

Association-Free Direct Filtering of Multi-Target Random Finite Sets with Set Distance Measures

Uwe D. Hanebeck and Marcus Baum
Intelligent Sensor-Actuator-Systems Laboratory (ISAS)
Institute for Anthropomatics and Robotics
Karlsruhe Institute of Technology (KIT), Germany
Uwe.Hanebeck@kit.edu, Marcus.Baum@kit.edu

Abstract—We consider association-free tracking of multiple targets without identities. The uncertain multi-target state and the uncertain measurements cannot be described by a random vector as this would imply a certain order. Instead, they are described by an unordered random finite set (RFS). Particle-based random finite set densities are used for characterizing the RFS in a simple and natural way. For recursive Bayesian filtering, optimal multi-target state estimates are calculated by systematically minimizing an appropriate set distance measure while directly operating on the particles. Although methods for calculating point estimates of random finite set densities based on appropriate distance measures are available in literature, the proposed recursive filtering is a novel contribution.

Keywords—Association-free tracking, multi-target tracking, random finite sets, direct filtering, empirical densities, set distance measures.

I. INTRODUCTION

A. Motivation

We consider association-free tracking of multiple targets without identities. The uncertain multi-target state cannot be described by a random vector as this would imply a certain order. Instead, it is described by a random finite set (RFS) [1], [2], which does not impose any order.

Realizations of a random finite set are again described by an unordered set of individual realizations of target locations. Hence, a very natural form of describing the probability density function of a random finite set is in the form of a particle density, i.e., a collection of weighted samples, where each sample is itself a weighted particle density in target space.

Measurements taken from the targets are also assumed to only comprise target locations. Identities are not revealed. Hence, uncertain measurements are also described by particle-based random finite set densities.

For recursive Bayesian filtering, we now face two problems: i) Prior density and measurement density are given by collection of particles only. ii) The association of targets between particle densities is unknown. In this paper, we propose an integrated approach for recursive filtering of random finite sets characterized by a collection of samples that is based on systematically minimizing a set distance measure for producing optimal estimates. Posterior

estimates are again characterized by particle-based random finite set densities.

Remark I.1. In order to illustrate the key concept of our approach, we investigate a simplified setup for multi-target tracking. First, we assume that the number of targets T is known a priori. Second, a simplified measurement model with an identity mapping and additive noise is considered. Third, clutter measurements or misdetections are ignored. ■

B. Related Work

The filtering of multi-target random finite sets is an important problem in multi-target tracking [3], [4], [5]. Most multi-target tracking algorithms such as the Joint Probabilistic Data Association Filter (JPDAF) [3] are based on enumerating all possible association hypotheses. But there are also *association-free* approaches such as the Probability Hypothesis Density (PHD) filter [1], [2], Symmetric Measurement Equation (SME) filter [6], [7], [8], and the Probabilistic Multihypothesis Tracker (PMHT) [9] that inherently avoid the explicit enumeration of association hypotheses.

Traditional multi-target tracking algorithms employ a joint state vector that contains the individual target states and aim at minimizing the Mean Squared Error (MSE) of the estimate. However, this might lead to undesired coalescence effects in case of closely-spaced objects, which is discussed in [10]: Due to the unlabeled measurements, the identities of the targets sooner or later get lost so that the MSE criterion becomes unsuitable. A systematic solution to this problem is to inherently ignore target identities. Technically, this can be achieved by performing estimation w.r.t. a metric for sets. The idea to determine point estimates with a set metric in a multi-target tracking context is discussed in [11], [12]. While [11] considers the Wasserstein metric [13], the work [12] focuses on the OSPA metric [14] (for sets with the same cardinality, the OSPA and Wasserstein metric coincide). In [12] and subsequently [15], [16], [17], [18], [19], [20], [21], [22], practical algorithms for calculating minimum mean OSPA (MMOSPA) estimates are presented. In [23], the kernel distance [24], [25] has been employed for determining point estimates in multi-target trackers.

Several multi-target tracking algorithms that optimize according to the OSPA distance (or its variants) have been developed, e.g., there is the Set-JPDAF [26], the

Set Multiple Hypothesis Tracker (MHT) [27], and tailored particle filters [28], [29], [30], [31]. All these approaches are based on explicitly performing a Bayes update with the multi-target likelihood, i.e., association hypotheses are enumerated.

A mathematically related concept are the so-called Wasserstein barycenters [32], [33], [34], [35], [36]. A Wasserstein barycenter extends the notion of a “mean” to a collection of probability measures. In [37], it is shown that the Wasserstein barycenter for point clouds is mathematically equivalent to the MMOSPA estimate for empirical distributions. Wasserstein barycenters recently gained significant interest in computer vision, see for example [32], [33], [34], [35], [36], [38], [39]. They also find applications for fusing particle distributions [40] and persistence diagrams [41], [42].

The filtering of multi-target random finite sets is an inherently nonlinear problem. A standard particle filter implementation requires the enumeration of all association hypotheses [43]. Our approach to this problem is inspired by the so-called Ensemble Kalman Filter (EnKF) [44], [45], which is a widely used data assimilation techniques (for standard vector-valued states). In the standard EnKF, prior state estimate and the data is represented by particles. The filtering is performed by pairwise combining the prior samples with data samples according to the Kalman filtering equations, where the Kalman gain is computed with the prior sample covariance matrix and the data covariance.

C. Organization of the Paper

In the next section, a rigorous formulation of the association-free tracking problem is given. A specific distance measure well suited for comparing particle sets and collections of particle sets is given in Sec. III. The association-free filtering method is derived in Sec. IV. Some application examples of calculating point estimates, measurement updates, and recursive filtering are described in Sec. V. Sec. VI concludes the paper.

II. PROBLEM FORMULATION

We consider T targets moving in an N -dimensional space at unknown true locations \underline{x}_i , $i = 1, \dots, T$. It is important to note that the targets are *indistinguishable* and thus, are not equipped with an identity.

We estimate the true target positions and denote the random variables used for describing the estimation by \underline{x}_i , $i = 1, \dots, T$. As the targets cannot be distinguished, it does not make sense to stack the targets in a vector as this would imply a certain order. Instead, we use a multi-target random finite set \mathcal{X} containing the individual states in an unordered way [1], [2]

$$\mathcal{X} = \{\underline{x}_1, \underline{x}_2, \dots, \underline{x}_T\} .$$

The random finite set \mathcal{X} is described by a random finite set density $\mathcal{X} \sim p(\mathcal{X})$. Realizations of this random finite set are again unordered sets of individual realizations of target locations. We describe the probability density function of a random finite set by a set of realizations, i.e., by a collection

of weighted samples f_j , $j = 1, \dots, P_x$, where each sample is itself a weighted particle density in target space with locations

$$\mathcal{X}_j = \{\underline{x}_{j1}, \underline{x}_{j2}, \dots, \underline{x}_{jT}\}$$

and associated weights w_{ji}^x for $j = 1, \dots, P_x$ and $i = 1, \dots, T$. The weights are positive and sum up to one. So, every realization of the random set \mathcal{X} is a Dirac mixture density

$$f_j(\underline{x}) = \sum_{i=1}^T w_{ji}^x \delta(\underline{x} - \underline{x}_{ji})$$

itself.

Producing a point estimate $\hat{\mathcal{X}}$ for $p(\mathcal{X})$ is not straightforward, as also in the random set realizations the targets do not have an identity. We just get location values, no identities. So, we do not know where a certain target, say target 1, is located in each random set sample. This corresponds to calculating an average measure of the given P_x samples. The resulting point estimate is another Dirac mixture density with T components!

We are now given T individual measurements of the targets that are corrupted by additive noise. If we draw samples from the measurement noise, we obtain the random set \mathcal{Y} that can be described by a set of samples g_j , $j = 1, \dots, P_y$ with locations

$$\mathcal{Y}_j = \{\underline{y}_{j1}, \underline{y}_{j2}, \dots, \underline{y}_{jT}\}$$

and associated weights w_{ji}^y for $j = 1, \dots, P_y$ and $i = 1, \dots, T$. The weights are positive and some up to one and every sample is described by the density

$$g_j(\underline{x}) = \sum_{i=1}^T w_{ji}^y \delta(\underline{x} - \underline{y}_{ji}) .$$

Instead of averaging measures, we are here interested in performing combination and filtering operations in a Bayesian fusion and estimation sense. Our goal is to perform recursive filtering by updating the multi-target set state \mathcal{X} based on incoming measurement sets. In order to avoid another index, time in this case, we just consider a single update step consisting of the combination of a predicted random set \mathcal{X}^p and a measurement \mathcal{Y} at a certain time step resulting in a new estimate \mathcal{X}^e .

III. DISTANCE MEASURE

We could use any distance measure between two Dirac mixture densities such as the Wasserstein [11] or OSPA distance [12]. Here, we proffer the one from [46] based on the Localized Cumulative Distributions [24] of the two Dirac mixture densities.

We consider two N -dimensional Dirac mixture densities \tilde{f} and f . \tilde{f} is given by

$$\tilde{f}(\underline{x}) = \sum_{i=1}^K w_i^y \delta(\underline{x} - \underline{y}_i) , \quad (1)$$

with positive weights $w_i^y > 0$ for $i = 1, \dots, K$, that sum up to one and K locations

$$\underline{y}_i = \left[y_i^{(1)}, y_i^{(2)}, \dots, y_i^{(N)} \right]^T$$

for $i = 1, \dots, K$. f is given by

$$f(\underline{x}) = \sum_{i=1}^L w_i^x \delta(\underline{x} - \underline{x}_i), \quad (2)$$

with positive weights $w_i^x > 0$ for $i = 1, \dots, L$, that sum up to one and L locations

$$\underline{x}_i = \left[x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(N)} \right]^T$$

for $i = 1, \dots, L$.

We denote the mean of \tilde{f} and f by \tilde{m} and m , respectively. Removing the mean gives $\tilde{y}_i = y_i - \tilde{m}$, $i = 1, \dots, L$ and $\tilde{x}_i = x_i - m$, $i = 1, \dots, L$. Based on the zero-mean samples, compared to [46] we obtain a slightly modified distance measure between the zero-mean samples plus a term accounting for different means

$$D(\tilde{f}, f) = D_{\underline{y}} - 2D_{\underline{x}\underline{y}} + D_{\underline{x}} + \gamma \|\tilde{m} - m\|_2^2, \quad (3)$$

with

$$\begin{aligned} D_{\underline{y}} &= \sum_{i=1}^K \sum_{j=1}^K w_i^y w_j^y \text{xlog} \left(\sum_{k=1}^N \left(\tilde{y}_i^{(k)} - \tilde{y}_j^{(k)} \right)^2 \right), \\ D_{\underline{x}\underline{y}} &= \sum_{i=1}^L \sum_{j=1}^K w_i^x w_j^y \text{xlog} \left(\sum_{k=1}^N \left(\tilde{x}_i^{(k)} - \tilde{y}_j^{(k)} \right)^2 \right), \\ D_{\underline{x}} &= \sum_{i=1}^L \sum_{j=1}^L w_i^x w_j^x \text{xlog} \left(\sum_{k=1}^N \left(\tilde{x}_i^{(k)} - \tilde{x}_j^{(k)} \right)^2 \right), \end{aligned}$$

with $\text{xlog}(z) = z \cdot \log(z)$ and a constant γ for adjusting the mean difference penalty.

A. Generalization: Several Given Densities

When several densities \tilde{f}_i , $i = 1, \dots, P$ are given, we calculate a mean distance by

$$MD(\{\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_P\}, f) = \sum_{i=1}^P \kappa_i D(\tilde{f}_i, f),$$

with $\kappa_i > 0$ and $\sum_{i=1}^P \kappa_i = 1$.

B. Generalization: Several Given Densities and Several Resulting Densities

When not only several densities \tilde{f}_i , $i = 1, \dots, P$ are given, but we also have several densities f_j , $j = 1, \dots, Q$, $Q \leq P$, we reduce the P given densities to the Q desired densities with the following mean distance measure

$$\begin{aligned} RD(\{\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_P\}, \{f_1, f_2, \dots, f_Q\}) \\ = \sum_{i=1}^P \kappa_i \sum_{j=1}^Q \lambda_j D(\tilde{f}_i, f_j), \end{aligned} \quad (4)$$

with $\kappa_i > 0$ and $\sum_{i=1}^P \kappa_i = 1$ and $\lambda_j > 0$ and $\sum_{j=1}^Q \lambda_j = 1$.

IV. MEASUREMENT UPDATE

For recursive filtering, we will now combine the given prior estimate \mathcal{X}^p with the measurement \mathcal{Y} in a Bayesian fusion sense. The key idea is combine all pairs of random set realizations of \mathcal{X}^p and \mathcal{Y} , i.e., the samples f_i^p and g_j for all $i = 1, \dots, P_x$ and $j = 1, \dots, P_y$. At first sight, this would result in a posterior estimate \mathcal{X}^e consisting of $P_x \cdot P_y$ random set realization. We will give several solutions for keeping the cardinality at a fixed level.

The combination is performed by calculating an average of the sample set realizations f_i and g_j by using the mean distance measure

$$MD(\{f_i^p, g_j\}, f_k^e) = \kappa_1 D(f_i^p, f_k^e) + \kappa_2 D(g_j, f_k^e), \quad (5)$$

with $\kappa_i > 0$, $i = 1, 2$ and $\kappa_1 + \kappa_2 = 1$. We set $\kappa_1 = \kappa$ and $\kappa_2 = 1 - \kappa$ for simplicity with $\kappa \in [0, 1]$ and explicitly denote the dependence of the distance on κ by $MD(\{f_i^p, g_j\}, f_k^e, \kappa)$. The new index k for the realizations f_k^e of the posterior estimate \mathcal{X}^e is given by $k = (i-1) \cdot P_x + j$ with $k = 1, \dots, P_x^e$ and new cardinality $P_x^e = P_x \cdot P_y$.

The result of the combination is obtained by performing a minimization of the mean distance measure as

$$f_k^{e,*}(\kappa) = \arg \min_{f_k^e} MD(\{f_i^p, g_j\}, f_k^e, \kappa).$$

There are several options for reducing the number of random set realizations describing the posterior estimate \mathcal{X}^e . We will take a look at the two extreme cases. The most accurate, but computationally most expensive option that easily works for different cardinalities P_x^p and P_y is to explicitly calculate all $P_x^p \cdot P_y$ realizations. Subsequently, these realizations are reduced to the desired number for further processing by minimizing (4). This could be the previous number of realizations, so $P_x^e = P_x^p$.

The simplest reduction method is obtained for the special case of equal numbers of random set realizations for prior estimate \mathcal{X}^p and measurement \mathcal{Y} , that is $P_x^p = P_y$. In that case a simple random selection can be performed similar to what is done in a particle filter. For a given random set realization f_i^p of \mathcal{X}^p , a random selection of a certain g_j is performed. When we sample without replacement, i.e., use all random set realizations in \mathcal{Y} , we can simply combine realizations with the same index, so we combine f_i^p with g_i for $i = 1, \dots, P_x^p$. Random selection can also easily be performed for different cardinalities P_x^p and P_y .

In any case, the result of the fusion of prior estimate \mathcal{X}^p and measurement \mathcal{Y} gives the posterior random set $\mathcal{X}^e(\kappa)$ that depends on the scalar κ and is characterized by P_x^e random set realizations $f_i^{e,*}(\kappa)$, $i = 1, \dots, P_x^e$. κ plays a similar role as the Kalman gain in a standard Kalman filter as it can be used to weight the random set with lower uncertainty more than the one with higher uncertainty when performing the fusion. In case of comparable uncertainties, a fixed selection of $\kappa = 0.5$ is possible and analog to selecting equal gains in the Kalman filter. In case of different uncertainties, κ should be selected systematically. Different uncertainties for example occur in recursive estimation,




Fig. 1. Results of calculating point estimates for particle-based multi-target random finite set density for nine targets drawn from isotropic Gaussian densities with different standard deviations σ . For standard deviations from $\sigma = 0.05$ to $\sigma = 0.20$, the random finite set density is shown in blue. For larger standard deviations, it is omitted. The point estimate is shown in red.

where the uncertainty of the multi-target state estimate shrinks over time with every measurement update. The optimal κ^* is obtained by assessing the extent of the posterior random set $\mathcal{X}^e(\kappa)$ in dependence on κ .

The question now, of course, is how to assess the extent of $\mathcal{X}^e(\kappa)$. The random set is inherently multimodal and we face the problem of missing identities of targets, so we cannot simply calculate of covariance matrix as a measure of extent as this is a concept that only makes sense for ordered random vectors with unimodal densities.

Here, we offer the following natural solution to this problem: For the posterior estimate $\mathcal{X}^e(\kappa)$, a point estimate $\hat{\mathcal{X}}^e(\kappa)$ is calculated and then the minimum distance measure between the original random set and its point estimate is calculated. The distance measure is given by

$$\begin{aligned}
 & MD(\{f_1^{e,*}(\kappa), f_2^{e,*}(\kappa), \dots, f_{P_x^e}^{e,*}(\kappa)\}, \hat{f}^e(\kappa)) \\
 &= \sum_{i=1}^{P_x^e} \lambda_i D(f_i^{e,*}(\kappa), \hat{f}^e(\kappa)) , \quad (6)
 \end{aligned}$$

with $\lambda_i > 0$, $i = 1, \dots, P_x^e$ and $\sum_{i=1}^{P_x^e} \lambda_i = 1$. The point estimate is obtained by minimization of the distance measure as

$$\begin{aligned}
 & \hat{f}^{e,*}(\kappa) \\
 &= \arg \min_{f^e(\kappa)} MD(\{f_1^{e,*}(\kappa), f_2^{e,*}(\kappa), \dots, f_{P_x^e}^{e,*}(\kappa)\}, \hat{f}^e(\kappa)) .
 \end{aligned}$$

The desired minimum distance measure is already calculated as a by-product when calculating the point estimate and is given by

$$\begin{aligned}
 & MD^*(\{f_1^{e,*}(\kappa), f_2^{e,*}(\kappa), \dots, f_{P_x^e}^{e,*}(\kappa)\}, \hat{f}^{e,*}(\kappa)) \\
 &= \min_{\hat{f}^e(\kappa)} MD(\{f_1^{e,*}(\kappa), f_2^{e,*}(\kappa), \dots, f_{P_x^e}^{e,*}(\kappa)\}, \hat{f}^e(\kappa)) . \quad (7)
 \end{aligned}$$

We now select the optimal κ^* as the κ that leads to the minimum distance measure between the posterior random set $\mathcal{X}^e(\kappa)$ and its point estimate $\hat{f}^{e,*}(\kappa)$ in (7). It is important to note that this distance measure is a convex function of the (scalar) $\kappa \in [0, 1]$. Hence, the optimal κ^* can be found with simple optimization methods.

As a final result, we obtain

- the optimal κ^* for the pairwise combination of the random set realizations of the prior estimate \mathcal{X}^p and the measurement \mathcal{Y} ,
- the resulting optimal posterior random set $\mathcal{X}^{e,*} = \mathcal{X}^e(\kappa^*)$,
- its optimal point estimate $\hat{f}^{e,*} = \hat{f}^{e,*}(\kappa^*)$,
- and the minimum distance measure between the optimal posterior random set $\mathcal{X}^{e,*}$ and its optimal point estimate $\hat{f}^{e,*}$ given by $MD^* = MD^*(\{f_1^{e,*}(\kappa^*), f_2^{e,*}(\kappa^*), \dots, f_{P_x^e}^{e,*}(\kappa^*)\}, \hat{f}^e(\kappa^*))$.

In summary, the distance measure (3) is used in two different contexts when fusing two random sets. In (5), it is used in a distance measure between pairs of random set

realization of prior estimate and measurement. In (6), it is used in a distance between the posterior random set and its corresponding point estimate in order to assess the extent of the estimate.

V. APPLICATION EXAMPLES

We start by calculating point estimates for a given particle-based multi-target random finite set density in $N = 2$ dimensions. We consider the nine target configuration from [37] with targets located at $\tilde{x}_{3,(i-1)+j} = [i-2, j-2]^T$, $i = 1, 2, 3$, $j = 1, 2, 3$, see Fig. 1. The random finite set density is characterized by a collection of $P = 200$ equally weighted particle sets of cardinality $T = 9$. These are obtained from sampling isotropic Gaussian densities with covariance matrices

$$C = \sigma^2 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (8)$$

where the standard deviation σ is used as parameter during the simulation. For standard deviations from $\sigma = 0.05$ to $\sigma = 0.20$, the random finite set density is shown in blue in Fig. 1. The point estimate is denoted by red dots. For small noise, the point estimate almost coincides with the true target locations. With increasing noise, the point estimates slightly move away from the true target locations and from each other.

We now perform a single measurement update for $T = 3$ targets in $N = 2$ dimensions. The true targets are located at $\tilde{x}_1 = [0, 2]^T$, $\tilde{x}_2 = [0, 0]^T$, and $\tilde{x}_3 = [0, -2]^T$. The prior multi-target state is described by samples from a Gaussian distribution with $m_1^p = [-2, 2]^T$, $m_2^p = [-2.5, 0]^T$, $m_3^p = [-2, -2]^T$ and equal covariance matrices

$$C^p = 0.2 \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

while the multi-target measurement is described by samples from a Gaussian distribution with $m_1^y = [2, 3]^T$, $m_2^y = [2, 0]^T$, $m_3^y = [2, -3]^T$ with equal covariance matrices

$$C^y = 0.05 \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

These multi-target random set realizations for \mathcal{X}^p and measurement \mathcal{Y} are then randomly permuted and shown in Fig. 2 as point sets in blue for \mathcal{X}^p and purple for \mathcal{Y} .

Realizations of the posterior random set estimates for the multi-target state are shown for fixed $\kappa \in \{0.25, 0.5, 0.75\}$ in red in Fig. 2. The point estimate $\hat{\mathcal{X}}^e$ corresponding to the random set is shown by black dots. It is obvious that the posterior estimate tends towards the measurement for small κ and towards the prior estimate for large κ . Of course, for $\kappa = 0$, the posterior estimate is equal to the measurement and for $\kappa = 1$, the posterior estimate is equal to the prior estimate.

The distance measure in (7) is visualized as a function of κ in Fig. 3 (left) and shows a clear minimum. The optimum κ^* is found by numerically minimizing the distance measure in (7) with respect to κ , which results in $\kappa^* = 0.32$. The posterior random set estimate for the optimum κ^* is shown

in Fig. 3 (right) in red together with the corresponding point estimate in black.

The next step is recursive filtering, where we come back to the nine target configuration. The initial random finite set density describing the prior estimate is described by samples from a Gaussian distribution with zero mean equal covariance matrices

$$C^p = 0.5 \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

The multi-target measurements for every time step are described by samples from a Gaussian distribution with zero mean and equal covariance matrices

$$C^y = 0.1 \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Starting with the prior estimate, 30 recursions are performed according to the filtering method given in Sec. IV. The results are shown in Fig. 4 for the initial estimate and for the results of recursion steps 2, 4, 6, 8, 10, 15, 20, and 30 together with the optimal κ^* . It is nice to see that κ^* increases with every recursion step, as the measurement is weighted less and less when the state estimate becomes more certain. The corresponding point estimates are denoted by black dots. There is obviously no bias on the state estimates calculated with the proposed filtering method.

VI. CONCLUSIONS

We proposed the first systematic method for association-free direct filtering of multi-target states based on a set distance measure. The uncertain multi-target states are described by a random finite set (RFS) instead of a random vector as the missing identities prohibit ordering. The random finite sets are characterized by the most natural density form, i.e., a collection of weighted samples, where each sample is itself a weighted particle density in target space.

Calculating point estimates for these random finite set densities is performed by minimizing an appropriate set distance measure that does without identities. Any set distance measure of this type could be used such as the Wasserstein [13] or the OSPA distance [14]. Here, we use a new distance measure based on Localized Cumulative Distributions [24].

Recursive filtering is directly performed on the random finite set densities describing estimates and measurements. We neither require a likelihood nor a multiplication with a likelihood. The same set distance measure is used to obtain optimal estimation results.

The proposed filtering approach also makes sense in the single target case as it allows the combination of estimates described by empirical measures.

REFERENCES

- [1] R. Mahler, "Statistics 101" For Multisensor, Multitarget Data Fusion," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 53–64, Jan. 2004.
- [2] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Boston, Mass.: Artech House, 2007.




Fig. 2. Result of fusing prior random set estimate \mathcal{X}^p (blue) with random set measurement \mathcal{Y} (purple) resulting in posterior random set estimate $\mathcal{X}^e(\kappa)$ (red) for (left) $\kappa = 0.25$, (middle) $\kappa = 0.5$, (right) $\kappa = 0.75$. The point estimate $\hat{\mathcal{X}}^e$ is shown as black dots.




Fig. 3. Fusing prior random set estimate \mathcal{X}^p with random set measurement \mathcal{Y} resulting in posterior random set estimate $\mathcal{X}^e(\kappa)$. (left) The distance measure D as a function of κ . (right) Result of the fusion for optimal κ^* . The point estimate $\hat{\mathcal{X}}^e$ is shown as black dots.

- [3] Y. Bar-Shalom, F. Daum, and J. Huang, "The Probabilistic Data Association Filter," *IEEE Control Systems Magazine*, vol. 29, no. 6, pp. 82–100, December 2009.
- [4] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: John Wiley & Sons, Inc., 2002.
- [5] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.
- [6] E. Kamen, "Multiple Target Tracking based on Symmetric Measurement Equations," *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 371–374, March 1992.
- [7] E. Kamen and C. Sastry, "Multiple Target Tracking using Products of Position Measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 29, no. 2, pp. 476–493, April 1993.
- [8] M. Baum and U. D. Hanebeck, "The Kernel-SME Filter for Multiple Target Tracking," in *Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*, Istanbul, Turkey, Jul. 2013.
- [9] P. Willett, Y. Ruan, and R. Streit, "PMHT: Problems and Some Solutions," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 738–754, Jul 2002.
- [10] H. Blom, E. Bloem, Y. Boers, and H. Driessen, "Tracking Closely Spaced Targets: Bayes Outperformed by an Approximation?" in *Proceedings of the 11th International Conference on Information Fusion (Fusion 2008)*, June 2008.
- [11] J. Hoffman and R. Mahler, "Multitarget Miss Distance and its Applications," in *Proceedings of the Fifth International Conference on Information Fusion*, vol. 1, Jul. 2002, pp. 149–155 vol.1.
- [12] M. Guerriero, L. Svensson, D. Svensson, and P. 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] C. Villani, *Optimal Transport*. Springer Verlag, 2009.
- [14] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A Consistent Metric for Performance Evaluation of Multi-Object Filters," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, 2008.
- [15] 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*, vol. 59, no. 7, pp. 3441–3447, Jul. 2011.
- [16] —, "Generalizations of Blom and Bloem's PDF Decomposition for Permutation-Invariant Estimation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2011, pp. 3840–3843.
- [17] D. F. Crouse, P. Willett, Y. Bar-Shalom, and L. Svensson, "Aspects of MMOSPA Estimation," in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, Orlando, FL, Dec. 2011*.
- [18] D. F. Crouse, P. Willett, M. Guerriero, and L. Svensson, "An Approximate Minimum MOSPA Estimator," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2011, pp. 3644–3647.




Fig. 4. Recursive association-free multi-tracking over 30 recursion steps. Shown are the results recursion steps 2, 4, 6, 8, 10, 15, 20, and 30. The multi-target state estimates are shown in red and their corresponding point estimates are denoted by black dots.

- [19] D. F. Crouse, P. Willett, L. Svensson, and Y. Bar-Shalom, "Comparison of Compressed Sensing, ML, and MMOSPA Estimation for Radar Superresolution," in *Proc. 45th Asilomar Conference on Signals, Systems and Computers*, 2011.
- [20] D. F. Crouse, "Advances in Displaying Uncertain Estimates of Multiple Targets," in *SPIE - Signal Processing, Sensor Fusion, and Target Recognition XXII*, vol. 8745, 2013, pp. 874 504–874 504–31.
- [21] M. Baum, P. Willett, and U. D. Hanebeck, "MMOSPA-based Direction-of-Arrival Estimation for Planar Antenna Arrays," in *Proceedings of the Eighth IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM 2014)*, A Coruña, Spain, Jun. 2014.
- [22] —, "Polynomial-Time Algorithms for the Exact MMOSPA Estimate of a Multi-Object Probability Density Represented by Particles," *IEEE Transactions on Signal Processing*, no. 10, pp. 2476 – 2484, May 2015.
- [23] M. Baum, P. Ruoff, D. Itte, and U. 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, Dec. 2012.
- [24] U. D. Hanebeck and V. Klumpp, "Localized Cumulative Distributions and a Multivariate Generalization of the Cramér-von Mises Distance," in *Proceedings of the 2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2008)*, Seoul, Republic of Korea, Aug. 2008, pp. 33–39.
- [25] J. M. Phillips and S. Venkatasubramanian, "A Gentle Introduction to the Kernel Distance," *CoRR*, vol. abs/1103.1625, 2011.
- [26] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, "Set JPDA Filter for Multi-Target Tracking," *IEEE Transactions on Signal Processing*, vol. PP, no. 99, p. 1, 2011.
- [27] D. F. Crouse, P. Willett, L. Svensson, D. Svensson, and M. Guerriero, "The Set MHT," in *Proceedings of the 14th International Conference on Information Fusion (Fusion 2011)*, Jul. 2011.

- [28] A. Garcia-Fernandez, M. Morelande, and J. Grajal, "Particle Filter for Extracting Target Label Information when Targets Move in Close Proximity," in *Proceedings of the 14th International Conference on Information Fusion (FUSION)*, 2011, pp. 1–8.
- [29] R. Georgescu, P. Willett, L. Svensson, and M. Morelande, "Two Linear Complexity Particle Filters Capable of Maintaining Target Label Probabilities for Targets in Close Proximity," in *15th International Conference on Information Fusion (FUSION)*, 2012, pp. 2370–2377.
- [30] E. Aoki, Y. Boers, L. Svensson, P. Mandal, and A. Bagchi, "The Rao-Blackwellized Marginal M-SMC Filter for Bayesian Multi-Target Tracking and Labelling," in *15th International Conference on Information Fusion (FUSION)*, 2012, pp. 90–97.
- [31] E. H. Aoki, Y. Boers, L. Svensson, P. Mandal, and A. Bagchi, "A Bayesian Look at the Optimal Track Labelling Problem," in *IET Data Fusion Target Tracking Conference (DF TT 2012): Algorithms Applications*, 2012, pp. 1–6.
- [32] M. Agueh and G. Carlier, "Barycenters in the Wasserstein Space," *SIAM Journal on Mathematical Analysis*, vol. 43, no. 2, pp. 904–924, 2011.
- [33] J. Bigot and T. Klein, "Consistent Estimation of a Population Barycenter in the Wasserstein Space," *ArXiv e-prints*, Dec. 2012.
- [34] M. Cuturi and A. Doucet, "Fast Computation of Wasserstein Barycenters," in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, Oct. 2014.
- [35] R. Bhattacharya and V. Patrangenaru, "Large Sample Theory of Intrinsic and Extrinsic Sample Means on Manifolds," *The Annals of Statistics*, vol. 31, no. 1, pp. 1–29, Feb. 2003.
- [36] G. Carlier, A. Oberman, and E. Oudet, "Numerical Methods for Matching for Teams and Wasserstein Barycenters," *Tech. Rep.*, May 2014.
- [37] M. Baum, P. Willett, and U. D. Hanebeck, "On Wasserstein Barycenters and MMOSPA Estimation," *IEEE Signal Processing Letters*, vol. 22, no. 10, pp. 1511–1515, Oct. 2015.
- [38] J. Solomon, R. Rustamov, L. Guibas, and A. Butscher, "Earth Mover's Distances on Discrete Surfaces," in *ACM Transactions on Graphics (SIGGRAPH)*, 2014.
- [39] H. Ding and J. Xu, "Finding Median Point-Set Using Earth Mover's Distance," in *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, 2014.
- [40] A. Bishop, "Information Fusion via the Wasserstein Barycenter in the Space of Probability Measures: Direct Fusion of Empirical Measures and Gaussian Fusion with Unknown Correlation," in *17th International Conference on Information Fusion*, Salamanca, Spain, 2014.
- [41] Y. Mileyko, S. Mukherjee, and J. Harer, "Probability Measures on the Space of Persistence Diagrams," *Inverse Problems*, vol. 27, no. 12, p. 124007, 2011.
- [42] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer, "Fréchet Means for Distributions of Persistence Diagrams," *ArXiv e-prints*, Jun. 2012.
- [43] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House Publishers, 2004.
- [44] G. Evensen, "The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation," *Ocean Dynamics*, vol. 53, no. 4, pp. 343–367, 2003.
- [45] —, *Data Assimilation: The Ensemble Kalman Filter*. Springer Science & Business Media, 2009.
- [46] U. D. Hanebeck, "Optimal Reduction of Multivariate Dirac Mixture Densities," *arXiv preprint: Systems and Control (cs.SY)*, Nov. 2014. [Online]. Available: <http://arxiv.org/abs/1411.4586>