JOINT MULTI-OBJECT AND CLUTTER RATE ESTIMATION WITH THE SINGLE-CLUSTER PHD FILTER

Isabel Schlangen, Vibhav Bharti, Emmanuel Delande, and Daniel E. Clark

School of Electrical and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK
Email: {is117, vb97, E.D.Delande, D.E.Clark}@hw.ac.uk

ABSTRACT

When working with real data, underlying parameters such as the detection or clutter rates are generally unknown and possibly varying over time, however the right parametrisation is crucial to extract proper statistics about the monitored objects. In this article, a single cluster Probability Hypothesis Density (PHD) filter is used to jointly estimate the location and number of a set of objects and the clutter rate over time. The algorithm is verified on a simulated scenario designed to emulate the challenging nature of Single-Molecule Localisation Microscopy (SMLM) imaging sequences and demonstrated on a similar scenario with real data.

***Index Terms*—** Clutter estimation, multi-target tracking, single-cluster PHD filter, single-molecule localisation microscopy.

1. INTRODUCTION

Multi-object estimation is important for many applications that require processing sequences of data. Research on this topic has developed since the 1970s, primarily motivated by aerospace applications. Earlier works include Multiple Hypothesis Tracking (MHT) [1] and Joint Probabilistic Data Association (JPDA) [2], and more recently, several Random Finite Set (RFS) approaches have been derived, primarily the PHD [3–5] and the Cardinalized PHD (CPHD) filters [6]. These techniques have lately been developed for a much wider range of applications including radar and sonar [7], tracking vehicle clusters [8] or even dolphin chirps [9].

Accurate clutter models are of particular importance for applications that are subject to high amounts of time-varying noise or to low background-foreground contrast, such as live-cell SMLM imaging which typically involves short-lived objects on complex background. Widely used tracking approaches for such applications are mostly heuristic [10–12], however those methods cannot cope with missed detections or false alarms. SMLM data is classically analysed in two steps: first, the molecules of interest are localised in the image frame using an image-based spot detection algorithm [13]. Based on these image coordinates per frame, the data is linked across time using a suitable multi-target tracking method. The first step is left out of the scope of this paper and the focus is on data linking.

Recent works have explored RFS methods for clutter estimation. In [14], the mean number of measurements is used for the clutter rate and the spatial clutter distribution is created fitting a Gaussian mixture onto the data; this method heavily relies on the constance of the clutter distribution and on the number of targets being significantly less than the number of false alarms. In [15], the target population is divided into actual targets and clutter generators and both estimated according to their own model. In [16], a similar concept of tracking clutter generators with a PHD filter is used; they also propose an Expectation-Maximisation approach to fit a Gaussian mixture on a non-homogeneous clutter distribution. A PHD filter with a negative binomial clutter model has been introduced in [17]. It can account for clutter models whose variance in the number of false alarms may be significantly higher than the mean, however the clutter parameters are not estimated but assumed given.

This paper presents a new approach for clutter estimation using the single-cluster PHD filter [18, 19]. The latter is already used in various applications such as camera calibration [20], Simultaneous Localisation And Mapping (SLAM) [21], telescope drift correction [22], or microscope drift estimation [23]. Similar approaches have been developed in [24] using different likelihood functions, and [25] involves explicit data association. Section 2 provides the formulation of the method, stating the PHD recursion together with the likelihood function and the variance of the filter update. In Section 3, the filter is tested on simulated and real data.

2. METHODOLOGY

2.1. The single-cluster PHD filter

The single cluster PHD filter which is the base of the proposed algorithm has been studied in [18]. The underlying idea is to

Isabel Schlangen is supported by the Edinburgh Super-Resolution Imaging Consortium (MR/K01563X/1). This work is supported by the MoD UDRC Phase 2 (EP/K014227/1). The authors thank Colin Rickman at Heriot-Watt University for providing the Photo-Activated Localisation Microscopy (PALM) images.

regard the population of objects as a single group that shows distinct global behaviour based on external influences on the sensor. Two common applications are SLAM [21], where the location of a moving sensor is determined using the features surrounding it, or camera calibration [20], where objects observed by two or more sensors give information about the relative position of the sensors. In this article, the single-cluster approach is used to estimate the false alarm rate from a series of SMLM images, assuming that point-like detections have been previously extracted from the raw data [13]. Note that several parameters could be estimated simultaneously within this framework (e.g. the probability of detection), leading to more complex algorithms. For the sake of simplicity, only the false alarm rate is considered in this paper.

In the following, let $\mathcal{S}$ be the sensor state space which describes the unknown clutter parameter, i.e. the false alarm rate. At time k, the multi-target configuration is represented by $X_k = \{x_k^1, \dots, x_k^{n_k}\} \subseteq \mathcal{X}$ where $\mathcal{X}$ is the target state space and n_k is the number of targets. Likewise, the collected observations at time k are represented by $Z_k = \{z_k^1, \dots, z_k^{m_k}\} \subseteq \mathcal{Z}$ where $\mathcal{Z}$ is the measurement space and m_k the number of measurements. The number of false alarms is described by some clutter parameter λ_k at time k (to be defined later). It is assumed that the clutter is generated by the sensor independently of the multi-target configuration so that the multi-target state and the clutter parameter can be jointly estimated with the joint distribution

$$p(X_k, \lambda_k | Z_{1:k}) = p(X_k | Z_{1:k}, \lambda_k) p(\lambda_k | Z_{1:k}). \tag{1}$$

The aim is to estimate two different kinds of uncertainty, namely the clutter parameter and the number and location of the targets. Since the object state estimation is dependent on the sensor output which is in turn affected by noise, the proposed algorithm can be formulated conveniently as a hierarchical structure of two random processes. The *parent process* estimates the clutter parameter λ_k while, conditioned on the clutter parameter, the *daughter process* estimates the multi-target state through a PHD filter.

2.2. The parent process: clutter estimation

The clutter parameter is assumed to be time-varying and evolving through some Markov transition function $t^{\mathcal{S}}_{k|k-1}$. The likelihood of collecting sensor observation Z at time k, given a clutter parameter λ, is described by the *multi-object likelihood* $\ell_k(Z|\lambda)$ which depends on the multi-object configuration estimated by the daughter process. With this, the parent process is described by the following Bayes recursion:

$$p_{k|k-1}(\lambda) = \int_{\mathcal{S}} t^{\mathcal{S}}_{k|k-1}(\lambda|\hat{\lambda}) p_{k-1}(\hat{\lambda}) \mathrm{d}\hat{\lambda}, \tag{2}$$

$$p_k(\lambda|Z_k) = \frac{\ell_k(Z_k|\lambda) p_{k|k-1}(\lambda)}{\int_{\mathcal{S}} \ell_k(Z_k|\hat{\lambda}) p_{k|k-1}(\hat{\lambda}) \mathrm{d}\hat{\lambda}}. \tag{3}$$

2.3. The daughter process: multi-target estimation

The PHD filter does not propagate the full distribution $p(X_k|Z_{1:k}, \lambda_k)$ but only the density of its first-order moment measure – or intensity, or Probability Hypothesis Density – namely $\mu(x_k|Z_{1:k}, \lambda_k)$. The intensity μ, integrated over an arbitrary region $B \subseteq \mathcal{X}$, gives the expected number of targets within B [3].

Let $p_{\mathrm{s},k}(x)$ denote the state-dependent probability of survival at a given time k and denote by $t_{k|k-1}$ the Markov transition of the target states from time $k-1$ to time k. The spontaneous birth of new targets at time k, assumed independent from the existing targets, is modelled with a Poisson point process with intensity $\mu_{\mathrm{b},k}$. Similarly, let us write $p_{\mathrm{d},k}(x)$ for the state-dependent probability of detection, and denote by $l_k(z|x)$ the single-target association likelihood of measurement z with target x at time k. The clutter process is assumed Poisson with intensity $\lambda_k s_{\mathrm{c},k}(z)$, where $s_{\mathrm{c},k}$ is the spatial distribution of the false alarms. The parameter λ_k, estimated by the parent process, is thus the average number of clutter points in the current time scan. In [17], a negative binomial clutter model was assumed instead.

With this, the prediction and update equations of the PHD filter, conditoned on some clutter rate λ, are given by [3]

$$\begin{aligned}\mu_{k|k-1}(x|\lambda) &= \mu_{\mathrm{b},k}(x) + \int_{\mathcal{X}} p_{\mathrm{s},k}(\hat{x}) t_{k|k-1}(x|\hat{x}) \mu_{k-1}(\hat{x}|\lambda) \mathrm{d}\hat{x}, \quad (4)\\ \mu_k(x|Z_k, \lambda) &= \mu_k^{\phi}(x|\lambda) + \sum_{z \in Z_k} \frac{\mu_k^z(x|\lambda)}{\lambda s_{\mathrm{c},k}(z) + \int_{\mathcal{X}} \mu_k^z(\hat{x}|\lambda) \mathrm{d}\hat{x}}, \quad (5)\end{aligned}$$

with missed detection and association terms

$$\mu_k^{\phi}(x|\lambda) = (1 - p_{\mathrm{d},k}(x)) \mu_{k|k-1}(x|\lambda), \tag{6}$$

$$\mu_k^z(x|\lambda) = p_{\mathrm{d},k}(x) l_k(z|x) \mu_{k|k-1}(x|\lambda), \tag{7}$$

for any measurement $z \in Z$.

Even though the PHD filter does not propagate higher-order moments on the target process describing the multi-target population, the *variance* of the updated target process-can be computed at any time step k as [26]

$$\begin{aligned}\mathrm{var}_k(B|Z_k, \lambda) &= \int_B \mu_k^{\phi}(x|\lambda) \mathrm{d}x + \sum_{z \in Z_k} \frac{\int_B \mu_k^z(x|\lambda) \mathrm{d}x}{\lambda s_{\mathrm{c},k}(z) + \int_{\mathcal{X}} \mu_k^z(x|\lambda) \mathrm{d}x} \quad (8)\\ &\quad \cdot \left(1 - \frac{\int_B \mu_k^z(x|\lambda) \mathrm{d}x}{\lambda s_{\mathrm{c},k}(z) + \int_{\mathcal{X}} \mu_k^z(x|\lambda) \mathrm{d}x}\right),\end{aligned}$$

in any region $B \subseteq \mathcal{X}$ of the state space. The statistics $\mu_k(B|\lambda)$ yields the mean value of the estimated number of target within B, while the statistics $\mathrm{var}_k(B|\lambda)$ gives the associated variance; both quantities will be used in the experiments to assess the accuracy of the PHD filter.

Finally, when the daughter process is a PHD filter, the multi-object likelihood function ℓ_k, used in the parent process update equation (3), is given by [18, 19, 21]

$$\ell_k(Z|\lambda) = \exp\left[-\lambda - \int_{\mathcal{X}} p_{\mathrm{d},k}(x)\mu_{k|k-1}(x|\lambda)\mathrm{d}x\right] \cdot \prod_{z\in Z}\left[\lambda s_{\mathrm{c},k}(z) + \int_{\mathcal{X}} \mu_k^z(x|\lambda)\mathrm{d}x\right]. \tag{9}$$

3. EXPERIMENTS

The experiments are set in the context of SMLM where a set of molecules of interest is observed through an optical microscope. Each target is described by its position and velocity coordinates in the image frame. From a filtering perspective, the observation process combines the image acquisition and the feature extraction in the manner of [13] so that the sensor observtion provides information on the position coordinates of the targets.

The parent process is implemented with a Sequential Monte Carlo (SMC) or particle filter following [18]. We further assume that the clutter parameter is constant over time, and thus the Markov kernel $t^{\mathcal{S}}_{k|k-1}$ is set to the identity function. Therefore, the Monte Carlo (MC) particles do not evolve and are initialised equidistantly on a range of admissible values for the clutter parameter λ. The daughter process is implemented with a Gaussian Mixture (GM) PHD filter [4].

3.1. Simulated data

The proposed method is first validated on simulated data. The following parameters are chosen arbitrarily without loss of generality since the same are used for simulation and filtering.

The target state space corresponds to a $50\,\mu\mathrm{m}$-wide square image frame. The target states evolve according to a Near-Constant Velocity (NCV) model, i.e. the object motion is subject to an acceleration noise with standard deviation $0.1\,\mu\mathrm{m\,s^{-2}}$ on both axes. The average number of newborn targets is set to 0.5 and their initial state follows a Gaussian distribution, centred on the image frame with no velocity with standard deviation of $25\,\mu\mathrm{m}$ and $0.5\,\mu\mathrm{m\,s^{-1}}$ on the position and velocity components, respectively. The probabilities of survival $p_{\mathrm{s},k}$ and detection $p_{\mathrm{d},k}$ are set to the constants 0.98 and 0.99, respectively. Each detected target produces a measurement, composed of the two position coordinates corrupted with Gaussian white noise with standard deviation $0.5\,\mu\mathrm{m}$ on both axes. The clutter rate λ is set to 10. The set of observations Z_k at time k is directly generated from the modelling parameters. The parent process is initialised with 100 MC particles, with equal weights and evenly spread on the interval $[1, 50]$, covering the admissible values for λ.

Fig. 1a depicts the estimation of the number of targets and the clutter rate across the scenario, averaged over 100 MC runs. We see that estimated clutter rate converges rapidly to the true value 10. In addition, the true number of targets stays within a 2σ-confidence region around the mean estimated value, where the standard deviation σ is given by the filter using Eq. (8). This suggests that the PHD filter is accurate and not overconfident in the estimation of the number of targets. Fig. 1b shows a typical output of the filter at time $k = 60$. It shows that at this stage in the scenario the probability mass function of the estimated clutter parameter λ is concentrated around the true value 10, consistent with the evolution of the estimated λ shown in Fig. 1a. The estimated number of targets is also close to the true value 14; as expected, a daughter process conditioned on a higher clutter rate yields a lower estimated number of targets since the PHD filter expects a higher number of false alarms among the collected measurements.

3.2. Real data

The real data was generated using PALM with Total Internal Reflection Fluorescence (TIRF), acquired on an Olympus Cell Excellence wide-field microscope fitted with a 512 px by 512 px EMCCD camera. The sample shows SNAP25 proteins labelled with PA-mCherry in human embryonic kidney cells. Fiducial markers are present in the image, generated by gold beads embedded in the cover slip. The pixel width is $106\,\mathrm{nm}$, resulting in a square field of view of $54.272\,\mu\mathrm{m}$ width. The images were captured at a 16.6 Hz sampling rate which corresponds to an exposure time of $60\,\mathrm{ms}$. One hundred frames were used and measurements were extracted by applying the à trous wavelet transform and finding the centroids of the enhanced blobs [13]. We cannot assess the output of the filter through a direct comparison with the ground truth since the latter is unavailable. Instead, we analyse the correlation between the variations in pixel intensity in the image and the variations in clutter intensity in the observation space. For this purpose, the image sequence was processed in a whole and cropped to a square subframe with a width of 200 px which corresponds to $21.2\,\mu\mathrm{m}$ (see Fig. 2).

For filtering purposes, the target motion model is described with a NCV model as in the previous section; because the SNAP25 proteins are expected to move very little, the acceleration noise is set to $1.06\,\mathrm{nm\,s^{-2}}$. The probability of survival is set to 0.95. The average number of newborn targets is set to 5 for the cropped frame and 20 for the whole frame. Their initial state follows a Gaussian distribution, centred on the image frame and with no velocity, with standard deviation of $10.6\,\mu\mathrm{m}$ and $27.136\,\mu\mathrm{m}$ on the position components for the cropped and whole frame, respectively, and $5.3\,\mathrm{nm\,s^{-1}}$ on the velocity components. The probability of detection is set to 0.95 and the measurement noise is $212\,\mathrm{nm}$ on both components. Because we have little knowledge on the true clutter rate, the parent process is initialised with 200 MC particles with equal weights and evenly spread on the interval $[1, 100]$, covering a wider range of admissible values for λ.

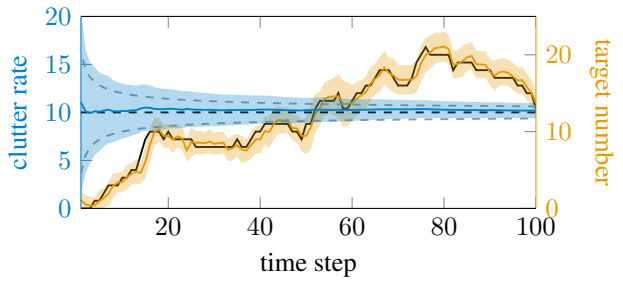

(a) Filter output over the whole scenario, averaged over 100 MC runs. The red plot shows the mean (—) and two standard deviations (■) of the estimated number of targets against the true number of targets (—). The blue plot shows the mean (—) and two standard deviations (■) of the estimated clutter rate against the true clutter rate (- - -). The dashed grey line shows the corresponding two standard deviation confidence region for the estimation of a Poisson distribution, computed from the Cramér-Rao Lower Bound.

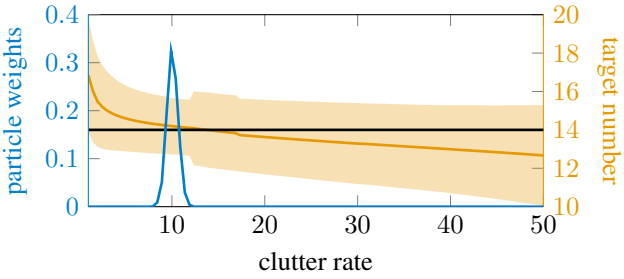

(b) Filter output at time $k = 60$. The red plot shows the mean (—) and standard deviation (■) of the estimated number of targets against the true number of targets (—). The blue plot shows the probability mass function of the estimated clutter rate (true clutter rate: 10).

Fig. 1: Simulation results.

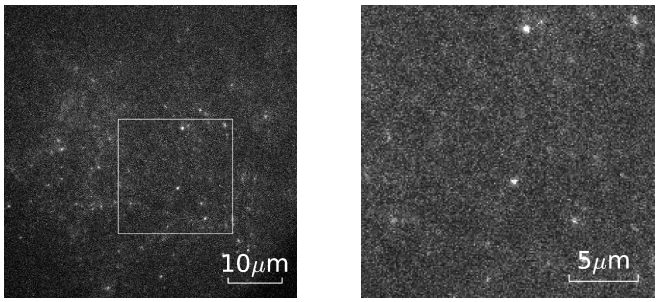

(a) The whole frame. (b) The cropped frame.

Fig. 2: Real data: PALM images showing SNAP25 activity.

Fig. 3 shows the results for the whole and the cropped frame. The fluctuating intensity μ_k illustrates the short lifespan of the molecules, rendering the discrimination between short-lived targets and false alarms challenging. The estimated clutter rate in the whole frame converges to 42.4 which corresponds to 0.014 false alarms per μm^2. On the other hand, the cropped sequence yields a clutter rate of 9.0 which corresponds to 0.020 false alarms per μm^2. This discrepancy is consistent with the localisation of the cropped image, situated in an area with slightly more activity than the image periphery due to illumination settings.

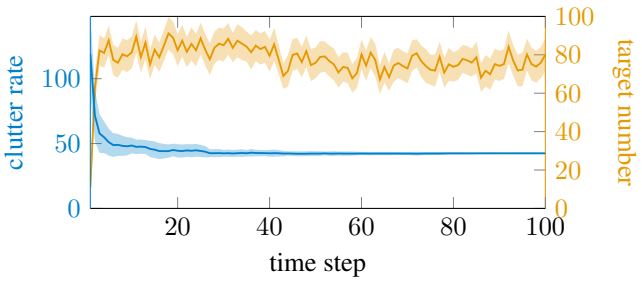

(a) Results for the whole frame.

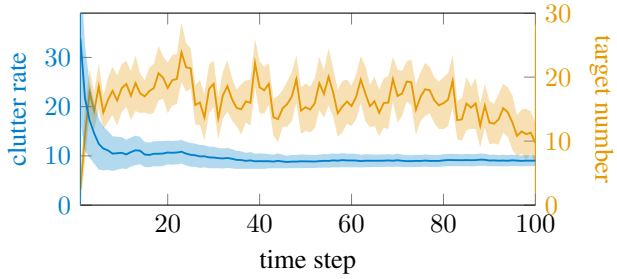

(b) Results for the cropped frame.

Fig. 3: Results for the real data. The red plot shows the mean (—) and standard deviation (■) of the estimated number of targets. The blue plot shows the mean (—) and standard deviation (■) of the estimated clutter rate.

4. CONCLUSION

This paper exploits a single-cluster PHD filter for the joint estimation of the multi-target configuration and the sensor clutter rate in multi-target detection and tracking. The algorithm is tested on simulated and real data in the context of Single-Molecule Localisation Microscopy, in which the number of clutter points is assumed Poisson with unknown rate. The results on simulated data show that the estimated clutter rate rapidly converges to the true value, while the underlying PHD filter provides an accurate estimation of the number of targets in the scene. While ground truth is not available for the assessment of the filter, the variation in the estimated clutter rate in the whole image and a cropped subframe appears consistent with the fluorescence activity.

5. REFERENCES

[1] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Autom. Control*, vol. 24, no. 6, pp. 843–854, Dec. 1979.

[2] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE Journal of Oceanic Engineering*, vol. 8, no. 3, pp. 173–184, July 1983.

[3] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, 2003.

[4] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4091–4104, 2006.

[5] B.-N. Vo, S. S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 4, pp. 1224–1245, 2005.

[6] R. P. S. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 4, pp. 1523–1543, Oct. 2007.

[7] W. Wu, J. Jiang, W. Liu, X. Feng, L. Gao, and X. Qin, "Augmented state GM-PHD filter with registration errors for multi-target tracking by Doppler radars," *Signal Processing*, vol. 120, pp. 117–128, 2016.

[8] S. A. Goli, B. H. Far, and A. O. Fapojuwo, "An accurate multi-sensor multi-target localization method for cooperating vehicles," in *Theoretical Information Reuse and Integration*, pp. 197–217. Springer, 2016.

[9] P. Gruden and P. R. White, "Automated tracking of dolphin whistles using Gaussian Mixture Probability Hypothesis Density (GM-PHD) filters," *J. Acoust. Soc. Am.*, pp. 1–47, Sept. 2016.

[10] I. K. Sethi and R. Jain, "Finding trajectories of feature points in a monocular image sequence," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 1, pp. 56–73, Jan. 1987.

[11] C. J. Veenman, M. J. T. Reinders, and E. Backer, "Resolving motion correspondence for densely moving points," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 1, pp. 54–72, Jan. 2001.

[12] E. Meijering, I. Smal, and G. Danuser, "Tracking in molecular bioimaging," *IEEE Signal Process. Mag.*, vol. 23, no. 3, pp. 46–53, 2006.

[13] I. Smal, M. Loog, W. Niessen, and E. Meijering, "Quantitative comparison of spot detection methods in fluorescence microscopy," *IEEE Trans. Med. Imag.*, vol. 29, no. 2, pp. 282–301, 2010.

[14] F. Lian, C. Han, and W. Liu, "Estimating Unknown Clutter Intensity for PHD Filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 4, pp. 2066–2078, Oct. 2010.

[15] R. P. S. Mahler, B.-T. Vo, and B.-N. Vo, "CPHD filtering with unknown clutter rate and detection profile," *IEEE Trans. Signal Process.*, vol. 59, no. 8, pp. 3497–3513, 2011.

[16] X. Chen, R. Tharmarasa, M. Pelletier, and T. Kirubarajan, "Integrated clutter estimation and target tracking using Poisson point processes," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 48, no. 2, pp. 1210–1235, Apr. 2012.

[17] I. Schlangen, E. Delande, J. Houssineau, and D. E Clark, "A PHD filter with negative binomial clutter," in *Information Fusion, 2016 19th International Conference on*. ISIF, 2016, pp. 658–665.

[18] A. Swain and D. E. Clark, "The single-group PHD filter: An analytic solution," in *Information Fusion, 2011 14th International Conference on*, 2011, pp. 1–8.

[19] B. Ristic, D. E. Clark, and N. Gordon, "Calibration of multi-target tracking algorithms using non-cooperative targets," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 390–398, 2013.

[20] J. Houssineau, D. E. Clark, S. Ivekovic, C.-S. Lee, and J. Franco, "A unified approach for multi-object triangulation, tracking and camera calibration," *IEEE Trans. Signal Process.*, vol. 64, no. 11, pp. 2934–2948, 2016.

[21] C. S. Lee, D. E. Clark, and J. Salvi, "SLAM with dynamic targets via single-cluster PHD filtering," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 543–552, 2013.

[22] O. Hagen, J. Houssineau, I. Schlangen, E. D. Delande, J. Franco, and D. E. Clark, "Joint estimation of telescope drift and space object tracking," in *Aerospace Conference, 2016 IEEE*. IEEE, 2016, pp. 1–10.

[23] I. Schlangen, J. Franco, J. Houssineau, W. T. E. Pitkeathly, D. E. Clark, I. Smal, and C. Rickman, "Marker-less stage drift correction in super-resolution microscopy using the single-cluster PHD filter," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 1, pp. 193–202, Feb. 2016.

[24] J. Mullane, B. N. Vo, M. D. Adams, and B. T. Vo, "A random-finite-set approach to Bayesian SLAM," *IEEE Trans. Robot.*, vol. 27, no. 2, pp. 268–282, Apr. 2011.

[25] E. Brekke, B. Kalyan, and M. Chitre, "A novel formulation of the Bayes recursion for single-cluster filtering," in *2014 IEEE Aerospace Conference*, Mar. 2014, pp. 1–16.

[26] E. D. Delande, M. Üney, J. Houssineau, and D. E. Clark, "Regional variance for multi-object filtering," *IEEE Trans. Signal Process.*, vol. 62, no. 13, pp. 3415–3428, July 2014.