University of Melbourne

Department of Electrical and Electronic Engineering

Submitted in partial fulfilment of the degree of Doctor of Philosophy

Efficient Algorithms for Autonomous Agents Facing Uncertainty

Author:
Daniel D. Selvaratnam
https://orcid.org/
0000-0003-3329-436X

Supervisors:
Dr. Iman Shames
Prof. Jonathan H. Manton
Dr. Branko Ristic

January 5, 2019

Abstract

This thesis considers the design and mathematical analysis of algorithms enabling autonomous agents to operate reliably in the presence of uncertainty. The algorithms are designed to preserve computational tractability, and to respect communication constraints. Four specific problems are addressed. First, the localisation of a signal source using only random binary measurements. A Bayesian estimation procedure is adopted that discretises the search space to achieve tractability. The effect of this discretisation on convergence is analysed rigorously, as well as the effect of relying on an inexact measurement model. Measurement locations are also optimised with respect to Fisher Information. In the second, the security of general quantized Bayesian estimators is analysed from the perspective of an adversary that sends false measurements to induce a misleading posterior. Fundamental limits on the set of posteriors that can be induced are derived, along with strategies to induce them. The third problem considers the design of control laws for maintaining reliable communication links between agents as they traverse the environment. Robustness to disturbances is established theoretically. Finally, optimisation problems are tackled involving cost functions and constraints that change unpredictably as new information becomes available. Performance bounds are provided for different classes of cost functions, and both first-order and gradient free methods are examined.

Declaration

This thesis comprises only my original work towards the degree of Doctor of Philosophy, except where indicated in the preface. Its length is fewer than 100 000 words, exclusive of tables, maps, bibliographies and appendices. Due acknowledgement has been made in the text to all other material used.

Daniel D. Selvaratnam

January 5, 2019

Preface

The work presented in this thesis was conducted in collaboration with my supervisors, Iman Shames, Jonathan H. Manton and Branko Ristic. They have also helped to review, edit and critique this thesis. In addition, Sections 1.2, 7.2 and Chapter 3 were produced in collaboration with Farhad Farokhi. I remain the primary author and investigator in all cases. Naomi Selvaratnam has helped to proofread sections of this thesis, but is unfamiliar with the academic disciples relating to its subject matter.

A preliminary version of the results in Chapter 2, Sections 1.1, 7.1, and Appendix A has appeared in

- D. D. Selvaratnam, I. Shames, D. V. Dimarogonas, J. H. Manton, and B. Ristic, "Co-operative Estimation for Source Localisation using Binary Sensors," in the proc. of the *56th IEEE Conference on Decision and Control (CDC)*, Melbourne, VIC, Australia, December 2017.

Material from Sections 1.4, 5.1, 5.2, 5.4, 5.5, 7.4 and Appendix D.1 has also appeared in

- D. D. Selvaratnam, I. Shames, J. H. Manton, M. Zamani, "Numerical Optimisation of Time-Varying Strongly Convex Functions Subject to Time-Varying Constraints", in the proc. of the *57th IEEE Conference on Decision and Control (CDC)*, Fontainebleau, Miami Beach, FL, USA, December 2018.

No portion of the work presented herein was conducted prior to my enrolment in the degree of Doctor of Philosophy at the University of Melbourne. Neither has any of it been submitted for any other qualification.

This research has been financially supported by DST Group under Collaborative Research Agreements MYIP #5922 and MYIP #5923. During my candidature, I was also supported by an RTP Scholarship from the University of Melbourne, and a PhD Support Grant from the Defence Science Institute as an initiative of the State Government of Victoria.

John: 'But I must think it is one or the other.'
Reason: 'By my father's soul, you must not –
until you have some evidence.
Can you not remain in doubt?'

The Pilgrim's Regress
C. S. Lewis

Acknowledgements

First, I am grateful to the Defence Science Institute (DSI), Defence Science and Technology Group (DST), the Australian Government, and the University of Melbourne, all of whom provided financial support during my candidature. Their willingness to support research students has kept me above subsistence level, and assisted with several opportunities for travel.

I was blessed with three remarkable supervisors. Their assistance in reviewing and critiquing all my work has been invaluable. Though rather a handful at times, I could not have asked for a better set of guides. Branko spent many hours with me as a fledgeling PhD student, patiently teaching me the theory that set the stage for all my subsequent research. In the maelstrom that is academia, he has always retained composure, gentleness and humility. I aim to follow his example, in character, technical expertise, and love of craft. Jonathan's insight and depth of understanding never cease to amaze me. He has an exceptional ability to deconstruct a topic, and teach it systematically by asking and answering the right questions. From him I not only learnt mathematics, but how to learn mathematics. Much has been absorbed by observing his approach, and the foundations he laid have enabled me to learn and make discoveries on my own. The impact Iman has had on my academic growth is difficult to overstate. He taught me to navigate the bewildering, isolating and often overwhelming landscape of research. Never one to shrink from an argument, he faithfully delivered forthright advice and challenged my perspectives. He identified the boundary of my reachable set, and (not without resistance) pushed me to converge to it. I am grateful for his commitment to my development, his foresight, and his generosity with both time and funding. This allowed me to travel, collaborate, and expand my horizons. Many wonderful people who became my friends along the way, were introduced to me by him.

Hoam Chung was my control systems lecturer in my final year of undergrad. His enthusiasm was infectious, and his delivery of the subject the reason I

chose to pursue a PhD in control and signal processing. Dimos Dimarogonas hosted my visit to KTH, and made time to discuss ideas which directed my work on connectivity preservation. Michael Pauley helped me grapple with probability theory, geometric mechanics, and guided me through a proof on estimator convergence that had been eluding me for several months. Farhad Farokhi first introduced me to security analysis, and proposed the premise of Chapter 3. He has been a wonderful collaborator, and a supportive friend. During a visit to ANU, Brian Anderson first pointed out to me the potential relationship between my early results on source localisation and KL divergence. His comments led to the stronger, more elegant results presented here. I am also grateful to Michael Cantoni for his insights into systems and control theory, and for the integrous example he sets to his students. Sergey, Andrei, Mark Fabbro, Adair, and Wences have been staunch colleagues. Their friendship has made me a better researcher and wiser person.

Nothing I can write will do justice to my parents. I lived with them for the duration of my PhD. They have been outstanding role models of love, hope, and integrity. They instilled in me the value of education, diligence, and taught me to pursue things of true value. They even read parts of this thesis, and provided practical suggestions to improve it. Mum (trained in medicine) was my first maths teacher, and to this day remains among my top three – she first taught me the field axioms and how to use them. Dad has consistently placed his family before himself, and been a faithful example of courage, endurance and strength. Both of them have sacrificed more for my development and well-being than I will probably ever understand. Though they will be slightly bewildered by the contents of this thesis, it owes its existence to them.

These people have made my PhD a rewarding experience, and a surprisingly spiritual one. I learned from them the value of rigour, sound argument, that truth is agnostic to ones feelings. They taught me humility, to question my intuition, and to mistrust the crowd. To honestly say 'I don't know', but still to seek out what I can know. My attempts to pursue this faithfully led me squarely back to a faith I was seeking to abandon. This may be the most unexpected result of my PhD, but it is for this I am most grateful. My final thanks and adoration therefore goes to God, who has lavished on me His love, and reveals to us glimpses of His truth. Every person mentioned here has borne His image and pointed me back to Him.

D.D.S.

Contents

Chapter 1

Introduction

> "An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators."
>
> *Artificial Intelligence: A Modern Approach*
> Russell and Norvig, 1995

> **autonomous:** *(adjective)* "Denoting or performed by a device capable of operating without direct human control."
>
> *Oxford English Dictionary*, 2018

Human-designed autonomous agents are playing an increasingly prominent role in modern life. Many of the world's foremost technology and automotive firms are racing to develop reliable autonomous vehicles, investing heavily in all forms of promising technology [1]. The political and moral implications of autonomous weapon systems have made them a focal point of public debate [2]. Autonomous 'bots' on social media stand accused of manipulating public opinions to an election-altering extent [3]. Beneficent or diabolical, the success of autonomous systems such as these, hinges on their ability to deal with the unknown. Environmental disturbances are unavoidable; external inputs are unpredictable; and the mathematical models used in design can only approximate reality. Estimation, optimisation and control theory establish a mathematical framework for decision-making under uncertainty, which will be exploited in this thesis. All autonomous agents have limited memory, computational resources and constraints on communication, which can preclude an ideal mathematical solution. Expanding these limits is costly, and it is often more economic to deploy a large team of cheap and dispensable

agents than to rely on a few high-performance individuals. Here, we design algorithms enabling autonomous agents to accomplish specific tasks in the presence of uncertainty. Limits on communication are taken into account, and some approximations made in the interests of computational efficiency. Their effects are analysed rigorously.

Four problems are considered. We begin with the localisation of a signal source using only random binary measurements. A Bayesian estimation procedure is adopted that discretises the search space to achieve tractability. We also examine the effect of relying on an inexact measurement model. The second problem introduces a different source of uncertainty: potentially deceptive sensors. The security of Bayesian estimators is analysed from the perspective of an adversary that seeks to induce a false belief. The third moves to designing control laws for maintaining reliable communication links between agents as they traverse the environment in the presence of disturbances. Optimisation algorithms are ubiquitous in autonomous systems, and our final problem tackles cost functions and constraints that change unpredictably with time as new information becomes available. These problems are introduced in greater depth below.

1.1 Source localisation

Chapter 2 considers the localisation of a stationary source using binary measurements obtained from a team of mobile agents. Source localisation involves estimating the location of a signal source using measurements from a set of available sensors. Regardless of its type, a single bit constitutes the minimum amount of information that can be extracted from a signal. For example, when dealing with chemical or radiological sources, this may correspond to detecting the presence or absence of particles of interest [4, 5]. In other situations, sensors may be required to process raw measurement data locally and report a binary outcome to a fusion centre [6, 7]. In general, any continuous or discrete-valued signal can be converted into a binary one via the use of a threshold. This is often desirable for applications with limited resources, because binary data demands less memory, communication bandwidth, and energy from the agents involved [8].

Consistent with standard practice, the measurements are treated as random variables taking values in $\{0, 1\}$, to capture the effects of sensor noise and environmental uncertainty. Given suitable models for the signal propagation and sensors, the probability of obtaining a detection (i.e. of measuring a 1)

becomes a well-defined function of the source and agent locations. Specific examples of such functions constructed for different types of signals and sensors can be found in [4, 8–11], as well as Section 2.6 of this thesis. A novel aspect of our analysis is that it assumes an arbitrary probability-of-detection function, subject to mild conditions. Thus, the algorithms and results of Chapter 2 apply to a large class of measurement models and localisation scenarios. The probability-of-detection function is initially assumed to be fully known, but this is later relaxed by analysing the performance of the algorithm when only an envelope for this function is known. Background false detection rates and missed detection probabilities are incorporated naturally into the framework.

Both Bayesian and classical parameter estimation techniques have been applied to solve the source localisation problem. Here we adopt the former, which has the advantage of incorporating prior knowledge about the source location, and of maintaining an entire posterior probability distribution rather than just a single estimate. Furthermore, a Bayesian framework permits the recursive addition of new measurements to update the posterior, without reprocessing past measurements. A disadvantage of this approach is that every iteration requires the computation of integrals that, in general, have no analytic solution. We obtain a tractable approximation by discretising the exploration region, thereby replacing the integrals with sums and generating a discrete posterior instead of a continuous one. This technique is well-known, however we believe the accompanying analysis to be novel. In particular, we explicitly consider the effects of finite discretisation by identifying points at which the discrete posterior is guaranteed to vanish asymptotically, and establish a relationship between the decay and Kullback-Leibler (KL) divergence. The analysis leads directly to conditions on measurement locations that guarantee sufficient information is being extracted by the agents. We then extend this by choosing measurement locations to maximise the determinant of the Fisher Information Matrix (FIM) [12, Section 4.3.3.1]. This is a widely adopted performance criterion known as *D-optimality* [13]. Focusing on the case where the probability of detection depends solely on distance, the resulting D-optimal geometries mirror the results of [14] for range-only sensors. We then show how the knowledge of these geometries can be exploited via a control strategy by guiding the agents into formation about an estimated source location.

Importance sampling is an alternative Bayesian technique which uses random sampling to numerically evaluate the required integrals [15, Chapter 14], [16]. The posterior is approximated by a weighted set of samples or *particles*, and well known results prove the convergence of this approximation to the true

posterior as the number of particles approaches infinity. In practice however, only a finite number of particles can ever be used. In Section 2.2.2, we establish that importance sampling with a finite number of particles can be treated as a special case of discretising the exploration region, under the appropriate choice of discretisation points and prior. Our approach is more general, because it allows the discretisation points to be chosen arbitrarily. Another key difference is that we analyse convergence over time, using only a finite number of particles. Particle filtering [16] extends important sampling to estimate a time-varying state based on an assumed dynamic model. Particle filtering does not fit within the framework of Chapter 2, because it requires the particles to be propagated according to the dynamic model, and re-sampled every time-step. However, when only a stationary source is involved, nothing is gained by performing these additional steps. Standard importance sampling/discretisation therefore remains a more appropriate choice for the problem at hand.

1.2 Security

Distributed statistical inference, which includes estimation, detection and filtering, is often performed over sensor networks [17]. While the use of multiple sensors has many obvious advantages, the distributed nature of such systems leaves them vulnerable to cyber-security attacks. Consider as an example the source localisation problem introduced in the previous section. There, the sensors transmit their binary decisions to a fusion centre which operates as a Bayesian estimator. If one of the agents is hijacked, an adversary can strategically feed misleading data to the fusion centre, hindering its ability to localise the source, or deliberately fostering certainty in a false location. In Chapter 3, we study the effect of such attacks on general discrete Bayesian estimators.

Recall that a Bayesian estimator constructs the posterior probability distribution of an estimand, given a collection of observations and a prior distribution. The goal of Chapter 3 is to specify fundamental limits on the ability of an adversary to manipulate the estimator's posterior distribution by injecting false data. The estimator is assumed to be oblivious to the adversary. Since noisy communication channels can only transmit quantized data, and discretisation of the parameter space is often performed for tractability, only Bayesian estimators with finite parameter and observation spaces are considered. Note that this framework also encompasses Bayesian hypothesis

testing. Although source localisation is used as the motivating example, the results of Chapter 3 are stated generally, and are therefore relevant to a variety of applications including target detection [18, 19], crowdsourcing [20], strategic misreporting [21], and human decision-making [22].

The novel contributions of Chapter 3 are as follows. We specify the set of posteriors that can be induced by the adversary in expectation. A stochastic strategy for the adversary to induce, on average, any posterior in this set is then proposed. We first consider the case where the estimator listens to the adversary alone, and then the case where the estimator also has access to honest measurements.

1.3 Connectivity preservation

Communication is a vital aspect of any multi-agent system. To enjoy the benefits of deploying a team, the individual agents must share information and strategically divide labour. Reliable communication is an obvious prerequisite, but guaranteeing this can be challenging when mobile agents are involved. The communication network topology at any point in time can be expressed as a graph. If this graph is connected, any two agents can exchange information by means of multi-hop communication. Chapter 4 is concerned with preserving network connectivity throughout the operation of the agents. Here, we adopt the standard assumption that communication links exist between two agents if and only if they are within some known range of each other. Thus, the connectivity preservation problem reduces to one of ensuring that some collection of inter-agent distances are not exceeded. This, it should be noted, presents a stronger requirement than set ISS [23], a point which is elaborated in Remark 4.3.1.

Connectivity preservation is rarely the primary goal of a multi-agent system. Network connectivity is typically only required as a means to an end. Thus, individual agents may receive additional control signals relating to their other tasks. Consider, for example, the source localisation problem of Chapter 2. In Section 2.6, a control law (2.65) is proposed to guide the agents to autonomously search for the source. However, the analysis of Chapter 2 presumes reliable communication between the agents and the fusion centre (which may itself be one of the agents), and this is not *a-priori* guaranteed. The proposed control law is piecewise continuous, bounded (see Remark 2.6.1), and depends on the measurements of the agents, which are stochastic. It is therefore desirable to develop a strategy for connectivity

preservation that accommodates such signals, without requiring prior knowledge of them. Any useful strategy must also offer robustness to uncertainty in the environment and the agent dynamics. Here, we design control-laws to preserve communication connectivity in the presence of piecewise continuous bounded disturbances, which can represent a large class of reference signals, environmental disturbances and un-modelled dynamics.

1.4 Time-varying optimisation

Recent literature on numerical optimisation has begun to explore problems involving cost functions, and even constraints, that change with time. These problems turn up within a variety of signal processing [9,24], robotics [25,26] and control [27, 28] algorithms for autonomous agents. Objectives or constraints often change in response to new information. For example, sequential Bayesian estimators (of the kind adopted in Chapter 2) construct a posterior probability distribution which they update in real time based on measurement data. The *maximum a-posteriori* (MAP) estimate is the mode of this distribution, and is therefore the solution of a time-varying optimisation problem. Similarly, the ML estimate maximises the likelihood function, which also changes on the arrival of new measurements. In general, allowing for temporal variations in the cost function and constraints offers scope for real-time optimisation in the presence of uncertainty.

An optimisation problem that changes at discrete instances in time can be treated as a sequence of optimisation problems. Assuming that every cost function and feasible set in this sequence is made available to the solver, one approach would be to solve each individual problem completely. This may not be tractable, depending on the rate at which new functions arrive. It is also unnecessary if the primary objective is to generate a sequence of iterates that 'track' the time-varying optimal points, or to achieve asymptotically low cost. Assuming there are bounds on the variation between consecutive cost functions and constraints, information from the current problem can be exploited to optimise its successor. A more efficient alternative, therefore, is to only solve each problem partially by limiting the number of iterations per cost function. Such approaches have been termed *running* methods in [29]. Chapters 5 and 6 consider the most extreme version of this, in which only a single iteration is performed per cost function. The machine learning literature refers to this as *online optimisation*, terminology which we adopt herein.

In Chapter 5, a first-order oracle is assumed, which provides gradient information about the current cost function. We rigorously analyse the performance of online projected gradient descent iterations, applied to different classes of time-varying smooth costs. Initially, unconstrained nonlinear cost functions are considered. Assuming a bounded increase between consecutive cost functions, we extend the well-known Zoutendijk result [30, Theorem 3.2] to obtain limit-inferior bounds on gradient magnitude. We then restrict attention to cost functions that satisfy the Polyak-Łojasiewicz (PL) inequality, first invoked by Polyak [31] to establish a linear convergence rate for gradient descent on time-invariant problems. We use the inequality to obtain steady-state sub-optimality bounds for time-varying problems. A detailed discussion of the PL inequality, and its relationship to other useful properties can be found in [32]. In particular, strong-convexity implies the PL inequality, and the PL inequality implies invexity. Finally, we focus specifically on strongly convex cost functions. The additional structure this affords allows us to include time-varying convex constraints in the analysis. Bounds on both sub-optimality, and distance to the minimiser (henceforth referred to as *tracking error*), are derived assuming a bounded shift between minimisers.

Chapter 6 then introduces an additional source of uncertainty. In place of a first-order oracle, a zeroth-order oracle is now assumed, which only provides the optimisation algorithm with the value of the cost function at the current time. The cost function is thereby treated as a black box, with a time-varying input-output map. This corresponds to practical scenarios in which derivatives are either unavailable [33, 34], or the cost of computing them is prohibitive. In such cases, the cost function directional derivatives can still be approximated via finite-differences. This method is adopted for time-invariant problems in [35], which relies on a two-point estimate of the directional derivative in a randomly chosen direction. Chapter 6 borrows techniques from [35] to deal with time-varying costs. In particular, we focus on strongly convex, unconstrained problems. Results are first obtained assuming a directional derivative oracle, and analysis for the zeroth-order case then proceeds along the same lines.

Part I

Estimation and Control

Chapter 2

Source Localisation with Minimal Information

> Something hidden. Go and find it.
> Go and look behind the Ranges—
> Something lost behind the Ranges.
> Lost and waiting for you. Go!
>
> *The Explorer*
> Rudyard Kipling

This chapter treats the localisation of a stationary signal source using a team of mobile agents that only take binary measurements. A Bayesian estimation algorithm, operating on a discretised search region, is adopted to solve this problem. We begin with a review of the relevant literature in Section 2.1. The problem is formulated mathematically in Section 2.2, and the estimation algorithm developed. This algorithm is analysed in Section 2.3, which rigorously examines the effects of discretisation. Section 2.4 derives D-optimal measurement locations, and Section 2.5 considers the implications of having inexact knowledge of the probability-of-detection function. A numerical example and simulation results are presented in Section 2.6, along with a control-law to autonomously guide the agents in their search.

2.1 Related works

A preliminary version of the results in this chapter appeared in the conference proceedings [36], which developed limited posterior convergence results fo-

cusing on two special cases: measurements taken at a single location with an arbitrarily located source, and measurements taken at a periodic sequence of locations assuming a source coincident with one of the chosen discretisation points. We now extend the latter to include an arbitrarily located source, and strengthen all the results to almost-sure convergence. Other new theoretical developments include the relationship with KL divergence, D-optimal location optimisation, and the analysis relating to inexact knowledge of the probability of detection. Below, we present a brief review of other relevant works in the literature, dividing them into Bayesian and classical approaches.

A Bayesian approach is adopted in [5] to localise a chemical source using a single mobile agent which detects the presence or absence of an odour. As in this thesis, the search region is discretised to approximate the posterior, however a theoretical convergence analysis is not offered. Rather, the focus of that paper is a search strategy based on maximising the rate of entropy reduction. Importance sampling is employed in [4] for source localisation with binary measurements, using a propagation model based on turbulent dispersion in the atmosphere. Their approach accommodates an unknown particle release rate by using Rao-Blackwellisation [37] to estimate it explicitly. The same Bayesian algorithm underpins [4,5,37], and the work herein. It should be emphasized that the contribution of this chapter is not to propose a new estimation algorithm, but rather to provide a rigorous treatment of the inevitable effects of discretisation, supplemented with numerical results. A search for multiple stationary targets is considered in [38], which considers a discrete environment to begin with, and assumes the agents directly observe the occupancy state of each cell with given false and missed detection probabilities. Since binary measurements are typically generated by means of a threshold, several works address the problem of designing threshold levels. These include [10], which studies the best achievable localisation accuracy using a binary sensor network, under a Gaussian plume propagation model. Threshold levels and sensor placement are investigated using the Bayesian Information Matrix (BIM), and the resulting theoretical error bounds are compared with the performance of the Metropolis-Hastings estimation algorithm. The tracking of a moving source using binary measurements is considered in [39], which uses particle filtering to estimate the source location, and proposes a heuristic for adaptively designing sensor threshold levels. This is extended to multi-bit measurements in [40], which focuses on adaptively designing quantisation thresholds based on the Bayesian Information Matrix.

Classical approaches treat the source location as a deterministic but unknown parameter, rather than a random variable. They tend to focus on

constructing estimators rather than maintaining a probability distribution. A maximum likelihood estimator is proposed in [9] for localising a diffusive source using binary measurements. That algorithm seeks to estimate a two-dimensional source location, time of signal emission, and several other model parameters via Fisher Scoring, a modified Newton method for maximising the likelihood function. Convergence guarantees are obtained as the number of sensors goes to infinity. Since each iteration requires reprocessing the entire batch of measurements, [9] also proposes a real-time approximate algorithm to avoid this. We compare the complexity and numerical performance of such maximum likelihood approaches with our own in Section 2.6.2. A set of different estimators are constructed in [41] without the use of any probability of detection model, but assuming noise free detections. Such model independent approaches clearly require less prior information, but typically display worse performance [11]. As in the Bayesian case, the design of binary quantisation thresholds based on the FIM is studied in [11]. Thresholds for multi-bit quantisation are studied in [42], which also compares the resulting theoretical error bounds with the performance of the maximum likelihood estimator and a second estimator that takes a weighted average of the sensor locations.

2.2 Problem formulation

We adopt the convention $\mathbb{N} = \{1, 2, ...\}$ and define $\mathbb{N}_k := \{1, ..., k\}$. Let $(a_k)_{k\in\mathbb{N}} \subset A$ to denote $a_k \in A$ for all $k \in \mathbb{N}$.

Consider a team of N agents exploring $\mathbb{R}^{\mathfrak{a}}$, where $\mathfrak{a} \in \{2, 3\}$. Let agent i have position $\mathbf{x}_i(t) \in \mathbb{R}^{\mathfrak{a}}$, which evolves in continuous time. We assume that all agents know their own position with respect to the same co-ordinate frame, and are equipped with identical sensors. The agents must search a compact region $S \subset \mathbb{R}^{\mathfrak{a}}$ for a source located at $\mathbf{s} \in S$. Together, the team of agents take a sequence of measurements $(d_k)_{k\in\mathbb{N}} \subset \{0, 1\}$ at a corresponding sequence of locations $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$. The measurement pairs $(\boldsymbol{\xi}_k, d_k)$ are transmitted in real-time to a fusion centre, where they are processed on arrival. The subscript k indexes the measurements according to the order in which are processed by the fusion centre. Note that the fusion centre is agnostic to the identity of the observing agent. Thus $\boldsymbol{\xi}_k \in \{\mathbf{x}_i(t_k) \mid i \in \mathbb{N}_N\}$, where $t_k \geq 0$ denotes the time at which reading d_k is taken.

We model $d_1, d_2, ...$ as random variables that are conditionally independent of each other, given the source location. We assume the probability of receiving

a detection is a known continuous function $\ell : \mathbb{R}^{\mathfrak{a}} \times \mathbb{R}^{\mathfrak{a}} \to (0,1)$ of the source and agent locations. Initially, we make no further assumptions about ℓ. Let $\ell(\mathbb{R}^{\mathfrak{a}}, \mathbb{R}^{\mathfrak{a}}) \subset (0,1)$ denote its image. Observe that there is always some non-zero probability of failing to detect the signal, as well as a non-zero background false detection probability regardless of where the source is. Having defined ℓ, the probability of obtaining the reading d_k from an agent at position $\boldsymbol{\xi}_k$ when the source location is $\mathbf{s}$, is given by the likelihood function

$$g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k) = \ell(\mathbf{s}, \boldsymbol{\xi}_k)^{d_k} \left[1 - \ell(\mathbf{s}, \boldsymbol{\xi}_k)\right]^{1-d_k}. \tag{2.1}$$

2.2.1 Discretised Bayesian framework

The estimation algorithm is developed in this section. We treat $\mathbf{s}$ as a random variable, drawn from some prior distribution p_0 over S. Bayesian techniques allow us to compute the posterior probability density of $\mathbf{s}$, given the history of measurements $d_{1:k} := (d_1, ..., d_k)$ and corresponding agent poses $\boldsymbol{\xi}_{1:k} = (\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_k)$. Bayes rule gives us a recursive description of this posterior density

$$p_k(\mathbf{s} \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = \frac{g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k) p_{k-1}(\mathbf{s} \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1})}{\displaystyle\int_S g(d_k \mid \mathbf{s}'; \boldsymbol{\xi}_k) p_{k-1}(\mathbf{s}' \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1}) d\mathbf{s}'}, \tag{2.2}$$

where the recursion is initialized with $p_0(\mathbf{s})$.

Although (2.2) is exact, the integrals involved do not, in general, have a closed-form, analytic solution. In order to work with arbitrary ℓ, the posterior must be approximated, and (2.2) computed numerically. To tackle this, we discretise S into a finite set of distinct points $\mathcal{C} := \{\mathbf{c}_1, ..., \mathbf{c}_M\}$, the elements of which we refer to as *centres*. If it is known that $\mathbf{s} \in \mathcal{C}$, then this yields a discrete version of the Bayes recursion (2.2),

$$\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) \hat{p}_{k-1}(i \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1})}{\displaystyle\sum_{j=1}^{M} g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k) \hat{p}_{k-1}(j \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1})}, \tag{2.3}$$

which is initialized with a discrete prior $\hat{p}_0(i)$. Without loss of generality, we assume $\hat{p}_0 : \mathbb{N}_M \to (0,1)$, noting that any $\mathbf{c}_i$ for which $\hat{p}_0(i) = 0$ can simply be omitted. For the more general case where $\mathbf{s} \in S$ is arbitrary, given

a particular choice of centres, we can define a set of cells $C_1, ..., C_M$, such that

1. each $C_i \subset S$ is connected, and $\mathbf{c}_i \in C_i$ for all i
2. $S = \bigcup_{i=1}^{M} C_i$
3. C_i and C_j are interior disjoint for all $i \neq j$.

This lends the following interpretation to the discrete posterior:

$$\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \approx \Pr(\mathbf{s} \in C_i \mid d_{1:k}).$$

As an example, the centres and cells can be chosen to form a grid or, more generally, a Voronoi diagram. Alternatively, the centres could be sampled from a probability distribution, such as the prior.

2.2.2 Relationship to importance sampling

Consider the posterior mean

$$\hat{\mathbf{s}}_k = \int_S \mathbf{s} p_k(\mathbf{s} \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) d\mathbf{s}.$$

This can be approximated numerically via

$$\hat{\mathbf{s}}_k \approx \frac{1}{M} \sum_{i=1}^{M} \mathbf{s}_i,$$

where each $\mathbf{s}_i \sim p_k(\mathbf{s} \mid d_{1:k}; \boldsymbol{\xi}_{1:k})$. However, we do not have a closed form expression for p_k, so we are unable to sample from it directly. Importance sampling assumes the ability to sample from some other more convenient density $q : S \to [0, \infty)$, known as the *importance density*. The importance density can be arbitrary, but its support must contain the support of p_k. The expectation can then be computed according to

$$\hat{\mathbf{s}}_k \approx \sum_{i=1}^{M} \hat{w}_k^i \mathbf{c}_i,$$

where $\mathbf{c}_i \sim q(\mathbf{s})$, and

$$w_k^i = \frac{p_k(\mathbf{c}_i \mid d_{1:k}; \boldsymbol{\xi}_{1:k})}{q(\mathbf{c}_i)}, \quad \hat{w}_k^i = \frac{w_i}{\sum_{j=1}^{M} w_j}.$$

Here, the $\mathbf{c}_i$ are referred to as *particles*, and the $\hat{w}_i$ as *weights*. Note that $\hat{w}_k^i$ is the normalized version of w_k^i. Recalling (2.2), and defining

$$\nu_k := \int_S g(d_k \mid \mathbf{s}'; \boldsymbol{\xi}_k) p_{k-1}(\mathbf{s}' \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1}) d\mathbf{s}',$$

we see that

$$p_k(\mathbf{c}_i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) p_{k-1}(\mathbf{c}_i \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1})}{\nu_k}.$$

The un-normalized weights therefore obey the recursive relationship

$$w_k^i = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) p_{k-1}(\mathbf{c}_i \mid d_{1:k-1}; \boldsymbol{\xi}_{1:k-1})}{\nu_k q(\mathbf{c}_i)} = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{\nu_k} w_{k-1}^i.$$

Applying the normalization,

$$\hat{w}_k^i = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) w_{k-1}^i}{\nu_k \sum_{j=1}^M \frac{g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k) w_{k-1}^j}{\nu_k}} = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) w_{k-1}^i}{\sum_{j=1}^M g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k) w_{k-1}^j}.$$

Finally, letting $W_k = \sum_{j=1}^M w_k^j$, we see that

$$\hat{w}_k^i = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) \frac{w_{k-1}^i}{W_{k-1}}}{\sum_{j=1}^M g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k) \frac{w_{k-1}^j}{W_{k-1}}} = \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k) \hat{w}_{k-1}^i}{\sum_{j=1}^M g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k) \hat{w}_{k-1}^j}.$$

Noting that this recursion is identical to (2.3), we see that the weight $\hat{w}_k^i$ obeys the same update rule as $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k})$. The initial weights are given by

$$\hat{w}_0^i = \frac{p_0(\mathbf{c}_i)}{q(\mathbf{c}_i) \sum_{j=1}^M \frac{p_0(\mathbf{c}_j)}{q(\mathbf{c}_j)}}. \tag{2.4}$$

In Section 2.2.1, no assumptions are made about how the centres are chosen, and an arbitrary discrete prior is assumed. If we choose the centres by sampling from an importance density, and initialize $\hat{p}_0(i) = \hat{w}_0^i$ as in (2.4), then the discretised approach of Section 2.2.1 is identical to importance sampling. Thus, importance sampling with a finite number of particles becomes a special case of discretisation. This is stated formally below.

Theorem 2.2.1 (Importance Sampling)**.** *Let* $\mathbf{c}_i \sim q(\mathbf{s})$ *for all* $i \in \mathbb{N}_M$*, where* $q : S \to [0, \infty)$ *is a probability distribution that satisfies*

$$q(\mathbf{s}) = 0 \implies \forall k \geq 0,\ p_k(\mathbf{s} \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = 0.$$

Furthermore, for all $i \in \mathbb{N}_M$*, let*

$$\hat{p}_0(i) = \frac{p_0(\mathbf{c}_i)}{q(\mathbf{c}_i) \sum_{j=1}^{M} \frac{p_0(\mathbf{c}_j)}{q(\mathbf{c}_j)}}.$$

Then $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = \hat{w}_k^i$ *for all* $i \in \mathbb{N}_M$ *and all* $k \geq 0$*.*

Remark 2.2.1. Theorem 2.2.1 guarantees that the remaining results of this chapter are also applicable to importance sampling with a finite number of particles. Any statements about the discretised posterior can be applied to the weights.

2.3 Consistency and convergence

Informally, the requirement of posterior consistency means that the posterior should become increasingly concentrated about the source location as $k \to \infty$. A precise definition and discussion of consistency can be found in [43, Section 4.1.1]. We adopt the following definition, specific to this problem.

Definition 2.1 (Posterior Consistency)**.** The posterior $\hat{p}_k$ is consistent if

$$\mathbf{s} \notin C_i \implies \lim_{k \to \infty} \hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = 0 \text{ a.s.}, \tag{2.5}$$

where C_i is the cell corresponding to centre $\mathbf{c}_i$.

If the source does not lie on the boundary between two cells, (2.5) is equivalent to

$$\mathbf{s} \in C_j \implies \lim_{k \to \infty} \hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = 1 \text{ a.s.}.$$

The set

$$\mathcal{O} := \left\{ i \in \mathbb{N}_M \mid \lim_{k \to \infty} \hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) = 0 \text{ a.s.} \right\}$$

is of clear interest, because consistency is also equivalent to $\mathcal{O} = \{i \in \mathbb{N}_M \mid \mathbf{s} \notin C_i\}$. Note that for any pair of cells (i, j), successive iterations of (2.3) up to time-step n gives us

$$\frac{\hat{p}_n(i \mid d_{1:n}; \boldsymbol{\xi}_{1:n})}{\hat{p}_n(j \mid d_{1:n}; \boldsymbol{\xi}_{1:n})} = \frac{\hat{p}_0(i)}{\hat{p}_0(j)} \prod_{k=1}^{n} \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k)}. \tag{2.6}$$

We accordingly define the likelihood ratio

$$Z_k^{(i,j)} := \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k)}, \tag{2.7}$$

which is the ratio of the probability of obtaining a reading d_k with the source at $\mathbf{c}_i$ to the probability of obtaining it with the source at $\mathbf{c}_j$. The expected value of the log-likelihood ratio conditioned on the source location

$$\mu_k^{(i,j)} := \mathbb{E}\left[\ln Z_k^{(i,j)} \mid \mathbf{s}\right] \tag{2.8}$$

$$= \sum_{d=0}^{1} \ln\left[\frac{g(d \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{g(d \mid \mathbf{c}_j; \boldsymbol{\xi}_k)}\right] g(d \mid \mathbf{s}; \boldsymbol{\xi}_k), \tag{2.9}$$

will play a key role in the subsequent analysis. Note that the value $\mu_k^{(i,j)}$ depends on $\mathbf{c}_i, \mathbf{c}_j, \mathbf{s}$ and $\boldsymbol{\xi}_k$. Following from (2.9), this relationship can be written as

$$\mu_k^{(i,j)} = \mu(\ell(\mathbf{c}_i, \boldsymbol{\xi}_k), \ell(\mathbf{c}_j, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k)), \tag{2.10}$$

where $\mu : (0, 1)^3 \to \mathbb{R}$,

$$\mu(x, y, z) := z \ln\left(\frac{x}{y}\right) + (1 - z) \ln\left(\frac{1 - x}{1 - y}\right) \tag{2.11}$$

$$= \ln\left(\frac{x^z (1 - x)^{(1-z)}}{y^z (1 - y)^{(1-z)}}\right). \tag{2.12}$$

In this problem, the expected log-likelihood ratio is intimately related to the Kullback-Leibler (KL) divergence $D(P||Q)$, which is a measure of the information lost when using a probability distribution Q to approximate another distribution P [44, Section 2.1]. Define

$$\mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_k) := D\Big(g(\cdot \mid \mathbf{s}; \boldsymbol{\xi}_k) \;\Big\|\; g(\cdot \mid \mathbf{x}; \boldsymbol{\xi}_k)\Big) \tag{2.13}$$

$$= -\sum_{d=0}^{1} g(d \mid \mathbf{s}; \boldsymbol{\xi}_k) \ln\left[\frac{g(d \mid \mathbf{x}; \boldsymbol{\xi}_k)}{g(d \mid \mathbf{s}; \boldsymbol{\xi}_k)}\right]$$

$$= -\mu(\ell(\mathbf{x}, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k)),$$

which is the KL divergence [45, Equation (2.26)] between the true conditional probability distribution for d_k, and a distribution which takes $\mathbf{x} \in \mathbb{R}^{\mathfrak{a}}$ as the source location. Noting the identity

$$\mu(x, z, z) - \mu(y, z, z) = \mu(x, y, z), \tag{2.14}$$

we obtain

$$\mu_k^{(i,j)} = \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_k) - \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_k). \tag{2.15}$$

Now consider a sequence of measurements $\mathbf{d} = (d_1, ..., d_n)$ taken at the corresponding locations $\boldsymbol{\xi}_{1:n} = (\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n)$. The probability distribution of $\mathbf{d}$, given source location $\mathbf{s}$ and measurement locations $\boldsymbol{\xi}_{1:n} \in \mathbb{R}^{n\mathfrak{a}}$, is

$$G(\mathbf{d} \mid \mathbf{s}; \boldsymbol{\xi}_{1:n}) = \prod_{k=1}^{n} g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k). \tag{2.16}$$

The extension of (2.13) to the distribution of a sequence of measurements is then given by

$$\mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_{1:n}) := D\Big(G(\cdot \mid \mathbf{s}; \boldsymbol{\xi}_{1:n}) \;\Big|\Big|\; G(\cdot \mid \mathbf{x}; \boldsymbol{\xi}_{1:n}) \Big) \tag{2.17}$$

$$= \sum_{k=1}^{n} \mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_k), \tag{2.18}$$

which follows from the additivity property of KL divergence for independent distributions (a corollary of [45, Theorem 2.5.3]). Thus

$$\sum_{k=1}^{n} \mu_k^{(i,j)} = \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:n}) - \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}). \tag{2.19}$$

2.3.1 General posterior convergence results

The first theoretical results can now be presented. We begin by establishing sufficient conditions for $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k})$ to decay to zero at index i, as $k \to \infty$.

Theorem 2.3.1. *Let $\mathbf{s} \in S$, and let $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$ be a bounded sequence of measurement locations. If there exists a pair of cells (i, j) and some $p > \frac{1}{2}$ such that*

$$\limsup_{n\to\infty} \left[\frac{1}{n^p} \sum_{k=1}^{n} \mu_k^{(i,j)} \right] < 0, \tag{2.20}$$

then under recursion (2.3), $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 0$ almost surely as $k \to \infty$.

Proof. Since we have assumed the co-domain of ℓ is $(0,1)$, this implies

$$\forall d \in \{0,1\},\ \forall \mathbf{s}, \boldsymbol{\xi} \in \mathbb{R}^{\mathfrak{a}},\quad g(d \mid \mathbf{s}; \boldsymbol{\xi}) > 0. \tag{2.21}$$

Recalling $\hat{p}_0(j) > 0$ for all j, recursion (2.3) together with (2.21) implies that $\hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) > 0$ for all $j \in \mathbb{N}_M$ and $k \in \mathbb{N}$. Thus (2.6) is well defined, and for any pair (i,j), we obtain

$$\frac{\hat{p}_n(i \mid d_{1:n}; \boldsymbol{\xi}_{1:n})}{\hat{p}_n(j \mid d_{1:n}; \boldsymbol{\xi}_{1:n})} = \frac{\hat{p}_0(i)}{\hat{p}_0(j)} \prod_{k=1}^{n} Z_k^{(i,j)}.$$

Since $\hat{p}_n(j \mid d_{1:n}; \boldsymbol{\xi}_{1:n}) \leq 1$, this yields the inequality

$$\hat{p}_n(i \mid d_{1:n}; \boldsymbol{\xi}_{1:n}) \leq \frac{\hat{p}_0(i)}{\hat{p}_0(j)} \prod_{k=1}^{n} Z_k^{(i,j)}. \tag{2.22}$$

Note that $Z_1^{(i,j)}, Z_2^{(i,j)}, ...$, are independent random variables because our measurements are independent. We also know $\mathbf{s}, \mathbf{c}_1, ..., \mathbf{c}_M \in S$, where $S \subset \mathbb{R}^{\mathfrak{a}}$ is compact. Furthermore, the sequence $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, ...$ is bounded, and therefore never leaves some compact subset $X \subset \mathbb{R}^{\mathfrak{a}}$. Since $\ell : \mathbb{R}^{\mathfrak{a}} \times \mathbb{R}^{\mathfrak{a}} \to (0,1)$ is continuous, it attains a minimum and maximum on $S \times X$ [46, Theorem 4.16]. Let

$$\ell_1 := \max \ell(S, X) < 1 \text{ and } \ell_0 := \min \ell(S, X) > 0. \tag{2.23}$$

Then (2.1) implies

$$\begin{aligned} \alpha := \min\left\{\frac{\ell_0}{\ell_1}, \frac{1-\ell_1}{1-\ell_0}\right\} &\leq \frac{g(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{g(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k)} \\ &\leq \max\left\{\frac{\ell_1}{\ell_0}, \frac{1-\ell_0}{1-\ell_1}\right\} =: \beta \end{aligned} \tag{2.24}$$

for all i, j, k. Thus for all i, j, k,

$$0 < \alpha \leq Z_k^{(i,j)} \leq \beta. \tag{2.25}$$

Now if condition (2.20) holds for some pair (i,j) and $p > \frac{1}{2}$, then by definition

$$\limsup_{n\to\infty} \frac{1}{n^p} \sum_{k=1}^{n} \mathbb{E}\left[\ln Z_k^{(i,j)}\right] < 0,$$

and applying the result of Lemma A.3, $\prod_{k=1}^{n} Z_k \to 0$ a.s. as $n \to \infty$. Equation (2.22) then implies $\hat{p}_n(i \mid d_{1:n}; \boldsymbol{\xi}_{1:n}) \to 0$ a.s.. □

Remark 2.3.1. This result is similar to [36, Theorem 1], which establishes convergence in probability of $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k})$ for $p = \frac{1}{2}$. Strengthening the requirement to $p > \frac{1}{2}$ allows us to obtain almost sure convergence.

Theorem 2.3.1 provides us with a sufficient condition ensuring that a cell index $i \in \mathcal{O}$, and this condition requires finding an index j for which $\sum_k \mu_k^{(i,j)}$ diverges at a sufficient rate.

Remark 2.3.2. Noting (2.19), if a pair (i, j) satisfy condition (2.20), then approximating the source location with $\mathbf{c}_j$ would asymptotically result in a lower KL divergence from $G(d_1, d_2... \mid \mathbf{s}; \boldsymbol{\xi}_1, \boldsymbol{\xi}_2, ...)$ than approximating the source with $\mathbf{c}_i$.

If the true source location coincides with some centre, Theorem 2.3.1 enables us to state a condition on the measurement location sequence $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, ...$ that guarantees posterior consistency.

Theorem 2.3.2 (Posterior Consistency)**.** *Let* $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$ *be bounded, and let* $\mathcal{C} = \{\mathbf{c}_1, ..., \mathbf{c}_M\} \subset S$*. Suppose* $\mathbf{c}_j = \mathbf{s}$ *for some* $j \in \mathbb{N}_M$*. If* $\forall i \neq j,\ \exists p > \frac{1}{2}$ *such that*

$$\liminf_{n\to\infty} \frac{1}{n^p} \sum_{k=1}^{n} (\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{c}_i, \boldsymbol{\xi}_k))^2 > 0, \tag{2.26}$$

then $\hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 1$ *a.s..*

Proof. For any $\mathbf{x} \in \mathbb{R}^{\mathfrak{a}}$, the total variation distance between distributions $g(\cdot|\mathbf{s}, \boldsymbol{\xi}_k)$ and $g(\cdot|\mathbf{x}, \boldsymbol{\xi}_k)$ is given by

$$\sup_{d\in\{0,1\}} |g(d|\mathbf{s}, \boldsymbol{\xi}_k) - g(d|\mathbf{x}, \boldsymbol{\xi}_k)| = |\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{x}, \boldsymbol{\xi}_k)|.$$

Pinsker's inequality [47, Lemma 2.5] is a lower bound on KL divergence, which then yields

$$\mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_k) \geq 2(\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{x}, \boldsymbol{\xi}_k))^2.$$

By assumption $\mathbf{c}_j = \mathbf{s}$, which implies $\mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_k) = 0$ for any $\boldsymbol{\xi}_k$. Thus, (2.18) and (2.19) imply that for all $n \in \mathbb{N}$ and $i \in \mathbb{N}_M$,

$$\sum_{k=1}^{n} \mu_k^{(i,j)} = -\mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}) \leq -\sum_{k=1}^{n} 2(\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{c}_i, \boldsymbol{\xi}_k))^2,$$

which in turn implies

$$\limsup_{n\to\infty} \frac{1}{n^p} \sum_{k=1}^{n} \mu_k^{(i,j)} \leq \limsup_{n\to\infty} \frac{-2}{n^p} \sum_{k=1}^{n} (\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{c}_i, \boldsymbol{\xi}_k))^2$$
$$= -\liminf_{n\to\infty} \frac{2}{n^p} \sum_{k=1}^{n} (\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{c}_i, \boldsymbol{\xi}_k))^2.$$

If for all $i \neq j$, there exists $p > \frac{1}{2}$ such that (2.26) holds, then the LHS of the above inequality is strictly negative. Theorem 2.3.1 then implies $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 0$ a.s. for all $i \neq j$. Since $\hat{p}_k$ is a probability distribution, this implies $\hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 1$ a.s.. □

Remark 2.3.3. If $\ell(\mathbf{c}_i, \boldsymbol{\xi}_k) = \ell(\mathbf{s}, \boldsymbol{\xi}_k)$, then having the source at $\mathbf{c}_i$ yields the same probability of detection at $\boldsymbol{\xi}_k$, as if the source was at $\mathbf{s}$. Thus, $\mathbf{c}_i$ cannot be distinguished from $\mathbf{s}$ using measurements taken at $\boldsymbol{\xi}_k$. Condition (2.26) ensures the agents take readings sufficiently often at locations which provide enough information to distinguish between cells.

Next, we consider what happens when the source does not coincide with any centre.

2.3.2 Periodic measurement locations

Analysing the general case where $\mathbf{s} \notin \mathcal{C}$ is difficult when considering completely arbitrary agent location sequences. We therefore restrict our attention to those that are periodic. Many bounded sequences of practical interest are either periodic, or converge to one that is. Furthermore, the effectiveness of any finite sequence $\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n$ can be analysed by considering a periodic sequence for which $\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n$ constitutes a single period.

Definition 2.2. A sequence $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}}$ is n-periodic iff $\boldsymbol{\xi}_k = \boldsymbol{\xi}_{k+n}$ for all k.

Examples of this include

1. a single agent moving in a periodic trajectory, taking measurements at the same locations every n time-steps
2. a team of n agents remaining stationary, and taking measurements in a fixed order.

Any n-periodic location sequence is fully specified by the vector $\boldsymbol{\xi}_{1:n} = (\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n) \in \mathbb{R}^{n\mathfrak{a}}$. For n-periodic trajectories, $\sum_{k=1}^{n} \mu_k^{(i,j)} < 0$ is a sufficient condition for (2.20).

Remark 2.3.4. The assumption of n-periodicity can be weakened. Partition an arbitrary sequence $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, ...$ into blocks of some fixed length $L \in \mathbb{N}$, and permute the measurement order within these blocks. This has no effect on condition (2.20). Thus, if there exists such a permutation which yields an n-periodic sequence, the results of this section (and Section 2.5.1) apply without modification to the original sequence.

Equation (2.19) then implies that, in the limit, the algorithm selects the indices of centres which, when treated as the source location, minimise KL divergence from the true measurement probability distribution. This is stated precisely below.

Theorem 2.3.3. *Let* $\mathbf{s} \in S$ *and let* $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$ *be* n*-periodic for some* $n \in \mathbb{N}$*. Then*

$$\mathcal{O}^c \subset \underset{i\in\mathbb{N}_M}{\arg\min}\, \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}) =: \mathcal{B}. \tag{2.27}$$

Proof. Consider any index $i \notin \mathcal{B}$. By definition, for any $j \in \mathcal{B}$, we have

$$\mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}) > \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:n}),$$

which implies $\sum_{k=1}^{n} \mu_k^{(i,j)} < 0$ by (2.19). Since $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}}$ is n-periodic,

$$\forall m \in \mathbb{N},\ \sum_{k=1}^{m} \mu_k^{(i,j)} = \left\lfloor \frac{m}{n} \right\rfloor \sum_{k=1}^{n} \mu_k^{(i,j)} + \sum_{k=1}^{m \bmod n} \mu_k^{(i,j)}.$$

Note that $\lim_{m\to\infty} m^{-1} \lfloor \frac{m}{n} \rfloor = \frac{1}{n}$ and $(m \bmod n) < n$ for any m. Therefore,

$$\begin{aligned}
&\lim_{m\to\infty} \frac{1}{m} \sum_{k=1}^{m} \mu_k^{(i,j)} \\
&\quad = \lim_{m\to\infty} \left(\frac{\lfloor \frac{m}{n} \rfloor}{m} \sum_{k=1}^{n} \mu_k^{(i,j)} + \frac{1}{m} \sum_{k=1}^{m \bmod n} \mu_k^{(i,j)} \right) \\
&\quad = \frac{1}{n} \sum_{k=1}^{n} \mu_k^{(i,j)} < 0.
\end{aligned}$$

Thus (2.20) is satisfied using $p = 1$, and $\hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 0$ a.s. by Theorem 2.3.1, which by definition implies $i \in \mathcal{O}$. We have shown $\mathcal{B}^c \subset \mathcal{O}$, which is equivalent to $\mathcal{O}^c \subset \mathcal{B}$. □

Remark 2.3.5. Equation (2.27) reveals that the posterior may fail to decay to zero only at indices which minimise $\mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n})$. The centres corresponding to these indices are solutions to $\arg\min_{\mathbf{x}\in\mathcal{C}} \mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_{1:n})$. They can be

considered approximate solutions to

$$\arg\min_{\mathbf{x}\in S} \mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}), \tag{2.28}$$

where the optimisation now takes place over the entire search region instead of the finite set $\mathcal{C}$.

The problem (2.28) is more amenable to analysis, as it does not depend on the choice of centres in $\mathcal{C}$, and its solution set provides us with additional insight.

Proposition 2.3.4. *Given* $\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n \in \mathbb{R}^{\mathfrak{a}}$ *and* $\mathbf{s} \in S$,

$$\min_{\mathbf{x}\in S} \mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}) = 0, \text{ and}$$

$$\arg\min_{\mathbf{x}\in S} \mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}) = \bigcap_{k=1}^{n} \{\mathbf{x} \in S \mid \ell(\mathbf{x},\boldsymbol{\xi}_k) = \ell(\mathbf{s},\boldsymbol{\xi}_k)\} := \mathcal{A}. \tag{2.29}$$

Therefore, $\mathbf{s} \in \arg\min_{\mathbf{x}\in S} \mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n})$.

Proof. Here we exploit the properties of KL divergence stated in [45, Theorem 2.6.3]. Since KL divergence is non-negative, we have $\mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}) \geq 0$ for all $\mathbf{x}$. From (2.16),

$$p(\mathbf{d} \mid \mathbf{s};\boldsymbol{\xi}_{1:n}) = \prod_{k=1}^{n} \ell(\mathbf{s},\boldsymbol{\xi}_k)^{d_k} \left[1 - \ell(\mathbf{s},\boldsymbol{\xi}_k)\right]^{1-d_k}.$$

If $\mathbf{x} \in \mathcal{A}$, then $\ell(\mathbf{x},\boldsymbol{\xi}_k) = \ell(\mathbf{s},\boldsymbol{\xi}_k)$ for all $k \in \mathbb{N}_n$, which implies $p(\mathbf{d} \mid \mathbf{s};\boldsymbol{\xi}_{1:n}) = p(\mathbf{d} \mid \mathbf{x};\boldsymbol{\xi}_{1:n})$ for all $\mathbf{d} \in \{0,1\}^n$. Recalling (2.17), this implies $\mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}) = 0$.

Recalling (2.18), if $\mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_{1:n}) = 0$, then $\mathcal{K}(\mathbf{s}||\mathbf{x};\boldsymbol{\xi}_k) = 0$ for all $k \in \mathbb{N}_n$. According to the definition (2.13), this holds if and only if

$$\forall k \in \mathbb{N}_n,\ \forall d \in \{0,1\},\ g(d \mid \mathbf{s},\boldsymbol{\xi}_k) = g(d \mid \mathbf{x};\boldsymbol{\xi}_k).$$

Finally referring to (2.1), choosing $d = 1$ implies $\ell(\mathbf{s},\boldsymbol{\xi}_k) = \ell(\mathbf{x},\boldsymbol{\xi}_k)$ for all $k \in \mathbb{N}_n$, and therefore $\mathbf{x} \in \mathcal{A}$. □

Thus (2.28) contains only the candidate locations that are indistinguishable from the source based on the entire history of measurements (see Remark 2.3.3). Given this characterisation, it is obviously desirable to define the n-periodic sequence $\boldsymbol{\xi}_{k\in\mathbb{N}}$ by choosing $\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n$ such that $\mathcal{A} = \{\mathbf{s}\}$. This is a useful requirement to impose when planning agent trajectories, as it guarantees there is sufficient information available from the measurements to uniquely identify the source.

Remark 2.3.6. Observe that $\mathcal{A} \subset \mathbb{R}^{\mathfrak{a}}$ is the solution set of n simultaneous non-linear equations. Therefore if $n > \mathfrak{a}$, typically only mild conditions on the measurement location geometry are required to guarantee $\mathcal{A} = \{\mathbf{s}\}$. Such conditions are developed in Section 2.3.4 for the case in which the probability of detection is purely a function of distance from the source.

If $\mathcal{A} = \{\mathbf{s}\}$ and $\mathbf{s} \in \mathcal{C}$, the posterior is consistent and estimation algorithm will eventually unambiguously identify the source index.

Corollary 2.3.5. *Let* $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$ *be* n*-periodic, and let* $\mathcal{C} = \{\mathbf{c}_1, ..., \mathbf{c}_M\} \subset S$*. Suppose* $\mathbf{c}_j = \mathbf{s}$ *for some* $j \in \mathbb{N}_M$*. If* $\mathcal{A} = \{\mathbf{s}\}$*, then* $\hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 1$ *a.s..*

Proof. Suppose $\mathcal{A} = \{\mathbf{s}\} = \{\mathbf{c}_j\} = \arg\min_{\mathbf{x}\in S} \mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_{1:n})$. Since $\mathbf{c}_j \in \mathcal{C} \subset S$, we have $\mathcal{B} = \{\mathbf{c}_j\}$. The centres are distinct and thus for all $i \neq j$, $\mathbf{c}_i \notin \mathcal{B}$, which implies $i \in \mathcal{O}$ by Theorem 2.3.3. Since $\hat{p}_k$ is a probability distribution, this implies $\hat{p}_k(j \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \to 1$ a.s.. □

If $\mathcal{A} = \{\mathbf{s}\}$ but $\mathbf{s}$ does not coincide with the centres, the support of the posterior (in the limit) will contain approximations to $\mathbf{s}$ in $\mathcal{C}$ that yield the lowest KL divergence from $G(\cdot \mid \mathbf{s}, \boldsymbol{\xi}_{1:n})$.
Remark 2.3.7. Note that cells having the lowest $\mathcal{K}$ value do not necessarily coincide with the centres closest to the source:

$$\|\mathbf{s} - \mathbf{c}_j\| \leq \|\mathbf{s} - \mathbf{c}_i\| \not\Longrightarrow \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:n}) \leq \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}).$$

According to the remark above, if the cells are based on the Voronoi decomposition of S, there is no guarantee of consistency (with respect to Definition 2.1) when $\mathbf{s} \notin \mathcal{C}$. However, simulation results in Section 2.6 indicate that the estimation procedure still performs well when the centres are chosen in a grid.

2.3.3 Bound on posterior rate of decay

We now investigate the rate at which the posterior decays to zero at a given cell index, and its relationship to KL divergence.

Proposition 2.3.6. *Let* $\mathbf{s} \in S$*, and let* $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}} \subset \mathbb{R}^{\mathfrak{q}}$ *be bounded. Suppose that the hypothesis of Theorem 2.3.1 is satisfied for some pair of cells* (i, j)*. Then there exists* $K > 0$ *such that for any* $\epsilon > 0$*,*

$$\Pr\{\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \ \geq \epsilon\} \leq \exp\left(-\frac{K\left(\vartheta - \gamma_m\right)^2}{m}\right) \tag{2.30}$$

for sufficiently large m*, where*

$$\gamma_m := \sum_{k=1}^{m} \mu_k^{(i,j)}, \ \ \vartheta := \ln\left[\frac{\epsilon \hat{p}_0(j)}{\hat{p}_0(i)}\right].$$

Therefore, $\Pr\{\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \geq \epsilon\} \to 0$ *as* $m \to \infty$*.*

Proof. Recall (2.22) from the proof of Theorem 2.3.1:

$$\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \leq \frac{\hat{p}_0(i)}{\hat{p}_0(j)} \prod_{k=1}^{m} Z_k^{(i,j)}.$$

Therefore

$$\left(\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \geq \epsilon\right) \implies \left(\frac{\hat{p}_0(i)}{\hat{p}_0(j)} \prod_{k=1}^{m} Z_k^{(i,j)} \geq \epsilon\right),$$

which implies that for any $\epsilon > 0$,

$$\Pr\left\{\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \geq \epsilon\right\} \leq \Pr\left\{\prod_{k=1}^{m} Z_k^{(i,j)} \geq \epsilon \frac{\hat{p}_0(j)}{\hat{p}_0(i)}\right\}.$$

Now since (2.20) holds for some $p > \frac{1}{2}$, this implies

$$\lim_{m\to\infty} \frac{\gamma_m}{\sqrt{m}} = -\infty. \tag{2.31}$$

For α, β defined in (2.24), let $K := 2(\ln\beta - \ln\alpha)^{-2}$. Given the bounds on $Z_k^{(i,j)}$ in (2.25), Lemma A.4 (Appendix A) can now be applied to obtain (2.30) for sufficiently large m. Furthermore, (2.31) implies $\frac{\gamma_m^2}{m} \to \infty$, by which

$$\lim_{m\to\infty} \frac{\left(\vartheta - \gamma_m\right)^2}{m} = \infty.$$

Therefore, (2.30) gives us $\Pr\{\hat{p}_m(i \mid d_{1:m}; \boldsymbol{\xi}_{1:m}) \geq \epsilon\} \to 0$. □

This proposition places a decaying bound on the probability that the posterior at a given index exceeds an arbitrary threshold, after a sufficiently large number of time-steps. Recalling (2.19), the rate at which this bound decays is clearly determined by the growth of

$$-\sum_{k=1}^{m} \mu_k^{(i,j)} = \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:m}) - \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:m})$$

as $m \to \infty$. If our location sequence is n-periodic, the finite sum $-\sum_{k=1}^{n} \mu_k^{(i,j)}$ fully determines the growth of the infinite sum. Furthermore, choosing j to minimise $\mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:m})$ yields the tightest bound on the decay available. Thus

$$\Gamma_i := \mathcal{K}(\mathbf{s}||\mathbf{c}_i; \boldsymbol{\xi}_{1:n}) - \mathcal{K}^\star \tag{2.32}$$

bounds the rate of decay of $\hat{p}_k$ at index i, where

$$\mathcal{K}^\star := \min_{j \in \mathbb{N}_m} \mathcal{K}(\mathbf{s}||\mathbf{c}_j; \boldsymbol{\xi}_{1:n}).$$

As Γ_i increases, the more rapidly this bound decays. In particular, we have $\Gamma_i = 0$ for every index $i \in \mathcal{B}$, and $\Gamma_i > 0$ for every other index. Note that if $\mathbf{s} \in \mathcal{C}$, then $\mathcal{K}^* = 0$.

2.3.4 Range dependent probability of detection

In many scenarios, the probability of detection will be purely a function of the distance from the source to an agent. We can then say ℓ is of the form $\ell(\mathbf{x}, \boldsymbol{\xi}) = \rho(\|\mathbf{x} - \boldsymbol{\xi}\|)$, where $\rho : [0, \infty) \to (0, 1)$. Typically ρ will be strictly decreasing, and therefore injective. We now consider the application of the previous results to this case.

We begin by examining the consistency requirement (2.26) for a general measurement location sequence in Theorem 2.3.2, when the source coincides with some centre.

Lemma 2.3.7. *Let $\mathbf{s} \in S$, let $(\boldsymbol{\xi}_k)_{k \in \mathbb{N}} \subset \mathbb{R}^{\mathfrak{a}}$ be a bounded sequence, and assume $\ell(\mathbf{x}, \boldsymbol{\xi}) = \rho(\|\mathbf{x} - \boldsymbol{\xi}\|)$, where $\rho : [0, \infty) \to (0, 1)$ is continuous and injective. Then*

$$\forall \delta > 0,\ \exists \epsilon \in (0, 1),$$
$$|\|\boldsymbol{\xi}_k - \mathbf{s}\| - \|\boldsymbol{\xi}_k - \mathbf{c}_i\|| \geq \delta \implies |\ell(\mathbf{s}, \boldsymbol{\xi}_k) - \ell(\mathbf{c}_i, \boldsymbol{\xi}_k)| \geq \epsilon.$$

Proof. Let $z := \sup\{\|\boldsymbol{\xi}_k - \mathbf{c}\| \mid k \in \mathbb{N},\ \mathbf{c} \in \mathcal{C} \cup \{\mathbf{s}\}\} < \infty$, and let $D := \rho([0, z])$. Let $\bar{\rho} : [0, z] \to D,\ \bar{\rho}(r) = \rho(r)$, which is a bijection because ρ is injective. By [46, Theorem 4.17], $\bar{\rho}^{-1}$ is continuous and therefore,

$$\forall l_1, l_2 \in D,\ \forall \delta > 0,\ \exists \epsilon > 0, |l_1 - l_2| < \epsilon \implies |\bar{\rho}^{-1}(l_1) - \bar{\rho}^{-1}(l_2)| < \delta. \quad (2.33)$$

Now for any $r_i \in [0, z]$, $l_i = \bar{\rho}(r_i) \in D$. Therefore (2.33) implies

$$\forall r_1, r_2 \in [0, z],\ \forall \delta > 0,\ \exists \epsilon > 0, |\bar{\rho}(r_1) - \bar{\rho}(r_2)| < \epsilon \implies |r_1 - r_2| < \delta,$$

which is in turn equivalent to

$$\forall r_1, r_2 \in [0, z],\ \forall \delta > 0,\ \exists \epsilon > 0, |r_1 - r_2| \geq \delta \implies |\bar{\rho}(r_1) - \bar{\rho}(r_2)| \geq \epsilon. \quad (2.34)$$

The image of $\bar{\rho}$ is contained in $(0, 1)$, which implies $\epsilon \in (0, 1)$. Now $\|\boldsymbol{\xi}_k - \mathbf{s}\|, \|\boldsymbol{\xi}_k - \mathbf{c}_i\| \in [0, z]$ by definition of z, and therefore $\bar{\rho}(\|\boldsymbol{\xi}_k - \mathbf{s}\|) = \ell(\mathbf{s}, \boldsymbol{\xi}_k)$ and $\bar{\rho}(\|\boldsymbol{\xi}_k - \mathbf{c}_i\|) = \ell(\mathbf{c}_i, \boldsymbol{\xi}_k)$. Thus choosing $r_1 = \|\boldsymbol{\xi}_k - \mathbf{s}\|$ and $r_2 = \|\boldsymbol{\xi}_k - \mathbf{c}_i\|$ in (2.34) completes the proof. □

Remark 2.3.8. Referring to the Lemma above, condition (2.26) is met for a particular cell i if $|\|\boldsymbol{\xi}_k - \mathbf{s}\| - \|\boldsymbol{\xi}_k - \mathbf{c}_i\|| \geq \delta$ occurs sufficiently often for the same $\delta > 0$. A simple way to guarantee this holds for every $\mathbf{c}_i \neq \mathbf{s}$ is to make sure the location sequence does not travel in (or converge to) a straight line indefinitely.

We now turn our attention to periodic locations sequences, allowing $\mathbf{s} \in S$ to be arbitrary. As discussed in Section 2.3.2, the basic requirement for an n-periodic location sequence is to ensure $\mathcal{A} = \{\mathbf{s}\}$. If enough readings are taken at location $\boldsymbol{\xi}$, then the probability of detection $\ell(\mathbf{s}, \boldsymbol{\xi})$ can be estimated from the ratio of hits to misses. When ρ is injective, this probability of detection can be mapped back to place $\boldsymbol{\xi}$ at a unique distance from $\mathbf{s}$. This suggests a strategy akin to trilateration will ensure $\mathbf{s}$ is the unique solution to (2.28), and the result below confirms our intuition. Let $\operatorname{aff}(X)$ denote the affine hull of some $X \subset \mathbb{R}^{\mathfrak{a}}$.

Proposition 2.3.8. *Let $\mathbf{s} \in S \subset \mathbb{R}^{\mathfrak{a}}$, where $\mathfrak{a} \in \{2, 3\}$. Let $\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n \in S$ and let ℓ be of the form $\ell(\mathbf{x}, \boldsymbol{\xi}) = \rho(\|\mathbf{x} - \boldsymbol{\xi}\|)$, where $\rho : [0, \infty) \to (0, 1)$ is continuous and injective. If $\dim \operatorname{aff}(\{\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_n\}) = \mathfrak{a}$, then*

$$\mathcal{A} := \bigcap_{k=1}^{n} \{\mathbf{x} \in S \mid \ell(\mathbf{x}, \boldsymbol{\xi}_k) = \ell(\mathbf{s}, \boldsymbol{\xi}_k)\} = \{\mathbf{s}\}.$$

Proof. Since ρ is injective,

$$\begin{aligned}\mathcal{A} &= \bigcap_{k=1}^{n} \{\mathbf{x} \in S \mid \rho(\|\mathbf{x} - \boldsymbol{\xi}_k\|) = \rho(\|\mathbf{s} - \boldsymbol{\xi}_k\|)\} \\ &= \bigcap_{k=1}^{n} \{\mathbf{x} \in S \mid \|\mathbf{x} - \boldsymbol{\xi}_k\| = \|\mathbf{s} - \boldsymbol{\xi}_k\|\},\end{aligned}$$

which is the intersection of n spheres in $S \subset \mathbb{R}^{\mathfrak{a}}$. For their intersection to be unique, it is sufficient for three of the $\boldsymbol{\xi}_k$ to not be collinear when $\mathfrak{a} = 2$, and four of the $\boldsymbol{\xi}_k$ to not be coplanar when $\mathfrak{a} = 3$. □

2.4 Optimal measurement locations

As noted in Remark 2.3.6, the requirement that $\mathcal{A} = \{\mathbf{s}\}$ typically imposes only mild constraints on the geometry of an n-periodic measurement location sequence if $n > \mathfrak{a}$. While this guarantees there is sufficient information available to uniquely identify the source, it makes no claim to optimality. In this section, we first attempt to optimise the measurement locations with respect to the determinant of the Bayesian Information Matrix (BIM) [12, Section 4.3.3.2], conditioned on the available information as in [48]. The inverse of the BIM is the Bayesian Cramer-Rao bound, a lower bound on the MSE of any estimator[1]. Thus, maximising the BIM determinant minimises a lower bound on the volume of the estimator's concentration ellipsoids [12, Section 4.3.2.1].

Recall we have a team of N agents, and suppose they each report one measurement in a fixed sequence every N time-steps. Consider the following question: given the information $(\boldsymbol{\xi}_k, d_k)_{k \in \mathbb{N}_n}$ received up to time n, what are the D-optimal agent locations $\boldsymbol{\Xi} := (\boldsymbol{\xi}_1^+, ..., \boldsymbol{\xi}_N^+) \in \mathbb{R}^{n\mathfrak{a}}$ at which to take the next N measurements $\mathbf{d}^+ := (d_1^+, ..., d_N^+) \in \{0, 1\}^N$?

The joint probability distribution for $\mathbf{d}^+$ and $\mathbf{s}$ conditioned on the information received is given by

$$p(\mathbf{d}^+, \mathbf{s} \mid d_{1:n}; \boldsymbol{\xi}_{1:n}, \boldsymbol{\Xi}) = G(\mathbf{d}^+ \mid \mathbf{s}; \boldsymbol{\Xi}) p_n(\mathbf{s} \mid d_{1:N}; \boldsymbol{\xi}_{1:N}),$$

where G is defined in (2.16). The Bayesian information matrix for estimating $\mathbf{s}$ from $\mathbf{d}^+$, given the information received up to time n, is then

$$\mathbf{J}_n(\boldsymbol{\Xi}) := -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ln p(\mathbf{d}^+, \mathbf{s} \mid d_{1:n}; \boldsymbol{\xi}_{1:n}, \boldsymbol{\Xi}) \mid d_{1:n}\right].$$

[1] subject to the bias conditions in [12, Equation (4.522)].

Let $\mathcal{X}_n \subset \mathbb{R}^{n\mathfrak{a}}$ denote the feasible set of agent locations at time n (which may, for example, incorporate motion constraints). Solving

$$\boldsymbol{\xi}_{n+1:n+N} = \underset{\boldsymbol{\Xi}\in\mathcal{X}_n}{\arg\max} \det \mathbf{J}_n(\boldsymbol{\Xi}) \tag{2.35}$$

yields D-optimal locations for the next N measurements. Although [48, 49] provide a recursive method for computing the BIM numerically, obtaining a direct solution to (2.35) is intractable. Instead, we propose a relaxed version of the problem which optimises the classical Fisher Information Matrix (FIM). The BIM can be written as

$$\mathbf{J}_n(\boldsymbol{\Xi}) = \mathbb{E}\left[\mathcal{J}(\mathbf{s};\boldsymbol{\Xi}) \mid d_{1:n}\right] + \mathbf{J}_n^{\ominus}(d_{1:n}, \boldsymbol{\xi}_{1:n}), \tag{2.36}$$

where

$$\mathbf{J}_n^{\ominus}(d_{1:n}, \boldsymbol{\xi}_{1:n}) := -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ln p_n(\mathbf{s} \mid d_{1:n}; \boldsymbol{\xi}_{1:n}) \mid d_{1:n}\right]$$

is the contribution of the information information already received, and

$$\begin{aligned}
\mathcal{J}(\mathbf{s};\boldsymbol{\Xi}) &= -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ \ln G(\mathbf{d}^+ \mid \mathbf{s};\boldsymbol{\Xi}) \mid \mathbf{s}\right] \\
&= -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \sum_{k=1}^{N} \ln g(d_k^+ \mid \mathbf{s}; \boldsymbol{\xi}_k^+) \mid \mathbf{s}\right] \\
&= -\sum_{k=1}^{N} \mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ln g(d_k^+ \mid \mathbf{s}; \boldsymbol{\xi}_k^+) \mid \mathbf{s}\right].
\end{aligned}$$

is the classical FIM conditioned on the source location [12, Equation (4.515)]. below. Both $\mathbf{J}_n^{\ominus}(d_{1:n}, \boldsymbol{\xi}_{1:n})$ and $\mathcal{J}(\mathbf{s};\boldsymbol{\Xi})$ are symmetric positive semi-definite [12, (4.519), (4.393)]. It then follows from the Minkowski determinant inequality that

$$\begin{aligned}
\det \mathbf{J}_n(\boldsymbol{\Xi}) &\geq \det \mathbb{E}\left[\mathcal{J}(\mathbf{s};\boldsymbol{\Xi}) \mid d_{1:n}\right] + \det \mathbf{J}_n^{\ominus}(d_{1:n}, \boldsymbol{\xi}_{1:n}) \\
&\approx \det \mathcal{J}(\hat{\mathbf{s}}_n;\boldsymbol{\Xi}) + \det \mathbf{J}_n^{\ominus}(d_{1:n}, \boldsymbol{\xi}_{1:n}),
\end{aligned} \tag{2.37}$$

where

$$\hat{\mathbf{s}}_n := \mathbb{E}[\mathbf{s} \mid d_{1:n}]$$

represents the mean of the posterior p_n. Since the second term of (2.37) does not depend on $\boldsymbol{\Xi}$,

$$\boldsymbol{\xi}_{n+1:n+N} = \underset{\boldsymbol{\Xi}\in\mathcal{X}_n}{\arg\max} \det \mathcal{J}(\hat{\mathbf{s}}_n;\boldsymbol{\Xi}), \tag{2.38}$$

is a relaxation of (2.35) that chooses measurement locations to maximise the determinant of the FIM evaluated at the expected source location. The FIM has the structure

$$\mathcal{J}(\mathbf{s};\Xi) = \sum_{k=1}^{N} J(\mathbf{s};\boldsymbol{\xi}_k^+), \tag{2.39}$$

where

$$J(\mathbf{s};\boldsymbol{\xi}_k) := -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ln g(d_k \mid \mathbf{s};\boldsymbol{\xi}_k)\right] \tag{2.40}$$

$$= \frac{\nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)^\top}{\ell(\mathbf{s},\boldsymbol{\xi}_k)[1-\ell(\mathbf{s},\boldsymbol{\xi}_k)]}. \tag{2.41}$$

is the FIM for a single reading taken at $\boldsymbol{\xi}_k$. Equation (2.41) is derived in Appendix B. We exploit this structure below to obtain an analytic solution to the relaxed problem (2.38) for a range-dependent probability of detection, under particular distance constraints.

Remark 2.4.1. It is interesting to note that the final expression for Fisher Information in (2.41) does not depend on the log-likelihood Hessian.

2.4.1 Range dependent probability of detection

Suppose the probability of detection is a smooth function of distance. We focus on localisation in the plane, letting $\mathbf{s} = (s_1, s_2)$ and $\boldsymbol{\xi}_k = (\xi_{k,1}, \xi_{k,2})$.

Proposition 2.4.1 (Fisher Information Matrix)**.** *Let $\mathbf{s} \in S \subset \mathbb{R}^2$, and assume $\ell(\mathbf{s},\boldsymbol{\xi}) = \rho(\|\mathbf{s}-\boldsymbol{\xi}\|)$, where $\rho : [0,\infty) \to (0,1)$ is continuously differentiable. Let $\boldsymbol{\xi}_{1:N} \in \mathbb{R}^{2N}$ be such that, $\forall k \in \mathbb{N}_N$, $\boldsymbol{\xi}_k \neq \mathbf{s}$. Define*

$$r_k := \|\mathbf{s}-\boldsymbol{\xi}_k\|, \quad \theta_k := \operatorname{atan2}(\xi_{k,2}-s_2, \xi_{k,1}-s_1).$$

Then

$$\mathcal{J}(\mathbf{s};\boldsymbol{\xi}_{1:N}) = \sum_{k=1}^{N} \frac{\rho'(r_k)^2}{\rho(r_k)[1-\rho(r_k)]} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix}. \tag{2.42}$$

Proof. Since $\ell(\mathbf{s},\boldsymbol{\xi}_k) = \rho(r_k)$, it follows from the chain rule that

$$\nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k) = \rho'(r_k)\nabla_{\mathbf{s}} r_k.$$

Equation (2.41) then implies

$$J(\mathbf{s};\boldsymbol{\xi}_k) = \frac{\rho'(r_k)^2}{\rho(r_k)[1-\rho(r_k)]} \nabla_{\mathbf{s}} r_k \nabla_{\mathbf{s}} r_k^\top.$$

Now $\nabla_{\mathbf{s}} r_k = \dfrac{\mathbf{s} - \boldsymbol{\xi}_k}{r_k}$, and note $\boldsymbol{\xi}_k \neq \mathbf{s} \iff r_k > 0$. By definition of θ_k, $\boldsymbol{\xi}_k - \mathbf{s} = r_k \begin{bmatrix} \cos\theta_k \\ \sin\theta_k \end{bmatrix}$, and therefore $\nabla_{\mathbf{s}} r_k = -\begin{bmatrix} \cos\theta_k \\ \sin\theta_k \end{bmatrix}$ for $r_k > 0$. Thus

$$J(\mathbf{s}; \boldsymbol{\xi}_k) = \frac{\rho'(r_k)^2}{\rho(r_k)[1-\rho(r_k)]} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix}, \tag{2.43}$$

and applying (2.39) yields (2.42). □

We now derive conditions for an optimal geometry, under the constraint that the agents be equidistant from the source location. Letting $r := r_1 = \ldots = r_N > 0$, the FIM becomes

$$\mathcal{J}(\mathbf{s}; \boldsymbol{\xi}_{1:N}) = \frac{\rho'(r)^2}{\rho(r)[1-\rho(r)]} \sum_{k=1}^{n} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix}. \tag{2.44}$$

It is clear from (2.44) that the optimal angles and radius can now be chosen independently.

Theorem 2.4.2 (Optimal Sensor Geometry)**.** *Assume $\mathbf{s} \in S \subset \mathbb{R}^2$, and let $\ell(\mathbf{s}, \boldsymbol{\xi}) = \rho(\|\mathbf{s} - \boldsymbol{\xi}\|)$, where $\rho : [0, \infty) \to (0, 1)$ is continuously differentiable. Constrain $\boldsymbol{\xi}_{1:N} \in \mathbb{R}^{2N}$ to be such that*

$$\forall k, m \in \mathbb{N}_N, \ \|\boldsymbol{\xi}_k - \mathbf{s}\| = \|\boldsymbol{\xi}_m - \mathbf{s}\| > 0. \tag{2.45}$$

Define $\theta_k := \operatorname{atan2}(\xi_{k,2} - s_2, \xi_{k,1} - s_1)$ and $r := \|\boldsymbol{\xi}_1 - \mathbf{s}\|$.

1. *For any fixed $r > 0$, $\det \mathcal{J}(\mathbf{s}; \boldsymbol{\xi}_{1:N})$ is maximised if and only if*

$$\sum_{k=1}^{N} \cos(2\theta_k) = 0 \text{ and } \sum_{k=1}^{N} \sin(2\theta_k) = 0. \tag{2.46}$$

2. *Apply the additional constraint $r \in [r_1, r_2]$ for some $0 < r_1 \leq r_2$. Then for any fixed $\theta_1, ..., \theta_N$, $\det \mathcal{J}(\mathbf{s}; \boldsymbol{\xi}_{1:N})$ is maximised if and only if*

$$r \in \arg\max_{x \in [r_1, r_2]} \frac{\rho'(x)^2}{\rho(x)[1-\rho(x)]}. \tag{2.47}$$

3. *Optimising jointly over $\theta_1, ..., \theta_N \in \mathbb{R}$ and $r \in [r_1, r_2]$, $\det \mathcal{J}(\mathbf{s}; \boldsymbol{\xi}_{1:N})$ is maximised if and only if* (2.46) *and* (2.47) *both hold.*

Proof. Under the constraint (2.45),

$$\det \mathcal{J}(\mathbf{s}; \boldsymbol{\xi}_{1:N}) = \frac{\rho'(r)^2}{\rho(r)[1-\rho(r)]} \det \left(\sum_{k=1}^{N} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix} \right)$$

$$= \frac{\rho'(r)^2}{\rho(r)[1-\rho(r)]} \det \left(\sum_{k=1}^{N} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix}^\top \right), \quad (2.48)$$

which has a form identical to the Fisher Information determinant for range-only measurements [14, Equation (13)]. It is also implied by [14, Theorem 2] that $\det \left(\sum_{k=1}^{N} \begin{bmatrix} \cos^2(\theta_k) & \frac{\sin(2\theta_k)}{2} \\ \frac{\sin(2\theta_k)}{2} & \sin^2(\theta_k) \end{bmatrix}^\top \right)$ is maximised if and only if (2.46) is satisfied, and this proves Statement 1. The coefficient $\frac{\rho'(r)^2}{\rho(r)[1-\rho(r)]}$ is a continuous function of r, and this guarantees the existence of a maximum on $[r_1, r_2]$. Statements 2 and 3 then follow immediately from (2.48). □

Remark 2.4.2. A particular type of geometry that satisfies condition (2.46) is to have the agents spaced out at equal angles about the source. This result is stated in [14, Proposition 2]. Other types geometries satisfying (2.46) can also be found in the same work.

The above result identifies optimal measurement locations with respect to the Fisher Information determinant, given the source location $\mathbf{s}$. In practice, of course, $\mathbf{s}$ is unknown. However, as emphasized at the beginning of Section 2.4, an approximate solution to (2.35) can be generated by optimising the Fisher Information determinant evaluated at the expected source location. For some $0 < r_1 \leq r_2$, let the feasible set at time n be

$$\mathcal{X}_n = \{\boldsymbol{\xi}_{1:N} \mid \forall k, m \in \mathbb{N}_N,\ \|\boldsymbol{\xi}_k - \hat{\mathbf{s}}_n\| = \|\boldsymbol{\xi}_m - \hat{\mathbf{s}}_n\| \in [r_1, r_2]\}.$$

It follows directly from Theorem 2.4.2 that an exact solution to (2.38) is given by

$$\forall k \in \mathbb{N}_N,\ \boldsymbol{\xi}_{n+k} = \hat{\mathbf{s}}_n + r \begin{bmatrix} \cos\theta_k \\ \sin\theta_k \end{bmatrix}, \quad (2.49)$$

where $r, \theta_1, ..., \theta_N$ satisfy (2.46) - (2.47).

2.5 Inexact probability of detection functions

In practice, the probability-of-detection function ℓ will not be known completely. Suppose that, instead, we have knowledge of a continuous function $\hat{\ell} : \mathbb{R}^{\mathfrak{a}} \times \mathbb{R}^{\mathfrak{a}} \to (0, 1)$ that satisfies

$$\forall \mathbf{s}, \mathbf{x} \in \mathbb{R}^{\mathfrak{a}},\ \hat{\ell}(\mathbf{s}, \mathbf{x}) \geq \ell(\mathbf{s}, \mathbf{x}). \tag{2.50}$$

That is, $\hat{\ell}$ is an envelope for ℓ. If the algorithm uses $\hat{\ell}$ in place of ℓ when computing (2.3), the corresponding version of the likelihood function is

$$\hat{g}(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k) = \hat{\ell}(\mathbf{s}, \boldsymbol{\xi}_k)^{d_k} \left[1 - \hat{\ell}(\mathbf{s}, \boldsymbol{\xi}_k)\right]^{1-d_k}. \tag{2.51}$$

By the same argument as (2.6), it is clear the convergence of the posterior depends on the ratio

$$\hat{Z}_k^{(i,j)} := \frac{\hat{g}(d_k \mid \mathbf{c}_i; \boldsymbol{\xi}_k)}{\hat{g}(d_k \mid \mathbf{c}_j; \boldsymbol{\xi}_k)}, \tag{2.52}$$

the logarithm of which has expected value

$$\hat{\mu}_k^{(i,j)} = \mathbb{E}\left[\ln \hat{Z}_k^{(i,j)} \mid \mathbf{s}\right] \tag{2.53}$$

$$= \mu(\hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k), \hat{\ell}(\mathbf{c}_j, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k)). \tag{2.54}$$

We emphasize that this expectation is taken with respect to the true distribution $g(\cdot \mid \mathbf{s}; \boldsymbol{\xi}_k)$, and remind the reader that μ is defined in (2.11). The KL divergence originally defined in (2.17) now generalizes to

$$\mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{x}, \hat{\ell}; \boldsymbol{\xi}_{1:n}) := D\left(G(\cdot \mid \mathbf{s}; \boldsymbol{\xi}_{1:n}) \,\middle\|\, \hat{G}(\cdot \mid \mathbf{x}; \boldsymbol{\xi}_{1:n})\right) \tag{2.55}$$

$$= -\sum_{k=1}^{n} \mu(\hat{\ell}(\mathbf{x}, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k)), \tag{2.56}$$

where $\hat{G}(\mathbf{d} \mid \mathbf{s}; \boldsymbol{\xi}_{1:n}) = \prod_{k=1}^{n} \hat{g}(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k)$. Using property (2.14), we then obtain

$$\sum_{k=1}^{n} \hat{\mu}_k^{(i,j)} = \sum_{k=1}^{n} \mu(\hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k), \hat{\ell}(\mathbf{c}_j, \boldsymbol{\xi}_k), \ell(\mathbf{s}, \boldsymbol{\xi}_k)) \tag{2.57}$$

$$= \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) - \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n}). \tag{2.58}$$

Remark 2.5.1. When $\hat{\ell} \neq \ell$, the convergence result in Theorem 2.3.1 holds when $\mu_k^{(i,j)}$ is replaced with $\hat{\mu}_k^{(i,j)}$. Similarly, Theorem 2.3.3 holds when $\mathcal{K}(\mathbf{s}||\mathbf{x}; \boldsymbol{\xi}_{1:n})$ is replaced with $\mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{x}, \hat{\ell}; \boldsymbol{\xi}_{1:n})$.

2.5.1 Periodic measurement locations

Once again, we restrict attention to n-periodic agent location sequences. In general, according to Theorem 2.3.3 and Remark 2.5.1, the posterior will decay to zero at every index outside the set

$$\mathcal{B}(\hat{\ell} \mid \ell) := \underset{i \in \mathbb{N}_m}{\arg\min}\, \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{x}, \hat{\ell}; \boldsymbol{\xi}_{1:n}). \tag{2.59}$$

Consider two cases. In both cases, the algorithm is run assuming the same envelope $\hat{\ell}$. The first case is a special case in which $\ell = \hat{\ell}$, and in the second case ℓ is arbitrary. Asymptotically, the support of the posterior is then contained in $\mathcal{B}(\hat{\ell} \mid \hat{\ell})$ and $\mathcal{B}(\ell \mid \hat{\ell})$ respectively. We now compare these two sets.

Lemma 2.5.1. *Suppose that,*

$$\forall k \in \mathbb{N}_n,\ \hat{\ell}(\mathbf{c}_j, \boldsymbol{\xi}_k) \geq \hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k). \tag{2.60}$$

Then

$$\begin{aligned} &\mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) < \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n}) \\ &\implies \mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) < \mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n}). \end{aligned} \tag{2.61}$$

Proof. For ease of notation, let $\hat{\omega}_k^i := \hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k)$, $z_k := \ell(\mathbf{s}, \boldsymbol{\xi}_k)$ and $\hat{z}_k := \hat{\ell}(\mathbf{s}, \boldsymbol{\xi}_k)$. Applying (2.58), we obtain

$$\begin{aligned} &\Big[\mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) - \mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n})\Big] \\ &- \Big[\mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) - \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n})\Big] \\ &= \sum_{k=1}^{n} \mu(\hat{\omega}_k^i, \hat{\omega}_k^j, \hat{z}_k) - \sum_{k=1}^{n} \mu(\hat{\omega}_k^i, \hat{\omega}_k^j, z_k) \\ &= \sum_{k=1}^{n} (\hat{z}_k - z_k) \ln \left[\frac{\hat{\omega}_k^i (1 - \hat{\omega}_k^j)}{\hat{\omega}_k^j (1 - \hat{\omega}_k^i)}\right]. \end{aligned} \tag{2.62}$$

By property (2.50), $\hat{z}_k \geq z_k$ for all k. Furthermore, the assumption (2.60) is equivalent to $\hat{\omega}_k^i \leq \hat{\omega}_k^j$ for all k, which then implies the RHS of (2.62) is non-positive. The result then follows. □

Proposition 2.5.2. *Let $(\boldsymbol{\xi}_k)_{k\in\mathbb{N}}$ be an n-periodic sequence such that,*

$$\begin{aligned} &\forall i \ \in \mathbb{N}_M,\ \exists j \in \mathcal{B}(\ell \mid \hat{\ell}),\ \forall k \in \mathbb{N}_n, \\ &\hat{\ell}(\mathbf{c}_j, \boldsymbol{\xi}_k) \ \geq \hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k). \end{aligned} \tag{2.63}$$

Then $\mathcal{B}(\hat{\ell} \mid \hat{\ell}) \subset \mathcal{B}(\ell \mid \hat{\ell})$.

Proof. Suppose $i \in \mathbb{N}_M \setminus \mathcal{B}(\ell \mid \hat{\ell})$. Then by assumption (2.63),

$$\exists j \in \mathcal{B}(\ell \mid \hat{\ell}),\ \forall k \in \mathbb{N}_n,\ \hat{\ell}(\mathbf{c}_j, \boldsymbol{\xi}_k) \geq \hat{\ell}(\mathbf{c}_i, \boldsymbol{\xi}_k).$$

Definition (2.59) implies,

$$\mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) < \mathcal{K}(\mathbf{s}, \ell \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n}).$$

We can then apply Lemma 2.5.1 to obtain

$$\mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_j, \hat{\ell}; \boldsymbol{\xi}_{1:n}) < \mathcal{K}(\mathbf{s}, \hat{\ell} \mid\mid \mathbf{c}_i, \hat{\ell}; \boldsymbol{\xi}_{1:n}),$$

which implies $i \notin \mathcal{B}(\hat{\ell} \mid \hat{\ell})$. We have shown $\mathcal{B}(\ell \mid \hat{\ell})^{\complement} \subset \mathcal{B}(\hat{\ell} \mid \hat{\ell})^{\complement}$, which is equivalent to the result. □

Remark 2.5.2. Proposition 2.5.2 can be interpreted as follows. Assuming the constraints (2.63) on $\boldsymbol{\xi}_k$ are satisfied, as long as $\hat{\ell}$ remains an envelope for ℓ, the limiting support of the posterior cannot shrink compared to the case where $\ell = \hat{\ell}$. Thus, the algorithm behaves conservatively in the limit. The result holds even if ℓ is time-varying.

Remark 2.5.3. When the envelope $\hat{\ell}$ is strictly decreasing with distance from the source, the assumption (2.63) requires all the agents to be closer to some centre in $\mathcal{B}(\ell \mid \hat{\ell})$ than to the other centres. This suggests a sensible strategy would be to drive the agents towards the current MAP estimate (i.e. a maximiser of the posterior).

2.6 Numerical example

In this section, we examine a concrete example involving the 2D localisation of an electromagnetic source. We present simulation results to supplement the analytical results of the previous sections, and to demonstrate the effectiveness of the Bayesian estimation algorithm when used in the loop with

a control law that guides the agents towards D-optimal measurement locations. Consider a source antenna located at ground level, transmitting an RF signal of wavelength λ, with input power P_T and effective area A_T. Suppose each agent is a UAV equipped with a receiving antenna of effective area A_R. Given source location $\left[\begin{smallmatrix}\mathbf{s}\\0\end{smallmatrix}\right] \in \mathbb{R}^3$, the power received by an agent at location $\left[\begin{smallmatrix}\mathbf{x}\\z\end{smallmatrix}\right] \in \mathbb{R}^3$ can be modelled by the Friis transmission formula [50]:

$$P_R(\mathbf{s}, \mathbf{x}) = \frac{A_R A_T P_T}{\lambda^2(\|\mathbf{s}-\mathbf{x}\|^2+z^2)}.$$

We simulate a team of four agents, constrained to fly at a constant altitude $z > 0$, with positions in the plane that evolve according to

$$\dot{\mathbf{x}}_i(t) = \mathbf{u}_i(t), \tag{2.64}$$

where $\mathbf{u}_i : [0, \infty) \to \mathbb{R}^2$ is the control signal applied to agent i. An agent at location $\boldsymbol{\xi}_k \in \mathbb{R}^2$ reports a binary measurement d_k by comparing the received power measured at time t_k with a threshold $\eta > 0$ according to

$$d_k = \begin{cases} 0, & P_R(\mathbf{s}, \boldsymbol{\xi}_k) + W_k < \eta \\ 1, & P_R(\mathbf{s}, \boldsymbol{\xi}_k) + W_k \geq \eta \end{cases},$$

where $W_k \sim \mathcal{N}(0, \sigma^2)$ accounts for sensor noise. The probability-of-detection function is therefore given by

$$\ell(\mathbf{s}, \mathbf{x}) = \mathcal{Q}\left(\frac{\eta - P_R(\mathbf{s}, \mathbf{x})}{\sigma}\right),$$

where

$$\mathcal{Q}(x) := \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{u^2}{2}} du$$

is the Q-function.

Suppose there is a maximum transmission delay of $\frac{\tau}{2} \geq 0$ seconds between the agents and the fusion centre. We ignore the effects of packet drop, and constrain the agents to take measurements synchronously every $T > \tau$ seconds. Recalling that subscripts are assigned according to the order in which the measurements arrive at the fusion centre, this implies

$$0 \leq t_1 = ... = t_4 < t_5 = t_6 = ...,$$

where $t_{k+4} - t_k = T$. The fusion centre processes all four measurements pairs $(\boldsymbol{\xi}_k, d_k), ..., (\boldsymbol{\xi}_{k+3}, d_{k+3})$ within the interval $[t_k, t_k + \frac{\tau}{2}]$. It then computes the posterior mean

$$\bar{\mathbf{s}}_k := \sum_{j=1}^{M} \hat{p}_k(i \mid d_{1:k}; \boldsymbol{\xi}_{1:k}) \mathbf{c}_j,$$

which it transmits back to the agents. All agents then receive $\bar{\mathbf{s}}_k$ no later than $t_k + \tau$, at which time they synchronously update their local copies of the mean $\hat{\mathbf{s}}_i(t_k + \tau) = \bar{\mathbf{s}}_k$. Thus, these local copies evolve in continuous-time according to

$$\hat{\mathbf{s}}_i(t) := \bar{\mathbf{s}}_{\kappa(t-\tau)},$$

where $\kappa : [0, \infty) \to \mathbb{N},\ \kappa(t) := \max\{k \mid t_k \leq t\}$. The posterior mean is used as an input to the controller proposed below:

$$\mathbf{u}_i(\mathbf{x}_i, \hat{\mathbf{s}}_i) = -\left(\mathbf{x}_i - \hat{\mathbf{s}}_i - \boldsymbol{\delta}_i\right) \tag{2.65}$$

where $\boldsymbol{\delta}_i := r \begin{bmatrix} \cos\theta_i \\ \sin\theta_i \end{bmatrix}$ is chosen according to (2.46) - (2.47). This drives the agents towards the optimal locations dictated by (2.49).

Remark 2.6.1. The control signal generated by (2.65) is a piecewise continuous function of time. Boundedness can also be guaranteed by imposing a saturation function, without fundamentally altering its behaviour.

For the simulations below, the search region S is a 75 m $\times$ 75 m planar region (at ground level), and the source location is sampled from the uniform distribution over S. We choose M centres, aligned in a uniform grid over a 100 m $\times$ 100 m region containing S at the centre. Parameter values are $A_R = A_T = 1$ m^2, $P_T = 1$ W, $\lambda = 1$ m, $z = 10$ m, $\eta = 5 \times 10^{-3}$ W, $\sigma = 2.5 \times 10^{-3}$ W, $T = 0.04$ s and $\tau = 0.02$ s, unless otherwise stated. Each agent was initialized as shown in Figure 2.1. A uniform prior for the source location was used to initialize the Bayesian updates.

To numerically examine the effects of discretisation on estimation performance, for every $M \in \{10^2, 20^2, ..., 50^2\}$ we run 100 Monte Carlo trials and compute the RMS estimation error e_k by averaging $\|\bar{\mathbf{s}}_k - \mathbf{s}\|$. A supplementary animation of a trial with $M = 30^2$ is available at `https://youtu.be/l8AwfOKCt4s`. The results are plotted in Figure 2.2a for up to $k = 1000$ measurements. While this captures the evolution of e_k over a finite time horizon, it does not necessarily reflect the asymptotic error e_∞ after the posterior converges to its steady-state. In practice, the entropy h_k of the posterior is a good indicator of convergence. We approximately evaluate e_∞ by computing the RMS error at $k = 1000$, averaging only the trials for which $h_{1000} < 1$ nat. These results are recorded in Table 2.1, and we observe that e_∞ decreases monotonically with grid spacing.

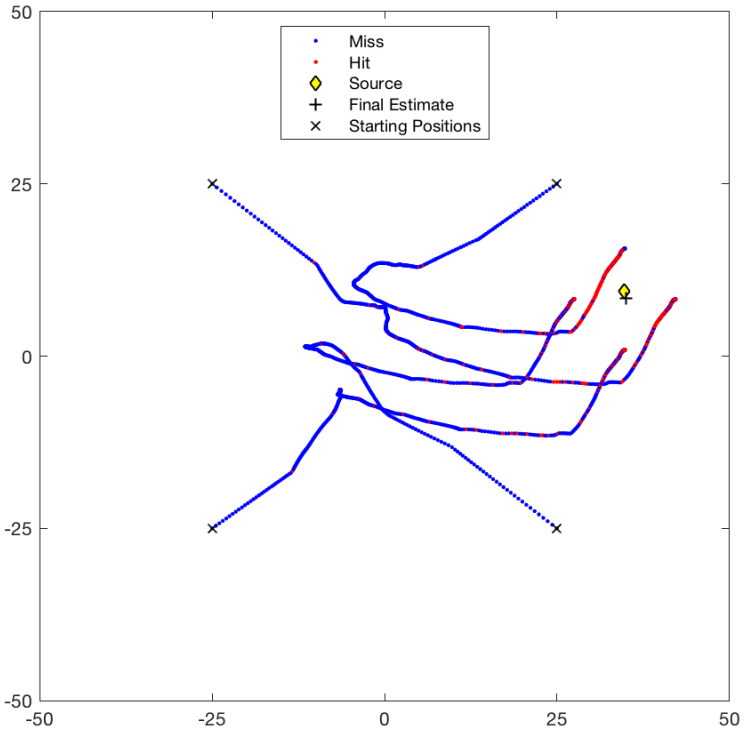

Figure 2.1: Sample history of measurement pairs ($M = 30^2$, units: m).

Table 2.1: Asymptotic estimation error.

$M =$	10^2	20^2	30^2	40^2	50^2
Grid Spacing (m)	10	5	3.33	2.5	2
Approx. e_∞ (m)	3.95	1.96	1.43	1.04	0.84
Trials with $h_{1000} < 1$	99 %	97 %	96 %	97 %	98 %

2.6.1 Computational complexity

Before proceeding to comparisons with other methods, we pause to consider the computational complexity of our estimation procedure. Suppose each evaluation of the likelihood function g takes at most L flops. Each iteration of (2.3) then takes $M(L+3)-1$ flops. If B bits of memory are required to store each number, then $MB\mathfrak{a}$ bits are required to store the centres, and MB bits to store $\hat{p}_k$. Finally, $B\mathfrak{a}+1$ bits must be transmitted to $N-1$ agents per measurement to encode the measurement pair $(d_k, \boldsymbol{\xi}_k)$.

2.6.2 Comparison with ML estimators

As mentioned in the introduction, we employ the same Bayesian estimation algorithm adopted in [4, 5, 37]. Here, we compare its performance with the maximum likelihood (ML) approach of [9]. The complexity of the full ML estimator in [9] is $O(\mathfrak{a}^3 k^2)$ for the kth measurement. The required memory grows linearly with k. Since this can be impractical for real-time processing, [9] also proposes a real-time approximation, which has complexity $O(\mathfrak{a}^3)$ per measurement and requires constant memory. In contrast, Section 2.6.1 establishes the complexity and memory requirements of our Bayesian method as both of $O(M)$, constant with respect to the number of measurements. Recall that M is the number of centres.

The RMS error obtained by implementing the full ML estimator in [9] is plotted in Figure 2.2b. This ML estimator is based on the Newton method, and since the log-likelihood function for this problem is non-concave, there are no convergence guarantees as $k \to \infty$. In the simulations, five Newton iterations are performed per measurement. To facilitate a meaningful comparison between the two approaches, the control-law (2.65) is also employed to direct the agents under the ML approach, but with the posterior mean $\hat{\mathbf{s}}$ replaced by the ML estimate.

Note that tracking the maximum likelihood estimate is precisely a time-

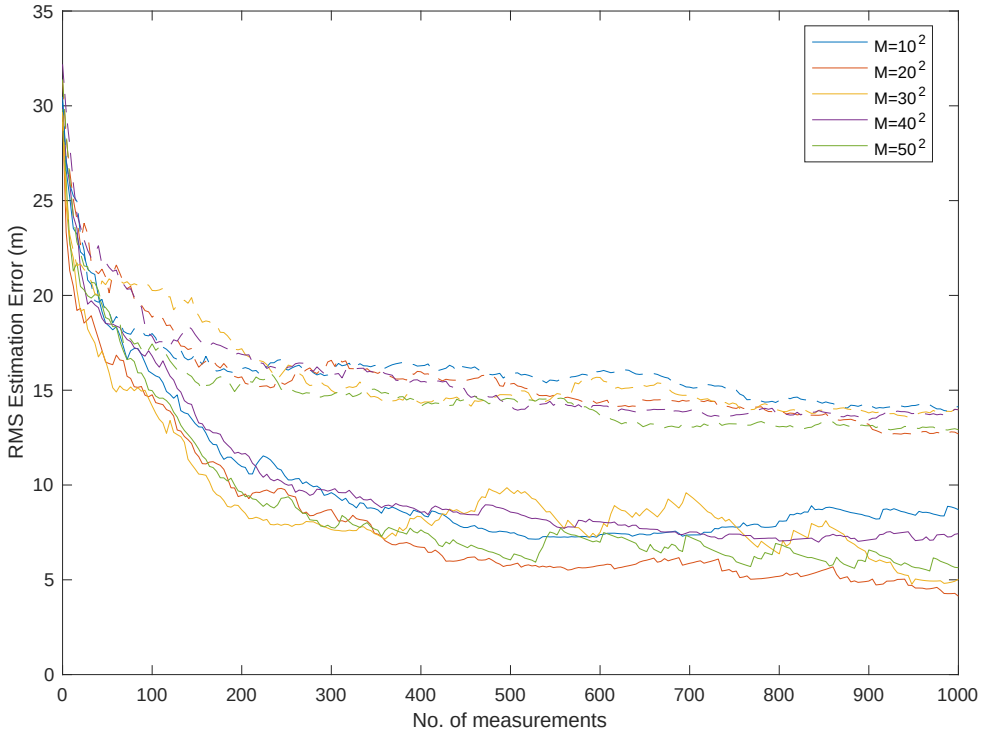

(a) Bayesian algorithm. (Solid lines = with controller. Dashed lines = without controller.)

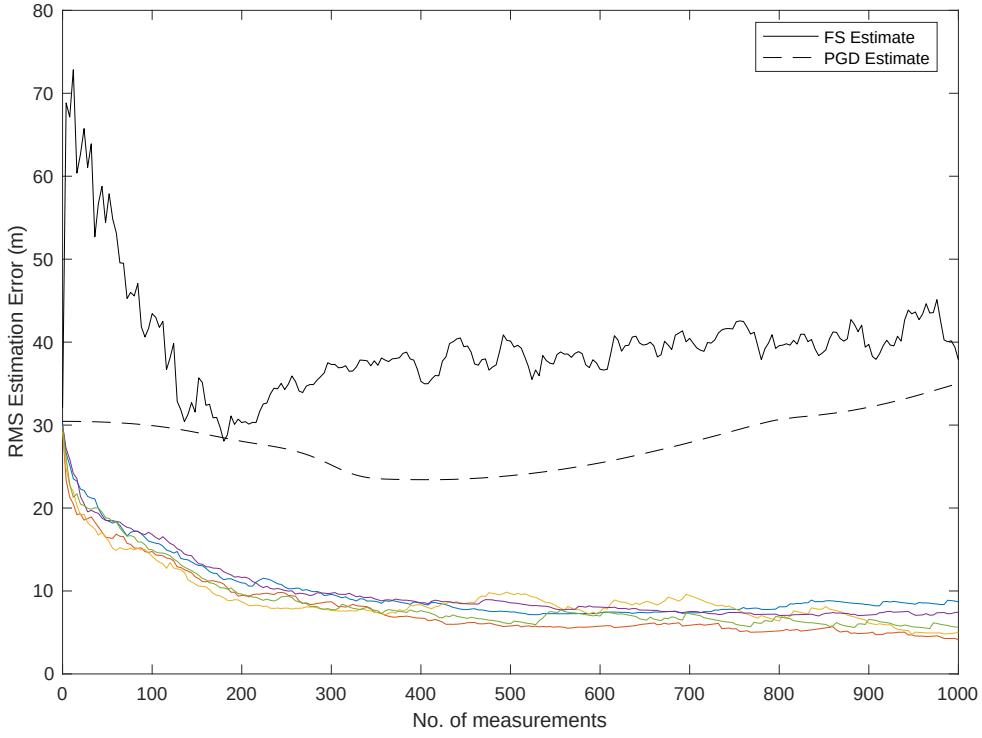

(b) Bayesian vs ML approach.

Figure 2.2: Estimator performance comparison

varying optimisation problem, of the type described in Section 1.4. For the sake of comparison, we have also implement projected gradient descent (PGD) to compute the ML estimate, using only one gradient iteration per new measurement. Here, the complexity and required memory both increase linearly with k. Results are presented in Figure 2.2b. The PGD algorithm is analysed in greater detail in Chapter 5.

2.6.3 Effect of control strategy

We now examine the effectiveness of the control strategy (2.65), which drives the agents into the D-optimal geometries defined by (2.49). In Figure 2.2a, the RMS estimation error of the Bayesian algorithm with the control-law in the loop is plotted against a scenario in which the agents remain fixed at their initial positions (which are evenly spaced throughout the environment). The motion of the agents significantly increases the rate at which the estimation error decays. A sample trajectory induced by the control law is plotted in Figure 2.1, along with the full history of measurement pairs $(\boldsymbol{\xi}_k, d_k)$.

2.6.4 Effect of inexact knowledge

To examine the effect of inexact knowledge of ℓ, we fix the assumed value of the transmitted power at $\hat{P}_T = 5W$, and vary the true value P_T between 1 W and 5 W. Thus the assumed probability-of-detection function $\hat{\ell}$ remains an envelope for ℓ. We use $M = 20^2$ grid points for the simulations. Consistent with the strategy proposed in Remark 2.5.3, we modify the control law (2.65) by replacing $\hat{\mathbf{s}}$ with the MAP estimate. This time, the angles θ_i are chosen according to (2.46), but we set $r = 2.5$ m so that the agents are driven to converge to points that are closer to the MAP estimate than the other centres. The results of 100 Monte Carlo trials are plotted in Figure 2.3. We observe a graceful degradation in RMS estimation error for $P_T \geq 3$ W, but the estimator ceases to be effective when the transmitted power falls to 2 W or less.

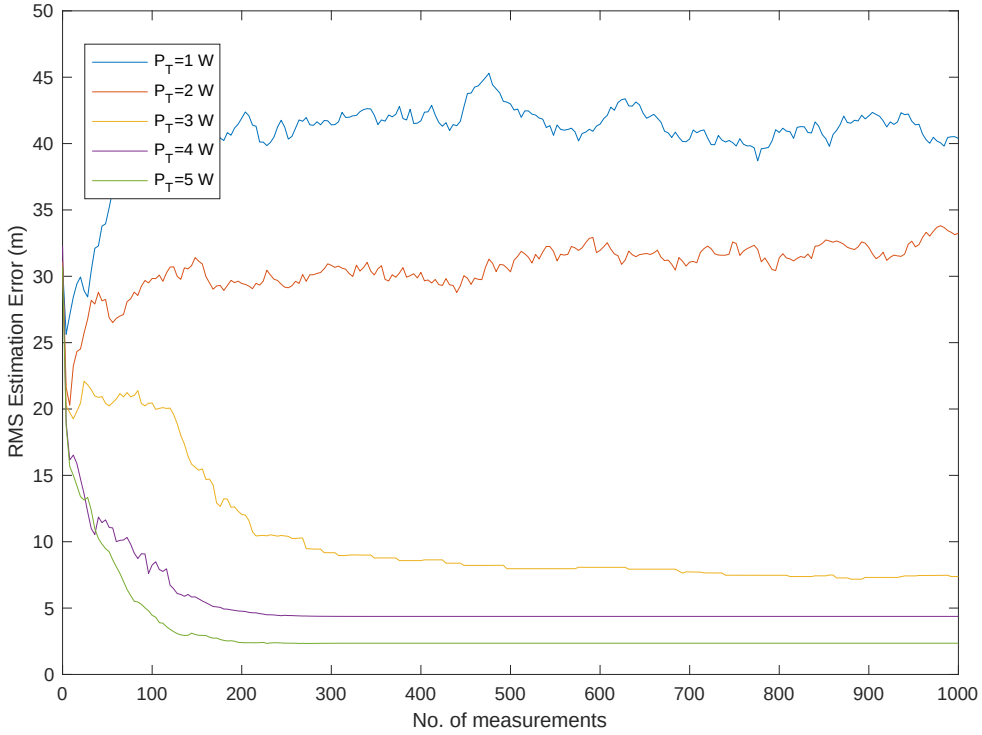

Figure 2.3: Effect of changing P_T. ($\hat{P}_T = 5$ W, $M = 20^2$)

Chapter 3

Security Analysis of Bayesian Estimators

> "I bid you come out before your doors and look abroad. Too long have you sat in the shadows and trusted to twisted tales and crooked promptings."
>
> Gandalf the Grey

As we saw in the previous chapter, a team of autonomous agents may be deployed to tackle co-operative estimation problems. Bayesian estimators can combine measurements from individual agents to generate a posterior probability distribution over possible values of the estimand. This chapter presents a security analysis of Bayesian estimators with finite parameter and observation spaces. We proceed from the perspective of an adversary that has 'hacked' one of the agents, and can feed false measurement data to the estimator. Assuming that the estimator is oblivious to the adversary, we derive the set of posteriors that can be induced by the adversary in expectation. A strategy for the adversary is also proposed, based on the posterior it desires to induce. The chapter begins with a brief review of the relevant literature in Section 3.1. The problem is formulated mathematically in Section 3.2, Section 3.3 derives the theoretical results, and simulation results are then presented in Section 3.4.

3.1 Related works

The security analysis herein has close connections to the signal processing literature on Byzantine attacks. The papers [51–55] all consider estimation problems with multiple sensors providing quantized measurements to an estimator. Some proportion of these sensors have been compromised, and traditionally, attention has focused on 'blinding' the estimator. That is, crafting attacks that render the set of measurements uninformative. The works [53–55] also consider possible mitigation strategies for the estimator to adopt. An analysis of fundamental limits on the set of posteriors that can be induced in a Bayesian estimator is missing from this body of literature, and we have endeavoured to fill this gap. Chapter 3 also relates to research on the security of cyber-physical systems. This ever-expanding field considers the impact of security flaws in networked devices on the underlying system dynamics. In particular, [56–58] study attacks on discrete-time linear dynamical systems from a control theoretic viewpoint. Finally, the economics literature on Bayesian persuasion [22, 59, 60] and cheap talk [61, 62] is also pertinent. There, a well-informed sender passes messages to an uninformed estimator with the intention of manipulating its decision. Those studies differ from ours by assuming that the estimator is always aware of the presence of the adversary, significantly restricting its ability to deceive.

3.2 Problem formulation

Let $\mathbb{P}\{\cdot\}$ denote the probability of an event. Consider the problem of computing the posterior distribution of a random variable X that takes values in a finite set $\mathbb{X}$, given information received in some message M that takes values in a finite set $\mathbb{M}$. For example, $\mathbb{X}$ may contain all the possible locations of a signal source, and $\mathbb{M}$ may be a binary set corresponding to whether the measured signal strength exceeds some threshold. The prior distribution of X is given by $p_X : \mathbb{X} \to [0, 1]$, where $p_X(x) := \mathbb{P}\{X = x\}$. The estimation objective is to compute the posterior $p_{X|M}(x|m) := \mathbb{P}\{X = x \,|\, M = m\}$. The probability distribution of the message conditioned on X is assumed to be $p_{M|X}(m|x)$ by the estimator. It accordingly follows Bayes' rule to compute the posterior:

$$p_{X|M}(x|m) = \frac{p_{M|X}(m|x)p_X(x)}{\sum_{x' \in \mathbb{X}} p_{M|X}(m|x')p_X(x')}. \tag{3.1}$$

Note that, since the message M is a random variable, the value of the conditional density function $p_{X|M}(x|M)$ at a given $x \in \mathbb{X}$ is itself a random variable.

Covert adversary

Consider the case where the adversary injects a message $M = m$, but because the estimator is unaware of this attack, it mistakenly follows (3.1). Can the adversary ensure the posterior generated by (3.1) matches a posterior that it desires to induce? The set of possible posteriors is given by $\mathcal{P} := \{p_m \mid m \in \mathbb{M}\}$, where $p_m : \mathbb{X} \to [0, 1]$,

$$p_m(x) := p_{X|M}(x|m). \tag{3.2}$$

The adversary may adopt a deterministic strategy that transmits m in order to induce $p_m \in \mathcal{P}$. However, the set cardinality $|\mathcal{P}| \leq |\mathbb{M}|$. It is therefore of interest to determine whether, by adopting a stochastic policy, the adversary can induce a larger set of posteriors in expectation. A deterministic policy is then a special case of this, with the relevant probabilities set to unity. Consider a family of functions $\{\pi_x\}_{x\in\mathbb{X}}$, where each π_x is the posterior the adversary seeks to induce when $X = x$.

Assumption 3.1 (Family of desired posteriors)**.** Let $\pi_x : \mathbb{X} \to [0, 1]$ satisfy $\sum_{x'\in\mathbb{X}} \pi_x(x') = 1$ for all $x \in \mathbb{X}$.

Assuming the adversary knows X, suppose it transmits stochastic messages M according to $\bar{p}_{M|X}(m|X)$.

Problem 3.1. Find a probability distribution $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0, 1]$ such that

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_M\{p_{X|M}(x'|M) \mid X = x\} = \pi_x(x'). \tag{3.3}$$

Note that, in the above equation, x' corresponds to a dummy variable for evaluating $\pi_x(x')$, and x corresponds to possible realisations of the random variable X.

Side-channel information

It is often the case that the estimator has access to some reliable side information, such as the measurements from honest sensors. A random variable $N \in \mathbb{M}$ with conditional probability function $p_{N|X}(n|x) := \mathbb{P}\{N = n \,|\, X = x\}$ is

used to show this side-channel information. It is assumed that, conditioned on the realisation of X, messages M and N are statistically independent, i.e., $\mathbb{P}\{M = m, N = n|X = x\} = \mathbb{P}\{M = m|X = x\}\mathbb{P}\{N = n|X = x\}$. In this case, the update rule in (3.1) must be adapted to

$$p_{X|M,N}(x|m,n) = \frac{p_{N|X}(n|x)p_{M|X}(m|x)p_X(x)}{\sum_{x'\in\mathbb{X}} p_{N|X}(n|x')p_{M|X}(m|x')p_X(x')}. \tag{3.4}$$

As before, the adversary may follow the conditional distribution $\bar{p}_{M|X}(m|X)$ to generate its messages, while the estimator unknowingly follows (3.4). This leaves us with the following problem.

Problem 3.2. Find a probability distribution $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0,1]$ such that

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_{M,N}\{p_{X|M,N}(x'|M,N) \mid X = x\} = \pi_x(x'). \tag{3.5}$$

With these problem formulations at hand, we are ready to present the main results.

3.3 Fundamental limitations

Without loss of generality, suppose $\mathbb{X} = \{1, ..., \mathfrak{a}\}$ and $\mathbb{M} = \{1, ..., \mathfrak{b}\}$. We now express the relevant marginal and conditional probabilities of the problem as vectors and matrices respectively. Let $\mathbf{x} \in \mathbb{R}^{\mathfrak{a}}$, $\mathbf{M}, \overline{\mathbf{M}}, \mathbf{N} \in \mathbb{R}^{\mathfrak{b}\times\mathfrak{a}}$, $\mathbf{\Pi} \in \mathbb{R}^{\mathfrak{a}\times\mathfrak{a}}$ and $\mathbf{P} \in \mathbb{R}^{\mathfrak{a}\times\mathfrak{b}}$, and define

$$\begin{aligned} &[\mathbf{x}]_i := p_X(i), && [\mathbf{M}]_{ij} := p_{M|X}(i|j), && [\overline{\mathbf{M}}]_{ij} := \bar{p}_{M|X}(i|j), \\ &[\mathbf{\Pi}]_{ij} := \pi_j(i), && [\mathbf{N}]_{ij} := p_{N|X}(i|j), && [\mathbf{P}]_{ij} := p_{X|M}(i|j), \end{aligned} \tag{3.6}$$

where $[\cdot]_i$ denotes the ith element of a vector, and $[\cdot]_{ij}$ the ijth element of a matrix.

3.3.1 Covert adversary

We first address Problem 3.1. The estimator's update rule (3.1) can be written as

$$[\mathbf{P}]_{ij} = \frac{[\mathbf{M}]_{ji}[\mathbf{x}]_i}{[\mathbf{M}\mathbf{x}]_j}. \tag{3.7}$$

Let $\mathrm{supp}(\cdot)$ denote the set-theoretic support of a function.

Proposition 3.3.1. *There exists a probability distribution* $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0, 1]$ *for which*

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_M\{p_{X|M}(x'|M)|X = x\} = \pi_x(x'), \tag{3.3}$$

if and only if *the following linear equation in* $\overline{\mathbf{M}}$ *admits a non-negative solution:*

$$\begin{bmatrix} \mathbf{P} \\ \mathbf{1}_{\mathfrak{b}}^\top \end{bmatrix} \overline{\mathbf{M}} = \begin{bmatrix} \mathbf{\Pi} \\ \mathbf{1}_{\mathfrak{a}}^\top \end{bmatrix}. \tag{3.8}$$

Proof.

$$\begin{aligned} \mathbb{E}_M\{p_{X|M}(i|M)|X = j\} &= \sum_{m \in \mathbb{M}} p_{X|M}(i|m)\bar{p}_{M|X}(m|j) \\ &= \sum_{m \in \mathbb{M}} [\mathbf{P}]_{im}[\overline{\mathbf{M}}]_{mj} = [\mathbf{P}\overline{\mathbf{M}}]_{ij}. \end{aligned} \tag{3.9}$$

Recalling (3.6) then yields the first row of (3.8). The second row is equivalent to the requiring $\sum_{m \in \mathbb{M}} \bar{p}_{M|X}(m|x) = 1$ for all $x \in \mathbb{X}$. □

Corollary 3.3.2. *The linear equation in Proposition 3.3.1 admits a solution only if* $\mathrm{supp}(\pi_x) \subseteq \mathrm{supp}(p_X)$ *for all* $x \in \mathbb{X}$.

Proof. If $[\mathbf{x}]_i = 0$, then (3.7) implies $[\mathbf{P}]_{ij} = 0$ for all $j \in \mathbb{X}$. If (3.8) holds, then $[\mathbf{\Pi}]_{ij} = 0$ for all $j \in \mathbb{X}$. □

If a solution to (3.8) exists, then $\overline{\mathbf{M}}$ provides the conditional probability distribution that the adversary should employ to generate messages. Corollary 3.3.2 states the obvious result that it is not possible to fool the Bayesian estimator to believe that an event with zero probability has happened.

Remark 3.3.1. The number of scalar linear equations in (3.8) is $\mathfrak{a}(\mathfrak{a}+1)$, while the number of unknowns is $\mathfrak{a}\mathfrak{b}$. Thus, intuitively, having a large observation space makes it easier to deceive the estimator.

The set of posteriors that can be induced in expectation via some stochastic policy can now be stated in terms of the convex hull, $\mathrm{conv}(\cdot)$.

Theorem 3.3.3. *There exists a probability distribution* $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0, 1]$ *for which*

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_M\{p_{X|M}(x'|M)|X = x\} = \pi_x(x'), \tag{3.3}$$

if and only if *every column of* $\mathbf{\Pi}$ *lies in the convex hull of the columns of* $\mathbf{P}$.

Proof. Let $\mathbf{p}_i$, $\boldsymbol{\pi}_i$ and $\bar{\mathbf{m}}_i$ denote the ith columns of $\mathbf{P}$, $\boldsymbol{\Pi}$ and $\overline{\mathbf{M}}$ respectively. Then (3.8) can be written as

$$\forall i \in \mathbb{X},\ \mathbf{P}\bar{\mathbf{m}}_i = \boldsymbol{\pi}_i,\ \mathbf{1}_{\mathfrak{b}}^\top \bar{\mathbf{m}}_i = 1,\ \bar{\mathbf{m}}_i \geq 0.$$

By the definition of convex hull [63, Section 2.1.4], $\boldsymbol{\pi}_i \in \text{conv}\{\mathbf{p}_1, ..., \mathbf{p}_{\mathfrak{b}}\}$ if and only if there exist $\theta_1, ..., \theta_{\mathfrak{b}} \geq 0$ such that $\bar{\mathbf{m}}_i = (\theta_1, ..., \theta_{\mathfrak{b}})$ satisfies the above equations. □

Remark 3.3.2. Each column $\boldsymbol{\pi}_i$ represents a desired posterior π_i. Each column $\mathbf{p}_j$ represents $p_j \in \mathcal{P}$, defined in (3.2). Thus, Theorem 3.3.3 states that the set of posteriors that can be induced in expectation via some $\bar{p}_{M|X}$ is conv$\mathcal{P}$. Clearly the closest achievable posterior to some desired π_x is its projection onto conv$\mathcal{P}$.

3.3.2 Side-channel information

Now, consider Problem 3.2. Receiving the messages N and M can be treated as receiving a single message $Z \in \mathbb{M}^2$. A matrix representation of $p_{Z|X}$ (as assumed by the estimator) is

$$\mathbf{Z} := \begin{bmatrix} \mathbf{n}_1 \otimes \mathbf{m}_1 & \mathbf{n}_2 \otimes \mathbf{m}_2 & \dots & \mathbf{n}_{\mathfrak{a}} \otimes \mathbf{m}_{\mathfrak{a}} \end{bmatrix} \in \mathbb{R}^{\mathfrak{b}^2 \times \mathfrak{a}},$$

where $\mathbf{n}_i$ and $\mathbf{m}_i$ are the columns of $\mathbf{N}$ and $\mathbf{M}$ respectively, and $\otimes$ is the Kronecker product. Each $(m, n) \in \mathbb{M}^2$ corresponds to some row k of $\mathbf{Z}$:

$$[\mathbf{Z}]_{ki} = [\mathbf{N}]_{ni}[\mathbf{M}]_{mi} = p_{N|X}(n|i)p_{M|X}(m|i). \tag{3.10}$$

Just as in (3.7), the update rule (3.4) can now be written as

$$[\mathbf{Q}]_{ik} := \frac{[\mathbf{Z}]_{ki}[\mathbf{x}]_i}{[\mathbf{Z}\mathbf{x}]_k} = p_{X|M,N}(i|m, n), \tag{3.11}$$

where $\mathbf{Q} \in \mathbb{R}^{\mathfrak{a} \times \mathfrak{b}^2}$. We now state an analogous result to Proposition 3.3.1.

Proposition 3.3.4. *There exists a probability distribution* $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0, 1]$ *for which*

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_{M,N}\{p_{X|M,N}(x'|M, N) \mid X = x\} = \pi_x(x'), \tag{3.5}$$

if and only if *the following linear equations in* $\{\bar{\mathbf{m}}_j\}_{j \in \mathbb{X}}$ *admit non-negative solutions:*

$$\forall j \in \mathbb{X},\ \mathbf{Q}(\mathbf{n}_j \otimes \mathbf{I}_{\mathfrak{b}})\bar{\mathbf{m}}_j = \boldsymbol{\pi}_j \tag{3.12}$$

$$\forall j \in \mathbb{X},\ \mathbf{1}_{\mathfrak{b}}^\top \bar{\mathbf{m}}_j = 1. \tag{3.13}$$

Proof. First recall there is a bijection $\sigma : \mathbb{M}^2 \to \{1, ..., \mathfrak{b}^2\}$ such that if $\sigma(m, n) = k$,

$$[\mathbf{N}]_{nj}[\overline{\mathbf{M}}]_{mj} = p_{N|X}(n|j)\bar{p}_{M|X}(m|j) = [\mathbf{n}_j \otimes \bar{\mathbf{m}}_j]_k.$$

Furthermore, $\mathbf{n}_j \otimes \bar{\mathbf{m}}_j = (\mathbf{n}_j \otimes \mathbf{I}_\mathfrak{b})\bar{\mathbf{m}}_j$. Together with (3.11), this implies

$$\begin{aligned}[\mathbf{Q}(\mathbf{n}_j \otimes \mathbf{I}_\mathfrak{b})\bar{\mathbf{m}}_j]_i = [\mathbf{Q}(\mathbf{n}_j \otimes \bar{\mathbf{m}}_j)]_i &= \sum_{k=1}^{\mathfrak{b}^2} [\mathbf{Q}]_{ik}[\mathbf{n}_j \otimes \bar{\mathbf{m}}_j]_k \\ &= \sum_{m,n \in \mathbb{M}} p_{X|M,N}(i|m, n)p_{N|X}(n|j)\bar{p}_{M|X}(m|j) \\ &= \mathbb{E}_{M,N}\{p_{X|M,N}(i|M, N) \mid X = j\}.\end{aligned}$$

The desired value for this is $[\boldsymbol{\pi}_j]_i$, which leads to (3.12). The requirement $\sum_{m \in \mathbb{M}} \bar{p}_{M|X}(m|j)$ for all j leads to (3.13). □

Applying same reasoning as in Theorem 3.3.3 leads us to another corresponding result.

Theorem 3.3.5. *There exists a probability distribution $\bar{p}_{M|X} : \mathbb{M} \times \mathbb{X} \to [0, 1]$ for which*

$$\forall x, x' \in \mathbb{X},\ \mathbb{E}_{M,N}\{p_{X|M,N}(x'|M, N) \mid X = x\} = \pi_x(x'), \tag{3.5}$$

if and only if, for all $j \in \mathbb{X}$, $\boldsymbol{\pi}_j$ lies in the convex hull of the columns of $\mathbf{Q}(\mathbf{n}_j \otimes \mathbf{I}_\mathfrak{b})$.

Remark 3.3.3. In Theorem 3.3.3, the set of posteriors achievable in expectation (conv$\mathcal{P}$) is independent of the realisation of X. In Theorem 3.3.5 this is no longer the case, because the estimator receives information about X from N. Thus, when $X = j$, then the set of posteriors achievable in expectation corresponds to the convex hull of the columns of $\mathbf{Q}(\mathbf{n}_j \otimes \mathbf{I}_\mathfrak{b})$.
Remark 3.3.4. Since $\mathbf{Q}(\mathbf{n}_j \otimes \mathbf{I}_\mathfrak{b}) \in \mathbb{R}^{\mathfrak{a} \times \mathfrak{b}}$ and $\bar{\mathbf{m}}_i \in \mathbb{R}^\mathfrak{b}$, (3.12) - (3.13) express $\mathfrak{a}(\mathfrak{a} + 1)$ equations in $\mathfrak{a}\mathfrak{b}$ unknowns. Thus, the addition of the side-channel information does not add to the complexity of crafting an attack.

The intuition behind Remark 3.3.4 is that the estimator can first update its original prior based on the side channel information, and then processes the adversary's measurement. The posterior of the estimator based on the side-channel information alone can therefore be treated as a new prior for the attack.

3.3.3 Attack strategies

To provide a complete answer to Problems 3.1 and 3.2, we must now identify the adversarial distribution $\bar{p}_{M|X}$ that will induce the desired family of posteriors. Theorems 3.3.3 and 3.3.5 show that an arbitrary π_x may not be achievable. However, the adversarial distribution inducing the projection of π_x onto the set of achievable posteriors can be generated by solving

$$\forall i \in \mathbb{X},\ \min_{\bar{\mathbf{m}}_i} \|\mathbf{A}_i \mathbf{m}_i - \boldsymbol{\pi}_i\| \quad \text{s.t.} \quad \begin{cases} \mathbf{1}_{\mathfrak{b}}^\top \bar{\mathbf{m}}_i = 1 \\ \bar{\mathbf{m}}_i \geq 0 \end{cases}, \tag{3.14}$$

where $\mathbf{A}_i \in \{\mathbf{P}, \mathbf{Q}(\mathbf{n}_i \otimes \mathbf{I}_{\mathfrak{b}})\}$ depends on the presence of side-channel information, and $\|\cdot\|$ is any norm of choice. These are a collection of $\mathfrak{a}$ convex problems, of state dimension $\mathfrak{b}$. They can all be solved offline by the adversary *a-priori*. Once the adversary knows the realisation of X, only the solution corresponding to $i = X$ need be implemented. Choosing the Euclidean norm yields quadratic programs, but in general other convex functions of $\bar{\mathbf{m}}_i$ (e.g. KL divergence) may be adopted.

3.4 Numerical example

We now present a numerical example, based on the source localisation problem of Chapter 2. For simplicity, we consider two stationary agents located at $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2 \in \mathbb{R}^2$. We adopt the measurement model of Section 2.6. A fusion centre (the estimator) receives a sequence of 5 binary measurements from each sensor. Starting with a uniform prior, the estimator uses the sequential Bayesian update rule (2.3) to generate a posterior over a discretised search region $\mathbb{X} = \mathcal{C} := \{-48, -44, \ldots, 44, 48\}^2$. Let $N, L \in \{0,1\}^5$ correspond to honest measurements from $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2$ respectively. That is,

$$N := (n_1, ..., n_5) \sim \prod_{k=1}^{5} g(n_k \mid X, \boldsymbol{\xi}_1), \quad L := (l_1, ..., l_5) \sim \prod_{k=1}^{5} g(l_k \mid X, \boldsymbol{\xi}_2),$$

where $X = \mathbf{s} \in \mathbb{X}$ is the source location and g is defined in (2.1). Now, suppose that $\boldsymbol{\xi}_2$ has been compromised by an adversary, and seeks to make the estimator believe the source location is $y \in \mathbb{X}$, regardless of the true value of X. Its family of desired posteriors are therefore $\pi_x(x') = \delta(x' - y)$ for all $x \in \mathbb{X}$, where δ is the discrete Dirac distribution. Let $M \in \{0,1\}^5$ denote the adversary's message, which it transmits according to the 2-norm solution of (3.14), accounting for the presence of the side-channel information N.

The posteriors induced in the estimator are plotted for two different cases:

(a) the estimator processes N, L, i.e., both sensors are honest

(b) the estimator processes N, M, i.e., one sensor is dishonest

Average posteriors (taking expectations with respect to measurements) are plotted in Figure 3.1, while posteriors corresponding to random realisations of M, L, N are plotted in Figure 3.2. The figures clearly demonstrate the effect of a sensor being 'hacked'. While the adversary does not succeed in inducing the desired Dirac distribution, it still has a significant and misleading effect on the estimators posterior.

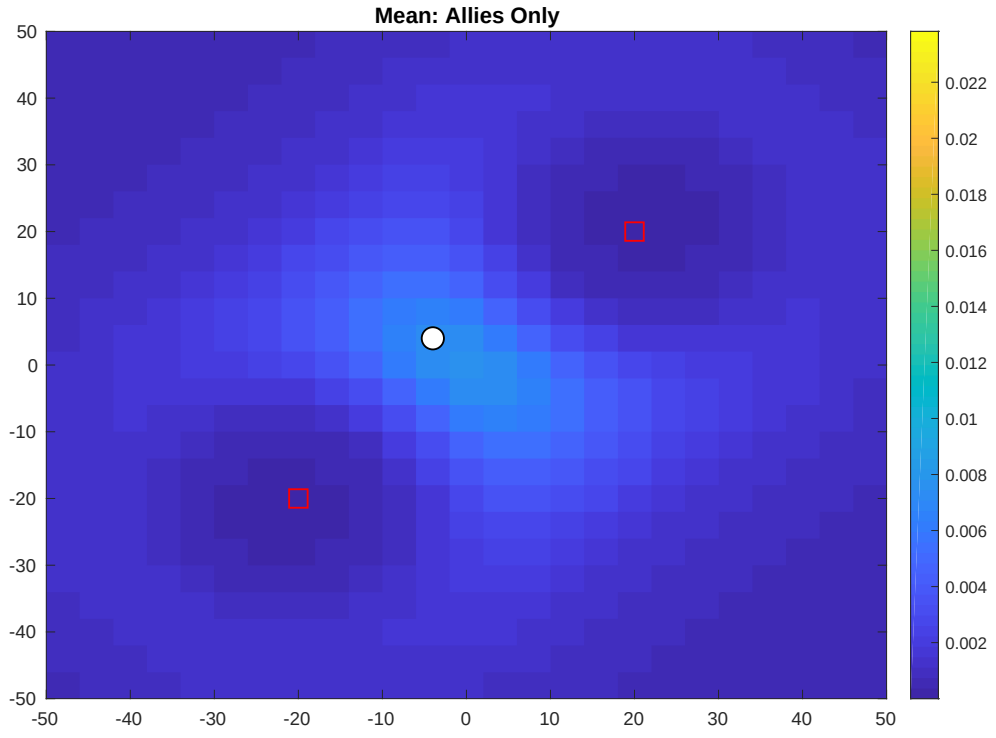

(a) $\mathbb{E}_{N,L}[p_{X|N,L}(x'|N,L) \mid X = x]$

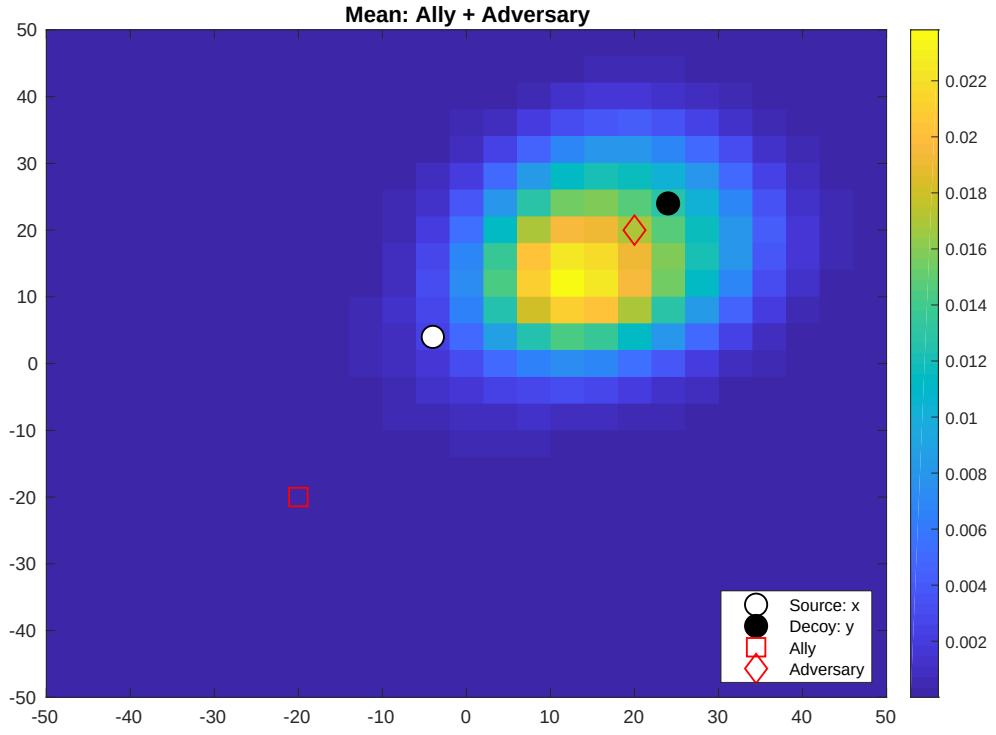

(b) $\mathbb{E}_{N,M}[p_{X|N,M}(x'|N,M) \mid X = x]$

Figure 3.1: Expected posteriors induced in estimator for $x = (-4, 4)$, $y = (24, 24)$, $\boldsymbol{\xi}_1 = (-20, -20)$, $\boldsymbol{\xi}_2 = (20, 20)$. (Units: m)

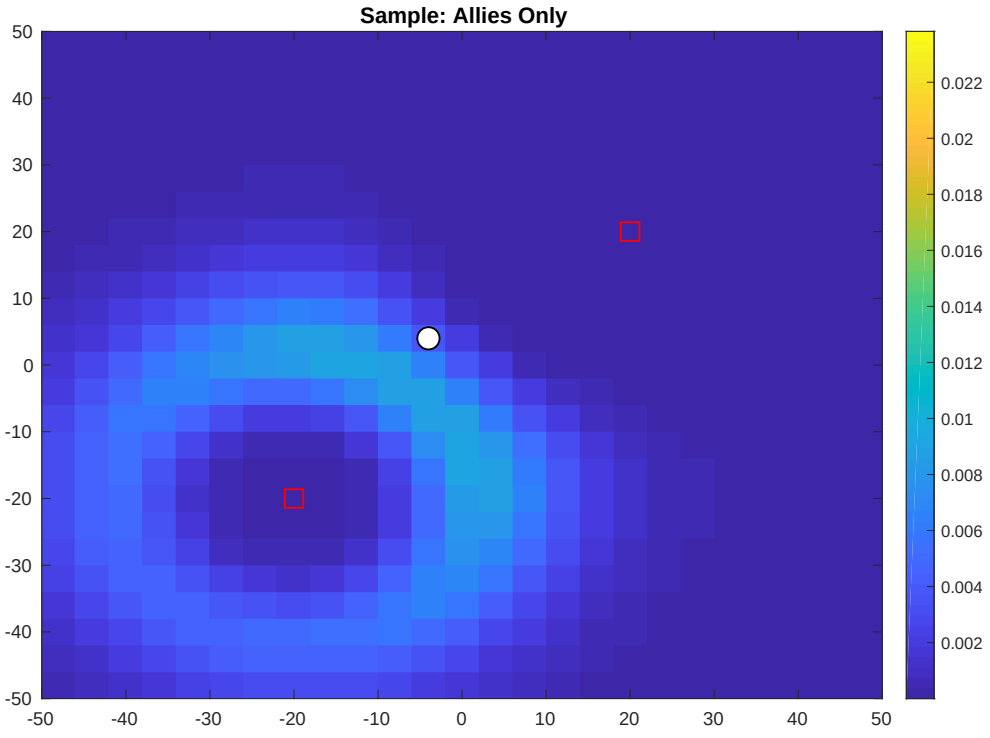

(a) $p_{X|N,L}(x'|N, L) \mid X = x$

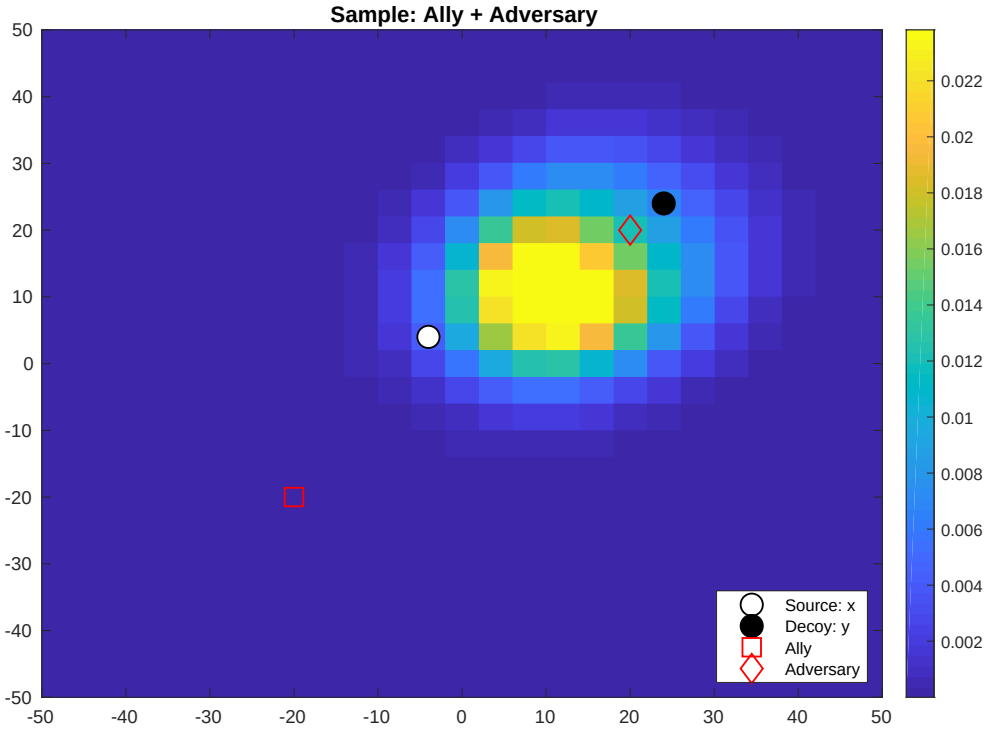

(b) $p_{X|N,M}(x'|N, M) \mid X = x$

Figure 3.2: Realisations of posteriors induced in estimator for $x = (-4, 4)$, $y = (24, 24)$, $\boldsymbol{\xi}_1 = (-20, -20)$, $\boldsymbol{\xi}_2 = (20, 20)$. (Units: m)

Chapter 4

Decentralised Robust Connectivity Preservation

> And the Lord said, "Indeed the people are one and they all have one language, and this is what they begin to do; now nothing that they propose to do will be withheld from them.
>
> Genesis 11:6

This chapter considers the design of control laws that guarantee connectivity preservation within a team of single-integrator agents subject to piecewise continuous bounded disturbances. Section 4.1 presents a review of the relevant literature. The problem is formulated mathematically and the controller is proposed in Section 4.2. Theoretical robustness guarantees are obtained in Section 4.3, and the chapter concludes with simulation results in Section 4.4.

4.1 Related works

Control laws for connectivity preservation are hardly a recent development. A large body of literature is devoted to achieving this under both fixed [64–66], and time-varying [67–69] topologies. The proposed control laws are often designed to achieve specific objectives such as rendezvous [64, 66–68],

flocking [69, 70], or formation control [68], while guaranteeing connectivity preservation throughout. Single integrator dynamics are commonly assumed, however more complex models such as double integrators [69–71], unicycles [72–74], and other structured nonlinear dynamics [67, 75], have not escaped attention. Strategies with bounded control inputs have also been developed in [72, 76, 77]. The above works all adopt gradient-following control laws that seek to minimise a potential function. Other, typically earlier, works rely on optimisation-based approaches to maximise the graph Fielder value [78, 79], or maintain positive-definiteness of the sum of powers of its adjacency matrix [80, 81]. The list of references above is not exhaustive, but we direct the reader to [82] for a more comprehensive review of the field. Only more recent works have considered robustness of the proposed control strategies to disturbances, and we focus on these below.

A collection of double integrator systems subject to external force disturbances is considered in [83]. The disturbances are assumed to be generated by a known autonomous linear system, implying that the disturbance dynamics are perfectly known. The authors design observer-based controllers to achieve leader-following rendezvous (i.e., all agents converging to the same reference trajectory), while preserving connectivity. This generalises their previous work [84], which assumes each agent is subject to the same disturbance. The same authors also design controllers to achieve connectivity preserving velocity consensus in [70], under the same disturbance assumptions, but without relying on a leader. The leader-following rendezvous work [71] extends [83] by considering both force and velocity disturbances, and allowing for parameter uncertainty in the assumed disturbance dynamics. Problem complexity is further increased in [67], which assumes double-integrator dynamics, but with an additional unknown nonlinear force $f_i(\cdot, \cdot)$, and smooth bounded force disturbances. A PID controller, with learning states to compensate for the unknown dynamics and disturbances, is designed to track a bounded reference trajectory. Smoothness features prominently in the assumptions of [67]: f_i is twice continuously differentiable, disturbances have bounded first and second derivatives, and the reference trajectory has bounded derivatives up to fourth order. Under these conditions, the controller guarantees that no edges will be lost. One-dimensional single integrator dynamics are assumed in [85], which proposes a discontinuous control law to achieve finite-time rendezvous in the presence of arbitrary continuous, bounded disturbances. A discontinuous sliding mode control law is also adopted in [64] to extend [67] by achieving leader-following rendezvous in finite time, under weaker smoothness assumptions. The problem of tracking the reference trajectory in formation, while preserving connectivity, is then addressed within the same framework.

In this chapter, we return to the basic problem of maintaining a fixed, connected network topology under single-integrator dynamics. We adopt a gradient-following control law, and show that the desired network topology plays an important role in proving robustness. The control law is decentralised, in that the individual agents compute their own actions, and only rely on information from their neighbours. In contrast to all the works mentioned in the previous paragraph, we do not explicitly assume additional objectives such as consensus or formation control. This allows us to provide connectivity preservation guarantees that are not restricted to a particular objective. Our framework still accommodates certain objectives via the choice of potential functions, desired network topology, and the addition of bounded control signals. Remarks 4.2.1 - 4.2.3 indicate how this may be achieved. Secondly, we do not assume a model for disturbances (c.f. [70, 71, 83, 84]), neither do we require them to be smooth (c.f. [67]). Our controller has a simpler structure that those of [67, 71, 83, 84] because it does not rely on the existence of a leader or reference trajectory, and does not require the construction of observers or learning algorithms. Note that [85] considers single-integrator agents subject to continuous bounded disturbances, as does the work herein. However, those results are only stated for one-dimensional agents, whereas our analysis allows the dimension to be arbitrary. Moreover, we achieve connectivity preservation with a smooth control law, thereby avoiding the technicalities of non-smooth analysis in [64,85]. We adopt the potential function approach of [68], which we generalize to a larger class of potential functions. That work did not establish robustness to disturbances. The primary contribution of Chapter 4 is to extend [68] by proving robust connectivity preservation under arbitrary piecewise continuous bounded disturbances, when the desired network topology is a tree.

4.2 Problem formulation

Consider a team of N mobile agents. Let $x_i \in \mathbb{R}^n, n \in \{2, 3\}$ denote the position of agent i, which evolves according to

$$\dot{x}_i = u_i,$$

where u_i is the control input to agent i. The vector $\mathbf{x} := (x_1, ..., x_N) \in \mathbb{R}^{Nn}$ fully specifies the configuration of the team. Let $\mathcal{G}(\mathbf{x}) = (\mathbb{N}_N, E(\mathbf{x}))$ be a simple graph with a node set $\mathbb{N}_N := \{1, ..., N\}$ indexing the team of agents, and an edge set $E(\mathbf{x})$, which contains all pairs of agents $\{i, j\}$ that can communicate with each other when in configuration $\mathbf{x}$. Assuming that

connectivity is equivalent to the agents being within some communication range $R > 0$ of each other, we define the edge set as

$$E(\mathbf{x}) := \{\{i, j\} \mid \|x_i - x_j\| < R,\ i, j \in \mathbb{N}_N\}. \tag{4.1}$$

Let E_0 be a desired set of edges that makes the graph $(\mathbb{N}_N, E_0)$ connected, and let

$$\mathcal{N}(i) := \{j \in \mathbb{N}_N \mid \{i, j\} \in E_0\}$$

denote the corresponding neighbour set of agent i.

Problem 4.1. Assuming the initial configuration $\mathbf{x}_0$ is chosen such that $E_0 \subset E(\mathbf{x}_0)$, design a control law $u_i(\mathbf{x})$ such that

1. $E_0 \subset E(\mathbf{x}(t))$ for all $t \geq 0$, in the presence of disturbances
2. the controller only uses local information. That is, computing $u_i(\mathbf{x})$ only requires information from agents in $\mathcal{N}(i) \cup \{i\}$.

Problem 4.1 can be solved via the use of an *edge-potential*, which can be thought of as a virtual tensile energy between neighbouring agents.

Definition 4.1 (Edge-potential)**.** An edge-potential is any function $w : [0, R^2) \to [0, \infty)$ that satisfies the following properties:

1. w is continuously differentiable
2. $w'(s)$ is locally Lipschitz
3. $w(s) \to \infty$ and $w'(s) \to \infty$ as $s \to R^2$.

Let the edges in E_0 be labelled $1, ..., m$, where $m = |E_0|$, and $k_{ij} = k_{ji} \in \mathbb{N}_m$ corresponds to edge $\{i, j\} \in E_0$. Consider a family of edge-potentials $\{w_k\}_{k=1}^m$. We define the total potential of the configuration as

$$T = \frac{1}{2} \sum_{i=1}^{N} \sum_{j \in \mathcal{N}(i)} w_{k_{ij}}(\|x_j - x_i\|^2). \tag{4.2}$$

Note that T is only defined for $\mathbf{x} \in \mathcal{R}$, where

$$\mathcal{R} := \{\mathbf{x} \in \mathbb{R}^{Nn} \mid \forall \{i, j\} \in E_0,\ \|x_j - x_i\| < R\}.$$

Note also that $\mathcal{R}$ is open and convex, but unbounded, and by definition,

$$E_0 \subset E(\mathbf{x}) \iff \mathbf{x} \in \mathcal{R}.$$

We adopt the control-law $\mathbf{u} : \mathcal{R} \to \mathbb{R}^{Nn}$,

$$\mathbf{u}(\mathbf{x}) := -\nabla_{\mathbf{x}} T(\mathbf{x}), \tag{4.3}$$

which drives the agents in the direction of steepest descent in potential. Note that

$$u_i(\mathbf{x}) = -\frac{\partial T}{\partial x_i} = \sum_{j \in \mathcal{N}(i)} w'_{k_{ij}}(\|x_j - x_i\|^2)(x_j - x_i),$$

and therefore Point 2 of Problem 4.1 is satisfied.

Remark 4.2.1. Certain additional objectives, such as rendezvous, formation control and even collision avoidance, can be incorporated through the appropriate design of the edge potentials and E_0. Specifically, local minimisers of $w_{k_{ij}}$ should correspond to desired distances between agents i and j.

Remark 4.2.2. The edge-tension functions proposed in [68, (14),(24)] satisfy the properties of an edge-potential. In the absence of disturbances, these induce rendezvous and formation control respectively.

Remark 4.2.3. While it will be shown the controller (4.3) guarantees connectivity preservation in the presence of additional piecewise continuous bounded control signals, further work is required to show that the additional signals still achieve their objective in the presence of (4.3). To this end, edge potentials can be designed such that $w_k(s) = 0$ for all $s \leq \delta^2$, where $0 \leq \delta < R$. Here, δ corresponds to some safe inter-agent distance, within which the connectivity preservation term does not interfere with the additional objective.

We conclude this section with a precise definition for piecewise continuity.

Definition 4.2 (Piecewise Continuous)**.** Let $f : \mathbb{R} \to \mathbb{R}^n$.

1. The function f is piecewise continuous on $[a, b] \subset \mathbb{R}$ if there exists a finite number of points $a = t_0 < t_1 < \ldots < t_k = b$ such that $f(t)$ is continuous on $t \in (t_{i-1}, t_i)$ for all $i \in \mathbb{N}_k$, and

$$\forall i \in \mathbb{N}_k, \lim_{t \to t_{i-1}^+} \|f(t)\| < \infty, \lim_{t \to t_i^-} \|f(t)\| < \infty.$$

2. Furthermore, f is piecewise continuous if it is piecewise continuous on every closed interval $[a, b] \subset \mathbb{R}$.

4.3 Existence and uniqueness

Consider the system

$$\dot{\mathbf{x}} = \mathbf{u}(\mathbf{x}) + \mathbf{d}(t), \tag{4.4}$$

where $\mathbf{d}$ is bounded and piecewise continuous in t. We assume that $\mathbf{d}$ is unknown.

Remark 4.3.1. Point 1 of Problem 4.1 is a stronger requirement than set ISS [23] with respect to $\mathcal{R}$. While set ISS asymptotically bounds the distance to $\mathcal{R}$ in terms of $\|\mathbf{d}\|_\infty$, connectivity preservation requires that $\mathbf{x}(t)$ always remain in the interior of $\mathcal{R}$ for any finite value of $\|\mathbf{d}\|_\infty$.

The controller $\mathbf{u} : \mathcal{R} \to \mathbb{R}^{Nn}$ is only locally Lipschitz, and its magnitude approaches infinity on the boundary of $\mathcal{R}$. A sufficient condition for the existence and uniqueness of solutions to (4.4) for all $t \geq 0$, is the forward invariance of some compact subset of $\mathcal{R}$ [86, Theorem 3.3]. In general, however, the trajectories of the agents $\mathbf{x}(t)$ need not be bounded, and therefore may not remain within any compact set. To circumvent this, we consider the dynamics of the edge vectors $x_j - x_i,\ \{i,j\} \in E_0$, which must be bounded for connectivity to be maintained. Suppose we arbitrarily assign a direction to each edge $1, ..., m$. Let $H \in \mathbb{R}^{N\times m}$ denote the (vertex-edge) incidence matrix of the resulting directed graph, and define $\mathbf{H} := H \otimes I_n \in \mathbb{R}^{Nn\times mn}$, where $\otimes$ is the Kronecker matrix product and I_n is the $n \times n$ identity matrix. The transformation

$$\mathbf{y} = \mathbf{H}^\top \mathbf{x}, \tag{4.5}$$

where $\mathbf{y} := (y_1, ..., y_m) \in \mathbb{R}^{mn}$, then yields $y_k = x_j - x_i$ for some $\{i,j\} \in E_0$. Note that

$$\mathbf{x} \in \mathcal{R} \iff \mathbf{H}^\top \mathbf{x} \in \mathcal{B}(R)^m, \tag{4.6}$$

where $\mathcal{B}(R) = \{y \in \mathbb{R}^n \mid \|y\| < R\}$. The total potential can then be expressed in terms of the edge vectors,

$$T = \frac{1}{2}\sum_{k=1}^{m} w_k(\|y_k\|^2), \tag{4.7}$$

only defined for $\mathbf{y} \in \mathcal{B}(R)^m$. Let $v_k : \mathcal{B}(R) \to \mathbb{R}^n$, where

$$\begin{aligned} v_k(y_k) &:= -w_k'(\|y_k\|^2)y_k \\ &= -\left(\frac{\partial T}{\partial y_k}\right)^\top. \end{aligned} \tag{4.8}$$

Defining $\mathbf{v}(\mathbf{y}) := (v_1(y_1), ..., v_m(y_m))$ then yields

$$\mathbf{v}(\mathbf{y}) = -\nabla_{\mathbf{y}} T(\mathbf{y}), \tag{4.9}$$

where $\mathbf{v} : \mathcal{B}(R)^m \to \mathbb{R}^{mn}$. An important relationship between $\mathbf{u}$ and $\mathbf{v}$ is derived below.

Lemma 4.3.1. *For any* $\mathbf{x} \in \mathcal{R}$, $\mathbf{u}(\mathbf{x}) = \mathbf{H}\mathbf{v}(\mathbf{H}^\top \mathbf{x})$.

Proof. Applying the chain rule to (4.3), the transformation (4.5) implies that

$$\begin{aligned} \forall \mathbf{x}' \in \mathcal{R},\ \mathbf{u}(\mathbf{x}') &= -\nabla_{\mathbf{x}} T(\mathbf{x}') \\ &= -\left[\left. \frac{\partial T}{\partial \mathbf{y}} \right|_{\mathbf{y}=\mathbf{H}^\top \mathbf{x}'} \left. \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}'} \right]^\top \\ &= -\mathbf{H} \nabla_{\mathbf{y}} T(\mathbf{H}^\top \mathbf{x}') \\ &= \mathbf{H}\mathbf{v}(\mathbf{H}^\top \mathbf{x}'). \end{aligned} \tag{4.10}$$

□

We now examine bounds on the magnitude of $\mathbf{u}(\mathbf{x})$. For $\epsilon \in [0, R]$, define

$$\overline{\mathcal{R}}_\epsilon := \{\mathbf{x} \in \mathbb{R}^{Nn} \mid \forall \{i, j\} \in E_0,\ \|x_j - x_i\| \leq R - \epsilon\}.$$

Just as in (4.6),

$$\mathbf{x} \in \overline{\mathcal{R}}_\epsilon \iff \mathbf{H}^\top \mathbf{x} \in \overline{\mathcal{B}}(R - \epsilon)^m, \tag{4.11}$$

where $\overline{\mathcal{B}}(\cdot)$ denotes the closure of $\mathcal{B}(\cdot)$ in $\mathbb{R}^n$.

Lemma 4.3.2. *For any* $r \in [0, R)$, $\mathbf{v} : \mathcal{B}(R) \to \mathbb{R}^{mn}$ *is bounded on* $\overline{\mathcal{B}}(r)$.

Proof. Since $\overline{\mathcal{B}}(r)$ is the finite Cartesian product of closed balls in $\mathbb{R}^n$, it is compact. If $r \in [0, R)$, then $\overline{\mathcal{B}}(r) \subset \mathcal{B}(R)$, and Point 1 of Definition 4.1 implies $\mathbf{v}$ is continuous on $\mathcal{B}(R)$. Thus, $\|\mathbf{v}(\overline{\mathcal{B}}(r))\|$ is compact by [46, Theorem 4.14], and therefore bounded. □

Corollary 4.3.3. *For any* $\epsilon \in (0, R]$, $\mathbf{u} : \mathcal{R} \to \mathbb{R}^{Nn}$ *is bounded on* $\overline{\mathcal{R}}_\epsilon$.

Proof. If $\mathbf{x} \in \overline{\mathcal{R}}_\epsilon$, then (4.11) implies $\mathbf{H}^\top \mathbf{x} \in \overline{\mathcal{B}}(R - \epsilon)^m$. The rest follows from Lemma 4.3.1 and Lemma 4.3.2. □

A self-contained equation for the edge vector dynamics is now derived.

Lemma 4.3.4 (Edge vector dynamics)**.** *Let $J \subset \mathbb{R}$ be an open interval. Suppose $\mathbf{x} : J \to \mathcal{R}$ is continuously differentiable, and $\mathbf{d} : J \to \mathbb{R}^{Nn}$ is continuous and bounded. Define $\mathbf{y}(t) := \mathbf{H}^\top \mathbf{x}(t)$. If*

$$\forall t \in J,\ \dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t)) + \mathbf{d}(t), \tag{4.12}$$

then

$$\forall t \in J,\ \dot{\mathbf{y}}(t) = \mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{y}(t)) + \mathbf{H}^\top \mathbf{d}(t). \tag{4.13}$$

Proof. The edge-vector dynamics can be written as

$$\begin{aligned}\dot{\mathbf{y}}(t) &= \mathbf{H}^\top \dot{\mathbf{x}}(t) \\ &= \mathbf{H}^\top \mathbf{u}(\mathbf{x}(t)) + \mathbf{H}^\top \mathbf{d}(t) \\ &= \mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{H}^\top \mathbf{x}(t)) + \mathbf{H}^\top \mathbf{d}(t),\end{aligned}$$

by Lemma 4.3.1. Substituting $\mathbf{y}(t) = \mathbf{H}^\top \mathbf{x}(t)$ yields (4.13). □

The existence and uniqueness of solutions to (4.13) under certain initial conditions and network topologies can be proved by appealing to [86, Theorem 3.3]. We restrict attention to trees, because the incidence matrix of any directed tree has full column rank [87, Lemma 2.5]. The proof of our main result relies on the following property.

Lemma 4.3.5. *Let $M \in \mathbb{R}^{N \times m}$ have full column rank, and let $v : [0, \tau) \to \mathbb{R}^m$, where $\tau \in (0, \infty]$. Then,*

$$\lim_{t \to \tau} \|v(t)\| = \infty \implies \lim_{t \to \tau} \|Mv(t)\| = \infty.$$

Proof. Since M has full column rank, $M^\top M \in \mathbb{R}^{m \times m}$ is symmetric positive definite and

$$\lambda_m \|v(t)\|^2 \leq v(t)^\top M^\top M v(t) = \|Mv(t)\|^2,$$

where $\lambda_m > 0$ is the smallest eigenvalue of $M^\top M$. □

Lemma 4.3.6. *If $M \in \mathbb{R}^{N \times m}$ has full column rank, then $M \otimes I_n$ has full column rank.*

Proof. This follows from the identity $\mathrm{rk}(M \otimes I_n) = \mathrm{rk}(M) \cdot \mathrm{rk}(I_n)$. □

We are now ready to establish the existence and uniqueness of solutions to the edge-vector dynamics. Solutions are defined in the sense of Caratheodory, and need only be piecewise continuously differentiable.

Proposition 4.3.7. *Let $\mathbf{d} : \mathbb{R} \to \mathbb{R}^{Nn}$ be a piecewise continuous bounded disturbance. Assume E_0 specifies a tree. Let $(t_0, \mathbf{x}_0) \in \mathbb{R} \times \mathcal{R}$. Then, there exists a unique solution $\mathbf{y} : [t_0, \infty) \to \mathcal{B}(R)^m$ to the initial value problem*

$$\dot{\mathbf{y}}(t) = \mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{y}(t)) + \mathbf{H}^\top \mathbf{d}(t),\ \mathbf{y}(t_0) = \mathbf{H}^\top \mathbf{x}_0. \tag{4.14}$$

Furthermore, $\exists r < R,\ \forall t \geq t_0,\ \mathbf{y}(t) \in \overline{\mathcal{B}}(r)^m$.

Proof. The RHS of (4.14) is piecewise continuous in t, and locally Lipschitz in $\mathbf{y}$ on the domain $\mathbb{R} \times \mathcal{B}(R)$. The assumption $\mathbf{x}_0 \in \mathcal{R}$ implies $\mathbf{y}(t_0) \in \mathcal{B}(R)^m$. As per [86, Section 3.1], there exists a right maximal interval of existence[1] $[t_0, \tau),\ \tau \in (t_0, \infty]$, over which (4.14) has a solution $\mathbf{y}(t)$. By definition, the solution is continuous and satisfies

$$\forall t \in [t_0, \tau),\ \mathbf{y}(t) \in \operatorname{dom} \mathbf{v} = \mathcal{B}(R)^m. \tag{4.15}$$

The derivative of the potential along the solution trajectory satisfies

$$\begin{aligned}
\dot{T} &= \frac{\partial T}{\partial \mathbf{y}} \dot{\mathbf{y}}(t) \\
&= -\mathbf{v}(\mathbf{y}(t))^\top \left[\mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{y}(t)) + \mathbf{H}^\top \mathbf{d}(t)\right] \\
&= -\|\mathbf{H}\mathbf{v}(\mathbf{y}(t))\|^2 - [\mathbf{H}\mathbf{v}(\mathbf{y}(t))]^\top \mathbf{d}(t) \\
&\leq -\|\mathbf{H}\mathbf{v}(\mathbf{y}(t))\| (\|\mathbf{H}\mathbf{v}(\mathbf{y}(t))\| - \|\mathbf{d}(t)\|),
\end{aligned} \tag{4.16}$$

for almost all $t \geq t_0$ (that is, except where $\mathbf{d}$ is discontinuous). Now, to obtain a contradiction, suppose

$$r := \sup\{\|y_k(t)\| |\ t \in [t_0, \tau), k \in \mathbb{N}_m\} \geq R. \tag{4.17}$$

Comparing this with (4.15), continuity implies

$$\exists j \in \mathbb{N}_m,\ \lim_{t \to \tau^-} y_j(t) = R.$$

Definition 4.1 and (4.7) then imply

$$\lim_{t \to \tau^-} T = \infty. \tag{4.18}$$

Since E_0 specifies a tree, H has full column rank [87, Lemma 2.5]. Lemma 4.3.6 then implies $\mathbf{H}$ has full column rank. Definition 4.1 and (4.8) imply $\|v_j(y_j(t))\| \to \infty$ as $t \uparrow \tau$. Therefore, $\|\mathbf{v}(\mathbf{y}(t))\| \to \infty$, and Lemma 4.3.5 then implies $\|\mathbf{H}\mathbf{v}(\mathbf{y}(t))\| \to \infty$, as $t \uparrow \tau$. Since $\mathbf{d}$ is bounded, (4.16) implies there

[1] see [88, Definition 8.31]

exists $t_1 \in [t_0, \tau)$ such that $\dot{T} < 0$ for almost all $t \in (t_1, \tau)$. This contradicts (4.18), proving the hypothesis (4.17) is false. Thus,

$$\forall j \in \mathbb{N}_m,\ \forall t \in [t_0, \tau),\ \|y_j(t)\| \leq r < R. \tag{4.19}$$

Therefore $\tau = \infty$ by [86, Theorem 3.3], which also establishes that $\mathbf{y}$ is the unique solution to (4.14). Finally, (4.19) is then equivalent to $\mathbf{y}(t) \in \overline{\mathcal{B}}(r)^m$ for all $t \geq t_0$. □

Although a unique solution to (4.14) is guaranteed, $\mathbf{H}^\top$ is rank deficient and cannot be inverted to directly obtain a solution to (4.4). The result below guarantees that such a solution can still be found via integration.

Lemma 4.3.8. *Let $(t_0, \mathbf{x}_0) \in \mathbb{R} \times \mathcal{R}$, and let $\mathbf{d} : \mathbb{R} \to \mathbb{R}^{Nn}$ be piecewise continuous and bounded. Consider the two initial value problems*

$$\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t)) + \mathbf{d}(t), \qquad \mathbf{x}(t_0) = \mathbf{x}_0, \tag{4.20}$$

$$\dot{\mathbf{y}}(t) = \mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{y}(t)) + \mathbf{H}^\top \mathbf{d}(t), \qquad \mathbf{y}(t_0) = \mathbf{H}^\top \mathbf{x}_0. \tag{4.14}$$

If $\mathbf{y} : [t_0, \infty) \to \mathcal{B}(R)^m$ is a solution to (4.14), then

$$\mathbf{x}(t) := \mathbf{x}_0 + \int_{t_0}^{t} \mathbf{H}\mathbf{v}(\mathbf{y}(\tau)) + \mathbf{d}(\tau) d\tau \tag{4.21}$$

is a solution to (4.20) on $t \geq t_0$.

Proof. Suppose $\mathbf{y} : [t_0, \infty) \to \mathcal{B}(R)^m$ is a solution to (4.14), implying that

$$\forall t \geq t_0,\ \mathbf{y}(t) \in \operatorname{dom} \mathbf{v} = \mathcal{B}(R)^m. \tag{4.22}$$

Defining $\mathbf{x}(t)$ as in (4.21), we obtain

$$\begin{aligned} \mathbf{H}^\top \mathbf{x}(t) &= \mathbf{H}^\top \mathbf{x}_0 + \int_{t_0}^{t} \mathbf{H}^\top \mathbf{H}\mathbf{v}(\mathbf{y}(\tau)) + \mathbf{H}^\top \mathbf{d}(\tau)\ d\tau \\ &= \mathbf{y}_0 + \int_{t_0}^{t} \dot{\mathbf{y}}(\tau) d\tau \\ &= \mathbf{y}(t), \end{aligned} \tag{4.23}$$

for all $t \geq t_0$. Let $\mathcal{D} \subset \mathbb{R}$ denote the set of points at which $\mathbf{d}$ is discontinuous, and let $\mathcal{T} := [t_0, \infty) \setminus \mathcal{D}$. Now (4.21) implies that $\mathbf{x}$ is differentiable on $\mathcal{T}$, and furthermore,

$$\begin{aligned} \forall t \in \mathcal{T},\ \dot{\mathbf{x}}(t) &= \frac{d}{dt} \int_{t_0}^{t} \mathbf{H}\mathbf{v}(\mathbf{y}(\tau)) + \mathbf{d}(\tau) d\tau \\ &= \mathbf{H}\mathbf{v}(\mathbf{y}(t)) + \mathbf{d}(t) \\ &= \mathbf{H}\mathbf{v}(\mathbf{H}^\top \mathbf{x}(t)) + \mathbf{d}(t), \end{aligned} \tag{4.24}$$

by (4.23). Now (4.6) and (4.22) together imply

$$\forall t \geq t_0,\ \mathbf{x}(t) \in \mathcal{R} = \operatorname{dom} \mathbf{u}.$$

Applying Lemma 4.3.1 to (4.24) then yields

$$\forall t \in \mathcal{T},\ \dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t)) + \mathbf{d}(t).$$

Noting that $\mathbf{x}(t_0) = \mathbf{x}_0$ completes the proof. □

The preceding results now come together to solve Problem 4.1.

Theorem 4.3.9. *Let $\mathbf{d} : \mathbb{R} \to \mathbb{R}^{Nn}$ be a piecewise continuous, bounded disturbance. Assume E_0 specifies a tree. Let $(t_0, \mathbf{x}_0) \in \mathbb{R} \times \mathcal{R}$. Then, there exists a unique solution $\mathbf{x} : [t_0, \infty) \to \mathcal{R}$ to the initial value problem*

$$\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t)) + \mathbf{d}(t),\ \mathbf{x}(t_0) = \mathbf{x}_0. \tag{4.20}$$

Furthermore, $\exists \epsilon > 0,\ \forall t \geq t_0,\ \mathbf{x}(t) \in \overline{\mathcal{R}}_\epsilon$.

Proof. Let $\mathbf{x}_0 \in \mathcal{R}$. Proposition 4.3.7 then establishes the existence of a unique $\mathbf{y} : [t_0, \infty) \to \mathcal{B}(R)^m$ that satisfies (4.14). Lemma 4.3.8 then implies that $\mathbf{x} : [t_0, \infty) \to \mathcal{R}$, as defined in (4.21), is a solution to (4.20). All that remains is to establish the uniqueness of this solution. Suppose that $\mathbf{x}_2 : [t_0, \infty) \to \mathcal{R}$ is also a solution to (4.20). Then, by definition, it satisfies the integral equation

$$\forall t \geq t_0,\ \mathbf{x}_2(t) = \int_{t_0}^{t} \mathbf{u}(\mathbf{x}_2(\tau)) + \mathbf{d}(\tau) + \mathbf{x}_0.$$

Let $\mathbf{y}_2(t) := \mathbf{H}^\top \mathbf{x}_2(t)$. Lemma 4.3.4 implies that $\mathbf{y}_2(t)$ is also a solution to (4.14). Since this solution is unique, $\mathbf{y}_2(t) = \mathbf{y}(t)$ for all $t \geq t_0$. Now applying Lemma 4.3.1,

$$\begin{aligned}
\forall t \geq t_0,\ \mathbf{x}_2(t) &= \int_{t_0}^{t} \mathbf{H}\mathbf{v}(\mathbf{H}^\top \mathbf{x}_2(\tau)) + \mathbf{d}(\tau) + \mathbf{x}_0 && (4.25)\\
&= \int_{t_0}^{t} \mathbf{H}\mathbf{v}(\mathbf{y}_2(\tau)) + \mathbf{d}(\tau) + \mathbf{x}_0 && (4.26)\\
&= \int_{t_0}^{t} \mathbf{H}\mathbf{v}(\mathbf{y}(\tau)) + \mathbf{d}(\tau) + \mathbf{x}_0 && (4.27)\\
&= \mathbf{x}(t), && (4.28)
\end{aligned}$$

by (4.21). Furthermore, since there exists $r < R$ such that $\mathbf{y}_2(t) = \mathbf{y}(t) \in \overline{\mathcal{B}}(r)^m$ for all $t \geq t_0$, (4.11) implies $\mathbf{x}_2(t) = \mathbf{x}(t) \in \overline{\mathcal{R}}_\epsilon$ for all $t \geq t_0$, where $\epsilon = R - r$. □

Not only does this guarantee that connectivity is preserved, but that it can also be achieved with a bounded control signal.

Corollary 4.3.10 (Bounded control effort)**.** *Under the dynamics* (4.20)*,* $\mathbf{u}(\mathbf{x}(t))$ *is bounded for all* $t \geq t_0$*.*

Proof. This follows from Corollary 4.3.3. □

4.4 Numerical simulation

We now present a numerical simulation of four single-integrator agents, choosing E_0 to be a path graph. The edge-potentials are designed as follows:

$$\forall k \in \mathbb{N}_3,\ w_k(s) = w(s) := \begin{cases} 0, & s \in [0, \delta^2] \\ \dfrac{(s-\delta^2)^2}{(s-R^2)^2}, & s \in (\delta^2, R^2) \end{cases}.$$

Agent i is subject to the disturbance

$$d_i(t, x_i) := U \,\mathrm{sat}[-x_i + p(t) + r_i(t)],$$

where $\mathrm{sat} : \mathbb{R}^n \to \mathbb{R}^n$,

$$\mathrm{sat}(x) := \frac{x}{\max\{\|x\|, 1\}},$$

both $p : \mathbb{R} \to \mathbb{R}^n$, $r_i : \mathbb{R} \to \mathbb{R}^n$ are piecewise constant, and $U = 1000$. This is an additional control term that directs all agents into a time-varying formation about the point $p(t) \in \mathbb{R}^n$. Specifically, agent i is displaced $r_i(t) \in \mathbb{R}^n$ from $p(t)$. Note that control law proposed in Section 2.6 takes this form. Choosing $\delta = 20$ m and $R = 30$ m, the edge lengths are plotted with and without imposing the control law (4.3) in Figure 4.1.

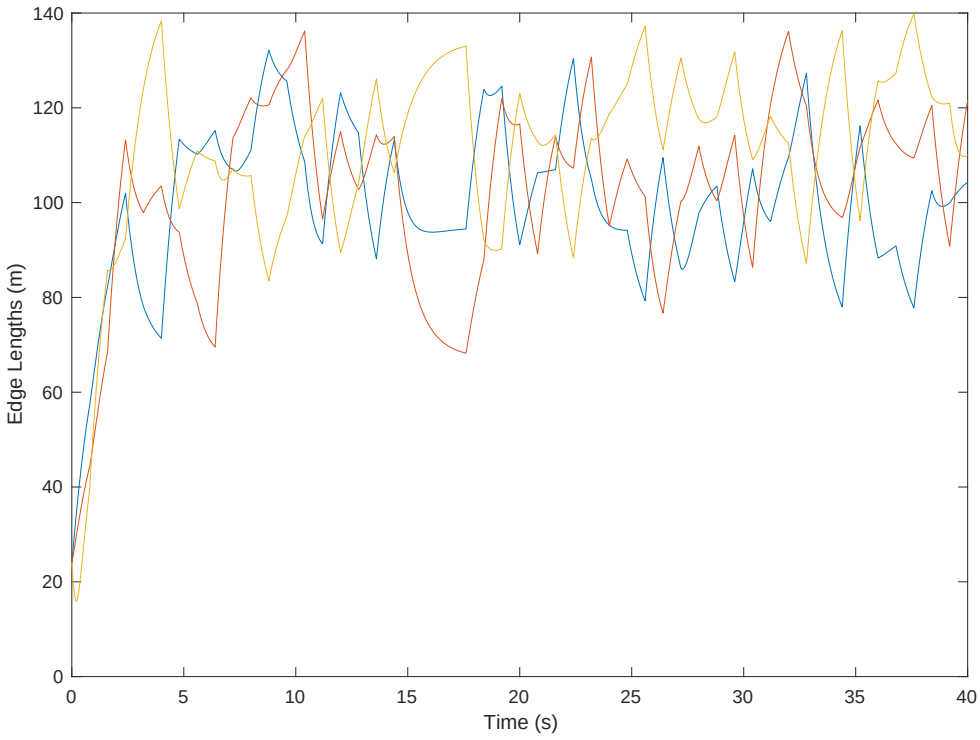

(a) Without controller

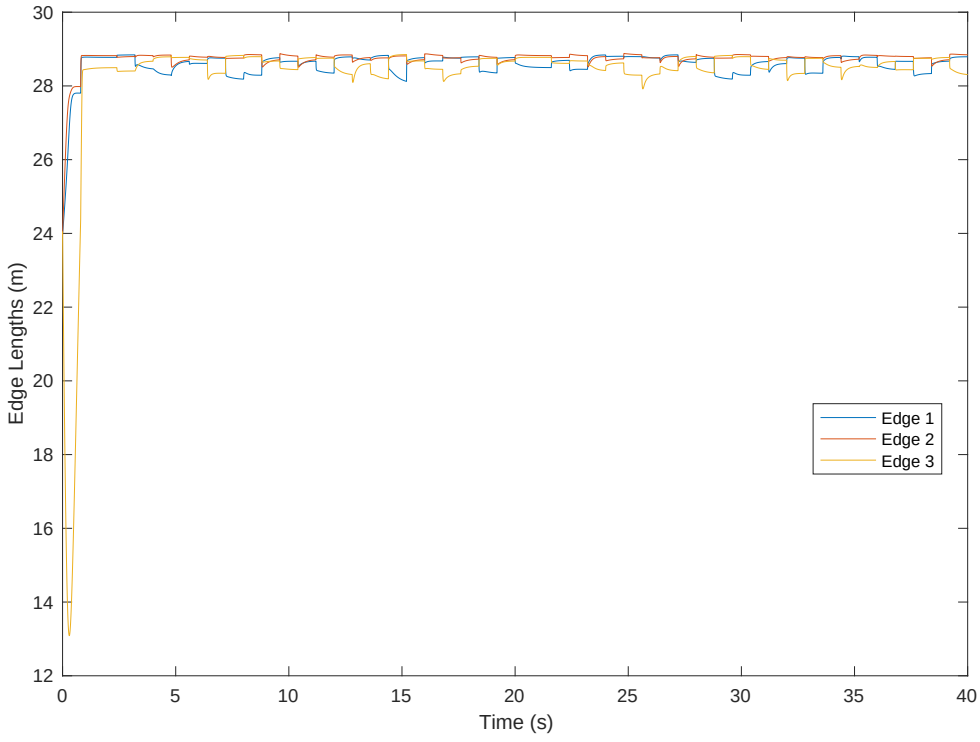

(b) With controller ($R = 30$)

Figure 4.1: Edge lengths

Part II

Time-Varying Optimisation Problems

Chapter 5

Online Gradient Descent

"We seek him here, we seek him there,
Those Frenchies seek him everywhere.
Is he in heaven? — Is he in hell,
That demmed, elusive Pimpernel?"

The Scarlet Pimpernel
Baroness Emmuska Orczy

This chapter considers optimisation problems with cost functions and constraints that change discretely with time[1]. We focus on obtaining performance bounds for online projected gradient descent iterations, when applied to different classes of cost functions. We begin in Section 5.1 with a brief review of the relevant literature on time-varying optimisation. The problem and its basic assumptions are then stated precisely in Section 5.2. Section 5.3 commences by considering general unconstrained nonlinear cost functions. The PL inequality is then invoked to obtain sub-optimality bounds in Section 5.3.1. Strongly convex cost functions are assumed in Section 5.4, which also incorporates time-varying convex constraints. Error and sub-optimality bounds are both derived, and Section 5.4.3 in particular tackles issues relating to constraints. Finally, two numerical examples are provided in Section 5.5.

[1]Material from Sections 1.4, 5.1, 5.2, 5.4, 5.5, 7.4 and Appendix D.1 has been accepted for publication in the proceedings of the *57th IEEE Conference on Decision and Control (CDC 2018)*, under the title 'Numerical Optimisation of Time-Varying Strongly Convex Functions Subject to Time-Varying Constraints'.

5.1 Related works

Gradient descent on a sequence of smooth, strongly convex functions is investigated in [89], with bounds on tracking error derived for unconstrained problems. Error bounds are also presented for cost functions that vary in continuous time, with the gradient descent iterates replaced by a gradient-based control law. A general framework for time-varying convex optimisation problems (with time-varying constraints) is proposed in [29], based on the theory of averaged operators. The authors develop error bounds on a variant of the Mann-Krasnosel'skii iterations with time-varying operators. These bounds are then used to prove the convergence of a variety of 'running' methods, including projected gradient descent, proximal-point, forward-backward splitting, dual ascent, dual decomposition, and ADMM.

Our own work on strongly convex costs extends the projected gradient descent results in [89] and [29] by exploring the consequences of time-varying constraints in greater detail. In particular, the tracking error bounds for gradient descent in [29] are stated under slightly different assumptions to [89], which requires the difference between consecutive cost function gradients to be bounded over the whole of $\mathbb{R}^n$. The latter is quite restrictive, given that strong convexity and an unconstrained domain are both assumed. In [29], a bounded distance between consecutive optimisers is assumed instead, with the results remaining valid under time-varying constraints. While this is more general, in practice it can be difficult to establish a useful bound on the distance between consecutive optimisers when time-varying constraints are involved (e.g. [90]). In Section 5.4, we derive such bounds under appropriate restrictions on the change between consecutive constraints. In addition, further conditions are stated which guarantee finite-time feasibility. We also present sub-optimality bounds, along with a more detailed analysis of the error than is available those works.

Time-varying convex optimisation problems are studied by the machine learning community in the context of Online Convex Optimisation (OCO) [91,92]. OCO frames the problem as a game, in which the iterates are actions selected by a player. In response to each action, an adversary selects a cost-function, which determines the cost of that action. In contrast to the aforementioned works, OCO is concerned with minimising the integrated cost rather than the asymptotic cost. Algorithmic performance is evaluated in terms of the *regret*, defined as the integrated difference between the observed cost and the cost at the best fixed decision variable. *Dynamic regret* is a generalisation which compares the observed cost with the instantaneous optimal

cost. OCO results are typically presented in terms of regret bounds, the goal being to achieve sub-linear regret. Dynamic regret is more relevant to the wider literature and to our own work, because it directly integrates the sub-optimality. Dynamic regret bounds for projected gradient descent are derived in [93] under the assumptions of strongly convex cost functions, and fixed, compact constraints. Compact constraints are a standard assumption in the OCO literature that is not required in more general formulations. Special cases of time-varying constraints are considered in [94], along with possibly non-smooth cost functions, however only static regret bounds are presented. In contrast, the work of Chapter 5 considers general time-varying (convex) constraints that need not be compact.

We now mention some literature on time-varying optimisation that considers a different class of iterations to projected gradient descent. In particular, [95–97] focus on Newton-type methods, with [97] discussing optimisation problems on manifolds with changing co-ordinate maps. Generalized equations are used in [98] to develop an augmented-Lagrangian method. Several other works have studied time-varying cost functions in the context of distributed optimisation over a network. The distributed structure of the solutions precludes the use of straightforward gradient descent. Of these, [99,100] both treat the entire problem in continuous time, focusing on quadratic costs and deriving control laws which track the optimum. Others present iterative solutions using a variety of methods such as consensus-based algorithms in [101, 102], and distributed gradient descent in [103]. Duality is exploited by [24, 25, 104] and [105], which develops online ADMM.

5.2 Problem formulation

We will work with a family of continuous functions

$$\mathcal{F} := \{f_k : \mathbb{R}^n \to \mathbb{R} \mid k \in \mathbb{N}_0\}.$$

Let $\mathcal{X} := \{X_k \subset \mathbb{R}^n \mid k \in \mathbb{N}_0\}$ be a family of closed, non-empty sets, which captures possible time-varying constraints. We are interested in generating solutions to the following sequence of optimisation problems

$$\min_{x \in X_k} f_k(x), \tag{5.1}$$

using iterates of the form

$$x_{k+1} = G_k(x_k), \tag{5.2}$$

where the operator $G_k : \mathbb{R}^n \to \mathbb{R}^n$ is restricted to using first order information and knowledge of the current feasible set X_k. This formulation reduces to a standard time-invariant optimisation problem when $\mathcal{F} = \{f\}$ and $\mathcal{X} = \{X\}$.

Our analysis will rely on the following assumptions.

Assumption 5.1 (Well-posedness)**.** For every $f \in \mathcal{F}$, the set of minima $\arg\min\{f(x) \mid x \in \mathbb{R}^n\}$ is non-empty.

Accordingly, define

$$f_k^\star := \min\{f(x) \mid x \in \mathbb{R}^n\},$$

and let

$$\phi_k := f_k(x_k) - f_k^\star \tag{5.3}$$

denote the level of sub-optimality.

Assumption 5.2 (Smoothness)**.** Every $f \in \mathcal{F}$ is twice continuously differentiable, and

$$\exists L > 0,\ \forall f \in \mathcal{F},\ \forall x, y \in \mathbb{R}^n,\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|.$$

Further assumptions will be made as required.

5.3 Nonlinear unconstrained problems

We begin by considering smooth, non-convex cost functions, and restrict attention to unconstrained problems, in which $\mathcal{X} = \{\mathbb{R}^n\}$. In particular, we examine iterates of the form (5.2), where

$$G_k(x) := x - \alpha_k p_k,$$

and $p_k \in \mathbb{R}^n$ is a descent direction; that is, $\nabla f_k(x)^\top p_k < 0$.

Zoutendijk's Theorem [30, Theorem 3.2] is a well-known classical result that relies on the *Wolfe Conditions* to establish the convergence of line-search methods to stationary points of a non-linear cost-function. The Wolfe Conditions [30, Section 3.1] place restrictions on the step size α_k to guarantee both sufficient decrease in cost and sufficient progress with each step. In order to extend Zoutendijk's theorem to time-varying problems, we restrict the change between consecutive cost functions as follows.

Assumption 5.3 (Bounded increase in cost)**.** For the family $\mathcal{F}$, there exists $\delta_0 \geq 0$ such that

$$\forall k,\ \forall x \in \mathbb{R}^n,\ f_{k+1}(x) - f_k(x) \leq \delta_0.$$

Theorem 5.3.1 (Modified Zoutendijk)**.** *Let $\mathcal{F}$ satisfy Assumptions 5.1 - 5.3. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = x_k + \alpha_k p_k,$$

where $\alpha_k > 0$ and $p_k \in \mathbb{R}^n$ are chosen such that

$$\forall k,\ p_k^\top \nabla f(x_k) < 0 \tag{5.4}$$

$$\exists c_1 \in (0,1),\ \forall k,\ f_k(x_{k+1}) \leq f_k(x_k) + c_1 \alpha_k \nabla f_k(x_k)^\top p_k \tag{5.5}$$

$$\exists c_2 \in (c_1,1),\ \forall k,\ \nabla f_k(x_{k+1})^\top p_k \geq c_2 \nabla f_k(x_k)^\top p_k. \tag{5.6}$$

Then,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{k=1}^{N} \|\nabla f_k(x_k)\|^2 \cos^2\theta_k \leq \frac{L}{C}\left(\delta_0 + \Delta\right), \tag{5.7}$$

where $\theta_k := \arccos\left(\frac{\nabla f_k(x_k)^\top p_k}{\|\nabla f_k(x_k)\|\|p_k\|}\right)$, $C := c_1(1-c_2)$ and

$$\Delta := \liminf_{k\to\infty} \frac{f_0^\star - f_k^\star}{k}. \tag{5.8}$$

Proof. Assumption 5.2 and (5.4) - (5.6) guarantee the hypotheses of [30, Theorem 3.2] are satisfied for any fixed f_k. Applying a single iteration yields

$$f_k(x_{k+1}) \leq f_k(x_k) - \tfrac{C}{L}\cos^2\theta_k \|\nabla f_k(x_k)\|^2,$$

which is derived in the proof of [30, Theorem 3.2]. Assumption 5.3 then implies

$$f_{k+1}(x_{k+1}) \leq f_k(x_k) - \tfrac{C}{L}\cos^2\theta_k \|\nabla f_k(x_k)\|^2 + \delta_0.$$

Applying this result recursively, we obtain

$$f_N(x_{N+1}) \leq f_0(x_0) - \sum_{k=0}^{N} \tfrac{C}{L}\cos^2\theta_k \|\nabla f_k(x_k)\|^2 + N\delta_0.$$

This is equivalent to

$$\begin{aligned} 0 \leq \phi_{N+1} &= f_N(x_{N+1}) - f_{N+1}^\star \\ &\leq f_0(x_0) - f_0^\star - \sum_{k=0}^{N} \tfrac{C}{L}\cos^2\theta_k \|\nabla f_k(x_k)\|^2 + N\delta_0 + f_0^\star - f_{N+1}^\star, \end{aligned}$$

which then yields

$$\frac{1}{N}\sum_{k=0}^{N}\cos^2\theta_k\|\nabla f_k(x_k)\|^2\leq \frac{L}{C}\left(\frac{\phi_0}{N}+\delta_0+\frac{f_0^\star-f_{N+1}^\star}{N}\right).$$

Taking the limit inferior of both sides as $N\to\infty$ yields the result. □

Remark 5.3.1. Conditions (5.5) - (5.6) are the Wolfe conditions, required to hold with the same constants c_1, c_2 for every $f\in\mathcal{F}$. For an arbitrary family $\mathcal{F}$, it is not *a-priori* clear that there exist choices of α_k, c_1, c_2 that satisfy (5.5) - (5.6). It is shown in Appendix D.2 that this existence is guaranteed for strongly convex cost functions when $p_k=-\nabla f_k(x_k)$.
Remark 5.3.2. The parameter Δ in (5.8) can be thought of as the average reduction in the optimal cost per iteration. It is possible for $\Delta<0$, which occurs if some subsequence of the $f_k^\star$ increases linearly. Note that if $\mathcal{F}$ is uniformly bounded below, then $\Delta=0$.

When applying Theorem 5.3.1 to gradient descent, (5.7) bounds the mean square gradient magnitude. This leads us to the following asymptotic result.

Proposition 5.3.2. *Let the hypotheses of Lemma 5.3.1 be satisfied. Choosing $p_k=-\nabla f_k(x_k)$ implies*

$$\liminf_{k\to\infty}\|\nabla f_k(x_k)\|^2\leq\frac{(\delta_0+\Delta)L}{C}. \tag{5.9}$$

Proof. Choosing $p_k=-\nabla f_k(x_k)$ implies $\cos^2\theta_k=1$. Equation (5.7) then yields

$$\lim_{N\to\infty}\frac{1}{N}\sum_{k=1}^{N}\|\nabla f_k(x_k)\|^2\leq\frac{(\delta_0+\Delta)L}{C}=:u. \tag{5.10}$$

To obtain a contradiction, suppose $l:=\liminf_{k\to\infty}\|\nabla f_k(x_k)\|^2>u$ and let $\epsilon:=\frac{l-u}{2}>0$. Thus $u+\epsilon<l$, and therefore

$$\exists K\in\mathbb{N},\ \forall k>K,\ \|\nabla f_k(x_k)\|^2>u+\epsilon.$$

Thus for any $N>K$,

$$\frac{1}{N}\sum_{k=1}^{N}\|\nabla f_k(x_k)\|^2>\frac{1}{N}\sum_{k=1}^{K}\|\nabla f_k(x_k)\|^2+\frac{(N-K)}{N}(u+\epsilon).$$

Taking $N\to\infty$ on both sides,

$$\lim_{N\to\infty}\frac{1}{N}\sum_{k=1}^{N}\|\nabla f_k(x_k)\|^2>u+\epsilon,$$

which contradicts (5.10). □

Remark 5.3.3. Proposition 5.3.2 guarantees the gradient sequence will return infinitely often to any ball of radius larger than $\sqrt{\frac{(\delta_0+\Delta)L}{C}}$, centred at the origin.

5.3.1 PL cost functions

We now restrict attention to families of functions that satisfy the PL inequality. We refer to such functions as PL functions.

Assumption 5.4 (Polyak-Łojasiewicz)**.** There exists $\mu > 0$ such that for all $f \in \mathcal{F}$,

$$\forall x \in \mathbb{R}^n,\ \frac{1}{2}\|\nabla f(x)\|^2 \geq \mu(f(x) - f^\star), \tag{5.11}$$

where $f^* = \min\{f(x) \mid x \in \mathbb{R}^n\}$.

Remark 5.3.4. If a function satisfies (5.11), then all its stationary points are global minima. That is, the PL inequality implies invexity.

Note that in Lemma D.3, it is shown that for sufficiently smooth functions, the gradient Lipschitz constant L is an upper bound for μ. Clearly, if the PL inequality is satisfied, Proposition 5.3.2 leads directly to limit inferior bounds on sub-optimality for online gradient descent.

Corollary 5.3.3. *Let the hypothesis of Proposition 5.3.2 be satisfied, along with Assumption 5.4. Then*

$$\liminf_{k\to\infty} \phi_k \leq \frac{(\delta_0 + \Delta)L}{2\mu C}.$$

Further analysis allows us to obtain limit superior sub-optimality bounds as well.

Lemma 5.3.4. *Let $\mathcal{F}$ satisfy Assumptions 5.1 - 5.4. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = x_k - \alpha_k \nabla f_k(x_k), \tag{5.12}$$

where $0 < \alpha_k \leq \frac{2}{L}$ for all $k \in \mathbb{N}_0$. Then,

$$\forall k \in \mathbb{N}_0,\ \phi_{k+1} \leq (1 - \eta_k)\phi_k + \delta_0 + f_k^\star - f_{k+1}^\star,$$

where $\eta_k := \alpha_k\mu\left(1 - \frac{\alpha_k L}{2}\right) \in (0, \frac{1}{2}]$.

Proof. By Taylor's Theorem,

$$\begin{aligned} f_k(x_{k+1}) - f_k(x_k) &\leq \nabla f_k(x_k)^\top (x_{k+1} - x_k) + \frac{L}{2}\|x_{k+1} - x_k\|^2 \\ &= \nabla f_k(x_k)^\top (-\alpha_k \nabla f_k(x_k)) + \frac{\alpha_k^2 L}{2}\|\nabla f_k(x_k)\|^2 \\ &= -\alpha_k \left(1 - \frac{\alpha_k L}{2}\right) \|\nabla f_k(x_k)\|^2 \\ &\leq -\alpha_k \mu \left(1 - \frac{\alpha_k L}{2}\right) (f_k(x_k) - f_k^\star) \\ &= \eta_k (f_k^\star - f_k(x_k)). \end{aligned}$$

Rearranging this, we obtain

$$f_k(x_{k+1}) - f_k^\star \leq (1 - \eta_k)(f_k(x_k) - f_k^\star).$$

Since $\mu \leq L$, note that

$$0 < \eta_k \leq \alpha_k L(1 - \tfrac{\alpha_k L}{2}) \leq \tfrac{1}{2}.$$

Now observe

$$\begin{aligned} f_{k+1}(x_{k+1}) - f_{k+1}^\star &= [f_{k+1}(x_{k+1}) - f_k(x_{k+1})] + [f_k(x_{k+1}) - f_k^\star] + [f_k^\star - f_{k+1}^\star] \\ &\leq \delta_0 + (1 - \eta_k)[f_k(x_k) - f_k^\star] + f_k^\star - f_{k+1}^\star. \end{aligned}$$

□

Corollary 5.3.5. *The choice of step-size that maximises η_k is $\alpha_k = \frac{1}{L}$, which yields $\eta_k = \frac{\mu}{2L}$.*

We now impose an additional assumption, which limits the decrease between consecutive optimal values.

Assumption 5.5. For the family $\mathcal{F}$, there exists $\delta^\star \geq 0$ such that

$$\forall k \in \mathbb{N}_0,\ f_k^\star - f_{k+1}^\star \leq \delta^\star.$$

Note this assumption implies $\Delta \leq \delta^\star$, where Δ is defined in (5.8). We use this to obtain a stronger result than Corollary 5.3.3.

Proposition 5.3.6. *Let $\mathcal{F}$ satisfy Assumptions 5.1 - 5.5. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = x_k - \frac{1}{L}\nabla f_k(x_k). \tag{5.13}$$

Then,

$$\limsup_{k\to\infty} \phi_k \leq \frac{2(\delta_0 + \delta^\star)L}{\mu}.$$

Proof. Lemma 5.3.4 and Corollary 5.3.5 imply

$$\forall k \in \mathbb{N}_0,\ \phi_{k+1} \leq (1 - \tfrac{\mu}{2L})\phi_k + \delta_0 + \delta^\star,$$

which yields the result. □

Remark 5.3.5. The definition of C in Lemma 5.3.1 implies that $C < \frac{1}{4}$. Thus,

$$\frac{(\delta_0 + \delta^\star)L}{2\mu C} > \frac{2(\delta_0 + \delta^\star)L}{\mu}.$$

The limit inferior bound in Corollary 5.3.3 is therefore made redundant by the limit superior bound in Proposition 5.3.6 under Assumption 5.5.

5.4 Strongly convex problems

We now turn attention to families of strongly convex functions. The additional structure this affords allows us to obtain stronger bounds, and to address time-varying convex constraints. Accordingly, we make the following assumptions.

Assumption 5.6 (Uniform strong convexity)**.** There exists $\sigma \in (0, L]$ such that

$$\forall f \in \mathcal{F},\ \forall x \in \mathbb{R}^n,\ \nabla^2 f(x) \succeq \sigma I.$$

Assumption 5.7 (Convex constraints)**.** Every $X \in \mathcal{X}$ is closed, convex and non-empty.

Together, these assumptions guarantee the existence of a unique global minimiser, which we define as

$$x_k^\star := \arg\min\{f_k(x) \mid x \in X_k\}.$$

Furthermore, any continuously differentiable, σ-strongly convex function satisfies the PL inequality with $\mu = \sigma$ [106, Lemma 1.3]. Thus, Assumptions 5.6 and 5.7 make Assumptions 5.1 and 5.4 redundant. For the remainder of this chapter, we employ online projected gradient descent to deal with time-varying constraints. This pertains to a specific choice of G_k that will be defined in the next subsection. In addition to making stronger assumptions and incorporating constraints, we also place different restrictions on the variation between cost functions. Specifically, instead of bounding the increase in cost, we bound the distance between consecutive global minimisers to obtain bounds on tracking error. We also show in Section 5.4.3 that the shift in minimisers can in turn be bounded by the change between consecutive cost function gradients.

5.4.1 Error bounds

Let $P_X : \mathbb{R}^n \to X$,

$$P_X(y) := \arg\min_{x \in X} \|x - y\|^2,$$

denote the projection operator, where $X \subset \mathbb{R}^n$ is a closed, convex set. The Projection Theorem guarantees the existence and uniqueness of $P_X(y)$. In online projected gradient descent, iterates evolve according to (5.2), where

$$G_k(x) := P_{X_k}\left(x - \alpha_k \nabla f_k(x)\right), \tag{5.14}$$

and $\alpha_k > 0$ is a sequence of step sizes.

We will consider the evolution of the *tracking error* $e_k := \|x_k - x_k^\star\|$, and the distance

$$\bar{e}_k := \|x_{k+1} - x_k^\star\|,$$

which we refer to as the *estimation error*, because x_{k+1} can be considered an estimate of $x_k^\star$ using all the information available up to time k. In [107, Proposition 6.1.8], it is established that, under Assumptions 5.2 – 5.7, G_k is a contraction map for $\alpha_k < \frac{2}{L}$. Specifically,

$$\forall x, y \in \mathbb{R}^n,\ \|G_k(x) - G_k(y)\| \leq \rho_k \|x - y\|, \tag{5.15}$$

where

$$\rho_k := \max\{|1 - \alpha_k L|, |1 - \alpha_k \sigma|\}. \tag{5.16}$$

Given the sequence α_k, we will also find it convenient to define

$$\rho := \sup\{\rho_k \mid k \in \mathbb{N}_0\}.$$

It is well known that $x_k^\star$ is a fixed point of G_k, and we present a simple proof of this below.

Lemma 5.4.1 (Fixed point)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and continuously differentiable. Let $X \subset \mathbb{R}^n$ be closed, convex and non-empty. Let $\alpha > 0$, and define*

$$G(x) := P_X(x - \alpha \nabla f(x)).$$

Then, $G(x^\star) = x^\star$ for any $x^\star \in \arg\min\{f(x) \mid x \in X\}$.

Proof. By definition of $x^\star$,

$$\forall x \in X,\ \nabla f(x^\star)^\top (x - x^\star) \geq 0. \tag{5.17}$$

Now the Projection Theorem [108, Theorem 1 of Section 3.12] implies

$$\forall x \in X,\ [x^\star - \alpha\nabla f(x^\star) - G(x^\star)]^\top [x - G(x^\star)] \leq 0.$$

Substituting $x = x^\star \in X$ into the above yields

$$\|x^\star - G(x^\star)\|^2 + \alpha\nabla f(x^\star)^\top [G(x^\star) - x^\star] \leq 0.$$

Referring to (5.17), both terms on the LHS are non-negative, and therefore both must vanish. □

Substituting $x = x_k$ and $y = x_k^\star$ into (5.15) then yields the inequality

$$\bar{e}_k \leq \rho_k e_k, \tag{5.18}$$

which we use to bound the tracking error evolution. For notational convenience, define $\prod_{j=N}^{N-1} \rho_j := 1$ for any $N \in \mathbb{N}$.

Theorem 5.4.2 (Tracking error dynamics)**.** *Let $\mathcal{F}$, $\mathcal{X}$ satisfy Assumptions 5.2, 5.6 and 5.7. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = P_{X_k}\left(x_k - \alpha_k \nabla f_k(x_k)\right), \tag{5.19}$$

where $0 < \alpha_k < \frac{2}{L}$ for all k. Then

$$\forall k \in \mathbb{N},\ e_{k+1} \leq \rho_k e_k + \|x_{k+1}^\star - x_k^\star\|. \tag{5.20}$$

Proof. Observe that

$$\begin{aligned} \|x_{k+1} - x_{k+1}^\star\| &= \|x_{k+1} - x_k^\star + x_k^\star - x_{k+1}^\star\| \\ &\leq \|x_{k+1} - x_k^\star\| + \|x_k^\star - x_{k+1}^\star\| \\ &\leq \rho_k \|x_k - x_k^\star\| + \|x_k^\star - x_{k+1}^\star\|, \end{aligned} \tag{5.21}$$

by application of (5.18). □

Corollary 5.4.3. *For all $N \in \mathbb{N}$,*

$$e_N \leq e_0 \prod_{j=0}^{N-1} \rho_j + \sum_{k=0}^{N-1} \|x_{k+1}^\star - x_k^\star\| \prod_{j=k+1}^{N-1} \rho_j. \tag{5.22}$$

Proof. Treating (5.21) as a dynamical system, we obtain

$$\forall N \in \mathbb{N},\ e_N \leq \Phi(N, 0) e_0 + \sum_{k=0}^{N-1} \Phi(N, k+1) \|x_{k+1}^\star - x_k^\star\|,$$

with state-transition operator $\Phi(N, k) := \prod_{j=k}^{N-1} \rho_j$. □

Remark 5.4.1. Theorem 5.4.2 reveals the tracking error is bounded by a discrete time dynamical system, where the shift in minimisers behaves as a disturbance input. Corollary 5.4.3 presents a closed-form expression for this bound.

If the shift in minimizers (i.e., the disturbance input) is bounded, this leads to a steady-state error bound under mild restrictions on the step sizes. We make the following assumption accordingly.

Assumption 5.8 (Bounded shift in minimiser)**.** For the pair $(\mathcal{F}, \mathcal{X})$, there exists $V \geq 0$ such that

$$\forall k \in \mathbb{N}_0,\ \|x_{k+1}^\star - x_k^\star\| \leq V.$$

Applying this leads directly to an alternative proof of [29, Corollary 7.1 (b)].

Corollary 5.4.4. *Let the hypotheses of Theorem 5.4.2 be satisfied, along with Assumption 5.8. If there also exists $\epsilon > 0$ and $\bar{\alpha} < \frac{2}{L}$ such that $\epsilon \leq \alpha_k \leq \bar{\alpha}$ for all k, then*

$$\limsup_{k\to\infty} e_k \leq \frac{V}{1-\rho} \tag{5.23}$$

$$\limsup_{k\to\infty} \bar{e}_k \leq \frac{V\rho}{1-\rho}. \tag{5.24}$$

Proof. Referring to (5.16), ρ_k only attains its maximum value of 1 over $\alpha_k \in [0, \frac{2}{L}]$ at $\alpha_k = 0$ and $\alpha_k = \frac{2}{L}$. Thus, for any $\epsilon \in (0, \frac{2}{L})$ and $\bar{\alpha} \in (\epsilon, \frac{2}{L})$, restricting $\alpha_k \in [\epsilon, \bar{\alpha}]$ for all k implies

$$\rho \leq \sup_{\alpha\in[\epsilon,\bar{\alpha}]} \Big(\max\{|1-\alpha L|, |1-\alpha\sigma|\} \Big) < 1.$$

Since ρ and V are upper bounds for ρ_k and $\|x_{k+1}^\star - x_k^\star\|$ respectively, we obtain

$$e_N \leq \rho^N e_0 + V \sum_{k=0}^{N-1} \rho^k$$

from (5.22). The result follows by taking limits, and applying (5.18). □

Remark 5.4.2. When Assumption 5.8 is satisfied, it is clear from (5.20), (5.23) and (5.24) that minimising ρ_k results in both the fastest rate of convergence of the bound, and its steady-state value. Referring to (5.16), this corresponds to a constant step size of $\alpha_k = \frac{2}{\sigma+L}$, which implies $\rho = \rho_k = \frac{L-\sigma}{L+\sigma}$ for all k.

Corollary 5.4.5. *If $\alpha = \frac{2}{\sigma+L}$ for all k, then*

$$\limsup_{k\to\infty} e_k \leq \frac{V(L+\sigma)}{2\sigma} \tag{5.25}$$

$$\limsup_{k\to\infty} \bar{e}_k \leq \frac{V(L-\sigma)}{2\sigma}. \tag{5.26}$$

Henceforth, the step size that minimizes some steady state bound on tracking error or sub-optimality will be referred to as the *optimal step size*.

Speed constraints

In certain situations, it may be of practical interest to bound the distance between consecutive iterates. For example, the x_k may be used to generate a sequence of way-points for a mobile robot, in which case bounds on $\|x_{k+1} - x_k\|$ correspond to speed constraints on the robot. Assume $\mathcal{X} = \{\mathbb{R}^n\}$, and suppose we require that

$$\|x_{k+1} - x_k\| = \alpha_k \|\nabla f_k(x_k)\| \leq U \tag{5.27}$$

for all k. It is then necessary that $\alpha_k \leq \frac{U}{\|\nabla f_k(x_k)\|}$, which may interfere with the optimal choice of step-size. Now if $\nabla f_k(x_k)$ is unbounded, then $\alpha_k \to 0$, which violates the hypotheses of Corollary 5.4.4. However, if the gradient sequence is bounded, then steady-state error bounds can be guaranteed while still ensuring (5.27).

5.4.2 Unconstrained sub-optimality bounds

We now turn our attention to the level of sub-optimality. Recall the definition of ϕ_k in (5.3), and similarly define

$$\bar{\phi}_k := f_k(x_{k+1}) - f_k(x_k^\star). \tag{5.28}$$

We refer to $\bar{\phi}_k$ as the *predicted sub-optimality*, which is relevant if the cost of action x_k is incurred before the function changes. Here, we consider bounds on ϕ_k and $\bar{\phi}_k$ for unconstrained problems. The effect of constraints will be considered in Section 5.4.4. The guarantee that $\nabla f_k(x_k^\star) = 0$ when $X_k = \mathbb{R}^n$ leads to the inequalities

$$\phi_k \leq \tfrac{L}{2} e_k^2 \tag{5.29}$$

$$\bar{\phi}_k \leq \tfrac{L}{2} \bar{e}_k^2, \tag{5.30}$$

which are proved in Lemma D.1 under Assumption 5.2. These allow our tracking and estimation error bounds to be translated directly into sub-optimality bounds. In particular, they yield the steady-state sub-optimality bounds below.

Corollary 5.4.6 (Steady-state sub-optimality)**.** *Suppose $\mathcal{X} = \{\mathbb{R}^n\}$ and $\mathcal{F}$ satisfies Assumptions 5.2, 5.6 and 5.8. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = x_k - \alpha \nabla f_k(x_k), \tag{5.31}$$

where $\alpha = \frac{2}{\sigma+L}$. Then

$$\limsup_{k\to\infty} \phi_k \leq \frac{LV^2(L+\sigma)^2}{8\sigma^2} \tag{5.32}$$

$$\limsup_{k\to\infty} \bar{\phi}_k \leq \frac{LV^2(L-\sigma)^2}{8\sigma^2}. \tag{5.33}$$

5.4.3 Issues specific to time-varying constraints

The error bounds of Section 5.4.1 apply equally well to constrained and unconstrained problems. However, the presence of constraints does warrant extra consideration. As yet, projected gradient descent does not guarantee $x_k \in X_k$, but only that $x_k \in X_{k-1}$. Establishing bounds on the shift in optima to satisfy Assumption 5.8 is also more challenging when constraints are involved. Furthermore, the sub-optimality bounds of Section 5.4.2 no longer apply because the constrained minima need not be stationary points. In this section, we present sufficient conditions to address the issues of feasibility, to bound the shift between minimisers, and to bound sub-optimality in the presence of time-varying constraints.

Finite-time feasibility

While projected gradient descent only guarantees feasibility with respect to the previous constraints, we can guarantee eventual feasibility under an additional assumption.

Assumption 5.9 (Sufficient overlap)**.** The pair $(\mathcal{F}, \mathcal{X})$ is such that

$$\forall k \in \mathbb{N}_0,\ \{x \in X_k \mid \|x - x_k^\star\| \leq R\} \subset X_{k+1},$$

for some $R > \frac{V\rho}{1-\rho}$.

Proposition 5.4.7 (Finite-time feasibility)**.** *Let Assumptions 5.2, 5.6 – 5.9 be satisfied. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = x_k - \alpha \nabla f_k(x_k),$$

where $\alpha \in (0, \frac{2}{L})$. Then,

$$\exists N \in \mathbb{N}_0,\ \forall k > N,\ x_k \in X_k.$$

Proof. For any $R > \frac{V\rho}{1-\rho}$, (5.24) implies there exists $N \in \mathbb{N}_0$ such that $\|x_{k+1} - x_k^\star\| \leq R$ for all $k \geq N$. Use of the projection operator P_{X_k} in (5.19) ensures $x_{k+1} \in X_k$, and Assumption 5.9 then guarantees $x_{k+1} \in X_{k+1}$ for all $k \geq N$. □

Bounds on the shift between minimisers

Here, we present alternative conditions that guarantee Assumption 5.8. These conditions impose restrictions on changes in the cost function gradients and the constraints between consecutive time steps.

Lemma 5.4.8. *Let $X_1, X_2 \subset \mathbb{R}^n$ be closed and convex. Suppose $f_1, f_2 : \mathbb{R}^n \to \mathbb{R}$ are twice continuously differentiable, and strongly convex with modulus $\sigma > 0$. Furthermore, assume $x_1^\star, x_2^\star \in X_1 \cap X_2$, where $x_i^\star = \arg\min\{f_i(x) \mid x \in X_i\}$ for $i = 1, 2$. If*

$$\exists \delta_1 \geq 0,\ \forall x \in X_1 \cap X_2,\ \|\nabla f_1(x) - \nabla f_2(x)\| \leq \delta_1,$$

then $\|x_2^\star - x_1^\star\| \leq \frac{\delta_1}{\sigma}$.

Proof. The result is trivial if $x_1^\star = x_2^\star$, so assume $v := x_2^\star - x_1^\star \neq 0$. Since f_1, f_2 are convex,

$$\forall i \in \{1, 2\},\ \forall x \in X_i,\ \nabla f_i(x_i^\star)^\top (x - x_i^\star) \geq 0. \tag{5.34}$$

Choosing $i = 2$ and $x = x_1^\star \in X_2$, we obtain

$$\nabla f_2(x_2^\star)^\top v \leq 0. \tag{5.35}$$

By the mean value theorem,

$$\nabla f_2(x_2^\star) = \nabla f_2(x_1^\star) + \nabla^2 f_2(x_1^\star + Tv)v,$$

for some $T = \text{diag}(t_1, ..., t_n)$, where $t_1, ..., t_n \in (0, 1)$. Substituting this result into (5.35),

$$\nabla f_2(x_1^\star)^\top v + v^\top \nabla^2 f_2(x_1^\star + Tv)v \leq 0,$$

which implies

$$v^\top \nabla^2 f_2(x_1^\star + Tv)v + \nabla f_1(x_1^\star)^\top v \leq \left[\nabla f_1(x_1^\star) - \nabla f_2(x_1^\star)\right]^\top v. \tag{5.36}$$

Now, choosing $i = 1$ and $x = x_2^\star \in X_1$, (5.34) also implies $\nabla f_1(x_1^\star)^\top v \geq 0$. Thus, by strong convexity,

$$\begin{aligned}\sigma\|v\|^2 &\leq v^\top \nabla^2 f_2(x_1^\star + Tv)v \\ &\leq v^\top \nabla^2 f_2(x_1^\star + Tv)v + \nabla f_1(x_1^\star)^\top v.\end{aligned}$$

Combining this with (5.36) then yields

$$\sigma\|v\|^2 \leq \left[\nabla f_1(x_1^\star) - \nabla f_2(x_1^\star)\right]^\top v \leq \delta_1\|v\|,$$

which implies the result. □

This result is a direct generalisation of [89, (A.16)] to the constrained case. We observe that two conditions are sufficient for bounding the shift in minimisers.

Assumption 5.10. There exists $\delta_1 \geq 0$ for the pair $(\mathcal{F}, \mathcal{X})$, such that

$$\forall k,\ \forall x \in X_k \cap X_{k+1},\ \|\nabla f_{k+1}(x) - \nabla f_k(x)\| \leq \delta_1.$$

Assumption 5.11. Let $(\mathcal{F}, \mathcal{X})$ be such that $x_{k-1}^\star, x_{k+1}^\star \in X_k$.

Remark 5.4.3. In practice, it is difficult to guarantee $x_{k+1}^\star \in X_k$, however, this follows automatically if $X_{k+1} \subset X_k$.

Proposition 5.4.9. *Let $(\mathcal{F}, \mathcal{X})$ satisfy Assumptions 5.2, 5.6, 5.7, 5.10 and 5.11. Then Assumption 5.8 is also satisfied with $V = \frac{\delta_1}{\sigma}$.*

Proof. Assumption 5.11 is equivalent to $x_k^\star, x_{k+1}^\star \in X_k \cap X_{k+1}$. The rest follows directly from Lemma 5.4.8. □

Applying Corollary 5.4.5 then yields the result below.

Corollary 5.4.10. *Let $(\mathcal{F}, \mathcal{X})$ satisfy Assumptions 5.2, 5.6, 5.7, 5.10 and 5.11. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = P_{X_k}\left[x_k - \frac{2}{\sigma + L}\nabla f_k(x_k)\right]. \tag{5.37}$$

Then,

$$\limsup_{k\to\infty} e_k \leq \frac{\delta_1(L+\sigma)}{2\sigma^2} \tag{5.38}$$

$$\limsup_{k\to\infty} \bar{e}_k \leq \frac{\delta_1(L-\sigma)}{2\sigma^2}. \tag{5.39}$$

Remark 5.4.4. Appendix C considers the special case of unconstrained strongly convex quadratics subject to Assumption 5.10. In Proposition C.1, a difference equation that governs the evolution of $x_k - x_k^\star$ exactly is derived for this case. The resulting bound in Corollary C.2 is the same bound implied by Theorem 5.4.2 and Proposition 5.4.9, which were derived under more general assumptions. This suggests that the error bounds of Sections 5.4.1 and 5.4.3 cannot be improved.

5.4.4 Sub-optimality bounds under compact constraints

One way of obtaining sub-optimality bounds in the presence of constraints is to impose a uniform Lipschitz property on the f_k over X_k. Given strong convexity, this implies the compactness of every feasible set.

Assumption 5.12 (Uniform Lipschitz cost)**.** For the pair $(\mathcal{F}, \mathcal{X})$, there exists $M > 0$ such that

$$\forall k \in \mathbb{N}_0,\ \forall x \in X_k,\ \|\nabla f_k(x)\| \leq M.$$

Lemma D.2 then gives us the inequalities

$$\phi_k \leq M e_k \tag{5.40}$$

$$\bar{\phi}_k \leq M \bar{e}_k. \tag{5.41}$$

Corollary 5.4.11 (Steady-state sub-optimality)**.** *Let $\mathcal{F}$, $\mathcal{X}$ satisfy Assumptions 5.2, 5.6 – 5.8 and 5.12. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = P_{X_k}\left(x_k - \alpha \nabla f_k(x_k)\right), \tag{5.42}$$

where $\alpha = \frac{2}{\sigma+L}$. Then

$$\limsup_{k\to\infty} \phi_k \leq \frac{MV(L+\sigma)}{2\sigma}. \tag{5.43}$$

$$\limsup_{k\to\infty} \bar{\phi}_k \leq \frac{MV(L-\sigma)}{2\sigma}. \tag{5.44}$$

A different bound on predicted sub-optimality can be obtained, if we choose step sizes $\alpha_k \leq \frac{1}{L}$.

Proposition 5.4.12. *Let $\mathcal{F}, \mathcal{X}$ satisfy Assumptions 5.2, 5.6, 5.7 and 5.12. Let $x_0 \in \mathbb{R}^n$, and consider iterations of the form*

$$x_{k+1} = P_{X_k}\left(x_k - \alpha_k \nabla f_k(x_k)\right),$$

where $\alpha_k \in (0, \frac{1}{L}]$ for all $k \in \mathbb{N}_0$. Then for all k,

$$\bar{\phi}_k \leq M\sqrt{M^2\alpha_k^2 + e_k^2} - M^2\alpha_k. \tag{5.45}$$

Proof. First note that, since $\alpha_k \in (0, \frac{1}{L}]$, [107, Proposition 6.1.6] implies [107, (6.14)] is satisfied. Thus, all the hypotheses of [107, Proposition 6.1.7] are satisfied. Applying [107, (6.17)],

$$\begin{aligned}\bar{e}_k^2 &= \|x_k^\star - x_{k+1}\|^2 \\ &\leq \|x_k^\star - x_k\|^2 - 2\alpha_k(f_k(x_{k+1}) - f_k(x_k^\star)) \\ &= e_k^2 - 2\alpha_k\bar{\phi}_k.\end{aligned}$$

The inequality (5.41) then implies

$$\bar{\phi}_k^2 \leq M^2(e_k^2 - 2\alpha_k\bar{\phi}_k),$$

which can be rearranged to form

$$\bar{\phi}_k^2 + 2M^2\alpha_k\bar{\phi}_k - M^2e_k^2 \leq 0.$$

Taking the non-negative solutions to the above inequality yields the result. □

Corollary 5.4.13. *If Assumption 5.8 is also satisfied, and $\alpha_k = \alpha \in (0, \frac{1}{L}]$ for all k, then*

$$\limsup_{k\to\infty} \bar{\phi}_k \leq M\sqrt{M^2\alpha^2 + \frac{V^2}{(1-\rho)^2}} - M^2\alpha, \tag{5.46}$$

where $\rho = \max\{|1 - \alpha L|, |1 - \alpha\sigma|\}$.

Proof. This follows from Corollary 5.4.4. □

Remark 5.4.5. The function $\sqrt{x^2 + b} - x$ is decreasing in x and increasing in b for $x, b > 0$. Furthermore, ρ is decreasing in α for $\alpha \in (0, \frac{1}{L}]$. This implies the bound in (5.46) is decreasing in α, and the optimal choice of step size is therefore $\alpha = \frac{1}{L}$.

Corollary 5.4.14. *In particular, if* $\alpha = \frac{1}{L}$*, then*

$$\limsup_{k\to\infty} \bar{\phi}_k \leq M\sqrt{\frac{M^2}{L^2} + \frac{V^2L^2}{\sigma^2}} - \frac{M^2}{L}. \tag{5.47}$$

Remark 5.4.6. The bounds in (5.44) and (5.47) are different bounds, which hold under different choices of step size. The minimum of the two depends on the parameters M, σ, δ_1, V and L.

5.5 Numerical experiments

The theoretical error and sub-optimality bounds obtained for strongly-convex functions are illustrated here by means of two numerical examples. In particular, a non-trivial example with time-varying constraints is provided that satisfies all the required assumptions.

5.5.1 Unconstrained example

In this example, we use a sequence of cost functions of the form

$$f_k(x) = \tfrac{1}{2}(x - s_k)^\top Q_k(x - s_k),$$

where $Q_k \in \mathbb{R}^{n\times n}$ and $s_k \in \mathbb{R}^n$ are randomly generated and satisfy

$$\sigma I \preceq Q_k = Q_k^\top \preceq LI \tag{5.48}$$
$$\|s_{k+1} - s_k\| \leq V, \tag{5.49}$$

for all k. Results for parameter values $n = 3,\ \sigma = 5,\ L = 10,\ V = 4$ are plotted in Figure 5.1, which compares actual performance with the bounds dictated by (5.20), (5.18), (5.29) and (5.30), for a step size $\alpha = \frac{2}{\sigma+L}$.

5.5.2 Constrained example

Here, we use a sequence of cost functions of the form

$$f_k(x) = \tfrac{1}{2}x^\top Qx + q_k^\top x$$

where $\sigma I \preceq Q = Q^\top \preceq LI$, and q_k is randomly generated to satisfy

$$\forall k,\ \|q_{k+1} - q_0\| \leq \delta. \tag{5.50}$$

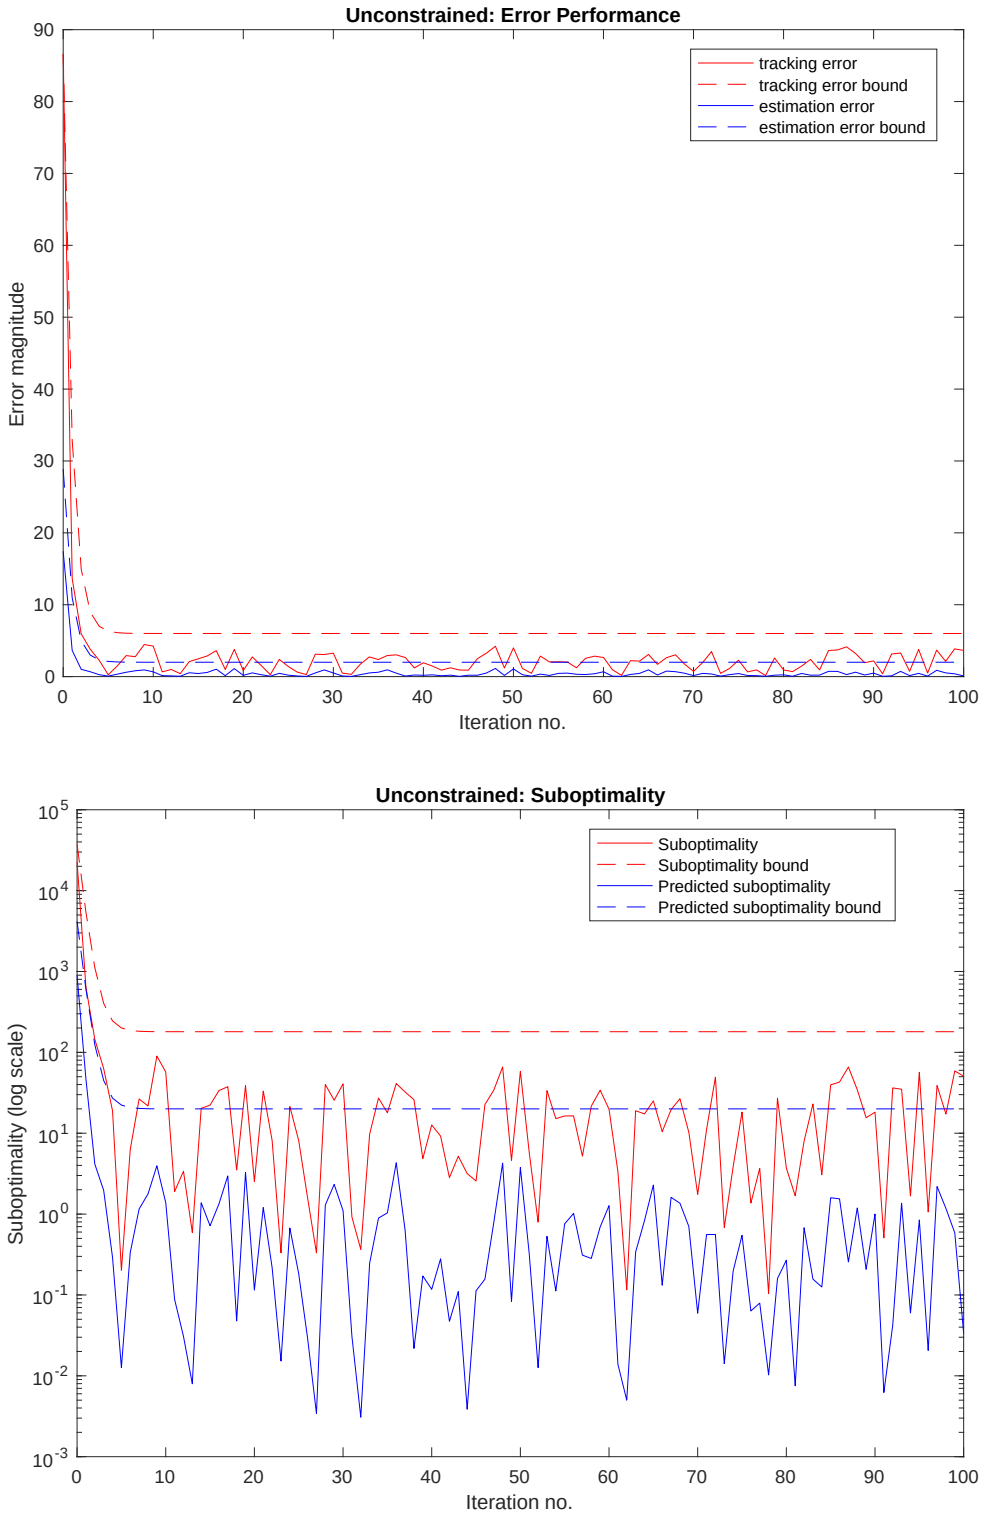

Figure 5.1: Unconstrained numerical example ($\alpha = \frac{2}{\sigma+L}$)

(a) Tracking and estimation error

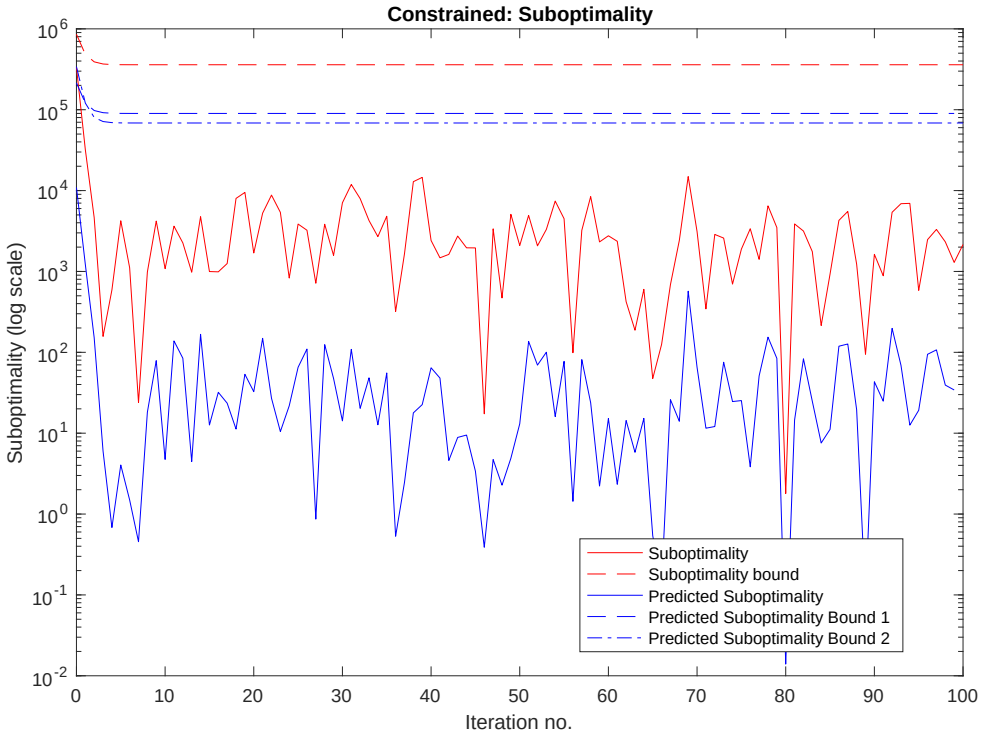

(b) Bound 1 corresponds to (5.41) and Bound 2 corresponds to (5.45).

Figure 5.2: Constrained numerical example ($\alpha = \frac{1}{L}$)

The sequence of constraints is of the form

$$X_0 = \{x \in \mathbb{R}^n \mid -u \leq x \leq u\} \tag{5.51}$$

$$\forall k \in \mathbb{N},\ X_k = \{x \in \mathbb{R}^n \mid A_k x \leq b_k,\ -u \leq x \leq u\}, \tag{5.52}$$

where $u \in \mathbb{R}^n$ defines the initial box constraints,

$$A_{k+1} = \begin{bmatrix} A_k \\ a_{k+1}^\top \end{bmatrix}, \quad b_{k+1} = \begin{bmatrix} b_k \\ \|a_{k+1}\|^2 \end{bmatrix}, \tag{5.53}$$

and $a_{k+1} \in \mathbb{R}^n$ is randomly generated based on the previous minimiser to satisfy

$$\|a_{k+1}\| \geq \|x_k^\star\| + R, \tag{5.54}$$

for some $R > \frac{2\rho\delta}{\sigma(1-\rho)}$.
Claim. Assumptions 5.2, 5.6 – 5.12 hold.

Proof. Assumptions 5.2, 5.6 and 5.7 are obviously true. Note that (5.50) implies Assumption 5.10 holds with constant $\delta_1 = 2\delta$. Now (5.53) implies $x_{k+1}^\star \in X_{k+1} \subset X_k$, and observe that

$$A_{k+1}x_k^\star = \begin{bmatrix} A_k x_k^\star \\ a_{k+1}^\top x_k^\star \end{bmatrix} \leq \begin{bmatrix} b_k \\ \|a_{k+1}\|\|x_k^\star\| \end{bmatrix} \leq \begin{bmatrix} b_k \\ \|a_{k+1}\|^2 \end{bmatrix} = b_{k+1},$$

by which $x_k^\star \in X_{k+1}$. This implies Assumption 5.11 is satisfied. By Proposition 5.4.9, Assumption 5.8 is then satisfied with $V = \frac{\delta_1}{\sigma}$. Let

$$y \in \{x \in X_k \mid \|x - x_k^\star\| \leq R\}.$$

Thus

$$\|y\| = \|y - x_k^\star + x_k^\star\| \leq R + \|x_k^\star\|.$$

Since $y \in X_k$,

$$\begin{aligned} A_{k+1}y = \begin{bmatrix} A_k y \\ a_{k+1}^\top y \end{bmatrix} &\leq \begin{bmatrix} b_k \\ \|a_{k+1}\|(R + \|x_k^\star\|) \end{bmatrix} \\ &\leq \begin{bmatrix} b_k \\ \|a_{k+1}\|^2 \end{bmatrix} = b_{k+1}, \end{aligned}$$

by (5.54), and thus $y_k \in X_{k+1}$. Noting that $R > \frac{V\rho}{1-\rho}$, this establishes Assumption 5.9. Finally, compactness of the initial box constraints guarantees

$$\|\nabla f_k(x)\| = \|Qx_k + q_k\| \leq L\|u\| + \|q_0\| + \delta := M,$$

thereby satisfying Assumption 5.12. □

Results for parameter values $n = 2,\ \sigma = 6,\ L = 8,\ \delta_1 = 300,\ u = 150 \cdot \mathbf{1} \in \mathbb{R}^n,\ R = 33.34$ are plotted in figure 5.2, which compares the actual performance with the bounds dictated by (5.20), (5.18), (5.40), (5.41) and (5.45), for a step size of $\alpha = \frac{1}{L}$. An animation of this can be found at `https://youtu.be/DV7Jb5IQDms`.

Chapter 6

Online Gradient-Free Optimisation

'Would you tell me, please, which way I ought to go from here?'
'That depends a good deal on where you want to get to,' said the Cat.
'I don't much care where—' said Alice.
'Then it doesn't matter which way you go'

Alice's Adventures in Wonderland
Lewis Carroll

This chapter continues our investigation of time-varying optimisation problems. In contrast to Chapter 5, we now focus on iterative algorithms that do not require full gradient information about cost function. In particular, the methods and analysis of Nesterov in [35] are adapted to time-varying problems. Note that while [35] obtains sub-optimality bounds in expectation, we derive bounds on the expected tracking error, focusing primarily on unconstrained problems. In Section 6.2 we assume access to the cost function directional derivatives, but in Section 6.3, only the cost function value is utilised. The notation of Chapter 5 is maintained throughout. We commence with a brief review of related works.

6.1 Related works

See [109] for a comprehensive review of time-invariant gradient-free optimisation. Here, we focus on reviewing the time-varying case, which has mainly been considered by the machine learning community. Gradient free OCO has been termed Bandit Convex Optimisation (BCO) because it corresponds to a multi-armed bandit problem [110] with convex losses. Static regret bounds for convex problems under fixed compact constraints are derived in [111,112], for iterative methods that rely on multi-point estimates of the gradient. Note that the two point version of this belongs to the class of iterations considered in [35]. Methods relying on a single-point estimate of the directional derivative are also developed in [113] and [110, Section 6.2]. A similar approach is applied to solve an energy pricing problem in [34]. The previous works all prove bounds on static regret. In contrast, the tracking error bounds derived herein relate more closely to dynamic regret, as explained in Section 5.1. Dynamic regret bounds for BCO are obtained in [33], which also incorporates time-varying compact constraints. However, the primal-dual saddle point iterations adopted in [33] do not belong to the class of iterations considered in this chapter. Our analysis also addresses unconstrained problems, which violate the compactness assumptions in BCO literature.

6.2 Directional derivative oracle

We start by assuming an oracle that returns the cost function directional derivative in some direction $u \in \mathbb{R}^n$. Define $g_k : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$,

$$g_k(x, u) := (\nabla f_k(x)^\top u)u.$$

Observe that

$$\nabla f_k(x)^\top g_k(x, u) = \nabla f_k(x)^\top (\nabla f_k(x)^\top u)u = (\nabla f_k(x)^\top u)^2,$$

and therefore $-g_k(x, u)$ is never an ascent direction of f_k for any choice of $u \in \mathbb{R}^n$. Consider iterations of the form

$$x_{k+1} = x_k - \alpha_k g_k(x_k, U_k), \tag{6.1}$$

where each $U_k \sim \mathcal{N}(0, I_n)$ is a i.i.d. random direction. It follows that the x_k are random variables obeying the Markov property.

Remark 6.2.1. Each x_{k+1} is conditionally independent of $x_0, ..., x_{k-1}$ given x_k.

It is clear that each step, in general, proceeds in a descent direction. It is further established in [35, (25)] that

$$\nabla f_k(x) = \mathbb{E}\left[g_k(x, U_k)\right], \tag{6.2}$$

which implies that, on average, (6.1) mimics gradient descent. In this chapter, we rely on Assumptions 5.2, 5.6, 5.7 and 5.8. Assumption 5.2 (smoothness), in particular, yields the inequality

$$\mathbb{E}[\|g_k(x, U_k)\|^2] \leq (n+4)\|\nabla f_k(x)\|^2, \tag{6.3}$$

also derived in [35, (32)]. Recall now the definitions of $x_k^\star$, e_k and $\bar{e}_k$ in Section 5.4. A useful property of general online line-search methods is now derived. Although this chapter is primarily concerned with unconstrained problems, constraints are included in the preliminary results below to set the stage for future work.

Lemma 6.2.1. *Let $\mathcal{F}$, $\mathcal{X}$ satisfy Assumptions 5.2, 5.6 and 5.7, and let $g_k \in \mathbb{R}^n$. If*

$$x_{k+1} = P_{X_k}[x_k - \alpha_k g_k],$$

where $\alpha_k > 0$, then

$$\bar{e}_k^2 \leq e_k^2 - 2\alpha_k g_k^\top (x_k - x_k^\star) + \alpha_k^2 \|g_k\|^2.$$

Proof. Since the projection operator is non-expansive,

$$\begin{aligned}
\bar{e}_k^2 &:= \|x_{k+1} - x_k^\star\|^2 \\
&= \|P_{X_k}(x_k - \alpha_k g_k) - x_k^\star\|^2 \\
&= \|P_{X_k}(x_k - \alpha_k g_k) - P_{X_k}(x_k^\star)\|^2 \\
&\leq \|x_k - \alpha_k g_k - x_k^\star\|^2 \\
&= \left[(x_k - x_k^\star) - \alpha_k g_k\right]^\top \left[(x_k - x_k^\star) - \alpha_k g_k\right] \\
&= e_k^2 - 2\alpha_k g_k^\top (x_k - x_k^\star) + \alpha_k^2 \|g_k\|^2.
\end{aligned}$$

□

Under Assumption 5.6 (strong convexity), each $f_k \in \mathcal{F}$ satisfies the restricted secant inequality [32, Appendix A]

$$\forall x \in X_k,\ \nabla f_k(x)^\top (x - x_k^\star) \geq \frac{\sigma}{2}\|x - x_k^\star\|^2, \tag{6.4}$$

which will be exploited in the forthcoming analysis.

Lemma 6.2.2. *Let $\mathcal{F}$, $\mathcal{X}$ satisfy Assumptions 5.2, 5.6 and 5.7. Furthermore, let $\mathcal{X} = \{X\}$. Choose $x_0 \in X$, and let*

$$x_{k+1} = P_X[x_k - \alpha_k g_k(x_k, U_k)],$$

where $U_k \sim \mathcal{N}(0, I_n)$ and $\alpha_k > 0$. Then for all $k \geq 0$,

$$\mathbb{E}[\bar{e}_k^2 \mid x_k] \leq (1 - \alpha_k \sigma) e_k^2 + \alpha_k^2 (n+4) \|\nabla f_k(x_k)\|^2.$$

Proof. First apply the result of Lemma 6.2.1. Then, taking conditional expectations,

$$\begin{aligned}
\mathbb{E}[\bar{e}_k^2 \mid x_k] &\leq e_k^2 - 2\alpha_k \mathbb{E}[g_k(x_k, U_k) \mid x_k]^\top (x_k - x_k^\star) + \alpha_k^2 \mathbb{E}\left[\|g_k(x_k, U_k)\|^2 \mid x_k\right] \\
&= e_k^2 - 2\alpha_k \nabla f_k(x_k)^\top (x_k - x_k^\star) + \alpha_k^2 \mathbb{E}\left[\|g_k(x_k, U_k)\|^2 \mid x_k\right] \\
&\leq e_k^2 - 2\alpha_k \nabla f_k(x_k)^\top (x_k - x_k^\star) + \alpha_k^2 (n+4) \|\nabla f_k(x_k)\|^2,
\end{aligned}$$

having applied both (6.2) and (6.3). Noting that $x_k \in X$ and applying the restricted secant inequality (RSI)

$$\nabla f_k(x_k)^\top (x_k - x_k^\star) \geq \frac{\sigma}{2} \|x_k - x_k^\star\|^2,$$

we obtain

$$\begin{aligned}
\mathbb{E}[\bar{e}_k^2 \mid x_k] &\leq e_k^2 - 2\alpha_k \frac{\sigma}{2} e_k^2 + \alpha_k^2 (n+4) \|\nabla f_k(x_k)\|^2 \\
&= (1 - \alpha_k \sigma) e_k^2 + \alpha_k^2 (n+4) \|\nabla f_k(x_k)\|^2.
\end{aligned}$$

□

This relationship between tracking and estimation error is analogous to (5.18). Assumption 5.8 further implies

$$\begin{aligned}
e_{k+1} &:= \|x_{k+1} - x_{k+1}^\star\| \\
&\leq \|x_{k+1} - x_k^\star\| + \|x_k^\star - x_{k+1}^\star\| \\
&\leq \bar{e}_k + V,
\end{aligned} \tag{6.5}$$

and therefore

$$\mathbb{E}[e_{k+1} \mid x_k] \leq \mathbb{E}[\bar{e}_k \mid x_k] + V. \tag{6.6}$$

6.2.1 Unconstrained strongly convex problems

Gradient magnitude is bounded by tracking error in unconstrained problems (see Lemma D.4):

$$\|\nabla f_k(x_k)\| \leq L e_k, \tag{6.7}$$

which leads us to the following result.

Theorem 6.2.3. *Let $\mathcal{X} = \{\mathbb{R}^n\}$, and let $\mathcal{F}$ satisfy Assumptions 5.2, 5.6, and 5.8. Choose $x_0 \in \mathbb{R}^n$, and consider the iterations*

$$x_{k+1} = x_k - \alpha_k g_k(x_k, U_k),$$

where $U_k \sim \mathcal{N}(0, I_n)$ and $\alpha_k > 0$. Then for all $k \geq 0$,

$$\mathbb{E}[e_{k+1} \mid x_k] \leq e_k \sqrt{1 - \alpha_k \sigma + (n+4)\alpha_k^2 L^2} + V,$$

and therefore

$$\mathbb{E}[e_{k+1}] \leq \mathbb{E}[e_k] \sqrt{1 - \alpha_k \sigma + (n+4)\alpha_k^2 L^2} + V. \tag{6.8}$$

Proof. Take the result of Lemma 6.2.2. Applying (6.7),

$$\begin{aligned} \mathbb{E}[\bar{e}_k^2 \mid x_k] &\leq (1 - \alpha_k \sigma) e_k^2 + \alpha_k^2 (n+4) L^2 e_k^2 \\ &= (1 - \alpha_k \sigma + (n+4)\alpha_k^2 L^2) e_k^2. \end{aligned}$$

It follows from Jensen's inequality that

$$\mathbb{E}[\bar{e}_k \mid x_k] \leq \sqrt{\mathbb{E}[\bar{e}_k^2 \mid x_k]} \leq e_k \sqrt{1 - \alpha_k \sigma + (n+4)\alpha_k^2 L^2},$$

and applying (6.6) we obtain

$$\mathbb{E}[e_{k+1} \mid x_k] \leq e_k \sqrt{1 - \alpha_k \sigma + (n+4)\alpha_k^2 L^2} + V,$$

and (6.8) follows by taking expectations with respect to x_k. □

The optimal step-size (with respect to the steady-state bound on tracking error) is

$$\alpha_k = \arg\min_{\alpha > 0} \left[1 - \alpha\sigma + (n+4)\alpha^2 L^2\right] = \frac{\sigma}{2(n+4)L^2}. \tag{6.9}$$

Corollary 6.2.4. *If* $\alpha_k = \frac{\sigma}{2(n+4)L^2}$ *for all* k*, then*

$$\limsup_{k\to\infty} \mathbb{E}[e_k] \leq \frac{V}{1-\sqrt{1-\frac{\sigma^2}{4(n+4)L^2}}}.$$

Proof. Substituting the optimal step-size into (6.8),

$$\mathbb{E}[e_{k+1}] \leq \mathbb{E}[e_k]\sqrt{1-\frac{\sigma^2}{4(n+4)L^2}} + V.$$

The result follows, noting that the coefficient of $\mathbb{E}[e_k]$ is less than one. □

6.3 Zeroth-order oracle

We now increase the difficulty of the problem by assuming an oracle that only provides the cost function value. In such cases, the directional derivative can still be approximated via finite differences. Define $g_{\eta,k} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$,

$$g_{\eta,k}(x,u) := \frac{f_k(x+\eta u) - f_k(x)}{\eta} u, \tag{6.10}$$

and consider iterations of the form

$$x_{k+1} = x_k - \alpha_k g_{\eta,k}(x_k, U_k), \tag{6.11}$$

where $U_k \sim \mathcal{N}(0, I_n)$. Observe that

$$g_k(x,u) = \lim_{\eta\to 0} g_{\eta,k}(x,u).$$

Hence, our analysis in the previous section corresponds to a limiting case of (6.11). It is shown in [35, (21)] that

$$\mathbb{E}[g_{\eta,k}(x,U_k)] = \nabla f_{\eta,k}(x), \tag{6.12}$$

where

$$f_{\eta,k}(x) := \mathbb{E}[f_k(x+\eta U_k)] \tag{6.13}$$

represents a smoothed version of f_k. The difference in gradients between the original and the smoothed cost is bounded by [35, (27)]

$$\|\nabla f_{\eta,k}(x) - \nabla f_k(x)\| \leq \frac{\eta}{2} L (n+3)^{\frac{3}{2}} \tag{6.14}$$

under Assumption 5.2, which also implies the magnitude bound [35, (35)]

$$\mathbb{E}[\|g_{\eta,k}(x,U_k)\|^2] \leq \frac{\eta^2}{2}L^2(n+6)^3 + 2(n+4)\|\nabla f(x)\|^2. \tag{6.15}$$

Remark 6.3.1. The bound in (6.3) is not recovered by setting $\eta = 0$ in (6.15). For this reason, performance under a zeroth-order oracle cannot yet be guaranteed to approach the performance of a directional derivative oracle as $\eta \to 0$.

6.3.1 Unconstrained strongly convex problems

Here we derive an analogous result to Theorem 6.2.3. An additional Lemma is first required.

Lemma 6.3.1. *Let $x, y, a, c \geq 0$, and $b \in \mathbb{R}$. Then*

$$x^2 \leq ay^2 + by + c \implies x \leq \sqrt{a}y + D,$$

where $D := \max\left\{\frac{b}{2\sqrt{a}}, \sqrt{c}\right\}$.

Proof. The definition of D implies both $2\sqrt{a}D \geq b$ and $D^2 \geq c$. Assuming the LHS of the implication holds,

$$\begin{aligned} x^2 &\leq ay^2 + by + c \\ &\leq ay^2 + 2\sqrt{a}Dy + D^2 \\ &= (\sqrt{a}y + D)^2. \end{aligned}$$

□

Theorem 6.3.2. *Let $\mathcal{X} = \{\mathbb{R}^n\}$, and let $\mathcal{F}$ satisfy Assumptions 5.2, 5.6, and 5.8. Choose $x_0 \in \mathbb{R}^n$, and consider the iterations*

$$x_{k+1} = x_k - \alpha_k g_{\eta,k}(x_k, U_k),$$

where $U_k \sim \mathcal{N}(0, I_n)$ and $\alpha_k > 0$. Then for all $k \geq 0$,

$$\mathbb{E}[e_{k+1} \mid x_k] \leq e_k\sqrt{c\alpha_k^2 - \sigma\alpha_k + 1} + a\alpha_k + V,$$

where $a := \frac{\eta L(n+6)^{3/2}}{\sqrt{2}}$ and $c := 2L^2(n+4)$. Therefore,

$$\mathbb{E}[e_{k+1}] \leq \mathbb{E}[e_k]\sqrt{c\alpha_k^2 - \sigma\alpha_k + 1} + a\alpha_k + V. \tag{6.16}$$

Proof. First, apply the result of Lemma 6.2.1 to obtain

$$\bar{e}_k^2 \leq e_k^2 - 2\alpha_k g_{\eta,k}(x_k, U_k)^\top (x_k - x_k^\star) + \alpha_k^2 \|g_{\eta,k}(x_k, U_k)\|^2.$$

Taking conditional expectations,

$$\mathbb{E}[\bar{e}_k^2 \mid x_k] \leq e_k^2 - 2\alpha_k \mathbb{E}[g_{\eta,k}(x_k, U_k) \mid x_k]^\top (x_k - x_k^\star) + \alpha_k^2 \mathbb{E}\left[\|g_{\eta,k}(x_k, U_k)\|^2 \mid x_k\right],$$

and then applying (6.12) and (6.15),

$$\mathbb{E}[\bar{e}_k^2 \mid x_k] \leq e_k^2 - 2\alpha_k \nabla f_{\eta,k}(x_k)^\top (x_k - x_k^\star) + 2\alpha_k^2 (n+4)\|\nabla f_k(x_k)\|^2 + \frac{(\eta\alpha_k L)^2}{2}(n+6)^3. \tag{6.17}$$

Equations (6.4) and (6.14) then imply

$$\begin{aligned}
\nabla f_{\eta,k}(x_k)^\top (x_k - x_k^\star) &= [\nabla f_{\eta,k}(x_k) - \nabla f_k(x_k)]^\top (x_k - x_k^\star) + \nabla f_k(x_k)^\top (x_k - x_k^\star) \\
&\geq \nabla f_k(x_k)^\top (x_k - x_k^\star) - \|\nabla f_{\eta,k}(x_k) - \nabla f_k(x_k)\| e_k \\
&\geq \frac{\sigma}{2} e_k^2 - \frac{\eta L}{2}(n+3)^{3/2} e_k.
\end{aligned}$$

Applying this result to (6.17), along with (6.7),

$$\begin{aligned}
\mathbb{E}[\bar{e}_k^2 \mid x_k] &\leq e_k^2 - \alpha_k \sigma e_k^2 + \alpha_k \eta L (n+3)^{3/2} e_k + 2\alpha_k^2 (n+4)\|\nabla f_k(x_k)\|^2 + \frac{(\eta\alpha_k L)^2}{2}(n+6)^3 \\
&\leq e_k^2 - \alpha_k \sigma e_k^2 + \alpha_k \eta L (n+3)^{3/2} e_k + 2\alpha_k^2 L^2 (n+4) e_k^2 + \frac{(\eta\alpha_k L)^2}{2}(n+6)^3 \\
&= [2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1] e_k^2 + [\eta L (n+3)^{3/2}\alpha_k] e_k + \frac{\eta^2 L^2 (n+6)^3}{2}\alpha_k^2.
\end{aligned} \tag{6.18}$$

Note that $\sigma \leq L < 2L\sqrt{(n+4)}$, and therefore

$$2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1 \geq \frac{1}{2} \tag{6.19}$$

for all α_k. Applying Lemma 6.3.1 to (6.18),

$$\mathbb{E}[\bar{e}_k \mid x_k] \leq \sqrt{\mathbb{E}[\bar{e}_k^2 \mid x_k]} \leq e_k \sqrt{2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1} + D_k$$

where

$$D_k := \eta\alpha_k L \max\left\{ \frac{(n+3)^{3/2}}{2\sqrt{2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1}}, \frac{(n+6)^{3/2}}{\sqrt{2}} \right\}.$$

Now (6.19) implies

$$2\sqrt{2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1} \geq \frac{2}{\sqrt{2}} = \sqrt{2},$$

which in turn implies that $D_k = \frac{\eta\alpha_k L(n+6)^{3/2}}{\sqrt{2}}$. Finally applying (6.6),

$$\mathbb{E}[e_{k+1} \mid x_k] \leq e_k\sqrt{2L^2(n+4)\alpha_k^2 - \sigma\alpha_k + 1} + D_k + V.$$

□

Corollary 6.3.3. *If $\alpha_k \in (0, \frac{\sigma}{c})$ for all k, then*

$$\limsup_{k\to\infty} \mathbb{E}[e_k] \leq \frac{a\alpha_k + V}{1 - \sqrt{c\alpha_k^2 - \sigma\alpha_k + 1}}.$$

Proof. If $0 < \alpha_k < \frac{\sigma}{c}$, then $c\alpha_k^2 - \sigma\alpha_k + 1 < 1$. □

It is straightforward but tedious to show that the optimal step size

$$\alpha_k = \arg\min_{\alpha\in\left(0,\frac{\sigma}{c}\right)} \left[\frac{a\alpha + V}{1 - \sqrt{c\alpha^2 - \sigma\alpha + 1}}\right] \tag{6.20}$$

satisfies

$$A\alpha_k^2 + B\alpha_k + C = 0, \tag{6.21}$$

where

$$\begin{aligned} A &:= (a\sigma + 2Vc)^2 - 4a^2c \\ B &:= -2V(a\sigma^2 + 2Vc\sigma + 4ac) \\ C &:= (V\sigma + 2a)^2 - 4a^2. \end{aligned}$$

The root of (6.21) in the interval $(0, \frac{\sigma}{c})$ should therefore be chosen.

6.4 Numerical experiments

The results of numerical experiments are now presented. The sequence of cost functions described in Section 5.5.1 is also adopted here, but with parameter values $\sigma = 1, L = 2, V = 0.5$ and $n = 3$. In Figure 6.1, the performance of the directional derivative iterations (6.1) is compared with the zeroth-order oracle iterations (6.11), for different values of η. The bounds (6.8) and (6.16) are also plotted. Observe that, although these bounds are only in expectation, in practice they also bound the performance of individual trials.

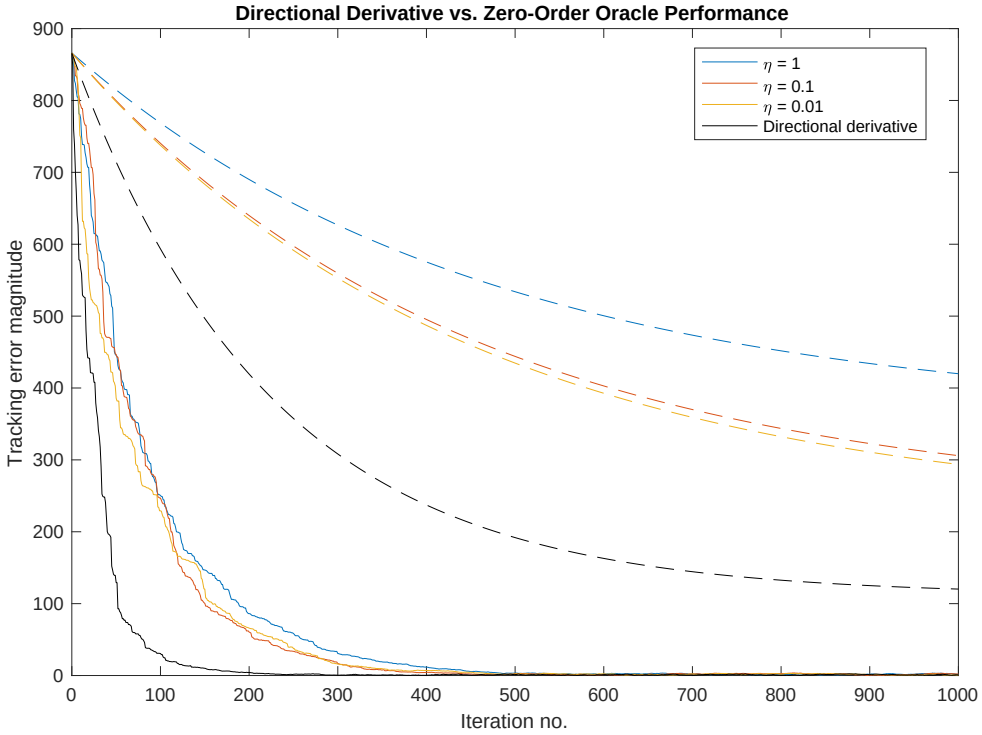

Figure 6.1: Algorithm tracking error (solid lines) compared with bounds in expectation (dashed lines).
Optimal step-sizes (6.9) and (6.20) are used in each trial.

Chapter 7

Conclusion

> My son, beware of anything beyond these.
> Of making many books there is no end,
> and much study is a weariness of the flesh.
>
> Ecclesiastes 12:12

7.1 Source localisation

The first problem considered in this thesis is the localisation of a stationary source using binary measurements. The adopted estimation procedure discretises the search region into a finite set of centres, and uses a Bayesian update rule to maintain a posterior over these centres. A theoretical analysis of this discrete posterior is presented. Conditions on the sequence of measurement locations are derived which guarantee posterior consistency when the source is coincident with a centre. The more general case of an arbitrarily located source is studied by restricting attention to periodic measurement location sequences. In this case, the algorithm asymptotically selects the indices of centres which minimise KL divergence from the true measurement probability distribution. The results described above hold for general, continuous probability-of-detection functions. Specific results are also derived for range-dependent probability-of-detection functions.

The design of D-optimal measurement locations with respect to the Bayesian Information Matrix is also formulated mathematically. Although obtaining an analytic solution is intractable, a relaxed version of the problem is proposed, which maximises the Fisher Information determinant about the

expected source location. The FIM for a range-dependent probability of detection is then derived, and a closed-form solution is established to a constrained version of the resulting optimisation problem. The effect of having inexact knowledge of the probability-of-detection function is examined by assuming knowledge of an envelope for the function. Under certain conditions, the asymptotic support of the posterior is shown to be no smaller than when the true probability-of-detection function coincides with the assumed envelope. Finally, a numerical example is simulated to supplement the theoretical results, with a control strategy proposed to guide the agents into the D-optimal measurement locations. The algorithm performance is also compared with the ML approach of [9] and the projected gradient descent algorithm in Chapter 5.

7.1.1 Future work

It would be useful to generalise the result of Chapter 2 to multi-bit measurements. While consistency is not guaranteed when the search space is partitioned into Voronoi cells, it is worth investigating whether other types of partitions can provide guarantees. Similarly, there may be modifications to the discrete Bayes recursion (2.3) that result in posterior convergence based on Euclidean distance rather than KL divergence. The general properties of Bayesian estimators with discretised parameter spaces and time-varying measurement probability distributions can also be studied without restricting attention to source localisation. Extending the estimation algorithm to deal with multiple targets is of practical importance. The closed-loop properties of the system under the control-law (2.65) should also be studied theoretically, and more advanced control strategies developed that are time-optimal or guarantee consistency. The works [114, 115] may also offer insight into designing iterative methods for optimising measurement locations with respect to the conditional BIM. Finally, a distributed implementation of the estimation algorithm should be considered, while incorporating the effects of transmission delay, asynchronous updates, and packet drop.

7.2 Security

Chapter 3 analyses the susceptibility of quantized Bayesian estimators to attack by an adversary. The adversary manipulates the measurements it sends to the estimator, with the aim of inducing a desired posterior. We

derive the set of achievable posteriors that can be induced in expectation, for both the case where the adversary is the only source of information, and for the case where the estimator has access to a truthful side-channel. The adversarial probability distribution inducing the projection of any desired posterior onto the set of achievable posteriors is posed as the solution of a convex optimisation problem. Simulation results then apply this to the source localisation problem of Chapter 2, and demonstrate that the adversary can have a powerful effect on the estimator, even in the presence of a side-channel.

7.2.1 Future work

Interesting extensions to the security analysis of Chapter 3 can be pursued by relaxing the finiteness assumption on the parameter and observation spaces. In particular, the asymptotic effect of an adversary should be considered, on a Bayesian estimator receiving an infinite sequence of measurements. The relationship between this and the the results of Chapter 3, which hold in expectation, needs to be explored. Obtaining results for continuous observation or parameter spaces are also of theoretical interest. Furthermore, the security analysis of Chapter 3 can be extended to Bayesian filtering by introducing a dynamic model for the state. Combining techniques from Chapters 2 and 3 in the context of source localisation, it should be possible to choose measurement locations that provide maximum security to the estimator by restricting the set of posteriors that can be induced by an adversary.

7.3 Connectivity preservation

Chapter 4 treats the design of decentralized control laws for guaranteeing connectivity preservation within a team of mobile agents, in the presence of piecewise continuous bounded disturbances. The goal is to maintain some fixed desired network topology by ensuring the required inter-agent distances are not exceeded. A potential function is defined based on the inter-agent distances, and a gradient-following control law is adopted to minimise it. Connectivity preservation guarantees are obtained when the desired topology is a tree. The strategy is also shown to result in a bounded control signal. Simulation results illustrate the performance of the controller on a team of agents subject to disturbances that appear as additional control signals.

7.3.1 Future work

The connectivity preservation guarantees in Chapter 4 hold when the desired network topology is a tree. Future work should focus on extending this to arbitrary network topologies, or on providing a counter-example involving a cyclic graph. The Łojasiewicz Inequalities [116, (0.1–0.2)] (not to be confused with the PL inequality) may prove useful to this end. Modifying the controller and extending the robustness guarantees to maintain connectivity through a time-varying network topology will provide more flexibility for the agents to complete other tasks. Robust connectivity-preserving control laws should also be designed for systems with more complicated dynamics, such as unicycles. Non-holonomic dynamics in particular pose a challenge to guaranteeing robustness, because the agents may not be able to move in the required direction instantaneously. However, theoretical results may still be possible if the set of allowable initial conditions is properly restricted.

7.4 Time-varying optimisation

The performance of online projected gradient descent on different classes of time-varying optimisation problems is considered in Chapter 5. Limit inferior bounds on gradient magnitude are first derived for unconstrained problems by modifying Zoutendijk's theorem and restricting the increase between consecutive cost functions. These are then strengthened to limit superior bounds on sub-optimality when the PL inequality is also assumed. Attention is then restricted to time-varying strongly convex cost functions subject to time-varying convex constraints. Tracking and estimation errors are defined, along with the analogous sub-optimality and predicted sub-optimality levels. Bounds on each of these quantities are derived assuming a bounded shift between consecutive minimisers. Conditions on the sequence of constraints are then presented that guarantee finite-time feasibility, and the shift between minimisers is also bounded in terms of the increment in consecutive cost functions gradients.

Chapter 6 adapts the analysis in [35] to time-varying optimisation problems. Tracking error bounds for unconstrained strongly convex problems are first derived for a randomised algorithm that relies on a directional derivative oracle. The analysis is then extended to a gradient free algorithm, which uses finite differences to approximate the directional derivative. The theoretical results in both chapters are illustrated using numerical examples.

7.4.1 Future work

Time-varying optimisation problems present researchers with boundless opportunities for placing bounds. Most of these remain unexplored, and a few possible directions are offered here. Chapters 5 and 6 both assume that only a single iteration is performed before new information arrives. While this is the least computationally intensive approach, the number of iterations per cost function need not, in general, be limited to one. The results herein can be generalised to permit multiple iterations. It is then of practical interest to calculate the number of iterations per cost function required to guarantee some desired convergence rate or steady-state error. The results of Section 5.4 rely on strong convexity, however relaxing this to strict convexity, and even convexity, will widen the scope for applications. Assumption 5.2 on smoothness can also be relaxed. In particular, this allows the performance of online sub-gradient methods on non-smooth convex functions to be analysed. The relationship between time-varying optimisation problems and their duals offers potential for new fundamental results. Note that the dual of any time-varying optimisation problem is convex, potentially non-smooth and time-varying, but subject to fixed non-negativity constraints. This therefore constitutes an important class of problems for research. Assumptions 5.3, 5.8 and 5.10 all limit the changes between consecutive cost-functions. It becomes important to identify conditions on the primal sequence that guarantee these assumptions are satisfied by the dual. The existing results can also be extended by considering alternatives to these assumptions.

The limit inferior bound on gradient magnitude in Section 5.3 is rather weak. Though Remark 5.3.3 suggests that it is not entirely useless, it remains to be determined whether stronger limit superior bounds can be obtained for general time-varying nonlinear cost functions. Distinct from convexity and the PL inequality, self concordance [117] is another useful function property that has not yet been exploited in the context of time-varying optimisation. Another obvious area of extension is to consider other types of optimisation algorithms, beyond just gradient descent and its zeroth-order analogue. The work of [97] considers the effect of changing co-ordinate maps on the performance of the Newton method for optimisation on manifolds. It would be interesting to further explore the relationship between changing co-ordinate maps and time-varying cost functions. This may lead to conditions under which well-designed optimisation algorithms achieve zero steady-state tracking error.

Finally, only a portion of the results in [35] have been applied to time-varying

problems in Chapter 6. Nesterov's results for constrained and non-smooth time-invariant problems remain to be extended. To this end, Lemma 6.2.1 and Lemma 6.2.2 provide a starting point for analysing constrained time-varying problems. So far, the bounds for directional derivative and gradient-free methods hold only in expectation. Further work is required to obtain steady-state bounds that hold almost surely. The work of Duchi *et al.* [118] on gradient-free stochastic time-invariant optimisation also offers useful analysis that can be exploited for time-varying problems.

Bibliography

> What has been is what will be, and what has been done is what will be done, and there is nothing new under the sun.
>
> Ecclesiastes 1:9

[1] "Autonomous vehicles are just around the corner," *The Economist*, Mar. 2018.

[2] "Russian disinformation distorts American and European democracy," *The Economist*, Feb. 2018.

[3] "Autonomous weapons are a game-changer," *The Economist*, Jan. 2018.

[4] B. Ristic, A. Gunatilaka, and R. Gailis, "Localisation of a source of hazardous substance dispersion using binary measurements," *Atmospheric Environment*, vol. 142, pp. 114–119, Oct. 2016.

[5] M. Vergassola, E. Villermaux, and B. I. Shraiman, "'Infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, pp. 406–409, Jan. 2007.

[6] J. F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Transactions on Signal Processing*, vol. 51, pp. 407–416, Feb. 2003.

[7] B. Chen, L. Tong, and P. K. Varshney, "Channel-aware distributed detection in wireless sensor networks," *IEEE Signal Processing Magazine*, vol. 23, pp. 16–26, July 2006.

[8] R. Niu and P. K. Varshney, "Target location estimation in wireless sensor networks using binary data," in *38th Annual Conference on Information Sciences and Systems*, 2004.

[9] S. Vijayakumaran, Y. Levinbook, and T. F. Wong, "Maximum Likelihood Localization of a Diffusive Point Source Using Binary Observations," *IEEE Transactions on Signal Processing*, vol. 55, pp. 665–676, Feb. 2007.

[10] B. Ristic, A. Gunatilaka, and R. Gailis, "Achievable accuracy in Gaussian plume parameter estimation using a network of binary sensors," *Information Fusion*, vol. 25, pp. 42–48, Sept. 2015.

[11] A. Shoari and A. Seyedi, "On Localization of A Non-Cooperative Target with Non-Coherent Binary Detectors," *IEEE Signal Processing Letters*, vol. 21, pp. 746–750, June 2014.

[12] H. L. Van Trees and K. L. Bell, *Detection Estimation and Modulation Theory, Part I (2nd Edition)*. Somerset, NJ, USA: John Wiley & Sons, 2013.

[13] K. Chaloner and I. Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995.

[14] A. N. Bishop, B. Fidan, B. D. Anderson, K. Doğançay, and P. N. Pathirana, "Optimality analysis of sensor-target localization geometries," *Automatica*, vol. 46, no. 3, pp. 479–492, 2010.

[15] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. Springer Texts in Statistics, New York, NY: Springer New York, 2004.

[16] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, Jan. 2004.

[17] V. V. Veeravalli and P. K. Varshney, "Distributed inference in wireless sensor networks," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Jan. 2012.

[18] M. Guerriero, L. Svensson, and P. Willett, "Bayesian Data Fusion for Distributed Target Detection in Sensor Networks," *IEEE Transactions on Signal Processing*, vol. 58, pp. 3417–3421, June 2010.

[19] P. K. Varshney, *Distributed Detection and Data Fusion*. New York, NY: Springer New York, 1997. OCLC: 853259779.

[20] J. Ok, S. Oh, J. Shin, and Y. Yi, "Optimality of Belief Propagation for Crowdsourced Classification," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, (New York, NY, USA), pp. 535–544, JMLR.org, 2016.

[21] F. Farokhi, A. M. H. Teixeira, and C. Langbort, "Estimation With Strategic Sensors," *IEEE Transactions on Automatic Control*, vol. 62, pp. 724–739, Feb. 2017.

[22] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.

[23] Y. Lin, E. Sontag, and Y. Wang, "Various results concerning set input-to-state stability," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 2, pp. 1330–1335 vol.2, Dec. 1995.

[24] F. Y. Jakubiec and A. Ribeiro, "D-MAP: Distributed Maximum a Posteriori Probability Estimation of Dynamic Systems," *IEEE Transactions on Signal Processing*, vol. 61, pp. 450–466, Jan. 2013.

[25] M. M. Zavlanos, A. Ribeiro, and G. J. Pappas, "Network Integrity in Mobile Robotic Networks," *IEEE Transactions on Automatic Control*, vol. 58, pp. 3–18, Jan. 2013.

[26] S. Y. Tu and A. H. Sayed, "Mobile Adaptive Networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, pp. 649–664, Aug. 2011.

[27] J. L. Jerez, P. J. Goulart, S. Richter, G. A. Constantinides, E. C. Kerrigan, and M. Morari, "Embedded Online Optimization for Model Predictive Control at Megahertz Rates," *IEEE Transactions on Automatic Control*, vol. 59, pp. 3238–3251, Dec. 2014.

[28] B. Gutjahr, L. Gröll, and M. Werling, "Lateral Vehicle Trajectory Optimization Using Constrained Linear Time-Varying MPC," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, pp. 1586–1595, June 2017.

[29] A. Simonetto, "Time-Varying Convex Optimization via Time-Varying Averaged Operators," *arXiv:1704.07338 [math]*, Apr. 2017.

[30] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer series in operations research and financial engineering, New York : Springer, c2006., 2006.

[31] B. T. Polyak, "Gradient methods for the minimisation of functionals," *USSR Computational Mathematics and Mathematical Physics*, vol. 3, pp. 864–878, Jan. 1963.

[32] H. Karimi, J. Nutini, and M. Schmidt, "Linear Convergence of Gradient and Proximal-Gradient Methods Under the Polyak-{\L}ojasiewicz Condition," *arXiv:1608.04636 [cs, math, stat]*, Aug. 2016.

[33] T. Chen and G. B. Giannakis, "Bandit Convex Optimization for Scalable and Dynamic IoT Management," *arXiv:1707.09060 [cs]*, July 2017.

[34] S. J. Kim and G. B. Giannakis, "An Online Convex Optimization Approach to Real-Time Energy Pricing for Demand Response," *IEEE Transactions on Smart Grid*, vol. 8, pp. 2784–2793, Nov. 2017.

[35] Y. Nesterov and V. Spokoiny, "Random Gradient-Free Minimization of Convex Functions," *Foundations of Computational Mathematics*, vol. 17, pp. 527–566, Apr. 2017.

[36] D. D. Selvaratnam, I. Shames, D. V. Dimarogonas, J. H. Manton, and B. Ristic, "Co-operative estimation for source localisation using binary sensors," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 1572–1577, Dec. 2017.

[37] B. Ristic, A. Gunatilaka, and Y. Wang, "Rao–Blackwell dimension reduction applied to hazardous source parameter estimation," *Signal Processing*, vol. 132, pp. 177–182, Mar. 2017.

[38] J. Hu, L. Xie, K. Y. Lum, and J. Xu, "Multiagent Information Fusion and Cooperative Control in Target Search," *IEEE Transactions on Control Systems Technology*, vol. 21, pp. 1223–1235, July 2013.

[39] M. Vemula, M. F. Bugallo, and P. M. Djuric, "Particle Filtering-Based Target Tracking in Binary Sensor Networks Using Adaptive Thresholds," in *2007 2nd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 17–20, Dec. 2007.

[40] O. Ozdemir, R. Niu, and P. K. Varshney, "Adaptive local quantizer design for tracking in a wireless sensor network," in *2008 42nd Asilomar Conference on Signals, Systems and Computers*, pp. 1202–1206, Oct. 2008.

[41] A. Shoari and A. Seyedi, "Localization of an uncooperative target with binary observations," in *2010 IEEE 11th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5, June 2010.

[42] R. Niu and P. K. Varshney, "Target Location Estimation in Sensor Networks With Quantized Data," *IEEE Transactions on Signal Processing*, vol. 54, pp. 4519–4528, Dec. 2006.

[43] J. K. Ghosh, M. Delampady, and T. Samanta, *An Introduction to Bayesian Analysis: Theory and Methods.* Springer texts in statistics, New York : Springer, c2006., 2006.

[44] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach.* Springer Science & Business Media, Dec. 2003.

[45] T. M. Cover and J. A. Thomas, *Elements of Information Theory.* Hoboken, N.J. : J. Wiley, 2005., 2005.

[46] W. Rudin, *Principles of Mathematical Analysis.* International series in pure and applied mathematics, New York, McGraw-Hill [1964], 1964.

[47] A. B. Tsybakov, *Introduction to Nonparametric Estimation.* Springer Series in Statistics, New York, NY: Springer New York, 2009.

[48] Y. Zheng, O. Ozdemir, R. Niu, and P. K. Varshney, "New Conditional Posterior Cramér-Rao Lower Bounds for Nonlinear Sequential Bayesian Estimation," *IEEE Transactions on Signal Processing*, vol. 60, pp. 5549–5556, Oct. 2012.

[49] L. Zuo, R. Niu, and P. K. Varshney, "Conditional Posterior Cramér–Rao Lower Bounds for Nonlinear Sequential Bayesian Estimation," *IEEE Transactions on Signal Processing*, vol. 59, pp. 1–14, Jan. 2011.

[50] H. T. Friis, "A Note on a Simple Transmission Formula," *Proceedings of the IRE*, vol. 34, pp. 254–256, May 1946.

[51] S. Marano, V. Matta, and Lang Tong, "Distributed Detection in the Presence of Byzantine Attacks," *IEEE Transactions on Signal Processing*, vol. 57, pp. 16–29, Jan. 2009.

[52] K. Agrawal, A. Vempaty, H. Chen, and P. K. Varshney, "Target localization in Wireless Sensor Networks with quantized data in the presence of Byzantine attacks," in *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, pp. 1669–1673, Nov. 2011.

[53] A. Vempaty, O. Ozdemir, K. Agrawal, H. Chen, and P. K. Varshney, "Localization in Wireless Sensor Networks: Byzantines and Mitigation Techniques," *IEEE Transactions on Signal Processing*, vol. 61, pp. 1495–1508, Mar. 2013.

[54] V. S. S. Nadendla, Y. S. Han, and P. K. Varshney, "Distributed Inference With M-Ary Quantized Data in the Presence of Byzantine Attacks," *IEEE Transactions on Signal Processing*, vol. 62, pp. 2681–2695, May 2014.

[55] J. Zhang, R. S. Blum, X. Lu, and D. Conus, "Asymptotically Optimum Distributed Estimation in the Presence of Attacks," *IEEE Transactions on Signal Processing*, vol. 63, pp. 1086–1101, Mar. 2015.

[56] Y. Mo and B. Sinopoli, "Secure control against replay attacks," in *2009 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 911–918, Sept. 2009.

[57] A. Teixeira, I. Shames, H. Sandberg, and K. H. Johansson, "A secure control framework for resource-limited adversaries," *Automatica*, vol. 51, pp. 135–148, Jan. 2015.

[58] H. Sandberg, S. Amin, and K. H. Johansson, "Cyberphysical Security in Networked Control Systems: An Introduction to the Issue," *IEEE Control Systems*, vol. 35, pp. 20–23, Feb. 2015.

[59] R. Alonso and O. Câmara, "Bayesian persuasion with heterogeneous priors," *Journal of Economic Theory*, vol. 165, pp. 672–706, 2016.

[60] A. Kolotilin, T. Mylovanov, A. Zapechelnyuk, and M. Li, "Persuasion of a Privately Informed Receiver," *Econometrica*, vol. 85, no. 6, pp. 1949–1964, 2017.

[61] V. P. Crawford and J. Sobel, "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, pp. 1431–1451, 1982.

[62] M. Battaglini, "Multiple referrals and multidimensional cheap talk," *Econometrica*, vol. 70, no. 4, pp. 1379–1401, 2002.

[63] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK ; New York: Cambridge University Press, 2004.

[64] C. Sun, G. Hu, L. Xie, and M. Egerstedt, "Robust finite-time connectivity preserving coordination of second-order multi-agent systems," *Automatica*, vol. 89, pp. 21–27, Mar. 2018.

[65] D. V. Dimarogonas and K. J. Kyriakopoulos, "Connectivity Preserving State Agreement for Multiple Unicycles," in *2007 American Control Conference*, pp. 1179–1184, July 2007.

[66] D. V. Dimarogonas and K. J. Kyriakopoulos, "Connectedness Preserving Distributed Swarm Aggregation for Multiple Kinematic Robots," *IEEE Transactions on Robotics*, vol. 24, pp. 1213–1223, Oct. 2008.

[67] Z. Feng, C. Sun, and G. Hu, "Robust Connectivity Preserving Rendezvous of Multi-robot Systems Under Unknown Dynamics and Disturbances," *IEEE Transactions on Control of Network Systems*, vol. PP, no. 99, pp. 1–1, 2016.

[68] M. Ji and M. Egerstedt, "Distributed Coordination Control of Multiagent Systems While Preserving Connectedness," *IEEE Transactions on Robotics*, vol. 23, pp. 693–703, Aug. 2007.

[69] M. M. Zavlanos, H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Hybrid control for connectivity preserving flocking," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2869–2875, 2009.

[70] Y. Dong and J. Huang, "Flocking with connectivity preservation of multiple double integrator systems subject to external disturbances by a distributed control law," *Automatica*, vol. 55, pp. 197–203, May 2015.

[71] Y. Su, "Leader-following rendezvous with connectivity preservation and disturbance rejection via internal model approach," *Automatica*, vol. 57, pp. 203–212, July 2015.

[72] A. Ajorlou and A. G. Aghdam, "Connectivity Preservation in Nonholonomic Multi-Agent Systems: A Bounded Distributed Control Strategy," *IEEE Transactions on Automatic Control*, vol. 58, pp. 2366–2371, Sept. 2013.

[73] D. V. Dimarogonas and K. J. Kyriakopoulos, "On the Rendezvous Problem for Multiple Nonholonomic Agents," *IEEE Transactions on Automatic Control*, vol. 52, pp. 916–922, May 2007.

[74] A. Ajorlou, A. G. Aghdam, and A. Jadbabaie, "A connectivity preserving containment control strategy for unicycles with static leaders," in *2012 American Control Conference (ACC)*, pp. 4186–4191, June 2012.

[75] C. D. Persis and R. Postoyan, "A Lyapunov Redesign of Coordination Algorithms for Cyber-Physical Systems," *IEEE Transactions on Automatic Control*, vol. 62, pp. 808–823, Feb. 2017.

[76] D. V. Dimarogonas and K. H. Johansson, "Bounded control of network connectivity in multi-agent systems," *IET control theory & applications*, vol. 4, no. 8, pp. 1330–1338, 2010.

[77] A. Ajorlou, A. Momeni, and A. G. Aghdam, "A Class of Bounded Distributed Control Strategies for Connectivity Preservation in Multi-Agent Systems," *IEEE Transactions on Automatic Control*, vol. 55, pp. 2828–2833, Dec. 2010.

[78] Y. Kim and M. Mesbahi, "On maximizing the second smallest eigenvalue of a state-dependent graph Laplacian," *IEEE Transactions on Automatic Control*, vol. 51, pp. 116–120, Jan. 2006.

[79] M. C. DeGennaro and A. Jadbabaie, "Decentralized Control of Connectivity for Multi-Agent Systems," in *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 3628–3633, Dec. 2006.

[80] M. M. Zavlanos and G. J. Pappas, "Controlling Connectivity of Dynamic Graphs," in *Proceedings of the 44th IEEE Conference on Decision and Control*, pp. 6388–6393, Dec. 2005.

[81] K. Srivastava and M. W. Spong, "Multi-agent coordination under connectivity constraints," in *2008 American Control Conference*, pp. 2648–2653, June 2008.

[82] M. M. Zavlanos, M. B. Egerstedt, and G. J. Pappas, "Graph-theoretic connectivity control of mobile robot networks," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1525–1540, 2011.

[83] Y. Dong and J. Huang, "Leader-Following Connectivity Preservation Rendezvous of Multiple Double Integrator Systems Based on Position Measurement Only," *IEEE Transactions on Automatic Control*, vol. 59, pp. 2598–2603, Sept. 2014.

[84] Y. Dong and J. Huang, "A leader-following rendezvous problem of double integrator multi-agent systems," *Automatica*, vol. 49, pp. 1386–1391, May 2013.

[85] J.-G. Dong, "Finite-time connectivity preservation rendezvous with disturbance rejection," *Automatica*, vol. 71, pp. 57–61, Sept. 2016.

[86] H. K. Khalil, *Nonlinear Systems.* Upper Saddle River, N.J. : Prentice Hall, 3rd ed., 2002.

[87] R. B. Bapat, *Graphs and Matrices.* Universitext, London: Springer, 2. ed ed., 2014. OCLC: 894735488.

[88] W. G. Kelley and A. C. Peterson, *The Theory of Differential Equations: Classical and Qualitative.* Universitext, New York: Springer, 2nd ed ed., 2010.

[89] A. Y. Popkov, "Gradient Methods for Nonstationary Unconstrained Optimization Problems," *Automation and Remote Control*, vol. 66, pp. 883–891, June 2005.

[90] J. W. Daniel, "Stability of the solution of definite quadratic programs," *Mathematical Programming*, vol. 5, pp. 41–53, Dec. 1973.

[91] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 928–936, 2003.

[92] E. Hazan, "Introduction to online convex optimization," *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2015. OCLC: 995130886.

[93] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online Optimization in Dynamic Environments: Improved Regret Rates for Strongly Convex Problems," *arXiv:1603.04954 [cs, math]*, Mar. 2016.

[94] M. J. Neely and H. Yu, "Online Convex Optimization with Time-Varying Constraints," *arXiv:1702.04783 [math]*, Feb. 2017.

[95] A. Simonetto, A. Mokhtari, A. Koppel, G. Leus, and A. Ribeiro, "A Class of Prediction-Correction Methods for Time-Varying Convex Optimization," *IEEE Transactions on Signal Processing*, vol. 64, pp. 4576–4591, Sept. 2016.

[96] A. Dontchev, M. Krastanov, R. Rockafellar, and V. Veliov, "An Euler–Newton Continuation Method for Tracking Solution Trajectories of Parametric Variational Inequalities," *SIAM Journal on Control and Optimization*, vol. 51, pp. 1823–1840, Jan. 2013.

[97] J. H. Manton, "A framework for generalising the Newton method and other iterative methods from Euclidean space to manifolds," *Numerische Mathematik*, vol. 129, pp. 91–125, Jan. 2015.

[98] V. Zavala and M. Anitescu, "Real-Time Nonlinear Optimization as a Generalized Equation," *SIAM Journal on Control and Optimization*, vol. 48, pp. 5444–5467, Jan. 2010.

[99] C. Sun, M. Ye, and G. Hu, "Distributed Time-Varying Quadratic Optimization for Multiple Agents Under Undirected Graphs," *IEEE Transactions on Automatic Control*, vol. 62, pp. 3687–3694, July 2017.

[100] M. Ye and G. Hu, "Distributed Seeking of Time-Varying Nash Equilibrium for Non-Cooperative Games," *IEEE Transactions on Automatic Control*, vol. 60, pp. 3000–3005, Nov. 2015.

[101] R. L. G. Cavalcante and S. Stanczak, "A Distributed Subgradient Method for Dynamic Convex Optimization Problems Under Noisy Information Exchange," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, pp. 243–256, Apr. 2013.

[102] A. Simonetto and G. Leus, "Double smoothing for time-varying distributed multiuser optimization," in *2014 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 852–856, Dec. 2014.

[103] C. Xi and U. A. Khan, "Distributed dynamic optimization over directed graphs," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 245–250, Dec. 2016.

[104] A. Simonetto and G. Leus, "Distributed asynchronous time-varying constrained optimization," in *2014 48th Asilomar Conference on Signals, Systems and Computers*, pp. 2142–2146, Nov. 2014.

[105] Q. Ling and A. Ribeiro, "Decentralized Dynamic Optimization Through the Alternating Direction Method of Multipliers," *IEEE Transactions on Signal Processing*, vol. 62, pp. 1185–1197, Mar. 2014.

[106] R. M. Gower, "Convergence Theorems for Gradient Descent." 2017.

[107] D. P. Bertsekas, *Convex Optimization Algorithms*. Athena Scientific Belmont, 2015.

[108] D. G. Luenberger, *Optimization by Vector Space Methods*. John Wiley & Sons, 1997.

[109] A. Conn, K. Scheinberg, and L. Vicente, *Introduction to Derivative-Free Optimization*. MOS-SIAM Series on Optimization, Society for Industrial and Applied Mathematics, Jan. 2009.

[110] S. Bubeck, "Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems," *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.

[111] A. Agarwal, O. Dekel, and L. Xiao, "Optimal Algorithms for Online Convex Optimization with Multi-Point Bandit Feedback.," in *COLT*, pp. 28–40, Citeseer, 2010.

[112] O. Shamir, "An optimal algorithm for bandit and zero-order convex optimization with two-point feedback," *Journal of Machine Learning Research*, vol. 18, no. 52, pp. 1–11, 2017.

[113] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: Gradient descent without a gradient," in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 385–394, Society for Industrial and Applied Mathematics, 2005.

[114] S. Liu, E. Masazade, X. Shen, and P. K. Varshney, "Adaptive non-myopic quantizer design for target tracking in wireless sensor networks," in *2013 Asilomar Conference on Signals, Systems and Computers*, pp. 1085–1089, Nov. 2013.

[115] E. Masazade, R. Niu, and P. K. Varshney, "An approximate dynamic programming based non-myopic sensor selection method for target tracking," in *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6, Mar. 2012.

[116] T. H. Colding and W. P. Minicozzi II, "Lojasiewicz inequalities and applications," *arXiv:1402.5087 [math]*, Feb. 2014.

[117] Y. Nesterov and A. Nemirovskii, *Interior-Point Polynomial Algorithms in Convex Programming*. Studies in Applied and Numerical Mathematics, Society for Industrial and Applied Mathematics, Jan. 1994.

[118] J. C. Duchi, M. I. Jordan, M. J. Wainwright, and A. Wibisono, "Optimal Rates for Zero-Order Convex Optimization: The Power of Two Function Evaluations," *IEEE Transactions on Information Theory*, vol. 61, pp. 2788–2806, May 2015.

[119] D. Williams, *Probability with Martingales*. Cambridge : Cambridge University Press, 1991., 1991.

[120] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American statistical association*, vol. 58, no. 301, pp. 13–30, 1963.

Appendices

Appendix A

Infinite products of random variables

This section contains some results concerning the convergence of infinite products of random variables, on which Chapter 2 relies.

Lemma A.1. *Let $(W_k)_{k\in\mathbb{N}}$ be a sequence of independent random variables such that $\mathbb{E}[W_k] = 0$ for all k. If the sequence is bounded, that is*

$$\exists M > 0 \text{ s.t. } \forall k,\ |W_k| < M,$$

then

$$\forall p > \frac{1}{2},\ \lim_{n\to\infty} \frac{1}{n^p} \sum_{k=1}^{n} W_k = 0 \text{ a.s..}$$

Proof. The sequence W_k is bounded, and therefore there exists $C > 0$ such that $\mathrm{Var}[X_k] \leq C$ for all k. Let $p > \frac{1}{2}$, and note that $\mathbb{E}\left[k^{-p}W_k\right] = 0$ and $\mathrm{Var}\left[k^{-p}W_k\right] \leq k^{-2p}C$. This implies

$$\forall n,\ \sum_{k=1}^{n} \mathrm{Var}\left[\frac{W_k}{k^p}\right] \leq \sum_{k=1}^{n} \frac{C}{k^{2p}}.$$

Now $2p > 1$, and therefore $\sum_{k=1}^{\infty} \mathrm{Var}\left[k^{-p}W_k\right] < \infty$ by [46, Theorem 3.28]. Applying [119, Theorem 12.2] yields

$$\sum_{k=1}^{\infty} \frac{W_k}{k^p} < \infty \text{ a.s.},$$

and Kroneckers' Lemma [119, Lemma 12.7] then implies the result. □

Corollary A.2. *Let $(X_k)_{k\in\mathbb{N}}$ be a sequence of independent, bounded random variables. Then*

$$\forall p > \frac{1}{2},\ \limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} X_k = \limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} \mathbb{E}[X_k] \ \text{a.s.}.$$

Proof. Letting $W_k := X_k - E[X_k]$, the result follows from Lemma A.1. □

Lemma A.3. *Let $(Z_k)_{k\in\mathbb{N}}$ be a sequence of independent random variables for which there exist $\alpha, \beta > 0$ such that $Z_k \in [\alpha, \beta]$ for all k. If there exists $p > \frac{1}{2}$ such that*

$$\limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} \mathbb{E}[\ln Z_k] < 0, \tag{A.1}$$

then

$$\lim_{n\to\infty} \prod_{k=1}^{n} Z_k = 0 \ \text{a.s.}$$

Proof. The random variable $\ln Z_k \in [\ln\alpha, \ln\beta] \subset (0, \infty)$. For any $p > \frac{1}{2}$,

$$\begin{aligned}\limsup_{n\to\infty} \frac{1}{n^p} \ln\left(\prod_{k=1}^{n} Z_k\right) &= \limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} \ln Z_k \\ &= \limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} \mathbb{E}[\ln Z_k] \text{ a.s.}\end{aligned}$$

by Corollary A.2. Suppose (A.1) holds, and fix a realisation $(Z_k)_{k\in\mathbb{N}}$ for which

$$c := \limsup_{n\to\infty} \frac{1}{n^p}\sum_{k=1}^{n} \mathbb{E}[\ln Z_k] < 0.$$

This implies there exists $N \in \mathbb{N}$ such that $\frac{1}{n^p}\ln\left(\prod_{k=1}^{n} Z_k\right) < \frac{c}{2} < 0$ for all $n > N$, and thus

$$\forall n > N,\ \ln\left(\prod_{k=1}^{n} Z_k\right) < \frac{n^p c}{2} < 0.$$

This in turn implies $\ln\left(\prod_{k=1}^{n} Z_k\right) \to -\infty$ as $n \to \infty$, which yields the result. □

Lemma A.4. *Let $(Z_k)_{k\in\mathbb{N}}$ be a sequence of independent random variables for which there exist $\alpha, \beta > 0$ such that $Z_k \in [\alpha, \beta]$ for all k. Let $\gamma_n := \sum_{k=1}^{n} \mathbb{E}[\ln Z_k]$. If*

$$\frac{\gamma_n}{\sqrt{n}} \to -\infty, \tag{A.2}$$

then for any $\epsilon > 0$, there exists $K \in \mathbb{N}$ such that

$$\forall n > K,\ \Pr\left(\prod_{k=1}^{n} Z_k \geq \epsilon\right) \leq \exp\left(-\frac{2\left(\ln\epsilon - \gamma_n\right)^2}{n(\ln\beta - \ln\alpha)^2}\right). \tag{A.3}$$

Proof. Let $X_k := \ln(Z_k)$, and define $S_n := \sum_{k=1}^{n} X_k$. The X_k are independent, and therefore

$$\mathbb{E}[S_n] = \sum_{k=1}^{n} \mathbb{E}\left[\ln Z_k\right] = \gamma_n.$$

Note that $\prod_{k=1}^{n} Z_k = \exp(S_n)$, and therefore

$$\left(\prod_{k=1}^{n} Z_k \geq \epsilon\right) \iff \left(S_n \geq \ln\epsilon\right),$$

where $\epsilon > 0$. Therefore for any $\epsilon > 0$,

$$\Pr\left(\prod_{k=1}^{n} Z_k \geq \epsilon\right) = \Pr\left(\frac{S_n - \gamma_n}{n} \geq \frac{\ln\epsilon - \gamma_n}{n}\right). \tag{A.4}$$

If (A.2) holds, then

$$\forall\epsilon > 0,\ \exists K \in \mathbb{N} \text{ s.t. } \forall n > K,\ \gamma_n < \ln\epsilon,$$

and therefore $\dfrac{\ln\epsilon - \gamma_n}{n} > 0$ for all $n > K$. Furthermore, $X_k \in [\ln(\alpha), \ln(\beta)]$ for all k. We can therefore apply Hoeffding's inequality [120, Theorem 2] for all $n \geq K$:

$$\Pr\left(\frac{S_n - \gamma_n}{n} \geq \frac{\ln(\epsilon) - \gamma_n}{n}\right) \leq \exp\left(-\frac{2(\ln\epsilon - \gamma_n)^2}{n(\ln\beta - \ln\alpha)^2}\right).$$

Combining this with (A.4) gives us (A.3). □

Appendix B

Fisher Information

A general expression for the FIM in Chapter 2 is derived below, based on the likelihood function g defined in (2.1). The log-likelihood gradient is given by

$$\begin{aligned}\nabla_{\mathbf{s}} \ln g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k) &= \frac{\nabla_{\mathbf{s}} g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k)}{g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k)} \\ &= \left[\frac{\nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{\ell(\mathbf{s}, \boldsymbol{\xi}_k)}\right]^{d_k} \left[\frac{\nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{\ell(\mathbf{s}, \boldsymbol{\xi}_k) - 1}\right]^{1-d_k} \\ &= \frac{\nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{(-1)^{1-d_k}[\ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{d_k}[1 - \ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{1-d_k}},\end{aligned}$$

and its Hessian,

$$\begin{aligned}&\nabla_{\mathbf{s}}^2 \ln g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k) \\ &= \frac{\partial}{\partial \mathbf{s}} \left[\nabla_{\mathbf{s}} \ln g(d_k \mid \mathbf{s}; \boldsymbol{\xi}_k)\right] \\ &= \left[\frac{\ell(\mathbf{s}, \boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^2 \ell(\mathbf{s}, \boldsymbol{\xi}_k) - \nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^\top \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{\ell(\mathbf{s}, \boldsymbol{\xi}_k)^2}\right]^{d_k} \\ &\quad \cdot \left[\frac{[\ell(\mathbf{s}, \boldsymbol{\xi}_k) - 1]\nabla_{\mathbf{s}}^2 \ell(\mathbf{s}, \boldsymbol{\xi}_k) - \nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^\top \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{[1 - \ell(\mathbf{s}, \boldsymbol{\xi}_k)]^2}\right]^{1-d_k} \\ &= \frac{\nabla_{\mathbf{s}}^2 \ell(\mathbf{s}, \boldsymbol{\xi}_k)}{(-1)^{1-d_k}[\ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{d_k}[1 - \ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{1-d_k}} - \frac{\nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)\nabla_{\mathbf{s}} \ell(\mathbf{s}, \boldsymbol{\xi}_k)^\top}{[\ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{2d_k}[1 - \ell(\mathbf{s}, \boldsymbol{\xi}_k)]^{2(1-d_k)}}.\end{aligned}$$

The FIM for a single reading is then

$$
\begin{aligned}
J(\mathbf{s};\boldsymbol{\xi}_k) &= -\mathbb{E}\left[\nabla_{\mathbf{s}}^2 \ln g(d_k \mid \mathbf{s};\boldsymbol{\xi}_k)\right] \\
&= -\sum_{d\in\{0,1\}} g(d \mid \mathbf{s};\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^2 \ln g(d \mid \mathbf{s};\boldsymbol{\xi}_k) \\
&= -\ell(\mathbf{s},\boldsymbol{\xi}_k)\left[\frac{\ell(\mathbf{s},\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^2\ell(\mathbf{s},\boldsymbol{\xi}_k) - \nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^\top\ell(\mathbf{s},\boldsymbol{\xi}_k)}{\ell(\mathbf{s},\boldsymbol{\xi}_k)^2}\right] \\
&\quad - [1-\ell(\mathbf{s},\boldsymbol{\xi}_k)]\left[\frac{[\ell(\mathbf{s},\boldsymbol{\xi}_k)-1]\nabla_{\mathbf{s}}^2\ell(\mathbf{s},\boldsymbol{\xi}_k) - \nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}^\top\ell(\mathbf{s},\boldsymbol{\xi}_k)}{[1-\ell(\mathbf{s},\boldsymbol{\xi}_k)]^2}\right] \\
&= \frac{\nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)\nabla_{\mathbf{s}}\ell(\mathbf{s},\boldsymbol{\xi}_k)^\top}{\ell(\mathbf{s},\boldsymbol{\xi}_k)[1-\ell(\mathbf{s},\boldsymbol{\xi}_k)]}.
\end{aligned}
$$

Appendix C

Time-varying quadratics

Here we consider a special case of Section 5.4 involving unconstrained strongly convex quadratic cost functions. Assume only first order changes over time.

Proposition C.1. *Let* $\mathcal{X} = \{\mathbb{R}^n\}$, *and*

$$f_k(x) = \frac{1}{2}x^\top Q x + q_k^\top x, \tag{C.1}$$

where $Q = Q^\top \succ 0$, *and* $q_k \in \mathbb{R}^n$ *for all* $k \geq 0$. *Let* $x_0 \in \mathbb{R}^n$, *and consider iterations of the form*

$$x_{k+1} = x_k - \alpha \nabla f_k(x_k).$$

Then for all k,

$$x_{k+1} - x_{k+1}^\star = (I - \alpha Q)(x_k - x_k^\star) + Q^{-1}(q_{k+1} - q_k).$$

Proof. The gradient $\nabla f_k(x) = Qx_k + q_k$. For the unconstrained problem, this implies

$$x_k^\star = -Q^{-1}q_k,$$

and therefore

$$q_k = -Qx_k^\star.$$

Applying the gradient descent iterations with constant step-size $\alpha > 0$,

$$\begin{aligned} x_{k+1} &= x_k - \alpha \nabla f_k(x_k) \\ &= x_k - \alpha(Qx_k + q_k) \\ &= x_k - \alpha Q x_k + \alpha Q x_k^\star. \end{aligned}$$

This yields the relationship

$$\begin{aligned} x_{k+1} - x^\star_{k+1} &= (I - \alpha Q)x_k + \alpha Q x^\star_k - x^\star_{k+1} \\ &= (I - \alpha Q)x_k + \alpha Q x^\star_k + Q^{-1}q_{k+1} \\ &= (I - \alpha Q)x_k + \alpha Q x^\star_k + Q^{-1}q_{k+1} - Q^{-1}q_k + Q^{-1}q_k \\ &= (I - \alpha Q)x_k + \alpha Q x^\star_k + Q^{-1}q_{k+1} - Q^{-1}q_k - x^\star_k \\ &= (I - \alpha Q)x_k + (\alpha Q - I)x^\star_k + Q^{-1}(q_{k+1} - q_k) \\ &= (I - \alpha Q)(x_k - x^\star_k) + Q^{-1}(q_{k+1} - q_k). \end{aligned}$$

□

Let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n > 0$ denote the eigenvalues of Q. Since Q is unitarily diagonalisable, it follows that

$$\|I - \alpha Q\| = \max\{|1 - \alpha\lambda_i| \mid i \in \mathbb{N}_n\} = \max\{|1 - \alpha\lambda_1|, |1 - \alpha\lambda_n|\},$$

and $\|Q^{-1}\| = \frac{1}{\lambda_n}$. Now $\mathcal{F}$ consists of cost functions of the form (C.1), which implies Assumptions 5.2 and 5.6 are satisfied with constants $L = \lambda_1$ and $\sigma = \lambda_n$ respectively.

Corollary C.2. *If $\|q_{k+1} - q_k\| \leq \delta_1$ for all k, then*

$$e_{k+1} \leq \rho e_k + \frac{\delta_1}{\sigma},$$

where $\rho := \max\{|1 - \alpha L|, |1 - \alpha\sigma|\}$.

Observe that exactly the same bound follows Theorem 5.4.2 and Proposition 5.4.9.

Appendix D

Technical results

D.1 Real analysis

Lemma D.1. *If $f : \mathbb{R}^n \to \mathbb{R}$ is twice continuously differentiable and satisfies*

$$\exists L > 0,\ \forall x, y \in \mathbb{R}^n,\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|,$$

then for all $x, x^\star \in \mathbb{R}^n$,

$$f(x) - f(x^\star) \leq \|\nabla f(x^\star)\|\|x - x^\star\| + \tfrac{1}{2}L\|x - x^\star\|^2.$$

Proof. Let $e := x - x^\star$. Then by Taylor's theorem, there exists $t \in (0, 1)$ such that

$$f(x) = f(x^\star) + \nabla f(x^\star)^\top e + \frac{1}{2} e^\top \nabla^2 f(x^\star + te)e.$$

The result then follows. □

Lemma D.2. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be continuously differentiable. Let $X \subset \mathbb{R}^n$ be convex and non-empty. Then*

$$\forall x, x^\star \in X,\ f(x) - f(x^\star) \leq M\|x - x^\star\|,$$

where $M := \sup\{\|\nabla f(x)\| \mid x \in X\}$.

Proof. By the mean value theorem, there exists $t \in (0, 1)$ such that

$$f(x) = f(x^\star) + \nabla f(tx + (1 - t)x^\star)^\top (x - x^\star).$$

□

Lemma D.3. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be twice continuously differentiable, and let ∇f be L-Lipschitz on $\mathbb{R}^n$. Furthermore, suppose f attains a minimum value $f^\star$. If*

$$\exists \mu > 0,\ \forall x \in \mathbb{R}^n,\ \frac{1}{2}\|\nabla f(x)\|^2 \geq \mu(f(x) - f^\star),$$

then $\mu \leq L$.

Proof. Choose $x \in \mathbb{R}^n \setminus \arg\min\{f(y) \mid y \in \mathbb{R}^n\}$, and define $x' := x - \frac{1}{L}\nabla f(x)$. By Taylor's Theorem,

$$\begin{aligned} f(x') &\leq f(x) - \frac{1}{L}\|\nabla f(x)\|^2 + \frac{1}{2}\|\tfrac{1}{L}\nabla f(x)\|^2 L \\ &= f(x) - \frac{1}{2L}\|\nabla f(x)\|^2. \end{aligned}$$

Therefore

$$\mu(f(x) - f^\star) \leq \frac{1}{2}\|\nabla f(x)\|^2 \leq L(f(x) - f(x')),$$

which implies

$$\frac{\mu}{L} \leq \frac{f(x) - f(x')}{f(x) - f^\star} \leq 1,$$

because $f^\star \leq f(x')$ by definition. □

Lemma D.4. *Let $X \subset \mathbb{R}^n$ be closed and convex. If $f : \mathbb{R}^n \to \mathbb{R}$ is continuously differentiable and satisfies*

$$\exists L > 0,\ \forall x, y \in X,\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|,$$

then for any $x^\star \in \arg\min\{f(x) \mid x \in X\}$ and $x \in X$,

$$|\|\nabla f(x^\star)\| - \|\nabla f(x)\|| \leq L\|x - x^\star\|$$

Proof. The result follows directly from the reverse triangle inequality. □

D.2 Wolfe condition guarantees

Lemma D.5 (Curvature Condition)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be continuously differentiable, and strongly convex with modulus $\sigma > 0$. Let $0 < \alpha < \frac{1}{\sigma}$, $x \in \mathbb{R}^n$ and $p = -\nabla f(x)$. Then, any $c_2 \in (0, 1 - \alpha\sigma]$ satisfies the curvature condition*

$$\nabla f(x + \alpha p)^\top p \geq c_2 \nabla f(x)^\top p.$$

Proof. By strong convexity of f, for any $x, x' \in \mathbb{R}^n$,

$$[\nabla f(x') - \nabla f(x)]^\top (x' - x) \geq \sigma \|x' - x\|^2.$$

Letting $x' = x + \alpha p$,

$$[\nabla f(x') - \nabla f(x)]^\top (\alpha p) \geq \sigma \|\alpha p\|^2 = \alpha^2 \sigma \|p\|^2.$$

If $p = -\nabla f(x)$, this can be written as

$$[\nabla f(x') - \nabla f(x)]^\top p \geq -\alpha\sigma \nabla f(x)^\top p,$$

which is equivalent to

$$\nabla f(x')^\top p \geq (1 - \alpha\sigma) \nabla f(x)^\top p.$$

Let $\alpha < \frac{1}{\sigma}$. If $c_2 \in (0, 1 - \alpha\sigma]$, then $1 - \alpha\sigma \geq c_2$, which implies

$$\nabla f(x')^\top p \geq c_2 \nabla f(x)^\top p,$$

as desired. □

Lemma D.6 (Armijo Condition)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be twice continuously differentiable, with a gradient ∇f that satisfies the Lipschitz condition*

$$\exists L > 0,\ \forall x, y \in \mathbb{R}^n,\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|.$$

Let $0 < \alpha < \frac{2}{L}$, $x \in \mathbb{R}^n$ and $p = -\nabla f(x)$. Then, any $c_1 \in (0, 1 - \frac{\alpha L}{2}]$ satisfies the Armijo condition

$$f(x + \alpha p) \leq f(x) + c_1 \alpha \nabla f(x)^\top p.$$

Proof. The result is obvious if $\nabla f(x) = 0$. Assume, now, that $\nabla f(x) \neq 0$. Using Taylor's theorem

$$\exists t \in (0, 1),\ f(x + \alpha p) = f(x) + \alpha \nabla f(x)^\top p + \frac{\alpha^2}{2} p^\top \nabla^2 f(x + t\alpha p) p.$$

Choose $\alpha \in (0, \frac{2}{L})$ and $c_1 \in (0, 1 - \frac{\alpha L}{2}]$, which implies

$$c_1 - 1 + \frac{\alpha L}{2} \leq 0. \tag{D.1}$$

Letting $H := \nabla^2 f(x + t\alpha p)$ and $p = -\nabla f(x)$,

$$\begin{aligned} f(x+\alpha p) &= f(x) + c_1\alpha\nabla f(x)^\top p + (1-c_1)\alpha\nabla f(x)^\top p + \frac{\alpha^2}{2}p^\top Hp \\ &= f(x) + c_1\alpha\nabla f(x)^\top p - (1-c_1)\alpha\|p\|^2 + \frac{\alpha^2}{2}\|p\|^2\frac{p^\top Hp}{p^\top p} \\ &\leq f(x) + c_1\alpha\nabla f(x)^\top p - (1-c_1)\alpha\|p\|^2 + \frac{\alpha^2}{2}\|p\|^2 L \\ &\leq f(x) + c_1\alpha\nabla f(x)^\top p + \alpha\|p\|^2\left(c_1 - 1 + \frac{\alpha L}{2}\right) \\ &\leq f(x) + c_1\alpha\nabla f(x)^\top p, \end{aligned}$$

where the last step is obtained by applying (D.1). □

Theorem D.7. *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is twice continuously differentiable, and satisfies*

$$\forall x \in \mathbb{R}^n,\ \sigma I \preceq \nabla^2 f(x) \preceq LI,$$

for some $\sigma, L > 0$. Define $\nu := \max\{\frac{L}{2}, \sigma\}$ and let $0 < \alpha < \frac{1}{\nu}$. Choose any $c_1 \in I_1 := (0, 1-\alpha\nu)$ and $c_2 \in I_2 := (c_1, 1-\alpha\sigma]$. Then

$$I_1, I_2 \neq \emptyset, \quad c_1 \in (0,1), \quad c_2 \in (c_1, 1).$$

Furthermore, for any $x \in \mathbb{R}^n$ and for $p = -\nabla f(x)$,

$$f(x+\alpha p) \leq f(x) + c_1\alpha\nabla f(x)^\top p \tag{D.2}$$

$$\nabla f(x+\alpha p)^\top p \geq c_2\nabla f(x)^\top p. \tag{D.3}$$

Proof. First note $\alpha < \min\{\frac{2}{L}, \frac{1}{\sigma}\}$, and therefore the hypotheses of Lemma D.5 and Lemma D.6 are satisfied. Furthermore, $0 < \alpha\nu < 1$ and therefore $I_1 \subset (0,1)$ is non-empty. Clearly $I_2 \subset (c_1, 1)$. Since $\nu \geq \sigma$, it also holds that $c_1 < 1 - \alpha\nu \leq 1 - \alpha\sigma$, and therefore I_2 is also non-empty. Finally, Lemma D.6 implies (D.2), and Lemma D.5 implies (D.3). □

Remark D.2.1. Theorem D.7 implies that if $\mathcal{F}$ satisfies Assumptions 5.2 and 5.6, and $\alpha_k = \alpha < \frac{1}{\nu}$ for all k, then there always exist c_1, c_2 that satisfy (5.5) - (5.6).