

Sequential Multi-Target Detection & Tracking of Low SNR Near-Earth Objects

Shahzad Virani^{a,1,*}, Marcus J. Holzinger^{a,2}, Vishnu Reddy^{b,3}

^aUniversity of Colorado - Boulder, 3775 Discovery Drive, Boulder, CO, 80303, USA

^bUniversity of Arizona, 1629 E University Blvd, Tucson, AZ 85721, USA

Abstract

This paper leverages the measurement likelihood ratio for electro-optical low SNR problem to motivate the use of target likelihood ratio recursion for detecting and tracking dim objects. A novel intensity marginalized likelihood ratio is proposed, which allows for simultaneous detection of low SNR targets in the entire surveillance region. This allows for quick computation of the probability of existence for all targets and any object with a probability of existence greater than a chosen threshold are handed-off to an image-based Multi-Bernoulli filter for track maintenance. This proposed method is implemented and tested on real data of Apophis asteroid and 2020 SO, which was initially cataloged as a NEO by the Minor Planet Center. In the 2020 SO dataset, the proposed filter was able to detect 5 more other objects than expected ^{††}. A follow-on journal paper will discuss the exact implementation details for computing the intensity marginalized likelihood function in real-time as well as rigorously compare the performance of the proposed method to other widely used low SNR detection algorithms.

Keywords: Intensity-Marginalized Likelihood, Track-before-Detect, Finite Set Statistics

1. Introduction

Traditional approaches for discovering Near-Earth Objects (NEOs) include a moving target indicator approach which involves identifying objects that move relative to the background stars in a sequence of optical images [1]. Automated software pipelines on major optical CCD surveys such as Pan-STARRS and Catalina Sky Survey (CSS) can produce tracklets of known and unknown asteroids in near real-time. Recently, these methods have been extended for detecting and tracking multiple-objects simultaneously [2]. These methods, which are characterized as Detect-before-Track, perform well when the apparent magnitude of the object of interest is sufficiently high compared to the noise floor and the limiting magnitude of the instrument. As the planetary defense community seeks to catalog dimmer NEOs over time, these methods are no longer sufficient as individual images of the object do not contain enough signal that can be extracted (typically under SNR 2), resulting in a miss-detection. This motivates the development of Track-before-Detect methods [3], [4], [5] in which all the data collected, including the pixel values, is used to simultaneously detect and track the objects as opposed to detecting the objects first by a thresholding approach.

In 2014, Shao et. al popularized the method of synthetic tracking for detecting NEOs. This well-established method, which has classically been used by NEO community, relies on acquiring many subsequent short exposures of a NEO of interest, thereby “freezing” their motion in an individual image

^{††}A video of the results can be accessed at shez.space/pdc2021

*Corresponding author

Email addresses: shez.virani@colorado.edu (Shahzad Virani), marcus.holzinger@colorado.edu (Marcus J. Holzinger), reddy@lpl.arizona.edu (Vishnu Reddy)

¹PhD Student, NDSEG Fellow, Ann & H.J. Smead Department of Aerospace Engineering Sciences

²Associate Professor, Ann & H.J. Smead Department of Aerospace Engineering Sciences

³Associate Professor, Department of Planetary Science, Lunar and Planetary Laboratory

[6]. An appropriate velocity vector is then computationally searched for by shifting successive frames relative to each other and then co-adding the shifted frames to create a long-exposure image as if a telescope were tracking the object. This method significantly decreases the losses due to a longer streaked image of the NEO since all of the signal is synthetically collected in a few pixels hence increasing the photometric signal-to-noise ratio (SNR) of the object. This batch-processing approach performs well if all of the images that are shifted-and-stacked contain some signal from the object, which is definitely not guaranteed during blind detection of NEOs.

In this paper, the problem of detecting low SNR NEOs is formulated using a Bayesian approach in which the position, velocity as well as the intensity of the NEO is estimated sequentially. A novel intensity-marginalized likelihood function is introduced as an effective “detector” that allows for computing and updating the probability of existence of a NEO as more data is collected. This approach also leverages Finite Set Statistics based Multi-Bernoulli filter to jointly detect and track multiple NEOs as well as model the birth and death (or appearance and disappearance) of NEOs in images. In blind detection and discovery of new NEOs, this becomes crucial since the position and velocities of new objects are not known a priori. The contributions from this paper are to introduce sequential detection and tracking of low SNR NEOs using intensity-marginalized likelihood function as well demonstrate the results on real data of Apophis and Centaur R/B 2020 SO using ground-based optical telescope observations collected by the Lunar and Planetary Laboratory at University of Arizona.

2. Theory

The theory discussed in this section was initially introduced by the authors in [7], but is summarized here for completeness. This section begins by describing the dynamics and measurement model of targets in electro-optical images for the low SNR detection and tracking problem. Then, the measurement likelihood function is to motivate the target likelihood ratio based recursion. The Bayesian filtering prediction and update equations are then transformed into likelihood ratio recursion. Additionally, an intensity marginalized likelihood ratio is introduced along with maximum likelihood estimate for intensity.

2.1. Dynamics Model

Consider a target of an unknown brightness/intensity moving in the $x - y$ image plane. The discrete dynamics of the target can be modeled as follows:

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k, \mathbf{v}_k) \quad (1)$$

where k denotes the time index, \mathbf{v}_k is the discrete-time process noise, $\mathbf{x}_k \in \mathbb{R}^5$ is the target state vector defined as follows

$$\mathbf{x}_k = [x_k \quad y_k \quad \dot{x}_k \quad \dot{y}_k \quad I_k]^T \quad (2)$$

Here the surveillance region is defined by the image and therefore (x_k, y_k) denote the position of the target in pixel space, (\dot{x}_k, \dot{y}_k) denote the velocity of the target in pixel space, and I_k denotes the target brightness/intensity in pixel counts. Additionally, if the target disappears from the surveillance region, it has a null state represented by ϕ . This is done to explicitly model the birth and death process. For convenience, also define the target's kinematic states as

$$\tilde{\mathbf{x}}_k = [x_k \quad y_k \quad \dot{x}_k \quad \dot{y}_k]^T \quad (3)$$

It is important to note that given the dynamics along with the statistics of the process noise is equivalent to knowing the transition density, $p(\mathbf{x}_{k+1}|\mathbf{x}_k)$. For additional information, see [8].

2.2. Measurement Model

The data collected are sequences of two-dimensional optical images of the surveillance region, in which each frame consists of $n \times m$ pixels of size $\Delta_x \times \Delta_y$. Let (i, j) $1 \leq i \leq n$, $1 \leq j \leq m$ denote the location of the center of each pixel. Additionally, let $\mathbf{y}_k = \{y_k^{(i,j)} : i = 1, \dots, n, j = 1, \dots, m\}$ denote the image observation received at time k and $\mathbf{Y}_{1:k} = \{\mathbf{y}_b, b = 1, \dots, k\}$ denote the set of measurements received up until time k . The intensity measured in pixel (i, j) can be modeled as follows:

$$y_k^{(i,j)} = \begin{cases} h_k^{(i,j)}(\mathbf{x}_k) + w_k^{(i,j)} & \text{if target present} \\ w_k^{(i,j)} & \text{if target absent} \end{cases} \quad (4)$$

where $h_k^{(i,j)}(\mathbf{x}_k)$ denotes the intensity in the pixel (i, j) contributed by a target with state \mathbf{x}_k , $w_k^{(i,j)}$ is the measurement noise in pixel (i, j) . In other words, the measurement received in a pixel will contain some signal + noise if a target exists, but only noise if no target exists. For this discussion, the noise is assumed to be Gaussian, i.e. $w_k^{(i,j)} \sim \mathcal{N}(0, \sigma^2)$. It is also important to note that unresolved imagery of space objects typically contains sky background and bright stars and are assumed to be removed from the image before detecting low SNR targets. This pre-processing step is discussed in detail in [9].

In the case of unresolved imagery, a target appears as a point source in the image and the photons from it are dispersed based on the point spread function (PSF) of the imaging system. This PSF is often approximated as a two-dimensional Gaussian function with circular symmetry and the expected signal in each pixel, $h_k^{(i,j)}(\mathbf{x}(k))$, is the integral of the PSF over the pixel (i, j) [3]. Therefore, the pixel intensity in (i, j) contributed by a target at position (x_k, y_k) with intensity I_k is

$$h_k^{(i,j)}(\mathbf{x}_k) \approx \frac{\Delta_x \Delta_y I_k}{2\pi \Sigma^2} \exp \left[-\frac{(i\Delta_x - x_k)^2 + (j\Delta_y - y_k)^2}{2\Sigma^2} \right] \quad (5)$$

where Σ is the blur factor, which is a function of the optics and seeing conditions and is assumed to be known. This is often estimated from stars within the image that are properly sampled. Several classical methods already exist to perform this step [10]. The results in this paper use Maximum Likelihood Estimation (MLE) to fit the PSF profile and the details will be discussed in a follow-on journal paper. Additionally, for convenience, also define the intensity normalized pixel contribution as

$$\begin{aligned} \tilde{h}_k^{(i,j)}(\tilde{\mathbf{x}}_k) &= \frac{\Delta_x \Delta_y}{2\pi \Sigma^2} \exp \left[-\frac{(i\Delta_x - x_k)^2 + (j\Delta_y - y_k)^2}{2\Sigma^2} \right] \\ h_k^{(i,j)}(\mathbf{x}_k) &= I_k \cdot \tilde{h}_k^{(i,j)}(\tilde{\mathbf{x}}_k) \end{aligned} \quad (6)$$

Since we assumed that a target appears as a two-dimensional circularly symmetric Gaussian function, it is evident from Eq. (5) that the pixel intensity contribution is only a function of the target's intensity and its positional states, (x_k, y_k) but not its velocity states. If the targets are moving fast relative to the sensor, then this assumption is no longer valid.

2.3. Measurement Likelihood Ratio

Since the measurement noise from pixel-to-pixel is assumed independent, the likelihood for the entire image at time t_k can be computed by the product of the individual pixel likelihoods. However, the individual pixel likelihood depends on whether a target exists in its vicinity or not. Hence, using Eq. (4), the measurement likelihood can be written as

$$\begin{aligned} p\left(y_k^{(i,j)} \mid \mathbf{x}_k\right) &= \mathcal{N}\left(y_k^{(i,j)}; h_k^{(i,j)}, \sigma^2\right) \quad (\text{if a target is in vicinity of pixel } (i, j)) \\ p\left(y_k^{(i,j)} \mid \phi_k\right) &= \mathcal{N}\left(y_k^{(i,j)}; 0, \sigma^2\right) \quad (\text{if no target is in vicinity of pixel } (i, j)) \end{aligned} \quad (7)$$

If the measurement likelihood is to be computed for a pixel (i, j) that is far from the target's hypothesized position (x_k, y_k) , then the target's intensity contribution, i.e. $h_k^{(i,j)}$ will approach zero. Therefore, it is trivial to see that the existence of a target will only affect the measurement likelihood in the pixels that are close to the target's state. Let $C_x(\mathbf{x}_k) = \{j : \|j\Delta_x - x_k\| \leq r\}$ and $C_y(\mathbf{x}_k) = \{i : \|i\Delta_y - y_k\| \leq r\}$ define the vicinity around a target with state \mathbf{x}_k of distance r , which is a user defined parameter based on the known PSF of the optics and seeing condition. Therefore, the total likelihood for the entire image at time t_k can be written as:

$$p(\mathbf{y}_k \mid \mathbf{x}_k) = \prod_{i \in C_y(\mathbf{x}_k)} \prod_{j \in C_x(\mathbf{x}_k)} p\left(y_k^{(i,j)} \mid \mathbf{x}_k\right) \prod_{i \notin C_y(\mathbf{x}_k)} \prod_{j \notin C_x(\mathbf{x}_k)} p\left(y_k^{(i,j)} \mid \phi_k\right) \quad (8)$$

This equation can be further simplified as follows:

$$p(\mathbf{y}_k \mid \mathbf{x}_k) = \prod_{i \in C_y(\mathbf{x}_k)} \prod_{j \in C_x(\mathbf{x}_k)} \underbrace{\frac{p\left(y_k^{(i,j)} \mid \mathbf{x}_k\right)}{p\left(y_k^{(i,j)} \mid \phi_k\right)}}_{l\left(y_k^{(i,j)} \mid \mathbf{x}_k\right)} \underbrace{\prod_{i=1}^n \prod_{j=1}^m p\left(y_k^{(i,j)} \mid \phi_k\right)}_{p(\mathbf{y}_k \mid \phi_k)} \quad (9)$$

In Bayesian framework, the update equation has a normalizing constant in the denominator as shown in Eq. (12). Hence, the measurement likelihood is only required to be known up to a proportionality constant. We can leverage this property and define the first term in the above equation as the pixel-wise measurement likelihood ratio which only depends on the vicinity of the hypothesized target state as $l(y_k^{(i,j)}|\mathbf{x}_k)$. Therefore, the measurement likelihood up to a proportionality constant is related to the measurement likelihood ratio as $p(\mathbf{y}_k|\mathbf{x}_k) \propto \frac{p(\mathbf{y}_k|\mathbf{x}_k)}{p(\mathbf{y}_k|\phi_k)} = l(\mathbf{y}_k|\mathbf{x}_k)$ and simplifying gives:

$$l(\mathbf{y}_k|\mathbf{x}_k) = \prod_{i \in C_y(\mathbf{x}_k)} \prod_{j \in C_x(\mathbf{x}_k)} \frac{p(y_k^{(i,j)}|\mathbf{x}_k)}{p(y_k^{(i,j)}|\phi_k)} = \exp \left\{ -\frac{1}{2\sigma^2} \left[\sum_{i \in C_y(\mathbf{x}_k)} \sum_{j \in C_x(\mathbf{x}_k)} h_k^{(i,j)}(\mathbf{x}_k) (h_k^{(i,j)}(\mathbf{x}_k) - 2y_k^{(i,j)}) \right] \right\} \quad (10)$$

This measurement likelihood ratio is convenient since it *only* depends on the pixels in the vicinity of the target state. This property is crucial in extending this method for multi-target tracking under the assumption that no two target PSFs overlap. Additionally, since Bayesian recursion only requires the measurement likelihood to be known up to a proportionality constant, this measurement likelihood ratio allows us to incorporate the data into the filter without having to compute the actual likelihood. This also motivates the reformulation of the Bayesian recursion in terms of the *target likelihood ratio density* as opposed to computing both the posterior target and null pdfs.

2.4. Bayesian Estimation of Low SNR Target

In the Bayesian framework, the objective is to compute the posterior pdf at time $k+1$, i.e. $p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k+1})$, given the posterior at time k , i.e. $p(\mathbf{x}_k|\mathbf{Y}_{1:k})$, and the measurement received, \mathbf{y}_{k+1} . This recursive estimation generally involves using the Chapman-Kolmogorov equation to compute the predicted posterior pdf for a discrete dynamical system with a known transition density, $p(\mathbf{x}_{k+1}|\mathbf{x}_k)$, and incorporating the measurements to compute the posterior using Bayes' rule, as shown below:

$$p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) = \int_{\mathcal{X}} p(\mathbf{x}_{k+1}|\mathbf{x}_k) p(\mathbf{x}_k|\mathbf{Y}_{1:k}) d\mathbf{x}_k \quad (11)$$

$$p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) = \frac{p(\mathbf{y}_{k+1}|\mathbf{x}_{k+1}) p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k})}{\int_{\mathcal{X}} p(\mathbf{y}_{k+1}|\mathbf{x}_{k+1}) p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) d\mathbf{x}_{k+1}} \quad (12)$$

2.4.1. Prediction and Update Equations for Target Likelihood Ratio

However, for general single/multi-target tracking, the target can appear into the surveillance region or target state-space, \mathcal{X} and disappear, i.e. the birth and death processes [11]. If the target disappears, its state is not defined and we can denote it using the null state, ϕ , and we augment it to the target state-space: $\mathcal{X} \cup \phi$. Accounting for the null state, the prediction equation, Eq. (11), becomes

$$p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) = \int_{\mathcal{X}} p(\mathbf{x}_{k+1}|\mathbf{x}_k) p(\mathbf{x}_k|\mathbf{Y}_{1:k}) d\mathbf{x}_k + p(\mathbf{x}_{k+1}|\phi_k) p(\phi_k|\mathbf{Y}_{1:k}) \quad (13)$$

$$p(\phi_{k+1}|\mathbf{Y}_{1:k}) = \int_{\mathcal{X}} p(\phi_{k+1}|\mathbf{x}_k) p(\mathbf{x}_k|\mathbf{Y}_{1:k}) d\mathbf{x}_k + p(\phi_{k+1}|\phi_k) p(\phi_k|\mathbf{Y}_{1:k}) \quad (14)$$

where $p(\mathbf{x}_{k+1}|\phi_k)$ and $p(\phi_{k+1}|\mathbf{x}_k)$ represent the birth and death of a target, respectively. If birth and death processes are modeled such that null state probability does not change [11], i.e. the probability mass exchange to and from ϕ and \mathcal{X} is balanced, the predicted null state probability integral simplifies to

$$p(\phi_{k+1}|\mathbf{Y}_{1:k}) = p(\phi_k|\mathbf{Y}_{1:k}) \quad (15)$$

Instead of keeping track of the target null state pdf, it is more convenient to compute the ratio of the target state pdf to the null state pdf, i.e. let the target likelihood ratio density be defined as

$$\Lambda(\mathbf{x}_k|\cdot) = \frac{p(\mathbf{x}_k|\cdot)}{p(\phi|\cdot)} \quad (16)$$

Note that $\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k})$ and $\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k+1})$ denote the predicted and posterior likelihood ratio densities, respectively. Also note that the this target likelihood ratio density is different than the measurement

likelihood ratio, which is defined in Eq. (10). Substituting the target likelihood ratio in the prediction equation and rewriting the predicted target likelihood ratio at t_{k+1} in terms of the posterior target likelihood ratio at t_k results in

$$\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) = \frac{p(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k})}{p(\phi_{k+1}|\mathbf{Y}_{1:k})} = \int_{\mathcal{X}} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \Lambda(\mathbf{x}_k|\mathbf{Y}_{1:k}) d\mathbf{x}_k + p(\mathbf{x}_{k+1}|\phi_k) \quad (17)$$

Similarly, accounting for the null state in the Bayesian update equation (12) and rewriting them using the target likelihood ratio gives

$$\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k+1}) = l(\mathbf{y}_{k+1}|\mathbf{x}_{k+1}) \Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) \quad (18)$$

Therefore, the Bayesian recursion shown in Eqs. (11) and (12) reformulated in terms of target likelihood ratio densities becomes as follows:

$$\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) = \int_{\mathcal{X}} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \Lambda(\mathbf{x}_k|\mathbf{Y}_{1:k}) d\mathbf{x}_k + p(\mathbf{x}_{k+1}|\phi_k) \quad (19)$$

$$\Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k+1}) = l(\mathbf{y}_{k+1}|\mathbf{x}_{k+1}) \Lambda(\mathbf{x}_{k+1}|\mathbf{Y}_{1:k}) \quad (20)$$

For more details, see [7]. Several sensors have multiple detectors that scan a region of the sky. In these instances, the total measurement likelihood ratio is just the product of the measurement likelihood ratio of the individual images from different detector, after being aligned using the celestial coordinates.

2.5. Intensity Marginalized Likelihood Ratio

An important property to note from Eq. (10) is that the likelihood ratio depends on both the target's position as well as the intensity. Using the intensity as a state in the filter adds another dimension to the states space, which can become computationally intensive. Additionally, it is difficult to jointly detect & track a low SNR target and estimate its brightness, especially for multiple targets in a wide field-of-view. To circumvent this problem, the target intensity can be marginalized, resulting in a likelihood ratio that only depends on the kinematic states of the target. By definition, based on Eq. (3), $l(\mathbf{y}_k|\mathbf{x}_k) = l(\mathbf{y}_k|\tilde{\mathbf{x}}_k, I_k)$. This results in

$$l(\mathbf{y}_k|\tilde{\mathbf{x}}_k) = \int_I l(\mathbf{y}_k|\mathbf{x}_k) \cdot p(I_k|\tilde{\mathbf{x}}_k) dI_k \quad (21)$$

Typically, the target intensity is not modeled as a function of the kinematic states. However, one could choose to do so based on the application. The general problem deals with detecting targets between two bounds of SNR, which does not depend on where the target is located in the surveillance region. Hence, the pdf of the intensity can be considered to be independent of the target's kinematic states, i.e. $p(I_k|\tilde{\mathbf{x}}_k) = p(I_k)$. The target intensity can also vary significantly between frames, especially for a low SNR object. Instead of choosing to model the intensity as a Markov process, this paper assumes a uniform distribution, i.e.

$$p(I_k) = \begin{cases} \frac{1}{I_{\max} - I_{\min}}, & I_{\min} \leq I_k \leq I_{\max} \\ 0, & \text{otherwise} \end{cases} \quad (22)$$

Therefore, computing the integral results in the intensity marginalized likelihood ratio as follows:

$$l(\mathbf{y}_k|\tilde{\mathbf{x}}_k) = \frac{1}{I_{\max} - I_{\min}} \int_{I_{\min}}^{I_{\max}} \exp \left\{ -\frac{1}{2\sigma^2} \left[\sum_{i \in \mathcal{C}_y(\mathbf{x}_k)} \sum_{j \in \mathcal{C}_x(\mathbf{x}_k)} h_k^{(i,j)}(\mathbf{x}_k) (h_k^{(i,j)}(\mathbf{x}_k) - 2y_k^{(i,j)}) \right] \right\} dI_k \quad (23)$$

$$= \Omega_3(\tilde{\mathbf{x}}_k, \mathbf{y}_k) \left[\operatorname{erf} \left(\frac{2I_{\max}\Omega_2(\tilde{\mathbf{x}}_k) - \Omega_1(\tilde{\mathbf{x}}_k, \mathbf{y}_k)}{\sqrt{8\sigma^2\Omega_2(\tilde{\mathbf{x}}_k)}} \right) - \operatorname{erf} \left(\frac{2I_{\min}\Omega_2(\tilde{\mathbf{x}}_k) - \Omega_1(\tilde{\mathbf{x}}_k, \mathbf{y}_k)}{\sqrt{8\sigma^2\Omega_2(\tilde{\mathbf{x}}_k)}} \right) \right] \quad (24)$$

where

$$\Omega_1(\tilde{\mathbf{x}}_k, \mathbf{y}_k) = \sum_{i \in \mathcal{C}_y(\tilde{\mathbf{x}}_k)} \sum_{j \in \mathcal{C}_x(\tilde{\mathbf{x}}_k)} 2y_k^{(i,j)} \cdot \tilde{h}_k^{(i,j)}(\tilde{\mathbf{x}}_k) \quad (25)$$

$$\Omega_2(\tilde{\mathbf{x}}_k) = \sum_{i \in \mathcal{C}_y(\tilde{\mathbf{x}}_k)} \sum_{j \in \mathcal{C}_x(\tilde{\mathbf{x}}_k)} (\tilde{h}_k^{(i,j)}(\tilde{\mathbf{x}}_k))^2 \quad (26)$$

$$\Omega_3(\tilde{\mathbf{x}}_k, \mathbf{y}_k) = \frac{\sqrt{\pi\sigma^2}}{(I_{\max} - I_{\min}) \sqrt{2\Omega_2(\tilde{\mathbf{x}}_k)}} \exp \left\{ \frac{[\Omega_1(\tilde{\mathbf{x}}_k, \mathbf{y}_k)]^2}{8\sigma^2\Omega_2(\tilde{\mathbf{x}}_k)} \right\} \quad (27)$$

This equation can then be used to compute the *intensity marginalized likelihood ratio* for an image y_k given a hypothesized target's kinematic state \tilde{x}_k . This equation might seem difficult to compute in real-time since it only allows for calculating the measurement likelihood ratio for each hypothesized target state \tilde{x}_k . However, careful implementation results in fast real-time computation of this likelihood ratio for the states in the surveillance region. Exact implementation details will also be further discussed in the follow-on journal paper.

2.6. Maximum Likelihood Estimate for Intensity

Since the new likelihood is based on marginalization of the target intensity, it is crucial to compute the intensity that maximizes the likelihood so as to pass the tracks to the MBF. This can easily be done using maximum likelihood estimation (MLE). It is important to note that the maximum likelihood estimate for the intensity is derived from the intensity based likelihood function shown in Eq. (10). The authors have previously shown in [7] that

$$\hat{I}_k = \operatorname{argmax}_{I_k} l(y_k | x_k) = \frac{\sum_{i \in C_y(\tilde{x}_k)} \sum_{j \in C_x(\tilde{x}_k)} y_k^{(i,j)} \cdot \tilde{h}_k^{(i,j)}(\tilde{x}_k)}{\sum_{i \in C_y(\tilde{x}_k)} \sum_{j \in C_x(\tilde{x}_k)} (\tilde{h}_k^{(i,j)}(\tilde{x}_k))^2} \quad (28)$$

From Eq. (28), an interesting property to note is that the maximum likelihood estimate for the intensity only depends on the pixel measurements and the blur factor - it *does not* directly depend on the statistics of the background noise.

3. Approach

This section discusses the overall approach that is used in this paper to detect low SNR objects. As noted earlier, to properly use the method proposed, it is crucial that the data is pre-processed such that the noise in the images is of known statistics (close to zero mean and known standard deviation). Typically, the sensor has hot pixels and bias, which can be removed using dark frame subtraction. Additionally, the variations in pixel-to-pixel sensitivity of the detector can also be mitigated using flat field correction. The sky brightness in ground-based sensors varies in an image which causes a bias in the pixel measurements. It must be estimated and subtracted as significantly affects the photometry of the objects. The success of the pre-processing step can be confirmed using the image histogram, as will be shown in the next section.

Once these effects are mitigated, the image only contains noise with known statistics and signals from point sources, including stars. It is not required to subtract the stars from the frame before detecting low SNR targets. However, it is important to note that this will result in many "detections," as stars are typically much brighter than the targets, which in turn increases the run-time significantly. To perform star subtraction, bright objects in the image are detected and compared against the Gaia DR2 Star catalog [12], [13]. This process is done using the widely popular open-source code, Astrometry.net [14].

The detected objects in the images are then associated with stars in the catalog to compute the inertial bearings, which is crucial for determining the right ascension and declination for the objects of interest. Additionally, the wind can cause the telescope to jitter while the data is being taken. Since this jitter is not explicitly modeled in the target dynamics, it is important that the frames in the sequence are registered to the inertial frame. To accomplish this task, the images are aligned by registering the detected stars from one frame to the subsequent frame. The pixels corresponding to the stars are also used to fit PSF using a MLE method. Alternatively, one could also use the traditional DAOPHOT method, which is widely used in the astronomy community [10]. The bright stars can also be removed from the image by simply masking out the associated pixels or iteratively subtracting the estimated signal from the stars.

The data is now pre-processed to correct for bias, pixel variation, sky background, stars, and telescope jitter due to the wind. Once the point-spread function is estimated, the data can then be used with the proposed method to detect and track low SNR objects. The likelihood ratio recursion along with the intensity marginalized likelihood ratio can be implemented in a particle-filter based framework. However, FiSSt based methodologies are more appropriate since they are inherently designed to track multiple targets [5]. These methods are typically robust and allow for modeling target dynamics of choice. However, they suffer from curse of dimensionality due to the particle-based implementations. Specifically,

the ability to detect low SNR targets heavily depends on the birth model. Since the targets can appear anywhere in the frame, this can be a quite expensive computation, especially for blind detection of targets using a wide field-of-view sensor.

As it was shown in the previous section, the intensity marginalized likelihood ratio only depends on the position of the targets in the surveillance region. This property can be leveraged to discretize the position states in the surveillance region and compute the intensity marginalized likelihood ratio in parallel. This essentially becomes a point-mass filter implementation of the likelihood ratio recursion, which allows for quick identification of possible targets. Any target track in the surveillance region with a likelihood ratio above a set threshold can then be passed to a Multi-Bernoulli Filter operating on image data. Ultimately, this methodology allows the point-mass filter to act as a temporary detector for the full Multi-Bernoulli filter. Therefore, under the assumptions made in this paper, the computational efficiency of the likelihood recursion is combined with the sensitivity of the Multi-Bernoulli filter giving a real-time implementation.

4. Results

The proposed method is tested on real dataset from Apophis asteroid and 2020 SO, Centaur R/B from Surveyor 2 mission that was initially cataloged as a NEO by the Minor Planet Center. The data sequences were collected by the Lunar and Planetary Lab at the University of Arizona.

4.1. Results on Apophis Asteroid Data

The images of the asteroid were collected during its closest approach on March 6th 2021 beginning at 03:23 UTC. The data was collected using a 0.5 m telescope and a FLI CCD sensor. Due to the asteroid's size and distance at the approach, the asteroid was very bright. Hence, 100 frames were taken with 0.5 second exposures to mimic a low SNR object. However, the object's photometric SNR was still ~ 3 in each frame. Additional noise was added in post-processing until the photometric SNR was below 1.4.

As discussed, the raw frames were processed by subtracting the dark frames and calibrating by using flat field correction. Since the exposure time was 0.5 seconds, the jitter in the telescope was quite noticeable. So the stars were detected, registered to the Gaia DR2 catalog, and used to align the frames to remove the telescope jitter due to the wind. Pixels associated with properly sampled stars were also used to fit the point spread function. Then the pixels associated with the stars were removed from the images.

Figure 1 shows a raw image from the dataset along with the histogram. Note that the histogram is not zero mean. Figure 2 shows the same image after correcting for dark frame, flat fields, sky background, and telescope jitter. Additionally, stars are removed and noise is added to decrease the photometric SNR below 1.4. The histogram shows a zero-mean image with higher standard deviation for the noise compared to the raw frame. Overlaid on the image is the (x, y) position estimate from the proposed filter.

Figure 3 shows the estimated probability of existence of the target for all the frames. The target was detected with a probability of existence > 0.9 at frame 5. After that frame, the filter sequentially maintains the track and tries to predict the death or disappearance of the target. Since the asteroid is present in all the images, the filter is expected to maintain the track with a probability of 1. However, due to the added noise in the frame, the probability of existence drops if the data at a particular time does not provide enough evidence to support the existence of a target at the predicted location.

4.2. Results on 2020 SO Data

After the 2020 SO was first detected by Pan-STARRS 1 in September 2020, the Minor Planet Center maintained its track. However, the object became too dim to be observable due to its vicinity around the Moon in October 2020. The Lunar and Planetary Lab at University of Arizona tracked the object and collected a sequence of 85 images on October 23 2020 beginning at 04:46 UTC. Since the object was quite dim, 60 second exposures were taken by centering the telescope boresight at the predicted 2020 SO location from the MPC track. Since the target was near the Moon in angular space, the sky background changed significantly over the dataset. To illustrate this, Figure 4 and 5 show the first and last frames from the raw dataset, respectively, with the same limits on the histogram ⁴.

⁴A video of the results can be accessed at shez.space/pdc2021

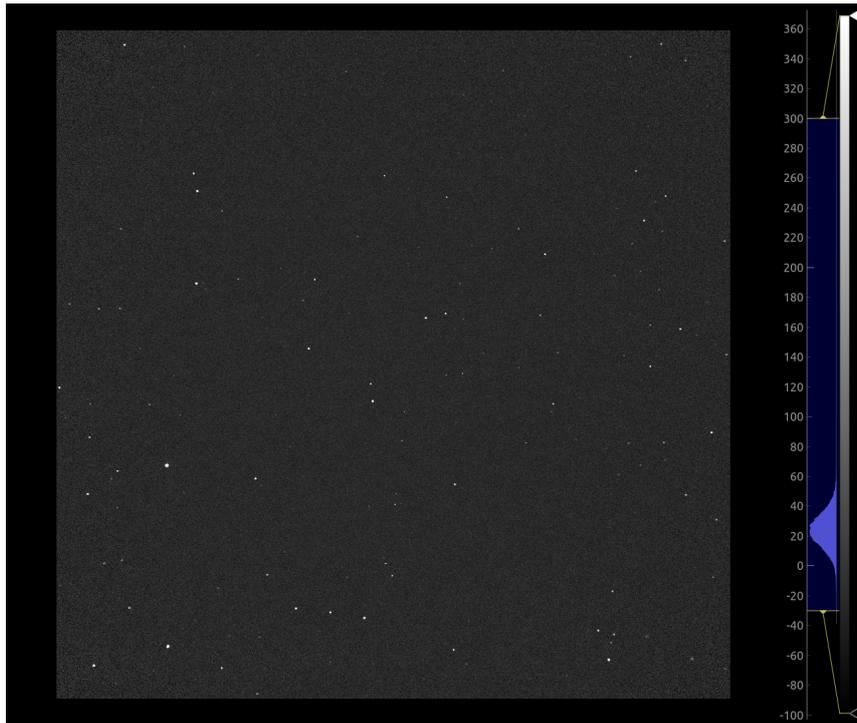


Figure 1: Apophis Dataset: Raw Image with Histogram [0.5s exposure].

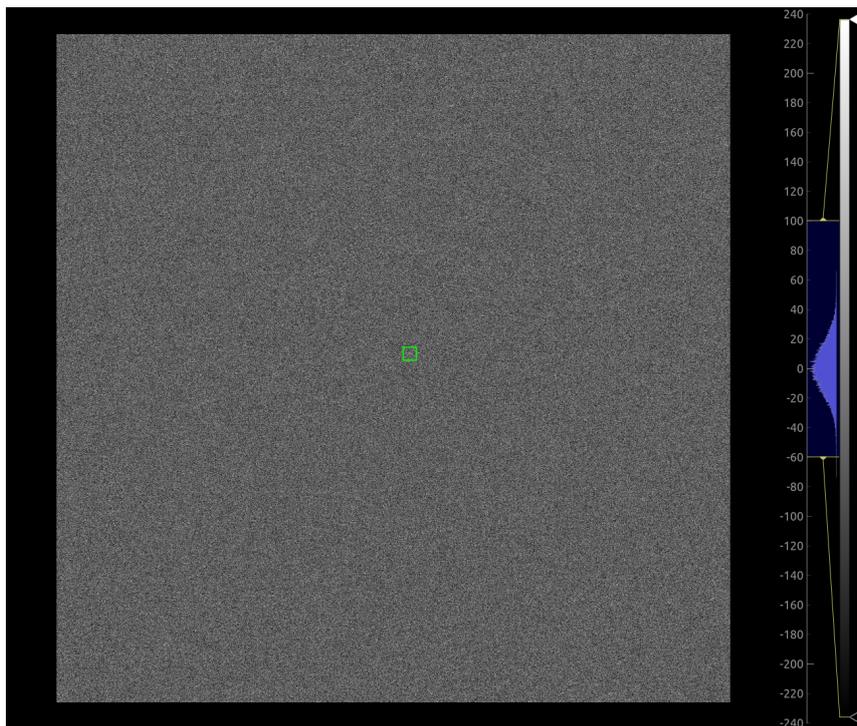


Figure 2: Apophis Dataset: Processed Image with Tracklet & Histogram [0.5s exposure].

The pre-processing step as described in the Approach section was performed on the entire dataset. Since the object was expected to have an apparent magnitude > 21 , it was crucial to remove all the sources corresponding to the stars. The images in the entire dataset were aligned and summed together to create a long exposure of the stars. The stars in this summed frame were then associated with the Gaia DR2 catalog as shown in Figure 6. The red markers indicate the stars that were success-

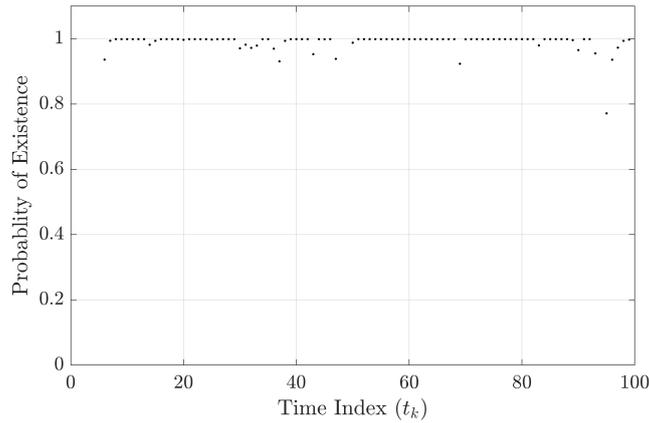


Figure 3: Apophis Dataset: Estimated Probability of Existence vs Time Index.

fully detected in the image and associated with stars in Gaia DR2 catalog. The dimmest star that was detected and identified had a magnitude of 21.02. However, some of the dimmer stars that were detected were not present in the star catalog. Therefore, a mask to subtract the stars was created based on the summed frame as opposed to removing only the stars that were associated with the star catalog.

The processed data was then used with the proposed filter and all objects with an estimated probability of existence greater than 0.85 were considered a detection. Even though the goal was to detect 2020 SO, the proposed filter was able to detect 5 other objects as shown in Figure 7. The position estimates for each object is shown using a green square marker along with an object ID number. Note that the 2020 SO corresponds to object number 2. The arrows indicate the approximate direction of the velocity vectors. Figure 8 shows sum of the 85 sub-frames shifted and added for each object based on its estimated position and velocity. This is shown to confirm the detections of the objects ⁵

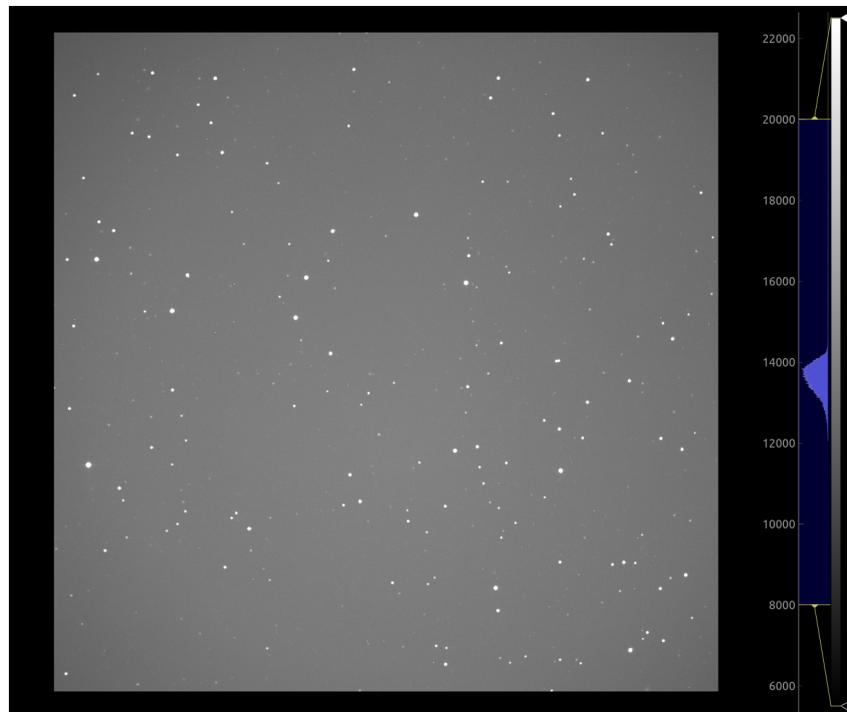


Figure 4: 2020 SO Dataset: Raw Frame 01 [60 sec exposure].

⁵A video of the results can be accessed at shez.space/pdc2021



Figure 5: 2020 SO Dataset: Raw Frame 85 [60 sec exposure].

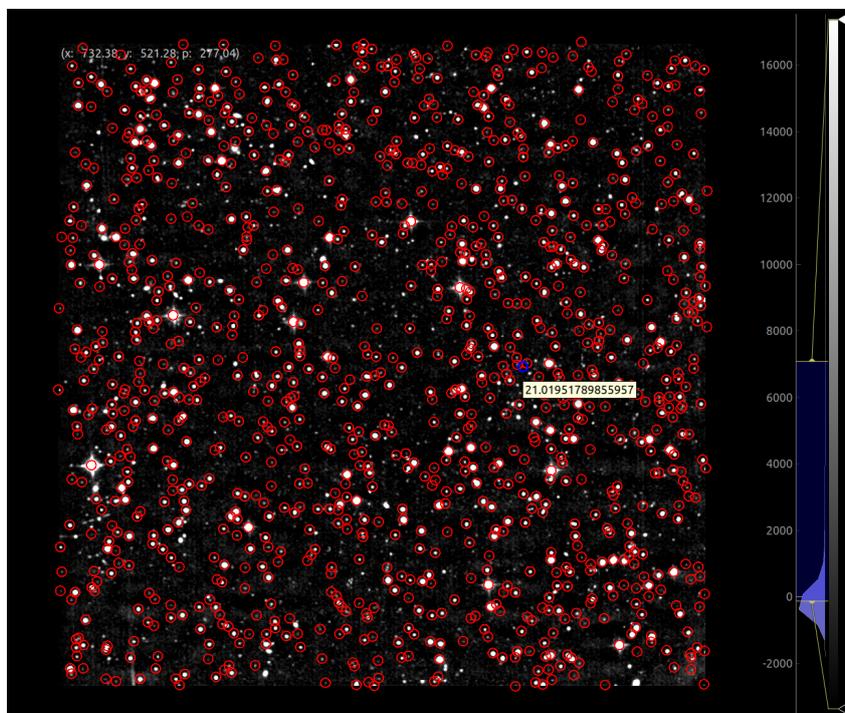


Figure 6: 2020 SO Dataset: Summed Frame with Associated Stars from Gaia DR2.

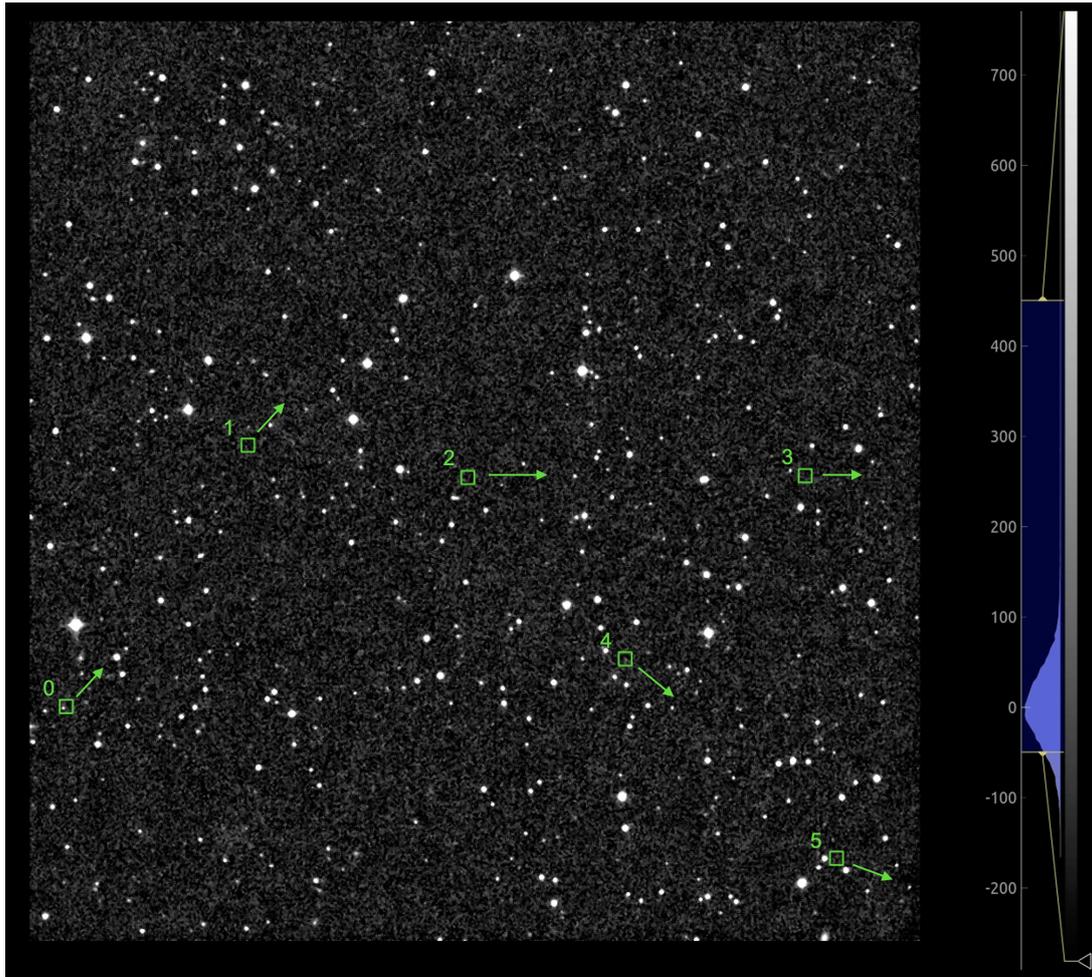


Figure 7: 2020 SO Dataset: Track Estimates from Filter. 2020 SO corresponds to object number 2.

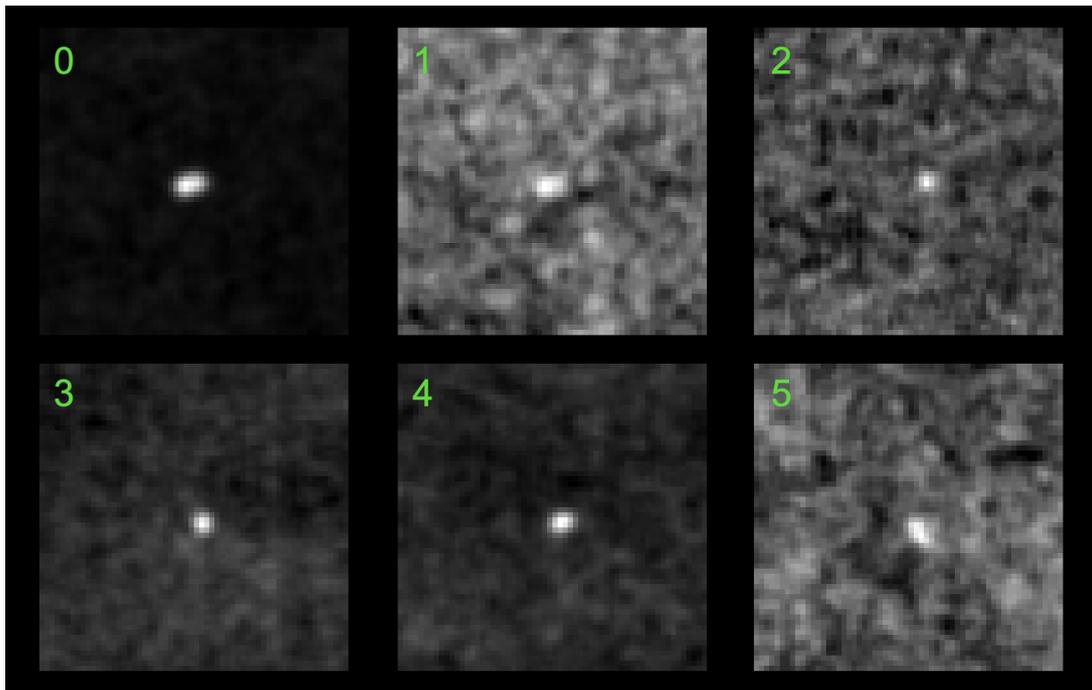


Figure 8: 2020 SO Dataset: Object sub-frames shifted-and-added to confirm detection

5. Conclusion

This paper formulates the detection and tracking problem for low SNR NEOs using Bayesian approach. The measurement likelihood function is leveraged to motivate the use of target likelihood ratio recursion for detecting and tracking dim objects. To overcome the difficulty of jointly estimating the object's intensity along with its kinematic states, a novel intensity-marginalized likelihood ratio is proposed. This allows for simultaneous detection of all low SNR targets in the surveillance region as well as estimate their probability of existence. Any object with an existence probability greater than a chosen threshold are handed-off to an image-based Multi-Bernoulli filter for track maintenance. This proposed method is implemented and tested on real data of Apophis asteroid and 2020 SO. In the 2020 SO dataset, the proposed filter was successfully able to detect 5 objects other than 2020 SO itself⁶. A follow-on journal paper will discuss the exact implementation details for computing the intensity marginalized likelihood function in real-time as well as rigorously compare the performance of the proposed method to other widely used low SNR detection algorithms.

Acknowledgments

This research was funded by the National Defense Science and Engineering Graduate (NDSEG) Fellowship awarded by the Department of Defense (DoD) and Air Force Office of Scientific Research (AFOSR). This work made use of the widely popular open-source code, Astrometry.net, and the authors would like to acknowledge the Astrometry.net Team. Additionally, this work has used data from the European Space Agency (ESA) mission *Gaia*⁷, processed by the *Gaia* Data Processing and Analysis Consortium (DPAC)⁸. Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

References

- [1] P. Vereš, N. Schmidt, Methods, Means and Governance of NEO Observation, in: Planetary Defense, Springer, 2019, pp. 49–70.
- [2] S. S. Blackman, Multiple hypothesis tracking for multiple target tracking, IEEE Aerospace and Electronic Systems Magazine 19 (2004) 5–18.
- [3] D. J. Salmond, H. Birch, A Particle Filter for Track-before-Detect, Proceedings of the American Control Conference 5 (2001) 3755–3760.
- [4] P. S. Gural, J. A. Larsen, A. E. Gleason, Matched filter processing for asteroid detection, The Astronomical Journal 130 (2005) 1951.
- [5] B. N. Vo, B. T. Vo, N. T. Pham, D. Suter, Joint Detection and Estimation of Multiple Objects from Image Observations, IEEE Transactions on Signal Processing 58 (2010) 5129–5141.
- [6] M. Shao, B. Nemati, C. Zhai, S. G. Turyshev, J. Sandhu, G. Hallinan, L. K. Harding, Finding very small near-earth asteroids using synthetic tracking, The Astrophysical Journal 782 (2014) 1.
- [7] S. Virani, M. J. Holzinger, Real-time Multi-target Detection and Tracking of Space Objects using FiSSt Methods, Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii (2020).
- [8] R. P. Mahler, Statistical Multisource-Multitarget Information Fusion, Artech House, Inc., 2007.
- [9] S. Virani, T. Murphy, M. J. Holzinger, B. A. Jones, Empirical Dynamic Data Driven Detection and Tracking Using Detectionless and Traditional FiSSt Methods, Advanced Maui Optical and Space Surveillance Technologies Conference (2017).
- [10] P. B. Stetson, Daophot: A computer program for crowded-field stellar photometry, Publications of the Astronomical Society of the Pacific 99 (1987) 191.
- [11] L. Stone, T. L. Corwin, C. A. Barlow, Bayesian Multiple Target Tracking.
- [12] T. Prusti, J. De Bruijne, A. G. Brown, A. Vallenari, C. Babusiaux, C. Bailer-Jones, U. Bastian, M. Biermann, D. Evans, L. Eyer, et al., The Gaia Mission, Astronomy & Astrophysics 595 (2016) A1.
- [13] A. Brown, A. Vallenari, T. Prusti, J. De Bruijne, C. Babusiaux, C. Bailer-Jones, M. Biermann, D. W. Evans, L. Eyer, F. Jansen, et al., Gaia Data Release 2-Summary of the contents and survey properties, Astronomy & astrophysics 616 (2018) A1.
- [14] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, S. Roweis, Astrometry.net: Blind astrometric calibration of arbitrary astronomical images, The astronomical journal 139 (2010) 1782.

⁶A video of the results can be accessed at shez.space/pdc2021

⁷<https://www.cosmos.esa.int/gaia>

⁸<https://www.cosmos.esa.int/web/gaia/dpac/consortium>