Search and Pursuit of Non-cooperative Targets with Unmanned Aerial Vehicles

THESIS PROPOSAL

Michael Dille
November 2011

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Sanjiv Singh, Chair
Benjamin Grocholsky
Maxim Likhachev
Paul Scerri
Thanasis Kehagias, Aristotle University, Greece

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Abstract

Across many rescue, surveillance, and scientific applications, there exists a broad need to perform wide-area reconnaissance and terrain surveys, for which unmanned aerial vehicles (UAVs) are increasingly popular. This thesis considers the task of using one or more UAVs to locate an object of interest, provide continuous viewing, and rapidly re-acquire tracking should it be lost for any reason.

For the common class of small field-launched UAVs considered, this is a difficult task due to a small sensor field of view, highly uncertain estimates of UAV pose, and limited maneuverability, requiring careful processing of observations and choosing flight paths to best find the object or keep it in view. Existing strategies for accomplishing this provide poor estimates of the object's location and rely on heuristic or computationally intensive trajectory generation for both pursuit and search.

This thesis proposes careful representation of both the environment structure and observation uncertainty as a means to simplify and better model the problem. Work to date has demonstrated substantial improvements in object location error through filter representations designed for high-uncertainty observations and explored the use of terrain constraints to improve pursuit performance. Proposed work expands this to greater exploitation of environmental structure, particularly the common case of networks of roads, to both permit treatment as a simpler discrete problem and enable mapping the problem to one of several well-studied pursuit-evasion abstractions. Ideas for extending this to more general environments are also considered.

Extensive field trials using widely-fielded vehicles have validated portions of the proposed ideas, laying the groundwork to evaluate and demonstrate algorithms to be developed on real aircraft and in simulation as practical. This work directly contributes to a strong practical need for greater automation of such UAVs to reduce operator workload and enable rapid completion of urgent missions with minimal resources.

Contents

CHAPTER 1

Introduction

1.1 Motivation

Wide-area reconnaissance and survey missions are a key element of many search and rescue, disaster response, military reconnaissance, law enforcement surveillance, and scientific exploration scenarios. Examples include searching for missing persons in the wilderness or at sea, detecting and tracking the spread of wildfires, locating and pursuing fleeing suspects, protecting a perimeter for a moving convoy, finding meteorites, mapping floating ocean debris, and identifying and tracking the movement of wildlife.

The character of such missions has evolved with social and computing trends, greatly affecting approach methodology. Typically, these are extremely data-driven and information rich, since a terrain map (e.g. from satellite imagery) and prior information on the relative priority of area regions are likely to be available, and mission results are frequently aggregated—often in real-time—for visualization or pattern detection. These missions are also extremely dynamic in that additional information—from witness reports, external sensors, or other responding agencies—may arrive at any time, and a well-executed mission plan must be able to react and adapt as necessary. Time-sensitivity is another probable characteristic, both because lives may be at stake and due to large operational costs, implying a need for rapidly field-deployed sensing, some form of optimal mission planning to minimize execution time, and automatic information passing. Finally, considerable risk-aversion is likely

(a) Wildfire detection and monitoring (b) Flood survivor identification (c) Tracking pirate activity

(d) Convoy perimeter protection (e) Pursuit of fleeing suspects

Figure 1.1: Example scenarios requiring substantial reconnaissance and survey components over a wide area.

to play a role, requiring that mission actions be thoroughly documented, follow strict procedures, and minimize the use of human intervention to avoid introducing errors or placing further lives at risk, all of which suggest extensive use of automation.

Particular commonalities of such scenarios may be further distilled as a step towards formalization. In general, these include one or more *observers* seeking some object of interest (henceforth simply referred to as the *target*). Potential prior knowledge about the target's possible location that can be updated with every observation (or lack of observation), and if mobile, the target (by virtue of being an inanimate object, wild animal, fleeing suspect, or an out-of-communication victim) will be non-cooperative and its motions non-deterministic. Once detected, multiple observations or persistent observation is desirable, and each observation may have some quality or error dependent on the relative pose of the observer and target (for instance due to terrain occlusion, range to the target, or intrinsic sensor properties). Meanwhile, the environment in which the mission takes place is likely to be hazardous and poorly traversable, but its structure well known from existing maps. Together, these properties suggest a general mission structure and desirable observer characteristics.

Independently, unmanned aerial vehicles (UAVs) have come to form the basis of a rapidly growing field of study and are seeing increasing real-world use, frequently

(a) RQ-4 GlobalHawk 35m wingspan 3,800kg empty | (b) RQ-8A Fire Scout 8m rotor diam. 940kg empty | (c) Scout 80cm rotor diam. 1.3kg typical | (d) Hummingbird 16cm wingspan 19g total

Figure 1.2: Several unmanned aircraft demonstrating a spectrum of portability roughly parameterized by the magnitude of vehicle size, weight, cost, feature set, flight altitude, mission duration, and launch facilities. These include (a) the Northrop Grumman RQ-4 Global Hawk, (b) the Northrop Grumman RQ-8 Fire Scout, (c) the Aeryon Scout, and (d) the AeroVironment Nano Hummingbird.

for such missions. Their viability may largely be attributed to advances in sensor miniaturization, computing power, battery chemistry, and composite materials, coinciding with increases in commercial production volume in each area for use in consumer electronics. Simultaneously, a shift in defense strategy towards unmanned vehicles, particularly aircraft, [74] has resulted in a major infusion of research funding in the area and the development of a considerable variety of UAVs by new and well-established companies. Combined, these advances and industry shifts have made available a wide spectrum of reliable, capable, and relatively affordable UAVs ranging in size from tens of grams to over ten thousand kilograms.

UAVs are particularly well suited to the aforementioned scenarios, having a number of advantages over ground vehicles. Most are rapidly deployable from a distance, without a need to assemble and transport human teams into an area—after verifying it is safe to do so—for a manned mission or to deploy an unmanned ground vehicle (UGV) within the affected area. Further, they are fast moving and thus able to reach the area of interest and cover a wide space quickly. By virtue of being airborne, UAVs are not subject to terrain obstacles, debris, and hazards below their flight altitude, while a wide area view is possible with minimal occlusion from obstacles ground vehicles could not see over. Additionally, given this lack of occlusion, reliable and long-range line-of-sight communication from a ground station or satellite is feasible.

Within the broad category of reconnaissance UAVs, a further distinction may be made along a spectrum of what might generally be termed portability—parameterized by increasing size, weight, cost, feature set, flight altitude, mission duration, and facilities required for deployment—depicted in Figure 1.2. At one extreme lies, for instance, the high-altitude Northrop Grumman RQ-4 Globalhawk, with a wingspan of 35 meters, 36-hour endurance, a complex sensor package including high-resolution Synthetic Aperture Radar, and a unit cost of approximately US$200 million, amortizing development expenses. [110] At the other end may be found the likes of the AeroVironment Nano Hummingbird, with a wingspan of 16 centimeters and weighing 19 grams, intended for short-duration hovering or perching surveillance in urban environments using only a low-resolution onboard camera. [3]

In between lies the class of aircraft considered as the focus of this thesis: small, human- or van-portable field-reconnaissance UAVs, having a wingspan of one to three meters, weighing one to ten kilograms, and flying at between 50 and several hundred meters above ground level. Commonly called small UAVs (SUAVs) or micro air vehicles (MAVs), these are most appropriate for a wide range of rapid-response reconnaissance missions in hazardous terrains for several reasons. First, they may be launched in minutes from a small clearing, with little to no infrastructure such as a runway or launch-propulsion equipment, very near the area of interest. Further, launch and operation need be performed by at most several lightly-trained operators using portable battery-powered ground stations. Finally, their relatively low cost—one to several tens of thousands of dollars—admits a substantial level of expendability, permitting particularly risky yet possibly high-reward missions over hostile or inaccessible terrain from which retrieval is impractical. However, design simplification necessary for weight, cost, and operator workload minimization results in a number of drawbacks. Control surfaces such as ailerons are often eliminated, limiting maneuverability and weakening control authority necessary to reject wind disturbances. Onboard autopilots may provide a simple (e.g. waypoint-based) high-level interface that benefits operators but which limits the complexity of algorithmically-generated flight paths. Onboard sensing is limited to lower-end (lightweight and low-cost) inertial measurement units (IMUs) and global positioning systems (GPSes) providing limited accuracy, with low- to medium-resolution cameras. This is aggravated by infrastructure-less recovery methods that are often little more than controlled crashes, the impacts of which disturb calibrated relative sensor mountings. Finally, limited processing power is available onboard, requiring that primarily compressed raw sensor data (commonly camera video) rather than smaller processed data (such as image pixel locations of detected targets) be sent over limited-bandwidth downlinks. Common designs therefore include relatively narrow-angle lensing to provide the highest-resolution imagery for a given area of ground, which, combined with low-altitude flight, produces a small sensor footprint. The result is a very distinct reconnaissance problem from that encountered with larger, higher-altitude UAVs: rather than needing to detect and identify objects of interest from within a wide sensor footprint (primarily a signal processing task), motion of the UAV itself to direct a smaller sensor footprint toward likely target locations or vantage points (an estimation and planning task) is required.

(a) ScanEagle
3m wingspan
13kg empty

(b) Joker 2
1.8m rotor diam.
8kg loaded

(c) RQ-16 T-Hawk
35cm diam.
8.4kg typical

(d) RQ-11 Raven
130cm wingspan
1.9kg typical

Figure 1.3: Examples of small field-reconnaissance UAVs considered primarily in this thesis, such as the Insitu ScanEagle catapult-launched and hook-retrieved tactical surveillance UAV, (a) the Maxi-Joker 2 camera-transport helicopter, (b) the Honeywell RQ-16 T-Hawk vertical takeoff and landing ducted-fan UAV, (c) and the hand-launched AeroVironment Raven local reconnaissance UAV. (d)

1.2 General Problem Statement

Using one or more UAVs to locate a target and provide persistent surveillance that generates an estimate of the target's location[1] is both an extremely desirable task

[1]The task of processing sensor observations (in sensor-local or vehicle-local coordinates) and converting them into a best estimate of a target's location in world (geographic) coordinates is commonly called *geolocation* and will be henceforth referred to as such.

in practice as well as the implicit objective of a considerable amount of literature on UAV search and tracking. Nevertheless, it appears to only have been first formally stated by Frew and Elston as the continuous cooperative search, acquisition, and tracking (CSAT) problem for active robot sensor networks. [36] In that framework, the three components are (1) an area search or coverage control task to traverse a region with a finite sensor footprint while optimizing a performance objective to detect one or more targets, (2) solving a task assignment or resource allocation task to assign each robot to a location somehow best covering all targets, and (3) performing cooperative tracking and estimation to both guide the motion of robots around detected targets and estimate target locations. Though phrased generically, it is most applicable to UAVs in that it considers robots with motion constraints, finite sensor footprints coupled to a robot's location, and inaccurate observations used to produce an estimate of target location.

As the focus of this thesis are the elements of search and tracking, the CSAT components defined by Frew of acquisition (target identification) and observer assignment will not be considered explicitly and instead treated as part of the transition between search and tracking. The problem may then be modeled as depicted in Figure 1.4, in which a target is first sought while in the *search* state, then, once detected, actively tracked via continuous *surveillance and geolocation* while observations are received. If tracking is lost (for instance, due to too much time elapsing since the last observation), the search state is re-entered to *recapture* the target. Owing to important differences between a target having completely unknown location that must be found and one that has been recently lost, recapture is treated as an independent mission state in the more formal problem decomposition of Section 2.2, however in its simplest form the problem may be viewed as having two high-level states: one in which the target is being observed, and another in which it is not and must be first located.

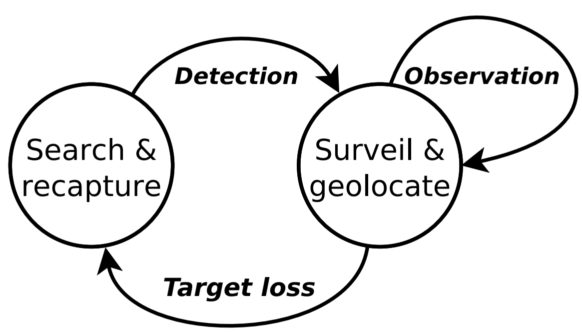

Figure 1.4: State diagram of the cooperative search, acquisition, and tracking problem, with the acquisition step folded into the transition between search and tracking.

Problem Subtasks

The high-level subtask loops of these two states and their interpretations in the context of this thesis, repeated continuously while in either state, may be stated as follows:

Search

1. Propagate prior knowledge of possible target position to present time.

   A prior map of possible target locations—perhaps stored as a map of relative probability or unsearched regions—may be available, and if so, it must be updated with the passage of time to reflect possible target motion into previously cleared or unlikely regions.

2. Move to optimize detection likelihood, prevent escape, or both.

   Flight paths are chosen to steer UAV sensor footprints to optimize one or more performance metrics, such as maximizing the expected number of targets detected in a given interval, minimizing the expected time until a target is detected, minimizing the probability of escape, or possibly guaranteeing that any targets in an area will be detected.

3. Process sensor readings to determine whether detected.

   As areas are traversed by a sensor footprint, a target detection or recognition mechanism must be executed to report whether a target is in view, initiating a transition to the tracking state.

4. Record searched areas and update possible target locations.

   Knowledge that a target is not present in a given region may affect future flight paths, for instance by marking the area as cleared or having reduced likelihood of containing a target, and a prior target location map indicating this may be provided to the next iteration of the search.

Tracking

1. Process sensor readings for continuous tracking.

   Sensor readings must be processed to determine whether the tracked target is currently in view, and if so, generate an observation.

2. Build and refine sensor detection or tracking model.

Given repeated observations of a target, it may be possible to improve detection accuracy or speed by automatically optimizing detection parameters for the currently tracked target, effectively learning its appearance, fingerprint, or signature that may change over time with the observer's perspective.

3. Filter observations to estimate target location.

   Every successful observation of a target may be used to improve or update its geolocation estimate, producing an estimate more accurate than any single observation.

4. Anticipate target motion between observations.

   The target's geolocation history may be used to predict its future location at the time of the next anticipated observation, for instance by computing and apply a velocity estimate or determining that the target is likely to be stationary.

5. Move to improve estimate certainty and maintain view.

   To maintain tracking, a path for each sensor footprint must be generated so as to produce good observations in the future, for instance by maintaining its placement over the target's predicted location.

Scope of Study

As the search and tracking problem as stated is incredibly broad and any practical implementation extremely complex, a number of assumptions and simplifications must be made for the purpose of this thesis. Except as noted, the following facets will be neglected for the remainder of this document and are assumed to be separately provided or resolved by existing or independent modules. Many have also received considerable attention in related literature. Key facets that will not be considered beyond references to related work in Section 3.2 include:

- Self-localization and environment mapping
- Automatic target recognition or data association
- Aerodynamics and low-level vehicle control
- Sensor aim or gimbal state planning
- Obstacle avoidance and agent deconfliction
- Distributed or decentralized estimation and control

Related Problems

The problem of continuous search and tracking for surveillance and geolocation is quite similar to but sharply distinct from a number of existing well-studied problems. Some of the most related include planar area coverage, (polygonal) region clearing, generalized pursuit-evasion, and so-called moving target indicator (MTI) tracking using radar. A brief statement of these problems and their differences is provided in Section 2.3.

1.3 Summary of Proposed Contributions

The general theme of this thesis is to simplify the search and pursuit problem by exploiting natural terrain constraints on a target that may be present, thereby reducing resource requirements and improving observable performance, in terms such as search time, geolocation error, frequency and duration of target loss, and recapture time. The form of such constraints considered is primarily that of roads within environments, though generalizations are also considered.

To accomplish this, it is proposed that results from several related fields be leveraged as toolboxes, with aspects and instances of this problem mapped to abstract problem instances for which theoretical results or tractable algorithms are known. These include the use of constrained and multi-modal estimation for geolocation from existing large-scale tracking and guidance problems (Section 5.3), discrete recapture and local search strategies from pursuit-evasion (Section 5.4), and wide-area constrained-target search from the field of graph search (Section 5.5).

High-level proposed contributions of this thesis include:

- A detailed study of the search and tracking problem including a decomposition into more specific problems as parameters are varied (Section 2.2)
- A study of and development of strategies for recapture as its own special case distinct from search or pursuit (Section 3.5)
- Improved geolocation representations and filters for unconstrained environments (Section 4.3)
- A practical open area search strategy optimized for high-level autopilots and slowly reacting target detectors (Section 4.2)
- Improved target pursuit by exploitation of road-network graph structure applying constrained and multi-hypothesis estimation (Section 5.3)

- Improved search and recapture also exploiting road-network graph structure applying graph search and pursuit techniques (Sections 5.4 and 5.5)
- Practical performance improvement through modeling and exploitation of target intent enabled by constrained environments (Section 5.3)
- Real-world implementation and field validation of many of the proposed ideas (Sections 4.1 and 5.8)

These ideas may be summarized in the proposed statement of this thesis:

> *By exploiting natural terrain constraints on a target's location and motion, the tractability and performance of UAV search and pursuit may be greatly improved.*

1.4 Document Outline

The remainder of this document is structured as follows:

Chapter 2 presents subcases of the search and tracking problem arising as parameters are varied and defines the task optimization problem in terms of observation inputs and control outputs.

Chapter 3 summarizes existing work in target search, geolocation, and pursuit and considers challenges both due to weaknesses in these approaches and fundamental to the tasks.

Chapter 4 covers work to date, including experiments providing a practical need for resulting improvements in search and pursuit representations, motivating further exploitation of environmental structure.

Chapter 5 proposes strategies for exploiting such structure by mapping tasks to simpler abstractions for which tractable, well-performing solutions are feasible, considers specific instances for the case of environments made of target-constraining roads and road networks, and suggests a means to validate these in future experiments.

Chapter 6 summarizes completed and proposed work, presents an anticipated timeline, and suggests possible future work that both may represent feasible thesis extensions and interesting ideas not likely within its scope.

CHAPTER 2

Problem Study

This chapter examines and more precisely defines the broad search and tracking problem studied thus far. Section 2.1 presents such a definition and phrases it as an abstract optimization task with varying objectives depending on the task phase and mission requirements. It is then treated as a general continuous sensor positioning problem that reduces to one of several simpler cases as properties of the target are varied in Section 2.2. Section 2.4 raises the importance of using environmental knowledge to guide the choice of approach given its effect on possible target position and motion. Finally, Section 2.5 chooses simple but largely realistic models for UAV sensing and motion to permit later formalization, algorithmic development, and results comparison.

2.1 Problem Definition

The search and tracking task considered may be concisely stated as:

> Locate a potentially evasive target in a bounded environment using one or more UAVs, provide persistent sensor coverage and continuous geolocation estimates once it is detected, and reacquire tracking when lost.

The first phase is a *search* task in which a trajectory for each UAV is sought that best (most quickly or with the greatest assurance) detects a target. The second is primarily a *pursuit* task in which trajectories are sought that maintain a sensor footprint over the target that provides good observations for geolocation. Re-acquisition attempts to *recapture* a target for which tracking was lost by performing a local search before uncertainty in the target's position resulting from possible motion grows too large. Not considered as an independent mission state in the original CSAT formulation, a recapture phase is considered explicitly throughout this thesis given important distinctions in target location knowledge from full-fledged search.

The general problem may be phrased as a UAV control optimization problem in which a flight path is chosen for each UAV with one of the following objective functions chosen based on phase and varying mission requirements.

Locating or re-acquiring:

- Maximize probability of detection over given search duration
- Minimize expected time until detection
- Minimize search time providing at least a given probability of detection (Important special case: p=1, or guaranteed search or recapture)

Tracking:

- Maximize expected time in view over a given tracking duration
- Minimize expected average or terminal geolocation uncertainty over a given tracking duration
- Minimize expected tracking time until geolocation uncertainty drops below a given threshold

Optimizing these objectives may be formally, albeit for now quite abstractly, stated. Let $\mathbf{x}_{\text{UAV}}$ be the full state of the UAV and $\mathbf{x}_{\text{tgt}}$ represent all information known about possible target position. Let $\mathcal{U}$ be the space of UAV control inputs at a given instant for all UAVs, elements of which depend on the precise motion model considered but which may for example be of the form of a bank angle to hold, a waypoint to head towards, a thrust vector to apply, or a multi-dimensional vector of subsystem commands (e.g., a rudder command and sensor pointing angle). Each implicitly contains a duration for which to apply the control command. Likewise, let $J(\mathbf{U}, \mathbf{x}_{\text{UAV}}, \mathbf{x}_{\text{tgt}})$ be the cost function evaluating a discrete-time control plan $\mathbf{U} \in \mathcal{U}^k$ of length k under a given objective. Then,

$$\mathbf{U}^* = \underset{\mathbf{U}}{\operatorname{argmin}}\, J(\mathbf{U}, \mathbf{x}_{\text{UAV}}, \mathbf{x}_{\text{tgt}})$$

is the optimal sequence of control inputs to apply. Depending on the objective, the sequence length k may be a fixed or free variable.

The UAV motion model defining $\mathcal{U}$ considered for the purpose of this thesis is given in Section 2.5 and specific cost functions for tracking and search are compared in Sections 3.3.3 and 3.4.2 respectively.

2.2 Problem Decomposition

An alternative view of the problem is as a continuous sensor positioning task with the immediate goal of providing useful observations of a target, implying the need to locate and reach it if not presently in view.

The three phases of the preceding section are then subcases of the same problem with differing certainty in target location. If a target's location is known—for instance because it is presently being tracked within the sensor footprint or had been very recently—the immediate objective of an observing UAV is *pursuit* of the target, moving so as to place the sensor footprint over the predicted future location of the target to maintain or quickly regain tracking and improve geolocation certainty. If a target's location is only roughly known—as lying within some small region or among a set of hypotheses—because it has not been observed recently, then *recapture* of the target is desired to prevent further growth in uncertainty. Lastly, if no estimate of a target's location is available or if it is at best contained within a large region, a wider-scale *search* is required to first locate the target to (re)commence pursuit. Characteristics of these cases and typical examples are summarized in Table 2.1, with intuitive graphical depictions given in Figure 2.1.

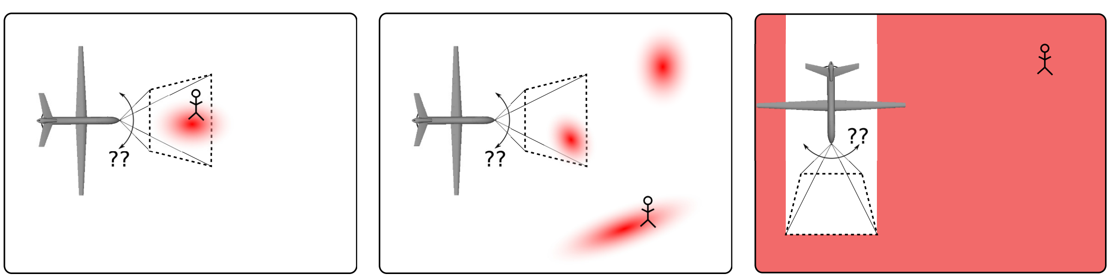

(a) Pursuit (b) Recapture (c) Search

Figure 2.1: Intuitive depictions of pursuit, recapture, and search, along with the need to steer the path of the observer. Red shading indicates the relative probability of target presence, and the trapezoidal region leading the UAV indicates the region visible to it (formalized in Section 2.5).

	Pursuit	Recapture	Search
Target location	Known (actively estimating)	High-uncertainty estimate	No prior or weak regional prior
Target observation	Recently observed	None recently	Possibly never
Examples	• Want to keep a target in view despite observer's or its motion • Target in view but location uncertainty is large or asymmetric • Target just left view and want to return to view quickly	• Pursued target was not in expected location after a brief loss of view • More targets than observers and must give up pursuit on one to check on another • Momentarily known location from beacon or remote intelligence and need to get back in view before loss	• Clearing an area containing an unknown number of targets • Identifying locations of targets known to exist within an area • Lost target not captured before uncertainty grew to entire area

Table 2.1: Characteristics of the three problem cases as target location certainty varies.

Another important aspect is knowledge of target motion, or simply target predictability. At one extreme, if a target is known to be *stationary* (or either has known velocity that may be directly subtracted or is cooperatively transmitting its location), the area simply need be covered until it is located, at which point the observer may choose flight paths around it providing the most desired observations. At the other, a completely *adversarial* or evasive target will presumably avoid detection and persistent surveillance, possibly requiring many UAVs executing carefully chosen trajectories to corner the target and block escape paths. In between lie *stochastically moving* targets, the motion of which is not precisely known but is modeled by an uncertainty distribution. This latter category accounts for non-cooperative targets that might otherwise act adversarially but are ignorant of or indifferent to their pursuers' presence.

Figure 2.2 depicts a further decomposition of the problem along this additional axis of target predictability. This exposes substantially more strictly defined and, in some cases, well-studied problems to which the general task reduces for a given level of certainty in target location and its predictability.

For an effectively stationary target, the problem is fundamentally an active sensor

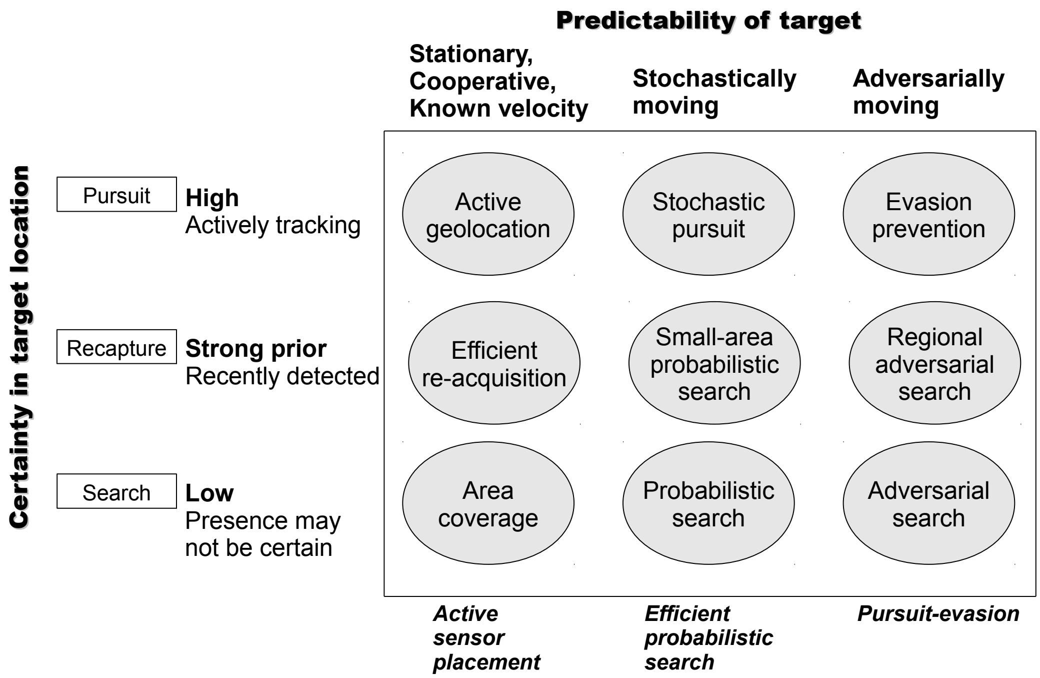

Figure 2.2: Decomposition of the unified search and tracking problem along the two axes of knowledge of target position and target predictability.

placement task with an unchanging world. For an actively tracked target, trajectories are chosen that are most likely to provide the best observations, for instance which provide lowest (possibly a-posteriori) geolocation uncertainty. Existing strategies for accomplishing this as the simplest case of pursuit are considered in Section 3.3.3. Meanwhile, a target in need of recapture in this context is one that view of which has been lost (perhaps due to wind-induced trajectory tracking error), and trajectories are required that quickly move the sensor footprint back into positions over the target providing similarly good observations. Typically, the same existing strategies apply to this case in that observations of most likely target locations are desired, and no concern is present for the target escaping. A target of unknown location, however, must first be located. To do this for a stationary target requires covering all locations the target might be, possibly prioritized by a relative likelihood, but again without concern for possible motion during the execution of the search that might recontaminate already-searched areas. Existing strategies for regular, uniform environments are reviewed in Section 3.4.1 and for irregularly-shaped or prioritized environments in Section 3.4.2, with an additional practical version contributed in Section 4.2.

The primary difference for a stochastically moving target is the need to continuously predict this motion and record resulting growth in uncertainty. For pursuit, this corresponds to predicting and diffusing likely future target locations when planning where to place sensor footprints for best observation. This corresponds to the more difficult general case of observation-predicting pursuit existing work described in Section 3.3.3 has considered. Extensions for improved geolocation accuracy and planning tractability for simple road-like environments and more general road networks are considered in Sections 4.4 and 5.3 respectively. Skipping to unknown-location targets requiring search, mere coverage is now insufficient because target motion may now, with varying likelihood, enter previously cleared areas. Existing strategies reviewed in Section 3.4.2 therefore consider probabilistic recontamination producing a changing irregular prior that is searched as for a stationary target with non-uniform location likelihood. Recapture of such targets has not previously been well considered, though fundamentally either a local-area probabilistic search or pursuit with a wide-area likelihood or multiple hypothesis may perform well. Intuitive possibilities are considered in Section 3.5, with more efficient strategies for road network environments proposed in Section 5.4.

Adversarial targets are of course the most challenging and are typically assumed to be able to move in any direction at any time, with either infinite or bounded speed, and surveilling them is fundamentally a pursuit-evasion task. Though a well-studied problem as a general abstract geometric problem, with several possibly-applicable abstractions considered in Section 3.1, this case has with few exceptions [85] not been considered for UAV pursuit and recapture. Given the applied nature of the task, this is likely due to the unrealism of omniscient high-speed targets, when in practice a real target might move a small fraction of the width of a sensor's field of view between observations. Search for fleeing or otherwise evasive targets by UAVs has received study and is often relevant simply when no confidence in any particular probabilistic motion model for a target is present. Existing methods are summarized in Section 3.4.1 for regularly-shaped environments and noted to be sparse for irregular environments in Section 3.4.2, while performance improvements for target-constraining environments are proposed in Section 5.5.

2.3 Related Problems

A number of existing problems that have received much previous attention address variations of aspects of the problem considered here. Some of the most related problems and a brief statement of differences include:

- Area coverage

  Generic algorithms to cover an area in the least time, such as producing a minimal carpeting through Boustrophedon or spanning-tree decomposition, may suffice to search for a non-adversarial target. However, a path these generate is intended as the path for the sensor footprint and must be translated to a UAV flight path, which may not be directly possible or require sharp turns a UAV cannot execute. Further, an open-loop coverage pattern must take into account both the finite width of the sensor footprint and the need to incorporate some overlap to address slight trajectory tracking errors due to wind. Several existing area coverage algorithms designed for UAV target search are considered in Section 3.4.1.

- Region clearing

  Generic formulations of clearing a polygon [46] or a discrete state space represented as a graph [75] have been studied as abstract search mechanisms and are discussed further in Section 3.1. Clearing may be useful as a search strategy when an unknown (and possibly zero) number of targets are present in an environment and the goal is to simply detect them all or verify none are actually present. A primary difference with these is that UAVs are not constrained to the same environment as the target—changing the problem rather completely. These formulations assume either a conical flashlight-like or (possibly bounded) omni-directional sensor model for the searcher that reaches outward into the environment containing a target, which is quite different from the sensing model here used for UAVs equipped with cameras or similar sensors defined in Section 2.5. In practice, area clearing is never guaranteed due to uncertainty in sensor pointing and target detection, requiring some accommodation for redundant observations to achieve a desired certainty threshold.

- Pursuit-evasion

  In general pursuit-evasion scenarios, also discussed further in Section 3.1, one or more pursuers attempts to reach or enclose a fleeing target, possibly well-modeling the case of a recapture search for a known target. A key difference from common formulations is that the position of a UAV's target is known with only potentially large uncertainty until viewed and that it is unlikely that a real target may be truly adversarial in its knowledge of pursuer position and intent. Further, pursuit-evasion typically concludes at capture and would not necessarily place pursuers in positions conducive to persistent tracking.

- GMTI tracking

Ground moving target indicators (GMTI) are extremely critical but little-publicized components of real-world large-scale targeting operations. Typically built on aerial Doppler radar sensors, they form the tracking portion of so-called look-down/shoot-down systems that modern defense forces worldwide have come to rely upon, with common implementations including the highly recognizable airborne early warning and control (AWACS) aircraft and radar nosecones of fighter jets shown in Figure 2.3. These systems typically operate at high-altitude (15km), can observe a wide-area of up to hundreds of square kilometers, and can report relative target locations with error in the tens of meters. The task of tracking with such sensors is largely one of target discrimination and data association, and the sheer scale, cost, and launch complexity is otherwise simply well beyond that considered for field-reconnaissance. General approaches to filtering geolocation observations using such sensors, however, remain largely applicable and are described further in Section 3.3.2.

(a) Boeing E-3 Sentry AWACS (b) MiG-31 Nosecone Radar

Figure 2.3: Examples of systems including GMTI sensing.

2.4 Effect of Environment on Problem Approach

Thus far, the problem has been stated independently of the environment in which the task is to be performed. Indeed, for a given point in the decomposed problem space presented, the basic approach strategy–be it covering an entire area by any efficient means to locate a stationary target or calculated sweeping to contain an adversarially moving target in a progressively shrinking region–remains the same regardless of the environment. However, the form and structure of an environment may critically affect both the difficulty of the task and the appropriateness of a given approach.

Possibly the two most important properties of an environment in this regard are its *sparsity* and *connectivity*, measuring the fraction taken up of the bounding convex hull in the continuous space in which it is embedded and the average number of locations to which a target may directly move in a short period of time from any given point, respectively. Although not particularly useful to quantify directly given that a numeric value of connectivity would be uncountably infinite for continuous environments, it is valuable to consider a spectrum of environment classes and their influence on choice of approach. General classes considered here are open and regularly shaped areas, target-obstacle ridden or irregularly shaped areas, collections of intersecting paths (such as road networks), and single paths. These are summarized in Table 2.2 with appropriate choices of representation of target location likelihood, which is most immediately affected by both factors given a need to define (at least implicitly) a possible target location set and a neighboring location set for each valid position. An accompanying intuitive graphical depiction of the class spectrum is shown in Figure 2.4.

	Open & regularly shaped	**Many obstacles or irregularly shaped**	**Intersecting paths**	**Single path**
Sparsity	Dense	Moderate	Very sparse	Arbitrary
Possible target motions	Continuous	Continuous	Discrete few	2 directions
Example likelihood store	• Continuous 2-D distribution • Dense grid • Unconstrained particles	• Densely connected cells • Constrained particles	• Multiple continuous 1-D distributions • Cells or particles lying on a graph	• Continuous 1-D distribution • Cell array • 1-D particles

Table 2.2: Characteristics of and appropriate target location uncertainty representations for general classes along the combined sparsity-connectivity spectrum.

The significant practical impact of these properties appears when evaluating the quality of placing sensor footprints at given future locations. This has two components. The first is, roughly speaking, the number of possible target locations that must be considered to determine whether each will be in view and the quality of observation that would be received if seen there, which depends on properties such

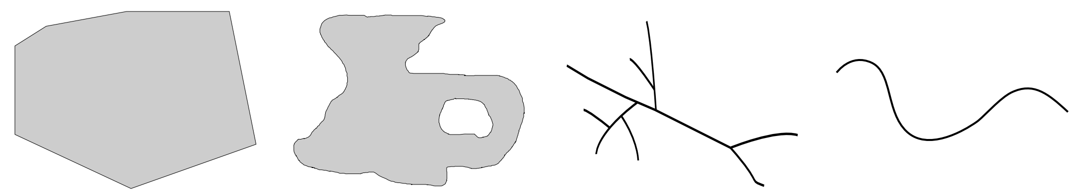

Figure 2.4: Intuitive graphical depictions of significant general classes along the combined sparsity-connectivity spectrum. Shaded regions or solid line segments indicate possible target locations.

as the dimensionality of this location space, the number of samples needed to cover it (if using a discretized representation), and whether these evaluations may be coalesced or somehow collectively approximated (for instance by computing the result of a single cumulative probability distribution function). The second component is the number of places the target can move in an instant, which arises when computing predicted target motion or area recontamination. This too is a similar measure of the size of a space, which in this case is the space of directions it may move or possible short-term destinations. Intuitively, smaller is better, suggesting wise use of efficient representations when sparsity and low-connectivity permit. This arises in formal terms in Section 3.3.3 and is the motivation of the proposed use of structure that enables low-dimensional representations for sparse, low-connectivity environments.

Finally, a number of other properties of the environment may affect the choice of approach, though not as severely, and are not a focus of this work. Beyond strict connectivity or binary location validity, relative traversability of regions of an area by a target, measured perhaps both as desirability and mobility of traversal, affecting the likelihood of a target moving into an area or a model of its motion within it. This is considered to some extent by existing methods for probabilistic search reviewed in Section 3.4.2 and is proposed as a means to improve performance in road networks in Section 5.3. Presence of occlusions to target observation such as tall buildings or trees might suggest the use of a higher-dimensional environment map representing directional visibility and is presently considered as an extension. Boundedness of the environment is another, as open areas from which the target might move infinitely far require additional reasoning for containment or require treatment as search and tracking in a progressively enlarging area. Finally, environment scale can have a substantial impact on the approach strategy. In the extremes, a very small area on the order of the scale of the sensor footprint can be trivially observed nearly continuously with little effort, while an extremely large environment may preclude the use of long-term trajectory optimization and permit only very locally optimal solutions. One major goal of the proposed exploitation of structure is to represent sparsely-connected subspaces of an otherwise large environment to avoid this very problem, but the scale spectrum and appropriate thresholds for choosing certain methods over others is not otherwise considered in this thesis.

2.5 Observer Modeling

UAVs differ from traditional ground-based or abstractly planar vehicles in a number of important ways. These are now enumerated, their models stated specifically for use in the remainder of this document, and their implications considered.

Sensing constraints

Whereas common models for a planar vehicle's sensor are generally of the form of a ray emitter with bounded range or bearing (if not omnidirectional), a UAV's field of view will generally form a sensor footprint on the ground that is a finite spot offset from the vehicle that is the result of intersecting similar ray emissions with the ground. The position and form of this spot may further be maneuver-dependent as UAV orientation changes, unless constant level flight is assumed for simplicity, and moving the spot requires potentially complex vehicle motions. This is further formalized in the following subsections and Section 3.3.1.

Additionally, given the long range to targets (50 to several hundred meters), limited onboard computation to process sensor readings, and the lossy nature of wireless transmission links, received raw sensor data will inevitably contain noise or be missing portions that cause likelihood of target detection given an observation substantially lower than for a ground vehicle that can get possibly arbitrarily close to a target, have large amounts of onboard computation, and need not wirelessly transmit sensor data for processing.

Motion constraints

All UAVs must maintain lift and are generally designed to operate within a region of their state space (appropriately called a flight envelope) within which an onboard autopilot is programmed to maintain this lift. For fixed-wing aircraft, one resulting constraint is that of a minimum speed that must be maintained, though with the positive effect that average speeds for fix-wing aircraft are much higher than those of a typical autonomous ground vehicle. For most aircraft, the flight envelope dictates a maximum turn rate corresponding to a minimum turning radius. Typically, there also exist dynamic constraints that couple acceleration or turn rate to vehicle attitude, including for instance the common need to bank (roll) to turn. This too is further formalized in a later subsection.

Terrain obstacles

While ground vehicles are generally incredibly vulnerable to obstacles–unable to see over those on the order of the vehicle height and unable to traverse obstacles above a small fraction of the vehicle height–UAVs, by virtue of their altitude, have in practice minor vulnerability to them. UAVs may pass over obstacles altitude-permitting, but it is much more critical that collisions be avoided than for ground vehicles, both because lightweight vehicle design implies great fragility and because the dynamic disturbance of a collision will result in a fatal excursion from the flight envelope. Likewise, again by virtue of altitude, UAVs can see over obstacles as the geometry allows (taller obstacles casting a longer shadow) but not beneath them as a ground vehicle might otherwise be able.

Vehicle design

Field-reconnaissance UAVs considered here are typically limited by fuel or battery capacity to relatively short flight durations of one half to several hours, whereas a ground vehicle need not continuously oppose gravity to maintain this weight at altitude. This has led to models of dynamically sized teams or explicit refueling subtasks to account for the limited longevity of a single UAV, but this aspect is not considered here. More importantly, as previously introduced, weight, cost, and size limitations that are much more stringent than those applicable to ground vehicles, internal state sensing or proprioception is substantially worse than that available on ground vehicles, with resulting detriment to observations relying on this state. Finally, such UAVs cannot meaningfully interact with targets given little available payload and thrust, so missions by necessity must be indefinite, as a target may simply continue to move once initially detected.

Taken together, UAVs possess substantial differences from ground vehicles, but it may be noted that these produce largely complementary properties. This motivates the possible use of both types for demanding missions, an idea that was the very focus of initial experimentation described in Section 4.1 providing motivation for practical representational improvements. Much room for further exploration in this area remains, however.

UAV Sensor Model

A great variety of external sensors have been either used or suggested for use aboard UAVs. Many, including laser scanners and active radar arrays are beyond the payload capacity of field-reconnaissance UAVs, but many more including especially passive varieties remain feasible. Visible-spectrum cameras (so-called electro-optical sensors) are a staple given that the use of autonomy on such UAVs is in its infancy, and human operators are still nearly always present. Passive infrared cameras, direction-finding radios (optionally with a wide-area transmitter component to detect emergency rescue reflectors), and detectors with specialized spectral sensitivity (e.g. radiological emission sensing) have likewise seen use or proposal. All of these have the property that they can be modeled as reporting a detection as a direction vector or spherical coordinate in vehicle-local coordinates, or as more commonly termed, a bearing-only sensor.

Here, visible-spectrum cameras (or any sensors having a sensing plane and an analogy for focal length) are assumed, but any sensor providing such a direction vector is equally relevant. A general observation at time t is defined as

$$\mathbf{z}_t = \begin{cases} \emptyset, & \text{no target detected} \\ \mathbf{x}_{\text{tgt}}^{\text{img}}, & \text{target detected} \end{cases}, \tag{2.1}$$

where $\mathbf{x}_{\text{tgt}}^{\text{img}}$ is the image pixel coordinate of (the center of) a detected target. According to the standard pinhole camera model, [108] this 2D image projection of a target is related to the target's 3D location in the reference frame of the camera via the 3x3 camera calibration matrix $\mathbf{M}$ by

$$\begin{bmatrix} \mathbf{x}_{\text{tgt}}^{\text{img}} \\ 1 \end{bmatrix} \equiv \mathbf{M}\mathbf{x}_{\text{tgt}}^{\text{cam}} \tag{2.2}$$

when expressed in homogeneous coordinates.

Inverting this relationship provides a ray parallel to the vector $\hat{\mathbf{v}}_{\text{tgt}}^{\text{cam}}$ in the camera reference frame given by

$$\hat{\mathbf{v}}_{\text{tgt}}^{\text{cam}} \equiv \mathbf{M}^{-1} \begin{bmatrix} \mathbf{x}_{\text{tgt}}^{\text{img}} \\ 1 \end{bmatrix}, \tag{2.3}$$

again expressed in homogeneous coordinates.

Finally, if $\mathbf{R}_{\text{UAV}}^{\text{world}}$ and $\mathbf{R}_{\text{cam}}^{\text{UAV}}$ are rotation matrices taking vectors from UAV to world and camera to UAV reference frames, respectively, then the target must lie along the ray in space given in world coordinates by

$$\hat{\mathbf{v}}_{\text{tgt}}^{\text{world}} = \mathbf{R}_{\text{UAV}}^{\text{world}} \mathbf{R}_{\text{cam}}^{\text{UAV}} \hat{\mathbf{v}}_{\text{tgt}}^{\text{cam}}. \tag{2.4}$$

Methods for transforming this ray into an estimate of target location–for instance by intersecting rays from multiple detections or intersecting a ray with a known terrain model on which a target must lie–are considered in Section 3.3.1.

UAV Motion Model

For simplicity, details of aircraft dynamics are not considered here. It is assumed, instead, that there exists a function computing a forward model of each UAV's dynamics so that if $\mathbf{x}_{\text{UAV}}(t)$ represents the complete state of the UAV at time t, then the predicted future state $\mathbf{x}_{\text{UAV}}(t')$ may be computed as

$$\mathbf{x}_{\text{UAV}}(t') = f_{\text{motion}}(\mathbf{x}_{\text{UAV}}(t), \mathbf{u}, t' - t), \tag{2.5}$$

where $\mathbf{u}$ represents a control input vector (such as a turn or thrust command) to be held for the duration $t' - t$.

Where it is necessary to consider an explicit model and for algorithmic comparison, a simple form of a discrete time kinematic model of a so-called Dubins car [30] is used as f_{motion}. The equations of motion are given by

$$\begin{aligned} x_{k+1} &= x_k + V \Delta t \cos \theta_k \\ y_{k+1} &= y_k + V \Delta t \sin \theta_k \\ \psi_{k+1} &= \psi_k + \omega_k \Delta t \\ &|\omega_k| \leq \omega_{max}, \end{aligned} \tag{2.6}$$

where $\mathbf{x}_{\text{UAV}} = [x_k, y_k, \psi_k]^T$ is the position and heading of a UAV in its motion plane (assuming constant altitude). A constant (or slowly varying) forward speed V is assumed, and the single control input is a steering or turn rate ω_k having bounded magnitude ω_{max} that is held for the entire duration of a timestep Δt. This is a common model in the literature on UAV motion planning that reasonably approximates fixed-wing aircraft and helicopters flying at fixed altitude.

As an additional realism improvement, a rough model of the need to bank to perform a turn is incorporated by coupling turning rate to roll angle so that turning at a given rate produces a roll angle resulting in a so-called coordinated turn, defined by

$$\phi = \text{atan}\left(\frac{V\omega}{g}\right), \tag{2.7}$$

where g is the gravitational constant.

Implications for Ground Vehicle Strategies

These distinctions from ground-based or abstract planar vehicles imply a need to treat UAVs somewhat differently, preventing the simple invocation of a search or tracking algorithm not designed explicitly for them for several reasons. First, the differing motion constraints from most ground vehicles renders infeasible trajectories that may call for reverse motion, rapid turns, or (for fixed-wing aircraft) low speeds. Further, trivial mappings of abstract trajectories to feasible UAV motions–such as substituting a slow orbital turn in place of a zero-radius turn–may result in differing-cost paths under a given objective function than expected and will produce sensor footprint trajectories with excursions from the intended path, undermining for instance a continuous sweep. Additionally, an algorithm assuming that an observer must lie in the same plane as a target will fail to take advantage of the ability of UAVs to fly above and see beyond obstacles, therefore providing paths that are necessarily sub-optimal. Lastly, the impracticality of reliably conveying large amounts of data over UAVs' lossy long-range links and hence need for efficient or distributed representations is easy overlooked in abstractions, though neither is this issue a focus of this thesis.

Most existing work in UAV search and tracking, summarized in the following chapter, has therefore focused on specialized or novel control strategies, approaching from the problem from an aerospace or estimation rather than abstract search or pursuit-evasion perspective. This thesis proposes an instance of that alternative perspective to attain superior performance where environment structure permits.

CHAPTER 3

Related Strategies for Search and Tracking

As a robust area of study, the disparate areas of UAV search and tracking have accumulated a large body of literature independently considering techniques for geolocation, control for pursuit, and search. This chapter considers this work in the context of the general CSAT task and presents a summary of existing approaches to each component. Section 3.1 first enumerates common abstractions of which the problem might be considered an instance and thereby have known properties or possibly available (albeit not necessarily immediately applicable) solutions. Related aspects of the problem that must be addressed in any practical implementation but not considered further here such as data association and aerodynamics are then briefly summarized in Section 3.2. Returning to relevant issues, approaches to producing geolocation observations, filtering these to best produce combined estimates in the face of potentially large error, and then using these estimates or raw observations to pursue targets with heuristic or optimized trajectories are then summarized in Section 3.3. These and more case-specific trajectory generation schemes for reaching and observing targets whose position is not well known–the search case–is next presented within a taxonomy of scenarios in Section 3.4. Finally, independent study of the recapture case for targets having large but clustered location uncertainty is advocated in Section 3.5, with several intuitive and potentially applicable existing strategies considered.

3.1 Common Abstractions for Search and Pursuit

Search and pursuit has been formally studied for many decades in a number of contexts, from the completely abstract to those with immediate mobile robotics or field applications. Traditionally, the problems of *search* and *pursuit-evasion*[1] have been treated quite independently. Search is commonly defined as the case in which a team of searchers wishes to clear an area of any targets that may be present, without necessarily knowing whether one is even present (thus also earning the name "one-sided search"). As a result, solutions often take the form of an offline precomputation. Pursuit-evasion typically takes the form of attempting to reach, encircle, or maintain coverage of a known target that may be trying to avoid capture and whose position is usually assumed to be known. As a pursuer's motion depends on a target's choice of motions, reactive online computations are generally required.

Separately but often tightly coupled to a given problem formulation, two types of targets are considered: adversarial and probabilistic. Adversarial targets actively attempt to evade detection or capture and may be granted infinite speed or complete omniscience of pursuer location, strategy, and intent. A strategy that is guaranteed to detect or capture an infinitely powerful target offers the additional desirable properties that no other motion model of the target is needed and that the temporal dimension is effectively removed (eliminating concern for physical scale of the environment or distances between points), however it may be extremely conservative both in resource requirements (such as number of agents) and the number of action steps or amount of time required to guarantee success. Probabilistic targets, on the other hand, are assumed to move with some stochastic model usually independent of searcher or pursuer motion, corresponding to a target that is oblivious of or indifferent to the presence of its pursuers. Strategies for detecting or capturing probabilistic targets, then, seek to perform optimally in expectation with respect to some metric such as minimum expected time to detection or capture, maximum probability of either within a fixed time interval, or maximum expected number of observations within an interval.

The search problem may be well divided between probabilistic and adversarial targets. Search for probabilistic targets formed the basis of the so-called theory of search in the operations research community, which has a long history dating back to naval warfare strategies of World War II and later formalized in the seminal works of Koopman. [64] Roughly, such a search seeks to maximize the probability of target

[1]The terms *search* and *pursuit-evasion* are used nearly interchangeably in some literature or with varying definitions between fields, to much confusion. Here, an attempt is made to simply provide a common, reasonable, and internally consistent definition.

detection given a prior distribution of possible target location using a fixed amount of available search effort and target detection likelihood per effort, producing a best search density schedule indicating where to apply given amounts of effort. This has over time grown to include more advanced models of target detection, variations in target motion, and more computationally complex solutions given advances in computing. Well-regarded surveys include those by Stone [99] and later by Benkoski. [12]

Search for adversarial targets, meanwhile, has been considered both for bounded continuous spaces and discrete environments whose connectivity between regions is represented as a graph. Adversarial search in continuous polygonal environments was most famously studied by Guibas et al., [46] providing an algorithm to clear a planar polygon possibly containing an arbitrarily fast target using a minimal team of searchers having omnidirectional unlimited-range sensors that cannot see through walls and showing the NP-hardness of the problem. Among other extensions, this has been generalized to searchers having limited-angle field of view [40] and bounded-speed targets [107] for improved realism and practicality in mobile robotics. For environments better represented as a graph of discrete locations, the game first defined by Parsons [75] and independently by Petrov [77] is that of locating a missing explorer in tunnels of a cave system by a minimal team of searchers, regardless of the motions of the explorer. This is the classical formulation of the problem known as *graph search* that typically assumes a completely omniscient adversary that is equivalently treated as a poisonous gas that must be corralled. The minimal searcher team size required for a given graph is called its *search number*, the determination of which (and therefore an optimal search schedule as well) has been shown to be NP-complete. [66] Though many problem variations exist, the class dichotomy considered here is that of a target that may lie only on an edge (requiring *edge search*)–the original and more commonly studied version–versus that of a target occupying only nodes (requiring *node search*). The case of node search is generally more applicable to practical robotics problems and has received attention in this context from Hollinger and Kehagias et al. [51, 59] for indoor search and from Kolling and Carpin [63] in a hybrid framework called GRAPH-CLEAR in which nodes represent areas to be cleared and edges are weighted by the number of searchers required to guard them. Edge search, however, may be more applicable to outdoor scenarios and forms the basis of several ideas in Chapter 5. Finally, traditional formulations of graph search typically permit searchers to instantaneously move between any two points in the graph (producing so-called teleporting or jump search), which is quite unrealistic for any physical application and is eliminated by constraining a search schedule to be an *internal search* in which searchers are themselves constrained to the graph topology as well. A stronger notion introduced by Barrière is that of *connected search* in which searchers must lie within a growing connected cleared subgraph, conceptually maintaining a safe region for supply lines from a start point. [10] A thorough survey of existing work in canonical forms of guaranteed graph search was performed by Fomin and Thilikos. [34]

Pursuit-evasion has likewise been considered in several separate continuous and discrete-space formulations. In continuous space, a pursuit-evasion problem is often treated as an instance of one of several simple games. One is the so-called *lion-and-man* game in which a lion (pursuer) attempts to reach the location of an evader having the same maximum speed within a closed geometric environment such as a circle or polygon. More recent research has considered unbounded or higher-dimensional areas and sensing limitations, and many results have over time been shown for variations of the problem. Several interesting ones include the facts that a single lion can eventually capture an evader in any simply-connected polygon, [55] that three can do so in any polygon, [13] and that $d+1$ lions can capture an evader in the open space $\mathbb{R}^d$ if and only if its initial location lies within their convex hull. [65] Somewhat of a generalization of the lion-and-man game is the so-called *princess-and-monster* game in which (in its original formulation [54]) a monster (pursuer) attempts to reach an evader of arbitrary speed in a closed environment in which neither player can see the other, providing extension to local-only sensing and fast targets. This is an instance of a more general class of differential games, of which the *homicidal chauffeur* problem [53] is more famous. In that problem, a pedestrian (evader) that can move arbitrarily at slow speed seeks to avoid an automobile (pursuer) that is much faster but is less maneuverable due to motion constraints (such as automobile or aircraft dynamics). An excellent summary of its history and known properties was compiled by Patsko and Turova. [76]

In discrete space, pursuit-evasion problems are also frequently represented as graphs. Given further discretization in time such that pursuer and evader move in alternating discrete rounds, the problem is now a discrete game commonly called *cops and robbers*. In typical formulations, pursuers attempt to reach the location of (the node containing) the evader, and both players receive total information as to the location of the other. In this game, the analogue to the search number is the *cop number*, denoting the number of pursuers required to guarantee capture in finite time for a given graph. The set of graphs requiring only one pursuer was characterized by Quilliot, [81] but it is EXPTIME-complete to determine whether k cops are sufficient for a given graph and initial conditions. [44] Full knowledge of target location may be realistic if its position is provided by a sensor network or by an aerial observer directing ground agents, but otherwise is impractical for most robotics applications. Adler et al. [2] consider the variant termed *hunter-and-rabbit* in which neither player knows the other's location or motion decisions (which is thus essentially the discrete form of the princess-and-monster game), for which only randomized pursuit strategies are shown to be useful and have tight asymptotic bounds on expected capture time.

For each of these problem formulations, many variations exist, relaxing or tightening assumptions on or capabilities of both the pursuers and evaders. Small changes of any form often change the problem structure substantially, and known results vary considerably. Though tantalizing similarities to practical physical tasks are present, remarkably little use of formal search and pursuit-evasion theory has been made in applied robotics literature. Chung et al. [21] present a review of several fundamental results and recent applications in mobile robotics. The immediate applicability of any given abstraction to the aerial search and tracking task is particularly unclear, especially because most abstractions model pursuers as restricted to the same space as the target and (for visibility-based search) as having a point sensor aimed outwardly within the planar environment. Two existing examples demonstrating some applicability include the use of discrete probabilistic search for a randomly moving evader by an air-ground team providing guaranteed finite capture time [112] and the use of continuous differential game theory to provide multi-UAV interception of a target whose current position but not future motions are known. [86] The prior provides no guarantees for adversarial targets or performance bounds for purposefully-moving targets, and the latter relies on the unrealistic assumption of precisely known target location while further lacking physical implementation. Several other instances of probabilistic search for non-adversarial targets without reference to formal search or pursuit-evasion and providing only locally optimal behavior are described in succeeding sections. A major goal of proposed work in Chapter 5 is the development of mappings to graph-based abstractions permitting formal reasoning for all three cases of CSAT.

3.2 Related Aspects not Considered

Many aspects of a practical implementation of or topics relevant to the search and tracking problem have been well studied previously and are neglected in this thesis, as much of the research into the design, control, and applications of UAVs comprised the bulk of original forays in the area as robotics or aerospace problems. Here, they are instead assumed to have been addressed as needed or remain fundamental challenges. Several of the most prominent include:

Self-localization and environment mapping

Accuracy of available vehicle state may be quite poor given the low-end sensors onboard field-reconnaissance UAVs, and there exists a sizeable body of work on attempting to reduce this error by augmenting these sensors with visual localization using an onboard camera likely already present for surveillance. Full simultaneous localization and mapping (SLAM) systems using Extended Kalman Filters (EKFs) were exhibited by Bryson and Sukkarieh [16] and Caballero et al. [17], in which the prior formulate their filter with bearing-only observations of tracked landmarks and the latter with 3D observations formed by projecting tracked landmarks and intersecting with a known ground plane. Madison et al. [69] demonstrate tracking multiple ground features over several tens of seconds to constrain the standard Structure-from-Motion 8-point algorithm, which solves for all the feature locations and vehicle poses. As an alternative to explicit landmark tracking, Andersen and Taylor [5] show that with a planar ground assumption, a homography-based visual odometry algorithm can be combined with GPS/INS sensors using an Unscented Kalman Filter (UKF) to improve the vehicle pose estimate. A major downside to most of these strategies is a need for stable visual tracking of landmarks over hundreds of video frames, which is difficult given low-resolution video of potentially featureless terrain from a fast-moving, wind-buffeted platform. Matching of live video to existing satellite imagery has been proposed by Conte and Doherty [24] and Sim et al. [95], but at significant computational cost and vulnerability to changes in terrain appearance. As any such strategy may leave substantial residual error, target geolocation with uncertain vehicle state is a related topic discussed in Section 3.3.2 and improved upon in Section 4.3.

Automatic target recognition or data association

Detection of a target, tracking it efficiently while in view, and distinguishing it from any other observable targets are themselves challenging problems, particularly using noisy sensor data. When using a camera, this is largely a computer vision problem, and has received much attention given its applicability to a huge range of applications. Automatic detection from a stored model or target classifier, automatic detection of moving targets, and operator-initiated or manual tracking are all means that have been previously proposed. Section 4.1 discusses several and describes the operator-initiated automatic visual tracking system used to motivate and evaluate proposed algorithmic improvements in field experiments. It is otherwise assumed that by some means, observations are received when a target is detected or tracked, and a negative observation received to indicate the lack of such. Section 2.5 defines this more precisely.

Aerodynamics and low-level vehicle control

The development and evaluation of small, low-cost control systems for SUAVs and MAVs has been a popular topic of study given the many constraints the design of such UAVs impose. As this is a complex and largely well-understood problem of its own, with many such UAVs now incorporating sophisticated controllers, the realistic assumption is made that an autopilot is present that provides high-level action primitives such as waypoints, orbit or hover points, or steering-like commands and which further deduces and cancels out effects of wind to the extent possible. This too is defined more precisely in Section 2.5.

Sensor aim or gimbal state planning

An onboard target-observing sensor such as a camera may have a means to alter its relative pointing orientation so that the sensor footprint on the ground may move without requiring motion of the UAV, for instance by mounting on an actuated gimbal. Though many field-reconnaissance UAVs may be equipped with a gimbal, explicit control of the gimbal is largely ignored for a number of reasons: it introduces additional complexities such as positioning limits and dynamical constraints of its own, an attempt is made to be applicable to a wider range of vehicles including those without one, an emphasis on required vehicle motions is desired instead, and because in many cases of flight planning it is possible to merely augment the vehicle control input with a state for gimbal. Therefore, either a fixed sensor or an automatic gimbal centered on the best estimate of target location is assumed. The prior is assumed unless otherwise indicated.

Obstacle avoidance and agent deconfliction

Though UAVs may fly over obstacles given sufficient altitude, it is typically critical that they avoid them, as collisions are unlikely to be as survivable as they may be for ground vehicles given greater airframe fragility and the need to maintain lift. An interesting problem that has received some study is the safe navigation of urban or mountainous terrain using low-altitude UAVs, however for simplicity the assumption will be made that any terrain obstacles will lie strictly below a fixed flight altitude. To some extent, this reduces the problem to a two-dimensional problem, though constraints such as the need to bank to turn and the altitude-dependent geometry of observations are maintained, covered further in Sections 2.5 and 3.3.1. Further, to simplify planning, it is assumed that even coincident flight paths will never result in collision, which may be enforced by the common practice of specifying a different altitude for each aircraft with a safe buffer distance between each.

Distributed or decentralized estimation and control

Given possibly large communication distances, limited bandwidth, feeble onboard computation and storage, and the nontrivial likelihood of losing contact with any one vehicle, it is desirable that practical implementations use some means to minimize the amount of data transferred and avoid single points of failure. This has been addressed by a large body of work on distributed estimation and decentralized control with much success. While such issues are not a focus of this work, concise representations of possible target location are sought that require little data to be transmitted or which are easily factorized for agent assignment *to enable* such strategies to be more easily applied. This also serves the vital purpose of greatly improving computational tractability, needed both in physical implementation and the example results considered. Levels of coordination and decentralization are touched on further in Section 3.3.3.

3.3 Tracking and Pursuit with UAVs

Given a current or recent detection of a target, the task of pursuing it is comprised of two components: processing the observations to produce the best geolocation estimate possible and choosing a near-term trajectory that aims to maintain or improve this estimate, typically implying the further need to keep the target in view. Here, this is split into a pipeline of three stages consisting of generating geolocation observations from target detections, filtering these observations and predicting target motion in their absence to produce a best target location estimate, and choosing flight paths most likely to maintain a useful stream of future observations or paths that specifically optimize expected observations.

3.3.1 General approaches to geolocation

Once a target is detected using an onboard bearing-only sensor such as a camera modeled in Section 2.5, two broad approaches for mapping a pixel coordinate to a ground location are found in the literature. One is to register live UAV video frames to previously geo-referenced satellite or aerial imagery of the area, [25] which immediately provides the world location for any pixel of the image without any dependence on potentially-erroneous UAV pose reported by telemetry. However, this approach requires such prior imagery, accurate image registration may not be possible if the terrain is highly repetitive or homogeneous, registration may fail entirely if significant changes have occurred in the area since the prior imagery was acquired (as could likely be the case with smoke and rubble when responding to a disaster or conflict), and the computational and memory burden necessary to store and manipulate large maps may be infeasible. More common, therefore, is a second approach that uses an estimate of the UAV's pose derived from on-board sensors and the identified pixel coordinate to generate a ray in world coordinates as in Equation 2.4 that is then intersected with previous observations or a prior model of the terrain to produce an observation of target location. Using this approach, a target's position in world coordinates $x_{\text{tgt}}^{\text{world}}$ must satisfy

$$\mathbf{x}_{\text{tgt}}^{\text{world}} = \mathbf{T}_{\text{UAV}}^{\text{world}} \mathbf{p}_{\text{cam}}^{\text{UAV}} + d\hat{\mathbf{v}}_{\text{tgt}}^{\text{world}}, d \in \mathbb{R}, \tag{3.1}$$

where, corresponding to Equation 2.4, $\mathbf{T}_{\text{UAV}}^{\text{world}}$ is the pose transformation taking coordinates from UAV to world frames, $\mathbf{p}_{\text{cam}}^{\text{UAV}}$ is the calibrated location of the camera with respect to the vehicle body, and d is the distance to the target along the ray $\hat{\mathbf{v}}_{\text{tgt}}^{\text{world}}$ emanating from the camera's center.

Examples intersecting rays from multiple observations to produce geolocation estimates include those by Madison et al. [69] and Quigley et al., [80] in which a best near-intersection (as errors in pixel localization and vehicle state will prevent exact intersection) of all rays is taken to be the target's location. This has the advantage that no dependence on any prior terrain imagery or model is required (and, indeed, targets need not be restricted to lying on the ground), but many observations may be needed before a meaningful target location is available, and applicability to moving targets is not completely straightforward. In contrast, the ray corresponding to a single target detection may be intersected with a prior model of terrain in the form of a 3D height-map or simpler Digital Terrain Elevation Database (DTED) available for much of the planet from satellite surveys. [42] As with matching to geo-referenced imagery, this provides an estimate of target location from a single detection but can operate with rapidly-changing or partially-obscured terrain, at the cost of a very strong dependence on accurate UAV state. For the case of locally flat terrain, the distance d to a target needed to compute its location using Equation 3.1 is simply

$$d = \frac{z_{\text{terrain}} - z_{\mathbf{p}_{\text{cam}}^{\text{world}}}}{z_{\hat{\mathbf{v}}_{\text{tgt}}^{\text{world}}}}, \tag{3.2}$$

where z_{terrain} is the local terrain height. This may be extended to account for jagged terrain and man-made obstacles through proper geometric intersection with a model.

Whatever the means used to produce an individual target observation, it is convenient to define a function in its place

$$\mathbf{x}_{\text{tgt}}^{\text{world}} = f_{\text{obs}}(\mathbf{s}_{\text{UAV}}, \mathbf{x}_{\text{tgt}}^{\text{img}}), \tag{3.3}$$

where $\mathbf{x}_{\text{tgt}}^{\text{img}}$ is the pixel coordinate of the center of the visible detected target and $\mathbf{s}_{\text{UAV}}$ represents a state vector containing all UAV state such as position, orientation, camera calibration parameters, etc. that may be necessary to produce the individual geolocation observation. This permits abstraction of any underlying details. In the remainder of this document, the output of f_{obs} is referred to as a *pseudo-observation* since observations are actually received as image pixel locations.

Different axes of target state contribute differently to error in geolocation observations, two of which are depicted in Figure 3.1. Under the ray intersection model, error in UAV (usually GPS-provided) or the camera's relative position contribute additively to geolocation error and typically produce errors at the level of up to several meters, while UAV and relative camera orientation has a trigonometric impact that may be very large as distances to targets may be on the order of several hundred meters. Worse, as IMUs onboard field reconnaissance UAVs were previously noted to have relatively low-bandwidth and low-accuracy, these are the most poorly sensed of UAV state dimensions. Consider the following example. Suppose a UAV lies 100m above and parallel to flat horizontal terrain and has a camera pitched downward by 30° centered on a target. The range to the target will then be 200m and the distance along the ground approximately 173m. If uncertainty in vehicle heading is modeled as a Gaussian distribution with $\sigma = 5°$, then with a rough heading error bound of $3\sigma = 15°$, uncertainty in resulting observations is bounded roughly by $173 \sin(15°) \approx 27$m, possibly encompassing quite a large area. For this reason, substantial effort in UAV geolocation literature has focused on filtering single observations to produce more accurate target location estimates.

3.3.2 Common geolocation filtering strategies

For stationary targets, direct averaging of observations produced by f_{obs} is an obvious first means of filtering, as it eliminates zero-mean noise in observations and provides some robustness to small fixed biases in UAV state estimates (such as in heading) that produce symmetric error that can be eliminated by averaging, for instance, over observations taken around an orbit of a target. Few error sources are symmetric in this way, let alone fixed, however. Redding et al. [84] perform essentially this,

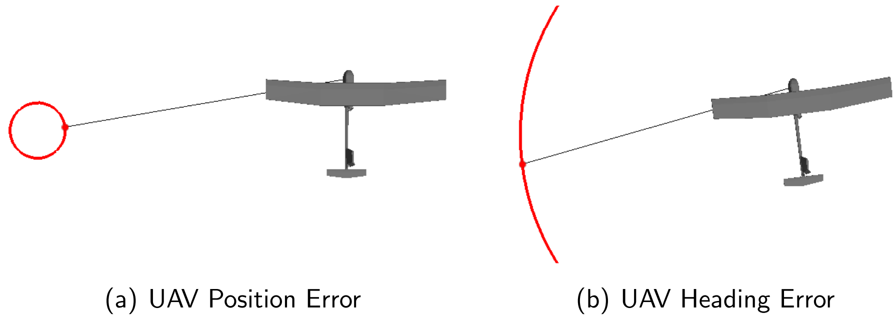

Figure 3.1: Graphical depictions of the impact of error in differing UAV state dimensions on an individual target geolocation observation using ray intersection with terrain. Here, a UAV using a side-mounted camera is geolocating a target on level terrain. In (a), if error in UAV position is introduced so as to trace a circle in the plane, an identically-scaled circle appears in geolocation estimates, indicating additive effect on error. If, however, error in UAV heading is introduced as in (b), geolocation error is greatly amplified and scales linearly with distance to the target, tracing out an arc and indicating large non-linear impact.

using linearization of the image projection function and recursive least-squares to efficiently compute a best target location while treating image projection as an uninvertible operation (conceptually permitting geolocation without terrain knowledge). Barber et al. [9] improve upon this by wrapping least-squares filtering with online optimization to solve for fixed biases in sensor yaw and pitch that minimize variance in (or equivalently maximize clustering of) individual observations. Overcoming large or varying offsets, however, remains challenging and was addressed in work to date detailed in Section 4.3 using an evidence-grid filter that greatly improves estimates in the presence of large, time-varying error.

Kalman filters provide an alternative strategy for geolocation filtering that permits explicit probabilistic modeling of uncertainty in particular UAV state dimensions and target motion. For these reasons and their instrinsic simplicity, they have received frequent mention (usually as EKFs) in the geolocation literature. [56, 79, 87, 114] Formulations typically belong to one of two varieties. One form treats observations as (correctly) lying in the space of pixel coordinates and explicitly includes the projection of target world to image location in the filter's observation function, while the other operates completely in world coordinates and receives pseudo-observations produced by f_{obs}. The latter is preferred here where possible as it permits the use of very simple filter models (typically, using the identity observation function) and encapsulates all aspects of observation generation within a single module. Either form requires an observation error propagation model to transform uncertainty in UAV state to uncertainty in observations, and therein lies the weakness of these filters. For the second model using pseudo-observations, standard Jacobian linearization of f_{obs} may be used to approximately convert Gaussian uncertainty in UAV state $\mathbf{s}_{\text{UAV}}$ having covariance $\mathbf{\Sigma}_{\text{UAV}}$ to Gaussian uncertainty in observed target location having covariance

$$\mathbf{\Sigma}_{\text{obs}} = \mathbf{J}_{f_{\text{obs}}} \mathbf{\Sigma}_{\text{UAV}} \mathbf{J}^T_{f_{\text{obs}}}, \tag{3.4}$$

where

$$\mathbf{J} = \frac{\partial f_{\text{obs}}}{\partial \mathbf{s}_{\text{UAV}}}. \tag{3.5}$$

Meanwhile, EKFs operating in the space of pixel coordinates implicitly perform identical linearization on the opposite, forward-projection, function. However, because the function representing either direction is highly nonlinear in a number of its parameters, particularly those pertaining to UAV orientation, Jacobian linearization may present a very poor approximation. For this reason, variations of UKFs or sigma-point Kalman filters have also been applied, with improved results. [18, 56] However, several problems remain. First, these too impose Gaussian uncertainty distributions on state dimension (often out of convenience), yet uniform or bounded uncertainty may be more appropriate for large errors or biases that may vary. Worse, covariances are often chosen entirely arbitrarily (and highly optimistically), producing filters that are extremely likely to enter an inconsistent state with overconfident, incorrect estimates. True uncertainty distributions resulting from typical state uncertainties are further explored in Section 4.3 along with filters using uncertainty having parametrically-representable bounded uncertainty that are suggested as one possible means of improvement.

Sampling, Monte-Carlo, or particle filter approaches permit approximate representation of arbitrary posterior target uncertainty distributions, and for this reason have also been considered in aerial target geolocation. [88] Here too, two forms may be considered. The first is the classical particle filter implementation, [6] in which given an observation $\mathbf{z}_t$, each particle $\mathbf{x}_i$ representing a possible target location may be projected to image coordinates using some means such as Jacobian linearization or the Unscented Transform to produce an uncertainty distribution in image coordinates and receives a weighting depending upon $p(\mathbf{z}_t|\mathbf{x}_i)$ under this distribution. Unfortunately, this requires a projection and uncertainty propagation per particle, which may be infeasible on low-end embedded hardware or where hypothetical future observations must be repetitively considered. A slightly differing approximation making use of pseudo-observations instead weights each particle under $p(\mathbf{x}_i|\mathbf{z}_{\text{pseudo}_t})$, requiring only a single projection and error propagation per observation update, greatly improving computational performance. Finally, discretized environments represented as grids may conceptually be treated as collections of particles as well, in which particles do not themselves move during process model execution but rather gain or lose probability, but otherwise receive identical Bayesian observation updates. [105]

Where a target is known to lie in an environment having discrete terrain types with differing characteristics, such as when a target is present on or near a road, constrained or multi-model geolocation provides a potential means to improve accuracy by making use of a most-specific target model offering the most accurate prediction at all times. Extensive use of such filters has been made in geolocation from high-altitude radar as mentioned in Section 2.3 on large vehicles such as those shown in Figure 2.3 to both automatically classify targets as on-road or off-road and improve geolocation accuracy when they are on-road by reducing the set of possible target locations to those lying on the road. Examples include the use of multiple-model Kalman filters [62] and variations of constrained particle filters. [31] Both have shown marked increases in estimate accuracy in this context, and for this reason and their representational compactness that they are proposed for use in road-constrained pursuit by field-reconnaissance UAVs. Present uses, however, differ substantially in that the purpose is to better identify road-traversing targets and improve geolocation estimates using already very accurate sensors. In such larger vehicle contexts, online pursuit is also nearly never considered, as given the omnidirectional nature of such radars and the typically huge distances to targets, even substantial motions on the part of the sensing aircraft can have negligible effect on the geometry or availability of observations. Further, missions for these aircraft are often pre-planned with little infrastructure in place to direct their paths based on the information needs of remote teams. This is in immense contrast to the smaller UAVs considered here that *must* continuously move to maintain view given small, finite sensor footprints providing low-accuracy observations and whose primary mission is active reconnaissance.

3.3.3 Control for pursuit

To maintain view of a target that was recently observed and to continue to reduce uncertainty in its location, flight paths that accomplish this must be chosen. This is of particular necessity for fixed-wing aircraft in constant motion for which turn decisions must be continuously made and when pursuing moving targets. Often closely coupled with geolocation schemes, a number of strategies for accomplishing this have been previously proposed and may be roughly separated in to two categories.

The first class of approaches are heuristic methods to maintain tracking using geometric or reactive guidance rules. Most rely on the intuition that keeping a pursued target as close to the center of the sensor footprint as possible will maintain view in the short-term and provide maximum robustness to disturbances or estimation errors that may place the center of view away from the target's true location. The simplest such means is a control law that directs the UAV to loiter near a target's estimated location and drives the sensor footprint's center to this estimate. For aircraft having a side-looking or gimbaled camera, a common strategy is to enter an orbit around the estimate of radius necessary to maintain that as the desired footprint center. [67, 80, 83] Alternative derivations without the explicit goal of an orbit but instead maintaining a constant relative viewing angle of a target are presented by Rysdyk [90] and Theodorakopoulos and Lacroix. [102] This has been extended to teams of multiple orbiting aircraft, in which by maintaining equal spacing between vehicles, maximally diverse viewpoints are provided. [37, 101, 115] This provides an informal solution to the problem of maximizing observation information diversity described next. For aircraft instead having downward or forward-pointing cameras, sinusoidal, square-wave, and flower-petal trajectories parameterized by a current target location estimate have been proposed to stay on top of targets of varying speeds. [68, 97] For targets having uncertain location, ellipsoidal orbits tracing level curves of a Gaussian uncertainty distribution providing reasonable coverage of an ellipsoidal region around a mean estimate have been suggested. [35, 61] For cases in which target geolocation estimates may be unavailable or particularly uncertain, visual-servoing has been used to direct UAV motion using control laws tied to the pixel coordinate location of a visible target, usually with the goal of driving it to the image center. [91, 92] When automatic detection is not available, similar control laws may be used that map user joystick input indicating desired displacement of the image center relative to the target to UAV bank commands. [80]

Conceptually, as long as the UAV can react sufficiently quickly to resulting control outputs, pursuit of moving targets is possible in addition to persistent coverage of a static target by simply adjusting control law set-points as new locations are observed. However, heuristic approaches provide no guarantee of successful persistent pursuit or even necessarily applicability of a particular form of trajectory to a given target scenario. By their nature, they also provide no intrinsic means for recovery if a target is not observed where it was expected to be present (or, with the few noted exceptions, handling of growth in target location uncertainty), since they typically consider only most-likely point locations of targets. Finally, they do not take into account the time required for a UAV to react to changes in target location, even if that is easily predicted from currently visible motion, instead just hoping that the settling time of a control law is small relative to the time required for a target to move a substantial fraction of the sensor footprint radius so as to maintain view.

A second class that addresses many of these weaknesses and provides formal reasoning about both uncertainty in target location and limited UAV maneuverability is that which may be generally described as predictive trajectory planning seeking to maximize expected observation quality. The essential principles underlying this are that uncertainty in a target observation is dependent on the pose of a UAV relative to the target (appearing, for instance, in Equation 3.4 as $\mathbf{J}$ varies with this relative pose), while a-posteriori geolocation filter uncertainty is necessarily dependent on the uncertainty of observations received (and thereby the relative UAV poses from which the observations were made). Intuitive graphical examples of these concepts are provided in Figure 3.2.

The trajectory optimization problem may first be defined formally, albeit generally. Suppose a present target location uncertainty distribution $p(\mathbf{x}_t|\mathbf{Z}_t)$, target motion model $p(\mathbf{x}_{t+1}|\mathbf{X}_t)$, and UAV sensor observation model $p(\mathbf{z}_t|\mathbf{x}_t)$ are known, where $\mathbf{Z}_t = \{\mathbf{z}_1, \ldots, \mathbf{z}_t\}$ is the observation sequence received so far and $\mathbf{X}_t = \{\mathbf{x}_1 \ldots, \mathbf{x}_t\}$ is a hypothetical target path. Further, let $C(\mathbf{x}_{\mathrm{UAV},t}, \mathbf{z}_t|\mathbf{X}_{\mathrm{UAV},t-1}, \mathbf{Z}_{t-1})$ be a scalar cost function indicating the value of receiving an observation $\mathbf{z}_t$ while having pose $\mathbf{x}_{\mathrm{UAV},t}$ and having previously followed the path $\mathbf{X}_{\mathrm{UAV},t-1} = \{\mathbf{x}_{\mathrm{UAV},1}, \ldots, \mathbf{x}_{\mathrm{UAV},t-1}\}$ and received the observation sequence $\mathbf{Z}_{t-1}$. Then, the optimal action to choose at time t having followed path $\mathbf{X}_{\mathrm{UAV},t}$ and having received observations $\mathbf{Z}_t$ is

$$\mathbf{u}_t^\star = \operatorname*{argmax}_{\mathbf{u_t}} \int_{\mathbf{z}_{t+1}} \int_{\mathbf{x}_{t+1}} \int_{\mathbf{x}_t} \begin{array}{l} C\left(f_{\mathrm{motion}}(\mathbf{x}_{\mathrm{UAV},t}, \mathbf{u}_t, \Delta t), \mathbf{z}_{t+1}|\mathbf{X}_{\mathrm{UAV},t}, \mathbf{Z}_t\right) \\ \cdot p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1}) \\ \cdot p(\mathbf{x}_{t+1}|\mathbf{X}_t) \\ \cdot p(\mathbf{x}_t|\mathbf{Z}_t) \end{array} \tag{3.6}$$

and may iterated repeatedly to evaluate the value of applying a control sequence $\mathbf{U}_t = \{\mathbf{u}_1, \ldots, \mathbf{u}_t\}$ and receiving observations $\mathbf{Z}_t$ (of which all but the first are hypothetical) if the target were to follow a hypothetical path $\mathbf{X}_t$ from initial UAV and target state estimates $\mathbf{x}_{\mathrm{UAV},1}$ and $p(\mathbf{x}_t)$ respectively. Note that choosing an optimal $\mathbf{U}_t^\star$ from an initial $\mathbf{x}_{\mathrm{UAV},1}$ and $p(\mathbf{x}_t)$ thus constitutes an additional maximization over all possible $\mathbf{X}_t$ and $\mathbf{Z}_t$ that may occur during its execution.

The cost metric $C(\cdot)$ may be defined a number of ways, so as to either maximize the probability of target visibility (intuitively keeping the target in view), maximize the certainty of or information provided by observations (intuitively seeking "good" observations), or maximize the so-called *information gain* of successive observations to minimize a-posteriori filter uncertainty (intuitively seeking "varied" observations).

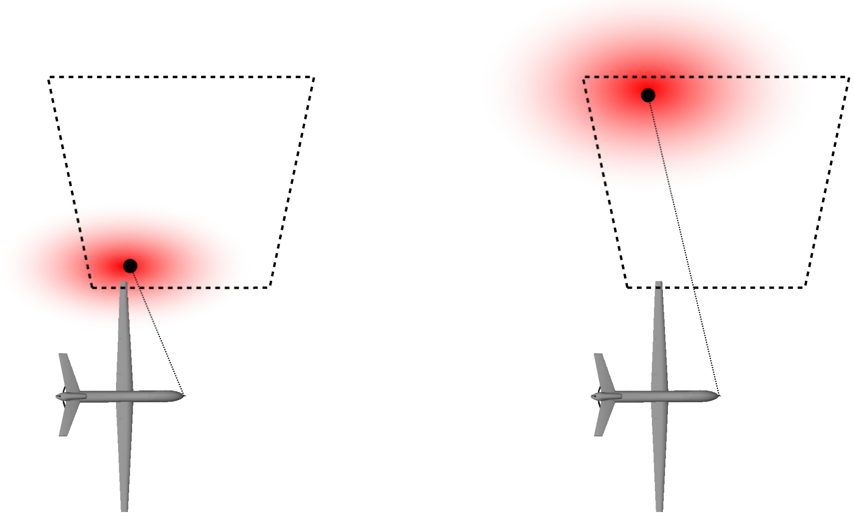

(a) Short-range observation

(b) Long-range observation

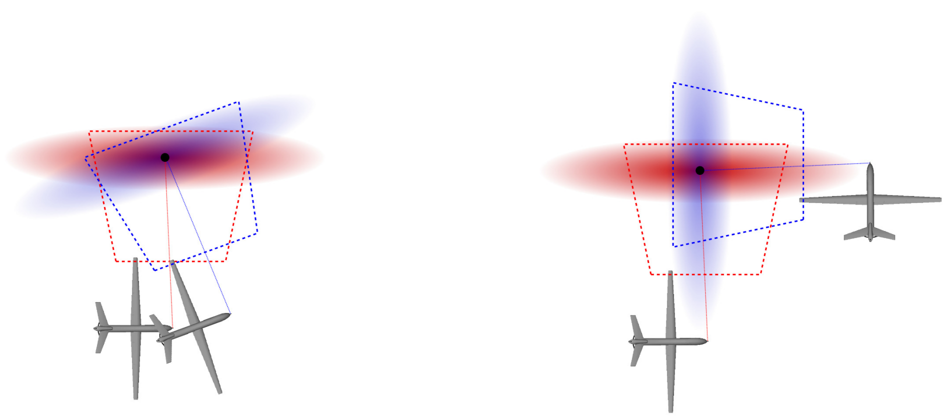

(c) Similar-pose observations

(d) Orthogonal-pose observations

Figure 3.2: Intuitive examples of varying UAV pose affecting the uncertainty of a single observation (here, that observations of closer targets (a) provide lower uncertainty in general than observations of more distant targets (b)) and the a-posteriori uncertainty of a filter receiving two observations from differing relative poses (here, that two observations having elongated uncertainty ellipses provide nearly overlapping uncertainty when taken nearby, (c) while the same observations may have much smaller overlap—admitting a much smaller set of possible target locations likely under both—when taken from orthogonal orientations (d)).

Scalar representations of observation or a-posteriori Gaussian uncertainty are commonly defined using varied matrix properties of either the so-called Fisher Information Matrix (FIM) or the relevant covariance matrix. Definitions and relative strengths of each matrix metric in the context of UAV tracking are thoroughly discussed by, for instance, Ponda [79] and Skoglar. [96] In practice, both anecdotally and as reported by some authors, [19] many of these produce very similar results. Finally, cost metrics may be defined in several ways over a trajectory, the two most common being the average value over expected observations (so as, for instance, to expect to view a target as frequently as possible or achieve consistently low observation error) and the terminal value (for instance, minimizing final a-posteriori uncertainty at the possible cost of potentially high interim uncertainty).

Quite clearly, directly optimizing Equation 3.6 is completely impractical, as it implicitly invokes hypothetical geolocation filter updates and each of the dimensions over which it integrates and optimizes is a continuous space. Thus, practical implementations make one or more substantial approximations. One of the simplest is to discretize any of the relevant spaces, such as considering only a small discrete set of possible UAV actions or perhaps a single most likely target motion given its current estimated state. Others include ignoring filter uncertainty so that $p(\mathbf{x}_t)$ is a delta function centered at the filter mean, ignoring vehicle pose uncertainty or assuming that only the most likely observation will occur (reducing $p(\mathbf{z}_t|\mathbf{x}_t)$ to a delta function), not running filter observation updates using predicted observations (eliminating the $p(\mathbf{x}_{t+1}|\mathbf{x}_t)$ term and effectively assuming a stationary target), planning assuming negative observations (which are usually faster as no camera-ray geometry need be computed) when maximizing the probability of seeing the target at least once during the trajectory, and metric-specific approximations such as roughly emulating information-gain by simply summing expected information matrices and applying the desired matrix metric at the terminal point of the trajectory.

Existing strategies applying such optimization for aerial pursuit fall into one of several categories. A common choice in the literature is to simply and greedily move along the gradient of the selected information metric, corresponding to a single-step optimization that may typically be computed quite efficiently. [19, 49, 50, 79] As expected, a major failing with these is a frequent rush to a local minimum, which may be particularly fatal for sensors such as cameras with finite field of view: since typically the lowest-uncertainty measurements are achieved when the UAV is as close as possible to the target, iterated optimization can easily cause the UAV to simply overfly the target and lose sight. Additionally, once a target is out of view, re-acquiring it is not necessarily guaranteed, as oscillatory behavior skirting the target or even flying directly away may result should a non-minimum phase trajectory have been required to reach it again. Slow direct search or cost-function optimization up to some horizon has also been used, [38, 39, 49, 89] but unsurprisingly, gross approximations, sparse discretizations, and very short horizons are required to maintain tractability. To improve upon this, pre-computation of optimal short-term trajectories given an estimated relative target pose using function optimization or dynamic programming have been proposed, [26, 82] greatly aiding tractability albeit still requiring substantial approximations.

A strongly related issue for multi-UAV teams is the level of collaboration between vehicles. [2] On one end of the spectrum, lies completely independent behavior, in which vehicles geolocate and plan independently, possibly reporting their independent estimates to one another or a central location for comparison or averaging. Substantially overlapping behavior and redundant observations may result, but all computation and communication may be completely decentralized and parallelized. Above this lies forms of ad-hoc cooperation, in which vehicles may asynchronously broadcast their observations, estimates, and intended motions so that each agent has the best possible information with which to choose its motions and may avoid areas or perspectives that are presently well covered by others, but without explicit group decision-making. A stronger level of collaboration still is that simply termed here cooperative behavior, also known as sequential allocation, in which agents synchronously report their intended motions so that each successive vehicle chooses its actions knowing what preceding ones will do in the near-future. This permits elimination of redundant coverage, but it may still produce arbitrarily poor behavior if the action cost metric is not chosen appropriately. A classical example is that of attempting to pursue two targets using two UAVs: if the metric chosen is to minimize the average per-vehicle distance between the sensor footprint center and each target's estimated location, both will choose to fly towards the midpoint of the target locations and fail to make any observations; if instead, the metric is maximizing average a-posteriori geolocation certainty, they will appropriately assign themselves to a different target. Finally, fully coordinated or joint motion planning, typically implemented as a central planner, treats the optimization problem as a higher-dimensional one in which an action is in fact a vector of actions, one for each vehicle. This permits proper consideration of arbitrarily complex joint behaviors, in which one UAV may be behaving in a highly suboptimally manner locally but yet contributes critically to reaching a global optimum. Clearly, progression along this spectrum is accompanied by greatly increased computational complexity, as action spaces merge and oppor-

[2]The terms *cooperation* and *coordination* have received varying definitions in the literature. Here, *cooperation* refers to loose levels of collaboration in which agents may assist one another, while *coordination* implies some form of strongly-coupled joint behaviors.

tunities for parallelism disappear. Two ways to mitigate this are to use sparser or lower-dimensional representations wherever possible to simply reduce problem complexity or to use separable or factorizable representations such as discrete hypotheses to which separate vehicle-to-target assignment may be applied.

Overall, pursuit represents a particularly challenging task, for several reasons. First, the finite size of a UAV's sensor footprint and limited maneuverability results in slow returns to view as the UAV must circle back to re-acquire an overflown target, producing so-called bursty observation dropouts. Uncertain target motion prediction increases the space of likely target locations with the passage of time, making long-horizon planning difficult and reducing the duration of dropouts that may be handled, while large uncertainty in vehicle state produces many possible observation outcomes for a given target location. Taken together, and despite many existing efforts, planning to produce best expected observations remains largely fundamentally intractable. Particularly agile (fast-moving or high-maneuverability) targets or those assumed to be taking evasive action only worsen this, as the space of likely target locations increases greatly between observations. A sure means to improve upon this is to reduce the space of possible target location or actions however possible, and one such way using a sparse and constraining environment representation is proposed in Section 5.3.

3.4 Search with UAVs

If instead a target's location is not known–either because it has never yet been observed or it successfully evaded tracking and has not been seen for some time–a search is required to locate it before pursuit. As small UAVs, particular early ones, have limited payload capacity (thereby restricting possible sensing to lightweight cameras) and maneuverability (limiting pursuit performance), search for targets or objects of interest is a natural application and has likewise received considerable attention. Roughly speaking, existing strategies may be classified based on the level of uniformity with which search attention is directed across the area of interest.

3.4.1 Search in regular areas with uniform prior

If an area being searched is in some sense regularly-shaped (for instance, tightly bounded by a convex simply-connected polygon) and uncertainty in target location is uniformly distributed throughout the area, the search task may be simply phrased–though it may not necessarily be easily accomplished–as covering the entire area using uniformly-directed search attention so as to have observed every point in the environment before a target possibly located in adjacent un-searched space would have had a chance to reach that point.

For stationary targets, this represents a relatively straightforward uniform coverage task. Ablavsky and Snorrason [1] provide commonly cited geometric strategies for UAVs modeled as Dubins cars using straight-line sweeps to form lawnmower, zamboni, and box-spiral patterns. In a similar vein, coverage using autopilot-provided orbit motion primitives has been considered in work to date described in Section 4.2. Extending this to multiple vehicles requires some form of region assignment to avoid redundant coverage or the possibility of vehicle collision. Enns et al. [32] suggest extending single-UAV sweeping to interleaved sweeps by multiple vehicles, addressing turning constraints. Beard and McLain [11] consider similar forms of sweeping while avoiding hazards and maintaining a minimum communication distance between vehicles should this be required. For larger areas that may not necessarily form a simply polygon, Maza and Ollero [70] propose a polygonal partitioning of an area, such that one partition is assigned to each UAV and covered using an existing simpler sweeping strategy. An alternative decomposition in which an area is decomposed into slightly-overlapping circular regions that are then covered by spiral coverage patterns is instead put forth by Girard et al. [43] Intriguingly, Jones [57] demonstrates aerial coverage by generating a planar coverage pattern using well studied methods for multiple ground vehicles and maps these to executable UAV waypoints. Overall, the case of stationary targets may be considered both well-studied and straightforward, however long-duration searches or many vehicles may be required for task completion in large areas.

Targets that may be moving adversarially in an otherwise regular environment having initially uniform location uncertainty must be treated somewhat differently, as target motion will recontaminate already-searched regions. Given a velocity bound on a target sufficiently below that of a UAV and modeling the UAV sensor footprint as a circle lying within the true footprint, McGee and Hedrick [71] show adversarial target search using containment by progressively shrinking irregular orbits. For larger areas and faster targets, Vincent and Rubin [113] suggest parallel sweeping modeling sensor footprints as small rectangles and state a relationship between the required number of UAVs and their minimum necessary speed to assure detection of a target having a particular maximum speed in an area of given size. Overall, the case of adversarial targets may be considered well-studied but is a difficult one as many fast UAVs may be required to assure detection of a target.

3.4.2 Search in irregular areas or with non-uniform prior

If the search area instead has highly irregular shape or substantial interior regions in which a target could not lie (such as obstacles or otherwise untraversable terrain), an outer bounding polygon or decomposition into few convex polygons may present a poor approximation of the area and result in needless coverage of areas in which a target could not lie. Similarly, if the prior uncertainty distribution on possible target location is highly non-uniform, it may be wasteful to cover low-likelihood areas at all and at least logical to search higher-likelihood areas first to achieve shorter expected detection time. Finally, for the case of a target that is moving non-adversarially and modeled stochastically, uniform coverage search as for a stationary target is inapplicable due to recontamination, yet adversarial search may be needlessly conservative, requiring excessive capture time or vehicle resources.

For stationary or stochastically-modeled targets, cellular observation-predicting trajectory planning is commonly applied, usually so as to choose a path for the searchers that maximizes the probability of detection or expected detection time along the anticipated coverage pattern. A first example that still results in uniform coverage of a region is that shown by Ahmadzadeh et al., [4] which plans paths subject to motion constraints and arbitrary initial conditions to maximize the average probability of observing any point in the environment during the path. Polycarpou et al. [78] present a framework later built upon by Flint [33] that plans paths for a team of UAVs searching for a stationary target in an environment with areas having varying likelihood of target presence as well as possibly moving threats, which is shown to provide substantially improved detection time over a uniform coverage search. For moving targets, it is common to discretize the environment by either decomposing it into cells containing a probability of target presence lying only within feasible regions or by distributing particles according to an initial prior distribution, running Bayesian probability updates as the search progresses as though for a cellular or particle geolocation filter from Section 3.3.2, and then optimizing paths that maximize probability of detection given expected coverage of cells or particles and their respective probability. Bourgault et al. [14] present a famous early example, in which one or more UAVs searches for a randomly diffusing target represented in a probability grid by moving along the local probability gradient. Despite this local behavior, by sharing observations, well-coordinated behaviors are demonstrated. Tisdale et al. [103, 104] implement and field-test a multi-UAV sequential allocation scheme using horizon planning to maximize the short-term probability of target detection using a similar grid or particle representation receiving Bayesian updates that fully take into account sensor uncertainty, detection likelihoods, and false-detection likelihoods. Yet another example includes that which Geyer [41] describes, using a similarly-updated particle representation for single-vehicle search for a target taking into account occlusion using a pre-computed visibility map and which contains an optimization to greatly speed up hypothetical particle filter updates given expected coverage for a given trajectory. Of particular note is his observation that memory-less planning (that does not perform hypothetical filter updates in response to expected observations at all but simply independently attempts to maximize the probability of target detection at each step) performs nearly as well as full planning. Conceptually, these forms of planning that optimize detection time or short-term probability of detection are fundamentally identical to the aforementioned trajectory planning for pursuit of Section 3.3.3, [50, 89] but can be perceived as having more target hypotheses and much weaker target location certainty, greatly worsening planning tractability and in practice requiring stronger approximations or suboptimal planning. Proposed work in Section 5.5 suggests that in some environments, these methods may still prove quite practical if the environment and target uncertainty can be represented compactly.

Meanwhile, search for adversarially moving targets having non-uniform initial location probability or in irregularly-shaped environments has not been well considered in the context of aerial search. For some environments, the aforementioned decomposition into regions that are then uniformly swept for adversarial targets may be feasible, but many details such as possible target motion between regions remain unclear. As reviewed in Section 3.1, adversarial search in arbitrary polygons has been well-studied for ground vehicles and serves as inspiration for proposed work mapping UAV search tasks to such abstractions.

3.5 Recapture

If observations are not received during pursuit for a long duration, tracking may be considered to have failed. At this point, the predicted location of the target is likely to have uncertainty covering a wide area, providing little use to short-term pursuit-level observer motion planning, as any trajectory covering a portion of the uncertain area will have nearly the same poor likelihood of detection. It may also happen that a target sighting is briefly reported, but no UAVs are near the reported location to begin pursuit. By the time any can reach the area, uncertainty in the target's location may have grown well beyond the point at which a UAV could expect to see the target at its reported location; rather, some form of local search is required. This case of recapture does not appear to have received attention in the aerial search and tracking literature, yet in practice it may arise frequently when attempting to track agile or evasive targets with poorly maneuverable, low-accuracy sensors. Resorting to initiation of a search of the surrounding area may be heavy-handed, possibly requiring the requisition of additional UAVs, substantial setup time, or the coverage of an unnecessarily large area by its conclusion.

Though the recapture scenario appears to have received little consideration in the existing literature on UAV search and tracking, several intuitive and some existing pursuit-evasion strategies may be applicable. For instance, if a target is lost during pursuit, varying possible target motions can be predicted to generate discrete hypotheses of increasing uncertainty. An optimal trajectory can be planned as for pursuit in expectation over all these possible motions, though likely requiring considerably more UAVs than for typical pursuit to achieve meaningful results. One example of this is the iteration of a process model along roads, splitting at branches to produce discrete hypotheses, [47] a concept built upon further in Section 5.4 in realization that the tractability of producing recapture plans depends critically on having few, moderately-accurate hypotheses of target location implying a strong need for great target predictability. This is of course difficult in open environments in which, subject to possible motion constraints, the target may move in nearly any direction. Thus, where having few hypotheses is impractical, it may be possible to simply diffuse possible target location using a speed estimate or bound, creating a progressively larger area that may be searched using any of the preceding methods for areas with non-uniform prior. Given a sufficient number of UAVs, it may also be feasible to optimally cover the target uncertainty distribution using sensor placement techniques [48] so that the target is most likely to be detected both instantaneously and once it moves. Finally, there may be cases in which a target's location is moderately actually known yet it cannot be continuously observed because its speed is comparable to that of pursuing UAVs, so that the target is in some sense constantly in need of recapture. This may be treated as more of a pursuit-evasion problem, having, as previously mentioned, limited mention in the aerial tracking literature.

CHAPTER 4

Work to date: Representational Improvements for Practical Search and Tracking

Effort to date has broadly focused on understanding the fundamental difficulties underlying the practical tasks of search and tracking through field experimentation. Section 4.1 describes in detail various aspects of these experiments, basic challenges encountered, and the essential conclusions that motivate improved representations for better performance. One simple but critical algorithmic extension providing a practical means of area coverage search is conveyed in Section 4.2, while Section 4.3 explains the need and several strategies for accurately representing and combining high-uncertainty observations in the face of a highly nonlinear observation function to produce geolocation estimates with otherwise impossible accuracy. As a first step towards exploitation of environmental structure to improve task performance, the use of a road constraint during pursuit is considered in Section 4.4 and is shown to have a number of obvious but inspirational benefits. Extension of this idea to road networks is motivated by the impracticality of a naive cellular representation in the style of existing work. Finally, conclusions to date are summarized in Section 4.5, strongly motivating the proposed use of environmental structure in the succeeding chapter to simplify the problem and permit the use of theoretically analyzable abstractions.

4.1 Experimental Motivation

Experiments to date have considered the use of stock or minimally-modified commercially available small robotic vehicles so as to build upon already highly refined low-level controllers and autopilots to permit high-level algorithmic study and demonstrate capability extension through field-installable software. Taking inspiration from the complementary nature of differences betwen air and ground vehicles noted in Section 2.5–higher fidelity sensing and larger computing payloads available from ground vehicles, but faster wide-area sensor coverage with little concern for terrain obstacles from aircraft–a heterogeneous experimental collaboration platform comprised of one or more UGVs and one or more UAVs was used.

The base system is comprised of a single UAV, a single UGV, and a common high-level control interface. The UAV platform selected for this system is the RQ-11B Raven produced by AeroVironment, Inc and pictured in Figure 4.1(b). Relatively inexpensive and easily transported by a single person, the Raven can be hand-launched from a small clearing, providing an excellent example of the class of UAVs considered. It is also possibly the currently most widely deployed UAV system in the world–with over 10,000 units delivered–so that applicable software extensions may have immediate worldwide impact. As the UGV platform, a stock iRobot Explosive Ordinance Disposal (EOD) Packbot with a development payload and additional sensing is used. Based upon a very widely used small ground robot that is easily transported in a Humvee and deployable by a single operator, it is pictured in Figure 4.1(a). Finally, the common control interface appears as software running on a ruggedized laptop computer (such as that shown in Figure 4.1(c)) for overall system portability.

Primarily only aspects of UAV search and tracking are considered here, though such hetereogeneous collaboration has shown itself [28] to be of great benefit to the completion of practical reconnaissance missions and is considered both as a possible extension within proposed work as well as a clearly fruitful area of future study.

4.1.1 Unified control infrastructure for fielded vehicles

As nearly all such vehicles, including those selected as the experimental platform, are in practice either manually controlled or directed using manually specified waypoints individually provided to each vehicle, a first step towards autonomous collaborative field-reconnaissance is a unified control interface. This was implemented in the form of an overhead satellite map window on which overall mission state is overlaid and a function-specific control window for each active vehicle, examples of which appear in Figure 4.2. Using this, an operator might for instance select an area to search using

(a) (b) (c)

Figure 4.1: The UGV, UAV (previously pictured in Figure 1.3), and computing hardware that comprises this heterogeneous collaborative surveillance system. All components are man-portable and battery operated.

the map window, designate an object of interest for tracking in live video presented in a vehicle's control window (perhaps an initial UAV detection), observe its geolocated position back in the map window, and direct another vehicle to the same location with a single high-level command (perhaps to direct a UGV for closer inspection). Described in further detail in a system-specific publication, [28] such interfaces are not a focus of this thesis, though they are a necessary component of an effective collaboration system using multiple vehicles performing coupled search and pursuit maneuvers. Additionally, with few existing examples, [52, 98] this is believed to be one of the first demonstrations of such unified air-ground mission execution.

4.1.2 Geolocation observations from visual tracking

In order to produce observations in the form of Equation 2.1 with which to generate geolocation observations, some means of continuous target detection or tracking is required. For the relatively wide angle camera and lossy video transmission of the UAV considered, coupled with the desire to track novel targets in rescue or surveillance scenarios, an operator-initiated tracking scheme is chosen over any form of automatic visual target recognition, though one such form is briefly considered. Under this scheme, as a UAV searches an area or attempts to recapture a target, targets visible in the video feed are designated by an operator click, these designated targets are then tracked through succeeding video frames without further operator intervention by one of several algorithms, and a recovery procedure is executed to attempt to re-acquire targets that have become occluded by terrain or that have left the view

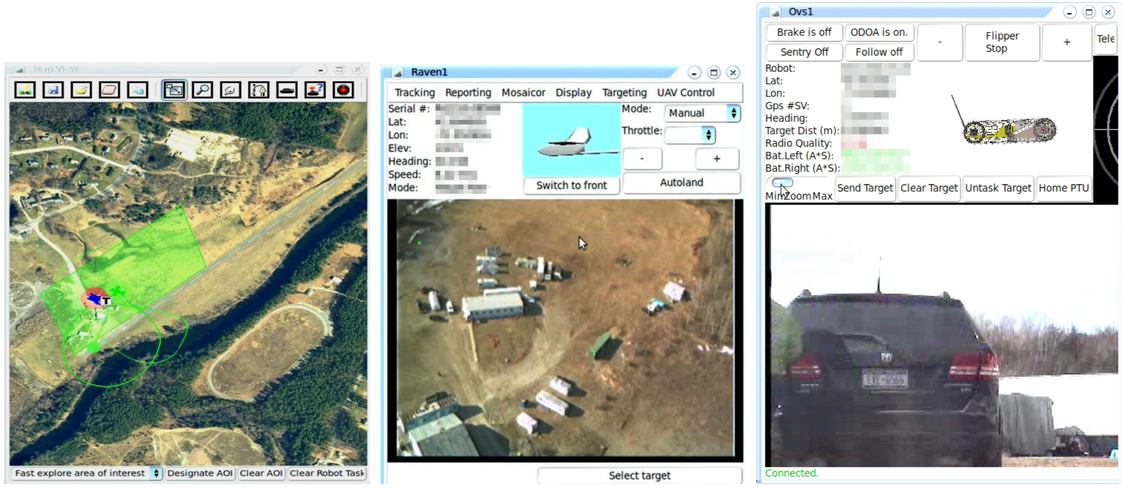

Figure 4.2: Screenshots of the overhead map view, UAV, and UGV control window components of the common operator interface for unified collaborative control.

of the camera. The center of this designated or tracked position is used as the image coordinate observation and is then used to generate a ray that is intersected with an estimate of the local groundplane derived from an assumed terrain database as described in Section 3.3.1 to produce a single geolocation observation in the space of target location.

The system developed handles high levels of camera motion caused by wind, performs out-of-view tracking recovery that is frequently required in practice, and handles tracking and recovery through video corruption or blackouts caused by communications glitches. As previously noted, wind is a particular problem for aircraft of such size as it causes rapid accelerations in the vehicle'ss attitude and hence considerable unwanted camera motion. This means an objects location in one video frame might be significantly different in subsequent frames, causing failures in many traditional tracking schemes and also making it difficult for a human user to initially designate a target in the video stream.

For the purposes of this work, it is assumed that the camera does not have variable zoom capability or a mechanical gimbal for pointing. The Raven has both electronic zoom and software panning abilities, however these are not used both for greater applicability to other UAVs and because these were not designed for automatic software control in the case of the Raven. The algorithms described here will operate perfectly on a vehicle with either and may be extended to take advantage of them.

The visual tracking pipeline implemented is described more completely in a concurrent report on the work. [73]. Briefly, it is divided into several segments. The first is image stabilization to assist a human operator's designation or an automatic tracking algorithm by eliminating effects of rapid unwanted camera motion caused by wind-induced attitude disturbances. Lucas-Kanade alignment [7] is one of the formative approaches. However, with large amounts of image motion, modification of this traditional approach to avoid local minima in its internal alignment search process was required. The next segment is the process of designation itself. Adopting the methodology of operator-assisted designation rather than highlighting a target in previously captured imagery or by using a joystick to continuously aim a gimbal or cursor, [80] the operator either draws a box around the target that is then automatically sized to best fit the identified target or clicks on the center of the target, causing an estimated box to be drawn around it that is then likewise resized. The last segment of the tracking process are the algorithms themselves used to track a designated target. For this purpose, both a mean-shift color template tracker (using a direct implementation of this concept [23] as well as a variation of an alternative [22] that continuously re-weights a set of features derived from functions of color values) and spatial appearance tracking using a modified form of the well-known Kanade-Lucas-Tomasi (KLT) algorithm. [106] Tracker selection remains a challenge given differing performance in varying scenarios highlighted in Figure 4.3. On average, the spatial template tracker performs best and is selected by default, though the operator may choose an alternate tracker at any time if appropriate.

(a) Changing perspective (b) Similar objects (c) Background in template

Figure 4.3: Examples showing two cases, (a) and (c), for which non-spatial tracking dominates, and another, (b), for which a spatial template is more appropriate.

Inevitably, a tracked target will often leave the view of a UAV's camera, especially when lacking a gimbal or if zoomed in, reducing the field of view. The described system attempts to recover from out-of-view tracking losses by estimating the camera motion by iterative image stabilization to predict the new location of the target, and when the predicted location re-enters the view, searches for the target inside a sizeable search window surrounding this expected location. In addition to out-of-view events, other factors may induce tracker failure. For instance, analog and digital radio communications loss causes corruption in the video stream. Failures in the tracker are heuristically indicated as occasions when the template no longer matches well with the search window, and then recovery is attempted once the period of corruption is over. An example of this procedure applied to both video corruption and an out-of-frame event is shown in Figure 4.4.

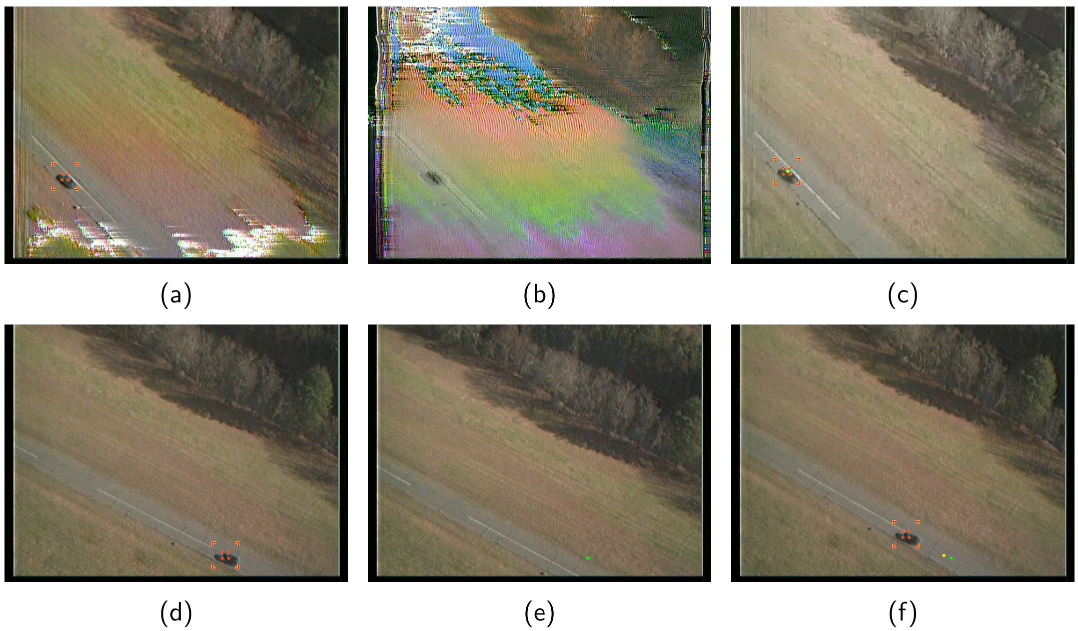

Figure 4.4: Different types of tracking recovery. The top row shows the tracking recovery after a period of communications corruption. The bottom row shows the tracking recovery after the target leaves the field of view.

4.1.3 Initial experimental conclusions

The described tracking pipeline, built on practical modifications of two popular template tracking methods, has proven incredibly reliable in practice, with tracking durations on the order of a minute without operator intervention not uncommon for unambiguous targets in still weather. Successful tracking has been demonstrated in open areas and of targets moving on roads and within clutter, examples of which are shown in Figure 4.5. Formal testing against standard tracking datasets remains to be performed, however.

(a) Open area (b) On-road (c) Cluttered area

Figure 4.5: Examples of tracking a target in the open, moving on a road, and in a cluttered area.

Overall experimentation has accumulated several tens of hours of testing and data collection in a variety of environments, in varying wind conditions, of stationary and varying-speed targets including several pursuit trials. Two main conclusions based on experimentation to date might be made.

The first is that individual target geolocation observations may be extremely inaccurate, primarily due to errors in vehicle orientation (particularly heading) caused by latency and jitter in the state data stream, wind-induced angular velocities exceeding the signal bandwidth of both onboard state estimation and the state data stream, and the lack of a reliable heading reference beyond Earth's weak magnetic field used by an onboard compass. In fact, common existing geolocation filtering methods chronicled in Section 3.3.2 provide only moderate improvement due to the magnitude of these errors and slowly varying biases. Section 4.3 presents several ways that have been considered to improve upon this.

Second is the observation that common autopilot implementations providing slow response to desired control inputs–as both latency in acting upon these and a low frequency cut-off in system dynamics–coupled with the typical prioritization of vehicle stability and safety over responsiveness results in large minimum turning radii (50 to one hundred meters) and high latency for a control input to take effect (up to several seconds). The practical impact is that if a target is overflown and the vehicle must turn around, the process may take on the order of half a minute, producing long dropouts in observations of a target, during which motion predictions coupled with a prior geolocation estimate are relied upon. This motivates the use of multiple UAVs to maintain persistent tracking during individual dropouts as well as the importance of accurate geolocation filtering and careful modeling of target motion to maximize the accuracy of predictions. Section 4.4 introduces the use of structure as a promising means to accomplish this.

4.2 Practical Open Area Coverage Search

When performing an area coverage search, a particular challenge arises in the detection or designation of targets in that a sufficiently clear image must be observed for a sufficient duration. This is difficult for the class of UAVs considered for several reasons. First, noisy or lossy video may require a number of viewings for a clear, recognizable image. Simultaneously, human operators are necessarily slow to react to seeing something of interest and designating it. Meanwhile, viewing from multiple angles might be necessary for automatic or human recognition of a target either given possibly asymmetric placement of prominent object features or unanticipated directional environmental occlusions. On top of this, if an object of potential interest is perceived but passed by, a lengthy vehicle turnaround time to return for re-inspection may be necessary given a large turning radius.

Therefore, a practical form of area coverage search was developed as a means to provide temporally-contiguous redundant observations of every point in an environment from multiple viewing directions. This is accomplished through the use of a coverage search based upon orbit motions rather than the more common use of linear sweeping described in Section 3.4.1. This search strategy requires only that an onboard autopilot provide an orbit-point control mode in which a side-looking camera is centered on the instantaneous orbit point center and is thus quite appropriate for small field-reconnaissance UAVs, unlike some of these aforementioned methods that require high-frequency steering updates.

Given a boustrophedon decomposition [20] of a polygonal search area $\mathcal{P}$ such as that in Figure 4.6 and a distance-optimal ordering of cell visitation, rather than directing the UAV to follow the plowing motions directly itself, the autopilot is fed a continuous stream of new desired orbit point (and hence sensor footprint) centers moving along this path. In practice, since a uniform area search suitable for a static target is unlikely to be particularly time sensitive given that the target must not be evasively escaping, all but large environment obstacles may be eliminated, greatly simplifying the decomposition and cell ordering problem. As an additional measure, obstacles causing plow segments to be shorter than two sensor footprint diameters are reduced or eliminated.

The "plow width" is chosen to be slightly less than the diameter of the circular area that would be covered by the rotating sensor footprint during a single orbit of a stationary point for slight overlap between adjacent path sweeps and thereby complete area coverage. Both the plow width and (to a greater extent) precession rate at which the orbit point is moved along the path are tunable parameters controlling the level of coverage redundancy desired. Nominally, a precession rate of one sensor

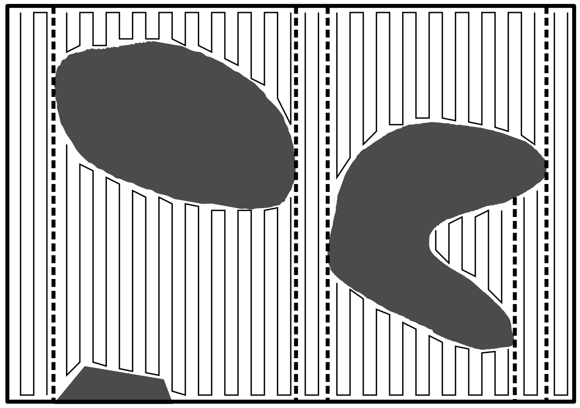

Figure 4.6: An example of a boustrophedon decomposition of an area containing two internal obstacles and an obstacle adjacent an area border, along with a possible coverage path through each region.

footprint diameter per time taken to complete an orbit assures coverage of every point from all viewing directions. New orbit points are provided to the autopilot at a rate of one every three seconds to ensure that the underlying orbit controller maximizes the time spent in its steady-state control region. A typical simple area search example is depicted in Figure 4.7, with examples of varying coverage redundancy parameters compared in Figure 4.8.

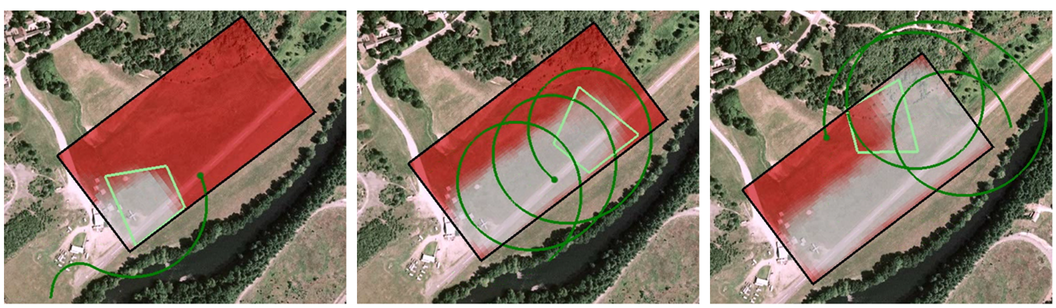

Figure 4.7: Several coverage snapshots from a search of a typical open rectangular area. The level of coverage is indicated by red shading, with completely uncovered areas displayed in the darkest shade of red, and sufficiently (exceeding an adjustable threshold on probability of detection) covered areas as white. The path of the UAV is drawn in dark green, its current location marked by a green dot, and the current sensor footprint as a light-green trapezoid.

As stated, the resulting trajectory is parameterized by time, resulting in a fixed path for the desired orbit point, independent of the UAV's actual success in main-

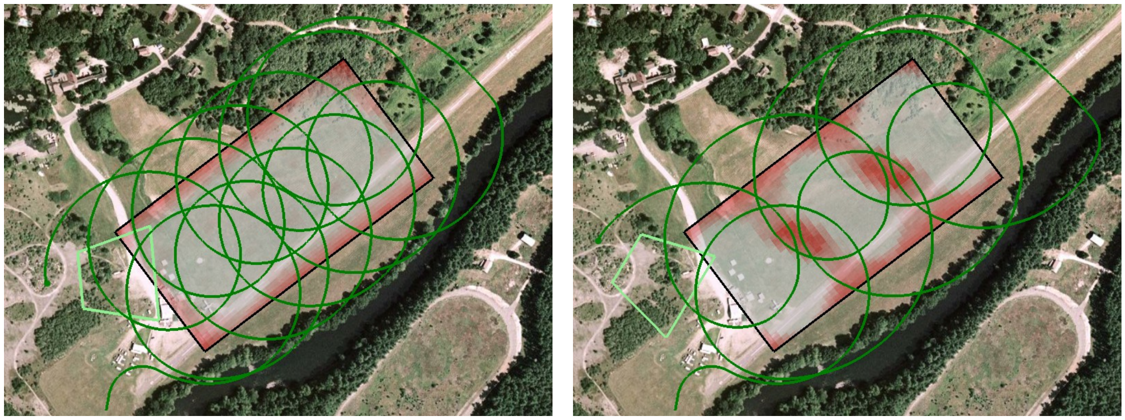

(a) High-redundancy coverage (b) Speed-optimized coverage

Figure 4.8: A comparison of two extremes in the resulting coverage pattern during a search, using (a) a high-redundancy (tight overlap) search and (b) a speed-optimized (low overlap) search providing possibly insufficient coverage of several interior regions. The level of coverage redundancy may be varied by changing a single parameter, the orbit-center precession rate. Colors and markings are as defined in Figure 4.7.

taining the intended orbit trajectories. In the presence of high wind, segments of a given orbit may take longer than nominal as the UAV faces a head-wind that reduces its ground-speed, while others less time given a tail-wind providing increased speed. Additionally, side-winds may induce trajectory tracking error that requires additional time to overcome. Therefore, an alternative parameterization providing some level of robustness is to continuously track actual angular position around each orbit and compute the matching point along the orbit point coverage path for the nominal case. Thus, in the extreme case, a hovering UAV paralyzed by wind would be perpetually provided an unchanging orbit point.

4.3 Improved Geolocation Under High Uncertainty

Once a target has been located and is being actively tracked, geolocation filtering such as using the existing methods reviewed in Section 3.3.2 is needed to provide an accurate estimate given a series of inaccurate individual observations. However, the nature of errors in observations produced by the described experimental platform presented particular challenges to these methods.

A typical plot of geolocation observations over a single orbit of a stationary target as shown in Figure 4.9 indicates the presence of a large, time-varying bias in addition to considerable per-observation error. In large part this is due to error in measurements of vehicle heading with which geolocation observations are made. Figure 4.10 depicts typical heading error over the course of several orbits, exhibiting error that is large, time-varying, and non-Gaussian in distribution. This violates typical existing filter assumptions of zero-mean or fixed-bias, Gaussian-distributed error. The immediate impact on uncertainty modeling, shown in Figure 4.11, is that Jacobian linearization of the highly nonlinear observation function in Equation 3.3 as used by most strategies in Section 3.3.2 can produce arbitrarily poor approximations of the actual crescent-shaped uncertainty distribution. Meanwhile, as also shown in the figure, imposing a Gaussian approximation of the true uncertainty distribution can at best provide an uninformative, over-conservative of little use to information-theoretic trajectory planning. Therefore, an effective filter representation must have two characteristics: it must handle the transformation of large vehicle state error through the highly nonlinear observation function and must represent target location uncertainty using a distribution that is non-Gaussian in this space.

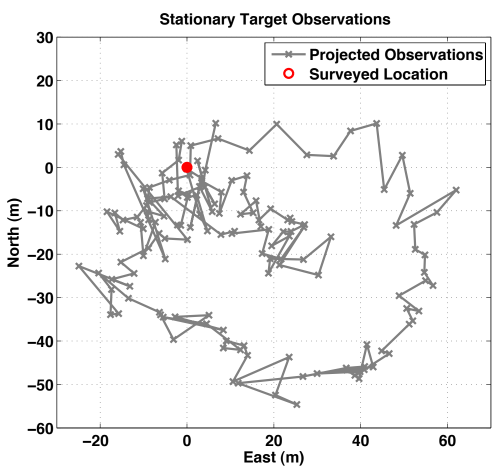

Figure 4.9: Observations of a stationary target over one orbit, exhibiting large, time-varying bias.

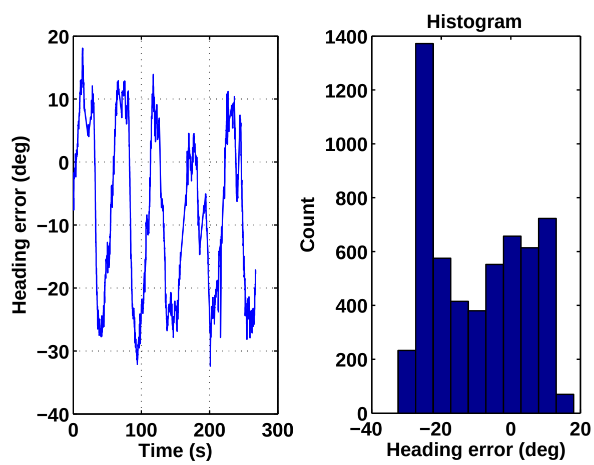

Figure 4.10: UAV heading error over several orbits, exhibiting large, time-varying, non-Gaussian distributed error.

While the magnitude of the errors described may represent a somewhat extreme example, the challenge of geolocation under high observer state uncertainty is a fundamental one and remains relevant with continued movement towards ever-smaller and inexpensive UAVs with little payload budget in every sense. Further, object localization with large, asymmetric error using a highly nonlinear observation function

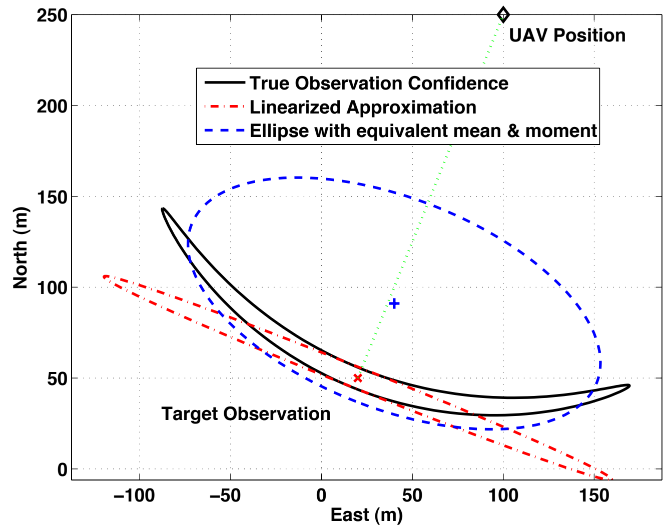

Figure 4.11: Comparison of the true observation uncertainty distribution with a Gaussian approximation generated by Jacobian linearization and a Gaussian approximation having the same mean and covariance as the true distribution.

represents a challenging nonlinear estimation problem in its own right.

To address this, three geolocation filters were developed that better handle this form and magnitude of uncertainty. The first uses an evidence-grid to accumulate sampled uncertainty over many observations of a stationary target. The latter two take inspiration from the highly polar nature of the uncertainty, one being a particle filter that receives observations in the space of range and bearing, and the second being a bounded recursive filter that operates in an over-parameterized space to parametrically represent uncertainty distributions having this crescent form. All have shown themselves to typically outperform existing methods on the geolocation data collected during experimentation, however thorough comparison between the filters and against existing methods remains to be performed.

4.3.1 Evidence-grid filter

To avoid modeling the complex crescent-shaped uncertainty distributions of each geolocation measurement parametrically, a discretized representation is considered that captures the shape of this distribution to produce a geolocation filter appropriate for stationary targets. [72]

The solution space – the location of the target in world coordinates – is first represented as a discrete grid, in which each cell contains the relative likelihood of the target being in that cell, much as a cellular probability grid receiving Bayesian likelihood updates during a search, however in this case cells do not contain explicit probabilities. Then, the essential idea is to sample the error distribution for each of the components of the vehicle pose and derive, after performing this and accumulating these distributions over a number of observations, the most likely target location. Formally, define

$$\mathbf{x}_{\mathrm{tgt,sample}}^{\mathrm{world}} = f_{\mathrm{sample}}(\mathbf{s}_{\mathrm{UAV}}, \mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}}, \mathbf{e}) = f_{\mathrm{obs}}(\mathbf{s}_{\mathrm{UAV}} + \mathbf{e}, \mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}}), \tag{4.1}$$

where f_{obs} is the pseudo-observation generation function defined in Equation 3.3 and $\mathbf{e}$ is a specific offset in the UAV state. The filter approach is to use $f_{\mathrm{sample}}(\cdot)$ to provide discrete geolocation hypotheses from a distribution of UAV state errors, $\mathbf{E}$, each comprised of many discrete state offsets, $\mathbf{e}$. The offsets that form $\mathbf{E}$ are sampled from the modeled error distribution of the UAV state. The error model comprises all dimensions of the vehicle pose, with the heading error being the dominant dimension in our case.

Next, let the pair of functions

$$\begin{aligned} (i, j) &= \mathrm{world2grid}(\mathbf{x}^{\mathrm{world}}) && (4.2) \\ \overline{\mathbf{x}}^{\mathrm{world}} &= \mathrm{grid2world}(i', j') && (4.3) \end{aligned}$$

define the mapping from any point in world frame to the integer index (i, j) in the grid, and from any integer index (i', j') in the grid to world coordinates, respectively.

Given this, the algorithm may be defined formally as shown in Algorithm 1. Starting with a blank grid representing the discretized world space comprised of $M \times N$ cells, all states of the vehicle $\mathbf{S}_{\mathrm{U}AV} = \{\mathbf{s}_{\mathrm{U}AV}\}$ and corresponding image pixel detections $\mathbf{X}_{\mathrm{tgt}}^{\mathrm{img}} = \{\mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}}\}$ at which the target was observed during a flight trajectory. For each of these, the algorithm loops through a discrete sampling of the space of possible error values $\mathbf{e}$ along the sensor or state dimensions constituting substantial error sources. For each of these, the observation function $f_{\mathrm{sample}}(\cdot)$ is applied to compute the world position of the target given this $\mathbf{e}$. To produce a best continuous-space estimate of the target's location, a kernel density estimate (KDE) [94] is computed by imposing a likelihood kernel function $K(\cdot)$ around each sampled world position. A two dimensional Gaussian kernel with standard deviation of one grid cell (representing the kernel bandwidth, here designated h) was selected in the implementation. The updates to the grid can be computed incrementally for each frame, making the filter efficient. The final estimate from the filter is computed as a weighted mean, $\mathbf{x}_{\mathrm{est}}$, over the grid.

Input: N, M, $\mathbf{S}_{\text{UAV}} = \{\mathbf{s}_{UAV}\}$, $\mathbf{X}_{\text{tgt}}^{\text{img}} = \{\mathbf{x}_{\text{tgt}}^{\text{img}}\}$, $\mathbf{E} = \{\mathbf{e}\}$, h
Output: Best geolocation estimate $\mathbf{x}_{\text{est}}$

grid$[1..M][1..N] = 0$
foreach $\{s_{\text{UAV}}, \mathbf{x}_{\text{tgt}}^{\text{img}}\} \in \{S_{\text{UAV}}, \mathbf{X}_{\text{tgt}}^{\text{img}}\}$ **do**
    foreach $\mathbf{e} \in \mathbf{E}$ **do**
        $\mathbf{x}' = f_{\text{sample}}(\mathbf{s}_{\text{UAV}}, \mathbf{x}_{\text{tgt}}^{\text{img}}, \mathbf{e})$
        for $i = 1$ *to* M **do**
            for $j = 1$ *to* N **do**
                $\mathbf{x} = \text{grid2world}(i, j)$
                $\text{grid}[i][j] = \text{grid}[i][j] + K_h(\mathbf{x}^{(k)} - \mathbf{x}'^{(k)})$

$\xi = \sum_{i=1}^{N} \sum_{j=1}^{M}(\text{grid}[i][j])$ `// Compute normalization factor`
$i_{\text{est}} = \sum_{i=1}^{N} \sum_{j=1}^{M}(i \times \text{grid}[i][j]/\xi)$
$j_{\text{est}} = \sum_{i=1}^{N} \sum_{j=1}^{M}(j \times \text{grid}[i][j]/\xi)$
$\mathbf{x}_{\text{est}} = \text{grid2world}(i_{\text{est}}, j_{\text{est}})$

Algorithm 1: Formal description of the evidence grid geolocation filter.

Specification of the samples in the error distribution $\mathbf{E}$ depends on the vehicle uncertainty model. In the experiments performed here, a conservative heading uncertainty distribution is chosen as a uniform distribution over ± 45 degrees. The errors of the other pose dimensions (roll, pitch, and position) are less profound and are noted empirically to be closer to Gaussian distributions. The specific values chosen are zero mean Gaussians with 3σ bounds of 5 degrees for roll and pitch, and 7 meters for position.

For the results depicted here, 2000 samples for each observation are chosen, with a grid size of 500m by 500m and each cell measuring 5m by 5m. An example of a sampling from a single observation from a flight test can be seen in Figure 4.12. There is an obvious trade-off between increasing the number of samples or the number of cells to improve accuracy and maintaining a tractable computational burden.

Figure 4.13 provides an example of the filter in operation though a series of snapshots from an actual geolocation experiment in which a single UAV orbited a stationary target. Table 4.1 provides a comparison between this filter, a reasonably tuned EKF using an appropriate constant-position process model and matching uncertainty covariances, and a naive Cartesian mean of all observation coordinates. Lastly, Figure 4.14 provide two illustrative examples from separate geolocation experiments (one in still winds and the other in extremely stiff winds) showing the relative superiority of this strategy.

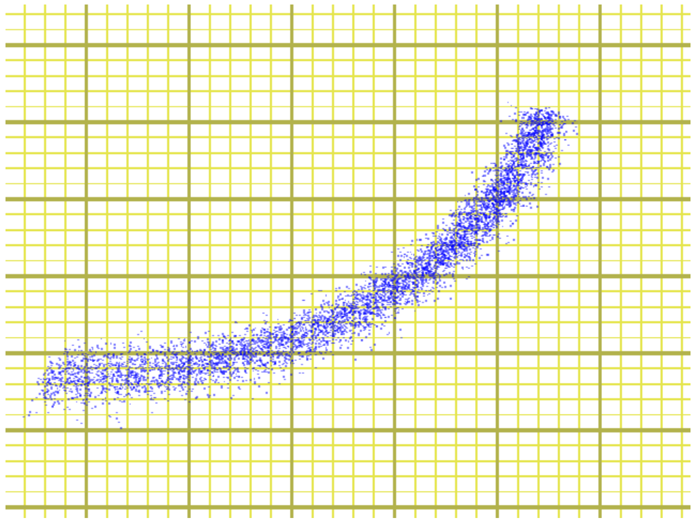

Figure 4.12: Illustration of a single observation, sampling the UAV uncertainty distribution and projecting geolocation estimates with each UAV state offset applied onto the ground plane. These are then accumulated over a series of observations. Scale: large grid cells are 50 × 50 meters.

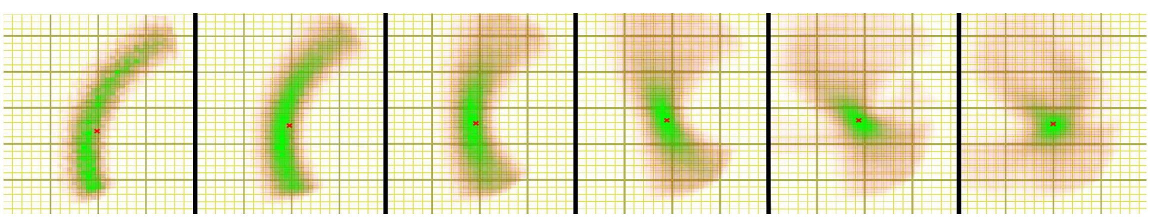

Figure 4.13: Visualization of the filter over a UAV's orbit of a target, in chronological order. The 2D evidence grid is represented as a grid of relative cell probabilities (the two axes are the North and East location of the object). Cells vary from bright green (high target probability) to red (low probability) to blank (none). Scale: large grid cells are 50×50 meters.

Flight #	Wind	Head Err	Alt Err	Mean	Kalman	Evidence-grid
1	High	Moderate	Moderate	17.5m	30.6m	10.6m
2	Moderate	High	Moderate	22.0m	28.1m	2.9m
3	Moderate	Very high	High	8.7m	15.3m	4.5m
4	Low	Moderate	Moderate	13.4m	19.7m	5.4m
5	High	Moderate	Moderate	29.6m	35.2m	8.0m
6	High	High	High	7.3m	15.4m	7.7m
7	Low	High	High	15.0m	22.2m	5.5m
Mean				16.2m	23.8m	6.4m

Table 4.1: Comparison of evidence grid filter against a reasonably tuned EKF and a naive Cartesian mean estimate over several flight tests

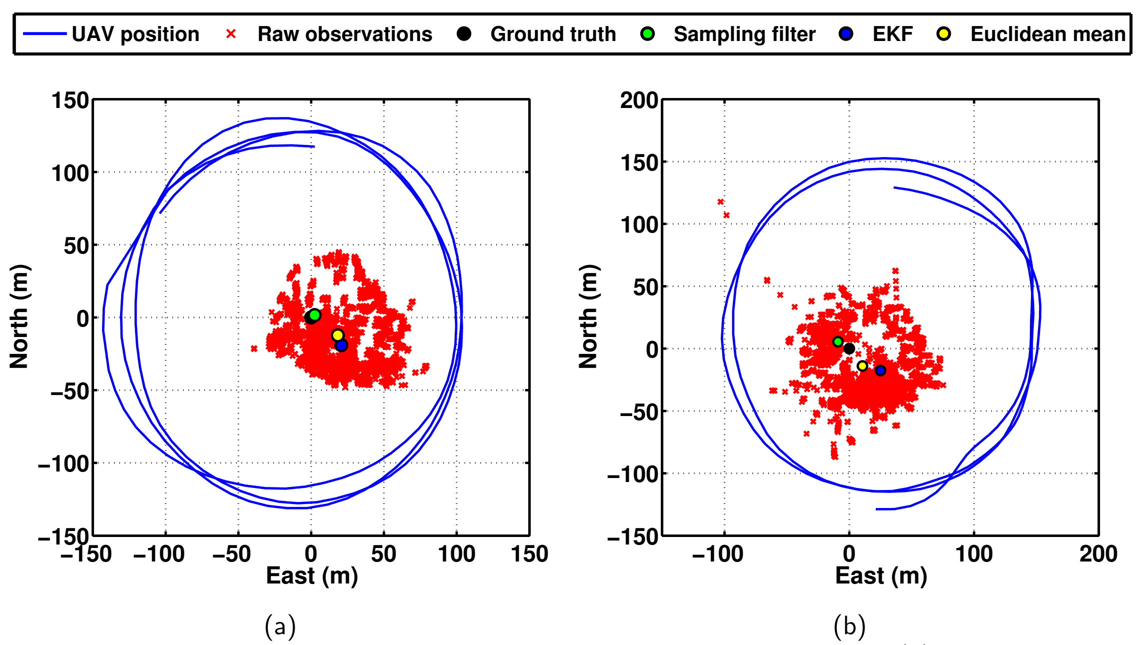

Figure 4.14: Plot of geolocation results from two test flights, under both (a) low to moderate wind and (b) high-wind conditions. The evidence-grid error-sampling filter outperforms an EKF and a naive Cartesian mean in both cases.

Both empirically and as the provided data suggest, this filter works very well in practice. Key advantages are that no assumptions need be made about the form of observation uncertainty, and that if the distribution is known, arbitrarily complex ones may be approximately represented. One disadvantage is that no clear quantitative confidence measure is available, and any chosen will necessarily be conservative given the additive nature of observation accumulation. Additionally, while its computational burden has not presented a problem in practice, greatly increasing the number of samples per observations or using a system having an extremely complicated f_{obs} function may become intractable. Finally, it is not easily extended to moving targets. Two obvious potential methods for such an extension include applying a diffusion kernel between observations and adding an additional dimension over discretized velocity vectors; both increase computational requirements substantially and are unlikely to perform well.

4.3.2 Range-bearing particle filter

Taking inspiration from such sampled representations of location uncertainty and the polar nature of a given observation's uncertainty similar to that present in range-only localization problems, [29] a particle filter taking observations parameterized as a polar range and bearing was also considered.

First, a particle filter essentially identical to that described in Section 3.3.2 using pseudo-observations is constructed. Then, given an observation $(s_{\mathrm{UAV}}, \mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}})$ and Gaussian uncertainties on UAV state having covariance Σ_{UAV}, $f_{\mathrm{obs}}(s_{\mathrm{UAV}}, \mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}})$ is applied as before to produce a location in world coordinates $\mathbf{x}_{\mathrm{tgt}}^{\mathrm{world}}$ (here, explicitly projected in 2D to a local ground-plane). However, rather than simply applying linearized Jacobian transformation of Σ_{UAV} to produce uncertainties Σ_{x} in the Cartesian observation location, it is first transformed into polar coordinates as

$$\begin{bmatrix} r \\ \beta \end{bmatrix} = f_{r\beta\ \mathrm{obs}}(s_{\mathrm{UAV}}, \mathbf{x}_{\mathrm{tgt}}^{\mathrm{img}}) = \begin{bmatrix} \sqrt{(x_{\mathrm{tgt}} - x_{\mathrm{UAV}})^2 + (y_{\mathrm{tgt}} - y_{\mathrm{UAV}})^2} \\ \mathrm{atan2}(y_{\mathrm{tgt}} - y_{\mathrm{UAV}}, x_{\mathrm{tgt}} - x_{\mathrm{UAV}}) \end{bmatrix}, \tag{4.4}$$

where $[x_{\mathrm{tgt}}, y_{\mathrm{tgt}}]^T = \mathbf{x}_{\mathrm{tgt}}^{\mathrm{world}}$ and $[x_{\mathrm{UAV}}, y_{\mathrm{UAV}}]^T$ is the projection of $\mathbf{x}_{\mathrm{UAV}}^{\mathrm{world}}$ onto the same local ground-plane. Then, a new linearized Jacobian transformation of UAV state uncertainty may be defined as

$$\Sigma_{r\beta} = \frac{\partial f_{r\beta\ \mathrm{obs}}}{\partial \mathbf{s}_{\mathrm{UAV}}} \Sigma_{\mathrm{UAV}} \frac{\partial f_{r\beta\ \mathrm{obs}}}{\partial \mathbf{s}_{\mathrm{UAV}}}^T. \tag{4.5}$$

Despite still being a linearization of the observation function, the range-bearing Gaussian approximation using this covariance can better model the shape of the true observation uncertainty than either Cartesian Jacobian linearization or the Unscented Transform, as Figure 4.15 indicates.

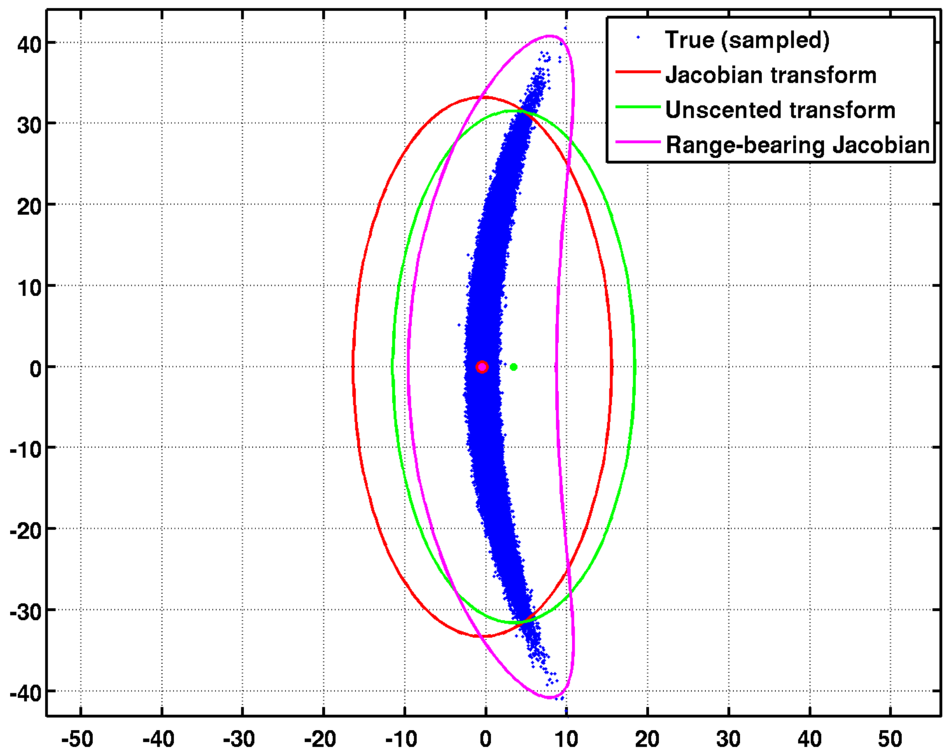

Figure 4.15: Comparison of uncertainty representation approximations compared with a monte-carlo sampling of the true distribution for a single observation with typical parameters (axis scale is meters), similar to Figure 4.11. Of these, a Gaussian uncertainty distribution represented in range-bearing space best matches the shape of the true distribution. The Unscented Transform in Cartesian space better captures the sample mean but remains constrained to a Gaussian.

Completing the implementation of the particle filter, as pseudo-observations are received as a polar position and uncertainty relative to the current UAV position, each particle receives a likelihood update using its likelihood under the observation Gaussian distribution as before except that now each particle is momentarily re-parameterized in polar coordinates using the latter half of Equation 4.4 for the evaluation.

A series of snapshots of this filter during pursuit of a target moving along a road is shown in Figure 4.16

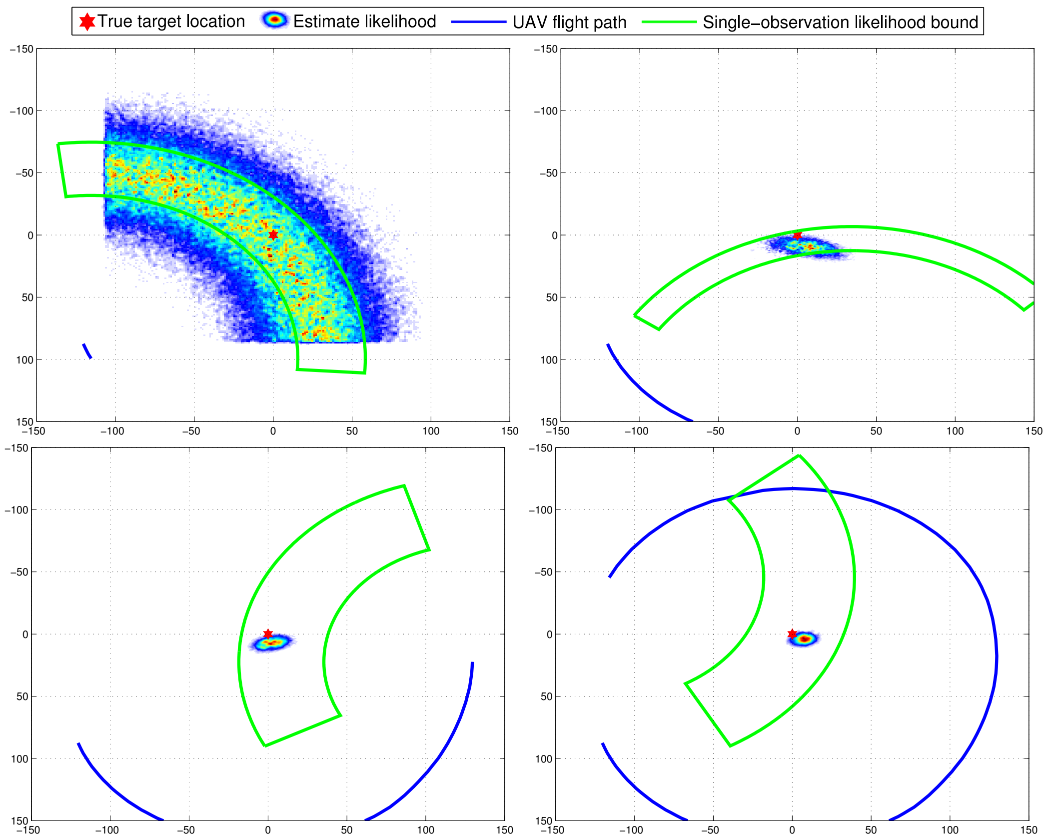

Figure 4.16: Snapshots during the execution of the range-bearing space particle filter during geolocation of a target. In this instance, observations are provided as Gaussian in range and uniformly distributed in bearing around the point observation. The estimate likelihood is here represented by a relative particle density. After approximately half an orbit, the filter converges to an estimate with approximately 12m error.

In practice, this filter performed approximately as well as the traditional form of a particle filter of Section 3.3.2 (in which f_{obs}^{-1} is linearized around each particle and the particle projected into image coordinates to compute a likelihood for each particle), but with much less computation, instead requiring only a single evaluation and linearization of f_{obs}. Thus, realtime execution with a typical particle set size of 5000 presented no difficulty.

While formal comparison against existing and the other developed filters remains to be performed, several conclusions may be made. First, though accuracy and validity of the confidence measure is better for moving targets than an EKF, this filter is still not nearly as accurate as the evidence-grid filter applied to stationary targets. Additionally, both severe sensitivity to the accuracy of the target process model (as slow but discontinuous drift in observation error may result in overconfident clustering at an incorrect location) and a strong need for delicate outlier rejection (as inlier observations may themselves be highly erroneous) were observed. Together, these imply a need for large numbers of particles (though still a tractable number) and careful tuning varying between target types. Finally, though again tractable on portable ground-station computing hardware, the substantial level of computation required exceeds what is practical to embed aboard aircraft in the foreseeable future. Likewise, as a representation comprised of a large set of hypotheses, the same difficulties performing observation-predictive planning as noted in Section 3.3.3 apply in full force.

4.3.3 Over-parameterized bounded filter

Given these ongoing weaknesses, yet another form of geolocation filter is motivated, with several properties standing out. First, while thus far the general crescent shape of observation uncertainty distributions can be approximately modeled parametrically as a range and bearing, the a-posteriori distribution does not retain this shape and has only been sampled, maintaining a desire for a filter with an exclusively analytical representation. Additionally, given the intended use in cooperating teams and as a continuous estimate to seed observation-predictive trajectory planning, a representation that is compact to transmit and rapid both to apply observation updates and to evaluate observation likelihoods is needed and would conveniently be provided by a small parametric representation. Separately, given that large, time-varying error is not well modeled by any particular probability distribution (as a bias may dwell near consistent values), a representation instead considering an error *bound* within which offsets must lie may prove more appropriate.

Therefore, a filter similar to one proposed for range-only SLAM [100] was applied, [45] combining two ideas to produce an analytical representation of observation and a-posteriori target location uncertainty in a recursive filter framework with only three state parameters.

First, a bounded recursive filter proposed by Schweppe [93] is selected, in which uncertain target location, target process model noise, and observation error are all represented by an individual bounded ellipsoidal set

$$\tilde{\mathbf{\Sigma}} = \left\{ \mathbf{x} : [\mathbf{x} - \hat{\mathbf{x}}]^T \mathbf{\Sigma}^{-1} [\mathbf{x} - \hat{\mathbf{x}}] \leq 1 \right\}, \tag{4.6}$$

where $\mathbf{x}$ is a true state, $\hat{\mathbf{x}}$ is an estimate, and $\mathbf{\Sigma}^{-1}$ is a positive semi-definite matrix analogous to the covariance of a Gaussian distribution.

Then, for a stochastic linear system using such bounds, such as

$$\mathbf{x}_{k+1} = \Phi\mathbf{x}_k + \mathbf{G}\mathbf{u} \tag{4.7}$$

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}, \tag{4.8}$$

where $\mathbf{\Phi}$ and $\mathbf{H}$ are transition and observation models and u and v represent process and observation noise having ellipsoidal bounds $\mathbf{Q}^{-1}$ and $\mathbf{R}^{-1}$ respectively, a recursive filter similar to the Kalman filter is defined as follows.

Predict:

$$\hat{\mathbf{x}}_{k+1}^{-} = \mathbf{\Phi}\hat{\mathbf{x}}_k \tag{4.9}$$

$$\mathbf{\Sigma}_{k+1}^{-} = \alpha_k[\mathbf{\Phi}\mathbf{\Sigma}_k\mathbf{\Phi}^T + \mathbf{Q}] \tag{4.10}$$

where $\alpha_k \in [1, 2]$.

Update:

$$\begin{aligned} \hat{\mathbf{x}}_k = {} & \hat{\mathbf{x}}_k^{-} + \rho_k[(\mathbf{\Sigma}_k^{-})^{-1} \\ & +\rho_k\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}]^{-1}\mathbf{H}^T\mathbf{R}^{-1}[\mathbf{z}_k \\ & -\mathbf{H}\hat{\mathbf{x}}_k^{-}] \end{aligned} \tag{4.11}$$

$$\mathbf{\Sigma}_k = \eta_k[(\mathbf{\Sigma}_k^{-})^{-1} + \rho_k\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}]^{-1} \tag{4.12}$$

$$\begin{aligned} \eta_k = {} & 1 + \rho_k - [\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^{-}]^T[\rho_k^{-1}\mathbf{R} \\ & +\mathbf{H}\mathbf{\Sigma}_k^{-}\mathbf{H}^T]^{-1}[\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_k^{-}] \end{aligned} \tag{4.13}$$

for non-negative ρ_k.

Conceptually these equations use set summation and intersection to perform estimation. The key algebraic difference to the Kalman filter is two scalar parameters α_k and ρ_k that adjust the set volume produced by sum and intersection operations. Known guaranteed constant values ($\alpha_k = 2$ and $\rho_k = 0.5$) prove too conservative in practice, so in implementation a scalar convex optimization is performed at each time step to find the bounded estimate with minimum set volume, requiring only several optimization iterations. Observations that do not intersect the predicted estimate are detected and rejected as outliers.

Next, the use of a higher-dimensional state space embedding is used as demonstrated by Hanebeck, [15] in which the 2D target geolocation estimate is mapped as

$$\mathbf{x}_{\text{tgt}}^{\text{world}} = \begin{bmatrix} x \\ y \end{bmatrix} \quad \rightarrow \quad \mathbf{x}^{\star} = \begin{bmatrix} x \\ y \\ x^2 + y^2 \end{bmatrix}. \tag{4.14}$$

Interestingly, a wide range of non-linear sensing problems resulting in non-linear, non-convex, and non-simply-connected estimate sets may be expressed in this space. Specifically relevant for this problem is the ease of natively representing crescent-shaped polar uncertainties.

Applying this to the geolocation estimation problem, a range-bearing observation is taken as in Equation 4.4, and $\Sigma_{r\beta}$ from Equation 4.5 is decomposed for the following equations as

$$\mathbf{\Sigma}_{r\beta} = \begin{bmatrix} \sigma_r^2 & \sigma_r\sigma_\beta \\ \sigma_r\sigma_\beta & \sigma_\beta^2 \end{bmatrix}. \tag{4.15}$$

Then, for each range-bearing observation, an approximate observation in the over-parameterized space is constructed as in [100] by combining annular and cylindrical sets that capture range and bearing bounds (respectively) as follows and illustrated in Figure 4.17.

Range Bound:

$$\begin{aligned} \mathbf{z}_{range}^{\star} &= r^2 - (x_{\text{uav}}^2 + y_{\text{uav}}^2 - \sigma_r^2) \\ \mathbf{H}_{range}^{\star} &= [-2x_{\text{uav}} \ -2x_{\text{uav}} \ 1] \\ \mathbf{R}_{range}^{\star} &= (2r\sigma_r)^2 \end{aligned} \tag{4.16}$$

Bearing Bound:

$$\begin{aligned} \mathbf{z}_{circle} &= \begin{bmatrix} x_{\text{uav}} \\ y_{\text{uav}} \end{bmatrix} + r \begin{bmatrix} \cos\beta \\ \sin\beta \end{bmatrix} e^{-\sigma_\beta^2/2} \\ \mathbf{R}_{circle} &= (2r\sin(\sigma_\beta/2))^2 \mathbf{I}_{2x2} \end{aligned} \tag{4.17}$$

Example snapshots from a typical execution of this filter are provided in Figure 4.18, showing the rapid convergence to an accurate estimate with small error bound. While here too thorough comparison against these other developed filters

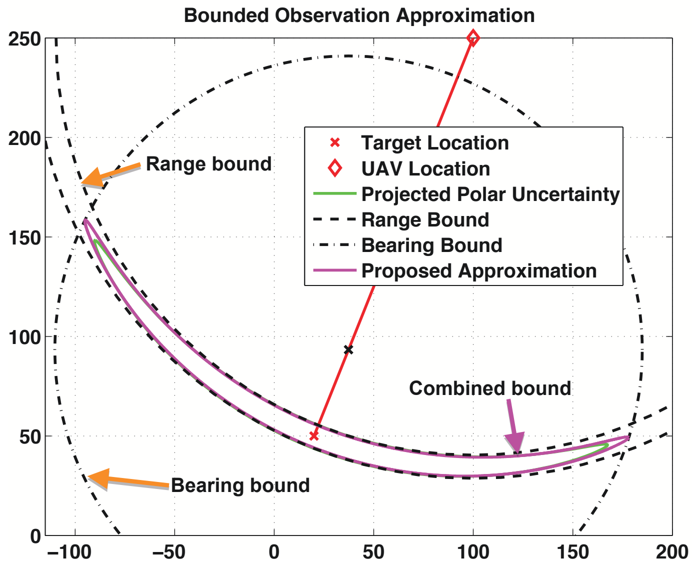

Figure 4.17: Construction of an approximate range-bearing observation uncertainty bound in the over-parameterized state space. Note the close match between Gaussian polar uncertainty and the combined approximate bound.

and existing methods remains to be performed, a number of preliminary conclusions may be made.

Advantages this filter provides include rapid estimate convergence from even extremely high-uncertainty observations given its set-intersect nature, no need to assume any particular error distribution in vehicle state, and a compact representation with only three state parameters and associated ellipsoid bound matrix need be stored or broadcast to convey estimate state. The small state space and linear recursive filter form result in rapid prediction and observation updates, while estimate set membership testing (analogous to target location likelihood evaluation) is also extremely fast. This greatly simplifies and improves the tractability of decentralized information-sharing and observation-predictive planning in which hypothetical filter branching and target likelihood evaluations are frequently performed.

Downsides include the absolute criticality of the validity of uncertainty bounds on observations given the use of an intersection operation, as the integration of an invalid measurement will result in rapid, if not immediate, filter failure. Careful outlier detection and rejection is therefore vital. The use of highly conservative error bounds for each observation improves robustness but provides a tradeoff against convergence rate and steady-state confidence. Additionally, where needed, representing the set

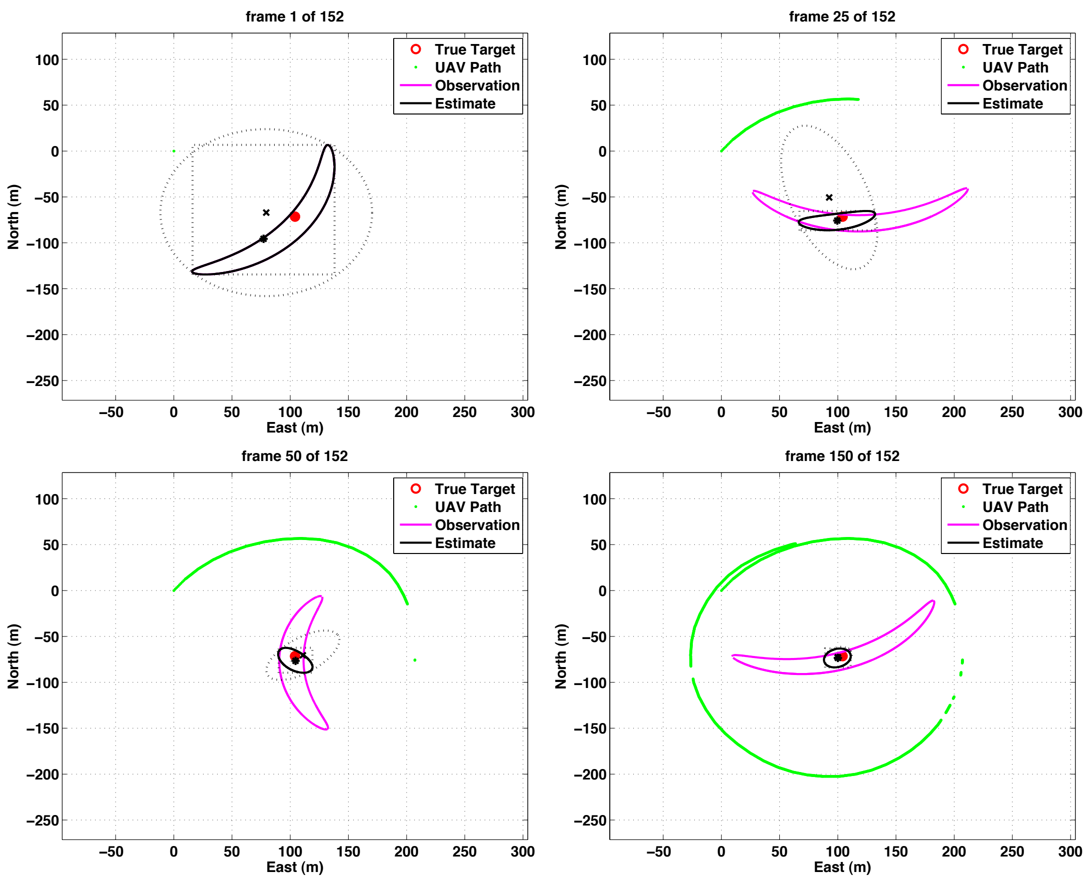

Figure 4.18: Sequence of snapshots during geolocation of a stationary target using the over-parameterized bounded filter. The dotted ellipse is a fast, conservative approximation of filter confidence; the dotted box is a more computationally intensive but tighter approximate confidence measure. The black crescent is a level-set of the true filter confidence estimate.

of 2D world locations lying within the higher-dimensional ellipsoidal bound is not completely straightforward. Membership testing (Equation 4.6) is fast to compute, and implicit level curves of the ellipsoidal projection suitable for visualization are conveniently a quartic polynomial having an analytic solution. For parametric representation, two options include simple geometric projection of the ellipsoid to two dimensions, providing a very fast but (possibly uselessly) over-conservative bound, and computing a smallest axis-aligned bounding box containing all components of the estimate set, which though substantially slower to compute provides a much tighter approximation. Examples of these options are depicted in Figure 4.18. Finally, a potentially serious limitation is that in order to retain the linearity of the model, only stationary or constant-velocity targets are representable. It remains an unanswered question as to whether EKF-style linearization that would enable arbitrary target process models retains numerical stability and acceptable levels of performance.

4.4 Improved Pursuit for Constrained Targets

The proposed methods for improved geolocation reduce estimate error and well model the nature of observation uncertainty found in small-UAV geolocation, but as they may still require a large number of observations to produce such an estimate, inadequacy for pursuit purposes remains. Further, given inevitable long-duration dropouts contemplated in Section 3.3.3 as a UAV circles back to an overflown target and limited onboard computation with which to consider best trajectories for possible target locations over a wide area, higher-certainty target prediction from a previously estimated location is needed. Together, these motivate a desire to limit both the space of locations in which target presence need be considered and the set of locations within this space to which a target can move at a given time. The means proposed in this thesis for accomplishing this is the use of pre-existing knowledge of environmental constraints producing such simplification.

4.4.1 Single-road pursuit

A simple but instructive and still highly relevant case to first consider is that of a target known to be moving along a single road segment. Given an existing map indicating the path of the road, the geolocation problem is now reduced to one dimension, regardless of the (possibly arbitrarily winding) actual path of the road. For the purpose of this exploration, a road represented as a contiguous sequence of linear line segments is assumed, as well as some means to choose the correct line segment given an observation (such as choosing the nearest one to the raw observation). Then, the problem is further reduced to the case of a series of raw observations and a single line segment.

As noted in Section 3.3.2, road-constrained geolocation has been previously considered extensively in tracking with high-altitude radar, and an enumeration of general benefits and distinctions from this context is warranted. The primary known benefit is simply greater geolocation accuracy, stemming from two sources: better accuracy from filtering of observations and improved prediction in the absence of observations. The first is due largely to the fact that relative to the surrounding open area, the set of possible target locations is much smaller (being one rather than two dimensional, the number of samples required for a uniform discrete sampling grows linearly rather than quadratically with the size of the region), such that even a completely uninformed uncertainty distribution occupies a physically small area. Better filter accuracy is also due to more accurate raw observations, which if also constrained to one dimension, likewise eliminate an entire uncertainty dimension. Meanwhile, the second source of overall geolocation accuracy–improved prediction–results from a similar dimensionality reduction in the space of target actions or destinations, in discrete form again having linear rather than quadratic growth with region size. The practical impact of such improved prediction is that reduced observation frequency (and, for moving targets, fewer additional observations more than stationary targets) is required.

In the context of pursuit with field-reconnaissance UAVs, however, road-constrained geolocation provides additional practical benefits that have not been considered in the aforementioned constrained estimation literature, primarily due to the mobile nature and limited field of view of sensors in this context. Clearly, overall improvement in geolocation accuracy fulfills an immediate CSAT mission goal. Given freedom in sensor placement not typically found in radar sensing applications, specific advantage of the target state space reduction may be taken so that the larger dimension of asymmetric observation uncertainty (as from large heading uncertainty) is aligned with the perpendicular axis of the road as frequently as possible and thereby eliminated, drastically improving raw observation quality. This is illustrated in Figure 4.19. Relatedly, the combined effect of improved estimate accuracy and improved predictiveness implies better robustness to observation dropouts during reacquisition of overflown targets, since the uncertainty region for reacquisition will be small. Taken together, all of these benefits improve task performance, but for a key reason specific to mobile sensing: the characteristic of having essentially a positive feedback loop on target location certainty, as a high-confidence location provides confident future sensor placement, whereas an uncertain location risks poor future placement. In terms of observable impact, such performance improvement translates to more accurate geolocation estimates and reduced frequency of target loss.

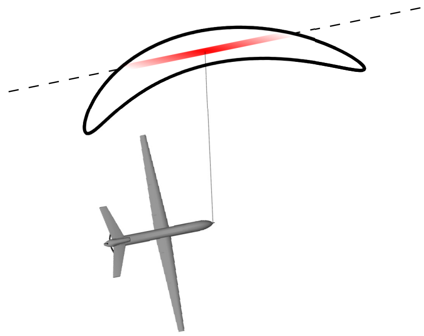

(a) Road-parallel observation

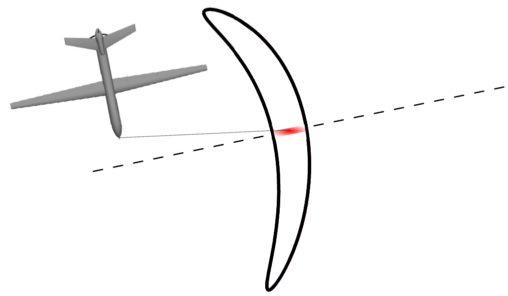

(b) Road-perpendicular observation

Figure 4.19: Illustration of relative value of placing a sensor footprint with asymmetric observation uncertainty (here, due to large heading uncertainty) parallel (a) or perpendicular (b) to the road vector. A likelihood bound (Gaussian level curve) for each observation is depicted as a black crescent, and one-dimensional projected observation uncertainty is shown in red.

In addition to performance, constrained geolocation also provides a drastic improvement in scalability. As both the spaces of target location and target actions are reduced, a radically simpler observation-predicting trajectory planning problem results through a reduced integration volume for Equation 3.6. Practically, when evaluating the value of placing the sensor footprint at a given position in the future (a single term of this integral), a much smaller set of possible observations need be tested or branched on. Indeed, if sampling a discrete set of possible observations, once again a number linear rather than quadratic in the size of the target region is needed. Additionally, reduced filter state dimension necessarily results in reduced memory and computational requirements when cloning current filter state and applying hypothetical target predictions and observation updates during UAV action search for planning. Altogether, such scalability improvement permits longer-horizon trajectory planning and more tractable increases in UAV team size.

To demonstrate and evaluate these benefits, a simulation study was performed [27] comparing the behavior of several filters and trajectory planning. First, given a road segment characterized by a line segment $\mathbf{p}_{\text{road}} + \hat{\mathbf{v}}_{\text{road}}t$ (for a 2D point $\mathbf{p}_{\text{road}}$ and unit vector $\hat{\mathbf{v}}_{\text{road}}$), a raw pseudo-observation $\mathbf{x}_{\text{tgt}}^{\text{world}}$ having Gaussian uncertainty with covariance $\Sigma_{\text{tgt}}^{\text{world}}$ is projected to the road as a new scalar pseudo-observation

$$x = (\mathbf{x}_{\text{tgt}}^{\text{world}} - p_{\text{road}}) \cdot \hat{\mathbf{v}}_{\text{road}} \tag{4.18}$$

having Gaussian uncertainty with scalar variance

$$\sigma^2 = \hat{\mathbf{v}}'_{\text{road}} \Sigma_{\text{tgt}}^{\text{world}} \hat{\mathbf{v}}_{\text{road}}. \tag{4.19}$$

Generation of the scalar observation is merely the geometric projection of a point onto a line, while variance is derived from Jacobian covariance transformation of the projection function and is exact because the projection is a linear function. Filter modification to incorporate the road constraint depends on its internal representation: for Kalman filters maintaining a Gaussian estimate, only a 1D Gaussian is used; grid representations are modified to only provide cells along the length of the road; and particle filters sample points only along the road's line segment. Trajectory planning for pursuit may be performed by any existing means, except that where possible observations need be enumerated or sampled, far fewer are needed. In experimentation here, sampling 1D Gaussians with 7 points proved adequate.

In the simulation scenario considered, the task is defined as pursuing a vehicle moving along a 2.5km long roadway using a single UAV. The target's nominal speed is 4.5m/s (10mph), and to roughly simulate evasiveness or confusion, it accelerates at up to 5.4m/s^2 (equivalent to accelerating from rest to 60mph in 5s), within the range 3.4 to 6.8m/s (5-15mph). This is roughly as fast of a target that a UAV traveling at 15m/s (33.6mph) can pursue while retaining any substantial maneuverability. Coordinated turns are enforced, and control commands received by the UAV are low-pass filtered with a cutoff frequency of 2Hz to somewhat simulate an autopilot optimizing for smooth flight. Uncertainties are modeled as uncoupled Gaussian uncertainties with standard deviations of 3m in lateral position, 5m in altitude, 3 degrees for roll and pitch, and 5 degrees in yaw. Camera parameters received uncertainties of several percent. For simulation, errors in measured values were drawn from these distributions. For consistent comparison, a trial is defined as 5 minutes of simulated elapsed time, during which time the target has moved approximately two thirds of the length of the roadway. Trials are stopped if a filter loses tracking (diverges or reports an incorrect estimate resulting in an irrecoverable trajectory) and logged as such. Each result presented is the average over several hundred trials with randomized initial conditions that place the target within the UAV's view, as though it had just detected the target during a search. An example visualization from the simulation is shown in Figure 4.20.

Figure 4.20: Visualization example from simulation showing a UAV pursuing a Jeep. The green trapezoid is the ground region visible to the camera. The orange dot indicates the geolocation filter's best estimate of the target's position at the center of the red smear denoting the surrounding uncertainty distribution.

As a baseline comparison, four unconstrained filters were implemented: a constant-velocity model Kalman filter using pseudo-observations with uncertainty generated by Jacobian linearization, a constant-velocity model Kalman filter using pseudo-observations with uncertainty generated by the Unscented Transform, a fixed-cell (1m in size) discrete Bayesian filter using a diffusing constant-position target model and pseudo-observations with uncertainty generated by the Unscented Transform, and the particle filter described in Section 4.3.2 using a constant-velocity target model and range-bearing pseudo-observations and corresponding Gaussian uncertainty. For the first set of experiments, the very simple control law of directing the UAV to orbit the filter's best estimate is used. Results summarized in Table 4.2 suggest first that a better strategy is sorely needed given the high rate of target loss (avoided by the Bayes filter given its large cells and slow diffusion). As expected, filtering using the Unscented Transform performs slightly better, and the Bayes grid filter performs slightly worse due to its low spatial resolution. The particle filter performed particularly well when it worked, but it suffered frequent particle collapse despite stratified resampling and outlier rejection (that may be improved upon by greatly increasing the particle set size and unreasonable scenario-specific process model tuning).

The same set of filters were again compared, this time incorporating the road constraint, and once again applying the simple control law of orbiting the best estimate. Results summarized in Table 4.3 show marked improvement in target loss rate and moderate improvement in accuracy. The particle filter continues to exhibit substantial fragility.

Finally, the experiment was repeated yet again now applying observation-optimizing

	KF + JT	KF + UT	Discrete Bayes	PF + RB JT
Trials losing target	27.0%	20.5%	0%	65.0%
Time in view	98.1%	97.5%	96.5%	98.3%
Avg view loss duration	1.8s	2.4s	1.6s	1.3s
Avg target error	10.5m	9.8m	15.5m	8.6m

Table 4.2: Simulation results for baseline filter comparison while orbiting the mean target location reported by each estimator.

	KF + JT	KF + UT	Discrete Bayes	PF + RB JT
Trials losing target	10.9%	9.5%	0%	20.5%
Time in view	97.4%	97.2%	95.9%	97.7%
Avg view loss duration	3.0s	3.3s	0.9s	2.1s
Avg target error	7.5m	7.4m	15.1m	7.9m

Table 4.3: Simulation results for road-constrained filter comparison while orbiting the mean target location reported by each estimator.

trajectory planning. For this, only the Kalman filter using uncertainty generated by the Unscented Transform was selected, as the overall best performing filter in the preceding experiments. Here, two optimization metrics were evaluated: maximizing the average probability of viewing the target over the course of the trajectory, and minimizing a-posteriori filter uncertainty. For simplicity, trajectory optimization was performed by brute-force breadth-first search in the action space of steering angles (held for 2s between action branches). For both the depth-1 (greedy) and depth-3 optimizations, planning in realtime provided no difficulty. Results summarized in Table 4.4 expose several interesting observations. First, unsurprisingly, greater look-ahead improves performance, however it does not do so by an immense margin for a moderate increase in depth, implying that given for instance a constrained onboard processor, it may be reasonable to use single-step look-ahead for at least this specific task. This is plausible because the growth in target location uncertainty during long-term look-ahead produces an uncertainty distribution that is flatter, and hence less informative. Particularly interestingly, one may observe that optimizing for even short-term minimum uncertainty produces better performance (primarily a much lower target loss proportion) than optimizing for more likely observations out to a greater horizon. This highlights the importance and power of accurate target motion prediction, as reduction of uncertainty (and hence, presumably error) in target velocity greatly aids prediction when it is not being viewed, ever more importantly for a target moving at a relatively consistent speed for even brief periods such as present in this scenario.

	Min Posterior Uncertainty		Max Viewing Likelihood	
	Depth 1	Depth 3	Depth 1	Depth 3
Trials losing target	2.5%	0.8%	15.0%	4.5%
Time in view	80.9%	88.4%	77.1%	86.6%
Avg view loss duration	3.6s	1.4s	1.6s	1.0s
Avg target error	7.3m	6.6m	12.1m	9.7m

Table 4.4: Simulation results for the road-constrained Kalman filter taking observations with Unscented Transform generated uncertainty, comparing search depths and metrics for trajectory planning.

As a useful direct illustration of the impact of the application of the road constraint, Figure 4.21 compares the overall trajectory of a UAV performing depth-3, posterior uncertainty minimizing planning with the constraint (as in the last simulation experiment) and without (using two-dimensional Gaussian filter uncertainty) it. The latter exhibits moderate elongation of portions perpendicular to the road, indicating a larger fraction of the total trajectory is spent collecting observations benefiting from the constraint and providing greater certainty.

4.4.2 Naive road-network decomposition

Simply extending road segment pursuit to more complex environments, such as a winding road with segments looping near one another or networks of roads as would likely be found in a common urban environment, is not practically straightforward given that the assumption of having a means to correctly select the true segment on which the target lies is unrealistic if ambiguity is presented by observation uncertainty overlapping multiple nearby segments. Additionally, if junctions providing multiple possible directions for a target to move are introduced, multi-modal prediction uncertainty results.

One possible means to address this is to use a cellular Bayesian estimator (used as a compared filter in the preceding experiments) and plan trajectories maximizing the average probability of observation, essentially executing a continuous local search as proposed by Ryan, et al. [89] (practical implementation of which is discussed further by Tisdale, et al. [103]) and previously referred to in Section 3.3. If target motion between cells is predicted during planning, the result is pursuit with the goal of maximizing target observation frequency. An alternative metric that may be applied

(a)

(b)

Figure 4.21: Comparison of a typical overall UAV trajectory having pursued a moving vehicle as in the preceding experiments without (a) and with (b) application of the road constraint, each performing trajectory planning to minimize expected short-term a-posteriori geolocation filter uncertainty. Although superficially quite similar, the trajectory using the road constraint makes approximately 22% fewer orbits in the same interval and exhibits approximately 20% greater elongation of each orbit (corresponding, very roughly, to that proportion greater viewing time from poses perpendicular to the road). Intuitively, this is due to attempts to acquire low-uncertainty observations, which as depicted in Figure 4.19 are more likely as the UAV moves along the perpendicular direction to the road, but which appear to be equally possible from any viewing direction if knowledge of the road is unavailable.

is to minimize expected a-posteriori uncertainty defined using the discrete weighted sample covariance over all cells.

Figure 4.22 provides a simulation example of live pursuit using such a cellular representation. Here, probability of observation is maximized using single-step (greedy) steering direction control, Gaussian diffusion informed by an assumed maximum speed is applied as prediction between observations, and for simulation the target executes a random policy at junctions, otherwise accelerating randomly as in the preceding experiments. In this example, a target moving along a straight segment slows briefly, requiring that the UAV turn to prevent overflying it. In the process, the target continues moving and leaves the field of view. Until its reacquisition, the filter smoothly tracks diffusion into adjacent road segments and incorporates negative observations as otherwise-likely locations are found to not contain the target. Upon detecting the target, the probability of cells rapidly re-aligns to a small uncertainty region around the true location.

Figure 4.22: An example of pursuit in a road network using cellular decomposition. The UAV anticipates loss of view (area within the green trapezoid) and begins to turn sharply. While out of view, it predicts diffusion in possible target location (red probability mass) and greedily steers to maximize the probability of re-detection.

To a large extent, then, such a cellular decomposition provides a total solution to pursuit in complex environments. Most importantly, it naturally handles multimodality that results from an observation ambiguously lying across multiple road segments–cells along each simply receive a high observation likelihood–as well as that resulting from predicting target motion near a junction–adjacent cells lying on connecting road segments will received increase likelihood through diffusion. It also naturally presents a solution to the subtasks of recapture (and, if required, search) since if a target is not re-acquired before it could have escaped into several distinct regions that cannot be swept in a single continuous motion, it provides a means to track remaining likely locations as each is visited.

At the same time, cellular decomposition is completely impractical for optimizing all but very coarse approximations of Equation 3.6. A large number of cells (each constituting a hypothesis) produces a large number of possible observations over which to take expectation or branch for planning. Likewise, the complexity of motion prediction and observation updates grows at least linearly with the number of cells. Further, a large number of cells *is* required, for several reasons: large cells provide poor geolocation resolution, a cell size at least on the order of possible inter-timestep target motion distance is necessary for meaningful discrete-time approximation of its motion, and additional dimensions for target speed or direction are needed for stateful estimation beyond unrealistic random motion. As a result, optimization by branching on hypothetical observations entails a large branching factor, substantial computation for applying these observations and ensuing target motion prediction, and copying a large amount of data to clone filter state. Worse, the typical simplification of assuming only negative observations will be received is inapplicable to an active pursuit scenario. Therefore, drastic measures such as assuming a stationary target probability distribution or planning to very short horizons are necessary. As an example, the simulation of Figure 4.22 utilized both these measures (greedy planning assuming a stationary target) and still provided considerably slower than realtime execution.

To overcome this, the following chapter proposes building upon the same existing constrained and multi-modal estimation literature from which the preceding road segment pursuit representation derives inspiration to cast complex road network pursuit as pursuit over a discrete collection of single road segments. In doing so, it is anticipated that similar performance benefits may be attained while motivating further exploitation of fundamental environment topology.

4.5 Conclusions to Date

A number of conclusions may be drawn from both study of related work and field and simulation experience thus far. First, with appropriate filtering, even highly uncertain observations may be used to derive an accurate instantaneous target location estimate, but many may be required to do so. Rapidly changing target state (as for agile moving targets) or estimation of higher-order states such as velocity (as observations provide only direct observability in location space) worsen this considerably, resulting in for instance particularly poor accuracy for freely moving targets in open areas. Meanwhile, observation-predictive trajectory planning is hindered where many possible observations exist, due either to high-dimensional target uncertainty, uncertainty simply spread over a wide area, or unconstrained target models permitting rapid divergence in target location. The practical effect of this is immense computational complexity when optimizing trajectories for expected observations. For target pursuit, limited UAV maneuverability (such as a long turn-around time) implies a need for many UAVs or highly accurate predictions to prevent target loss, as without these, the result is poor robustness to periodic observation drop-outs of an individual UAV. For search having the goal of guaranteed capture, large environments or fast targets result in rapid recontamination requiring many UAVs, imposing high resource demands and poor scalability. As a more positive observation, the use of environment structure to provide additional information about possible target position, motion, or capability has proven fruitful, motivating its use wherever possible. Proposed work attempts to specifically overcome the aforementioned limitations in limited but common scenarios providing such structure.

CHAPTER 5

Proposed Approach: Improved Search and Tracking by Exploiting Environmental Structure

Although considerable practical performance improvements may be achieved by careful representation of target observation uncertainty and accounting for limitations in observer motion, search and tracking with field-reconnaissance UAVs remains a fragile, computationally-demanding task requiring many vehicles to guarantee search success and persistent pursuit. This chapter proposes exploring the use of environmental structure to alleviate many of these difficulties in limited but common cases where this is possible. Section 5.1 motivates the case of roads and road networks specifically, as both a realistic environment and one lending itself to particularly fruitful representation as a graph. Common graph-based search and pursuit abstractions are briefly summarized in Section 5.2. Possible solutions that are more tractable and better-performing resulting from application of these abstractions to pursuit (Section 5.3), recapture (Section 5.4), and search (Section 5.5) are then considered. From there, Section 5.6 suggests conceptual generalizations to mixed environments containing both roads and open areas. Finally, overall benefits of applying these abstractions are summarized in Section 5.7, and Section 5.8 proposes continued field and simulation experimentation to evaluate these ideas.

5.1 Motivation for Road-network Constraints

If an area of interest contains primarily roads, such as in an urban area or as a series of roads through otherwise untraversable terrain, the structure or topology of these roads is particularly informative for a search or pursuit task. The three key elements of such information include the restriction of a target's possible location to a small subset of the environment (assuming a target must lie on a road), the imposition of strong continuity (and hence short term predictability) on its motion during the typical case of motion along a road segment, and limitation of the branching factor in target paths to splitting at junctions (or, less probably, U-turns).

Such constraints are immensely useful to UAV observers by directly easing many of the task challenges enumerated in the preceding chapter. As explored therein, sharp constraints as provided by roads greatly mitigate large uncertainty in individual observations, improved predictiveness as provided by limited target paths better survives frequent observation dropouts and slow returns to view, and fewer target location hypotheses and lower-dimensional uncertainty distributions reduces both the set of possible observations to evaluate for predictive planning and the computational and memory burden of running or branching filters based on hypothetical observations.

At the same time, while superficially representing a very small segment of the space of possible environments, road networks constitute a realistic and worthwhile case to consider. Primarily, as the developed proportion of land area worldwide continues to expand, any rescue, surveillance, perimeter protection, or suspect pursuit mission is likely to take place in or near a population center. Indeed, already by 2007 half the world's population lived in urban areas, and this is expected to reach 70% by 2050. [109] This presents an immediate practical demand. Additionally, the tracking and guidance community, previously noted as having extensively studied constrained and multi-modal geolocation on road networks using high-altitude radar, has long recognized the value of incorporating such terrain knowledge in heavily fielded systems and thus provides precedence. Finally, it is certainly the case that for many scenarios, operation over an equivalently-sized open area may simply be infeasible with a given number of vehicles and a given level of computation, so a practical solution to these strictly stretches the operating envelope of such systems.

5.2 Graph-based Search and Pursuit Abstractions

As a collection of road segments (representable by a set E) having some number of intersections (likewise representable as a set V), a road network may be represented as an abstract graph $G = (V, E)$. Doing so permits abstraction of any spatial details (such as road segment curves, or even length) and allows instead consideration of only the discrete connectivity or topology of the environment. This suggests the use of principles of discrete pursuit-evasion, such as searching for or pursuing a target moving on a graph. Many such problems are well-studied and have tractable solutions, which if applied to the CSAT task within road networks may prove much more feasible than existing general trajectory planning in continuous space.

Of the common search and pursuit abstractions referenced in Section 3.1, the two primary types considered here are so-called graph search and cops & robbers problems. As previously defined, graph search is generally phrased as having searchers perform a series of discrete motions in the graph (such as sweeps of edges) that by its completion would assuredly corner a target or verify that none is present, while cops & robbers considers the game-theoretic problem of moving a team of searchers between nodes to assure that a searcher will reach the same node as an evader moving in turn in finite time or with maximum probability within a fixed amount of time. Variations of these problems may be considered along the axes of (among many others) target visibility (whether its location is known), observability or capture distance, relative speed, level of adversariality, and whether targets are assumed to occupy nodes or edges.

Graph search typically assumes an infinitely capable and knowledgeable target and thus is extremely over-conservative in any practical problem, however this trivially provides applicability to continuous problems having finite-length edges with possibly-varying traversability as it has no dependence on the execution time of actions. If a target is treated as lying only along edges and occupies a junction for only an instant, the problem is one of *edge search*, which conveniently was previously noted as the commonly studied case (over *node search*, the opposite problem). In the graph search abstraction, as long as the number of searchers exceeds the *search number* of the graph, a *search schedule* may be computed providing an ordered list of searcher edge transitions that, when swept in sequence, guarantees capture of any target that may be present by its completion.

Scenario variations using other parameters for the target (visible, non-omniscient, or having finite speed) would fall into some class of cops & robbers problem, with a target now lying on a node (either by treating road segments as nodes and junctions as a set of edges connecting these or by inserting nodes along road segments). In this abstraction, the most commonly considered case is that of a visible, adversarial, finite-speed target. Except in some scenarios considered in Section 5.4, this case is not immediately applicable, as a target having known location is likely in view (or already "cornered" for UAV-observer purposes), and it is rather its future motions that need to be anticipated. Instead, the most applicable abstract case would be that of an invisible, possibly-adversarial, finite-speed target, for which some simple solutions also exist.

A vital fact when attempting to apply any such search or pursuit abstraction to real-world tasks as considered here is that UAVs do not behave as typical searchers in either abstraction. Most importantly, they are not themselves restricted to the graph and may be able to move between edges sharing no incident node and vice-versa. However, at the same time, motion constraints may limit instantaneous connectivity, for instance preventing a sharp turn or hovering at a location to block passage through a point. Additionally, none of the varying gross abstractions of sensor models such as a radius around current searcher position, infinite or finite-range rays in the plane emanating at this current location, or total visibility within the currently occupied edge or node are applicable. Instead, the true sensor footprint is of course a finite spot at a (possibly-fixed) location relative to current searcher pose that may cover incomplete portions of zero or more nodes or edges.

Therefore, mapping a UAV search or pursuit task to a graph-based abstraction to reap its benefits is itself by no means straightforward and is a major focus of proposed work. Towards this end, a possible progression of such mappings is proposed, in increasing order of level of integration between the UAV and target state spaces, that is also in increasing order of complexity and uncertainty of successful completion.

Separate target and UAV spaces. As the simplest case, target and UAV state spaces may be treated independently as has been done so far, with the target moving in a graph and UAVs along continuous paths unconstrained except for vehicle dynamics. Graph topology is used to constrain target location and possible motion, but UAV trajectory generation remains a continuous-space optimization problem. As an incremental alternative, searcher or pursuer motions may be generated in the graph abstraction and then a continuous-space trajectory generated that moves a (possibly joint) sensor footprint along the intended path within the graph. For example, to execute a graph motion corresponding to a right-angle turn in world space using two fixed-wing aircraft having fixed forward-facing cameras, a joint motion is needed in which one UAV moves towards and passes over the turn, while the other approaches from the orthogonal direction timed so as to provide a contiguous joint footprint.

UAV action primitives. Considering only the target graph, primitive actions for UAVs may be defined corresponding to needed abstract graph motions. The two simplest include an edge sweeping primitive (defined as laterally aligning a

UAV so that the sensor footprint overlaps the edge and following its length) and a guard primitive (defined, when possible, as hovering over or orbiting a point so that a desired finite region continuously lies within the sensor footprint). Key difficulties are that the time required to execute a given primitive is not constant (such as for sweeping edges of varying lengths) and that transitions between a given pair of primitives may take a varying amount of time or may not be possible without an intermediate motion taking the sensor footprint away from the required continuous path (such as overflying a location and turning around slowly to follow an otherwise too-sharp turn). One means to address this is to simply incorporate additional guard actions by other UAVs during such transitions so that possible gaps in an intended continuous coverage trajectory are filled in, though clearly this requires additional vehicles and is at times unnecessarily conservative. This may be improved upon by exploiting a lower bound on the factor by which UAV speed exceeds a known bound in target speed to model the problem as one in which searchers execute n such primitives in the time taken for a target to move between a single pair of graph nodes. Better would be to (pre)compute switching time between primitives and whether a guard is required for a given such transition and generate discrete primitive action plans given a known target speed bound. As stated, such primitives primarily provide a means to perform *edge search*.

Pose-dependent adjacency matrix. The road graph may be augmented with nodes added along edges at most one UAV sensor footprint diameter apart, at which point a UAV having any portion of the graph at the center of its field of view may be said to have a location within the graph (such as at the graph node most within the field of view). Treating the sensor footprint as a steerable object within the graph, a pose-dependent adjacency matrix for searchers may be defined for nodes in the graph based on the presence of a local plan to move the sensor footprint from one node to an adjacent one without leaving the graph. Possibly, such plans may include the use of an additional vehicle to fill in sensing gaps that may otherwise arise during (for instance) sharp turns. Use of such a representation permits planning discrete searcher motions directly within the graph by additionally tracking the orientation (or any needed higher-order states) of each UAV to generate pose-dependent neighborhoods, thereby providing a possible *node search* or pursuit in the space of nodes.

Combined target and UAV graph. At least conceptually, in addition to augmenting the graph with nodes along edges, nodes representing a sampling of the UAV state space having positions near the graph may also be added. These nodes could be redundant in that a single point in the original graph may receive several nodes, each representing a distinct sample of the UAV configuration space at which this point in the graph lies within the sensor footprint of a UAV with this configuration. A target may therefore lie at multiple points in the graph at once, and adjacency for searcher motion again depends on the presence of a local plan between any two configurations. Very roughly, this takes inspiration from the probabilistic roadmap class of motion planning planning, [58] in which otherwise intractable non-convex continuous search problems are reduced to discrete path planning in a graph by randomly sampling the configuration space and building an adjacency graph by finding local paths between nearby points. Such graph augmentation provides unification of target and UAV state spaces, but it is not immediately clear how to tractably generate the augmented graph, whether the resulting graph will be sparse enough to benefit from graph-based search and pursuit abstractions, and how easily methods for these can be extended to a target occupying multiple nodes (effectively, multiple targets moving under a coupled model).

5.3 Pursuit in Structured Environments

As a first step in making use of a graph constraint during pursuit, a single road segment was explored as inspiration in Section 4.4.1, showing substantial improvement in target geolocation accuracy, loss frequency, and trajectory planning tractability. Meanwhile, the naive cellular decomposition of a more general road network studied in Section 4.4.2, while providing a comprehensive solution that benefits from a reduction in target state space and motion branching factor, remains impractical due to limited tractability and ability to represent purposeful target motion.

As generalization to road networks representable as a graph, it is instead proposed that just as for a single road segment, principles of existing constrained and multi-model estimation used in existing work in geolocation from high-altitude radar reviewed in Section 3.3.2 may be applied to reap specific benefits for small-UAV pursuit. Specifically, road-constrained estimation as defined in Section 4.4.1 may be implemented in a multi-model framework such that a one-dimensional geolocation estimator is run for several road segments in parallel each having a relative likelihood of being the correct hypothesis (each defined as a model in this context), observations received are applied to each filter, relative likelihoods are recomputed given each filter's agreement with each observation, and an overall best estimate is derived through a likelihood-weighted average of each estimator's mean. A common means for accomplishing this generically is via the Interacting Multi-Model (IMM) filter, [8] implemented in GMTI radar applications as the variable structure IMM (VS-IMM). [62]

Alternatively, this may be implemented as a particle filter by any of several means referenced in Section 3.3.2, such as either running a distinct particle filter for each model wrapped by an overall IMM filter as above, using one single set of particles constrained to the graph in which the process model produces multi-model behavior by choosing an outgoing road for each particle at junctions, or through an unconstrained set of particles (having two-dimensional state) on which a combined observation likelihood is produced through a weighted sum over models. These too have been previously explored in a radar tracking context. [31]

However implemented, the result is a small number of hypotheses—represented as a set of road-constrained filter estimates or particles lying on one or more road segments—each individually deriving the accuracy benefits of a road-constrained filter. If pursuit is performed using an objective function branching on or taking expectation over this set of hypotheses, only a linear factor larger (by the size of this set) additional computation is required by a fixed hypothesis set over a single segment, and in practice unpruned hypotheses over only up to several road segments at a time seem likely. During trajectory planning, the expected hypothesis set grows exponentially with the number of junctions that may be reached by a target within the horizon considered, but for any typical environment with road segments at least one hundred meters long, this number too is small. Thus, the result is a tractable means to retain the accuracy and robustness provided by road-segment-constrained pursuit while expanding to larger environments.

Using a metric such as maximum total expected reduction in uncertainty, a team of UAVs will even likely self-distribute among nearby hypotheses as the information benefit of observing two targets with one UAV each exceeds that of observing one target with two UAVs due to diminishing returns. Such assignment may also be performed explicitly, though choosing an optimal assignment non-heuristically may be computationally demanding.

Possession of a compact representation of target location uncertainty in the form of several, moderately accurate hypotheses with limited options for short-term motion admits a number of interesting possible extensions. One is that while evaluating UAV trajectories, if a large step increase in entropy occurs for all joint trajectories considered as a target is predicted to reach a junction, it may be assumed that growth in uncertainty (from a progressive increase in hypotheses) has exceeded the available resources of the UAV team that would be necessary to be poised to most quickly disambiguate between future hypotheses, for instance by effectively assigning a UAV to each hypothesis. In such cases, which may occur for instance when the number of total hypotheses has exceeded the number of UAVs, an appropriate course of action may be to stop planning at this horizon since any actions considered beyond that point will be relatively equally poorly performing in expectation, while simply placing sensors to best take advantage of this jump in uncertainty—for instance by simply planning to have the relevant intersection in view when the target is expected to reach it—will provide eventual disambiguation.

Another now-feasible extension is the use of knowledge of intent in a target process model. If some road segments are known to be more preferable than others to a target—for instance based on level of occlusion, congestion, safety, or traversal speed—then predicted target motions may be more heavily weighted towards more likely path choices. Similarly, if one or more intended destinations for the target are known (such as safe-houses, political boundaries, or rescue locations), prediction may be guided towards these by predicting only along best paths (under one of the preceding metrics) or at least by avoiding predicting along segments in clearly opposing directions. Combining knowledge of path preference with a known destination permits an easily accomplished search for shortest paths a target that is rational and well aware of its environment is most likely to take these destinations.

Yet another such extension is the use of future total path branching factors to more heavily weight hypotheses that could more easily escape. Consider, for instance, pursuing a target nearing a junction with two otherwise equally likely branches, one leading to a dead end and the other to another intersection having many branches through which it could escape. Given that the cost of mistakenly better covering the segment leading to many branches lies far below that of the opposite (as, in this case, if a target is not seen where expected, the other hypothesis must lie in the dead end), artificially inflating the likelihood of the target choosing the better escape route may be in order. Indeed, this need not be completely heuristic but rather represent an explicit model of an evasive target with knowledge of the environment itself as having intent leave as many escape routes open as possible.

Finally, it may be noted that the proposed strategies retain a number of limitations. First, though the representation of a target is simplified through constraints and discretization (of hypotheses), UAV motion remains a continuous planning problem. As this aspect of the problem is fundamentally one of anticipating short-term target motion, it is not immediately clear how this may be improved upon in a graph-theoretic framework. Additionally, only stochastically-moving rather than truly omniscient adversarial targets (having fully knowledge of pursuer state, algorithm, and intent) have been considered in that trajectory optimization fundamentally relies on maximizing objective expectation over target location uncertainty and possible motions. However, capture of adversarial targets is a common goal of graph-based pursuit-evasion abstractions, suggesting that though a sufficiently adversarial target may evade persistent pursuit, good strategies for repeated recapture may exist. Finally, the expressiveness and tractability of the approach of a discrete set of hypotheses falters as the size of this set increases. Common approximations listed in Section 3.3.3 or explicit assignment of UAVs to targets may be performed to somewhat mitigate growth in computational demands, but beyond a point, the problem is better characterized as a search, for which recapture strategies that also exploit the environment's graph structure are desired.

5.4 Recapture in Structured Environments

If a target within a road network is lost during pursuit or was reported to have been at a location at some previous time, the graph structure of the environment is immensely useful for recapture because it limits the target to a small discrete set of paths from the point of loss or any last believed locations. For simplicity, and with only minimal loss of generality, the assumption will be made that a target will never make U-turns, either at a junction or at any point along an edge; in other words, the only time a target will change from moving in a given direction along an edge is at a junction, by choosing another edge distinct from the one by which it entered. Two general recapture strategies for an escaping target are proposed that exploit these limited routes of escape, followed by consideration of a general abstraction.

In the simplest case, if a target of initially known location is moving at known constant speed v, then at any point in the future t, its possible location is one of any of the points in the graph having distance vt from its initial location, along its original direction of motion. The subset of the graph reachable in time t may be explicitly represented as a tree rooted at the initial location, with duplicates of original graph nodes added as necessary to account for possible backtracking due to cycles. Now, the recapture problem is one of eliminating a set of point hypotheses in a tree by visiting each in some sequence (a discrete moving target assignment problem) and is equivalent to a traveling salesman problem with moving cities. Though well known to be NP-hard, its intractability arises from the number of cities, not distances between them, and as hypotheses are only added slowly as each passes through a junction, this remains an easily solved problem. A first extension is to consider a target having instead initially uncertain location, which produces a set of equally-sized uncertainty distributions moving through the tree, rather than points. This may be accommodated by scheduling short road sub-segment sweeps rather than point visitations. Likewise, if the target's speed is not known precisely (but instead estimated during prior pursuit), the hypothesis set is now a set of uncertainty distributions growing in size with integrated speed uncertainty requiring increasing-length sweeps. Finally, if uncertain detection is included in the sensor model, finite rather than instantaneous-duration visitation is required, for instance by scheduling brief orbits or repeated sweeps of expected locations.

As a generalization from known or estimated speed, simply bounded speed may also be considered. In this case, a target moving at speed at most $v_{\max}$ may lie at any point in the graph having distance up to $v_{\max}t$ from the initial location, again along its original direction of motion. In this case, the earliest time at which each branch at a junction may occur may be enumerated by tracing the spread of a hypothetical target moving at speed $v_{\max}$. Through linear search in this list, it may be rapidly determined whether the N UAVs can together outrun all hypotheses prior to a given branch (by optimal assignment and execution of a single shortest direct path to the assigned hypothesis) before the number of hypotheses exceeds N. If so, then a guaranteed recapture strategy is to reach the most distant possible location of each hypothesis just before this branch and sweep backward, up the tree.

In general, the most applicable abstraction for recapture is the cops & robbers problem on a graph. The actual case here is that of invisible (unknown or uncertain location) robbers that is not well studied in the discrete pursuit-evasion literature but represents an interesting avenue for theoretical exploration. One worthwhile aspect to consider is that of the dependence of the cop number of a graph on the relative speed of searchers and a target. Clearly, an infinite-speed cop is equivalent to having a stationary target, providing the useful reminder than for UAVs of sufficiently high relative speed, local coverage search may provide an effective solution. Conversely, an infinite-speed target could have escaped to any location in the road network at the moment it was lost and is better considered in the context of search. In between lie a number of special cases, one of which is that of a stochastically-moving target having known location but comparable speed to the UAVs (hence preventing the trivial recapture strategy of moving directly to the target's location), to which the problem abstraction sometimes referred to as cops & drunk robbers may be applied, for which algorithms exist at least for small graphs. [60]

Overall, a graph structure provides a number of benefits for the task of recapture. Most importantly, it permits a principled approach to considering possible target paths in the absence of observations without resorting to heuristic-driven or brute-force local area search. Given a discrete set of such paths, scalability to larger areas for a given number of UAVs is enabled by a greatly reduced search space compared to an open area, and in some cases guaranteed recapture is possible that would have otherwise required a large number of vehicles for area containment rather than simply blocking a small number of escape paths. Fundamentally, however, recapture can still be viewed as a local search in a small but growing subgraph, suggesting as a possible area of exploration the question of whether a viable recapture strategy is to block a well chosen set of edges to fix the size of this subgraph and then execute a search strategy from the following section within it. In any event, if this subgraph grows beyond some threshold, it may be said that the window of opportunity for recapture has passed and that the problem is now a search task.

5.5 Search in Structured Environments

Some of the most interesting UAV behavior has the potential to be produced by an appropriately designed search exploiting environmental structure. The two primary pieces of information a road network provides this are that a target must lie within a greatly reduced surface area that needs to be searched and that a target can only move into previously searched areas along known road corridors that tightly channel the spread of recontamination. As the literature in graph-theoretic search and pursuit makes a strong distinction between adversarial targets requiring conservative *guaranteed* search and stochastically-moving targets for which strategies performing optimally in expectation instead provide an *efficient* search, the same will be made here. Separately, independently of the true intent of a target, a mission planner will likely have some form of stochastic model for any target and may either need to assure detection if present or simply wishes to try to locate it as quickly as possible, already imposing a decision as to which type of search strategy is to be used.

For a given environment, existing work in graph search provides strategies for guaranteed capture of a target if sufficient searchers are available. Employing classical graph search, a clearing strategy can be generated for a minimum number of abstract graph-bound searchers that guarantees capture of an adversarial target of any speed. This may then be applied to UAV searchers as suggested in Section 5.2 using either a direct mapping or by modifying the environment graph to account for feasible UAV motions. Doing so, however, may require an unreasonable number of UAVs given over-conservatism from handling targets of any speed as well as suboptimality in the mapping from abstract to UAV searchers. In practice, targets are unlikely to have speed even substantially faster than the UAVs themselves, and a conservative bound on this speed can be estimated. Unfortunately, abstract graph search for a target of unknown location but finite speed is not a well studied problem. One promising avenue for possible modification is the IGNS algorithm by Kehagias et al. [59] providing a guaranteed node search strategy for targets having specifiable per-timestep reachability matrix.

In many cases it is likely that fewer UAVs are available than are required to attain a capture guarantee, and in such cases or when it is simply desirable to locate a target as quickly as possible, efficient search may be a reasonable secondary goal. Two possible objective phrasings include minimizing expected capture time and maximizing probability of capture in a given search duration. Mapping the UAV search task to a graph representation for this purpose is less straightforward as there exists the fundamental challenge that a probability of target presence for every point in the environment must somehow be stored, yet targets lie at a continuous point along finite graph edges rather than only at nodes or across an entire edge. This is essentially the same difficulty encountered when representing target location uncertainty for the case of pursuit, and here too, decomposing edges into cells and searching using an existing cellular representation from Section 3.4.2 presents a solution. However, for the same reasons as observed in Section 4.4.2—primarily intractability from requiring large numbers of cells—this remains impractical, evermore so because many more cells over a wider area may contain significant probability than for pursuit. At the same time, the multi-hypothesis solution proposed for pursuit is not directly applicable to search since, until detected, no positive target observations are available with which to maintain distinct geolocation filters. Nevertheless, inspiration may be derived from both. It is proposed instead that target location probability be represented by particles lying along edges drawn from and approximating the true continuous distribution, which are then updated using a particle filter and a search planner maximizing a short-term probability of detection as proposed by Geyer [41] or Tisdale et al. [105] Though still fundamentally a discretization of the environment, it carries a number of tractability and expressiveness advantages over a cellular one. First, fewer samples are needed since particle presence in low-probability areas will be sparse, and areas containing no particles are simply not considered during planning or observation updates, contributing nothing to computational load. The number of samples required must still be sufficient to meaningfully approximate (and varies with the shape of) the current target probability distribution but may otherwise be varied depending on the desired fidelity of approximation to reduce resource requirements if needed. Additionally, applying target diffusion for recontamination corresponds to running a motion model for each particle rather than an operation every discrete point in the environment and thus grows in complexity with the number of samples rather than the size of the environment. Finally, incorporating purposeful diffusion only entails adding higher-order dimensions (such as speed and direction) to the particle state and motion model, only expanding the number of samples needed at locations already being sampled (rather than at every point in the environment), and continuous distributions in higher-order state may themselves also be approximated by sampling rather than inevitably coarse cellular discretization.

Likewise, as suggested for pursuit, limited paths for target motion enable much more effective use of knowledge of target intent or capability than for an open area. Target preference for or avoidance of certain edges or intent to reach one or more destinations may be built into the recontamination or diffusion process model, possibly reducing expected capture time or the number of UAVs required to achieve a given expected capture time since target presence in or motion towards certain areas is less likely. Another possible extension worth consideration is that during a pursuit, if a sufficient number of idle UAVs are available, they may be placed at good starting locations for a guaranteed search strategy in a small surrounding subgraph that is continuously recomputed to provide rapid guaranteed recapture in the event of target loss.

Generally, exploitation of a road network's presence to map abstract graph search methodologies to the UAV search task provides scalability improvement over search in an open area as less area need be covered and fewer escape routes blocked, detection guarantees using fewer searchers for the same reasons, and a principled means of determining whether a guarantee is possible for a given number of UAVs (by comparing this to the search number of the corresponding abstract graph, accounting for additional guard UAVs that may be required by the abstraction mapping used). Though due to inefficiency in this mapping the number of UAVs determined to be needed for searching an area may exceed the number that is strictly *necessary*, applying a graph search abstraction provides a tractable means of determining a *sufficient* number that may be considerably less than that generated by a guaranteed sweep of the equivalent open area.

5.6 Generalization to Mixed Environments

Though as stated, the proposed improvements rely on environments easily represented as a graph, generalization to a broader class of environments is possible with some ease. One such means, proposed here, is to apply a hierarchical strategy in which pinch points or short boundaries between otherwise open regions are first identified, producing a set of disjoint regions with inter-connectivity defined by shared boundaries that may be blocked to prevent target passage between the two regions. This may at one level be treated as a graph, in which nodes are the open regions and edges are their boundaries. At a lower level, each region may be seen as a small open area reasonably efficiently covered by any existing guaranteed sweeping strategy. Several examples of such decompositions of increasing complexity are shown in Figure 5.1.

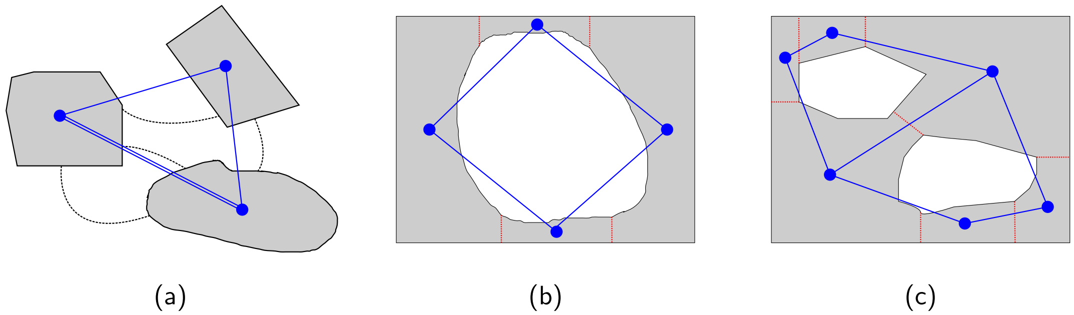

Figure 5.1: Three example environments to which the proposed decomposition is applied, as open areas connected by boundaries that may then be treated as nodes and edges, respectively, depicted as a superimposed graph in blue. The first (a) presents an ideal example in which several areas each requiring few searchers have only a few narrow (easily-guarded) interconnections, while the total area is large and would require many searchers if treated as an open area. The second (b) provides an alternative case of an otherwise large area containing an obstacle occupying most of the area. Here too, if treated as an open area, many searchers would be required, whereas it may instead be decomposed into several regions each requiring few searchers and connected by short boundaries. The last (c) points out that though decompositions of complex environments may be treated as a graph, it may not be advantageous to do so. In this case, the number of UAVs required to guard boundaries and perform regional searches exceeds that required to simply sweep the enclosing rectangle, as may be seen from the presence of two large open areas and many boundaries, the length of even several of which is approximately the width of the entire enclosing rectangle.

Using such a decomposition, the outer virtual graph may, conceptually, be searched using any applicable aforementioned means, with several alterations. These entail mapping the blocking of an edge to blocking an entire regional boundary, while clearing or guarding a node constitutes executing an appropriate open-area sweep within the region (if not already performed since the last opportunity for recontamination) and blocking all its entrances. Given sufficiently many UAVs, environments of considerable size and complexity may now be considered. Similarly, the case of recapture may be likewise extended, operating at the level of the virtual graph, though once a target has possibly reached a large open area, it must be searched, quickly growing in required resources to the number of vehicles required for a full-fledged search. The case of pursuit is more challenging, as the proposed improvements rely on the presence of narrow corridors such as may be present between regions, and options for generalization remain an open point of consideration.

The proposed decomposition reduces a large class of environments that otherwise occupy a large area and which present an infeasible challenge down to a distinct set of simpler environments joined as a graph, permitting formal reasoning using their connectivity. Clearly, for sufficiently densely connected environments such as those containing only sparsely distributed obstacles, the decomposition provides little benefit and may produce a graph requiring an unnecessarily large number of UAVs to search it. It is argued that by and large, such environments reach levels of dense internal connectivity that they are better considered as nearly completely open areas that simply remain a difficult problem.

5.7 Summary of Benefits

In any environment that contains large areas that are untraversable by a target that must pass narrowly between them, such as the case of road networks considered here, knowledge of the environment provides constraints that reduce the target domain, limit target motion between observations, decrease available escape paths once lost, and leave few enumerable paths to possible destinations that may be given varying likelihoods given additional knowledge of target intent or capability.

Such simplifications permit efficient, constrained representations of possible target location and motion. Compared to more general representations required for arbitrary environments, improvements are achievable in accuracy of target geolocation estimates, the computation time required for estimation updates and vehicle trajectory planning, the number of agents required to achieve a detection guarantee or desired level of geolocation accuracy, the expected time to capture or expected proportion of time a target is in view using a given number of vehicles, maximum target speed or maneuverability that may be pursued with a given size and capability of UAV team, and the amount of data that needs to be transferred between vehicles to complete a search or pursuit task. In practice, these benefits will produce an immediately observable improvement in task performance.

5.8 Validation Strategy

Field experimentation to date described in Section 4.1 has built up the architectural and software groundwork to evaluate any of the proposed ideas on a variety of real vehicles having high-level autopilots as access to such hardware permits, with minimal modification required between vehicle types. A presently ongoing project makes use of teams of one or more Maxi-Joker helicopters pictured in Figure 1.3 as scouts to search in obstacle-ridden terrain for landing sites for a larger helicopter, which may be seen as an efficient or guaranteed search task using knowledge of terrain boundary and obstacles to plan paths avoiding impossible areas. Additional opportunities within projects focusing on rescue and surveillance are likewise expected.

For initial algorithmic evaluation and where impractical to perform tests on real vehicles–as may be the case for highly repetitive, risky, or large-team scenarios–experimentation in simulation is also anticipated. Much of the geolocation and pursuit work described in Chapter 4 was performed in a simple simulation environment developed for this purpose. Improvements in the modularity and speed of this software are underway. Using this, it is expected that rapid testing of parameter and team-size variations may be performed with minimal overhead in experimental effort.

For consistent comparison in simulation, testing on a substantial set of graph-like environments is anticipated. To date, maps for several actual locations have been generated as shown in Figure 5.2, and more may be straightforwardly derived for a portion of any city in the United States using geo-referenced satellite imagery and so-called TIGER (Topologically Integrated Geographic Encoding and Referencing system) [111] data published by the U.S. Census Bureau. Generation of additional synthetic environments is expected, to evaluate the effects of varying sparsity (proportion of the total bounding region that possible target locations take up), connectivity density (degree and frequency of junctions, representable as the average length of road segments), and observation ambiguity (average proximity of nearby roads). Clearly, these properties are strongly coupled, but each may be individually quantified.

(a) (b) (c)

Figure 5.2: Several real-world locations representable as a graph used in simulation experimentation to date: the well-known McKenna MOUT site at Ft. Benning, GA (a) and portions of the South-side (b) and Squirrel Hill (c) residential areas within the city of Pittsburgh, PA. In all cases, the areas are assumed to be bounded, for instance through coverage of the perimeter by a separate ground-based team.

CHAPTER 6

Conclusion

The proposed thesis considers the problem of continuous search and tracking with teams of unmanned aircraft, which has been noted as being most closely related to the previously defined cooperative search, acquisition, and tracking (CSAT) problem, disregarding elements of cooperation and acquisition for the purposes of scope. Decomposed in Chapter 2, a given portion of a mission may be classified as a pursuit, recapture, or search task, each of which may be extremely hard problems themselves.

A substantial body of existing work in aerial geolocation, pursuit, search was covered in Chapter 3, yet practical performance using small field-reconnaissance UAVs remains poor given their limited sensing accuracy, maneuverability, and ground-station or on-board computation. Trajectory planning tractability is a particular challenge given that a myriad of possible target actions, UAV motions, and erroneous observations given sensing uncertainty must be considered.

Work to date described in Chapter 4 has demonstrated practical improvements in visual tracking to produce observations and geolocation filtering of these, while exploring the performance benefits of exploiting constraints on target position and actions. Proposed work outlined in Chapter 5 seeks to build on this by specifically considering road-network environments, more fully exploiting the accuracy and predictability benefits that representations enabled by such environments provide, while mapping relevant search and pursuit abstractions onto the UAV task to leverage existing solutions that are tractable and well-performing, even if strictly suboptimal given sacrifices in producing such mappings.

Proposed workplan and schedule

The completion of proposed work is anticipated according to the schedule in Figure 6.1. Total completion time is estimated to be 18 months, concluding Spring 2013. The timing of the proposed validation phase is subject to hardware availability and supporting project funding, however this is not considered a concern as such availability is frequent and expected.

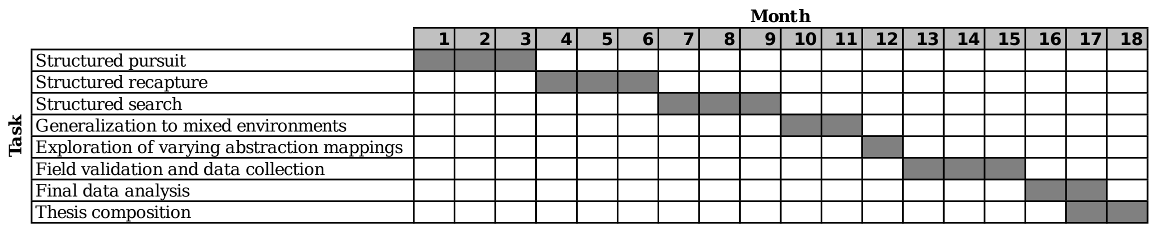

Figure 6.1: Proposed schedule and workplan. Phases correspond to sections of proposed work described in Chapter 5.

Possible extensions and future work

A number of closely related issues or avenues for extension remain unaddressed by the proposed work. Several of these include:

Target occlusion In practice, environments may have obstacles of sufficient height to shadow targets nearby. Reasoning about these effects can eliminate a large proportion of unexpected observation dropouts.

Ground agents Work to date on practical air-ground collaboration has shown clear value in pairing air and ground vehicles given their complementary capabilities. Including ground agents in proposed strategies provides observations having a very complementary sensing model, permits physical interaction with targets, and enables direct application of pursuit-evasion abstractions intended for planar (ground) vehicles.

Multiple targets As stated, a single target is generally assumed, though many representations naturally handle multiple targets as additional hypotheses or by implicitly resulting in the assignment of observers to particular targets. In an ad-hoc fashion, teams might be divided explicitly between targets, and existing work in task allocation may be applicable to formalize this.

Decentralized operation In practice, failures of individual agents are inevitable, and communication bandwidth is highly limited. Applying and extending existing strategies such as decentralized estimation and anonymous cooperation for reducing reliance on a central ground station and the amount of data that must be transmitted by each vehicle would improve both the robustness and practical applicability of these systems.

If an opportunity or clear need arises, any of these may be included in the eventual work to be performed. Otherwise, they present a list of possible directions for future work in this area.

Acknowledgements

Much of the work presented in Chapter 4 is joint work as part of several projects in the Field Robotics Center. Portions of the described visual tracking and static geolocation effort were led by Stephen Nuske, and nonlinear estimation strategies were led by Ben Grocholsky.

The support of US Army RDECOM-ARDEC in collaboration with iRobot Corporation, the United Technologies Research Center, and a DoD National Defense Science and Engineering Graduate (NDSEG) fellowship are gratefully acknowledged.

Finally, the support and guidance of the thesis committee in the preparation and improvement of this document, especially the enthusiastic efforts by Thanasis and many thoughtful discussions with Ben, are deeply appreciated.

Bibliography

[1] Vitaly Ablavsky and Magnús Snorrason. Optimal search for a moving target: A geometric approach. In *AIAA Conference on Guidance, Navigation, and Control*, 2000.

[2] Micah Adler, Harald Räcke, Naveen Sivadasan, Christian Sohler, and Berthold Vöcking. Randomized pursuit-evasion in graphs. *Combinatorics, Probability and Computing*, 12:225–244, May 2003.

[3] AeroVironment, Inc. AeroVironment develops worlds first fully operational life-size hummingbird-like unmanned aircraft for DARPA. Press Release, 2011.

[4] Ali Ahmadzadeh, James Keller, George J. Pappas, Ali Jadbabaie, and Vijay Kumar. Multi-UAV cooperative surveillance with spatio-temporal specifications. In *IEEE Conference on Decision and Control*, 2006.

[5] E.D. Andersen and C.N. Taylor. Improving MAV pose estimation using visual information. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 3745–3750, November 2007.

[6] M. Sanjeev Arulampalam, Simon Maskell, and Neil Gordon. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.

[7] Simon Baker and Iain Matthews. Lucas-Kanade 20 years on: A unifying framework. *International Journal of Computer Vision*, 56(3):221–255, 2004.

[8] Yaakov Bar-Shalom and X. Rong Li. *Estimation and Tracking: Principles, Techniques, and Software.* Artech House, 1993.

[9] D. Blake Barber, Joshua D. Redding, Timothy W. Mclain, Randal W. Beard, and Clark N. Taylor. Vision-based target geo-location using a fixed-wing miniature air vehicle. *J. Intell. Robotics Syst.*, 47(4):361–382, 2006.

[10] Lali Barrière, Pierre Fraigniaud, Nicola Santoro, and Dimitrios M. Thilikos. Searching is not jumping. In *Workshop on Graph Theoretic Concepts in Computer Science*, pages 34–45, 2003.

[11] Randal W. Beard and Timothy W. McLain. Multiple UAV cooperative search under collision avoidance and limited range communication constraints. In *IEEE Conference on Decision and Control*, 2003.

[12] Stanley J. Benkoski, Michael G. Monticino, and James R. Weisinger. A survey of the search theory literature. *Naval Research Logistics*, 38(4):469–494, August 1991.

[13] Deepak Bhadauria and Volkan Isler. Capturing an evader in a polygonal environment with obstacles. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.

[14] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Coordinated decentralized search for a lost target in a Bayesian world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.

[15] K Briechle and U Hanebeck. Localization of a mobile robot using relative bearing measurements. In IEEE *Transactions on Robotics and Automation*, volume 20(1), pages 36–44, 2004.

[16] Mitch Bryson and Salah Sukkarieh. Bearing-only SLAM for an airborne vehicle. In *Australasian Conf. on Robotics and Automation*, 2005.

[17] F. Caballero, L. Merino, J. Ferruz, and A. Ollero. Vision-based odometry and SLAM for medium and high altitude flying UAVs. *J. Intell. Robotics Syst.*, 54(1-3):137–161, 2009.

[18] Mark E. Campbell and William W. Whitacre. Cooperative tracking using vision measurements on SeaScan UAVs. *IEEE Transactions on Control Systems Technology*, 15(4):613–626, 2007.

[19] D.W. Casbeer, Pengcheng Zhan, and A.L. Swindlehurst. A non-search optimal control solution for a team of MUAVs in a reconnaissance mission. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2006.

[20] Howie Choset and Philippe Pignon. Coverage path planning: The boustrophedron decomposition. In *Proceedings of Field and Service Robotics (FSR)*, Canberra, Australia, December 1997.

[21] Timothy H. Chung, Geoffrey A. Hollinger, and Volkan Isler. Search and pursuit-evasion in mobile robotics. *Autonomous Robots*, 31(4):299–316, 2011.

[22] Robert T. Collins and Yanxi Liu. On-line selection of discriminative tracking features. In *IEEE Conf. on Computer Vision (ICCV)*, pages 346–352, 2003.

[23] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 142–149, 2000.

[24] Gianpaolo Conte and Patrick Doherty. Vision-based unmanned aerial vehicle navigation using geo-referenced information. *EURASIP J. Adv. Signal Process*, 2009:1–18, 2009.

[25] Gianpaolo Conte, Maria Hempel, Piotr Rudol, David Lundström, Simone Duranti, Mariusz Wzorek, and Patrick Doherty. High accuracy ground target geo-location using autonomous micro aerial vehicle platforms. In *AIAA Conference on Guidance, Navigation, and Control*, August 2008.

[26] Jeffrey B. Corbets and Jack W. Langelaan. Parameterized trajectories for target localization using small and micro unmanned aerial vehicles. In *American Control Conference*, 2008.

[27] Michael Dille, Ben Grocholsky, and Stephen Nuske. Persistent visual tracking and accurate geo-location of moving ground targets by small air vehicles. In *AIAA Infotech@Aerospace Conference*, March 2011.

[28] Michael Dille, Ben Grocholsky, Stephen Nuske, Mark Moseley, and Sanjiv Singh. Air-ground collaborative surveillance with human portable hardware. In *AUVSI Unmanned Systems North America*, August 2011.

[29] Joseph Djugash, Sanjiv Singh, and Benjamin P. Grocholsky. Modeling mobile robot motion with polar representations. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, October 2009.

[30] L. Dubins. On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79:497–516, 1957.

[31] Mats Ekman and Egils Sviestins. Multiple model algorithm based on particle filters for ground target tracking. In *International Conference on Information Fusion*, 2007.

[32] Dale Enns, Dan Bugajski, and Steve Pratt. Guidance and control for cooperative search. In *American Control Conference*, 2002.

[33] Matthew D. Flint. *Cooperative Unmanned Aerial Vehicle (UAV) Search in Dynamic Environments Using Stochastic Methods*. PhD thesis, University of Cincinnati, 2004.

[34] Fedor V. Fomin and Dimitrios M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theoretical Computer Science*, 399(3):236–245, June 2008.

[35] Eric W. Frew. Cooperative standoff tracking of uncertain moving targets using active robot networks. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2007.

[36] Eric W. Frew and Jack Elston. Target assignment for integrated search and tracking by active robot networks. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2008.

[37] Eric W. Frew, Jack Langelaan, and Maciej Stachura. Adaptive planning horizon on information velocity for vision-based navigation. In *AIAA Conference on Guidance, Navigation, and Control*, 2007.

[38] Eric W. Frew and Stephen M. Rock. Trajectory generation for constant velocity target motion estimation using monocular vision. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[39] Brian R. Geiger, Joseph F. Horn, Anthony M. DeLullo, Lyle N. Long, and Albert F. Niessner. Optimal path planning of UAVs using direct collocation

with nonlinear programming. In *AIAA Conference on Guidance, Navigation, and Control*, 2006.

[40] Brian P. Gerkey, Sebastian Thrun, and Geoff Gordon. Visibility-based pursuit-evasion with limited field of view. *International Journal of Robotics Research*, 25(4):299–315, April 2006.

[41] Christopher Geyer. Active target search from UAVs in urban environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2008.

[42] D. Gibbins, P. Roberts, and L. Swierkowski. A video geo-location and image enhancement tool for small unmanned air vehicles (UAVs). In *Intelligent Sensors, Sensor Networks and Information Processing Conference*, pages 469–473, Dec. 2004.

[43] Anouck R. Girard, Adam S. Howell, and J. Karl Hedrick. Border patrol and surveillance missions using multiple unmanned air vehicles. In *IEEE Conference on Decision and Control*, 2004.

[44] Arthur S. Goldstein and Edward M. Reingold. The complexity of pursuit on a graph. *Theoretical Computer Science*, 143(1):93 – 112, 1995.

[45] Ben Grocholsky, Michael Dille, and Stephen Nuske. Efficient target geolocation by highly uncertain small air vehicles. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, September 2011.

[46] Leonidas Guibas, Jean-Claude Latombe, Steven Lavalle, David Lin, and Rajeev Motwani. Visibility-based pursuit-evasion in a polygonal environment. In *Workshop on Algorithms and Data Structures*, pages 17–30, 1997.

[47] Ruijie He, Abraham Bachrach, and Nicholas Roy. Efficient planning under uncertainty for a target-tracking micro-aerial vehicle. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8, 2010.

[48] Tamir A. Hegazy. *A Distributed Approach to Dynamic Autonomous Agent Placement for Tracking Moving Targets with Application to Monitoring Urban Environments*. PhD thesis, Georgia Institute of Technology, 2004.

[49] Marcel L. Hernandez. Optimal sensor trajectories in bearings-only tracking. In *International Conference on Information Fusion*, 2004.

[50] Gabriel M. Hoffmann, Steven L. Waslander, and Claire J. Tomlin. Distributed cooperative search using information-theoretic costs for particle filters, with quadrotor applications. In *AIAA Conference on Guidance, Navigation, and Control*, 2006.

[51] Geoffrey A. Hollinger. *Search in the Physical World*. PhD thesis, Carnegie Mellon University, 2010.

[52] Mong–ying A. Hsieh, Luis Chaimowicz, A. Cowley, Ben Grocholsky, Jim Keller, Vijay Kumar, C. J. Taylor, Y. Endo, Ronald Arkin, Boyoon Jung, Denis F. Wolf, Gaurav S. Sukhatme, and Doug MacKenzie. Adaptive teams of autonomous aerial and ground robots for situational awareness. *Journal of Field Robotics*, 24(11):991–1014, 2007.

[53] Rufus Isaacs. Games of pursuit. Technical Report P-257, RAND Corporation, Santa Monica, CA, November 1951.

[54] Rufus Isaacs. *Differential Games*. John Wiley, 1965.

[55] Volkan Isler, Sampath Kannan, and Sanjeev Khanna. Randomized pursuit-evasion in a polygonal environment. *IEEE Transactions on Robotics*, 21(5):875–884, October 2005.

[56] Gregory F. Ivey and Eric N. Johnson. Investigation of methods for target state estimation using vision sensors. In *AIAA Conference on Guidance, Navigation, and Control*, 2005.

[57] Phillip J. Jones. *Cooperative Area Surveillance Strategies Using Multiple Unmanned Systems*. PhD thesis, Georgia Institute of Technology, 2009.

[58] Lydia Kavraki, Petr Svestka, Jean-Claude Latombe, and Mark Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

[59] Athanasios Kehagias, Geoffrey Hollinger, and Sanjiv Singh. A graph search algorithm for indoor pursuit/evasion. *Mathematical and Computer Modeling*, 50:1305–1317, 2009.

[60] Athanasios Kehagias and Pawel Pralat. Some remarks on cops and drunk robbers. arXiv:1106.2414v1 [cs.DM], 2011.

[61] Jongrae Kim and Yoonsoo Kim. *Optimal circular flight of multiple UAVs for target tracking in urban areas*. IN-TECH, 2009.

[62] Thiagalingam Kirubarajan, Yaakov Bar-Shalom, Krishna R. Pattipati, and Ivan Kadar. Ground target tracking with variable structure IMM estimator. *IEEE Transactions on Aerospace and Electronic Systems*, 36(1):26–46, January 2000.

[63] Andreas Kolling and Stefano Carpin. Pursuit-evasion on trees by robot teams. *IEEE Transactions on Robotics*, 26(1):32–47, 2010.

[64] Bernard O. Koopman. The theory of search, part III: The optimum distribution of searching effort. *Operations Research*, 5(5):613–626, October 1957.

[65] Swastik Kopparty and Chinya V. Ravishankar. A framework for pursuit evasion games in Rn. *Information Processing Letters*, 96(3):114–122, November 2005.

[66] Andrea S. LaPaugh. Recontamination does not help to search a graph. *Journal of the Association for Computing Machinery*, 40(2):224–245, 1993.

[67] Dale A. Lawrence, Eric W. Frew, and William J. Pisano. Lyapunov vector fields for autonomous UAV flight control. *AIAA Journal of Guidance, Control, and Dynamics*, 31(5):1220–1229, September 2008.

[68] Jusuk Lee, Rosemary Huang, Andrew Vaughn, Xiao Xiao, J. Karl Hedrick, Marco Zennaro, and Raja Sengupta. Strategies of path-planning for a UAV to track a ground vehicle. In *Autonomous Intelligent Networks and Systems Conference*, June 2003.

[69] R. Madison, P. DeBitetto, A. Rocco Olean, and M. Peebles. Target geolocation from a small unmanned aircraft system. In *IEEE Aerospace Conference*, pages 1–19, March 2008.

[70] Ivan Maza and Anibal Ollero. Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms. In *International Symposium on Distributed Autonomous Robotic Systems*, June 2004.

[71] Timothy G. McGee and J. Karl Hedrick. Guaranteed strategies to search for mobile evaders in the plane. In *American Control Conference*, 2006.

[72] Stephen Nuske, Michael Dille, Ben Grocholsky, and Sanjiv Singh. Representing substantial heading uncertainty for accurate target geolocation by small UAVs. In *AIAA Conference on Guidance, Navigation, and Control*, August 2010.

[73] Stephen Thomas Nuske, Michael Dille, and Ben Grocholsky. Robust automatic visual tracking from small UAVs. Technical Report CMU-RI-TR-11-?? (Forthcoming), Carnegie Mellon Robotics Institute, Pittsburgh, PA, November 2011.

[74] Office of the U.S. Secretary of Defense. Unmanned aerial vehicles roadmap 2005-2030, 2005.

[75] Torrence Parsons. Pursuit-evasion in a graph. In *Theory and applications of graphs*, pages 426–441. Springer, 1976.

[76] Valerii Semenovich Patsko and Varvara Leonidovna Turova. Homicidal chauffeur game: history and modern studies. Technical report, Russian Academy of Sciences, Ural Branch, Institute of Mathematics and Mechanics, Ekaterinburg, Russia, February 2009.

[77] N. N. Petrov. A problem of pursuit in the absence of information on the pursued. *Differentsial'nye Uravneniya*, 18:1345–1352, 1982.

[78] Marios M. Polycarpou, Yanli Yang, and Kevin M. Passino. A cooperative search framework for distributed agents. In *IEEE International Symposium on Intelligent Control (ISIC)*, 2001.

[79] Sameera S. Ponda. Trajectory optimization for target localization using small unmanned aerial vehicles. Master's thesis, Massachusetts Institute of Technology, 2008.

[80] M. Quigley, M.A. Goodrich, S. Griffiths, A. Eldredge, and R.W. Beard. Target acquisition, localization, and surveillance using a fixed-wing mini-UAV and gimbaled camera. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2600–2605, April 2005.

[81] A. Quilliot. *Problémes de jeux, de point fixe, de connectivité et de représentation sur des graphes, des ensembles ordonnés et des hypergraphes.* These d´Etat, Université de Paris VI, 1983.

[82] Steven A. P. Quintero, Francesco Papi, Daniel J. Klein, Luigi Chisci, and J. P. Hespanha. Optimal UAV coordination for target tracking using dynamic programming. In *IEEE Conference on Decision and Control*, 2010.

[83] Fahd Rafi, Saad M. Khan, Khurram H. Shafiq, and Mubarak Shah. Autonomous target following by unmanned aerial vehicles. In *SPIE Defense and Security Symposium*, Orlando, USA, 2006.

[84] Joshua Redding, Timothy W. Mclain, Randal W. Beard, and Clark Taylor. Vision-based target localization from a fixed-wing miniature air vehicle. In *American Control Conference*, 2006.

[85] Johan Reimann and George Vachtsevanos. UAVs in urban operations: Target interception and containment. *Journal of Intelligent Robotics Systems*, 47(4):383–396, 2006.

[86] Johan M. Reimann. *Using Multiplayer Differential Game Theory to Derive Efficient Pursuit-Evasion Strategies for Unmanned Aerial Vehicles*. PhD thesis, Georgia Institute of Technology, 2007.

[87] James A. Ross, Brian R. Geiger, Gregory L. Sinsley, Joseph F. Horn, Lyle N. Long, and Albert F. Niessner. Vision-based target geolocation and optimal surveillance on an unmanned aerial vehicle. In *AIAA Conference on Guidance, Navigation, and Control*, August 2008.

[88] Allison Ryan. Information-theoretic tracking control based on particle filter estimate. In *AIAA Conference on Guidance, Navigation, and Control*, 2008.

[89] Allison Ryan, Hugh Durrant-Whyte, and J. Karl Hedrick. Information-theoretic sensor motion control for distributed estimation. In *ASME International Mechanical Engineering Congress and Exposition*, November 2007.

[90] Rolf Rysdyk. UAV path following for constant line-of-sight. In *AIAA Conference on Guidance, Navigation, and Control*, 2003.

[91] Jeffrey Saunders and Randal W. Beard. Reactive vision based obstacle avoidance with camera field of view constraints. In *AIAA Conference on Guidance, Navigation, and Control*, 2008.

[92] Andrey V. Savkin and Hamid Teimoori. Bearings-only guidance of a unicycle-like vehicle following a moving target with a smaller minimum turning radius. *IEEE Transactions on Automatic Control*, 55(10):2390–2395, 2010.

[93] F C Schweppe. Recursive state estimation: unknown but bounded errors and system inputs. In IEEE *Transactions on Automatic Control*, volume AC(13)-1, pages 22–28, 1968.

[94] D. W. Scott and S. R. Sain. *Multi-Dimensional Density Estimation*, pages 229–263. Elsevier, Amsterdam, 2004.

[95] Dong Gyu Sim, Rae Hong Park, Rin Chul Kim, Sang Uk Lee, and Ihn Cheol Kim. Integrated position estimation using aerial image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1):1–18, 2002.

[96] Per Skoglar. *Planning Methods for Aerial Exploration and Ground Target Tracking*. Licentiate thesis, Linköping University, 2009.

[97] Stephen Spry, Andy Vaughn, Xiao Xiao, and J. Karl Hedrick. A vehicle following methodology for UAV formations. In *International Conference on Cooperative Control and Optimization*, 2003.

[98] Anthony Stentz, Alonzo Kelly, Herman Herman, Peter Rander, Omead Amidi, and Robert Mandelbaum. Integrated air/ground vehicle system for semi-autonomous off-road navigation. In *AUVSI Unmanned Systems Symposium*, July 2002.

[99] Lawrence D. Stone. What's happened in search theory since the 1975 Lanchester prize? *Operations Research*, 37(3):501–506, June 1989.

[100] Ethan Stump, Ben Grocholsky, and Vijay Kumar. Extensive representations and algorithms for nonlinear filtering and estimation. In *Workshop on Algorithmic Foundations in Robotics (WAFR)*, 2006.

[101] P.B. Sujit, Derek Kingston, and Randy Beard. Cooperative forest fire monitoring using multiple uavs. In *Conference on Decision and Control*, December 2007.

[102] Panagiotis Theodorakopoulos and Simon Lacroix. A strategy for tracking a ground target with a UAV. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008.

[103] John Tisdale, Hugh Durrant-Whyte, and J. Karl Hedrick. Path planning for cooperative sensing using unmanned vehicles. In *ASME International Mechanical Engineering Congress and Exposition*, November 2007.

[104] John Tisdale, Zu Kim, and J. Karl Hedrick. An autonomous system for cooperative search and localization using unmanned vehicles. In *AIAA Conference on Guidance, Navigation, and Control*, 2008.

[105] John Tisdale, Allison Ryan, Zu Kim, David Törnqvist, and J. Karl Hedrick. A multiple UAV system for vision-based search and localization. In *American Control Conference*, 2008.

[106] Carlo Tomasi and Takeo Kanade. Detection and tracking of point features. Technical Report CMU-CS-91-132, Carnegie Mellon Robotics Institute, Pittsburgh, PA, 1991.

[107] Benjamin Tovar and Steven M. LaValle. Visibility-based pursuit-evasion with bounded speed. *International Journal of Robotics Research*, 27(11-12):1350–1360, November/December 2008.

[108] Emanuele Trucco and Alessandro Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice-Hall, 2002.

[109] United Nations Population Fund. UN state of the world population. `http://web.unfpa.org/swp/2007/english/chapter_1/urbanization.html`, 2007.

[110] United States Air Force. RQ-4 Global Hawk printable fact sheet. `http://www.af.mil/information/factsheets/factsheet.asp?id=13225`, 2009.

[111] United States Census Bureau. Topologically integrated geographic encoding and referencing (TIGER) system. `http://www.census.gov/geo/www/tiger/`, September 2011.

[112] Rene Vidal, Omid Shakernia, H. Jin Kim, David Hyunchul Shim, and Shankar Sastry. Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation. *IEEE Transactions on Robotics and Automation*, 18(5):662–669, 2002.

[113] Patrick Vincent and Izhak Rubin. A framework and analysis for cooperative search using UAV swarms. In *ACM Symposium on Applied Computing*, 2004.

[114] Yoko Watanabe, Patrick Fabiani, and Guy Le Besnerais. Air-to-ground target tracking in a gps-denied environment using optical flow estimation. In *AIAA Conference on Guidance, Navigation, and Control*, 2009.

[115] Richard A. Wise and Rolf T. Rysdyk. UAV coordination for autonomous target tracking. In *AIAA Conference on Guidance, Navigation, and Control*, 2006.