Intelligent Sensor Tasking for Minimum Time Space Object Acquisition

Trevor N. Wolf, Brandon A. Jones
The University of Texas at Austin

ABSTRACT

We consider the problem and associated challenges of optimally allocating tasks with limited sensor resources to rapidly acquire a newly detected space object. Optical telescopes or radars typically provide novel measurements to constitute a detection, however, taken independently, lack sufficient information to characterize the state of the space object. Gaining custody of the object then requires either sufficient geometric diversity between multiple measurements or a combination of different sensing modalities. Generating a second measurement for either case may then require searching a large feasible measurement space that can, in principle, take considerable time and, as a result, produce a null measurement. This work introduces a new method for efficiently directing a sensor or multiple sensors to search a feasible state set. Our approach is built around a finite set statistics framework so that the set of states associated to the untracked object is represented by a cardinalized probability hypothesis density (CPHD) based description of the multi-target state. The objective of the tasking agent is to minimize an objective taken to be the expected differential entropy of the CPHD intensity conditioned on the directed task. We show that effectiveness of our approach in a closed-loop tasking simulation for two representative cases of importance to the space situational awareness community; a geostationary orbit (GEO) and geostationary transfer orbit (GTO).

1. INTRODUCTION

Gaining custody of newly detected, high priority space objects is an open challenge in the Space Situational Awareness (SSA) community. While optical telescopes or radars provide partial state information in singular measurement tracks, taken alone they lack sufficient information to fully characterize the state of the space object. Gaining custody of the object then requires either geometric diversity between multiple measurements or a combination of different sensing modalities. To facilitate data association between multiple measurement tracks, physics-informed assumptions allow for generating a representation of the feasible set of states from single measurements known as the admissible region [1, 2]. In [3] the authors represent the admissible region by a Gaussian Mixture Model (GMM) and perform initial orbit determination by forecasting the evolution of this probability density function (PDF) at future secondary measurement times. Generating an admissible region that ensures a positive detection at a future times requires well constrained bounds for semimajor axis and eccentricity, which are most often not available. Relaxing these bounds provides higher confidence that the object's true state is contained in the admissible region but drastically increases the hypervolume of dynamically feasible states associated to the target. The issue is compounded in many real cases where physical occlusions, shadowing, or limited sensor resources exclude the availability of timely follow-up observations, between which the feasible search space may grow and become too cumbersome to search with naive approaches.

Previous work has investigated the problem of efficiently searching an admissible region with optical telescopes. In [4], the authors introduce a method for rapidly searching a feasible space based on the divergence of the search volume. They take the area over time as the performance metric in forming a time optimal schedule to search the entirety of feasible space. The authors in [5] extend this by considering Monte-Carlo Tree Search as an approach to direct tasks for rapidly covering the area associated to the feasible states set. The work in [6] incorporates information contained in null measurements to help guide follow-up telescope pointing towards regions containing a higher likelihood of detection.

This work introduces a closed-loop sensor allocation scheme for efficiently searching a feasible state set. The approach is formulated as a Markov Decision Process (MDP). A Finite Set Statistics (FISST) framework is employed so that the set of feasible states associated to the untracked target is represented as the Cardinalized Probability Hypothesis Density (CPHD). This approach operates on the first moment of a multi-target state distribution, known as the intensity function, and a cardinality distribution describing the number of targets within a search domain. The tasking agent's

objective is to minimize an information theoretic objective taken to be the expected differential entropy of the CPHD intensity, conditioned on a potential directed task. Other work has investigated information theoretic measures for sensor allocation applied to catalog maintenance in SSA applications (e.g., [7], [8], [9], [10]), however, to the best of our knowledge, ours is the first to consider these measures in application to rapid space object acquisition.

The rest of this work is organized as follows: in Section 2 we introduce the CPHD formulation for representing the feasible state set and the tasking agent's objective function. Section 3 describes a regularized particle filter implementation for approximating the CPHD intensity as well as a weighted k-nearest neighbors entropy approximation for the tasking agent's objective. Section 4 presents several synthetic test case scenarios demonstrating closed-loop follow-up tasking with a singular sensor for a geostationary orbit (GEO) and a geostationary transfer orbit (GTO). We also include a test case demonstrating simultaneous coordinated follow-up tasking from two ground-based telescopes. We perform a Monte-Carlo analysis to examine the statistical spread in the number of measurements scans required to obtain a follow-up true positive detection for the GTO example. We summarize the contributions of this work and provide suggestions for our future work in Section 5.

2. PROBLEM FORMULATION

In this section we describe the framework we use for tasking a sensor to acquire a newly detected object. Our methods are built around a Random Finite Set (RFS) formulation in which a multi-target state and measurement are

$$\boldsymbol{X}_k = \{\boldsymbol{x}_{k,1}, \cdots, \boldsymbol{x}_{k,n_k}\} \in \mathscr{F}(\mathscr{X}), \tag{1}$$

$$\boldsymbol{Z}_k = \{\boldsymbol{z}_{k,1}, \cdots, \boldsymbol{z}_{k,m_k}\} \in \mathscr{F}(\mathscr{Z}), \tag{2}$$

respectively, where n_k is the number of tracked targets, and m_k is the number of returned measurements at a current time step k. This formulation provides several distinct advantages for our application. For instance, it allows for the ingestion of negative information contained in a null detection of an object, characterization of false alarms (i.e., clutter), and information contained in a true positive detection. In some instances, a user may be provided information on the number of targets as well (e.g., the existence of at least 3 targets within a search space), which can be properly leveraged with RFS theory. Among the many tools stemming from the RFS paradigm, we use the Cardinalized Probability Hypothesis Density (CPHD) [11] for our application, and adopt the notation used in [12]. The CPHD operates on the first-order moment of the multi-target state, $f_{k|k}(\cdot)$, which is given by

$$D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k}) = \int f_{k|k}\big(\boldsymbol{X}|\boldsymbol{Z}_{1:k}\big)\,\delta\boldsymbol{X}, \tag{3}$$

and the cardinally distribution of the true number of targets, n, within a search space,

$$\rho(n|\boldsymbol{Z}_{1:k}) = \frac{1}{n!}\int f_{k|k}\big(\{\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n\}|\boldsymbol{Z}_{1:k}\big)d\boldsymbol{x}_1 \cdots d\boldsymbol{x}_n. \tag{4}$$

The function $D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k})$ is often referred to as the intensity function, and satisfies the property

$$\sum_{n=1}^{\infty} n\rho(n|\boldsymbol{Z}_{1:k}) = \int D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k})d\boldsymbol{x}. \tag{5}$$

In other words, the expected number of targets within the search space is the total integral of the intensity function.

In this work we make several simplifying assumptions, but the reader should note in general these are not required. First, we assume that only a single target exists within our search space and we assume that there is no clutter in the detection model. Additionally, we assume that there are no target births or deaths, so the single target is persistent throughout. As such, the cardinally distributions for the object number and clutter are

$$\rho(n) = \begin{cases} 1 & \text{if } n = 1 \\ 0 & \text{if } n \neq 1 \end{cases} \tag{6}$$

and,

$$\rho_{\kappa}(n_{\kappa}) = \begin{cases} 1 & \text{if } n_{\kappa} = 0 \\ 0 & \text{if } n_{\kappa} \neq 0 \end{cases}, \tag{7}$$

respectively, where n is the number of objects, and n_κ is the number of clutter returns. At each measurement time, we can then assume two cases: $|\boldsymbol{Z}_k| = 0$ or $|\boldsymbol{Z}_k| = 1$ (i.e., there is a null detection or the singular target is detected). While we omit the derivation here for brevity, it follows then that the posterior intensity is

$$D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\emptyset) = \big(1 - P_D(\boldsymbol{x};\boldsymbol{p})\big)\frac{D_{k|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1})}{\big\langle 1 - P_D(\boldsymbol{x}), D_{k|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1})\big\rangle}, \tag{8}$$

for the case that $|\boldsymbol{Z}_k| = 0$, and

$$D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\boldsymbol{z}_k) = P_D(\boldsymbol{x};\boldsymbol{p})\frac{g_k(\boldsymbol{z}_k|\boldsymbol{x})D_{k|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1})}{\big\langle P_D(\boldsymbol{x})D_{k|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1}), g_k(\boldsymbol{z}_k|\boldsymbol{x})\big\rangle}, \tag{9}$$

for the case that $|\boldsymbol{Z}_k| = 1$. Here, $\langle\cdot,\cdot\rangle = \int f(\boldsymbol{x})g(\boldsymbol{x})d\boldsymbol{x}$ is the function convolution, $g(\boldsymbol{z_k}|\boldsymbol{x})$ is measurement likelihood function, and $P_D(\boldsymbol{x})$ is the probability of detection given as

$$P_D(\boldsymbol{x};\boldsymbol{p}) = \begin{cases} \tilde{P}_D(\boldsymbol{x};\boldsymbol{p}) & \text{if } \boldsymbol{h}\big(\boldsymbol{x};A_k^\pi(\cdot)\big) \in \mathscr{V} \\ 0 & \text{if } \boldsymbol{h}\big(\boldsymbol{x};A_k^\pi(\cdot)\big) \notin \mathscr{V} \end{cases}. \tag{10}$$

In the above, $\tilde{P}_D(\boldsymbol{x};\boldsymbol{p})$ is the probability of detection for the object given it is contained in the sensor's detectable set, which we denote $\mathscr{V}$. It can depend on a number of parameters, $\boldsymbol{p}$, which may not all be characterized well including atmospheric seeing, optical cross section, reflectivity, etc. For the remainder of the discussion, we omit the use of $\boldsymbol{p}$ for brevity. The set $\mathscr{V}$ may be the field of view (FOV) for an optical telescope or the time gate of photon arrival for a optical laser ranging system. The function $\boldsymbol{h}(\boldsymbol{x};A_k^\pi(\cdot))$ is the transformation from the object's state to the measurement plane of the sensor. Note that the measurement function implicitly depends on the action of the sensor network, $A_k^\pi(\cdot)$, at the current time step, k.

A tasking agent would like to choose an action $A_k^\pi(\cdot)$ associated to a policy π that will maximize a desired utility; in our case, to acquire the space object quickly. The acquisition problem presents a unique challenge because it is not known *a priori* whether a measurement will result in a detection and as such, the number of measurements required to obtain a true positive detection. We choose to express the problem as a receding horizon optimization which is most generally given as

$$\min_\pi \mathbb{E}^\pi\Big\{\sum_{k'=0}^{T}\gamma^{k'}C_{k'}^\pi\big(S_{k'},A_{k'}^\pi(S_{k'})\big)\Big\}. \tag{11}$$

We use the symbol k' here to denote the future time steps within the window. T is the length of the horizon considered, $\gamma \in [0,1]$ is a discount factor, $C_{k'}^\pi(\cdot)$ is a utility function associated to the policy π, and $S_{k'}$ is the state of the system at time step, k'. This type of optimization problem is most commonly employed in dynamic programming and reinforcement learning [13], however, instead of learning a policy π offline, in this work we directly optimize Equation 11 at each time step k.

We take the expected differential entropy of the posterior intensity function to be the utility function $C_{k'}^\pi(\cdot,\cdot)$. The differential entropy of an arbitrary distribution $p(\boldsymbol{x})$ is given as

$$\mathscr{H}\big(p(\boldsymbol{x})\big) = -\int p(\boldsymbol{x})\ln\big(p(\boldsymbol{x})\big)d\boldsymbol{x}. \tag{12}$$

For our problem, to determine the expected entropy value for a single sensor follow-up scenario, we must consider the effect of the two potential aforementioned measurement cardinality cases on the predicted CPHD intensity within time horizon. The associated probability of occurrence for these cases are

$$P(|\boldsymbol{Z}_{k'}| = 1) = \int P_D(\boldsymbol{x})D_{k'|k'-1}\big(\boldsymbol{x}|\boldsymbol{Z}_{1:k'-1}\big)d\boldsymbol{x}, \tag{13}$$

and $P(|\boldsymbol{Z}_{k'}| = 0) = 1 - P(|\boldsymbol{Z}_{k'}| = 1)$. With this, the expected differential entropy at the time step k' is

$$\mathbb{E}\big[\mathscr{H}\big(D_{k+k'|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\overline{\boldsymbol{Z}}_{k:k+k'})\big)\big] = \sum_{p=1}^{2^{k'}}\Big(\prod_{j=0}^{k'}P(|\overline{\boldsymbol{Z}}_{k+j}^p|)\Big)\mathscr{H}\big(D_{k+k'|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\overline{\boldsymbol{Z}}_{k:k+k'}^p)\big). \tag{14}$$

The above can be understood as the expected entropy over all possible enumerations of the measurement sets up to time step k', $\{\overline{\boldsymbol{Z}}^p_{k:k+k'}\}_{p=1}^{2^{k'}}$. Here, p indexes a potential measurement set outcome, $P(|\overline{\boldsymbol{Z}}^p_{k+j}|)$ is the probability of cardinality for the measurement at the future time step $k+j$, and $D_{k+k'|k-1}(\cdot)$ is the predicted posterior intensity function given a measurement set outcome $\overline{\boldsymbol{Z}}^p_{k:k+k'}$. For cases where the cardinality is one, we sequentially update the intensity function with the expected measurement $\bar{\boldsymbol{z}}_{k'}$ given as

$$\bar{\boldsymbol{z}}_{k'} = \mathbb{E}\big[\boldsymbol{z}_{k'}|A^{\pi}_{k'}(\cdot)\big] = \frac{\int \boldsymbol{h}\big(\boldsymbol{x};A^{\pi}_{k'}(\cdot)\big)P_D(\boldsymbol{x})D_{k+k'-1|k-1}\big(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\overline{\boldsymbol{Z}}_{k:k+k'-1}\big)d\boldsymbol{x}}{\int P_D(\boldsymbol{x})D_{k+k'-1|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\overline{\boldsymbol{Z}}_{k:k+k'-1})d\boldsymbol{x}} \tag{15}$$

In the multi-target tracking community, Equation 15 is often referred to as the predicted ideal measurement (PIM) model [11]. Equation 11 then becomes

$$\min_{\pi}\sum_{k'=0}^{T}\gamma^{k'}\mathbb{E}\big[\mathscr{H}\big(D_{k+k'|k-1}(\boldsymbol{x}|\boldsymbol{Z}_{1:k-1},\overline{\boldsymbol{Z}}_{k:k+k'})\big)\big]. \tag{16}$$

Using entropy as the objective function has a few advantages for our application. First, it directly quantifies the expected uncertainty associated to the target intensity function. This leads a sensing agent to visit regions of the feasible state space that will result in the greatest uncertainty reduction, regardless of whether the region has previously been observed. This is particularly useful in situations where $\tilde{P}_D(\boldsymbol{x})$ can not be assumed unity. Secondly, in considering follow-up tasking with multiple sensor resources, entropy is interpretable by any sensor within the network which lends well to simultaneous cooperative tasking between agents.

Equation 14 presents a significant computational challenge because the number of measurement set combinations grows geometrically with the length of the horizon, T. The issue is only exacerbated in considering a multi-sensor tasking as the base of the exponential becomes $2\times r$ where r denotes the number of sensors. For this work, we only consider a single measurement set look-ahead (i.e., $T=1$), however, our future work intends to manage this combinatorial problem by training a sensing agent to learn an optimal policy π off-line.

3. PARTICLE IMPLEMENTATION OF THE CPHD RECURSION FOR CLOSED-LOOP SENSOR TASKING

3.1 Particle-Based CPHD Filter and Entropy Approximation

In this section, we detail our implementation for a closed-loop tasking agent. To approximate the CPHD intensity function, we choose an ensemble of weighted particles so that

$$D_{k|k}(\boldsymbol{x}|Z_{1:k}) \approx \sum_{i=1}^{N} w^i_{k|k}\boldsymbol{\delta}_{\boldsymbol{x}^i_{k|k}}(\boldsymbol{x}). \tag{17}$$

Here, N is the total number of particles in the ensemble, $w^i_{k|k}$ is the weight associated to each particle i, and $\boldsymbol{\delta}_{(\cdot)}$ is the Dirac delta function. In this form, the particles and associate particle weights postulate the expected number of targets within a volume of space. For the case $|\boldsymbol{Z}_k|=0$, the posterior particle weights are

$$w^i_{k|k} = \frac{\big(1-P_D(\boldsymbol{x}^i_{k|k-1})\big)w^i_{k|k-1}}{\sum_{j=1}^{N}\big(1-P_D(\boldsymbol{x}^j_{k|k-1})\big)w^j_{k|k-1}}, \tag{18}$$

and for the case $|\boldsymbol{Z}_k|=1$,

$$w^i_{k|k} = \frac{P_D(\boldsymbol{x}^i_{k|k-1})g_k(\bar{\boldsymbol{z}}_k|\boldsymbol{x}^i_{k|k-1})w^i_{k|k-1}}{\sum_{j=1}^{N}P_D(\boldsymbol{x}^j_{k|k-1})g_k(\bar{\boldsymbol{z}}_k|\boldsymbol{x}^j_{k|k-1})w^j_{k|k-1}}. \tag{19}$$

The cardinality probabilty and expected measurement are approximated as

$$P(|\boldsymbol{Z}_k|=1) \approx \sum_{i=1}^{N}P_D(\boldsymbol{x}^i_{k|k-1})w^i_{k|k-1} \tag{20}$$

and

$$\bar{\boldsymbol{z}}_k \approx \frac{\sum_{i=1}^{N} \boldsymbol{h}(\boldsymbol{x}_{k|k-1}^{i}; A_k^{\pi}) P_D(\boldsymbol{x}_{k|k-1}^{i}) w_{k|k-1}^{i}}{\sum_{i=1}^{N} P_D(\boldsymbol{x}_{k|k-1}^{i}) w_{k|k-1}^{i}}, \tag{21}$$

respectively.

We approximate the expected differential entropy of the weighted particle ensemble using the weighted k-nearest neighbors sample entropy approximation provided in [14]. This is given as

$$\widehat{\mathscr{H}}\big(D_{k|k}(\boldsymbol{x}|\boldsymbol{Z}_{1:k})\big) = -\sum_{i=1}^{N} \frac{\sum_{j\in\mathscr{S}_\ell^i} w_{k|k}^{j}}{\ell} \log \frac{\sum_{j\in\mathscr{S}_\ell^i} w_{k|k}^{j} \Gamma(n_x/2+1)}{\varphi_{i,\ell}^{n_x} \pi^{n_x/2}} + \log(\ell) - \Psi(\ell). \tag{22}$$

To not be confused with the time index k previously introduced, we use ℓ to denote the traditionally used k^{th} nearest neighbors index. The set of indices $\mathscr{S}_\ell^i$ correspond to the $\ell-1$ closest neighbors to particle i. The function $\Gamma(\cdot)$ is the gamma function, n_x is the dimension of the state, and $\varphi_{i,\ell}$ is the Euclidean distance between particle i and its ℓ^{th} nearest neighbor. $\log(\cdot)$ is the natural logarithm, and $\Psi(\cdot)$ is the digamma function. The choice of ℓ affects the bias in this estimate. In [15], the authors show that increasing the value of ℓ works to reduce this bias. For our application, we choose $\ell = 15$ and find this to be sufficient.

The particle-based entropy approximation in Equation 22 is advantageous because it does not require importance sampling of the intensity function $D_{k|k}(\cdot)$. This means that the entropy estimate is a deterministic function for optimization. Secondly, most of the arguments can be computed off-line (i.e., the particle neighbors and particle neighbor distances) which makes it computationally tractable in our application.

3.2 Closed-Loop Tasking Implementation

Here, we detail the steps involved with implementing and simulating a closed-loop tasking scenario for rapid target acquisition. To initialize a simulation, we propagate a space object's orbit until it is contained in the field of regard of a desired ground-based sensor resource. In this case, we assume that this sensor provides a novel detection of the space object. The sensor obtains a measurmenet, and we use the admissible region to instantiate the search domain for follow-up observations. A GMM represents the feasible dynamic states of the admissible region. To generate this, a uniformly sampled grid is placed over the range/rate-rate admissible space with an uncorrelated homoskedastic covariance at each grid point defined in the spherical coordinate domain (i.e., angles/angles-rates and range/range-rate). The unscented transform converts this GMM into the Cartesian state space of the orbit.

The transformed GMM is randomly sampled and the samples are forecast using a Vinti analytic propagator [16, 17] to find potential follow-up times in our sensor network. The Vinti propagator includes the first few zonal spherical harmonic gravity potential terms by transforming and propagating the Cartesian state in oblate spheroidal coordinates. For our case, this includes J_2 and J_3. We choose this propagator for computational efficiency, however, the reader should note that any desired fidelity force-model and analytical/numerical integrator can substitute our approach. Our future work plans to balance the trade between computational efficiency and speed in this particular application by leveraging multi-fidelity methods for this step [18].

A potential follow-up observation exists if 1) there exists a line-of-sight between the sensor and the target, 2) the target is illuminated, and 3) the sensor is not illuminated. To determine these three conditions, we use the `SIGHT` and `LIGHT` algorithms outlined in [19]. The JPL design ephemeris generates the positional states of the sun in determining these calculations [20]. A desired feasible feasible measurement interval for a particular sensor is selected, and the forecast particles are projected into the field of regard of the sensor. We use an R-Tree [21] to store the projected particles, which allows for efficient spatial queries of the particle ensemble as the FOV is translated over the search domain during optimization. We use simulated annealing to optimize the sensor pointing direction at each time step. Specifically, we use the `SciPy` implementation of dual simulated annealing [22].

Upon convergence of the optimizer, we simulate the directed sensor's task and update the CPHD intensity. If the projected true target state is contained within the sensor's directed FOV, we sample a binomial distribution with probability of success given as $\tilde{P}_D(\boldsymbol{x})$ with a single trial. In doing so, we simulate a real follow-up detection. If successful, we update the weights of the particle ensemble according to Equation 19, otherwise the particle weights are updated with Equation 18. In the positive measurement case, we process a randomly drawn measurement sample given as

$$\tilde{\boldsymbol{z}}_k \sim \mathscr{N}(\boldsymbol{h}(\boldsymbol{x}_{\text{true}}; A_k^{\pi}), R), \tag{23}$$

where R is the measurement noise covariance of the sensor.

We re-sample a regularized particle ensemble after the particle weight update. For our implementation, we use a regularization kernel based on the k-nearest neighbors empirical covariance estimate. That is, we re-sample from GMM given as

$$\boldsymbol{x}_{k|k} \sim \sum_{i=1}^{N} w_{k|k}^{i} \mathcal{N}(\hat{\boldsymbol{\mu}}_{k|k}^{i}, \hat{P}_{k|k}^{i}), \tag{24}$$

where,

$$\hat{\boldsymbol{\mu}}_{k|k}^{i} = \sum_{j \in \mathscr{D}^{i}} w_{k|k}^{j} \boldsymbol{x}_{k|k-1}^{j}, \tag{25}$$

$$\hat{P}_{k|k}^{i} = \beta \left[\frac{4}{\ell(n_x+2)} \right]^{\frac{2}{n_x+4}} P_{k|k}^{i}, \tag{26}$$

and,

$$P_{k|k}^{i} = \sum_{i \in \mathscr{D}^{i}} w_{k|k}^{i} (\boldsymbol{x}_{k|k-1}^{i} - \boldsymbol{\mu}_{k|k}^{i})(\boldsymbol{x}_{k|k-1}^{i} - \boldsymbol{\mu}_{k|k}^{i})^{\top}. \tag{27}$$

This approach has some commonalities with the methods introduced in [23]. Here, the set $\mathscr{D}$ are the set of indices corresponding to the ℓ nearest neighbors to particle i. Equation 26 is a kernel estimator based off Silverman's rule of thumb [24], and we apply a multiplicative factor $\beta \in [0,1]$ to account for the non-Gaussian nature of the posterior $D_{k|k}(\cdot)$. In this work, we take $\beta = 0.5$. To summarize the steps involved in this work, we include a flowchart illustrating the closed-loop tasking simulation in Fig. 1.

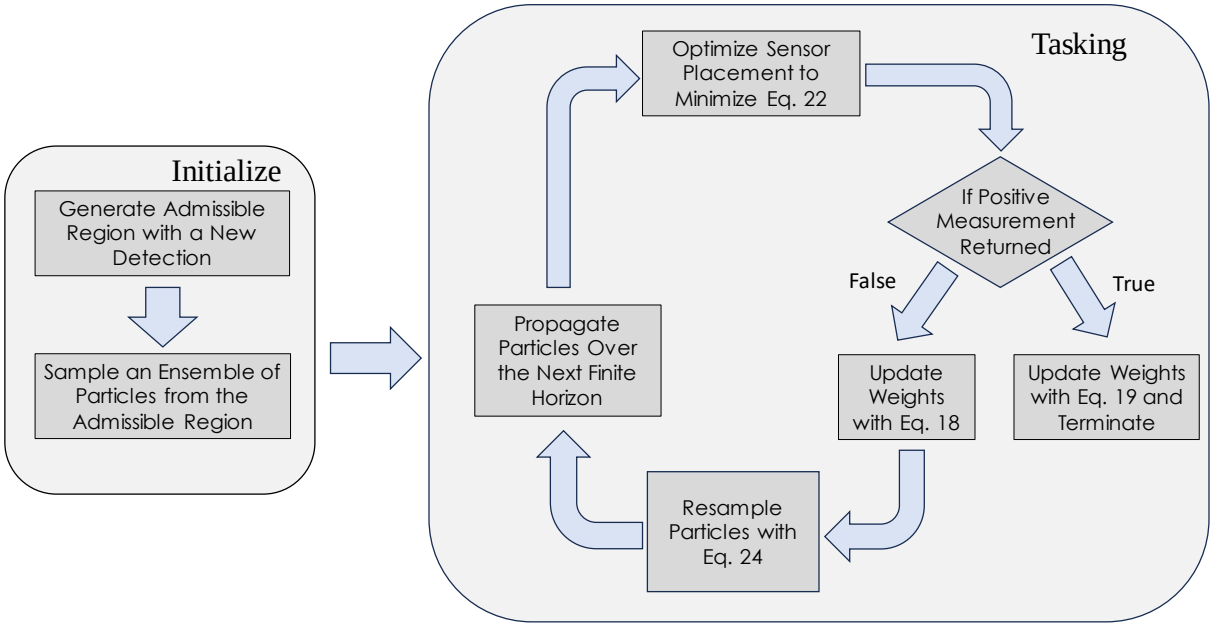

Fig. 1: Flowchart for follow-up sensor tasking simulations.

4. COORDINATING FOLLOW-UP TASKING WITH AN OPTICAL SENSOR NETWORK

Here we show simulated scenarios in which a network of sensors is tasked with acquiring a newly detected space object. Our focus is exclusive to a network of ground-based optical telescopes, however, the methods introduced in this work are more generally applicable to both space- and ground-based sensors of multiple sensing modalities.

The scenarios consists of 15 resources at primary SSA sensor locations which are provided in [25]. We list these locations, along with the associated geodetic latitudes and longitudes in Tab. 1. Note that the sensor modalities used in these simulations may not be consistent with the real world (e.g., some sites may use radars). We take the

Table 1: Optical sensor network configuration for the simulations used in this section.

Location	Latitude [deg]	Longitude [deg]
Eglin	30.57	−86.21
Flyingdales	54.36	−0.67
Cape Cod	41.75	−70.54
Beale	39.13	−121.35
Thule	76.57	−68.30
Cavalier	48.72	−97.90
Ascension	−7.91	−14.40
Clear	64.30	−149.19
Kwajalein	9.40	167.48
Millstone	42.62	−71.49
Vandenberg	34.74	−120.57
Maui	20.70	−156.30
Socorro	34.06	−106.89
Vardo	70.37	31.11
Diego Garcia	−7.32	72.42

probability of detection, $\tilde{P}_D(\boldsymbol{x})$, to be 0.9 for all telescopes. The FOV for each telescope is a 4×4 degree array centered along the telescope's optical axis. The measurement noise standard deviation is $\sigma_{\alpha,\delta} = 3$ arcsec in the angles and $\sigma_{\dot{\alpha},\dot{\delta}} = 0.3$ arcsec/sec for the angular rates.

This work considers two target orbits: a GEO and a GTO. The osculating orbital elements at the start of the simulation are listed in Tab. 2. Three different follow-up tasking scenarios are considered. The first two use the GEO case with

Table 2: Osculating orbital elements for the orbits used in our scenarios.

Test Case	a [km]	e	i [deg]	Ω [deg]	$\nu+\omega$ [deg]	$\nu+\Omega$ [deg]
GEO	42259.0	0.001	0.0	0.0	0.0	0.0
GTO	25447.5	0.66	0.0	0.0	0.0	0.0

one and two follow-up sensors, respectively. For both, an admissible region is instantiated by a telescope located at Ascension Island. The third case uses the GTO example for follow-up with a single sensor. For this case, the admissible region is again instantiated by a telescope located at Ascension Island. Tab. 3 summarizes these test cases. Fig. 2 shows the schedule for the feasible follow-up times based on the initial detection for both the GEO and GTO cases. The epoch time for all scenarios is 2459386.75 JD. A contour plot for the GMM representation of the admissible

Table 3: Description for generating the optical telescope follow-up scenarios. The bounds correspond to those used for generating the admissible region.

Test Case	Detection	Follow-Up	a [10^3 km]	e	r_p [km]
GEO (single sensor)	Ascension, 0.0 hr	Millstone, 2.0 hr	$10 \leq a \leq 45$	$0.0 \leq e \leq 0.7$	$r_p \geq 6778$
GEO (two sensors)	Ascension, 0.0 hr	Millstone & Cape Code, 2.0 hr	$10 \leq a \leq 45$	$0.0 \leq e \leq 0.7$	$r_p \geq 6778$
GTO (single sensor)	Ascension, 0.0 hr	Kwajalein, 1.25 hr	$20 \leq a \leq 30$	$0.5 \leq e \leq 0.7$	$r_p \geq 6778$

region for the GEO cases and GTO case are shown in Fig. 3.

Figure 4 shows evolution of the target's CPHD intensity function as a single telescope searches the feasible state set. Measurements are collected at a 30 second cadence. The blue rectangles in each panel correspond to the directed FOV of the telescope per time step. Each red rectangle is a previous measurement that resulted in a null detection. In these cases, the particles inside the FOV are de-weighted. Upon re-sampling, we see that these regions are populated with

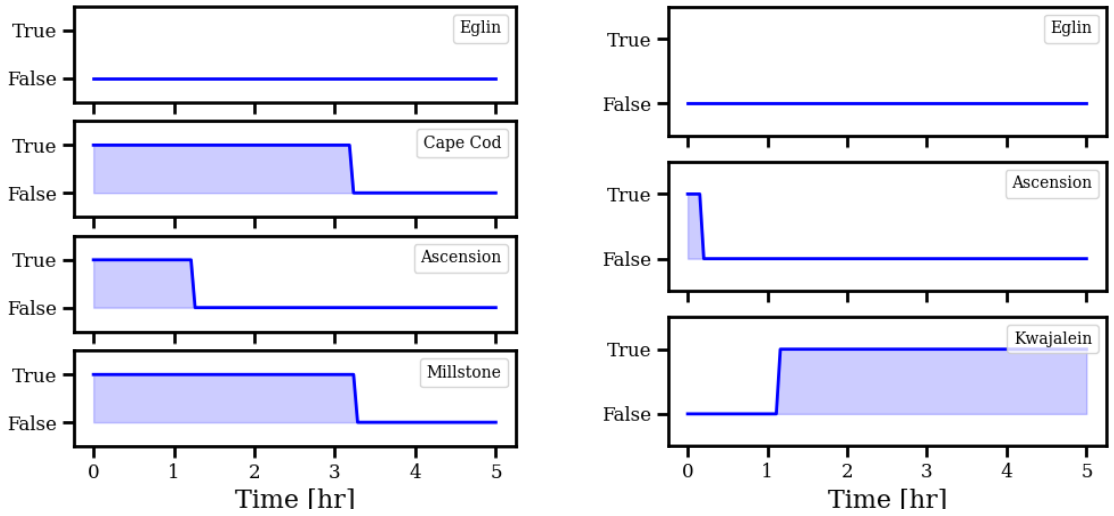

(a) GEO Case (a) GTO Case

Fig. 2: Feasible observation windows for the network of ground-based sensors for (left) GEO and (right) GTO cases. Most sensor do not have a feasible observability window, like the Eglin site. We exclude these here for brevity.

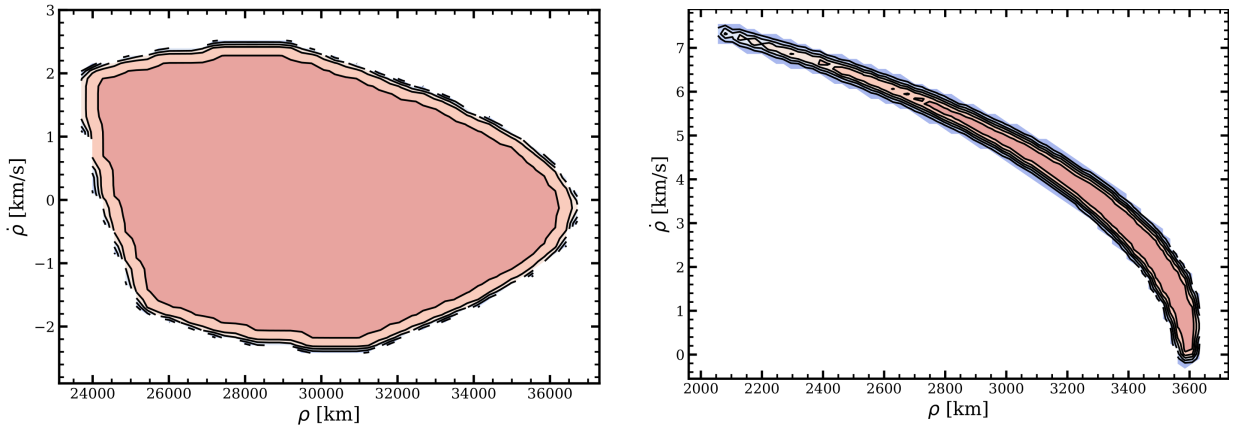

Fig. 3: Optical admissible region for the (left) GEO and (right) GTO test cases contained in Tab. 3. The admissible regions are instantiated by angles and angles rates measurements returned from an optical telescope located on Ascension Island.

fewer, but a non-zero number of particles, as the probability of detection is not unity. Regions not yet visited or visited multiple time-steps prior are subsequently populated with more particles. This effect is sometimes referred to as the "spooky effect" after the quantum mechanics phenomenon "spooky action at distance" [26].

The blue diamond in each panel is the sampled true right ascension and declination of the target. The last panel demonstrates the collapse in target uncertainty resulting from a follow-up true positive detection. For this particular simulation realization, a total of 7 measurement iterations are required before a positive detection is returned.

A very similar plot is shown in Fig. 5 but for the GTO test case in this study. Just 75 minutes after target detection, the initial admissible region diffuses into a sizable portion of the follow-up sensor's field of regard despite the conservative

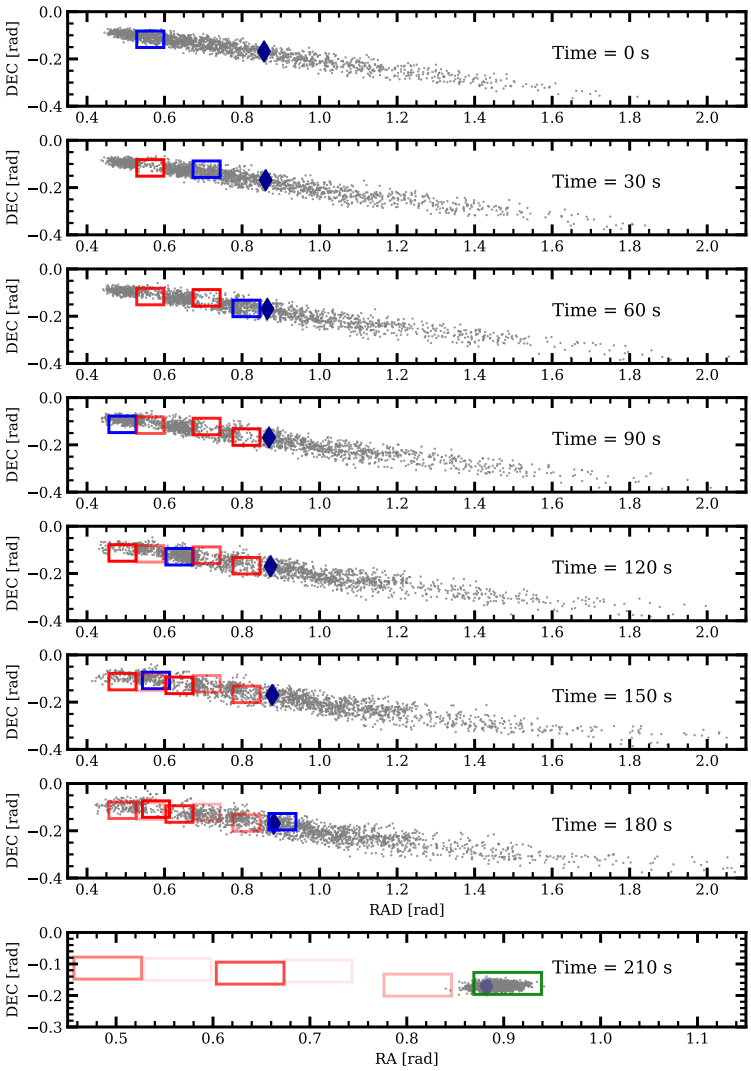

Fig. 4: Search space evolution with a single-sensor follow-up for the GEO test case. Each panel shows the planned action by the blue FOV, and the previous actions that resulted in a null detection by the red FOV. Based on new information, the sensing agent plans next action to minimize the expected entropy of the posterior CPHD intensity. The time displayed in each panel corresponds to the time since the start of a pass. A true positive detection is shown by the green rectangle in panel 8.

admissible region constraints used in Tab. 3. Future work intends to explore how to properly partition this region when it may overlap with non-observable portions from the sensor's point-of-view. For the particular scenario realization, a total of 9 measurements are required before receiving a positive detection of the object. For this test case, we also perform a Monte-Carlo analysis to characterize the distribution in the number of measurement scans require to return a positive detection of the space object. 100 Monte-Carlo trials are used, of which, 75 return a positive detection within the first 30 scans, equating to 15 minutes of telescope time. We show a histogram of this result in Fig. 6. The mean number of scans for these cases is 11.6.

In Fig. 7 we see the evolution of the search space in the plane of two sensors performing a coordinated search. In this case, two optical telescopes located at Millstone, MA and Cape Code, MA. These two sensors are geospatially close, so the feasible search space appears nearly identical for the two, however, this is not necessarily the case for all sensors and is not required for our approach. This particular scenario realization requires 4 observations from each telescope. The Cape Cod telescope receives a positive detection at 90 seconds, collapsing the uncertainty for both sensing agents

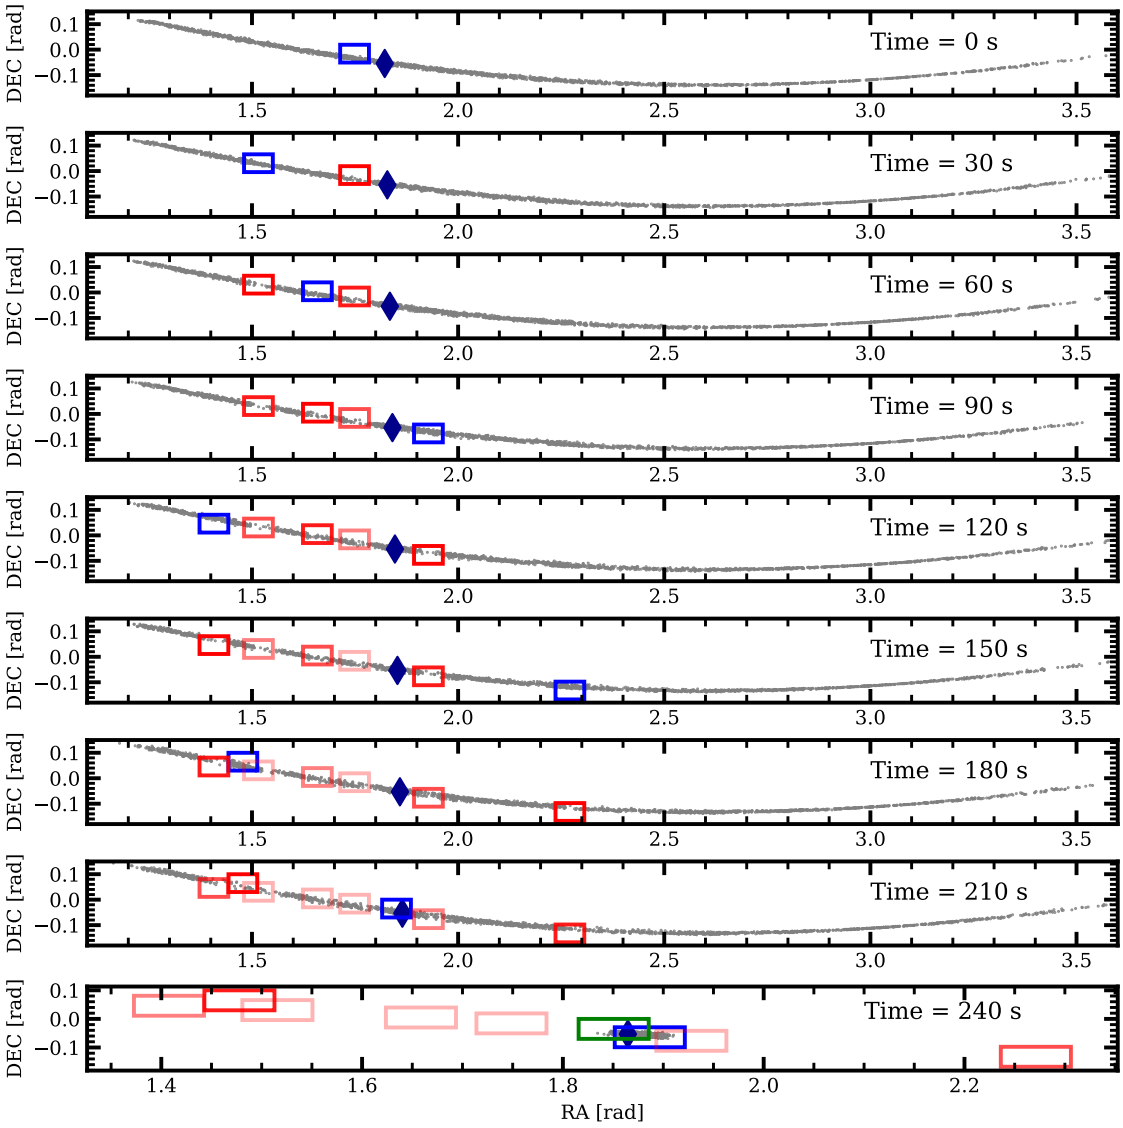

Fig. 5: Search space evolution with a single-sensor follow-up for the GTO test case. Each panel shows the planned action by the blue FOV, and the previous actions that resulted in a null detection by the red FOV. Based on new information, the sensing agent plans next action to minimize the expected entropy of the posterior CPHD intensity. The time displayed in each panel corresponds to the time since the start of a pass. A true positive detection is shown by the green rectangle in panel 9.

at time step 5.

5. CONCLUSIONS

This work introduces a method for searching a feasible state set with one or multiple sensors to rapidly gain custody of a previously detected space object. We use an RFS framework so that the set of feasible states is represented by the CPHD intensity function. A tasking agent optimizes its sensor placement to result in minimum expected differential entropy of the posterior CPHD intensity over a finite-time horizon. Differential entropy directly quantifies the uncertainty associated to the target state which leads a tasking agent to explore regions most informative, regardless of whether the said region has been previously visited. This is particularly useful when the probability of detection is less than unity. A true positive detection cannot be assumed for every action and so determining the expected

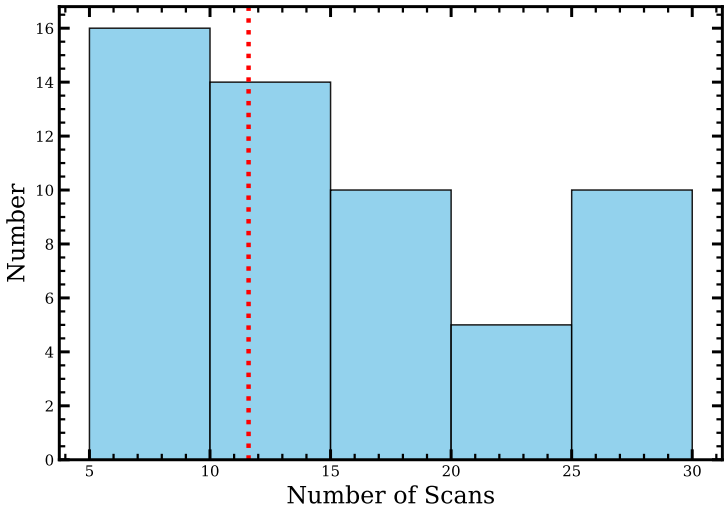

Fig. 6: Histogram for the number of measurement scans required to receive a positive detection for the GTO test case. Of the 100 Monte-Carlo trials, 75 produced a positive detection within the first 30 scans, which equates to 15 minutes of telescope time. The red line indicates the mean number of scans required of the 75 that produced a positive detection which is 11.6.

differential entropy requires enumerating the predicted effect of all measurement cardinality combinations on the posterior CPHD intensity function within the finite time horizon. To manage this computational challenge, in this work, we use a single step look-ahead strategy.

We use a sequential Monte-Carlo implementation of the CPHD filter for closed-loop sensor tasking. For this, the CPHD intensity function is represented by an ensemble of particles and particle weights. In each iteration, upon optimizing the sensor placement, the sensor collects a measurement, and the particle weights are updated in accordance to whether a positive or negative detection is obtained. The particle ensemble is re-sampled with a regularization kernel based on Silverman's rule-of-thumb.

The procedure is tested for two orbital cases of relevance to the SSA community: a GEO and a GTO case. We test the GEO case with a one- and two-sensor follow-up scenario. The GTO case is tested with a single sensor follow-up. For all cases, the simulation terminates with a true-positive detection of the target space object.

There are a few potential avenues that we intend to explore in future work. First, we would like to loosen the assumptions placed in the problem formulation, that is, consider situations in which there is an unknown number of targets contained in a search domain and a non-zero distribution of clutter returns. As mentioned previously, calculating the expected differential entropy presents a computational challenge for direct real-time optimization. One avenue that we intend to explore is the application of reinforcement learning to this problem. Finally, while we only use ground-based optical telescopes in all simulations presented in this work, future work will relax this to consider scenarios consisting of ground- and space-based sensor resources with multiple sensing modalities. One scenario that we are particularly interested in is a ground-based telescope co-located with a satellite laser ranging system.

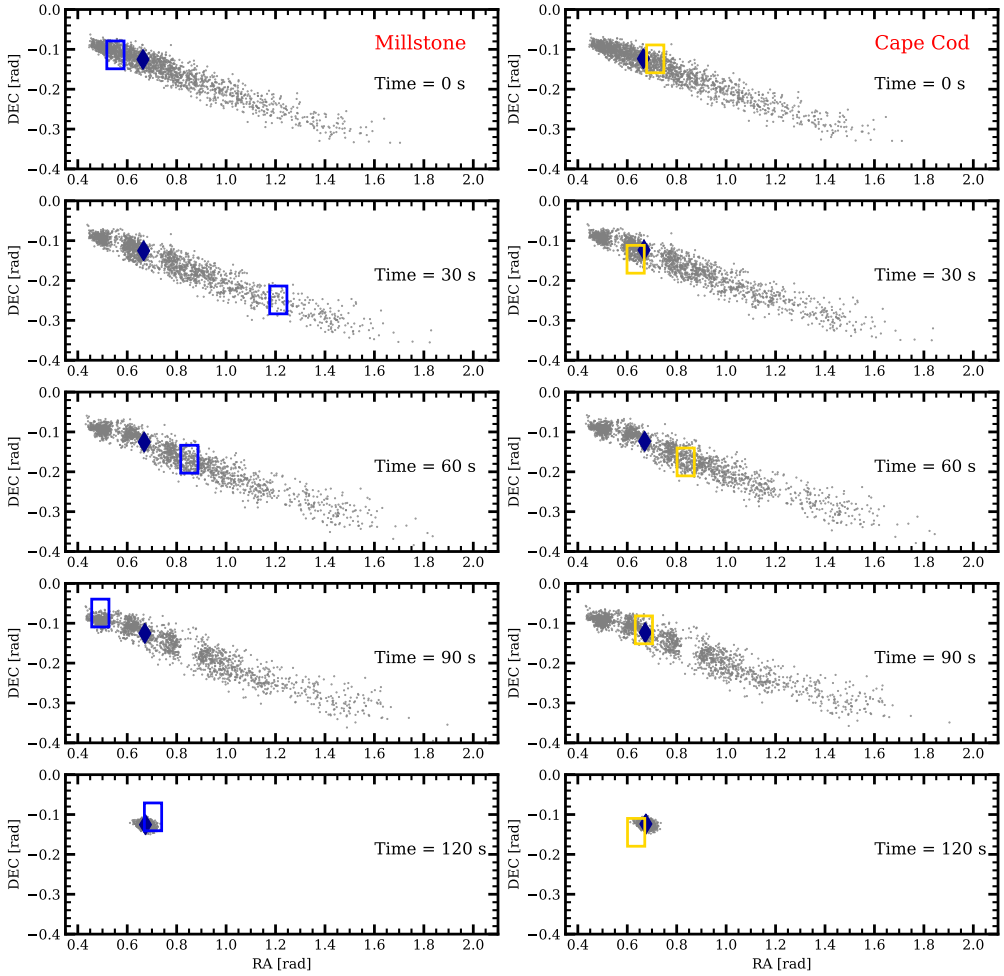

Fig. 7: Search space evolution with a two-sensor follow-up for the GEO test case. The left panels correspond to the search space visible by a telescope at Millstone, and the right is visible from Cape Cod. Each panel shows the planned action by the blue or gold FOV. Based on new information obtained from both agents, each sensing agent then plans next action to minimize the expected entropy of the posterior CPHD intensity. The time displayed in each panel corresponds to the time since the start of a pass. A true positive detection is obtained by Cape Cod after 4 measurment cycles.

ACKNOWLEDGEMENTS

This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA9453-21-2-0064. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements , either expressed or implied, of Air Force Research Laboratory (AFRL) and or the U.S. Government.

REFERENCES

[1] Andrea Milani, Giovanni F Gronchi, Mattia De' Michieli Vitturi, and Zoran Knežević. Orbit determination with very short arcs. i admissible regions. *Celestial Mechanics and Dynamical Astronomy*, 90:57–85, 2004.

[2] Kohei Fujimoto and Daniel J. Scheeres. Applications of the admissible region to space-based observations. *Advances in Space Research*, 52(4):696–704, 2013.

[3] Kyle J. DeMars and Moriba K. Jah. Probabilistic initial orbit determination using gaussian mixture models. *Journal of Guidance, Control, and Dynamics*, 36(5):1324–1335, 2013.

[4] T. S. Murphy and M. J. Holzinger. Generalized Minimum-Time Follow-up Approaches Applied to Tasking Electro-Optical Sensor Tasking. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 54, January 2017.

[5] Samuel J. Fedeler, Marcus J. Holzinger, and William W. Whitacre. Tasking and Estimation for Minimum-Time Space Object Search and Recovery. *Journal of the Astronautical Sciences*, 69(4):1216–1249, August 2022.

[6] Tyler A. Hobson and I. Vaughan L. Clarkson. A particle-based search strategy for improved space situational awareness. In *2013 Asilomar Conference on Signals, Systems and Computers*, pages 898–902, 2013.

[7] Steven Gehly, Brandon Jones, and Penina Axelrad. Sensor allocation for tracking geosynchronous space objects. *Journal of Guidance, Control, and Dynamics*, 41(1):149–163, 2018.

[8] Steven Gehly, Brandon A. Jones, and Penina Axelrad. Search-detect-track sensor allocation for geosynchronous space objects. *IEEE Transactions on Aerospace and Electronic Systems*, 54(6):2788–2808, 2018.

[9] Nagavenkat Adurthi, Puneet Singla, and Manoranjan Majji. Mutual information based sensor tasking with applications to space situational awareness. *Journal of Guidance, Control, and Dynamics*, 43(4):767–789, 2020.

[10] Nicholas Ravago and Brandon A. Jones. Risk-aware sensor scheduling and tracking of large constellations. *Advances in Space Research*, 68(6):2530–2550, 2021.

[11] Ronald Mahler. Phd filters of higher order in target number. *IEEE Transactions on Aerospace and Electronic Systems*, 43(4):1523–1543, 2007.

[12] B. Ristic, D. Clark, Ba-Ngu Vo, and Ba-Tuong Vo. Adaptive target birth intensity for phd and cphd filters. *IEEE Transactions on Aerospace and Electronic Systems*, 48(2):1656–1668, 2012.

[13] Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

[14] Jiří Ajgl and Miroslav Šimandl. Differential entropy estimation by particles. *IFAC Proceedings Volumes*, 44(1):11991–11996, 2011. 18th IFAC World Congress.

[15] Nikolai Leonenko, Luc Pronzato, and Vippal Savani. A class of Rényi information estimators for multidimensional densities. *The Annals of Statistics*, 36(5):2153 – 2182, 2008.

[16] John P. Vinti. New method of solution for unretarded satellite orbits. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, page 105, 1959.

[17] Ashley D. Biria and Ryan P. Russell. Analytical Solution to the Vinti Problem in Oblate Spheroidal Equinoctial Orbital Elements. *Journal of the Astronautical Sciences*, 67(1):1–27, May 2019.

[18] Brandon A. Jones and Ryan Weisman. Multi-fidelity orbit uncertainty propagation. *Acta Astronautica*, 155:406–417, 2019.

[19] David A Vallado. *Fundamentals of astrodynamics and applications*, volume 12. Springer Science & Business Media, 2001.

[20] William M. Folkner, James G. Williams, and Dale H. Boggs. The planetary and lunar ephemeris DE421. IPN Progress Report 42-178, Jet Propulsion Laboratory, California Institute of Technology, http://ipnpr.jpl.nasa.gov/progress_report/42-178/178C.pdf, August 2009.

[21] Antonin Guttman. R-trees: A dynamic index structure for spatial searching. 14(2):47–57, jun 1984.

[22] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[23] Andrey A Popov and Renato Zanetti. Ensemble gaussian mixture filtering with particle-localized covariances. In *2023 26th International Conference on Information Fusion (FUSION)*, pages 1–7. IEEE, 2023.

[24] C. D. Kemp. Density Estimation for Statistics and Data Analysis. *Journal of the Royal Statistical Society Series D: The Statistician*, 36(4):420–421, 12 2018.

[25] Eun-Jung Choi, Sungki Cho, Jung Hyun Jo, Jang-Hyun Park, Taejin Chung, Jaewoo Park, Hocheol Jeon, Ami Yun, and Yonghui Lee. Performance Analysis of Sensor Systems for Space Situational Awareness. *Journal of Astronomy and Space Sciences*, 34:303–314, December 2017.

[26] Ba Tuong Vo and Ba Ngu Vo. The para-normal bayes multi-target filter and the spooky effect. In *2012 15th International Conference on Information Fusion*, pages 173–180, 2012.