control subject to delays satisfying an h_∞ norm bound. In *53rd IEEE Conference on Decision and Control*, pages 4006–4013. IEEE, 2014.

[102] Laurent Lessard. State-space solution to a minimum-entropy h_∞-optimal control problem with a nested information constraint. In *53rd IEEE Conference on Decision and Control*, pages 4026–4031. IEEE, 2014.

[103] C. A. Rösinger and C. W. Scherer. Structured controller design with applications to networked systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 4771–4776, Dec 2017.

[104] M. Ahmadi, M. Cubuktepe, U. Topcu, and T. Tanaka. Distributed synthesis using accelerated admm. In *2018 Annual American Control Conference (ACC)*, pages 6206–6211, June 2018.

[105] James Anderson, John C Doyle, Steven H Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 2019.

[106] Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. Localized lqr optimal control. In *53rd IEEE Conference on Decision and Control*, pages 1661–1668. IEEE, 2014.

[107] Yuh-Shyang Wang and Nikolai Matni. Localized lqg optimal control for large-scale systems. In *2016 American Control Conference (ACC)*, pages 1954–1961. IEEE, 2016.

[108] Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. Separable and localized system-level synthesis for large-scale systems. *IEEE Transactions on Automatic Control*, 63(12):4234–4249, 2018.

[109] John C. Doyle, Nikolai Matni, Yuh-Shyang Wang, James Anderson, and Steven Low. System level synthesis: A tutorial. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, Dec 2017.

[110] James Anderson, John C. Doyle, Steven H. Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, 2019.

[111] Munther A. Dahleh and Mustafa H. Khammash. Controller design for plants with structured uncertainty. *Automatica*, 29(1):37–56, Jan 1993.

[112] M. Khammash and J.B. Pearson. Performance robustness of discrete-time systems with structured uncertainty. *IEEE Transactions on Automatic Control*, 36(4):398–412, Apr 1991.

[113] Jing Shuang (Lisa) Li and John C. Doyle. Distributed robust control for systems with structured uncertainties. Submitted to IEEE Conference on Decision and Control 2022.

[114] Yuh-Shyang Wang, Nikolai Matni, and John C. Doyle. Separable and localized system-level synthesis for large-scale systems. *IEEE Transactions on Automatic Control*, 63(12):4234–4249, Dec 2018.

[115] Stephen Boyd. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.

[116] Olle Kjellqvist and John C. Doyle. ν-analysis: A new notion of robustness for large systems with structured uncertainties. Submitted to IEEE Conference on Decision and Control 2022.

[117] Nikolai Matni, Yuh-Shyang Wang, and James Anderson. Scalable system level synthesis for virtually localizable systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3473–3480. IEEE, 2017.

[118] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.

[119] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.

[120] M. E. Csete. Reverse engineering of biological complexity. *Science*, 295(5560):1664–1669, Mar 2002.

[121] Andrew Packard and John C. Doyle. The complex structured singular value. *Automatica*, 29(1):71–109, 1993.

[122] Jackson Liang, Shruthi Balachandra, Sang Ngo, and Lucy Erin O'Brien. Feedback regulation of steady-state epithelial turnover and organ size. *Nature*, 548(7669):588–591, Aug 2017.

[123] Tomislav Domazet-Lošo, Alexander Klimovich, Boris Anokhin, Friederike Anton-Erxleben, Mailin J. Hamm, Christina Lange, and Thomas C.G. Bosch. Naturally occurring tumours in the basal metazoan hydra. *Nature Communications*, 5(1), Jun 2014.

[124] R. Cathriona Millane, Justyna Kanska, David J. Duffy, Cathal Seoighe, Stephen Cunningham, Günter Plickert, and Uri Frank. Induced stem cell neoplasia in a cnidarian by ectopic expression of a pou domain transcription factor. *Development*, 138(12):2429–2439, Jun 2011.

[125] Donald F. Squires. Neoplasia in a coral? *Science*, 148(3669):503–505, Apr 1965.

[126] Michael J. Abrams, Ty Basinger, William Yuan, Chin-Lin Guo, and Lea Goentoro. Self-repairing symmetry in jellyfish through mechanically driven reorganization. *Proceedings of the National Academy of Sciences*, 112(26), Jun 2015.

[127] David A. Gold and David K. Jacobs. Stem cell dynamics in cnidaria: are there unifying principles? *Development Genes and Evolution*, 223(1–2):53–66, Nov 2012.

[128] Claudia M. Arenas Gómez, Keith Z. Sabin, and Karen Echeverri. Wound healing across the animal kingdom: Crosstalk between the immune system and the extracellular matrix. *Developmental Dynamics*, 249(7):834–846, May 2020.

[129] Zhihai Qin, Hye-Jung Kim, Jens Hemme, and Thomas Blankenstein. Inhibition of methylcholanthrene-induced carcinogenesis by an interferon receptor–dependent foreign body reaction. *Journal of Experimental Medicine*, 195(11):1479–1490, Jun 2002.

[130] Jerry E. Chipuk, Tomomi Kuwana, Lisa Bouchier-Hayes, Nathalie M. Droin, Donald D. Newmeyer, Martin Schuler, and Douglas R. Green. Direct activation of bax by p53 mediates mitochondrial membrane permeabilization and apoptosis. *Science*, 303(5660):1010–1014, Feb 2004.

[131] Xiao Tian, Jorge Azpurua, Christopher Hine, Amita Vaidya, Max Myakishev-Rempel, Julia Ablaeva, Zhiyong Mao, Eviatar Nevo, Vera Gorbunova, and Andrei Seluanov. High-molecular-mass hyaluronan mediates the cancer resistance of the naked mole rat. *Nature*, 499(7458):346–349, Jun 2013.

[132] Julia C. Liu, Xiao Guan, Jeremy A. Ryan, Ana G. Rivera, Caroline Mock, Vishesh Agarwal, Anthony Letai, Paul H. Lerou, and Galit Lahav. High mitochondrial priming sensitizes hescs to dna-damage-induced apoptosis. *Cell Stem Cell*, 13(4):483–491, Oct 2013.

[133] Jonathan D. G. Jones and Jeffery L. Dangl. The plant immune system. *Nature*, 444(7117):323–329, Nov 2006.

[134] Matthew L. Meizlish, Ruth A. Franklin, Xu Zhou, and Ruslan Medzhitov. Tissue homeostasis and inflammation. *Annual Review of Immunology*, 39(1):557–581, Apr 2021.

[135] Lucy MacCarthy-Morrogh and Paul Martin. The hallmarks of cancer are also the hallmarks of wound healing. *Science Signaling*, 13(648), Sep 2020.

[136] J Burkert, NA Wright, and MR Alison. Stem cells and cancer: an intimate relationship. *The Journal of Pathology*, 209(3):287–297, 2006.

[137] A. J. Becker, E. A. McCulloch, and J. E. Till. Cytological demonstration of the clonal nature of spleen colonies derived from transplanted mouse marrow cells. *Nature*, 197(4866):452–454, Feb 1963.

[138] Alfred G. Knudson Jr. Mutation and cancer: Statistical study of retinoblastoma. *Proceedings of the National Academy of Sciences*, 68(4):820–823, Apr 1971.

[139] Douglas Hanahan and Robert A. Weinberg. Hallmarks of cancer: The next generation. *Cell*, 144(5):646–674, Mar 2011.

[140] Tannishtha Reya, Sean J. Morrison, Michael F. Clarke, and Irving L. Weissman. Stem cells, cancer, and cancer stem cells. *Nature*, 414(6859):105–111, Nov 2001.

[141] Cristian Tomasetti and Bert Vogelstein. Variation in cancer risk among tissues can be explained by the number of stem cell divisions. *Science*, 347(6217):78–81, Jan 2015.

[142] Pedro M. Aponte and Andrés Caicedo. Stemness in cancer: Stem cells, cancer stem cells, and their microenvironment. *Stem Cells International*, 2017:1–17, 2017.

[143] A Mantovani, S Sozzani, M Locati, P Allavena, and A Sica. Macrophage polarization: tumor-associated macrophages as a paradigm for polarized m2 mononuclear phagocytes. *Trends in Immunology*, 23(11):549–555, Nov 2002.

[144] Richard Peto. Epidemiology, multistage models, and short-term mutagenicity tests, 1977.

[145] Andrei Seluanov, Vadim N. Gladyshev, Jan Vijg, and Vera Gorbunova. Mechanisms of cancer resistance in long-lived mammals. *Nature Reviews Cancer*, 18(7):433–441, Apr 2018.

[146] Néstor J. Oviedo and Wendy S. Beane. Regeneration: The origin of cancer or a possible cure? *Seminars in Cell Developmental Biology*, 20(5):557–564, Jul 2009.

[147] John J. Hopfield. Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity, Oct 1974.