



# Tasking and Estimation for Minimum-Time Space Object Search and Recovery

Samuel J. Fedeler<sup>1</sup> · Marcus J. Holzinger<sup>1</sup> · William W. Whitacre<sup>2</sup>

Accepted: 9 June 2022 / Published online: 15 July 2022  
© The Author(s), under exclusive licence to American Astronautical Society 2022

## Abstract

A critical challenge in Space Domain Awareness (SDA) is gaining custody of newly detected or maneuvering objects. The dynamically feasible location of one such object may be represented as an admissible region or reachable set. This paper proposes techniques for exploring search sets using sample-based planning. Several methods for sampling feasible regions of measurement space are developed, and Monte Carlo Tree Search (MCTS) is applied over a limited horizon to determine time-optimal sets of actions. Tree search aids in overcoming local minima in solutions found with traditional optimization methods. With tasking methodologies in place, a second contribution of this paper explores estimation methodologies in the presence of null detections; these developments are critical to the search and recovery problem, in which the target probability density projected to measurement space is generally much larger than the sensor field of view. A special focus is applied to the admissible region, and given a Gaussian mixture representation of the admissible region, a novel methodology for splitting mixands in an arbitrary measurement space is presented. A mixand weight update is derived for the key scenario in which no detection is made at a measurement epoch. Merging methodologies are applied, and the resultant Gaussian sum filter is demonstrated for a representative case in which follow-on tracking of a geostationary object is desired.

**Keywords** Space domain awareness · Sensor tasking · Monte Carlo Tree Search · Gaussian sum filter · Negative information

---

This article belongs to the Topical Collection:  
Advanced Maui Optical and Space Surveillance Technologies (AMOS 2021)  
Guest Editors: Lauchie Scott, Ryan Coder, Paul Kervin, Bobby Hunt.

---

✉ Samuel J. Fedeler  
samuel.fedeler@colorado.edu

Extended author information available on the last page of the article

## 1 Introduction

Space domain awareness may be defined as “actionable knowledge required to predict, avoid, deter, operate through, recover from, and attribute cause to the loss and degradation of space capabilities and services” [1]. To maintain and gain knowledge of the local space environment, ground and space-based observers must make timely observations, and critically, decisions on what to observe. The number of space objects (SOs) for which to maintain tracks is quickly increasing, especially in oft-used regions such as the geostationary belt. The European Space Agency maintains tracks on almost 50,000 objects as of 2019 [2], and this figure is expected to greatly increase on the order of 200,000 objects over a period of years as satellite deployment costs continue to decrease. Additionally, a growing concern in the geostationary region is the limited availability of orbital slots necessary for satellites to be resolved by ground-based observers to maintain custody. Combined, these new challenges in SDA motivate the efficient utilization of a limited portfolio of observation tools.

Generally, methodologies for driving sensor tasking can be categorized depending on the objective considered. First, one may wish to maintain existing estimates, informing knowledge on a catalog of SOs. A variety of strategies have been proposed assuming a priori knowledge on state estimates and uncertainties. Erwin et al apply linear optimization to form a tasking solution and propose useful quantities for interpreting the value of a tasking decision [3]. This work is extended by Williams et al, using Lyapunov exponents to probe the stability of SO estimates [4]. A variety of approaches have also taken inspiration from the machine learning literature, with techniques such as stochastic gradient ascent [5], asynchronous actor-critic methods [6], and Monte Carlo Tree Search [7]. In each of these methods, the driving goal is determination of an optimal policy for decision making given a large set of candidate observations.

Alternatively, one may wish to generate new state estimates, expanding the set of SOs studied by searching for natural objects, orbiting satellites, or debris. Wide-ranging techniques for this objective exist in literature. Often, long-period stares over an optical field are performed, acting as a sweep through orbital parameter space [8]. Striping methodologies may also be formed in measurement space, and this strategy accommodates optimization [9].

It is also important to note that detections made with optical sensors generally do not fully observe the object state; as a result of this “Too Short Arc” problem, an admissible region (AR) [10] of unobservable ranges  $\rho$  and range rates  $\dot{\rho}$  may be formed. This admissible region is a two-dimensional manifold of feasible pairs  $(\rho, \dot{\rho})$  that may be projected into the six-dimensional state space. Note that this region may be uniformly distributed in range and range rate or probabilistic [11] if measurement uncertainty is incorporated. Gehly et al. leverage the AR methodology in tandem with Finite Set Statistics to approach the tracking problem, representing the admissible region as a Gaussian mixture to be ingested by a Cardinalized Probability Hypothesis Density (CPHD) filter [12]. Methodologies for generating Gaussian mixture representations of admissible regions are introduced

by DeMars and Jah [13]. AR pairs over longer observation intervals may be used for initial orbit determination [14, 15]. These methodologies are not typically used in an online manner, but rather consider large populations of admissible regions generated from detected tracklets over several observation campaigns.

This research's approach to the admissible regions problem is largely motivated by [16] and [17]. In [17], Murphy considers the direct follow-up tasking problem for an admissible region, representing the region as a feasible set that has grown over time in state space. The area of the admissible region is computed over time using high order Taylor series expansions, and the idea of minimizing the search set using the divergence of the set as an observational metric is considered. Simulated annealing in combination with this observational metric is found to achieve some success in minimizing the search space in a time-optimal manner. This approach is effective in situations in which the initial measurement can be considered an uninformative prior that isn't necessarily sufficient to inform generation of a state estimate.

Note that this methodology can be considered analogous to a variety of scenarios, such as the problem of searching a reachable set. Here, an object may have made a maneuver at some prior time, leading to loss of custody of said SO. The region of state space the object can reach over the time horizon considered can then be projected into measurement space in a similar manner to an admissible region. Additionally, one may desire to track an object that already has a large uncertainty. When projected into measurement space, this uncertainty is larger than the sensor field of view, and multiple observations must be taken to ensure the object is detected. Therefore, in this paper the region of interest will be referred to as the search set to generalize for any context. However, several relevant constraints for admissible regions will also be outlined.

This paper extends the developments made by [17], more closely examining optimization methodologies for search set minimization. Monte Carlo Tree Search is utilized as a means for efficient generation of viable action trajectories. It can be applied as a flexible framework that may consider questions of how to best explore the search set with a limited number of actions and how to exhaust the search set in a time-optimal manner.

An additional goal of this research is development of an efficient estimation scheme to be utilized in tandem with the tasking methodology. Of particular interest is whether the spooky effect is apparent, as may be seen in multi-target filters such as the Gaussian mixture CPHD filter [18]. This behavior describes the impact of a missed or null detection in one region of measurement space on the probability hypothesis density arbitrarily far away. A logical extension of this effect, then, is determining how a null detection in a subset of the projected feasible set may affect knowledge on other regions within that volume.

To further introduce this research, Sect. 2 will establish principles of Monte Carlo Tree Search used to probe the search set. Section 3 will develop the mathematical formulations necessary to frame this sample-based search strategy, and Sect. 4 will outline the estimation process and methods for ingesting negative information. Finally, Sect. 5 shall demonstrate the methodology for a scenario in which follow up observation is needed for a geostationary target.

## 2 Monte Carlo Tree Search

Briefly, critical concepts for the Monte Carlo Tree Search approach to this work are outlined. Further reading on the subject can be found in [7, 19–23]. Generally, MCTS can be applied to any sequential decision process, with special focus in game theory literature. This work narrows the problem context to partially-observable Markov decision processes (POMDPs). In this regard, the underlying state studied is the search set considered, and decisions are observations enacted on a subset of the projection of the region into measurement space at a given time.

### 2.1 Partially-Observable Markov Decision Processes

A POMDP can be formally represented by the 7-tuple  $(S, A, T, R, O, H, \Gamma)$  with the following definitions. The problem is defined over a state space  $S$ ; this is a representation of a discrete or continuous space in which the studied system may evolve over  $\mathfrak{R}^n$ . Decisions are made over an action space  $A$ ; again, this space may be either discrete or continuous, with dimension  $\mathfrak{R}^p$ . States evolve with transition probabilities  $T : \mathfrak{R}^n \times \mathfrak{R}^p \rightarrow \mathfrak{R}^n$ . Generally,  $T$  represents the propagation of both states and uncertainties over time, but actions taken may also impact the evolution of states. A reward  $R : \mathfrak{R}^n \times \mathfrak{R}^p \rightarrow \mathfrak{R}^1$  applies an arbitrary objective function to determine the value of an associated change in state and action. The system is observed over the observation space  $O$ , with probabilities defined by  $H$  acting on the current state. As with the state and action spaces,  $O$  may be discrete or continuous over the domain  $\mathfrak{R}^m$ .  $H : \mathfrak{R}^n \rightarrow \mathfrak{R}^m$  is simply a measurement function incorporating uncertainty in some manner. Immediate rewards are favored over distant rewards by the discount factor  $\Gamma \in [0, 1]$ .

For the application of follow-up sensor tasking, these definitions are as follows. The search set is allowed to evolve over six-dimensional state space  $S$  according to two-body orbital dynamics ( $T$ ). Actions taken  $A$  are exposures by an optical observer with a known field of view, centered on a chosen right ascension and declination ( $O$ ) that is feasible given observer state and constraints. Two reward functions  $R$  are considered in this research; one can either choose the decrease in the projected area of the search set in measurement space as a result of an action or the increase in cumulative probability of detection as an objective. This research assumes that an object in the observer field of view shall be detected, but measurement likelihoods may also be incorporated, assuming a probability of detection derived from knowledge of object shape and albedo properties. As the total change in search set area or cumulative probability of detection over the observation period is desired, a discount factor  $\Gamma = 1$  is selected. This choice assigns equal value to future actions, as opposed to some prioritization of the immediate action.

The ultimate goal of a POMDP is the determination of an optimal policy  $\pi$  (sequence of actions  $a_{1:N} \in A$ ) such that the discrete time Bellman equation  $V$  [24] is maximized given an initial belief in states  $b_0$ , where

$$V^\pi(b_0) = \sum_{t=0}^{\infty} \Gamma^t E[R(\mu_t, a_t) | b_0, \pi] \quad (1)$$

and  $\mu_t$  is the state at time  $t$  and  $a_t$  is the associated action. The set of actions may terminate or occur over an infinite horizon. Note the differentiation between immediate reward  $R$  and the value function  $V$  that describes reward over a potentially infinite horizon. Traditional methods of solving decision processes such as value iteration [25] or grid-based algorithms become challenging in practice as observation and belief spaces become large. This led to the development of MCTS and other sample-based planning based methodologies.

## 2.2 Monte Carlo Tree Search

Monte Carlo evaluation allows random actions to be simulated until a valuable result is reached. A search tree is built on which actions branch over the horizon studied. The search tree tracks value of action sets and runs in an anytime manner, allowing for simulation over a desired interval. The following concepts are critical to the understanding of the MCTS methodology in that they provide a brief background on methods for search tree generation and evaluation.

1. **Nodes** Following general data structures terminology we refer to an arbitrary index in the search tree as a node. The search tree is initialized by a root node. Any node may have zero to many child nodes, and a node with no children is referred to as a leaf node. Other than the root node, any node must have a parent node. We also differentiate between action nodes, where a new action is sampled, and observation nodes, where an observation is generated and associated with an action.
2. **Rollout-based planning** Generally, MCTS is applied over a set depth or until the problem at hand is resolved to a terminal state. To explore a large decision space, a methodology to select new actions must be determined. A rollout heuristic as the means to generate a new set of actions from a leaf node in a search tree. The rollout heuristic can take a variety of forms; fully random sequences could be chosen, or system knowledge can be applied to inform the relative value of actions. In any case, actions must be generated until the terminal state or maximal search depth is reached. If a new action is not needed, a prior child node is selected and tree search is recursively simulated from that node. In this problem, several initial rollout heuristics are considered. First, the relative density of the projected set in measurement space is considered. Using a particle representation of the admissible region, it is trivial to approximate a portion of the probability distribution being captured by a given action and apply weights  $\omega_i$  by studying what particles would lie in the associated field of view  $i$ . This methodology, described when referenced further as the probability of detection heuristic, is visualized in Fig. 1a. If a Gaussian mixture representation is utilized, this methodology is easily modified. A Gaussian integral may be computed over the field of view using methodologies like Genz integration [26] or expectation propagation [27].

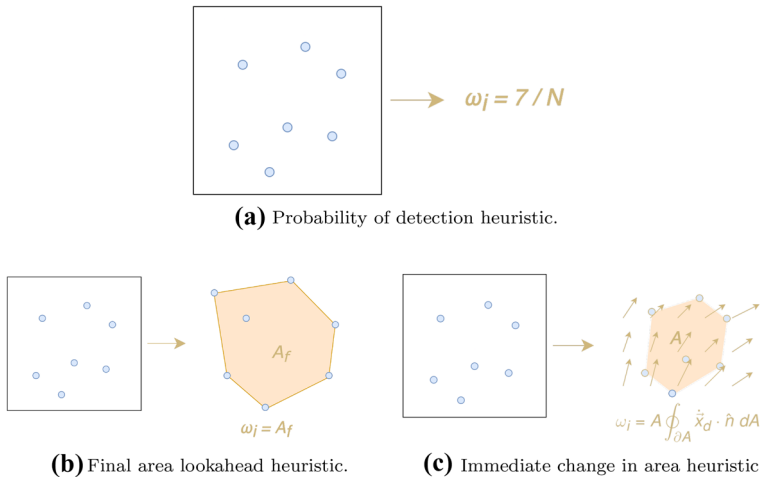


Fig. 1 Rollout heuristic visualizations for MCTS weighting

Alternatively, one can weight actions by considering how a region may change in the future. Two approaches are developed in this regard. Given a desired action, the projected area of all unobserved particles within the field of view at the end of the search period may be considered as a weight from which to sample. This process is shown in Fig. 1b, and in further reference, is described as the lookahead heuristic. In addition, the more immediate change in area of these particles may be considered. The immediate rate of change in area of the region studied may be computed using the divergence of the measurement velocity vector field. This rate of change is then reweighted by the area of the studied region to form a third set of sampling weights, hereby referred to as the immediate change in area heuristic. This weighting scheme is shown in Fig. 1c. Note that this final methodology leads to additional challenges, in which the area rate of change is not strictly positive. Because of this, it is useful to here consider a deterministic rollout method, in which at a given timestep, the  $n$ th new action chosen is the  $n$ th "best" as evaluated by the rollout policy. Further consideration for these policies is given after development of a mathematical basis in the next section. These methodologies may be modified for a Gaussian mixture representation in several ways. First, the mean states of each mixand may be considered a particle for the purposes of sampling. Second, the final area lookahead heuristic may be evaluated to an appropriate level of variance, again utilizing Gaussian integrals.

3. **Backpropagation** - Backpropagation is defined as the means by which immediate rewards simulated by leaf nodes in the search tree impact the estimated value at parent nodes. Given a rollout or simulation routine that returns the discounted cumulative reward  $R_i$  sampled for a sequence of actions, one must determine how to revalue the immediate action taken. Generally, the average reward returned for an immediately sampled action  $a \in A$

$$V(a) = V(a) + \frac{R_i - V(a)}{N(a)}$$

is utilized, but other statistical measures may also be incorporated, especially if there are concerns about the variance of rewards. In this research, since it is assumed that an object lying within the field of view will be detected, there is no stochasticity inherent to that observation. As such, in considering the value of an immediate action, we may instead recursively utilize the maximal value returned from the set of child trajectories through the search tree, at each iteration applying

$$V(a) = \max (V(a), R_i). \quad (2)$$

4. **Selection** If a new action is not generated, one must determine what previously sampled action to take. Generally, the selection method must balance more detailed exploration of simulated actions with high expected value with further exploration of undersampled actions. As such, a deterministic score function is applied for selection such that the child node maximizing

$$sc_n(i) = V(i) + \sqrt{\frac{f(N)}{N(i)}} \quad (3)$$

is selected.  $N(i)$  represents the number of times child node  $i$  was previously selected, and  $f(N)$  is an arbitrary non-decreasing map from  $\mathfrak{R}^1$  to  $\mathfrak{R}^1$ . This second term is derived from multi-armed bandit literature, and can be related to a confidence interval for the true value of an action [28]. Generally, the natural logarithm is utilized, but other methods such as polynomial exploration have been applied [29]. The polynomial map  $f(N) = N^\beta$  is utilized in this work, where  $\beta \in (0, 1)$ .

5. **Progressive widening** Generally, large state and action spaces can lead to curses of dimensionality in decision processes. When state and observation spaces are large or continuous, curses of history can also occur. As actions lead to transitions described by generative models, search trees can become infinitely wide after a single transition; that is, an arbitrary action will lead to a different representation of belief for each associated observation that is sampled. As such, in order to limit the breadth of the search tree, one must artificially limit the number of actions explored, as well as the number of observations associated with each sampled action. This so-called arm-increasing rule or progressive widening is analyzed in [30]. Widening is applied for MCTS by [31] with success; this is the first example of double progressive widening, in which the search tree breadth is slowly widened for both generation of new action sequences and state transition or observation generation. Generally, whether progressive widening is allowed is determined by a rule as a function of visits to the parent node  $i$

$$|i| \leq N(i)^{\beta_d}$$

such that the number of child nodes are upper bounded by a power law  $\beta_d \in (0, 1)$ . Note that  $\parallel$  is utilized to describe the number of children at an arbitrary node.

With these concepts outlined, the tree search routine from a root node is outlined as follows. A detailed visualization of the search algorithm is given in [7]. First, an action is determined using progressive widening. If the tree is allowed to widen, a new action is sampled; otherwise, a previous action is chosen utilizing the selection criterion. If a new observation node is generated, the rollout model is applied. The simulation process is then recursively completed from the selected child node. Finally, cumulative rewards from the rollout or recursive simulation are utilized to update expected reward, and total cumulative reward for the search iteration is returned.

### 3 Search Set Behavior

A more extensive mathematical basis for the evolving behavior of search sets is outlined in Murphy et al. [17]. This section provides a general overview of that work as needed for analysis of the search sets and development of the metrics utilized in this work. The intention of this section is to first consider in a general manner how a projected subset of state space might evolve over time. This analysis is then applied to an optical observer, and knowledge of search set behavior is used to inform sampling methods for MCTS tasking.

Assume that the dynamical system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) \tag{4}$$

is associated with the flow function  $\phi$ .

$$\mathbf{x}(t_1) = \phi(t_1; \mathbf{x}(t_0), t_0) \tag{5}$$

We assume the flow function may be applied to a subset of states,  $\mathcal{S}$ , defined by a  $p$ -dimensional vector of constraint equations,  $\mathbf{c}$ , in the global set of state dimension  $m$ ,  $\mathcal{X} = \mathcal{R}^m$ .

$$\mathcal{S}(t) = \phi(t; \mathcal{S}(t_0), t_0) \tag{6}$$

$$\mathcal{S}(t_0) = \{ \mathbf{x} \in \mathcal{X} : c_i(\mathbf{x}) \leq 0, i = 1 : p \} \tag{7}$$

Sets are analyzed in the measurement space  $\mathcal{H}_o$  of dimension  $n$  associated with observer  $o$ . At a given time, the search set may be projected onto the field of regard of the observer using the measurement function

$$\mathbf{h} : \mathcal{R}^m \rightarrow \mathcal{R}^n \tag{8}$$

Note that the full state  $\mathbf{x}$  may be partitioned into a determined subset  $d$  and unobservable subset  $u$ .



$$\mathbf{x} = [\mathbf{x}_d \ \mathbf{x}_u] \quad (9)$$

The partition  $\mathbf{x}_d$  is representative of the subset of the full state that is observable through  $\mathbf{h}$ . The observable portion of the set is described as

$$\mathcal{S}_d = \{\mathbf{x}_d : \mathbf{x}_d = \mathbf{h}(\mathbf{x}), \mathbf{x} \in \mathcal{S}\} \quad (10)$$

Of particular interest for this work is the area of the projected set and how this changes over time. Leveraging the representation of the projected set as a vector field, one may apply Gauss's theorem. For an arbitrary vector field  $\mathbf{F}$  defined over a compact subset of  $\mathcal{R}^n$ ,  $\mathcal{T}$ , with a piecewise smooth boundary  $\partial\mathcal{T}$ , the theorem relates the flux of that vector field through the closed surface to the divergence of the vector field within the region. Note that a  $n$ -dimensional integral over a subset is designated a volume integral  $dV$ , while a  $n - 1$ -dimensional integral over the boundary of that subset is designated a surface integral  $dS$ .

$$\int_{\mathcal{T}} \nabla \cdot \mathbf{F} dV = \oint_{\partial\mathcal{T}} (\mathbf{F} \cdot \hat{\mathbf{n}}) dS \quad (11)$$

Using this theorem, one may consider how the projected area of the search set will change over time. The area of the region of interest may be expressed as the total integral,

$$A_{\mathbf{h}}(\mathcal{S}) = \int_{\mathcal{S}_d} dV = \frac{1}{n} \int_{\mathcal{S}_d} \nabla \cdot \mathbf{x}_d dV = \frac{1}{n} \oint_{\partial\mathcal{S}_d} (\mathbf{x}_d \cdot \hat{\mathbf{n}}) dS \quad (12)$$

where  $n$  is the measurement space dimensionality and  $\partial\mathcal{S}_d$  is the boundary of the projected set. Following this result, the Leibniz integral rule may be applied to take time derivatives to arbitrary order. In general, the following result is found.

$$\frac{d^n}{d^n t} A_{\mathbf{h}}(\mathcal{S}) = \oint_{\partial\mathcal{S}_d} \left( \sum_{i=1}^n \mathbf{x}_d^{(i)} \cdot \frac{\partial \mathbf{x}_d^{(n-i)}}{\partial \mathbf{x}_d} \right) \cdot \hat{\mathbf{n}} dS \quad (13)$$

This result may be utilized to describe the area of a subset of the projected set  $\mathcal{S}$  at an arbitrary epoch. Murphy discusses additional considerations that must be made for this projected set. Generally the projection itself is not necessarily an injective function, and  $\mathcal{S}$  may be "folded" in a variety of manners, such that multiple points in the unobservable set may map to the same point in  $\mathcal{S}_d$ . The supremum of flux out of a boundary associated with this set of feasible states is necessarily applied to describe changes in the search set over time. Fortunately, at least in the case of admissible regions, the projection may be expected to be injective in the short term, when follow-up observation is desired.

### 3.1 Optical Applications

The developed background is now applied for an optical observer. It is assumed that an admissible region-based search set  $\mathcal{A}$  is previously defined at time  $t_0$ , and right ascension and declination measurements are desired at time  $t$ . This admissible region is a two-dimensional manifold in the range ( $\rho$ ) and range-rate ( $\dot{\rho}$ ) half plane projected into six-dimensional state space. Note that the admissible region is formed from an attributable vector with initially determined angles and angular rates. It is most useful to define the problem as follows. Right ascension is defined as  $\alpha$  and declination is defined as  $\delta$ .

$$\mathbf{x}_d = [\alpha \ \delta]^T \tag{14}$$

$$\mathcal{A}_d(t) = \left\{ [\alpha \ \delta]^T : [\alpha \ \delta]^T = \mathbf{h}(\phi(t; \mathbf{x}, t_0); \mathbf{o}(t)), \mathbf{x} \in \mathcal{A} \right\} \tag{15}$$

The divergence of the velocity vector field for this sensor  $\dot{\mathbf{x}}_d = [\dot{\alpha} \ \dot{\delta}]^T$  may be expressed as

$$\nabla \cdot \dot{\mathbf{x}}_d = \frac{d\dot{\alpha}}{d\alpha} + \frac{d\dot{\delta}}{d\delta} \tag{16}$$

The area of the set in the sensor field of regard is computed as

$$A_{\mathbf{h}}(\mathcal{A}) = \frac{1}{2} \oint_{\partial \mathcal{A}_d} (\mathbf{x}_d \cdot \hat{n}) d\mathcal{A} = \frac{1}{2} \oint_{\partial \mathcal{A}_d} -\delta \cos(\delta) d\alpha + \alpha d\delta \tag{17}$$

Assuming the set is projected as a grid, set of particles or triangulation, this line integral may be computed numerically over the surface bounds. The area rate of change reflects a similar pattern and is computed as

$$\frac{d}{dt} A_{\mathbf{h}}(\mathcal{A}) = \oint_{\partial \mathcal{A}_d} (\dot{\mathbf{x}}_d \cdot \hat{n}) d\mathcal{A} = \oint_{\partial \mathcal{A}_d} -\dot{\delta} \cos(\delta) d\alpha + \dot{\alpha} d\delta \tag{18}$$

These results may be extended to compute higher order time derivatives of the search set area. It is noted that these higher order derivatives are explicitly dependent on the dynamical system associated with the problem, and become more challenging to compute. The second order derivative is generalized in Eq. 19 as

$$\frac{d^2}{dt^2} A_{\mathbf{h}}(\mathcal{A}) = \oint_{\partial \mathcal{A}_d} \frac{\partial}{\partial t} (\dot{\mathbf{x}}_d \cdot \hat{n}) d\mathcal{A} = \oint_{\partial \mathcal{A}_d} \left( \ddot{\mathbf{x}}_d \cdot \hat{n} + \dot{\mathbf{x}}_d \cdot \frac{\partial \dot{\mathbf{x}}_d}{\partial \mathbf{x}_d} \cdot \hat{n} \right) d\mathcal{A}. \tag{19}$$

Here, both  $\ddot{\mathbf{x}}_d$  and  $\dot{\mathbf{x}}_d \cdot \frac{\partial \dot{\mathbf{x}}_d}{\partial \mathbf{x}_d}$  are a function of the full state at any point in the projected set.  $\ddot{\mathbf{x}}_d$  may be found analytically, but requires acceleration information from the full state, and further analysis is needed for the second term. The variation of the determined subset at time  $t$  with respect to initial admissible region coordinates may be expressed as

$$\frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}_u(t_0)} = \frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}(t)} \frac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_0)} \frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \Big|_{\mathbf{x}_u(t_0)} \quad (20)$$

$$= H\Phi(t, t_0) \frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \Big|_{\mathbf{x}_u(t_0)} \quad (21)$$

where  $H$  is the measurement Jacobian and  $\Phi(t, t_0)$  is the state transition matrix. Similarly,

$$\frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}_u(t_0)} = \frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}(t)} \frac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_0)} \frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \Big|_{\mathbf{x}_u(t_0)} \quad (22)$$

$$= \dot{H}\Phi(t, t_0) \frac{\partial \mathbf{x}(t_0)}{\partial \mathbf{x}_u(t_0)} \Big|_{\mathbf{x}_u(t_0)} \quad (23)$$

and  $\dot{H}$  is the Jacobian of the derivative of the measurement function. The desired matrix may then be found as

$$\frac{\partial \dot{\mathbf{x}}_d}{\partial \mathbf{x}_d} = \frac{\partial \dot{\mathbf{x}}_d(t)}{\partial \mathbf{x}_u(t_0)} \left( \frac{\partial \mathbf{x}_d(t)}{\partial \mathbf{x}_u(t_0)} \right)^{-1} \Big|_{\mathbf{x}_u(t_0), t} \quad (24)$$

These derivatives may then be applied to a Taylor series expansion to provide an analytic solution as an infinite series or an approximate series for the area of the search set projection over time. [17] provides analysis for the error in area computation as a function of the number of terms utilized in the expansion, as well as a generalization for higher order derivatives.

The strategies in this paper for tracking the growth of the search set over time benefit from this theory and apply a particle representation of the admissible set. Equation 18 may be explicitly computed for subsets of the region using knowledge of particles over time, and long term growth of subsets of the region may be observed by tracking particles over time. In using this methodology, this work diverges from the strategy in [17] requiring computation of high order derivatives, but recognizes the clear application of these analytic concepts to developing search heuristics.

In a particle formulation, the rollout heuristics previously discussed may be implemented in a relatively straightforward manner. Given sufficient sampling, the particle representation is representative of the initial uniform probability distribution over the unobservable subset of state space. In propagating particles over time, one may determine how this distribution and as a result, the projection into measurement space, evolves. The net probability of detection for a given action may be determined by the number of unobserved particles captured in the region chosen. Determining areas and rates of change in area for particle sets requires formulations of shape representations of particles. A broad set of algorithms are available for this problem, and here, alpha shapes are utilized to better account for non-convex deformations of the search set as observations are made and particles within the observed region are removed from the tracked shape. An alpha shape covers a set of points,

and may be conceptualized as the intersection of a set of closed discs with a specified  $\alpha$ -radius containing each point [32].

Once particles are determined to be within the field of view associated with a specified action, the projection of these particles at the end of the search period may be used to form an alpha shape, and the area of the shape may be used to drive weighting for MCTS. When a shape is generated, the boundary edges are found and used to compute instantaneous rates of change of search regions.

## 4 Evolving Estimates of Search Sets

In order to instantiate an estimation scheme, one must first consider how to represent the probability density function (PDF) over the feasible set within which the target lies. The probability density over this region may be uniform [10] or non-uniform, as in the case of the probabilistic admissible region [11] or a reachable set generated with some a priori knowledge on maneuver probability. It is clear that it is illogical for the probability density to be represented by a univariate Gaussian, and that the density is typically quite non-Gaussian. Particle representations may be considered, but these methods suffer from a curse of dimensionality and quickly become computationally expensive. Thus, the DeMars method [13] for instantiating a Gaussian mixture representation of probability density for an admissible region is considered in further detail.

### 4.1 Gaussian Mixture Approximation for Admissible Regions

Broadly, the AR approximation process may be considered an extension of the problem of approximating a univariate uniform distribution. This approximation may be solved as a root finding problem with several constraints on the structure of mixands. Assuming equal variance  $\sigma^2$ , equal weights  $\omega = \frac{1}{L}$  for  $L$  mixands, and evenly distributed mixand means  $\mu$ , an optimization may be performed to determine optimal variances. Specifically, the  $L_2$  distance [13] between the mixture PDF  $q$  and a uniform PDF  $p$  with support on the closed interval  $[a, b]$  is

$$L_2[p||q] = \frac{1}{b-a} + \frac{\omega^2}{2\sqrt{\pi}\sigma} \sum_{i=1}^L \sum_{j=1}^L \exp\left(-\frac{1}{4}\left(\frac{\mu_i - \mu_j}{\sigma}\right)^2\right) - \frac{\omega}{b-a} \sum_{k=1}^L (\text{erf}(B_k) - \text{erf}(A_k)) \tag{25}$$

and

$$A_k = \frac{a - \mu_k}{\sqrt{2}\sigma} \quad B_k = \frac{b - \mu_k}{\sqrt{2}\sigma} \tag{26}$$

The derivative of the  $L_2$  distance with respect to standard deviation may then be explicitly or numerically computed using any optimizer. In this work, the

Levenberg-Marquardt algorithm is utilized with support from the Eigen template library.

To extend this methodology to the two-dimensional problem of approximation for admissible regions successive one-dimensional approximations are applied. First, consider that at any range within the AR, the marginalized density over range rate is uniform. However, the range-marginal PDF, computed as the integral over range rate, is generally not uniform. To account for this, a further optimization may be performed on mixand weights once the univariate uniform distribution has been approximated over the support of the range-marginal PDF. In order to perform this optimization,  $M$  ranges must be sampled over the support of the range-marginal PDF. One may then compute the  $i \times j$  likelihood matrix  $\Lambda$ , where  $\Lambda(i, j)$  is the likelihood range  $i$  is drawn from mixand  $j$ . The vector  $\mathbf{p}$  is then computed by evaluating the range-marginal PDF at each sampled range. A least-squares problem may then be formulated as

$$\min J = \|\mathbf{p} - \Lambda\boldsymbol{\omega}\| \text{ subject to } \boldsymbol{\omega} \geq \mathbf{0} \text{ and } \mathbf{1}^T\boldsymbol{\omega} = 1 \quad (27)$$

where  $\boldsymbol{\omega}$  is the concatenated set of weights.

Finally, taking the mean range at each mixand in the approximated range-marginal PDF, range rate may be incorporated. Considering range rate as an independent random variable, standard deviation in range rate may be optimized using the same process. Mixands may be augmented with no correlation between range and range rate uncertainties. Note that care is taken to determine the number of mixands  $L$  to utilize in each case as a function of minimum desired variance in range and range rate. Further detail on this process may be found in [13]. The resultant mixands, augmented with the measurement utilized to generate the AR  $\mathbf{y}$  with measurement uncertainty  $R$ , may then be linearly transformed into state space. Note that the ensuing subsection is meant as a summary of this methodology.

## 4.2 Measurement updates in the presence and absence of observations

Given the generated set of mixands, one may now consider how mixands are updated when tasking decisions are made and measurements are taken. The simpler scenario is that of a newly made detection. Here, a typical measurement update for a Gaussian sum filter may be performed.

It is important to first note the measurement model utilized in this work. Optical sensors make angular detections on the celestial sphere, and during long exposures, angular rates may also be determined. Right ascension and declination measurements are computed as

$$\alpha = \arctan\left(\frac{\rho_y}{\rho_x}\right). \quad (28)$$

$$\delta = \arcsin\left(\frac{\rho_z}{\|\boldsymbol{\rho}\|}\right) \quad (29)$$

and

$$\boldsymbol{\rho} = \mathbf{r} - \mathbf{o} \tag{30}$$

is the expected relative position in an inertial frame.  $\mathbf{r}$  is the object position, and  $\mathbf{o}$  is the observer position. Since the angular measurements are explicitly a function of  $\boldsymbol{\rho}$ , one may apply the measurement jacobian  $H = \frac{d\mathbf{y}}{d\boldsymbol{\rho}}$ , where  $\mathbf{y} = [\alpha \ \delta]$ , to compute angular rates. Then,

$$[\dot{\alpha} \ \dot{\delta}] = H\dot{\boldsymbol{\rho}}$$

$$H = \begin{bmatrix} -\frac{\rho_y}{|\rho_{xy}|^2} & \frac{\rho_x}{|\rho_{xy}|^2} & 0 \\ -\frac{\rho_x\rho_z}{|\boldsymbol{\rho}|^2|\rho_{xy}|} & -\frac{\rho_y\rho_z}{|\boldsymbol{\rho}|^2|\rho_{xy}|} & \frac{|\rho_{xy}|}{|\boldsymbol{\rho}|^2} \end{bmatrix} \tag{31}$$

$$|\rho_{xy}| = (\rho_x^2 + \rho_y^2)^{\frac{1}{2}}$$

and the expected relative velocity vector  $\dot{\boldsymbol{\rho}}$  is easily computed given knowledge of the observer trajectory and the object state estimate. With angular uncertainty  $R_\theta$ , angular rate uncertainty is explicitly  $R_{\dot{\theta}} = \frac{2R_\theta}{\Delta t^2}$  for exposure time  $\Delta t$  [7].

Given measurement and measurement uncertainty

$$\mathbf{y} = [\alpha \ \delta \ \dot{\alpha} \ \dot{\delta}]^T, \quad R = \begin{bmatrix} R_\theta & 0 \\ 0 & R_{\dot{\theta}} \end{bmatrix} \tag{32}$$

the traditional Gaussian sum filter updates may be performed [33]. In this work, unscented measurement updates are utilized [34], with likelihood-based weight updates in addition such that

$$\omega_i = \frac{\omega'_i \mathcal{L}_i(\mathbf{y}, R)}{\sum_{j=1}^L \omega'_j \mathcal{L}_j(\mathbf{y}, R)} \tag{33}$$

and

$$\mathcal{L}_i(\mathbf{y}, R) = \mathcal{N}(\mathbf{y} - \mathbf{h}(\mathbf{x}_i), HP_iH^T + R). \tag{34}$$

Updating the PDF when no detection is made introduces further complexities. First, one must consider the probability that the target SO lies in a given field of view (FOV) of an optical sensor. If mixand uncertainties are sufficiently small in range space relative to the range between the mixand and the optical sensor, one may assume that the probability of detection  $p_D$  is uniform over the mixand. Then, the cumulative likelihood of observing the target is

$$P(\mathbf{y} \neq 0) = \int_{\text{FOV}} p_D(\mathbf{z})P(\mathbf{z})d\mathbf{z} \tag{35}$$

Applying the mixand representation of the PDF, a distinct probability of detection may be assumed for each mixand, and

$$P(\mathbf{y} \neq 0) = \int_{\text{FOV}} \sum_{i=1}^L p_{D,i} \omega_i P(\mathbf{z}|k=i) d\mathbf{z} \quad (36)$$

One must then evaluate the observation likelihood  $P(\mathbf{z}|k=i)$  conditioned on mixand  $i$ . Note that the transformation from state space to measurement space is locally linear with the same assumptions on range. Therefore, the mixand density projected into measurement space is still Gaussian, with

$$P(\mathbf{z}|k=i) \approx \mathcal{N}(\mathbf{h}(\mu_i), HP_i H^T). \quad (37)$$


It is also clear that the integral is separable, and thus,

$$P(\mathbf{y} \neq 0) = \sum_{i=1}^L \left( p_{D,i} \omega_i \int_{\text{FOV}} P(\mathbf{z}|k=i) d\mathbf{z} \right) \quad (38)$$

To compute the likelihood of observation, one must then evaluate the Gaussian integral in measurement space over the rectangular FOV. Methods such as Genz integration [26] or expectation propagation [27] may be applied for this purpose. Given this result, one must now consider how a null detection may affect existing mixands. It is clear that the probability a null detection occurs is the complement of Eq. 38. Then, one must consider how to apply this result to knowledge on each mixand. Working from first principles, we may apply Bayes rule.

$$P(\mathbf{x}|k=i, \mathbf{y} = \emptyset) = \frac{P(\mathbf{y} = \emptyset|\mathbf{x}, k=i)P(\mathbf{x}|k=i)}{P(\mathbf{y} = \emptyset)} \quad (39)$$

Immediately, challenges arise when considering the term  $P(\mathbf{y} = \emptyset|\mathbf{x}, k=i)$ , the probability a null detection is made, conditioned on the SO state captured by mixand  $i$ . Consider the PDF for mixand  $i$  in further detail, with the temporary assumption that the projection of probability density into measurement space is larger in spread than the sensor field of view. At any point within the support of the projected PDF outside of the field of view, the probability of a null detection must be unity, since that point simply cannot be captured during the observation. This leads to a scenario in which a subset of the mixand is scaled as a function of the probability of detection, while the remainder is unaffected. The structure of this update is clearly non-Gaussian, and is further illustrated in Fig. 2.



**Fig. 2** Non-Gaussianity in a negative information update

This behavior may be broadly categorized into distinct groups. First, the density captured in the field of view can be relatively small; this commonly occurs when the mixand is either a large normalized distance away from the sensor FOV in measurement space, or quite large in spread in measurement space as compared to the field of view. In this case, a null detection’s impact on the PDF is negligible, since the probability  $P(\mathbf{y} = \emptyset | \mathbf{x}, k = i)$  is effectively unity. Alternatively, the probability density may be almost entirely captured by the sensor FOV, in which case the entire mixand is rescaled and the null detection probability for the given mixand nears zero. The third case, in which the sensor field of view overlays a significant portion of the mixand, but not the entirety, requires further consideration. To resolve this case, a splitting method is proposed to reduce the projected spread of mixands in measurement space and ensure negative information updates remain Gaussian.

### 4.3 Oriented Gaussian splitting

Consider a mixand with mean  $\mu$  and covariance  $P$  in state space  $\mathcal{S}$ . This mixand may be defined relative to an observer  $\mathcal{O}$  with measurement function  $\mathbf{h}$ . Letting the measurement function be differentiable, the local behavior of  $\mathbf{h}$  may be examined utilizing the gradient. For each scalar measurement, this leads to a tangent that may be normalized and considered as the direction in state space leading to a maximal change in the associated scalar measurement. The resultant set of tangent vectors forms the basis of a tangent space of dimension  $n$ , where  $n$  is the rank of the measurement Jacobian. Each measurement tangent vector may be explicitly computed as

$$\hat{l}_i = \frac{\frac{\partial h_i}{\partial \mathbf{x}} |^T}{\left| \frac{\partial h_i}{\partial \mathbf{x}} |^T \right|} \mu \tag{40}$$

It is desired to split the mixand into a set of mixands with unknown means and equivalent covariance  $P$ , while enforcing that the combined PDF of the resultant mixands captures the same first and second moment of the original mixand. Additionally, it is desired that the observer is also taken into consideration, such that the new mixands are perturbed about the defined measurement bases. Without loss of generality, let the new set contain  $2m + 1$  mixands, where  $m$  is the dimension of the measurement space. Let one mixand be placed at the original mean, with another pair of mixands evenly distanced along the the tangent vector associated with each scalar measurement with distance  $a_k P^{\frac{1}{2}}$ , where  $P^{\frac{1}{2}}$  is the matrix square root of the original mixand covariance. Note that this methodology is typical when considering transformations on Gaussians, exemplified by the work of Havlak and Campbell [35]. The matrix square root is incorporated to ensure that similar mixand structure is in place, and to enforce positive-definiteness. Additionally, let each new mixand have equivalent weight  $\omega = \frac{1}{2m+1}$ . The mean of the resultant distribution is then



$$\begin{aligned} \mu_{\text{TOT}} &= \sum_{k=0}^{2m} \omega_k \mu_k \\ &= \frac{1}{2m+1} \left( \mu + \sum_{k=1}^m \left( \mu + a_k P^{\frac{1}{2}} \hat{l}_k \right) + \sum_{k=1}^m \left( \mu - a_k P^{\frac{1}{2}} \hat{l}_k \right) \right) = \mu. \end{aligned} \tag{41}$$

The covariance of the new distribution must also be determined. For a Gaussian mixture, the total covariance is

$$P_{\text{TOT}} = \sum_{i=1}^N \omega_i P_i + \sum_{i=1}^N \omega_i (\mu_i - \mu_{\text{TOT}})(\mu_i - \mu_{\text{TOT}})^T. \tag{42}$$

In this case, then, we find

$$\begin{aligned} P_{\text{TOT}} &= \sum_{k=0}^{2m} \omega_k P^* + \sum_{k=0}^{2m} \omega_k (\mu_k - \mu)(\mu_k - \mu)^T. \\ &= P^* + 2 \sum_{k=1}^m \frac{1}{2m+1} a_k^2 P^{\frac{1}{2}} \hat{l}_k \hat{l}_k^T P^{\frac{T}{2}} = P \end{aligned} \tag{43}$$

With this result, one can determine the updated covariance

$$P^* = P - 2 \sum_{k=1}^m \frac{1}{2m+1} a_k^2 P^{\frac{1}{2}} \hat{l}_k \hat{l}_k^T P^{\frac{T}{2}} \tag{44}$$

With this result in mind, it is still important to consider whether  $P^*$  is positive definite. The key determination is whether the eigenvalues of  $P^*$  are all positive. It is possible to left and right multiply  $P^*$  by  $P^{-\frac{1}{2}}$  and maintain the definiteness of the matrix such that

$$\begin{aligned} P_{\text{NORM}} &= P^{-\frac{1}{2}} P^* P^{-\frac{T}{2}} \\ &= I - 2 \sum_{k=1}^m \frac{1}{2m+1} a_k^2 \hat{l}_k \hat{l}_k^T \end{aligned} \tag{45}$$

Now, any eigenvector for the summation must also be an eigenvector of  $P_{\text{NORM}}$ . For each eigenvector  $\mathbf{v}_i$ , the associated eigenvalue of the summation is  $\lambda_i$ , and the eigenvalue of  $P_{\text{NORM}}$  must be enforced to be strictly greater than zero such that

$$\lambda_p = 1 - \lambda_i > 0 \tag{46}$$

This allows for gains  $a_k$  to be chosen as a function of the structure of the cumulative outer product of the tangent space. First, consider the outer product  $\hat{l}_k \hat{l}_k^T$ . Since the tangent vectors are normalized, this matrix is symmetric and positive semi-definite with a single eigenvalue at unity. This offers an upper bound on the eigenvalues of the summation. If each tangent vector is collinear, the summation will then have a single nonzero eigenvalue

$$\lambda^* = \frac{2ma_k^2}{2m + 1} \tag{47}$$

that must be less than unity. Gains must in general then be no greater than

$$a_k < \sqrt{\frac{2m + 1}{2m}}. \tag{48}$$

Note that these gains may be increased if there is further knowledge of the measurement space. If two unit vectors are orthogonal, one may infer that the sum of outer products of these vectors has two eigenvalues at unity in addition to zero eigenvalues. This argument may be expanded, considering the full set of tangent vectors in the space. The maximum eigenvalue of the summed matrix must be no greater than, assuming gain is held constant,

$$\lambda^* = \frac{2a_k^2}{2m + 1}(m + 1 - \text{rank}(H)) \tag{49}$$

where  $\text{rank}(H)$  is the rank of the measurement Jacobian, which is equivalent to the rank of the set of tangent vectors.

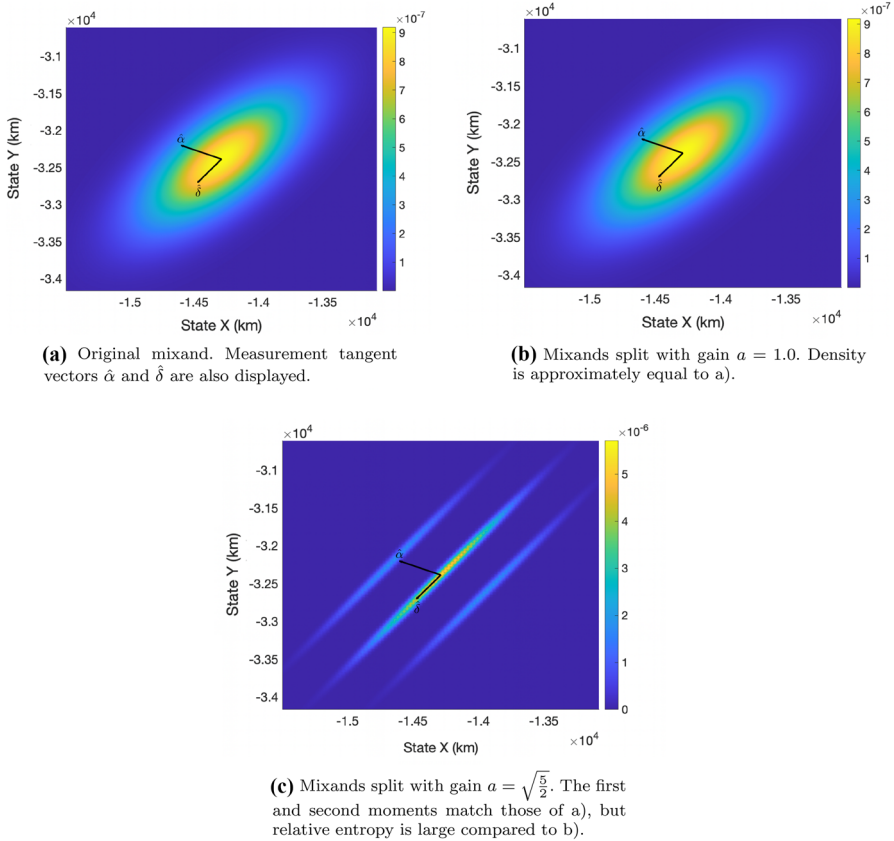
This result may explicitly be demonstrated for an optical case in which right ascension and declination measurements are taken. This is the critical case for negative information updates, because projected mixands must be split in angular space to ensure the update remains Gaussian. For this case, the dimension of the measurement is  $m = 2$ . It is also known that the gradients of right ascension and declination are orthonormal in state space. Therefore, we have

$$\lambda^* = \frac{2a_k^2}{5} < 1 \tag{50}$$

and

$$a_k < \sqrt{\frac{5}{2}}. \tag{51}$$

It is then ensured that newly generated mixands continue to have positive definite covariances as splitting occurs. Note that while this gain is the theoretical maximal bound to ensure positive definite covariances, it is not advisable to use. As gain nears this quantity, the projected PDF of the split mixands will become drastically small in the eigendirection associated with the eigenvalues nearing zero. This behavior is visualized in Fig. 3, where the PDF is presented for a mixand before a split, with a split using reasonable gains, and a split with gains nearing the theoretical maximum. With this projection, it is most clear that mixand uncertainties become quite small in the  $\hat{\alpha}$  direction, and they also become quite small in the  $\hat{\delta}$  direction (which is largely out of the plane). More specifically, as the gain approaches the theoretic maximum, the projections of uncertainties into



**Fig. 3** Mixand densities in a subset of position space before and after splitting

measurement space collapse into discrete points. This behavior is demonstrated in Fig. 4, where mixands are projected into measurement space.

Understanding this behavior helps visualize a trade in gain selection. Increased gain ensures a reduction in spread for the mixands utilized, but also increases relative entropy between the original mixand and resultant mixand. Care must be taken to allow new mixands to become as small as needed while maintaining a distribution sufficiently close to the original.

### 4.4 Updating Gaussians

Now that a formulation for splitting Gaussians such that they will be sufficiently small in measurement space is outlined, a criterion for determining whether a mixand shall be split must be established. It is logical to incorporate some measure of offset from the center of the sensor FOV in measurement space and the comparative spread of uncertainty to the sensor FOV. A critical Mahalanobis distance may be defined, such that a mixand shall only be split if

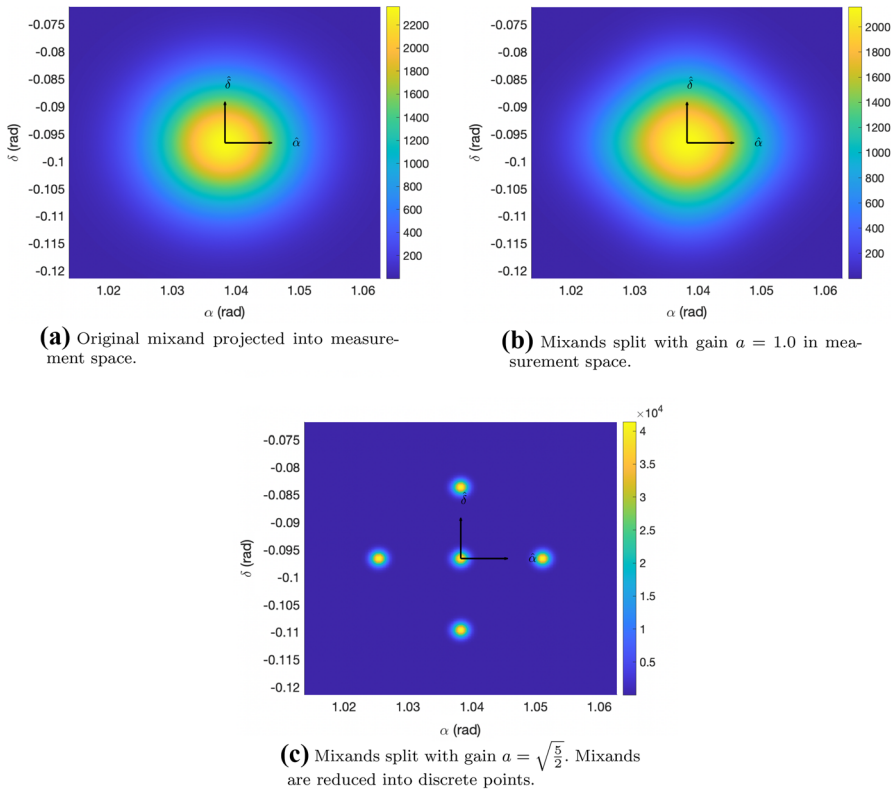


Fig. 4 Mixand densities in measurement space before and after splitting

$$D_M(\mu_i, P_i; \mathcal{O}) < d^*. \tag{52}$$

Additionally, the angular spread may be defined as the square root of the maximal diagonal value of the projected covariance trace

$$s = \sqrt{\max(HPH^T)} > s^*. \tag{53}$$

This may be compared with a critical value that is a function of the diagonal field of view of the sensor.

Once mixands are rescaled such that they are either sufficiently distant in measurement space from the sensor FOV or of comparable size to the sensor FOV, the negative information update may be considered in further detail. Revisiting Eq. 38, the problematic term may now be considered approximately discrete in that the mixand is either fully covered by the sensor FOV or is sufficiently far from the sensor. As such, it is now logical to consider Eq. 38 as a weight update on each mixand in much the same manner as a particle filter, using the intermediate density

$$\begin{aligned}
 g(\mathbf{y}|k = i) &= \frac{P(\mathbf{y} = \emptyset|\mathbf{x}, k = i)}{P(\mathbf{y} = \emptyset)} \\
 &= \frac{P(\mathbf{y} = \emptyset|\mathbf{x}, k = i)}{\sum_{j=1}^L P(\mathbf{y} = \emptyset|\mathbf{x}, k = j)}
 \end{aligned}
 \tag{54}$$

The denominator can simply be considered a normalization, while the numerator may be approximately evaluated as

$$P(\mathbf{y} = \emptyset|\mathbf{x}, k = i) \approx \begin{cases} 1 & i \notin \text{FOV} \\ 1 - p_D & i \in \text{FOV} \end{cases} .
 \tag{55}$$

Note that it is still useful to explicitly compute the Gaussian integrals over the FOV because of the stopping criterion on splitting; these results may be scaled by the tail probabilities computed.

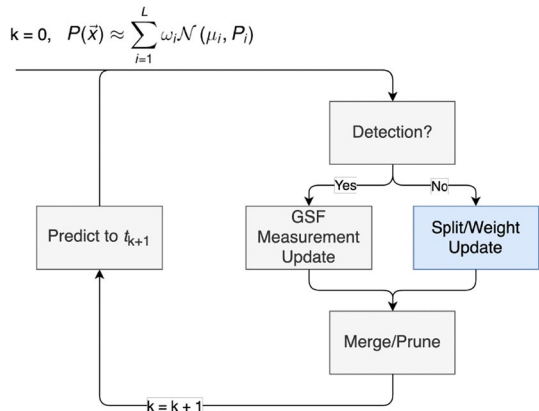
### 4.5 Merging and Filter Outline

With the filter update fully expressed, one now must ensure there is no hypothesis explosion in mixands so that the filter remains computationally efficient. With a goal of minimizing Kullback–Liebler divergence during the merging process, the well-known Runnall’s method is utilized [36]. A discrimination bound,

$$B(i, j) = \frac{1}{2} [(\omega_i + \omega_j) \log |P_{ij}| - \omega_i \log |P_i| - \omega_j \log |P_j|]
 \tag{56}$$

may be iteratively computed with the merged covariance for mixands  $i$  and  $j$ ,  $P_{ij}$ . Merging is iteratively performed until a threshold maxima of mixands is reached. Pruning may also be applied if weights are sufficiently small, but care must be taken to ensure that this does not disregard mixands split during the negative information update. With this methodology in place, the full filter is outlined in Fig. 5.

**Fig. 5** Gaussian Sum Filter diagram. Major contribution highlighted in blue



### 5 Results

The methodology is now presented for a case in which a prior detection is made and follow-up observation is desired. The object tracked has the true initial state (in kilometers and kilometers per second)

$$\mathbf{x} = [-27100 \ -32300 \ -100 \ 2.36 \ -1.98 \ 0] \tag{57}$$

in the Earth-centered inertial (ECI) frame. An observation is made by an observer at the initial ECI position (in kilometers)


$$\mathbf{o} = [517.859 \ -5281.538 \ 3526.190] . \tag{58}$$

An admissible region is then formed from knowledge of observer state and the attributable vector (in radians and radians per second) as


$$\mathbf{y} = [\alpha \ \delta \ \dot{\alpha} \ \dot{\delta}] = [-2.36716 \ -0.093581 \ 7.30762e-5 \ -1.53752e-9] . \tag{59}$$

Angular uncertainties of 5 arcseconds and angular rate uncertainties of 1.825 arcseconds per second are assumed in the attributable vector generation. The admissible region is first instantiated in the unobservable range range-rate halfplane. To constrain the admissible region, several assumptions are made on feasible orbits and applied in this space. The assumption is made that the target SO is on an elliptic trajectory following two body dynamics, with energy less than 0. A maximal eccentricity of 0.3 is applied. Finally a minimum radius of periapsis of 6500 km is assumed, ensuring the target will not collide with Earth. The admissible region is approximated using a mixture representation, again in range and range rate. The resultant admissible region is visualized in Fig. 6. In the Figure, the feasible set is presented in the range-range rate half plane. For further use, the admissible region is linearly transformed into the full state space (Cartesian position-velocity space). The admissible region may then propagated forward in time until the next possible observation

**Fig. 6** GMM representation of a geostationary admissible region in the range-range rate halfplane



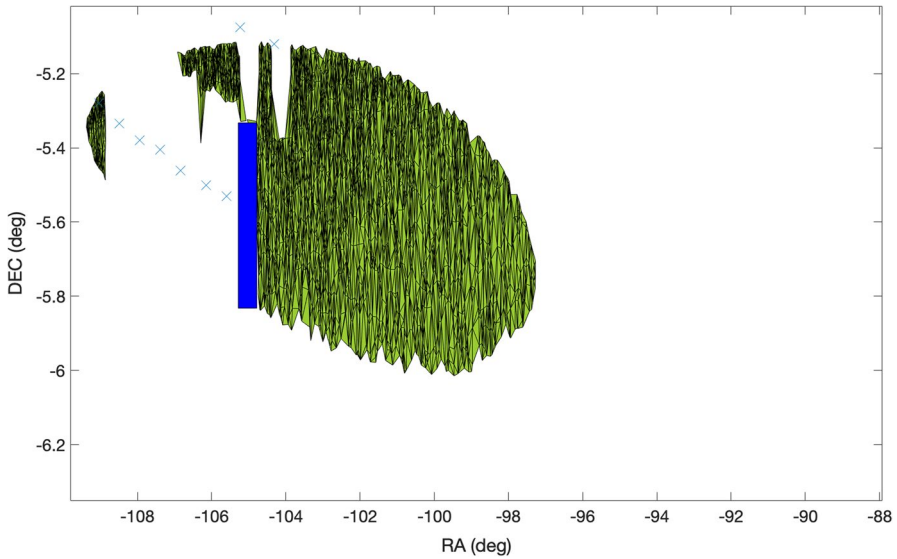
**Fig. 7** An admissible region for a GEO object propagated over time and projected into the field of regard of a ground-based observer



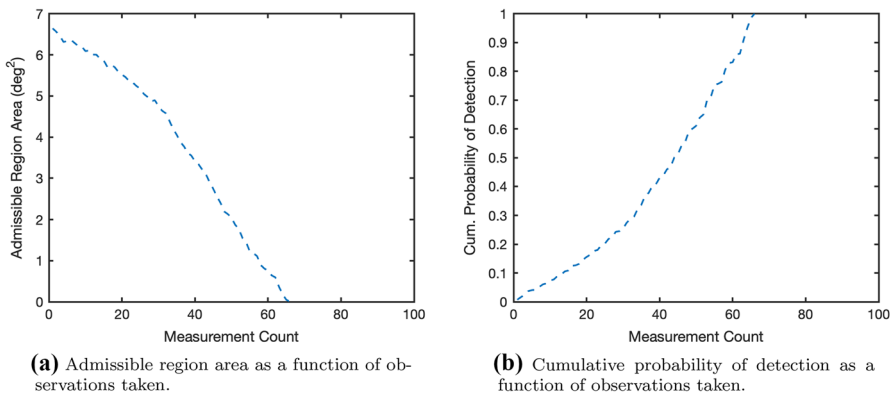
period, here assumed to be 2 h. Finally, the propagated admissible region may then be projected into the optical observer field of regard in Fig. 7.

Several characteristics may be noted from Figs. 6 and 7. First, probability density is uniform over the unobservable subspace, as no information is available on the SO range or range rate. Therefore, each feasible range and range rate pair has equal probability. However, this is not true when this feasible set is projected into another space. One may consider that the admissible region manifests as a point in measurement space, where it was fully observed. As the admissible region is propagated over time, the feasible region then grows in measurement space through a combination of dynamical uncertainty and rotation of the field of view of the observer, allowing the state to eventually become fully observable. Because different subsets of the admissible region grow at different rates, and because the projection of the admissible region into measurement space may still be concentrated around the mean angular rates over the set, the projection is decidedly not uniform in measurement space. Figure 7 demonstrates this fact, and this behavior is important to note when considering searching for the SO over the region. A successful tasking methodology captures two objectives in approaching this goal. First, it is critical to consider what regions of the feasible set are growing quickly in measurement so that it is possible to exhaust the admissible region over time. Second, it is critical to search over subregions with high probability density to maximize probability of detection.

Before search heuristics are considered, it is critical to provide an overview of the structure of the admissible region over time. As a baseline, it is worth noting an approximate end projected area of the admissible region of  $15 \text{ deg}^2$ . With a sensor field of view of  $0.25 \text{ deg}^2$ , at least 60 observations of the region would be required at this point. This is quite a conservative estimate, as this makes assumptions about the region at its largest point, but additional challenges arise in that the region is dynamic, and in that the set covering problem is NP-hard. As such, naive or randomized search methodologies are expected to struggle. With these points in mind, it is also important to consider the structure of the region over the course of a typical search. It is apparent that an effective search strategy must also carefully handle leftover regions of search space, in order to minimize cleanup after the vast majority



**Fig. 8** The admissible region after 10 observations have been taken



**(a)** Admissible region area as a function of observations taken.

**(b)** Cumulative probability of detection as a function of observations taken.

**Fig. 9** Admissible region area and probability of detection over time using a naive scanning pattern

of the feasible search region is exhausted. Figure 8 presents a visualization of the tasking process midway through a search routine. Note that particle sets are utilized in this visualization, unlike in Fig. 7.

In order to provide a comparison to the MCTS methodologies, a scanning strategy is developed, utilizing scenario-specific knowledge that the fastest-growing subset of the search space is that where right ascension is largest. A striping method is applied, moving along increasing declination and decreasing right ascension, hoping to capture these quickly growing subsets early in the search scenario. Results for this point of comparison are shown in Fig. 9, demonstrating how this methodology



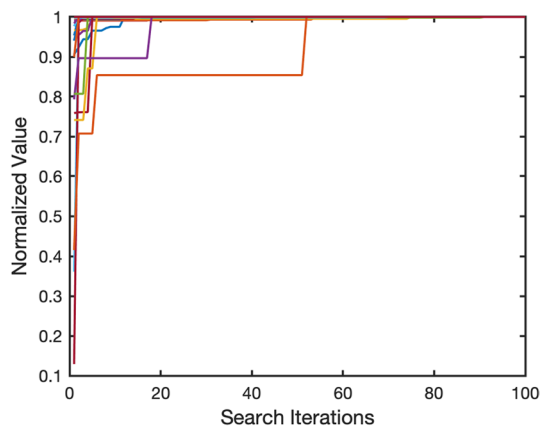
reduces admissible region area and increases probability of detection over time. As the reward functions utilized in MCTS, admissible region area and probability of detection will be the primary indicators of the relative merits of differing search strategies.

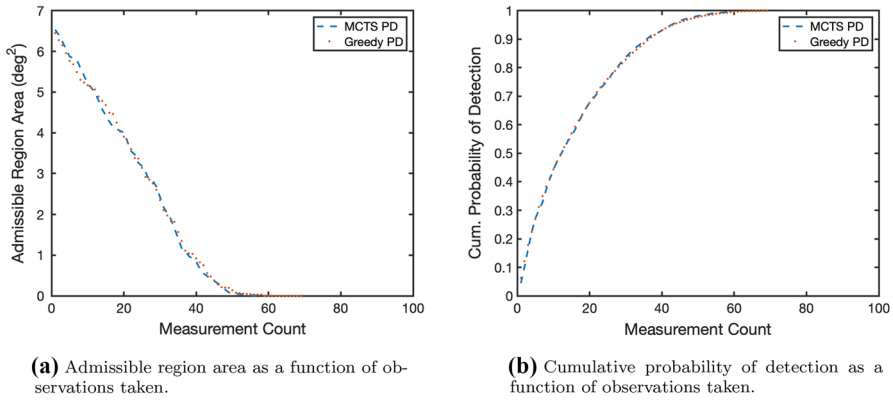
This baseline scanning pattern is able to complete the search campaign in the allocated period of 100 observations at a 15 s observation cadence (25 min of instrument time). 66 observations are required, and this result may be assumed as a minimum goal for the MCTS methodology. With this knowledge, tree search results are presented in two stages. For each rollout heuristic used, a greedy version is considered to determine that the heuristic is logical in application to the problem at hand. Then, MCTS is run to demonstrate the advantages of a search over a receding horizon, illustrating the combinatoric challenges of the problem. MCTS is run for 100 iterations down the search tree at each time step, with a search depth of 40 observations. As the search period nears the end of the 100 allocated observations, this depth is reduced accordingly such that only 100 observations are considered across the scope of the problem. Figure 10 gives some intuition in terms of the intrinsic value of further search over time. In the Figure, we observe that the rollout heuristics (in this case, the probability of detection heuristic) can very quickly achieve decision quality close to the maximum value found. However, this gain is still significant, especially for what is a relatively low number of search iterations.

The first tree search strategy to demonstrate considers probability of detection as a rollout heuristic and reward. The search space is discretized with a covering set of potential observational actions, and for each possible action, a value is determined by any previously unobserved particles lying in the field of view. The greedy and MCTS results are presented in Fig. 11.

Both the greedy and MCTS methods here offer significant gains as compared to the scanning strategy. Notably, the methods used are much more efficient early in the search period, achieving an almost constant rate of decrease in remaining search area over the first 40 measurements used. After this period, the comparative advantages of the MCTS method are revealed, and that method begins to achieve increased value as less of the search area remains. Logically, this is the result of

**Fig. 10** Normalized value estimated over a large set of tree search runs

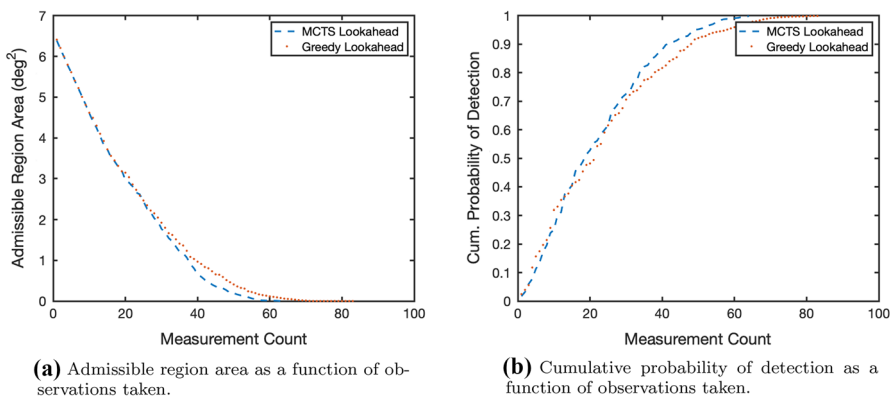




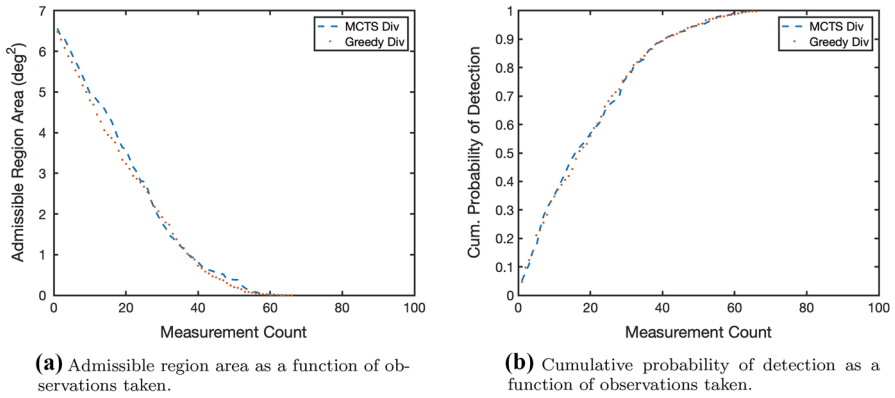
**Fig. 11** Admissible region area and probability of detection over time using the probability of detection heuristic

improved planning early in the search period; the greedy method is inefficient in that it leaves more small, disjoint regions that must be individually observed late in the search period in order to exhaust the search. This is, in effect, a gain of 10 observations that may be utilized for different sensor tasking objectives. The MCTS method here requires 59 observations to complete the campaign, and the greedy method takes 69 observations.

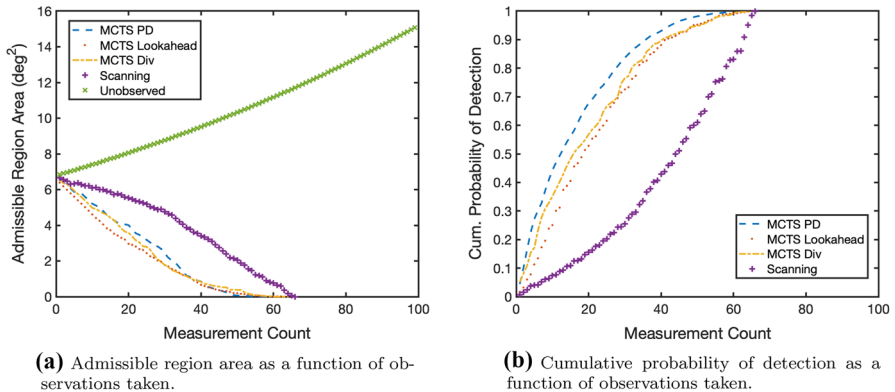
Next, the final projected area of a subset captured in an observation is considered using the lookahead heuristic. Cumulative probability of detection is again utilized as a reward. The search space is discretized with a covering set of potential observational actions, and for each possible action, the final projected area of particles within the field of view is applied as an observation weight. The greedy and MCTS results for this scenario are presented in Fig. 12.



**Fig. 12** Admissible region area and probability of detection over time using the lookahead heuristic



**Fig. 13** Admissible region area and probability of detection over time using the immediate change in area heuristic



**Fig. 14** Comparisons between each MCTS scenario and the naive scanner

There are several differences from the prior case noted in this result. With this heuristic, there seems a clear trade between gaining early decreases in search set area while also noting a slightly smaller probability of detection. It is likely that this behavior was influenced by the reduction in area reward function utilized in tandem with the lookahead heuristic. Interestingly, the MCTS version of the lookahead methodology offers vast improvements on the greedy version, requiring 64 observations as compared to 83. One explanation for this is that it seems this methodology requires significant cleanup at the end of search, and MCTS is able to minimize that need.


Finally, more immediate behavior of an observed subset is considered using the immediate change in area heuristic. The greedy and MCTS results for this scenario are presented in Fig. 13.

The structure of these results is much the same as the other two MCTS scenarios. Similar gains are made in this case, as compared to the scanning result, and as such it is most useful to make comparisons between each MCTS case. In Fig. 14, one may observe that this scenario seems to offer a middle ground in that it trails the lookahead heuristic in reducing area and it trails the probability of detection heuristic in achieving cumulative probability of detection. Each method offers clear advantages over the scanning methodology. The general structure of the growth of the unobserved region over time is also shown.



With these results outlined, MCTS is then applied in combination with the developed estimation paradigm. As in the initial search case, observations are taken at a 15 s cadence until the search region is exhausted; the first follow-up observation is performed 2 h after the initial detection. The admissible region, with initial area in measurement space of approximately  $7 \text{ deg}^2$ , is exhausted with a sequence of 66 observations, using a sensor with a square field of view of  $0.25 \text{ deg}^2$ . At the time of the 66th observation, the projected admissible region has an area of approximately  $12 \text{ deg}^2$ . Over the course of this tasking solution, a single observation is made at timestep  $t_{15} = 7410 \text{ s}$  when the observer is pointing at the angular coordinates  $\alpha = -1.82939 \text{ rad}$ ,  $\delta = -0.093562 \text{ rad}$ . In each other case, no detection is made, and the negative information is processed. Figure 15 visualizes the effects of processing negative information 7 observations into the tasking scenario. Note that probability density in the observed subset of measurement space has been greatly reduced.

In Fig. 15, the most recent observation is represented by a shaded rectangle, and the true state, at  $\alpha = -1.8346 \text{ rad}$  and  $\delta = -0.09318 \text{ rad}$ , is marked by a plus sign. Observed regions are most visible from  $\alpha = -1.87$  to  $-1.85$  radians and  $\delta = -0.1$  to  $-0.09$  radians. Because of these reductions, unobserved subsets of the projected admissible region are now comparatively more likely. Indeed, the density at the true state is approximately 10 percent higher than prior to the processing of any negative information. This result is quite comparable in essence to the spooky effect described by Franken et al [18]. The negative information in this case describes missed detections on a subset of mixands in the ensemble,

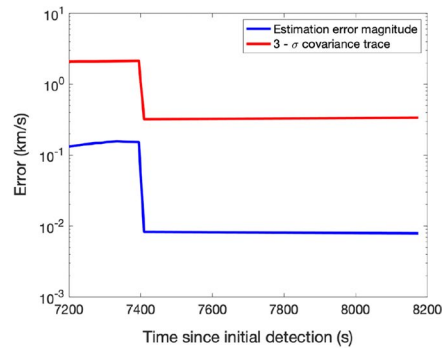
**Fig. 15** The admissible region projected into measurement space after 7 null detections



**Fig. 16** The admissible region projected into measurement space after a follow up detection is made. Uncertainty in angular space is equivalent to measurement uncertainty



**(a)** Positional estimation error over the course of the simulated observation campaign.



**(b)** Velocity estimation error over the course of the simulated observation campaign.

**Fig. 17** Estimation errors before and after the detection is made

increasing the likelihood that each other mixand is “truth” and may be associated with the true state.

It is also important to consider the behavior of the filter when a measurement is received. Figure 16 demonstrates the reduction of the mixture when this occurs at time  $t_{15}$ . Here, the projected area of the mixture is now reduced to that of measurement uncertainty, assumed to be on the order of 5 arcseconds in this simulation. After this observation is made, there is negligible effect on the state estimate through further processing of negative information, but this is still critical to do in real scenarios, when the likelihood of false alarm measurements is non-negligible. Finally, Fig. 17 visualizes estimation error over the course of the simulation. Note that the estimation error remains within the covariance bounds throughout the simulation, and is greatly reduced when the follow-up observation is received.

The covariance bounds are a bit conservative in this context because of the relatively short time between observations for a geostationary object. Interestingly, the full state is only weakly observable, and though positional uncertainties are reduced

by an order of magnitude, there is still significant uncertainty in the sensor line of sight direction after the measurement, on the order of  $2000 \text{ km } 3 - \sigma$ . This a reflection of the Too Short Arc problem [37], and the unobservable subset of state space has not fully rotated into the field of regard of the observer. However, uncertainty is sufficiently reduced such that only a small number of mixands are needed to effectively represent probability density after the measurement. There is little impact on the processing of negative information on estimation error, but this is as expected. Before a detection is made, the PDF is still quite large in state space; while some mixands are eliminated through processing negative information, the ensemble mean is not very useful for such a large PDF, and in this case, it just happens to be that the original mean is somewhat close to the true state. It would be interesting to determine whether the ensemble mean becomes more useful in cases where most of the feasible region is exhausted before a follow-up detection is made.

## 6 Discussion

Generally, all of the methods developed performed very well. An interesting avenue of further research would be in the process of discretizing the search space. A relatively efficient gridding method is utilized, but it is likely that further improvements could be seen in treating the sub problem of determining feasible actions as a covering or packing problem. That problem is NP-hard, but it is likely that utilization of more rigorous partitioning of the action space using approximate covering algorithms could lead to more efficient options, especially when the search region becomes somewhat sparse.

Also of interest for further study is the extension of this research to probabilistic admissible regions (PARs) and reachable sets. Incorporation of a non-unity observation probability when actions are taken is a useful consideration for challenging cases with dim and distant, near-cislunar objects, as well as for flexibility if comparatively poor sensors are in operation. Searching reachable sets is a critical problem as well, and this methodology is quite applicable for maintaining state estimates on maneuvering objects. It is important to note that the tasking process is essentially analogous for these extensions, but that particle weights may change according to the probability density of the PAR or a priori knowledge on a maneuver or unmodeled force.

Additional avenues may also be considered for the second contribution of this work. While the logical focus for processing negative information in this paper is the context of optical observations, this methodology could also be extended for utilization with radar measurements and a variety of novel observational techniques. The key use case is any situation in which the spread of the projected state estimate exceeds the sensor field of regard in measurement space.

Theoretic bounds on the size of Gaussian mixand splits are outlined, but it is noted that further perturbation along measurement axes comes with an increase in Kullback–Liebler divergence from the original mixture. Future work may aim to formally pose this splitting methodology as an optimization in which Kullback–Liebler

divergence is minimized with an additional cost objective comparing the spread of resultant mixands to a target value (for example, a sensor field of view).

This splitting methodology is utilized to ensure that the weight update remains Gaussian, and it is important to make further comparisons to particle representations of the state estimate, as this state representation is considered for instantiation of the tasking contributions. Particles may also be utilized with likelihood updates on particles within a sensor field of view when no detection is made, but the Gaussian mixture representation may be considered advantageous in this context for several reasons. First, it allows for a smooth representation of probability density over the support of the state estimate. This is quite useful in regions of measurement space where particles may be quite sparse. Additionally, there are computational advantages to utilizing Gaussian mixture representations, in that a large set of particles need not be sampled; this is especially critical when particles are sampled from a high-dimensional state space. In the context of admissible regions, this is not the case, but a curse of dimensionality becomes rather important when considering search over reachable sets in Cartesian position-velocity space.

In general, results demonstrate advantages in utilization of negative information before follow-up detections are made. Extensions of these results will be made considering a variety of target orbits, utilization of more interesting dynamics, and follow-up tracking of maneuvering objects.

Finally, it is important to note that the end goal of this research is incorporation in an online manner. The problem of tracklet association for multiple generated admissible regions is solved, but this research is beneficial when there is immediate need for follow-up observation. It would be quite interesting to consider whether this tasking objective may be leveraged in tandem with other objectives such as pure search and catalog maintenance. In future work, a primary goal is to consider the multi-objective optimization problem inherent to this question.

## 7 Conclusion

In this work, new methodologies for space object search and recovery are developed using Monte Carlo Tree Search. A variety of heuristics useful for object reacquisition are considered, and the developed techniques are applied to a scenario representative of modern sensor tasking needs. These methods enable optical recovery in the contexts of initial orbit determination and the tracking of maneuvering targets. Estimation is performed along side these tasking methodologies, and special considerations are made such that a Gaussian Sum Filter may ingest negative information when no detection is made. A novel Gaussian splitting methodology is developed to ensure this update remains Gaussian, and the resultant filter is demonstrated to converge over the course of search and recovery. The developed methodologies augment the initial orbit determination literature, offering a new perspective that is especially useful for instantiating and maintaining knowledge on high-priority space objects.

**Acknowledgements** Research on this project has been supported by the National Science Foundation Graduate Research Fellows Program Grant Number DGE 1650115 and the Draper Scholars.

## Declarations

**Conflict of interest** On behalf of all authors, the corresponding author states that there is no conflict of interest.

## References


1. Holzinger, M.J., Jah, M.K.: Challenges and potential in space domain awareness. *J. Guid. Control. Dyn.* **41**(1), 15–18 (2018). <https://doi.org/10.2514/1.G003483>
2. ESA Space Debris Office: ESA's Annual Space Environment Report, No. July, 2019, pp. 1–78
3. Erwin, R.S., Albuquerque, P., Jayaweera, S.K., Hussein, I.: Dynamic sensor tasking for space situational awareness. In: Proceedings of the 2010 American Control Conference, ACC 2010, pp. 1153–1158 (2010). <https://doi.org/10.1109/acc.2010.5530989>
4. Williams, P.S., Spencer, D.B., Erwin, R.S.: Coupling of estimation and sensor tasking applied to satellite tracking. *J. Guid. Control. Dyn.* **36**(4), 993–1007 (2013). <https://doi.org/10.2514/1.59361>
5. Sunberg, Z., Chakravorty, S., Erwin, R.S., Member, S.: Information space receding horizon control for multisensor tasking problems. *IEEE Trans. Cybern.* **46**(6), 1325–1336 (2016). <https://doi.org/10.1109/TCYB.2015.2445744>
6. Linares, R., Furfaro, R.: Dynamic sensor tasking for space situational awareness via reinforcement learning. In: Advanced Maui Optical and Space Surveillance Technologies Conference, pp. 1–10 (2016)
7. Fedeler, S., Holzinger, M., Whitacre, W.: Sensor tasking in the cislunar regime using Monte Carlo Tree Search. *Adv. Space Res.* **70**, 1–19 (2022). <https://doi.org/10.1016/j.asr.2022.05.003>
8. Africano, J., Schildknecht, T., Matney, M., Kervin, P., Stansbery, E., Flury, W.: A geosynchronous orbit search strategy. *Space Debris* **2**, 357–369 (2000). <https://doi.org/10.1023/B:SDEB.0000030025.04930.08>
9. Frueh, C., Fielder, H., Herzog, J.: Heuristic and optimized sensor tasking observation strategies with exemplification for geosynchronous objects. *J. Guid. Control. Dyn.* **41**(5), 1036–1048 (2018). <https://doi.org/10.2514/1.G003123>
10. Milani, A., Sansaturio, M.E., Tommei, G., Arratia, O., Chesley, S.R.: Multiple solutions for asteroid orbits: computational procedure and applications. *Astron. Astrophys.* **431**(2), 729–746 (2005). <https://doi.org/10.1051/0004-6361:20041737>
11. Worthy, J.L., Holzinger, M.J.: Use of uninformative priors to initialize state estimation for dynamical systems. *Adv. Space Res.* **60**(7), 1373–1388 (2017). <https://doi.org/10.1016/j.asr.2017.06.040>
12. Gehly, S., Jones, B.A., Axelrad, P.: Search-detect-track sensor allocation for geosynchronous space objects. *IEEE Trans. Aerosp. Electron. Syst.* **54**(6), 2788–2808 (2018). <https://doi.org/10.1109/TAES.2018.2830578>
13. DeMars, K.J., Jah, M.K.: Probabilistic initial orbit determination using Gaussian mixture models. *J. Guid. Control. Dyn.* **36**(5), 1324–1335 (2013). <https://doi.org/10.2514/1.59844>
14. Fujimoto, K., Alfriend, K.T.: Optical short-arc association hypothesis gating via angle-rate information. *J. Guid. Control. Dyn.* **38**(9), 1602–1613 (2015). <https://doi.org/10.2514/1.G000927>
15. Schumacher, P.W., Gaebler, J.A., Roscoe, C.W., Wilkins, M.P., Axelrad, P.: Parallel initial orbit determination using angles-only observation pairs. *Celest. Mech. Dyn. Astron.* **130**(9), 1–20 (2018). <https://doi.org/10.1007/s10569-018-9852-6>
16. Hobson, T.A., Clarkson, I.V.L.: A particle-based search strategy for improved space situational awareness. In: 2013 Asilomar Conference on Signals, Systems and Computers, pp. 898–902 (2013). <https://doi.org/10.1109/ACSSC.2013.6810418>
17. Murphy, T.S., Holzinger, M.J.: Generalized minimum-time follow-up approaches applied to tasking electro-optical sensor tasking. In: Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference (2017)
18. Fränken, D., Schmidt, M., Ulmke, M.: “Spooky action at a distance” in the cardinalized probability hypothesis density filter. *IEEE Trans. Aerosp. Electron. Syst.* **45**(4), 1657–1664 (2009). <https://doi.org/10.1109/TAES.2009.5310327>
19. Chang, H.S., Fu, M.C., Hu, J., Marcus, S.I.: An adaptive sampling algorithm for solving Markov decision processes. *Oper. Res.* **53**(1), 126–139 (2005). <https://doi.org/10.1287/opre.1040.0145>



20. Coulom, R.: Efficient selectivity and backup operators in Monte-Carlo tree search. In: Fürnkranz, J., Scheffer, T., Spiliopoulou, M. (eds.) *Machine Learning: ECML 2006*. ECML 2006. Lecture Notes in Computer Science, vol. 4212, pp. 282–293. Springer, Berlin, Heidelberg (2006). [https://doi.org/10.1007/11871842\\_29](https://doi.org/10.1007/11871842_29)
21. Kocsis, L., Szepesvari, C.: Bandit based Monte-Carlo Planning. *Lect. Notes Comput. Sci.* **4212**, 282–293 (2006). <https://doi.org/10.1109/scc.2004.1358005>
22. Silver, D., Veness, J.: Monte-Carlo planning in large POMDPs. In: *Advances in Neural Information Processing Systems (NIPS)*, pp. 1–9(2010)
23. Sunberg, Z., Kochenderfer, M.J.: Online algorithms for POMDPs with continuous state, action, and observation spaces. In: *Twenty-Eighth International Conference on Automated Planning and Scheduling* (2018)
24. Ross, S.: *Introduction to Stochastic Dynamic Programming*. Academic Press, New York (1983)
25. Pineau, J., Gordon, G., Thrun, S.: Point-based value iteration: an anytime algorithm for POMDPs. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 1025–1030 (2003)
26. Genz, A., Trinh, G.: Numerical computation of multivariate normal probabilities using bivariate conditioning. *J. Comput. Graph. Stat.* **1**, 141–149 (1992)
27. Cunningham, J.P., Hennig, P., Lacoste-Julien, S.: Gaussian probabilities and expectation propagation. *arXiv e-prints*. [arXiv:1111.6832](https://arxiv.org/abs/1111.6832) (2011)
28. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* **47**, 235–256 (2002)
29. Auger, D., Couëtoux, A., Teytaud, O.: Continuous upper confidence trees with polynomial exploration - consistency. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 194–209(2013)
30. Wang, Y., Audibert, J.Y., Munos, R.: Algorithms for infinitely many-armed bandits. In: *Advances in Neural Information Processing Systems 21 - Proceedings of the 2008 Conference*, pp. 1729–1736(2009)
31. Couëtoux, A., Hoock, J.-B., Sokolovska, N., Teytaud, O.: Continuous upper confidence trees. In: Coello, C.A.C. (eds.) *Learning and Intelligent Optimization. LION 2011. Lecture Notes in Computer Science*, vol. 6683, pp. 433–445. Springer, Berlin, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-25566-3\\_32](https://doi.org/10.1007/978-3-642-25566-3_32)
32. Edelsbrunner, H., Kirkpatrick, D., Seidel, R.: On triangulations of a set of points in the plane. *IEEE Trans. Inf. Theory* **29**(4), 228–240 (1983)
33. Alspach, D.L., Sorenson, H.W.: Nonlinear bayesian estimation using Gaussian sum approximations. *IEEE Trans. Autom. Control* **17**(4), 439–448 (1972). <https://doi.org/10.1109/TAC.1972.1100034>
34. Julier, S.J., Uhlmann, J.K.: New extension of the Kalman filter to nonlinear systems. In: *Signal Processing, Sensor Fusion, and Target Recognition VI. Proc. SPIE 3068* (1997). <https://doi.org/10.1117/12.280797>
35. Havlak, F., Campbell, M.: Discrete and continuous, probabilistic anticipation for autonomous robots in Urban environments. *IEEE Trans. Robotics* **30**(2), 461–474 (2014). <https://doi.org/10.1109/TRO.2013.2291620>
36. Runnalls, A.R.: Kullback-Leibler approach to Gaussian mixture reduction. *IEEE Trans. Aerosp. Electron. Syst.* **43**(3), 989–999 (2007). <https://doi.org/10.1109/TAES.2007.4383588>
37. Milani, A., Gronchi, G.F., Vitturi, M.D.M., Knežević, Z.: Orbit determination with very short arcs. I admissible regions. *Celest. Mech. Dyn. Astron.* **90**(1–2), 59–87 (2004). <https://doi.org/10.1007/s10569-004-6593-5>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Samuel J. Fedeler<sup>1</sup>  · Marcus J. Holzinger<sup>1</sup> · William W. Whitacre<sup>2</sup>

Marcus J. Holzinger  
marcus.holzinger@colorado.edu

William W. Whitacre  
wwhitacre@draper.com

<sup>1</sup> Smead Aerospace Engineering Sciences, University of Colorado Boulder, Boulder, USA

<sup>2</sup> The Charles Stark Draper Laboratory, Cambridge, USA