(Preprint) AAS 16-453

DIRECT IMAGE-TO-LIKELIHOOD FOR TRACK-BEFORE-DETECT MULTI-BERNOULLI FILTER

Timothy S. Murphy ; Marcus J. Holzinger ; Brien Flewelling;

This paper aims to apply the random finite set-based multi-Bernoilli filter to frame-to-frame tracking of space objects observed in electro optical imagery for space domain awareness applications. First, this paper will review random finite set filters applied to frame to frame tracking and their applications to space. A new likelihood function for space based imagery will be presented, based on the matched filter. A more educated birth model will be proposed which better models potential SO using observer characteristics and object dynamics. Simulation results will explore the range of objects that can be tracked. The final algorithm is able to perform completely uncued detection down to a total object SNR of 5.6 and a per pixel SNR of 1.5. Promising but inconclusive results are shown for total object SNR of 3.35 and per pixel SNR of 0.7.

INTRODUCTION

Part of space domain awareness (SDA) is the task of maintaining sufficient knowledge of the space domain to inform decisions as they relate to space assets, national security, and commercial ventures. Space situational awareness (SAA) is the task of providing all relevant knowledge to a particular mission at a particular time. Both directives necessitate orbit determination efforts for small space objects which often produce only low magnitude electro-optical signatures. In particular, methods must allow for uncued identification of objects, or identification cued with only partial orbital knowledge. Rate tracking is impossible for object discovery, and blindly searching image series for an unknown signal is inefficient. This will necessitate a method which incorporates, but does not require, prior knowledge in a general way, while rigorously incorporating all information from measurement sources for completely uncued detection. Furthermore, the required method should search a state space in a smart way to find low SNR signals.

Due to the non-linear nature of space, general methods derived from Bayesian filtering [13] have become popular for space object tracking and orbit determination [3],[2]. When considering the high dimensionality of space, such particle based filters can become difficult to implement [1]. This paper will instead looking at the more simple case of frame-to-frame tracking in images obtained from EO sensors. This allows a reduced dimensionality in the state space by tracking only position and velocity in the image plane. Such a method should be as general as possible, allowing for all relevant objects to be tracked. Much work has been done on finite set statistics (FISST) based filters

*Graduate Student, The Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Georgia Institute of Technology North Ave NW, Atlanta, GA 30332.

†Assistant Professor, The Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Georgia Institute of Technology North Ave NW, Atlanta, GA 30332. AIAA Senior Member

‡Research Aerospace Engineer, Space Vehicles Directorate, Air Force Research Laboratory, 3550 Aberdeen Ave. SE, Kirtland AFB, NM, USA.

in the past for a variety of applications [8]. In particular, recent pushes in FISST filter theory in the field of computer vision has looked multi-target tracking in images [6], [5], [4]. In particular, this work operates directly on pixel data and requires no detection algorithm. Very recently, the SDA community has begun to look at FISST filters [7] and detectionless frame-to-frame tracking [4].

The first problem this paper will analyze is the likelihood function. The current state of the art likelihood function was proposed by Vo et. al. [14]. This likelihood function is intended for randomly moving point sources. For space object, movement happens in a predictable way. Furthermore, object signal has structure based on object movement. Instead, this paper will use the matched filter predicted by a particular particle to evaluate measurement likelihood [15] [10]. This likelihood function should effectively search for low SNR signals, because the matched filter is the SNR optimal linear image filter.

Another area this paper will improve upon is the particle birth model. Typical Bernoulli filters generate new particles with a random uniform distribution. However, the intention of the birth model is that it should be catered to a particular problem. This paper will therefore develop a method for bounding the types of signals that can occur based on sensor and algorithm requirements and orbital mechanics.

This paper will be presented as follows. Section will review the basics of Bernoulli and multi-Bernoulli filtering. Section , subsection will discuss the likelihood function based on a matched filter. Section , subsection will discuss a variety of ways to constrain a birth model for frame-to-frame tracking. First some preliminary observations on birth models will be explained. This paper will briefly explain how to use partial prior orbital knowledge as a birth model. Then uncued detection birth models will be discussed. Simulations will implement the likelihood function and some of the birth models for a particular problem. Computational tractability analysis will be discussed, primarily as a future work analysis.

BACKGROUND ON RANDOM FINITE SET FILTERING FOR FRAME TO FRAME TRACKING

Dynamics

Consider the standard filtering problem with discrete dynamics and measurement models

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k) + \mathbf{w}_1 \tag{1}$$

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{w}_2 \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{z} \in \mathbb{R}^m$, $\mathbf{w}_1 \in \mathbb{R}^n \sim f_{w_1}$, and $\mathbf{w}_2 \in \mathbb{R}^m \sim f_{w_2}$. $\mathbf{x}$ can be thought of as the state to be estimated, while $\mathbf{z}$ can be thought of as the measurements. For frame-to-frame image tracking, the measurement $\mathbf{z}_k$ is a series of m pixels, $\{z_i\}_{i=1}^m$. The state, $\mathbf{x}$, will be modeled as position and velocity, in the frame, measured in pixels.

$$\mathbf{x} = [x,\ y,\ \dot{x},\ \dot{y}]^T \tag{3}$$

Multi-Bernoulli Filter

This paper will use the multi-Bernoulli filter used in [14]. A brief review of the theory and application will be presented for the readers convenience. The Bernoulli filter starts from Bayesian Filtering equations

$$f(\mathbf{x}_k|\mathbf{z}_{k-1}) = \int f(\mathbf{x}_k|\mathbf{x}_{k-1}) f(\mathbf{x}_{k-1}|\mathbf{z}_{k-1}) \qquad \text{(Prediction)} \tag{4}$$

$$f(\mathbf{x}_k|\mathbf{z}_k) = \frac{f(\mathbf{z}_k|\mathbf{x}_k) f(\mathbf{x}_k|\mathbf{z}_{k-1})}{f(\mathbf{z}_k|\mathbf{z}_{k-1})} \qquad \text{(Update)} \tag{5}$$

where $\mathbf{z}_k$ is the time series of measurements up until time step k. Equation 4 is the prediction step while Equation 5 is the update step.

For a Bernoulli filter, the state is modeled as a Bernoulli random finite set (BRFS). The BRFS is a set $\mathcal{S}_k$ containing a RFS which contains 1 object with probability r_k and empty with probability $1 - r_k$. The RFS position is described by the PDF, $p(\mathbf{x}_k)$. In other words, the object is described by a PDF and probability of existence.

$$\mathcal{S} = \{p(\mathbf{x}_k), r_k\} \tag{6}$$

The Bernoulli filter estimates the PDF with equations (4) and (5) and updates the probability of existence with update equations that can be found in [14].

Now, the update equations for a multi-target tracking scheme are

$$\pi_{k|k-1}(\mathcal{X}_k|\mathbf{z}_{1:k-1}) = \int f_{k|k-1}(\mathcal{X}_k|\mathcal{X}) \pi_{k-1}(\mathcal{X}|\mathbf{z}_{1:k-1}) \partial\mathcal{X} \qquad \text{(Prediction)} \tag{7}$$

$$\pi_k(\mathcal{X}_k|\mathbf{z}_{1:k}) = \frac{g(\mathbf{z}_k|\mathcal{X}_k) \pi_{k|k-1}(\mathcal{X}_k|\mathbf{z}_{1:k-1})}{\int g(\mathbf{z}_k|\mathcal{X}_k) \pi_{k|k-1}(\mathcal{X}|\mathbf{z}_{1:k-1}) \partial\mathcal{X}} \qquad \text{(Update)} \tag{8}$$

where $g(\cdot|\cdot)$ is the likelihood function. The distribution π_k is now a multi-object belief function. In essence, π_k is similar to a PDF in that higher density means higher probability of an object existing in that area. However, π_k does not need to integrate to 1; instead, the total mass of an area is the expected number of objects in that area.

For the multi-object filter, $\mathcal{X}_k$ is a random finite set used to represent the multi-object tracking problem where the number of objects is unknown. The multi-Bernoulli filter represents the multi object state as the union of a series of BRFSs

$$\mathcal{X}_k = \bigcup \mathcal{S}_k \tag{9}$$

This allows the implementation of a Multi-Bernoulli filter to reduce to a series of Bernoulli filters. There are particle implementations of Bernoulli filters [12], and the union of these filters is relatively straight forward to compute. This paper will focus on applying the Multi-Bernoulli filter which operates on image frame-to-frame tracking to EO sensors tracking. The primary change will be defining a likelihood function $g(\cdot|\cdot)$ which marries the space imagery and Multi-Bernoulli filter. A more in depth, general analysis of FISST can be seen in [8].

THEORETICAL RESULTS

Likelihood Function

The current established likelihood function for tracking point sources in images is from [14]. This has already been used in SO tracking [4].

This method first defines a series of pixels, $T(\mathbf{x})$, which are predicted to have signal from an object $\mathbf{x}$. Pixels are predicted to have mean contribution $h_i(\mathbf{x})$ from a state x. This gives

$$p(z_i|\mathbf{x}) = \left\{ \begin{array}{ll} \psi_i(z_i) & i \in T(\mathbf{x}) \\ \phi_i(z_i) & i \notin T(\mathbf{x}) \end{array} \right\} \tag{10}$$

One common assumed distribution for ψ and ϕ [14] is

$$\psi(z_i) = \mathcal{N}(z_i; 0, \sigma^2) \tag{11}$$

$$\phi(z_i, \mathbf{x}) = \mathcal{N}(z_i; h_i(\mathbf{x}), \sigma^2) \tag{12}$$

where $\mathcal{N}(\cdot; \mu, \sigma^2)$ is a Gaussian with mean μ and variance σ^2. The likelihood function then becomes

$$g(z|X) = \left(\prod_{\mathbf{x} \in X} \prod_{i \in T(\mathbf{x})} \psi_i(z_i, \mathbf{x}) \right) \left(\prod_{i \notin \cup T(\mathbf{x})} \phi_i(z_i) \right) \tag{13}$$

The probability of existence update is formulated as a particle-wise relative likelihood.

$$g_z(\mathbf{x}) = \prod_{i \in T(\mathbf{x})} \frac{\phi(z_i, \mathbf{x})}{\psi(z_i)} \tag{14}$$

This likelihood function is general enough to work for the case of arbitrary point tracking. This paper will propose a new likelihood function based on a hypothesis test on the matched filter. Such a likelihood function has been used on a particle filter by the authors already [11].

The matched filter is a SNR optimal linear image filter when the form of the signal is known. Similar to the already shown likelihood function, the measured signal is assumed to be zero mean noise and predicted signal $h_i(\mathbf{x})$ in a series of pixels $T(\mathbf{x})$. The matched filter is the the weighted sum of the measured pixels, z_i, weighted by the predicted values, $h_i(\mathbf{x})$.

$$z_{MF} = \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) z_i \tag{15}$$

Note that the SNR of the matched filter is scale invariant with respect to the predicted values $h_i(\mathbf{x})$. This is because an increase in scale equally elevates both the signal and noise. For the matched filter, $h_i(\mathbf{x})$ need only describe the relative values predicted for an image. This is an important innovation, because for SO object discovery, actual object brightness is completely unknown. Under the prevision distribution assumptions, z_{MF} is distributed as follows

$$\mathbb{E}[z_{MF}] = \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) h_i(\mathbf{x}_{true})$$
$$\text{Var}[z_{MF}] = \mathbb{E}\left[\left(\sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) w_i\right)^2\right] = \sigma_w \sum_{i \in T(\mathbf{x})} h_i^2$$
$$z_{MF} \sim \mathcal{N}(\mathbb{E}[z_{MF}], \alpha \sigma_w^2) \tag{16}$$

where w_i is the noise in pixel i, and α is a scaling factor dependent on the actual weights, h_i. Note that because h_i is scale invariant, the values of h_i can be chosen such that $\alpha = 1$. Note that if the pixel in $T(\mathbf{x})$ are all noise only, then the expected value of $h_i(\mathbf{x}_{true})$ over $T(\mathbf{x})$ will be zero. Otherwise, the expected value for the MF will be greater than zero. A hypothesis test is desired for determining if there is signal present which is effecting z_{MF}. This equivalent to asking if z_{MF} has a mean of zero or a mean greater than zero. The matched filter will be assumed to exist in one of two distributions

$$\tilde{\psi}(z_{MF}) = \mathcal{N}(z_{MF}; 0, \sigma_{MF}^2) \tag{17}$$
$$\tilde{\phi}(z_{MF}) = \mathcal{N}(z_{MF}; \beta, \sigma_{MF}^2), \ \beta > 0 \tag{18}$$

where β will be set equal to the value calculated from a particular matched filter. This allows the definition of the following null and test hypotheses for a binary hypothesis test

$$H_0 : z_{MF} \sim \tilde{\psi}(z_{MF})$$
$$H_1 : z_{MF} \sim \tilde{\phi}(z_{MF}) \tag{19}$$

In essence, this test assumes the matched filter was performed on pure zero mean noise, and asks whether the calculation gives significant evidence of underlying signal. For binary hypothesis testing, a probability of false alarm is set, p_{FA}, which in term defines an integration threshold, z_{TH}, based on the null hypothesis PDF.

$$p_{FA} = \int_{z_{TH}}^{\infty} \tilde{\psi}(z) dz \tag{20}$$

The test hypothesis PDF is then integrated over, giving a probability of detection.

$$p_{Detect} = \int_{z_{TH}}^{\infty} \tilde{\phi}(z) dz \tag{21}$$

For more on this subject, see [9]. This hypothesis test can be used to determine if there is significant evidence that the predicted signal exists in the predicted location. Because the matched filter

gives a SNR gain, this test should maximize p_{Detect}, though an explicit proof of this claim has not yet been shown.

The Probability of existence update can also be formulated in terms of the matched filter. The relative likelihood can be calculated from the two distribution in Equation (18). This leads to the particle-wise relative likelihood

$$g_z(\mathbf{x}) = \frac{\tilde{\phi}(z_{MF})}{\tilde{\psi}(z_{MF})} \tag{22}$$

Birth Model

As the Multi-Bernoulli filter runs, multiple single Bernoulli filters run in parallel. New Bernoulli filters are constantly added into the system to promote exploration of the design space. These new filters are initially sampled from a birth model, which attempts to predict the kind of states a new track could have. Birth models have been proposed and used in the past. This paper will develop a particular birth model for searching for a SO. Specifically, it will define a series of constraints which will allow stars and improbable orbits to be ignored. This birth model should also incorporate previous partial orbit knowledge.

Theoretically, the birth model should enforce a series of constraints in the form of subsets of $\mathbb{R}^6$, that is, the space of possible orbits represented as a position and velocity vector, $\mathbf{r}, \dot{\mathbf{r}}$. If a series of c constraint sets are defined, $\{\mathcal{S}_i\}_{i=1}^{c}$, the birth model would then just be an uninformed prior over the intersection of those sets.

$$\mathcal{S} = \bigcap_{i=1}^{n} \mathcal{S}_i \tag{23}$$

$$\Gamma = \begin{cases} \frac{1}{\int \mathcal{S}} & \text{if } \mathbf{x}' \in \mathcal{S} \\ 0 & \text{if } \mathbf{x}' \notin \mathcal{S} \end{cases} \tag{24}$$

where Γ is the birth model to sample from, and $\mathbf{x}' \in \mathbb{R}^6$ is an orbit.

The first assumption on our birth model will be of an earth orbiting object. This will be done with a maximum on orbital specific energy, $\mathcal{E} < 0$, and a minimum on radius of periapse constraint $r_p > r_{Earth}$. The second assumption is that the object exists in the field of view of the sensor. This is effectively a subset of $\mathbb{R}^3$ defining possible positions an object could exist at. This birth model does not yet restrict the problem very much. An Earth orbiting object could, as defined above, have a huge range of possible velocities. This means that the above model will sample too large a space to be useful. Another problem that is immediately obvious, is that sampling a subset of $\mathbb{R}^6$ is a very difficult task. Instead, it is more efficient to take constraints and map them into $\mathbb{R}^4$ Equations (23) and (24) still apply, but the series of constraints, $\{\mathcal{S}_i\}_{i=1}^{c}$, are defined as subsets of $\mathbb{R}^4$.

Partial Prior Knowledge In some cases, the orbit of an object or objects, while not completely known, may have some prior knowledge. This includes a previous observation and associated admissible region, a known event (such as a break up event), or on object that actively maneuvered from a previous orbit. In all these cases, if the prior information is restrictive enough, the prior can be sampled and used as a birth model.

For example, consider a satellite break-up event. If a known object break up can be traced back to a specific time, it can be assumed that all objects from the break up had a known position at a known time. Furthermore, a restrictive set of velocities can be calculated (excluding ballistic trajectories and estimating possible post break up velocities). Through such a method, a good birth model for searching for a specific class of objects can be synthesized.

This paper will focus on when prior knowledge is not available, so this topic will not be explored further in this paper. Similar topics are explored in [10]. Future work by the authors will look into this case more closely.

No Prior Knowledge When no prior model is available, there is still one major innovation available for a birth model. In order for a Bernoulli filter to reliably track an object, the object must exist in the frame of the image for a minimum number of frames. If an object is in the top left corner of an image with a velocity headed out of the frame, the filter will never be able to track it. Instead, a minimum number of observations should be enforced within the birth model This defines a birth model within the image as a uniform distribution over position and a uniform distribution over constrained range of velocities. The equations in x and y are independent and so the birth model will be derived in one dimension.

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

$$x \in [x_{min},\ x_{max}] \tag{25}$$

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

$$\dot{x} \in \left[\frac{x_{min} - x}{n_0 \Delta t},\ \frac{x_{max} - x}{n_0 \Delta t}\right] \tag{26}$$

Note that the size of this interval is

$$\frac{x_{max} - x}{n_0 \Delta t} - \frac{x_{min} - x}{n_0 \Delta t} = \frac{x_{max} - x_{min}}{n_0 \Delta t} \tag{27}$$

This is important because it is independent of location. This implies that all possible values of x will have equal probability mass and can therefore be sampled independently. This birth model can be implemented by independently sampling $[x,\ y]$ from (25), then sampling $[\dot{x},\ \dot{y}]$ from an appropriately formed (26).

A further constraint will be used in this paper. Assume the observer is located on earth. This allows the question, what is the smallest possible relative angular velocity that an object will move at with respect to the observer? This should be an object that is as far away as possible moving at the lowest velocity possible. This will correspond to a geostationary transfer orbit (GTO).

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

$$v_a = \sqrt{\frac{\mu}{r_a}}\sqrt{\frac{2}{1 + \frac{r_a}{r_p}}} = 1.586\frac{\text{km}}{\text{s}} \tag{28}$$

Table 1. Orbital Elements for Simulated Objects

	a	e	Om	i	om	f
Obj. 1	42165	0.0044	0	0.0017	3.1163	1.5699
Obj. 2	42165	0.0087	4.6862	0.5237	1.5734	4.7063
Obj. 3	42165	0.0087	1.5206	0.1747	4.7417	4.7063

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

$$\tan^{-1}\left(\frac{1.568 \text{ km/s}}{42642 \text{ km}}\right) = 0.002107^\circ/\text{s} \tag{29}$$

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

Under the right circumstances, the minimum velocity constraint on the birth model can remove a large amount of the velocity space that needs to be sampled. In the simulations that will be done in this paper, this constraint will remove approximately 1/3 of the possible velocities. Just as importantly, this constraint intelligently does not contain stars in the space of hypothesized signals. Through re-sampling, velocities may drift In cases with a different field of view and resolution, this constraint may be less effective.

SIMULATION RESULTS

Problem Set Up

The multi-Bernoulli filter is shown in this section tracking objects in simulated data. In the simulation, a sensor is modeled as a simple pinhole camera with a 2 degrees field of view and 512 pixel resolution in both x and y. Three objects are simulated with various orbital characteristics shown in Table 1. Note that small amounts of eccentricity and inclination are added to all orbits to give a range of in plane headings and disambiguate orbital elements. An observer is simulated at Georgia Tech at time $UTC = [2015, 11, 12, 20, 00, 00]$ performing a fixed stare at $\mathbf{r} = [0; -42164; 0]$. This is effectively a sidereal stare tasking.

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

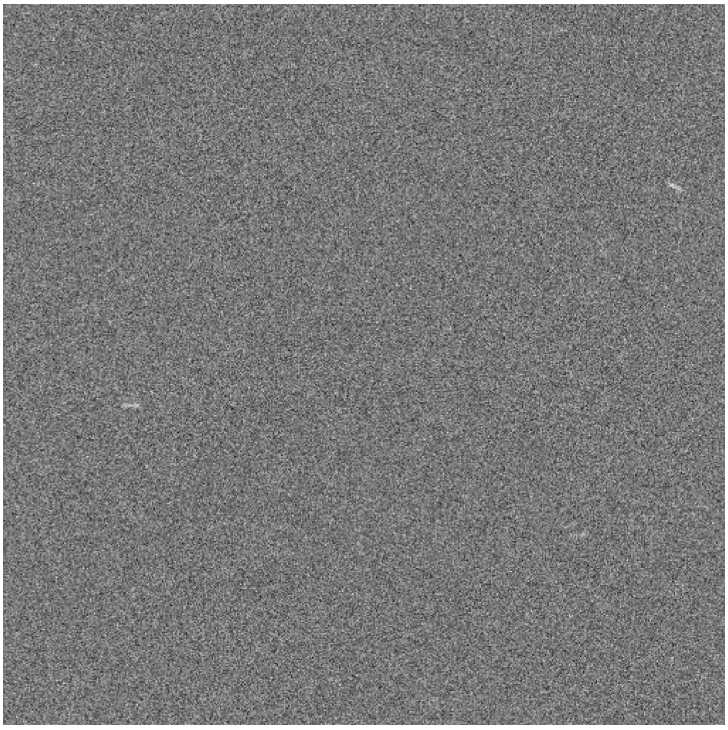

Figure 1. Simulated Image

pixels) can be calculated via [11]

$$\begin{aligned} \text{SNR}\left(\sum y_j(\mathcal{T})\right) &= \frac{\mathbb{E}[\sum y_j(\mathcal{T})]}{\sqrt{\sum \mathbb{E}[(y_j(\mathcal{T}) - \mathbb{E}[y_j(\mathcal{T})])^2]}} \\ &= \frac{1}{\sqrt{n}} \frac{F}{\sigma_w} \end{aligned} \tag{30}$$

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

Multi-Bernoulli Filter Implementation Notes

This section will outline specific heuristics and design choices used for the Bernoulli filter in order to help it converge. The major heuristic edit was added to avoid multiple filters converging on a single object. The assumption was made that the probability of multiple objects existing in the same place at the same time is zero. This assumption is not a novel innovation of this paper

[14]. Based on this assumption, the likelihood function for each Bernoulli filter is set to zero in the area around every other Bernoulli filter's maximum a posteriori estimate from the previous iteration. Once one Bernoulli filter is tracking one object with good accuracy, this assumption will make it impossible for other filters to track the same object. However, this assumption can lead to some instability when multiple filters are trying to track one object, but none are tracking accurately enough to zero out the correct pixels in the likelihood calculation. This can also lead to problems when multiple tracks overlap, as this is the exact case that breaks the given assumption.

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

Multi-Bernoulli Filter on GEO level Objects

Figures 2(a), 2(b), and 2(e) show the error and Figures 2(c), 2(d), and 2(f) show the probability of existence of all Bernoulli filters which track a particular object. Most objects initially have multiple filters which track them at least somewhat accurately. Eventually, one filter tracks an object accurately enough, and all other Bernoulli filters diverge. Figures 2(a) and 2(c) show object one which has the highest SNR and is in the field of view from time 0 to 340. This object is tracked with relative ease. Figures 2(e) and 2(f) show object three which has the lowest SNR and is in the field of view from time 180 to the end. It takes some time for a filter to accurately converge on this object but once it does, the probability of existence quickly rises to almost 1. Figure shows the truth tracks of each object in black with the most successful filters overlaid in gray. Specifically, the maximum a posteriori estimate of the PDF, or the point of highest probability density after updating, is plotted. Iterations of the filters before and after they successfully track an object are not included, as these steps consist of the maximum a posteriori estimate wandering through the image.

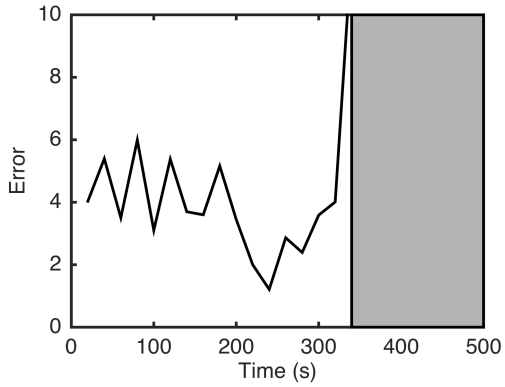

(a) Object 1 Error (pixels)

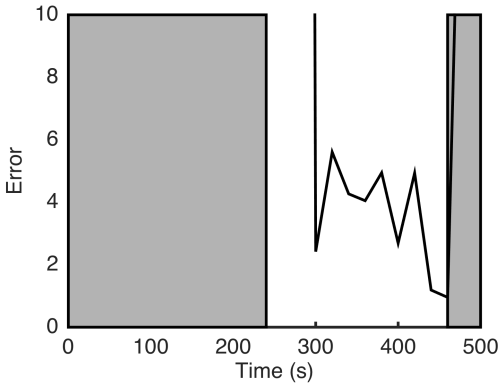

(b) Object 2 Error (pixels)

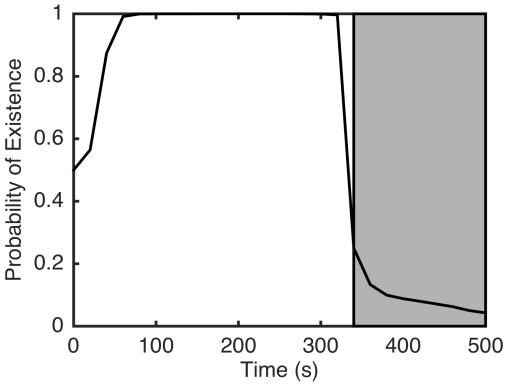

(c) Object 1 Probability of Existence

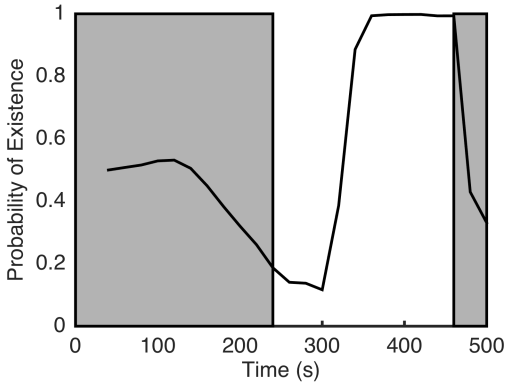

(d) Object 2 Probability of Existence

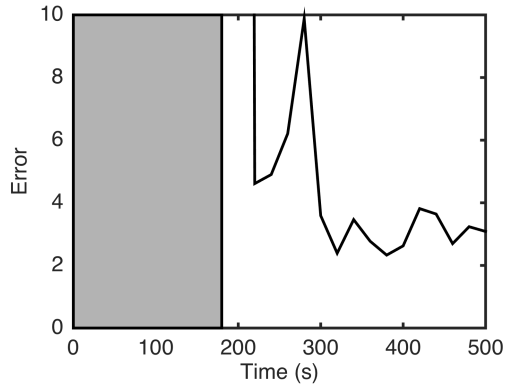

(e) Object 3 Error (pixels)

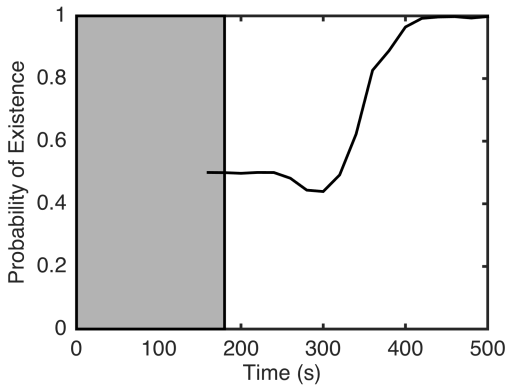

(f) Object 3 Probability of Existence

Figure 2. Results

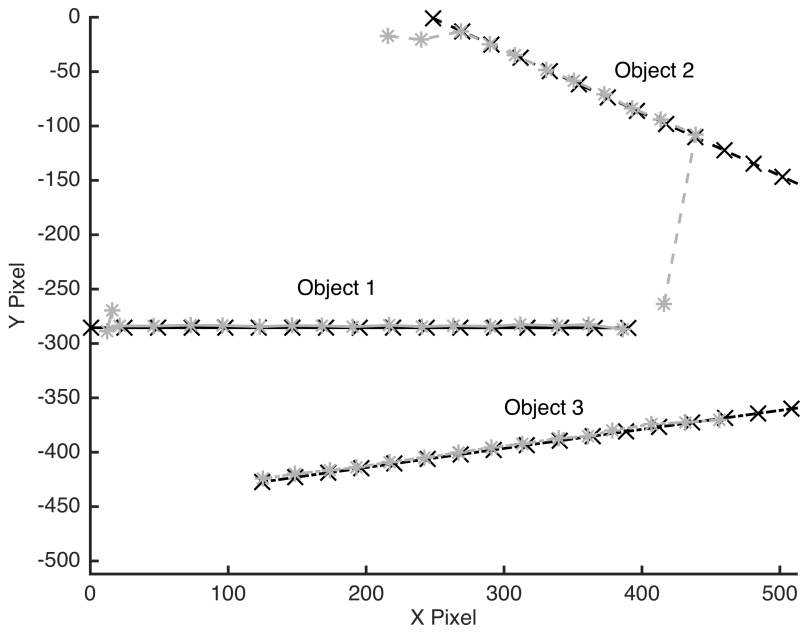

Figure 3. Object Tracks

Limiting SNR of Multi-Bernoulli Filter

The ultimate goal of this research is to detect as low SNR objects as possible. A similar test case is used with three objects over 500 s. All objects spend the entire duration in the field of view, have the same velocity, and vary only in starting position. Objects one, two, and three have total object SNR of 4.47, 3.35, and 2.24, and max per pixel SNR of 1.0, 0.7, and 0.56, respectively. Figures 4(a), 4(b), 4(c), and 4(d), show the error and probability of existence of all Bernoulli filters which track a particular object. Both object 1 and object 2 have multiple filters begin to track the object but none of the filters manage to fully converge on the object. In general, the objects are accurately tracked by the algorithm in terms of the PDF. The probability of existence update is too weak to provide sufficient evidence to the filter in all test cases. While the spacial PDF converges, the probability of existence diverges causing the filter to fail. No filter succeeds at tracking object 3. These results are promising though, as they imply the spacial likelihood is effective at tracking objects down to a total object SNR of 3.35 and per pixel SNR of 0.7. The probability of existence update with be explored further in future work.

CONCLUSION

There are two primary goals of this research arc. The first is to implement a multi-Bernoulli filter for tracking SOs which focuses on the uncued detection of as dim as possible space objects. This paper introduces the majority of the innovations the authors will use. The only problem with respect to the first goal is one of too big a search space and too few particles. Therefore future work will look at smarter implementation of a more restrictive birth model. Future work will also look into any assumptions that can be worked into the likelihood function to help with convergence.

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

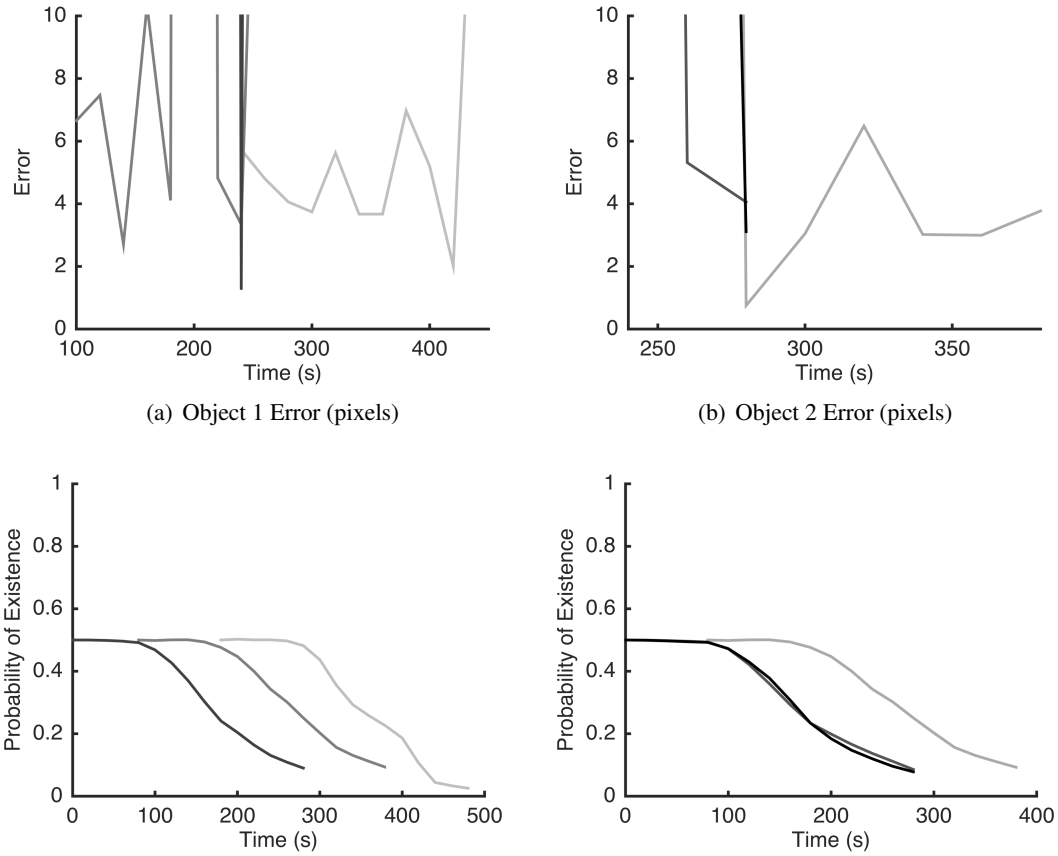

(a) Object 1 Error (pixels)

(b) Object 2 Error (pixels)

(c) Object 1 Probability of Existence

(d) Object 2 Probability of Existence

Figure 4. Limiting SNR Test Results

possible. The first method that will be investigated is marginalization of the Bernoulli filters. Other methods to be considered will be a variable number of particles, and combined Kalman filter and Bernoulli filter multi-target tracking.

ACKNOWLEDGMENT

We would like to acknowledge the National Defense Science and Engineering Graduate Fellowship for supporting this work. We would like to acknowledge and thank the Air Force Research Laboratory Space Scholars Program for supporting this project. We would also like to acknowledge Dr. Fujimoto whose discussions helped direct parts of this paper.

REFERENCES

[1] Fred Daum and Jim Huang. Curse of dimensionality and particle filters. In *Aerospace Conference, 2003. Proceedings. 2003 IEEE*, volume 4, pages 4_1979–4_1993. IEEE, 2003.

[2] 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.

[3] Kyle J DeMars, Moriba K Jah, and Paul W Schumacher. Initial orbit determination using short-arc angle and angle rate data. *Aerospace and Electronic Systems, IEEE Transactions on*, 48(3):2628–2637, 2012.

[4] Kohei Fujimoto, Masahiko Uetsuhara, and Toshifumi Yanagisawa. Statistical track-before-detect methods applied to faint optical observations of resident space objects. In *Advanced Maui Optical and Space Surveillance Technical Conference*, 2015.

[5] Reza Hoseinnezhad, Ba-Ngu Vo, and Ba-Tuong Vo. Visual tracking in background subtracted image sequences via multi-bernoulli filtering. *IEEE Transactions on Signal Processing*, 61(2):392–397, jan 2013. URL: `http://dx.doi.org/10.1109/tsp.2012.2222389`, `doi:10.1109/tsp.2012.2222389`.

[6] Reza Hoseinnezhad, Ba-Ngu Vo, Ba-Tuong Vo, and David Suter. Visual tracking of numerous targets via multi-bernoulli filtering of image data. *Pattern Recognition*, 45(10):3625–3635, oct 2012. URL: `http://dx.doi.org/10.1016/j.patcog.2012.04.004`, `doi:10.1016/j.patcog.2012.04.004`.

[7] Islam Hussein, Kyle J DeMars, Carolin Früh, Richard S Erwin, Moriba K Jah, et al. An aegis-fisst integrated detection and tracking approach to space situational awareness. In *Information Fusion (FUSION), 2012 15th International Conference on*, pages 2065–2072. IEEE, 2012.

[8] Ronald PS Mahler. *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.

[9] Douglas C Montgomery and George C Runger. *Applied statistics and probability for engineers*. John Wiley & Sons, 2010.

[10] Timothy S. Murphy, Marcus J Holzinger, and Brien Flewelling. Orbit determination for partially understood object via matched filter bank. In *AAS/AIAA Astrodynamics Specialists Meeting*, 2015.

[11] Timothy S. Murphy, Marcus J Holzinger, and Brien Flewelling. Space object detection in images using matched filter bank and bayesian update (submitted). *Journal of Guidance, Control, and Dynamics*, 2016.

[12] Branko Ristic. *Particle Filters for Random Set Models*. Springer Science + Business Media, 2013. URL: `http://dx.doi.org/10.1007/978-1-4614-6316-0`, `doi:10.1007/978-1-4614-6316-0`.

[13] Branko Ristic, Sanjeev Arulampalam, and Neil James Gordon. *Beyond the Kalman filter: Particle filters for tracking applications*. Artech house, 2004.

[14] Ba-Ngu Vo, Ba-Tuong Vo, Nam-Trung Pham, and David Suter. Joint detection and estimation of multiple objects from image observations. *IEEE Transactions on Signal Processing*, 58(10):5129–5141, oct 2010. URL: `http://dx.doi.org/10.1109/tsp.2010.2050482`, `doi:10.1109/tsp.2010.2050482`.

[15] J Chris Zingarelli, Eric Pearce, Richard Lambour, Travis Blake, Curtis JR Peterson, and Stephen Cain. Improving the space surveillance telescope's performance using multi-hypothesis testing. *The Astronomical Journal*, 147(5):111, 2014.