Statistical Track-Before-Detect Methods Applied to Faint Optical Observations of Resident Space Objects

Kohei Fujimoto
Utah State University, United States

Masahiko Uetsuhara
Independent research engineer, Japan

Toshifumi Yanagisawa
Japan Aerospace Exploration Agency, Japan

ABSTRACT

In this paper, we apply a statistically rigorous track-before-detect (TBD) method, the Bernoulli particle filter, to actual imagery of resident space objects (RSOs). Robust methods in this realm will lead to better space domain awareness (SDA) while reducing the cost of sensors and optics. We focus on estimating sensor-level kinematics of RSOs for low signal-to-noise ratio (SNR) short-arc observations. For four minute arcs of 29 images, targets with SNR as low as 2.6983 are tracked to 3 arcsec accuracy in image plane position and 0.07 arcsec/sec in velocity.

1. INTRODUCTION

Automated detection and tracking of faint objects in optical, or bearing-only, sensor imagery is a topic of immense interest in space surveillance [7, 32, 40, 48, 56, 62]. Robust methods in this realm will lead to better space domain awareness (SDA) while reducing the cost of sensors and optics [5]. They are especially relevant in the search for high area-to-mass ratio (HAMR) objects, as their apparent brightness can change significantly over time [45]. A track-before-detect (TBD) approach has been shown to be suitable for faint, low signal-to-noise ratio (SNR) images of resident space objects (RSOs). TBD does not rely upon the extraction of feature points within the image based on some thresholding criteria, but rather directly takes as input the intensity information from the image file [8, 44]. Not only is all of the available information from the image used, TBD avoids the computational intractability of the conventional feature-based line detection (i.e.,"string of pearls") approach to track detection for low SNR data.

Implementation of TBD rooted in finite set statistics (FISST) theory has been proposed recently by Vo, et al. [36,37,58] Compared to other TBD methods applied so far to SDA, such as the stacking method or multi-pass multi-period denoising, the FISST approach is statistically rigorous and has been shown to be more computationally efficient, thus paving the path toward on-line processing. In this paper, we apply a Bernoulli filter to actual CCD imagery of RSOs [42]. The Bernoulli filter can explicitly account for the birth and death of a target in some measurement arc. TBD is achieved via a sequential Monte Carlo (SMC) implementation. For four minute arcs of 29 images, targets with SNR as low as 2.6983 are tracked to 3 arcsec accuracy in image plane position and 0.07 arcsec/sec in velocity.

Although the advent of fast-cadence scientific CMOS sensors have made the automation of faint object detection a realistic goal, it is nonetheless a difficult goal, as measurements arcs in space surveillance are often both short and sparse [16, 50]. FISST methodologies have been applied to the general problem of SDA by many authors, but they generally focus on tracking scenarios with long arcs or assume that line detection is tractable [3, 4, 10, 11, 19, 20, 24–26, 28–30, 33]. We will instead focus this work on estimating sensor-level kinematics of RSOs for low SNR too-short arc observations. Once said estimate is made available, track association and simultaneous initial orbit determination may be achieved via any number of proposed solutions to the too-short arc problem, such as those incorporating the

admissible region [12, 17, 39, 43, 47, 51, 54, 61]. The benefit of combining FISST-based TBD with too-short arc association goes both ways; i.e., the former provides consistent statistics regarding bearing-only measurements, whereas the latter makes better use of the precise dynamical models nominally applicable to RSOs in orbit determination.

The outline of this paper is as follows. We first motivate the use of TBD for space surveillance with optical sensors (Section 2). Then, we discuss the current literature on FISST applications to SDA and TBD (Section 3). After a brief introduction of the Bernoulli particle filter (Section 4), the "fine print" of TBD encountered during the implementation are detailed, with a focus on improving filter performance through tunable parameters and the measurement likelihood function (Section 5). Finally, filter results are analyzed for three tracks of varying SNR (Section 6).

2. SENSOR-LEVEL TRACKING IN SDA

Traditionally, the detection of astronomical objects in telescope imagery has been achieved through a detect-before-track (DBT) approach, in which a group of pixels above a user-defined intensity threshold are assigned to a corresponding hypothesized object [46]. In more modern algorithms, one may additionally require that the pixels follow some point spread function (PSF) [64] or assign hypothesized objects based on the signal's local energy instead of a uniform threshold [32]. If some subset of these hypothesized objects follow an admissible kinematic motion – nominally, linear constant velocity motion for RSOs – then the objects are said to form a track or tracklet. Although this "string of pearls" approach has been very successful thus far and certainly has its benefits, such as drastically compressing the information from each image to a small set of coordinates, it also could become a computational bottleneck in space surveillance should either the total number of images or hypothesized objects in a track drastically increase. Quantitatively, the run-time complexity of DBT is order $O(n^m)$ if n hypothesized objects are assigned in each image for a track of m images. An increase in both m and n are realistic concerns in SDA as the sensitivity and, especially with the advent of scientific CMOS sensors, the temporal resolution of image sensors improve rapidly in the coming years. Furthermore, there is an inherent loss of information by treating detection as a boolean decision, rather than a probabilistic one, as well as by separating it from subsequent SDA tasks such as tracking, characterization, and the delivery of actionable information.

An alternative approach is track-before-detect (TBD), in which intensity information from entire images are processed instead of solely their characteristic points. There are three major advantages to TBD. First, the run-time complexity is order $O(m)$; i.e., it is no longer dependent on the number of objects in the image and only scales linearly with respect to the number of images in the track. It is possible to reap the benefits of this linear growth even with small values of m. Most notably, the dimensionality of the kinematics is four (two position states and two velocity states in the image plane) instead of the six in the three-dimensional orbit problem, making highly robust implementations of TBD, like SMC, a viable option. In fact, although only proposed as future work in this paper, we envision real-time processing of fast-cadence imagery of RSOs by employing general-purpose computing on graphics processing units (GPGPU), to which TBD is highly amenable.

Second, the state uncertainty of RSOs is rigorously characterized at the time of measurement. This information may be carried on throughout the SDA workflow via numerous techniques studied in recent years [9, 18, 21, 22, 27, 31, 38, 57], as well as inform the development of better physical models for measurement error. Such consistent quantification of RSO state uncertainty starting at the sensor-level is integral to "improv(ing) space safety." [15]

Finally, as TBD methods infer the existence of objects based on multiple images rather than assign hypothesized objects per image, they have been shown to excel when SNR is low [8, 44, 58]. A study of the Pareto frontier for a Raven-class telescope suggested that "the SNR threshold required by the detection algorithm largely influences the overall detection capability of the system." [5] Improved detection at low SNR will also enable minimalist passive sensors to make substantial contributions to SDA [48]. Therefore, TBD will greatly enhance the utility of both existing and future SDA optical sensors without significant, if any, hardware investment.

Initial work in TBD for optical observations of RSOs emerged from the stacking method proposed by Yanagisawa, et al., where low SNR RSOs are surfaced through a brute-force search for its motion in the image plane [63]. Images in a track are successively displaced by some constant amount. If an RSO in the image plane moves at the same rate as the displacements, then the median intensity value in its vicinity should be higher than the background. In order to reduce the number of displacement hypotheses that must be tested per track, the method has been implemented

on an FPGA [62] as well as combined with *a priori* information from RSO population models [55]. More recently, a Bayesian estimation approach was taken in [56], where the measurement likelihood function of RSOs was first determined by a genetic algorithm. RSOs were subsequently tracked with a particle filter. Independently, the banked matched filter (MF) was studied by Murphy, et al. [40]. Here, beginning with partial *a priori* state information, such as a short-arc measurement, one may generate a bank of templates regarding what one can expect to image at a later time. When a convolution with the actual observations is computed, templates that resemble the true state result in an increase in SNR; a subsequent hypothesis test gives the probability of detection. As these templates are inherently assigned to only a localized area in the sensor frame, the banked MF is computationally efficient. Furthermore, the aforementioned probability of detection can directly be inserted as a PDF into a Bayesian estimator. Another approach proposed by Dao, et al. is based on the concept of random sampling and consensus (RANSAC) [7]. Detection is achieved by searching for a statistically significant series of characteristic points, chosen according to user-defined criteria on the PSF and intensity, which align linearly in a data cube.

The goal in this work is to combine the stacking method's *a priori* state free generality with the statistical rigor of filtering approaches. We apply a new family of estimation algorithms called random finite set (RFS) filters. In the next several sections, we outline their benefits and how they have been applied to SDA thus far. We then focus on one particular flavor: the Bernoulli filter.

3. FISST AS APPLIED TO SDA AND TBD

One important assumption in Bayesian estimation, which has been the workhorse of orbit determination since the beginning of spaceflight [53], is that the measurements being processed are of and only of the object being tracked. Therefore, data association is usually treated by a separate heuristic preprocessor. In many scenarios pertinent to space surveillance, however, this assumption is no longer valid; not all RSOs are bright enough to produce a consistent signal for the entirety of a track, for instance. In addition, data association by itself can become computationally intractable unless association hypotheses are properly gated [1, 52].

RFS filters, which operate on unordered finite sets, are one proposed solution to these problems. RFSs are fully specified by the distribution of the number of elements in the set (cardinality) and the joint distribution of the elements conditioned upon the cardinality. As such, RFS filters can explicitly account for multiple tracking targets or, more importantly for TBD, the complete lack thereof. Furthermore, one may retain notions such as probability density functions (PDFs), their statistical moments, and Bayes rule for RFSs via the FISST framework; e.g., the FISST PDF of an RFS $\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ comprised of n elements is [42]

$$f(\mathbf{X}) = n! \cdot \rho(n) \cdot p_n(\mathbf{x}_1, \ldots, \mathbf{x}_n), \tag{1}$$

where ρ is the cardinality distribution and p_n is the joint distribution of the elements. We refer readers to a short list of existing literature for further details [34–36, 42].

Perhaps the most direct application of FISST-based methods to SDA is to assign an RSO's state to each set element. As the exact multi-Bayesian RFS filter is often intractable, many approximations exist and have been applied to SDA [11, 24], including but not limited to the GM-PHD [10], GM-CPHD [19, 29], ISP [4], and GLMB filters [30]. If we look to survey papers, comparisons between different RFS filters and competing multi-target multi-sensor tracking methodologies are presented in [3, 20, 33], and challenges specific to the SDA problem are outlined in [28]. Finally, solving the joint sensor management and allocation problem using FISST is discussed in [26, 65, 66], and the joint classification problem in [25].

The Bernoulli filter is a RFS filter for which the cardinality distribution is given as a Bernoulli distribution; that is, we specify and subsequently estimate the probability q that a target exists [42]. It thus allows for an explicit distinction in the measurement likelihood function based on whether or not an RSO is in the field of view. Furthermore, it may be extended to a multi-target problem by defining a multi-Bernoulli or labeled multi-Bernoulli RFS [35, 59, 60]. The family of Bernoulli filters has successfully been applied to TBD [23, 41, 58]; this proven track record is another motivation for us to apply the Bernoulli filter to actual CCD imagery of RSOs.

Finally, we note here that, in this paper, we choose to focus on sensor-level tracking as we are reminded of the lack of astrometric information one gains from short-arc optical observations of RSOs [16]. Even as sensor technology

advances, survey designs that emphasize coverage will continue to produce tracks that cover only a minute fraction of an object's orbit. On the other hand, it has been shown that additional information such as the estimated rate-of-change of the sensor bearing [12, 17, 54, 61], integral invariants of the orbital motion [39], or special solutions to the Lambert problem [43, 47, 51] allows one to greatly constrain the measurement likelihood function. As such, instead of a "top-down" approach where the immediate purpose of RFS filtering is to holistically track and characterize the RSO population, we take a "bottom-up" approach where we initially aim to distribute measurement information to a limited yet sufficient set of states. As time passes and more measurement or dynamical information becomes available, the problem can be constrained in higher dimensions so as to add more states and parameters to our knowledge of RSOs. We present RFS filters, then, as a means to effectively conduct this initial distribution of information – effective both in terms of computational efficiency and low SNR performance – as well as gain probabilistic metrics on said effectiveness.

4. BERNOULLI PARTICLE FILTERS

We introduce the prediction and update equations for the Bernoulli filter; a more thorough treatment is given in [42]. Assumptions made in the derivation are listed below

- The measurements are assumed to be statistically independent of the state RFS
- The state RFS dynamics is a Markov process with a known transitional density
- The initial state PDF, initial target existence probability, and target birth / survival probabilities are known

Given that measurements taken at time t_0 through t_{k-1} have been processed to obtain state $\mathbf{x}'$, the prediction and update steps to obtain the probability of existence of a target object q and a PDF $s(\mathbf{x})$ for the state $\mathbf{x}$ is

Prediction

$$q_{k|k-1} = p_b(1 - q_{k-1|k-1}) + (1 - p_s)q_{k-1|k-1} \tag{2}$$

$$s_{k|k-1}(\mathbf{x}) = \frac{p_b(1 - q_{k-1|k-1})b_{k|k-1}(\mathbf{x})}{q_{k|k-1}} + \frac{p_s q_{k-1|k-1} \int \pi_{k|k-1}(\mathbf{x}|\mathbf{x}')s_{k-1|k-1}(\mathbf{x}')\mathrm{d}\mathbf{x}'}{q_{k|k-1}}, \tag{3}$$

where p_b is the probability of a target being born during t_{k-1} to t_k, p_s is the probability of a target surviving during the same timeframe, $b(\mathbf{x})$ is the PDF of the target birth process, and $\pi(\mathbf{x}|\mathbf{x}')$ is the Markov process transitional density.

Update

$$l_k(\mathbf{z}_k|\mathbf{x}_k) = \prod_{i=1}^{n} \frac{g_1^{(i)}\left[z_k^{(i)}|\mathbf{x}_k\right]}{g_0^{(i)}\left[z_k^{(i)}\right]} \tag{4}$$

$$I_k = \int l_k(\mathbf{z}_k|\mathbf{x})s_{k|k-1}(\mathbf{x})\mathrm{d}\mathbf{x} \tag{5}$$

$$q_{k|k} = \frac{q_{k|k-1}I_k}{1 - q_{k|k-1} + q_{k|k-1}I_k} \tag{6}$$

$$s_{k|k}(\mathbf{x}) = \frac{\prod_{i=1}^{n} g_1^{(i)}\left[z_k^{(i)}|\mathbf{x}\right] s_{k|k-1}(\mathbf{x})}{\int \prod_{i=1}^{n} g_1^{(i)}\left[z_k^{(i)}|\mathbf{x}\right] s_{k|k-1}(\mathbf{x})\mathrm{d}\mathbf{x}}, \tag{7}$$

where the n measurement elements $z_k^{(1)}, \ldots, z_k^{(n)}$, nominally individual pixel intensities, at time t_k are expressed as a vector $\mathbf{z}_k$, $g_1^{(i)}$ is the likelihood function of element $z_k^{(i)}$ when a target exists at $\mathbf{x}_k$, and $g_0^{(i)}$ is that when a target does not exist.

We note that the integrals in the equations above generally cannot be computed analytically. The most robust quadrature is an SMC, in which PDFs are expressed as an ensemble of weighted sample points. Suppose that the *a posteriori*

state PDF at time t_{k-1} is approximated as

$$\hat{s}_{k-1|k-1}(\mathbf{x}) = \sum_{i=1}^{N} w_{k-1}^{(i)} \delta_{\mathbf{x}_{k-1}^{(i)}}(\mathbf{x}), \tag{8}$$

where $w_{k-1}^{(i)}$ is the weight and $\mathbf{x}_{k-1}^{(i)}$ the state of sample i, respectively, and $\delta_{\mathbf{x}_{k-1}^{(i)}}(\mathbf{x})$ is the Dirac delta function centered at $\mathbf{x}_{k-1}^{(i)}$. We'd like to express the predicted state PDF as

$$\hat{s}_{k|k-1}(\mathbf{x}) = \sum_{i=1}^{N} w_{k|k-1}^{(i)} \delta_{\mathbf{x}_{k|k-1}^{(i)}}(\mathbf{x}). \tag{9}$$

Of the N samples, B samples are used to represent the object birth density; i.e., the first term in eq. (3). If we index these samples as $i = N - B + 1, \ldots, N$,

$$w_{k|k-1}^{(i)} = \begin{cases} \dfrac{p_s\, q_{k-1|k-1}}{q_{k|k-1}} \dfrac{\pi_{k|k-1}\left[\mathbf{x}_{k|k-1}^{(i)} | \mathbf{x}_{k-1}^{(i)}\right]}{\rho_k\left[\mathbf{x}_{k|k-1}^{(i)} | \mathbf{x}_{k-1}^{(i)}, \mathbf{z}_k\right]} w_{k-1}^{(i)} & i = 1, \ldots, N - B \\ \dfrac{p_b(1 - q_{k-1|k-1})}{q_{k|k-1}} \dfrac{\pi_{k|k-1}\left[\mathbf{x}_{k|k-1}^{(i)} | \mathbf{x}_{k-1}^{(i)}\right]}{\rho_k\left[\mathbf{x}_{k|k-1}^{(i)} | \mathbf{x}_{k-1}^{(i)}, \mathbf{z}_k\right]} \dfrac{1}{B}, & i = N - B + 1, \ldots, N \end{cases} \tag{10}$$

where $\rho_k\left[\mathbf{x}_{k|k-1}^{(i)} | \mathbf{x}_{k-1}^{(i)}, \mathbf{z}_k\right]$ is the importance sampling proposal distribution of $\mathbf{x}_{k|k-1}$, assuming that the birth density samples undergo the same dynamics as the rest. Similarly, for the update step,

$$I_k = \sum_{i=1}^{N} l_k\left[\mathbf{z}_k | \mathbf{x}_{k|k-1}^{(i)}\right] w_{k|k-1}^{(i)}, \tag{11}$$

and $q_{k|k}$ may be computed via eq. (6). The integral in the state PDF update is, in essence, a normalization factor

$$\hat{s}_{k|k}(\mathbf{x}) = \sum_{i=1}^{N} w_{k|k}^{(i)} \delta_{\mathbf{x}_k^{(i)}}(\mathbf{x}) \tag{12}$$

$$w_{k|k}^{(i)} = \frac{\left\{\prod_{j=1}^{n} g_1^{(j)}\left[z_k^{(j)} | \mathbf{x}\right]\right\} w_{k|k-1}^{(i)}}{\sum_{i=1}^{N}\left\{\prod_{j=1}^{n} g_1^{(j)}\left[z_k^{(j)} | \mathbf{x}\right]\right\} w_{k|k-1}^{(i)}}, \tag{13}$$

where $N - B$ samples of $\mathbf{x}_k^{(i)}$ are chosen from $\mathbf{x}_{k|k-1}^{(1)}, \ldots \mathbf{x}_{k|k-1}^{(N)}$ with a probability corresponding to their individual weights, and the remaining B samples are drawn from the birth density. Standard methods in SMC to avoid degeneracy are applicable to the SMC implementation of the Bernoulli filter as well, such as regularization. Again, we refer readers to a short list of existing literature on SMC for further details [2, 13, 14].

5. PROCESSING TELESCOPE DATA USING BERNOULLI FILTERS

To test the performance of the Bernoulli filter in processing RSOs, we process three optical tracks of varying SNR taken during the night of October 21, 2011 at the National Central University Lulin Observatory in Taiwan [56]. We shall refer to them as Tracks 1, 2, and 3. Images in the tracks have been preprocessed to remove limb-darkening, bias, and bright stars, then cropped such that they contain exactly one RSO target. Finally, the intensity values are normalized such that the median of the maximum intensity of each image corresponds to an intensity of 1. The visual magnitudes of the targets are 16.70 in Track 1, 17.68 in Track 2, and 18.22 in Track 3. In terms of SNR as defined through the coefficient of variation, they are 9.0095, 4.4797, and 2.6983, respectively. Relevant specifications and parameters of the measurement data are listed in Table 1, and representative images are in Figures 1 and 2. Additional information regarding the data set can be found in [56].

The target state we'd like to estimate are the kinematic position and velocity $[x, y, \dot{x}, \dot{y}]$ in the image plane. We assume that RSOs appear to move in a linear constant velocity motion with a constant acceleration, unbiased process noise of

Table 1: Relevant instrument specifications and data parameters. (Originally in [56])

Spec	**Value**
Detector type	CCD (white light)
(# of Images)/track	29
Size	301 px × 301 px
Resolution	3 arcsec/px
FOV	0.25 deg
Exposure time	5.9 sec
Readout time	2.66 sec
Mount tracking	Off (Earth-fixed)

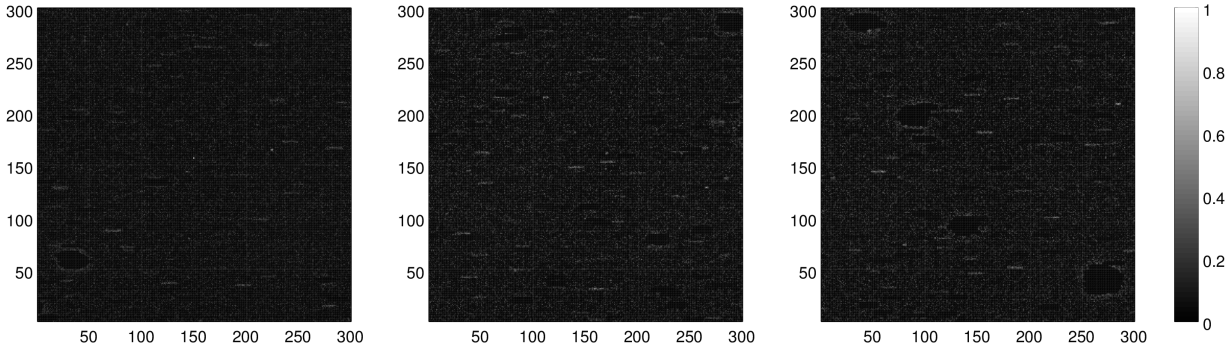

Figure 1: The first image frame from each track. From left to right, Track 1, 2, and 3.

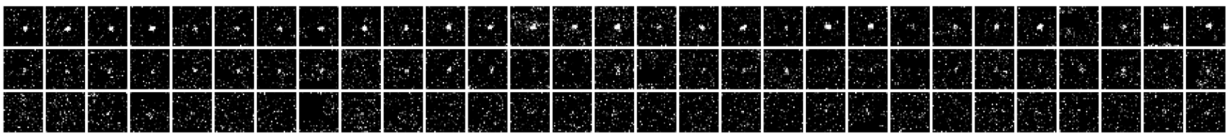

Figure 2: RSOs cropped out from Tracks 1 (top row), 2 (middle row), and 3 (bottom row). Time elapses from left to right. (Originally in [56])

0.1 px/s^2 standard deviation. We assume that there is no *a priori* information regarding the initial state of the target nor its existence probability; thus, we set the initial state PDF to be uniform with a support of $[0,300]$ px in position and $[-2,2]$ px/sec in velocity. We similarly assume no *a priori* information for the target birth process. The target birth probability is set to 0.05, and the survival probability to 1.

We base our measurement likelihood function on the one proposed in [58]

$$g_1^{(i)}\left[z_k^{(i)}|\mathbf{x}_k\right] = \mathcal{N}\left[z_k^{(i)}; h_k^{(i)}(\mathbf{x}_k), \sigma_0^2\right] \tag{14}$$

$$g_0^{(i)}\left[z_k^{(i)}\right] = \mathcal{N}\left[z_k^{(i)}; \mu_0, \sigma_0^2\right], \tag{15}$$

where $\mathcal{N}(x;\mu,\sigma^2)$ is a normal distribution over state x with mean μ and variance σ^2, μ_0 is the noise bias, and σ_0 is the noise standard deviation. $h_k^{(i)}$ is the point spread function

$$h_k^{(i)}(\mathbf{x}_k) = \frac{I_k}{2\pi\sigma_h^2}\exp\left[-\frac{(i_x - x)^2 + (i_y - y)^2}{2\sigma_h^2}\right], \tag{16}$$

where I_k is the source intensity, σ_h is the blurring factor, and (i_x, i_y) are the image plane coordinates of pixel (i). I_k determines the sensitivity of the conversion from pixel intensity values in the vicinity of a particle to the particle weight. The likelihood ratio is

$$l_k(\mathbf{z}_k|\mathbf{x}) = \exp\left[\frac{\left\{\mu_0 - h_k^{(i)}(\mathbf{x}_k)\right\}\left\{\mu_0 - 2\mathbf{z}_k + h_k^{(i)}(\mathbf{x}_k)\right\}}{2\sigma_0^2}\right]. \tag{17}$$

Although functions better suited for raw CCD intensity output certainly exist – e.g., those with a strictly positive support – the one proposed is fast to compute. Further cautionary notes regarding the likelihood function are listed here.

Parameter selection Nominally, one would choose parameters σ_0, μ_0, σ_h, and I_k based upon statistics of the measurement data. Such values, however, make eq. (4) much larger than can be handled even with extended precision arithmetic. Thus, in this paper, we estimate only μ_0, which are enumerated in Table 2, by taking the mean intensity over each image. The rest are treated as tuning parameters: $\sigma_h = 1$ for all tracks, whereas σ_0 and I_k are set as in Table 2. To elevate the Bernoulli filter to space surveillance operations, methods to automatically choose these parameters should be studied.

Clutter It is apparent from Figures 1 and 2 that not all of the signal in the tracks originate from RSOs. The likelihood function, then, must account for spurious inputs such as bright stars that are only partially removed, dim stars that are not removed, and false detections due to cosmic rays. As the target RSOs and exposure times chosen in the current dataset are such that RSOs do not appear as streaks, we amend eq. (14) such that

$$g_1^{(i)}\left[z_k^{(i)}|\mathbf{x}_k\right] = \begin{cases} \mathcal{N}\left[z_k^{(i)}; h_k^{(i)}(\mathbf{x}_k), \sigma_0^2\right] & (p_t < p_{t,\min} \text{ and } p_k < p_{k,\min}) \\ g_0^{(i)}\left[z_k^{(i)}\right]. & (\text{otherwise}) \end{cases} \tag{18}$$

For each particle, a Student's t test for normality is conducted on the intensity values for the fourth to ninth closest pixels for all horizontal scan lines in the image. If the p-value is less than some user-defined minimum $p_{t,\min}$, then the signal is deemed to be so spread out that it is due to a streaking star rather than an RSO. Similarly, a normality hypothesis test based on kurtosis is conducted on the intensity values for the nine closest pixels for all horizontal scan lines in the image [6]. If the p-value is less than some user-defined minimum $p_{k,\min}$, then the signal is deemed to be so leptokurtic, i.e., "peaked," that it is due to a cosmic ray rather than an RSO. $p_{k,\min}$ is set to 5×10^{-5} for all tracks, whereas $p_{t,\min}$ is set as in Table 2. The higher $p_{t,\min}$ is set, the more aggressively the filter will ignore signal that are spread out.

Much more work is required to devise a statistically rigorous yet computationally tractable likelihood function suitable for RSOs in all orbital regimes, and thus will be the focus of future development. We are encouraged by the results in [40], which suggest that even a small amount of *a priori* information can strongly constrain both the search volume and shape of RSO imagery.

Table 2: Filter parameters chosen for each track.

	Track 1	**Track 2**	**Track 3**
μ_0 (Estimated)	0.079066	0.10309	0.10722
I_k	0.375	0.50309	0.60722
$p_{t,\min}$	0.01	0.1	0.2
N	10,000	20,000	20,000

The particular flavor of SMC implemented for this paper is a post-regularized particle filter using the optimal bandwidth for a Gaussian kernel [13]. All computation was conducted in GNU Octave on a quad-core Intel Core i5 (Haswell, 3.40 GHz) workstation with 16 GB of RAM. The number of particles in the SMC was largely constrained by RAM; Octave became unstable if more than 20,000 particles were processed. The number of particles N used to process each track is given in Table 2. Computation time was approximately 2 frames/minute for every 10^4 particles in the SMC.

6. RESULTS

Figures 3 – 6 are results from successful Bernoulli filter runs. For each track, Figures 3 – 5 show the evolution of the SMC particle cloud along with the estimated probability of target existence, and Figure 6 contain the time histories of the estimated mean accuracy and covariance. Note that for Track 3, the filter did not converge reliably with the number of particles that fit in memory. A "hot start" was thus implemented; the support of the initial state PDF in position space was reduced to a 100 px $\times$ 100 px region about the true position.

All runs shown converge to within 1 px of the true position and 0.2 px/frame in velocity in both the x and y directions, which correspond to 3 arcsec and 0.07 arcsec/sec error in angle / angle-rate. The estimated covariances are consistent with the actual estimation errors. As noted previously, these statistics may directly be handed down the SDA workflow, be it initial orbit determination or conjunction assessment algorithms. Additionally, the proposed likelihood function is effective in distributing SMC particles only in the vicinity of signal from RSOs.

One "lesson learned" that emerged from these results is that even in solving for the projected motion of the RSO in the image plane, the number of particles required for the filter to converge became rather large for low SNR observations. Since we have ignored shape parameters of targets in this work, all of the instantaneous information from images are in the position space. Velocity space information arises solely from the dynamics; particles are heavily weighed over multiple images if their velocity is close to truth. Therefore, the particle cloud after the update step must have enough instantiations in the velocity space for an object to be consistently tracked. There are multiple paths we can take to alleviate this issue. For example, we are motivated to conduct future work in a compiled language with multi-core or GPGPU support to drastically increase the number of particles in the SMC while retaining fast computational turnaround [49]. Alternatively, we can improve the implementation of the Bernoulli particle filter; e.g., add labeled RFSs [23, 59] or change the regularization technique [13].

7. CONCLUSIONS

In this paper, a Bernoulli particle filter that directly processes sensor output from electro-optical systems is implemented so as to simultaneously detect and track resident space objects. Compared to current operational practices in ground-based optical space surveillance, the proposed method has the potential to reduce computational burden for high-sensitivity fast-cadence imagery, provide consistent uncertainty quantification from the time of observation through the delivery of actionable information, and lower the signal-to-noise ratio required to maintain object custody. Approaches rooted in finite set statistics, like the Bernoulli filter, are advantageous in their ability to explicitly account for multiple targets or the lack thereof. In the implementation, filter parameter selection and clutter rejection through the measurement likelihood function were identified as having the largest influence on filter performance.

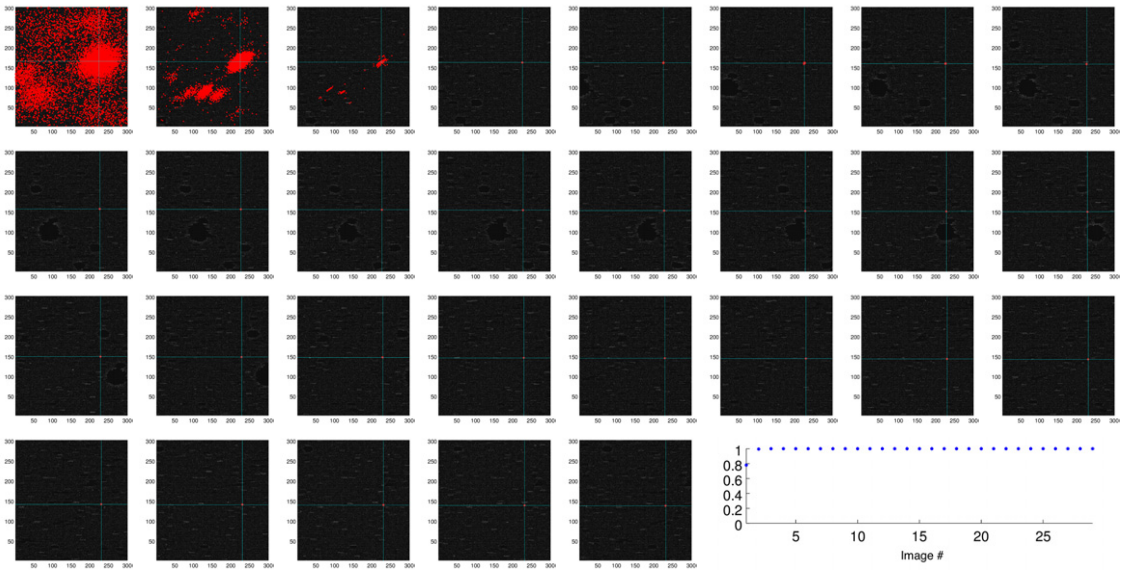

Figure 3: Evolution of particle cloud for Track 1 in red. Time flows from left to right, then top to bottom. Contour of pixel intensity in grayscale. True position state indicated by cyan lines. In the bottom right, the estimated probability of target existence as a function of time.

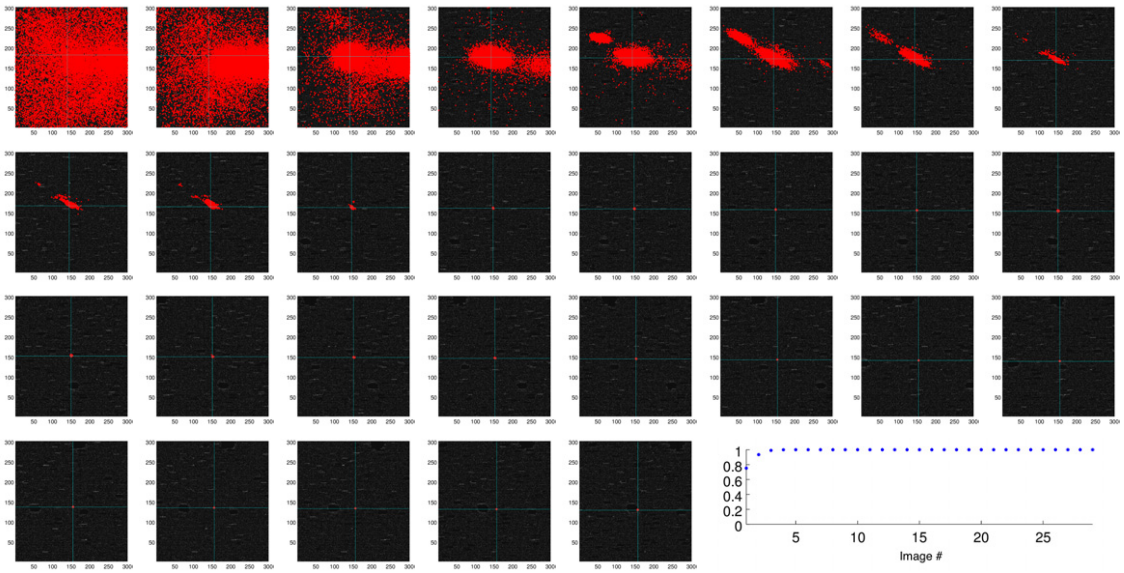

Figure 4: Evolution of particle cloud for Track 2 in red. Time flows from left to right, then top to bottom. Contour of pixel intensity in grayscale. True position state indicated by cyan lines. In the bottom right, the estimated probability of target existence as a function of time.

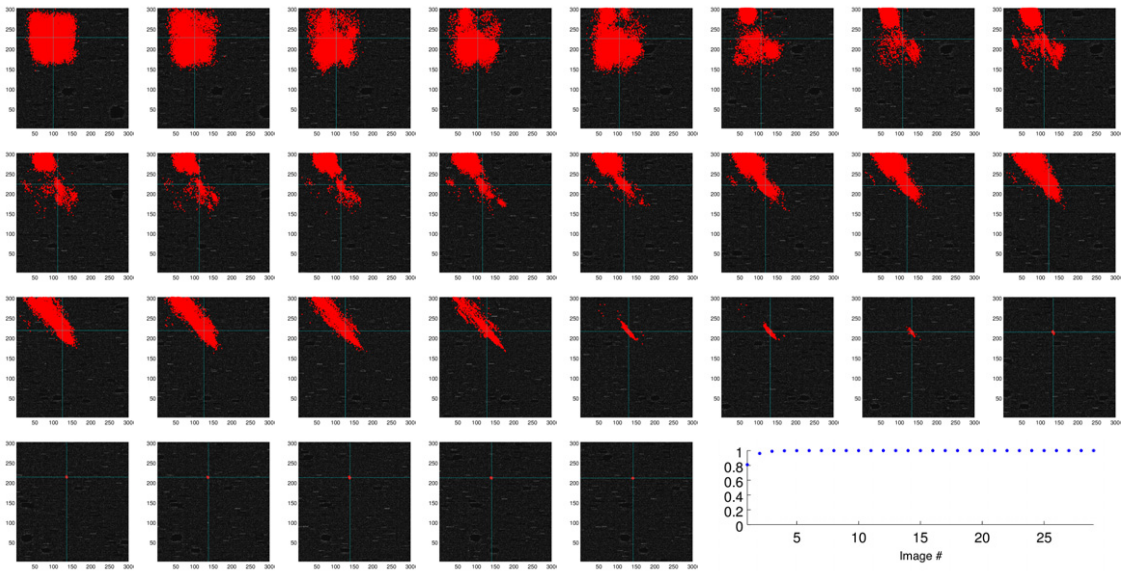

Figure 5: Evolution of particle cloud for Track 3 in red. Time flows from left to right, then top to bottom. Contour of pixel intensity in grayscale. True position state indicated by cyan lines. In the bottom right, the estimated probability of target existence as a function of time.

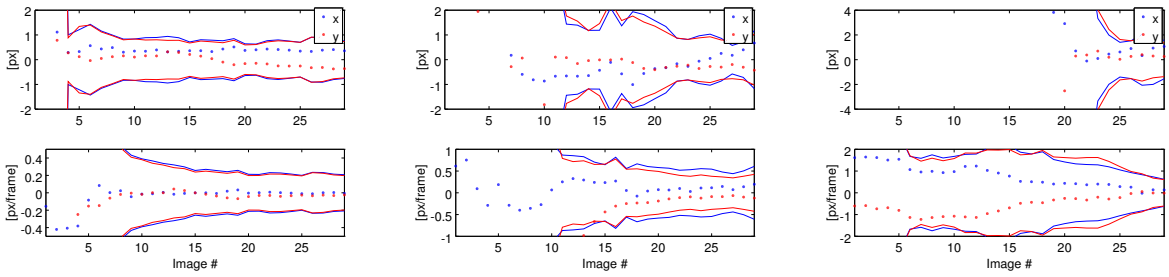

Figure 6: Dots indicate the error in the estimated mean position (top) and velocity (bottom) as a function of time. Lines are the ± 3-σ bounds as derived from the estimated covariance.

The Bernoulli filter was successful in detecting and tracking objects in actual data with signal-to-noise ratio as low as 2.6983, although a "hot start" was necessary as only 20,000 filter particles fit in memory for our particular setup. Further improvements can be made in the likelihood function, parallelizing the filter, automating filter parameter choice, and implementing a more appropriate flavor of the Bernoulli or SMC filters.

ACKNOWLEDGEMENTS

The authors would like to thank Andrew Wang for coordinating telescope operations. We also thank Marcus Holzinger, Tim Murphy, and Brien Flewelling for sharing their insight on the matched filter. Finally, we thank input and help from our colleagues in preparing this manuscript; namely, Moriba Jah and Islam Hussein.

References

[1] J. M. Aristoff, J. T. Horwood, N. Singh, A. B. Poore, C. Sheaff, and M. K. Jah. Multiple hypothesis tracking (MHT) for space surveillance: theoretical framework. 2013. Presented at the *2013 AAS/AIAA Astrodynamics Specialist Conference*. AAS 13-705.

[2] M. S. Arulampalam, S. Maskell, N. J. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, February 2002.

[3] Y. Cheng, C. Frueh, and K. J. DeMars. Comparisons of PHD filter and CPHD filter for space object tracking. In *Proceedings of the 2013 AAS/AIAA Astrodynamics Specialist Conference*, volume 150 of *Advances in the Astronautical Sciences*, 2013. AAS 13-770.

[4] D. Clark, E. Delande, and C. Frueh. Multi-object filtering for space situational awareness. Technical report, AFRL/AFOSR/EOARD, 2014.

[5] R. D. Coder and M. J. Holzinger. Sizing of a raven-class telescope using performance sensitivities. In *Proceedings of the 2013 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2013.

[6] R. B. D'Agostino, A. Belanger, and R. B. D'Agostino Jr. A suggestion for using powerful and informative tests of normality. *The American Statistician*, 44(4):316–321, November 1990.

[7] P. Dao, R. Rast, W. Schlaegel, V. Schmidt, S. Gregory, and A. Dentamaro. Track-before-detect algorithm for faint moving objects based on random sampling and consensus. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.

[8] S. J. Davey, M. G. Rutten, and B. Cheung. A comparison of detection performance for several track-before-detect algorithms. *EURASIP Journal on Advances in Signal Processing*, 2008(428036), 2008.

[9] K. J. DeMars. *Nonlinear Orbit Uncertainty Prediction and Rectification for Space Situational Awareness*. PhD thesis, Faculty of the Graduate School of The University of Texas at Austin, 2010.

[10] K. J. DeMars, I. I. Hussein, C. Frueh, M. K. Jah, and R. S. Erwin. Multiple-object space surveillance tracking using finite-set statistics. *Journal of Guidance, Control, and Dynamics*, 2015. In press.

[11] K. J. DeMars, I. I. Hussein, M. K. Jah, and R. S. Erwin. The cauchy-schwarz divergence for assessing situational information gain. In *Proceedings of the 15th International Conference on Information Fusion*, 2012.

[12] K. J. DeMars and M. K. Jah. Probabilistic initial orbit determination using gaussian mixture models. *Journal of Guidance, Control, and Dynamics*, 36(5):1324–1335, 2013. doi: 10.2514/1.59844.

[13] A. Doucet, N. de Freitas, and N. J. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, NY, 2001.

[14] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later, 2011.

[15] Committee for the Assessment of the U.S. Air Force's Astrodynamic Standards. Continuing Kepler's quest: Assessing air force space command's astrodynamics standards. Technical report, National Research Council, 2012.

[16] K. Fujimoto and K. T. Alfriend. Optical short-arc association hypothesis gating via angle-rate information. *Journal of Guidance, Control, and Dynamics*, 2015. In press.

[17] K. Fujimoto and D. J. Scheeres. Correlation of optical observations of earth-orbiting objects and initial orbit determination. *Journal of Guidance, Control, and Dynamics*, 35(1):208–221, 2012. doi: 10.2514/1.53126.

[18] K. Fujimoto, D. J. Scheeres, and K. T. Alfriend. Analytical nonlinear propagation of uncertainty in the two-body problem. *Journal of Guidance, Control, and Dynamics*, 35(2):497 – 509, 2012.

[19] S. Gehly, B. Jones, and P. Axelrad. An AEGIS-CPHD filter to maintain custody of GEO space objects with limited tracking data. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.

[20] S. Gehly, B. A. Jones, and P. Axelrad. Comparison of multitarget filtering methods as applied to space situational awareness. In *Proceedings of the 2013 AAS/AIAA Astrodynamics Specialist Conference*, 2013. AAS 13-765.

[21] D. Giza, P. Singla, and M. Jah. An approach for nonlinear uncertainty propagation: Application to orbital mechanics. 2009. Presented at the *2009 AIAA Guidance, Navigation, and Control Conference*, Chicago, IL. AIAA 2009-6082.

[22] J. Horwood, N. D. Aragon, and A. B. Poore. Gaussian sum filters for space surveillance: Theory and simulations. *Journal of Guidance, Control, and Dynamics*, 34(6):1839 – 1851, 2011.

[23] R. Hoseinnezhad, B-N Vo, B-T Vo, and D. Suter. Visual tracking of numerous targets via multi-Bernoulli filtering of image data. *Pattern Recognition*, 45:3625–3635, 2012.

[24] I. I. Hussein, K. J. DeMars, C. Frueh, R. S. Erwin, and M. K. Jah. An AEGIS-FISST integrated detection and tracking approach to space situational awareness. In *Proceedings of the 15th International Conference on Information Fusion*, 2012.

[25] I. I. Hussein, C. Frueh, R. S. Erwin, and M. K. Jah. An AEGIS-FISST algorithm for joint detection, classification, and tracking. In *Proceedings of the 2013 AAS/AIAA Astrodynamics Specialist Conference*, 2013.

[26] I. I. Hussein, Z. Sunberg, S. Chakravorty, M. K. Jah, and R. S. Erwin. Stochastic optimization for sensor allocation using AEGIS-FISST. In *Proceedings of the 2014 AAS/AIAA Spaceflight Mechanics Meeting*, 2014.

[27] B. Jia, M. Xin, and Y. Cheng. Sparse-grid quadrature nonlinear filtering. *Automatica*, 48(2):327–341, 2012.

[28] B. A. Jones, D. S. Bryant, B-T Vo, and B-N Vo. Challenges of multi-target tracking for space situational awareness. 2015. Presented at the *International Conference on Information Fusion - Fusion 2015*, Washington, D.C.

[29] B. A. Jones, S. Gehly, and P. Axelrad. Measurement-based birth model for a space object cardinalized probability hypothesis density filter. In *Proceedings of 2014 SPACE Conferences & Exposition*, 2014. AIAA 2014-4311.

[30] B. A. Jones and B-N Vo. A labeled multi-Bernoulli filter for space object tracking. 2015. Presented at the *2015 AAS/AIAA Spaceflight Mechanics Meeting*, Williamsburg, VA. AAS 15-413.

[31] B. J. Jones, A. Doostan, and G. H. Born. Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos. *Journal of Guidance, Control, and Dynamics*, 36(2):430 – 444, 2013.

[32] D. Koblick, A. Goldsmith, M. Klug, P. Mangus, B. Flewelling, M. K. Jah, J. Shanks, R. Piña, J. Stauch, J. Baldwin, J. Campbell, and T. Blake. Ground optical signal processing architecture for contributing space-based SSA sensor data. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.

[33] Q. M. Lam and J. L. Crassidis. Probability hypothesis density filter based design concept: A survey for space traffic modeling and control. In *Proceedings of Infotech@Aerospace 2012*, 2012. AIAA 2012-2566.

[34] R. Mahler. "Statistics 101" for multisensor, multitarget data fusion. *IEEE A & E Systems Magazine*, 19(1):53–64, January 2004.

[35] R. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2007.

[36] R. Mahler. "Statistics 102" for multisource-multitarget detection and tracking. *IEEE Journal on Selected Topics in Signal Processing*, 7(3):376–389, June 2013.

[37] R. Mahler. *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2014.

[38] M. Majji, R. Weisman, and K. T. Alfriend. Solution of the liouville's equation for keplerian motion: Application to uncertainty calculations. 2012. Presented at the *22nd AAS/AIAA Spaceflight Mechanics Meeting*, Charleston, SC.

[39] A. Milani, G. Tommei, D. Farnocchia, A. Rossi, T. Schildknecht, and R. Jehn. Correlation and orbit determination of space objects based on sparse optical data. *Mon. Not. R. Astron. Soc.*, 417:2094–2103, 2012. doi: 10.1111/j.1365-2966.2011.19392.x.

[40] T. S. Murphy, M. J. Holzinger, and B. Flewelling. Orbit determination for partially understood object via matched filter bank. 2015. Presented at the *2015 AAS/AIAA Astrodynamics Specialist Conference*, Vail, CO. AAS 15-726.

[41] F. Papi, B-N Vo, B-T Vo, C. Fantacci, and M. Beard. Generalized labeled multi-Bernoulli approximation of multi-object densities. *IEEE Transactions on Signal Processing*, 2015. In press.

[42] B. Ristic, B-T Vo, B-N Vo, and A. Farina. A tutorial on Bernoulli filters: Theory, implementation and applications. *IEEE Transactions on Signal Processing*, 61(13):3406–3430, 2013.

[43] C. Roscoe, P. W. Schumacher, Jr., and M. Wilkins. Parallel track initiation for optical space surveillance using range and range-rate bounds. 2013. Presented at the *2013 AAS/AIAA Astrodynamics Specialist Conference*, Hilton Head, SC. AAS 13-767.

[44] M. G. Rutten, B. Ristic, and N. J. Gordon. A comparison of particle filters for recursive track-before-detect. In *Proceedings of the 7th International Conference on Information Fusion*, 2005.

[45] T. Schildknecht, C. Früh, J. Herzog, A. Hinze, and A. Vananti. AIUB efforts to survey, track, and characterize small-size objects at high altitudes. 2010. Presented at the *Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea-Maui, HI.

[46] T. Schildknecht, M. Ploner, and U. Hugentobler. The search for debris in GEO. *Advances in Space Research*, 28(9):1291–1299, 2001.

[47] P. W. Schumacher, Jr., A. Segerman, A. Hoskins, and M. P. Whittaker. Range bounds for angle-based satellite tracking. 2013. Abstract submitted to the *2013 AAS/AIAA Space Flight Mechanics Meeting*, Honolulu, HI. AAS 13-246.

[48] B. Sease, T. S. Murphy, B. Flewelling, M. J. Holzinger, and J. Black. Enabling direct feedback between initial orbit determination and sensor data processing for detection and tracking of space objects. In *Proceedings of the 2015 SPIE Defense + Commerical Sensing*, 2015.

[49] H. Shen, C. Karlgaard, R. P. Russell, V. Vittaldev, and E. Pellegrini. Parallel sigma point and particle filters for orbit determination. 2012. Presented at the *9th US/Russian Space Surveillance Workshop*, Krestovaya Pad, Listvyanka, Russia.

[50] J. Silha, T. Schildknecht, and T. Flohrer. Improved space object orbit determination using CMOS detectors. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.

[51] J. A. Siminski, O. Montenbruck, H. Fiedler, and T. Schildknecht. Short-arc tracklet association for geostationary objects. *Advances in Space Research*, 53(8):1184–1194, 2014. doi: 10.1016/j.asr.2014.01.017.

[52] N. Singh, J. T. Horwood, J. M. Aristoff, and A. B. Poore. Multiple hypothesis tracking (MHT) for space surveillance: Results and simulation studies. 2013. Presented at the *2013 Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea-Maui, HI.

[53] B. D. Tapley, B. E. Schutz, and G. H. Born. *Statistical Orbit Determination*. Elsevier Academic Press, Burlington, MA, 2004. pp. 159-284.

[54] G. Tommei, A. Milani, and A. Rossi. Orbit determination of space debris: admissible regions. *Celestial Mechanics and Dynamical Astronomy*, 97:289–304, 2007. doi: 10.1007/s10569-007-9065-x.

[55] M. Uetsuhara, T. Hanada, H. Yamaoka, T. Yanagisawa, H. Kurosaki, and Y. Kitazawa. Strategy to search fragments from breakups in GEO. *Advances in Space Research*, 49:1151–1159, 2012.

[56] M. Uetsuhara and N. Ikoma. Faint debris detection by particle based track-before-detect method. In *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.

[57] V. Vittaldev and R. P. Russell. Collision probability for space objects using gaussian mixture models. 2013. Presented at the *2013 AAS/AIAA Spaceflight Mechanics Meeting*, Kauai, HI. AAS 13-351.

[58] B-N Vo, B-T Vo, N-T Pham, and D. Suter. Joint detection and estimation of multiple objects from image observations. *IEEE Transactions on Signal Processing*, 58(10), October 2010.

[59] B-T Vo and B-N Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13):3460–3475, July 2013.

[60] B-T Vo, B-N Vo, and A. Cantoni. The cardinality balanced multi-target multi-Bernoulli filter and its implementations. *IEEE Transactions on Signal Processing*, 57(2):409–423, February 2009.

[61] J. L. III Worthy and M. J. Holzinger. Incorporating uncertainty in admissible regions for uncorrelated detections. *Journal of Guidance, Control, and Dynamics*, 2015. In press.

[62] T. Yanagisawa, H. Kurosaki, H. Banno, Y. Kitazawa, M. Uetsuhara, and T. Hanada. Comparison between four detection algorithms for geo objects. In *Proceedings of the 2012 Advanced Maui Optical and Space Surveillance Technologies Conference*, 2012.

[63] T. Yanagisawa, H. Kurosaki, and A. Nakajima. "The Stacking Method": the technique to detect small size of GEO debris and asteroids. Technical report, Japan Aerospace Exploration Agency, February 2008.

[64] T. Yanagisawa, H. Kurosaki, H. Oda, and M. Tagawa. Ground-based optical observation system for LEO objects. *Advances in Space Research*, 56:414–420, 2015.

[65] A. Zatezalo, A. El-Fallah, R. Mahler, R. K. Mehra, and B. James. Dispersed and disparate sensor management for tracking low earth orbit satellites. In Ivan Kadar, editor, *Proceedings of Signal Processing, Sensor Fusion, and Target Recognition XVIII*, volume 7336. SPIE, 2009.

[66] A. Zatezalo, A. El-Fallah, R. Mahler, R. K. Mehra, and K. Pham. Joint search and sensor management for geosynchronous satellites. In Ivan Kadar, editor, *Proceedings of Signal Processing, Sensor Fusion, and Target Recognition XVII*, volume 6968. SPIE, 2008.