

INTELLIGENT INFORMATION GATHERING: USING CONTROL FOR SENSING AND DECISION MAKING

Thesis by
Timothy H. Chung

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



California Institute of Technology
Pasadena, California

2007
(Defended May 18, 2007)

© 2007

Timothy H. Chung

All Rights Reserved

I humbly dedicate my thesis to my family, who have always been my greatest teachers:
my sister, who inspires me with her strength and courage to persevere in my
journey through academics as well as in life,
my mother, who elevates me with her love and support to the highest of my potential,
and my father, who guides me with his memory and spirit to becoming a great man myself.

– THC

Acknowledgements

If I have seen further it is by standing on the shoulders of giants.

– Isaac Newton

I wish to extend my warmest thanks to the various people who have made my graduate career such a rewarding and positive experience. My greatest gratitude goes to my advisors, Joel Burdick and Richard Murray. Through their mentorship and encouragement, my time at Caltech has been accentuated by many great opportunities for my academic as well as personal development.

The guidance given by my thesis committee members, Erik Antonsson and Pietro Perona, has been invaluable in probing and uncovering new ideas and avenues, and in conjunction with my advisors, have helped make this thesis possible. A special thanks goes to Maria Koeper for her company and her chocolates, as well as her many efforts on my behalf.

A great number of collaborators during my time at Caltech have enriched my knowledge as well as friendships, and to all of them, I owe a great debt of thanks for their enthusiasm, inspiration and patience. Furthermore, SOPS (The Society of Professional Students) and the great network of friends in Thomas building have been a wonderful family to me, and have been an unending source of fond memories.

Last, but most definitely not least, I thank my family for their endless support and faith in me. They have truly been the giants in my life, upon whose shoulders I am proud to stand and see so much.

Abstract

Information is everywhere and evolving, which necessitates both deliberate and efficient processing to acquire a good understanding of the dynamic situation, environment, or system of interest. Intelligent agents such as autonomous mobile sensors can control the way they gather information and thereby take advantage of feedback to improve the quality of that information. This approach reflects a shift from traditional “sensing for control” notions to “control for sensing” methods for addressing information-based objectives. This thesis presents several algorithms for distributed sensing tasks in the context of a team of mobile sensing agents. Applications of these types of mobile sensor networks include target tracking, dynamic environment monitoring, and distributed classification. These methods point beyond the use of sensory data for control and toward a framework for using control to improve information-based decisions made by intelligent agents. The sequential decision-theoretic framework presented herein has relevant applications in engineered systems such as search and rescue using a robotic team, as well as potential connections to natural systems including search strategies in the human vision system.

Contents

Acknowledgements	iv
Abstract	v
1 Introduction	1
1.1 Motivation and Scope of Thesis	1
1.2 Background and Relevant Works	2
1.3 Summary of Contributions	4
1.4 Overview of Thesis	6
2 Distributed Sensing and Estimation	8
2.1 Distributed Estimation of a Dynamic Process	9
2.1.1 Evolution of Dynamic Processes	9
2.1.2 Measurement Process	10
2.1.2.1 Heterogeneous Measurement Models	11
2.1.2.2 Observation of Multiple Dynamic Processes	12
2.1.3 Case Study: Distributed Target Tracking	12
2.2 Distributed Sensor Fusion	15
2.2.1 Fusion of Local Observations	17
2.2.2 Fusion of Locally Filtered Estimates	17
2.2.3 Case Study: Distributed Target Tracking	19
2.2.4 Construction of the Distributed Sensing Cost Function	19
2.3 Gradient Analysis	20
2.3.1 Gradient-Descent Control Law	21
2.3.2 Gradient Analysis for Tracking Multiple Dynamic Processes	22
2.3.3 Case Study: Distributed Target Tracking, Revisited	24
2.4 Extension to Imperfect Communication	26
2.4.1 Case Study: Distributed Target Tracking with Imperfect Communication . .	29
2.4.2 Relationship between Sensing and Communication	31

2.5	Extension to Classification	32
2.5.1	Classification Probability	33
2.5.1.1	Gaussian noise and class distributions	34
2.5.2	Relationship between Estimation and Classification	35
3	Control for Decision Making	38
3.1	Dynamics of Decision Making	40
3.1.1	Sequential Probability Ratio Test	40
3.1.1.1	Measures of Performance of Decisions	41
3.1.2	Application of Sequential Decision Making to Search	43
3.2	Probabilistic Search as a Decision	44
3.2.1	Problem Setup	44
3.2.2	Search is a Decision Problem	45
3.2.3	Detection of the Target	46
3.2.4	Case Study: Search on a Grid	47
3.3	Discrete Bayesian Formulation	48
3.3.1	Recursive Computation of the Cell Belief Probabilities	48
3.3.2	Closed-Form Expressions for Uniform PDFs	50
3.3.3	Case Study: Search on a Grid, Revisited	53
3.3.3.1	Random Search	53
3.3.3.2	Sweeping Search	54
3.4	Control Strategies for Decision Making	56
3.4.1	Look Ahead Search	57
3.4.2	“ <i>Drosophila</i> -Inspired” Search	60
3.4.3	“Saccadic” Search	61
3.4.4	Discussion on Control Strategies	63
4	Conclusions and Future Directions	65
4.1	Discussion and Summary	65
4.2	Future Directions	66
A	Derivations of Matrix Gradients	68
A.1	Gradient expression for J_1	68
A.2	Gradient expression for J_2	69
A.3	Gradient expression for $J_{3,j}$	70

B Analysis of the Cost Function	72
B.1 Cost Function Expression for Two Sensors	72
B.2 Cost Function Expression for Three Sensors	75
Bibliography	78

List of Figures

2.1	Illustration of the distributed target tracking problem.	13
2.2	Schematic of range and bearing observation model.	14
2.3	Illustration of the benefit of fusing multiple observations.	19
2.4	Three sensors track a (biased) randomly walking target.	25
2.5	Evolution of the cost for an illustrative simulation run.	26
2.6	Effect of communication constraints on sensing.	30
2.7	Performance trade-off a function of communication parameter, μ	31
2.8	Performance trade-off as function of sensing parameter, a_1	32
3.1	Schematic of the evolution of three sample decision processes.	42
3.2	Example of discretized search region, \mathcal{A}	45
3.3	Illustration of Gaussian target prior probability distribution.	47
3.4	Illustration of uniform target prior probability distribution.	51
3.5	Belief function evolution as a function of false alarm rate, α	52
3.6	Belief function evolution as a function of missed detection rate, β	53
3.7	Effect of receiving a nonnull observation on the belief evolution.	54
3.8	Belief probability distribution evolution: random walk search	55
3.9	Belief probability distribution evolution: sweeping search	56
3.10	Simulation: look ahead search strategy for varying window sizes	59
3.11	Simulation: “ <i>Drosophila</i> -inspired” search strategy	61
3.12	Simulation: “saccadic” search strategy	62
B.1	Relative bearing dependence of cost function for two sensors.	76
B.2	Relative bearing dependence of cost function for three sensors.	77

Chapter 1

Introduction

1.1 Motivation and Scope of Thesis

Autonomy requires the ability to acquire information, process it appropriately and efficiently, and reason about its content and context. For autonomous robotic agents, increasingly intelligent systems can begin to escape their confinement to tasks traditionally relegated to robots called the “Three D’s,” i.e., tasks that are dirty, dull, and/or dangerous. Equipped with greater abilities to gather useful information, perceive relevant features, and execute efficient actions, autonomous and intelligent agents can elevate humanity. In this way, autonomous systems can serve as focusing elements of a vast amount of information that is present in the environment.

Given the importance of the information gathering capabilities to autonomous systems, efforts to improve the quality of their acquired information merit further study to further augment. The desire to understand the role of feedback control on enhancing the information gathering process motivates the research presented in this thesis.

The general theme of this research is the use of feedback control to improve the gathering of information as a mechanism for endowing autonomous systems with greater intelligence. Information has many representations and therefore require investigations into the role of feedback from various perspectives, whether at the level of sensory data (e.g., range and bearing measurements) or at the higher level of decisions (e.g., presence/absence of a target).

In particular, of interest is the class of problems where mobility is the control mechanism by which the quality of the information is improved. For many types of robotic systems, motion control is the primary means of actuation, and its role in information gathering tasks is relevant for a variety of applications, such as target detection and tracking, autonomous search, robotic mapping, cooperation, etc.

1.2 Background and Relevant Works

Given the diversity and ubiquity of information gathering tasks, the topic of using control to improve the quality of sensing is a rapidly evolving field of research, and as such, has engendered much attention in recent years. The advent of low cost sensing, actuation and computation has further accelerated this line of research, spanning the various areas of active sensing, mobile sensor networks and multi-agent approaches for dynamic target tracking. The work presented in Chapter 2 draws on elements of all these topics.

The field of active sensing rests on the general principle that the coupling between actuation and sensing components in systems should be taken advantage of, balancing the benefit of improved information gain against the added cost of controlling the system. Relevant to many robotic applications, such as vision [1, 2, 3], manipulation [4], and exploration [5, 6, 7], active sensing in the context of this thesis corresponds to the use of mobility to maneuver a team of mobile sensor agents to improve their sensing performance. Related to the notions of optimal experiment design, Ucin-ski [8, 9] proposes a mathematical framework for addressing optimal sensor trajectories (of which optimal sensor placement [10, 11] is a special case) utilizing the Fisher Information Matrix. The inverse of this information matrix, which is simply the estimation covariance matrix, is widely used to define a cost function to govern the optimization process, as done in the work presented in this thesis.

Closely related but more specific to the task of improved information gathering of dynamic targets, extensive research has been done in a diverse array of approaches. Built on the foundations of target tracking work by Bar-Shalom and Fortman [12], Bar-Shalom and Li [13], and Rao, Durrant-Whyte, and Sheen [14], these approaches examine different but intrinsically related measures of the estimation process, applying different methodologies for the aggregation and propagation of the information. Spletzer and Taylor [15] recognize the computational challenges of optimization of motion trajectories for multiple mobile sensors, and require the use of approximate particle filter methods to solve the problem. The information-theoretic approach presented by Logothetis [16] optimizes trajectories for bearings-only sensing by maximizing the information content in observations. Grocholsky [17] and his collaborators [6] employ a decentralized framework, the Decentralized Data Fusion algorithm [18], for gathering and exchanging information quantities among a team of mobile sensors. These methods illustrate both the complexity of the distributed estimation task as well as the desire to find distributed or decentralized solutions for applicability to teams of autonomous mobile sensing agents.

Additionally, communication in the network of mobile sensors adds another component to the challenges of sensing and control [19, 20]. The previous approaches for multi-robot systems do not address practical issues of communication between agents, given that computation of a distributed

motion control law for improved sensing is difficult enough. While mobile sensor networks have received considerable focus from the networking community, much of the emphasis has been on managing only the communication performance (e.g., transmission characteristics, network connectivity, or energy efficiency) via node mobility without specific consideration for utilizing motion to improve the quality of the sensing. In contrast, work by Mostofi and others [21, 22] recognize that a cross-layer (between communication and actuation layers) design approach is necessary, and investigates the balancing act of maintaining acceptable sensing performance while being constrained by communication limitations. The resulting representation of optimal sensor positions are achieved using a variable step-size (discrete) motion controller. Investigation of other communication constraints such as network connectivity in relation to agent motion is conducted in Spanos, Olfati-Saber, and Murray [23].

These previous works acknowledge the hurdles of computation and communication for utilizing control to improve the quality of the gathered sensing information. While elegant in formulations as optimization problems, global solutions to these generally nonlinear optimal control problems require resorting to numerical techniques [24] and are impractical for implementation in real systems. In contrast, the control laws presented in this thesis are given by closed-form expressions, thereby enabling efficient realizations of the motion control strategies. Although solutions found by this analytic gradient approach are only guaranteed to be locally optimal [25, 26], the immense practical advantages in computation provide solid motivation and validation for their use. Additionally, these results extend naturally to incorporate communication constraints, such as that of fading wireless links between agents, preserving their closed-form structure as well as their usefulness in physical systems. Furthermore, the analysis leads to an intuitive understanding of the behavior of multi-agent sensing systems, and enables greater cognizance of the role of control in improving the performance of these distributed sensing and estimation approaches. The intuition, in turn, can be used to facilitate better design and application of these sensing, computation, and communication systems. This insight and the general closed-form expressions for the motion control laws comprise some of the main contributions of this thesis.

Research in decision making systems possesses a diverse heritage in robotics and artificial intelligence [27]. Various approaches for modeling the search process have been developed in the context of physical search of a target in a search region as developed in Chapter 3.

Search theory as a formal field of study developed out of the defining works of Koopman [28, 29, 30, 31] in the area of operations research. In the context of naval search and rescue [32], these works developed a probabilistic foundation of search theory incorporating notions of optimal allocation of resources and probability of detection of the target. In addition to these seminal works, Washburn [33], Stone [34], and many others [35] have continued to develop search theory in a probabilistic Bayesian context further, investigating classes of search paths, various models for

detection (e.g., radar, sonar, vision) and methods for addressing uncertainty in the system. Various models for target motion and behavior have also been studied [36]. More recently, advances in computation capabilities led Bourgault and his colleagues [37, 38] to reconstruct the problem of searching for lost targets, using the probability of detection of these targets as the objective function for optimal search trajectory generation. The Bayesian filtering approach [39, 40] of these works offers an advantage over other methods for maintaining and updating all information relevant to the search (e.g., the target probability density function), and as such serves as the backbone of the computation machinery presented in this thesis.

Unlike the physical search problem investigated above, the study of the role of feedback in such systems in a sequential decision-theoretic context is a relatively nascent area, and the construction and examination of the search problem in this framework is another main contribution of this thesis. Sequential notions in decision theory stem from the methods developed by Wald [41] and Wolfowitz [42] in the Sequential Probability Ratio Test (SPRT). These works laid the theoretical groundwork for iterative methods in experiment design and hypothesis testing [43], whose principles have been used in a wide variety of decision making applications. The nature of the SPRT is recognized, as presented in Chapter 3, to represent an intrinsic notion of dynamics or evolution of the decision process, for which the principle of feedback can be utilized to drive the process in a controlled manner. For search, this feedback occurs in the form of motion control of the searcher’s trajectory.

Motivated by the Bayesian construction and the dynamics of the SPRT structure, a decision-theoretic approach to the physical search problem is realized. Due to the formulation as a decision, the evolution of the probability distributions relevant to the search task can be succinctly represented by a closed-form update expression, representing another major contribution of this work. In particular, the derived analytic expression allows for greater computational efficiency, traditionally a severely limiting factor in many Bayesian filter applications. Furthermore, the formulation presented in this thesis enables the generalization of the physical search task to a broader class of search-related objectives. Such search problems can be found in engineered systems, as in the physical search task examined in the previous works above, but also in natural systems, such as animal search behaviors [44] and human visual search [45], emphasizing the desire for a unifying approach to these search objectives. The methodology presented in this thesis provides investigations and insight into these problems and serves as a candidate approach for this generalization.

1.3 Summary of Contributions

Previous research on methods for distributed sensing and estimation have generally recognized the need for approximate approaches [15], as the optimal solution for the motion trajectories of multiple

mobile sensors is computationally expensive. These optimal control or numerical approaches are computationally prohibitive for implementation on practical systems, as well as lacking in their ability to provide insight into the role of motion control on improving the quality of the sensing information. The main contributions of Chapter 2 include:

- Derivation of closed-form expressions for a motion control law which locally minimize the overall uncertainty in observations of a dynamic target by a team of mobile sensors.
- Extension of the analytic formulas to address imperfect communication (e.g., fading wireless channels) between agents. The modified expressions preserve the general structure and highlight the trade-offs required to address sensing and communication objectives.
- Statement and proof of an “estimation-classification” duality principle, which relates the performance in distributed sensing tasks to that of classification. By employing the expressions for improving the sensing quality, the duality principle ensures that probability of correct classification is also improved.

The insight provided by these contributions enables a richer understanding of the distributed sensing task, as well as demonstrates the ease and efficiency for implementation of these analytic motion control laws in physical systems.

Similarly, the problem of physical search has also largely been restricted to numerical methodologies, incorporating assumptions on detection models that are mostly relevant for specific tasks, such as naval search applications. Furthermore, additional works rest on the special case of no false alarms [38] or other simplifications that reduce the computational challenge of updating the probabilistic representation of the target location. Additionally, given their focus on autonomous robotic search applications, these previous works do not easily generalize to a broader class of search tasks, such as that exhibited by animals or in human visual search. Chapter 3 presents several main contributions to this problem of search:

- Formulation of the physical search problem as a decision, enabling the use of sequential decision-theoretic notions for representing the belief probability of the target’s presence or absence within a search region.
- Derivation of closed-form expressions for the temporal evolution of the belief probability as a function of the sequential observations attained by a searcher as it moves through the search space.
- Presentation of two novel search control strategies – “*Drosophila*-inspired” and “saccadic” methods – for governing the searcher’s search trajectory as novel approaches to accomplish the search task, demonstrating the generality of the proposed framework in addressing a diverse set of search-related problems.

The construction of this search objective as a sequential decision enables the intuitive characterization of various search metrics, such as the time-till-decision, and illuminates the role of controlling the searcher’s motion trajectory in improving the decision making performance in the search application.

1.4 Overview of Thesis

Chapter 2 investigates the case where a team of sensors is given the task to gather and process data regarding the state of dynamic targets in the presence of a noisy environment. The benefits of cooperation amongst multiple sensing agents are evident in the reduction of overall uncertainty of the combined measurements.

Section 2.1 considers how to use motion control to improve the performance of distributed multi-agent sensing and estimation tasks. Combination of estimates via sensor fusion is described in Section 2.2. The fusion process guides the design of a cost function which captures the overall uncertainty in the fused estimates. Optimization of this cost function is discussed in Section 2.3, where the main result of closed-form expressions for the gradient of the cost are derived. These analytic formulas define the locally-optimal motion control laws that govern the motions of the team of mobile sensing agents. Extension of these control laws to the case where communication constraints between sensors are present is conducted in Section 2.4. This section also considers the trade-off between sensing and communication performance in the cooperating team. Furthermore, foreshadowing the desire to accomplish higher-level tasks with mobile sensors, the relationship between improved sensing and improved classification is studied and explicitly proved under mild assumptions on the underlying probability distributions.

Chapter 3 investigates a different aspect of control in intelligent systems, where mobility is used to influence the decision making process of mobile vehicles. In particular, for target detection in a search region, motion of the searcher is utilized to improve the performance of the search task, such as reducing the average time it takes to make that determination.

The relevance of decision theory as a tool for examining the role of feedback control on decision making is discussed in Section 3.1 where the notions of sequential testing is introduced along with measures of the performance of a decision process. Section 3.2 presents a formal framework which casts the search problem as a decision between binary hypotheses given the decision-maker’s current knowledge. In the specific case typically considered in the physical search theory literature, the decision reflects a belief of whether or not the target of interest is present in the search region. The temporal evolution of the decision by means of Bayesian filtering methods is described in Section 3.3, where novel formulas that govern the decision evolution process are derived. In Section 3.4, the explicit role of motion control of the searcher is investigated via several search control strategies. Several control strategies including two new strategies, are considered in this section. The

novel “*Drosophila*-inspired” and “saccadic” control strategies presented in Section 3.4 suggest the generalized framework presented in this thesis may have application to a broad set of problems. The analysis of the search strategies is enabled by the closed-form expressions presented in Section 3.3.

Chapter 2

Distributed Sensing and Estimation

Information exists everywhere and evolves endlessly. In order to understand the dynamic situation, surroundings, or system of interest, it is essential that this information be gathered and processed in a deliberate and efficient manner. Complicating this effort is the fact that in many settings, the information can best be captured by use of multiple perspectives, which augment the overall understanding, but which are obtained at the potential cost of additional confusion, overhead, or expenditure of energy.

This chapter examines some of these information-rich, distributed, dynamic processes, including the modeling of the distributed sensing framework, a measure of sensing performance, and a means of increasing this performance by use of feedback control.

Section 2.1 describes a general distributed sensing problem where the state of a dynamic process is estimated using multiple observations. This section develops a framework by which control can be introduced into the distributed sensing problem for purposes of improving the quality of the observed information.

Sensor fusion methods for combining these multiple measurements are discussed in Section 2.2, where the relationship between individual sensor control actions and their effect on the overall sensing performance are investigated. This examination results in a global representation and measure of the quality of information.

Optimization of this metric enables an improvement in the accuracy and understanding of the global system, and an analytic framework for this optimization via gradient methods is developed Section 2.3. In this section, the main contribution of this work is presented in the derivations of closed-form expressions for efficient control laws which collectively drive the sensor states to locally optimal configurations.

Finally, given the relevance of wireless communication in these distributed systems, a model for imperfect communication is examined in Section 2.4 and extend the gradient methods of the

previous section to address these communication constraints.

2.1 Distributed Estimation of a Dynamic Process

The distributed estimation task requires the use of multiple sources of information to estimate the state of a dynamic process. In order to accomplish this objective, one must have an understanding or model of the system state to be estimated, as well as a representation of the manner in which the information is acquired and processed.

In this section, a distributed estimation formulation is outlined in which multiple observations are used to estimate the state of a simple dynamical system.

2.1.1 Evolution of Dynamic Processes

Given that the world is constantly changing, it should come as no surprise that systems that change and evolve in time are of great interest. In examining the states of such a system, a model for the dynamics of the system is required, which is assumed to take the following general state-space form:

$$\mathbf{x}[k + 1] = \mathbf{f}(\mathbf{x}[k]) + \mathbf{w}[k],$$

where $\mathbf{x}[k] \in \mathbb{R}^n$ contains the state variables of interest at discrete time index k , and $\mathbf{f}(\mathbf{x})$ represents a nonlinear model of the dynamics by which the state of the system evolves. Further, the process noise, $\mathbf{w}[k] \in \mathbb{R}^n$, captures the fact that uncertainty exists in the environment as well as in the modeling of the process itself. For the moment, no restrictions on the types of allowable noises (e.g., bounded versus unbounded) are assumed, with $\mathbf{Q}[k]$ denoting the positive-definite process noise covariance.

Given the varied and versatile tools available for linear systems, it is useful to examine the linearization of the nonlinear process model above

$$\mathbf{x}[k + 1] = \mathbf{F}\mathbf{x}[k] + \mathbf{w}[k],$$

where $\mathbf{F} \in \mathbb{R}^{n \times n}$ is the linearized dynamics matrix. This dynamic process model holds for a multitude of systems, including, for example, the motion of a vehicle in space, where the state variables might consist of the vehicle's position and velocity components as a function of time. This example of a dynamic vehicle will be used throughout the development of the theory as a case study for the distributed sensing methodology, but note that the theory holds for observation of general dynamic processes.

2.1.2 Measurement Process

The necessity for accurate modeling of the observation process should be obvious, given the fact that the objective is to gather and process sensory information. For the distributed estimation task, one is particularly interested in integrating various sources of information to estimate the given dynamic process. Consider the case where M observations are taken at each time step to estimate the process state. Each of these observations obeys:

$$\mathbf{y}_i[k] = \mathbf{h}_i(\mathbf{x}[k]) + \mathbf{v}_i[k], \quad (2.1.1)$$

where $\mathbf{y}_i[k] \in \mathbb{R}^m$ ($i = 1, \dots, M$) is the observation of the i^{th} sensor, and $\mathbf{h}_i(\mathbf{x})$ is the nonlinear measurement function for this observation. Note that \mathbf{h}_i is generally a function of the dynamic process state. Furthermore, the i^{th} measurement is assumed to be perturbed by additive measurement noise, $\mathbf{v}_i[k] \in \mathbb{R}^m$ with its uncertainty represented by covariance $\mathbf{R}_i[k] \in \mathbb{R}^{m \times m}$. The measurement noise processes for different observations are assumed independent, i.e., $\mathbb{E}[\mathbf{v}_i \mathbf{v}_j^T] = 0$, where $\mathbb{E}[\cdot]$ and $(\cdot)^T$ are the expectation and transpose operators, respectively. In the case that these multiple observations are made using a team of M distinct sensors, \mathbf{y}_i corresponds to the observation generated by the i^{th} sensing agent.

As done previously, the linearized form of the measurement model is used

$$\mathbf{y}_i[k] = \mathbf{H}_i \mathbf{x}[k] + \mathbf{v}_i[k],$$

where $\mathbf{H}_i \in \mathbb{R}^{m \times m}$ is the linearized measurement matrix.

As embedded technologies combining sensing, computation and actuation become more prevalent, it is likely that the observation process will also be dependent on the state of the sensor. In other words, the i^{th} sensing system is also a dynamic process itself, whose state $\mathbf{z}_i[k] \in \mathbb{R}^s$ can also evolve as

$$\mathbf{z}_i[k+1] = \mathbf{\Phi}_i \mathbf{z}_i[k] + \mathbf{\Gamma}_i \mathbf{u}_i[k]. \quad (2.1.2)$$

The sensor state can be affected by application of the control signal $\mathbf{u}_i[k]$. In laying the groundwork for the distributed estimation problem, the case where the sensor system is linear and unperturbed by disturbances is initially investigated.

An example of sensor-actuator coupling is the control of a pan-tilt camera settings which are adjusted to improve the photograph. Another such case is the fine adjustments of the knob on a radio to tune to the right frequency for improved reception. In both of these examples, the objective is to use feedback control to improve the sensing performance.

In this manner, the observation model defined above is also a function of the sensor state, in

addition to the dynamic state to be estimated:

$$\mathbf{y}_i[k] = \mathbf{h}_i(\mathbf{x}[k], \mathbf{z}_i[k]) + \mathbf{v}_i[k].$$

In other words, by appropriate choice of control signal \mathbf{u}_i , the measurement process can be tailored to the sensing objective.

As formulated above, the observation function, and not the measurement noise covariance, can be affected by controlling the i^{th} sensor's state, \mathbf{z}_i . However, the uncertainty in the observation is typically defined by the measurement noise covariance [46, 47]. By an appropriate choice of coordinate transformation, the dependence on process and sensor states (\mathbf{x} and \mathbf{z}_i , respectively) and the shape of the measurement uncertainty can be captured in a single matrix quantity. In other words, by application of the transformation matrix $\mathbf{T}_i \in \mathbb{R}^{m \times m}$, the observation model becomes

$$\mathbf{T}_i \mathbf{y}_i[k] = \mathbf{T}_i \mathbf{h}(\mathbf{x}[k], \mathbf{z}_i[k]) + \mathbf{T}_i \mathbf{v}_i[k].$$

The coordinate transformation matrix, \mathbf{T}_i , will depend on the sensor state as well, thereby enabling control inputs to affect the transformed measurement noise covariance given by $\mathbf{T}_i \mathbf{R}_i[k] \mathbf{T}_i^T$.

Hence, by injecting control in the measurement process, the shape of the uncertainty profile can be manipulated. In the case of improved sensing, such actions reduce the effective size (e.g., volume or major axis) of the uncertainty envelope. The statement of the distributed sensing task can now be stated as follows: *Given M controllable sensing agents, each providing an observation of the dynamic process, the distributed sensing objective is to generate a coordinated control law such that the M agents maneuver to reduce an overall state estimation error.* In this way, feedback control is used to improve the quality of sensing.

2.1.2.1 Heterogeneous Measurement Models

The observation models for each of the M sensors need not be the same. The ensuing algorithms only require that the observations from these different sensors are compatible (i.e., measurements are of the same dynamic state). In other words, the framework described here allows for sensors with various sensing modalities to interact and collectively gather and share information in a straightforward manner.

For example, many different sensors for measuring a vehicle's spatial position exist, such as sonar range-finders, laser-scanning range devices, radar, and stereoscopic camera systems. Any and all combinations of these sensors may be utilized in the distributed sensing task, offering benefits of different resolution and sensor characteristics. This feature of the team sensing approach highlights the general nature of the framework and reflects its applicability to a variety of realistic systems.

2.1.2.2 Observation of Multiple Dynamic Processes

Additionally, one can model the evolution of N multiple dynamic processes, each with its own state, $\mathbf{x}_j \in \mathbb{R}^{n_j}$, simply by augmenting the state vector to comprise the multiple states. Independence of these different states facilitate an analytically simple structure (i.e., diagonal dynamics and noise covariance matrices), but coupled or constrained states can easily be incorporated in this framework.

Similarly, the observation process can also be modified to represent the simultaneous measurement of the multiple dynamic processes by augmentation of the measurement vectors. For independent observations, the uncertainty and respective transformation matrices can be redefined as block diagonal matrices:

$$\mathbf{R}_i \longrightarrow \text{diag}(\mathbf{R}_{i,1}, \dots, \mathbf{R}_{i,N}), \quad \mathbf{T}_i \longrightarrow \text{diag}(\mathbf{T}_{i,1}, \dots, \mathbf{T}_{i,N}),$$

for N simultaneously observed processes.

A challenge that arises when multiple processes are observed is that of data association. Namely, confusion may exist as to which measurement corresponds to which dynamic process, requiring a mechanism for identifying each process uniquely. As the theory and implementation of data association methodologies comprise an active area of research, we assume that simple but reliable methods, such as the use of identifying markers (e.g., color, visual patterns, and RFID), are employed in the observation of the multiprocess systems. The interested reader is referred to [12, 13] and the references therein for a survey of various data association approaches.

2.1.3 Case Study: Distributed Target Tracking

To illustrate the modeling approach described in the previous sections, the example problem of estimating the state of a planar dynamic vehicle can be considered, where $\mathbf{x} \in \mathbb{R}^2$ is the Cartesian position of the target vehicle's center of mass. Assume M multiple mobile sensing agents, each equipped with a (possibly different) range-and-bearing sensor such as sonar or LADAR (a.k.a. laser radar) units commonly used on many robotic platforms for target tracking tasks. A pictorial representation of the this system is given in Figure 2.1.

Let the state of the target vehicle evolve as

$$\mathbf{x}[k+1] = \mathbf{F}\mathbf{x}[k] + \mathbf{w}[k],$$

where the target has arbitrary dynamics (such as omnidirectional, car-like, or differential drive) represented by the linear dynamics matrix, $\mathbf{F} \in \mathbb{R}^{2 \times 2}$ and the process noise, $\mathbf{w} \in \mathbb{R}^2$, are disturbances (e.g., zero-mean, white and Gaussian) on the target's (x, y) position as it moves in the plane.

The observation model for typical target-tracking sensors is nonlinear, due to the coordinate




Figure 2.1: Illustration of the distributed target tracking problem. A team of $M = 2$ mobile sensing platforms take range-and-bearing measurements of the target position in \mathbb{R}^2 .

change from relative (i.e., body-fixed) to global coordinates. These sensors typically output noisy range and bearing measurements, r and b , respectively, to the target with respect to the local frame. Using $\|\cdot\|_2$ and $\angle(\cdot)$ to denote the Euclidean distance and relative angle in the sensor frame, the observation model is given by (e.g., [48]):

$$\begin{aligned} \mathbf{y}_i[k] &= \begin{pmatrix} r_i \\ b_i \end{pmatrix} = \begin{pmatrix} \|\mathbf{x} - \mathbf{z}_i\|_2 \\ \angle(\mathbf{x} - \mathbf{z}_i) \end{pmatrix} + \mathbf{v}_i[k] \\ &= \begin{pmatrix} \sqrt{(\mathbf{x}^x - \mathbf{z}_i^x)^2 + (\mathbf{x}^y - \mathbf{z}_i^y)^2} \\ \text{atan}\left(\frac{\mathbf{x}^y - \mathbf{z}_i^y}{\mathbf{x}^x - \mathbf{z}_i^x}\right) + \mathbf{z}_i^\theta \end{pmatrix} + \mathbf{v}_i[k], \end{aligned}$$

where the state of the i^{th} sensor, $\mathbf{z}_i \in \mathbb{R}^3$, comprises the Cartesian coordinates of the sensing agent's position, $(\mathbf{z}_i^x, \mathbf{z}_i^y)$, and its heading angle, \mathbf{z}_i^θ , and its evolution is given by the dynamic system described by Equation 2.1.2.

The linearization of the measurement function, evaluated at the sensor's current configuration,

is given by the measurement matrix, $\mathbf{H}_i \in \mathbb{R}^{2 \times 2}$:

$$\begin{aligned} \mathbf{H}_i &= \begin{pmatrix} \frac{\mathbf{x}^x - \mathbf{z}_i^x}{\sqrt{(\mathbf{x}^x - \mathbf{z}_i^x)^2 + (\mathbf{z}_i^y - \mathbf{z}_i^y)^2}} & \frac{\mathbf{x}^y - \mathbf{z}_i^y}{\sqrt{(\mathbf{x}^x - \mathbf{z}_i^x)^2 + (\mathbf{z}_i^y - \mathbf{z}_i^y)^2}} \\ -\frac{\mathbf{x}^y - \mathbf{z}_i^y}{(\mathbf{x}^x - \mathbf{z}_i^x)^2 + (\mathbf{z}_i^y - \mathbf{z}_i^y)^2} & \frac{\mathbf{x}^x - \mathbf{z}_i^x}{(\mathbf{x}^x - \mathbf{z}_i^x)^2 + (\mathbf{z}_i^y - \mathbf{z}_i^y)^2} \end{pmatrix} = \begin{pmatrix} \frac{\mathbf{x}^x - \mathbf{z}_i^x}{r_i} & \frac{\mathbf{x}^y - \mathbf{z}_i^y}{r_i} \\ -\frac{\mathbf{x}^y - \mathbf{z}_i^y}{r_i^2} & \frac{\mathbf{x}^x - \mathbf{z}_i^x}{r_i^2} \end{pmatrix} \\ &= \begin{pmatrix} \cos(b_i + \mathbf{z}_i^\theta) & \sin(b_i + \mathbf{z}_i^\theta) \\ -\frac{1}{r_i} \sin(b_i + \mathbf{z}_i^\theta) & \frac{1}{r_i} \cos(b_i + \mathbf{z}_i^\theta) \end{pmatrix}, \end{aligned}$$

where it can easily be seen how the observation can be changed by controlling the sensor state variables.

No conditions on the measurement noise, \mathbf{v}_i , are required in the general formulation, but for the sake of this example, let the observation be corrupted by zero-mean, white Gaussian noise, whose uncertainty envelope, consisting of both range and bearing errors, is shown in Figure 2.2.




Figure 2.2: Schematic of the observation model, capturing the uncertainty in range-and-bearing measurements of the target's position in \mathbb{R}^2 .

Note that the measurement uncertainty is represented in body-fixed coordinates. Consistent with standard range-and-bearing models [48] for target tracking in the plane, the measurement noise covariance matrix in this reference frame has the following two-dimensional diagonal structure:

$$\mathbf{R}_i = \begin{pmatrix} f_i^r(\mathbf{x}, \mathbf{z}_i) & 0 \\ 0 & f_i^b(\mathbf{x}, \mathbf{z}_i) \end{pmatrix} = \begin{pmatrix} (\sigma_i^r)^2 & 0 \\ 0 & (\sigma_i^b)^2 \end{pmatrix},$$

where $(\sigma_i^r)^2$ and $(\sigma_i^b)^2$ are the range and bearing measurement noise variances for the i^{th} sensor's observation, respectively. Note that the model above can represent a large class of observation uncertainties by allowing for the noise variances to possibly depend on the sensor and target states (e.g., range-to-target) via functions, f_i^r and f_i^b . Such state-dependent uncertainties are present in a

wide range of practical sensing modalities (e.g., atmospheric corrections in radar observations).

To transform the local (i.e., range and bearing) observations to a global (Cartesian) coordinate system (to facilitate compatibility of measurements from different sensors) use the (rotation) transformation matrix,

$$\mathbf{T}_i = \begin{pmatrix} \cos(\mathbf{z}_i^\theta) & -\sin(\mathbf{z}_i^\theta) \\ \sin(\mathbf{z}_i^\theta) & \cos(\mathbf{z}_i^\theta) \end{pmatrix}.$$

Application of this coordinate transformation to the linearized measurement model yields

$$\mathbf{T}_i \mathbf{y}_i[k] = \mathbf{T}_i \mathbf{H}_i \mathbf{x}[k] + \mathbf{T}_i \mathbf{v}_i[k],$$

such that

$$\mathbf{T}_i \mathbf{H}_i = \begin{pmatrix} \cos(b_i) & \sin(b_i) \\ -\frac{1}{r_i} \sin(b_i) & \frac{1}{r_i} \cos(b_i) \end{pmatrix},$$

is the transformed linear observation matrix which no longer has the explicit dependence on the sensor state it had previously. Further, the transformed measurement noise covariance is now given by $\mathbf{T}_i \mathbf{R}_i[k] \mathbf{T}_i^T$, given by:

$$\begin{pmatrix} \cos(\mathbf{z}_i^\theta) & -\sin(\mathbf{z}_i^\theta) \\ \sin(\mathbf{z}_i^\theta) & \cos(\mathbf{z}_i^\theta) \end{pmatrix} \begin{pmatrix} (\sigma_i^r)^2 & 0 \\ 0 & (\sigma_i^b)^2 \end{pmatrix} \begin{pmatrix} \cos(\mathbf{z}_i^\theta) & \sin(\mathbf{z}_i^\theta) \\ -\sin(\mathbf{z}_i^\theta) & \cos(\mathbf{z}_i^\theta) \end{pmatrix},$$

which now captures the shape of the measurement uncertainty envelope in the global coordinate system. Thus, we see that by controlling the states (specifically, \mathbf{z}_i^θ) of the M sensing agents, we can change the uncertainty associated with the observation process. The objective is to design the control inputs so as to “shrink” the resulting uncertainty envelope as much as possible. The questions of incorporating multiple observations and reducing the uncertainty are addressed in the following sections, after which the examination of this target-tracking case study will be continued.

2.2 Distributed Sensor Fusion

The question of combining measurements from multiple information sources is that of sensor fusion. In essence, the fusion of observations is an information distillation process, by which large quantities of information are brought together to form a reduced but refined representation. The inherent benefits of fusion include the redundancy present in multiple observations of the same state as well as the differences in spatial or temporal resolution that capture a wider perspective of the system behavior. The natural result of these features is the reduction of uncertainty in the overall fused observation.

As discussed in Chapter 1, significant resources have been dedicated over the past half century

to the development of sensor fusion techniques and theory. More specifically, as distributed systems have become more prevalent, methods for distributed or decentralized sensor fusion have spurred much interest, and it is this subset of data fusion research that is most relevant to the distributed sensing objective considered here.

Since the focus of this work is to examine the role of control in distributed sensing tasks, the goal is not to develop new distributed sensor fusion algorithms, but instead to demonstrate how control of sensor states enters into the sensor fusion process. To start, consider the simple but general fusion approach of combining information quantities weighted by the inverse of their uncertainties. Fundamentally, this method represents the simplified theoretical backbone of many fusion algorithms, such as the inverse covariance Kalman filter [49], and many other Kalman-filter variants (e.g., covariance intersection [50] or fusion of Gaussian observations [51]). Henceforth, the formulation will use the “inverse covariance” fusion rules to illuminate this relationship between control of sensor states and fusion of sensor measurements, and rigorous studies of the effect of control in more sophisticated sensor fusion algorithms (e.g., [18]) remain for future investigation.

These “inverse covariance” fusion relations, for combining M independent observations, are given by:

$$\mathbf{P}_{\text{fused}}^{-1} \hat{\mathbf{x}}_{\text{fused}} = \sum_i^M \mathbf{P}_i^{-1} \hat{\mathbf{x}}_i, \quad \mathbf{P}_{\text{fused}}^{-1} = \sum_i^M \mathbf{P}_i^{-1},$$

where $\hat{\mathbf{x}}_i$ and \mathbf{P}_i , are the local estimate vectors and estimate error covariance matrices of the i^{th} sensor, $\mathbf{P}_{\text{fused}} \in \mathbb{R}^{n \times n}$ is the fused estimate error covariance matrix, and $\hat{\mathbf{x}}_{\text{fused}} \in \mathbb{R}^n$ is the fused estimate of the dynamic process state, \mathbf{x} . Note that the time index k is implied. The term “estimate” is defined as the sensing agent’s representation of the state, and any ambiguity in its use in the context of the distributed sensing task will be cleared away shortly. The intuition behind these expressions is that the best fused estimate is a weighted sum of all individual estimates, where each estimate is scaled by the inverse of the uncertainty.

One must note that these relations assume the agents’ estimates are uncorrelated with one another. Given that the M observations are of the same dynamic process(es), this independence simplification is not strictly true [52, 53]. However, this assumption leads to only minimal estimation performance penalties [53] for a broad class of systems, and the reader is referred to the myriad methods in the references cited previously for addressing fusion of correlated estimates. Again, as motivated above, the goal is to illuminate the role of feedback control in the distributed estimation task, and such simplifying assumptions help provide the desired insight.

As mentioned previously, an estimate of the dynamic process state, \mathbf{x} , is any quantity reflecting the sensing agent’s representation of the state. In this manner, we investigate two representations, each with its own merits, of the process state, which can each be used for generating a fused estimate using the fusion rule above. One representation is simply the observation, $\mathbf{y}_i[k]$, itself as per the

measurement model given by Equation 2.1.1. Another method for representing the estimate is to process measurements locally (i.e., by each sensing agent individually) using a Kalman filter, which uses knowledge of the system parameters to generate an optimal estimate given the local observations.

2.2.1 Fusion of Local Observations

In this setting, each sensing agent produces its own observation and shares it, along with the corresponding measurement error covariance matrix, with all other agents. The advantage of this approach is that no additional processing needs to be performed at the sensor node. For implementations that make use of small and/or cheap sensor platforms, where computation capability may be severely limited, this method of fusing measurements is computationally efficient and easily utilized. This fusion strategy also provides immediate insight into the structure of the distributed sensor fusion process, without the need for implementing additional data-processing methods.

The fusion equations for the estimate and its error covariance in the case of shared observations becomes:

$$\mathbf{P}_{\text{fused}}^{-1} \hat{\mathbf{x}}_{\text{fused}} = \sum_i^M \mathbf{P}_i^{-1} \hat{\mathbf{x}}_i = \sum_i^M (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T)^{-1} \mathbf{y}_i,$$

$$\mathbf{P}_{\text{fused}}^{-1} = \sum_i^M \mathbf{P}_i^{-1} = \sum_i^M (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T)^{-1}.$$

This expression shows that multiple observations, transformed into a common global reference frame, can be combined in a straightforward manner. Application of the transformation matrix, \mathbf{T}_i , enables the use of control of the sensor state to influence the global fused estimate of the dynamic process state.

2.2.2 Fusion of Locally Filtered Estimates

Alternatively, sensor observations can be processed locally by each mobile sensor using a Kalman filter to reduce the effect of uncertainty in measurements. The filtered result (i.e., the state estimate and estimate error covariance) is then shared and fused with those of other sensing agents.

This approach addresses some of the challenges faced by the simpler fusion of local observations. Due to the recursive structure of the Kalman filter, the history of measurements and knowledge of the dynamic process model parameters are used to help improve the state estimate. Further, when data sets are large due to data-intensive observations (e.g., high resolution scans, and image/video data), the communication bottleneck may prevent transmission of the raw measurements; instead, the measurements can be processed into a more manageable size in the form of a state estimate, which can then be transmitted to the sensing teammates.

Recall the general Kalman filter equations [54, 55],

$$\begin{aligned}
\hat{\mathbf{x}}[k]^- &= \mathbf{F}\hat{\mathbf{x}}[k-1], \\
\mathbf{P}[k]^- &= \mathbf{F}\mathbf{P}[k-1]\mathbf{F}^T + \mathbf{Q}, \\
\mathbf{K}[k] &= \mathbf{P}[k]^- \mathbf{H}^T (\mathbf{H}\mathbf{P}[k]^- \mathbf{H}^T + \mathbf{R})^{-1}, \\
\hat{\mathbf{x}}[k] &= \hat{\mathbf{x}}[k]^- + \mathbf{K}[k] (\mathbf{y}[k] - \mathbf{H}\hat{\mathbf{x}}[k]^-), \\
\mathbf{P}[k] &= (\mathbf{I} - \mathbf{K}[k]\mathbf{H})\mathbf{P}[k]^-,
\end{aligned}$$

where \mathbf{K} is the Kalman estimator gain.

Thus, the i^{th} sensor generates its local estimate, $\hat{\mathbf{x}}_i$, and estimate error covariance, \mathbf{P}_i according to the following equations (after application of the transformation matrix \mathbf{T}_i):

$$\begin{aligned}
\hat{\mathbf{x}}_i[k]^- &= \mathbf{F}\hat{\mathbf{x}}_i[k-1], \\
\mathbf{P}_i[k]^- &= \mathbf{F}\mathbf{P}_i[k-1]\mathbf{F}^T + \mathbf{Q}, \\
\mathbf{K}_i[k] &= \mathbf{P}_i[k]^- \mathbf{H}_i^T \mathbf{T}_i^T (\mathbf{T}_i \mathbf{H}_i \mathbf{P}_i[k]^- \mathbf{H}_i^T \mathbf{T}_i^T + \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T)^{-1}, \\
\hat{\mathbf{x}}_i[k] &= \hat{\mathbf{x}}_i[k]^- + \mathbf{K}_i[k] (\mathbf{T}_i \mathbf{y}_i[k] - \mathbf{T}_i \mathbf{H}_i \hat{\mathbf{x}}_i[k]^-), \\
\mathbf{P}_i[k] &= (\mathbf{I} - \mathbf{K}_i[k] \mathbf{T}_i \mathbf{H}_i) \mathbf{P}_i[k]^- = \mathbf{\Delta}_i - \mathbf{K}_i \mathbf{T}_i \mathbf{H}_i \mathbf{\Delta}_i,
\end{aligned}$$

where $\mathbf{\Delta}_i \triangleq \mathbf{P}_i[k]^- = \mathbf{F}\mathbf{P}_i[k-1]\mathbf{F}^T + \mathbf{Q}$, and exchanges these quantities with the other sensing agents. Note that while $\mathbf{\Delta}$ is independent of current sensor states, the Kalman gain \mathbf{K}_i , through the dependence on the sensor state of the transformed measurement noise covariance matrix, $\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T$, can be affected by appropriate control inputs.

In this manner, both methods – fusion of local observations and of locally filtered estimates – demonstrate how the global (i.e., fused) sensing process can be influenced by local sensor state controls.

It merits mentioning that by using a decentralized sensor fusion algorithm, each sensor need not receive all data from all other sensors in order to compute a fused estimate and uncertainty. Clearly, the result of this calculation using observations from only a subset of the team will not be the true global fused values; however, this framework allows for general networks of sensors with varied communication topologies, not necessarily fully connected as often required for optimal centralized solutions. Further, this fact also means that the calculation of the control laws discussed in the following sections is robust to the addition or deletion of sensor nodes. This robustness is particularly relevant and useful in practical implementations of the methods presented in this chapter.

2.2.3 Case Study: Distributed Target Tracking

The components of the distributed sensing discussed above can be further illuminated by revisiting the illustrative example of tracking a target in the plane using multiple range-and-bearing sensors. Employing the sensor fusion algorithms described above, it is possible to attain a fused or aggregate observation of the target position that possesses lower overall uncertainty than each of the range and bearing measurements individually. Figure 2.3 depicts a simulated example of the benefit of fusing two independent local observations. Panels (a) and (b) show multiple observations taken individually by sensor 1 and 2, respectively. Panel (c) illustrates the result of employing the sensor fusion method described in Section 2.2.1, resulting in a smaller uncertainty envelope for the fused measurements.




Figure 2.3: Illustration of the benefit of fusing multiple observations. Two (stationary) sensing agents make independent, noisy range and bearing observations of a planar target, which are fused to yield estimates with reduced uncertainty.

As can be seen, the resulting uncertainty of the fused observations, represented by the uncertainty ellipse, is smaller than that of the individual observations. This reduction of uncertainty in this example highlights the desire to use distributed agents for gathering information.

2.2.4 Construction of the Distributed Sensing Cost Function

For purposes of control, a measure of performance of the ensemble of sensing agents is required. This measure will enable the design of control laws that drive the individual sensor agents to states which improve overall sensing performance.

The estimate error covariance generated by the sensor fusion process provides a representation of the uncertainty present in the estimate. Said another way, $\mathbf{P}_{\text{fused}}$ captures the overall “shape” of the fused estimate uncertainty. Various metrics can be defined which use a function of this error covariance matrix, and a particular metric can be selected according to the optimization task.

Definition of the uncertainty size metric defines a cost function, J , which can be minimized via

control. The remainder of this chapter will explore the use of the following cost function:

$$J = \det \mathbf{P}_{\text{fused}} = \det \left(\sum_i^M \mathbf{P}_i^{-1} \right)^{-1}.$$

The determinant of the estimate error covariance matrix represents the volume of uncertainty, and the proposed optimization seeks to reduce this volume. Known as D -optimal design [4], the determinant is chosen (instead of the trace, maximum eigenvalue, etc.) to facilitate the development of the theoretical framework in Section 2.3.

Define two cost functions, J_1 and J_2 , for the two approaches outlined in this section; namely, fusion of local observations and fusion of locally processed estimates, respectively. By simple substitution of the expressions for the fused estimate error covariance computations, the cost functions of interest are given by:

$$J_1 = \det \mathbf{P}_{\text{fused}} = \det \left(\sum_i^M (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T)^{-1} \right)^{-1}, \quad (2.2.1)$$

$$J_2 = \det \mathbf{P}_{\text{fused}} = \det \left(\sum_i^M (\Delta_i - \mathbf{K}_i \mathbf{T}_i \mathbf{H}_i \Delta_i)^{-1} \right)^{-1}. \quad (2.2.2)$$

By design, J_1 and J_2 are functions of the sensor states, carried intrinsically in the transformed measurement noise covariances. Thus, by varying the states of the sensors, the estimate error covariance can be influenced in a manner which reduces the cost. The question posed here is how to do so in a decentralized way, given a team of M sensing agents.

2.3 Gradient Analysis

Given an expression for the cost function, J , a method for determining the minimal cost configuration for the states of the sensing agents is desired. In general, determination of the optimal trajectories for the evolution of the sensor states is computationally expensive, and becomes prohibitive as the number of the dimensions of the process states increases.

In contrast to the computational cost of calculating optimal sensor state evolutions, gradient-descent methods for finding minimums have the feature that only local information is used in directing the optimization. While one can only be guaranteed to find a local minimum, the computational benefits alone merit investigation into gradient-based approaches. An additional advantage of the gradient is its decentralizing effect on functions that possess summations by distributing the differentiation [56], such as the cost functions J_1 (Equation 2.2.1) and J_2 (Equation 2.2.2). In other words, the resultant expression for the gradient of such functions depends only on the states involved in the differentiation (while holding other states constant) and terms containing only other states

vanish in the calculation. Consequently, the gradient approach results in closed-form analytic expressions which make it an appealing choice for insight into the distributed state estimation problem in general, as well as for implementation in practical physical systems.

The computation of the gradients of J_1 and J_2 employs the following standard matrix calculus identities [57]:

$$\frac{\partial}{\partial z} h(\mathbf{A}(z)) = \text{tr} \left[\frac{\partial h}{\partial \mathbf{A}} \frac{\partial \mathbf{A}}{\partial z} \right], \quad (\text{matrix chain rule}), \quad (2.3.1)$$

$$\frac{\partial}{\partial \mathbf{A}} \det(\mathbf{A}) = |\mathbf{A}| \mathbf{A}^{-T} = |\mathbf{A}| \mathbf{A}^{-1}, \quad (\text{derivative of the determinant}), \quad (2.3.2)$$

$$\frac{\partial}{\partial z} \mathbf{A}^{-1} = -\mathbf{A}^{-1} \left(\frac{\partial \mathbf{A}}{\partial z} \right) \mathbf{A}^{-1}, \quad (\text{derivative of the inverse}), \quad (2.3.3)$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric, positive-definite matrix, $h : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ is a real-valued matrix function, and $z \in \mathbb{R}$ is a scalar variable.

The derivatives of the cost function with respect to each component of the i^{th} sensor's state, $\mathbf{z}_i = (z_i^1, \dots, z_i^l, \dots, z_i^s)^T$, which help to determine the locally cost-minimizing path for sensor i , are given in the following proposition.

Proposition 1 *The gradient of J_1 and J_2 with respect to the l^{th} coordinate of the i^{th} sensor state, given by z_i^l , has the general form:*

$$\frac{\partial J}{\partial z_i^l} = |\mathbf{P}_{fused}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused} \right],$$

where

$$\mathbf{\Pi} = \begin{cases} \mathbf{T}_i \mathbf{R}_i^{-1} \mathbf{T}_i^T, & \text{for } J_1, \\ \mathbf{P}_i^{-1} \mathbf{K}_i, & \text{for } J_2. \end{cases}$$

Proof: Proof is realized by application of Equations 2.3.1-2.3.3 to either Equation 2.2.1 or Equation 2.2.2, and the complete derivation can be found in Appendix A. ■

2.3.1 Gradient-Descent Control Law

The gradient formulas can be used to define a control law, $\mathbf{u}_i(\mathbf{z}_i)$, for the control of the i^{th} sensor's state. This control law defines the inputs to the sensor state evolution expression, which has already been shown to influence the overall sensing performance of the coordinated team. The direction to achieve the steepest slope of the cost function for the i^{th} sensor is given by the gradient:

$$\nabla_{\mathbf{z}_i} J(\mathbf{z}_1, \dots, \mathbf{z}_M, \mathbf{x}) = \frac{\partial J}{\partial z_i^1} e_{z_i^1} + \dots + \frac{\partial J}{\partial z_i^s} e_{z_i^s} = \sum_l^s \left(\frac{\partial J}{\partial z_i^l} e_{z_i^l} \right),$$

where $e_{z_i^l}$ is the unit vector associated with coordinate z_i^l .

Construction of the gradient-based control law can simply be done by taking the components of the gradient expression above (negative for steepest descent or minimization of the cost):

$$\mathbf{u}_i(\mathbf{z}_i) = \left(-\frac{\partial J}{\partial z_i^1} \Big|_{\mathbf{z}_i} \quad \cdots \quad -\frac{\partial J}{\partial z_i^s} \Big|_{\mathbf{z}_i} \right)^T,$$

evaluated at the current state of the i^{th} sensor. Due to the decentralizing effect of the gradient on the cost function, $\mathbf{u}_i(\mathbf{z}_i)$ is an explicit function of *only* the i^{th} sensor's state. In other words, as can be seen in the derivation, all other terms corresponding to other sensing agents vanish when computing the derivative.

As a result, the i^{th} sensor's control law can be computed from just the sensor's own current state, its sensor model and the fused estimate error covariance. This simple result occurs because all information from other sensing agents relevant for computation of the control law is already captured implicitly in the fused estimate error covariance, $\mathbf{P}_{\text{fused}}$. In this manner, the optimization for motion control is conducted in a completely decentralized fashion. Furthermore, by construction, the control law is the instantaneously (i.e., locally) optimal action to reduce the uncertainty in the system.

The above statement highlights an important feature of this formulation, in that each sensor can operate independently of all other sensors. In other words, knowledge of the dynamics, measurement and uncertainty models of other sensors is not required in order to compute the local control law. Importantly, information regarding the dynamic states of all other sensors is also unnecessary, which reduces the amount of data required to be transmitted at every time step.

2.3.2 Gradient Analysis for Tracking Multiple Dynamic Processes

In addressing the tracking of multiple dynamic processes, a straightforward extension to the computation of the gradient exists [47]. Recall the augmentation approach used for the dynamics, observation, and corresponding transformation matrices. Due to this structure, the tracking of multiple processes appears in the gradient computation simply as derivatives over additional independent coordinates, while maintaining the same form of the closed-form expression.

Proposition 2 *The gradient of J_1 and J_2 with respect to the i^{th} sensor's coordinate, z_i^l , relative to the j^{th} dynamic process has the general form given by:*

$$\frac{\partial J}{\partial z_{i,j}^l} = |\mathbf{P}_{\text{fused}}| \operatorname{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_{i,j} \mathbf{R}_{i,j} \mathbf{T}_{i,j}^T) \mathbf{\Pi}^T \mathbf{P}_{\text{fused}} \right],$$

where

$$\mathbf{\Pi} = \begin{cases} \mathbf{T}_{i,j} \mathbf{R}_{i,j}^{-1} \mathbf{T}_{i,j}^T, & \text{for } J_1, \\ \mathbf{P}_{i,j}^{-1} \mathbf{K}_{i,j}, & \text{for } J_2. \end{cases}$$

Proof: This result follows immediately from the calculation for Proposition 1, using the augmented process, measurement noise covariance and transformation matrices defined for tracking multiple processes. The reader is referred again to Appendix A for the complete derivation. ■

The computation of the gradient with respect to process j is independent of those for all other processes, following from the independence of these N processes. The cost function now incorporates the additional variables of each process, thereby resulting in the following gradient expression:

$$\nabla_{\mathbf{z}_i} J(\mathbf{z}_1, \dots, \mathbf{z}_M, \mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_l^s \left(\sum_j^N \frac{\partial J}{\partial z_{i,j}^l} \right) e_{z_i^l}.$$

Note that the number of derivative computation scales with the number of dynamic processes the system needs to track. In general, the tracking problem becomes computationally challenging as more processes are tracked; however, the analytic expressions for the gradient enable significantly more efficient computations. Consequently, the closed-form formulas allow for examination of systems with large numbers of processes as well as sensing agents.

The modified control law for the i^{th} sensor observing N target processes is given by:

$$\mathbf{u}_i(\mathbf{z}_i) = \sum_j^N \mathbf{u}_{i,j}(\mathbf{z}_i),$$

where the vector sum is over the control signals for the j^{th} process is:

$$\mathbf{u}_{i,j}(\mathbf{z}_i) = \left(-\frac{\partial J}{\partial z_{i,j}^1} \Big|_{\mathbf{z}_i} \quad \cdots \quad -\frac{\partial J}{\partial z_{i,j}^s} \Big|_{\mathbf{z}_i} \right)^T.$$

The gradient-based control law for tracking multiple processes is simply the sum of the control laws for each process. With these expressions, multiple sensing agents are able to observe and estimate multiple dynamic processes, maneuvering individually in a locally optimal fashion using the gradient descent control laws presented here. Additionally, this summation of control laws implies that rather than measuring all N processes, it is possible to observe only a subset of these states and generate a control law that is only relevant for this subset. This feature demonstrates the flexibility of this framework for examining the appropriate distribution of sensing resources (i.e., subteams of sensors) for different processes. While this problem of resource allocation brings up the challenge of data association and task assignment (already indicated to be active research areas for multi-agent systems), the distributed framework presented herein may serve as a foundation from which

to conduct further investigations.

Stability of the distributed gradient feedback control law about the local minimum (i.e., equilibrium point) can be examined using tools from Lyapunov stability [25, 58], graph theory [59], or artificial potential function approaches [60, 26]. Explicit calculation of the multiple sensor states resulting in local minima of the cost function is difficult in general. Investigation of these minima for the simplified case of planar observations using two and three mobile sensors is conducted in Appendix B.

2.3.3 Case Study: Distributed Target Tracking, Revisited

Recall the example of tracking a dynamic target vehicle using a team of M mobile sensors. Assume that these mobile platforms can move omnidirectionally in the plane (i.e., agents are fully actuated in \mathbb{R}^2) and communicate completely with all team members. Further, given the observation model presented previously, consider as an illustrative example a measurement uncertainty model where the measurement noise in range is quadratic in the distance, $r_{i,j}$, from the i^{th} sensor to the j^{th} target. In other words,

$$f_i^r(\mathbf{x}, \mathbf{z}_i) = f_i^r(r_{i,j}) = a_2 (r_{i,j} - a_1)^2 + a_0,$$

where a_0, a_1, a_2 are constant coefficients. This model corresponds to the notion of a “sweet spot” in sensing, located at a distance a_1 from the target, where uncertainty in measurements is minimal [46]. Furthermore, let the measurement noise in bearing simply be a fixed multiple α of the range noise variance, such that $f_i^b(r_{i,j}) = \alpha f_i^r(r_{i,j})$.

In this case of range and bearing observations, the derivative term in the gradient-based control law can be computed explicitly using the uncertainty model given above. In terms of the i^{th} sensor’s range and bearing (i.e., polar) coordinates, z_i^r and z_i^b , respectively, with respect to a single target, differentiation of the transformed covariance matrices yields:

$$\frac{\partial}{\partial z_i^r} \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T = \mathbf{T}_i \begin{pmatrix} \frac{\partial f^r}{\partial z_i^r} & 0 \\ 0 & \frac{\partial f^b}{\partial z_i^r} \end{pmatrix} \mathbf{T}_i^T,$$

where, for this illustrative example using the given quadratic range-dependent uncertainty model, the derivatives are:

$$\frac{\partial f^r}{\partial z_i^r} = 2a_2(r_i - a_1), \quad \frac{\partial f^b}{\partial z_i^r} = 2\alpha a_2(r_i - a_1),$$

and further

$$\begin{aligned} \frac{\partial}{\partial z_i^b} \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T &= \frac{\partial \mathbf{T}_i}{\partial z_i^b} \mathbf{R}_i \mathbf{T}_i^T + \mathbf{T}_i \mathbf{R}_i \frac{\partial \mathbf{T}_i^T}{\partial z_i^b} \\ &= \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T + \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}^T. \end{aligned}$$

Since the gradient is computed in polar coordinates (i.e., z_i^r, z_i^b) centered about the target, the direction of steepest slope is given by

$$\nabla_{z_i^r, z_i^b} J(z_1^r, z_1^b, \dots, z_M^r, z_M^b) = \frac{\partial J}{\partial z_i^r} e_{z_i^r} + \frac{1}{r_i} \frac{\partial J}{\partial z_i^b} e_{z_i^b}.$$

Hence, the gradient-based control law for the i^{th} sensor is given in these polar coordinates by

$$\mathbf{u}_i(z_i^r, z_i^b) = \left(-\frac{\partial J}{\partial z_i^r} \quad -\frac{1}{r_i} \frac{\partial J}{\partial z_i^b} \right)^T.$$

All that remains is to convert the control signal from polar to Cartesian coordinates, done easily by a simple rotation:

$$\mathbf{u}_i(z_i^r, z_i^b) = \mathbf{T}_i \mathbf{u}_i(z_i^x, z_i^y) = \mathbf{T}_i \mathbf{u}_i(\mathbf{z}_i) \quad \Rightarrow \quad \mathbf{u}_i(\mathbf{z}_i) = \mathbf{T}_i^T \mathbf{u}_i(z_i^r, z_i^b).$$

The resulting configuration, for the case of a team of homogeneous sensors, (i.e., $f_i^r(\cdot) = f^r(\cdot), \forall i$) with uncertainty model parameters $a_0 = 0.1528, a_1 = 15.625, a_2 = 0.0008$ and bearing scaling parameter, $\alpha = 5$, is depicted in Figure 2.4, where the motion of the agents are governed by cost functions J_1 and J_2 (shown in (a) and (b), respectively). Identical initial positions of the sensors of arbitrary placement were used in the two simulations, with the target starting at the origin and randomly drifting along the positive x -axis.




Figure 2.4: Three sensors track a (biased) randomly walking target using the proposed gradient-descent control law for costs (a) J_1 , (b) J_2 .

The sensors converge to an intuitively symmetric configuration, where the three similar sensors are uniformly distributed about the target at their optimal sensing distance. In the case of $M = 3$ agents, this configuration corresponds to the globally optimal sensor placement, with separation between agents of 120° (or equivalently, due to the symmetry of the uncertainty, 60°) on a circle of radius a_1 centered about the target. (See Appendix B.)




Figure 2.5: Evolution of the cost for an illustrative simulation run. Comparison of the algorithms shows that mobile sensors outperform stationary ones (i.e., have lower overall sensing uncertainty). Selection of the fusion approach (either sharing local observations or locally filtered estimates) is a design choice between performance and computation constraints.

A comparison of the performance under the two cost functions for this illustrative simulation run is shown in Figure 2.5, and is contrasted to the scenario where sensors fixed at their initial positions track the same moving target. It can be seen that mobility aids in reducing the uncertainty in the state estimate on the average, independent of initial conditions. Further, as expected, the local processing of measurements provides significant improvement over simply sharing observations, due to the filter’s ability to incorporate previous measurements and predict target motion. If onboard computation is limited, however, motion control based on simple observation fusion may serve as a less computationally intensive alternative.

2.4 Extension to Imperfect Communication

One of the challenges faced by distributed systems is the need to have a communication infrastructure in place to enable the sharing of gathered information. This communication overhead is precisely the subject of much scrutiny in distributed and networked control systems research, which includes,

for example, analysis of estimation processes when scheduling of observations is necessary due to communication constraints.

The distributed framework presented thus far has assumed that complete communication is possible without any degradation in the transmitted information. In practice, this assumption does not always hold, especially in dynamic environments. Moreover, many distributed systems connect information gathering sources via wireless networks, such as the case of teams of mobile sensing platforms where wired communication links are impractical. In this manner, the objective of improving the quality of sensed information merits the examination of the distributed sensing task in the presence of imperfect wireless communication between agents, which is the subject of this section.

As seen previously, the state of the sensor influences the quality of the observation. Similarly, the communication quality can also be modeled to be affected by the sensor state. However, instead of depending on the relationship between the sensor and the target process, the uncertainty in communication is tied to the differences in state between sensing agents. As will be seen in Section 2.4.2, these differing dependencies will yield a situation which requires a trade-off between control for improved sensing and control for improved communication.

The homogeneous sensor observation model can be easily modified to include the effects of noisy communication links, such as fading wireless channels [22]. Sensor i' broadcasts its transformed observation, $\mathbf{T}_{i'}\mathbf{y}_{i'}$, which is received by the i^{th} sensor according to the following process:

$$\hat{\mathbf{y}}_{i,i'}[k] = \mathbf{T}_{i'}\mathbf{y}_{i'}[k] + \mathbf{v}_{i,i'}^c[k] = \mathbf{T}_{i'}\mathbf{H}_{i'}\mathbf{x}[k] + \mathbf{T}_{i'}\mathbf{v}_{i'}[k] + \mathbf{v}_{i,i'}^c[k],$$

where $\mathbf{v}_{i,i'}^c[k] \in \mathbb{R}^m$ is the communication noise in the received observation. In other words, as in the perfect communication case, the i^{th} sensor receives estimates from some or all of its teammates, but in contrast, these estimates are perturbed by additive communication noise. Assume that the communication noise across links is symmetric, i.e., $\mathbf{v}_{i,i'}^c = \mathbf{v}_{i',i}^c$, and that there is no noise for self-transmissions (i.e., $\mathbf{v}_{i,i}^c = 0$).

Denote the covariance matrix of the communication noise vector, $\mathbf{v}_{i,i'}^c$, as $\mathbf{C}_{i,i'} \in \mathbb{R}^{m \times m}$. This communication noise is often dependent on a variety of issues such as transmission power and environmental effects [61, 22], and as such, is modeled to have a spatial dependence of the general form $g^c(\mathbf{x}_i, \mathbf{x}_{i'})$, which explicitly highlights the possible dependence on the states of the receiving and transmitting agents. Given this model, the communication noise covariance matrix is:

$$\mathbf{C}_{i,i'} = \begin{pmatrix} g^c(\mathbf{z}_i, \mathbf{z}_{i'}) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & g^c(\mathbf{z}_i, \mathbf{z}_{i'}) \end{pmatrix} = \begin{pmatrix} (\sigma_{i,i'}^c)^2 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & (\sigma_{i,i'}^c)^2 \end{pmatrix},$$

where the uncertainty in the received observation due to communication noise is denoted $(\sigma_{i,i'}^c)^2$.

Given the measurement and communication channel models, the modified fusion equation for the covariance of the estimate error is [22]:

$$\mathbf{P}_{i,\text{fused}}^{-1} = \sum_{i'}^M (\mathbf{P}_{i'} + \mathbf{C}_{i,i'})^{-1},$$

where $\mathbf{P}_{i,\text{fused}}$ is the i^{th} sensor's fused estimate error covariance matrix. Note that in the previous perfect communication case (and assuming complete communication between all agents), the sensor fusion process results in identical fused estimates and covariance matrices for all agents, due to the fact that each agent is essentially acting as the central fusion node. However, in the current scenario, since the information collected by each sensor is no longer identical to that of its teammates because of the communication noise, $\mathbf{P}_{i,\text{fused}}$ will be different for each sensor, thereby requiring the additional identifying subscript i .

The cost function $J_{3,i}$ to be minimized is once again the determinant of the fused uncertainty matrix, given as:

$$J_{3,i} = \det \mathbf{P}_{i,\text{fused}} = \det \left(\sum_{i'}^M (\mathbf{P}_{i'} + \mathbf{C}_{i,i'})^{-1} \right)^{-1}. \quad (2.4.1)$$

Similar to the cost functions J_1 and J_2 , this cost function is also dependent on the sensor states, following along the lines of the derivations of the previous section to yield the following result:

Proposition 3 *The gradient of $J_{3,i}$ with respect to the i^{th} sensor's coordinate, z_i^l , is given by:*

$$\begin{aligned} \frac{\partial J_{3,i}}{\partial z_i^l} &= |\mathbf{P}_{i,\text{fused}}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{i,\text{fused}} \right] \\ &\quad + \sum_{i' \neq i}^M |\mathbf{P}_{i,\text{fused}}| \text{tr} \left[\mathbf{\Phi}_{i,i'} \frac{\partial \mathbf{C}_{i,i'}}{\partial z_i^l} \mathbf{\Phi}_{i,i'}^T \mathbf{P}_{i,\text{fused}} \right], \end{aligned}$$

where $\mathbf{\Phi}_{i,i'} \triangleq (\mathbf{P}_i + \mathbf{C}_{i,i'})$, $\mathbf{\Pi}$ is chosen according to whether simple observations (i.e., $\mathbf{P}_i = \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T$) or locally filtered measurements (i.e., $\mathbf{P}_i = \mathbf{\Delta}_i - \mathbf{K}_i \mathbf{T}_i \mathbf{H}_i \mathbf{\Delta}_i$) are fused, and the derivative of the communication noise covariance matrix is given by:

$$\frac{\partial \mathbf{C}_{i,i'}}{\partial z_i^l} = \begin{pmatrix} \frac{\partial}{\partial z_i^l} g_{i,i'}^c(\mathbf{z}_i, \mathbf{z}_{i'}) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{\partial}{\partial z_i^l} g_{i,i'}^c(\mathbf{z}_i, \mathbf{z}_{i'}) \end{pmatrix}.$$

Proof: The matrix calculus identities are employed once again, as done previously, noting that the communication noise term vanishes when considering self-transmission, and persists otherwise.

The reader is referred again to Appendix A for the detailed algebraic calculation. ■

Given this expression for the gradient, the gradient-based control law can easily be constructed for the case of imperfect communication. Taking the partial derivative components, construct the control vector as follows:

$$\mathbf{u}_i(\mathbf{z}_i) = \left(-\frac{\partial J_{3,i}}{\partial z_i^1} \Big|_{\mathbf{z}_i} \quad \cdots \quad -\frac{\partial J_{3,i}}{\partial z_i^s} \Big|_{\mathbf{z}_i} \right)^T,$$

evaluated at the current state of the i^{th} sensor.

The first term in the derivative term is exactly the expression derived for the control law for the perfect communication case. The second term represents the effect of the communication noise from each of the i^{th} sensor's teammates. The component of the control law in the l^{th} coordinate (i.e., the derivative computed in Proposition 3) comprises these two terms:

$$u_i^l(\mathbf{z}_i) \triangleq -\frac{\partial J_{3,i}}{\partial z_i^l} \Big|_{\mathbf{z}_i} = \underbrace{|\mathbf{P}_{i,\text{fused}}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{i,\text{fused}} \right]}_{(u_i^l)^{\text{sensing}}} + \underbrace{\sum_{i' \neq i}^M |\mathbf{P}_{i,\text{fused}}| \text{tr} \left[\mathbf{\Phi}_{i,i'} \frac{\partial \mathbf{C}_{i,i'}}{\partial z_i^l} \mathbf{\Phi}_{i,i'}^T \mathbf{P}_{i,\text{fused}} \right]}_{(u_i^l)^{\text{comm}}},$$

such that the control vector can be decomposed easily into

$$\mathbf{u}_i(\mathbf{z}_i) = (\mathbf{u}_i)^{\text{sensing}} + (\mathbf{u}_i)^{\text{comm}}.$$

The sensors' control inputs are clearly governed by two influences, namely that of applying control to improve sensing quality combined with applying control to improve communication performance. Note that these effects, as they appear in the context of the control signal, are decoupled. This point is examined more closely in the Section 2.4.2.

2.4.1 Case Study: Distributed Target Tracking with Imperfect Communication

In the previous construction of the distributed target tracking example, perfect transmission of observations was assumed. This section examines the behavior of the system in the presence of imperfect communication.

Many different models exist for the internodal communication noise function, $g^c(\cdot)$, which may include the effects of transmission distance and power, frequency, quantization sizes, and number of bits per transmission over the channel (see [22] and references therein). As an illustrative example,

choose a noise model where the signal-to-noise ratio (SNR) obeys a general inverse-square law [61]:

$$g^c(\mathbf{x}_i, \mathbf{x}_j) = (\sigma_{i,j}^c)^2 = \frac{1}{SNR}, \quad \text{where } SNR = \frac{\mu}{d_{i,j}^2},$$

for $\mu > 0$ a constant communication noise parameter and $d_{i,j} \triangleq \|\mathbf{x}_i - \mathbf{x}_j\|_2$ is the distance between sensor i and sensor j .

Applying this communication noise model, the derivative component in the control signal is dependent on the distance, $d_{i,i'}$, given by

$$\frac{\partial \mathbf{C}_{i,i'}}{\partial z_i^l} = \begin{pmatrix} \frac{\partial}{\partial z_i^l} g_{i,i'}^c(d_{i,i'}) & 0 \\ 0 & \frac{\partial}{\partial z_i^l} g_{i,i'}^c(d_{i,i'}) \end{pmatrix},$$

with

$$\frac{\partial}{\partial z_i^l} g_{i,i'}^c(d_{i,i'}) = \frac{2}{\mu} \frac{\partial d_{i,i'}}{\partial z_i^l} = \frac{2}{\mu} (z_i^l - z_{i'}^l).$$

The performance of the modified gradient-based control law for the task of target tracking with two mobile sensors is depicted in Figure 2.6, with communication parameter $\mu = 1000$ and measurement uncertainty model given in the perfect communication example (see Section 2.3.3). The sensors start at arbitrary initial conditions, some distance away from the randomly walking target. The resulting configuration of the sensors under communication constraints is shown overlaying the outcome under the same initial conditions without communication noise.




Figure 2.6: Two sensors track a randomly walking target: With imperfect communication (dark), or with perfect communication (light).

Clearly, the communication constraint biases the mobile agents to remain closer to one another

than in the case where perfect inter-agent communication is assumed. This behavior results in a compromise of sensing performance. Note that, given the above sensing and communication models, the sensing agents maintain their relative bearing at $\pm\frac{\pi}{2}$ radians for optimal sensing. Different communication models may yield configurations where the agents maintain their optimal sensing distances and instead reduce their relative bearing to the target [22].

2.4.2 Relationship between Sensing and Communication

One can next investigate how performance, measured by the cost function Equation 2.4.1, is related to communication and sensing parameters. Varying the parameter μ (which is an aggregate measure of communication characteristics) results in Figure 2.7 for the case where multiple sensors observe a single target. As μ becomes small (communication noise becomes large) no useful information is exchanged amongst the sensing agents. Thus, each agent tends to act independently, and none of the benefits of cooperative sensing is realized. For high values of μ (the communication link quality is very good), overall cost is reduced by the use of an increasing number of sensors. Interesting behavior appears to occur in the intermediate range, where there is a trade-off between the sensing and communication objectives.




Figure 2.7: Sensing and communication performance tradeoff as a function of communication parameter, μ .

Further insight into this trade-off can be found by varying the optimal sensing distance (a.k.a. “sweet spot”), a_1 . As shown in Figure 2.8, reducing the optimal range to a target has the effect of improving the overall system performance – the smaller this distance is, the closer the cooperating sensors can operate while maintaining a wide-enough perspective. Conversely, while large separations between sensors provide sufficiently different views of the target, communication noise increases

commensurately. Clearly, the choice or design of sensors plays a role in the overall performance of the mobile sensing network.




Figure 2.8: Sensing and communication performance tradeoff as a function of sensing sweet spot, a_1

2.5 Extension to Classification

Estimation, as presented in the previous sections, is at the level of processing of sensory data. As interest in endowing sensing agents with more intelligence increases, a logical extension of the state estimation problem is that of classification of processes.

A wide class of situations require that collected data be used to identify the category of the observed process, including applications of target identification and shape classification. As done in the context of improving the state estimate by use of mobility, an analogous examination of utilizing motion control to improve the classification performance can be conducted.

Hence, the problem statement for the classification task presented in this section can be captured as follows. *Given measurements of the state of a target, how should the sensors move in order to improve their overall ability to correctly classify a single target?* To address this question, first a cost function is developed which evaluates the quality of the classification. The intuitive choice is the probability of error in classification, denoted p_e , which describes the likelihood that the categorization made from measurements is not the true class of the target.

2.5.1 Classification Probability

Formulation of the expression for p_e is done in a general manner, but emphasis will be on the binary hypothesis problem, where there are two classes (e.g., heads or tails, 0 or 1, friend or foe). Much literature already exists for this class of problems (e.g., see [62]), and further, multiple alternative problems are often treated as extensions of the two-class scenario [62, 63].

Let \mathcal{C} denote the set of N classes, i.e., $\mathcal{C} = \{c_1, \dots, c_N\}$. For instance, in the binary hypothesis scenario, $N = 2$ represents examination of two classes such that c_1 and c_2 might correspond to, for example, “friend” and “foe,” respectively.

Following [63], the different hypotheses are represented as:

$$c_j : p(y|x_j) = \mathcal{N}(x_j, \Sigma_j). \quad (2.5.1)$$

Here, y represents the single resulting fused estimate arising from the sensor fusion process, such as the inverse covariance Kalman filter described in the previous section. Describing Equation 2.5.1 in words, if the target is a member of class c_j , the measurements are distributed normally with covariance Σ_j and centered about the true state x_j . In addition, x_j itself is also a random variable whose distribution may depend on the class c_j . Consider in particular the case where the probability distributions of members of a class are also random variables:

$$p(x_j) \equiv p(x|c_j) = \mathcal{N}(\bar{x}_j, \Lambda_j), \quad (2.5.2)$$

where \bar{x}_j and Λ_j represent the mean and covariance of the distribution of members in class c_j .

Classification decision rules are assumed to be known in advance. The rules defining the boundaries between classes can be learned, or determined from first principles. In either case, denote the region in measurement space where c_j is the correct class by $\Omega_{y,j}$ and its complement by $\bar{\Omega}_{y,j}$.

With the above definitions in mind, following [62] the probability of error in classification can be expressed as:

$$\begin{aligned} p_e &= \sum_{j=1}^N p(\text{error}|c_j)p(c_j) \\ &= \sum_{j=1}^N p(c_j) \left(1 - \int_{\Omega_{y,j}} p(y|c_j) dy \right) \\ &= 1 - \sum_{j=1}^N \int_{\Omega_{y,j}} p(y|c_j)p(c_j) dy, \end{aligned}$$

where $p(c_j)$ represents the *a priori* probability of the target being in class c_j , and $\sum_{j=1}^N p(c_j) = 1$.

Note that the probability of correct classification, p_c , is related to p_e by $p_e = 1 - p_c$. In the case

of many classes, it is easier to investigate the probability of correct classification. When the *a priori* class distributions are independent of the measurements, the utility function to be maximized is:

$$p_c = \sum_{j=1}^N p(c_j) \int_{\Omega_{y,j}} p(y|c_j) dy. \quad (2.5.3)$$

The goal is to find an expression $p(y|c_j)$, which relates the probability distribution of the measurements with a particular class. By definition of marginal probabilities,

$$p(y|c_j) = \int p(y, x|c_j) dx = \int p(y|x, c_j)p(x|c_j) dx, \quad (2.5.4)$$

where the second equation comes from simple application of Bayes' rule.

For a particular class c_j , Equation 2.5.4 can be rewritten to incorporate the relationship between target state x and the j^{th} class c_j :

$$p(y|c_j) = \int p(y|x_j)p(x_j) dx_j. \quad (2.5.5)$$

where x_j is introduced as the nuisance parameter [63]. Thus Equations 2.5.3 and 2.5.5, which depend implicitly upon the sensors' states, define a utility function whose maximization over the set of possible sensor motions leads to the best classification performance of a team of cooperating mobile sensors.

2.5.1.1 Gaussian noise and class distributions

In the particular case where the sensing noise is Gaussian and the class distributions are Gaussian, the following expressions can be obtained from Equations 2.5.1 and 2.5.2, respectively:

$$p(y|x_j) = \frac{1}{(2\pi)^{\frac{n}{2}} |\Sigma_j|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(y - x_j)^T \Sigma_j^{-1}(y - x_j)\right),$$

$$p(x_j) = \frac{1}{(2\pi)^{\frac{n}{2}} |\Lambda_j|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x_j - \bar{x}_j)^T \Lambda_j^{-1}(x_j - \bar{x}_j)\right).$$

such that Equation 2.5.4 becomes:

$$\begin{aligned} p(y|c_j) &= \int \mathcal{N}(x_j, \Sigma_j) \mathcal{N}(\bar{x}_j, \Lambda_j) dx_j \\ &= \mathcal{N}_y(\bar{x}_j, (\Sigma_j + \Lambda_j)), \end{aligned}$$

utilizing the fact that the integral of the product of Gaussians is itself a Gaussian distribution [64].

Thus in this specialized case, the utility function to be maximized (the probability of correct

classification) is given by:

$$p_c = \sum_{j=1}^N p(c_j) \int_{\Omega_{y,j}} \mathcal{N}_y(\bar{x}_j, (\Sigma + \Lambda_j)) dy. \quad (2.5.6)$$

2.5.2 Relationship between Estimation and Classification in the Gaussian Case

In general, the computation of the probability of correct classification is challenging, even when done numerically [64]. Optimization over sensor positions in order to determine the best trajectories of the sensors further renders the computation more challenging.

However, an understanding of the behavior of Equation 2.5.6 can be attained by investigating its dependence on the sensor states. Note that the covariance of the target estimate is affected by sensor motion, due to the spatially dependent measurement noise. Before proceeding, the following theorem is presented as it forms the basis for a *duality principle* that simplifies the task of maximizing the utility function.

Theorem 2.5.1 *Let the scalar-valued function $f : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ that operates on a matrix $A \in \mathbb{R}^{n \times n}$ be given as:*

$$f(A) = \sum \lambda \int \mathcal{N}(A) dz,$$

where λ is a constant, $z \in \mathbb{R}^n$, and

$$\mathcal{N}(A) = \frac{1}{(2\pi)^{\frac{n}{2}} |A|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} z^T A^{-1} z\right).$$

Then, given that A and $B \in \mathbb{R}^{n \times n}$ are positive definite, the following scalar-valued inequality holds:

$$f(A+B) < f(A).$$

Proof: Examination of the function, f , applied to the perturbed matrix $A+B$ yields that

$$\begin{aligned} \mathcal{N}(A+B) &= \frac{1}{(2\pi)^{\frac{n}{2}} |A+B|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} z^T (A+B)^{-1} z\right) \\ &< \frac{1}{(2\pi)^{\frac{n}{2}} |A|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} z^T (A+B)^{-1} z\right), \end{aligned}$$

noting that scalar determinants in the denominators are related by $|C+D| > |C|$ for C and D positive definite [57]. Further, use of the identity

$$(C^{-1} + D^{-1})^{-1} = C - C(C+D)^{-1}C,$$

and the fact that A and B are nonsingular allow further simplification:

$$\begin{aligned} \mathcal{N}(A+B) &< \frac{1}{(2\pi)^{\frac{n}{2}} |A|^{\frac{1}{2}}} e^{-\frac{1}{2}z^T (A^{-1}-A^{-1}(A^{-1}+B^{-1})^{-1}A^{-1})z} \\ &< \mathcal{N}(A) e^{\frac{1}{2}z^T (A^{-1}(A^{-1}+B^{-1})^{-1}A^{-1})z} \\ &< \mathcal{N}(A). \end{aligned}$$

The last inequality is seen by observing that a quadratic form is always positive, and hence, the exponential factor must be greater than or equal to identity. Noting that f is the integral over positively valued \mathcal{N} , the proof is complete. ■

Theorem 2.5.1 leads to the conclusion that the probability of correct classification, p_c , is inversely related to the sum of Σ_j and Λ_j , which are both positive definite. Hence, the following corollary follows immediately:

Corollary 1 *In the case of Gaussian sensor noise and Gaussian class distributions, maximization of the probability of correct classification, p_c , is achieved by minimization of the determinant of the estimate error covariance matrix $|\Sigma_j|$.*

Proof: Note that the class distribution covariance Λ_j is independent of sensor positions, and so examination can focus on changes in Σ_j due to the motion of the sensors. Define A greater than B (i.e., $A > B$) for $A, B \in \mathbb{R}^{n \times n}$ and positive definite if and only if the matrix $A - B$ is also positive definite. Hence, with this relationship defined for A, B positive definite matrices,

$$|A| > |B| \Leftrightarrow A > B,$$

or in other words, decreasing a positive definite matrix, such as Σ_j , decreases its determinant. Thus, the inverse relationship between p_c and $|\Sigma_j|$ from application of Theorem 2.5.1 is immediate. ■

Note that the objective for distributed optimal estimation tasks is to choose the collective team motions to minimize the estimation cost function, commonly given by the determinant of the estimate error covariance matrix. For this reason, we term the statement of Corollary 1 the *estimation-classification duality principle*.

Thus, for the special case of Gaussian noise and class distributions, there exists a direct connection between the classification problem and the task of estimating the target state. This result simplifies the distributed classification objective to one of distributed sensing, for which many possible approaches exist, including the gradient-based motion control laws developed earlier. These methods yield the optimal sensor motion paths and configurations for estimating or tracking targets,

which are, for the given formulation, the same paths and configurations for optimal classification of targets.

Chapter 3

Control for Decision Making

Beyond the collection and processing of sensory data, intelligent agents may be endowed with the ability to make decisions. Decisions represent a distillation of information contained in observations, and enable the autonomous system to perform a greater variety of tasks.

One can extend the notions of the previous chapter from using feedback control to improve the quality of sensing to that of using control to improve the quality of decision making. As before, the goal is to define both the characteristics of the decisions to be made as well as a measure of the quality of those decisions.

Traditionally, decision making systems have simply taken sensory information and reasoned about their environment, enabling them to follow this with an action of some sort. Of interest is an understanding of how such a system can take an action in order to improve the decision.

In order to begin this investigation, formulation of the decision problem such that it depends on the current mode or state of the system is necessary. Furthermore, the temporal notion of decisions should be captured by this formulation to characterize the “dynamics” of these decision making processes.

A simple example of using feedback to drive the decision making process is the task of changing lanes. Consider a driver who wishes to move to an adjacent lane. The driver may first glance (i.e., the control action) at the rear-view mirror to see if any vehicles are present within that view. Depending on what is visible in that view, she can surmise the presence or absence of an obstacle, or determine that another observation is necessary. Upon deeming another perspective necessary, she may take a look at the side-view mirror, again followed by the decision of whether it is safe to change lane or not. Finally if uncertainty remains, an additional action of checking the vehicle’s “blind spot” by turning her head may be taken to make another observation. The decision to change lanes can now be made with high degree of confidence.

Note that in the above problem, there are various aspects which make this problem interesting. First, the question of the sequence of looks is related to the type of information attained and the speed of the action needed. Further, it is possible to mistakenly believe a vehicle is present when in

fact it is not, and vice versa, due to uncertainty in observations (e.g., dirty windows, poor visibility, or limited visual acuity). Additionally, depending on how conservative the driver is, she may be less willing to accept greater uncertainty in her decision regarding the safety of the lane change, characterized by the need to take more observations. A more reckless driver, for example, may make a premature decision after simply glancing in the rear-view mirror.

This example highlights the various components of decision making that are sought in this research: the sequential nature of the problem, the fact that the uncertain environment leads to imperfect observations, and the role of the confidence in the decision. Once a proper understanding of the relationship between these factors is achieved, one can begin to examine how the decisions are affected by appropriate injection of control into the decision making system.

This chapter develops a framework for examining decision making systems that facilitate a natural study of the role of feedback in these processes, and presents a formulation of a search problem as a canonical example for this investigation. Delving into the relationship between decisions and feedback control yields a demonstration of the generality of this framework, as it pertains to a diverse array of search-related problems.

Section 3.1 lays the foundation for examining how decisions are made in a sequential fashion. Information is accumulated and integrated over time, providing a manner of dynamics for these decision processes. Available and applicable tools from statistical learning methods, such as the Sequential Probability Ratio Test, provide an approach to modeling decisions in a sequential framework, which facilitates the introduction of an example of a search problem as it pertains to decision making.

A class of search tasks is cast in a probabilistic context in Section 3.2, motivated by the fact that knowledge is imperfect due to uncertainties in the environment, modeling and information gathering modalities. The structure of this search problem is presented by developing the search context (e.g., the environment), a general model of target detection that accounts for imperfect information, and the probabilistic representation of these components, which enable the formulation of the search objective as a decision task.

Section 3.3 develops the Bayesian filter methodology that takes advantage of the sequential nature of the search task. The construction of the search problem as a decision (as presented in the previous section) facilitates the derivation of analytic expressions which represent the sequential evolution of the decision process as a function of the system parameters such as the detection model and the initial state of the decision. The closed-form expressions encourage and enable further analyses of the decision's evolution, illuminating various facets of this search problem, thereby enabling the researcher to attain additional insight.

The work to develop the framework presented in the previous section culminates in Section 3.4, which investigates the role of employing different strategies for controlling or directing the search.

The benefit of utilizing information in a feedback manner is clearly demonstrated in several contexts. Furthermore, these search strategies demonstrate the applicability and generality of this decision making construction, as it is shown to be relevant to conventional robotic search objectives, such as mobile agent search and rescue, as well as those not traditionally treated in a robotic context, such as modeling of behavior and human visual search.

3.1 Dynamics of Decision Making

Decision theory allows one to examine the process of decision making given observations. As described in the introduction, this process has been studied extensively in various communities, ranging from statistical learning and pattern recognition to signal detection and estimation. In these fields, the canonical construction is to inspect the decision of a binary hypothesis, where either an affirmative or null result is possible.

The process of evaluating and classifying a set of observations has traditionally been based on a binary hypothesis test, which assumes that a large set of observations has already been obtained. In essence, the computation of the decision rule or threshold, which is used to determine the class of subsequent observations, is a batch process after collection of many samples from the experiment or process. Various treatments of classical decision theory are readily found in [62, 63].

However, this requirement of performing a batch computation is a departure from the desired objective of gathering information sequentially and reasoning about the observation in real time. As motivated previously, this notion of sequential or iterative decision making is necessary for addressing dynamic systems where information changes over time. This necessity led to the development of decision-theoretic tools, including the Sequential Probability Ratio Test (SPRT), which provides a mechanism for integrating information over time and making a decision upon aggregation of sufficient information. In this section, a review of the SPRT and its construction are presented, along with a discussion on the metrics of performance of the resulting decision. These sequential decision making notions are used to introduce and motivate an examination of the search problem.

3.1.1 Sequential Probability Ratio Test

Developed by Wald and Wolfowitz [41, 42], the SPRT considers the decision of an alternative choice against the null hypothesis. As in the classical binary decision case, consider the task of choosing either hypothesis $H = 1$ or $H = 0$. In the lane-changing example, this decision problem might be equivalent to determining whether it is safe ($H = 1$) or not ($H = 0$) to move to the adjacent lane. Given a set of n observations, $D^n = \{d^1, \dots, d^n\}$, the classical approach is to represent the joint probability distribution of the measurements, $Pr(d^1, \dots, d^n)$, in the standard ratio test of a

hypothesis:

$$L(D^n) \triangleq \frac{Pr(d^1, \dots, d^n | H = 1)}{Pr(d^1, \dots, d^n | H = 0)} \underset{(H = 0)}{\overset{(H = 1)}{\geq}} \bar{L},$$

where $Pr(d^1, \dots, d^n | H = 1)$ represents the probability of obtaining the data given that the hypothesis is true, and similarly for $Pr(d^1, \dots, d^n | H = 0)$ given it is false. The constant, \bar{L} , represents the decision rule in the classical setting, such that a value of this ratio greater than the threshold corresponds to a decision accepting the hypothesis and vice versa. Notions of optimality in these types of decisions pertain to the appropriate selection of this threshold, \bar{L} , such as in the Neyman-Pearson criterion [62].

In the context of the Sequential Probability Ratio Test, however, *two* decision thresholds are considered such that given $t < n$ observations:

$$B < \frac{Pr(d^1, \dots, d^t | H = 1)}{Pr(d^1, \dots, d^t | H = 0)} < A, \quad (3.1.1)$$

where A and B ($B < A$) are positive constants. For ratio values such that the above expression is true (i.e., $B < L(D^t) < A$), an additional observation is required, whereas for $L(D^t) < B$ or $L(D^t) > A$, no additional measurements are taken and the decision process is completed with acceptance or rejection, respectively, of the alternative hypothesis, $H = 1$.

Continuing with the analysis, independence of the observations (and taking the logarithm) yields

$$\begin{aligned} \log L(D^t) &= \log \frac{Pr(d^1 | H = 1) \times \dots \times Pr(d^t | H = 1)}{Pr(d^1 | H = 0) \times \dots \times Pr(d^t | H = 0)}, \\ &= \sum_{i=1}^t \log \frac{Pr(d^i | H = 1)}{Pr(d^i | H = 0)} \triangleq \sum_{i=1}^t z_i, \end{aligned}$$

such that the decision task described by Equation 3.1.1 becomes

$$\log B < \sum_{i=1}^t z_i < \log A.$$

Thus, as each observation is obtained in sequence, the cumulative sum of the probability ratios is computed and tested against the thresholds, (A, B) , resulting either in termination of the decision process or the necessity to obtain an additional measurement.

3.1.1.1 Measures of Performance of Decisions

Given the parameters of the decision test, the sequential decision formulation provides insights into the relationship between the power of the test (i.e., error probability rates), the confidence of the

resulting decision, and the average number of observations required until a decision is reached. These characteristics define a multitude of metrics to evaluate the performance of the sequential decision task, such as time till decision, asymptotic confidence thresholds, rate of change of the belief function [41].

Given the power of the test and the thresholds at which decisions occur, one can determine the average or expected number of observations required to achieve a decision. This time till decision, denoted t_D , makes for an obvious measure of performance, as many dynamic tasks often require decisions and subsequent actions to be made in a timely manner. Other metrics mentioned above can be related to t_D , as a trade-off relationship exists between an increase in decision confidence levels or a reduction of error rates with a likely increase in the time until the decision is made. In the context of the decision thresholds, (A, B) , of the sequential decision process, the time till decision, t_D , is given by

$$t_D = \min \{ \arg(B(t) = A), \arg(B(t) = B) \}.$$

As time till decision will be the most relevant measure in the dynamic decision tasks in this chapter, it will hereafter be used to quantify the performance of the decision system given the parameter set. In the context of search and detection problems, this metric of time-to-decision is akin to the time-to-detection measure, as seen in the search theory literature [34].




Figure 3.1: Schematic of the evolution of three decision processes, with varying values of time till decision. Feedback can be used to improve the performance, i.e., reduce the time until a decision is made.

An illustration of this metric for schematic evolutions of different decision processes are given in Figure 3.1. All three example decision processes result in the selection of the null hypothesis at varying times, $t_{D1} < t_{D2} < t_{D3}$. One mechanism for improved performance (defined by a reduction in the time necessary to make a decision) is to inject appropriate feedback control into the dynamic decision process. Mobility of the searcher is used to apply this control in focusing *where* to search, but

other forms of feedback, e.g., modulation of sensor characteristics, can be used to answer additional questions such as *how* to search. The analysis and application of such feedback is the objective of the research presented in this chapter.

3.1.2 Application of Sequential Decision Making to Search

As introduced in the lane change example, a general application of sequential decision making is that of search, where the objective is to determine the presence of a target and its location in a search region, or determine that it is otherwise absent from the space. Search is intrinsically a sequential decision making problem, as illustrated by the lane-changing example, that requires observations at different times and places in order to gather information relevant to making a decision about the state of the environment.

The goal in a physical search problem is to generate the search paths in uncertain environments that best enable the searcher to locate a target (perhaps among other objects) using one or more mobile sensor platforms, possibly under resource constraints [28, 29, 30]. As discussed in Chapter 1, the above search problem serves a wide variety of applications and has thus been extensively studied by a variety of research communities.

The task of detecting an object or target in a region may be categorized by different aspects of the search problem [35]. A classification according to the model of the target dynamics can be represented by three distinct categories. A type I search represents the case where the target remains stationary, such as an injured person in the case of a search-and-rescue objective or land mines in demining operations. Alternatively, the target may have some dynamics, where the target's motion is not intentionally designed to either help or hinder the detection process by the searcher. These Type II search tasks include scenarios where the target is unaware it is being sought after (as in the case of searching for vehicles in adjacent lanes in the lane-changing example), or scenarios where the target's motion is due to environmental disturbances (e.g., a life raft drifting in the ocean) [38]. Classification as a Type III search is applicable for situations where the target is adversarial and employs motion to delay or prevent detection by the searcher [65]. Such tasks may include the detection of an intruder by a security guard in a building or exploration by a predator in search of prey. This taxonomy of search problems by the target dynamics is helpful for examining the diverse applications of search theory in general. Investigations of Type I and II search scenarios are explicitly presented in this thesis, noting further that the construction of the analytic formulas and numerical studies focuses on the case of a single searcher of these types so as to develop an intuition and understanding of the analysis. Study of noncooperative search as well as multiple searchers, in addition to other extensions remains a promising area of future research; however, the general search framework developed herein facilitates a straightforward means of examining this latter case, requiring only that an appropriate probabilistic representation of these extensions be

provided. Further discussion on these future directions can be found in Chapter 4.

The search problem described here may be generalized to a broader class of problems where the choice of observations is controlled to best search for an object or outcome. For instance, the scheduling of individual sensor nodes in a wireless sensor network [66] or the selection of the focus-of-attention in visual systems [67, 45] can also be formulated as search problems, where a control policy is generated to improve the information obtained by observations. The aim of the research presented in this chapter is to provide a unifying framework for modeling, analyzing, and understanding search in the context of decision-theoretic notions.

3.2 Probabilistic Search as a Decision

This section formulates the search task as a decision problem, examining the target’s presence or absence in a hypothesis-testing framework. As will be shown, this framework incorporates the various components of the search solutions proposed in previous works [34, 38] while enabling study of a broader class of search problems.

3.2.1 Problem Setup

Consider a single searcher, denoted S , located at position x_S in an environment \mathcal{A} , which is represented by a discretized grid of $|\mathcal{A}|$ cells, such as the example depicted in Figure 3.2. Such a discretization can be employed to capture either system characteristics such as limited range of the detector or physical partitions of the environment such as rooms in a building. The theoretical formulation generalizes to continuous representations of the environment; however, an emphasis on the discrete framework is motivated by its intuitive appeal for sequentially arriving observations as well as its practical advantage for implementation in applied systems.

Knowledge of the searcher’s whereabouts is assumed perfect to within the resolution of the search space discretization. Each cell in the discretized search environment will be indexed by a ($a \in \mathbb{Z}^+ \leq |\mathcal{A}|$), such that $x_S = a$ denotes that the searcher is currently in the a^{th} cell. The searcher is able to make a single detection observation in a given time step in a given cell. The mobile searcher’s motion, which consists of either remaining at its present location or moving to a single other cell within \mathcal{A} , is governed by its control strategy, which will be subject of study in Section 3.4.

Denote the target’s location by x_T , such that the expression “ $x_T \in \mathcal{A}$ ” reflects the presence of the target in the search space, and similarly “ $x_T \notin \mathcal{A}$ ” represents its absence. Furthermore, “ $x_T = a$ ” and “ $x_T \neq a$ ” signify that the target is and is not, respectively, located specifically in the a^{th} cell. Additionally, a model of the target dynamics is captured by the target state transition matrix, which represents the probability of the target transitioning from a given cell to another cell (see Section 3.2.4 for more details).




Figure 3.2: Example of discretized search region, \mathcal{A} . The locations of the target, T , and the searcher, S , are denoted x_T and x_S , respectively.

3.2.2 Search is a Decision Problem

The fundamental question the searcher is trying to answer is whether or not the target is present in the search region. This task clearly poses a binary decision problem with the alternative ($H = 1$) and null ($H = 0$) hypotheses defined as follows:

$$H = \begin{cases} 0, & \text{if } x_T \notin \mathcal{A}, \\ 1, & \text{if } x_T \in \mathcal{A}. \end{cases}$$

The objective is now to determine the probability that the affirmative hypothesis is true, i.e., $Pr(H = 1)$. This probability is a measure of the *aggregate belief* (ranging between 0 and 1) that the target is present somewhere in the region \mathcal{A} . In a similar manner, the individual cell belief probability, $Pr(x_T = a)$, represents the probability that the target is specifically located in cell a . This latter probability expression addresses the secondary question of identifying the location of the target within \mathcal{A} . Noting that the target can only occupy a single cell at any given instant, it can be immediately seen (by the Law of Disjoint Probabilities) that the overall or aggregate probability that the target is present in the search region is given by:

$$Pr(H = 1) = Pr(x_T = 1 \vee \dots \vee x_T = |\mathcal{A}|) = \sum_{a=1}^{|\mathcal{A}|} Pr(x_T = a).$$

Before a search has begun and before any measurements have been taken, the prior probability that the target is in \mathcal{A} is defined to be $Pr(H = 1) \triangleq \delta$, for $0 \leq \delta \leq 1$. For example, a prior belief value of $\delta = 1$ indicates that the target is certainly present in the region, and the search problem reduces to simply localizing the target, once detected, within the region. Nonunity values of δ ,

however, allow for the possibility that the target is not in the region to begin with. It is assumed that the target does not enter or exit the search region throughout the duration of the search, i.e., the fact that $x_T \in \mathcal{A}$ ($x_T \notin \mathcal{A}$) at the start of the search implies that the target is (not) in \mathcal{A} for all time. The confidence in this prior belief is reflected in the evolution of the probability as a function of observations (e.g., decreasing as null observations are received), and the ability to capture this initial uncertainty is one of the advantages of the framework presented in this chapter.

3.2.3 Detection of the Target

A major reason for employing a decision-theoretic approach is because it admits the use of detection sensor models that capture false alarms and missed detections. Such an approach is warranted as detection of a target can only be guaranteed in a probabilistic sense because measurements are generally taken in the presence of noise in practice. Consider d a binary random variable capturing the detection state, which can take values of either 0 or 1 with some probability. This detection variable represents the decision of whether or not the detection of a specific target has been made locally within a given cell.

Augmenting the notation to address the search problem, define $d_{a_l}^t$ to be the detection measurement at discrete time step t taken in cell a . This subscript term, a_l , is included when referring to the l^{th} detection measurement (allowing for multiple measurements in a single cell) of the specified cell a , and will be omitted otherwise when there is no ambiguity.

In this manner, the following detection model can be constructed for an imperfect detection measurement process, given the presence or absence of the target in the cell:

$$Pr(d_a^t | x_T) : \begin{cases} Pr(d_a^t = 0 | x_T = a) &= \beta, \\ Pr(d_a^t = 1 | x_T = a) &= 1 - \beta, \\ Pr(d_a^t = 0 | x_T \neq a) &= 1 - \alpha, \\ Pr(d_a^t = 1 | x_T \neq a) &= \alpha, \end{cases}$$

where α and β are the detection error probabilities for false alarms and missed detections, respectively. These error rates quantify the noise characteristics of the observation, and for a given sensor (α, β) can be determined experimentally or by sensor specifications. In the parlance of decision theory, the detector model is nothing more than the likelihood function of receiving a measurement given the state or hypothesis, and the error probabilities (α, β) represent the power of the decision test for the sensor.

The statement of the search task can now be stated as follows: *Given the detector model (α, β) and the prior belief δ , determine the evolution of the belief that the target is present in region \mathcal{A} as a function of the observations made until time t , i.e., $Pr(H = 1 | D^t)$, where $D^t = \{d^1, \dots, d^t\}$. This*

belief evolution ultimately governs the decision of whether (and if so, where) the target is located in \mathcal{A} , thereby completing the search.

3.2.4 Case Study: Search on a Grid

An example of the search problem described above can be constructed as follows. This example will be used repeatedly throughout this chapter. Consider the search space \mathcal{A} given by a square 10×10 grid (see Figure 3.3), such that $|\mathcal{A}| = 100$. A single mobile searcher is initially located in cell $(i, j) = (5, 8)$, and is equipped with a detector, characterized by its error rates (α, β) of false alarms and missed detections, respectively, as per Section 3.2.3. The searcher's initial belief or representation of the prior probability that the target is present in \mathcal{A} at the onset of search is given by $Pr(H = 1) = \delta$, which is divided up among the $|\mathcal{A}|$ cells according to some prior distribution.

For illustrative purposes, consider the example prior distribution depicted in Figure 3.3, which is modeled as a discrete approximation to a Gaussian distribution, centered about the cell $(i, j) = (1, 3)$ with covariance values $\sigma_i = 7$ and $\sigma_j = 15$, and appropriately normalized such that $\sum_{a=1}^{|\mathcal{A}|} Pr(x_T = a) = \delta$. However, it is important to note that any form for the initial belief probability density function (PDF), including multimodal and nonsmooth distributions, can be utilized due to the general Bayesian formulation of this framework.




Figure 3.3: The search problem on a 10×10 grid, with prior distribution modeled as a discretization of a Gaussian distribution with mean $(1, 3)$ and covariances $\sigma_i = 7$ and $\sigma_j = 15$. The initial positions of the searcher and target are $(5, 8)$ and $(4, 3)$, respectively.

Consider further a single target located in cell $(4, 3)$ in the prescribed search region. Note that in this example the initial belief, given by the distribution of δ in \mathcal{A} , allocates the belief prior probability incorrectly by attributing greater likelihood of finding the target in cells which do not contain the

target, and vice versa. The robustness of the decision-based belief framework presented in this paper is demonstrated by the eventual correction of this “bad guess” by sequential accumulation of information.

The target dynamics, as described earlier, can be represented probabilistically in terms of its state transition matrix. A stationary target (i.e., a search problem of Type I) is simply represented by an identity matrix of dimension $|\mathcal{A}| \times |\mathcal{A}|$, such that the probability of transitioning to any cell other than its current location is zero (requiring a self-transition probability of unity). A dynamic target of Type II will have more complex transition matrices, which will depend on the discretization of the environment as well as on the constraints on target motion. For example, a target which moves simply from one cell to the next following cell (defined by some ordering of the cells) is represented by the transition probability matrix, $\mathbf{\Pi}$:

$$\mathbf{\Pi} = [p_{ij}] = \begin{pmatrix} 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \\ 1 & 0 & \cdots & 0 \end{pmatrix},$$

where p_{ij} is the probability of transitioning from cell i to cell j in the next time step.

As can be seen by this example, deterministic models for the target motion can be captured as easily as probabilistic ones by these state transition probability matrix representations.

3.3 Discrete Bayesian Formulation

The use of Bayesian filters for autonomous applications is ubiquitous in the robotics community [40]. The primary reasons for their popularity are their applicability to general (i.e., non-Gaussian) probability density functions, as well as their inherent recursive formulation, such that arbitrary probability density functions (PDFs) can be maintained and updated in a relatively efficient recursion algorithm.

3.3.1 Recursive Computation of the Cell Belief Probabilities

Introduced in the previous section, the probability distribution of interest is the belief function, $B(t)$, defined by

$$B(t) \triangleq Pr(H = 1|D^t) = \sum_{a=1}^{|\mathcal{A}|} Pr(x_T = a|D^t). \quad (3.3.1)$$

This expression shows that in order to compute the aggregate belief probability, given a sequence of t observations, one must compute each of the individual cell belief probabilities. Note that $Pr(x_T)$

represents the PDF of the location of the target, x_T .

Computation of the cell belief probability is done using Bayesian filtering [40, 39], which incorporates the propagation of the target PDF with an update step after an observation is received. The prediction component of the discrete filter utilizes the discrete analog of the Chapman-Kolmogorov equation (e.g., see [38]), given by

$$Pr(x_T^t|D^{t-1}) = \sum_i Pr(x_T^t|x_T^{t-1} = i)Pr(x_T^{t-1} = i|D^{t-1}),$$

which multiplies the process model of the target, captured by $Pr(x_T^t|x_T^{t-1} = i)$, with the target PDF (i.e., the belief) of the previous time step.

The update step is performed by application of Bayes rule over the target probability distribution, given the latest observation:

$$Pr(x_T^t|D^t) = \frac{Pr(d_k^t|x_T^t, D^{t-1})}{Pr(d_k^t|D^{t-1})}Pr(x_T^t|D^{t-1}).$$

As can be seen by Equation 3.3.1, the computation of the observation update of the belief probability function in an individual cell is required and it is given by:

$$Pr(x_T = a|D^t) = \frac{Pr(d_k^t|x_T = a, D^{t-1})}{Pr(d_k^t|D^{t-1})}Pr(x_T = a|D^{t-1}),$$

where one can recognize the numerator term to be the detector model (i.e., likelihood function), and $Pr(x_T^t = a|D^{t-1})$ is the individual cell belief value determined at the previous time step, which provides the recursion in the filter. The term in the denominator, $Pr(d_k^t|D^{t-1})$, is the marginalization of the measurement, and can be computed in closed form by

$$Pr(d_k^t|D^{t-1}) = \sum_{H=\{0,1\}} Pr(d_k^t|H, D^{t-1})Pr(H|D^{t-1}).$$

Making use of the fact that

$$Pr(d_k^t|H = 1, D^{t-1}) = \frac{\sum_{b=1}^{|\mathcal{A}|} Pr(d_k^t|x_T^t = b, D^{t-1}) Pr(x_T^t = b|D^{t-1})}{\sum_{b=1}^{|\mathcal{A}|} Pr(x_T^t = b|D^{t-1})},$$

and that

$$Pr(d_k^t|H = 0, D^{t-1}) = Pr(d_k^t|x_T^t \neq b, D^{t-1}) \forall b,$$

the final expression for the update of the individual cell belief probabilities is determined to be:

$$Pr(x_T^t = a|D^t) = \frac{Pr(d_k^t|x_T^t = a, D^{t-1}) Pr(x_T^t = a|D^{t-1})}{Pr(d_k^t|x_T^t = k, D^{t-1})Pr(x_T^t = k|D^{t-1}) + Pr(d_k^t|H=0, D^{t-1})(1 - Pr(x_T^t = k|D^{t-1}))}. \quad (3.3.2)$$

Introducing the definitions

$$\begin{aligned} \Psi &\triangleq (1 - d^t)\beta + d^t(1 - \beta), \\ \Phi &\triangleq (1 - d^t)(1 - \alpha) + d^t\alpha, \end{aligned}$$

Ψ and Φ represent detections for when the target is present and absent, respectively, in the currently observed cell. The expression, Equation 3.3.2, for the individual cell belief probability update becomes

$$Pr(x_T^t = a|D^t) = \frac{Pr(d_k^t|x_T^t = a, D^{t-1}) Pr(x_T^t = a|D^{t-1})}{\Psi Pr(x_T^t = k|D^{t-1}) + \Phi (1 - Pr(x_T^t = k|D^{t-1}))}. \quad (3.3.3)$$

The terms in this expression can be described in the following intuitive manner. The numerator represents the probability of receiving a detection measurement at time t in cell k times the prior probability of the target's presence, whereas the denominator contains the normalization constant over all detection types. This recursive expression provides a compact and efficient way to update the belief function at every time step as the searcher observes a sequence of unexplored and/or previously visited cells. The benefit of its simplicity is seen both in analytic extensions and in algorithmic implementations demonstrated in the following sections.

3.3.2 Closed-Form Expressions for Uniform PDFs

Simplification of the update expressions can be made under the assumption that the initial belief probability density function is uniformly distributed. Given no prior knowledge of the distribution on the target's location, one can use uniform distribution to represent maximal uncertainty. In this case where each of the initial (i.e., prior to any observations being taken) cell belief probabilities are uniform (constant), it is possible to generate closed-form expressions for the evolution of the belief function, parametrized by the number of measurements taken. These functions provide insight into the search process, as well as form the foundation for computationally efficient algorithms.

For a sequence of null detections ($D^t = 0$) and for uniform prior probabilities, $Pr(x_T = a) = \frac{\delta}{|\mathcal{A}|}, \forall a$, the overall belief probability is given by:

$$Pr(H = 1|D^t = 0) = \frac{t\beta\delta + (1 - \alpha)(|\mathcal{A}| - t)\delta}{t\beta\delta + (1 - \alpha)(|\mathcal{A}| - t\delta)}. \quad (3.3.4)$$

Recall the problem setup proposed in Section 3.2.4. The objective is to search for a target in the 10×10 search space with the initial belief prior probability of $\delta = 0.5$. Consider the special case of a uniform prior distribution of this belief probability, as shown in Figure 3.4. The belief evolution governed by Equation 3.3.4 can be studied as a function of the error rates, made possible by the closed-form expression.




Figure 3.4: Illustration of uniform prior probability distribution. The initial belief (i.e., aggregate probability) value is $\delta = 0.5$, such that all cells contain $\frac{\delta}{|\mathcal{A}|}$ probability.

Figure 3.5 plots, for different values of the false alarm rate α , the evolution of the belief function for a searcher with a constant 10% missed detection rate (i.e., $\beta = 0.1$) while it searches a 10×10 array of cells. Alternatively, the dependence on the missed detection rate β is illustrated for this sequence of null detections in Figure 3.6, with $\alpha = 0.5$ held constant. Both Figures 3.5 and 3.6 are enabled by the derived closed-form formulas, which give insight into how the belief function evolves as a function of time and sensor characteristics. For example, the limiting case of no false alarms in Figure 3.5 demonstrates that the possibility of having missed the target in the sweep of the search space limits the total performance of the search. Similarly, increasing the rate of missed detections reduces the confidence in the received null detections, as seen in Figure 3.6.

In the case of uniform priors, the derivative of the above expression can be taken explicitly to determine the rate-of-change of the belief as a function of time step t :

$$\frac{\partial}{\partial t} Pr(H = 1 | D^t = 0) = \frac{|\mathcal{A}|(1 - \alpha)(\alpha + \beta - 1)(1 - \delta)\delta}{[|\mathcal{A}|(1 - \alpha) + t\delta(\alpha + \beta - 1)]^2}.$$

The belief evolution for a consecutive sequence of null measurements followed by a nonnull measure-




Figure 3.5: Evolution of $B(t)$, given a sequence of null observations, $D^t = 0$, for varying α , constant $\beta = 0.1$, $\delta = 0.5$, and grid size, $|\mathcal{A}| = 100$.

ment is often of interest (e.g., [40, 68]). This belief is given succinctly by

$$\begin{aligned} Pr(H = 1 | d^t = 1, D^{t-1} = 0) \\ = \frac{\alpha\beta(t-1)\delta + (1-\beta)(1-\alpha)\delta + \alpha(1-\alpha)(|\mathcal{A}| - t)\delta}{\alpha\beta(t-1)\delta + (1-\beta)(1-\alpha)\delta + \alpha(1-\alpha_D)(|\mathcal{A}| - t\delta)}, \end{aligned} \quad (3.3.5)$$

where again the advantage of such a closed-form expression is seen by the fact that its derivative is easily computed:

$$\begin{aligned} \frac{\partial}{\partial t} Pr(H = 1 | d^t = 1, D^{t-1} = 0) \\ = \frac{|\mathcal{A}|\alpha^2(1-\alpha)(\alpha+\beta-1)(1-\delta)\delta}{[|\mathcal{A}|(1-\alpha) + (1-t\alpha)\delta(\alpha+\beta-1)]^2}. \end{aligned}$$

Figure 3.7 depicts the belief evolution for an illustrative example where a nonnull observation is received by the detector during the course of the search (at $t = 65$), where the searcher sequentially enters a new cell at each step. As can be seen, the perfect detector (i.e., $\alpha = \beta = 0$) can immediately decide that the target is present upon arrival of a nonnull observation, and thus the belief in this case jumps to unity, which ends the search task. In contrast, since false alarms are possible with an imperfect detector, only an incremental increase in the belief is registered with a positive detection, and the search necessarily must continue as the desired confidence to make a decision has not yet been reached.

The special case of uniform prior distributions is presented to simply demonstrate the structure




Figure 3.6: Evolution of $B(t)$, given a sequence of null observations, $D^t = 0$, for varying β , constant $\alpha = 0.5$, $\delta = 0.5$, and grid size, $|\mathcal{A}| = 100$.

of the Bayesian framework; however, nonuniform prior belief probabilities are easily and efficiently incorporated into the decision-based formulation in the general case, as is demonstrated by the example in the following section.

3.3.3 Case Study: Search on a Grid, Revisited

Consider again the search problem constructed in Section 3.2.4, where the searcher moves through the search region of size $|\mathcal{A}| = 100$ seeking to determine the location of a single target if it is present. Given the expressions for computing the evolution of the searcher's belief as detections are sequentially processed, the performance of different motions through the search region can be investigated. These motion strategies govern the searcher's trajectory (i.e., the sequence of visited cells), thereby affecting the belief probability distribution through its observation updates. Two basic but naïve approaches presented below are random search, where the next cell to visit is chosen randomly from adjacent cells, and sweeping search, which executes parallel search tracks in exhaustively covering the search region.

3.3.3.1 Random Search

Serving as a baseline strategy for comparison, a searcher employing a random walk strategy provides a lower bound on the time to decision for any complete search, which is one that guarantees visiting all cells at least once [34]. The selection of the next cell to visit, denoted k^t , is chosen at random, such




Figure 3.7: Effect of receiving a nonnull observation at $t = 65$ on the evolution of the belief for the perfect and imperfect detector.

that the distance between the current and next cell, $\Delta k = \|k^t - k^{t-1}\|$ is within the reachable set of cells, $\Delta k < \kappa$. The constraint κ may reflect dynamic constraints on the mobile searcher, representing a limitation on the maximal cell-to-cell distance attainable in a single time step. Choice of the next cell represents the searcher's motion *control signal*, such that for the random walk strategy it is given by:

$$k^t = \underset{\Delta k < \kappa}{\text{rand}}(k) \quad (\text{Strategy I}),$$

where the *rand* function denotes the random selection of the next cell, given that it is a reachable cell. Note that this strategy does not require knowledge of the belief distribution. Variations on this strategy include biased random walk methods, where, on the average, the searcher may tend toward some preferred direction such as the peak location of the prior belief distribution. Such an approach may be useful for low computation, local sensing of the belief gradient, which may be relevant in foraging scenarios or static source (e.g., odor, chemical plume) localization tasks.

Figure 3.8 depicts several instances of the evolution of the belief probability distribution using the random search method for the example search task outlined previously. As can be seen, the random approach leads to inefficient searches resulting in long waits until a decision is made. While a decision can eventually be reached (e.g., in this example, after over 1000 time steps), it is evident that different approaches can be utilized to reduce this time till decision.

3.3.3.2 Sweeping Search

The sweeping approach, where the searcher moves from one cell to the next in an incremental fashion, represents the classical Boustrophedan searching strategy used in traditional robotic coverage algorithms [69]. The sweeping search, like the random search, also does not use the distribution of the belief in selecting the next cell to visit. The sweeping search strategy is represented by the




Figure 3.8: Illustration of belief probability distribution evolution for a single searcher randomly walking through the example search space to determine the presence of a stationary target. The random walk search strategy results in long times-till-decision but can be used as a limiting case in examining the performance of different search strategies.

following control signal

$$k^t = k^{t-1} + 1 \quad (\text{Strategy II}),$$

where the next cell to visit is simply the next cell according to some ordering of the cells (e.g., cells in a grid numbered sequentially top-to-bottom, left-to-right). The changes in the belief probability distribution for a simulation using the sweeping search strategy is shown in Figure 3.9. The problem setup is, once again, the same as presented in Section 3.2.4.

As in the random search, the sweeping approach discussed above is also inefficient in search, as much time is spent examining cells with low chance of containing the target. The result is that many observations are required before a decision regarding the target's presence can be made, such that the large time till decision is impractical for most decision making contexts

One advantage of the sweeping approach is that the coverage of the region is done in linear time and distributes the number of repeated visits to a specific cell in the context of fairness (i.e., a cell will be revisited only after all other cells have been visited). However, in the presence of noisy detections, multiple measurements in particular cells, resulting in a biased allocation of observations, may often be beneficial for reducing the uncertainty present in those cells.

The random walk and sweeping strategies generate search paths that ignore any available current or updated information about the belief of the target's presence. As such, these approaches do not offer much in the way of performance, as the time until decision are orders of magnitude larger when compared to strategies (see Section 3.4) which do take advantage of this information in a feedback




Figure 3.9: Illustration of belief probability distribution evolution for a single searcher executing sweeping search in example search space to determine the presence of a stationary target. The sweeping search strategy, like random search, also results in long times-till-decision.

manner. Instead, the random walk and sweeping methods can serve as lower bounds on performance (measured by the time until a decision is made) for the class of strategies that do take the current belief into account in determining search trajectories [33, 34]. Clearly, the use of feedback to control and guide the searcher’s motion should affect how quickly the search task can be completed, which motivates the investigation into the role of control in improving the decision making process.

3.4 Control Strategies for Decision Making

The framing of the search problem as a sequential decision problem enables analysis of the decision’s temporal evolution. However, the analysis thus far has only considered the role of the searcher’s state on the search process passively, rather than providing a method for governing the dynamics of the searcher’s state in future time. In other words, the question of *how* to modify the search trajectory to improve the quality or performance of the decision making task requires application of feedback control.

Using the formulas for belief evolution, one can evaluate different *search control policies* with respect to their ability to improve the decision making performance. The task of generating this policy defines the search path control problem. This section examines several approaches to address this task, ranging from intuitive and known methods to several novel search strategies which are facilitated by the decision making framework presented in this chapter. In addition to being relevant to the target search problem studied here, these novel search strategies also offer insight into other

search problems that arise in a variety of fields, including visual search and sensor networks.

Of importance in any sequential decision task is the specification of the termination criteria. An additional advantage of the decision framework proposed in this chapter is that it intrinsically defines the conditions for terminating the search. More specifically, the search continues as long as the value of the belief function, as defined by Equation (3.3.1), lies within threshold values, A and B ($0 \leq B < A \leq 1$). A belief threshold value of $A = 1$ ($B = 0$) requires that the search terminate only when the searcher is absolutely certain that the target is present (absent). These decision thresholds offer additional degrees of freedom in the search problem definition. The construction of the search problem has the form of the sequential hypothesis-testing methodology described by the Sequential Probability Ratio Test (as described in Section 3.1.1) developed by Wald and Wolfowitz [42]. In the SPRT-like framework, observations are gathered as long as there is ambiguity in the acceptance or rejection of the hypothesis, which, in the case of the search task, is the determination of the presence or absence of the target within region \mathcal{A} according to the thresholds A, B .

These thresholds for specifying the end of the decision process lead to a practical and ubiquitous measure of the decision task performance given by the time until a decision is made. In many applications, this measure may represent the goal of making a decision in minimum time. Using this time till decision as the metric for the proposed search task, a comparison of different control strategies can be conducted, perhaps weighed against other characteristics such as their computation requirements and complexity.

Each search control strategy described below specifies which cell, k^t , to visit next at time t , given the sequence of previous observations up to time $t - 1$, D^{t-1} . Where relevant, dynamic constraints on the searcher may limit the reachability of cells and require that the next cell be within some distance κ from the current cell, such that $\Delta k \triangleq \|k^t - k^{t-1}\| < \kappa$ for a given choice of k^t , where $\|\cdot\|$ denotes an appropriate distance metric (e.g., 1-norm or Euclidean distance). Furthermore, since, by construction, the entire probability distribution function, $Pr(x_T | D^{t-1})$, at time, $t - 1$, is known, the cell containing the maximal cell belief value, k_{max}^{t-1} at time, $t - 1$, can easily be determined and used to guide the selection of k^t .

3.4.1 Look Ahead Search

In optimal path-planning, the goal is to generate the search trajectory along which an objective function is maximized. As is well known, the solution to such problems is often computationally expensive, if not intractable [70]. In fact, the problem of search on a discrete grid to minimize the time to detection is *NP*-hard [71]. Choosing probability of detection as the objective function to be optimized, as in, e.g., [40], can also be shown to be *NP*-complete. Hence, given the computational complexity of the task, suboptimal solutions (e.g., approximation or heuristic methods) are sought to solve the search path-planning problem within a practically relevant time.

Table 3.1: Pseudocode for the look ahead algorithm with finite horizon w

```

function LookAhead(current cell  $k^{t-1}$ , belief  $B_{t-1}$ , window size  $w$ )
  returns: optimal policy,  $\pi^{t-1}$ , and utility,  $U(\pi^{t-1})$ .

  Define the set of allowable actions (e.g., list of reachable cells)
  For each action,  $k^t$ 
    Take a detection measurement,  $d^t$ , in cell  $k^t$ 
    Compute belief function,  $B_t$ , over all cells (Equation 3.3.3)
    If not at end of window
      LookAhead( $k^t$ ,  $B_t$ ,  $w - 1$ )  % depth-first optimization
    Else
      Store belief in  $k^t$  as utility value
  end

  Identify the utility-optimizing action,  $\bar{k}^t$ 
  Return augmented policy  $\pi^t$  and utility,  $U(\pi^t)$ 

```

A common approach is to use a “look ahead” window of w time steps, over which the optimization of the path is performed. Also known in optimal control as *receding horizon control* [72, 24], with finite horizon w , this approach generates a sequence of actions, or a policy, for the next w steps, denoted $\pi = \{k_1, \dots, k_w\}$, and finds the optimal policy, π^* , which maximizes the objective function along the path over the horizon window, i.e.,

$$\pi^* = \arg \max_{\pi = \{k_1, \dots, k_w\}} \sum_j^w Pr(x_T = k_j | D^t, d_{k_1}, \dots, d_{k_{j-1}}).$$

The first step of this finite horizon optimal policy is then selected to be the next cell to be visited in the next time step. Implementation of the look ahead search algorithm in pseudocode is described in Table 3.1. At the next time step, the optimal trajectory over the horizon is recomputed. Note that for $w = 1$, the one-step look ahead strategy is the discrete steepest gradient approximation of the belief, which is simply the greedy algorithm for maximally increasing the objective function. More generally, the planning window size w can be chosen to provide an approximation to the optimal solution while maintaining computational feasibility. The closed-form expressions derived in Section 3.3 offer substantial implementation advantages by enabling search path plans over longer planning horizons.

Investigation of the performance of this look ahead strategy can be done to evaluate the role of feedback in the decision making process. As an illustration of the decision making framework, statistical analyses of the strategy can be conducted in the context of the example scenario constructed in Sections 3.2.4 and 3.3.3. Recall that the task is for a searcher to maneuver throughout the search region \mathcal{A} according to some search strategy, which will be taken to be the look ahead strategy presented above.

Figure 3.10 illustrates the histogram plot of the decision times for 5000 simulation trials employing the look ahead strategy for varying window sizes, $w = \{1, 2, 3\}$. Each run utilized the same initial conditions for the locations of the searcher and target, as well as the initial belief probability, $\delta = 0.5$, and detection model, $\alpha = 0.2, \beta = 0.1$, such that the distribution of the decision times is a result of the random uncertainty present in the detection observations in each cell.




Figure 3.10: Histogram plots of 5000 simulation studies of search employing the look ahead strategy for window sizes, $w = \{1, 2, 3\}$. The resulting mode values for the time till decision are $t_D = \{57, 51, 48\}$, respectively, showing a reduction in average search time for increasing search horizon length.

From this plot, the mode (i.e., the most frequent value of the histogram) of the decision time distributions can easily be determined for each of the window sizes to be:

$$\text{decision time, } t_d : \begin{cases} 57, & w = 1, \\ 51, & w = 2, \\ 48, & w = 3. \end{cases}$$

As expected, there is a slight improvement in performance (i.e., a shorter time to decision) for increasing window sizes, which is due to the fact that, for longer planning horizons, additional information is taken into account in generating the search path. However, given the fact that the performance of the look ahead algorithm is greatly dependent on initial conditions (e.g., initial location of the searcher relative to the target, error in the initial prior belief probability distribution, or size of the environment), the cost of computational burden seems to outweigh the improvement in performance. Instead, a one-step look ahead control law (i.e., steepest gradient ascent), albeit suboptimal, may serve as a reasonable compromise between performance and complexity. Observing

that the random-walk and sweeping strategies are extremely inefficient in their searches, significant advantages in performance can be found by utilizing even simple feedback laws such as look ahead search with small window sizes.

3.4.2 “*Drosophila*-Inspired” Search

Search problems are not restricted only to engineered problems. In fact, the search problem presents itself in many components of natural systems, as relevant notions are involved in tasks ranging from foraging for resources to hunting and predation in a territory.

An example of search is exhibited in walking *Drosophila melanogaster*, a.k.a. the fruit fly, in search of food. Study of such ethology of biological systems is of significant interest in the efforts to relate neural processing of motor and sensory feedback systems to exhibition of macroscopic animal behaviors (e.g., [73, 74]).

As studied and presented in [75], the trajectories of individual fruit flies are characterized by the behavior of identifying a particular location of interest, pursuing a straight-line trajectory (with minimal deviation) toward this goal, and upon arrival, finding another location and following another straight-line path to this updated goal location. The insect’s motion is presumed to be in search of food and evidence suggests that its behavior is mediated by visual sensory feedback.

Inspired by the ability of the fruit fly to execute search behavior that goes beyond naïve random search, this section develops a search strategy which identifies the cell containing the maximum belief value (which is simple to determine as the searcher maintains the belief distribution over all cells) and specifies a goal location for the searcher. The computational simplicity of this approach offers an advantage over more computationally expensive, planning-oriented approaches, such as the look ahead strategy examined in Section 3.4.1. This strategy is presented as a low computation expense yet high information gain option, whose performance will be quantified in what follows.

In this algorithm, the searcher determines and records the cell number containing the highest belief value, k_{max} . Recall that κ represents a constraint on how far the searcher can travel in a single time step, and can either be given by the searcher dynamics or specified by the desired resolution of the search. The next cell to observe is determined by:

$$k^t = \arg \min_{\Delta k < \kappa} (\|k_{max} - k^{t-1}\|) \quad (\text{Strategy IV}).$$

For fruit flies walking on a flat surface (i.e., in continuous \mathbb{R}^2 space), the use of the 2-norm (i.e., Euclidean distance) offers a reasonable approximation to their behavior, and $\|\cdot\|$ is also used and shown in [38] to address scenarios of premature search termination due to sensitivities in the optimization routine.

The results of simulation studies are shown in Figure 3.11, which represents the histogram of

time till decision values for 5000 trials, given the same setup as treated previously. The 1-norm is used to define distances between cells, appropriate for the lattice structure of the grid.




Figure 3.11: Histogram of times-till-decision for 5000 trials employing the “*Drosophila*-inspired” search strategy. The average time to decide accurately that the target was present is given by the mode value, $t_D = 89$.

The “*Drosophila*-inspired” search strategy yielded an average search time of $t_D = 89$, which is longer than the average search times of the greedy algorithm examined previously. This fact is certainly expected, as the goal location is only updated once the searcher arrives there (which may take multiple time steps), unlike the recomputation of optimal paths at every time step done in look ahead search. However, the computational simplicity of the flylike search strategy offers significant advantages without sacrificing much in terms of performance in the decision making task. Clearly this approach is more practical in cases where computation resources are limited, as in individual sensor nodes in a mobile sensor network, or in deployments of microscale aerial vehicles, as well as, arguably, in the brain of the fruit fly. The “*Drosophila*-inspired” search strategy demonstrates that even small feedback interactions in the search decision making process offer substantial performance enhancements over those without such feedback mechanisms.

3.4.3 “Saccadic” Search

Further motivated by examples of search that occur naturally, this section presents an approach motivated by the human (primate) visual system. The strategy derives its name from the way the eye in the human visual system *saccades*, or executes a rapid jump in its focus of attention, from one salient feature in its visual field to another [76]. In visual search, the task is to find, if it is present, an object or feature within a scene (e.g., an image) [45, 77, 78, 79].

One can readily envision the applicability of the search decision framework presented in this chapter to this problem of visual search. The “saccadic” search strategy developed in this section simply focuses the search on the cell containing the maximal belief probability at every time step,

executing a “saccade”-like jump from peak to peak in the belief probability distribution.

Analogous to the control signal representations of the previous search strategies, the next cell to visit under the “saccadic” strategy is determined according to:

$$k^t = k_{max}^{t-1} \triangleq \arg \max_k (Pr(x_T = k | D^{t-1})) \quad (\text{Strategy V}).$$




Figure 3.12: Histogram of times-till-decision for 5000 trials employing the “saccadic” search strategy. The average time to decide accurately that the target was present is given by the mode value, $t_D = 61$.

The average time till decision for the “saccadic” strategy in the context of the example problem setup was observed to be a value of $t_D = 61$. The distribution of these decision times over 5000 simulation trials is illustrated in Figure 3.12. In terms of performance, this strategy offers comparable average times until a decision is made as that of the look ahead search approach, under the caveat that jumps between cells may not be dynamically feasible, e.g., for a mobile sensor platform with finite velocity. However, this strategy provides a method for examining the “inertialess” or damped case of these physical search strategies. Another interpretation is to consider the fact that the transit time to get to the goal cell must be integrated, representing a possible method for weighting search trajectories by this additional cost. Future investigations into this and other notions can provide further insight into the search mechanisms as well as the role of feedback control in decision making tasks.

The appearance of additional smaller peaks in the histogram plots of both the “saccadic” and the “*Drosophila*-inspired” strategies can be explained. These secondary peaks correspond to distinct “modes” in the search evolution, which depend on the sequence of observations of the cell containing the target. As shown in the distributions in Figures 3.11 and 3.12, the primary peak corresponds to the trials where the first visit to the target cell resulted in a positive detection. Subsequent peaks correspond to the case where a missed detection occurs during the first visit to the cell, which leads to revisits of the target cell.

As suggested by its namesake, the “saccadic” search strategy may offer some insight into (or may provide a model for understanding) the mechanisms governing human visual search [45, 77], in which information is likely gathered sequentially using foveated glimpses of patches of the visual field, which can be modeled as observations of cells in the presented framework. Discussion of this line of future research can be found in Chapter 4.

Other relevant uses include, as an example, sensor networks for event detection applications, where sensor nodes (i.e., cells) are triggered by possible events (such as intrusion). In these applications, a belief-dependent decision or action is executed in the context of search. The advantage of the decision making framework presented in this chapter is that it generalizes these various search problems into a single objective of making a decision regarding the presence or absence of a target. Whether the search is for a physical entity such as a person or robot, or it is for a visual feature or entity within a visual field, the underlying goal is the same.

3.4.4 Discussion on Control Strategies

As demonstrated by the search strategies presented in the preceding sections, the utilization of feedback greatly improves the performance of the decision objective, as measured by the time taken to reach a conclusive decision. The flexibility of the decision making framework, as it pertains to probabilistic search, enables it to address a wide variety of search-related objectives.

The approach of computing the optimal path over a finite horizon, as presented in the look ahead strategy, is one that offers high performance, in terms of making a decision in a minimal number of time steps but at a computational expense. However, the decision making formulation of the search task facilitated the closed-form solutions derived in Section 3.2, which greatly enable more efficient computation of these search trajectories.

Two additional strategies provide computationally inexpensive alternatives for conducting search, which serve well in a host of applications where computation power is limited, e.g., sensors embedded on microaerial vehicles or simple nodes in a sensor network. Specifically, while the look ahead strategies offer slightly better performance, a more comprehensive metric of performance which accounts for computing time of search paths would immediately illuminate the significant advantages of the “saccadic” and “*Drosophila*-inspired” strategies. In the case of search using mobile sensors, it may be more prudent to have simpler, fast-moving detectors rather than necessitating computationally powerful (which usually translates to more hardware and less dynamic capability) sensor platforms. Furthermore, unlike the optimization-based look ahead schemes, these strategies have the added benefit that their performance is largely independent of the initial conditions, and they also avoid the issues of getting stuck in local minima of the objective function which limit most optimization routines.

The results of the simulation studies presented in the previous section are summarized in Ta-

Table 3.2: Performance comparison of different search control strategies.

Search Control Strategy	Time till decision t_D (mode)	Time/Observation Δt_C (msec)	Total Effective Time t_C (msec)
Random Walk	2833	0.2778	787.0074
Sweeping	3150	0.2820	888.3
Look ahead ($w = 1$)	57	0.6491	36.9987
Look ahead ($w = 2$)	51	3.4800	177.48
Look ahead ($w = 3$)	48	31.9576	1533.9648
“ <i>Drosophila</i> -inspired”	89	0.3784	33.6776
Random Jump	3496	0.3024	1057.1904
“Saccadic”	61	0.3271	19.9531

ble 3.2. The average time taken per observation is a measure of the computational cost of each calculation cycle, which includes the time for the observation, the belief distribution update, and the determination of the next cell to visit. Hence, the total effective computation time, t_C , reflects the performance of the decision task (i.e., time till decision) weighted by the computational expense of the search strategy employed. As can be seen, each observation and belief update is computed in minimal time for the naïve random walk and sweeping strategies due to their simplicity, but the disadvantage of lacking feedback in the search strategy is demonstrated in the significantly large number of observations required to complete the search task. Conversely, the look ahead strategy can arrive at a decision in the fewest number of observations but incurs large penalties in computation, especially as the window size increases. Since the effective performance of the look ahead strategy further deteriorates with a larger feasible action set, the advantage of this approach for many applications may be outweighed by the computational burden imposed. The “*Drosophila*-inspired” strategy provides an efficient and effective alternative, as seen by the minimal total effective time necessary to complete the search task.

Additionally, comparison of the “saccadic” strategy to a random jump strategy (where the next cell to visit is chosen randomly from all cells in the search region) further shows that utilizing information about the belief distribution in a feedback manner offers significant performance benefits. Such comparisons and insights are both useful and necessary for the application of search to practical implementations.

Chapter 4

Conclusions and Future Directions

4.1 Discussion and Summary

The principle of using feedback to augment the capabilities of dynamic systems is well established, as outlined in the previous section. Recognizing that information gathering is itself a dynamic and evolving process motivates the underlying theme to the contributions of this thesis, which is how the use of feedback can enhance the information gathering process.

In Chapter 2, a distributed estimation framework was presented for studying how a team of mobile agents can cooperatively use motion to improve the performance of the sensing and tracking of a dynamic process. Given a representation for the overall uncertainty of the team's estimates, a gradient-based decentralized motion control law for each mobile agent was derived in closed form. These analytic expressions generate local uncertainty-minimizing trajectories for the (possibly heterogeneous) sensors as they take measurements of (possibly multiple) dynamical processes, embodied in this work as dynamic targets. The advantage of these specific formulas are apparent in their efficient implementations, making it possible for realizations in practical distributed sensing mobile agent systems. Investigation of the effect of a noisy communication environment, where disturbances also depend on inter agent distances, resulted in a straightforward extension of the analytic motion control law, which illuminated the trade-off relationship between sensing and communication performance. This insight can provide additional guidance for the design of distributed mobile sensing networks.

Additionally, the “estimation-classification duality principle” derived in Section 2.5 describes the fact that the use of motion to improve estimation has the added benefit of improving the performance in certain classification tasks, under assumptions of underlying Gaussian distributions. As computation of classification probabilities are typically challenging, this result enables the use of simpler motion control laws, such as the gradient-based method proposed in this thesis, to minimize the probability of incorrect classification.

Chapter 3 formulated a probabilistic decision-theoretic model for the task of physical search of

a target in a search region. The Bayesian framework represents this search as a decision problem, further enabling the use of imperfect sensors, which have not been well addressed in the literature on robotic coverage. Investigation of the evolution of the belief of the target’s presence in the search region resulted in generation of analytic closed-form expressions that govern the belief function. These functions serve two main purposes: they offer a more general understanding of the search problem (e.g., notions of rate-of-change of belief), and they facilitate an easier implementation of search strategies in physical search systems.

Several different strategies were investigated and proposed for generating search paths. These studies found that the decision framework enables the use of a variety of measures of the search performance, including the time until a decision is made about the presence or absence of the target. Different search approaches were evaluated using this metric in simulation studies, observing that slight increases in performance are often not worth the computational burden, which is an important consideration for many applications of interest.

The two approaches proposed – “saccadic” and “*Drosophila*-inspired” search – highlight the relevance and general notion of the search problem to other classes of search-related tasks. The decision-based formulation discussed in this chapter offers a common framework for examining these types of problems, ranging from human visual search to efficient event detection in sensor networks to the search ethology exhibited by insects, and perhaps more interestingly, may provide a method for uncovering the mechanisms which govern the behavior of systems in these search tasks.

4.2 Future Directions

Immediate future work for active sensing includes a thorough analysis of the computational complexity of this framework as a function of the number of sensors or targets, as scalability issues are a limitation of many distributed approaches. Other desired results are an understanding of the convergence properties of the gradient-based motion control, in terms of estimation error bounds, as well as the corresponding steady-state sensor configurations for different classes of target motions.

Additionally, further examination of the proposed algorithms can be conducted under relaxations of the assumptions made. For example, the uncertainty profile of the sensor model or the communication link may not be known *a priori*, and hence collected and shared data can be used to determine these profiles in real time. Further, fully connected communication networks, in practice, are impractical; an additional area of future research is how different network topologies can be addressed in terms of robustness, performance and stability of the distributed estimation process [23, 59].

The use of more complex models for both the target and sensor agents is a relatively straightforward extension, and should be investigated. Other target-tracking methods, such as pursuit-evasion concepts, offer approaches where targets actively try to confound the mobile sensors. This scenario

requires that the sensor agents consider the target *behavior* beyond simply a dynamical motion model. Incorporation of these types of ideas may offer new perspectives on target-tracking and other distributed sensing applications.

In the context of sequential decision-theoretic methods for search, there are many avenues for future research. Of immediate interest is the investigation of other measures of performance, including notions of achievable confidence in the choice of limits on the number of observations as well as the role of (possibly changing) error rates α and β on the evolution of the belief. These concepts can be best studied by definition and derivation of the expected value of the belief function, due to the randomness introduced by uncertainty in the detections. Utilizing additional tools from the sequential decision theory of the SPRT may offer avenues for more theoretical developments.

A wide range of interesting problems exist for application of these methods to mobile search teams, including the use of the decision framework for examining target identification, switching task allocation (e.g., whether to continue the search or to perform the rescue), and coordination of a team of sensors which may fuse decisions rather than use sensor-based estimates. Another particularly intriguing extension to this work is the search for antagonistic targets, which maneuver to hinder the progress of the search. Game-theoretic notions, such as *levels of rationality* (which are related to look ahead strategies), are certainly relevant as applied to search games [65, 68, 80], and thus merit further research into these areas.

A final promising avenue of future work lies in the extension of the search models to more realistic human visual search scenarios. The properties of foveated vision [81, 82], where observations occur in a spatially diffuse manner rather than only at the location of the focus of attention [76, 83], can be captured in the search framework presented in this thesis in a straightforward manner. Extension of this framework may provide a viable model to enable further study of visual search systems, as well as potentially illuminate the relationship between top-down (i.e., task-driven) and bottom-up (i.e., feature-driven) attentional mechanisms [77, 78].

Appendix A

Derivations of Matrix Gradients

The expressions for the gradients of the cost function, as given by Propositions 1, 2, and 3, are derived in this section.

Recall the matrix calculus identities

$$\begin{aligned} \frac{\partial}{\partial z} h(\mathbf{A}(z)) &= \text{tr} \left[\frac{\partial h}{\partial \mathbf{A}} \frac{\partial \mathbf{A}}{\partial z} \right], & \text{(matrix chain rule),} \\ \frac{\partial}{\partial \mathbf{A}} \det(\mathbf{A}) &= |\mathbf{A}| \mathbf{A}^{-T} = |\mathbf{A}| \mathbf{A}^{-1}, & \text{(derivative of the determinant),} \\ \frac{\partial}{\partial z} \mathbf{A}^{-1} &= -\mathbf{A}^{-1} \left(\frac{\partial \mathbf{A}}{\partial z} \right) \mathbf{A}^{-1}, & \text{(derivative of the inverse),} \end{aligned}$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric, positive-definite matrix, $h : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ is a real-valued matrix function, and $z \in \mathbb{R}$ is a scalar variable.

A.1 Gradient expression for J_1

The cost J_1 represents the sensing cost for distributed estimation of a single dynamic process using the fusion of local observations, as described in Section 2.2.1. The gradient of J_1 is given by:

$$\frac{\partial J_1}{\partial z_i^l} = |\mathbf{P}_{fused}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused} \right], \quad \mathbf{\Pi} \triangleq \mathbf{T}_i \mathbf{R}_i^{-1} \mathbf{T}_i^T.$$

Proof: The above result follows from applying Equation 2.3.1 (matrix chain rule) to find

$$\frac{\partial J_1}{\partial z_i^l} = \frac{\partial}{\partial z_i^l} |\mathbf{P}_{fused}| = \text{tr} \left[\frac{\partial |\mathbf{P}_{fused}|}{\partial \mathbf{P}_{fused}} \frac{\partial \mathbf{P}_{fused}}{\partial z_i^l} \right],$$

where, by the derivative of the determinant, we get

$$\frac{\partial |\mathbf{P}_{fused}|}{\partial \mathbf{P}_{fused}} = |\mathbf{P}_{fused}| \mathbf{P}_{fused}^{-1},$$

and

$$\begin{aligned}
\frac{\partial \mathbf{P}_{fused}}{\partial z_i^l} &= \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M (\mathbf{T}_{i'} \mathbf{R}_{i'} \mathbf{T}_{i'}^T)^{-1} \right)^{-1}, \\
&= \mathbf{P}_{fused} \left(\frac{\partial}{\partial z_i^l} \sum_{i'}^M (\mathbf{T}_{i'} \mathbf{R}_{i'} \mathbf{T}_{i'}^T)^{-1} \right) \mathbf{P}_{fused}, \\
&= \mathbf{P}_{fused} \left(\sum_{i'}^M \frac{\partial}{\partial z_i^l} (\mathbf{T}_{i'} \mathbf{R}_{i'} \mathbf{T}_{i'}^T)^{-1} \right) \mathbf{P}_{fused}, \\
&= \mathbf{P}_{fused} \left(\sum_{i'}^M \mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_{i'} \mathbf{R}_{i'} \mathbf{T}_{i'}^T) \mathbf{\Pi}^T \right) \mathbf{P}_{fused}, \\
&= \mathbf{P}_{fused} \mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused},
\end{aligned}$$

where the summation is eliminated by the fact that $\frac{\partial \mathbf{R}_{i'}}{\partial z_i^l} = 0, \forall i \neq i'$.

Substitution of these terms yields

$$\begin{aligned}
\frac{\partial J_1}{\partial z_i^l} &= \text{tr} \left[\left(|\mathbf{P}_{fused}| \mathbf{P}_{fused}^{-1} \right) \left(\mathbf{P}_{fused} \left(\mathbf{\Pi} \frac{\partial}{\partial z_i^l} \mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T \right) \mathbf{\Pi}^T \mathbf{P}_{fused} \right) \right], \\
&= \text{tr} \left[|\mathbf{P}_{fused}| \mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused} \right], \\
&= |\mathbf{P}_{fused}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused} \right].
\end{aligned}$$

This completes the proof. ■

A.2 Gradient expression for J_2

The cost J_2 represents the sensing cost for distributed estimation of a single dynamic process using the fusion of locally filtered estimates, as described in Section 2.2.2. The gradient of J_2 is given by:

$$\frac{\partial J_2}{\partial z_i^l} = |\mathbf{P}_{fused}| \text{tr} \left[\mathbf{\Pi} \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{\Pi}^T \mathbf{P}_{fused} \right], \quad \mathbf{\Pi} \triangleq \mathbf{P}_i^{-1} \mathbf{K}_i.$$

Proof: The proof proceeds as done previously, such that application of the matrix chain rule and the derivative of the determinant yields

$$\begin{aligned}
\frac{\partial \mathbf{P}_{fused}}{\partial z_i^l} &= \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M \mathbf{P}_{i'}^{-1} \right)^{-1} = \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M (\boldsymbol{\Delta}_{i'} - \mathbf{K}_{i'} \mathbf{T}_{i'} \mathbf{H}_{i'} \boldsymbol{\Delta}_{i'})^{-1} \right)^{-1} \\
&= -\mathbf{P}_{fused} \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M \mathbf{P}_{i'}^{-1} \right) \mathbf{P}_{fused} \\
&= -\mathbf{P}_{fused} \left(\sum_{i'}^M \frac{\partial}{\partial z_i^l} \mathbf{P}_{i'}^{-1} \right) \mathbf{P}_{fused} \\
&= \mathbf{P}_{fused} \left(\sum_{i'}^M \mathbf{P}_{i'}^{-1} \frac{\partial \mathbf{P}_{i'}^{-1}}{\partial z_i^l} \mathbf{P}_{i'}^{-1} \right) \mathbf{P}_{fused} \\
&= \mathbf{P}_{fused} \left(\sum_{i'}^M \mathbf{P}_{i'}^{-1} \frac{\partial}{\partial z_i^l} (\boldsymbol{\Delta}_{i'} - \mathbf{K}_{i'} \mathbf{T}_{i'} \mathbf{H}_{i'} \boldsymbol{\Delta}_{i'}) \mathbf{P}_{i'}^{-1} \right) \mathbf{P}_{fused} \\
&= -\mathbf{P}_{fused} \left(\sum_{i'}^M \mathbf{P}_{i'}^{-1} \frac{\partial \mathbf{K}_{i'}}{\partial z_i^l} (\mathbf{T}_{i'} \mathbf{H}_{i'} \boldsymbol{\Delta}_{i'}) \mathbf{P}_{i'}^{-1} \right) \mathbf{P}_{fused} \\
&= \mathbf{P}_{fused} \left(\mathbf{P}_l^{-1} \mathbf{K}_l \frac{\partial}{\partial z_i^l} (\mathbf{T}_i \mathbf{R}_i \mathbf{T}_i^T) \mathbf{K}_l^T \mathbf{P}_l^{-1} \right) \mathbf{P}_{fused}
\end{aligned}$$

■

A.3 Gradient expression for $J_{3,j}$

$$\begin{aligned}
\frac{\partial J_{3,l}}{\partial z_i^l} &= |\mathbf{P}_{l,f}| \text{tr} \left[\boldsymbol{\Pi} \frac{\partial \mathbf{R}_l}{\partial z_i^l} \boldsymbol{\Pi}^T \mathbf{P}_{l,f} \right] \\
&\quad + \sum_{i \neq l}^M |\mathbf{P}_{l,f}| \text{tr} \left[\boldsymbol{\Phi}_{l,i} \frac{\partial \mathbf{C}_{l,i}}{\partial z_i^l} \boldsymbol{\Phi}_{l,i} \mathbf{P}_{l,f} \right],
\end{aligned}$$

where $\boldsymbol{\Phi}_{l,i} \triangleq (\mathbf{P}_i + \mathbf{C}_{l,i})$, and $\boldsymbol{\Pi}$ is chosen according to whether simple observations (i.e., $\mathbf{P}_l = \mathbf{R}_l$) or locally filtered measurements (i.e., $\mathbf{P}_l = \boldsymbol{\Delta}_l - \mathbf{K}_l \mathbf{H}_l \boldsymbol{\Delta}_l$) are fused.

Proof: Again, (2.3.1) is used, where

$$\begin{aligned}
\frac{\partial \mathbf{P}_{l,f}}{\partial z_i^l} &= \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M \Phi_{l,i}^{-1} \right)^{-1} = \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M (\mathbf{P}_i + \mathbf{C}_{l,i})^{-1} \right)^{-1} \\
&= -\mathbf{P}_{l,f} \frac{\partial}{\partial z_i^l} \left(\sum_{i'}^M \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f} \\
&= -\mathbf{P}_{l,f} \left(\sum_{i'}^M \frac{\partial}{\partial z_i^l} \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f} \\
&= \mathbf{P}_{l,f} \left(\sum_{i'}^M \Phi_{l,i}^{-1} \frac{\partial \Phi_{l,i}}{\partial z_i^l} \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f} \\
&= \mathbf{P}_{l,f} \left(\sum_{i'}^M \Phi_{l,i}^{-1} \frac{\partial}{\partial z_i^l} (\mathbf{P}_i + \mathbf{C}_{l,i}) \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f} \\
&= \mathbf{P}_{l,f} \left(\Phi_{l,l}^{-1} \frac{\partial}{\partial z_i^l} (\mathbf{P}_l + \mathbf{C}_{l,l}) \Phi_{l,l}^{-1} \right) \mathbf{P}_{l,f} \\
&\quad + \mathbf{P}_{l,f} \left(\sum_{i \neq l}^M \Phi_{l,i}^{-1} \frac{\partial}{\partial z_i^l} (\mathbf{P}_i + \mathbf{C}_{l,i}) \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f} \\
&= \mathbf{P}_{l,f} \left(\mathbf{P}_l^{-1} \frac{\partial \mathbf{P}_l}{\partial z_i^l} \mathbf{P}_l^{-1} \right) \mathbf{P}_{l,f} + \mathbf{P}_{l,f} \left(\sum_{i \neq l}^M \Phi_{l,i}^{-1} \frac{\partial \mathbf{C}_{l,i}}{\partial z_i^l} \Phi_{l,i}^{-1} \right) \mathbf{P}_{l,f}
\end{aligned}$$

■

Appendix B

Analysis of the Cost Function

B.1 Cost Function Expression for Two Sensors

An analytic expression for the cost function, J_1 , though difficult to come by in general, is desirable for better understanding of the cost minimization process as well as the generation of the sensor control laws. Recall that cost function J_1 represents the fusion of local observations, or in other words, the fused covariance matrix is given by:

$$\mathbf{P}_{fused}^{-1} = \sum_i^M \mathbf{R}_i^{-1} = \sum_i^M (\mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T)^{-1}$$

such that the cost function is

$$J_1 \triangleq \det \mathbf{P}_{fused} = \det \left(\sum_i^M \mathbf{R}_i^{-1} \right)^{-1}$$

Consider the case where $M = 2$, i.e., only two sensors are taking measurements. For $\mathbf{P}_{fused} \in \mathbb{R}^{2 \times 2}$, define the two measurement covariance matrices, \mathbf{R}_1 and \mathbf{R}_2 , to be:

$$\mathbf{R}_1 = \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix}, \quad \mathbf{R}_2 = \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix}.$$

Then, taking the inverses of these matrices yields

$$\mathbf{R}_i^{-1} = \frac{1}{a_i d_i - b_i c_i} \begin{pmatrix} d_i & -b_i \\ -c_i & a_i \end{pmatrix},$$

such that the sum of these inverses becomes

$$\begin{aligned} \sum_{i=1}^2 \mathbf{R}_i^{-1} &= \sum_{i=1}^2 \begin{pmatrix} \frac{d_i}{a_i d_i - b_i c_i} & \frac{-b_i}{a_i d_i - b_i c_i} \\ \frac{-c_i}{a_i d_i - b_i c_i} & \frac{a_i}{a_i d_i - b_i c_i} \end{pmatrix} \\ &= \begin{pmatrix} \frac{d_1}{a_1 d_1 - b_1 c_1} + \frac{d_2}{a_2 d_2 - b_2 c_2} & \frac{-b_1}{a_1 d_1 - b_1 c_1} + \frac{-b_2}{a_2 d_2 - b_2 c_2} \\ \frac{-c_1}{a_1 d_1 - b_1 c_1} + \frac{-c_2}{a_2 d_2 - b_2 c_2} & \frac{a_1}{a_1 d_1 - b_1 c_1} + \frac{a_2}{a_2 d_2 - b_2 c_2} \end{pmatrix} \triangleq \begin{pmatrix} \bar{a} & \bar{b} \\ \bar{c} & \bar{d} \end{pmatrix}. \end{aligned}$$

Now, taking the inverse of this expression yields the fused covariance matrix:

$$\mathbf{P}_{fused} = \left(\sum_{i=1}^2 \mathbf{R}_i^{-1} \right)^{-1} = \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}} \begin{pmatrix} \bar{d} & -\bar{b} \\ -\bar{c} & \bar{a} \end{pmatrix} = \begin{pmatrix} \frac{\bar{d}}{\bar{a}\bar{d} - \bar{b}\bar{c}} & \frac{-\bar{b}}{\bar{a}\bar{d} - \bar{b}\bar{c}} \\ \frac{-\bar{c}}{\bar{a}\bar{d} - \bar{b}\bar{c}} & \frac{\bar{a}}{\bar{a}\bar{d} - \bar{b}\bar{c}} \end{pmatrix}.$$

The cost function is simply the determinant of this matrix.

$$J_1 = \det \mathbf{P}_{fused} = \frac{1}{(\bar{a}\bar{d} - \bar{b}\bar{c})^2} (\bar{d}\bar{a} - \bar{b}\bar{c}) = \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}},$$

where the fact that $|\alpha A| = \alpha^n |A|$, $A \in \mathbb{R}^{n \times n}$ is employed.

Substitution of the variables yields

$$\begin{aligned} J_1 &= \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}} \\ &= \frac{1}{\left(\frac{d_1}{a_1 d_1 - b_1 c_1} + \frac{d_2}{a_2 d_2 - b_2 c_2} \right) \left(\frac{a_1}{a_1 d_1 - b_1 c_1} + \frac{a_2}{a_2 d_2 - b_2 c_2} \right) - \left(\frac{b_1}{a_1 d_1 - b_1 c_1} + \frac{b_2}{a_2 d_2 - b_2 c_2} \right) \left(\frac{c_1}{a_1 d_1 - b_1 c_1} + \frac{c_2}{a_2 d_2 - b_2 c_2} \right)}, \\ &= \frac{1}{\left(\frac{d_1}{|\mathbf{R}_1|} + \frac{d_2}{|\mathbf{R}_2|} \right) \left(\frac{a_1}{|\mathbf{R}_1|} + \frac{a_2}{|\mathbf{R}_2|} \right) - \left(\frac{b_1}{|\mathbf{R}_1|} + \frac{b_2}{|\mathbf{R}_2|} \right) \left(\frac{c_1}{|\mathbf{R}_1|} + \frac{c_2}{|\mathbf{R}_2|} \right)}, \\ &= \frac{1}{(d_1 |\mathbf{R}_2| + d_2 |\mathbf{R}_1|) (a_1 |\mathbf{R}_2| + a_2 |\mathbf{R}_1|) - (b_1 |\mathbf{R}_2| + b_2 |\mathbf{R}_1|) (c_1 |\mathbf{R}_2| + c_2 |\mathbf{R}_1|)}, \\ &= \frac{1}{|\mathbf{R}_2|^2 |\mathbf{R}_1| + |\mathbf{R}_1| |\mathbf{R}_2| (a_1 d_2 + a_2 d_1 - b_2 c_1 - b_1 c_2) + |\mathbf{R}_1|^2 |\mathbf{R}_2|}, \\ &= \frac{1}{|\mathbf{R}_2|^2 |\mathbf{R}_1| + |\mathbf{R}_1| |\mathbf{R}_2| (|\mathbf{R}_1 + \mathbf{R}_2| - |\mathbf{R}_1| - |\mathbf{R}_2|) + |\mathbf{R}_1|^2 |\mathbf{R}_2|}, \\ &= \frac{|\mathbf{R}_1| |\mathbf{R}_2|}{|\mathbf{R}_1 + \mathbf{R}_2|}, \end{aligned}$$

where the matrix representation of the determinant has been used.

Now, for the given structure of our problem, the measurement covariance matrix is given by a

diagonal matrix transformed via rotation matrices as below:

$$\mathbf{R}_i \triangleq \mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T = \begin{pmatrix} \cos(\theta_i) & -\sin(\theta_i) \\ \sin(\theta_i) & \cos(\theta_i) \end{pmatrix} \begin{pmatrix} \alpha_i & 0 \\ 0 & \beta_i \end{pmatrix} \begin{pmatrix} \cos(\theta_i) & \sin(\theta_i) \\ -\sin(\theta_i) & \cos(\theta_i) \end{pmatrix}.$$

Expansion of the terms of \mathbf{R}_i gives

$$\begin{aligned} \mathbf{R}_i &= \begin{pmatrix} \alpha_i \cos(\theta_i) & -\beta_i \sin(\theta_i) \\ \alpha_i \sin(\theta_i) & \beta_i \cos(\theta_i) \end{pmatrix} \begin{pmatrix} \cos(\theta_i) & \sin(\theta_i) \\ -\sin(\theta_i) & \cos(\theta_i) \end{pmatrix} \\ &= \begin{pmatrix} \alpha_i \cos^2(\theta_i) + \beta_i \sin^2(\theta_i) & (\alpha_i - \beta_i) \cos(\theta_i) \sin(\theta_i) \\ (\alpha_i - \beta_i) \cos(\theta_i) \sin(\theta_i) & \alpha_i \sin^2(\theta_i) + \beta_i \cos^2(\theta_i) \end{pmatrix}, \\ &\triangleq \begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix}. \end{aligned}$$

Furthermore, the determinant can be found to be

$$|\mathbf{R}_i| = |\mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T| = |\mathbf{T}_i| |\bar{\mathbf{R}}_i| |\mathbf{T}_i^T| = |\bar{\mathbf{R}}_i| = \alpha_i \beta_i,$$

because the determinant of matrix products is the product of determinants and the determinant of a rotation matrix is unity.

So ultimately, the main point of interest is the matrix summation term, $|\mathbf{R}_1 + \mathbf{R}_2|$, in that all other terms are scalar terms that do not contain the dependent variables, θ_i 's. Substitution of terms yields

$$\begin{aligned} \mathbf{R}_1 + \mathbf{R}_2 &= \begin{pmatrix} \alpha_1 \cos^2(\theta_1) + \beta_1 \sin^2(\theta_1) & (\alpha_1 - \beta_1) \cos(\theta_1) \sin(\theta_1) \\ (\alpha_1 - \beta_1) \cos(\theta_1) \sin(\theta_1) & \alpha_1 \sin^2(\theta_1) + \beta_1 \cos^2(\theta_1) \end{pmatrix} \\ &\quad + \begin{pmatrix} \alpha_2 \cos^2(\theta_2) + \beta_2 \sin^2(\theta_2) & (\alpha_2 - \beta_2) \cos(\theta_2) \sin(\theta_2) \\ (\alpha_2 - \beta_2) \cos(\theta_2) \sin(\theta_2) & \alpha_2 \sin^2(\theta_2) + \beta_2 \cos^2(\theta_2) \end{pmatrix}, \end{aligned}$$

Investigation of each element of this matrix will hopefully yield simplification of the expression.

Consider the (1,1) component of the matrix:

$$\begin{aligned} R_{11} &= \alpha_1 \cos^2(\theta_1) + \beta_1 \sin^2(\theta_1) + \alpha_2 \cos^2(\theta_2) + \beta_2 \sin^2(\theta_2), \\ &= \frac{1}{2}(\alpha_1 + \beta_1 + \alpha_2 + \beta_2) + \frac{1}{2}(\alpha_1 - \beta_1) \cos(2\theta_1) + \frac{1}{2}(\alpha_2 - \beta_2) \cos(2\theta_2), \end{aligned}$$

and the (2, 2) component

$$\begin{aligned} R_{22} &= \alpha_1 \sin^2(\theta_1) + \beta_1 \cos^2(\theta_1) + \alpha_2 \sin^2(\theta_2) + \beta_2 \cos^2(\theta_2), \\ &= \frac{1}{2}(\alpha_1 + \beta_1 + \alpha_2 + \beta_2) + \frac{1}{2}(\beta_1 - \alpha_1) \cos(2\theta_1) + \frac{1}{2}(\beta_2 - \alpha_2) \cos(2\theta_2), \end{aligned}$$

where the trigonometric identities, $\sin^2(\theta) = \frac{1-\cos(2\theta)}{2}$ and $\cos^2(\theta) = \frac{1+\cos(2\theta)}{2}$, have been used.

Computation of the determinant yields

$$\begin{aligned} |\mathbf{R}_1 + \mathbf{R}_2| &= \left(\frac{1}{2}(\alpha_1 + \beta_1 + \alpha_2 + \beta_2) + \frac{1}{2}(\alpha_1 - \beta_1) \cos(2\theta_1) + \frac{1}{2}(\alpha_2 - \beta_2) \cos(2\theta_2) \right) \times \\ &\quad \left(\frac{1}{2}(\alpha_1 + \beta_1 + \alpha_2 + \beta_2) + \frac{1}{2}(\beta_1 - \alpha_1) \cos(2\theta_1) + \frac{1}{2}(\beta_2 - \alpha_2) \cos(2\theta_2) \right) \\ &\quad - ((\alpha_1 - \beta_1) \cos(\theta_1) \sin(\theta_1) + (\alpha_2 - \beta_2) \cos(\theta_2) \sin(\theta_2))^2, \\ &= \frac{1}{2}(\alpha_2\beta_1 + 2\alpha_2\beta_2 + \beta_1\beta_2 + \alpha_1\alpha_2 + 2\alpha_1\beta_1 + \alpha_1\beta_2) \\ &\quad - \frac{1}{2}(\alpha_1 - \beta_1)(\alpha_2 - \beta_2) \cos(2(\theta_1 - \theta_2)), \\ &= \rho_1 - \rho_2 \cos(2(\theta_1 - \theta_2)). \end{aligned}$$

Thus, in order to **minimize** the cost function, the quantity $|\mathbf{R}_1 + \mathbf{R}_2|$ needs to be **maximize**, which is done for $\gamma_{12} \triangleq \theta_1 - \theta_2 = \pm \frac{\pi}{2}$.

In this way, parameterizing the cost function by γ_{12} yields

$$J_1 = \det \mathbf{P}_{fused} = \frac{|\mathbf{R}_1||\mathbf{R}_2|}{|\mathbf{R}_1 + \mathbf{R}_2|} = \frac{\alpha_1\beta_1\alpha_2\beta_2}{\rho_{02} - \rho_{12} \cos(2\gamma_{12})},$$

where

$$\begin{aligned} \rho_{02} &\triangleq \frac{1}{2}((\alpha_1 - \beta_1)(\alpha_2 - \beta_2) + 2(\alpha_1 + \alpha_2)(\beta_1 + \beta_2)), \\ \rho_{12} &\triangleq \frac{1}{2}(\alpha_1 - \beta_1)(\alpha_2 - \beta_2). \end{aligned}$$

This dependence on the relative bearing of the sensors is illustrated in Figure B.1 for homogeneous sensors (i.e., $\alpha_i = 1, \beta_i = 5, \forall i$).

B.2 Cost Function Expression for Three Sensors

A similar analysis can be conducted for $M = 3$ sensors in the plane, such that

$$\begin{pmatrix} \bar{a} & \bar{b} \\ \bar{c} & \bar{d} \end{pmatrix} \triangleq \begin{pmatrix} \sum_{i=1}^M \frac{d_i}{|\mathbf{R}_i|} & \sum_{i=1}^M \frac{b_i}{|\mathbf{R}_i|} \\ \sum_{i=1}^M \frac{c_i}{|\mathbf{R}_i|} & \sum_{i=1}^M \frac{a_i}{|\mathbf{R}_i|} \end{pmatrix}$$




Figure B.1: Dependence of the cost function (for two sensors) on their relative bearings with respect to the target. This shows that in order to minimize the cost, the two agents should be separated by $\pm \frac{\pi}{2}$ radians.

As before, the fused covariance matrix is given by

$$\mathbf{P}_{fused} = \left(\sum_{i=1}^M \mathbf{R}_i^{-1} \right)^{-1} = \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}} \begin{pmatrix} \bar{d} & -\bar{b} \\ -\bar{c} & \bar{a} \end{pmatrix} = \begin{pmatrix} \frac{\bar{d}}{\bar{a}\bar{d} - \bar{b}\bar{c}} & \frac{-\bar{b}}{\bar{a}\bar{d} - \bar{b}\bar{c}} \\ \frac{-\bar{c}}{\bar{a}\bar{d} - \bar{b}\bar{c}} & \frac{\bar{a}}{\bar{a}\bar{d} - \bar{b}\bar{c}} \end{pmatrix},$$

and the cost function is simply the determinant of this matrix.

$$J_1 = \det \mathbf{P}_{fused} = \frac{1}{(\bar{a}\bar{d} - \bar{b}\bar{c})^2} (\bar{d}\bar{a} - \bar{b}\bar{c}) = \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}},$$

employing the fact that $|\alpha A| = \alpha^n |A|$, $A \in \mathbb{R}^{n \times n}$.

Substitution of the variables yields

$$\begin{aligned} J_1 &= \frac{1}{\bar{a}\bar{d} - \bar{b}\bar{c}} \\ &= \frac{1}{\left(\sum_{i=1}^3 \frac{d_i}{|\mathbf{R}_i|} \right) \left(\sum_{i=1}^3 \frac{a_i}{|\mathbf{R}_i|} \right) - \left(\sum_{i=1}^3 \frac{b_i}{|\mathbf{R}_i|} \right) \left(\sum_{i=1}^3 \frac{c_i}{|\mathbf{R}_i|} \right)}, \\ &= \frac{1}{\left(\sum_{i=1}^3 \sum_{j=1}^3 \frac{d_i a_j}{|\mathbf{R}_i| |\mathbf{R}_j|} \right) - \left(\sum_{i=1}^3 \sum_{j=1}^3 \frac{b_i c_j}{|\mathbf{R}_i| |\mathbf{R}_j|} \right)}, \\ &= \frac{|\mathbf{R}_1| |\mathbf{R}_2| |\mathbf{R}_3|}{|\mathbf{R}_3| |\mathbf{R}_1 + \mathbf{R}_2| - |\mathbf{R}_1| |\mathbf{R}_2| + |\mathbf{R}_2| |\mathbf{R}_1 + \mathbf{R}_3| - |\mathbf{R}_1| |\mathbf{R}_3| \\ &\quad + |\mathbf{R}_1| |\mathbf{R}_2 + \mathbf{R}_3| - |\mathbf{R}_2| |\mathbf{R}_3|}. \end{aligned}$$

Similar expansion as done in the previous section shows that minimization of the cost function

requires maximization of the denominator term, which becomes

$$\begin{aligned} & |\mathbf{R}_3| |\mathbf{R}_1 + \mathbf{R}_2| - |\mathbf{R}_1| |\mathbf{R}_2| + |\mathbf{R}_2| |\mathbf{R}_1 + \mathbf{R}_3| - |\mathbf{R}_1| |\mathbf{R}_3| + |\mathbf{R}_1| |\mathbf{R}_2 + \mathbf{R}_3| - |\mathbf{R}_2| |\mathbf{R}_3| \\ &= \rho_{03} - \rho_{12} \cos(2\gamma_{12}) - \rho_{23} \cos(2\gamma_{23}) - \rho_{13} \cos(2(\gamma_{12} + \gamma_{23})), \end{aligned}$$

where the constant coefficient terms can be shown to be

$$\begin{aligned} \rho_{03} &= \left(\frac{1}{2} \alpha_1 \alpha_2 \alpha_3 \beta_1 + \frac{1}{2} \alpha_1 \alpha_2 \alpha_3 \beta_2 + 2 \alpha_1 \alpha_2 \beta_1 \beta_2 + \frac{1}{2} \alpha_1 \alpha_3 \beta_1 \beta_2 \right. \\ &\quad + \frac{1}{2} \alpha_2 \alpha_3 \beta_1 \beta_2 + \frac{1}{2} \alpha_1 \alpha_2 \alpha_3 \beta_3 + \frac{1}{2} \alpha_1 \alpha_2 \beta_1 \beta_3 + 2 \alpha_1 \alpha_3 \beta_1 \beta_3 \\ &\quad + \frac{1}{2} \alpha_2 \alpha_3 \beta_1 \beta_3 + \frac{1}{2} \alpha_1 \alpha_2 \beta_2 \beta_3 + \frac{1}{2} \alpha_1 \alpha_3 \beta_2 \beta_3 + 2 \alpha_2 \alpha_3 \beta_2 \beta_3 \\ &\quad \left. + \frac{1}{2} \alpha_1 \beta_1 \beta_2 \beta_3 + \frac{1}{2} \alpha_2 \beta_1 \beta_2 \beta_3 + \frac{1}{2} \alpha_3 \beta_1 \beta_2 \beta_3 \right), \\ \rho_{12} &= \frac{1}{2} \alpha_3 \beta_3 (\alpha_1 - \beta_1) (\alpha_2 - \beta_2), \\ \rho_{23} &= \frac{1}{2} \alpha_1 \beta_1 (\alpha_2 - \beta_2) (\alpha_3 - \beta_3), \\ \rho_{13} &= \frac{1}{2} \alpha_2 \beta_2 (\alpha_1 - \beta_1) (\alpha_3 - \beta_3). \end{aligned}$$

The maximal values of denominator term occurs at $(\gamma_{12}, \gamma_{23})$ pairs resulting in separations of $\pm \frac{\pi}{3}$ (or equivalently, $\pm \frac{2\pi}{3}$) radians, as illustrated in Figure B.2.




Figure B.2: Dependence of the cost function (for three sensors) on their relative bearings with respect to the target. Minimization of cost requires that agents be separated by $\pm \frac{\pi}{3}$ (or by symmetry, $\pm \frac{2\pi}{3}$) radians.

Bibliography

- [1] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, “Active vision,” *Int. J. Computer Vision*, vol. 1, pp. 333–356, 1988.
- [2] P. Mowforth, “Active sensing for mobile robots,” in *Proc. Intl. Conf. on Control*, vol. 2, 1991, pp. 1141–1146.
- [3] S. Hutchinson, G. Hager, and P. Corke, “A tutorial on visual servo control,” *Robotics and Automation, IEEE Transactions on*, vol. 12, no. 5, pp. 651–670, Oct 1996.
- [4] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. D. Schutter, “Active sensing for robotics - a survey,” in *Proc. of the 5th Intl. Conf. on Numerical Methods and Applications*, Borovets, Bulgaria, August 2002, pp. 316–324.
- [5] J. S. Bay, “A fully autonomous active sensor-based exploration concept for shape-sensing robots,” *IEEE Trans. Systems, Man, and Cybernetics*, vol. 21, no. 4, pp. 850–860, July/August 1991.
- [6] F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte, “Information based adaptive robotic exploration,” in *Proc. Of 2002 Intl. Conf. On Intelligent Robots and Systems*, vol. 1, Lausanne, Switzerland, 30 Sept.-5 Oct. 2002 2002, pp. 540 –545.
- [7] F. Lorussi, A. Marigo, and A. Bicchi, “Optimal exploratory paths for a mobile rover,” in *Proc. of IEEE Conf. on Robotics and Automation*, 2001, pp. 2078–2083.
- [8] D. Ucinski, *Measurement Optimization for Parameter Estimation in Distributed Systems*. Technical University Press, 1999.
- [9] —, *Optimal Measurement Methods for Distributed Parameter System Identification*. Boca Raton: CRC Press, 2005.
- [10] L. E. Navarro-Serment, J. M. Dolan, and P. K. Khosla, “Optimal sensor placement for cooperative distributed vision,” in *Proc. of the 2004 IEEE Conf. on Robotics and Automation*, New Orleans, LA, USA, April 2004, pp. 939–944.

- [11] A. Arbel, "Sensor placement in optimal filtering and smoothing problems," *IEEE Trans. Automatic Control*, vol. 27, no. 1, pp. 94–98, Feb 1982.
- [12] Y. Bar-Shalom and T. E. Fortman, *Tracking and Data Association*. Academic Press, 1988.
- [13] Y. Bar-Shalom and X. R. Li, *Estimation and Tracking: Principles, Techniques, and Software*. Norwood, MA: Artech House, Inc, 1993.
- [14] B. Rao, H. Durrant-Whyte, and J. Sheen, "A fully decentralized multi-sensor system for tracking and surveillance," *Intl. Journal of Robotics Research*, vol. 12, pp. 20–44, 1993.
- [15] J. R. Spletzer and C. J. Taylor, "Dynamic sensor planning and control for optimally tracking targets," *International Journal of Robotics Research*, vol. 22, no. 1, pp. 7–20, January 2003.
- [16] A. Logothetis, A. Isaksson, and R. Evans, "An information theoretic approach to observer path design for bearings-only tracking," in *Proc. of the 36th Conference on Decision and Control*. San Diego, CA, USA: CDC, December 1997, pp. 3132–3137.
- [17] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, U. Sydney, 2002.
- [18] B. Rao and H. Durrant-Whyte, "A decentralized bayesian algorithm for identification of tracked targets," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 23, no. 6, pp. 1683–1698, Nov/Dec 1993.
- [19] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey of sensor networks," *IEEE Communications Magazine*, vol. 38, pp. 102–114, August 2002.
- [20] C. Y. Chong and S. P. Kumar, "Sensor networks: Evolution, opportunities, and challenges," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1247–1256, August 2003.
- [21] Y. Mostofi and R. Murray, "Distributed sensing and estimation under communication constraints," in *Proceedings of 45th IEEE Conference on Decision and Control (CDC)*, San Diego, CA, Dec 2006.
- [22] Y. Mostofi, T. H. Chung, R. M. Murray, and J. W. Burdick, "Communication and sensing trade-offs in decentralized mobile sensor networks: A cross-layer design approach," in *Intl. Conf. on Information Processing in Sensor Networks*, 2005.
- [23] D. Spanos, R. Olfati-Saber, and R. Murray, "Distributed sensor fusion using dynamic consensus," in *Proceedings of IFAC*, 2005.
- [24] W. B. Dunbar and R. M. Murray, "Distributed receding horizon control for multi-vehicle formation stabilization," *Automatica*, vol. 42, no. 4, pp. 549–558, 2006.

- [25] R. R. III and J. E. Hurtado, “Stability and control of collective systems,” *Journal of Intelligent and Robotic Systems*, vol. 39, pp. 43–55, 2004.
- [26] N. E. Leonard and E. Fiorelli, “Virtual leaders, artificial potentials and coordinated control of groups,” in *Proc. 40th IEEE Conf. Decision and Control*, 2001, pp. 2968–2973.
- [27] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, ser. Artificial Intelligence. Prentice-Hall, 1995.
- [28] B. O. Koopman, “The theory of search, part I. kinematic bases.” *Operations Research*, vol. 4, pp. 324–346, 1956.
- [29] —, “The theory of search, part II. target detection.” *Operations Research*, vol. 4, pp. 503–531, 1956.
- [30] —, “The theory of search, part III. the optimum distribution of searching effort.” *Operations Research*, vol. 5, pp. 613–626, 1957.
- [31] —, “Search and its optimization,” *The American Mathematical Monthly*, vol. 86, no. 7, pp. 527–540, Aug-Sept 1979.
- [32] J. R. Frost and L. D. Stone, “Review of search theory: Advances and applications to search and rescue decision support,” U.S. Coast Guard Research and Development Center, Tech. Rep. CG-D-15-01, September 2001.
- [33] A. R. Washburn, *Search and Detection*, 4th ed., ser. Topics in Operations Research Series. INFORMS, 2002.
- [34] L. D. Stone, *Theory of Optimal Search*, 2nd ed., J. Kettelle, Ed. Academic Press, 1989.
- [35] S. J. Benkoski, M. G. Monticino, and J. R. Weisinger, “A survey of the search theory literature,” *Naval Research Logistics*, vol. 38, pp. 469–494, 1991.
- [36] L. C. Thomas, “Optimization and search,” in *Proc. of Intl. Conf. on Systems, Man and Cybernetics*, 1993.
- [37] F. Bourgault, T. Furukawa, and H. Durrant-Whyte, *Field and Service Robots IV*, ser. Springer Tracts in Advanced Robotics (STAR). Springer-Verlag, 2006, vol. 24, ch. Optimal Search for a Lost Target in a Bayesian World, pp. 209–222.
- [38] E. M. Wong, F. Bourgault, and T. Furukawa, “Multi-vehicle bayesian search for multiple lost targets,” in *Proc. of Intl. Conf. on Robotics and Automation*, April 2005, pp. 3169– 3174.
- [39] Z. Chen, “Bayesian filtering: From kalman filters to particle filters, and beyond,” McMaster University, Tech. Rep.

- [40] S. Thrun, W. Burgard, , and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [41] A. Wald, *Sequential Analysis*. New York, NY: John Wiley and Sons, Inc., 1947.
- [42] A. Wald and J. Wolfowitz, “Bayes solutions of sequential decision problems,” *PNAS*, vol. 35, pp. 99–102, 1949.
- [43] D. A. Castanon, “Optimal search strategies in dynamic hypothesis testing,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 25, pp. 1130–1138, 1995.
- [44] K. G. Götz, *Neural Basis of Behavioural Adaptations*. Gustav Fischer Verlag, 1994, vol. 39, ch. Exploratory Strategies in *Drosophila*, pp. 47–59.
- [45] J. Najemnik and W. S. Geisler, “Optimal eye movement strategies in visual search,” *Nature*, vol. 434, pp. 387–391, 2005.
- [46] T. H. Chung, V. Gupta, J. W. Burdick, and R. M. Murray, “On a decentralized active sensing strategy using mobile sensor platforms in a network,” in *Proc. of the IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, Dec 2004.
- [47] T. H. Chung, J. W. Burdick, and R. M. Murray, “Decentralized motion control of mobile sensing agents in a network,” in *Intl. Conference on Robotics and Automation*. ICRA, May 2006.
- [48] K. Ramachandra, *Kalman Filtering Techniques for Radar Tracking*. New York, NY: Marcel Dekker, Inc., 2000.
- [49] P. Maybeck, *Stochastic Models, Estimation and Control*. Academic Press, 1979, vol. 1.
- [50] S. Julier and J. Uhlmann, “A non-divergent estimation algorithm in the presence of unknown correlations,” in *Proc. of 1997 American Control Conference*, vol. 4, Albuquerque, NM, USA, June 1997, pp. 2369–2373.
- [51] A. Stroupe, M. C. Martin, and T. Balch, “Merging probabilistic observations for mobile distributed sensing,” Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-00-30, 2000.
- [52] Y. Bar-Shalom, “On the track-to-track correlation problem,” *IEEE Trans. on Automatic Control*, vol. 26, no. 2, pp. 571–572, April 1981.
- [53] Y. Bar-Shalom and L. Campo, “The effect of the common process noise on the two-sensor fused-track covariance,” *IEEE Trans. on Aerospace and Electronic Systems*, vol. 22, no. 6, pp. 803–805, 1986.
- [54] R. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

- [55] G. Welch and G. Bishop, "Introduction to the kalman filter," University of North Carolina at Chapel Hill, Tech. Rep. TR 95-041, 2001.
- [56] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation : Numerical Methods*. Prentice Hall, 1989.
- [57] D. S. Bernstein, *Matrix Mathematics*. Princeton, NJ: Princeton University Press, 2005.
- [58] P. Ogren, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 8, pp. 1292–1302, August 2004.
- [59] A. Fax, "Optimal and cooperative control of vehicle formations," Ph.D. dissertation, California Institute of Technology, November 2002.
- [60] R. Olfati-Saber and R. M. Murray, "Distributed cooperative control of multiple vehicle formations using structural potential functions," in *15th IFAC World Congress*, Barcelona, Spain, July 2002.
- [61] J. D. Gibson, *The Mobile Communications Handbook*, 2nd ed. CRC Press and IEEE Press, 1999.
- [62] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. John Wiley & Sons, Inc., 2001.
- [63] R. D. Hippenstiel, *Detection Theory: Applications and Digital Signal Processing*. CRC Press, 2002, §4.9.
- [64] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, H. Booker and N. DeClaris, Eds. Academic Press, Inc., 1972, § 3.2.
- [65] S. Gal, *Search Games*. Academic Press New York, 1980.
- [66] V. Gupta, T. H. Chung, B. Hassibi, and R. M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and dynamic sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, February 2006.
- [67] M. I. Posner and S. E. Petersen, "The attention system of the human brain," *Annu. Rev. Neurosci.*, vol. 13, pp. 25–42, 1990.
- [68] R. Vidal, O. Shakernia, H. Kim, D. Shim, and S. Sastry, "Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 662–669, Oct 2002.

- [69] E. U. Acar, H. Choset, Y. Zhang, and M. Schervish, “Path planning for robotic demining: Robust sensor-based coverage of unstructured environments and probabilistic methods,” *The International Journal of Robotics Research*, vol. 22, no. 7-8, pp. 441–466, July-August 2003.
- [70] D. Bertsekas and S. Shreve, *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, 1978.
- [71] K. E. Trummel and J. R. Weisinger, “The complexity of the optimal searcher path problem,” *Operations Research*, vol. 34, no. 2, pp. 324–327, Mar. - Apr. 1986.
- [72] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, pp. 789–814, 2000.
- [73] M. N. Shadlen and J. I. Gold, *The Cognitive Neurosciences*. MIT Press, 2004, ch. The Neurophysiology of Decision-Making as a Window on Cognition.
- [74] K. Götz, *Neurobiology of Sensory Systems*. New York: Plenum Press, 1989, ch. Search and choice in *Drosophila*, pp. 139–153.
- [75] A. A. Robie and M. H. Dickinson, “Walking *Drosophila* demonstrate shape preference during local search but not global search.” Society for Neuroscience, November 2005.
- [76] L. Itti and C. Koch, “A saliency-based search mechanism for overt and covert shifts of visual attention,” *Vision Research*, vol. 40, pp. 1489–1506, 2000.
- [77] V. Navalpakkam and L. Itti, “Search goal tunes visual features optimally,” *Neuron*, vol. 53, pp. 605–617, February 2007.
- [78] A. Torralba, A. Oliva, M. S. Castelhana, and J. M. Henderson, “Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search,” *Psychological Review*, vol. 113, pp. 766–786, 2006.
- [79] J. Wolfe, *Attention*. University College London Press, 1998, ch. Visual search, pp. 13–73.
- [80] J. Hespanha, M. Prandini, and S. Sastry, “Probabilistic pursuit-evasion games: A one-step nash approach,” in *Proc. Of the 39th Conf. On Decision and Control*, vol. 3, 2000, pp. 2272–2277.
- [81] M. Bolduc and M. Levine, “A review of biologically motivated space-variant data reduction models for robotic vision,” *Computer vision and image understanding*, vol. 69, no. 2, pp. 170–184, 1998.
- [82] G. Sandini and V. Tagliasco, “An anthropomorphic retina-like structure for scene analysis,” *Computer Vision, Graphics and Image Processing*, vol. 14, no. 3, pp. 365–372, 1980.

- [83] V. Isler, S. Khanna, J. Spletzer, and C. Taylor, "Target tracking with distributed sensors: The focus of attention problem," *Computer Vision and Image Understanding Journal*, vol. Special Issue on Attention and Performance in Computer Vision, p. accepted, 2005, accepted.