

MMOSPA-based Track Extraction in the PHD filter – A Justification for k -Means Clustering

Marcus Baum*, Peter Willett*, and Uwe D. Hanebeck[§]

*Department of Electrical and Computer Engineering, University of Connecticut, USA.

[§] ISAS Laboratory, Karlsruhe Institute of Technology (KIT), Germany.

Email: baum@engineer.uconn.edu, willett@engr.uconn.edu, uwe.hanebeck@ieee.org

Abstract—Displaying tracks is an essential part of a multi-target tracking system. Recently, it was proposed to extract tracks with respect to the *Optimal Sub-Pattern Assignment (OSPA)* metric, i.e., the traditionally used squared error loss is replaced with an OSPA loss, which leads to the so-called *Minimum Mean OSPA (MMOSPA)* estimate. So far, work concentrated on traditional trackers that maintain probability densities for the targets. In this paper, we aim at extracting the MMOSPA estimate from a *Probability Hypothesis Density (PHD)* as used within the PHD filter. We elaborate that the PHD in general does not contain enough information to determine the exact MMOSPA estimate. However, we then show that if the loss function has a specific form, it is indeed possible to extract point estimates from a PHD that are optimal w.r.t. the underlying unknown random finite set. We discuss two specific loss functions that fulfill this condition and are potentially close to the OSPA loss, a nearest neighbor loss and a kernel distance loss. It turns out that track extraction based on the nearest neighbor loss can be performed with the well-known k -means algorithm. Simulations show when the estimates based on the nearest neighbor and the kernel loss are close to the MMOSPA estimate.

I. INTRODUCTION

The *Probability Hypothesis Density (PHD)* filter [15] is a popular multi-target tracking algorithm due to its low computational complexity and its capability to estimate the number of targets. As indicated by its name, the PHD filter is a recursive updating scheme for the first-order moment of a *Random Finite Set (RFS)*, the so-called PHD function. Intuitively, the integral of the PHD function over a specific area corresponds to the expected number of targets in that area. Hence, in contrast to classical multi-target trackers such as the *Joint Probabilistic Data Association Filter (JPDAF)* [1], the estimated tracks are not explicitly available from the PHD. In order to obtain reasonable track estimates, clustering techniques such as k -means or *Expectation Maximization (EM)* have been proposed [6], [14], [23], [24] in literature. However, it is a common opinion that these clustering techniques are heuristics [18], [19]. In order to overcome the shortcomings of k -means and EM, several alternatives have been proposed such as CLEAN [22], [24], mean shift [21], and [5], [16], [18], [19], [25].

In classical multi-target trackers such as JPDAF, tracks are directly given by the means of the probability density functions for the targets. However, it recently turned out

that the mean may result in undesired coalescence effects in the case of closely-spaced targets. In order to prevent these coalescence effects, it was proposed to extract tracks with respect to the *Optimal Sub-Pattern Assignment (OSPA)* [20], which is the standard metric for evaluating multi-target trackers. In this manner, multi-target trackers that optimized according to the same criterion that is used for evaluation can be developed [3], [7]–[10].

This paper is built upon the idea to extract MMOSPA estimates just from PHDs in case the underlying *Random Finite Set (RFS)* density is unknown. We will first show that the PHD in general does not contain enough information to determine the exact MMOSPA estimate. Subsequently, we discuss two ad-hoc approaches for reconstructing an approximate MMOSPA estimate. Finally, we present a systematic means to extract point estimates from PHDs. We show that if the loss function has a specific form it is indeed possible to extract a point estimate from a PHD that is optimal w.r.t. the underlying (unknown) RFS. We discuss two specific loss functions that fulfill this condition, i.e., a nearest neighbor loss and a kernel distance loss. Surprisingly, we will show that the point estimate based on the nearest neighbor loss can be obtained by means of the k -means algorithm. In this way, k -means should no longer be considered as a heuristic approach as it is optimal w.r.t. a reasonable loss function and hence, is mathematically well-grounded. The shortcomings of k -means reported in the literature can be attributed to the fact that the nearest neighbor loss is in general a poor approximation of the OSPA loss.

Remark 1. Throughout the entire paper, we assume that the number of targets is given. The results are applicable to a PHD filter for an unknown number of targets by a two step procedure: (1) Estimate the target number using the *Expected a Posteriori (EAP)* criterion [15], i.e., the mass under the PHD. (2) Use the techniques for given number of tracks as introduced in this paper in order to extract the tracks.

II. BACKGROUND: RANDOM SETS AND PHD FILTER

Traditional multi-target tracking algorithms such as the *Joint Probabilistic Data Association Filter (JPDAF)* [1] stack the state vectors of the targets at time k in a single

joint vector

$$\underline{x}(k) = [\underline{x}_1(k)^T, \dots, \underline{x}_n(k)^T]^T, \quad (1)$$

where $\underline{x}_i(k) \in \mathbb{R}^d$ denotes the state vector of target $i \in \{1, \dots, n\}$ and n is the number of targets.

The tracking algorithm maintains a probability density function for the joint state vector $\underline{x}(k)$ conditioned on the available measurements

$$p(\underline{x}(k) | \mathcal{Y}(k)), \quad (2)$$

where $\mathcal{Y}(k)$ denotes all available measurement data up to time k .

Remark 2. Note that the standard JPDAF assumes that the targets are independent. There is also an extension – the coupled JPDAF [1] – that works with the complete joint density.

The joint state vector (1) inherently imposes a labeling on the targets, i.e., it specifies which target is $\underline{x}_1(k)$ etc. A representation that does not incorporate labeling is given by a set

$$X(k) = \{\underline{x}_1(k), \dots, \underline{x}_n(k)\} \in \mathcal{S}_n,$$

where $\mathcal{S}_n$ denotes the set of all subsets of $\mathbb{R}^d$ with n elements. A probability distribution over sets with finite elements can be specified by a so-called *Random Finite Set (RFS)* density [15]. If the number of targets is known to be n (as in our case) and a joint density (2) for the targets is given, the corresponding RFS density is

$$f(\{\underline{x}_1(k), \dots, \underline{x}_n(k)\} | \mathcal{Y}(k)) = \sum_{\pi \in \Pi_n} p(P_\pi(\underline{x}(k)) | \mathcal{Y}(k)), \quad (3)$$

where Π_n denotes all permutations of the set $\{1, \dots, n\}$, and $P_\pi(\underline{x}(k))$ permutes the single target states in $\underline{x}(k)$ according to π , i.e., $P_\pi(\underline{x}(k)) = [\underline{x}_{\pi(1)}(k)^T, \dots, \underline{x}_{\pi(n)}(k)^T]^T$. It is not possible to reconstruct the joint density based on the RFS density as the labeling information is not available.

The *Probability Hypothesis Density (PHD)* filter [15] recursively maintains the PHD, which results from (3) according to

$$\begin{aligned} D(\underline{z}(k) | \mathcal{Y}(k)) &= \int_{\tilde{X} \in \mathcal{S}_{n-1}} f(\underline{z}(k) \cup \tilde{X}(k)) d\tilde{X}(k) = \\ &= \frac{1}{n-1} \int f(\{\underline{z}(k), \underline{x}_2(k), \dots, \underline{x}_n(k)\} | \mathcal{Y}(k)) \\ &\quad d\underline{z}(k), d\underline{x}_2(k), \dots, d\underline{x}_n(k) \end{aligned} \quad (4)$$

with $\tilde{X}(k) = \{\underline{x}_2(k), \dots, \underline{x}_n(k)\}$.

Above, we used the definition of the set integral [15] for sets with a fixed number of elements

$$\int_{X \in \mathcal{S}_n} f(X) dX = \frac{1}{n!} \int f(\{\underline{x}_1(k), \underline{x}_2(k), \dots, \underline{x}_n(k)\}) d\underline{x}_1(k), d\underline{x}_2(k), \dots, d\underline{x}_n(k). \quad (5)$$

The PHD (4) is the first moment of the RFS density (3). In general, the RFS density (3) cannot be reconstructed from

the PHD (4). Note that the integral over the PHD gives the (expected) number of targets, i.e., $n = \int D(\underline{z}(k)) d\underline{z}(k)$.

Remark 3. In this paper, we concentrate on a single point in time. Hence, we can make the following notational simplification:

- The time index can be omitted, i.e., instead of $\underline{z}(k)$, we write $\underline{z}$.
- The conditioning on the available data $\mathcal{Y}(k)$ is not explicitly written, i.e., instead of $p(\underline{x}(k) | \mathcal{Y}(k))$, we write $p(\underline{x})$.

III. TRACK EXTRACTION FROM AN RFS DENSITY

In this section, we address the question of how to extract point estimates, i.e., tracks, from an RFS density (3). In Bayesian estimation, point estimates are chosen to minimize the expected loss with respect to a specific loss function as described in the following definition.

Definition 1 (Point Estimate). A point estimate $\hat{X}^l \in \mathcal{S}_n$ for the RFS density $f(X)$ in (3) minimizes the expected loss

$$\hat{X}^l := \arg \min_{\hat{X} \in \mathcal{S}_n} E\{l(\hat{X}, X)\}_{f(X)}, \quad (6)$$

where $l : \mathcal{S}_n \times \mathcal{S}_n \rightarrow \mathbb{R}$ is a loss function that measures how “close” $\hat{X}$ is to X .

As Definition 1 deals with sets, the squared error loss for vectors usually used in Bayesian estimation is not suitable. In this context, [12] proposed to extract point estimates with respect to the (squared) *Optimal Sub-Pattern Assignment (OSPA)* [20] distance, which is a distance metric on sets.

Definition 2 (OSPA). The *Optimal Sub-Pattern Assignment (OSPA)* [20] distance between the two sets $X = \{\underline{x}_1^T, \dots, \underline{x}_n^T\} \in \mathcal{S}_n$ and $\hat{X} = \{\hat{\underline{x}}_1^T, \dots, \hat{\underline{x}}_n^T\} \in \mathcal{S}_n$ is defined as

$$\text{OSPA}(\hat{X}, X)^2 = \frac{1}{n} \min_{\pi \in \Pi_n} \sum_{i=1}^n \|\underline{x}_i - \hat{\underline{x}}_{\pi(i)}\|^2. \quad (7)$$

The above (squared) OSPA distance can directly be used as the loss function in Definition 1: $l(\hat{X}, X) := \text{OSPA}(\hat{X}, X)^2$. The point estimate with respect to OSPA distance $\hat{X}^{\text{OSPA}}$ is the *Minimum Mean OSPA (MMOSPA)* estimate.

Remark 4. In [12], it was shown that the MMOSPA estimate coincides for all joint density functions (2) that specify the same RFS density (3). In this manner, the MMOSPA estimate can be uniquely defined for a joint density.

IV. PROBLEM FORMULATION

In this work, we consider the problem of extracting the MMOSPA estimate from the PHD (4). From our point of view, the PHD $D(\underline{z})$ is given, e.g., obtained from the PHD filter recursions. However, the true underlying RFS density $f(X)$ is unknown. In this way, the PHD gives us a hint on how the true RFS density $f(X)$ might look.

Ultimately, we are interested in $\hat{X}^{\text{OSPA}}$ from $f(X)$. The question is, can we get $\hat{X}^{\text{OSPA}}$ from $D(\underline{z})$? And if not, can we get an approximation?

Fig. 1 and Fig. 2 show already the answer to the first question. In general, each different RFS density that is consistent with the PHD gives a different MMOSPA estimate. Too much information is lost in the PHD. Hence, the remainder of this work is concerned with the second question.

V. FIRST ATTEMPT

A reasonable attempt is to say: if we need the true underlying RFS density corresponding to $D(\underline{z})$ but we do not have it, then we just pick a reasonable RFS density that is consistent with $D(\underline{z})$. A natural way is to associate a Poisson RFS density with the PHD. As we assume the number of targets to be given by n , we would obtain

$$f(X) = n! \cdot \prod_{m=1}^n \frac{D(\underline{x}_m)}{n}, \quad (8)$$

which says each target location is given by the probability density $\frac{1}{n}D(\underline{z})$.

Unfortunately, it turns out that the MMOSPA estimate for (8) is useless as it coalesces the targets, see Fig. 2.

VI. SECOND ATTEMPT

In a second attempt, we consider the special case that each target is distributed according to a Gaussian distribution, i.e., the joint density (2) is Gaussian

$$p(\underline{x}) = \mathcal{N}(\underline{x} - \underline{m}, \mathbf{C}) \quad (9)$$

with mean $\underline{m} = [\underline{m}_1^T, \dots, \underline{m}_n^T]^T$ and covariance matrix $\mathbf{C} = \text{diag}(\mathbf{C}_1, \dots, \mathbf{C}_n)$. The corresponding RFS density is

$$f(X) = \sum_{\pi \in \Pi_n} \prod_{i=1}^n \mathcal{N}(\underline{x}_{\pi(i)} - \underline{m}_{\pi(i)}, \mathbf{C}_{\pi(i)}) . \quad (10)$$

In this case, the PHD (4) is the Gaussian mixture density

$$D(\underline{z}) = \sum_{i=1}^n \mathcal{N}(\underline{z} - \underline{m}_i, \mathbf{C}_i) . \quad (11)$$

Based on the PHD (11), $\underline{m}$ and $\mathbf{C}$ can be recovered up to permutation using, e.g., *Expectation Maximization (EM)*. Having reconstructed the RFS density (or the joint density), standard MMOSPA estimation algorithms can be used to extract the MMOSPA estimate, see [12]. Of course, this approach can also be used if the true underlying joint density is not Gaussian. However, all told, only an approximation of the MMOSPA estimate is obtained.

VII. SYSTEMATIC APPROACH

A systematic approach is to choose a loss function that exploits only information about the joint density that is also available in the PHD, so that we can extract the point estimate from the PHD. It is by no means obvious that there exists a meaningful (non-trivial) loss function with this property. In the following, we first derive a condition that loss functions must satisfy for this purpose. Second, we discuss reasonable choices for loss functions, which are supposed to be close to the OSPA loss function.

Theorem 1. *If the loss function in (6) is of the following form*

$$l(\hat{X}, X) = l_1(\hat{X}) + \sum_{i=1}^n l_2(\hat{X}, \underline{x}_i) + l_3(X) , \quad (12)$$

where $l_1 : \mathcal{S}_n \rightarrow \mathbb{R}$, $l_2 : \mathcal{S}_n \times \mathbb{R}^d \rightarrow \mathbb{R}$, and $l_3 : \mathcal{S}_n \rightarrow \mathbb{R}$ are functions, then

$$\hat{X}^l = \arg \min_{\hat{X}} \mathbb{E}\{l(\hat{X}, X)\}_{f(X)} \quad (13)$$

$$= \arg \min_{\hat{X}} l_1(\hat{X}) + \int l_2(\hat{X}, \underline{z}) D(\underline{z}) d\underline{z} , \quad (14)$$

and hence, the point estimate $\hat{X}^l$ can be obtained from the PHD $D(\underline{z})$ no matter what $f(X)$ is.

Proof:

$$\begin{aligned} \mathbb{E}\{l(\hat{X}, X)\}_{f(X)} &= \int l(\hat{X}, X) f(X) dX = \\ &= \int (l_1(\hat{X}) + \sum_{i=1}^n l_2(\hat{X}, \underline{x}_i) + l_3(X)) f(X) dX = \\ &= l_1(\hat{X}) + \sum_{i=1}^n \int l_2(\hat{X}, \underline{x}_i) f(X) dX + c = \\ &= l_1(\hat{X}) + n \int l_2(\hat{X}, \underline{z}) f(X) d\{\underline{z}, \underline{x}_2, \dots, \underline{x}_n\} + c = \\ &= l_1(\hat{X}) + n \frac{1}{n!} \int l_2(\hat{X}, \underline{z}) \\ &\quad f(\{\underline{z}, \underline{x}_2, \dots, \underline{x}_n\}) d\underline{z}, d\underline{x}_2, \dots, d\underline{x}_n + c = \\ &= l_1(\hat{X}) + \int \int l_2(\hat{X}, \underline{z}) \\ &\quad f(\{\underline{z}, \underline{x}_2, \dots, \underline{x}_n\}) d\{\underline{x}_2, \dots, d\underline{x}_n\} d\underline{z} + c = \\ &= l_1(\hat{X}) + \int l_2(\hat{X}, \underline{z}) \int \\ &\quad f(\{\underline{z}, \underline{x}_2, \dots, \underline{x}_n\}) d\{\underline{x}_2, \dots, d\underline{x}_n\} d\underline{z} + c , \quad (15) \end{aligned}$$

where c is a constant independent of $\hat{X}$. ■

The OSPA loss function obviously does not satisfy (12). Hence, the last resort is to modify OSPA (as little as possible) in order to obey (12) while we still hope that it is close to the OSPA. For this purpose, we will discuss two meaningful loss functions in the following.

A. Nearest Neighbor Loss

A “simplified” variant of the OSPA loss that we will call the *Nearest Neighbor (NN)* loss is given by

$$\text{NN}(\hat{X}, X) = \sum_{i=1}^n \min_j \|\hat{\underline{x}}_j - \underline{x}_i\|^2 , \quad (16)$$

which results from the OSPA distance by allowing multiple assignments to the same target, i.e., each target in X is assigned to the nearest target in $\hat{X}$. It is essential to note that (16) is not a metric on *sets* because it is not symmetric,




Fig. 1: Example 1: A PHD for two one-dimensional targets and two possible RFS densities. The MMOSPA estimates for the RFS densities (blue crosses) do not coincide.




Fig. 2: Example 2: A PHD for two one-dimensional targets and two possible RFS densities. The MMOSPA estimates for the RFS densities (blue crosses) are significantly different.

i.e., in general $\text{NN}(\hat{X}, X) \neq \text{NN}(X, \hat{X})$. Also, it is obvious that (16) satisfies (12).

Plugging (16) into (13) results in

$$\hat{X}^{\text{NN}} = \arg \min_{\hat{X}} \int \min_j \|\hat{x}_j - \underline{z}\|^2 D(\underline{z}) d\underline{z} . \quad (17)$$

It can be shown that – in the case of a particle representation of the PHD – the solution of (17) minimizes the so-called *Within-Cluster Sum of Squares (WCSS)*, which is used in the *k*-means clustering algorithm [4], [11].

For this purpose, assume the PHD is represented with N_p equally weighted particles according to

$$D(\underline{z}) = \frac{n}{N_p} \sum_{l=1}^{N_p} \delta(\underline{z} - \underline{z}^{(l)}) , \quad (18)$$

where $\underline{z}^{(l)}$ denotes the l -th particle. In this case (17) becomes

$$\hat{X}^{\text{NN}} = \arg \min_{\hat{X}} \sum_{l=1}^{N_p} \min_j \|\hat{x}_j - \underline{z}^{(l)}\|^2 . \quad (19)$$

If we define the clusters

$$\mathcal{S}_j = \left\{ l \mid j = \arg \min_m \|\hat{x}_m - \underline{z}^{(l)}\|^2 \right\} , \quad (20)$$

(19) can be written as

$$\hat{X}^{\text{NN}} = \arg \min_{\hat{X}} \sum_{j=1}^{N_p} \sum_{l \in \mathcal{S}_j} \|\hat{x}_j - \underline{z}^{(l)}\|^2 , \quad (21)$$

which is exactly the *Within-Cluster Sum of Squares (WCSS)*. As a consequence, the *k*-means algorithm can be used to perform the optimization (17). The findings above give a theoretical justification for using *k*-means for track extraction in the PHD filter. In contrast to what seems to be a common opinion in literature, *k*-means is *not* ad-hoc: it minimizes a reasonable loss function – the *Nearest Neighbor (NN)* loss (16). Of course, it is up to debate how close the NN loss is to the OSPA loss. And most likely the shortcomings of the *k*-means algorithm for track extraction discussed in literature result from the poor approximation quality.

B. Kernel Distance

Another loss function that satisfies (12) is the squared kernel distance defined as follows (see [2] for a discussion of the kernel distance in the context of multi-target tracking).

Definition 3. The (squared) kernel distance between the two sets $X = \{\underline{x}_1^T, \dots, \underline{x}_n^T\} \in \mathcal{S}_n$ and $\hat{X} = \{\hat{\underline{x}}_1^T, \dots, \hat{\underline{x}}_n^T\} \in \mathcal{S}_n$




Fig. 3: Illustration of the (four-dimensional) joint density for two two-dimensional targets used in the evaluation: Intuitively, one target is approximately at location $[-0.1, -0.1]^T$ and the other one is at $[0.4, -0.1]^T$. Also note that the correlations are not visible from the marginal densities. The correlations exclude the case that the two targets are at the same location. We use $\underline{x}_1^{(1)}$ and $\underline{x}_2^{(2)}$ to denote the first and second component of the vector $\underline{x}_1$ (same notation for $\underline{x}_2$).

is defined as

$$\text{Kernel}(\hat{X}, X)^2 = \sum_{i=1}^n \sum_{j=1}^n K^b(\hat{x}_i - \hat{x}_j) - 2 \sum_{i=1}^n \sum_{j=1}^n K^b(\hat{x}_i - \underline{x}_j) + \sum_{i=1}^n \sum_{j=1}^n K^b(\underline{x}_i - \underline{x}_j), \quad (22)$$

where $K^b(\cdot) = \mathcal{N}(\cdot - \underline{0}, b\mathbf{I}_d)$ is (here) a Gaussian kernel function with kernel width b .

The kernel distance is indeed a distance metric on sets. Intuitively, it interprets the set of targets $\hat{X}$ as a continuous functions by placing the kernel at each target location, and then calculates the L_2 distance between the functions.

At the first blush, the kernel loss seems to be superior to the NN loss as it is based on a true distance. However, in general, the kernel distance may significantly differ from the OSPA distance. Hence, if we argue that the goal is to extract MMOSPA estimates, the kernel loss may not necessarily be preferable for extracting point estimates from a PHD. We will investigate this issue in the following section.

Remark 5. In case of known number of targets the OSPA distance coincides with the Wasserstein distance [13]. Both the Wasserstein distance and the kernel distance are well-known and widely used distance measures for sets in various areas such as computer vision [17]. To the authors' knowledge, there are no quantitative results on how well the kernel distance might approximate the Wasserstein distance.

VIII. EVALUATION

In this section, we provide a comparison of the NN and kernel loss function for track extraction. Specifically, we investigate how precise the resulting point estimates from the PHD approximate the true MMOSPA estimate from the

(unknown) multi-target joint density (2). For this purpose, we consider the specific Gaussian mixture multi-target joint density for two two-dimensional targets as illustrated in Fig. 3. This is a typical joint density for two well-separated targets. In our experiment, we decrease the distance between the two targets (meaning the modes of the Gaussian) and investigate how the NN and kernel loss function based estimates behave with respect to the MMOSPA and MMSE estimate (the mean) of the true underlying joint density. Two different kernel loss functions are employed – one with a small kernel width 0.1 and one with a large kernel width 0.5 – in order to find out the influence of the kernel width.

Technically, the point estimates are obtained by sampling the joint density/PHD and performing an explicit optimization. This also means we do not use the k -means algorithm itself, we only optimize the *Within-Cluster Sum of Squares* (WCSS) used by k -means. In this manner, we seek to avoid local minima in the optimization.

The results are shown in Fig. 4. Clearly, the NN loss nearly always coincides with the OSPA loss. The kernel estimate gives also a good approximation to the OSPA estimate in case of a small kernel width. If the kernel width is larger, the approximation quality decreases.

These results should be taken with care as they are only valid for the specific example joint densities; however, they demonstrate that the NN estimate can give a reasonably precise approximation of the MMOSPA estimates. Also, we want to stress that in case of highly multimodal joint densities, the NN estimate may visually be significantly different from the MMOSPA estimate. So far, there is no general quantitative assessment of the approximation quality available.

IX. CONCLUSION AND FUTURE WORK

To date, methods for track extraction in the PHD filter are often used without any strong theoretical justifications. In this work, we presented a systematic and mathematically well-grounded concept for extracting point estimates from a PHD. The underlying idea is to use a loss function that only exploits information from the underlying RFS that is also available in the PHD. Based on this concept, and with the goal of deriving MMOSPA estimates, we were able to provide a theoretical and intuitive justification for a widely-used technique for track extraction – k -means clustering.

Future work will be concerned with finding loss functions that are even closer to the OSPA loss than the NN loss. So far we made the simplifying assumption that the number of targets is given by the EAP estimate. Further improvements might be possible, e.g., for the *Cardinalized Probability Hypothesis Density* (CPHD) filter [15].




Fig. 4: PHD for the joint density in Fig. 3, where the distance between the modes is varied. The mean and MMOSPA estimate of the joint density, the k -means estimate, and the kernel distance estimate (for kernel width 0.2 and 0.4).

ACKNOWLEDGMENTS

Marcus Baum and Peter Willett were supported by the Office of Naval Research under N00014-13-1-0231.

REFERENCES

- [1] Yaakov Bar-Shalom, Peter K. Willett, and Xin Tian. *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.
- [2] Marcus Baum, Patrick Ruoff, Dominik Itte, and Uwe D. Hanebeck. Optimal Point Estimates for Multi-target States based on Kernel Distances. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC 2012)*, Maui, Hawaii, USA, December 2012.

- [3] Marcus Baum, Peter Willett, and Uwe D. Hanebeck. Calculating Some Exact MMOSPA Estimates for Particle Distributions. In *Proceedings of the 15th International Conference on Information Fusion (Fusion 2012)*, Singapore, July 2012.
- [4] P. Berkhin. A Survey of Clustering Data Mining Techniques. In Jacob Kogan, Charles Nicholas, and Marc Teboulle, editors, *Grouping Multidimensional Data*, pages 25–71. Springer Berlin Heidelberg, 2006.
- [5] A.O. Bozdogan, M. Efe, and R. Streit. Reduced Palm Intensity for Track Extraction. In *16th International Conference on Information Fusion (FUSION)*, pages 1243–1250, July 2013.
- [6] D.E. Clark and J. Bell. Multi-Target State Estimation and Track Continuity for the Particle PHD Filter. *IEEE Transactions on Aerospace and Electronic Systems*, 43(4):1441–1453, October 2007.
- [7] D. Crouse, M. Guerriero, P. Willett, R. Streit, and D. Dunham. A Look at the PMHT. In *Proceedings of the 12th International Conference on Information Fusion (Fusion 2009)*, July 2009.
- [8] D. F. Crouse, P. Willett, and Y. Bar-Shalom. Developing a Real-Time Track Display That Operators Do Not Hate. *IEEE Transactions on Signal Processing*, 59(7):3441–3447, July 2011.
- [9] D. F. Crouse, P. Willett, and Y. Bar-Shalom. Generalizations of Blom and Bloem’s PDF Decomposition for Permutation-Invariant Estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3840–3843, May 2011.
- [10] David Frederic Crouse, Peter Willett, Marco Guerriero, and Lennart Svensson. An Approximate Minimum MOSPA Estimator. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3644–3647, May 2011.
- [11] Chris Ding and Xiaofeng He. K-Means Clustering via Principal Component Analysis. In *Proceedings of the Twenty-first International Conference on Machine Learning, ICML ’04*, pages 29–, New York, NY, USA, 2004. ACM.
- [12] Marco Guerriero, Lennart Svensson, Daniel Svensson, and Peter Willett. Shooting Two Birds with Two Bullets: How to Find Minimum Mean OSPA Estimates. *Proceedings of the 13th International Conference on Information Fusion (Fusion 2010)*, 2010.
- [13] J.R. Hoffman and R.P.S. Mahler. Multitarget Miss Distance and its Applications. In *Proceedings of the Fifth International Conference on Information Fusion*, volume 1, pages 149–155 vol.1, July 2002.
- [14] Weifeng Liu, Chongzhao Han, Feng Lian, and Hongyan Zhu. Multitarget State Extraction for the PHD Filter using MCMC Approach. *IEEE Transactions on Aerospace and Electronic Systems*, 46(2):864–883, April 2010.
- [15] Ronald P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Boston, Mass., 2007.
- [16] K. Panta, B.N. Vo, and S. Singh. Novel Data Association Schemes for the Probability Hypothesis Density Filter. *IEEE Transactions on Aerospace and Electronic Systems*, 43(2):556–570, April 2007.
- [17] Jeff M. Phillips and Suresh Venkatasubramanian. A Gentle Introduction to the Kernel Distance. *CoRR*, abs/1103.1625, 2011.
- [18] B. Ristic. *Particle Filters for Random Set Models*. Springer, 2013.
- [19] B. Ristic, D. Clark, and Ba-Ngu Vo. Improved SMC implementation of the PHD filter. In *2010 13th Conference on Information Fusion (FUSION)*, pages 1–8, July 2010.
- [20] D. Schuhmacher, B.-T. Vo, and B.-N. Vo. A Consistent Metric for Performance Evaluation of Multi-Object Filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, 2008.
- [21] Roy L. Streit. *Poisson Point Processes: Imaging, Tracking, and Sensing*. Springer US, 2010.
- [22] X. Tang and P. Wei. Multi-target State Extraction for the Particle Probability Hypothesis Density Filter. *IET Radar, Sonar Navigation*, 5(8):877–883, Oct 2011.
- [23] M. Tobias and A.D. Lanterman. Probability Hypothesis Density-based Multitarget Tracking with Bistatic Range and Doppler Observations. *IEEE Proceedings Radar, Sonar and Navigation*, 152(3):195–205, June 2005.
- [24] M. Tobias and A.D. Lanterman. Techniques for Birth-Particle Placement in the Probability Hypothesis Density Particle Filter Applied to Passive Radar. *IET Radar, Sonar & Navigation*, 2(5):351–365, 2008.
- [25] Hongyan Zhu, Chongzhao Han, and Yan Lin. Particle Labeling PHD Filter for Multi-target Track-Valued Estimates. In *2011 Proceedings of the 14th International Conference on Information Fusion (FUSION)*, pages 1–8, July 2011.