INTEGRATED TASKING, PROCESSING, AND ORBIT DETERMINATION FOR OPTICAL SENSORS IN A SPACE SITUATIONAL AWARENESS FRAMEWORK

A Dissertation
Presented to
The Academic Faculty

By

Timothy S. Murphy

In Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology

May 2018

INTEGRATED TASKING, PROCESSING, AND ORBIT DETERMINATION FOR OPTICAL SENSORS IN A SPACE SITUATIONAL AWARENESS FRAMEWORK

Approved by:

Dr. Marcus J. Holzinger, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Dr. E. Glenn Lightsey
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Eric N. Johnson
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Brien R. Flewelling
Chief SSA Architect
ExoAnalytic Solutions

Dr. Chris Sabol
Researcher
Air Force Research Lab

Date Approved: March 28, 2018

“If we accept the person we are when we fall, the journey ends. That failure becomes our destination. To love the journey is to accept no such end. I have found, through painful experience, that the most important step a person can take is always the next one.”

Brandon Sanderson

To Shelly.

You've kept me surprisingly sane during this half-decade or so.

ACKNOWLEDGEMENTS

There are several people who deserve acknowledgement here. First, of course, I would like to thank my advisor, Marcus, whom for some reason decided to take me under his wing. I hope I'm proving it wasn't a bad decision. I also need to thank my second mentor, Brien, whose influence has kept my work from wandering too far into the weeds more than once. My fellow grad students, Ryan S. Coder, Andres, Johnny "The Worst" Worthy, Shahezad and Julian have made the day in day out life of a researcher not just tolerable but actually enjoyable. Shelly, you already have a dedication so I don't think I am going to acknowledge you too. Don't be greedy. I think they already know this, but I owe everything I have accomplished to the constant guidance of my parents, David and Barbara Murphy. Finally, my real best friend Belladonna has been an amazing companion during my many months of writing this document.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Space situational awareness (SSA) depends on having networks of telescopes taking large numbers of observations of myriad space objects to perform orbit estimation and object characterization. This involves many engineering problems, including choosing search patterns to look for objects, choosing telescope parameters to take the optical data, processing images to look for objects, and turning that information into orbit estimates. In the current state of the art, these technical problems are often treated as insular challenges, and are solved without regard to each other. This dissertation looks at telescope tasking, image processing, and statistical orbit estimation as a single complete problem. This work first analyzes the telescope tasking problem on advanced astrodynamic priors. Next, telescope tasking is analyzed with respect to specific image processing and detection methods. A specific estimation algorithm is used to optimize detection sensitivity. Finally, advanced astrodynamic priors are used to inform further detection methods. The contributions in this work find ways to leverage knowledge of the full SSA problem to make significant improvements on modeling and solving the individual problems. The end result is a more complete SSA system and data stream.

CHAPTER 1
INTRODUCTION

1.1 Motivation

As the use of space continues to grow, more man-made debris will accumulate in earth orbit. This has lead to the growth of research in Space situational awareness (SSA), defined as obtaining all knowledge relevant to the space domain as needed to allow spacecraft and their operators to make informed decisions [1]. One primary mission of SSA is to efficiently identify orbits of any and all objects, regardless of size, orbit regime, and location. The Joint Space Operations Center (JSpOC) under U.S. Strategic Command operates the Space Surveillance Network (SSN) which tracks, as of writing this, upwards of 22,000 space objects of diameters greater than 10 cm. This number is increasing all the time. The SSN takes around 400,000 observations each day with radar, optical, and space-based sensors [2] [1]. The current and future growth of radar systems will allow much smaller objects in low altitude orbits to be tracked, greatly expanding the current catalog of objects [3] [2]. In particular, the space fence is expected to track objects the size of a marble and increase the number of objects tracked in low Earth orbit (LEO) by an order of magnitude [4]. At higher altitudes, radar is difficult to use effectively, so operators rely on electro-optical sensors (EOS) to form orbit estimates [5]. These range from large, expensive, highly sensitive EOS to smaller, more cost efficient networks of EOS [6]. There is also a growing interest in utilizing space-based optical sensors as a new means of collecting data [7] [8]. This includes the Space-Based Space Surveillance (SBSS) system and GEO Space Situation Awareness Program (GSSAP) which are satellite constellations designed specifically to track other satellites. This dissertation primarily focuses on optical sensors, but many

[1] space-track.org (accessed March, 2018)

[2] https://www.leolabs.space/about (accessed March, 2018)

of the techniques can be extended to other measurement frameworks. This work in general focuses on probing questions about how optical systems can and should be used, as opposed to proposing specific techniques to solve specific situations. The growing number of sensors in SSA can solve some of the many problems we face, but brute force coverage should not outweigh fundamental research on how these sensors should be utilized.

The importance of fundemental research into how to perform SSA cannot be understated. Events like the Iridium-Cosmos collision highlight not only the possibility of future collisions but the potential for new space debris to be created [9]. At higher altitudes, small regions of advantageous orbits, in particular geostationary orbits (GEO), continue to be prioritized. Debris at GEO altitude is viewed as a major problem due to the unpredictable dynamics of high area-to-mass ratio (HAMR) objects [10], the dense arrangement of satellites, and the high economic value of satellite orbit slots [11]. The importance of this work is tied to how the United States of America, along with much of the rest of the world, relies on satellites for communication, weather forecasting, agriculture, and fundamental Earth science [11]. Because these space regions are highly important and can be easy to disrupt, there exists an obvious military interest in this research as well [12].

However, the current SSN system has several areas in need of expansion. The Air Force, among others, has expressed a desire to move toward "Trusted, Highly Autonomous Decision-Making Systems" [13]. Much of the previous design work on optical systems has come out of a collaboration of astrodynamics specialists and astronomers, while the community of researchers working solely in SSA is a small but growing community. This means that the optimal arrangement for an optical system designed specifically for SSA is still being actively researched [14]. Such design problems rely on and inform research and analysis into how optical sensors can, could, and should be used to achieve objectives. For EOS, the problem is not lack of detections but excess of detections. Exquisite sensors scan the sky and collect detections which do not match an object in the SSN, due to both the clutter of space objects and the difficulty of correlation [15]. However, these observations

do not provide full orbits, but rather large admissible regions in which a hypothesized orbit could exist. The operator can either task a sensor on this object for a large period of time while an orbit is estimated, or, more often, leave the detection as an uncorrelated track (UCT). Attempting to correlate two UCTs blindly is an active area of research but often requires high performance computing [16].

The other side of this problem involves the types of sensors that are used. On one end, there are exquisite, expensive sensors at premium locations such as the 3.6 meter telescope, Advanced Electro-Optical System (AEOS), at AMOS [17, 18]. While such sensors can often detect many exceedingly small objects, there will never be enough of these sensors to make the follow-up problem practical. On the other end, there are inexpensive sensors and sensors in areas with light pollution. The follow-up performance of less exquisite sensors can be improved with sufficiently precise priors. Specifically, Raven-class telescopes are built with commercial off-the-shelf parts and deployed across the globe in varying conditions for comparatively low costs [6]. Such sensors can be deployed in large networks, but can struggle to detect exceedingly dim objects at a reliable and precise cadence. An advantage provided by large networks of sensors is the ability to quickly obtain range estimation through geometrically diverse sensors [19]. Obtaining this range estimate requires enabling the tasking of two sensors on a single unknown object without an a priori range estimate.

Coordination of the diverse sensors available in SSA is difficult. At the most basic level, this is knowing when and how follow-up can and should be performed. For less exquisite EOS, this means maximizing detection sensitivity in follow up situations. For highly exquisite EOS, this means optimizing how the sensor is tasked to maximize detection opportunities. Understanding follow-up tasking requires knowing how orbit estimation is initialized and how statistical priors evolve over time. Maximizing follow-up detection sensitivity requires understanding how tasking a sensor interacts with detection methods. All of this work requires understanding how to make sensitive detections on point sources

in images, with considerations as to what SSA images look like.

What this leads to is a vision for SSA operations that relies on an interconnected framework of sensor tasking, image processing, astrodynamics, and estimation. While a wealth of research in each of these topics applied to SSA exists, insufficient research has been done considering the problem as a whole. Detection methods from the field of computer vision have been applied to SSA images, but have not been specialized to the unique properties of a satellite point source in an image. EOS tasking that inherently considers the dynamics of a statistical prior of a space object is a new area of research. EOS data is taken without consideration to the detection methods that extract information from an image. This dissertation proposes and solves a more complete SSA problem.

1.1.1 Broad Categories of Search

This section provides a broad definition of search. Search in this context is the process of looking for an object, or group of objects, which is thought to exist within a set of orbits, known as the "search set". The set of orbits could be defined through a variety of methods. In a lost object scenario, the prior probability density function (PDF) that takes into account conservative uncertainty in dynamics, can have a search set defined directly by bounding the PDF. Compared to assuming a Gaussian distribution, this method is preferred on dynamics which are difficult to quantify. Set-based methods, such as an admissible region, easily define a search set [20]. Objects which have performed an unknown maneuver can be used to define a set of orbits by providing a limiting bound on total fuel expenditure [21] [22]. Object maneuver limits can also be applied to other anomalous events such as a satellite break up. Finally, a search set can be defined more broadly as a region of orbits where objects may exist for object discovery methods. Furthermore, this definition is not restricted to man-made space objects; search for particular classes of asteroids, search in air or sea domains, and a variety of other problems all fit within this framework. Here is acknowledged two categories in which the search described in this dissertation would be

useful.

The first is the problem of how to search a volume of state space to observe both known and unknown objects in it. A common example is to attempt to maintain custody of objects in or near GEO, by tasking a sensor to constantly scan through that region. An example of a more unique version of this set, which is explored in Chapter 6, is asset protection. This is done by defining the set of intercept orbits and searching it to assure it does not contain a hostile threat. There are a variety of other ways to pose this problem, but the common theme is constantly searching a large set of orbits to find previously unknown objects.

The second version of search is to perform follow-up on an object(s) when its location is not entirely known. A prime example of this is the follow-up on an object with a prior characterized by an admissible region. This can be broadened to searching for debris after a break-up event, searching for an object that has maneuvered, and reacquiring a catalog object with sufficiently large uncertainties. Note that such uncertain priors tend to spatially diverge over time, causing a cascading problem of difficult-to-model uncertainty growth. If the search region for this object is large, the search strategies may be similar to that of the custody problem, while for a small search region it becomes trivial.

1.2 Thesis Statement

> Combined integration of advanced modeling methods for astrodynamic priors, high-sensitivity point source detection methods, and optical data collection strategies improves each individual method by leveraging an analysis of the combined problem.

To summarize, this dissertation explores the following. Coordinating a large network of EOS for an effective SSN requires several novel capabilities. The first piece of this work involves the search set, or the set of orbits which may contain relevant object(s). A tasking algorithm is needed to optimally choose a sensor trajectory through the sky that a sensor

should take to search a set of orbits. Because of the malleable definition of a set, this allows EOS to be tasked on a wide variety of search and reacquisition problems. Next, when taking actual data the sensor exposure time, slew rate, and campaign length need to be chosen to optimize the quality of image data. These tasking parameters are chosen with respect to the detection and estimation algorithms themselves, which all relate back to a maximum likelihood method. Next, the object detection algorithms should be as sensitive as possible. This enables a larger network of lower cost telescopes to be deployed, and ensures that performance is robust to light pollution, enabling new telescope locations. These types of networks are needed to allow the kinds of sensor architectures which support interesting handoff and reacquisition problems. Finally, to make proper telescope communication, hand-off, and long term reacquisition possible, detection algorithms should utilize any prior information (search set) on a particular object or class of objects for more sensitive and efficient detection. This supports hand-off between arbitrary locations and longer delay times before reacquisition from the detection side of the problem. Figure 1.1 gives an outline of how these systems fit into one larger plan.

1.3 Related Work

This section details prior work on the different subjects explored in this dissertation.

A variety of sensor tasking strategies exist, which look at tasking in terms of catalog upkeep and track maintenance [23–26]. These techniques tend to look at the strategic tasking problem, that is, how to use a network of sensors to look at a catalog of objects. These methods often do not account for shape or size of the region or prior being searched. This dissertation looks to analyze the tactical tasking problem: given a prior on an object and a sensor, how should the sensor search for the object? In particular, this work extends a particle-based search strategy proposed by Hobson [27–29]. Existing methods are based on a greedy maximum probability observation technique, which attempts observing the densest region of probability. This dissertation expands this methodology by analyzing

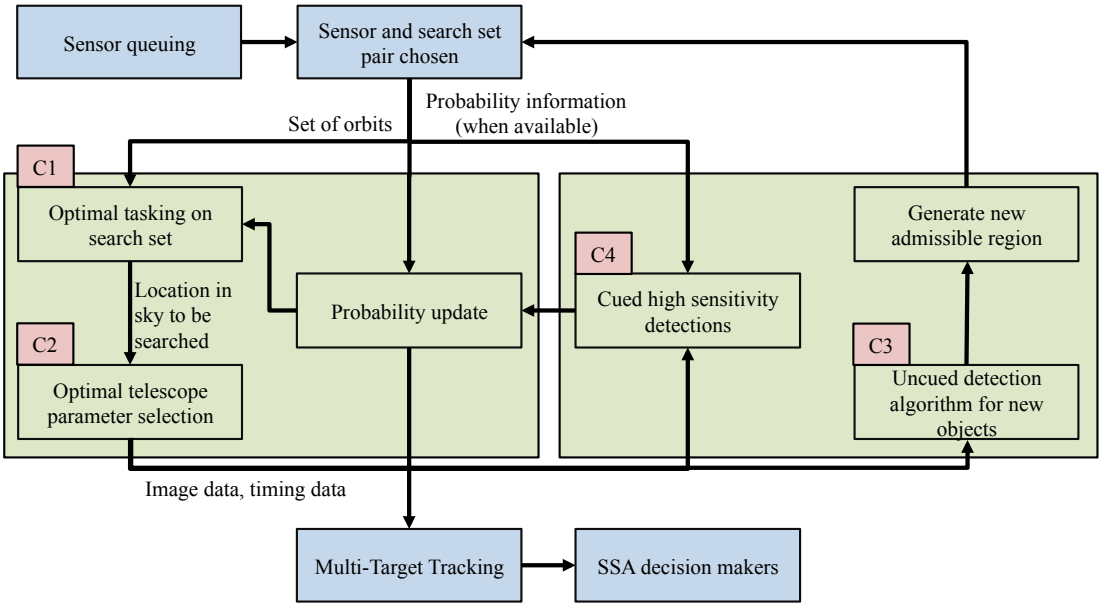

Figure 1.1: Dissertation overview. Green boxes are work included in this dissertation while blue boxes represent exterior methods which would need to work in conjunction with this work. Red boxes label each of the four contributions.

a time-optimal version. Once an observation is made, there exist a variety of techniques for correlating observations with a point in the search set [30, 31]. This problem is also relevant to allowing non-collocated EOS to observe a single object simultaneously for fast orbit determination [19]. This work can also be combined with modern models for prior information in space objects such as sequential Bayesian filters [32, 33].

In typical SSA data a series of stationary or moving point sources exist in an image over an integration time. The image shows each point or line blurred by a point spread function, a combination of the airy disc diffraction pattern and other random processes such as atmospheric blur, camera jitter, and photo-electronic bleeding [34–36]. These processes can produce varied spatial distributions depending on the specific situation, but are often approximated as Gaussian [37, 38]. Astrometry is the process of locating where a point source exists in the field of view of a sensor. This is a relevant problem in SSA, astronomy, microscopy, and other fields, and is accomplished through various techniques including thresholding, calculating a centroid, or maximum likelihood estimation. In particular, the maximum likelihood detection method is seen as a statistically rigorous method for making

sensitive detection [39–41] and has grown in popularity recently for SSA applications [42]. This method incorporates pixel-to-pixel data in the images and works along the same principles as a matched filter [43]. In astronomy, from which many SSA methods are derived, object motion is virtually nonexistent and limits on exposure time are less critical. For SSA, good estimates on both space object position and velocity in the image plane are desired. In this work, astrometry then refers to measuring the position of an object in an image and the velocity of an object through the image. Image data can also be used for categorization and intent characterization of space objects. This work looks at point sources, where the only information is how much light is received from an object. Light curve inversion is a newer method in SSA which looks at identifying the shape and properties on a space object-based on the light curve [44, 45]. These techniques are based on methods developed for identifying the approximate shape of asteroids [46, 47]. This can look at trying to estimate the shape model of a completely unknown object. Previous work also suggests that with a shape model, a light curve can indicate the type of mission the object may have [48]. For all of these missions, the data input into the estimation technique is time-resolved flux information. These methods should be improved by both denser time resolution of the flux and more accurate estimation of the flux values. The optimization in this work focuses on these two factors.

There are a multitude of methods for detecting and tracking low signal-to-noise ratio (SNR) space objects. Examples include multi-object filtering [49], multiple hypothesis testing [50], multi-frame matched filters [51], track-before-detect, [52], shift and add methods [53], and more. Due to the non-linear equations of motion and perturbations, general methods derived from Bayesian filtering [54] have become popular for space object tracking and orbit determination [32, 33]. Space is represented as a 6 dimension or higher state space, making such particle-based filters computationally difficult to implement [55]. This work instead looks at kinematic frame-to-frame tracking in images which allows a reduced dimensionality in the state space. The visual tracking method in this work is de-

rived from finite set statistics (FISST) based filters [49]. In particular, recent pushes in FISST filter theory in the field of computer vision have looked at multi-target tracking in images [42, 56, 57]. The likelihood operates directly on pixel data and requires no explicit detection algorithm. Very recently, the SSA community has begun to look at FISST filters [58] and visual tracking [42]. FISST methods used in this dissertation are intended to build on existing work in the field and to better complement the needs in SSA.

With prior information on an object, methods like matched filtering can enable lower SNR detections [59]. The matched filter (MF) originates in image and signal processing as a way to maximize SNR [60]. It has been applied to dim asteroid detection and multi-frame moving target detection [51, 61]. A MF correlates a known or hypothesized signal structure with a measured image to provide a single value result. Work on a MF incorporating a binary hypothesis test is the basis for the likelihood function proposed in this dissertation [62]. The velocity filter hypothesizes signal structure with linearized dynamics over an EOS exposure time [59]. If the underlying signal is unknown, a brute force bank of velocity filters can be used [63]. The first problem with the past use of the MF in SSA applications is the dynamics agnostic approach. Linearization is not correct when sensors contain significant rotation or distortion. More importantly, prior knowledge is treated with a binary availability. Many SOs which either do not exist in the space object catalog or have not received updating measurements for long time periods may not be accurately described by their catalog entry or other prior information. Instead, modern models for prior knowledge via the admissible region and Gaussian mixture models should be incorporated into the MF framework [33, 64, 65]. These prior PDFs can be non zero over large ranges of orbital positions, implying any MF-based on them should be a bank of plausible matched filters. This statistically rigorous framework can and should be the basis for hypothesis signals in an MF or MF bank. This work also draws from a different branch of mathematics to partition admissible region-like priors. A good discussion of metric-based search and partition methods can be found in the text *Similarity Search: The Metric Space Approach* by

Zezula et al. [66]. The broad approach falls into two categories, ball partition (BP) methods and generalized hyperplane tree (GHT) methods [67, 68]. The literature is rich with implementations of these methods [69, 70]. Work has also been done to incorporate hypothesis testing into SSA MFs, which is included [62]. The work in this dissertation shows how to use partial orbit knowledge to generate hypothesized signals expected in images.

The remaining discussion here looks at ongoing and near-term future work that supports this dissertation.

The admissible region is talked about throughout this introduction and the entire dissertation. In particular it is a major driving inspiration and justification of developing and using a search set or set of orbits. The admissible region when used for estimation techniques is not a PDF but an uniformative prior [71]. This implies that doing a direct likelihood update on an admissible region requires carefully defining what is being updated. Concurrent work to this dissertation has shown that a promising way to update an admissible region is through Dempster Schafer belief and plausibility [72]. As of writing this, these techniques are in an early stage of development and cannot be wrapped into this dissertation. However there are some interesting avenues for future work, such as how belief and plausibility interact with a detectionless likelihood, and how images which only partially observe the search set update belief and plausibility. Alternatively, many of these problems can be sidestepped by using admissible region intersection techniques [30].

The tasking problem in this paper is treated as a tactical tasking problem, that is, given a sensor and target, how should the sensor search for the target. In reality, this problem is informed by, and directly affects, the strategic tasking problem wherein a decision maker chooses from many options which sensor observes which target. In this situation, concurrent work is also on-going [25]. The work in this dissertation provides analytical tools which would support how the strategic tasking problem can be solved. The strategic and tactical tasking form one interconnected problem which, due to the developing nature of both, cannot be sufficiently detailed in this dissertation. This provides an excellent avenue

for future work.

1.4 Specific Contributions

1.4.1 Optical Sensor Tasking on an Advanced Astrodynamic Prior (C1)

This contribution involves both advanced astrodynamic models and sensor tasking. *Sensors must be tasked to utilize advanced astrodynamics models with orbit determination in mind.* Some strategies are based on maximizing information gain in a single image; many of these strategies ignore the case where an entire prior cannot be observed in a single image. This work builds off a particle-based search strategy proposed by Hobson [27] by creating a tasking framework for set-based priors. In particular this can represent tasking an EOS to reacquire an object observed by a previous EOS at an arbitrary but different location. This approach is a generalization of the follow-up problem, allowing any desired geometry to be utilized, but also incorporates more general search problems in the same framework. A prior, represented as discrete point-wise approximation, can be analyzed in terms of how it evolves in the measurement frame of the sensor over time. The search set is analyzed by modeling set area growth with a Taylor series and analytic derivatives. The work here goes on to achieve the tasking in a time optimal manor, by targeting areas of maximum growth. This contribution then looks ahead at how updates occur, by partitioning the prior into an observed and unobserved part. This partial update closes the loop, allowing the tasking strategy to feed into the orbit update.

1.4.2 Optical Sensor Parameter Optimization (C2)

Sensor tasking and detection methods are combined by optimizing the observation parameters to maximize astrometric and photometric information. When observing a certain region of orbits, the slew rate, exposure time, number of frames, and possibly other sensor parameters must be chosen. This chapter develops a framework of how to choose tasking parameters for taking electro-optical SSA data. Specific reward functions are developed

for position, velocity, and brightness of a point sources. This includes a novel derivation of point source velocity information and time varying photometry aliasing bias. Each of these reward functions are extended to a streaking object as well. Timing uncertainty is incorporated into the reward functions. The gradients with respect to exposure times are developed for all reward functions to ease analysis and optimization. Several constraints on how data can be taken are included to ensure optimization solutions are realistic. Finally, simulation work provides a blueprint on how to solve such optimization problems while provided insight into what general properties an optimal campaign has. The reward functions are based on a maximum likelihood estimator which is closely tied to the likelihood used in subsequent chapters for detection.

1.4.3 High Sensitivity Detections with Finite Set Statistics (C3)

Advanced statistical techniques, in particular finite set statistics, are leveraged to enable dim object detection in images in an uncued manor. The computer vision community has developed a range of more sensitive tracking methods [73], which are referred to as visual tracking or detectionless tracking, which have recently been adapted by the SSA community [42]. The third contribution in this thesis takes detectionless visual tracking methods and adapts them to SSA images. This contribution develops a likelihood function which allows such detectionless methods to be used for SSA missions. These methods, in implementation, are non-linear particle-based methods which can be done blindly or initiated with a prior. Analysis is performed to develop a birth model for the visual tracking technique which directly support SSA requirements. The technique is extensively validated on real image data of varying quality.

1.4.4 High Sensitivity Detections on an Advanced Astrodynamic Prior (C4)

Performance of detection algorithms can be increased when an advanced astrodynamic prior is used [74], which the final contribution analyzes. *Prior knowledge in the form*

of a set of orbits can be used to interrogate an image for a particular object, allowing highly sensitive tracking and direct orbit determination. This last chapter uses the prior set to directly define a likelihood function in an image. Set partitioning methods are used to minimize the number of templates used to search an image. The SNR limits of this detection method are explored in simulation, and are shown to be much lower than uncued techniques in the previous chapter. The likelihood is also used as a direct update on the prior set, allowing direct orbit determination.

1.4.5 List of Major Contributions

The completed work provides insight into how to vertically integrate an SSA system. The first contribution (C1) is the design of a sensor tasking strategy based on advanced astrodynamics techniques, providing the connection seen in Figure 1.2. This work provides a rigorous and general approach, based on optimization techniques, to task a sensor to search for a particular object or region. The second contribution (C2) presents a method that integrates detection methods with how sensors are tasked, as shown in Figure 1.2. Specifically, this contribution develops a multi-objective optimal approach to choosing sensor tasking parameters. The third contribution (C3) describes how to adapt modern visual tracking methods into an SSA detection algorithm. This contribution provides a building block for follow-on research. The final contribution (C4) describes how modern detection methods can be integrated with advanced astrodynamics methods to obtain increased performance for both the detection method and the orbit determination process. This contribution is the link between high sensitivity detection used in contribution 1 and orbit parameter estimation, seen in Figure 1.2.

1.4.6 List of Publications

The peer reviewed journal articles associated with this dissertation are as follows:

J1 Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Space Object Detection in Im-

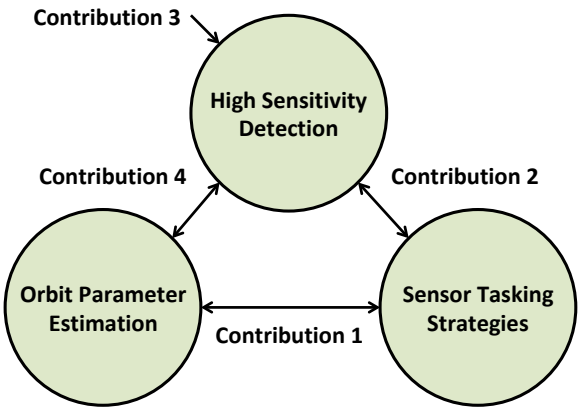

Figure 1.2: Primary contributions solve the broad and interconnected problem of detection, estimation and tasking

ages Using Matched Filter Bank and Bayesian Update", *AIAA Journal of Guidance, Control, and Dynamics*, vol. 40, no. 3, 2017, pp. 497-509

J2 Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Visual Tracking Methods for Improved Sequential Image-based Object Detection", *AIAA Journal of Guidance, Control, and Dynamics*, vol. 41, no. 1, 2018, pp. 74-87

J3 Murphy, T. S., Holzinger, M. J., Luu, K. K., and Sabol, C., "Generalized Minimum-Time Follow-up Approaches Applied to Electro-Optical Sensor Tasking", *IEEE Transactions on Aerospace and Electronic Systems*, 2018, [Submitted]

J4 Murphy, T. S., and Holzinger, M. J., "Optimal Selection of Telescope Parameters for Space Situational Awareness Astrometry and Photometry", *IEEE Transactions on Aerospace and Electronic Systems*, 2018, [to be Submitted, April 2018]

The conference papers and presentations associated with this dissertation are as follows:

C1 Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Particle and Matched Filtering Using Admissible Regions", *AAS/AIAA Spaceflight Mechanics Meeting*, Williamsburg, VA, Feb, 2015

C2 Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Orbit Determination for Partially Understood Object via Matched Filter Bank", *AAS/AIAA Astrodynamics Specialists Conference*, Vail, CO, Aug, 2015

C3 Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Direct Image-to-Likelihood for Track-Before-Detect Multi-Bernoulli Filter", *AAS/AIAA Spaceflight Mechanics Meeting*, Napa, CA, Feb, 2016

C4 Murphy, T. S., and Holzinger, M. J., "Uncued Low SNR Detection with Likelihood from Image Multi-Bernoulli Filter", *Advanced Maui Optical and Space Surveillance Technical Conference*, Wailea, HI, Sep, 2016

C5 Murphy, T. S., Holzinger, M. J., Luu, K. K., and Sabol, C., "Optical Sensor Follow-up Tasking on High Priority Uncorrelated Track", *AAS/AIAA Spaceflight Mechanics Meeting*, San Antonio, TX, Feb, 2016

C6 Murphy, T. S., and Holzinger, M. J., "Generalized Minimum-Time Follow-up Approaches Applied to Tasking Electro-Optical Sensor Tasking", *Advanced Maui Optical and Space Surveillance Technical Conference*, Wailea, HI, Sep, 2017

C7 Murphy, T. S., and Holzinger, M. J., "Optimal Selection of Telescope Parameters for Space Situational Awareness Astrometry and Photometry", *Advanced Maui Optical and Space Surveillance Technical Conference*, Wailea, HI, Sep, 2018 [to be submitted]

The contributions and associated background literature can be seen in Table 1.1.

1.5 Unifying Example

The work in this dissertation in many parts is exploratory. Because of this, there is no unifying simulation work which could realistically wrap all of the sub-problems together; this would require large amounts of extra research work. Instead, at the end of each chapter of

Table 1.1: Visualization of background material as it relates to each major contribution.

Citations	Generalized Minimum-Time Follow-up Approaches Applied to Electro-Optical Sensor Tasking	Optimal Selection of Telescope Parameters for Space Situational Awareness Astrometry and Photometry	Visual Tracking Methods for improved Sequential Image-based Object Detection	Space Object Detection in Images Using Matched Filter Bank and Bayesian Update
[23–26] [27–29]	Modern sensor tasking strategies Particle-based and set-based search			
[34–38] [39–41] [44–48].		CCD astronomy references Maximum likelihood estimation of images Photometric light curves and light curve inversion		
[49–53]			Methods for detecting and tracking low SNR objects	
[49, 58, 75, 76]			Finite Set Statistics and the multi-Bernoulli filter	
[42, 56, 57]			Detectionless likelihood function tracking	
[51, 59, 62]				Matched filters as applied to detection of space objects
[33, 64, 65]				Admissible region theory as they apply to SSA
[66–70]				Metric-based partitioning of sets
Contributions				
Chapter 3 (J3)	Create a methodology to optimally select EOS tasking trajectory through measurement space			
Chapter 4 (J4)		Define reward functions for information content in SSA data and optimally select telescope parameters		
Chapter 5 (J2)			Synthesize new image likelihood function for SSA and test in Mulit-Bernoulli framework	
Chapter 6 (J1)				Construct a matched filter bank to directly interrogate image based on prior and perform Bayesian update

this document, the chapter's work will be explained within the context of a single example. The building blocks provided in this dissertation form together, through this example, to solve a real application problem in an integrated loop.

The problem starts when an exquisite sensor with a large aperture detects an unknown Geosyncronous level space object. At the current time, follow-up observations are of limited use; for such an object, observability of the full orbit is not attainable until well past 10 minutes after the first observation [71]. Furthermore, this dimly lit space object is detectable through the single exquisite sensor but may be difficult to detect through a less exquisite follow-up sensors. Finally, there are generally too few sensors to follow-up on every object. Because of these reasons, most of these objects are not tracked immediately after first detection and cannot be found again when the orbit is fully observable.

The contributions in this dissertation will, in order, do the following. First, an admissible region will be formed on the detected unknown object. Contribution 1 shows how to take this admissible region, and search through it at a later time. Through this, a follow-up sensor can effectively delay reacquisition until observability is met. Contribution 2 shows how the specific data acquisition parameters can be chosen for each region searched in contribution 1. Contribution 2 feeds a series of images into contribution 3, which in turn searches through the images for any signals from a star, object, or other source. If contribution 3 fails to find the space object associated with the prior, contribution 4 searches through the images using the prior provided by the admissible region, enabling much higher sensitivity. Furthermore, contribution 4 uses the detections on the images to perform orbit determination on the original admissible region.

CHAPTER 2
MATHEMATICAL BACKGROUND

2.1 Astrodynamics of Space Situational Awareness

First this work requires defining the general dynamics framework. Define $\mathcal{X}$ as the set of all states which could be position and velocity or some alternative representation like orbit elements. Given a dynamic system,

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t); t) \tag{2.1}$$

an object can be propagated forward in time via a flow function

$$\mathbf{x}(t_1) = \boldsymbol{\phi}(t_1; \mathbf{x}(t_0), t_0) \tag{2.2}$$

which simply propagates a state via the dynamics in (2.1). This work considers the propagation of not just a single state, but a set, referred to as the search set. Consider a set of dynamic states, $\mathcal{S}(t_0) \subseteq \mathcal{X}$, which is defined with a series of constraints $\{g_a(\mathbf{x})\}_{a=1}^{b}$

$$\mathcal{S} = \{\mathbf{x} \in \mathcal{X} : g_a(\mathbf{x}) \leq 0 \ \forall \ a\} \tag{2.3}$$

Because each element in this set is dynamic, the set can be propagated as

$$\mathcal{S}(t) = \boldsymbol{\phi}(t; \mathcal{S}(t_0), t_0) \tag{2.4}$$

Note that this region is time varying, according to the dynamics; the time notation is dropped unless explicitly necessary.

In the case of orbit estimation it is useful to think in terms of object and observer

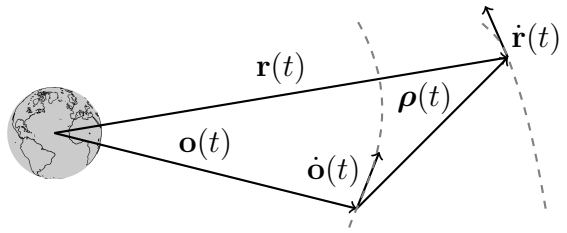

Figure 2.1: Geometry of SSA problem. An observer and space object are represented with time varying vectors. The line of sight vector is $\boldsymbol{\rho}(t)$.

position and velocity. For an object, position and velocity are $\mathbf{r}(t)$ and $\dot{\mathbf{r}}(t)$, seen by an optical observer at position and velocity $\mathbf{o}(t)$ and $\dot{\mathbf{o}}(t)$ along line-of-sight vector, $\boldsymbol{\rho}(t)$, defined in an inertial frame. At an instantaneous time, an EOS measures a unit vector $\hat{\boldsymbol{\rho}}(t)$. It should be noted that while this work assumes $\mathbf{x}(t) \in \mathbb{R}^6$ throughout as position and velocity, this can be generalized to any formulation of an orbit as well. The observer itself could be either ground-based or space-based. The geometry is illustrated in Figure 2.1.

2.2 Background on Modeling and Statistics of Images

This section defines the statistical models for an image that are used throughout the dissertation. An image is made up of a series of measurements broken into some deterministic signal content and some zero mean noise of known distribution (assuming background subtraction has been performed [77]),

$$z_i = s_i + w_i \tag{2.5}$$

$$w_i \sim p(w),\ \mathbb{E}[w_i] = 0 \tag{2.6}$$

where $p(w)$ is some PDF or probability mass function (PMF) of known form, and i is an indexing denoting a particular pixel.

The distribution of w_i in Equation (2.5) can be arbitrary, but often is assumed to be of a specific form. A common way to model the statistics of a pixel is with a series of Poisson distributions [78]. The pixel is composed of some signal, which is Poisson distributed and

referred to as shot noise, and some dark noise and read noise both of which are Poisson distributed, and can therefore be combined into one Poisson distributed random variable, w_i. These leads to a pixel, $z_i = s_i + w_i$, made up of signal, $s_i \sim \mathcal{P}(\lambda_{s,i})$, and other noise sources, $w_i \sim \mathcal{P}(\lambda_{w,i})$. However, it is more convenient to redefine a background subtracted z_i. In this dissertation, background subtraction means subtraction of known statistical outliers, such as hot pixels, and the background mean [77].

$$s_i = \lambda_{s,i} \tag{2.7}$$

$$w_i \sim \{\mathcal{P}(\lambda_{w,i} + \lambda_{s,i}) - (\lambda_{w,i} + \lambda_{s,i})\} \tag{2.8}$$

This new definition of z_i allows the problem to be posed as the sum of a deterministic signal and zero mean random noise. A Poisson distribution approaches a Gaussian distribution as λ increases, which is explored further in Lemma 5.3.2. Values of λ larger than 100 are typically sufficient. If such a simplification is applicable, the following more simplified definition of w_i can be used.

$$w_i \sim \mathcal{N}(0, \lambda_{w,i} + \lambda_{s,i}) \tag{2.9}$$

Furthermore, for bright objects the contribution of shot noise on total system randomness is often ignored and read noise is mostly constant within a small region of pixels. Then, a more simplified definition of w_i is

$$w_i \sim \mathcal{N}(0, \lambda_w) \tag{2.10}$$

Note that in this formulation, s_i is equivalent to the mean of the measurement z_i, which is Gaussian distributed. It is important to understand that measurements taken over the kth

integration time are dependent on the length of the integration

$$\lambda_w = GQ\delta^2 F_o t_{I,k} = BF_o t_{I,k} \tag{2.11}$$

where G is the sensor's gain, Q is the quantum efficiency, δ^2 is the area of each pixel, and F_o is the background flux rate. This implies that the design variable $\mathbf{t}_I$ is able to affect the information content directly through changing the background brightness. In the final line, the constants are wrapped into the parameter $B = \delta^2 GQ$ to ease notation.

The pixel-wise SNR for a particular pixel, z_i, is defined in Equation (2.12).

$$\begin{aligned} \text{SNR}(z_i) &= \frac{\mathbb{E}[z_i]}{\sqrt{\mathbb{E}[(z_i - \mathbb{E}[z_i])^2]}} \\ &= \frac{s_i}{\sigma_{w,i}} \end{aligned} \tag{2.12}$$

In astronomy, the total pixel-to-pixel structure of an object may exist in multiple pixels, and therefore the per pixel SNR does not fully capture the information available. Total object SNR is calculated as

$$\begin{aligned} \text{SNR}\left(\sum z_i\right) &= \frac{\mathbb{E}[\sum z_i]}{\sqrt{\mathbb{E}[([\sum z_i]^2 - \mathbb{E}[\sum z_i])^2]}} \\ &= \frac{\sum s_i}{\sqrt{\sum \sigma_{w,i}^2}} \qquad (2.13) \\ &\approx \sqrt{N_z}\frac{\bar{s}_i}{\sigma_w} \qquad (2.14) \end{aligned}$$

where N_z is the number of pixels which contain some signal. Equation (2.13) assumes only that pixels are uncorrelated. It is convenient to also define an approximate total object SNR, in Equation (2.14), which assumes all pixels have equal variance. This assumption allows $\bar{s}_i$ to be defined as the average signal value over the pixels being considered, and σ_w as the approximate standard deviation in each pixel.

The two types of SNR warrant further discussion. First, neither values are any less valid for non-Gaussian distributions, and in particular work equally well for Poisson distributions. Second, both methods say something different but useful about the signal. Pixel SNR shows how easily an object can be picked out with the naked eye. Photometric SNR combines all the evidence for an object in the numerator with all the noise in the denominator and therefore indicates how well a statistical analysis should preform at detecting an object. Finally, choosing m_z pixels which contain signal from an object is a nebulous concept. Typically, there are pixels which may or may not contain signal, and especially in low SNR cases, ambiguity exists. This is an active are a research with no clear solutions [78].

2.3 Admissible Region

A typical optical measurement of a space object contains good knowledge of angle and angle rates known as observable states, $\mathbf{x}_d \in \mathbb{R}^4$. The variable $\mathbf{x}_d$ is the vector of determined states, also known as the attributable vector. In the case of an optical observer,

$$\mathbf{x}_d = [\alpha, \delta, \dot{\alpha}, \dot{\delta}]^T \tag{2.15}$$

An orbit requires six disparate data types to be fully constrained. Historically, the final two data types are angle accelerations or range and range rate (from radar data). This AR formulation uses range and range rate as the final two states, known as undetermined states $\mathbf{x}_u \in \mathbb{R}^2$. It is assumed that no information can be reliably used to determine these states from measurements. The following notation was developed by Worthy et al [20]. Using the fact that the measurement cannot be dependent on $\mathbf{x}_u$, the following measurement function can be written.

$$\mathbf{y}(t) = \mathbf{h}(\mathbf{x}_d; \mathbf{k}, t) \tag{2.16}$$

Furthermore, this implies a bijective relationship between $\mathbf{y}(t)$ and $\mathbf{x}_d$

$$\mathbf{x}_d = \mathbf{h}^{-1}(\mathbf{y}; \mathbf{k}, t) \tag{2.17}$$

An AR is then created by enforcing a series of constraints of the form

$$g_i(\mathbf{x}_d, \mathbf{x}_u; \mathbf{k}, t) \leq 0 \tag{2.18}$$

$$g_i(\mathbf{h}^{-1}(\mathbf{y}; \mathbf{k}, t), \mathbf{x}_u; \mathbf{k}, t) \leq 0 \tag{2.19}$$

It should be noted that these constraints can be thought of as hypotheses. The AR is then the space where all hypotheses are true. Next, we define an admissible region, $\mathcal{AR}_i \in \mathbb{R}^2$, predicated on g_i

$$\mathcal{AR}_i := \left\{ \mathbf{x}_u | g_i(\mathbf{h}^{-1}(\mathbf{y}; \mathbf{k}, t), \mathbf{x}_u; k, t) \leq 0 \right\} \tag{2.20}$$

In practice, an admissible region predicated on n constraints is used.

$$\mathcal{AR} = \bigcap_{i=1}^{n} \mathcal{AR}_i \tag{2.21}$$

Next, the two most common constraints are defined. For optical observer AR, the primary constraint typically considered is that of Earth-orbiting SO, defined as $g_{\mathcal{E}}$. There exists a derivation to show the following equations represent energy as a function of $[\rho, \dot{\rho}]$ and known parameters. The following results were originally derived for Earth objects by Tommei et al [79]. $\mathcal{E}$ is the orbital energy, which must be negative for an Earth-orbiting

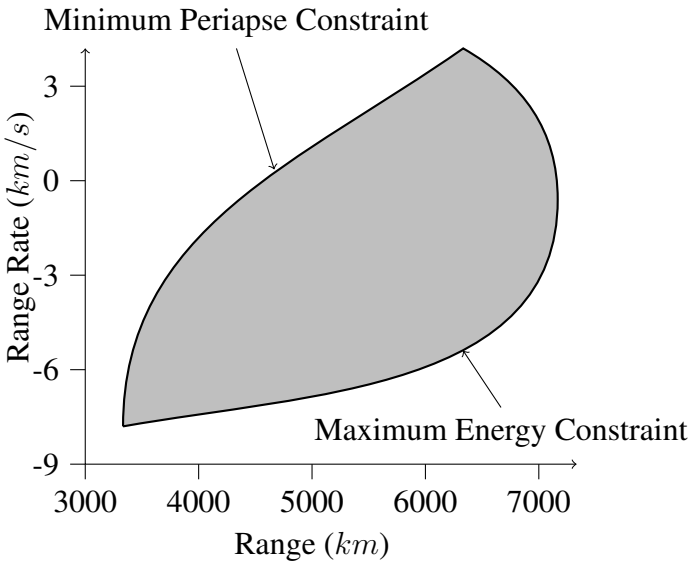

Figure 2.2: Sample admissible region

SO.

$$2\mathcal{E} = g_{\mathcal{E}}(\mathbf{h}^{-1}(\mathbf{y};\mathbf{k},t),\mathbf{x}_u;\mathbf{k},t) \leq 0 \tag{2.22}$$

$$2\mathcal{E} = \dot{\rho}^2 + w_1\dot{\rho} + T(\rho) - \frac{2\mu}{\sqrt{S(\rho)}} \leq 0 \tag{2.23}$$

$$T(\rho) = w_2\rho^2 + w_3\rho w_4 \tag{2.24}$$

$$S(\rho) = \rho^2 + w_5\rho + w_0 \tag{2.25}$$

where the constants w_{0-5} are functions of the parameter vector, $\mathbf{k}$, defined as

$$w_0 = ||\mathbf{o}||^2 \tag{2.26}$$

$$w_1 = 2\dot{\mathbf{o}} \cdot \hat{\boldsymbol{\rho}} \tag{2.27}$$

$$w_2 = \dot{\alpha}^2 \cos^2 \delta + \ddot{\delta} \tag{2.28}$$

$$w_3 = 2(\dot{\alpha}\dot{\mathbf{o}} \cdot \hat{\boldsymbol{\rho}}_{\boldsymbol{\alpha}} + \dot{\delta}\dot{\mathbf{o}} \cdot \hat{\boldsymbol{\rho}}_{\boldsymbol{\delta}}) \tag{2.29}$$

$$w_4 = ||\dot{\mathbf{o}}||^2 \tag{2.30}$$

$$w_5 = 2\mathbf{o} \cdot \hat{\boldsymbol{\rho}} \tag{2.31}$$

where $\mathbf{o}$ and $\dot{\mathbf{o}}$ are the position and velocity of the observer, $\hat{\boldsymbol{\rho}}$ is the unit vector from the angles, and $\hat{\boldsymbol{\rho}}_\alpha$ and $\hat{\boldsymbol{\rho}}_\delta$ are given by

$$\hat{\boldsymbol{\rho}} = [\cos\alpha\cos\delta \;\; \sin\alpha\cos\delta \;\; \sin\delta]^T \tag{2.32}$$

$$\hat{\boldsymbol{\rho}}_{\boldsymbol{\alpha}} = [-\sin\alpha\cos\delta \;\; \cos\alpha\cos\delta \;\; 0]^T \tag{2.33}$$

$$\hat{\boldsymbol{\rho}}_{\boldsymbol{\delta}} = [-\cos\alpha\sin\delta \;\; -\sin\alpha\sin\delta \;\; \cos\delta]^T \tag{2.34}$$

A second commonly used constraint is g_r, a constraint on the radius of perigee. There exists an analytic derivation for the following [65]. It should be noted that $\mathbf{D}$, $\mathbf{E}$, $\mathbf{F}$, and $\mathbf{G}$ are vector quantities.

$$g_r(\rho, \dot{\rho}, \mathbf{x}_d; \mathbf{k}) = (r_{min}^2 - \|\mathbf{D}\|^2)\dot{\rho} - P(\rho)\dot{\rho} - U(\rho) + r_{min}^2 T(\rho) - \frac{2r_{min}\mu}{\sqrt{S(\rho)}} \leq 0 \tag{2.35}$$

$$P(\rho) = 2\mathbf{D} \cdot \mathbf{E}\rho^2 + 2\mathbf{D} \cdot \mathbf{F}\rho + 2\mathbf{D} \cdot \mathbf{G} - r_{min}^2 w_1 \tag{2.36}$$

$$U(\rho) = \|\mathbf{E}\|^2\rho^4 + 2\mathbf{E} \cdot \mathbf{F}\rho^3 + (2\mathbf{E} \cdot \mathbf{G} + \|\mathbf{F}\|^2)\rho^2 + 2\mathbf{F} \cdot \mathbf{G}\rho + \|\mathbf{G}\|^2 - 2r_{min}\mu \tag{2.37}$$

where the $\mathbf{D}$, $\mathbf{E}$, $\mathbf{F}$, and $\mathbf{G}$ are defined as

$$\mathbf{D} = \mathbf{q} \times \hat{\boldsymbol{\rho}} \tag{2.38}$$

$$\mathbf{E} = \hat{\boldsymbol{\rho}} \times (\dot{\alpha}\hat{\boldsymbol{\rho}}_{\alpha} + \dot{\delta}\hat{\boldsymbol{\rho}}_{\delta}) \tag{2.39}$$

$$\mathbf{F} = \mathbf{q} \times (\dot{\alpha}\hat{\boldsymbol{\rho}}_{\alpha} + \dot{\delta}\hat{\boldsymbol{\rho}}_{\delta}) + \hat{\boldsymbol{\rho}} \times \dot{\mathbf{q}} \tag{2.40}$$

$$\mathbf{G} = \mathbf{q} \times \dot{\mathbf{q}} \tag{2.41}$$

More ways exist to further constrain an admissible region. Beyond perigee and energy, any further restriction typically requires an additional assumption about the object. For example, if the object being observed is known to be near a GEO, constraints could be placed on semi major axis or eccentricity. This can enable prior orbital knowledge to be used to enhance convergence.

CHAPTER 3
GENERALIZED MINIMUM-TIME FOLLOW-UP APPROACHES APPLIED TO ELECTRO-OPTICAL SENSOR TASKING

3.1 Introduction

This chapter proposes a methodology for tasking sensors to search an area of state space for a particular object, group of objects, or class of objects. This chapter creates a general unified mathematical framework for analyzing reacquisition, search, scheduling, and custody operations. In particular, this chapter looks at searching for unknown space object(s) with prior knowledge in the form of a set. For example, an uncorrelated track (UCT), obtained from an EOS, can be used to define a set of orbits consistent with the observations, also known as an admissible region [20,64]. The sets used in this framework can also be defined through reachability [21], hypothesized regions of state space, or other set-based methods. This set of orbits can be propagated and projected into the field of regard any sensor, so follow-up can be tasked from any location and time. In general, this region can be much larger than the field of view of a sensor and requires multiple observations [29]. This chapter analyzes the area of a search region over time to inform a time-optimal search method. Simulations look at analyzing search regions relative to a particular sensor, and testing a tasking algorithm to search through the region. The tasking algorithm is also validated on a reacquisition problem with a telescope system at Georgia Tech.

The problem facing this technique is how to choose an optimal trajectory when performing a search, and what type of optimality is desirable. Existing methods are based on a greedy maximum probability observation technique, which attempts observing the densest region of probability [27]. It becomes clear through the work in this chapter that time-optimality is the highest priority. Many SSA sensor systems have a problem of too

many objects and not enough sensors, making time on high sensitivity sensors a priority. Furthermore, when searching a set of orbits, there is necessarily associated probability information. If probability density information is available, a greedy maximum probability observation may take significantly more time to fully search the region. Because an EOS can be prone to false detections, it is typically more wise to search the entire search set than to trust the first detection that is made. Therefore time-optimal tasking is pursued in this chapter.

This optimization problem, then, is the covering salesman problem (CSP), a variation of the traveling salesman problem [80]. The CSP looks for the time or distance optimal path to take between a series of points such that every point, or a point within a limiting distance, is visited. The tasking problem outlined above has a further complication over the classic CSP, in that the points to be visited have dynamics, and typically spread out as time passes. This problem is in general NP-hard, and obtaining complete optimal solutions is difficult. This chapter therefore both attempts the time-optimal optimization, but also develops tools that can efficiently analyze and compare different search strategies.

3.1.1 Methodology

The first requirement is an analytic formulation of the problem. A region which must be searched is modeled as a set of orbits. An EOS tasking scheme is modeled as a series of observation sets centered at fixed angular coordinates, at certain time steps. These discrete task locations can consist of any number of observations taken at any locations in any order. Quantities like exposure time and number of observations taken at each location, as well as characteristics of the sensor, are explored in the next chapter, so this particular work fixes these quantities for simplicity. These ideas are combined to propose an observed set and an unobserved set, two partitions of the original search set which describe the progress of a search and remaining set to be searched, respectively.

An analysis on the area of a search set, or subset, is presented by looking at set area in

the measurement space. This area is presented as a line integral over the boundary of the region, in the measurement space. Time derivatives of an arbitrary order are then derived as line integrals. The calculation of these sets on admissible regions is then shown in more detail. This section provides a method to analyze search sets relative to a particular sensor, giving estimates of search campaign length and search feasibility. Section 3.2 details this work.

An optimization framework is then proposed as a way to choose a trajectory. This chapter looks to minimize the amount of time spent searching for the object, while attempting to scan the entire search set. In previous work, a variation on simulated annealing [81] was tried with some success. Due to the high dimensional and non-convex search space, this optimization can be difficult to successfully implement. Instead, this chapter uses an optimization greedy on search area growth; it looks to maximize the size of the observed area at some future time. A region of high divergence grows quickly, increasing the total number of observations needed to observe it. The implication is that a time-optimal algorithm prioritizes high divergence. This cost function is suited for a finite time horizon control approach. Section 3.3 details this work.

In order to test the tasking algorithm and illustrate its performance and behavior to the reader, two simulated test cases are presented. The first test case looks at a same location handoff. An admissible region, formed from a short arc observation, is propagated an hour into the future and reacquisition is attempted. This test case shows off many of the strengths of the work in this chapter. Second, a search set is defined to contain objects in geostationary transfer orbits which intersect a particular GEO slot over a given time interval. This search set and subsequent tasking simulation showcase the variety of scenarios to which these techniques can be applied. Section 3.4 details this work.

3.2 Theory: Search Regions

This section introduces the concept of a search set, a set of orbits that a sensor or group of sensors must search through. Then a series of analytic tools are presented for analyzing a search set and the projection of the search set into a measurement subspace. The intention is to provide insight into the nature of these sets and the implications on their searchability.

3.2.1 General Problem Framework

First this chapter requires defining the general dynamics framework. This discussion is an expansion of the preliminary discussion in Section 2.1. To review, define $\mathcal{X}$ as the set of all states which could be position and velocity or some alternative representation like orbit elements. Given a dynamic system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) \tag{3.1}$$

which is a reproduction of Equation (2.1), an object can be propagated forward in time via a flow function in Equation (2.2) which simply propagates a state via the dynamics. This chapter in particular considers the propagation of not just a single state, but a set, referred to as the search set. Consider a set of dynamic states, $\mathcal{S}(t_0) \subseteq \mathcal{X}$,

$$\mathcal{S} = \{\mathbf{x} \in \mathcal{X} : g_a(\mathbf{x}) \leq 0 \ \forall \ a\} \tag{3.2}$$

Because each element in this set is dynamic, the set can be propagated through Equation 2.4. Note that this region is time varying, according to the dynamics; the time notation is dropped unless explicitly necessary.

The space in which this set exists, $\mathcal{X}$, is in effect a phase space representation. In accordance with Liouville's theorem, the measure, $|\cdot|$, of this set must stay constant under the given dynamics [82]. However, note that many of the sets used in this dissertation are sub-dimensional manifolds in $\mathcal{X}$ and therefore have zero measure. To review, measure is a

generalization of volume used to quantify the size of a set [83]. In this chapter volume and area are used to describe the size of sets which exist in an appropriate space.

This work requires analyzing sets of objects within the measurement space defined by a particular sensor, $\mathcal{H}_o$, where this space has dimension h. A state, $\mathbf{x}$, can be partially represented with the states determined from a measurement, $\mathbf{x}_d$. In the case of inherent unobservability, there is also an unobservable subspace, so the state can be partitioned into two sub-components, where $\mathbf{x_u}$ is part of a state unobserved by a particular observer, based on notation used in [20].

$$\mathbf{x} = [\mathbf{x_d}^T, \mathbf{x_u}^T]^T \tag{3.3}$$

This requires a measurement function, $\mathbf{h}_o$, which maps the full state space into the measurement subspace,

$$\mathbf{x_d} \in \mathcal{H}_o \tag{3.4}$$

$$\mathbf{h}_o : \mathcal{X} \rightarrow \mathcal{H}_o \tag{3.5}$$

The o subscript refers to which observer is being considered, but is dropped unless multiple observers are being explicitly used. Similarly, a search set can be projected into the measurement space,

$$\mathcal{S}_d = \mathbf{h}(\mathcal{S}), \mathcal{S}_d \subseteq \mathcal{H} \tag{3.6}$$

Note that this mapping is a projection onto a subspace, so multiple elements of $\mathcal{S}$ may map to a single component of $\mathcal{S}_d$. The nominal evolution of both $\mathcal{S}$ and $\mathcal{S}_d$ are illustrated in Figure 3.1. Note that while $|\mathcal{S}(t_0)| = |\mathcal{S}(t)|$ in accordance with Louisville's Theorem, in general $|\mathcal{S}_d(t_0)| \neq |\mathcal{S}_d(t)|$ as they exist in a subspace of the phase space.

An observation can be modeled as a subset of the search set. A sensor is tasked to

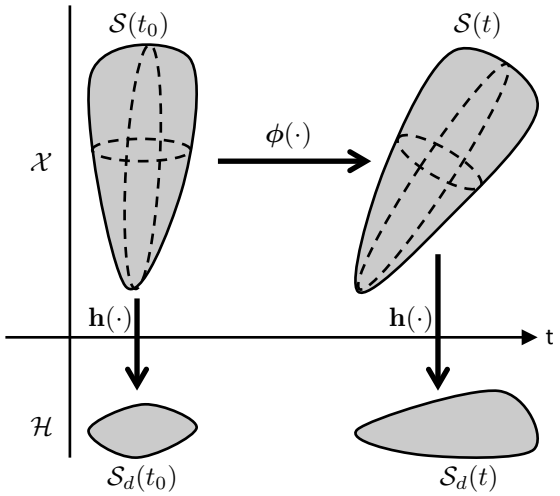

Figure 3.1: A visualization of how the search space evolves over time, in both $\mathcal{X}$ and $\mathcal{H}$.

observe a particular region of the measurement space and corresponding area of the state space at a particular time step represented by $\mathcal{F}(t_k) \subset \mathcal{X}$.

A requirement for successful search is that every part of $\mathcal{S}$ is observed. At each time step, a section of the measurement space is observed through $\mathcal{F}$. Given a search space at time step k, $\mathcal{S}(t_k)$, the subset observed at that time is $\mathcal{O}_k(t_k) = \mathcal{S}(t_k) \cap \mathcal{F}(t_k)$ while the unobserved part is $\mathcal{U}(t_k) = \mathcal{S}(t_k) \setminus \mathcal{O}_k(t_k)$. By these definitions, $\mathcal{S}(t)$ is the total search set,

$$\mathcal{O}(t) = \bigcup_{k}^{k=N} \mathcal{O}_k(t) \tag{3.7}$$

is the observed portion of the search set, and $\mathcal{U}(t)$ is the unobserved portion. Once a part of the set has been observed, it no longer needs to be observed and should therefore no longer be part of the search space. Therefore, the dynamics of $\mathcal{U}(t)$ are of the highest interest and the goal of a search is to drive $\mathcal{O} \to \mathcal{S}$ and $\mathcal{U} \to \emptyset$, which would imply a complete search. The sets $\mathcal{O}, \mathcal{U}$ can be projected into the measurement space as $\mathcal{O}_d, \mathcal{U}_d$, similar to other sets. Note that the set $\mathcal{O}$ does not exclude object existence in that set, but that the part of the space has been searched. Missed detections may lead to such a situation.

The above problem set up assumes a discrete set of observations indicative of a "stop and stare" type campaign. This problem could instead be treated as a continuous time

problem which would more closely resemble a time-delay-integrate (TDI) [84]. In such a case, the continuous time behavior of $\mathcal{U}(t)$ is of interest

$$\mathcal{U}(t) = \mathcal{S}(t) \setminus \mathcal{O}(t) \tag{3.8}$$

where $\mathcal{S}(t)$ is a known quantity and

$$\mathcal{O}(t) = \{\mathbf{x}(t) \in \mathcal{S}(t) : \mathbf{x}(\tau) \in \mathcal{O}(\tau), \tau \in [0, t]\}. \tag{3.9}$$

This analytic set up provides challenges which are not explored further in this dissertation, but provide a rich area for future work.

3.2.2 Search Region Area

The area of a search region at a given time is a useful quantity in the analysis of search regions. However, there are several different measures of area, when dealing with projections of high dimensional sets. This chapter analyzes the spatial area of a projected search region, $|\mathcal{S}|_{\mathcal{H}}$, as the spatial area of the set $\mathcal{S}$ when projected into the measurement space. Note that the spatial area is not the full measure of the set, as it measures the projected set and even then does not include the velocity states which often exist in the measurement space. The spatial area, $|\cdot|_{\mathcal{H}}$, is useful for quantifying the area a sensor can measure. In particular, the amount of a search space an optical sensor can measure is defined by the field of view of the sensor. For ease of notation, this area from now on is represented as

$$A_{\mathbf{h}}(\mathcal{S}) = |\mathcal{S}|_{\mathcal{H}} \tag{3.10}$$

This area, then, is a function of the search set which is itself a function of time, $A_{\mathbf{h}}(\mathcal{S}(t))$. However, a set does not necessarily have a clean analytic definition, and the propagation of the set over time has no analytic solution, so this area is not always well

defined. In the worst case scenario, a set of orbits can be brute force represented with particles which can be used to approximate the area.

The area of this search region has implications on the searchability of a set. In most realistic scenarios, different velocities within a search set cause the set to expand spatially. Conversely, an observer, with each observation, reduces the remaining area that must be searched. A successful search requires the observations to reduce the search area faster than the expansion is causing the area to rise.

3.2.3 Search Area and Divergence

This section looks to determine both the over-all growth rate of a search region and the local regions of expansion and contraction, as seen by a particular sensor. The main motivation for this section is that the search region over time is not analytically known. Assume that an observer, located at $\mathbf{o}(t_o)$, wants to take an observation in this search space at time t_o. First, the search set must be propagated to the appropriate time, via Equation (2.4). Next, the search set must be projected into the field of regard of a sensor, through some measurement function $\mathbf{h}$,

$$\mathbf{x}_d = \mathbf{h}(\mathbf{x}) \tag{3.11}$$

$$\mathcal{S}_d = \{\mathbf{x}_d : \mathbf{x}_d = \mathbf{h}(\mathbf{x}), \mathbf{x} \in \mathcal{S}\}. \tag{3.12}$$

The function $\mathbf{h}$ maps to $\mathbf{x}_d$, the component an observer can measure.

A clean derivation of area can be calculated using the divergence theorem. To review, the divergence theorem states that for a vector field, $\mathbf{F}$, defined over a region $\mathcal{S}$, with boundary $\partial\mathcal{S}$,

$$\int_{\mathcal{S}} \nabla \cdot \mathbf{F} \mathrm{d}V = \oint_{\partial\mathcal{S}} (\mathbf{F} \cdot \hat{\mathbf{n}}) \mathrm{d}S \tag{3.13}$$

where ∇ is the vector of derivatives with respect to $\mathbf{x}_d$. By setting $\mathbf{F}$ equal to $\mathbf{x}_d$, this

integral implies that the area of a particular region of state space can be calculated as the total integral

$$\begin{aligned} A_{\mathbf{h}}(\mathcal{S}) &= \int_{\mathcal{S}_d} \mathrm{d}V = \frac{1}{\mathrm{dim}} \int_{\mathcal{S}_d} \nabla \cdot \mathbf{x}_d \mathrm{d}V \\ &= \frac{1}{\mathrm{dim}} \oint_{\partial \mathcal{S}_d} (\mathbf{x}_d \cdot \hat{\mathbf{n}}) \mathrm{d}S \end{aligned} \tag{3.14}$$

where dim is the dimensionality of the space, used to normalize $\nabla \cdot \mathbf{x}_d$. The lower dimensional version of this equation is known as Green's theorem and is discussed in the following section.

Note that the bounds and differentials on this integral are functions of time, so to analyze the immediate expansion or contraction of a search region, the Leibniz integral rule can be used to take the first derivative as

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} A_{\mathbf{h}}(\mathcal{S}) =& \frac{1}{\mathrm{dim}} \frac{\mathrm{d}}{\mathrm{d}t} \int_{\mathcal{S}_d} \nabla \cdot \mathbf{x}_d \mathrm{d}V \\ =& \frac{1}{\mathrm{dim}} \int_{\mathcal{S}_d} \frac{\partial}{\partial t} (\nabla \cdot \mathbf{x}_d) \mathrm{d}V \\ &+ \frac{1}{\mathrm{dim}} \oint_{\partial \mathcal{S}_d} (\nabla \cdot \mathbf{x}_d)(\dot{\mathbf{x}}_d \cdot \hat{\mathbf{n}}) \mathrm{d}S \\ =& \frac{2}{\mathrm{dim}} \oint_{\partial \mathcal{S}_d} (\dot{\mathbf{x}}_d \cdot \hat{\mathbf{n}}) \mathrm{d}S. \end{aligned} \tag{3.15}$$

This value is easily calculable and only involves the velocity at the boundary of the region. The expansion is mathematically captured by the divergence of velocity vector field,

$$\mathrm{div}(\dot{\mathbf{x}_d}) = \nabla_{\mathbf{x}_d} \cdot \dot{\mathbf{x}_d} \tag{3.16}$$

which when combined with Equation (3.13), yields

$$\oint_{\partial \mathcal{S}_d} (\dot{\mathbf{x}}_d \cdot \hat{\mathbf{n}}) \mathrm{d}S = \int_{\mathcal{S}_d} \nabla \cdot \dot{\mathbf{x}}_d \mathrm{d}V. \tag{3.17}$$

The implication of this equation is that the contribution to the expansion (or contraction) of the region at any particular point is due solely to the divergence of the velocity vector field. Taking the derivative of Equation (3.15)

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\oint_{\partial\mathcal{S}_d}(\dot{\mathbf{x}}_d\cdot\hat{\mathbf{n}})\mathrm{d}S &= \oint_{\partial\mathcal{S}_d}\frac{\partial}{\partial t}\left(\dot{\mathbf{x}}_d\cdot\hat{\mathbf{n}}\mathrm{d}S\right)\\ &= \oint_{\partial\mathcal{S}_d}\left(\ddot{\mathbf{x}}_d\cdot\hat{\mathbf{n}}+\dot{\mathbf{x}}_d\cdot\frac{\partial\dot{\mathbf{x}}_d}{\partial\mathbf{x}_d}\cdot\hat{\mathbf{n}}\right)\mathrm{d}S\end{aligned} \tag{3.18}$$

which requires calculation of the partials of the velocity vector field with respect to position. To accurately project area of a region forward in time, higher order derivatives should be used in the Taylor series expansion. Below are the derivatives of area from the second to the fifth time derivative. The general form is identical to higher order time derivatives of $f(t) = x(t)\cdot y(t)$, where $x(t)$ is $\dot{\mathbf{x}}_d$ and $y(t)$ is $\frac{\partial\mathbf{x}_d}{\partial\mathbf{x}_d}$

$$\frac{\mathrm{d}^2}{\mathrm{d}^2t}A_{\mathbf{h}}(\mathcal{S}) = \oint_{\partial\mathcal{S}_d}\left(\ddot{\mathbf{x}}_d+\dot{\mathbf{x}}_d\cdot\frac{\partial\dot{\mathbf{x}}_d}{\partial\mathbf{x}_d}\right)\cdot\hat{\mathbf{n}}\mathrm{d}S \tag{3.19}$$

$$\frac{\mathrm{d}^3}{\mathrm{d}^3t}A_{\mathbf{h}}(\mathcal{S}) = \oint_{\partial\mathcal{S}_d}\left(\dddot{\mathbf{x}}_d+2\ddot{\mathbf{x}}_d\cdot\frac{\partial\dot{\mathbf{x}}_d}{\partial\mathbf{x}_d}+\dot{\mathbf{x}}_d\cdot\frac{\partial\ddot{\mathbf{x}}_d}{\partial\mathbf{x}_d}\right)\cdot\hat{\mathbf{n}}\mathrm{d}S \tag{3.20}$$

which gives a general formula for the nth derivative,

$$\frac{\mathrm{d}^n}{\mathrm{d}^nt}A_{\mathbf{h}}(\mathcal{S}) = \oint_{\partial\mathcal{S}_d}\sum_{i=1}^{n}\mathbf{x}_d^{(i)}\cdot\frac{\partial\mathbf{x}_d^{(n-i)}}{\partial\mathbf{x}_d}\cdot\hat{\mathbf{n}}\mathrm{d}S. \tag{3.21}$$

Because a search region of arbitrary dimension and shape is being projected through a non-injective function, multiple orbits with multiple velocities could all project to a single $\mathbf{x}_d$. Even if $\dim(\mathcal{S}) = \dim(\mathcal{S}_d)$, the manifold of $\mathcal{S}$ might be folded in such a way that two or more parts of the original manifold map to the same part of $\dim(\mathcal{S}_d)$. The vector field is then ill-defined. Therefore, this chapter considers the 3 possible cases; the first is when the projection is one to one, the second is when the projection is a discrete number of points to one (folding), and the third is when a continuum of points in the original manifold map to

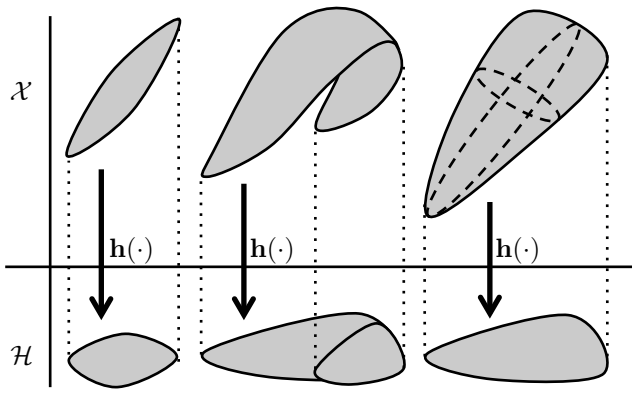

Figure 3.2: A manifold can project into the search space in a variety of ways. This shows such a projection for the one-to-one, discrete-to-one, and continuum-to-one case.

a single point in the projected manifold. These three cases are illustrated in Figure 3.2. For the one-to-one case, Equation (3.15) can be calculated directly, as can future derivatives. For the discrete-to-one case, $\dot{\mathbf{x}}_d$ has multiple values at a single $\mathbf{x}_d$. Equation (3.15) has several disparate values on each level of the original manifold. For any analytic purposes, these sub-manifolds have completely unique properties and need no special treatment.

If the projection is such that a continuum of points are mapped to a single point in the measurement space, a continuum of velocity vector fields are defined over the search set. Calculating area and divergence becomes difficult as there are multiple values for divergence at any given point. This work proposes defining an upper bound to the area rate-based on maximum possible line integral from the vector fields. Given a series of vector fields, define the set of $\dot{\mathbf{x}}_d$ defined at a particular $\mathbf{x}_d$,

$$S_{\dot{\mathbf{x}}_d}(\mathbf{x}_d) = \left\{ \dot{\mathbf{x}}_d = \frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{h}(\mathbf{x})) : \mathbf{x}_d = \mathbf{h}(\mathbf{x}) \right\} \tag{3.22}$$

This set definition now enables the following inequality,

$$\oint_{\partial \mathcal{S}_o} \dot{\mathbf{x}}_d \cdot \hat{\mathbf{n}} \mathrm{d}S \leq \oint_{\partial \mathcal{S}_o} \sup_{\mathbf{y} \in S_{\dot{\mathbf{x}}_d}} (\mathbf{y} \cdot \hat{\mathbf{n}}) \, \mathrm{d}S \tag{3.23}$$

where the supremum chooses the velocity with the largest outward orthogonal component

at any given point along the boundary. This integral is larger than any other piece-wise combination of vector fields. To understand the implication of this inequality, consider

$$\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{S} : \mathbf{h}(\mathbf{x}_1) = \mathbf{h}(\mathbf{x}_2) = \mathbf{x}_d \in \partial\mathcal{S}_d. \tag{3.24}$$

Now, note that

$$\dot{\mathbf{x}}_{d,1}(t) \cdot \hat{\mathbf{n}} > \dot{\mathbf{x}}_{d,2}(t) \cdot \hat{\mathbf{n}} \rightarrow \mathbf{x}_{d,2}(t + \delta t) \notin \partial\mathcal{S}_d \tag{3.25}$$

because the boundary $\partial\mathcal{S}_d$ must be moving outward at at a minimum rate of $\dot{\mathbf{x}}_{d,1} \cdot \hat{\mathbf{n}}$. In other words, if $\mathcal{S}_d$ has a continuum of velocity vector fields, the internal motion still does not effect the growth of the region, and the expansion at the boundary is fully described by $\sup_{\mathbf{y} \in S_{\dot{\mathbf{x}}_d}} (\mathbf{y} \cdot \hat{\mathbf{n}})$.

Finally, it is worth considering the behavior of partitions of a search set. The search space $\mathcal{S}_d$ can be partitioned into m sub regions,

$$\mathcal{S}_d = \bigcup^{m} \mathcal{S}_d^i \tag{3.26}$$

$$\emptyset = \mathcal{S}_d^a \cap \mathcal{S}_d^b \; \forall \; a \neq b \tag{3.27}$$

where $\mathcal{S}_d^i$ are the non-overlapping partitions. The area over time, and area growth, can be calculated for each partition to give insight into local areas of growth within the search space. Note that unless the manifold mapping into the search space is one-to-one, non-overlapping partitions in $\mathcal{H}$ may overlap in the future.

3.2.4 Admissible Region-based Search Area

This section focuses on how to apply the previous analysis to follow up on an admissible region by an optical observer. Consider an admissible region, $\mathcal{AR}$, defined at some previous

time t_0. Assume that an optical observer, located at $\mathbf{o}(t)$, wants to search $\mathcal{R}$ at time t_f.

$$\mathbf{x}_d = [\alpha, \delta]^T = \mathbf{h}(\mathbf{x}; \mathbf{o}) \tag{3.28}$$

$$\mathcal{S}_d = \{[\alpha, \delta]^T : [\alpha, \delta]^T = \mathbf{h}\left(\phi(t_f; \mathbf{x}, t_0); \mathbf{o}(t_f)\right), \mathbf{x} \in \mathcal{R}\}. \tag{3.29}$$

The divergence for this optical sensor is

$$\text{div}(\mathbf{x}_d) = \frac{\mathrm{d}\dot{\alpha}}{\mathrm{d}\alpha} + \frac{\mathrm{d}\dot{\delta}}{\mathrm{d}\delta}. \tag{3.30}$$

The area can then be calculated as

$$A_{\mathbf{h}}(\mathcal{S}) = \frac{1}{2} \oint_{\partial \mathcal{S}_d} (\mathbf{x}_d \cdot \hat{\mathbf{n}}) \mathrm{d}S = \frac{1}{2} \oint_{\partial \mathcal{S}_d} -\delta \cos(\delta) \mathrm{d}\alpha + \alpha \mathrm{d}\delta \tag{3.31}$$

where the signs and arrangement of α and δ account for the normal vector. The area rate of change is

$$\frac{\mathrm{d}}{\mathrm{d}t}(A_{\mathbf{h}}(\mathcal{S})) = \oint_{\partial \mathcal{S}_d} -\dot{\delta} \cos(\delta) \mathrm{d}\alpha + \dot{\alpha} \mathrm{d}\delta \tag{3.32}$$

Note that any occurrence of $\mathrm{d}\alpha$ has an additional factor of $\cos(\delta)$. This is because a differential in angular space such as $\mathrm{d}S$ or $\mathrm{d}A$ must account for the curved nature of the angular space. This also implies that all angles must be represented in radians, and the area being calculated is in steradians. These may also be difficult or impossible to calculate near the singularities at $\delta = \pm\pi/2$.

The higher order derivatives of $A_{\mathbf{h}}$ exist and can be calculated. An infinite Taylor series can be used to approximate the area over time to very high accuracy. It can be assumed that sufficient accuracy is achievable by increasing the number of terms in the Taylor series approximation.

The admissible region from an optical observer is a two-dimensional manifold of orbits.

Any search set therefore is either entirely one-to-one or has some folding that causes multiple orbits to map to a single α and δ pair. Qualitatively, there is no guarantee that folding does not occur, but in practical experience over short periods of time after the inception of $\mathcal{R}$, folding tends to not occur. The structure of $\mathcal{R}$ makes the analysis on partitioning particularly powerful; because $\mathcal{R}$ is a continuously defined region, partitioned regions do not overlap outside of folding.

If no folding occurs for a two-dimensional manifold, the original boundary of the manifold is associated with the projected boundary. When considering an admissible region, the projected boundary of the region then corresponds to the boundary of $\mathcal{R}$ in the original space.

Equation (3.18) requires calculating $\frac{\partial \dot{\mathbf{x}}_d}{\partial \mathbf{x}_d}(t)$. This quantity must be calculated along the manifold created by the propagated admissible region, $\mathcal{S}_d(t)$. This manifold, at a particular time t, does not have a closed form definition. Instead the derivative must be tied back to the original definition of the manifold, that is, the original admissible region. The variation of a location in $\mathcal{H}$ with respect to admissible region coordinates is

$$\begin{aligned}\frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}_u(t_0)} &= \frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}(t)}\frac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_0)}\frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \\ &= \frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}(t)}\Phi(t, t_0; \mathbf{x}(t_0))\frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)}\end{aligned} \tag{3.33}$$

where $\frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}(t)}$ is the Jacobian of the measurement function, $\mathbf{h}$, and Φ is the state transition matrix. Similarly,

$$\begin{aligned}\frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}_u(t_0)} &= \frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}(t)}\frac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_0)}\frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \\ &= \frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}(t)}\Phi(t, t_0; \mathbf{x}(t_0))\frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)}\end{aligned} \tag{3.34}$$

where, similarly, $\frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}(t)}$ is the Jacobian of the derivative of the measurement function. Then

the final quantity is

$$\frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}_d(t)} = \frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}_u(t_0)} \left(\frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}_u(t_0)} \right)^{-1}. \tag{3.35}$$

Note that although this definition requires a matrix inverse, the final matrix is an easily invertible two-by-two. The above calculation does not require an understanding of the manifold orbit to calculate; instead for a particular point in the sky, the exact point in the admissible region is needed. The given $\rho, \dot{\rho}$ pair defines an orbit with the original measurement, which defines all the inputs to Equation (3.35). The same process as above can be used to calculated higher order derivatives, by replacing $\dot{\mathbf{x}}_d(t)$ with $\ddot{\mathbf{x}}_d(t)$.

The matrix in the above inverse is not guaranteed to be full rank. Qualitatively, the inverse not existing refers to the case where a small change in $\rho, \dot{\rho}$ provides no change in α, δ. The only case where moving along this two dimensional manifold provides no change in the measurement space is when the manifold (in $\mathcal{X}$) is directly orthogonal to the measurement subspace, $\mathcal{H}$. This occurs in certain special cases, for example when $t = t_0$. More importantly, when folding occurs in the measurement space, this undefined inverse occurs exactly along the fold. This fact is both good and bad; this calculation provides a useful way of predicting when a manifold is folding, but makes the area acceleration undefined along that boundary.

3.2.5 Analysis of the Search Set

This section focuses on applying the above analysis to the observed and unobserved sets, $\mathcal{O}$ and $\mathcal{U}$. Given a particular sensor, $\mathbf{o}$, over every integration, t_I, at time t_k a known area of the sky can be observed, $\mathcal{O}_k(t_k)$. Then the total search set area can be calculated at each

time step as the Taylor series

$$\begin{aligned} A_{\mathbf{h}}(\mathcal{S}(t_k)) &= |\mathcal{S}(t_k)|_{\mathbf{h}} \\ &\approx \sum_{i=0}^{\infty} \frac{\mathrm{d}^i}{\mathrm{d}^i t} A_{\mathbf{h}}(\mathcal{S}(t_0)) \frac{(t_k - t_0)^i}{i!} \end{aligned} \tag{3.36}$$

while the observed area can similarly be represented as

$$\begin{aligned} A_{\mathbf{h}}(\mathcal{O}(t_k)) &= \left| \bigcup_{l=0}^{k} \mathcal{O}_l(t_k) \right|_{\mathbf{h}} \\ &\approx \sum_{i=0}^{\infty} \frac{\mathrm{d}^i}{\mathrm{d}^i t} A_{\mathbf{h}} \left(\bigcup_{l=0}^{k} \mathcal{O}_l(t_0) \right) \frac{(t_k - t_0)^i}{i!} \end{aligned} \tag{3.37}$$

where l is a counter between time step 0 and the current time step, k. Note that, with the Taylor series expansions, all analysis on the search set can be done at the initial time step. Equation (3.36) is a smooth function which represents the gradual expansion and contraction of the search set in the measurement space. Equation (3.37) is instead a step-like function with respect to t_k; at each t_k a new observation is incorporated into Equation (3.7), instantaneously expanding $\mathcal{O}$. The more interesting set is the set of unsearched orbits,

$$A_{\mathbf{h}}(\mathcal{U}(t_k)) = A_{\mathbf{h}}(\mathcal{S}(t_k) \setminus \mathcal{O}(t_k)), \tag{3.38}$$

as this represents the extent of the current search problem. This equation is explicitly dependent on the proposed search strategy. Note that if $\mathbf{h}$ is injective and continuous,

$$A_{\mathbf{h}}(\mathcal{U}(t_k)) = A_{\mathbf{h}}(\mathcal{S}(t_k)) - A_{\mathbf{h}}(\mathcal{O}(t_k)), \tag{3.39}$$

These equations can be turned into a stepwise form,

$$A_{\mathbf{h}}(\mathcal{U}(t_{k+1})) = A_{\mathbf{h}}(\mathcal{U}(t_k)) + \frac{\mathrm{d}}{\mathrm{d}t} A_{\mathbf{h}}(\mathcal{U}(t_k))(t_{k+1} - t_k) - A_{\mathbf{h}}(\mathcal{O}_{k+1}(t_{k+1})) \tag{3.40}$$

where the first term is the unsearched area at the previous time step, the second term is the unsearched area expansion over the time step, and the third term is the reduction of $\mathcal{U}$ by observing a new region of measurement space. If in general

$$\frac{\mathrm{d}}{\mathrm{d}t} A_{\mathbf{h}}(\mathcal{U}(t_k))(t_{k+1} - t_k) > A_{\mathbf{h}}(\mathcal{O}_{k+1}(t_{k+1})) \tag{3.41}$$

the current search strategy does not close; the region is expanding faster than it is being observed. This provides a simple and efficient calculation for evaluating whether a search is feasible. Furthermore, the amount the area changes each time step provides insight into approximately how long a search should take.

The first derivative of area is dependent on the divergence, giving a tool for defining short term high priority search areas. Figure 3.3 shows the divergence changing over a region. Given multiple choices for $\mathcal{O}_k$, a higher divergence subset, given equal time, expands into a larger, harder to search area at a future time. For example, in Figure 3.3, the region of high divergence on the left should be prioritized over the region of low divergence on the right because at some final time t_N

$$A_{\mathbf{h}}(\mathcal{O}_k^a(t_N)) > A_{\mathbf{h}}(\mathcal{O}_k^b(t_N)) \tag{3.42}$$

which, combined with 3.40, implies a smaller $\mathcal{U}$ at a future time step. This is the basis of a finite horizon control method presented in Section 3.3.3.

Remember that the set $\mathcal{O}_k$ is dependent on sensor location and field of view, which is given when solving a particular search. An additional application of this work is to also

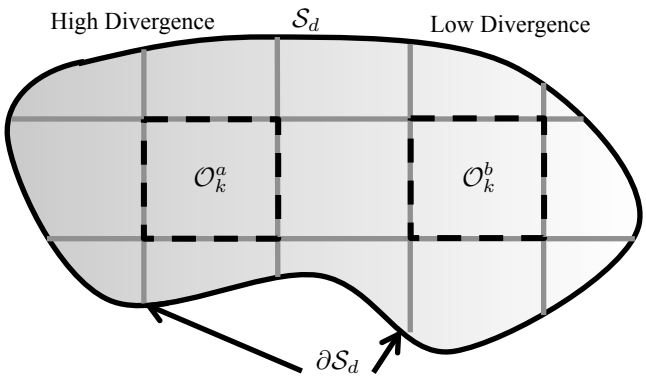

Figure 3.3: The divergence of $\dot{\mathbf{x}}_d$ determines which areas of a search set are high priority.

be used as a sensor design tool. Varying sensor location and field of view gives different results in 3.40 allowing interesting comparisons.

Equation (3.36) contains an infinite Taylor series expansion in it, which in practice requires sufficient terms to guarantee small errors. This can be done by defining both a time horizon, t_f, and a relative tolerance, ϵ_{rel}, measured in steradians. Recall that a steradian is a unit of angular area. Then the required N terms for an approximation out to a certain time horizon satisfies

$$\epsilon_{rel} > \frac{\mathrm{d}^N}{\mathrm{d}^N t} A_{\mathbf{h}}(\mathcal{S}(t_0)) \frac{(t_f - t_0)^N}{N!}. \tag{3.43}$$

It should be noted that the divergence is PDF agnostic. When operating on a pure set or a PDF whose effective boundaries are used to define a set, the method gives the same result. This is particularly relevant to admissible region theory, where a set is used to define a PDF. Divergence methods operate identically on either representation of the admissible region, avoiding analytic confusion.

Search regions can become sufficiently large that a region extends out of the field of regard on an EOS. This may even extend to a ring of orbits around Earth, or all of the unit sphere, $\mathbf{S}(2)$. In these cases, the total region divergence may no longer be a useful measure, but divergence of partitions is very much still a useful measure.

3.2.6 Reachability Considerations

Consider a search set, $\mathcal{S}(t_0) \subseteq \mathcal{X}$, as a set of initial conditions for an orbit. These orbits have well defined dynamics, and can therefore define a control problem. In this framework, the set at a future time step, $\mathcal{S}(t)$, can be written in terms of a reachability problem

$$\mathcal{S}(t) = \mathcal{R}(t, t_0, \mathbf{u}_{max}; \mathcal{S}(t_0)) \tag{3.44}$$

where $\mathcal{R}$ is the reachable set of orbits at time t, based on the initial condition set, $\mathcal{S}(t_0)$, and the maximum control effort $\mathbf{u}_{max}$. The problem discussed in above sections can then be cast as a reachability problem, where the search set at a given time step is the reachable set-based on the initial condition set, $\mathcal{S}(t_0)$, and a maximum control effort of $\mathbf{u}_{max} = \mathbf{0}$.

This generalization allows the theoretical framework in this chapter to incorporate unknown control effort. Using any in a variety of control distance analysis frameworks [21], a reachability set can be calculated and used as the search set. In general, $\mathcal{R}$ is not a manifold but rather full dimension subset of $\mathcal{X}$. Like $\mathcal{S}$ in sections above, $\mathcal{R}$ can be amended with constraints to make the search set more manageable. Calculation of the derivatives of $A_{\mathbf{h}}$ are based on an uncontrolled trajectory. Therefore the search set $\mathcal{S}_d$ grows faster than what is predicted by the search area divergence calculations. A full analysis of the technical challenges involved in reacquiring maneuverable objects is out of the scope of this dissertation.

3.3 Theory: Optimal Search

3.3.1 The Control of Search on a Set with an Optical Sensor

The search region elements are assumed to be orbits consisting of position and velocity at a given time, $\mathbf{x}(t) = [\mathbf{r}^T(t), \mathbf{v}^T(t)]^T \in \mathbb{R}^6$. For an optical observer, the determined states

are the inertial angle pair, α and δ, with the undetermined states as follows

$$\mathbf{x}_d = [\alpha, \delta]^T \tag{3.45}$$

$$\mathbf{x}_u = [\dot{\alpha}, \dot{\delta}, \rho, \dot{\rho}]^T. \tag{3.46}$$

Note that while angular rates are not instantaneously observed, they are often inferred from a series of measurements. From a tasking perspective, it is assumed that regardless of the angular rate of the space object, it can be detected if it is in the field of view. If a sensor can approximately rate track the observed area, $\mathcal{O}(t_k)$, the assumption is good so long as the rates do not vary dramatically over the field of view. If they do vary dramatically, the expansion of the search set is often prohibitively high for a complete search.

The observed part of the measurement space for an optical sensor is just the field of view of the sensor. The observed state space is then fully defined by the state of the observer, $\mathbf{o}, \dot{\mathbf{o}}$, and three angles, $[\alpha_u, \delta_u, \theta_u]$, where α_u and δ_u define pointing and θ_u defines rotation about the line of sight vector. For analytic and computational simplicity, the bore sight rotation, θ_u, is ignored. The full observation campaign is represented as a pair of vectors of pointing angles, and $\mathcal{O}$ is represented as an explicit function of these controls,

$$\mathbf{U} = [\boldsymbol{\alpha}_u, \boldsymbol{\delta}_u] \tag{3.47}$$

$$\mathcal{O}_k(t_k) = \mathcal{O}_k(\boldsymbol{\alpha}_u(k), \boldsymbol{\delta}_u(k); t_k) \tag{3.48}$$

where the entire campaign consists of N observations. This is illustrated in Figure 3.4. An EOS makes an observation by taking a series of images in a certain location in the sky. The time spent attempting a detection at time step k, is represented by $t_I(t_k)$. Once these series of images have been taken, the EOS slews to a new location. Because every sensor mount is different, an arbitrary function, $f_t(\alpha_1, \delta_1, \alpha_2, \delta_2)$, represents the time it takes a particular sensor to slew from location 1 to location 2.

If the search set is defined similar to an admissible region, that is as a series of con-

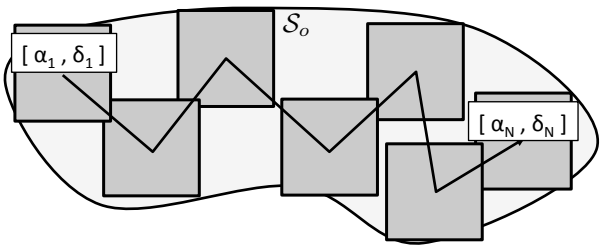

Figure 3.4: Example of a tasking trajectory over a portion of the sky.

straints, the removal of pieces of search set can be analytically incorporated into admissible region theory. An admissible region is an intersection of a series of sets, each defined in terms of a constraint,

$$\mathcal{U} = \bigcap \mathcal{U}_k \tag{3.49}$$

$$\mathcal{U}_k = \{\mathbf{x}_u \in \mathbb{R}^u : g_k(\mathbf{x}_u; \mathbf{x}_d, \dot{\mathbf{x}}_d, \mathbf{k}) \leq 0\} \tag{3.50}$$

where g_k is a constraint and $\mathbf{k}$ is a parameter vector. The removal of observed parts of state space defines the constraint set,

$$\mathcal{U}_k = \{\mathbf{x}_u \in \mathbb{R}^u : \phi(t_j; \mathbf{x}_u, \mathbf{x}_d, t_k) \notin \mathcal{F}(t_j), \forall j\} \tag{3.51}$$

which in essence removes whatever section(s) of a set that intersect a particular observation.

3.3.2 Time-Optimal Cost Function

Many search regions have positive over all divergence, though there can exist situations where this is not the case. Time-optimal search is built on the assumption that sensor time is a valuable commodity, and the best way to optimally use a sensor is to spend as little time accomplishing a task as possible. The cost function is then the total time spent over

the observation campaign,

$$f_{time}(\boldsymbol{\alpha}, \boldsymbol{\delta}) = \sum_{k=1}^{N} [t_{I,k} + f_t(\alpha(t_k), \delta(t_k), \alpha(t_{k+1}), \delta(t_{k+1}))] \tag{3.52}$$

where f_t is the slew time of the sensor between two points. The constraint on this optimization is to reduce the size of the unobserved set, $|\mathcal{U}|_{\mathcal{H}}$, to zero. The constraint is difficult to enforce directly as it depends upon the intersection of time varying sets. One way to incorporate it is as a penalty function. It can be enforced as a penalty function

$$f_{cost}(\boldsymbol{\alpha}, \boldsymbol{\delta}) = f_{time}(\boldsymbol{\alpha}, \boldsymbol{\delta}) + f_{penalty}(\boldsymbol{\alpha}, \boldsymbol{\delta}; \mathcal{S}) \tag{3.53}$$

Specifically, this is an exterior penalty function [85]; as the tasking scheme becomes inadmissible, a large cost is added to the function forcing the constraint to be enforced. Penalty functions have a known problem in that they can lead to inadmissible solutions to the optimization. The constraint in practice is enforced with a discrete point-wise approximation of the search set, $\mathcal{S}(t)$, to approximate the percentage of the set contained in the observation.

$$f_{penalty}(\mathcal{U}) = |\mathcal{U}(t_f)|_{\mathbf{h}} \tag{3.54}$$

where $\mathcal{U}$ is defined through Equation (3.7) and Equation (3.8).

3.3.3 Divergence Greedy Cost Function

Because of the computational difficulty of doing a full optimization, a second optimization method is being proposed. This method is based on the idea that search should be done quickly to make best use of a sensor. Furthermore, it is assumed that there is no PDF to define the location of an object, only a set. This is particularly relevant to follow up on an admissible region. Consider Equation (3.36). At a given time step, k, the best control is the one which minimizes $\mathcal{S}_o(t)$ for all time. This implies that the observed space should be as

large as possible, but it also implies the divergence is as large as possible. By maximizing the divergence in the observed space, the rate of change of area at future time steps is minimized. This optimization can be solved at each time step but is better posed in a finite time horizon manner. Again, assume a series of observations defined in Equation (3.47), but assume further that the total number of observations, N, defines a finite horizon over which the optimizer searches. A cost function that looks to search the largest area possible with the largest divergence possible would then be

$$f_{Div} = -\sum_{k=1}^{N} \frac{\mathrm{d}}{\mathrm{d}t} |\mathcal{F}_o(\mathbf{u}(t_k)) \cap \mathcal{S}_o(t(t_k))| \tag{3.55}$$

keeping in mind that $\mathcal{S}_o(t_k)$ changes each time step based on both dynamics and sections being removed by previous observations. Note that the total divergence is proportional to the area which is integrated over, so this cost function prefers large, highly divergent regions.

3.3.4 Optimization Methods Considered

This section discusses a few of the optimization methods attempted on the above problems and reports on the success of each method.

A basic descent method is considered for the full time-optimal cost function [85]. The performance of the descent method is well below acceptable, primarily due to two considerations. First, the region itself is highly multi-modal, and fairly poor local optima are often convergent. Second, the nature of the problem formulation requires a discrete point approximation of the set. The constraint on the optimization, in the form of a penalty function, acts as a step function when calculating the derivatives making an accurate approximation difficult.

Because a stochastic optimization method is clearly needed, a genetic algorithm could provide a solution. However, the genetic algorithm proves insufficient, primarily due to the

high dimensional nature of the search space and the computational requirements to run a robust enough genetic algorithm on this problem. Further work may provide more success, but it is beyond this research arc.

Finally a simulated annealing-like algorithm is considered. Over all, this algorithm has the best success; it is able to search the complicated region, avoid the many local minima, avoid derivatives, and generally behave as a descent method would. However, even the success of the simulated annealing is only a partial success. A global minimum is still difficult to find, and stochastic search methods are inherently unrepeatable. The algorithm, after fine tuning, can find an acceptable solution reliably, but often creates obviously suboptimal solutions. More details on this algorithm can be found in [86]. The difficulties of this search method motivate this analysis that could inform the search before an optimization takes place, and enable a smaller scope optimization like the receding horizon method.

3.4 Simulation Results

3.4.1 Colocated Reacquisition

This test case demonstrates the results of the analytic work and control implementation on a situation with a small size, two-dimensional manifold. Because the manifold is 2 dimensional, the higher order area derivatives are all calculable. This allows the full power of this work to be demonstrated on a best-case, but still useful, scenario.

In the test case presented one observer, located at Georgia Institute of Technology observes, loses and attempts to reacquire an single object. The space object is in a Geosynchronous orbit in view of the observer. The UTC time is such that the Earth-centered Earth-fixed, and Earth-centered inertial frames are aligned at $t = 0$, for simplicity. First, at time $t_0 = 0$ the observer receives a short arc observation generating an admissible region containing the object of interest. The admissible region is assumed to have no uncertainties, and defines the search set, $\mathcal{S}(t_0)$. The ground station (GS) is at a latitude and longitude of 33.78° and -84.40°. The space object (SO) state truth at t_0 is shown in the Table 3.1.

Table 3.1: Space object state in the Earth-centered inertial frame.

SO ECI location, $\mathbf{r}$ (km)	SO ECI velocity, $\dot{\mathbf{r}}$ (km/s)
$\begin{bmatrix} -27.1 \\ -32.3 \\ -0.1 \end{bmatrix} \times 10^3$	$\begin{bmatrix} 2.36 \\ -1.98 \\ 0.0 \end{bmatrix}$

At t_0 the search set is all located in a single point in the sky relative to the observer. To create an interesting test case, one hour is allowed to pass, such that $t_1 = 1$ hour, and the admissible region is propagated and projected in order to calculate $\mathcal{S}_{\mathbf{h}}(t_1)$. The timing and geometry of this test case are chosen such that the resulting search is plausible, difficult to fully search, but still solvable. Reacquisition is attempted on this region.

First, an analysis of the region area as a function of time is presented. The search set is propagated over the course of a one hour interval to $t_2 = 2$ hours, starting at t_1. The search set in the measurement space, $\mathcal{S}_d$, has a size of 0.54 millisteradians (msr)at t_1, and a size of 14 millisteradians at t_2, calculated through propagation. The area is then also predicted out over the same time period using the area time derivatives in Equation (3.21). Figure 3.5 presents the true and estimated areas between t_1 and t_2 for first through fifth order Taylor series approximations. Figure 3.6 presents the log of the percent error. Note that the error goes negative over a short time span before diverging towards positive, accounting for the behavior on the smaller times in Figure 3.6. Percent error can be reduced further with higher order approximations.

Next, an actual observation control trajectory, $\mathbf{U}$, is calculated. The optimization presented is a final time area greedy algorithm; at each time step, t_i, the observation is chosen whose area at some horizon time, t_f, is maximum, $|\mathcal{O}_{k,d}(t_f)|_{\mathcal{H}}$. For this test case, a 100 observation time horizon is used to create a trajectory that fully searches the region, using an exhaustive search. Only 83 observations are required to have complete search.

The trajectory is calculated for 83 observations over which the region is observed. Fig-

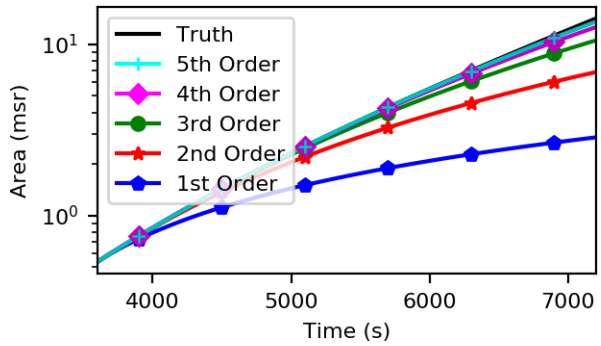

Figure 3.5: Absolute error for various Talyor series area approximations.

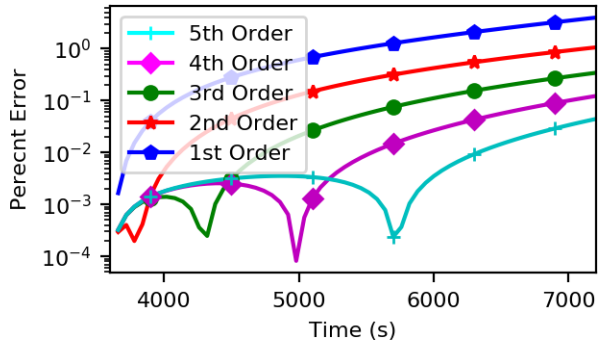

Figure 3.6: Log of percent error for various Talyor series area approximations.

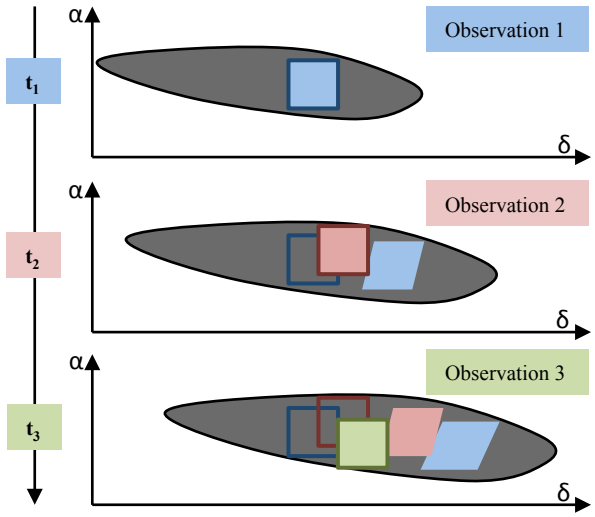

Figure 3.7: Illustration of the search set's movement and growth's affect on observation campaign.

Table 3.2: Observation campaign parameters.

Field of View	Dwell time	Slew time
0.25°	15 seconds	0 seconds

ure 3.7 gives an intuitive illustration of the form of the results. Each observation is taken in a particular part of the sky, which corresponds to a particular region of the search set. The region moves in the positive right ascension direction, while simultaneously expanding and internally shifting over the course of the simulation. This change means that future observations may overlap with previous observations in the sky, and observations taken at different times do not overlap perfectly like the square observations could.

Figure 3.8 shows snapshots of the trajectory after 30, and 83 observations. The black x's are the locations of the various observations. The black square is the current location of the sensor field of view. The final trajectory tends to move in the negative right ascension direction, against the movement of the region. Because of these movements in opposite directions, the trajectory appears to dwell in the same area for an extended period of time, but the observations are all looking at distinct areas in the search set. The original region location is shown on the left side of the figure. The dark regions of the search set are the unobserved set, $\mathcal{U}$, while the light regions are the observed set, $\mathcal{O}$.

Figure 3.9 shows the area of the original region, and the area of the unobserved region over time. The dashed line represents the area over time of the original region, while the solid line represents the unobserved area. In the lower figure, the unobserved region is projected out over a a short time at each time step to give perspective on the remaining area growth rate. The figure illustrates the difficulty of this test case; the original region is growing at a substantial rate, which is constantly fighting against the observation schemes ability to scan the region. A simple scan of the region is also shown in Figure 3.9. This scan starts on the right side of the region (positive RA), and creates columns of measurements in the DEC direction. Each column of observations is one field of view size to the left of the

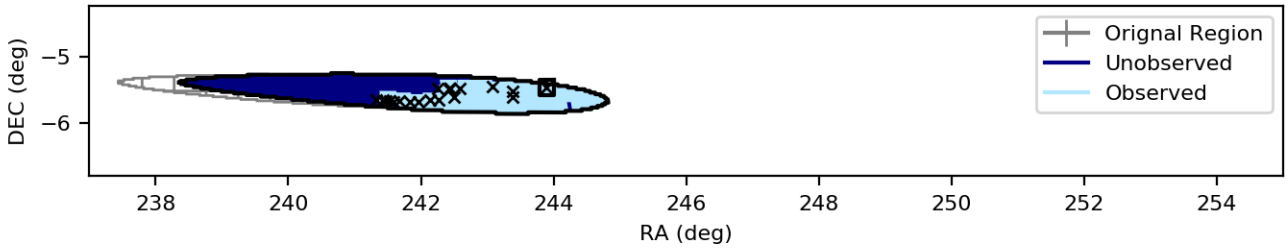

(a) Trajectory after 20 Observations.

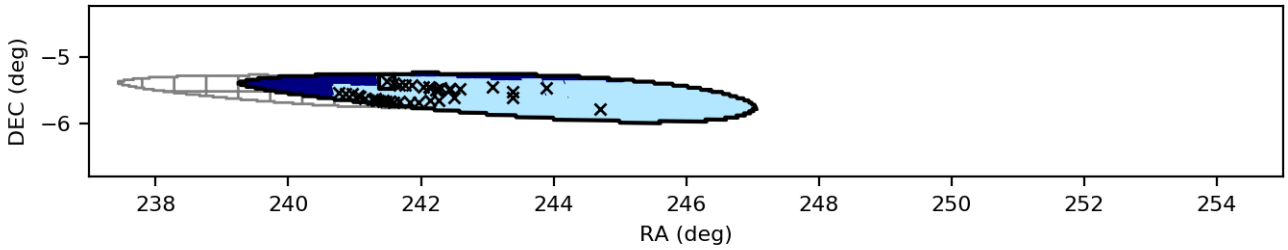

(b) Trajectory after 40 Observations.

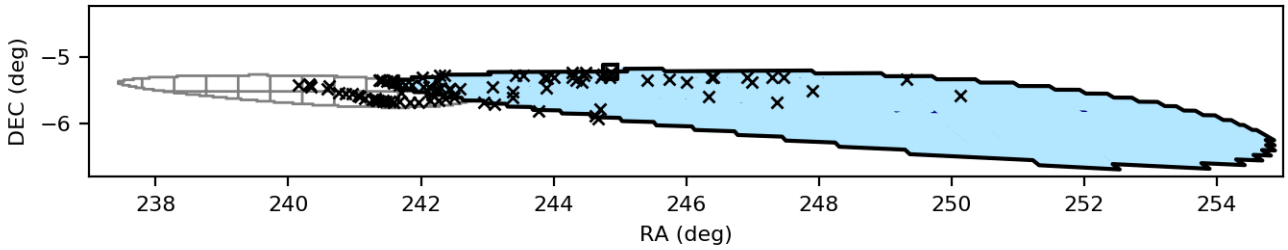

(c) Completed Trajectory after 83 Observations.

Figure 3.8: Final search campaign after 20, 40, and the full 85 observations. Observation locations marked with x's.

previous column. The observations also take into account the approximate motion of the region. This scan attempts to leave no "holes" in the region, instead trying to thoroughly search from right to left. This example is intended to show the pitfall involved in analyst-designed raster scans, which are commonly employed in searches.

3.4.2 Geostationary Slot Protection

The final test case demonstrates the method operating on a more complicated set. This set is defined as the objects in a geostationary transfer orbit (GTO), which reach apogee in a

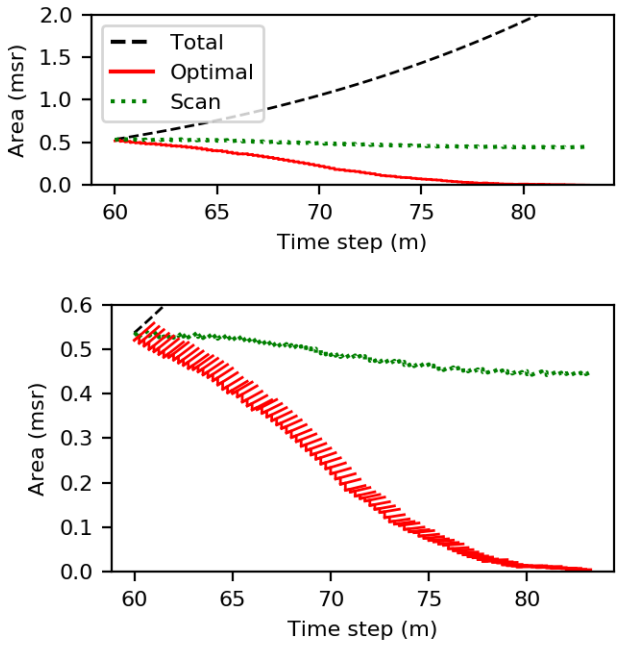

Figure 3.9: Total area and unobserved area over time for same location hand-off.

particular geostationary slot. Specifically, the scalar radius of perigee is set as constant, $r_p = 7000$km, the three-dimensional position at apogee is any where in a cube of 73 km sides centered around a particular geostationary object (approximating a "GEO slot"), the inclination is between -5 and +5 degrees, and the time until intercepting GEO is between 10 and 40 minutes. Note that the above definition uses negative inclination to represent objects coming up from below the GEO plane (e.g. $i = -5$ is really an $i = +5$ but with true anomaly, argument of perigee, and argument of ascending node rotated by $180°$). These six quantities (absolute perigee, position at apogee, inclination, time to intercept) fully constrain an orbit, and the ranges on all quantities except perigee define a set of orbits. Note that the position at apogee is defined around a geostationary orbit at a particular time defined by the time to intercept. The resulting region is a trapezoidal set of orbits which get closer together as they approach the target GEO slot. The geometry of this problem is nominally shown in Figure 3.10.

The exact same method as in previous simulations is used to calculate a trajectory with which to scan through the region, which is shown in Figure 3.11. Because of the nature of

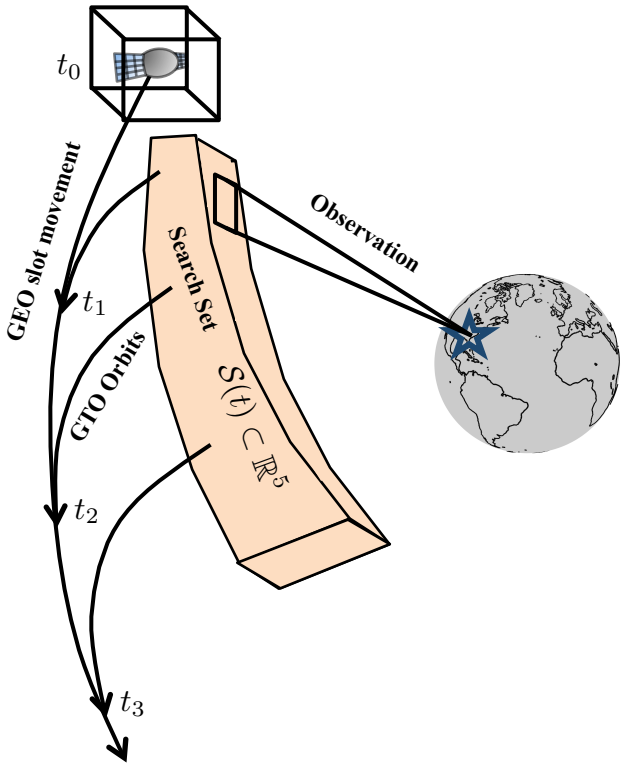

Figure 3.10: Illustration of the 3 dimensional GEO slot, and 5 dimensional search set of intersecting GTO orbits.

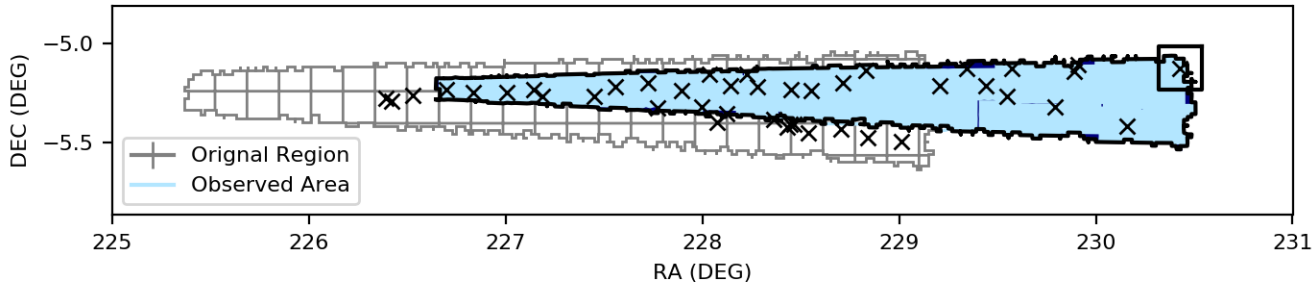

Figure 3.11: GTO region at beginning and end of search. Observation locations marked with x's.

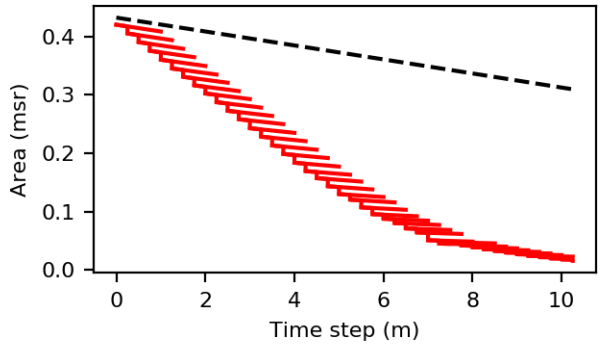

Figure 3.12: Total area and unobserved area over time for GTO Search.

this set, it is actually getting smaller as time goes on meaning the area rate is negative. The total area and unobserved area are shown in Figure 3.12. Note that the rate of contraction of the remaining search set appears to get smaller as the simulation runs on. This is due primarily to the remaining area of the region; the rate of contraction of a smaller area is, by definition, smaller.

3.4.3 Further Discussion of Results

This final section is intended to discuss the results over all, including many of the practical problems in implementation. In all test cases, a grid on points over RA and DEC at t_0 is calculated, along with all higher derivatives of α and δ. All area over time calculations require only to pick a series of grid points and calculate a few line integrals. This methodology still requires propagating the grid to future time steps, in order to determine what grid

points are inside a potential $\mathcal{O}_k$ at time step k. To lighten this burden, the N time steps are pre-chosen and the grid is pre-propagated to each t_k. After this no propagation is required. The derivatives of area can be ignored in favor of pure particle propagation, but no major computational benefits are reaped either way.

The greedy optimization can lead to a problem where, because sequential observed sets, $\mathcal{O}_k$, may not align along their boundaries, small holes in the observed set can form. The beginning of this problem can be seen notionally in Figure 3.7. At the end of the campaign, these holes must be filled one at a time, causing the final search sets to be small, and lengthening the total campaign time. The results in this section used a heuristic to mitigate this problem. When choosing $\mathcal{O}_k$ in the optimization algorithm, the field of view is reduced by 10% on each side before being input into the cost function. Then, once an optimal location is chosen, it is expanded on each side to determine the true $\mathcal{O}_k$. This encourages overlap between observed sets and reduces holes in the final campaign. However, aggressive reduction of these holes can be counter productive. Small holes and slivers in the unobserved set, $\mathcal{U}$, tend to stay small over long periods of time. Attempting to leave no holes or slivers is the basis for the scan shown in Figure 3.9, which can lead to an infeasible or highly inefficient search. Ultimately, an exhaustive optimization search for $\mathbf{U}$ would avoid this problem, but the curse of dimensionality makes it infeasible, at least within the scope of this chapter.

The results in Figure 3.9, which compare the presented tasking method to a raster type scan, warrant further discussion. The scan fails to make sufficient progress primarily due to the inability to leave gaps. When the search region height is divisible by the field of view height, the observations stack nicely and both methods have comparable performance. When the search set is, for example, 2.1 field of views wide, two observations are used efficiently but a third observation only utilizes 10% of the sensor field of view. This is compounded with the inefficiency involved in not prioritizing high divergence, which is discussed in depth in the theory section. The raster scan shown in this chapter is not nec-

essarily representative of all test cases and analysts may be able to design better scans than the tasking algorithm in some situations. However, the automated way in which this algorithm operates, along with the analytic tools developed in the theory, provide the value of this work.

3.5 Conclusion

This chapter provides a framework for SSA tasking strategies. Reacquisition, scheduling, custody, and search operations can all be represented with a search set. As such, all can be analyzed with divergence methods.

A search set is best analyzed when projected into the measurement space of a particular observer. The evolution of these regions over time is dominated by the vector field of velocity and higher order derivatives and the divergence of those vector fields, along the boundary of the search set. This allows quick analytic solutions for the evolution of the search set over time. The area over time makes for a suitable metric for both analyzing and predicting tasking schemes on a search set. This metric also is suitable for use as a control scheme cost function.

Many search sets are still infeasible to search, and this method does little to solve that problem. Rather, it provides tools to analyze search sets, and shows how to make search feasible. This is not always easy or possible. Search sets can be challenging to define in an explicit, analytic way, making the analytic formulations involved in this work difficult to calculate.

3.5.1 Unifying Example

At this point in our example, an admissible region has been formed for an unknown detected object. According to prior work, observability is achieved past ten minutes and closer to one hour [71]. Reacquisition should then take place somewhere between 30 minutes and one hour after the first detection (in order to provide a full orbit estimate). Figure

3.8 shows the reacquisition process after an hour for the exact situation being described. Because of the size of the search region, reacquisition after an hour is shone to be difficult so reacquisition after 30 minutes, or with a larger field-of-view, will provide more reasonable results. This first contribution outputs a series of locations and times to search the night sky for a reacquisition sensor. It also, notably, provides a dwell time to spend in each location searching.

CHAPTER 4
OPTIMAL SELECTION OF TELESCOPE PARAMETERS FOR ASTROMETRIC AND PHOTOMETRIC PERFORMANCE

4.1 Introduction

This work proposes a methodology for optimally selecting telescope mount and camera parameters when taking data for space situational awareness applications with various objectives. The primary objectives examined in this work are maximizing information of object location and speed in an image, and information on object brightness over time. The design variables are the telescope integration times over a series of images, the mount angular rates, and camera gain. Constraints include the data downlink rate, pixel saturation, and star field clutter. The final problem is a multi-objective optimization problem with a series of constraints and a variable length vector in the form of the exposure times. The simulation work analyzes the various reward functions individually and the combined Pareto surface.

Telescope operators, when performing SSA missions, have a variety of parameters which must be chosen. This includes exposure time, gain and gamma, exact pointing and slew rate, and in some cases bore sight rotation [34]. These parameters can have large impact on data efficacy, but are often chosen in an ad-hoc manner [87]. From the perspective of research, the telescope is often modeled academically as a sensor which is pointed towards a location in the sky, and obtains measurements on sufficiently bright objects. To the authors knowledge, telescope parameters have never been rigorously explored for SSA specific problems.

This chapter treats the telescope tasking problem as a multi-objective optimization problem. The telescope parameters in question are the design variables for this problem,

while the objectives are "good" astrometry and photometry. Constraints are introduced by calling out situations which would corrupt data beyond use or be physically impossible.

4.1.1 Methodology

In section 4.2 the optimization problem is set up. First, the various available design variables are defined and discussed. The statistics of image and point source data form the groundwork of this chapter and are succinctly introduced. A few other background topics are discussed as well.

Section 4.3 introduces the three main reward functions in this chapter. First, maximum likelihood estimation on an image is introduced along with a general formula for Fisher information. The astrometric information for both position and velocity is presented. Both these functions are developed for both a stationary and streaking object. Timing uncertainty is also incorporated into these reward functions. A formulation for photometry information is developed, which incorporates flux bias from rapidly changing signals.

Section 4.4 finishes defining the optimization problem by discussing various constraints. This includes side constraints for variables and physical constraints for a telescope system such as pixel saturation and pixel bleeding. More interestingly, star background clutter is analyzed as a constraint, providing an analytic way to trade off information gain from long exposures with potential information loss from objects intersecting stars.

Section 4.5 discusses the various optimization techniques used in this chapter. First, gradients for each of the reward functions are developed, which provide both insight into how the reward functions behave and a faster way to solve the optimization problem. The Karush Kuhn Tucker conditions are briefly discussed in this section. Finally, a loose outline of how the single objective and multi-objective optimizations are solved is presented. This includes a discussion of some of the hurdles and pitfalls in the optimization process.

Finally, Section 4.6 analyzes the various reward functions in simulation. First each reward function is solved by itself. Discussions on the general shape of solutions along

with intuitive analysis of the behavior of the reward functions are included. The combined multi-objective optimization process is also analyzed, with the inclusion of both a Pareto surface and sample solutions along the surface.

4.2 Optimization Set Up

This section, in order, defines the design variables, objectives, and constraints of the optimization problem.

4.2.1 Available Design Variables

Before the definitions and behaviors of the various SSA objectives can be analyzed, a list of design variables is needed.

First, the camera integration time is defined. Integration time, t_I, also known as exposure time, is the amount of time spent collecting photons for a single image. Longer integration time implies more photons are collected which implies higher informational content. Longer exposure time also leads to higher probability of detection in a particular frame. Shorter integrations leads to more images over a set period of time which gives more temporal information. Some photometry missions attempt reconstruction of highly variable light curves over time, while velocity information in images needs some element of temporal information.

In general, this chapter assumes multiple images may be taken, each with its own integration time. This forms a rather unique design variable, a variable length vector, $\mathbf{t} \in \mathbb{R}^{N_t+1}$. To aid this discussion, refer to Figure 4.1 for clarification. This chapter uses the start times of each integration, t_k, as the design variable and models the rest of the timing quantities around it. Sensors require some time to upload pixel information between integrations, which is modeled as a constant dead time, t_d, seen in Figure 4.1. Then the kth

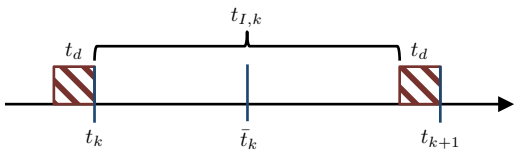

Figure 4.1: Various variables associated with time during a series of integrations.

integration is defined as

$$t_{I,k} = t_{k+1} - t_k - t_d \tag{4.1}$$

Note that in this formulation the first element t_1 is always zero and a final t_{N_t+1} must also be included. Even though there are $N_t + 1$ elements, the first and last are always fixed so there are only $N_t - 1$ variables to optimize over. As mentioned, the problem involves a variable length vector in that N_t can be varied and whatever optimization method used must be able to handle this. Finally, the center time of each integration is useful to define for future discussions as the midpoint of a particular integration,

$$\bar{t}_k = (t_k + t_{k+1} - t_d)/2 \tag{4.2}$$

which can also be seen in Figure 4.1.

Next, two continuous variables, $\alpha(t)$ and $\delta(t)$ are defined. This chapter treats these variables as the inertial angles, right ascension and declination, but they could just as easily represent a unit vector or local frame (e.g. azimuth and elevation). These variables need to be defined over the exposure time taken by a sensor. In many cases, this leads to defining four parameters $[\alpha(t_0), \delta(t_0), \dot{\alpha}(t_0), \dot{\delta}(t_0)]$. This can be further simplified by moving the variables into the image plane as $[x(t_0), y(t_0), \dot{x}(t_0), \dot{y}(t_0)]$. Finally, it is worth noting that $x(t_0)$ and $y(t_0)$ are mostly immaterial as design variable because they only determine where in the image an object shows up; they do not affect information on the brightness or shape of the object. This chapter assumes the object is in the field of view and there is insignificant distortion in the sensor to affect information, which may not be true for very

wide fields of view. Better rate tracking of an object concentrates more light into fewer pixels, increasing probability of detection and photometry information. Astrometry, however, can be improved in certain situations by imperfect rate tracking, depending on the exact mapping between the point spread function and pixel space.

Next, the gain map must be discussed. This consists of an overall gain, G, and a gamma filter, Γ. These effect the mapping of electrons on a CCD to pixel values. If chosen incorrectly pixel saturation occurs distorting and biasing photometry information. However, on the other end, very dim signals can have very poor discretization.

4.2.2 Image Statistics

This section reviews statistical models for an image that are used throughout this chapter. An image is made up of a series of pixel measurements broken into deterministic signal content and zero mean noise as seen in Equation (2.5) where i is an index denoting a particular pixel.

A good way to model the statistics of a pixel is with a series of Poisson distributions [78]. The pixel is composed of signal (shot noise), dark noise, and read noise which are all Poisson distributed. It is convenient to factor out the mean of the shot noise as signal, s_i, and subtract out the mean of the remaining noise, which gives Equation (2.7), However, this chapter deals with a more simplified model. A Poisson distribution approaches a Gaussian distribution as λ increases [88]. Furthermore, for bright objects the contribution of shot noise on total system randomness is often ignored. The final formulation of a pixel, which is used throughout this chapter, is then Equation (2.10) Note that the Fisher information for Poisson distributed pixels is derived in the same way as that of Gaussian pixels [40]. The work presented in this chapter can therefore by generalized to explicitly Poisson noise as well.

It is important to understand that measurements taken over the kth integration time are dependent on the length of the integration time along with a series of other factors. The

expected value of the noise, which is equal to the variance due to the underlying Poisson distribution, is

$$\lambda_{w,i} = \sigma_i^2 = GQ\delta^2 F_o t_{I,k} = BF_o t_{I,k} \tag{4.3}$$

where Q is the quantum efficiency, δ^2 is the area of each pixel, and F_o is the background and dark current flux rate. In the final line, the constants are wrapped into the parameter $B = \delta^2 GQ$ to ease notation. This implies that the design variable vector $\mathbf{t}$ affects the information content directly through changing the background variance.

4.3 Reward Functions for SSA

4.3.1 Maximum Likelihood Estimator

For the purposes of this chapter, an estimator needs to be constructed for the flux, position and velocity of a space object in the image frame at some time $\theta = (F, x^*(t^*), y^*(t^*), \dot{x}^*(t^*), \dot{y}^*(t^*))$. Note that the position and velocity shown above are treated as the location and velocity of the center (mean) of a point light source in the image plane. A point source object spreads over pixels through a spatial PDF, $h(\mathbf{x}_{i,k}|t)$, called a point spread function (PSF), where $\mathbf{x}_{i,k} = [x_{i,k}, y_{i,k}]^T$ is a specific pixel coordinate. Under ideal conditions the PSF is an Airy disc [36], but is often approximated as Gaussian for analysis purposes. In reality, this function is disturbed by atmospheric distortion and mount jitter as well, further distributing the light. In practical applications, the PSF is estimated directly from image data, either as an average over several frames or directly in a frame-to-frame manner. The PSF spatial

distribution is then

$$h(\mathbf{x}_{i,k}|t_k) = \frac{1}{\sqrt{(2\pi w^2)^2}} e^{-(\xi_{i,k}^2+\eta_{i,k}^2)/2w^2} \tag{4.4}$$

$$\xi_{i,k} = (x_{i,k} - x_k^*),\ \eta_{i,k} = (y_{i,k} - y_k^*) \tag{4.5}$$

$$x_k^* = x^*(t^*) + \dot{x}^*(t^*)(t_k - t^*) \tag{4.6}$$

$$y_k^* = y^*(t^*) + \dot{y}^*(t^*)(t_k - t^*) \tag{4.7}$$

where w^2 is the variance of the point spread function. Note that the location of the point source in the image is defined through a series of constants to be estimated, $x^*(t^*)$, $y^*(t^*)$, $\dot{x}^*(t^*)$, $\dot{y}^*(t^*)$, and the time t_k. This final PSF $h(\mathbf{x})$ describes the spatial probability density of photons hitting the ccd plane.

The photons coming off of an object, entering the aperture of a telescope, and hitting the CCD can be described as a Poisson process and distribution, $P(\lambda)$. Each photon goes through a series of random events, including hitting a random place on the CCD (PSF, $h(\mathbf{x})$), having a random chance at not being detected (Quantum Efficiency, Q), and being converted in batches to pixel counts (Gain, G). Each pixel can be then be described as a Poisson distribution, and approximated as Gaussian as in Equation (2.10). The expected number of pixel counts (parameter of the Poisson distribution) for pixel i and timestep k is

$$\begin{aligned} s_{i,k} = \lambda_{i,k} &= G\,Q\,F\,t_{I,k} \int_{\text{pixel}} h(\mathbf{x})\mathrm{d}\mathbf{x} \\ &\approx G\,Q\,F\,t_{I,k}\,\delta^2 h(\mathbf{x}_{i,k}) \\ &= \frac{B t_{I,k} F}{\sqrt{(2\pi w^2)^2}}\, e^{-(\xi_{i,k}^2+\eta_{i,k}^2)/2w^2} \end{aligned} \tag{4.8}$$

where F is the (average) photon rate off of the object entering the aperture in photons per second and the integral calculates the probability a given photon hits pixel i. In reality, the PSF should be integrated over the pixel, but the approximation in line 2 is a common

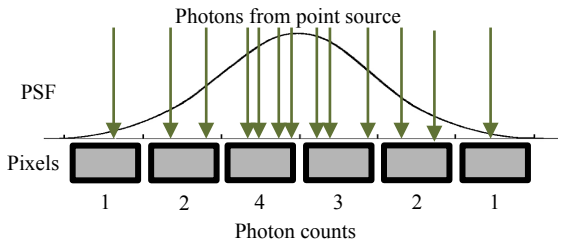

Figure 4.2: Photons from a point source are randomly distributed across pixels according to a point spread function. Number of photons is determined randomly via a Poisson distribution.

simplification [40] and effectively a linear approximation of the integral.

Each pixel also has further corruptions from dark current and read noise which can both be modeled as Poisson and approximated as zero mean Gaussian. Each pixel has a distribution of

$$p(z) = \frac{1}{\sqrt{2\pi\sigma_{i,k}^2}} e^{-(z-s_{i,k})^2/2\sigma_{i,k}^2} \tag{4.9}$$

where $s_{i,k}$ is the expected signal value. Note that it is common to approximate this distribution as Gaussian, so long as the expected value is high enough. The entire process can be seen, nominally, in Figure 4.2.

Precise astrometry can be calculated in a variety of ways, but this chapter looks at the commonly used maximum likelihood estimator [40]. The astrometry can be evaluated through how well the covariance matrix of the measurement is minimized. Similarly, this can be seen as maximizing the Fisher information in the estimator [89]. The derivation of the Fisher information starts with the log likelihood function, $l(\theta) = \log(L(\theta|\mathbf{z})) = \log(p(\mathbf{z}|\theta))$, which for this example is

$$l(\theta) = \sum_{k=1}^{N_t} \sum_{i=1}^{N_k} \log\left(\frac{1}{\sqrt{2\pi\sigma_{i,k}^2}}\right) + \left(\frac{-(z_{i,k} - s_{i,k})^2}{2\sigma_{i,k}^2}\right) \tag{4.10}$$

Note that the summation over k and i are combined and dropped in future equations for

ease of notation, but in general all summations occur over both k and i. Also note that N_k is a function of the time step k. Then the first derivative is generally

$$\frac{\partial l(\theta)}{\partial \theta_j} = \sum \frac{1}{\sigma_{i,k}^2}(z_{i,k} - s_{i,k})\frac{\partial s_{i,k}}{\partial \theta_j} \tag{4.11}$$

The Fisher information is found by taking the expected value of the square of the gradient above. Future sections calculate Fisher information for a variety of values of θ_j.

4.3.2 Streaking Object Point Spread Function

This section quantifies the modified point spread function, $\tilde{h}(\mathbf{x}_i)$, for a streaking object. The goal is to modify the PSF in Equation (4.8) for a streaking object. The primary difference is explicitly integrating over the integration time rather than just including $t_{I,k}$,

$$\begin{aligned} t_{I,k}\tilde{h}(\mathbf{x}_{i,k}|t) &= \int_{t_k}^{t_{k+1}-t_d} h(\mathbf{x}_{i,k}|t)\mathrm{d}t \\ &= \int_{t_k}^{t_{k+1}-t_d} \frac{1}{\sqrt{(2\pi w^2)^2}} e^{-(\xi_{i,k}^2(t)+\eta_{i,k}^2(t))/2w^2}\mathrm{d}t \end{aligned} \tag{4.12}$$

where $\xi_{i,k}$ and $\eta_{i,k}$ are now time dependent implying h must be explicitly integrated. Note that $t_{I,k}$ must be included as part of the definition, because it is calculated through the integral.

There is a useful shortcut in the evaluation of the integral in Equation (4.12). In order ease this calculation, this integral can be reparameterized into the distance parallel and orthogonal to the motion of the Gaussian, $[\dot{x}^*, \dot{y}^*]^T$, which are internal parameters in h. Note that while the parallel distance, $\tilde{\xi}$, is time varying the orthogonal distance, $\tilde{\eta}$, is time

invariant. Then the same integral is

$$t_{I,k}\tilde{h}(\mathbf{x}_{i,k}|t) = \frac{e^{-\tilde{\eta}_{i,k}^2/2w^2}}{\sqrt{(2\pi w^2)^2}} \int_{t_{I,k}}^{t_{I,k+1}} e^{-\tilde{\xi}_{i,k}^2(\tilde{x}(t))/2w^2} \mathrm{d}t \tag{4.13}$$

where one direction of the Gaussian is removed from the integral and the other is calculated from a simple 1-D Gaussian integral evaluation.

4.3.3 Astrometry Reward for Space Object Orbit Determination

In order to synthesize a reward function for tasking parameters with respect to astrometry, the Fisher information is calculated in this section. Astrometry for space objects in this case is $\theta = (x^*(t^*), y^*(t^*), \dot{x}^*(t^*), \dot{y}^*(t^*))$. The derivation starts with evaluating Equation (4.11) for the above θ, which leads to

$$\begin{aligned}
\frac{\partial l}{\partial x^*} &= \frac{1}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \xi_{i,k} s_{i,k} \\
\frac{\partial l}{\partial y^*} &= \frac{1}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \eta_{i,k} s_{i,k} \\
\frac{\partial l}{\partial \dot{x}^*} &= \frac{1}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \xi_{i,k} s_{i,k} (\bar{t}_k - t^*) \\
\frac{\partial l}{\partial \dot{y}^*} &= \frac{1}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \eta_{i,k} s_{i,k} (\bar{t}_k - t^*)
\end{aligned} \tag{4.14}$$

The Fisher information is found by taking the expected value of the square of the gradient above. First note that in this form, the only random variable is the pixel measurements, $z_{i,k}$, which appears once. While the square of Equation (4.14) is the square of a sum, $\mathbb{E}[(z_{i,k} - s_{i,k})(z_{j,k} - s_{j,k})] = 0$ for $i \neq j$ so long as noise is uncorrelated. For the rest of the

terms, $\mathbb{E}[(z_{i,k} - s_{i,k})^2] = \sigma_{i,k}^2$, giving Fisher information results of

$$I_{xx} = \frac{1}{w^4} \sum \frac{\xi_{i,k}^2 s_{i,k}^2}{\sigma_{i,k}^2} \tag{4.15}$$

$$I_{yy} = \frac{1}{w^4} \sum \frac{\eta_{i,k}^2 s_{i,k}^2}{\sigma_{i,k}^2}$$

$$I_{\dot{x}\dot{x}} = \frac{1}{w^4} \sum \frac{\xi_{i,k}^2 s_{i,k}^2}{\sigma_{i,k}^2} (\bar{t}_k - t^*)^2 \tag{4.16}$$

$$I_{\dot{y}\dot{y}} = \frac{1}{w^4} \sum \frac{\eta_{i,k}^2 s_{i,k}^2}{\sigma_{i,k}^2} (\bar{t}_k - t^*)^2$$

Next, the same calculation is made for a moving point spread function. The moving point spread function in Equation (4.12) can be plugged directly into the second line of Equation (4.8). Referencing Equation (4.11), the same derivative is calculated which for position and velocity information is unaffected by the presence of an integral over time.

$$\frac{\partial l}{\partial x^*} = \frac{BF}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \int_{t_k}^{t_{k+1}-t_d} \xi_{i,k} h_{i,k} \mathrm{d}t \tag{4.17}$$

$$\frac{\partial l}{\partial \dot{x}^*} = \frac{BF}{w^2} \sum \frac{(z_{i,k} - s_{i,k})}{\sigma_{i,k}^2} \int_{t_k}^{t_{k+1}-t_d} \xi_{i,k} h_{i,k} (t - t^*) \mathrm{d}t \tag{4.18}$$

which turns into the position and velocity information

$$I_{xx} = \frac{B^2 F^2}{w^4} \sum \frac{1}{\sigma_{i,k}^2} \left(\int_{t_k}^{t_{k+1}-t_d} \xi_{i,k} h_{i,k} \mathrm{d}t \right)^2 \tag{4.19}$$

$$I_{\dot{x}\dot{x}} = \frac{B^2 F^2}{w^4} \sum \frac{1}{\sigma_{i,k}^2} \left(\int_{t_k}^{t_{k+1}-t_d} \xi_{i,k} h_{i,k} (t - t^*) \mathrm{d}t \right)^2 \tag{4.20}$$

In effect, the streak information only requires calculating the modified expected value of each pixel. These integrals can be easily calculated for each pixel through a Riemann approximation.

4.3.4 Photometry Reward for Space Object Characterization

First, the information on the estimate of the photometric flux, that is, how accurately the flux is estimated in an image, is shown. Starting with Equation (4.11),

$$\frac{\partial l(F)}{\partial F} = \sum \frac{1}{\sigma_{i,k}^2}(z_{i,k} - s_{i,k})Bt_{I,k}h(\mathbf{x}_{i,k}) \tag{4.21}$$

where the F in the denominator cancels with the F in $s_{i,k}$. To get the information, the expectation of the square of Equation (4.21) is required,

$$I_{FF} = B^2 \sum \frac{t_{I,k}^2 h(\mathbf{x}_{i,k})^2}{\sigma_{i,k}^2} \tag{4.22}$$

Photometry Bias

For a time varying photometric flux signature, the point spread function must be integrated over time.

$$\begin{aligned} s_{i,k} &= G \int_{t_k}^{t_{k+1}-t_d} F(t) \int_{\text{pixel}} h(\mathbf{x})\mathrm{d}\mathbf{x}\mathrm{d}t \\ &\approx Bt_{I,k}\bar{F}h(\mathbf{x}_{i,k}) \end{aligned} \tag{4.23}$$

The flux is only observed through $\bar{F}$, the average flux over the integration time. Because average flux over a time period is a difficult measurement to incorporate into an estimation scheme, the measurement is often treated as the instantaneous flux at the central time, $\bar{t}_k$. For flux that varies linearly over the integration time, this assumption is accurate, but many real scenarios can contain significantly varying signals. Furthermore, the variations in flux are often an important piece of information in object characterization. These variations provide an inherent bias, ϵ_F, in the photometric measurements, which is illustrated in Figure 4.3.

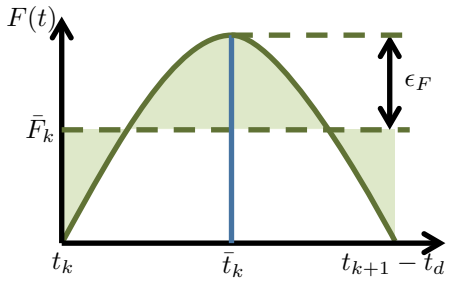

Figure 4.3: An inherent bias, ϵ_F, is created when observing a fluctuating signal.

The expected value of the pixels are then

$$s_{i,k} = Bt_{I,k}(F(\bar{t}_k) + \epsilon_F)h(\mathbf{x}_{i,k}) \tag{4.24}$$

where $\bar{F}$ is split into the true term to be estimated and a bias error term, ϵ_F. The bias is dependent on the shape of the signal itself, thus to make further progress a model for the flux is proposed. Typically, the kind of quick flux changes which occur during an integration are due to periodic variations on flux. The flux is modeled as a sum of sinusoidal signals. It is also assumed that if the highest frequency signal is sampled sufficiently, all lower frequency signals are also sufficiently sampled. Therefore, the goal of photometry is to sufficiently sample

$$F(t) = F_0 + A\sin(\omega t + \theta) \tag{4.25}$$

where F_0 is the average flux, θ is the phase shift, and A and ω are the amplitude and frequency of the highest frequency time varying component. The bias is then

$$\begin{aligned}\epsilon_F &= F(\bar{t}_k) - \frac{1}{t_{I,k}} \int_{t_k}^{t_{k+1}-t_d} F(t)\mathrm{d}t \\ &= A\sin(\omega\bar{t}_k + \theta) - \frac{1}{t_{I,k}} \int_{t_k}^{t_{k+1}-t_d} A\sin(\omega t + \theta)\mathrm{d}t\end{aligned} \tag{4.26}$$

For further analysis, it would be useful to quantify the distribution of ϵ_F. The unknown phase shift, θ, is uniformly distributed over $\mathbb{R}$ but, due to the cyclical nature of the sine function, can be assumed as uniformly distributed over $[0, 2\pi]$. For the timing $\bar{t}_k$, t_k, and t_{k+1} can be assumed as any given choice simply shifts $\omega\bar{t}_k + \theta$, which is still uniform over one cycle of sine. Applying these assumptions and a change in variable, $\tau = \omega t$, gives

$$\begin{aligned}\epsilon_F &= A\left(\sin(\theta) - \frac{1}{t_{I,k}}\int_{-t_{I,k}/2}^{t_{I,k}/2}\sin(\omega t + \theta)\mathrm{d}t\right)\\ &= A\left(\sin(\theta) - \frac{1}{\omega t_{I,k}}\int_{-\omega t_{I,k}/2}^{\omega t_{I,k}/2}\sin(\tau + \theta)\mathrm{d}\tau\right)\end{aligned} \tag{4.27}$$

where $\omega t_{I,k}$ forms a time factor. As $\omega t_{I,k} \to 0$, the right side of the equation becomes a very accurate estimate of $\sin(\theta)$ which drives $\epsilon_F \to 0$. As $\omega t_{I,k} \to \infty$, the integral becomes cyclical between -1 and +1, and so $\epsilon_F \to \sin(\theta)$. This integral can be evaluated directly, giving

$$\begin{aligned}\epsilon_F =& A\Bigg(\sin(\theta) + \frac{1}{\omega t_{I,k}}\cos(\theta + \omega t_{I,k}/2)\\ &- \frac{1}{\omega t_{I,k}}\cos(\theta - \omega t_{I,k}/2)\Bigg)\end{aligned} \tag{4.28}$$

which can be simplified with the half angle formula into

$$\begin{aligned}\epsilon_F &= A\sin(\theta)\left(1 - \frac{2}{\omega t_{I,k}}\sin(\omega t_{I,k}/2)\right)\\ &= \alpha\sin(\theta)\end{aligned} \tag{4.29}$$

where α is the modified amplitude

$$\alpha = A\left(1 - \frac{2}{\omega t_{I,k}}\sin(\omega t_{I,k}/2)\right) \tag{4.30}$$

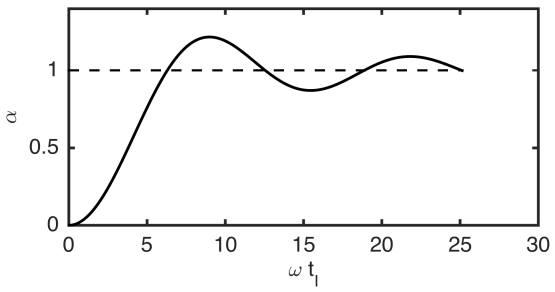

Figure 4.4: Variation of α with respect to the time constant $\omega t_{I,k}$ with constant $A = 1$.

Note that the phase still heavily affects the bias through the $\sin(\theta)$ and the results as $\omega t_{I,k}$ goes to zero or infinity still hold. Finally, if θ is uniformly distributed,

$$p_{\text{asin}}(\epsilon_F; \alpha) = \frac{1}{\pi\sqrt{\alpha^2 - \epsilon_F^2}} \tag{4.31}$$

giving a PDF for the flux bias. This distribution is the symmetric case of the arcsine distibution. The expected value is zero which intuitively makes sense, and the variance is

$$\begin{aligned} \text{Var}(\epsilon_F) &= \int \epsilon_F^2 \frac{1}{\pi\sqrt{\alpha^2 - \epsilon_F^2}} \text{d}\epsilon_F \\ &= \alpha^2/2 \end{aligned} \tag{4.32}$$

which can be analytically calculated.

High Frequency Photometry

Now the photometric information can be calculated with the bias, ϵ_F, taken into account. Because the bias is inherently stochastic and unknown, the best available estimate of the expected value in each pixel is still Equation (4.8), where F is the instantaneous flux.

However, each pixel has an added stochastic term

$$\begin{aligned} z_{i,k} &= B(F+\epsilon_F)t_{I,k}h(\mathbf{x}_{i,k}) + w_{i,k} \\ &= s_{i,k} + B\epsilon_F t_{I,k}h(\mathbf{x}_{i,k}) + w_{i,k} \end{aligned} \tag{4.33}$$

$$\epsilon_F \sim p_{\text{asin}}(\alpha), w_{i,k} \sim \mathcal{N}(0, \sigma_{i,k}^2) \tag{4.34}$$

$$\text{Var}(z_{i,k}) = \sigma_{i,k}^2 + (B^2 t_{I,k}^2 \alpha^2 h_{i,k}^2)/2 \tag{4.35}$$

In order to incorporate this into the Fisher information, Equation (4.11) is evaluated for the new distribution,

$$\begin{aligned} \frac{\partial l(F)}{\partial F} &= B\sum \frac{(z_{i,k}-s_{i,k})}{\text{Var}(z_{i,k})} t_{I,k}h_{i,k} \\ \frac{\partial l(F)}{\partial F} &= B\sum \frac{(w_{i,k}+B\epsilon_{F,k}h_{i,k})}{\text{Var}(z_{i,k})} t_{I,k}h_{i,k} \end{aligned} \tag{4.36}$$

noting that $\epsilon_{F,k}$ is invariant over index i. To get information, the expectation of the square of the above derivative is required, but unlike before, cross terms are not necessarily uncorrelated,

$$I_{FF} = B^2\mathbb{E}\left[\left(\sum_k\sum_i (w_{i,k}+B\epsilon_{F,k}h_{i,k})\frac{t_{I,k}h_{i,k}}{\text{Var}(z_{i,k})}\right)^2\right] \tag{4.37}$$

where $w_{i,k}$ and $\epsilon_{F,k}$ are the only stochastic terms. Note that while the various $w_{i,k}$ are uncorrelated with each other and $w_{i,k}$ is uncorrelated with $\epsilon_{F,k}$, there are $\epsilon_{F,k}\epsilon_{F,l}$ which are correlated based on their time offset by the term $\cos(\omega(\bar{t}_k - \bar{t}_l))$. In reality this creates a quadruple sum reward function over two time step dimensions and two pixel dimensions. Furthermore, the correlation can be both positive and negative and is based on an initially unknown frequency. For the simulation work, these correlation terms are not included for

this reward function giving a simplified information of

$$I_{FF} = B^2 \sum_{i,k} \frac{t_{I,k}^2 h_{i,k}^2}{\text{Var}(z_{i,k})} \tag{4.38}$$

This version of the reward inflates the variance of each pixel with the added uncertainty of ϵ_F.

4.3.5 Streaking Photometry

The streaking point spread function is straight forward to calculate. The primary difference is in substituting $\tilde{h}$ in Equation (4.12) for h,

$$I_{FF} = B^2 \sum_{i,k} \frac{\left(\int_{t_k}^{t_{k+1}-t_d} h_i(t) \mathrm{d}t \right)^2}{\text{Var}(z_{i,k})} \tag{4.39}$$

Note that $\tilde{h}$ also must be substituted into the variance equation.

4.3.6 Timing Uncertainty

When considering real optical sensors, there are a variety of other parameters, such as timing uncertainty, which may have inherent uncertainty and therefore affect the information content. This section serves as a template on how to incorporate additional uncertainties into the above reward functions. Additional uncertainty only needs to be introduced into the stochastic measurement variables, $z_{i,k}$, which typically takes the form of inflating the final variance, $\sigma_{i,k}^2$. In the case of timing uncertainty, zero mean Gaussian timing error is added at each time step,

$$\tilde{t}_k = t_k + \epsilon_{t,k}, \qquad \epsilon_{t,k} \sim \mathcal{N}(0, w_t^2) \tag{4.40}$$

where w_t is the known timing variance. Equations (4.15) and (4.16) are valid then, because the timing uncertainty does not change the expected value of the pixels, only the variance.

The pixel variance needs to be updated for the new value of $\sigma^2_{i,k}$ which incorporates timing uncertainty. The primary effect of timing uncertainty occurs through changing the integration time. Using the $\tilde{\cdot}$ notation to represent variables which have an internal timing uncertainty,

$$\begin{aligned}\tilde{t}_{I,k} &= t_{k+1} + \epsilon_{t,k+1} - t_k - \epsilon_{t,k} - t_d \\ &= t_{I,k} + \sqrt{2}\epsilon_{t,k}\end{aligned} \tag{4.41}$$

which gives

$$\begin{aligned}z_{i,k} &= \tilde{s}_{i,k} + w_{i,k} \\ \tilde{s}_{i,k} &= BFh(\mathbf{x}_{i,k})\tilde{t}_{I,k} \\ &= s_{i,k} + \sqrt{2}BFh(\mathbf{x}_{i,k})\epsilon_t \\ z_{i,k} &= s_{i,k} + w_{i,k} + \sqrt{2}BFh(\mathbf{x}_{i,k})\epsilon_t.\end{aligned} \tag{4.42}$$

For a streaking object, Equation (4.12) must be used as the basis for describing pixels, so the uncertainty is part of the integral bounds. Considering uncertainty in only t_k for ease on derivation, the integral can be broken into the original component and a small stochastic component,

$$\begin{aligned}z_{i,k} &= \tilde{s}_{i,k} + w_{i,k} \\ \tilde{s}_{i,k} &= BF \int_{t_k+\epsilon_{t,k}}^{t_{k+1}-t_d} h_{i,k}(t)\mathrm{d}t \\ &= s_{i,k} + BF \int_{t_k+\epsilon_{t,k}}^{t_k} h_{i,k}(t)\mathrm{d}t.\end{aligned} \tag{4.43}$$

If $\epsilon_{t,k}$ is small, which typically is the case, the integral can be linearized around t_k

$$\tilde{s}_{i,k} \approx s_{i,k} + BFh_{i,k}(t_k)\epsilon_{t,k} \tag{4.44}$$

where $h_{i,k}(t_k)$ is the point spread function at the time t_k. This linearization also depends on a small in plane velocity; the assumption is that the point spread function is effectively stationary over the range of the timing error. Taking into account uncertainty at t_{k+1} as well gives

$$\tilde{s}_{i,k} \approx s_{i,k} + BF(h_{i,k}(t_k)\epsilon_{t,k} + h_{i,k}(t_{k+1})\epsilon_{t,k+1}). \tag{4.45}$$

First note that the timing uncertainty affects pixels proportionally to the point spread function at that time, $h_{i,k}(t)$. This implies that only pixels near the end points are affected by timing uncertainty.

The timing uncertainty can be incorporated into any of the above reward functions by inflating the variance of each pixel by the additive Gaussian term in Equation (4.42) or Equation (4.45) for a stationary or moving PSF respectively. The timing can also affect the information in a second way: by changing exposure time, the mean background radiation is affected as well. This uncertainty is a gray area because effective background subtraction implies that any such variation in background is estimated and removed.

4.4 Constraining Parameters

This section outlines the constraints being used for this optimization problem. The standard notation for inequality constraints is as functions $g(x) \leq 0$.

4.4.1 Physical Constraints

The first constraint is total time allocated for images, T. The series of N_t images each take their respective integration times to be captured. The camera must upload the pixel

information between integrations. This dead time, t_d, occurs $N_t - 1$ times and is a function of the image size and the data transfer rate of the camera. Because of the structure of $\mathbf{t}_I$, this can be implemented as a constraint on the last element of the vector, t_{I,N_t+1}, and is effectively an equality constraint that must be zero.

$$g_t(\mathbf{t}_I) = (N_t - 1)t_d + \sum_{k=1}^{N_t} t_{I,k} - T \tag{4.46}$$

$$g_t(t_{I,N_t+1}) = t_{I,N_t+1} - T \tag{4.47}$$

The next constraint is based on pixel saturation and bleeding. When a CCD integration is sufficiently high such that a pixel reaches its maximum number of counts, or FWC (full well capacity), the pixel begins to "bleed" its counts into adjacent pixels [90]. Pixel bleeding adds extra signal over vast areas of an image, with poorly understood statistical distributions, rendering those areas unusable. This chapter treats bleeding as unacceptable; even if bleeding in an image does not affect the target being observed, secondary objects may move through an image but be obscured by bleeding. Bleeding is dependent on the number of photons hitting a pixel and is therefore unaffected by gain. It is also dependent on the flux of the brightest star in the image, F_{max}, which must therefore be predicted or estimated. This leads to

$$g_b(t_{I,k}) = \frac{F_{\text{max}}\, t_{I,k}\, \delta^2}{\sqrt{(2\pi w^2)^2}} - \text{FWC} \tag{4.48}$$

where $\sqrt{(2\pi w^2)^2}$ is the PSF's maximum value (center). Pixel saturation can occur in the same way, by reaching FWC for a pixel, but this is uncommon for space objects relevant to SSA. Instead, the gain for a camera can be set too high, so when the analog-to-digital conversion occurs, the pixel's bit depth, BD, is reached. This saturation makes photometry impossible and can have unpredictable affects on astrometry, but does not lead to bleeding.

Similar to Equation (4.48), the constraint to avoid saturating relevant pixels is

$$g_s(t_{I,k}, G) = \frac{G\,QE\,F\,t_{I,k}\,\delta^2}{\sqrt{(2\pi w^2)^2}} - \text{BD} \tag{4.49}$$

where F is the brightness of either the space object being tracked or the brightest possible space object. Note that the negative effect of saturation on astrometry may not be sufficient to always require this constraint, but that is case dependent.

4.4.2 Star Field Clutter

In many cases of SSA missions, the intuitive tasking involves rate tracking a space object. Intuitively, the reward functions in this chapter often prioritize longer exposure times. The reason exposure time is often not maximized is due to star field clutter. Rate tracking an object causes stars to streak through an image, and when many stars are present this can clutter a large percentage of the image. Methods exist to subtract out the stars and preserve data on other light sources [91] but because star photons are a Poisson process, they still lead to large spikes in uncertainty, especially in photometry. Furthermore, SSA missions often detect previously unknown objects as uncorrelated tracks (UCT); with high levels of star clutter this becomes difficult.

This constraint is formulated as a maximum acceptable percentage of pixels obfuscated by stars, $\%_{\text{max}}$.

$$l = \frac{t_{I,k}\dot{\theta}}{\text{iFoV}} \tag{4.50}$$

$$N_{\text{star}} = \#_{\text{star}}(9\pi w^2 + 6lw) \tag{4.51}$$

$$g_{\%}(t_{I,k}, \dot{\theta}) = \frac{N_{\text{px}}}{N_{\text{total}}} - \%_{\text{max}} \tag{4.52}$$

where l is the length of star streaks in the image, $\dot{\theta}$ is the tracking rate, $\#_{\text{star}}$ is the number of star in the field of view, w is the standard deviation of the point spread function, and

N_{total} is the total number of pixels in the image.

4.4.3 Probability of Detection

An important consideration is the probability of detection of an object. This idea can be used as two possible constraints, one on detecting a specific object of interest and one on detecting various objects of opportunity in the background. Object of opportunity refers to an object which may unexpectedly pass through the field of view during a data campaign. SSA missions often desire that as many objects as possible be detected and characterized to increase overall situational awareness. Mathematically, this objective is best quantified as the total photometric evidence for an object, while considering the total corruption of that evidence by noise. This is then measured through total photometric SNR of an object in a particular image

$$\text{SNR}\left(\sum z_i\right) = \frac{\mathbb{E}\left[\sum z_i\right]}{\sqrt{\mathbb{E}\left[\left(\left[\sum z_i\right] - \mathbb{E}\left[\sum z_i\right]\right)^2\right]}} \tag{4.53}$$

$$\approx \frac{GQ\bar{F}t_I}{\sqrt{N_k}\sigma} \tag{4.54}$$

which encapsulates in one number the ratio of information to corruption. Note that $GQ\bar{F}t_I$ is the sum of Equation (4.8) over all pixels for a given time step and N_k is the number of pixels over which the signal is spread. The number of pixels N_k is dependent on the size of the point spread function and the length over which the object streaks. For tracking a known object, this number is already calculated for the summations in the information equations while for objects of opportunity this is inherently unknown. For objects of opportunity a nominal value must be chosen for N_k. This gives an equality constraint of

$$g_{\text{SNR}}(t_{I,k}) = \text{SNR}_{\text{min}} - \text{SNR}\left(\sum z_i\right) \tag{4.55}$$

where SNR_{min} is some minimum value of SNR, dependent on the detection method being used. An aggressive minimum SNR would be 3, which pushes the limits of current detection methods [88].

4.4.4 Side Constraints

Many of the design variables have side constraints which may be considered.

$$t_{I,\text{min}} \leq t_{I,k} \leq t_{I,\text{max}} \tag{4.56}$$

$$g_{\text{FoR}}(\alpha, \delta) \leq 0 \tag{4.57}$$

$$g_{\text{slew}}(\dot{\alpha}, \dot{\delta}) \leq 0 \tag{4.58}$$

In Equation (4.56), integration time is usually bounded by minimum and maximum values which are dependent on the specific camera itself. In Equation (4.57), observatories sometimes have physical constraints on what azimuth and elevation angle pairs are achievable, which again can be used on a case by case basis. In Equation (4.58), the telescope mount may require a constraint on slew rate, which is dependent on the mount in question. The constraint must be tailored to each sensor mount, and may not be necessary for campaigns involving slow moving objects.

4.5 Optimization Techniques

This section discusses how to calculate gradients on the various reward functions, how to treat some of the design variables, and what optimization methods are used in this work.

4.5.1 Gradient on Stationary Point Source

This section develops gradients for the reward functions with respect to the timing vector, $\mathbf{t}_I$. The gradients here are only shown for the simpler stationary point spread function, while the more complicated moving point spread function gradients are shown in the next

section. In order to efficiently search for optima, it is useful to have analytic gradients of the reward functions with respect to the integration time vector. Before proceeding, note the derivatives of the integration time and center times from Equation (4.1) and Equation (4.2) for the exposures

$$\frac{\partial t_{I,k}}{\partial t_k} = -1, \qquad \frac{\partial t_{I,k}}{\partial t_{k-1}} = 1 \tag{4.59}$$

$$\frac{\partial \bar{t}_k}{\partial t_k} = 1/2, \qquad \frac{\partial \bar{t}_k}{\partial t_{k-1}} = 1/2 \tag{4.60}$$

For spatial information, Equation (4.15) must be expanded using Equations (4.8) and (2.11),

$$\frac{\partial I_{xx}}{\partial t_k} = \frac{\partial}{\partial t_k}\left(\frac{BF^2}{F_o w^4}\sum_k t_{I,k}\sum_i \xi_{i,k}^2 h_{i,k}^2\right) \tag{4.61}$$

where $t_{I,k}$ is defined in Equation (4.1). The differentiation results in

$$\frac{\partial I_{xx}}{\partial t_k} = \frac{BF^2}{F_o w^4}\left(\sum_i \xi_{i,k}^2 h_{i,k}^2 - \sum_i \xi_{i,k-1}^2 h_{i,k-1}^2\right). \tag{4.62}$$

This equation represents the trade-off between lengthening the $k-1$th integration while shortening the kth integration. Consequently, if the shape of the point spread function is constant from frame to frame, the two summations yield a derivative of zero. This implies for astrometric position information, when the exposure is taken does not matter, only the amount of exposure; if, say, 20 seconds of exposure will be taken this time can be spread between any number of measurements of any lengths and achieve equal information. In realistic circumstances, the presence of dead time implies that longer exposure times are better, as less time is wasted.

Similarly for the velocity astrometry information,

$$\begin{aligned}\frac{\partial I_{\dot{x}\dot{x}}}{\partial t_k} &= \frac{\partial}{\partial t_k}\frac{BF^2}{F_o w^4}\sum_k t_{I,k}(\bar{t}_k - t^*)^2 \sum_i \xi_{i,k}^2 h_{i,k}^2 \\ &= \frac{BF^2}{F_o w^4}\Big(t_{I,k}(\bar{t}_k - t^*) - (\bar{t}_k - t^*)^2\Big)\sum_i \xi_{i,k}^2 h_{i,k}^2 \\ -\Big(t_{I,k-1}(\bar{t}_{k-1} - t^*) &- (\bar{t}_{k-1} - t^*)^2\Big)\sum_i \xi_{i,k-1}^2 h_{i,k-1}^2 \end{aligned} \tag{4.63}$$

where $\bar{t}_k$ is defined in Equation (4.2). This equation contains an interesting interplay which provides qualitative insight on velocity astrometry. The $-(\bar{t}_k - t^*)^2$ term, which is always negative for the t_k integration and positive for the t_{k-1} integration, represents how increasing t_k reduces one exposure time while increasing another. Conversely, the $t_{I,k}(\bar{t}_k - t^*)$ term tends to push the t_k away from t^* to maximize the baseline. These two parts of the gradient oppose each other and the optimum lies somewhere in between.

The photometry information gradient is more complicated due to the various instances of $t_{I,k}$ in the denominator, but it is calculable. First note the expanded version of I_{FF},

$$\begin{aligned} I_{FF} &= B^2 \sum_k \sum_i \frac{t_{I,k}^2 h_{i,k}^2}{BF_o t_{I,k} + B^2 t_{I,k}^2 h_{i,k}^2 \alpha_k^2/2} \\ &= B\sum_k \sum_i \frac{t_{I,k} h_{i,k}^2}{F_o + B t_{I,k} h_{i,k}^2 \alpha_k^2/2} \end{aligned} \tag{4.64}$$

where α_k is dependent on $t_{I,k}$. The derivative of this expression with respect to t_k is then a function of the kth and $k-1$th time steps,

$$\begin{aligned}\frac{\partial I_{FF}}{\partial t_k} &= B^3 \sum_i \frac{-t_{I,k}^2 h_{i,k}^2}{\text{Var}(z_{i,k})^2}\left(F_o + B t_{I,k}^2 h_{i,k}^2 \alpha_k \frac{\partial \alpha_k}{\partial t_k}\right) \\ - B^3 \sum_i &\frac{-t_{I,k-1}^2 h_{i,k-1}^2}{\text{Var}(z_{i,k-1})^2}\left(F_o - B t_{I,k-1}^2 h_{i,k-1}^2 \alpha_{k-1} \frac{\partial \alpha_{k-1}}{\partial t_k}\right) \end{aligned} \tag{4.65}$$

where the derivative of α_k is

$$\begin{aligned}\frac{\partial \alpha_k}{\partial t_k} &= A\left(\frac{1}{t_{I,k}}\cos\left(\frac{\omega t_{I,k}}{2}\right) - \frac{2}{\omega t_{I,k}^2}\sin\left(\frac{\omega t_{I,k}}{2}\right)\right) \\ &= \frac{A}{t_{I,k}}\left(\cos\left(\frac{\omega t_{I,k}}{2}\right) - 1\right) + \frac{\alpha_k}{t_{I,k}}\end{aligned} \tag{4.66}$$

This gradient gives less intuition than the previous ones. In practice its performance depends on α as to whether it tries to drive the exposures to be small or large.

4.5.2 Gradient on Streaking Point Source

Position information on a streak object can be reformulated to give

$$\frac{\partial I_{xx}}{\partial t_k} = \frac{\partial}{\partial t_k}\left(\frac{BF^2}{F_o w^4}\sum_k \frac{1}{t_{I,k}}\sum_i C_{i,k}^2\right) \tag{4.67}$$

with

$$C_{i,k} = \int_{t_k}^{t_{k+1}-t_d} \xi_i(t)h_i(t)\mathrm{d}t \tag{4.68}$$

Included for notation ease. This leads to

$$\begin{aligned}\frac{\partial I_{xx}}{\partial t_k} =& \frac{BF^2}{F_o w^4}\left(\sum_i \frac{C_{i,k}^2}{t_{I,k}^2} - \frac{2C_{i,k}}{t_{I,k}}\xi_i(t_k)h_i(t_k)\right. \\ &\left. - \sum_i \frac{C_{i,k-1}^2}{t_{I,k-1}^2} - \frac{2C_{i,k-1}}{t_{I,k-1}}\xi_i(t_k - t_d)h_i(t_k - t_d)\right)\end{aligned} \tag{4.69}$$

where $t_{I,k}$ is defined in Equation (4.1). Similarly for velocity astrometry,

$$\begin{aligned} I_{\dot{x}\dot{x}} &= \frac{BF^2}{F_o w^4} \sum_k \frac{1}{t_{I,k}} \sum_i D_{i,k}^2 \\ \frac{\partial I_{\dot{x}\dot{x}}}{\partial t_k} &= \frac{BF^2}{F_o w^4} \Bigg(\sum_i \frac{D_{i,k}^2}{t_{I,k}^2} - \frac{2D_{i,k}}{t_{I,k}} d_i(t_k) \\ &\quad - \sum_i \frac{D_{i,k-1}^2}{t_{I,k-1}^2} - \frac{2D_{i,k-1}}{t_{I,k-1}} d_i(t_k - t_d) \end{aligned} \tag{4.70}$$

where

$$D_{i,k} = \int_{t_k}^{t_{k+1}-t_d} \xi_i(t) h_i(t)(t - t^*) \mathrm{d}t \tag{4.71}$$

$$d_i(t) = \xi_i(t) h_i(t)(t - t^*) \tag{4.72}$$

is another constant substituted for ease of notation. For both of these formulations, the information in the y direction is identical save for η replacing ξ.

For streaking photometry information

$$\begin{aligned} I_{FF} &= B^2 \sum_{i,k} \frac{\left(\int_{t_k}^{t_{k+1}-t_d} h_i(t) \mathrm{d}t \right)^2}{\mathrm{Var}(z_{i,k})} \\ &= B^2 \sum_{i,k} \frac{\left(\int_{t_k}^{t_{k+1}-t_d} h_i(t) \mathrm{d}t \right)^2}{BF_o t_{I,k} + B^2 \alpha_k^2 \int_{t_k}^{t_{k+1}-t_d} h_i(t) \mathrm{d}t} \end{aligned} \tag{4.73}$$

The derivative is then

$$\begin{aligned} \frac{\partial I_{FF}}{\partial t_k} &= B^2 \sum_i \Bigg(\frac{-2\mathrm{Var}(z_{i,k}) E_{i,k} h_i(t_k) + E_{i,k}^2 e_{i,k}(t_k)}{\mathrm{Var}(z_{i,k})^2} \\ &- B^2 \sum_i \frac{-2\mathrm{Var}(z_{i,k-1}) E_{i,k-1} h_i(t_k - t_d) + E_{i,k-1}^2 e_{i,k}(t_k - t_d)}{\mathrm{Var}(z_{i,k-1})^2} \end{aligned} \tag{4.74}$$

where the derivative of α_k is defined in Equation (4.66) and $E_{i,k}$ and $e_{i,k}$ are

$$E_{i,k} = \int_{t_k}^{t_{k+1}-t_d} h_i(t)\mathrm{d}t \tag{4.75}$$

$$e_{i,k}(t) = (BF_o + B^2\alpha_k^2 E_{i,k} h_i(t) - B^2\alpha_k E_{i,k}^2 \frac{\partial \alpha_k}{\partial t_k}) \tag{4.76}$$

Note that the (t) in $e_{i,k}$ refers to the fact that $h_i(t)$ is evaluated at a specific time, t_k for time step k and $t_k - t_d$ for times step $k-1$.

4.5.3 Karush Kuhn Tucker Conditions

Many optimization problems with analytic gradients can be solved analytically. In the general case allowing for constraints, the Karush-Kuhn-Tucker conditions are a series of equations which provide a framework for analytic solutions [92]. The basis of such an analysis involves setting the gradient of the reward function equal to zero and solving for the design variables. Such an analysis on Equation (4.62) gives

$$\sum_i \xi_{i,k}^2 h_{i,k}^2 = \sum_i \xi_{i,k-1}^2 h_{i,k-1}^2 \tag{4.77}$$

which is not a function of the design variables and in general always true. Assuming the above equality holds, the velocity information in Equation (4.63) gives

$$\begin{aligned} &t_{I,k}(\bar{t}_k - t^*) - (\bar{t}_k - t^*)^2 \\ &\quad = t_{I,k-1}(\bar{t}_{k-1} - t^*) - (\bar{t}_{k-1} - t^*)^2 \end{aligned} \tag{4.78}$$

which is an equation with three unknowns, t_{k-1}, t_k, and t_{k+1}. The full gradient gives one equation for every unknown and is therefore solvable, though non-linearity may provide multiple solutions. The photometry information gradient is much more complicated, due to multiple instances of $t_{I,k}$ in the gradient numerator and denominator. Finally, though the constraints discussed in Section 4.4 may not be active in many of the simulations, their

inclusion implies the need for a Lagrange multiplier analysis. Due to the complicated nature of these gradients and the ultimate desire for Pareto type analysis, this chapter does not spend any further time discussing analytic optimization solutions. The prospect of further analysis here marks a potent area for future work.

4.5.4 Optimization Methods

The primary optimization method used in this work is a descent method. Because analytic gradients exist, this is an efficient method which reliably solves the optimization problem. In practice, the optimization seems well posed with no obvious problems involving discrete variables, constraints, or local minima. These comments are purely observations and no explicit analysis of global optima has been done. Constraints are handled with an exterior penalty function which is easy to implement and has shown no problems in practice [92]. The time vector is of a variable length, which corresponds to different numbers of exposure times. This is currently solved by running the optimization for all possible lengths, N_t, and choosing the solution with the best reward. There may be ways to solve this more efficiently, which should be explored in future work.

The multi-objective problem is well posed and can be solved with a simple weighted sum method [92]. The actual Pareto surface is typically non-convex, but the Pareto surface for a fixed value of N_t appears convex in practice. The weighted sum method gives well posed Pareto surfaces for each value of N_t and the resulting series of Pareto surfaces are combined into a single final surface. This can be seen in the follow results section in Figures 4.16 and 4.18.

4.6 Optimal Solutions for Control of Telescope

4.6.1 Discussion and Analysis of Reward Functions

This section gives an intuitive explanation of the various reward functions developed in this chapter. The main design variables that affect the image information content are the

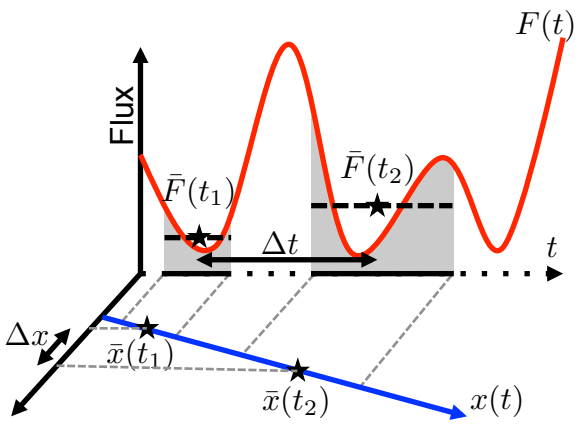

Figure 4.5: Astrometry gets more information from long exposure times, and a long time baseline. Photometry gets more information from short exposure times, to reduce uncertainty from time varying signals.

exposure times. Figure 4.5 illustrates the interplay of photometry and astrometry. Figure 4.5 shows two exposures centered at two times, t_1 and t_2, and how they appear with respect to both position and brightness. Exposure time in an image is generally proportional to information content, and so t_2 provides more raw position and brightness information. This is exactly true for position astrometry, but velocity astrometry is dependent on both Δx and Δt as well so separating two long exposures is also useful. Constant flux photometry would also want long exposures, but the curvature of $F(t)$ in Figure 4.5 shows why long exposures are a liability due to the potential bias.

In the next few subsections, the reward functions and subsequent Pareto surfaces are explored over a particular nominal test case. The test case has parameters defined in Table 4.1. Total time refers to the total length of the campaign and is equivalent to the final element of the exposure time vector, t_{N_t+1}. The primary design variable being explored is the exposure time vector. The slew variable is simplified as object speed (px/s) and object heading (degrees) in the pixel plane, though the object speed is set to zero for many of these simulations to maximize the intuition in the results. Gain, G, quantum efficiency, Q, and PSF size, w, play only a small role in information content and are set to be round values to minimize their impact. The flux, background noise, and amplitude are chosen as

Table 4.1: Various parameters and constants used for simulated observation campaign.

G	Q	F_o	δ	w
1.0	1.0	100. ph/(s px^2)	1.0 px	1.5 px
Total time	t_d	Mean Flux	ω	A
20 s	0.5 s	1000 ph/s	$\pi/4$ rad/s	10, 100 ph/s

nominal values. The relative values of A (sinusoidal amplitude) and F_o (background noise flux) tend to determine whether photometry prioritizes avoiding bias through ϵ_F versus raw number of photons. Two values of A are used in the simulation section so each case can be separately analyzed.

Before fully analyzing the reward functions, the aliasing error discussed in Section 4.3 is validated here. Measurement values are all generated by generating actual streaks (with no noise) and summing up the pixels. For the time span and parameters defined in Table 4.1, the different exposure times are shown in Figure 4.6. The exposures have three different lengths, one well above the Nyquist frequency, one which experiences high levels of aliasing, and one with only small errors.

The velocity cost function pushes exposure center time $\bar{t}_k$ away from the campaign center time, t^*. This results in a single exposure that occurs over the entire time span which necessitates $\bar{t}_k = t^*$ and zero information. Furthermore, while the parameter t^* is set as the center time of the campaign, as this is typically when velocity is estimated, it could be set to other values which can give very different answers. These examples highlight the limitations of the above analysis; the reward functions are the trace of a single part of the Fisher information matrix, don't account for nuisance parameters and don't account for every possible factor. The reward functions are surrogates to actual information and have a good deal of room for improvement.

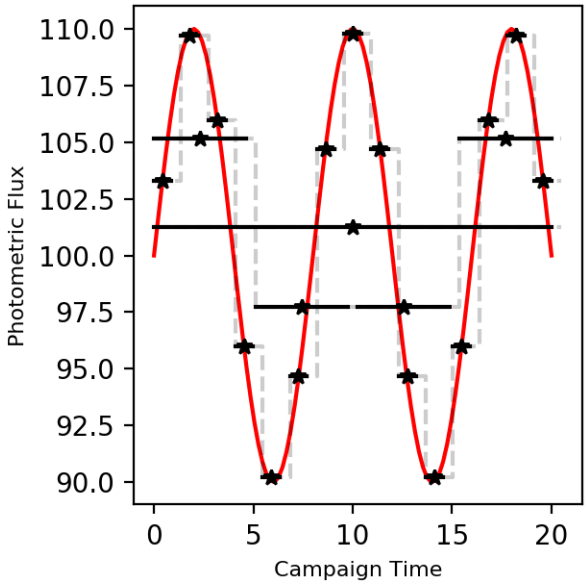

Figure 4.6: Flux over time as measured by different exposure times. The smooth red line is truth while the solid black lines represent measured flux for exposure times of $t_{I,k} = 20$s, $t_{I,k} = 4.6$s, and $t_{I,k} = 0.86$s. Stars are paced at the center time of each exposure, $\bar{t}_k$.

4.6.2 Optimal Velocity Astrometry Solution

More photons lead to better statistical significance for position astrometry in an image or series of images. For velocity astrometry, more photons still help, but the time baseline over which observations are taken is important as well. This leads to prioritizing long exposure times in the first and last observation of an image set. There is a natural counterweight to this, because longer first and last exposure times push the first and last center times, $\bar{t}_k$ closer towards the center. As mentioned, the optimal arrangement of exposure times is solved with a descent algorithm, which is performed over a range of values for the number of exposure times, N_t. Because each number of exposures is solved independently, each solution can be shown and respectively analyzed. In Figure 4.7, each solution integration set is visually represented. Most of the solutions regardless of N_t have the same general form, prioritizing the first and last exposure. As Equation (4.16) implies, the integrations furthest away from the center time, $t^* = 10$s, have the most information value. This is better

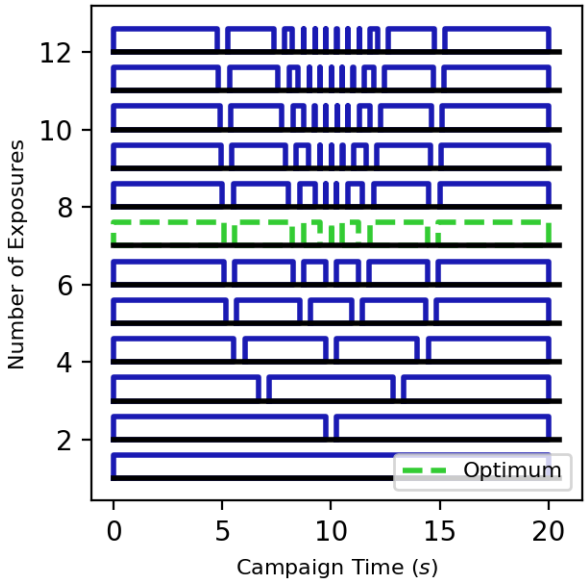

Figure 4.7: For the velocity information reward function, this shows each optimal set of integrations for a respective total number of integrations. The total number of integration which provides the highest reward is highlighted with a green dashed line.

illustrated in Figure 4.8 which shows the cumulative information after each exposure as a function of time. The integrations very close to 10s have little impact on the information content, and morph to support maximizing information from the first and last integrations. This leads to all of the solutions with more than 4 integrations having similar first and last integrations, while the integrations in the middle of the set vary greatly. The information content in each solution in Figure 4.7 is shown in Figure 4.9. Note how the information content for all solutions past 6 exposures is effectively the same. Figure 4.9 also compares the solutions to campaigns of constant exposures, to illustrate the advantages.

4.6.3 Optimal Photometry Solution

Constant flux photometry requires maximizing the number of photons, similar to position astrometry, and therefore wants a long exposure time. Contrary to the other reward functions, photometry on a variable flux signal requires sufficiently short exposure times in order to minimize the flux bias, ϵ_F. This requirement creates a reward function that wants

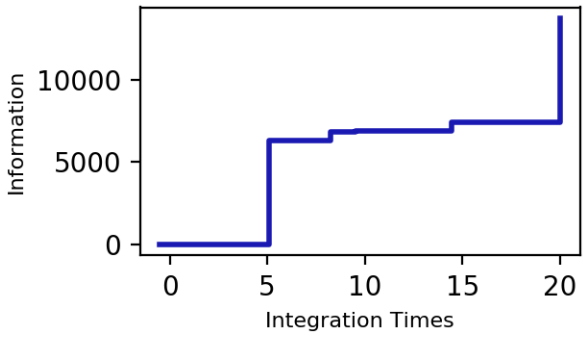

Figure 4.8: For the velocity information reward function, this shows the cumulative information over time for the optimal solution. Information is focused in the first and last exposures.

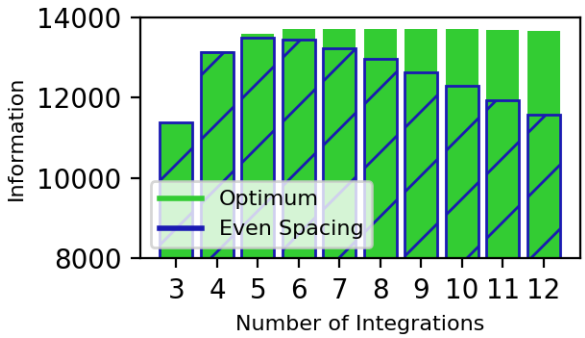

Figure 4.9: For the velocity information reward function, this shows the information content for each possible number of integrations (value of N_t). Information is shown for a constant exposure time at each value of N_t for comparison.

both short exposures and as little dead time as possible. To highlight the duality of this reward function, Figure 4.10 shows solutions for $A = 100$, while Figure 4.13 shows solutions for $A = 10$. In general there is no reason to vary the exposure time so the solutions typically have constant t_I. Figure 4.10 illustrates why this is in the solutions for a low number of exposures; the solution involves minimizing as many exposures as possible and leaving one large exposure with a large photometry bias. The primary difference between $A = 100$ and $A = 10$ is the optimal exposure time. The integration time in the solution does not have intuitive meaning without the context provided by the frequency of the oscillating flux. The choice of integration time combined with $\omega = \pi/4$ gives a value on Figure 4.4, associating with a certain amount of error. Figure 4.13 is skewed towards ignoring the variable flux and prioritizing as few exposures as possible, and so the solution occurs as 4 integrations of about 4.5 seconds each with $\alpha = 0.44A$. Figure 4.10 is skewed towards prioritizing the variable flux as a large source of error, and so the solution occurs as 8 integrations of about 2.0 seconds each with $\alpha = 0.10A$. Recall that the value of α directly defines the variance in Equation (4.32), which is in the denominator of Equation (4.38). The optimization chooses values of t_I only in that they directly define the error due to ϵ_F. Figures 4.11 and 4.14 show the information over time for the optimal campaign, which takes the form of a simple step pattern. Figures 4.12 and 4.15 show the information for each exposure time compared against a constant exposure time campaign. The two only diverge for low numbers of exposures corresponding with the non constant exposure campaigns in Figures 4.10 and 4.13.

4.6.4 Velocity Astrometry Versus Photometry Pareto Surface

Because this is multi-objective optimization, there is no single solution but a set of solutions which all attempt to balance the two reward functions. This set of solutions can be seen in Figures 4.16 and 4.18, for $A = 100$ and $A = 10$ respectively, as a curve through the space defined by the two reward functions. Because the rewards have different units, both

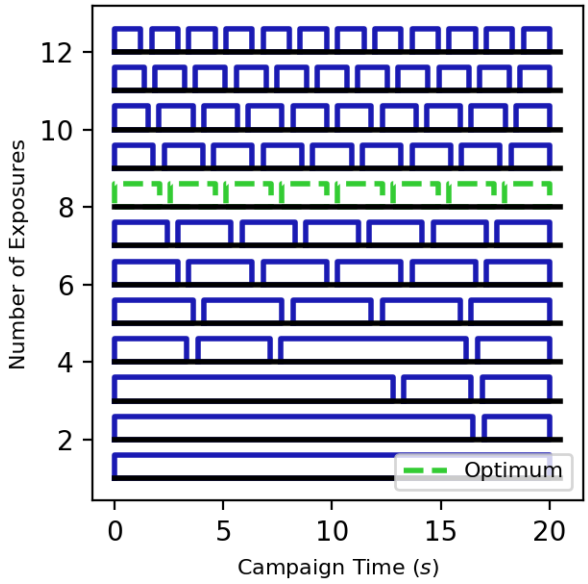

Figure 4.10: For the photometry information reward function and $A = 100$, this shows each optimal set of integrations for a respective total number of integrations. The over all optimum is highlighted with a green dashed line.

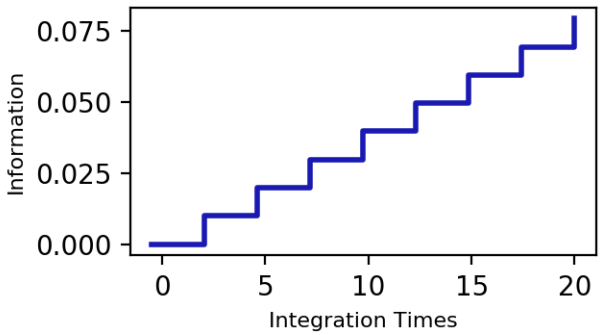

Figure 4.11: For the photometry information reward function and $A = 100$, this shows the cumulative information over time for the optimal solution. Information is focused in the first and last exposures.

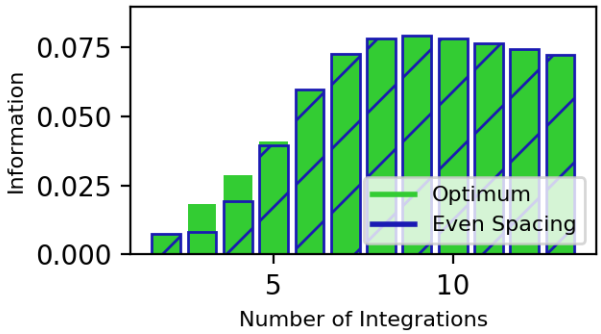

Figure 4.12: For the photometry information reward function and $A = 100$, this shows the information content for each possible number of integrations (value of N_t). Information is shown for a constant exposure time at each value of N_t for comparison.

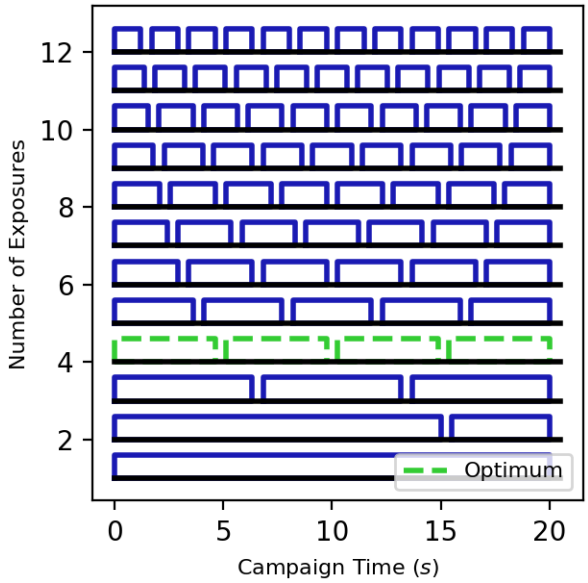

Figure 4.13: For the photometry information reward function and $A = 10$, this shows each optimal set of integrations for a respective total number of integrations. The over all optimum is highlighted with a green dashed line.

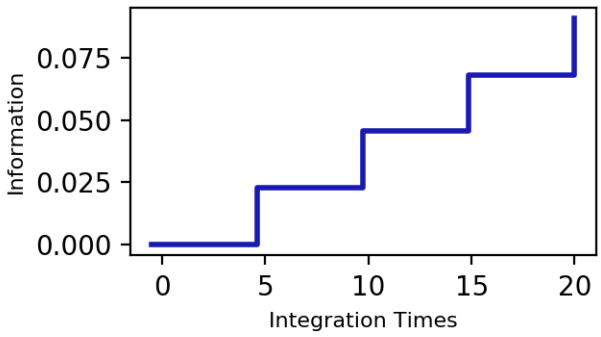

Figure 4.14: For the photometry information reward function and $A = 10$, this shows the cumulative information over time for the optimal solution. Information is focused in the first and last exposures.

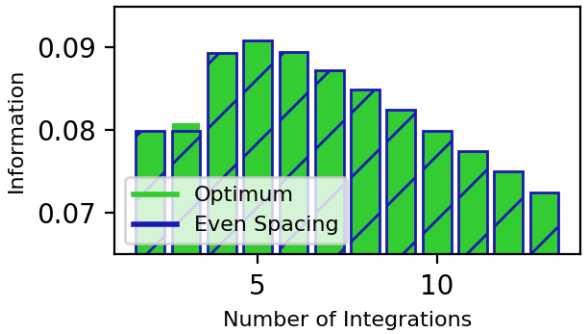

Figure 4.15: For the photometry information reward function and $A = 10$, this shows the information content for each possible number of integrations (value of N_t). Information is shown for a constant exposure time at each value of N_t for comparison.

solutions are normalized by the individual optima for each reward function. This set of individual optima is known as the utopia point. As discussed in Section 4.5, the actual Pareto Surfaces are composed of a series of visually convex surfaces corresponding to different numbers of exposures, N_t. The surface in Figure 4.16 goes from $N_t = 8$ to $N_t = 7$ to $N_t = 6$ and then back to $N_t = 7$ for the very last few solutions, though it is hard to distinguish visually. The surface in Figure 4.18 is comprised of a series f disjoint surfaces from $N_t = 4$ to $N_t = 7$. To complement the Pareto surfaces, a series of solutions along the Pareto frontier are provided in Figures 4.17 and 4.19. These solutions show the steady progression between a velocity astrometry optimal solution in Figure 4.7 at the bottom to a photometry optimal solution in Figures 4.10 and 4.13 at the top. To provide context for a single solution, the solution closest to the utopia point, $[1, 1]^T$, in terms of Euclidean distance is highlighted on all plots. The highlighted solutions in the middle of Figures 4.17 and 4.19 provide a campaign which intuitively balances photometry and astrometry by taking the optimal first and last integrations for astrometry and inserting the optimal photometry tasking in between.

4.7 Conclusion

This chapter provides a framework for analyzing how to select telescope parameters for SSA data collection. Reward functions are developed for the primary two goals of SSA, astrometry and photometry. Gradients are developed with respect to the main design variable considered, exposure time. This provides results for an optimal exposure time schedule and Pareto frontiers between the two reward functions.

One of the main lessons from this work is the trade-off between the various reward functions. Photometry comes in two varieties, long exposures on an unchanging flux versus high frame rate of varying flux. Velocity astrometry by itself has an trade-off between long exposures and maximization of the time baseline. When combined together, the various reward functions can have unexpected interconnections as well, which is shown in their

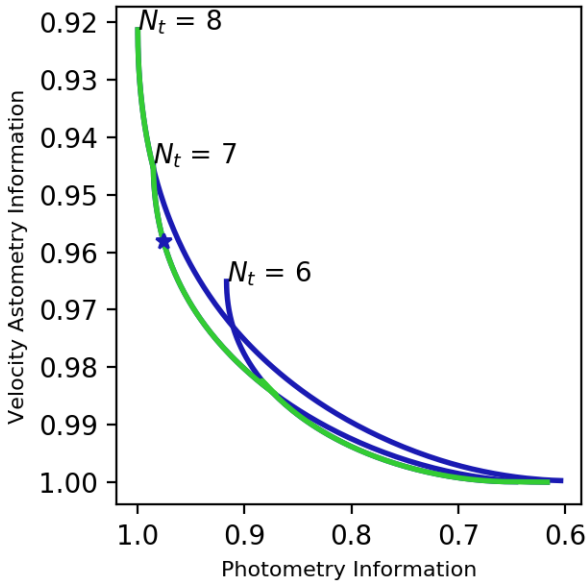

Figure 4.16: Demonstration of the Pareto surface between velocity astrometry and photomtery reward functions for $A = 100$. The axes are normalized by the individual optima for each reward function, or utopia point.

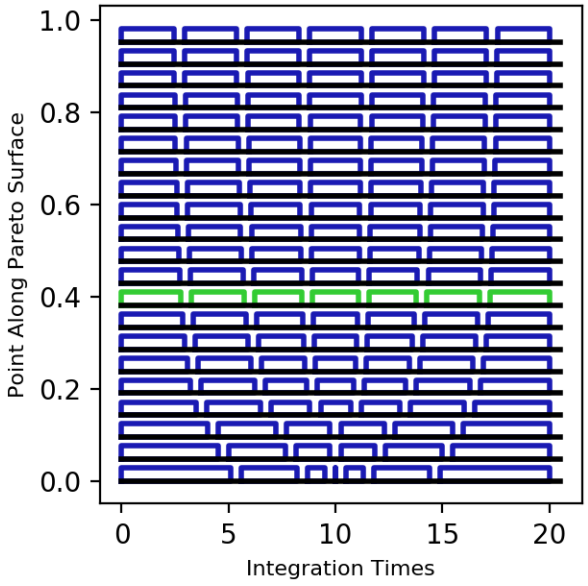

Figure 4.17: Each optimal set of integrations for a respective point along the Pareto surface for $A = 100$. Solutions vary from the optimal velocity solution in Figure 4.7 to the optimal photometry solution in Figure 4.10.

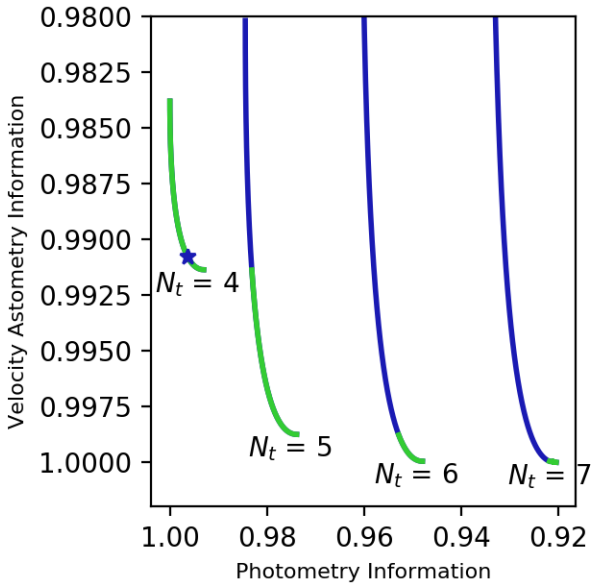

Figure 4.18: Demonstration of the Pareto surface between velocity astrometry and photomtery reward functions for $A = 10$. The axes are normalized by the individual optima for each reward function, or utopia point.

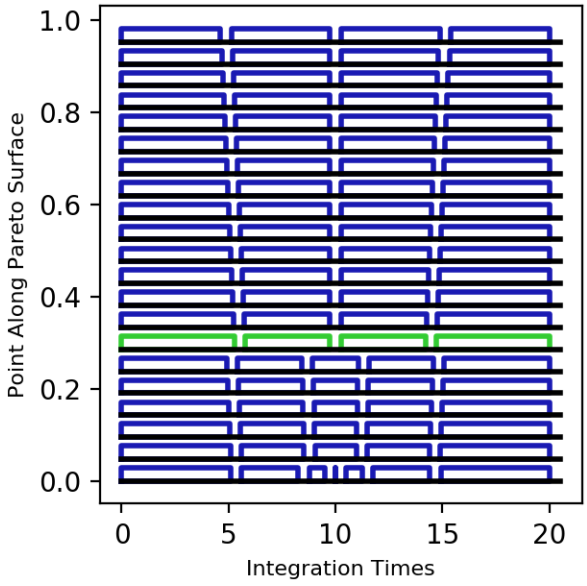

Figure 4.19: Each optimal set of integrations for a respective point along the Pareto surface for $A = 10$. Solutions vary from the optimal velocity solution in Figure 4.7 to the optimal photometry solution in Figure 4.13.

Pareto surfaces. While many data campaigns with a single goal may stay the same, taking data efficiently always has value, which is the purpose of this work.

4.7.1 Unifying Example

At this point in our example, we have a series of locations to be search along with a start and end time for data collection. We implicitly also have a nominal orbit for which we are searching, provided by the admissible region, for each location we are searching. The work in this contribution first commands that the sensor rate track that nominal orbit. The sensor is really tasking on a small set of orbits, which all cannot be rate tracked, so a mean orbit must be chosen. The way in which to effectively task on a set is a source of future work for this paper. Next, a Pareto analysis is performed for every data collection period, according to the work in this chapter. Depending on the sensor operator's goals, the final data collection solution would either solely focus on astrometry or balance astrometry and photometry; both strategies have merits. This second contribution outputs a series of images for each area of the admissible region, which must now be analyzed. Implicitly, there is also prior knowledge of the object motion through the image, from the admissible region.

CHAPTER 5
VISUAL TRACKING METHODS FOR IMPROVED SEQUENTIAL IMAGE-BASED OBJECT DETECTION

The first problem this chapter analyzes is the likelihood function. The current state-of-the-art likelihood function is based on modeling the statistics of an image [41]. This likelihood function has been around for a long time, and has been used in recent years for detectionless tracking. [73]. The same likelihood can be applied to moving point sources as well as complex shapes such as tracking of people [56]. More critically, this method doesn't make any detections. Instead the data in the image is used produce a likelihood for a given state. If an object's signal deviates significantly from noise and follows a set of given dynamics, the filter obtains enough evidence to say a detection has been made.

A different method, based on a hypothesis test on the matched filter predicted by a particular particle to evaluate measurement likelihood, is used in this chapter [62] [74]. This newer likelihood provides a series of improvements, including eliminating the requirement to predict object brightness, and a built-in way to modulate probability of false alarm. This chapter shows a fundamental similarity between the two likelihoods; the only substantial difference is the use of a hypothesis test. A simulation-based analysis gives evidence for why one algorithm may be better than another in a certain situation.

Another area this chapter improves upon is the particle birth model. The Bernoulli filters cited in this chapter generate new particles using a random uniform distribution. The birth model of a random finite set-based tracker is intended to be catered to particular problems. No such birth model has been designed for SSA applications, and this chapter proposes some guidelines for creating one. The discussion in this chapter analyzes how to search a particular orbit regime, how to incorporate prior knowledge, and how to efficiently search a space given no prior knowledge.

5.1 Methodology

This chapter applies the random finite set-based multi-Bernoilli filter with a detection-less likelihood function to frame-to-frame tracking of space objects observed in electro-optical imagery for space domain awareness applications. First, this chapter reviews multi-Bernoulli filters applied to frame-to-frame tracking, image statistics, and matched filters. A likelihood function for space-based imagery is analyzed in comparison to the previously used likelihood function. A birth model is proposed which better models potential space objects using observer characteristics and object dynamics. In simulation, the final algorithm is able to perform completely uncued detection down to a total photometric Signal-to-Noise Ratio (SNR) of 5.6 and a per pixel SNR of 1.5. Promising results are shown for total photometric SNR of 3.35 and per pixel SNR of 0.7. The algorithm is also applied to empirical data, tracking of low SNR, Geostationary objects in images taken with a 0.5m Raven-class telescope.

To reiterate, this chapter makes the following contributions. First, state-of-the-art multi-object tracking methods for space object discovery in a series of images are presented. This chapter proves that the current state-of-the-art likelihood function is inherently related to a matched filter and therefore has SNR optimality. The new likelihood function is analyzed for different noise models that can apply to electro-optical images. The previous state of the art likelihood and the new likelihood are compared to show which filter is better suited for what situations. This chapter goes on to explore a series of specialized birth models which can apply to SSA Imagery. This includes pure object discovery, search in a specific orbit regime, and search primed by previous orbit knowledge. Finally, the validation of proposed methods is presented on empirical data taken with the Georgia Tech Space Object Research Telescope.

This chapter is presented as follows. Section 5.2 reviews the basics of multi-Bernoulli filtering. Section 5.3 presents a comprehensive discussion of the theoretical and quantita-

tive aspects of the two likelihood functions. Section 5.4 motivates and derives the various birth models proposed in this work. Section 5.5 explains the particular implementation details involved in this filter. Section 5.6 shows simulation results which provide insight into limitting SNR values for the algorithm. Section 5.7 presents the results of operating the algorithm on empirical data taken with a 0.5m Raven-class telescope.

5.2 Multi-Bernoulli Filtering for Frame-to-Frame Tracking

5.2.1 Dynamics

Consider the standard filtering problem with discrete dynamics and measurement models

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k) + \mathbf{w} \tag{5.1}$$

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v} \tag{5.2}$$

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{z} \in \mathbb{R}^m$, $\mathbf{w} \sim f_w : \mathbb{R}^n \to \mathbb{R}$, and $\mathbf{v} \sim f_v : \mathbb{R}^m \to \mathbb{R}$. $\mathbf{x}$ can be thought of as the state to be estimated, while $\mathbf{z}$ can be thought of as the measurements. For frame-to-frame image tracking, the measurement is a series of m pixels, $\mathbf{z}_k := \{z_{k,i}\}_{i=1}^m$, where k represents the time step associated with a particular frame.

5.2.2 Multi-Bernoulli Filter

This chapter uses the multi-Bernoulli filter used in [73]. Various other variations of this filter, including the Labeled Multi-Bernoulli Filter [76], can be used as well. This chapter's primary focus is not on which multi-target particle filter is best, but on how to implement such filters on sequences of images. A brief review of the theory and application is presented for the reader's convenience. The Bernoulli filter starts from Bayesian filtering

equations [54]

$$p(\mathbf{x}_k|\mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{x}_{k-1}|\mathbf{z}_{1:k-1})\mathrm{d}\mathbf{x}_{k-1} \qquad \text{(Prediction)} \qquad (5.3)$$

$$p(\mathbf{x}_k|\mathbf{z}_{1:k}) = \frac{p(\mathbf{z}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{z}_{1:k-1})}{p(\mathbf{z}_k|\mathbf{z}_{1:k-1})} \qquad \text{(Update)} \qquad (5.4)$$

where $\mathbf{z}_{1:k}$ is the time series of measurements up until times step k. Equation 5.3 is the prediction step while Equation 5.4 is the update step.

For the multi-object filter, $\mathcal{X}_k$ is defined as a random finite set (RFS) used to represent the multi-object tracking problem where the number of objects is unknown. In essence, an RFS is a set of objects represented as random vectors, but in which the number of objects itself is also a random variable. For a single Bernoulli filter, the state is modeled as a Bernoulli random finite set (BRFS). The BRFS is a special set $\mathcal{S}_k$ containing a probability of existence, r_k, and a RFS which contains 1 object with probability r_k and is empty with probability $1 - r_k$. The object state is described by the PDF, $p(\mathbf{x}_k)$. In other words, the object is described by a PDF and probability of existence shown in Equation (5.5).

$$\mathcal{S}_k = \{p(\mathbf{x}_k), r_k\} \qquad (5.5)$$

The Bernoulli filter estimates the PDF with equations (5.3) and (5.4) and updates the probability of existence with update equations that can be found in [73].

A multi-Bernoulli filter is a bank of Bernoulli filters used to process the full multi-object tracking problem. The prediction and update equations for a multi-target tracking scheme are

$$\pi_{k|k-1}(\mathcal{X}_k|\mathcal{Z}_{1:k-1}) = \int p_{k|k-1}(\mathcal{X}_k|\mathcal{X}_{k-1})\pi_{k-1}(\mathcal{X}_{k-1}|\mathcal{Z}_{1:k-1})\delta\mathcal{X} \qquad \text{(Prediction)} \qquad (5.6)$$

$$\pi_k(\mathcal{X}_k|\mathcal{Z}_{1:k}) = \frac{g(\mathcal{Z}_k|\mathcal{X}_k)\pi_{k|k-1}(\mathcal{X}_k|\mathcal{Z}_{1:k-1})}{\int g(\mathcal{Z}_k|\mathcal{X}_k)\pi_{k|k-1}(\mathcal{X}_k|\mathcal{Z}_{1:k-1})\delta\mathcal{X}} \qquad \text{(Update)} \qquad (5.7)$$

where $g(\cdot|\cdot)$ is the likelihood function, a subject of novel work in this chapter. Keep in mind that these are FISST set integrals, not standard integrals. In essence, π_k is similar to a PDF in that higher density means higher probability of an object existing in that area. However, π_k does not need to integrate to 1; instead, the integral over an area is the expected number of objects in that area.

The multi-Bernoulli filter represents the multi-object state as the union of a series of BRFSs

$$\mathcal{X}_k = \bigcup_{l=1}^{n_s} \mathcal{S}_k^l \tag{5.8}$$

This allows the implementation of a Multi-Bernoulli filter to reduce to a series of Bernoulli filters. There are particle implementations of Bernoulli filters [75], and the union of these filters is relatively straight forward to compute. A more in depth, general analysis of FISST can be seen in [49]. The next few sections focus on applying the Multi-Bernoulli filter which operates on image frame-to-frame tracking to EO sensors tracking. The primary novel change is defining a likelihood function $g(\cdot|\cdot)$ which combines space imagery and Multi-Bernoulli filter. The likelihood function is used to update particle weights and is where a measurement function is defined. Because in images, the measurement is the whole image, special considerations should be taken to properly define this function.

5.3 Analysis of Likelihood Functions Suitable for Detection of Space Objects

5.3.1 Statistics of an Image

This section defines the statistical models for an image that are used throughout the chapter. This includes a review of the information in Section 2.2. In general there are a variety of ways to model the signal and noise an image, and this chapter discusses results with respect to multiple models.

An image is made up of a series of measurements taken at the same time. In practice,

each pixel z_i is broken into some deterministic signal content s_i and some zero mean noise w_i of known distribution $p(w)$ (assuming background subtraction has been performed) as seen in Equation (2.5). While z_i represents the i^{th} pixel, this same quantity can be defined as a vector of N_z pixels, $\mathbf{z} = [z_1 \dots z_{N_z}]^T$.

The distribution of w_i in Equation (2.5) can be arbitrary, but often is assumed to be of a specific form. A common way to model the statistics of a pixel is with a series of Poisson distributions [78]. The pixel is composed of some signal (shot noise), dark noise and read noise. It is convenient to redefine a background subtracted z_i in Equation (2.7). This new definition of z_i allows the problem to be posed as the sum of a deterministic signal and zero mean random noise. The Poisson distributions approaches a Gaussian distribution as the parameter, λ, increases. Further assumptions can be made to assume the dark noise and read noise dominate the shot noise, are equivariant in all pixels, and are independent of each other, giving Equation 2.10

Note that in this formulation, s_i is equivalent to the mean of the measurement z_i, which is Gaussian distributed. It is worth noting that these measurements are taken over an integration time, of which they are functions

$$\mathcal{T} = [t_0, t_0 + t_I] \tag{5.9}$$

$$z_i = z_i(\mathcal{T}) \tag{5.10}$$

While this notation is not included in this chapter, timing and integration error is a further source of uncertainty in pixels.

SNR is a measure of the amount of signal present relative to the strength of the noise and a common measure of signal quality. The SNR for a particular pixel, z_i, is defined in Equation (2.12). Total photometric SNR for an object that exists in N_z pixels (assuming independent noise between pixels) is calculated in Equation (2.13).

5.3.2 Matched Filters

The matched filter (MF) is an image filter which is used to define a likelihood function in this chapter. The MF is predicated on correlating a hypothesized signal structure with a series of pixels in an image. Given a state, $\mathbf{x}$, in some state space, define the set and function, $T(\mathbf{x})$ and $h_i(\mathbf{x})$, which predict the location of an object in an image, and the pixel values produced by the object, respectively.

$$h_i(\mathbf{x}) = s_i' \tag{5.11}$$

$$T(\mathbf{x}) = \{i : s_i' > 0\} \tag{5.12}$$

In other words, if an image has m pixels, and there exists a hypothesized state $\mathbf{x}$, then $T(\mathbf{x})$ is a list of pixels which contain some signal given $\mathbf{x}$, and $h_i(\mathbf{x})$ is the predicted value in pixel i given $\mathbf{x}$. Also note that the cardinality of $(T(\mathbf{x}))$ is N_z, which is directly related to the N_z in Equation (2.14). These two quantities, h_i and T, are referred to as the template in matched filter theory. If the object has no apparent motion in an image, $T(\mathbf{x})$ and $h_i(\mathbf{x})$ are defined by the point spread function in the image, the flux of the photons coming off the object, and the exposure time. If the object has apparent motion, the template is generated via an integration of a moving point spread function along the dynamics. See [74] for more details.

Given the most general formulation of z_i in Equation (2.5), the matched filter is a linear image filter which seeks to maximize the post filter SNR. Consider an arbitrary linear filter, $\mathbf{m} := \{m_i\}_{i=1}^{N_z}$

$$z_{MF} = \sum_{i \in T(\mathbf{x})} m_i z_i = \mathbf{m}^T \mathbf{z} = \mathbf{m}^T \mathbf{s} + \mathbf{m}^T \mathbf{w} \tag{5.13}$$

$$\text{SNR}_{MF} = \frac{\mathbf{m}^T \mathbf{s}}{\sqrt{\mathbb{E}[\mathbf{m}^T \mathbf{w}]^2}} \tag{5.14}$$

where $\mathbf{s}$ is the pure signal content and $\mathbf{w}$ is the pure zero mean noise content. At this point we can define the covariance matrix of $\mathbf{w}$, $\mathbf{R}_w = \mathbb{E}[\mathbf{w}\mathbf{w}^T]$. From matched filter theory the optimal choice for $\mathbf{m}$, to maximize the SNR (5.14) of the final result, is

$$\mathbf{m} = a\mathbf{R}_w^{-1}\mathbf{s} \tag{5.15}$$

where a is an arbitrary scaling parameter [93] [94].

If the underlying state $\mathbf{x}$ is known, then $\mathbf{s} = \mathbf{h}(\mathbf{x})$. For a hypothesized state $\mathbf{x}$, then $\mathbf{s} = \mathbf{h}(\mathbf{x})$ is still a reasonable method for testing the validity of that hypothesis, but produces a sub-optimal filter. If the noise vector $\mathbf{w}$ is independent and equivariant, $\mathbf{R}_w^{-1}$ can be absorbed into the arbitrary scaling parameter, a, giving $\mathbf{m} = a\mathbf{s}$. The matched filter makes no assumptions about the distribution of the noise. In other words, if the expected value of the measurements ($\mathbf{h}(\mathbf{x})$) and the covariance matrix of the noise can be estimated ($\mathbf{R}_w$), the matched filter is SNR optimal.

The final matched filter equation for some predicted state, $\mathbf{x}$, and i.i.d. noise is then

$$z_{MF} = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) z_i \tag{5.16}$$

5.3.3 Likelihood Ratio

The current established likelihood function for tracking point sources in images is a likelihood ratio [73], [41]. This has already been used in SO tracking [42].

Each pixel is assumed to belong to one of two distributions, depending on whether the pixel is predicted to have signal.

$$p(z_i|\mathbf{x}) = \begin{cases} \psi_i(z_i) & i \in T(\mathbf{x}) \\ \phi_i(z_i) & i \notin T(\mathbf{x}) \end{cases} \tag{5.17}$$

Commonly assumed distribution for ψ and ϕ is

$$\phi(z_i) = \mathcal{N}(z_i; 0, \sigma^2) \tag{5.18}$$

$$\psi(z_i, \mathbf{x}) = \mathcal{N}(z_i; h_i(\mathbf{x}), \sigma^2) \tag{5.19}$$

where $\mathcal{N}(\cdot; \mu, \sigma^2)$ is a Gaussian with mean μ and variance σ^2. Note that this method does not rely on a Gaussian assumption. Like the matched filter, this method should work for any distributions that can be properly modeled. The total likelihood for the entire filter is then

$$g(z|\mathcal{X}_k) = \left(\prod_{\mathbf{x}_k \in \mathcal{X}_k} \prod_{i \in T(\mathbf{x}_k)} \psi_i(z_i, \mathbf{x}_k) \right) \left(\prod_{i \notin \cup T(\mathbf{x}_k)} \phi_i(z_i) \right) \tag{5.20}$$

The total likelihood is a quantity used in multi-target filters. The per pixel likelihood is calculated as the product of relative likelihoods of a pixel belonging to either the distribution in Equation (5.17).

$$g_z(\mathbf{x}) = \prod_{i \in T(\mathbf{x})} \frac{\psi(z_i, \mathbf{x})}{\phi(z_i)} \tag{5.21}$$

The probability of existence update is formulated with the same relative likelihood. This likelihood function is general enough to work for the case of arbitrary point tracking.

The current likelihood function, in the case of Gaussian noise, is nothing more than the ratio of the measured matched filter result tested against a zero mean and assumed variance.

Lemma 5.3.1. *Using the definition of z_i in Equation* (2.10)*, the likelihood ratio defined in Equation* (5.21) *can be rewritten as a likelihood ratio of the result of a matched filter.*

Proof. Using the definition of a Gaussian PDF

$$
\begin{aligned}
g_z(\mathbf{x}) &= \prod_{i \in T(\mathbf{x})} \frac{e^{-\frac{1}{2}(z_i - s_i)^2}}{e^{-\frac{1}{2}(z_i - 0)^2}} \\
&= \prod_{i \in T(\mathbf{x})} e^{-\frac{1}{2}(z_i^2 - 2z_i s_i + s_i^2)} e^{\frac{1}{2} z_i^2} \\
&= e^{\sum (z_i s_i - \frac{1}{2} s_i^2)}
\end{aligned}
\tag{5.22}
$$

Next, define the following constant scaling parameter as the root sum square of the per pixel photometric intensity.

$$
\alpha = \sqrt{\sum_{i \in T(\mathbf{x})} (s_i^2)} \tag{5.23}
$$

Recalling the matched filter defined in Equation (5.13), the likelihood ratio becomes,

$$
\begin{aligned}
g_z(\mathbf{x}) &= e^{\sum (z_i s_i - \frac{1}{2} s_i^2)} \\
&= e^{\frac{z_{MF}}{\alpha}\alpha - \frac{1}{2}\alpha^2 + \left(\frac{1}{2}\frac{z_{MF}}{\alpha}\right)^2 - \left(\frac{1}{2}\frac{z_{MF}}{\alpha}\right)^2} \\
&= \frac{e^{-\frac{1}{2}\left(\frac{z_{MF}}{\alpha} - \alpha\right)^2}}{e^{-\frac{1}{2}\left(\frac{z_{MF}}{\alpha}\right)^2}}
\end{aligned}
\tag{5.24}
$$

Note the identical form between Equations (5.22) and (5.24). □

Note that the quantity z_{MF}/α is a scaled matched filter result with an expected value of α. The likelihood ratio is therefore a comparison of the matched filter result on the true signal and pure noise.

The primary difficulty in using this canonical likelihood ratio is the required estimate of the photometric root mean square of the pixel intensities. In space object tracking, and especially in object discovery, the actual photon flux coming off a space object is difficult or impossible to know before a detection is made. Conversely, with a good estimate of a point spread function, the relative photometric intensities between the pixel values can be

predicted; this is only dependent on the dynamics and point spread function. The relative likelihood cannot be correctly computed without the expected signal values in each pixel. Previous efforts to implement this likelihood estimate the mean of the signal [42]. However, the matched filter is invariant to an arbitrary scaling parameter; the SNR gain is optimal so long as the expected relative values are correct. Because of this, the matched filter can be used to define a likelihood function which doesn't need a photon flux estimate.

5.3.4 Hypothesis Test Likelihood

This chapter proposes a new likelihood function for detectionless multi-Bernoulli filters based on a hypothesis test on the matched filter. Such a likelihood function has been used for a general particle filter [88].

Similar to the already shown likelihood function, the measured signal is assumed to be zero mean noise and predicted signal $h_i(\mathbf{x})$ in a series of pixels $T(\mathbf{x})$. The matched filter is the weighted sum of the measured pixels, z_i, weighted by the predicted values, $h_i(\mathbf{x})$, shown in Equation (5.16). Under an arbitrary distribution of w_i assumed in Equation (2.5), the distribution of z_{MF} can be calculated.

$$p_{MF} = p_{h_1(\mathbf{x})z_1} * p_{h_2(\mathbf{x})z_2} * \ldots * p_{h_{nz}(\mathbf{x})z_{nz}} \tag{5.25}$$

where $*$ represents the convolution operator. This can be difficult to calculate for arbitrary complex distributions, but is possible.

Because the noise, w_i, is assumed to be mean zero, if the predicted region contains no signal, $s_i = 0 \; \forall \, i \in T(\mathbf{x})$, the resulting distribution is mean zero.

$$\mathbb{E}[z_{MF} | s_i = 0] = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) \, \mathbb{E}[w_i] = 0 \tag{5.26}$$

By a similar argument, if there is signal content in the pixels on which the matched filter operates, the result has a mean greater than zero. If the matched filter operates with a

perfect template, the distribution has a known expected value.

$$\mathbb{E}[z_{MF}|s_i = h_i(\mathbf{x}_{true})] = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) h_i(\mathbf{x}_{true}) = a\alpha^2 \tag{5.27}$$

Note that $a\alpha^2$ is typically not achieved in particle filtering, due to the random nature of particle sampling. The matched filter, z_{MF}, can be assumed to exist in one of two distributions

$$\tilde{\phi}_{MF}(z) = p_{MF}(z; 0) \tag{5.28}$$

$$\tilde{\psi}_{MF}(z; \beta) = p_{MF}(z; \beta), \ \beta > 0 \tag{5.29}$$

where the distribution $p_{MF}(z; \beta)$ has an expected value of β and is otherwise an arbitrary PDF calculated in Equation (5.25). These distributions in Equations (5.17) and (5.29) are mirrors of each other; Equation (5.17) defines the two possible distributions for a pixel given signal or no signal, while Equation (5.29) defines the two possible distributions for a matched filter given signal or no signal.

Recall that the Equation (5.16) assumes that noise is i.i.d., which is violated by shot noise which explicitly dependent on the magnitude of the signal. Because signal magnitude is unknown, shot noise must either be estimated or ignored. For some low SNR signals, shot noise is dominated by other noise sources and can be ignored, such as when read noise is high. It should be noted that this is not always true, in which case shot noise must be estimated. If estimated, it requires calculation of the covariance matrix $\mathbf{R}_w$ in Equation (5.15). This is a limitation to both likelihoods analyzed in this chapter. Note that pixel SNR in Equation (2.12) is $\lambda_s / \sqrt{\lambda_w + \lambda_s}$. Choose a constant, γ, such that in order to say shot noise is dominated, $\lambda_w > \gamma \lambda_s$. Also choose an SNR below which this method is being

used. Then if

$$\lambda_w > (\text{SNR})^2 \frac{1+\gamma}{\gamma^2} \tag{5.30}$$

holds, either the SNR is sufficiently high that is does not matter, or $\lambda_w > \gamma\lambda_s$.

Next, the following null and alternate hypotheses for a binary hypothesis test are defined

$$\begin{aligned} H_0 &: z \sim \tilde{\phi}(z) \\ H_1 &: z \sim \tilde{\psi}(z;\beta); \beta = z_{MF} \end{aligned} \tag{5.31}$$

In essence, this test assumes no signal content in the pixels defined by $T(\mathbf{x})$, and asks whether the measurement gives significant evidence of underlying signal (mean of matched filter greater than zero). For binary hypothesis testing, a probability of false alarm is set, p_{FA}, which in turn defines an integration threshold, z_{TH}, based on the null hypothesis cumulative density function (CDF).

$$p_{FA} = \int_{z_{TH}}^{\infty} \tilde{\phi}(z)dz \tag{5.32}$$

A probability of detection can be calculated with the following integral.

$$p_D = \int_{z_{TH}}^{\infty} \tilde{\psi}(z)dz \tag{5.33}$$

To calculate this probability of detection, the arbitrary distribution $p(z;\beta)$ in Equation (5.31) must be known. For more on binary hypothesis testing, see [95]. This hypothesis test determines if there is significant evidence that the predicted signal exists in the predicted location. Because the matched filter gives an SNR gain, this test should maximize p_D, though an explicit proof of this claim has not yet been shown.

The probability of existence update can also be formulated in terms of the matched

filter. The relative likelihood can be calculated from the two distributions in Equation (5.29). This leads to the particle-wise relative likelihood

$$g_z(\mathbf{x}) = \frac{\tilde{\psi}(z_{MF})}{\tilde{\phi}(z_{MF})} \tag{5.34}$$

This backtracking to a relative likelihood is necessary for the probability of existence update, which does not operate correctly without a relative likelihood. The simulation results shown in Section 5.6 highlight the need for further research into this likelihood ratio. The actual test distributions of z_{MF}, $\tilde{\phi}(z)$ and $\tilde{\psi}(z)$, are not trivial to calculate in general. If z_i is assumed to be Gaussian distributed as shown in Equation (2.10), and the noise is again assumed to be i.i.d., z_{MF} is distributed as

$$\begin{aligned} \text{Var}[z_{MF}] &= \mathbb{E}\left[\left(a\sum_{i\in T(\mathbf{x})} h_i(\mathbf{x})w_i\right)^2\right] = a^2\sigma_w^2\sum_{i\in T(\mathbf{x})} h_i^2(\mathbf{x}) \\ z_{MF} &\sim \mathcal{N}(\alpha^2, a^2\alpha^2\sigma_w^2) \end{aligned} \tag{5.35}$$

where α is defined in Equation (5.23). Recalling Equation (5.15), the matched filter has an arbitrary scaling parameter, a. First, in preprocessing, an image can be scaled by $1/\sigma_w$ to reduce the variance in each pixel to 1. If the value $a = 1/\alpha$ is chosen, the variance is reduced to 1. Alternatively, a factor of $1/\sigma_w$ can be inserted into a, rather than in preprocessing. Recall that the expected value in Equation (5.35) can only be achieved with a perfect template, which is typically not the case. Regardless, the test distributions in this special case become

$$\tilde{\phi}_{MF}(z) = \mathcal{N}(z; 0, 1) \tag{5.36}$$

$$\tilde{\psi}_{MF}(z) = \mathcal{N}(z; \beta, 1),\ 0 < \beta \leq \alpha \tag{5.37}$$

For Poisson distributed z_i as seen in Equation (2.7), Equation (5.29) takes one of two

forms. If the expected value of the Poisson distribution is small, the test distributions are two discrete PMFs. If the expected value is high enough, the test distributions become approximately Gaussian.

For a Poisson distribution, it is easy to show

$$\begin{aligned} w \sim \frac{\lambda^k e^{-\lambda}}{k!}, k = 0, 1, 2, 3 \ldots \implies \\ c_1 w + c_2 \sim \frac{\lambda^{(k-c_2)/c_1} e^{-\lambda}}{((k-c_2)/c_1)!}, k = +c_2, c_1+c_2, 2c_1+c_2, 3c_1+c_2 \ldots \end{aligned} \tag{5.38}$$

The random variable $c_1 w + c_2$ has a PMF with mass locations dependent on the scalars c_1 and c_2. The matched filter, z_{MF}, can be represented by a PMF, because it is a weighted sum of discrete random variables with known PMFs. While Poisson random variables technically have infinite permutations, for sufficiently low values of λ, sufficient mass can be captured with a computationally bounded number of permutations.

This method of permutations can be used to exactly calculate the PMF in Equation (5.29), and evaluate the hypothesis test. The number of permutations grows proportionately with the cardinality of $T(\mathbf{x})$, and the Poisson parameter of the noise λ. When not computationally restrictive, the permutations of the PMF should be used to evaluate the hypothesis test; otherwise, a Gaussian approximation can be used. This motivates the following Lemma.

Lemma 5.3.2. *For a series of Poisson distributed random variables, $w_i \sim \mathcal{P}(\lambda_i)$, the random variable $z_{MF} = \sum_{i=1}^{N_z} c_i w_i$ is approximately Gaussian for sufficiently high values of either λ_i or N_z.*

Proof. First, the mean and variance of z is

$$\mathbb{E}[z_{MF}] = \sum_{i=1}^{N_z} c_i \lambda_i \tag{5.39}$$

$$\sigma^2 = \sum_{i=1}^{N_z} c_i^2 \lambda_i \tag{5.40}$$

Next we must define a new random variable, which is just the matched filter normalized by its mean and standard deviation.

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \frac{\sum_{i=1}^{N_z} c_i w_i - \sum_{i=1}^{N_z} c_i \lambda_i}{\sqrt{\sum_{i=1}^{N_z} c_i^2 \lambda_i}} \tag{5.41}$$

It is convenient to also define $w_i' = w_i/\lambda_i$ so that $\mathbb{E}[w_i'] = 1$ for the next part of this proof.

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \frac{\sum_{i=1}^{N_z} c_i \lambda_i w_i' - \sum_{i=1}^{N_z} c_i \lambda_i}{\sqrt{\sum_{i=1}^{N_z} c_i^2 \lambda_i}} \tag{5.42}$$

Next, we need to simply send all the Poisson parameter to be sufficiently high, $\lambda_i \to \infty$ $\forall i$. This implies that all λ_i are approximately equal to λ which gives the following form

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \sqrt{\lambda} \frac{\sum_{i=1}^{N_z} c_i w_i' - \sum_{i=1}^{N_z} c_i}{\sqrt{\sum_{i=1}^{N_z} c_i^2}} \tag{5.43}$$

By the central limit theorem, this random variable becomes a standard normal random variable. Similarly,

$$\begin{aligned} \frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} &= \frac{\sum_{i=1}^{N_z} c_i w_i - \sum_{i=1}^{N_z} c_i \lambda_i}{\sqrt{\sum_{i=1}^{N_z} c_i^2 \lambda_i}} \\ &= \sqrt{N_z} \frac{\sum_{i=1}^{N_z} c_i \frac{w_i}{N_z} - \sum_{i=1}^{N_z} c_i \frac{\lambda_i}{N_z}}{\sqrt{\sum_{i=1}^{N_z} c_i^2 \frac{\lambda_i}{N_z}}} \end{aligned} \tag{5.44}$$

which again, by the central limit theorem approaches a normal random variable as $N_z \rightarrow \infty$. □

This lemma has the following implications. As a signal is spread out over more pixels, either through a long streaking object or a large PSF, the hypothesis distributions should approach Gaussian. Additionally, as the photometric magnitude increases, the hypothesis distribution approaches Gaussian. In testing, the Gaussian approximation appears to be appropriate in most situations.

5.3.5 Qualitative Comparison of Likelihoods

To test the effectiveness of these two methods, the following Monte Carlo analysis is shown. The noise model is additive zero mean Gaussian, as in Equation (2.10), and the signal shape is assumed to be perfectly known. Each likelihood function is tested on an image containing signal with a perfect matched filter (no error in template or total magnitude), giving g_{s+w}. Each likelihood function is also tested on an image containing no signal, only noise, giving g_{0+w}. Because a Bernoulli filter uses a PDF, all particles are normalized each frame. Because of this, only the relative value between the two likelihoods affects the performance of the filter. To evaluate the performance of each likelihood function, a likelihood gain is calculated, g_{s+w}/g_{0+w}.

The likelihood gain is calculated for both filters over a variety of SNR values. It is also calculated over a variety of errors in predicted signal strength for the likelihood ratio function, and a variety of probability of false alarm values for the hypothesis test likelihood function. Each likelihood gain is calculated over 10,000 test runs, and a mean and variance of the likelihood gain are calculated.

Figures 5.1(a) and 5.1(b) show the shape of the gain with respect to the two variables of interest. The hypothesis test likelihood function can achieve higher gain values with a diminishing probability of false alarm, but at the expense of increased variance. Figures 5.2(a) and 5.2(b) show the shape of the variance with respect to the two variables of interest,

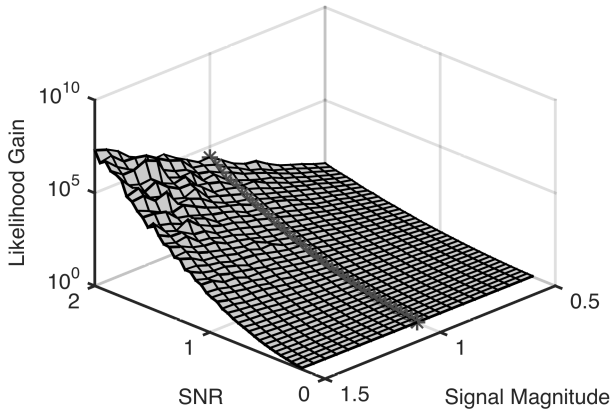

(a) Likelihood ratio function gain

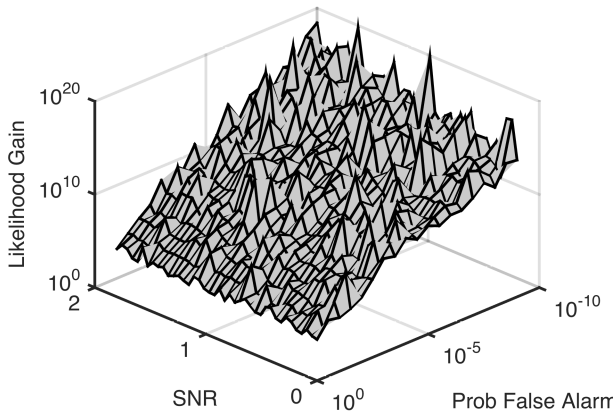

(b) Hypothesis test likelihood function gain

Figure 5.1: Comparison of gain achieved when using each likelihood function.

which appears to vary directly with gain. Figure 5.3 gives the best reason for why the new likelihood is better; while variance in gain is high, the new method preforms better at low SNR. Note that both likelihoods are dependent on knowing the exact shape of the signal defined through Equations (5.11) and (5.12). One of the implications of Lemma 5.3.1 is that both the likelihood ratio and hypothesis test likelihood are inherently based on a matched filter and therefore degrade in performance with an incorrect signal model.

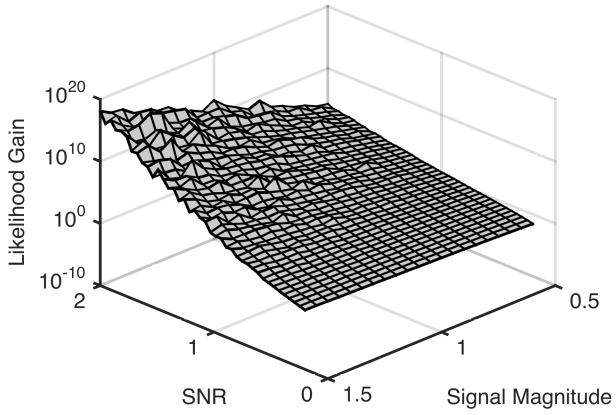

(a) Likelihood ratio function variance

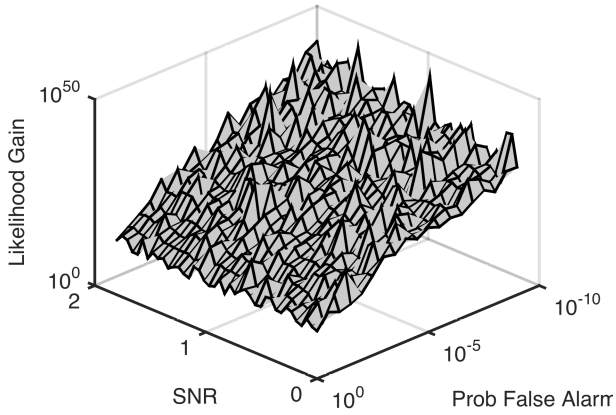

(b) Hypothesis test likelihood function variance

Figure 5.2: Comparison of the variance obtained when using each likelihood function.

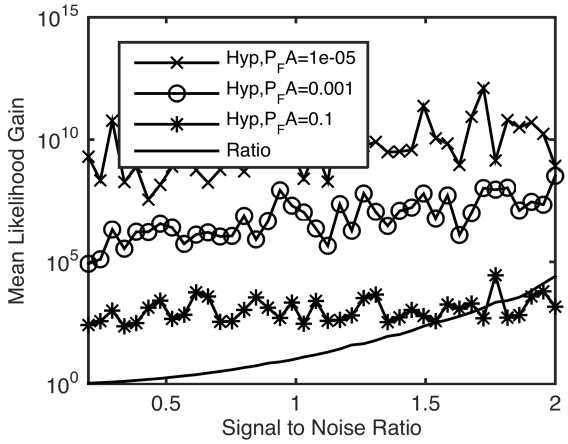

Figure 5.3: Likelihood gain comparison between methods.

5.4 Birth Model

As the Multi-Bernoulli filter runs, multiple single Bernoulli filters run in parallel. New Bernoulli filters are constantly added into the system to promote exploration of the state space. These new filters are initially sampled from a birth model, which attempts to predict the kind of states a new track could have. Birth models have been proposed and used in the past [75]. This chapter develops a series of birth models for searching for an SO.

The birth model should enforce a series of constraints in the form of sets of possible states $\mathbf{x} = [\mathbf{r}^T, \dot{\mathbf{r}}^T]^T \in \mathbb{R}^6$. If a series of c constraint sets are defined, $\{\mathcal{A}_i\}_{i=1}^{c}$, the birth model would then just be an uninformed prior over the intersection of those sets.

$$\mathcal{A} = \bigcap_{i=1}^{c} \mathcal{A}_i \tag{5.45}$$

$$\Gamma = \begin{cases} \frac{1}{\int \mathcal{A} d\mathbf{x}} & \text{if } \mathbf{x} \in \mathcal{A} \\ 0 & \text{if } \mathbf{x} \notin \mathcal{A} \end{cases} \tag{5.46}$$

In Equation (5.46), Γ is the birth model PDF to sample from, and $\mathbf{x}$ is some orbit state. Note that $\mathcal{A}$ must be a bounded set, or the uninformed prior is ill posed. This thought process is directly inspired by work done on admissible regions [64], [20], [96].

The first assumption on our birth model is of an Earth orbiting object. This is done with a maximum on orbit specific energy, $\mathcal{E} < 0$, and a minimum on radius of periapse constraint $r_p > r_{Earth}$. The second assumption is that the object exists in the field of view of the sensor. This is effectively a subset of $\mathbb{R}^3$ defining possible positions at which an object could exist. This birth model does not sufficiently restrict the sample space. The above model attempts to sample what could be all possible Earth orbiting objects. This is computationally infeasible.

This motivates the next few subsections. Often, it is beneficial to design constraints in $\mathbb{R}^4$. Equations (5.45) and (5.46) still apply, but the series of constraints, $\{\mathcal{A}_i\}_{i=1}^{c}$, are de-

fined as subsets of $\mathbb{R}^4$. In some situations, a beneficial constraint may exist in $\mathbb{R}^6$, that is the space of orbits. In such a case, the constraint can be mapped to $\mathbb{R}^4$ through Equation (5.2). Note that this chapter presents several possible birth models which may be appropriate for certain observation campaigns.

5.4.1 Partial Prior Knowledge

In some cases, the orbit of an object or objects, while not completely known, may have some prior knowledge. This includes a previous observation and associated admissible region, a known event (such as a break up event), or objects that actively maneuvered from a previous orbit. If the prior PDF or set can be reasonably represented with a discrete particle approximation, the prior can be used as a birth model.

For example, consider a satellite break-up event at time t_0. There exists a known object at location $\mathbf{r}_0(t_0) \in \mathbb{R}^3$ and velocity $\mathbf{v}_0(t_0) \in \mathbb{R}^3$, known to have broken up at time t_0. It can be assumed that object i from the break up had the same position at the original object's position $\mathbf{r}_0$ at time t_0, and a velocity from a distribution based on a break up model.

$$\mathbf{r}_i(t_0) = \mathbf{r}_0(t_0), \;\; \dot{\mathbf{r}}_i(t_0) = \dot{\mathbf{r}}_0(t_0) + \epsilon \tag{5.47}$$

$$\epsilon \sim p_\epsilon(\mathbf{v}) \tag{5.48}$$

For example, $p_\epsilon(\mathbf{v})$ could be a zero mean Gaussian distribution, or a more complicated distribution, approximating the spread of objects from the break up. Furthermore, a restrictive set of velocities can be calculated by excluding ballistic trajectories. This model can be sampled, as a birth model; all particles should then be propagated through orbital mechanics and a measurement model for the electro-optical sensor to create particles in the correct state space. An admissible region would be done in an identical manner. See [74] for more details about propagation and projection of particles.

This chapter focuses on when prior knowledge is not available, so this topic is not

explored further in this chapter.

5.4.2 No Prior Knowledge

For a Bernoulli filter to reliably track an object, the object must exist in the frame of the image for a minimum number of frames. If an object is in the top left corner of an image with a velocity headed out of the frame, it is impossible for the filter to track it; several frames are typically needed for filter convergence. Instead, a minimum number of observations should be enforced within the birth model. Such a constraint can be related to the to photometric SNR of the object over multiple frames. An object with sufficiently high photometric SNR in a single frame may be identified as existing in that single frame. An object with low photometric SNR in a single frame, may have sufficient photometric SNR over several frames to be identified. Through an assumed minimum per frame photometric SNR, and a minimum multi-frame photometric SNR for detection, a minimum number of needed detections could be defined. Further analysis is out of the scope of this chapter.

This defines a birth model within the image as a uniform distribution over set of admissible position and velocity states. The birth model can be derived in one dimension, because x and y can be sampled independently.

The object position must exist within some range dependent on the size of an image.

$$x \in [x_{min},\ x_{max}] \tag{5.49}$$

Assume that the filter requires a minimum number of observations, n_0, and observations are taken at equal intervals of Δt. Then velocities for a particular position must be within range of values defined in Equation (5.50).

$$\dot{x} \in \left[\frac{x_{min} - x}{n_0 \Delta t},\ \frac{x_{max} - x}{n_0 \Delta t}\right] \tag{5.50}$$

The size of this interval is is shown in Equation (5.51).

$$\frac{x_{max} - x}{n_0 \Delta t} - \frac{x_{min} - x}{n_0 \Delta t} = \frac{x_{max} - x_{min}}{n_0 \Delta t} \tag{5.51}$$

This equation is important because interval size is independent of location. The implication is that all possible values of x have equal probability mass and can therefore be sampled independently of $\dot{x}$. The birth model can be implemented by independently sampling $[x,\ y]$ from Equation (5.49), then sampling $[\dot{x},\ \dot{y}]$ from an appropriately formed Equation (5.50).

5.4.3 Rejection of Inertially fixed Objects

Assume the observer is located on Earth. This allows the question, is there a way to lower bound the relative angular velocity that an object can move at with respect to the observer? Note that relative angular velocity refers to the rotation of the vector between the object and observer relative to the inertial. This analysis enables deliberately omitting stars from the considered state space. This should be an object that is as far away as possible moving at the lowest velocity possible. For an observation campaign looking at the geostationary belt, this corresponds to a geostationary transfer orbit (GTO). For an observation campaign looking at the low Earth orbits, it is a ballistic trajectory.

Under the right circumstances, the minimum velocity constraint on the birth model can remove a large amount of the velocity space that needs to be sampled. In the simulations that are presented in this chapter, this constraint removes approximately 1/3 of the possible velocities. Just as importantly, this constraint does not contain star dynamics in the space of hypothesized signals. Star identification and subtraction are recommended, but any errors in the star identification or subtraction leave residual signal in an image, which inherently provide likelihood to particular particles. In implementation on real data, star identification and subtraction have proven difficult, and filters have often converged to stars. It is therefore

beneficial to ignore stars in the birth model and likelihood function. An example of the implementation of such a model is given in Section 5.5.2.

5.4.4 Hypothesized Orbits

This section briefly describes the process of informing a Bernoulli filter birth model based on an assumed orbit regime. If observations are taken with the intention of looking for a particular class of space objects, the dynamics of those space objects can be assumed, and used as a birth model for a Multi-Bernoulli filter.

As an example consider the Geostionary (GEO) regime. The GEO regime in particular contains debris and man made objects at a range of brightnesses and relative angular rates which are extremely dissimilar to other populated orbits. These objects exist in a range of realistic values for semi major axis and eccentricity which determine the relative angular speed with respect to a ground-based observer. The inclination can also exist within a bounded range of values, which modulates the relative direction of the speed of objects. In this way, a range of relative velocity rates and headings can be defined, and used in a birth model. This analysis is partly based on work done in admissible regions [79]. The results shown on real data in Section 5.7 use this type of birth model.

Other orbit regimes can be analyzed in a similar way.

5.5 Implementation Notes

5.5.1 Multi-Bernoulli Filter Implementation Notes

This section outlines specific heuristics and design choices used for the Bernoulli filter in order to help it converge. The major heuristic edit was added to avoid multiple filters converging on a single object. The assumption is made that the probability of multiple objects existing in the same place at the same time is zero. This assumption is not a novel innovation of this chapter [73]. Based on this assumption, the likelihood function for each Bernoulli filter is set to zero in the area around every other Bernoulli filter's maximum a

posteriori estimate from the previous iteration. Once one Bernoulli filter is tracking one object with good accuracy, this assumption makes it impossible for other filters to track the same object. However, this assumption can lead to some instability when multiple filters are trying to track one object, but none are tracking accurately enough to zero out the correct pixels in the likelihood calculation. This can also lead to problems when multiple tracks overlap, as this is the exact case that breaks the given assumption.

The Multi-Bernoulli filter is started with only 1 filter running. Every three iterations, a new Bernoulli filter is added (sampled entirely from the birth model) and all Bernoulli filters below a probability of existence of 0.05 are deactivated. The first assumption in this section, when combined with too many filters, leads to filter competition which hurts convergence. The gradual addition of Bernoulli filters allows each new filter to have time to search for a new object without competition. The relatively small number of simultaneous filters keep computation time manageable.

The state space chosen for frame-to-frame tracking has 2 dimensions in both image position and velocity. The state, $\mathbf{x}$, is modeled as position and velocity, in the frame, measured in pixels.

$$\mathbf{x} = [x,\ y,\ \dot{x},\ \dot{y}]^T \tag{5.52}$$

The dynamics chosen for frame-to-frame tracking are a simple double integrator.

$$\mathbf{x}_{k+1} = \begin{bmatrix} I_{2\times2} & I_{2\times2}\Delta t \\ 0_{2\times2} & I_{2\times2} \end{bmatrix} \mathbf{x}_k + \mathbf{v} \tag{5.53}$$

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}) \tag{5.54}$$

The process noise for the dynamics uses a covariance matrix of $diag([5, 5, 1, 1]^T)$ in pixels and pixels per second.

Because of computational restriction, the Bernoulli filter typically cannot fully sample

the state space, as defined by the birth model. Particles randomly sampled from a birth model near the true space object in position space typically don't have the correct velocity state, due to particle impoverishment. By artificially inflating the dynamics covariance, state space exploration is encouraged. This leads to better state space exploration and convergence for Bernoulli filters, but worse final tracks generated by the algorithm. There is a trade off which is not further explored in this chapter.

The filter bank uses 20,000 particles per Bernoulli filter. The filters typically perform well for this number of particles, and computation times are bearable for implementation on a personal laptop.

5.5.2 Example of Minimum Velocity Birth Model

A geostationary transfer orbit, in the most extreme case, has a periapsis and apoapsis radius of $r_p = 6471$ km and $r_a = 42164$ km. This gives an apoapsis velocity of

$$v_a = \sqrt{\frac{\mu}{r_a}}\sqrt{\frac{2}{1+\frac{r_a}{r_p}}} = 1.586\frac{\text{km}}{\text{s}} \tag{5.55}$$

where $\mu = 398600\text{km}^3/\text{s}^2$ is the specific gravitation parameter of Earth. By varying the observer's location on Earth, the instantaneous orthogonal component of the velocity changes are under 2%. The distance between the object and observer changes more significantly. The worst case scenario is a distance of 42642 km and orthogonal velocity component of 1.568 km/s. The instantaneous relative angular rate is then

$$\tan^{-1}\left(\frac{1.568\text{ km/s}}{42642\text{ km}}\right) = 0.002107^\circ/\text{s} \tag{5.56}$$

The sensor used in the simulations in this chapter has a field of view of 2 degrees and a resolution of 512 pixels in both directions. In this example, this gives a minimum angular velocity of 0.539 pixels/s.

Table 5.1: Orbital elements for simulated objects.

	$a(km)$	e	$\Omega(deg)$	$i(deg)$	$\omega(deg)$	$f(deg)$
Obj. 1	42165	0.0044	0	0.0017	3.1163	1.5699
Obj. 2	42165	0.0087	4.6862	0.5237	1.5734	4.7063
Obj. 3	42165	0.0087	1.5206	0.1747	4.7417	4.7063

5.6 Simulation Results

5.6.1 Problem Set Up

The multi-Bernoulli filter is shown in this section tracking objects in simulated data. In the simulation, a sensor is modeled as a simple pinhole camera with a 2 degrees field of view and 512 pixel resolution in both x and y. The point spread function (PSF) is approximated as a 5px by 5px Gaussian distribution [97], [98]. Note that for simulation purposes, the measurement images are simulated with the same PSF as the matched filter templates. Three objects are simulated with various orbit characteristics shown in Table 5.1. Note that small amounts of eccentricity and inclination are added to all orbits to give a range of in plane headings and disambiguate orbital elements. An observer is simulated at Georgia Tech $(84.39W, 33.78N)$ on November 12, 2015, at 20:00:00 UTC performing a fixed stare at $\mathbf{r} = [0; -42164; 0]$. This is effectively a sidereal stare tasking.

The three objects are simulated with a constant total flux per exposure. The fluxes for each object are $F_1 = 1000$, $F_2 = 750$, and $F_3 = 500$. The unit and scale of these numbers is ultimately insignificant; the SNR is the ultimate deciding factor for detection, which is unitless, and measured as a relative scale. The images include per pixel normally distributed noise with variance of $\sigma_w^2 = 100$, $w \sim \mathcal{N}(0, 100)$. The total object SNR (for when an object's signal is spread over multiple pixels) can be calculated via [88]

$$\begin{aligned} \text{SNR}\left(\sum y_j(\mathcal{T})\right) &= \frac{\mathbb{E}[\sum y_j(\mathcal{T})]}{\sqrt{\sum \mathbb{E}[(y_j(\mathcal{T}) - \mathbb{E}[y_j(\mathcal{T})])^2]}} \\ &= \frac{1}{\sqrt{N_z}} \frac{F}{\sigma_w} \end{aligned} \tag{5.57}$$

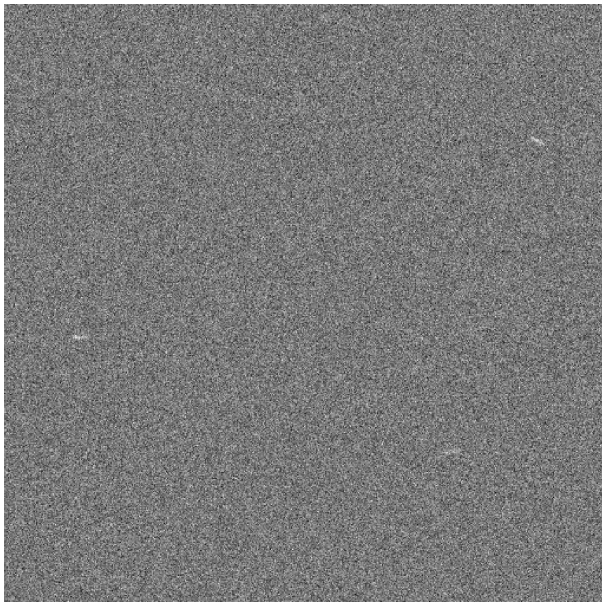

Figure 5.4: Example of a simulated image.

where N_z is the total number of pixels being considered, and F is the total signal from the SO. All objects spread over approximately 80 pixels in the simulations, giving $\text{SNR}_1 = 11.18$, $\text{SNR}_2 = 8.39$, and $\text{SNR}_3 = 5.59$. Note that by the above calculation, the SNR in a single pixel is much lower than the above values. The actual value depends of the spreading of the signal over the pixels. Per pixel, the SNR for each object is around 2.7, 2.2, and 1.5 respectively. Exposures are 10 seconds long and taken every 20 seconds. This flux is blurred over multiple pixels by the movement of the object over the exposure and then by a 5 by 5 normal distribution kernel. The system is simulated for 25 frames (500 s). Object 1 begins in the image and leaves at frame 17 (340 s). Object 2 enters in the frame 12 (240 s) and leaves in frame 23 (460 s). Object 3 enters at frame 9 (180 s) and does not leave. There is no significance to the enter and exit times beyond testing a variety of conditions. Figure 5.4 shows am example of a simulated image.

The simulations uses a combination of two birth models. First, a minimum detections birth model is used, with a minimum number of detections of $n_0 = 10$, outlined in Section 5.4.2. Second, a minimum angular rate is enforced, shown in Section 5.5.2

5.6.2 Multi-Bernoulli Filter on GEO level Objects

Figures 5.5(a), 5.6(a), and 5.7(a) show the error, Figures 5.5(b), 5.6(b), and 5.7(b) show the velocity error, and Figures 5.5(c), 5.6(c), and 5.7(c) show the probability of existence of all Bernoulli filters which track a particular object. Initially, multiple filters may track a single object, due to the random nature of particle filters. Eventually, one filter tracks an object accurately enough, and all other Bernoulli filters diverge; it is explicitly enforced that no two filters can have equal states, as assumed in Section 5.5.1. Figures 5.5(a), 5.5(b), and 5.5(c) show object one which has the highest SNR and is in the field of view from time 0 to 340. This object is tracked with relative ease. Figures 5.7(a), 5.7(b), and 5.7(c) show object three which has the lowest SNR and is in the field of view from time 180 to the end. It takes some time for a filter to accurately converge on this object but once it does, the probability of existence quickly rises to almost 1. Figure 5.8 shows the truth tracks of each object in black with the most successful filters overlaid in gray. Error is calculated in pixels as the norm of the difference between the true position and the of the maximum a posteriori position estimate. The maximum a posteriori estimate, or point of maximum probability density, is subject to pixel level errors due to the difficulty in calculating it; keep in mind that Bernoulli filters in general is multi-modal, so the mean cannot be used. Iterations of the filters before and after they successfully track an object are not included, as these steps consist of the maximum a posteriori estimate wandering through the image.

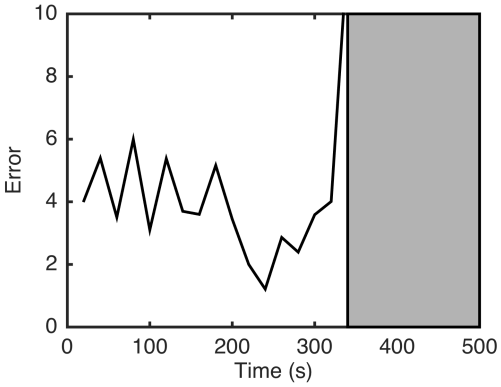

(a) Object 1 Error (pixels) (SNR = 11.18).

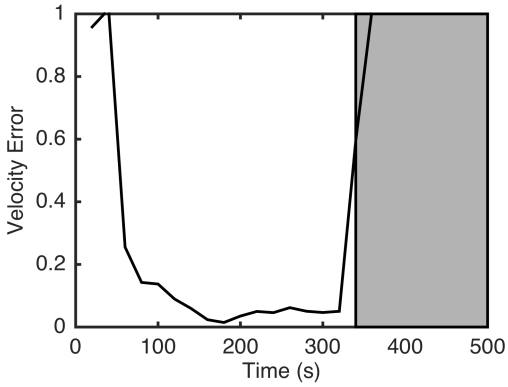

(b) Object 1 Velocity Error (pixels/s) (SNR = 11.18).

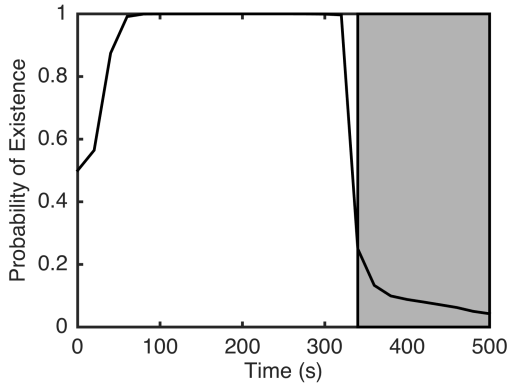

(c) Object 1 Probability of Existence (SNR = 11.18).

Figure 5.5: Tracking results for object one.

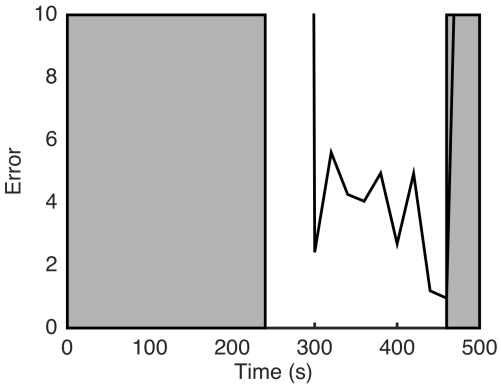

(a) Object 2 Error (pixels) (SNR = 8.39).

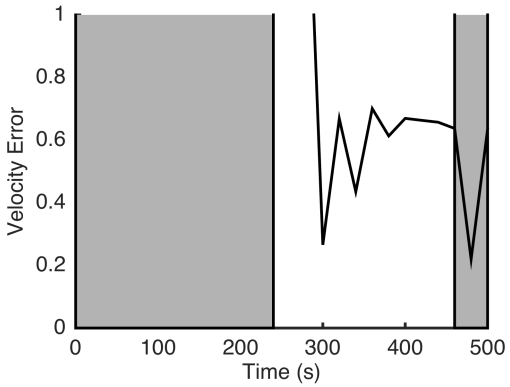

(b) Object 2 Velocity Error (pixels/s) (SNR = 8.39).

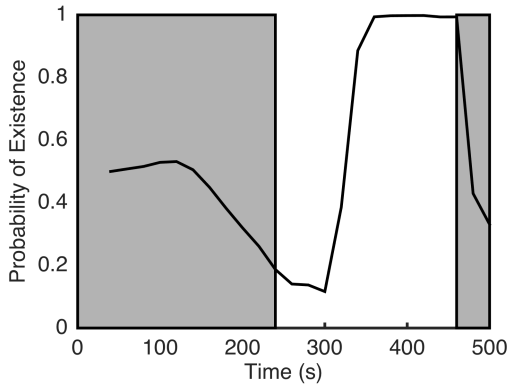

(c) Object 2 Probability of Existence (SNR = 8.39).

Figure 5.6: Tracking results for object two.

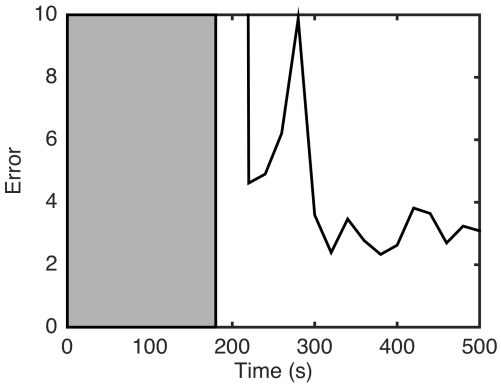

(a) Object 3 Error (pixels) (SNR = 5.59).

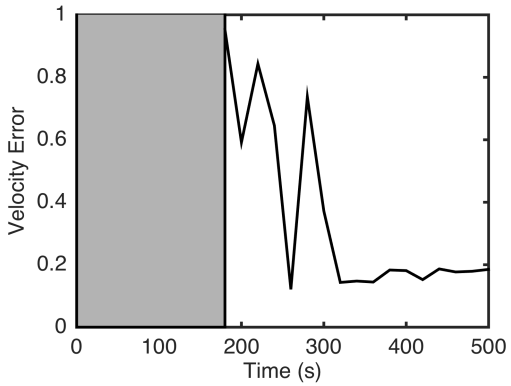

(b) Object 3 Velocity Error (pixels/s) (SNR = 5.59).

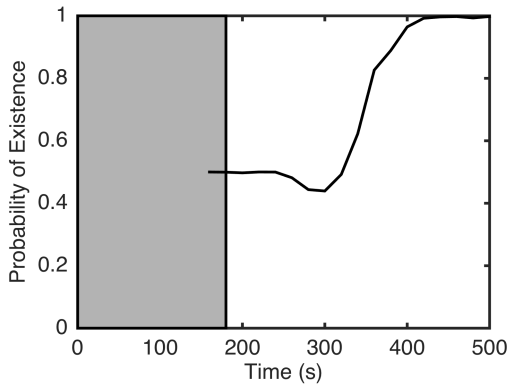

(c) Object 3 Probability of Existence (SNR = 5.59).

Figure 5.7: Tracking results for object three.

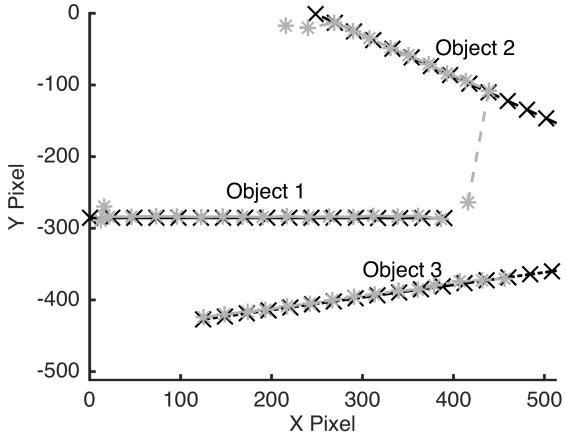

Figure 5.8: Tracking various entering and exiting objects.

5.6.3 Exploration of Limiting SNR of Multi-Bernoulli Filter

The ultimate goal of this research is to detect as low SNR objects as possible. A similar test case is used with three objects over 500 s. All objects spend the entire duration in the field of view, have the same velocity, and vary only in starting position. Objects one, two, and three have total object SNR of 4.47, 3.35, and 2.24, and max per pixel SNR of 1.0, 0.7, and 0.56, respectively. Figures 5.9(a), 5.9(b), 5.10(a), and 5.10(b), show the error and probability of existence of all Bernoulli filters which track objects 1 and 2. Both object 1 and object 2 have multiple filters begin to track the object but none of the filters manage to fully converge on the object. The objects are accurately tracked by the algorithm in terms of the PDF. The probability of existence update, as it is formulated, does not increase from frame-to-frame. No filter succeeds at tracking object 3. The hypothesis likelihood tracks objects down to a total photometric SNR of 3.35. The probability of existence update warrants further work.

Previous work with this likelihood confirms the capability to track down to a total photometric SNR of less than 1 [74], [88]. This previous work was done with a Bayesian particle filter and a space object with prior knowledge. Given enough particles and a constrained enough birth model, such performance should be possible with a multi-Bernoulli filter, though computational feasibility is a concern.

5.7 Empirical Data Results

5.7.1 Geostationary Object, small Field of View

This section presents the validation of the multi-Bernoulli filter on real telescope data. The data was taken by GT-SORT, a 0.5m Raven-class telescope located at Georgia Tech. The field of view is approximately 14 by 11 (arc minutes). The sensor has a resolution of 2736 by 2192 pixels. This specific data set was taken by performing sidereal stare on a region known to contain GEO objects. Images were taken with 1 second exposures, at a rate

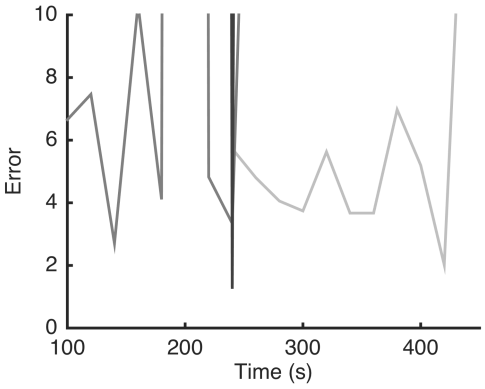

(a) Object 1 Error (pixels) (SNR = 4.47).

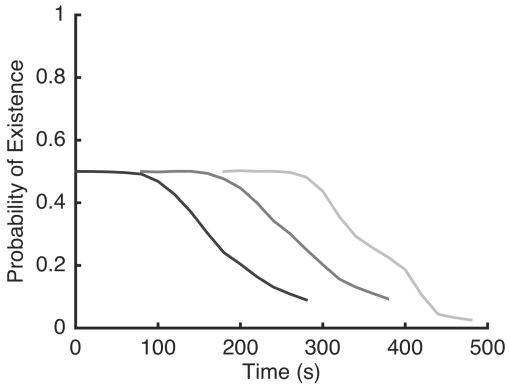

(b) Object 1 Probability of Existence (SNR = 4.47).

Figure 5.9: Tracking results for object four.

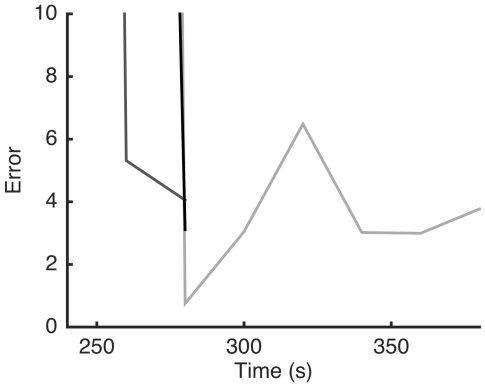

(a) Object 2 Error (pixels) (SNR = 3.35).

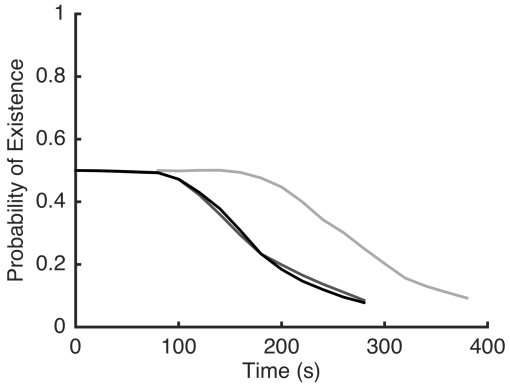

(b) Object 2 Probability of Existence (SNR = 3.35).

Figure 5.10: Tracking results for object five. This is the limiting SNR test case.

of approximately 1 image every 1.6 seconds. Twenty five consecutive images are used for these results. Over the twenty five images, an GEO object starts in the center of the image and eventually exits at the bottom of the image. A second object enters from the top of the image for 12 frames. The zero inclination GEO orbit plane runs approximately vertically through the image. The PSF is estimated from a bright star in the image, and assumed to be invariant between frames and location in the image. The actual object, as it appears in the data, can be seen in Figure 5.13. The images, including Figure 5.13, are re-scaled between $\pm$ 3 standard deviations of the background noise, in order to better visualize outlying signals. Two different birth models are used. The first birth model uses the estimated velocity of a zero inclination geostationary object, to better capture expected objects. The second birth model uses the estimated velocity of a any geostationary object, but doesn't assume a direction of motion.

Using Equations (2.12) and (2.13), the per pixel and total photometric SNRs of this object can be estimated from the images. The per pixel SNR is approximately 0.18, while the total photometric SNR is approximately 6.9. These values should be taken as rough estimates; for objects of 0.18 per pixel SNR, noise is pervasive, and it is difficult to determine which pixels contain signal. Figure 5.13 shows the zoomed in point spread of the object.

The resulting two tracks for the assumed heading birth model are plotted over an image from the data set in Figure 5.11. Similarly, the results for the no assumed heading birth model are plotted over an image from the data set in Figure 5.12. When a Bernoulli filter is initialized, due to particle impoverishment and the inherent randomness of the system, the maximum a posteriori estimate randomly moves around the image. Both estimates are random location in the image for a few frames before the Bernoulli filters converge; this is normal behavior. This is due to the low signal SNR; noise randomly mimics object signal enough to cause random hits occasionally, but the true object signal also follows the dynamic model, allowing convergence. There are also around 5 other Bernoulli filters operating in parallel, which are not shown, due to lack of significant convergence. Stars in

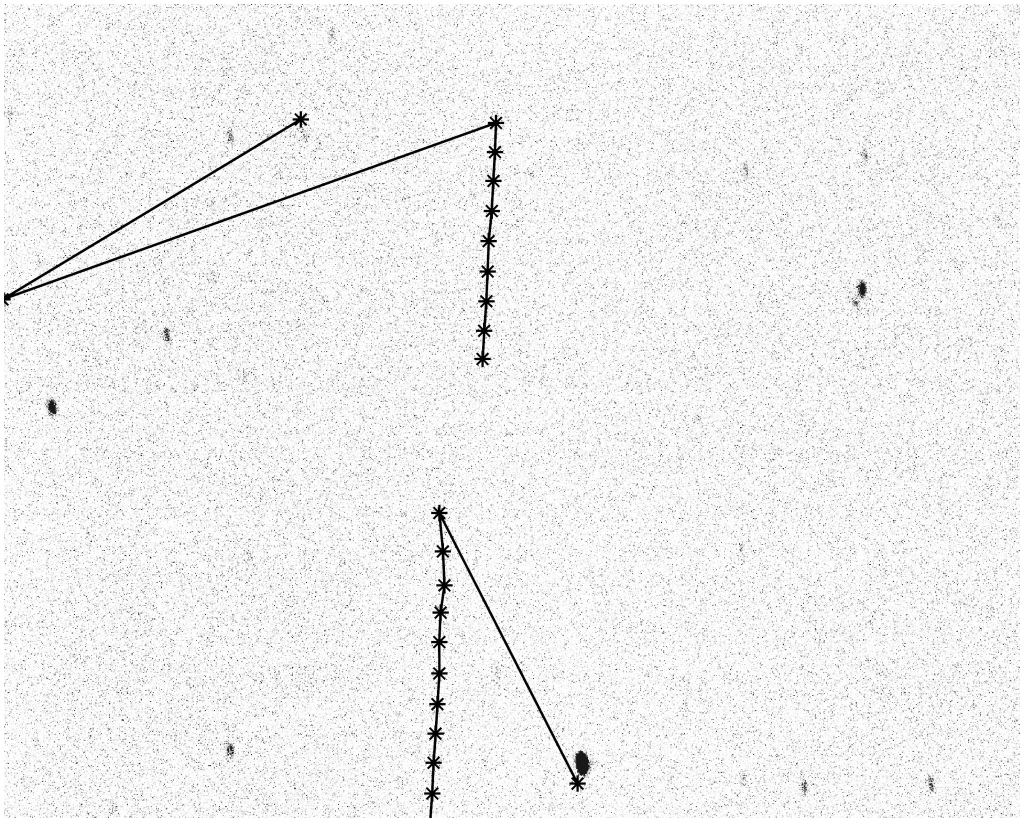

Figure 5.11: Object tracks overlayed on image. 2192 by 2736 pixels. Assumed heading birth model.

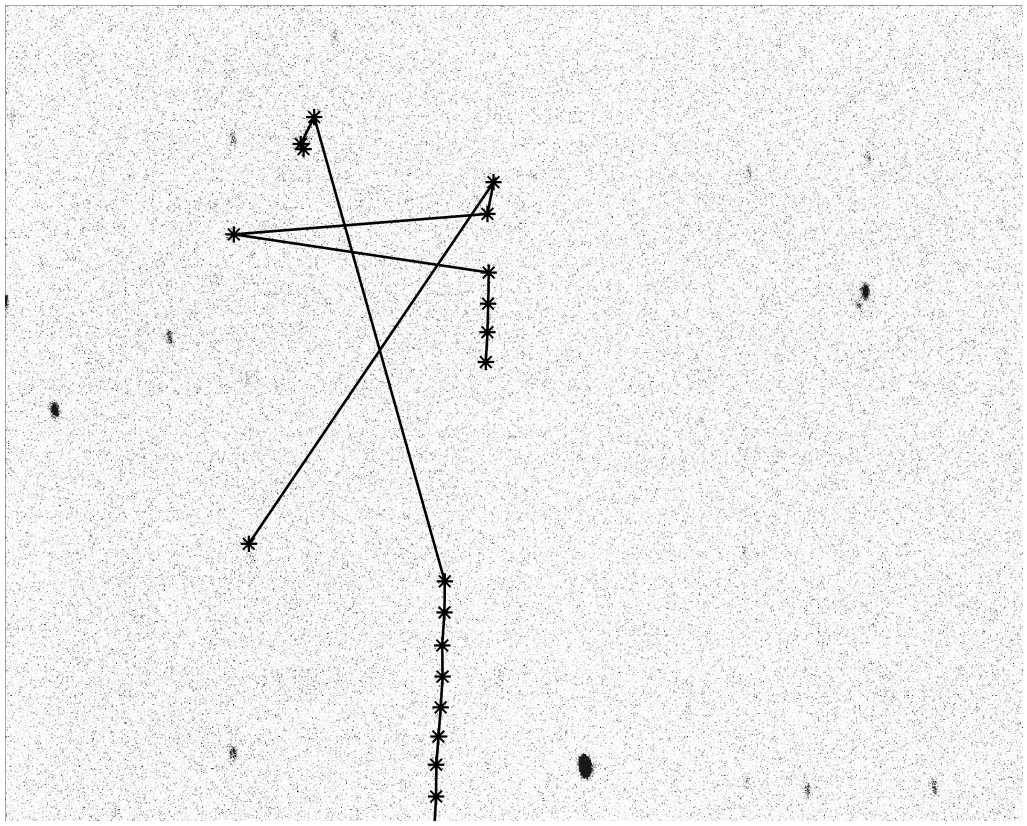

Figure 5.12: Object tracks overlayed on image. 2192 by 2736 pixels. No assumed heading birth model.

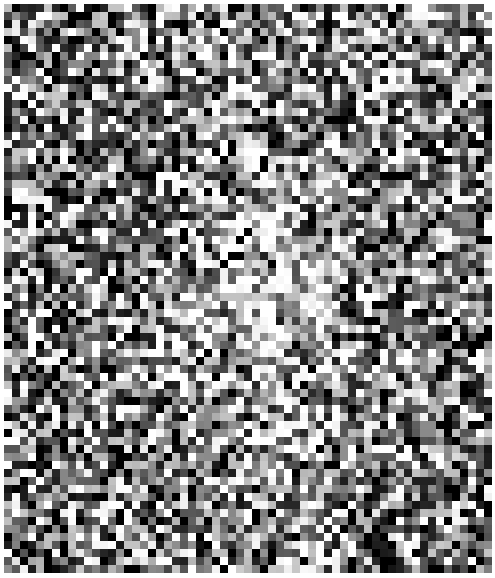

Figure 5.13: Object size and signal strength.

the images are inertially fixed from frame-to-frame. The stars are detected in preprocessing and zeroed out in the image. Stars may appear to have a vertical oblong shape; this is due to a combination of vibrations in the telescope mount, wind buffeting, and atmospheric turbulence. The objects appear somewhere between 8 or 9 arcminutes apart in the images.

5.7.2 Tumbling Low Earth Orbit, Wide Field of View

On the 26th of March, 2016, the Hitomi Astro-H spacecraft lost communication with the ground and experienced a break up due to what is thought to be an unstable spin rate. On the 29th of March, 2016, a series of images were taken with a wide field of view sensor of the Hitomi as it passed over the Georgia Tech observatory. The sensor has a field of view of about 20 degrees, and a resolution of 1024 by 1280 pixels. The object is seen for about 26 frames, over which large variations in photometric magnitude are observered, as the object tumbles. Because observations were taken multiple days after the breakup, no debris objects were seen in the images.

In order to better validate the work in this chapter, the multi-Bernoulli Filter was used to detect objects in this data set. Figure 5.14 shows the observed track laid over a sample

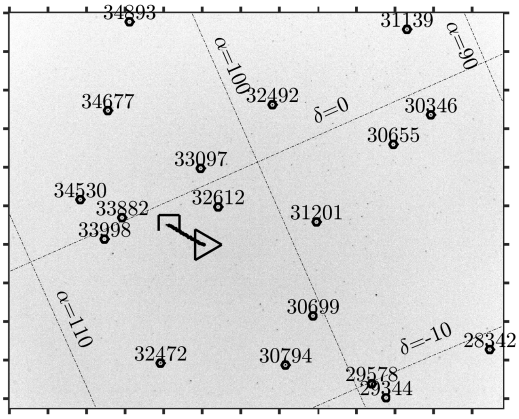

Figure 5.14: Track of Hitomi Astro-H. Triangle is beginning of track.

image. The space object total photometric SNR variations over the iterations are shown in Figure 5.15. The minimum total photometric SNR is 1.5, while the total photometric SNR falls below 3 for at most three consecutive iterations. Star identification has been performed for this data set with the open source astrometry.net algorithm [1].

Filter performance is not hampered by object signal variations. Object probability of existence converges to 1 after approximately 3 iterations, and never decreases. This example highlights an important property of the hypothesis test likelihood; variations in brightness do not affect performance at all, because no approximation of photometric brightness is required.

5.8 Conclusion

The first objective of this work is to implement a multi-Bernoulli filter for tracking SOs which focuses on the uncued detection of as dim as possible space objects. This chapter uses a likelihood function which attempts to allow detection of as dim as possible signals. The birth model is also tailored towards being as restrictive as possible, in an attempt to help particle convergence.

The second goal is to use every possible method to make the multi-Bernoulli filter as

[1] astrometry.net, May 3rd, 2016

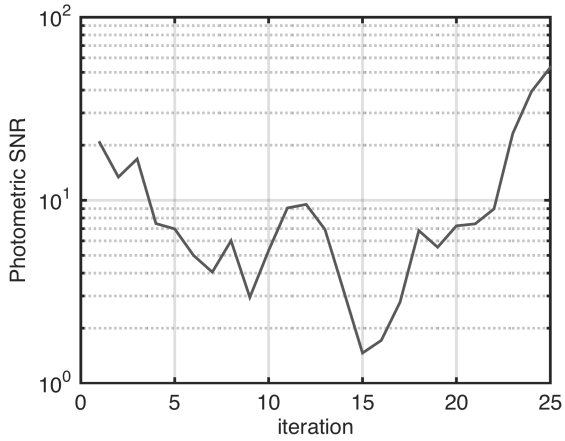

Figure 5.15: SNR variations of Hitomi Astro-H over observation.

efficient as possible. The restrictive birth model attempts to cut down the state space as much as possible, which helps with sample density. Ultimately there is always a trade-off between performance and computation time, as modulated by number of particles used. The primary direction of further research should be in making these methods as computationally tractable as possible.

5.8.1 Unifying Example

At this point in our example, we have a series of images for each of several locations spread across the admissible region prior. The goal now is to search through each image to identify all sources of light. First, trivially, all stars must be tracked. Tracking stars both provides the astrometry analysis with a better basis for star identification and identifies stars singals so they can be ignored by the object identification algorithms. These objects can be located by simply setting the birth model velocity in this chapter's contribution to track stars. Next, the multi-Bernoulli filter is run on the data sets to search for any space objects. Any objects found can easily be associated with the prior admissible region [99]. Detections not associated with the prior are useful to further update the space object catalog. If this method does not detect the space object associated with the prior, contribution 4 provides a more sensitive framework.

CHAPTER 6

SPACE OBJECT DETECTION IN IMAGES USING MATCHED FILTER BANK AND BAYESIAN UPDATE

Electro-optical sensors, when used to track space objects, are often used to produce detections of space objects for some orbit determination scheme. Instead, this chapter proposes a series of methods to use electro-optical images directly in orbit determination. This work uses the SNR optimal image filter, called a matched filter, to search for partially known space objects.

When considering an EOS attempting to measure a partially known object with low SNR, two problems are defined which are addressed in this chapter. First, how can an EOS signal be predicted by partial orbit knowledge and how can that information be used to detect an SO? This chapter approaches this problem by modeling a partially known SO as a set of possible orbits. The set of possible orbits is mapped into the frame of an EOS, providing a framework to rigorously search for an object. By defining a metric for measuring matched filter template similarity, a bank of matched filters is efficiently defined by partitioning the prior knowledge set [66]. Once partitioned sets are known, the matched filter bank can be localized to regions of the sky. A method for hypothesis testing the result of a matched filter for a space object is developed. Finally, a framework for orbit determination based on the matched filter result is developed. Simulation shows that the analytic results enable a better framework for implementing matched filters for low SNR object detection.

Second, how can the results of this MF Bank be incorporated into the existing SSA framework? The detectionless likelihood from the previous chapter is applied in a sequential Bayesian update formulation on the prior information, reformulated as a PDF. Another result which this chapter presents involves reducing the computational burden of the MF

bank. In general, the MF is a discrete correlation of a known signal template with the measurement image. The MF need only be calculated in certain localized areas as defined by a partition. The chapter develops precisely how to define these local search regions, ensuring no wasted computation.

6.1 Methodology

This chapter's primary contributions are as follows. First, this chapter develops a matched filter based on true space object dynamics. This chapter defines a metric which measures how dissimilar two EOS signals are. The topic of set partitioning based on metrics is introduced to the SSA community and applied to create a matched filter bank. A localized matched filter bank is presented to enable search on specific areas of an image based on the partitioning. Matched filter results are then used as the measurement likelihood for a Bayesian update. Finally, the methods in this chapter are applied to detect low SNR and dim space objects with an electro-optical sensor.

The chapter organization is as follows. Section 6.2 introduces the concepts behind an MF which the reader should be familiarized with. Section 6.3.1 defines the problem set up and some of the functions this chapter uses. Section 6.3.2 defines the framework for the partitioning process and the measurement dissimilarity metric. Section 6.3.3 develops a matched filter-based partition for searching for partially known SO. Section 6.3.4 derives the measurement dissimilarity metric. Section 6.3.5 uses the measurement dissimilarity metric to develop an implementation for the partitioning process. Section 6.3.6 develops the localized matched filter process. Section 6.4 gives intuition into how to implement this methodology. Section 6.5 simulates the matched filter bank to demonstrate strengths.

6.2 Matched Filter

This section reviews the concepts of matched filtering, image generation and processing, and mapping of continuous distributions to discrete images.

This chapter assumes an effective background subtraction has been performed on all images, as is common in pre-processing [77]. When effective background subtraction is performed, the mean of the noise is subtracted out, meaning that what is left is zero mean. Note that this can lead to negative pixels; this does not mean that negative photons were received, but rather, less photons were received than the expected value of the noise. The read noise in images can be effectively modeled by a Poisson distribution. When a Poisson distribution's parameter becomes large enough, the distribution becomes approximately Gaussian in shape [95]. Pixel values are integers while a Gaussian distribution is continuous. Using a Gaussian distribution to model pixel noise is therefore an approximation, which becomes more correct with larger pixel values.

First, define a measurement as $\mathbf{Z}(t_0, t_I) \in \mathbb{R}^{d_1 \times d_2}$ which is a matrix of pixel values, which are measured in analogue to digital units (ADU), for an EOS. Assume that the pixel z_i consists of some signal s_i and zero mean Gaussian read noise w_i as defined in Section 2.2 and Equation (2.5).

The values t_0 and t_I represent the beginning time and length of the observer exposure, respectively. This allows the definition of the integration time set,

$$\mathcal{T} = [t_0, t_0 + t_I]. \tag{6.1}$$

Each pixel is an integration of photons (or some unit of measure such as ADU) from a particular direction in the sky, and is therefore a function of the integration time set, $z_i(\mathcal{T})$. The matrix of pixels can be broken into signal and noise

$$\begin{aligned} \mathbf{Z}(\mathcal{T}) &= \mathbf{M}(\mathcal{T}) + \mathbf{W}(\mathcal{T}) \\ \mathbf{M}(\mathcal{T}) \in \mathbb{R}^{d_1 \times d_2},&\ \ \mathbf{W}(\mathcal{T}) \in \mathbb{R}^{d_1 \times d_2}. \end{aligned} \tag{6.2}$$

For SNR, which is discussed and refered to throughout this chapter, a brief discussion is included in Section 2.2, Equations (2.12), (2.13), and (2.14). The MF is predicated on a

hypothesized signal structure, which is referred to as a template, $\mathbf{M}_0(\mathcal{T}) \in \mathbb{R}^{d_3 \times d_4}$, which is a hypothesis of the signal contained in $\mathbf{M}(\mathcal{T})$. $\mathbf{M}_0(\mathcal{T})$ is a smaller or equal dimension matrix than $\mathbf{M}(\mathcal{T})$, that is $d_1 \geq d_3$ and $d_2 \geq d_4$ due to the constraint that the observed true signal cannot be larger than the entire optical frame. It should be noted that $\mathbf{M}_0(\mathcal{T})$ is not always a function of the exposure time, but could be predicated on any number of factors. For optical observation, the signal is dependent on the timing of the exposure making $\mathbf{M}_0(\mathcal{T})$ dependent on the timing of the exposure. A matched filter can then be defined as a mapping $g_{MF} : \mathbb{R}^{d_1 \times d_2} \times \mathbb{R}^{d_3 \times d_4} \rightarrow \mathbb{R}^{d_1 \times d_2}$

$$\mathbf{Z}'(\mathcal{T}) = g_{MF}(\mathbf{Z}(\mathcal{T}), \mathbf{M}_0(\mathcal{T})). \tag{6.3}$$

The matched filter gives an optimal SNR gain for linear image filters [93] [94]. In other words, if the predicted signal $\mathbf{M}_0(\mathcal{T})$ is identical to the measured signal in $\mathbf{M}(\mathcal{T})$, the resulting SNR is maximized. In practice, the function g_{MF} is a discrete correlation of $\mathbf{M}_0(\mathcal{T})$ centered on every pixel in $\mathbf{Z}(\mathcal{T})$. These results may be normalized to make perfect correlation positive one and perfect anti-correlation negative one.

6.3 Theory

6.3.1 Definition of Problem and Dynamics

This subsection presents the background material for the chapter. Consider a space object, with a state $\mathbf{x}(t) \in \mathcal{X}$ composed of position and velocity $\mathbf{r}(t)$ and $\dot{\mathbf{r}}(t)$, seen by an optical observer at position and velocity $\mathbf{o}(t)$ and $\dot{\mathbf{o}}(t)$ along line-of-sight vector, $\boldsymbol{\rho}(t)$, defined in an inertial frame. It should be noted that while this chapter assumes $\mathbf{x}(t) \in \mathbb{R}^6$, as position and velocity, this can be generalized to any formulation of an orbit as well. The observer itself could be either ground-based or space-based.

At an instantaneous time, an EOS measures a unit vector $\hat{\boldsymbol{\rho}}(t)$. This allows the following

partition of the state, derived from the notation used in [20],

$$\begin{aligned}
\mathbf{x}'(t) &= \left[\mathbf{x}_d^T(t)\ \dot{\mathbf{x}}_d^T(t)\ \mathbf{x}_u^T(t)\right]^T \\
\mathbf{x}_d(t) &= \hat{\boldsymbol{\rho}}(t) \in \mathbf{S}(2) \\
\dot{\mathbf{x}}_d(t) &= \dot{\hat{\boldsymbol{\rho}}}(t) \in \mathbb{R}^2 \\
\mathbf{x}_u(t) &= [\rho(t)\ \dot{\rho}(t)]^T \in \mathbb{R}^+ \times \mathbb{R}.
\end{aligned}$$

Note that $\boldsymbol{\rho} = \rho\hat{\boldsymbol{\rho}}$. As a reminder, the unit sphere, $\mathbf{S}(2)$, is being used to represent the set of directions in $\mathbb{R}^3$. It can be represented by unit vectors and can also be thought of as the set of points that lie on the unit sphere or celestial sphere. It is easy to see that $\mathbf{x}'$ is simply a representation of $\mathbf{x} = [\mathbf{r}^T\ \mathbf{v}^T]^T$ in an observer-centric, spherical coordinate frame. Therefore, there exists a transformation between the two representations of the state

$$\mathbf{x}'(t) = \mathbf{m}(\mathbf{x}(t); \mathbf{k}_o(t)) \tag{6.4}$$

where $\mathbf{k}_o(t)$ is a parameter vector containing $\mathbf{o}_o(t)$ and $\dot{\mathbf{o}}_o(t)$, observer o, and any other information needed on the observer. Multiple measurements or a non-zero exposure time allow inference of the rate of change of the unit vector, allowing the following further definition

$$\mathbf{x}'_d(t) = [\mathbf{x}_d^T(t)\ \dot{\mathbf{x}}_d^T(t)]^T. \tag{6.5}$$

In general, $\dot{\mathbf{x}}_d$ is inferred from two $\mathbf{x}_d$ measured at different times. It is useful to define the following instantaneous measurement function, which maps an orbit at a particular time to a measurement as seen by a particular observer, o in Equation (6.6).

$$\mathbf{x}_d(t) = \mathbf{h}_o(\mathbf{x}(t); \mathbf{k}_o(t)) \tag{6.6}$$

This chapter works heavily with sets of orbits which are mapped to a particular EOS

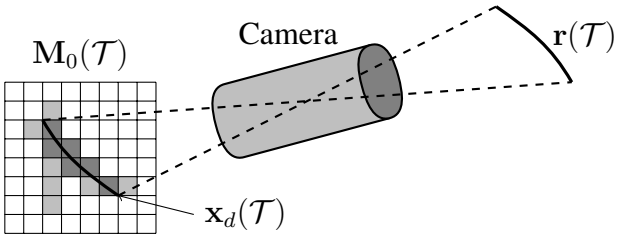

Figure 6.1: Modeling an optical observer and measurement arc in Equation (6.8).

measurement space. In general, if $\mathcal{S}$ is a set of orbits at time t, then

$$\mathcal{S}_o(t) = \mathbf{h}_o(\mathcal{S}(t); \mathbf{k}_o(t)). \tag{6.7}$$

gives these orbits mapped into points in $\mathbf{S}(2)$, centered around $\mathbf{o}_o$. For ease of notation, the parameter vector $\mathbf{k}$ may be dropped from Equation (6.6), but is always implied.

Matched filtering relies on predicting the image that an optical observer produces based on the presence of a particular SO. The actual measurement obtained is the integration of a time-varying line-of-sight unit vector, $\hat{\boldsymbol{\rho}}(t)$, over an exposure time $t \in \mathcal{T} = [t_0, t_0 + t_I]$. Each bin on the EOS obtains photons from a small subset of $\mathbf{S}(2)$. As the line-of-sight vector $\hat{\boldsymbol{\rho}}(t)$ changes over the exposure, it illuminates different pixels, producing an image that the user sees. This process, which can be seen in Figure 6.1, is defined as the mapping

$$\mathbf{M}(\mathcal{T}) = \tilde{\mathbf{h}}_o(\mathbf{x}'_d(t); \mathcal{T}, \mathbf{k}_o(t)). \tag{6.8}$$

which represents the model for a particular observer's mapping from angular space into pixel values. This transformation may include further observer parameters which are also contained in the observer parameter vector, k_o.

The prediction of objects also requires modeling the movement of objects through space. This chapter uses the following flow function to map $\mathbf{x}(t)$ through time

$$\mathbf{x}(t) = \boldsymbol{\phi}(t; \mathbf{x}(t_0), t_0). \tag{6.9}$$

For a typical SO this is practically done by integrating the nonlinear dynamics with a differential equation solver.

The admissible region (AR) is used throughout this chapter as an example to illustrate parts of the discussion.

A typical optical measurement of a space object contains good knowledge of angle and angle rates known as observable states, $\mathbf{x}_d \in \mathbb{R}^4$. An orbit requires six disparate data types to be fully constrained. This AR formulation uses range and range rate as the final two states, known as undetermined states $\mathbf{x}_u \in \mathbb{R}^2$. It is assumed that no information can be reliably used to determine these states from measurements. An AR is created by enforcing a series of constraints of the form

$$g_o(\mathbf{x}_d, \mathbf{x}_u; \mathbf{k}, t) \leq 0 \tag{6.10}$$

$$g_o(\mathbf{h}^{-1}(\mathbf{y}; \mathbf{k}, t), \mathbf{x}_u; \mathbf{k}, t) \leq 0. \tag{6.11}$$

It should be noted that these constraints can be thought of as hypotheses. The AR is then the space where all hypotheses are true. When an AR is represented as a set, the notation $\mathcal{AR}$ is used. Next, we define an admissible region, $\mathcal{AR}_o \in \mathbb{R}^2$, predicated on g_o

$$\mathcal{AR}_o := \left\{ \mathbf{x}_u | g_o(\mathbf{h}^{-1}(\mathbf{y}; \mathbf{k}, t), \mathbf{x}_u; k, t) \leq 0 \right\}. \tag{6.12}$$

The most common hypotheses are a closed earth orbit [79] and a minimum radius of periapse to avoid ballistic trajectories [65].

6.3.2 General Matched Filter Primed by Previous Orbital Knowledge

The first contribution of this chapter is proposing an MF for orbital tracking applications that does not require a constant velocity assumption. Assume there exists a hypothesis orbit of an SO, $\mathbf{x}(t_0)$. This hypothesis can then be mapped to the measurement frame of an observer o via Equation (6.6). This hypothesis can also be mapped forward in time using

the known flow function from Equation (6.9), giving

$$\mathbf{x}_d(t) = \mathbf{h}_o(\boldsymbol{\phi}(t; \mathbf{x}(t_0), t_0), \mathbf{k}_o(t)). \tag{6.13}$$

In reality, the result of interest is to predict the object's position trajectory over a known integration time. By varying t, a hypothesized orbit, $\mathbf{x}(t)$, can be mapped to any point in the integration time of an observer. This hypothesis can then be used to calculate the orbit arc $\mathcal{A}_r$ through $\mathbb{R}^3 \times \mathcal{T}$,

$$\begin{aligned} \mathcal{A}_r = \{&[\mathbf{r}^T(t), t]^T \in \mathbb{R}^3 \times \mathcal{T} : \\ &\mathbf{r}(t) = [I_{3\times 3}\ 0_{3\times 3}]\ \boldsymbol{\phi}(t; \mathbf{x}(t_0), t_0), t \in \mathcal{T}\}. \end{aligned} \tag{6.14}$$

This is simply the position trajectory of the orbit over the integration time, with associated time history. $\mathcal{A}_r$ is labeled in Figure 6.2. In reality, an optical observer does not observe the true arc of the orbit, but rather the arc through angular space, or $\mathbf{S}(2)$. Similarly, $\mathcal{A}_m$ can be defined as an arc through $\mathbf{S}(2)$ in Equation (6.15).

$$\begin{aligned} \mathcal{A}_m = \{&[\mathbf{x}_d^T(t), t]^T \in \mathbf{S}(2) \times \mathbb{R} : \\ &\mathbf{x}_d(t) = \mathbf{h}_o(\boldsymbol{\phi}(t; \mathbf{x}(t_0), t_0), \mathbf{k}(t)), t \in \mathcal{T}\} \end{aligned} \tag{6.15}$$

This measurement arc is the arc $\mathcal{A}_r$ mapped through Equation (6.6). The measurement arc and this mapping process can be visualized in Figure 6.2. While the notation of mapping $\mathcal{A}_r$ to $\mathcal{A}_m$ is useful for visualization purposes, the ultimate purpose of this discussion is to map an orbit to a measurement arc. The following mapping is defined

$$\mathcal{A}_m = \mathbf{h}_o(\mathbf{x}(t); \mathcal{T}) \tag{6.16}$$

which maps an orbit, $\mathbf{x}(t)$, with an associated integration time, $\mathcal{T}$, to a measurement arc,

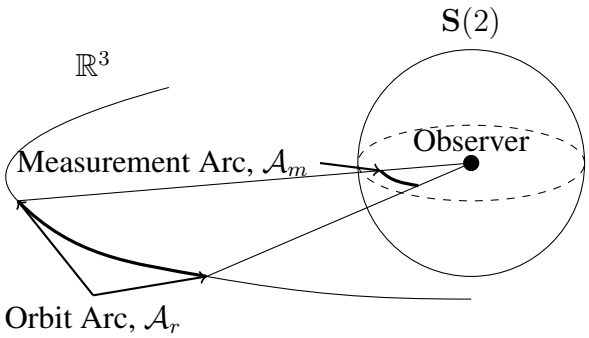

Figure 6.2: Mapping of orbit to measurement arc.

$\mathcal{A}_m$. The use of $\mathbf{h}_o$ in Equation (6.16) is a slight abuse of notation; it is assumed that the orbit is integrated over the exposure time to allow this new mapping. It is assumed that when $\mathbf{h}_o$ is used, time history is preserved alongside the position history. The time history is needed in Subsection 6.3.4.

In addition, the measurement arc is ultimately mapped to a template which exists as predicted EOS pixel values so Equation (6.8) should be redefined as

$$\mathbf{M}(\mathcal{T}) = \tilde{\mathbf{h}}_o(\mathcal{A}_m, \mathbf{k}_o). \tag{6.17}$$

In practice, $\mathcal{A}_m$ should be discretized over the integration time, and each discrete point mapped into a pixel. It is also worth noting that $\mathcal{A}_m$ is sensor specific. Perhaps a better notation would be $\mathcal{A}_{m,i}$, where o denotes the sensor. It should be clear which sensor is being considered throughout this chapter, so the o notation is omitted. Modeling phenomenology such as atmospheric blurring and distortion models is advisable but not required. This MF template is built directly from the hypothesized SO signal structure, accounting for observer motion, and has no built in linearization.

Equation (6.16) represents the culmination of the results in this subsection. This measurement arc, $\mathcal{A}_m$ is a mathematically rigorous equivalent to the measurement an observer would take. This subsection lays down a framework for creating templates based on orbits. The final result is an image that models what an EOS should see. In general, there are more

considerations that should be included in template generation. This can include but is not limited to SO apparent brightness and subpixel localization. These processes are discussed in Section 6.4.

This chapter discusses sets of measurement arcs. Typically, these result from mapping sets of orbits through Equation (6.15). This chapter therefore defines $\mathbb{A}$ for a given t_0, t_I as the set of all possible measurement arcs

$$\mathbb{A} = \{\mathcal{A}_m = \mathbf{h}_o(\mathbf{x}(t); \mathcal{T}) : \mathbf{x} \in \mathcal{X}\}. \tag{6.18}$$

This is convenient for defining subsets of $\mathbb{A}$ in Subsection 6.3.4.

6.3.3 Matched Filter Bank Primed by Prior Distribution

This subsection motivates the need for a bank of matched filters, provides an explanation for how partitions can correctly define a bank of matched filters, and finalizes the definition of this partitioning. It should be noted that, without loss of generality, a single observer is being considered from here on, so the o notation is dropped.

Consider prior information in the form of a set of orbits, $\mathcal{S}_0$, which is referred to as the prior set. By creating s non-overlapping subsets which cover $\mathcal{S}_0$, a partitioning can be created for the prior set, seen in Equations (6.19), (6.20).

$$\mathcal{S}_0 = \bigcup_{k=1}^{s} \mathcal{S}_k \tag{6.19}$$

$$\mathcal{S}_l \cap \mathcal{S}_k = \emptyset, \forall l \neq k. \tag{6.20}$$

It is not necessary that the partitioning be non-overlapping, and certain partitioning schemes may benefit from this condition being relaxed. The intention is to define each partition set, $\mathcal{S}_k$, such that all orbits from this set produce templates nearly identical to

some representative template, $\mathbf{M}_k$

$$\mathcal{S}_k = \{\mathbf{x}(t) : \mathbf{M}_k \approx \tilde{\mathbf{h}}(\mathbf{h}(\mathbf{x}(t); \mathcal{T}, \mathbf{k}))\}. \tag{6.21}$$

It is generally the case that a single subset $\mathcal{S}_k \subseteq \mathcal{S}_0$ does not contain all possible orbits from a prior distribution implying a series of sets, $\{\mathcal{S}_k\}_{k=1}^{k=s}$. This is the basis on which a bank of matched filters is defined.

Next, an analysis of the contents of $\mathcal{S}_k$ is presented. Consider a nominal orbit, $\bar{\mathbf{x}}$, and its particular measurement arc, $\bar{\mathcal{A}}_m$, with arbitrary time and observer parameters

$$\bar{\mathcal{A}}_m = \mathbf{h}(\bar{\mathbf{x}}(t); \mathcal{T}). \tag{6.22}$$

The measurement mapping defined in Equation (6.6) maps a position in $\mathbb{R}^3$ to the two dimensional space $\mathbf{S}(2)$ through the equation

$$\boldsymbol{\rho} = \mathbf{r} - \mathbf{o} \tag{6.23}$$

$$\hat{\boldsymbol{\rho}} = \frac{\boldsymbol{\rho}}{\|\boldsymbol{\rho}\|} \tag{6.24}$$

where $\|\cdot\|$ is the 2-norm. Note that while $\hat{\boldsymbol{\rho}}$ is a vector in $\mathbf{R}^3$, it only has two free dimensions as an element of $\mathbf{S}(2)$. Because this is the mapping on which the measurement arc is based, different $\boldsymbol{\rho}$ inputs which vary only in magnitude all map to a single $\hat{\boldsymbol{\rho}}$. Therefore there must exist a non-trivial set of orbits, $\bar{\mathcal{S}}$, with corresponding orbit arcs which all map to the same measurement arc

$$\bar{\mathcal{S}} = \{\mathbf{x} \in \mathcal{X} : \bar{\mathcal{A}}_m = \mathbf{h}(\mathbf{x}(t); \mathcal{T})\}. \tag{6.25}$$

This is an inversion of the admissible region approach; instead of having a measurement and creating a set of consistent orbits, consistent orbits are being grouped together to

predict a measurement. A single partition, defined in this way, is illustrated in Figure 6.3. Small enough variations in $\bar{\mathcal{A}}_m$ provide effectively identical templates. In order to have an effective matched filter, the template does not necessarily need to be perfect; it may benefit a user to have fewer templates with an increased risk of imperfect matching but decreased computation time. It is possible to define a new partition set,

$$\mathcal{S}_k = \{\mathbf{x} \in \mathcal{S}_0 : d(\bar{\mathcal{A}}_m, \mathbf{h}(\mathbf{x}(t); \mathcal{T})) \leq d_{max}\}. \tag{6.26}$$

where $d : \mathbb{A} \times \mathbb{A} \rightarrow \mathbb{R}^+$ defines a metric which evaluates the difference between two measurement arcs. Assuming an appropriate metric exists, all orbits in $\mathcal{S}_k$ produce measurement arcs similar enough to $\bar{\mathcal{A}}_m$.

Prior knowledge is often some probability density function (PDF), available from an a previous estimate. In the extreme case of object discovery, this PDF may be large, consisting of an admissible region, or a large Gaussian mixture or discrete particle distribution. For this analysis, it is necessary to define a prior set of orbits, $\mathcal{S}_0$. In the case of a PDF, this could be thought of as the interior of an iso-probability surface of the PDF which is used to create this prior set, $\mathcal{S}_0$

$$\mathcal{S}_0 \triangleq \{\mathbf{x} \in \mathcal{X} : f(\mathbf{x}) > f_{min}\} \tag{6.27}$$

where the function $f : \mathcal{X} \rightarrow \mathbb{R}^+$ is the prior PDF over the state space.

The next subsection goes on to define a suitable metric, d, for Equation 6.26.

6.3.4 Measurement Dissimilarity Metric

A bank of MF are desirable for large prior distributions, but due to computation time it is also useful to choose a minimal number of templates. This subsection is concerned with creating a framework to allow similar templates to be grouped together. Conversely, templates which are sufficiently different in shape should not be grouped together. Therefore,

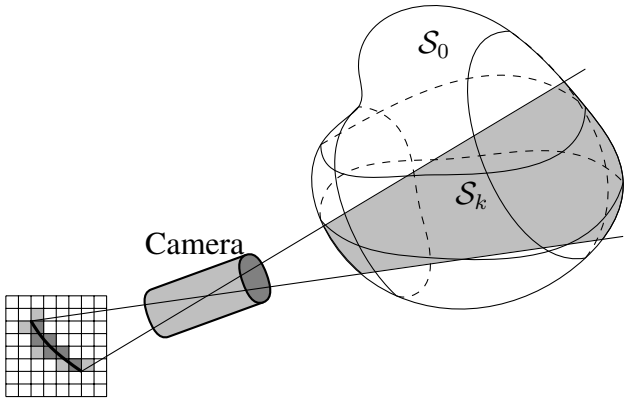

Figure 6.3: A non trivial region of orbits all map to a single MF template.

a measure of similarity of templates is necessary. As it has been shown before, the measurement arc $\mathcal{A}_m$ defined in Equation (6.15) is the continuous space equivalent to the MF template, $\mathrm{M}_0(t)$. There could be many ways to compare similarity of measurement arcs, but a metric has many desirable properties. The most important feature of metrics which motivates this discussion is the existence of a large body of work on partitioning metric spaces. Previous work on partitioning metric spaces is discussed in Subsection 6.3.5, after the metric to be used has been defined [66].

The first step is to choose an existing metric to compare arc similarity. First, consider the similarity of two points on the celestial sphere, $\hat{\boldsymbol{\rho}}_a(t)$ and $\hat{\boldsymbol{\rho}}_b(t)$. The distance between points on the surface of a sphere, along the surface of the sphere, is called the great-circle distance or orthodromic distance and is defined as

$$\Delta\theta_{a,b}(t) = 2\sin^{-1}\left(\frac{||\hat{\boldsymbol{\rho}}_a(t) - \hat{\boldsymbol{\rho}}_b(t)||}{2}\right). \tag{6.28}$$

This particular formulation of the orthodromic distance assumes $\hat{\boldsymbol{\rho}}_a(t)$, $\hat{\boldsymbol{\rho}}_b(t)$ are both unit vectors. The unit vector assumption makes orthodromic distance equivalent to the angle between the unit vectors, in radians. A metric which compares two measurement arcs is desirable. When considering the entire measurement arc, integrating $\Delta\theta_{a,b}(t)$ over a predefined exposure time, $[t_0, t_0 + t_I]$ takes into account how two measurement arcs differ

over the relevant time interval. Thus the measurement dissimilarity metric (MDM) for a given $[t_0, t_0 + t_I]$ is defined in Equation (6.29).

$$d_{MDM}(\hat{\boldsymbol{\rho}}_a(t), \hat{\boldsymbol{\rho}}_b(t)) = \int_{t_0}^{t_0+t_I} \Delta\theta_{a,b}(\tau)d\tau \tag{6.29}$$

This metric can be wrapped into the measurement arc notation defined in Subsection 6.3.2. Given a particular measurement arc, $\mathcal{A}_{m,a}$, defines a time-varying unit vector, $\hat{\boldsymbol{\rho}}_a(t)$ over a predefined time interval $[t_0, t_0 + t_I]$. Therefore, the MDM can be directly computed between two measurement arcs,

$$d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,b}) = \int_{t_0}^{t_0+t_I} \Delta\theta_{a,b}(\tau)d\tau. \tag{6.30}$$

Lemma 6.3.1. *The measurement dissimilarity metric defined in Equation* (6.30) *over a fixed time interval,* $[t_0, t_0 + t_I]$*, is a metric on* $\mathbb{A}$*, the set of measurement arcs, for any pair* $\hat{\boldsymbol{\rho}}_a(t) \in \mathcal{A}_{m,a}, \hat{\boldsymbol{\rho}}_b(t) \in \mathcal{A}_{m,b}$*, that is, it satisfies the following:*

1. $d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,b}) \geq 0$ (non-negativity)
2. $d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,b}) = 0$ *if and only if* $\hat{\boldsymbol{\rho}}_a(t) = \hat{\boldsymbol{\rho}}_b(t)$ (coincidence)
3. $d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,b}) = d_{MDM}(\mathcal{A}_{m,b}, \mathcal{A}_{m,a})$ (symmetry)
4. $d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,c}) \leq d_{MDM}(\mathcal{A}_{m,a}, \mathcal{A}_{m,b}) + d_{MDM}(\mathcal{A}_{m,b}, \mathcal{A}_{m,c})$ (triangle inequality)

Proof. Non negativity follows directly from the range of the functions involved,

$$
\begin{aligned}
||\hat{\boldsymbol{\rho}}_a - \hat{\boldsymbol{\rho}}_b|| &\in [0, 2] \\
\sin^{-1}\left(\frac{||\hat{\boldsymbol{\rho}}_a(t) - \hat{\boldsymbol{\rho}}_b(t)||}{2}\right) &\in [0, \pi/2] \\
\Delta\theta_{a,b}(t) &\geq 0 \\
\int_{t_0}^{t_0+t_I} \Delta\theta_{a,b}(t)dt &\geq 0.
\end{aligned} \tag{6.31}
$$

For coincidence, first note that $\int_{t_1}^{t_2} f(t)dt = 0 \iff f(t) = 0,\ \forall t \in [t_1, t_2]$, for any non-negative continuous function $f(t)$ and arbitrary values of t_1 and t_2. Next, note that over the range of inputs, $x \in [0, 1], \sin^{-1}(x) = 0 \iff x = 0$, and is non-negative as shown above. This directly implies 2. This is equivalent to saying the only way the measurement dissimilarity metric can be zero is if the measurement arcs are identical.

For symmetry, note that $||\hat{\boldsymbol{\rho}}_a - \hat{\boldsymbol{\rho}}_b|| = ||\hat{\boldsymbol{\rho}}_b - \hat{\boldsymbol{\rho}}_a||$ for an arbitrary pair of unit vectors. 3 follows immediately.

The shortest distance along the surface of a sphere between two points is the orthodromic distance. Therefore $\Delta\theta_{a,c}(t) \leq \Delta\theta_{a,b}(t) + \Delta\theta_{b,c}(t)$ must be true, because $\Delta\theta_{a,b}(t) + \Delta\theta_{b,c}(t)$ represents another path along the surface of a sphere from a to c. Remember the following property of integrals

$$
f_a(t) \leq f_b(t) \Rightarrow \int_{t_1}^{t_2} f_a(t)dt \leq \int_{t_1}^{t_2} f_b(t)dt.
$$

which is true for any t_1 and t_2, so

$$
\int_{t_0}^{t_0+t_I} \Delta\theta_{a,c}(t)dt \leq \int_{t_0}^{t_0+t_I} \Delta\theta_{a,b}(t)dt + \int_{t_0}^{t_0+t_I} \Delta\theta_{b,c}(t)dt. \tag{6.32}
$$

This gives the triangle inequality. □

Thus, the MDM is a metric on $\mathbb{A}$, the set of measurement arcs, allowing partition methods to be applied to $\mathbb{A}$. The desired partitioning, however, occurs on orbits which exist as elements in $\mathcal{X}$. The MDM can also be defined over a set of orbits because of the onto mapping from orbits to measurement arcs, Equation (6.16). This can also be shown by acknowledging that $\hat{\boldsymbol{\rho}}_a(t)$ is a function of $\mathbf{x}(t)$, so

$$\begin{aligned} \Delta\theta_{a,b}(t) &= 2\sin^{-1}\left(\frac{||\hat{\boldsymbol{\rho}}_a(\mathbf{x}(t)) - \hat{\boldsymbol{\rho}}_b(\mathbf{x}(t))||}{2}\right) \\ d_{MDM}(\mathbf{x}_a, \mathbf{x}_b) &= \int_{t_0}^{t_0+t_I} \Delta\theta_{a,b}(\tau)d\tau. \end{aligned} \tag{6.33}$$

The following corollary shows that the MDM is a pseudo-metric on $\mathcal{X}$, or any set of orbits.

Corollary 6.3.2. *The measurement dissimilarity metric defined in Equation* (6.29) *over a fixed time interval, $t \in [t_0, t_0 + t_I]$, is a pseudo-metric on $\mathcal{X} \times \mathbb{R}$ for any pair of orbits $\mathbf{x}_a(t), \mathbf{x}_b(t)$, that is, it satisfies the following:*

1. $d_{MDM}(\mathbf{x}_a(t), \mathbf{x}_b(t)) \geq 0$ (non - negativity)
2. $d_{MDM}(\mathbf{x}_a(t), \mathbf{x}_b(t)) = d_{MDM}(\mathbf{x}_b(t), \mathbf{x}_a(t))$ (symmetry)
3. $d_{MDM}(\mathbf{x}_a(t), \mathbf{x}_c(t)) \leq d_{MDM}(\mathbf{x}_a(t), \mathbf{x}_b(t)) + d_{MDM}(\mathbf{x}_b(t), \mathbf{x}_c(t))$ (triangle inequality)

Proof. Non negativity and symmetry follow from Lemma 6.3.1. Every orbit maps to a measurement arc which exists on $\mathbb{A}$.

First fix three orbits, $\mathbf{x}_a$, $\mathbf{x}_b$, and $\mathbf{x}_c$. Because the mapping from orbits to measurement arcs is onto, these orbits map to single time-varying unit vectors, $\hat{\boldsymbol{\rho}}_a(t)$, $\hat{\boldsymbol{\rho}}_b(t)$, and $\hat{\boldsymbol{\rho}}_c(t)$. The triangle inequality follows from Lemma 6.3.1. □

The pseudo-metric does not have a coincidence property because the mapping from orbits to measurement arcs is not one-to-one. However, Subsection 6.3.5 discusses metrics

rather than pseudo-metrics. In order to enable metric space partitioning based on pseudo-metrics, coincident objects are considered as one object [68]. This is equivalent to saying all orbits that produce the identically same streak are the same object. Such a simplification is the inversion of the admissible region argument. Instead of having a measurement and creating a set of consistent orbits, consistent orbits are being grouped together to predict a measurement. Each $\mathcal{AR}$ is an individual object in the space to be partitioned. This definition is a mathematical necessity, but does not change the partition process in the next subsection. A prior set could be redefined by grouping admissible regions analytically, but such an analysis is unnecessary and not included in this chapter.

One question remains: is the MDM robust to the presence of sensor dynamics and distortion? If a sensor is rotating, all objects in a measurement frame appear to arc and rotate through the measurement frame. Furthermore, this rotation can be different depending on where an object is in the sensor frame. Sensor distortion can alter streaks as a function of where they occur in the frame. Different templates are needed for different localized areas. Equation (6.17) includes an assumed dependence on $\mathbf{k}$, the parameter vector which includes $\mathbf{o}, \dot{\mathbf{o}}$. The parameter vector can be expanded to include parameters on local sensor distortion and observer rotation to address this. For typical distortion and rotation models, the number of distinct templates should be minimal. The error produced by not including a correction should be small for most narrow field of view sensors. The localization analysis in Subsection 6.3.6 should further reduce this problem. An in-depth analysis of these effects and how best to localize them is out of the scope of this chapter.

6.3.5 Metric-Based Partitioning

The properties of a metric seem obvious and trivial at first glance, but in fact a wide variety of partition methods exist. A good discussion of metric-based search and partition methods can be found in the text *Similarity Search: The Metric Space Approach* by Zezula et al. [66]. An overview of the concepts and methods is discussed here. The broad approach falls

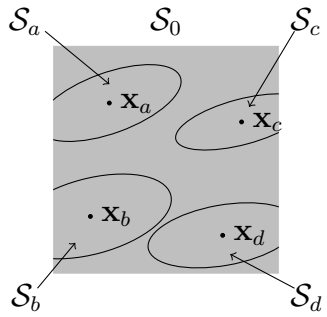

Figure 6.4: Ball partition and generalized hyperplane methods [66].

into two categories, ball partition (BP) methods and generalized hyperplane tree (GHT) methods [67] [68]. The literature is rich with implementations of these methods [69], [70]. Most variations of methods are attempts to improve methods for particular situations or overcoming general pitfalls. The bulk of these methods are simply a reworked BP or GHT method. This subsection introduces the BP method, which is used for the results of this chapter.

BP is based around a predetermined maximum distance, d_{max}. It can be visualized in Figure 6.4. First, a maximum dissimilarity of two measurements, d_{max}, is defined as the sufficient difference for templates to be unique. Once a d_{max} is set, BP is straightforward. BP is implemented by choosing a series of pivot points, $\mathbf{x}_k \in \mathcal{S}_0$. Each partition $\mathcal{S}_a$ is a ball centered at a respective pivot, $\mathbf{x}_a$, with a radius d_{max},

$$\mathcal{S}_k = \{\mathbf{x} \in \mathcal{S}_0 : d_{MDM}(\mathbf{x}, \mathbf{x}_k) \leq d_{max}\}. \tag{6.34}$$

In practice, the set is represented as a discrete point-wise approximation and the points are sorted. This can be done as a single run, with a large number of pivots, or recursively, gradually adding pivot points.

In this process, a few questions still exist: what should d_{max} be set as and why, and how fine does the discrete point-wise approximation need to be?

A higher value of d_{max} results in fewer partitions, while a lower d_{max} gives more par-

titions. The accuracy of the templates is dependent on how large the chosen d_{max} is. The computation time is directly proportional to the number of templates chosen as well. The correct value of d_{max} is also highly dependent on the characteristics of the sensor being used. Therefore, a "best" d_{max} is difficult to define and problem dependent.

A proposed way is to assure that the two measurements with $d_{MDM} \leq d_{max}$ appear essentially the same to an observer. The choice of d_{max} should be based on the pixel size of the sensor $\delta_{\mathbf{x}}$, the exposure time, t_{exp}, and length of the measurement arc L. The MDM as derived gives arc dissimilarity integrated over a time period. The threshold should instead give weight to longer streaks. What is desired is that, at most, a single pixel has an error of length of one pixel. This leads to the formulation, which normalizes the MDM by the average object speed in the image, $\frac{L}{t_{exp}}$, and upper bounds it be the pixel size,

$$d_{MDM}(\hat{\boldsymbol{\rho}}_1(t), \hat{\boldsymbol{\rho}}_2(t)) \left(\frac{L}{t_{exp}}\right) < \delta_{\mathbf{x}}. \tag{6.35}$$

L is the approximate length of the expected streak, in number of the pixels, which allows d_{max} to be normalized for larger or smaller length streaks. This value in practice is very restrictive, and for this reason has not be implemented in the simulation results section. More analysis on how to analytically define a threshold for this process is needed.

The methods described thus far rely on a discrete point-wise approximation. As with any such approximation, sufficient particles are needed to allow the approximation to be representative of the set. This implies highly dimensional sets are difficult to model. Exploiting sub-spaces of lower dimensionality is recommended, such as admissible regions. Insufficient particles can also lead to poorly defined partition set boundaries.

6.3.6 Localized Matched Filtering

The matched filter involves a discrete correlation of the template over the measurement image. The matched filter can be applied over the entire measurement image, but is a

computationally costly process. If certain templates can be skipped or searched in only a local area, large computation times are saved. This subsection defines a method by which the partition sets can be used to localize the matched filter search.

When a particular hypothesis orbit is mapped into the sensor frame, $\mathcal{A}_m$ has an associated location in that frame. Each partition $\mathcal{S}_k$, as defined, is a set of hypothesis orbits. If a particular anchor time, t^*, is defined, a particular partition can be mapped into $\mathbf{S}(2)$, so

$$\mathcal{Q}'_k = \{\mathbf{x}_d \in \mathbf{S}(2) : \mathbf{x}_d = \mathbf{h}(\boldsymbol{\phi}(t^*; t_0, \mathbf{x}(t_0)), \mathbf{k}(t_0)), \mathbf{x}(t_0) \in \mathcal{S}_k, \}. \tag{6.36}$$

For this particular partition's template to be necessary, the observer must be looking at the object during the exposure time

$$\mathcal{O}_o = \{\mathbf{x}_d \in \mathbf{S}(2) : \mathbf{x}_d \text{ in observer frame at } t^*\}. \tag{6.37}$$

When an MF template is used to search an image, it is tiled over the pixels. This is done by choosing an anchor pixel in the template, typically the center pixel. This anchor pixel is aligned with each pixel in the image and then the correlation is evaluated. $\mathcal{S}'_k$ is the set of all unit vectors which can be generated from the partition set $\mathcal{S}_k$ at t^*. If the anchor pixel is defined as the location of the object at t^*, a search space can be defined

$$\mathcal{Q}_k = \mathcal{O}_o \cap \mathcal{Q}'_k. \tag{6.38}$$

$\mathcal{Q}_k$ represents the search space for template k in an image taken by observer o. This can be seen in Figure 6.5, where each template has local areas where it should be evaluated.

This method yields one of two results. First, $\mathcal{Q}_k$ can be empty, in which case the template does not require searching for a particular image. Second, $\mathcal{Q}_k$ is non-empty, and the template need only be searched over a subset of $\mathcal{O}_o$. More generally, this illustrates a fundamental problem of searching for partially known SO: the search space defined by

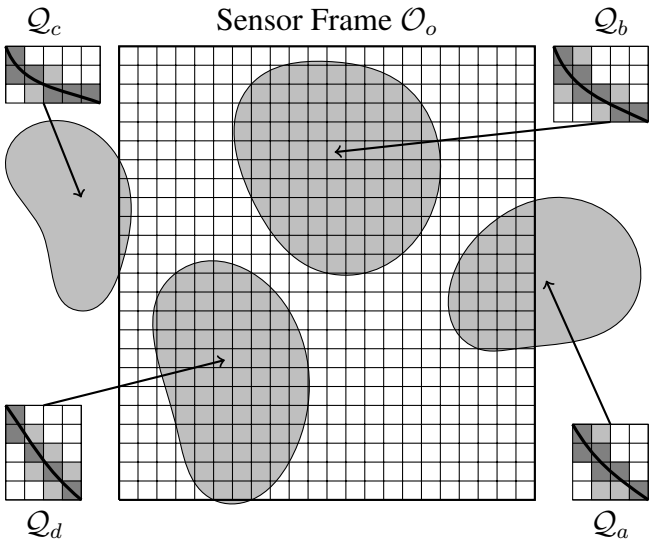

Figure 6.5: Localized searching of the sensor frame for partitions a through d.

prior knowledge can be larger than what can be captured by a single measurement. The prior set, $\mathcal{S}_0$, and the subsets, $\mathcal{Q}'_k$ can inform sensor tasking. For example, a tasking scheme could be to take a minimum number of observations to observe the entire set, $\mathcal{S}_0$. Tasking schemes for observing partially known space objects based on this framework is considered future work. Sensor pointing and orientation must be known in order to implement this method. In the case of significant uncertainty in pointing, search areas should be expanded appropriately (e.g. one pixel of point uncertainty implies one pixel of padding on search spaces is needed).

6.3.7 Probability of Detection and Likelihood Map from Matched Filter

Hypothesis testing on matched filter results is not a new concept [62]. Multiple methods exist and this chapter makes no claims about which is best or worst. Instead, this subsection presents one possible formulation.

Consider a matched filter template which consists of N_z non-zero pixels a_i. When calculating the matched filter for a particular position, the template, $\mathbf{M}_0$, is correlated with a subsection of the image, $\mathbf{Z}$. This appears as a weighted sum, where each pixel in the template, a_i, is associated with a particular measured pixel value, z_i. Then the matched

filter correlation is calculated by

$$z_{MF} = \sum_{i=1}^{N_z} a_i z_i. \tag{6.39}$$

As per Section 6.2, assume that the pixels consist of some signal, s_i, and zero mean i.i.d. Gaussian read noise, w_i, shown in Equations (2.5) and (2.6). Note that in this formulation, s_i is equivalent to the mean of the measurement z_i. In order to have zero mean read noise in practice, an effective background subtraction method is needed to remove artifacts including dark current and hot pixels. Under these assumptions, z_{MF} is distributed as follows:

$$\begin{aligned} \mathbb{E}[z_{MF}] &= \sum_{i=1}^{N_z} a_i s_i \\ \text{Var}[z_{MF}] &= \mathbb{E}\left[\left(\sum_{i=1}^{N_z} a_i w_i\right)^2\right] = \sigma_w^2 \sum_{i=1}^{N_z} a_i^2 \\ z_{MF} &\sim \mathcal{N}(\Sigma a_i s_i, \alpha\sigma_w^2) \end{aligned} \tag{6.40}$$

where $\alpha = \sum_{i=1}^{N_z} a_i^2$ is a scaling factor dependent on the actual weights, a_i. A hypothesis test is desired for determining if there is signal present which is effecting z_{MF}. This is equivalent to asking if z_{MF} has a mean of zero or a mean greater than zero. Consider the null and test hypotheses for a binary hypothesis test, for a matched filter

$$\begin{aligned} H_0 &: z_{MF} \sim \mathcal{N}(0, \alpha\sigma_w^2) \\ H_1 &: z_{MF} \sim \mathcal{N}(\beta, \alpha\sigma_w^2),\ \beta > 0. \end{aligned} \tag{6.41}$$

For binary hypothesis testing, a probability of false alarm is set, p_{FA}, which in term

defines an integration threshold, z_{TH}, based on the null hypothesis PDF, $f_0(z)$,

$$p_{FA} = \int_{z_{TH}}^{\infty} f_0(z)dz. \tag{6.42}$$

The test hypothesis PDF, $f_1(z)$, is then integrated over, giving a probability of detection,

$$p_{Detect} = \int_{z_{TH}}^{\infty} f_1(z)dz. \tag{6.43}$$

For more on this subject, see [95]. Note that the this analysis relies on the Gaussian noise assumption. If the noise is non-Gaussian, this methodology is still usable, but requires further work to model the distributions for the null and test hypotheses. By using a continuous Gaussian distribution, an inherent error is incurred based on the difference between the true CDF and approximate Gaussian CDF. Under the assumption that the discrete function is approximately Gaussian in shape, this continuous CDF approximation should be a smooth function that interpolates through the discrete CDF. As the resolution of the discretization goes up, the approximation becomes better. This hypothesis test can be used to determine if there is significant evidence that the predicted signal exists in the predicted location. Because the matched filter gives an optimal SNR gain, this test should maximize p_{Detect}, though an explicit proof of this claim is out of the scope of this chapter.

6.3.8 Orbit Determination and Sequential Bayesian Filter

The main motivating case for this research is an admissible region hand off for fast and accurate orbit determination. Admissible regions hand off refers to attempting to finding an object with an associated admissible region, with a sensor that was not used to generate the admissible region. Work has been done on this problem in the past [32]. The basic case of this is trivial; if two observers with sufficiently diverse locations observe an SO at the same time, a good orbit determination can be calculated through triangulation. For

more on this subject, see [19]. For the case of observers at different times and positions attempting to perform orbit determination on a particular space object, more analysis is needed. First, consider the aforementioned matched filter bank, performed on an admissible region detected by observer o, $\mathcal{AR}_o$. This admissible region can be propagated forward to time t_0, forming $\mathcal{S}_0$ and mapped into the sensor frame of a second observer p. If a statistically significant detection is made by the matched filter bank, a measurement can be made. This measurement, made by observer p, consists of the determinable states for an EOS:

$$\mathbf{x}'_{d,p}(t_0) = [\alpha(t_0), \delta(t_0), \dot{\alpha}(t_0), \dot{\delta}(t_0)]^T. \tag{6.44}$$

If the calculated $\mathbf{x}'_{d,p}(t_0)$ is not consistent with the location and rate of the template used to detect it, then the measurement can be disregarded; the object that has been detected is not consistent with any hypothesis generated from $\mathcal{S}_0$. Assuming, the measurement and template are consistent, an admissible region, $\mathcal{AR}_p$, can be constructed. Now, there exists two sets $\mathcal{S}_0$ and $\mathcal{AR}_p$ both represented by orbits at a particular time, t_0. Furthermore, it is known that both admissible regions contain an orbit consistent with a particular measurement, $\mathbf{x}'_{d,p}(t_0)$. This implies a non zero intersection of the two admissible regions, that is,

$$\mathcal{S}_0 \cap \mathcal{AR}_p \neq \emptyset. \tag{6.45}$$

Fujimoto and Scheeres showed that if two arbitrary admissible region intersect, there is a diminishingly low probability that they were not generated from the same object [31]. This statement implies that whatever orbit in $\mathcal{S}_0$ that could create the measurement $\mathbf{x}_{d,p}(t_0)$, are the intersection of these two admissible regions, and therefore the correct orbit. The Fujimoto and Scheeres chapter works with exact admissible regions and draws the conclusion that the intersection must necessarily be a single point in the state space. This can be

expanded to include uncertainties in the measurements of $\mathcal{S}_0$ and $\mathcal{R}_p$, implying a non-point intersection,

$$\mathcal{S}_0' = \mathcal{S}_0 \cap \mathcal{R}_p. \tag{6.46}$$

where $\mathcal{S}_0'$ is a new set of orbits. The new set, $\mathcal{S}_0'$, can also be written as $\mathcal{R}_o \cap \mathcal{R}_p$. Note that by definition $\mathcal{S}_0 \supseteq \mathcal{S}_0 \cap \mathcal{R}_p$, implying that this calculation consists only of deleting parts of $\mathcal{S}_0 = \mathcal{R}_o$. The knowledge obtained from $\mathbf{x}'_{d,p}(t_0)$ can be used to rate track the object or inform an online matched filter [100].

The problem with this approach is that new measurements can either be significantly noisy or the probability of detection may be below one. In these cases, a sequential Bayesian update scheme, in the form of a particle filter, can be formulated to operate directly on hypothesis test results. At this point, the prior set must be instead thought of as a prior PDF. The uncertainty in $\mathcal{S}_0$ can either be represented by a known prior PDF or with a uniform prior similar to an admissible region, The new measurement defines a Bayesian update

$$P(\mathbf{x}|\mathbf{x}'_{d,p}(t_0)) = \frac{P(\mathbf{x}'_{d,p}(t_0)|\mathbf{x})P(\mathbf{x}|\mathcal{S}_0)}{P(\mathbf{x}'_{d,p}(t_0))}. \tag{6.47}$$

The measurement likelihood map, $P(\mathbf{x}'_{d,p}(t_0)|\mathbf{x})$, can be calculated from the localization of $x_{d,p}(t_0)$. The localization of $x_{d,p}(t_0)$ must necessarily be calculated from the matched filter result, which is represented by the hypothesis test in Equation (6.41).

$$P(\mathbf{x}'_{d,p}(t_0)|\mathbf{x}) = P(z_p > 0 \ \forall \ j \in [1, \ldots, n]) \tag{6.48}$$

The right hand side of Equation (6.48) is equivalent to the hypothesis test in Equation (6.41). The matched filter can be built into a Bayesian update, allowing the matched filter bank to be built into a particle-based filter. It should be noted that while this chapter's

results only include a basic particle filter, this likelihood can apply to a wide range of filters including random finite set and multi-target filters.

There is a statistical hurdle that occurs in certain cases when performing this update; the image which provides the likelihood update may not cover the entire prior set, meaning some particles simply cannot be updated. In such a case, the authors recommend partitioning the prior set $\mathcal{S}_0$ into two pieces, the observered partition, $\mathcal{S}_0^{\mathcal{O}}$, and the unobserved partition, $\mathcal{S}_0^{\backslash\mathcal{O}}$. Because no information is obtained on the unobserved partition, the probability density must stay the same, $P_{\mathcal{S}_0^{\backslash\mathcal{O}}}(\mathbf{x}) = P_{\mathcal{S}_0^{\backslash\mathcal{O}}}(\mathbf{x}|\mathbf{x}_{d,p}(t_0))$. Conversely, Bayes rule can be applied to only the observed set,

$$P_{\mathcal{S}_0^{\mathcal{O}}}(\mathbf{x}|\mathbf{x}'_{d,p}(t_0)) = \frac{P_{\mathcal{S}_0^{\mathcal{O}}}(\mathbf{x}'_{d,p}(t_0)|\mathbf{x})P_{\mathcal{S}_0^{\mathcal{O}}}(\mathbf{x}|\mathcal{S}_0^{\mathcal{O}})}{P(\mathbf{x}'_{d,p}(t_0))}. \tag{6.49}$$

This particular case of partial likelihood is incomplete and deserves a great deal more work. Further analysis is out of the scope of this chapter.

6.4 Implementation of Theory

In Figure 6.6, the logic for an implementation of a matched filter bank based on the work in this chapter is mapped out. Boxes 1 and 2 define the prior information for the process. Box 3 maps a prior set of orbits to measurement arcs, which in practice are what is used to partition and generate templates. In Box 5, the pivot point $\mathbf{x}_k$ can be chosen in a few different ways. The simulations in this chapter choose the first pivot point randomly and choose all following pivots as the point furthest in MDM from the previous pivot. Box 6 represents comparing the distance between each point and the pivot in MDM. The measurement arcs should be pre-calculated for each point to expedite this. The time density of the measurement arcs should be chosen based on processing power. The partitioning module should pass out both the pivot point and the associated partition set. The subsection labeled "Local Matched Filter" represents how to process the matched filters. $\mathcal{L}$ is the intersection

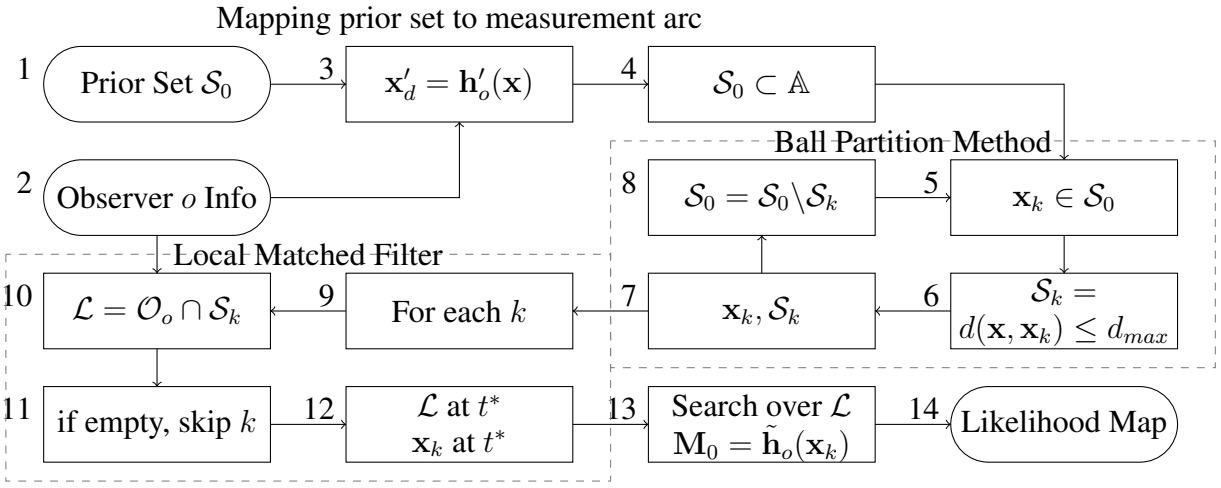

Figure 6.6: Code flow diagram for implementing MF Bank.

of the measurement frame and a particular partition. In practice, this can be generated by mapping the measurement arcs into the measurement frame with the desired camera model and locating all pixels with any signal. It is useful to pad the localization matrix to expand the search area by a few rows and columns. This assures that if the truth appears near the boundary of the search area, the matched filter correlation peak is captured fully. If $\mathcal{L}$ is empty, this template can be skipped. If $\mathcal{L}$ is not empty, the non-empty pixels defines the locations where the signal should be checked.

The likelihood can be used in a variety of ways: each particle's likelihood can be evaluated for a particle filter update, the peak likelihood can be directly considered as a measurement, or the likelihood itself can used seen as the final product. Another extra processing technique which is worth considering is point spread function (PSF) fitting [98]. This technique involves optimizing over the underlying signal's end points and PSF at sub-pixel accuracy. The matched filter template with the best correlation can be used an initial guess for the PSF fit. The PSF fit result can then be used as the detection, which should lead to sub-pixel accurate localization of endpoints.

Table 6.1: Orbital elements for simulated object.

a (km)	e	Ω (rad)	i (rad)	ω (rad)	f (rad)
1.2658e+04	0.2198	2.3863	0.3489	1.2698	0.3724

6.5 Simulation Results

6.5.1 Problem Geometry

A medium Earth orbiting object, with parameters in Table 6.1 is simulated to demonstrate the admissible region hand off. The object is simulated at the UTC, November 12, 2011, at 20:34:07.

Two EOS observers, $\mathbf{o}_1$ and $\mathbf{o}_2$, located in Colorado and Hawai'i respectively, take observations of the space object at times $t_1 = 0$s and $t_2 = 120$s. Consider observer 1 to be a high sensitivity sensor able to detect this space object via conventional methods. Consider observer 2 to be a low sensitivity sensor which would normally be unable to detect this space object. The geometry of the problem is shown in Figure 6.7. The SO and observer position and velocities at the time of the observations are:

$$\begin{aligned}
\mathbf{r}(t_1) &= [-5906.3 \ \ -7313.5 \ \ 3410.0]^T (km) \\
\dot{\mathbf{r}}(t_1) &= [5.0791 \ \ -4.7364 \ \ -0.01212]^T (km/s) \\
\mathbf{r}(t_2) &= [-5280.5 \ \ -7860.5 \ \ 3398.8]^T (km) \\
\dot{\mathbf{r}}(t_2) &= [5.3445 \ \ -4.3765 \ \ -0.17361]^T (km/s) \\
\mathbf{o}_1 &= [-1207.0 \ \ -4746.7 \ \ 4072.0]^T (km) \\
\dot{\mathbf{o}}_1 &= [0.34519 \ \ -0.08777 \ \ 0.0000]^T (km/s) \\
\mathbf{o}_2 &= [-5446.2 \ \ -2431.8 \ \ 2252.0]^T (km) \\
\dot{\mathbf{o}}_2 &= [0.17337 \ \ -0.39759 \ \ 0.0]^T (km/s).
\end{aligned}$$

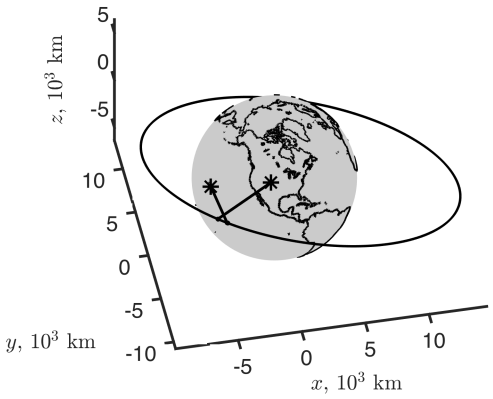

Figure 6.7: Location of observers and SO.

6.5.2 Observations

The first observer takes an image which leads to a detection and measurement by conventional methods. The admissible region throughout these simulation results is represented by a discrete point-wise approximation of 20,000 particles. The second observer, $\mathbf{o}_2$, takes low SNR observation of the SO, also at $t_2 = 120$s. The image simulation is done without star signals, to prevent false detections. Stars present a problem in implementation on real data, but can be avoided with good star identification and subtraction. The image is generated with a 10 second exposure time, field of view of 10 degrees, and sidereal stare pointing mode. The total object SNR in this image is approximately 6.8, while the maximum pixel-wise SNR in this image is approximately 1.0. Conventional methods, such as thresholding and corner detection, fail to detect this object. Because the only orbital knowledge on this object in the admissible region from observer 1, conventional velocity filters cannot be implemented.

6.5.3 Partitioning

In order to give insight into the computational cost of the ball partitioning, the simulation has been run for a range of MDM threshold values. For these multi-value test runs, object

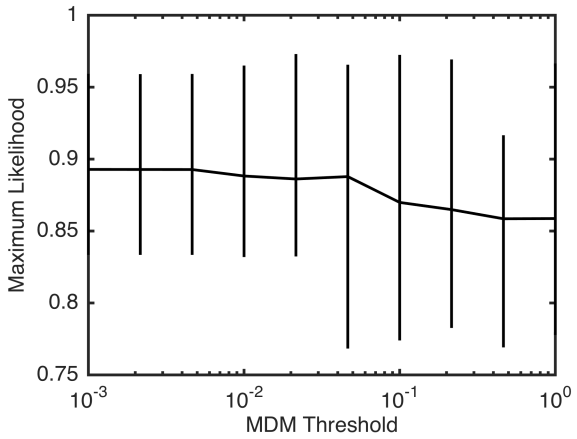

Figure 6.8: Maximum likelihood results from varying MDM threshold with 3 sigma bars.

SNR is set to approximately 0.25, the probability of false alarm threshold is set to 0.999, and 10,000 particles are used to represent the AR. In this particular simulation, only the 2 dimensional admissible region need be represented by a discrete set of particles. For such space, 10,000 particles is typically sufficient. Another note on particle density is that in order for the partitions to be well formed, sufficient particles per subset are desired. A good rule of thumb is around 100 particles per partition, but as always, more particles help. In Figure 6.8, the maximum likelihood is shown to vary more significantly with a higher MDM threshold. Because raw measurement SNR is constant, a perfect template should always provide the same likelihood. The sporadic behavior at high MDM threshold is a result of the inability of a high MDM threshold to consistently provide effective templates. The number of templates and computation time are plotted against MDM threshold on a log-log scale in Figure 6.9. Computation time follows a predictable curve; more templates proportionately leads to more computation time.

Next, a particular partitioning is shown in more detail. In remaining simulations, object SNR is set to approximately 1, the probability of false alarm threshold is set to 0.999, and 20,000 particles are used to represent the AR. A maximum MDM value of $d_{max} = 0.1\ rad \cdot s$ is chosen. The measurement arc is calculated for each particle, represented by a time series of unit vectors. A ball partitioning algorithm based on approximate integration over

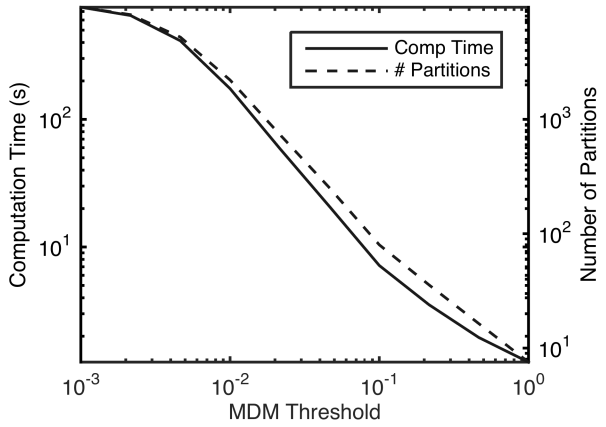

Figure 6.9: Computational time results from varying MDM threshold.

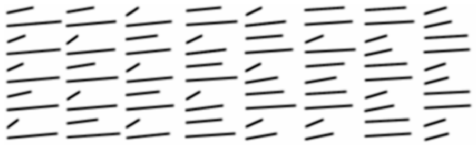

(a) Mosaic of Templates

(b) Temp A (c) Temp B (d) Temp C

Figure 6.10: Sampling of matched filter templates from partitioning. Colors are Inverted.

the time series sorts the particles into subsets containing consistent orbits. The partitioning provides 80 distinct partitions with corresponding templates. Figure 6.10 shows a sample of the types of templates which are generated. A diverse set of templates are created by the partitioning, corresponding to the range of possible signals that could be produced by the space object, given it exists in the admissible region. Figure 6.11 shows the partitions on the original admissible region. Each template in Figure 6.10 corresponds to a partition shown in 6.10. When the evaluating the matched filters, these partitions define search spaces for the associated templates. Three particular partitions, labeled A, B, and C, are followed more closely through the remaining process. Partition C is the most correct partition, and contains truth somewhere in it.

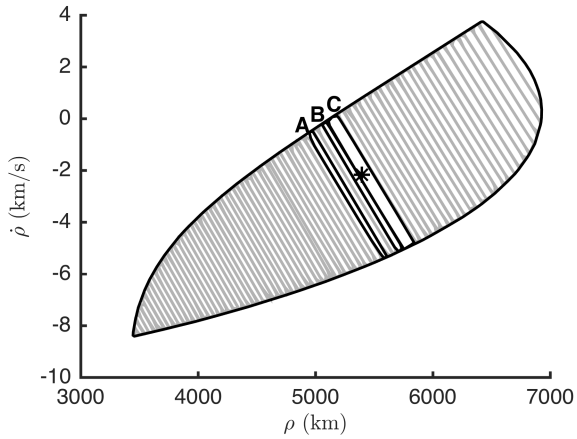

Figure 6.11: Sets generated from partitioning, shown on original AR.

6.5.4 Observation Search

The templates with corresponding partitions are searched in the image. Figure 6.12 shows the image with overlaid partitions. The matched filter gives a value for each pixel in the image which the template is centered at. Each matched filter result is then tested via the hypothesis test in Equation (6.41). This gives a likelihood function over the partition, which can evaluated for all particles. Figure 6.13 shows the correlation maps for three templates. Figure 6.13(a) shows an incorrect template which doesn't produce significant SNR gain. Figure 6.13(b) shows a template with an imperfect template, illustrating the kind of SNR gain possible in real operation. Figure 6.13(c) shows a template primed with perfect orbital knowledge, simply to illustrate what the ideal case can look like.

6.5.5 Orbit Update

The final step is calculate a Bayesian update on the space object. Each particle's likelihood is evaluated within its partition. Some particles do not fall within the EOS image which poses a problem for the update. As previously discussed, Equation (6.49) is used to perform a likelihood update on only the particles which are observed, while the particles outside the image are kept at constant weight. Using the new measurements, the PDF over the original

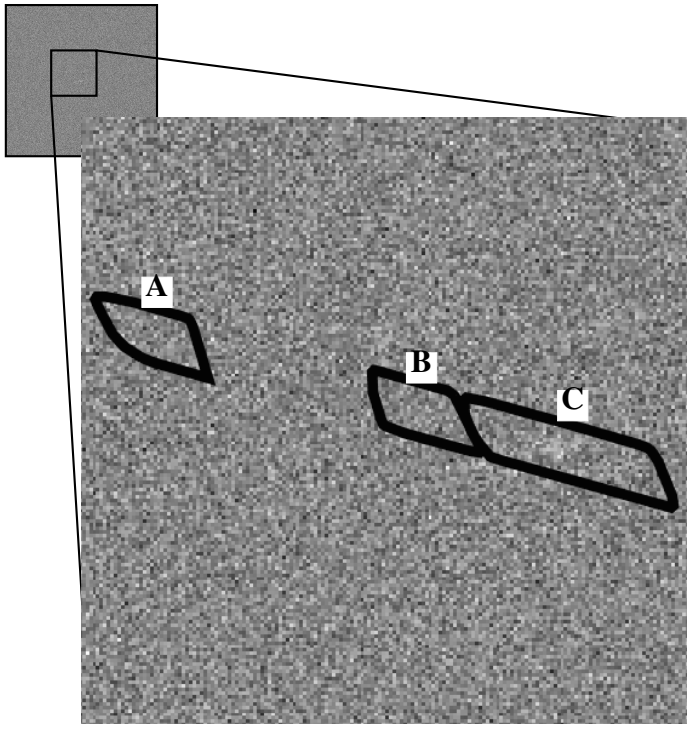

Figure 6.12: Measurement image taken by observer 2.

(a) (b) (c)

Figure 6.13: Correlation results from a variety of templates.

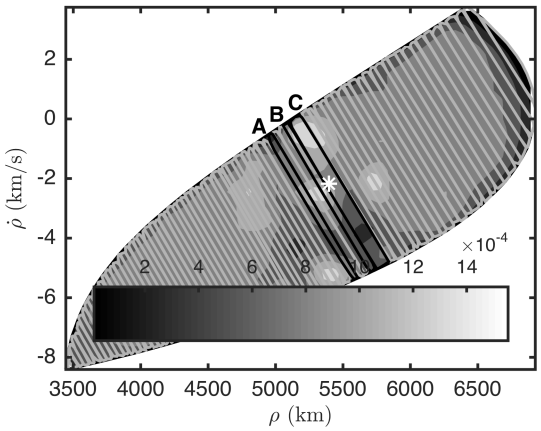

Figure 6.14: Final PDF update from matched filter likelihood.

admissible region can be updated, point-wise, similar to a particle filter update. Figure 6.14 shows the results of a single Bayesian update. It is important to note that this method does not make a hard update decision and should not. When objects at a low SNR are observed, it can be impossible to confirm that a statistically significant "detection" has been made. Many false alarms mimic the signal giving the false peaks seen in the PDF. Only through multiple detections, which show both statistical significance in the image and adherence to a dynamic model, can a conclusive statement be made. Ultimately, this would be used in a filtering scheme such as a particle filter or multiple hypothesis filter, as the likelihood function.

6.5.6 Sequential Particle Filter

As of now, the likelihood function update at a given time only makes sense if it can be evaluated for all objects within the PDF. The following results proceeds under the assumption that the measurement likelihood can be evaluated for all particles. Two observers in Georgia and a MEO object are simulated. Similar to the previous simulation, the first observer takes an observation of the object and an admissible region is created. The AR is passed off to the second observer, which then takes a series of 9 observations. Each observation is a 10 second exposure, sidereal stare observation, which are initiated at 20 second intervals.

The following initial parameters are used

$$\begin{aligned}
\mathbf{r}(t_1) &= [-3420.2 \;\; -8138.0 \;\; 4698.5]^T (km) \\
\dot{\mathbf{r}}(t_1) &= [6.5260 \;\; -2.3753 \;\; -0.0121]^T (km/s) \\
\mathbf{o}_1 &= [652.95 \;\; -5284.4 \;\; 3499.5]^T (km) \\
\dot{\mathbf{o}}_1 &= [0.3843 \;\; 0.0475 \;\; 0.0000]^T (km/s) \\
\mathbf{o}_2 &= [522.51 \;\; -5281.8 \;\; 3525.1]^T (km) \\
\dot{\mathbf{o}}_2 &= [0.3841 \;\; 0.0380 \;\; 0.0]^T (km/s).
\end{aligned}$$

The final posterior PDF of the particle filter can be seen in Figure 6.15. The standard deviations in range and range rate space can be seen plotted over iterations in Figure 6.16. The final values of the particle filter are highly dependent on factors such as process noise, image pixel size, object SNR, and more. In addition, the noise at such low SNR values provides a large numbers of false alarms. These plots are meant to demonstrate that particle filter convergence is possible. Object SNR over the series of images is shown in Figures 6.17 and 6.18. Figure 6.17 shows the total object SNR calculated by Equation (2.13). The total object SNR, which is the SNR of all pixels which the object appears in, gives an idea of the total available information for an object. Figure 6.18 shows the maximum per pixel SNR calculated by Equation (2.12). The maximum per pixel SNR gives a measure of the ability of classical detection methods (thresholding, edge detection) to detect this object. It is worth noting that this particle filter scales gracefully for different object SNR values. This means that this method can be used as a general object detection method which can capture dim space objects.

6.6 Conclusion

A methodology has been developed for obtaining low SNR detections of space objects when uncertain prior information is available. Unlike typical matched filters, exact knowl-

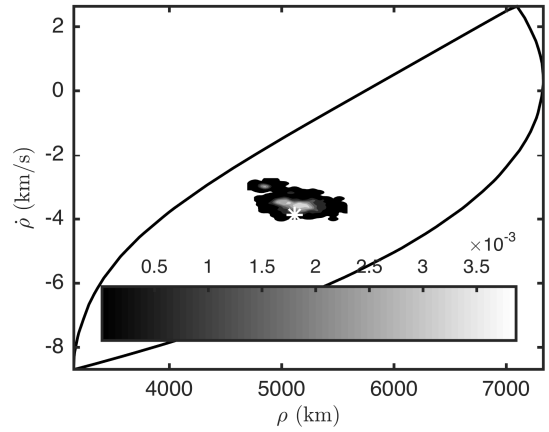

Figure 6.15: Final PDF update from multi-frame matched filter likelihood particle filter.

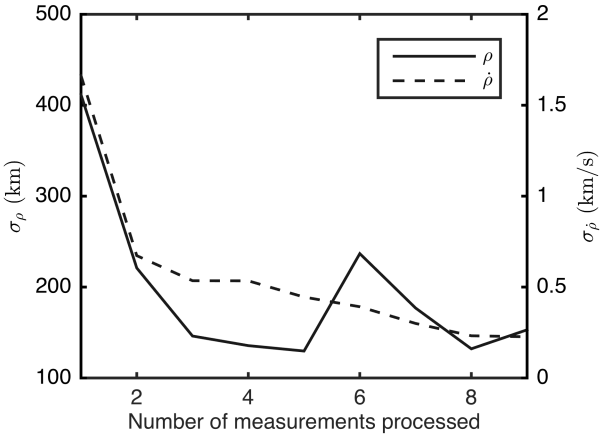

Figure 6.16: Standard deviation of PF over iterations.

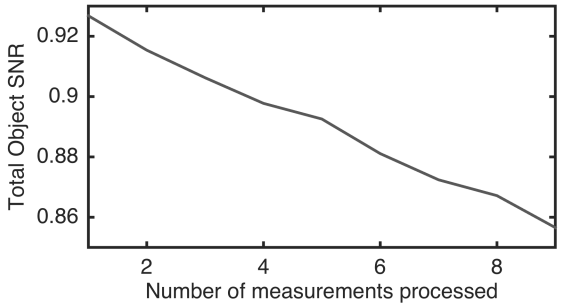

Figure 6.17: Object SNR over observation series.

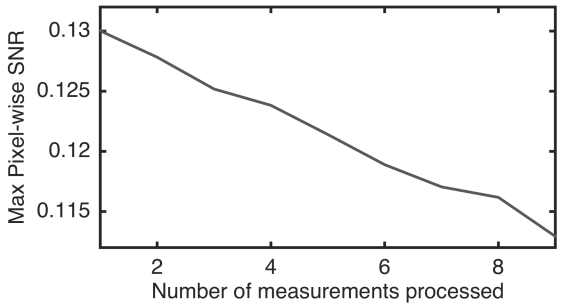

Figure 6.18: Pixel SNR over observation series.

edge is not necessary with this matched filter bank. The partitioning method that is proposed in this chapter enables efficient and mathematically rigorous generation of the matched filter bank. A novel metric is first defined in this chapter, upon which the partition process is defined. The resulting partitions are then used for an analytic result which allows the localization of a matched filter. This MF bank enables a statistical update to an orbit to be derived directly from image data. Reliable low SNR detection methods have the ability to decrease require costs for performing SSA with electro-optical sensors. Instead of building more monolithic telescopes, large numbers of smaller sensors can be used to gather important SSA data. Large, exquisite sensors can then be focused on object discovery rather than catalog upkeep. Furthermore, this methodology can be generalized to a wide variety of sensors and data types.

6.6.1 Unifying Example

At this point in our example, we still have a series of images for each of several locations spread across the admissible region prior. All reasonable bright lights sources are also identified and checked for association with the prior. Unlike contribution 3, the goal now is to search through each image to perform orbit determination directly on the admissible region. The example provided in the simulation section of this chapter is exactly the same as the last piece of this unifying example. Each image set provided by contributions one and two are searched with matched filters provided by the underlying admissible region.

The likelihood provided by the matched filters is used to update the admissible region, provided the final orbit determination.

CHAPTER 7
CONCLUSIONS

In moving toward an ever improving Space Surveillance Network, there are myriad different interconnecting techniques which all affect our ability to track, estimate and characterize space objects. Some of the most important techniques concern data acquisition, which requires both where to point a sensor and what qualities the data should have. Just as important is what to do with the data once it is taken, both for tracking objects and turning these tracks into meaningful information on orbits. This work is broken into four distinct contributions; the first looks at how a sensor should scan a set in order to take data, the second shows how to choose optimal telescope parameters when taking data, the third looks at image processing algorithms for uncued detections, while the fourth looks at cuing detections with a set and performing orbit determination on the set.

The first major pillar of this work is showing the interconnectivity between orbit statistics and sensor tasking. One of the foundational motivations of this work, and a major new development in SSA, is admissible region theory. Admissible regions, as seen throughout this work, are an inherently set-based method of storing prior knowledge on a space object. While admissible regions have been incorporated into orbit estimation techniques, this work looks at how they fit into sensor tasking and image processing. The first major contribution in this dissertation constructs an entire tasking algorithm based on reacquiring a object with a set as prior knowledge. This work leverages the set by analyzing the boundaries expansion rate and working this knowledge directly into the optimization algorithm. The fourth contribution in this work looks at how the admissible region set can be used to inform detection methods. This piece of work is able to achieve extremely sensitive detections which are attainable only through prior information.

The second major pillar of this work is the integration of sensor tasking and image

processing techniques. Another major motivation of this work, the maximum likelihood estimator, is introduced as a popular and thorough technique for detection and analysis in SSA. This estimator is used to define statistical reward functions which allow the comparison of image data. This allows sensor tasking strategies and data campaigns to be chosen based on the reward inherent to the data taken. In Lemma 5.3.1, the inherent relationship is shown between the maximum likelihood estimator and the hypothesis test likelihood function used for detection and estimation algorithms.

The final major pillar of this work is the combined integration of image processing and orbit statistics. This piece is built upon the use of a detectionless likelihood function, in particular one designed for SSA. This is a likelihood function which allows statistical conclusions to be drawn directly from an image. The third contribution develops and analyzes this likelihood function purely as a detection technique. The Multi-Bernoulli Filter is equipped with this likelihood function and specialized for SSA purposes. The fourth contribution brings in admissible region theory by formulating a method by which orbit determination can be performed directly with the detectionless likelihood function.

The work in this thesis provides an array of new research ideas. The detection methods provided in this work push the current boundaries of image detection sensitivity. The derivation on velocity astrometry information provides a novel way to choose how data is taken. The search set, an analytic concept that is used throughout this dissertation, provides a new way of thinking about orbits and unifies several disparate concepts. Admissible regions are used extensively in SSA but are often modeled as PDF's, reachability sets have been explored but never fully utilized, and search regions are often erroneously modeled as areas of the sky. Furthermore, the set which is required to formally define a PDF is typically ignored in SSA applications. This dissertation represents a new way of conceptualizing orbits; all of these constructs are modeled first as a set of orbits. The search set, and the methods in this dissertation which utilize it, provide a unifying framework around which SSA systems can and should be built. This work also supports a broad push away

from simple orbit models, such as TLEs. This dissertation also branches out to currently popular estimation methods such as maximum likelihood estimation and finite set statistics, broadening the applicability. Finally, it should be noted that multiple pieces of this dissertation have inspired direct future work internally at Georgia Tech. The multi-Bernoulli filter detection technique has an ongoing follow-up effort [101] and work has begun on a follow-up effort for the optimal sensor tasking techniques.

The search set concept has the potential to broadly alter how the SSN is currently operated. A core goal of this dissertation is to enable specific space objects to be reacquired regardless of the time, location, and quality of the follow-up sensor or sensors. This works towards enabling reacquisition of poorly known, maneuverable, and difficult-to-detect objects regardless of circumstance. This has the potential to enable a variety of new SSA missions. Tracking small objects from multiple locations provides an enabling technology for networks of inexpensive and unconventional sensors to support space object catalog expansion. Operators now have a way to search for maneuverable and dynamic objects through reachability sets. Finally, these techniques build toward military missions such as asset protection, uncooperative target tracking, and intercept missions.

The mathematics in this dissertation are purposefully written to be as broadly applicable as possible. The search set, and dynamics it obeys, are broad enough to apply to any number of situations from sea domain awareness to tracking individuals for security purposes. The search set technique, if used for other domains, now has a dissertation's worth of supporting techniques attached to it. The detection methods are all closely related to computer vision where similar techniques are being employed. The new likelihood function and matched filter bank in this dissertation represent steps forward for the computer vision community.

The four chapters of this dissertation, as a whole, forecast the future of the Space Surveillance Network. Each chapter shows how new ideas in SSA, like detectionless filters and admissible regions, can be combined together to form powerful new techniques. Highly sensitive low SNR detections, orbit determination from multiple sensor locations

and times, reacquisition on barely detectable objects with large priors. In many ways, the work in this document only scratches the surface. Many of the results in this dissertation are only preliminary, and further work will provide stronger methods.

7.1 Joint Tasking and Estimation of a Search Region

There exist a variety of methods to handle the scheduling problem [23–25], as discussed in the introduction. These methods can all be used in conjunction with the work described in this dissertation. A scheduling algorithm can be used to decide between multiple prior distributions and multiple sensors, while the tasking algorithms in this dissertation prescribes how a sensor searches and acquires data. Area growth rate estimation can be used as a foundation for a scheduling algorithm, which is ongoing future work [102]. Finally, it is worth mentioning that the scheduling problem is a special case for the work in this dissertation, and the combined scheduling tasking problem can be considered one problem within the framework of this work. The computational tractability of this problem, which is in essence a second CSP wrapped around the CSP considered in this chapter, is not practical.

A discussion about future estimation schemes which make sense to be used in conjunction with these algorithms is included. This paper works under a Bayesian assumption; in the case that prior information on a single object is represented with a PDF, a fairly standard Bayesian filter is sufficient. This takes the form of a particle filter for orbit determination directly on an admissible region. For a robust implementation on real data, special considerations must be taken in order to account for false and missed detections. For operating on an admissible region, there are two primary options. The first method, which is commonly used is the literature and is used here, makes a uniform prior distribution assumption and uses Bayesian estimation, similar to above [32]. A second, newer method exists which makes no uniform prior distribution assumption, instead modeling the set with an uninformative prior and using Dempster-Shafer theory to perform estimation [71]. This method of uninformative priors has not been explicitly extended to a general set for orbit estimation,

but such a method should work and would apply to all cases in this dissertation. Neither of these methods consider how estimation should be handled when a measurement only interrogates a subset of the state-space, rather than the whole space. However, Dempster-Shafer theory's use of belief and plausibility as bounds on the true probability may provide an interesting framework to directly account for partial state space interrogation. This is a primary area of future work.

7.2 Opportunities for Future Work

The first contribution on search through prior sets has several interesting new directions. Area over time is a simplistic bulk measure for search sets prediction. One area of future work should be to provide more complex predictions that better take into account search space shape and how it intersects with sensor field of view. If such an analysis can be worked into a cost function, it creates a powerful optimization tool. Sensor design is an obvious extension of this work. The analyses provided take both sensor location and sensor parameters into consideration, and can be easily turned around to look at what those factors should be to provide certain search capabilities. The actual estimation of objects which exist in a set may prove to be more complicated than one might think. Methods based on finite set statistics and Dempster-Shafer theory may provide the best avenue. Such a method should allow for a solution of no objects, one object, or many objects existing in the set.

The second contribution on optimal telescope parameters was chronologically the last piece of this document to be researched and written. This chapter ended with a wealth of interesting research problems that could not fit into a single chapter. The reward function gradients and optimization results are all centered on the exposure time vector as the primary design variable. Some minor analysis has been performed on slew rate, which lead to the intuitive result that rate tracking generally leads to better information. The primary results treat slew rate as a fixed variable determined by the feasibility of rate tracking.

Especially when considering some constraints such as star clutter and reliability of star identification, there may be interesting situations when choosing a non-rate track trajectory could be advantageous. Another variable which is not discussed in detail in this dissertation is the ability of modern sensors to sample a smaller "window" inside an image, or sub-sample blocks of pixels into single pixels. These processes, known as windowing and binning, provide a way to exchange data quality for data down-link rate. In the case of wanting high frequency data, an interesting case can be made to include these processes as full design variables. Related to windowing and binning is the pointing uncertainty of a telescope, which is typically found through star identification. Not only could this inherently have an affect on the various information metrics, but it is dependent on factors such as slew rate, field of view, and astrometry information of stars.

The results in this dissertation are rather theoretical, providing a framework of how to design data campaigns with only minimal analysis of the actual data campaigns. Part of the ongoing and future work associated with this dissertation is validation on simulated and actual data campaigns. While the rewards functions here are based on the maximum likelihood estimator, there are a variety of detection, tracking and estimation techniques for both astrometry and photometry. Data can take a variety of forms depending on telescope field of view, aperture, location, etc. An extensive set of validations would be best to explore how the reward functions in this dissertation, as well as future reward functions, perform.

In the case of a signal with a known frequency or already well known light curve, the only engineering decision is choosing what point on the Pareto surface is desired. This is complicated when the signal frequency and amplitude are unknown. Furthermore, signals may not be sinusoidal in shape and can have interesting artifacts such as glints. This dissertation sidesteps this design decision by selecting a worst case scenario frequency and amplitude and solving accordingly. The intuitive result for such a situation is to initially have exposure times as short as possible to allow dynamic estimation of frequency content,

then switch to an Pareto solution suitable to the detected frequency. This is at odds with the astrometry which calls for starting the campaign with a single long exposure. Solving these issues with a dynamic algorithm that updates its solution based on new data provides a wealth of new engineering problems.

As of right now, the velocity astrometry rewards function registers zero information for a single rate tracking image, and has a problematic parameter, t^*. Further analysis into how to develop velocity astrometry may provide an easy, intuitive solution to this. The current astrometry rewards assume that position and velocity information are the only realistic objectives. However, given a long enough campaign, full orbit information becomes observable. Future work should therefore look at how to extend or replace the astrometry reward functions with one which maximizes information on full orbit determination. Photometry, similarly, is solved in a simplistic way by attempting to minimize the uncertainty of the flux over time. In reality, photometry is part of a broader problem of object classification, whose objectives may require a wholly new reward function or functions. All of the reward functions developed in thi should not be viewed as final results, but instead as a blueprint on how to develop a reward function for whatever specific objects are necessary in a given situation. It is the author's hope that this work spawns a multitude of follow-on projects analyzing new ways to maximize SSA data.

The third contribution on a detectionless filter is part of a much more crowded field with currently ongoing work. One critical broad piece of of ongoing and future work involves the computational benefits of Gaussian and Gaussian mixture methods. The detectionless filter requires a full particle filter even though many of the final distributions appear Gaussian. Finite set statistics filters often have Gaussian mixture forms, but the current likelihood is not compatible with a Gaussian update as the measurement function can be very non-linear. Future work should look at finding cases where Gaussian estimation can be used alongside detectionless filtering. The best possible outcome would be a way to effectively marry Gaussian mixtures with the full detectionless filter without sacrificing quality, though the

feasibility of this is unknown.

BIBLIOGRAPHY

[1] Blake, T., Sánchez, M., Krassner, J., Georgen, M., and Sundbeck, S., "Space Domain Awareness," *Advanced Maui Optical and Space Surveillance Technical Conference*, 2011.

[2] O'Connor, B., "Handbook for Limiting Orbital Debris. NASA Handbook 8719.14," *National Aeronautics and Space Administration, Washington, DC*, 2008.

[3] Morales, M. D., "Space Fence program awards contracts for concept development," July 2009.

[4] Schiff, B., Foster, J., McShane, W., and Simon, K., "Attaining Situational Understanding in the Space Domain," *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2017.

[5] Flury, W., Massart, A., Tchildknecht, T., Hugentobler, U., Kuusela, J., and Sodnik, Z., "Searching for small debris in the geostationary ring," *ESA bulletin*, Vol. 104, 2000, pp. 92–100.

[6] Kervin, P., Hoo, V. S., Nishimoto, D., and Liang, D., "Rapidly deployable Raven-class systems SSA Support in the Field," Tech. rep., DTIC Document, 2009.

[7] Flohrer, T., Krag, H., Klinkrad, H., and Schildknecht, T., "Feasibility of performing space surveillance tasks with a proposed space-based optical architecture," *Advances in space research*, Vol. 47, No. 6, 2011, pp. 1029–1042.

[8] Sease, B., Flewelling, B., and Xu, Y., "Catalog-free angular rate estimation and online detecton of resident space objects," *24th AAS/AIAA Space Flight Mechanics Meeting*, 2014.

[9] Kelso, T., "Analysis of the iridium 33 cosmos 2251 collision," Tech. rep., Center for Space Standards & Innovation, 2009.

[10] Skinner, M. A., Kelecy, T. M., Gregory, S. A., Toth, J. P., Liang, D., Yamanaka, D., Kent, S., Tjoelker, R., Margineantu, D., Allison, A. L., et al., "Commercial space situational awareness: An investigation of ground-based ssa concepts to support commercial geo satellite operators," *Proceedings of the 2013 AMOS Technical Conference, Wailea, Hawaii*, 2013.

[11] on the Peaceful Uses of Outer Space, U. N. C., *Report of the scientific and technical subcommittee on its fifty-second session,*, United Nations Publications, 2015.

[12] Saunders, P. C. and Lutes, C. D., "China's ASAT Test: Motivations and Implications," Tech. rep., National Defense Univ Washington DC Inst for National Strategic Studies, 2007.

[13] Scientist, U. A. F. C., "Report on Technology Horizons: A Vision for Air Force Science and Technology During 2010–2030," Tech. rep., Volume 1, Technical Report AF/ST-TR-10-01-PR, Department of the Air Force, Washington, DC, 2010.

[14] Coder, R. D. and Holzinger, M. J., "Multi-objective design of optical systems for space situational awareness," *Acta Astronautica*, Vol. 128, nov 2016, pp. 669–684. doi:10.1016/j.actaastro.2016.07.008.

[15] Jones, B. A., Bryant, D. S., Vo, B.-T., and Vo, B.-N., "Challenges of multi-target tracking for space situational awareness," *Information Fusion (Fusion), 2015 18th International Conference on*, IEEE, 2015, pp. 1278–1285.

[16] Milani, A., Tommei, G., Farnocchia, D., Rossi, A., Schildknecht, T., and Jehn, R., "Correlation and orbit determination of space objects based on sparse optical data," *Monthly Notices of the Royal Astronomical Society*, Vol. 417, No. 3, 09 2011, pp. 2094–2103. doi:10.1111/j.1365-2966.2011.19392.x.

[17] Vigil, M. L., Witte, D. J., LeVan, P. D., Wallentine, P. J., Briscoe, D. E., and Anderson, D. L., "Sensor suite for the Advanced Electro-Optical System (AEOS) 3.6-m telescope," *SPIE's 1996 International Symposium on Optical Science, Engineering, and Instrumentation*, International Society for Optics and Photonics, 1996, pp. 151–169.

[18] Africano, J., Stansbery, E., and Kervin, P., "The optical orbital debris measurement program at NASA and AMOS," *Advances in Space Research*, Vol. 34, No. 5, 01 2004, pp. 892–900. doi:10.1016/j.asr.2003.02.022.

[19] Sease, B., Schmittle, K., and Flewelling, B., "Multi-Observer Resident Space Object Discrimination and Ranging," *AAS/AIAA Space Flight Mechanics Meeting*, 2015.

[20] Worthy, J. L. and Holzinger, M. J., "Incorporating Uncertainty in Admissible Regions for Uncorrelated Detections," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 9, sep 2015, pp. 1673–1689. doi:10.2514/1.g000890.

[21] Holzinger, M. J., Scheeres, D. J., and Alfriend, K. T., "Object Correlation, Maneuver Detection, and Characterization Using Control Distance Metrics," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 4, jul 2012, pp. 1312–1325. doi:10.2514/1.53245.

[22] Jaunzemis, A. D., Mathew, M. V., and Holzinger, M. J., "Control Cost and Mahalanobis Distance Binary Hypothesis Testing for Spacecraft Maneuver Detection," *Journal of Guidance, Control, and Dynamics*, Vol. 39, No. 9, sep 2016, pp. 2058–2072. doi:10.2514/1.g001616.

[23] Erwin, R. S., Albuquerque, P., Jayaweera, S. K., and Hussein, I., "Dynamic sensor tasking for Space Situational Awareness," *Proceedings of the 2010 American Control Conference*, Institute of Electrical & Electronics Engineers (IEEE), jun 2010. doi:10.1109/acc.2010.5530989.

[24] Sunberg, Z., Chakravorty, S., and Erwin, R., "Information space sensor tasking for Space Situational Awareness," *2014 American Control Conference*, Institute of Electrical and Electronics Engineers (IEEE), jun 2014. doi:10.1109/acc.2014.6858922.

[25] Jaunzemis, A. D., Holzinger, M. J., and Luu, K. K., "Sensor Tasking for Spacecraft Custody Maintenance and Anomaly Detection Using Evidential Reasoning," *Journal of Aerospace Information Systems*, feb 2018, pp. 1–26. doi:10.2514/1.i010584.

[26] Gehly, S., Jones, B., and Axelrad, P., "Sensor Allocation for Tracking Geosynchronous Space Objects," *Journal of Guidance, Control, and Dynamics*, aug 2016, pp. 0–0. doi:10.2514/1.g000421.

[27] Hobson, T. A. and Clarkson, I. V. L., "A particle-based search strategy for improved Space Situational Awareness," *2013 Asilomar Conference on Signals, Systems and Computers*, Institute of Electrical & Electronics Engineers (IEEE), nov 2013. doi:10.1109/acssc.2013.6810418.

[28] Hobson, T., Gordon, N., Clarkson, I., Rutten, M., and Bessell, T., "Dynamic steering for improved sensor autonomy and catalogue maintenance," *Advanced Maui Optical and Space Surveillance Technologies Conference, Wailea, HI*, 2014.

[29] Hobson, T., Clarkson, I., Bessell, T., Rutten, M., Gordon, N., Moretti, N., and Morreale, B., "Catalogue Creation for Space Situational Awareness with Optical Sensors," *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2016.

[30] Worthy III, J. L., Holzinger, M. J., and Scheere, D. J., "An Optimization Based Approach to Corellation of Observations with Uncertainty," *AIAA/AAS Space Flight Mechanics Conference*, 2016.

[31] Fujimoto, K. and Scheeres, D., "Correlation of optical observations of earth-orbiting objects and initial orbit determination," *Journal of guidance, control, and dynamics*, Vol. 35, No. 1, 2012, pp. 208–221.

[32] DeMars, K. J., Jah, M. K., and Schumacher, P. W., "Initial orbit determination using short-arc angle and angle rate data," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 48, No. 3, 2012, pp. 2628–2637.

[33] DeMars, K. J. and Jah, M. K., "Probabilistic initial orbit determination using gaussian mixture models," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 5, 2013, pp. 1324–1335.

[34] Howell, S. B., *Handbook of CCD Astronomy*, Cambridge University Press, 2nd ed., 2006. doi:10.1017/cbo9780511807909.

[35] King, I. R., "The Profile of a Star Image," *Publications of the Astronomical Society of the Pacific*, Vol. 83, apr 1971, pp. 199. doi:10.1086/129100.

[36] Vanvliet, L. J., *Grey-scale measurements in multi-dimensional digitized images*, Ph.D. thesis, Delft University, 10 1993.

[37] Mullikin, J. C., van Vliet, L. J., Netten, H., Boddeke, F. R., van der Feltz, G., and Young, I. T., "Methods for CCD camera characterization," *Image Acquisition and Scientific Imaging Systems*, edited by H. C. Titus and A. Waks, SPIE, may 1994. doi:10.1117/12.175165.

[38] Zhang, B., Zerubia, J., and Olivo-Marin, J.-C., "Gaussian approximations of fluorescence microscope point-spread function models," *Applied Optics*, Vol. 46, No. 10, mar 2007, pp. 1819. doi:10.1364/ao.46.001819.

[39] Lindegren, L., "Photoelectric astrometry: a comparison of methods for precise image location," *International Astronomical Union Colloquium*, Vol. 48, 1978, pp. 197–217. doi:10.1017/S0252921100074157.

[40] Hagen, N. and Dereniak, E. L., "Gaussian profile estimation in two dimensions," *Applied Optics*, Vol. 47, No. 36, dec 2008, pp. 6842. doi:10.1364/ao.47.006842.

[41] Davey, S. J., Rutten, M. G., and Cheung, B., "A Comparison of Detection Performance for Several Track-before-Detect Algorithms," *EURASIP Journal on Advances in Signal Processing*, 2008. doi:10.1155/2008/428036.

[42] Fujimoto, K., Uetsuhara, M., and Yanagisawa, T., "Statistical Track-Before-Detect Methods Applied to Faint Optical Observations of Resident Space Objects," *Advanced Maui Optical and Space Surveillance Technical Conference*, 2015.

[43] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Visual Tracking Methods for Improved Sequential Image-Based Object Detection," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 1, jan 2018, pp. 74–87. doi:10.2514/1.g002238.

[44] Linares, R., Jah, M. K., Crassidis, J. L., and Nebelecky, C. K., "Space Object Shape Characterization and Tracking Using Light Curve and Angles Data," *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 1, jan 2014, pp. 13–25. doi:10.2514/1.62986.

[45] Coder, R. D., Holzinger, M. J., and Linares, R., "Three-Degree-of-Freedom Estimation of Agile Space Objects Using Marginalized Particle Filters," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 2, feb 2018, pp. 388–400. doi:10.2514/1.g001980.

[46] Kaasalainen, M., "Optimization Methods for Asteroid Lightcurve Inversion I. Shape Determination," *Icarus*, Vol. 153, No. 1, sep 2001, pp. 24–36. doi:10.1006/icar.2001.6673.

[47] Kaasalainen, M., "Optimization Methods for Asteroid Lightcurve Inversion II. The Complete Inverse Problem," *Icarus*, Vol. 153, No. 1, sep 2001, pp. 37–51. doi:10.1006/icar.2001.6674.

[48] Coder, R. D., Wetterer, C. J., Hamada, K. M., Holzinger, M. J., and Jah, M. K., "Inferring Active Control Mode of the Hubble Space Telescope Using Unresolved Imagery," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 1, jan 2018, pp. 164–170. doi:10.2514/1.g002223.

[49] Mahler, R. P., *Statistical multisource-multitarget information fusion*, Artech House, Inc., 2007.

[50] Blackman, S., "Multiple Hypothesis Tracking For Multiple Target Tracking," *IEEE A and E Systems Magazine*, 2004.

[51] Reed, I., Gagliardi, R., and Shao, H., "Application of three-dimensional filtering to moving target detection," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. AES-19, No. 6, 1983, pp. 898–905. doi:10.1109/taes.1983.309401.

[52] Tonissen, S. and Evans, R., "Peformance of dynamic programming techniques for Track-Before-Detect," *IEEE Trans. Aerosp. Electron. Syst.*, Vol. 32, No. 4, 1996, pp. 1440–1451. doi:10.1109/7.543865.

[53] Zhai, C., Shao, M., Nemati, B., Werne, T., Zhou, H., Turyshev, S. G., Sandhu, J., Hallinan, G., and Harding, L. K., "Detection of a Faint Fast-Moving Near-Earth Asteroid using the Synthetic Tracking Technique," *ApJ*, Vol. 792, No. 1, aug 2014, pp. 60. doi:10.1088/0004-637x/792/1/60.

[54] Ristic, B., Arulampalam, S., and Gordon, N. J., *Beyond the Kalman filter: Particle filters for tracking applications*, Artech house, 2004.

[55] Daum, F. and Huang, J., "Curse of dimensionality and particle filters," *Aerospace Conference, 2003. Proceedings. 2003 IEEE*, Vol. 4, IEEE, 2003, pp. 4_1979–4_1993.

[56] Hoseinnezhad, R., Vo, B.-N., Vo, B.-T., and Suter, D., "Visual tracking of numerous targets via multi-Bernoulli filtering of image data," *Pattern Recognition*, Vol. 45, No. 10, oct 2012, pp. 3625–3635. doi:10.1016/j.patcog.2012.04.004.

[57] Hoseinnezhad, R., Vo, B.-N., and Vo, B.-T., "Visual Tracking in Background Subtracted Image Sequences via Multi-Bernoulli Filtering," *IEEE Transactions on Signal Processing*, Vol. 61, No. 2, jan 2013, pp. 392–397. doi:10.1109/tsp.2012.2222389.

[58] Hussein, I., DeMars, K. J., Früh, C., Erwin, R. S., Jah, M. K., et al., "An AEGIS-FISST integrated detection and tracking approach to Space Situational Awareness," *Information Fusion (FUSION), 2012 15th International Conference on*, IEEE, 2012, pp. 2065–2072.

[59] Levesque, M., "Automatic reacquisition of satellite positions by detecting their expected streaks in astronomical images," *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2009, p. E81.

[60] North, D. O., "An analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems," *Proceedings of the IEEE*, Vol. 51, No. 7, 1963, pp. 1016–1027. doi:10.1109/proc.1963.2383.

[61] Mohanty, N., "Computer tracking of moving point targets in space," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, , No. 5, 1981, pp. 606–611.

[62] Zingarelli, J. C., Pearce, E., Lambour, R., Blake, T., Peterson, C. J., and Cain, S., "Improving the Space Surveillance Telescope's Performance Using Multi-Hypothesis Testing," *The Astronomical Journal*, Vol. 147, No. 5, 2014, pp. 111.

[63] Dragovic, M., "Velocity filtering for target detection and track initiation," Tech. rep., DTIC Document, 2003.

[64] Milani, A., Gronchi, G. F., Vitturi, M. d., and Knežević, Z., "Orbit determination with very short arcs. I admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 90, No. 1-2, 2004, pp. 57–85.

[65] Farnocchia, D., Tommei, G., Milani, A., and Rossi, A., "Innovative methods of correlation and orbit determination for space debris," *Celestial Mechanics and Dynamical Astronomy*, Vol. 107, No. 1-2, 2010, pp. 169–185. doi:10.1007/s10569-010-9274-6.

[66] Zezula, P., Amato, G., Dohnal, V., and Batko, M., *Similarity search: the metric space approach*, Vol. 32, Springer Science & Business Media, 2006. doi:10.1007/0-387-29151-2.

[67] Uhlmann, J. K., "Satisfying general proximity/similarity queries with metric trees," *Information processing letters*, Vol. 40, No. 4, 1991, pp. 175–179. doi:10.1016/0020-0190(91)90074-r.

[68] Chávez, E., Navarro, G., Baeza-Yates, R., and Marroquín, J. L., "Searching in metric spaces," *ACM computing surveys (CSUR)*, Vol. 33, No. 3, 2001, pp. 273–321. doi:10.1145/502807.502808.

[69] Batko, M., Novak, D., and Zezula, P., "MESSIF: Metric Similarity Search Implementation Framework," *Digital Libraries: Research and Development*, Springer Science, Business Media, 2007, pp. 1–10. doi:10.1007/978-3-540-77088-6_1.

[70] Gionis, A., Indyk, P., Motwani, R., et al., "Similarity search in high dimensions via hashing," *VLDB*, Vol. 99, 1999, pp. 518–529.

[71] Worthy, J. L. and Holzinger, M. J., "Use of uninformative priors to initialize state estimation for dynamical systems," *Advances in Space Research*, Vol. 60, No. 7, oct 2017, pp. 1373–1388. doi:10.1016/j.asr.2017.06.040.

[72] Worthy III, J. L. and Holzinger, M. J., "Dempster-Shafer Theory Applied to Admissible Regions," *AIAA/AAS Space Flight Mechanics Conference*, 2017.

[73] Vo, B.-N., Vo, B.-T., Pham, N.-T., and Suter, D., "Joint Detection and Estimation of Multiple Objects From Image Observations," *IEEE Transactions on Signal Processing*, Vol. 58, No. 10, 10 2010, pp. 5129–5141. doi:10.1109/tsp.2010.2050482.

[74] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Orbit Determination for Partially Understood Object via Matched Filter Bank," *AAS/AIAA Astrodynamics Specialists Meeting*, 2015.

[75] Ristic, B., *Particle Filters for Random Set Models*, Springer Science + Business Media, 2013. doi:10.1007/978-1-4614-6316-0.

[76] Reuter, S., Vo, B.-T., Vo, B.-N., and Dietmayer, K., "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, Vol. 62, No. 12, jun 2014, pp. 3246–3260. doi:10.1109/tsp.2014.2323064.

[77] Benezeth, Y., Jodoin, P., Emile, B., Laurent, H., and Rosenberger, C., "Review and evaluation of commonly-implemented background subtraction algorithms," *2008 19th International Conference on Pattern Recognition*, Institute of Electrical & Electronics Engineers (IEEE), dec 2008, pp. 1–4. doi:10.1109/icpr.2008.4760998.

[78] Sanson, F. and Frueh, C., "Noise Quantification in Optical Observations of Resident Space Objects for Probability of Detection and Likelihood," *AAS/AIAA Astrodynamic Specialist Conference, Vail, CO*, 2015, pp. 15–634.

[79] Tommei, G., Milani, A., and Rossi, A., "Orbit determination of space debris: admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 97, No. 4, 2007, pp. 289–304. doi:10.1007/s10569-007-9065-x.

[80] Gutin, G. and Punnen, A. P., editors, *The Traveling Salesman Problem and Its Variations*, Springer US, 2007. doi:10.1007/b101971.

[81] Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P., "Optimization by Simulated Annealing," *Science*, Vol. 220, No. 4598, may 1983, pp. 671–680. doi:10.1126/science.220.4598.671.

[82] Chavel, I., *Riemannian Geometry*, Cambridge University Press, 2006. doi:10.1017/cbo9780511616822.

[83] Stromberg, K., *An Introduction to Classical Real Analysis*, American Mathematical Society, oct 2015. doi:10.1090/chel/376.

[84] Lepage, G., Bogaerts, J., and Meynants, G., "Time-Delay-Integration Architectures in CMOS Image Sensors," *IEEE Transactions on Electron Devices*, Vol. 56, No. 11, nov 2009, pp. 2524–2533. doi:10.1109/ted.2009.2030648.

[85] Vanderplaats, G. N., "Multidiscipline Design Optimization," *Applied Mechanics Reviews*, Vol. 41, No. 6, jun 1988, pp. 257. doi:10.1115/1.3151897.

[86] Murphy, T. S., Holzinger, M. J., Luu, K. K., and Sabol, C., "Optical Sensor Follow-Up Tasking on High Priority Uncorrelated Track," *AAS/AIAA Space Flight Mechanics Meeting*, 2017.

[87] Massey, P. and Jacoby, G. H., "CCD data: The good, the bad, and the ugly," *Astronomical CCD observing and reduction techniques*, Vol. 23, 1992, p. 240.

[88] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Space Object Detection in Images Using Matched Filter Bank and Bayesian Update," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 3, mar 2017, pp. 497–509. doi:10.2514/1.g001934.

[89] Kay, S. M., *Fundamentals of Statistical Signal Processing: Estimation Theory*, Prentice Hall, 1993.

[90] Howell, S. B., "CCD manufacturing and operation," *Handbook of CCD Astronomy*, Cambridge University Press, 2006, pp. 8–35. doi:10.1017/cbo9780511807909.004.

[91] Sease, B., Flewelling, B., and Black, J., "A class of convex optimization problems for template-based star subtraction," *26th AAS/AIAA Space Flight Mechanics Meeting*, 2016.

[92] Arora, J., *Introduction to optimum design*, Elsevier, 2004.

[93] Pratt, W. K., *Digital Image Processing*, Wiley-Blackwell, jan 2007. doi:10.1002/0470097434.

[94] Turin, G. L., "An introduction to matched filters," *IRE Transactions on Information Theory*, Vol. 6, No. 3, jun 1960, pp. 311–329. doi:10.1109/tit.1960.1057571.

[95] Montgomery, D. C. and Runger, G. C., *Applied statistics and probability for engineers*, John Wiley & Sons, 2010.

[96] Worthy III, J. L. and Holzinger, M. J., "Use of Uninformative Priors to Initialize State Estimation for Dynamical Systems," *AIAA/AAS Astrodynamics Specialist Conference*, 2015.

[97] Mortari, D., Bruccoleri, C., La Rosa, S., and Junkins, J. L., "CCD Data Processing Improvements for Star Cameras," *International Conference on Dynamics and Control of Systems and Structures in Space*, 2002.

[98] Kouprianov, V., "Distinguishing features of CCD astrometry of faint GEO objects," *Advances in Space Research*, Vol. 41, No. 7, 2008, pp. 1029–1038. doi:10.1016/j.asr.2007.04.033.

[99] Worthy, J. L., Holzinger, M. J., and Scheeres, D. J., "An optimization approach for observation association with systemic uncertainty applied to electro-optical systems," *Advances in Space Research*, mar 2018. doi:10.1016/j.asr.2018.02.041.

[100] Sease, B., Murphy, T., Flewelling, B., Holzinger, M. J., and Black, J., "Enabling direct feedback between initial orbit determination and sensor data processing for detection and tracking of space objects," *SPIE Defense+ Security*, International Society for Optics and Photonics, 2015, pp. 94690M–94690M. doi:10.1117/12.2181980.

[101] Virani, S., Murphy, T., Holzinger, M., and Jones, B., "Empirical Dynamic Data Driven Detection and Tracking Using Detectionless and Traditional FiSSt Methods," *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2017.

[102] Jaunzemis, A., Minotra, D., Holzinger, M. J., Feigh, K. M., Chane, M. W., and Shenoyf, P. P., "Judicial Evidential Reasoning for Decision Support Applied to Orbit Insertion Failure," *1 st IAA Conference on Space Situational Awareness (ICSSA)*, 2017.

VITA

Timothy Murphy is a Mechanical and Aerospace Engineer who received his Bachelors of Science in mechanic engineering from The George Washington University in 2013. At Georgia Tech, he took an interest in the areas of astrophysics, mathematical analysis, and statistical estimation and decided to create his PhD thesis around this foundation. Much of his work and his personal research style comes from the mentorship of Marcus Holzinger and Brien Flewelling.

Timothy originally comes from Drexel Hill, a small suburban town outside of Philadelphia, Pennsylvania. After graduation, he plans to settle in Maryland near Washington, D.C. He has an extensive interest in boardgames, roleplaying games, and analog gaming in general. His secret life goal is to design and publish a boardgame of his own. He also has interests in ballroom and latin dance, cooking, and travel.