Uncued Low SNR Detection with Likelihood from Image Multi Bernoulli Filter

Timothy S. Murphy, Marcus J. Holzinger

The Guggenheim School of Aerospace Engineering, Georgia Institute of Technology

Both SSA and SDA necessitate uncued, partially informed detection and orbit determination efforts for small space objects which often produce only low strength electro-optical signatures. General frame to frame detection and tracking of objects includes methods such as moving target indicator, multiple hypothesis testing, direct track-before-detect methods, and random finite set based multi-object tracking. This paper will apply the multi-Bernoilli filter to low signal-to-noise ratio (SNR), uncued detection of space objects for space domain awareness applications. The primary novel innovation in this paper is a detailed analysis of the existing state-of-the-art likelihood functions and a likelihood function, based on a binary hypothesis, previously proposed by the authors. The algorithm is tested on electro-optical imagery obtained from a variety of sensors at Georgia Tech, including the GT-SORT 0.5m Raven-class telescope, and a twenty degree field of view high frame rate CMOS sensor. In particular, a data set of an extended pass of the Hitomi Astro-H satellite approximately 3 days after loss of communication and potential break up is examined.

I. Introduction

Part of space domain awareness (SDA) is the task of maintaining sufficient knowledge of the space environment to inform decisions as they relate to space assets, national security, and commercial ventures. Space situational awareness (SSA) is the task of providing all relevant knowledge to a particular mission at a particular time [1]. This work looks to improve methods for uncued identification of objects, or identification cued with only partial or poor orbital knowledge. When considering object discovery, rate tracking is impossible and blind search of images is inefficient. Work has been done on partially cued detection but more generalized methods are needed [2], [3]. Detection algorithms for SDA should incorporate, but not require, prior knowledge in a general way. At the same time, these algorithms should allow uncued detections to be made. Finally, algorithms should incorporate all information from measured image to ensure no information is wasted. The required method should search a state space in a smart way to find low Signal-to-Noise Ratio (SNR) signals. It should be noted that this problem also has a broader impact, including detection of near earth asteroids, computer vision, and information fusion.

There are a multitude of methods for detecting and tracking low SNR space objects. Examples include multi-object filtering [4], multiple hypothesis testing [5], multi-frame matched filters [6], Track-Before-Detect, [7], Shift and add methods [8], and more. This paper presents a new method for efficiently searching through images in a sequential manner. The methods presented here are meant to compliment existing methods.

Due to the non-linear equations of motion and perturbations, general methods derived from Bayesian filtering [9] have become popular for space object tracking and orbit determination [3],[10]. Space is represented as a 6 dimension or higher state space, making such particle based filters can become computationally difficult to implement [11]. This paper instead looks at the more simple case of kinematic frame-to-frame tracking in images obtained from Electro-Optical (EO) sensors. This allows a reduced dimensionality in the state space by tracking only position and velocity in the image plane. Much work has been done on finite set statistics (FISST) based filters in the past for a variety of applications [4]. In particular, recent pushes in FISST filter theory in the field of computer vision have looked multi-target tracking in images [12], [13], [14]. This previous work operates directly on pixel data and requires no dedicated detection algorithm. Instead, per pixel information from an image is used as evidence for the existence or lack of existence of an object. Very recently, the SDA community has begun to look at FISST filters [15] and detectionless frame-to-frame tracking [14].

The main analytic problem this paper analyzes is the likelihood function used for these filters. The current state-of-the-art likelihood function is based on modeling the statistics of an image [16]. This likelihood function is not unique to this problem, but has been adapted in recent years for detectionless tracking, or tracking that does not require an explicit detection algorithm. [17]. This likelihood can be applied to moving point sources as well as complex shapes such as tracking of people [12]. A different method, based on a hypothesis test on the matched filter predicted by a particular particle to evaluate measurement likelihood, is used in this paper [18] [2]. This newer likelihood provides a series of improvements, including eliminating the requirement to predict object brightness, and a built-in way to modulate probability of false alarm. This paper shows a fundamental similarity between the two likelihoods; the only

substantial difference in the use of a hypothesis test. A simulation based analysis gives evidence for why one algorithm may be better than another in a certain situation.

This paper also applies the algorithm to data taken with the optical resources at Georgia Tech. First, the algorithm is tested on tracking geostationary orbits as they move through a sidereal stare campaign. The results illustrate the algorithms ability to blindly track space objects. The algorithm is then tested on data from a 10 degree field of view sensor mounted on a tripod. The Hitomi Astro-H Satellite experienced a spin malfunction and break up on March 26th, 2016. The data set tested observes Hitomi on March 29th, 2016, at which point the satellite has a very obvious spin, and rapidly changing brightness. The algorithm tracks through the varying light curve with no problems.

To reiterate, this paper makes the following contributions. First, state-of-the-art multi-object tracking methods for space object discovery in a series on images are presented. This paper proves that the current state-of-the-art likelihood function is inherently related to a matched filter and therefore has SNR optimality. The new likelihood function is analyzed for different noise models that can apply to electro-optical images. The previous state of the art likelihood and the new likelihood are compared to show which filter is better suited for what situations. This includes pure object discovery, search in a specific orbital regime, and search primed by previous orbital knowledge. Finally, the validation of proposed methods is presented on empirical data taken with the Georgia Tech Space Object Research Telescope.

This paper is presented as follows. Section II reviews the basics of multi-Bernoulli filtering. Section III presents a comprehensive discussion of the theoretical and quantitative aspects of the two likelihood functions. Section IV discusses specific implementation decisions made to generate the results for this paper. Section V presents the results of operating the algorithm on empirical data taken with a 0.5m raven class telescope and wide field of view sensor.

II. Multi-Bernoulli Filtering for Frame-to-Frame Tracking

A. Dynamics

Consider the standard filtering problem with discrete dynamics and measurement models

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

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

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{z} \in \mathbb{R}^m$, $\mathbf{w} \sim f_w : \mathbb{R}^n \to \mathbb{R}$, and $\mathbf{v} \sim f_v : \mathbb{R}^m \to \mathbb{R}$. $\mathbf{x}$ 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 is a series of m pixels, $\mathbf{z}_k := \{z_{k,i}\}_{i=1}^m$.

B. Multi-Bernoulli Filter

This paper uses the multi-Bernoulli filter used in [17]. Various other variations of this filter, including the Labeled Multi-Bernoulli Filter [19], can be used as well. This paper's primary focus is not on which multi-target particle filter is best, but on how to use these filters for low SNR SO detection. A brief review of the theory and application is presented for the reader's convenience. The explanation of the Bernoulli filter starts from the single target Bayesian filtering equations [9]

$$p(\mathbf{x}_k|\mathbf{z}_{k-1}) = \int p(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{x}_{k-1}|\mathbf{z}_{k-1})\mathrm{d}\mathbf{x}_{k-1} \qquad \text{(Prediction)} \tag{3}$$

$$p(\mathbf{x}_k|\mathbf{z}_k) = \frac{p(\mathbf{z}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{z}_{k-1})}{p(\mathbf{z}_k|\mathbf{z}_{k-1})} \qquad \text{(Update)} \tag{4}$$

where $\mathbf{z}_k$ is the time series of measurements up until time step k. Equation 3 is the prediction step while Equation 4 is the update step.

For the multi-object filter, $\mathcal{X}_k$ is a random finite set (RFS) used to represent the multi-object tracking problem where the number of objects is unknown. In essence, a RFS is a set of objects represented as random vectors, but in which the number of objects itself is also a random variable. For a single Bernoulli filter, the state is modeled as a Bernoulli random finite set (BRFS). The BRFS is a special set $\mathcal{S}_k$ containing a probability of existence, r_k, and a RFS which contains 1 object with probability r_k and is empty with probability $1 - r_k$. The object position is described by the PDF, $p(\mathbf{x}_k)$. In other words, the object is described by a PDF and probability of existence shown in Equation (5).

$$\mathcal{S}_k = \{p(\mathbf{x}_k), r_k\} \tag{5}$$

The Bernoulli filter estimates the PDF with equations (3) and (4) and updates the probability of existence with update equations that can be found in [17].

A multi-Bernoulli filter is a bank of Bernoulli filters used to process the full multi-object tracking problem. The analytic prediction and update equations for a multi-target tracking scheme are

$$\pi_{k|k-1}(\mathcal{X}_k|\mathcal{Z}_{1:k-1}) = \int p_{k|k-1}(\mathcal{X}_k|\mathcal{X})\pi_{k-1}(\mathcal{X}|\mathcal{Z}_{1:k-1})\delta\mathcal{X} \qquad \text{(Prediction)} \qquad (6)$$

$$\pi_k(\mathcal{X}_k|\mathcal{Z}_{1:k}) = \frac{g(\mathcal{Z}_k|\mathcal{X}_k)\pi_{k|k-1}(\mathcal{X}_k|\mathcal{Z}_{1:k-1})}{\int g(\mathcal{Z}_k|\mathcal{X}_k)\pi_{k|k-1}(\mathcal{X}|\mathcal{Z}_{1:k-1})\delta\mathcal{X}} \qquad \text{(Update)} \qquad (7)$$

where $g(\cdot|\cdot)$ is the likelihood function, a subject of novel work in this paper. Keep in mind that these are FISST set integrals, not standard integrals, and a full explanation of the nuance in these equations is not in the scope of this paper. 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.

The multi-Bernoulli filter represents the multi object state as the union of a series of BRFSs

$$\mathcal{X}_k = \bigcup_{l=1}^{n_s} \mathcal{S}_k^l \qquad (8)$$

This allows the approximate implementation of a FISST filter to reduce to a series of Bernoulli filters. There are particle implementations of Bernoulli filters [20], and the union of these filters is relatively straight forward to compute. A more in depth, general analysis of FISST can be seen in [4]. The next few sections focus on applying the Multi-Bernoulli filter which operates on image frame-to-frame tracking to EO sensors tracking. The primary novel change is defining a likelihood function $g(\cdot|\cdot)$ which combines space imagery and Multi-Bernoulli filter. The likelihood function is used to update particle weights and is where a measurement function is defined. Because in images, the measurement is the whole image, special considerations should be taken to properly define this function.

III.   Analysis of Likelihood Functions suitable for Detection of Space Objects

A.   Statistics of an Image

This section defines the statistical models for an image that are used throughout the paper. In general there are a variety of ways to model the signal and noise an image, and this paper discusses results with respect to multiple models.

An image is made up of a series of measurements taken at the same time. In practice, each pixel is broken into some deterministic signal content and some zero mean noise of known distribution (assuming background subtraction has been performed).

$$z_i = s_i + w_i \qquad (9)$$

$$w_i \sim p(w), \ \mathbb{E}[w_i] = 0 \qquad (10)$$

where $p(w)$ is some PDF or Probability Mass Function (PMF) of known form, and i is an indexing denoting a particular pixel. While z_i represents the i^{th} pixel, this same quantity can be defined as a vector of n pixels, $\mathbf{z} = [z_1 \ldots z_n]^T$. In general every scalar with subscript i is assumed to belong to a vector of pixels.

The distribution of w_i in Equation (9) can be arbitrary, but often is assumed to be of a specific form. A common way to model the statistics of a pixel is with a series of Poisson distributions [21]. The pixel is composed of some signal, which is Poisson distributed and referred to as shot noise, and some dark noise and read noise both of which are Poisson distributed, and can therefore be combined into one Poisson distributed random variable, w_i. These leads to a pixel, $z_i = s_i + w_i$, made up of signal, $s_i \sim \mathcal{P}(\lambda_{s,i})$, and noise, $w_i \sim \mathcal{P}(\lambda_{w,i})$. However, it is more convenient to redefine a background subtracted z_i. In this paper, background subtraction means subtraction of known statistical outliers, such as hot pixels, and the background mean [22].

$$s_i = \lambda_{s,i} \qquad (11)$$

$$w_i \sim \{\mathcal{P}(\lambda_{w,i} + \lambda_{s,i}) - (\lambda_{w,i} + \lambda_{s,i})\} \qquad (12)$$

This new definition of z_i allows the problem to be posed as the sum of a deterministic signal and zero mean random noise. A Poisson distribution approaches a Gaussian distribution as the parameter, λ, increases. Values of λ larger

than 100 are typically sufficient. If such a simplification is applicable, the following more simplified definition of z_i can be used.

$$w_i \sim \mathcal{N}(0, \lambda_{w,i} + \lambda_{s,i}) \tag{13}$$

Further assumptions can be made to assume the dark noise and read noise dominate the shot noise, are equivariant in all pixels, and are independent of each other, giving

$$w_i \sim \mathcal{N}(0, \lambda_w) \tag{14}$$

Note that in this formulation, s_i is equivalent to the mean of the measurement z_i, which is Gaussian distributed. It is worth noting that these measurements are taken over an integration time, of which they are functions

$$\mathcal{T} = [t_0, t_0 + t_I] \tag{15}$$
$$z_i = z_i(\mathcal{T}) \tag{16}$$

While this notation is not included in this paper, timing and integration error is a further source of uncertainty in pixels.

SNR is a measure of the amount of signal present relative to the strength of the noise and a common measure of signal quality. The SNR for a particular pixel, z_i, is defined in Equation (17).

$$\text{SNR}_{pz}(z_i) = \frac{\mathbb{E}[z_i]}{\sqrt{\mathbb{E}[(z_i^2 - \mathbb{E}[z_i])^2]}} = \frac{s_i}{\sigma_{w,i}} \tag{17}$$

In astronomy, the total pixel-to-pixel structure of an object may exist in multiple pixels, and therefore the per pixel SNR does not fully capture the information available. Total photometric SNR for an object that exists in n_z pixels (assuming independent noise between pixels) is calculated as

$$\text{SNR}_{ph}\left(\sum z_i\right) = \frac{\mathbb{E}[\sum z_i]}{\sqrt{\mathbb{E}[([\sum z_i]^2 - \mathbb{E}[\sum z_i])^2]}}$$
$$= \frac{\sum s_i}{\sqrt{\sum \sigma_{w,i}^2}} \tag{18}$$
$$\approx \sqrt{n_z}\frac{\bar{s}_i}{\sigma_w} \tag{19}$$

where $\bar{s}_i$ is the average signal value over the pixels being considered, and σ_w is some approximate per pixel noise.

The two types of SNR warrant further discussion. First, neither values are any less valid for non-Gaussian distributions, and in particular work equally well for Poisson distributions. Second, both methods say something different but useful about the signal. Photometric SNR shows how well an algorithm (e.g. matched filter) should preform on detecting an object. Pixel SNR shows how easily an object can be picked out with the naked eye. Finally, choosing n_z pixels which contain signal from an object is a nebulous concept. Typically, there are pixels which may or may not contain signal, and especially in low SNR cases, ambiguity exists. This is an active are a research with no clear solutions [21].

B. Matched Filters

The matched filter (MF) is an image filter which is used to define a likelihood function in this paper. The MF is predicated on correlating a hypothesized signal structure with a series of pixels in an image. Given a state, $\mathbf{x}$, in some state space, there must exist a set and function, $T(\mathbf{x})$ and $h_i(\mathbf{x})$, which predict the location of an object in an image, and the pixel values produced by the object, respectively.

$$h_i(\mathbf{x}) = s_i' \tag{20}$$
$$T(\mathbf{x}) = \{i : s_i' > 0\} \tag{21}$$

In other words, if an image has n pixels, and there exists a hypothesized state $\mathbf{x}$, then $T(\mathbf{x})$ is a list of pixels which contain some signal given $\mathbf{x}$, and $h_i(\mathbf{x})$ is the predicted value in pixel i given $\mathbf{x}$. Also note that the cardinality of $(T(\mathbf{x}))$ is n_z, which is directly related to the n_z in Equation (19). These two quantities, h_i and T, are referred to as

the template in matched filter theory. If the object has no apparent motion in an image, $T(\mathbf{x})$ and $h_i(\mathbf{x})$ are defined by the point spread function in the image, the flux of the photons coming off the object, and the exposure time. If the object has apparent motion, the template is generated via an integration of a moving point spread function along the dynamics. See [2] for more details.

Given the most general formulation of z_i in Equation (9), the matched filter is a linear image filter which seeks to maximize the post filter SNR. Consider an arbitrary linear filter, $\mathbf{m} := \{m_i\}_{i=1}^{n_z}$

$$z_{MF} = \sum_{i \in T(\mathbf{x})} m_i z_i = \mathbf{m}^T \mathbf{z} = \mathbf{m}^T \mathbf{s} + \mathbf{m}^T \mathbf{w} \tag{22}$$

$$\text{SNR}_{MF} = \frac{\mathbf{m}^T \mathbf{s}}{\sqrt{\mathbb{E}[\mathbf{m}^T \mathbf{w}]^2}} \tag{23}$$

where $\mathbf{s}$ is the pure signal content and $\mathbf{w}$ is the pure zero mean noise content. At this point we can define the covariance matrix of $\mathbf{w}$, $\mathbf{R}_w = \mathbb{E}[\mathbf{w}\mathbf{w}^T]$. From matched filter theory the optimal choice for $\mathbf{m}$, to maximize the SNR (23) of the final result, is

$$\mathbf{m} = a\mathbf{R}_w^{-1}\mathbf{s} \tag{24}$$

where a is an arbitrary scaling parameter [23] [24].

If the underlying state $\mathbf{x}$ is known, then $\mathbf{s} = \mathbf{h}(\mathbf{x})$. For a hypothesized state $\mathbf{x}$, then $\mathbf{s} = \mathbf{h}(\mathbf{x})$ is still a reasonable method for testing the validity of that hypothesis, but produces a sub-optimal filter. If the noise vector $\mathbf{w}$ is independent and equivariant, $\mathbf{R}_w^{-1}$ can be absorbed into the arbitrary scaling parameter, a, giving $\mathbf{m} = a\mathbf{s}$. The matched filter makes no assumptions about the distribution of the noise. In other words, if the expected value of the measurements ($\mathbf{h}(\mathbf{x})$) and the covariance matrix of the noise can be estimated ($\mathbf{R}_w$), the matched filter is SNR optimal.

The final matched filter equation for some predicted state, $\mathbf{x}$, and i.i.d. noise is then

$$z_{MF} = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) z_i \tag{25}$$

C. Likelihood Ratio

The current established likelihood function for tracking point sources in images is a likelihood ratio [17], [16]. This has already been used in SO tracking [14].

Each pixel is assumed to belong to one of two distributions, depending on whether the pixel is predicted to have signal.

$$p(z_i|\mathbf{x}) = \begin{cases} \psi_i(z_i) & i \in T(\mathbf{x}) \\ \phi_i(z_i) & i \notin T(\mathbf{x}) \end{cases} \tag{26}$$

Commonly assumed distribution for ψ and ϕ is

$$\phi(z_i) = \mathcal{N}(z_i; 0, \sigma^2) \tag{27}$$
$$\psi(z_i, \mathbf{x}) = \mathcal{N}(z_i; h_i(\mathbf{x}), \sigma^2) \tag{28}$$

where $\mathcal{N}(\cdot; \mu, \sigma^2)$ is a Gaussian with mean μ and variance σ^2. Note that this method does not rely on a Gaussian assumption. Like the matched filter, this method should work for any distributions that can be properly modeled. The total likelihood for the entire filter is then

$$g(z|\mathcal{X}_k) = \left(\prod_{\mathbf{x}_k \in \mathcal{X}_k} \prod_{i \in T(\mathbf{x}_k)} \psi_i(z_i, \mathbf{x}_k) \right) \left(\prod_{i \notin \cup T(\mathbf{x}_k)} \phi_i(z_i) \right) \tag{29}$$

The total likelihood is a quantity used in multi-target filters. The per pixel likelihood is calculated as the product of relative likelihoods of a pixel belonging to either the distribution in Equation (26).

$$g_z(\mathbf{x}) = \prod_{i \in T(\mathbf{x})} \frac{\psi(z_i, \mathbf{x})}{\phi(z_i)} \tag{30}$$

The probability of existence update is formulated with the same relative likelihood. This likelihood function is general enough to work for the case of arbitrary point tracking.

The current likelihood function, in the case of Gaussian noise, is nothing more than the ratio of the measured matched filter result tested against a zero mean and assumed variance.

Lemma III.1. *Using the definition of z_i in Equation (14), the likelihood ratio defined in Equation (30) can be rewritten as a likelihood ratio of the result of a matched filter.*

Proof. Using the definition of a Gaussian PDF

$$g_z(\mathbf{x}) = \prod_{i \in T(\mathbf{x})} \frac{e^{-\frac{1}{2}(z_i - s_i)^2}}{e^{-\frac{1}{2}(z_i - 0)^2}} \tag{31}$$

$$= \prod_{i \in T(\mathbf{x})} e^{-\frac{1}{2}(z_i^2 - 2 z_i s_i + s_i^2)} e^{\frac{1}{2} z_i^2}$$

$$= e^{\sum (z_i s_i - \frac{1}{2} s_i^2)}$$

Next, define the following constant scaling parameter as the root sum square of the per pixel photometric intensity.

$$\alpha = \sqrt{\sum_{i \in T(\mathbf{x})} (s_i^2)} \tag{32}$$

Recalling the matched filter defined in Equation (22), the likelihood ratio becomes,

$$g_z(\mathbf{x}) = e^{\sum (z_i s_i - \frac{1}{2} s_i^2)}$$

$$= e^{\frac{z_{MF}}{\alpha} \alpha - \frac{1}{2} \alpha^2 + \left(\frac{1}{2} \frac{z_{MF}}{\alpha}\right)^2 - \left(\frac{1}{2} \frac{z_{MF}}{\alpha}\right)^2}$$

$$= \frac{e^{-\frac{1}{2}\left(\frac{z_{MF}}{\alpha} - \alpha\right)^2}}{e^{-\frac{1}{2}\left(\frac{z_{MF}}{\alpha}\right)^2}} \tag{33}$$

Note the identical form between Equations (31) and (33). □

Note that the quantity z_{MF}/α is a scaled matched filter result with an expected value of α. The likelihood ratio is therefore a comparison of the matched filter result on the true signal and pure noise.

The primary difficulty in using this canonical likelihood ratio is the required estimate of the photometric root mean square of the pixel intensities. In space object tracking, and especially in object discovery, the actual photon flux coming off a space object is difficult or impossible to know before a detection is made. Conversely, with a good estimate of a point spread function, the relative photometric intensities between the pixel values can be predicted; this is only dependent on the dynamics and point spread function. The relative likelihood cannot be correctly computed without the expected signal values in each pixel. Previous efforts to implement this likelihood estimate the mean of the signal [14]. However, the matched filter is invariant to an arbitrary scaling parameter; the SNR gain is optimal so long as the expected relative values are correct. Because of this, the matched filter can be used to define a likelihood function which doesn't need a photon flux estimation.

D. Hypothesis Test Likelihood

This paper proposes a new likelihood function for detectionless multi-Bernoulli filters based on a hypothesis test on the matched filter. Such a likelihood function has been used for a general particle filter [25].

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, $h_i(\mathbf{x})$, shown in Equation (25). Under an arbitrary distribution of w_i assumed in Equation (9), the distribution of z_{MF} can be calculated.

$$p_{MF} = p_{h_1(\mathbf{x})z_1} * p_{h_2(\mathbf{x})z_2} * \dots * p_{h_{n_z}(\mathbf{x})z_{n_z}} \tag{34}$$

where $*$ represents the convolution operator. This can be difficult to calculate for arbitrary complex distributions, but is possible.

Because the noise, w_i, is assumed to be mean zero, if the predicted region contains no signal, $s_i = 0 \ \forall \, i \in T(\mathbf{x})$, the resulting distribution is mean zero.

$$\mathbb{E}[z_{MF}|s_i = 0] = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x}) \, \mathbb{E}[w_i] = 0 \tag{35}$$

By a similar argument, if there is signal content in the pixels on which the matched filter operates, the result has a mean greater than zero. If the matched filter operates with a perfect template, the distribution has a known expected value.

$$\mathbb{E}[z_{MF}|s_i = h_i(\mathbf{x}_{true})] = a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x})h_i(\mathbf{x}_{true}) = a\alpha^2 \tag{36}$$

Note that $a\alpha^2$ is typically not achieved in particle filtering, due to the random nature of particle sampling. The matched filter, z_{MF}, can be assumed to exist in one of two distributions

$$\tilde{\phi}_{MF}(z) = p_{MF}(z;0) \tag{37}$$
$$\tilde{\psi}_{MF}(z;\beta) = p_{MF}(z;\beta), \ \ \beta > 0 \tag{38}$$

where the distribution $p_{MF}(z;\beta)$ has an expected value of β and is otherwise an arbitrary PDF calculated in Equation (34). These distributions in Equations (26) and (38) are mirrors of each other; Equation (26) defines the two possible distributions for a pixel given signal or no signal, while Equation (38) defines the two possible distributions for a matched filter given signal or no signal.

Recall that the Equation (25) assumes that noise is i.i.d., which is violated by shot noise which explicitly dependent on the magnitude of the signal. Because signal magnitude is unknown, shot noise must either be estimated or ignored. Typically, for low SNR signals, shot noise is dominated by other noise sources and can be ignored. If estimated, it will require calculation of the covariance matrix $\mathbf{R}_w$ in Equation (24). This is a limitation of the current method.

Next, the following null and alternate hypotheses for a binary hypothesis test are defined

$$H_0 : z \sim \tilde{\phi}(z)$$
$$H_1 : z \sim \tilde{\psi}(z;\beta); \beta = z_{MF} \tag{39}$$

In essence, this test assumes no signal content in the pixels defined by $T(\mathbf{x})$, and asks whether the measurement gives significant evidence of underlying signal (mean of matched filter greater than zero). For binary hypothesis testing, a probability of false alarm is set, p_{FA}, which in turn defines an integration threshold, z_{TH}, based on the null hypothesis cumulative density function (CDF).

$$p_{FA} = \int_{z_{TH}}^{\infty} \tilde{\phi}(z)dz \tag{40}$$

A probability of detection can be calculated with the following integral.

$$p_D = \int_{z_{TH}}^{\infty} \tilde{\psi}(z)dz \tag{41}$$

To calculate this probability of detection, the arbitrary distribution $p(z;\beta)$ in Equation (39) must be known. For more on binary hypothesis testing, see [26]. This hypothesis test determines 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_D, 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 distributions in Equation (38). This leads to the particle-wise relative likelihood

$$g_z(\mathbf{x}) = \frac{\tilde{\psi}(z_{MF})}{\tilde{\phi}(z_{MF})} \tag{42}$$

This backtracking to a relative likelihood is necessary for the probability of existence update, which does not operate correctly without a relative likelihood. The actual test distributions of z_{MF}, $\tilde{\phi}(z)$ and $\tilde{\psi}(z)$, are not trivial to calculate in general. If z_i is assumed to be Gaussian distributed as shown in Equation (14), and the noise is again assumed to be i.i.d., z_{MF} is distributed as

$$\text{Var}[z_{MF}] = \mathbb{E}\left[\left(a \sum_{i \in T(\mathbf{x})} h_i(\mathbf{x})w_i\right)^2\right] = a^2\sigma_w^2 \sum_{i \in T(\mathbf{x})} h_i^2(\mathbf{x})$$
$$z_{MF} \sim \mathcal{N}(\alpha^2, a^2\alpha^2\sigma_w^2) \tag{43}$$

where α is defined in Equation (32). Recalling Equation (24), the matched filter has an arbitrary scaling parameter, a. First, in preprocessing, an image can be scaled by $1/\sigma_w$ to reduce the variance in each pixel to 1. If the value $a = 1/\alpha$ is chosen, the variance is reduced to 1. Alternatively, a factor of $1/\sigma_w$ can be inserted into a, rather than in preprocessing. Recall that the expected value in Equation (43) can only be achieved with a perfect template, which is typically not the case. Regardless, the test distributions in this special case become

$$\tilde{\phi}_{MF}(z) = \mathcal{N}(z; 0, 1) \tag{44}$$

$$\tilde{\psi}_{MF}(z) = \mathcal{N}(z; \beta, 1), \ 0 < \beta \leq \alpha \tag{45}$$

For Poisson distributed z_i as seen in Equation (11), Equation (38) will take one of two forms. If the expected value of the Poisson distribution is small, the test distributions are two discrete PMFs. If the expected value is high enough, the test distributions become approximately Gaussian.

For a Poisson distribution, it is easy to show

$$w \sim \frac{\lambda^k e^{-\lambda}}{k!}, k = 0, 1, 2, 3 \ldots \implies$$

$$c_1 w - c_2 \sim \frac{\lambda^{(k+c_2)/c_1} e^{-\lambda}}{((k+c_2)/c_1)!}, k = -c_2, c_1 - c_2, 2c_1 - c_2, 3c_1 - c_2 \ldots \tag{46}$$

The random variable $c_1 w - c_2$ has a PMF with mass locations dependent on the scalars c_1 and c_2. The matched filter, z_{MF}, can be represented by a PMF, because it is a weighted sum of discrete random variables with known PMFs. While Poisson random variables technically have infinite permutations, for sufficiently low values of λ, sufficient mass can be captured with a computationally bounded number of permutations.

This method of permutations can be used to exactly calculate the PMF in Equation (38), and evaluate the hypothesis test. The number of permutations grows proportionately with the cardinality of $T(\mathbf{x})$, and the Poisson parameter of the noise λ. When not computationally restrictive, the permutations of the PMF should be used to evaluate the hypothesis test; otherwise, a Gaussian approximation can be used. This motivates the following Lemma.

Lemma III.2. *For a series of Poisson distributed random variables, $w_i \sim \mathcal{P}(\lambda_i)$, the random variable $z_{MF} = \sum_{i=1}^{m} c_i w_i$ is approximately Gaussian for sufficiently high values of either λ_i or m.*

Proof. First, the mean and variance of z is

$$\mathbb{E}[z_{MF}] = \sum_{i=1}^{m} c_i \lambda_i \tag{47}$$

$$\sigma^2 = \sum_{i=1}^{m} c_i^2 \lambda_i \tag{48}$$

Next we must define a new random variable, which is just the matched filter normalized by its mean and standard deviation.

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \frac{\sum_{i=1}^{m} c_i w_i - \sum_{i=1}^{m} c_i \lambda_i}{\sqrt{\sum_{i=1}^{m} c_i^2 \lambda_i}} \tag{49}$$

It is convenient to also define $w_i' = w_i / \lambda_i$ so that $\mathbb{E}[w_i'] = 1$ for the next part of this proof.

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \frac{\sum_{i=1}^{m} c_i \lambda_i w_i' - \sum_{i=1}^{m} c_i \lambda_i}{\sqrt{\sum_{i=1}^{m} c_i^2 \lambda_i}} \tag{50}$$

Next, we need to simply send all the Poisson parameter to be sufficiently high, $\lambda_i \to \infty \; \forall i$. This implies that all λ_i are approximately equal to λ which gives the following form

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \sqrt{\lambda} \frac{\sum_{i=1}^{m} c_i w_i' - \sum_{i=1}^{m} c_i}{\sqrt{\sum_{i=1}^{m} c_i^2}} \tag{51}$$

By the central limit theorem, this random variable becomes a standard normal random variable. Similarly,

$$\frac{z_{MF} - \mathbb{E}[z_{MF}]}{\sigma} = \frac{\sum_{i=1}^{m} c_i w_i - \sum_{i=1}^{m} c_i \lambda_i}{\sqrt{\sum_{i=1}^{m} c_i^2 \lambda_i}}$$

$$= \sqrt{m} \frac{\sum_{i=1}^{m} c_i \frac{w_i}{m} - \sum_{i=1}^{m} c_i \frac{\lambda_i}{m}}{\sqrt{\sum_{i=1}^{m} c_i^2 \frac{\lambda_i}{m}}} \tag{52}$$

which again, by the central limit theorem approaches a normal random variable as $m \to \infty$. $\qquad\square$

This lemma has the following implications. As a signal is spread out over more pixels, either through a long streaking object or a large PSF, the hypothesis distributions should approach Gaussian. Additionally, as the photometric magnitude increases, the hypothesis distribution will approach Gaussian. In testing, the Gaussian approximation appears to be appropriate in most situations.

E. Qualitative Comparison of Likelihoods

To test the effectiveness of these two methods, the following Monte Carlo analysis is shown. The noise model is additive zero mean Gaussian, as in Equation (14), and the signal shape is assumed to be perfectly known. Each likelihood function is tested on an image containing signal with a perfect matched filter (no error in template or total magnitude), giving g_{s+w}. Each likelihood function is also tested on an image containing no signal, only noise, giving g_{0+w}. Because a Bernoulli filter uses a PDF, all particles are normalized each frame. Because of this, only the relative value between the two likelihoods affects the performance of the filter. To evaluate the performance of each likelihood function, a likelihood gain is calculated, g_{s+w}/g_{0+w}.

The likelihood gain is calculated for both filters over a variety of SNR values. It is also calculated over a variety of errors in predicted signal strength for the likelihood ratio function, and a variety of probability of false alarm values for the hypothesis test likelihood function. Each likelihood gain is calculated over 10,000 test runs, and a mean and variance of the likelihood gain are calculated.

Figures 1a and 1b show the shape of the gain with respect to the two variables of interest. The hypothesis test likelihood function can achieve higher gain values with a diminishing probability of false alarm, but at the expense of increased variance. Figures 2a and 2b show the shape of the variance with respect to the two variables of interest, which appears to vary directly with gain. Figure 3 gives the best reason for why the new likelihood is better; while variance in gain is high, the new method preforms better at low SNR.

IV. Implementation Notes

A. Multi-Bernoulli Filter Birth Model

As the Multi-Bernoulli filter runs, multiple single Bernoulli filters run in parallel. New Bernoulli filters, sampled from a birth model, Γ, are constantly added into the system to promote exploration of the state space. Though the birth model is not a focus of this paper, a review of some options will be briefly presented See [27] for more information.

The most basic birth model is a completely uniform distribution, over the considered state space. A more advanced birth model can enforce a series of constraints in the form of sets of possible states, with a uniform distribution used on this constrained space. More advanced methods can define a non uniform PDF either a priori, or with a learning algorithm [20].
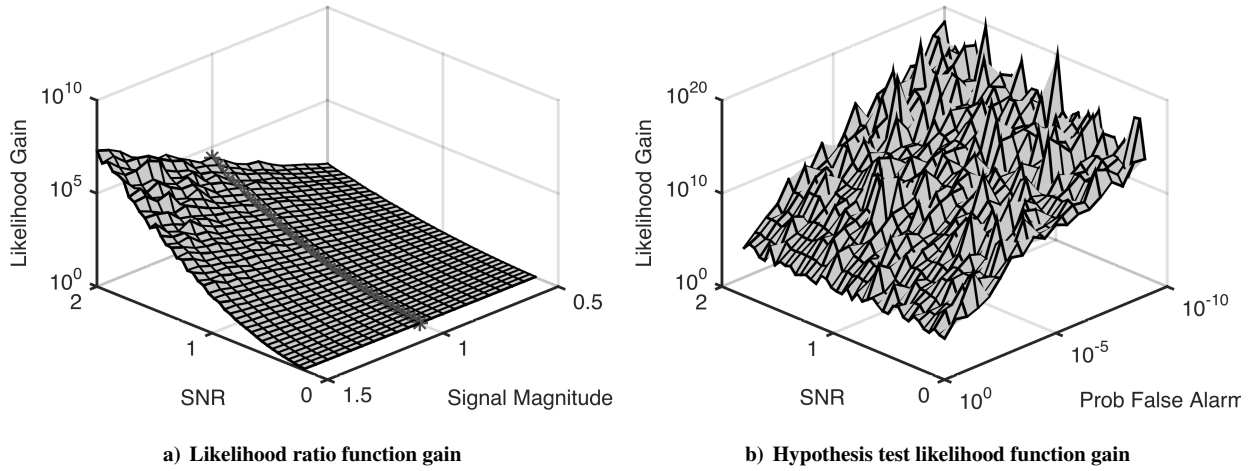
a) Likelihood ratio function gain

b) Hypothesis test likelihood function gain

Figure 1. comparison of likelihood function gain

a) Likelihood ratio function variance

b) Hypothesis test likelihood function variance
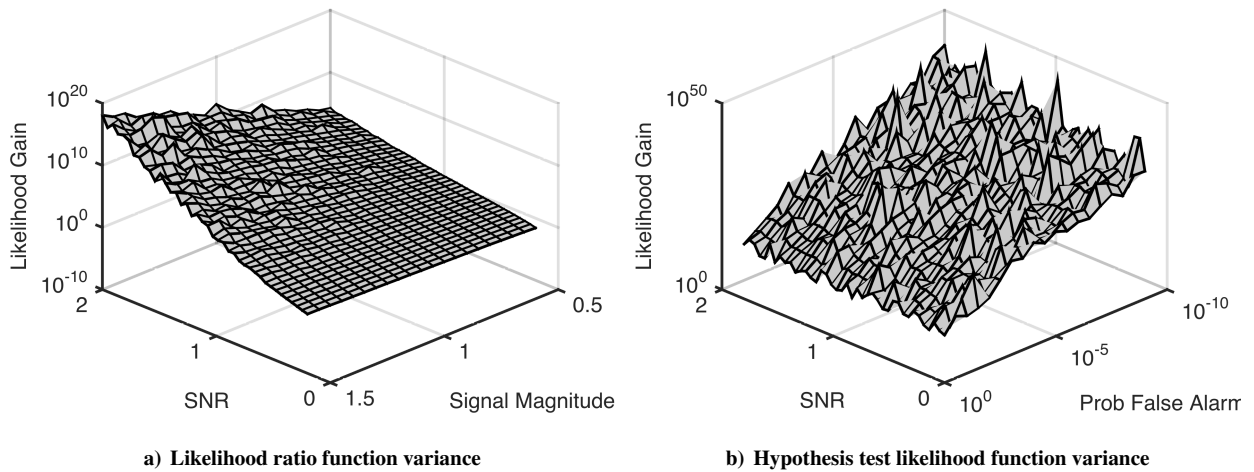
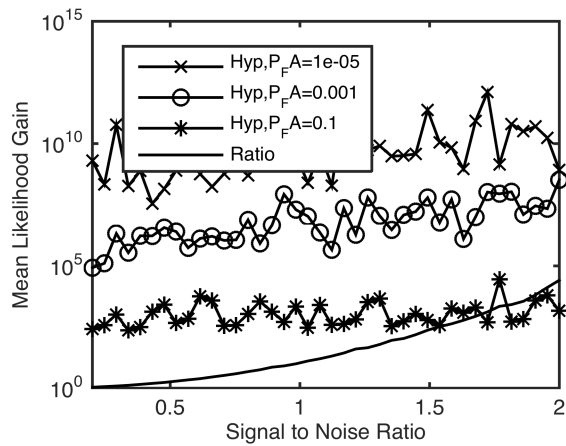**Figure 2. comparison of likelihood function variance**

Figure 3. Likelihood gain comparison between methods

The first advanced birth model that can be used is based on partial prior knowledge, such as an admissible region, or prior PDF. These prior models can be sampled either as a uniform distribution of the set of all possible orbits, or directly from a prior PDF. In the case of a prior in the full orbit state space, these sampled particles can then be projected into the reduced image state space. See [2] for more details about propagation and projection of particles. This paper focuses on when prior knowledge is not available, so this topic is not explored further in this paper.

In tracking with no prior information, two basic constraints are possible. First, a minimum number of timesteps in which the object is in a sensors field of view an be used to upper bound velocity. Second, orbital mechanics can be used to lower bound velocity based on assuming non ballistic trajectories. These topics are explored in detail in [2].

Finally, regions of state space can be hypothesized to further constrain the problem. If observations are taken with the intention of looking for a particular class of space objects, this may imply only certain trajectories through the field of view are possible, allowing explicit constraints on the birth model. One of the results of this paper, looking at geostationary objects, compares uniform birth model with a model that assumes low inclination orbits, with the latter performing better. Any assumptions on the orbit will come at the cost of failing to track objects not within the assumed model.

B.   Multi-Bernoulli Filter Implementation Notes

This section outlines 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 is 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 [17]. 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 makes 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.

The state space chosen for frame to frame tracking is a 2 dimensional in image position and velocity The state, $\mathbf{x}$, is modeled as position and velocity, in the frame, measured in pixels.

$$\mathbf{x} = [x, \; y, \; \dot{x}, \; \dot{y}]^T \tag{53}$$

The dynamics chosen for frame to frame tracking are a simple double integrator.

$$\mathbf{x}_{k+1} = \begin{bmatrix} I_{2\times2} & I_{2\times2}\Delta t \\ 0_{2\times2} & I_{2\times2} \end{bmatrix} \mathbf{x}_k + \mathbf{v} \tag{54}$$

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}) \tag{55}$$

The process noise for the dynamics uses a covariance matrix of $diag([5, 5, 1, 1]^T)$, measured in pixels and pixels/second.

Because of computational restriction, the Bernoulli filter typically cannot fully sample the state space, as defined by the birth model. Particles randomly sampled from a birth model near the true space object in position space typically don't have the correct velocity state, due to particle impoverishment. By artificially inflating the dynamics covariance, state space exploration is encouraged. This leads to better state space exploration and convergence for Bernoulli filters, but worse final tracks generated by the algorithm. There is a trade off which is not further explored in this paper.

The filters typically use 20,000 particles per Bernoulli filter. The filters typically perform well for this number of particles, and computation times are bearable for implementation on a person laptop.

V.   Empirical Data Results

A.   Geostationary Object, small Field of View

This section presents the validation of the multi-Bernoulli filter on real telescope data. The data was taken by GT-SORT, a 0.5m Raven-class telescope located at Georgia Tech. The field of view is approximately 14 by 11 (arc

minutes). The sensor has a resolution of 2736 by 2192 pixels. This specific data set was taken by performing sidereal stare on a region known to contain GEO objects. Images were taken with 1 second exposures, at a rate of approximately 1 image every 1.6 seconds. Twenty five consecutive images are used for these results. Over the twenty five images, an GEO object starts in the center of the image and eventually exits at the bottom of the image. A second object enters from the top of the image for 12 frames The zero inclination GEO orbit plane runs approximately vertical through the image. The PSF is estimated from a bright star in the image, and assumed to be invariant between frames and location in the image. Two different birth models are used. The first birth model uses the estimated velocity of a zero inclination geostationary object, to better capture expected objects. The second birth model uses the estimated velocity of a any geostationary object, but doesn't assume a direction of motion. This second birth model is the equivalent of circular geostationary altitude orbits with no assumed inclination.

Using Equations (17) and (18), the per pixel and total photometric SNRs of this object can be estimated from the images. The per pixel SNR is approximately 0.18, while the total photometric SNR is approximately 6.9. These values should be taken as rough estimates; for objects of 0.18 per pixel SNR, noise is pervasive, and it is difficult to determine which pixels contain signal. Figure A shows the zoomed in point spread of the object.


Figure 4. Object tracks overlayed on image. 2192 by 2736 pixels. Assumed heading birth model.

The resulting two tracks for the assumed heading birth model are plotted over an image from the data set in Figure 4. Similarly, the results for the no assumed heading birth model are plotted over an image from the data set in Figure 5. When a Bernoulli filter is initialized, due to particle impoverishment and the inherent randomness of the system, the maximum a posteriori estimate randomly moves around the image. Both estimates are random location in the image for a few frames before the Bernoulli filters converge; this is normal behavior. This is due to the low signal SNR; noise randomly mimics object signal enough to cause random hits occasionally, but the true object signal also follows the dynamic model, allowing convergence. There are also around 5 other Bernoulli filters operating in parallel, which are not shown, due to lack of significant convergence. Stars in the images are inertially fixed from frame to frame. The stars are detected in preprocessing and zeroed out in the image. Stars may appear to have a vertical oblong shape; this is due to a combination of vibrations in the telescope mount, wind buffeting, and atmospheric turbulence. The objects appear somewhere between 8 or 9 arcminutes apart in the images.

Figure 5. Object tracks overlayed on image. 2192 by 2736 pixels. No assumed heading birth model.

Figure 6. Object size and signal strength.

B.   Tumbling Low Earth Orbit, Wide Field of View

On the 26th of March, 2016, the Hitomi Astro-H spacecraft lost communication with the ground and experienced a break up due to what is though to be an unstable spin rate. On the 29th of March, 2016, a series of images was taken with a wide filed of view sensor of the Hitomi as it passed over the Georgia Tech Observatory. The sensor has a field of view of about 20 degrees, and a resolution of 1024 by 1280 pixels. The object is seen for about 26 frames, over which large variations in photometric magnitude are observered, as the object tumbles. Because observations were taken multiple days after the breakup, no debris object were seen in the images. Figure 7 shows the observed track laid over a sample image. The space object total photometric SNR variations over the iterations is shown in Figure 8. The minimum total photometric SNR is 1.5, while the total photometric SNR falls below 3 for at most three consecutive iterations. Star identification has been performed for this data set with the open source astrometry.net algorithm [a].

Filter performance is not hampered by object signal variations. Object probability of existence converges to 1 after approximately 3 iterations, and never decreases. This example highlights an important property of the hypothesis test likelihood; variations in brightness do not affect performance at all, because no approximation of photometric brightness is required.


**Figure 7. Track of Hitomi Astro-H. Triangle is beginning of track.**

Figure 8. SNR variations of Hitomi Astro-H over observation.

VI.   Acknowledgment

---

[a]astrometry.net, May 3rd, 2016

References

[1] Blake, T., "Space Domain Awareness (SDA)," Tech. rep., DTIC Document, 2011.

[2] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "ORBIT DETERMINATION FOR PARTIALLY UNDERSTOOD OB-JECT VIA MATCHED FILTER BANK," *AAS/AIAA Astrodynamics Specialists Meeting*, 2015.

[3] DeMars, K. J., Jah, M. K., and Schumacher, P. W., "Initial orbit determination using short-arc angle and angle rate data," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 48, No. 3, 2012, pp. 2628–2637.

[4] Mahler, R. P., *Statistical multisource-multitarget information fusion*, Artech House, Inc., 2007.

[5] Blackman, S., "Multiple Hypothesis Tracking For Multiple Target Tracking," *IEEE A and E Systems Magazine*, 2004.

[6] Reed, I., Gagliardi, R., and Shao, H., "Application of three-dimensional filtering to moving target detection," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. AES-19, No. 6, 1983, pp. 898–905. doi:10.1109/taes.1983.309401.

[7] Tonissen, S. and Evans, R., "Peformance of dynamic programming techniques for Track-Before-Detect," *IEEE Trans. Aerosp. Electron. Syst.*, Vol. 32, No. 4, 1996, pp. 1440–1451. doi:10.1109/7.543865.

[8] Zhai, C., Shao, M., Nemati, B., Werne, T., Zhou, H., Turyshev, S. G., Sandhu, J., Hallinan, G., and Harding, L. K., "DETEC-TION OF A FAINT FAST-MOVING NEAR-EARTH ASTEROID USING THE SYNTHETIC TRACKING TECHNIQUE," *ApJ*, Vol. 792, No. 1, aug 2014, pp. 60. doi:10.1088/0004-637x/792/1/60.

[9] Ristic, B., Arulampalam, S., and Gordon, N. J., *Beyond the Kalman filter: Particle filters for tracking applications*, Artech house, 2004.

[10] DeMars, K. J. and Jah, M. K., "Probabilistic initial orbit determination using gaussian mixture models," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 5, 2013, pp. 1324–1335.

[11] Daum, F. and Huang, J., "Curse of dimensionality and particle filters," *Aerospace Conference, 2003. Proceedings. 2003 IEEE*, Vol. 4, IEEE, 2003, pp. 4_1979–4_1993.

[12] Hoseinnezhad, R., Vo, B.-N., Vo, B.-T., and Suter, D., "Visual tracking of numerous targets via multi-Bernoulli filtering of image data," *Pattern Recognition*, Vol. 45, No. 10, oct 2012, pp. 3625–3635. doi:10.1016/j.patcog.2012.04.004.

[13] Hoseinnezhad, R., Vo, B.-N., and Vo, B.-T., "Visual Tracking in Background Subtracted Image Sequences via Multi-Bernoulli Filtering," *IEEE Transactions on Signal Processing*, Vol. 61, No. 2, jan 2013, pp. 392–397. doi:10.1109/tsp.2012.2222389.

[14] Fujimoto, K., Uetsuhara, M., and Yanagisawa, T., "Statistical Track-Before-Detect Methods Applied to Faint Optical Obser-vations of Resident Space Objects," *Advanced Maui Optical and Space Surveillance Technical Conference*, 2015.

[15] Hussein, I., DeMars, K. J., Früh, C., Erwin, R. S., Jah, M. K., et al., "An AEGIS-FISST integrated detection and tracking approach to Space Situational Awareness," *Information Fusion (FUSION), 2012 15th International Conference on*, IEEE, 2012, pp. 2065–2072.

[16] Davey, S. J., Rutten, M. G., and Cheung, B., "A Comparison of Detection Performance for Several Track-before-Detect Algorithms," *EURASIP Journal on Advances in Signal Processing*, Vol. 2008, 2008, pp. 1–11. doi:10.1155/2008/428036.

[17] Vo, B.-N., Vo, B.-T., Pham, N.-T., and Suter, D., "Joint Detection and Estimation of Multiple Objects From Image Observa-tions," *IEEE Transactions on Signal Processing*, Vol. 58, No. 10, oct 2010, pp. 5129–5141. doi:10.1109/tsp.2010.2050482.

[18] Zingarelli, J. C., Pearce, E., Lambour, R., Blake, T., Peterson, C. J., and Cain, S., "Improving the Space Surveillance Tele-scope's Performance Using Multi-Hypothesis Testing," *The Astronomical Journal*, Vol. 147, No. 5, 2014, pp. 111.

[19] Reuter, S., Vo, B.-T., Vo, B.-N., and Dietmayer, K., "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, Vol. 62, No. 12, jun 2014, pp. 3246–3260. doi:10.1109/tsp.2014.2323064.

[20] Ristic, B., *Particle Filters for Random Set Models*, Springer Science + Business Media, 2013. doi:10.1007/978-1-4614-6316-0.

[21] Sanson, F. and Frueh, C., "Noise Quantification in Optical Observations of Resident Space Objects for Probability of Detection and Likelihood," *AAS/AIAA Astrodynamic Specialist Conference, Vail, CO*, 2015, pp. 15–634.

[22] Benezeth, Y., Jodoin, P., Emile, B., Laurent, H., and Rosenberger, C., "Review and evaluation of commonly-implemented background subtraction algorithms," *2008 19th International Conference on Pattern Recognition*, Institute of Electrical & Electronics Engineers (IEEE), dec 2008. doi:10.1109/icpr.2008.4760998.

[23] Pratt, W. K., *Digital Image Processing*, Wiley-Blackwell, jan 2007. doi:10.1002/0470097434.

[24] Turin, G. L., "An introduction to matched filters," *IRE Transactions on Information Theory*, Vol. 6, No. 3, jun 1960, pp. 311–329. doi:10.1109/tit.1960.1057571.

[25] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "Space Object Detection in Images Using Matched Filter Bank and Bayesian Update (Submitted)," *Journal of Guidance, Control, and Dynamics*, 2016.

[26] Montgomery, D. C. and Runger, G. C., *Applied statistics and probability for engineers*, John Wiley & Sons, 2010.

[27] Murphy, T. S., Holzinger, M. J., and Flewelling, B., "DIRECT IMAGE-TO-LIKELIHOOD FOR TRACK-BEFORE-DETECT MULTI-BERNOULLI FILTER," *AAS/AIAA Space Flight Mechanics Meeting*, 2016.

[28] Milani, A., Gronchi, G. F., Vitturi, M. d., and Knežević, Z., "Orbit determination with very short arcs. I admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 90, No. 1-2, 2004, pp. 57–85.

[29] Worthy III, J. L. and Holzinger, M. J., "Incorporating Uncertainty in Admissible Regions for Uncorrelated Detections," *Journal of Guidance, Control, and Dynamics*, Vol. 0, No. 0, 2015, pp. 1–17. doi:10.2514/1.g000890.

[30] Worthy III, J. L. and Holzinger, M. J., "Use of Uninformative Priors to Initialize State Estimation for Dynamical Systems," *AIAA/AAS Astrodynamics Specialist Conference*, 2015.

[31] Tommei, G., Milani, A., and Rossi, A., "Orbit determination of space debris: admissible regions," *Celestial Mechanics and Dynamical Astronomy*, Vol. 97, No. 4, 2007, pp. 289–304. doi:10.1007/s10569-007-9065-x.