Symmetrizing Measurement Equations for Association-free Multi-Target Tracking via Point Set Distances

Uwe D. Hanebeck[a], Marcus Baum[b], and Peter Willett[c]

[a]Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany
[b]Institute of Computer Science, University of Goettingen, Germany
[c]Department of Electrical and Computer Engineering, University of Connecticut, USA

ABSTRACT

We are tracking multiple targets based on noisy measurements. The targets are labeled, the measurements are unlabeled, and the association of measurements to targets is unknown. Our goal is association-free tracking, so the associations will never be determined as this is costly and impractical in many scenarios. By employing a permutation-invariant and differentiable point set distance measure, we derive a modified association-free multi-target measurement equation. It maintains the target identities but is invariant to permutations in the unlabeled measurements. Based on this measurement equation, we derive an efficient sample-based association-free multi-target Kalman filter. The proposed new approach is straightforward to implement and scalable.

Keywords: Multi-target tracking, association-free tracking, point set distance, localized cumulative distribution, symmetric measurement equation, random finite sets, sample-based nonlinear Kalman filter, deterministic sampling

1. INTRODUCTION

Motivation: Multi-target tracking is concerned with maintaining an estimate of the locations (and often velocities) of several moving targets. The uncertainty of the estimates increases over time, so that it is necessary to regularly update the locations by measurements that are also uncertain. When the measurements are unlabeled, the updating process becomes a main challenge – especially for closely spaced targets – as it is not clear which measurement stems from which target.

Related Work: Many of the algorithms for multi-target tracking perform an enumeration of all possible target-measurement associations,[1,2] which is combinatorial in the number of targets and the number of measurements.

For example, the Multiple Hypothesis Tracker (MHT)[3] maintains association hypotheses over many time steps and continuously updates them, which is in general intractable so that approximate methods have to be employed in order to reduce computational complexity. However, even these approximations are typically computationally intensive. The Joint Probabilistic Data Association (JPDA) filter[4,5] also involves the calculation of marginal association probabilities, i.e., all feasible target-measurement associations have to be enumerated.

The Probability Hypothesis Density (PHD) filter[6] prevents the explicit enumeration of associations by maintaining the first-order moment (this is what the PHD is) of the posterior random finite set. However, random finite set methods[7,8] that work with multi-Bernoulli representations involve the explicit enumeration of associations.

A simpler alternative is to estimate a hard association at every time step. One option is to calculate an acceptance region based on the prior multi-target state estimate and associate the closest measurements within

Further author information: (Send correspondence to U.D.H.)
U.D.H.: E-mail: Uwe.Hanebeck@kit.edu, Telephone: +49 (0)721 608 43909
M.B.: E-mail: Marcus.Baum@cs.uni-goettingen.de, Telephone: +49 (0)551 39 172024
P.W.: E-mail: Peter.Willett@uconn.edu, Telephone: +1 860 486 2195

this region to (the means of) the given tracks. This option is called local nearest neighbor association[2] and has a very low computational complexity. Its disadvantage is that a single measurement might be associated with several targets or not be associated at all. Another option that alleviates this problem is to perform a global nearest neighbor neighbor association,[2] which finds the optimal permutation of measurements that minimizes the sum of distances (or squared distances) between targets and measurements.

Another approach operating at one time step that avoids the enumeration of all possible associations between target states and measurements is to perform a symmetric transformation of the measurement equation, which leads to the Symmetric Measurement Equation (SME) filter.[9,10] The resulting measurement equation is permutation-invariant with respect to the order of the states and the order of the measurements, so that an association hypothesis is not required. As pointed out before,[11] the symmetric measurement equations in [9] and [10] are highly nonlinear, do not scale well with the number of targets and measurements, and a suitable transformation is nontrivial to find in dimensions larger than one. In addition, when the targets are labeled, a measurement equation that is permutation-invariant with respect to the target states is not required and not even desirable as it makes maintaining the target labels across the measurement update step more difficult.

The Kernel-SME filter[11] aims at solving these issues by using a kernel transformation, e.g., with Gaussian kernels. By this means, a set is interpreted as a Gaussian mixture function. Unfortunately, the Gaussian mixture function must be discretized/sampled in order to obtain an explicit measurement equation. The number and locations of the discretization points strongly influences the estimation quality. In [12], the Kernel-SME filter is generalized to correlated targets.

When the only information about both the target state and the measurements is given by sets of unordered particles, Association-Free Direct Filtering[13] can be employed. This approach is similar to the Ensemble Kalman Filter (EnKF)[14,15] for standard vector-valued states, where the prior state samples are pairwise combined with samples from the given measurement and its uncertainties. The EnKF uses a Kalman gain computed from the prior sample covariance matrix and the (given) measurement covariance. The big difference with Association-Free Direct Filtering is that the samples are unordered and, without an association between the samples themselves and between targets and measurements, are difficult to combine. In addition, no covariance matrix can be calculated as this is only possible for vector-valued quantities. The key idea is to use a permutation-invariant distance measure[16,17] to calculate weighted barycenters for all pairs of samples from target states and measurements, which results in a set of unordered samples constituting an updated multi-target state. The weights are calculated by iteratively minimizing an extension measure of the updated multi-target state that does not need an association between samples.

Contribution: In this paper, we consider the tracking of multiple labeled targets based un unlabeled noisy measurements. We also use the permutation-invariant distance measure from [16], [17] between multi-target states and measurements. However, we use its gradient to obtain an explicit modified multi-target measurement equation that is permutation-invariant with respect to the measurements but accepts a vector-valued multi-target state and maintains target identities. As a result, no association needs to be calculated, maintained, or updated. This approach is different from the SME approach, where the measurement equation is permutation-invariant with respect to both target states and measurements.

Organization of the Paper: The paper is structured as follows. In the next section, a detailed formulation of the problem of tracking labeled targets based on unlabeled measurements is given. Sec. 3 derives an association-free multi-target measurement equation based on a certain point set distance measure and its gradient. Based on that measurement equation, an efficient sample-based nonlinear Kalman filter for association-free multi-target tracking is derived in Sec. 4. Implementation details are provided in Sec. 5. The performance of the proposed new approach is demonstrated by means of simulations in Sec. 6. Sec. 7 concludes the paper. An outlook to future work is given in Sec. 8.

2. PROBLEM FORMULATION

We consider T targets that are moving in a D-dimensional space and, at a certain point in time, are located at unknown true locations $\underline{\tilde{x}}_i$, $i = 1, \ldots, T$. The time index is omitted for the sake of simplicity. The number of targets T is known and the targets are equipped with an identity, i.e., are labeled, so they are not simply interchangeable. Thanks to this order, we can collect the true target states in a true multi-target state vector $\underline{\tilde{x}} = \left[(\underline{\tilde{x}}_1)^\top, (\underline{\tilde{x}}_2)^\top, \ldots, (\underline{\tilde{x}}_T)^\top\right]^\top$.

A prior estimate in form of a random vector

$$\underline{\boldsymbol{x}}^p = \left[(\underline{\boldsymbol{x}}_1^p)^\top, (\underline{\boldsymbol{x}}_2^p)^\top, \ldots, (\underline{\boldsymbol{x}}_T^p)^\top\right]^\top$$

of the multi-target state is available consisting of a mean vector

$$\underline{\hat{x}}^p = \left[(\underline{\hat{x}}_1^p)^\top, (\underline{\hat{x}}_2^p)^\top, \ldots, (\underline{\hat{x}}_T^p)^\top\right]^\top \tag{1}$$

and a covariance matrix

$$\mathbf{C}^p \in \mathbb{R}^{DT \times DT} \ . \tag{2}$$

For each target, we obtain $P \geq 1$ measurements, where in this paper we assume the number of measurements per target to be equal. To again simplify notation, we first assume $P = 1$, i.e., one measurement per target with a detection probability of one. The extension to $P > 1$ measurements will be shown in a subsequent paper, see Sec. 8. The approach proposed in this paper can cope with arbitrary nonlinear relations between the target states and the measurements. However, for the exposition here we focus on an identity mapping, so the measurement are simply direct observations of the target locations. Hence, the measurements are given by random vectors $\underline{\boldsymbol{y}}_i$, $i = 1, 2, \ldots, T$ with mean vectors $\underline{\hat{y}}_i$, $i = 1, 2, \ldots, T$ and corresponding covariance matrices $\mathbf{C}_i^y$, $i = 1, 2, \ldots, T$. It is important to note that the *association of measurements to targets is unknown*, so the indexing of the measurements is arbitrary and bears no relation to the indexing of the targets*. It is assumed that all the measurements are uncorrelated with the predicted multi-target state estimate. For simplicity, the measurements themselves are all assumed to be mutually uncorrelated in this paper.

Without a specific ordering, it does not make sense to collect the measurement information in a vector, unless this vector is processed in a permutation-invariant way as we will see later. So we have to be content with describing it as a random finite set $\boldsymbol{\mathcal{Y}}$ given by $\boldsymbol{\mathcal{Y}} = \{\underline{\boldsymbol{y}}_1, \underline{\boldsymbol{y}}_2, \ldots, \underline{\boldsymbol{y}}_T\}$, i.e., a set of random vectors that in the considered special case is of fixed cardinality. When we assume different uncertainty quantifications, they have to be associated with the individual measurements. Here, we assume Gaussian measurement noise and the measurement set is written as set of tuples

$$\boldsymbol{\mathcal{Y}} = \left\{ \left(\underline{\boldsymbol{y}}_1, \underline{\hat{y}}_1, \mathbf{C}_1^y\right), \left(\underline{\boldsymbol{y}}_2, \underline{\hat{y}}_2, \mathbf{C}_2^y\right), \ldots, \left(\underline{\boldsymbol{y}}_T, \underline{\hat{y}}_T, \mathbf{C}_T^y\right) \right\} \ .$$

A specific realization of this random finite set at a certain time step again is an unordered set, but now we get specific measurement vectors collected in a set $\mathcal{Y} = \{\underline{y}_1, \underline{y}_2, \ldots, \underline{y}_T\}$.

In this paper, we focus on updating a prior target state estimate $\underline{\boldsymbol{x}}^p$ based on the given set of noisy measurements $\mathcal{Y}$ in order to obtain an enhanced posterior estimate $\underline{\boldsymbol{x}}^e = \left[(\underline{\boldsymbol{x}}_1^e)^\top, (\underline{\boldsymbol{x}}_2^e)^\top, \ldots, (\underline{\boldsymbol{x}}_T^e)^\top\right]^\top$ of the target state with mean vector $\underline{\hat{x}}^e = \left[(\underline{\hat{x}}_1^e)^\top, (\underline{\hat{x}}_2^e)^\top, \ldots, (\underline{\hat{x}}_T^e)^\top\right]^\top$ and covariance matrix $\mathbf{C}^e \in \mathbb{R}^{DT \times DT}$. This constitutes a single multi-target measurement update for a given measurement set at a certain time step. A recursive filter is simply obtained by alternating between the multi-target measurement update and the multi-target time update based on the system model given by equation (4) in [11].

The individual target states will typically be correlated after a measurement update, so $\mathbf{C}^e$ will be a fully populated matrix even when starting with a block-diagonal initial covariance matrix $\mathbf{C}^p$ of the prior multi-target state estimate.

*This will become obvious in the simulation of measurements in Fig. 2 (top).

3. ASSOCIATION-FREE MULTI-TARGET MEASUREMENT EQUATION

Let us begin with the simplest case of known associations of measurements to targets. The measurements can be ordered and we come up with a multi-dimensional measurement equation. When the relation between target states and measurements is linear, we obtain a set of linear measurement equations and a standard Kalman filter can directly be used for updating the prior multi-target estimate based on the measurements.

When the associations are not known, according to our discussion above, many methods first estimate some form of optimal association. In a second step, a set of (linear) measurement equations is then set up based on the estimated association and used for subsequent filtering. For the given association, this set of measurement equations is used to form a measurement hypothesis based on the given multi-target state estimate. Based on the difference between the obtained measurements and the measurement hypothesis together with the uncertainties of both prior target and measurement uncertainties, the target state is corrected in such a way that this difference becomes smaller. The difference between given and hypothesized measurements is calculated by means of standard vector subtraction.

Here, we want to perform *association-free* multi-target tracking, which means that we want to perform the multi-target measurement update without ever having to calculate the associations from measurements to targets. For that purpose, we need a measurement equation relating targets to measurements that is invariant to permutations in the measurements. Hence, we need to somehow reformulate the original measurement problem with unknown associations in such a way that the order of the measurements does not matter but the target identities are still maintained. But how do we correct the target states so that filtering pulls them closer to the measurements?

The key idea is the following: We set up a modified measurement equation in such way that target states are attracted by the measurements during filtering without performing an association. Of course, in order to tell when the target states are close to the measurements, we need an appropriate distance measure.

For that purpose, we use the point set distance measure $D(.,.)$,[16,17] which is in fact a generalized Cramér-von Mises distance based on Localized Cumulative Distributions (LCDs) specialized to Dirac mixtures. It measures the distance between two (weighted) point sets interpreted as Dirac mixtures without the need to know or find an explicit association. The distance measure is employed to measure the distance between the random finite set describing the multi-target state $\boldsymbol{\mathcal{X}} = \{\underline{\boldsymbol{x}}_1, \underline{\boldsymbol{x}}_2, \ldots, \underline{\boldsymbol{x}}_T\}$ and the random finite set describing the measurements $\boldsymbol{\mathcal{Y}} = \{\underline{\boldsymbol{y}}_1, \underline{\boldsymbol{y}}_2, \ldots, \underline{\boldsymbol{y}}_T\}$ according to

$$\boldsymbol{d} = D(\boldsymbol{\mathcal{Y}}, \boldsymbol{\mathcal{X}}) = D(\{\underline{\boldsymbol{y}}_1, \underline{\boldsymbol{y}}_2, \ldots, \underline{\boldsymbol{y}}_T\}, \{\underline{\boldsymbol{x}}_1, \underline{\boldsymbol{x}}_2, \ldots, \underline{\boldsymbol{x}}_T\}) \ .$$

The order of the measurements and also, for now, of the target states does not matter and no association between targets and measurements is required.

When applied to multi-target tracking, a significant advantage of this point set distance measure now is that we do not have to minimize it explicitly. As it is differentiable, we can use its gradient and move the target states in such a way that we come close to its stationary point, i.e., the gradient should be zero when the target states are as close as possible to the measurements[†]. There is another advantage that we will see shortly.

The gradient of the association-free point set distance measure $D(.,.)$[16] can be calculated in closed form.[17] In abstract form, it is given by

$$\underline{\boldsymbol{g}} = \frac{\partial D(\boldsymbol{\mathcal{Y}}, \boldsymbol{\mathcal{X}})}{\partial \underline{\boldsymbol{x}}} = \underline{G}(\boldsymbol{\mathcal{Y}}, \underline{\boldsymbol{x}}) \ . \tag{3}$$

We will now use this equation as a multi-target measurement equation. It is important to note that it is given in an implicit form, i.e., the target states and measurements enter the equation on the right-hand-side. As a result, the equation can be viewed a (nonlinear) constraint for the target states based on the given measurements.

[†]A gradient of zero not only indicates the desired global minimum of the point set distance measure $D(.,.)$ but also its local minima/maxima or saddle points. In the evaluation, we never found it necessary to distinguish the global minimum from the local extrema or saddle points. However, when this should become an issue, there is a simple mechanism to achieve this, see the end of Remark 3.1.

For filtering, we desire the left-hand-side $\underline{\boldsymbol{g}} \in \mathbb{R}^{DT}$ of this implicit multi-target measurement equation to be as close to the zero vector as possible. It is exactly zero, when the multi-target state is as close as possible to the measurement, which for filtering, we express with a pseudo-measurement $\underline{g} = \underline{0}$. When given a prior estimate for the multi-target state, the filter will position the posterior multi-target state estimate in such a way that the left-hand-side of (3) comes closer to the pseudo-measurement $\underline{g} = \underline{0}$ (but is typically not zero). This of course depends on the relative uncertainties of the measurements and the prior multi-target state estimate.

On the right-hand-side, we see the second advantage mentioned before: The gradient depends upon the measurements, but is permutation-invariant with respect to them and it depends upon the target states but for them, the order is important. This is caused by taking the derivative with respect to the components of the multi-target state vector. Hence, we use the multi-target state random vector $\underline{\boldsymbol{x}} = \left[\underline{\boldsymbol{x}}_1^\top, \underline{\boldsymbol{x}}_2^\top, \ldots, \underline{\boldsymbol{x}}_T^\top\right]^\top$ as an input to the multi-target measurement equation instead of just the set $\boldsymbol{\mathcal{X}}$.

You could argue that this distance-based procedure is not so much different from minimizing a standard distance, such as for global nearest neighbors (GNN), but it is! For GNN, a hard association is determined just based on the expected values without considering the uncertainties in target states and measurements. Plus, the association is already fixed before the filter starts working its magic. In our case, the target states are shifted towards the measurements directly by the filter in order to minimize the set point distance measure.

REMARK 3.1. *In some cases, it is possible to use the point set distance measure $D(.,.)$ directly instead of its gradient. However, this comes with two restrictions. First, it requires one to calculate the pseudo-measurement corresponding to a location of the target states that minimize the distance $D(.,.)$ between the target states and the measurements without considering the prior knowledge about the targets. The pseudo-measurement is zero only in very simple scenarios such as the one considered in this paper. Even for slightly more complex scenarios with a fixed number of (still unlabeled) measurements from each target, the point set distance measure $D(.,.)$ obviously typically does not become zero even in the global minimum, as the target states then have a non-zero distance from the measurements. For more complex scenarios, see Sec. 8.*

Second, the set distance measure $D(.,.)$ is permutation-invariant to both the measurements and the targets, so that it is more difficult to maintain the target identities. Third, the pseudo-measurement is scalar, while the state is TD-dimensional, so that less information is fed back during the measurement update compared to using the gradient leading to worse estimation results.

However, there is an application of directly using the point set distance in this paper. When it is required to distinguish the global minimum from the local extrema or saddle points, which cannot be done with the gradient in (3) alone, the point set distance can be used to additionally augmenting the measurement in (12).

4. ASSOCIATION-FREE MULTI-TARGET FILTERING

We now want to use the multi-target measurement equation (3) to derive an appropriate filter for updating the prior multi-target state estimate based on the given noisy measurements. Here, we focus on a sample-based nonlinear Kalman filter or Linear Regression Kalman Filter (LRKF),[18] such as the Unscented Kalman Filter (UKF)[19,20] or the Smart Sampling Kalman Filter (S2KF).[21–23] The measurement equation can, however, also be used to develop more advanced filters such as particle filters.

The general idea of sample-based nonlinear Kalman filters is to take samples of all input random variables based on their uncertainties, propagate them through the measurement equation in order to obtain hypothetical measurement outputs, and calculate the mean and covariance matrix of the joint state and measurement vector. Then, an update of the prior state estimate is performed accordingly. In our case, we consider the output of the multi-target measurement equation that is given in *implicit* form, which makes filtering a bit more complicated. In addition, sampling of the measurement noise is more involved as the measurements are unlabeled.

Let us first take a coarse look on determining an estimate for the statistics of the joint multi-target prior and pseudo-measurement state. Deterministically sampling the prior multi-target state estimate is the easy part and is done the standard way with many sampling methods are available.[24,25] We use the prior mean (1) and the prior covariance matrix (2), assume a Gaussian distribution, and generate P samples $(\underline{x}^p)^{(i)}$ for $i = 1, 2, \ldots, P$ with a method[22] that allows to generate an arbitrary number of sample points. For the measurements, we have to

sample the random finite set $\boldsymbol{\mathcal{Y}}$, which results in a set of P random finite set samples $\left\{\mathcal{Y}^{(1)}, \mathcal{Y}^{(2)}, \ldots, \mathcal{Y}^{(P)}\right\}$. The absolute order of the measurements does not matter and also the different samples can have different orders with respect to each other as order is not used at all in the measurement equation. Then we plug the P samples into the measurement equation (3) and obtain P hypothesized pseudo-measurement samples $\underline{g}^{(i)}$, $i = 1, 2, \ldots, P$ that uniquely correspond to multi-target state samples $\underline{x}^{(i)}$ for $i = 1, 2, \ldots, P$. Stacking the samples gives the desired samples in joint space and we can calculate the desired statistics.

Now we give a detailed account on how to perform the measurement update. The calculations will be rewritten in vector/matrix-form, which is more compact and in most implementations also more efficient. For that purpose, we exploit the fact that the order of the measurements or their indexing does not affect the result of the multi-target measurement equation (3). We can select an arbitrary indexing scheme and temporarily collect all the specific measurements at one time step in a vector

$$\underline{y} = \left[(\underline{y}_1)^\top, (\underline{y}_2)^\top, \ldots, (\underline{y}_T)^\top\right]^\top \in \mathbb{R}^{DT} \tag{4}$$

with a corresponding block-diagonal covariance matrix

$$\mathbf{C}^y = \begin{bmatrix} \mathbf{C}_1^y & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_2^y & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{C}_T^y \end{bmatrix} \in \mathbb{R}^{DT \times DT} \ .$$

We stack the prior multi-target state mean vector and the measurement vector to form an augmented random vector

$$\underline{\boldsymbol{z}} = \begin{bmatrix} \underline{\boldsymbol{x}}^p \\ \underline{\boldsymbol{y}} \end{bmatrix} \in \mathbb{R}^{2DT} \tag{5}$$

with mean

$$\hat{\underline{z}} = \begin{bmatrix} \hat{\underline{x}}^p \\ \underline{y} \end{bmatrix}$$

and covariance matrix

$$\mathbf{C}^z = \begin{bmatrix} \mathbf{C}^p & \mathbf{0} \\ \mathbf{0} & \mathbf{C}^y \end{bmatrix} \in \mathbb{R}^{2DT \times 2DT} \ .$$

The off-diagonal block matrices of $\mathbf{C}^z$ are zero as the predicted multi-target states and the measurements are uncorrelated. For obtaining samples from the random vector $\underline{\boldsymbol{z}}$ with mean $\hat{\underline{z}}$ and covariance matrix $\mathbf{C}^z$, we use the matrix square root $\boldsymbol{\mathcal{C}}^z$ with $\boldsymbol{\mathcal{C}}^z(\boldsymbol{\mathcal{C}}^z)^\top = \mathbf{C}^z$, where $\boldsymbol{\mathcal{C}}^z$ is a lower-triangular real matrix with positive diagonal entries. It can be obtained by Cholesky decomposition.[26] We obtain a matrix $\mathbf{Z}$ containing the desired P samples of $\underline{\boldsymbol{z}}$ obviously by

$$\mathbf{Z} = \boldsymbol{\mathcal{C}}^z \cdot \mathbf{S} + \hat{\underline{z}} \cdot (\underline{1}_P)^\top \ , \tag{6}$$

where $\mathbf{S}$ is a matrix containing samples $\underline{s}_i$, $i = 1, 2, \ldots, P$ from a $(2DT)$-dimensional standard normal distribution as columns according to

$$\mathbf{S} = \left[\underline{s}_1, \underline{s}_2, \ldots, \underline{s}_P\right] \in \mathbb{R}^{2DT \times P} \tag{7}$$

and $\underline{1}_P$ is a vector of length P containing all ones. The columns of $\mathbf{Z}$ are the desired samples, where every sample $\mathbf{Z}(:, i)$[‡] for $i = 1, 2, \ldots, P$ has an upper part $\mathbf{Z}(1 : DT, :)$ corresponding to the prior multi-target state estimate and a lower part $\mathbf{Z}(DT + 1 : 2DT, :)$ corresponding to the multi-target pseudo-measurements. Again, for this lower part, the order of the DT-dimensional subvectors in the entire vector is arbitrary and does not affect the multi-target measurement equation.

To simplify the notation, we overload the multi-target measurement equation so that it accepts a single vector composed according to the discussion above. Furthermore, we can overload it in such a way that it accepts an entire matrix composed of these vectors, such as $\mathbf{Z}$, calculates an output vector for each of these columns, and

[‡]We use the indexing notation from [26] that is also used in MATLAB.

returns an output matrix containing these output vectors as columns in the same column order as the vectors in the input matrix. Hence, we get

$$\mathbf{G} = \left[\underline{g}^{(1)}, \underline{g}^{(2)}, \ldots, \underline{g}^{(P)}\right] = \underline{G}(\mathbf{Z}) \tag{8}$$

with sample mean

$$\hat{\underline{g}} = \text{Mean}(\mathbf{G}) \ . \tag{9}$$

The matrix containing the desired samples in the joint space of the prior multi-target estimate $\underline{\boldsymbol{x}}^p$ and the output of the multi-target measurement equation $\underline{\boldsymbol{g}}$ is given by stacking $\mathbf{Z}(1:DT,:)$ and $\mathbf{G}$. The sample covariance matrix of these joint space samples is then given by

$$\hat{\mathbf{C}}_{\ g}^{p} = \text{Cov}\left(\begin{bmatrix}\mathbf{Z}(1:DT,:)\\ \mathbf{G}\end{bmatrix}\right) \ .$$

This sample covariance matrix can be partitioned into the required sample covariance matrices required for the measurement update

$$\begin{bmatrix}\hat{\mathbf{C}}^{pp} & \hat{\mathbf{C}}^{pg}\\ \hat{\mathbf{C}}^{gp} & \hat{\mathbf{C}}^{gg}\end{bmatrix} = \hat{\mathbf{C}}_{\ g}^{p} \ ,$$

where $\hat{\mathbf{C}}^{pp}$ corresponds to the prior multi-target state estimate. We do not need to use the sample covariance matrix $\hat{\mathbf{C}}^{pp}$ as we already have the prior covariance matrix $\mathbf{C}^p$ available. The sample covariance matrix $\hat{\mathbf{C}}^{gg}$ corresponds to the output of the multi-target measurement equation (3). The sample covariance matrices $\hat{\mathbf{C}}^{pg}$ and $\hat{\mathbf{C}}^{pg}$ with $\hat{\mathbf{C}}^{pg} = (\hat{\mathbf{C}}^{pg})^\top$ are the cross-variance matrices between the multi-target state and the output of the multi-target measurement equation (3).

Given the required covariance matrices, the multi-target state estimate can now be updated according to the standard Kalman filter equations with a pseudo-measurement $\underline{g}$ of zero

$$\hat{\underline{x}}^e = \hat{\underline{x}}^p + \hat{\mathbf{C}}^{pg}\left(\hat{\mathbf{C}}^{gg}\right)^{-1}(\underline{g} - \hat{\underline{g}}) \text{ with } \underline{g} = \underline{0} \tag{10}$$

and

$$\mathbf{C}^e = \mathbf{C}^p - \hat{\mathbf{C}}^{pg}\left(\hat{\mathbf{C}}^{gg}\right)^{-1}\hat{\mathbf{C}}^{gp} \ . \tag{11}$$

The advantage of maintaining target labels is that we can calculate the joint covariance matrices of the multi-target estimates, e.g., $\mathbf{C}^e$. This allows us to capture the correlations between the target states. $\mathbf{C}^e$ becomes a fully populated covariance matrix even when starting with uncorrelated targets.

5. IMPLEMENTATION DETAILS

Generation of samples: The sample-based nonlinear Kalman filter described in Sec. 4 for performing the measurement update requires the set of samples in (7). We could simply use random samples for that purpose. However, for a small number of samples, this leads to poor results. A large number of random samples leads to a high computational load. In addition, random samples give non-reproducible results.

Deterministic samples, on the other hand, provide reproducible results and typically cover the state space in a more homogeneous way so that less samples are required. As the filter works similar to a standard Unscented Kalman Filter (UKF), we could use its corresponding deterministic samples capturing the first two moments of Gaussian distributions, i.e., two samples per dimension and an optional center sample.

Better results are obtained with deterministic sampling schemes that provide an arbitrary prespecified number of samples for a given Gaussian distribution. Here, we use the methods in [23] and[27]. We restrict ourselves to samples of a standard Gaussian distribution that are subsequently transformed by (6) to yield samples with desired mean and covariance matrix. This allows the use of a library of precomputed and stored sample sets, which are indexed by the number of dimensions of the underlying random vector and the number of samples P. For Gaussian distributions with arbitrary covariance matrices, precomputing and storing the samples is not possible as the covariance matrices then depend on continuous parameters. For sampling $\underline{\boldsymbol{z}}$ in (5), we require $2DT$-dimensional samples. The number of samples is set to $P = 6(2DT) + 1$.

Augmentation of measurement equation: The performance of the sample-based nonlinear Kalman filter described in Sec. 4 can be increased by augmenting the measurement equation (3) in Sec. 3 with appropriate summary statistics of measurements and multi-target state. It is important that the statistics for the measurements be permutation-invariant.

For simplifying the calculation of these permutation-invariant statistics, we represent the prior multi-target mean in (1) and the temporary measurement vector in (4) in a different form by using the vec operator. The vec operator with $\mathrm{vec} : \mathbb{R}^{M \times N} \to \mathbb{R}^{M \cdot N}$ in general stacks the column of a matrix $\mathbf{U} \in \mathbb{R}^{M \times N}$ into a single vector $\underline{u} \in \mathbb{R}^{M \cdot N}$ as

$$\underline{u} = \begin{bmatrix} \underline{u}_1 \\ \underline{u}_2 \\ \vdots \\ \underline{u}_N \end{bmatrix} = \mathrm{vec}_{M,N}\left(\left[\underline{u}_1, \underline{u}_2, \ldots, \underline{u}_N\right]\right) = \mathrm{vec}_{M,N}\left(\mathbf{U}\right) \text{ with } \underline{u}_i \in \mathbb{R}^M \text{ for } i = 1, 2, \ldots, N \ .$$

The inverse of the vec operator is given by $\mathrm{vec}^{-1}_{M,N} : \mathbb{R}^{M \cdot N} \to \mathbb{R}^{M \times N}$ that constructs the matrix $\mathbf{U}$ from the vector $\underline{u}$ as $\mathbf{U} = \mathrm{vec}^{-1}_{M,N}(\underline{u})$.

The prior multi-target mean in (1) can now be represented in matrix form as

$$\hat{\mathbf{X}}^p = \mathrm{vec}^{-1}_{D,T}(\hat{\underline{x}}^p) = \left[\hat{\underline{x}}^p_1, \hat{\underline{x}}^p_2, \ldots, \hat{\underline{x}}^p_T\right] \in \mathbb{R}^{D \times T} \ .$$

which also works for the temporary measurement vector in (4) as

$$\mathbf{Y} = \mathrm{vec}^{-1}_{D,T}(\underline{y}) = \left[\underline{y}_1, \underline{y}_2, \ldots, \underline{y}_T\right] \in \mathbb{R}^{D \times T} \ .$$

Permutation-invariant summary statistics are obtained by comparing summaries of the rows of $\hat{\mathbf{X}}^p$ and $\mathbf{Y}$. For example, unbiasedness conditions are obtained by comparing row means[§]

$$\underline{c}_1 = \frac{1}{T}\hat{\mathbf{X}}^p \, \underline{1}_T - \frac{1}{T}\mathbf{Y} \, \underline{1}_T \stackrel{!}{=} \underline{0} \ ,$$

where $\underline{1}_T$ is a column vector containing T ones. The same holds for conditions based on quadratic terms

$$\underline{c}_2 = \frac{1}{T}\left(\hat{\mathbf{X}}^p \odot \hat{\mathbf{X}}^p\right) \underline{1}_T - \frac{1}{T}\left(\mathbf{Y} \odot \mathbf{Y}\right) \underline{1}_T \stackrel{!}{=} \underline{0} \ ,$$

where $\odot$ denotes the Hadamard product, i.e., entrywise multiplication. The augmentation of (3) is now performed by

$$\bar{g} = \begin{bmatrix} \underline{g} \\ \underline{c}_1 \\ \underline{c}_2 \end{bmatrix} \ . \tag{12}$$

When performing this augmentation, we remove the means of the target states and the measurements before feeding them into the distance measure and the gradient in (3). This leads to slightly modified expressions for the gradient, where we also have to subtract its means.

Handling dependencies: The multi-target measurement equation (3) generates output vectors $\underline{\boldsymbol{g}}$ with dependent components, as shifting one target requires appropriate shifting of others. As a result, the sample covariance matrix $\hat{\mathbf{C}}^{gg}$ in (8) corresponding to the output $\underline{\boldsymbol{g}}$ is singular. Its rank is given by $\mathrm{rank}(\hat{\mathbf{C}}^{gg}) = D(T-1)$ with D being the number of dimensions. However, we need $\hat{\mathbf{C}}^{gg}$ to be regular as its inverse is used in the multi-target measurement update (10) and (11).

There are two options to handle this singularity. The first option is to investigate the sample mean $\hat{\underline{g}}$ in (9) and extract its $D(T-1)$ independent components. The second option is to perform a regularization of the sample covariance matrix $\hat{\mathbf{C}}^{gg}$ by adding a diagonal matrix with very small diagonal entries as $\hat{\mathbf{C}}^{gg} = \hat{\mathbf{C}}^{gg} + \epsilon \, \mathbf{I}$ with $\mathbf{I}$ an identity matrix of appropriate dimensions and ϵ a small constant. Both options showed identical results, so we selected the latter one for its simplicity.

[§]The symbol $u \stackrel{!}{=} v$ means "u shall be made equal to v".

6. APPLICATION EXAMPLE

We consider $T = 3$ targets, say aircraft, that move along piecewise straight line segments on a two-dimensional plane (D=2). Here, we have to introduce a time index k as we now consider target motion and recursive tracking. The initial (true) locations are

$$\underline{\tilde{x}}_{k,1} = \begin{bmatrix} -1.0 \\ -1.0 \end{bmatrix} \ , \ \underline{\tilde{x}}_{k,2} = \begin{bmatrix} -\sqrt{2} \\ -1.0 \end{bmatrix} \ , \ \underline{\tilde{x}}_{k,3} = \begin{bmatrix} -1.0 \\ 1.0 \end{bmatrix}$$

that are stacked to form the true multi-target state vector $\underline{\tilde{x}}_k = \left[(\underline{\tilde{x}}_{k,1})^\top, (\underline{\tilde{x}}_{k,2})^\top, (\underline{\tilde{x}}_{k,3})^\top\right]^\top \in \mathbb{R}^6$.

The true motion of the targets is given by a discrete-time dynamic system or motion model

$$\underline{\tilde{x}}_{k+1} = \underline{\tilde{x}}_k + \Delta \underline{x}_k$$

with a time-variant motion increment $\Delta \underline{x}_k$. The number of time steps S is set to $S = 500$. The motion of the three aircraft is shown in Fig. 1 (top). Their velocities are all equal so that they move side by side along the parallel motion segments.

We use the following probabilistic multi-target system model that relates random multi-target vectors between time steps

$$\underline{\boldsymbol{x}}_{k+1} = \underline{\boldsymbol{x}}_k + \Delta \underline{x}_k + \underline{\boldsymbol{w}}_k \ , \tag{13}$$

with $\underline{\boldsymbol{w}}_k$ zero-mean Gaussian noise with time-invariant covariance matrix $\mathbf{C}^w = (\sigma^w)^2 \operatorname{diag}\left(\left[1,1,1,1,1,1\right]^\top\right)$ and $\sigma^w = 0.03$.

Assuming an identity relation between targets and measurements, noisy measurements are generated as

$$\underline{\boldsymbol{y}}_{k,i} = \underline{\tilde{x}}_{k,i} + \underline{\boldsymbol{v}}_k$$

for $i = 1, 2, 3$ with $\underline{\boldsymbol{v}}_k$ zero-mean Gaussian noise with time-invariant covariance matrix $\mathbf{C}^v = (\sigma^v)^2 \operatorname{diag}\left(\left[1,1\right]^\top\right)$. For specific noise samples, we obtain three measurement samples per time step $\underline{\hat{y}}_{k,i}$ for $i = 1, 2, 3$. Now, we draw a random permutation $\pi(.)$ from all possible permutations to obscure the association of measurements to targets and obtain $\underline{\hat{y}}_{k,\pi(i)}$ for $i = 1, 2, 3$.

We apply the new multi-target tracker by alternating between multi-target measurement updates using the measurements at a given time step k and a multi-target time update by propagating the multi-target state estimate through the system model (13). Initial multi-target state estimates are set to the true locations with a diagonal covariance matrix with equal standard deviations of 0.05. The estimated target tracks for a low measurement noise level $\sigma^v = 0.02$ are depicted in Fig. 1 (bottom).

For comparison, a tracker with a two-step measurement update that first estimates a global association between target states and measurements and then uses a Kalman filter step to update the prior multi-target state estimate (global nearest neighbor tracker) is applied to the same measurement data. The time update step is the same and again based on the system model (13). The tracking results are shown in Fig. 1 (middle). Both trackers exhibit a similar performance.

We now consider a very high measurement noise level with $\sigma^v = 0.08$. The received measurements are depicted in Fig. 2 (top). The colors demonstrate the random obfuscation of the measurement indexes. Measurements received with index 1 are shown in red, with index 2 in green, and with index 3 in blue. It is obvious that this indexing scheme does not relate the measurements to their corresponding original targets. The results of the global nearest neighbor tracker are shown in Fig. 2 (middle), while the results of the proposed association-free multi-target tracker are shown in Fig. 2 (bottom). The results again are rather similar. Both trackers confuse the target identities after a while.

In the simulations, we only showed representative tracking results of the new multi-target tracker compared to a Kalman filter based tracker with global nearest neighbor (GNN) association. A more in-depth evaluation will be presented in a forthcoming paper.

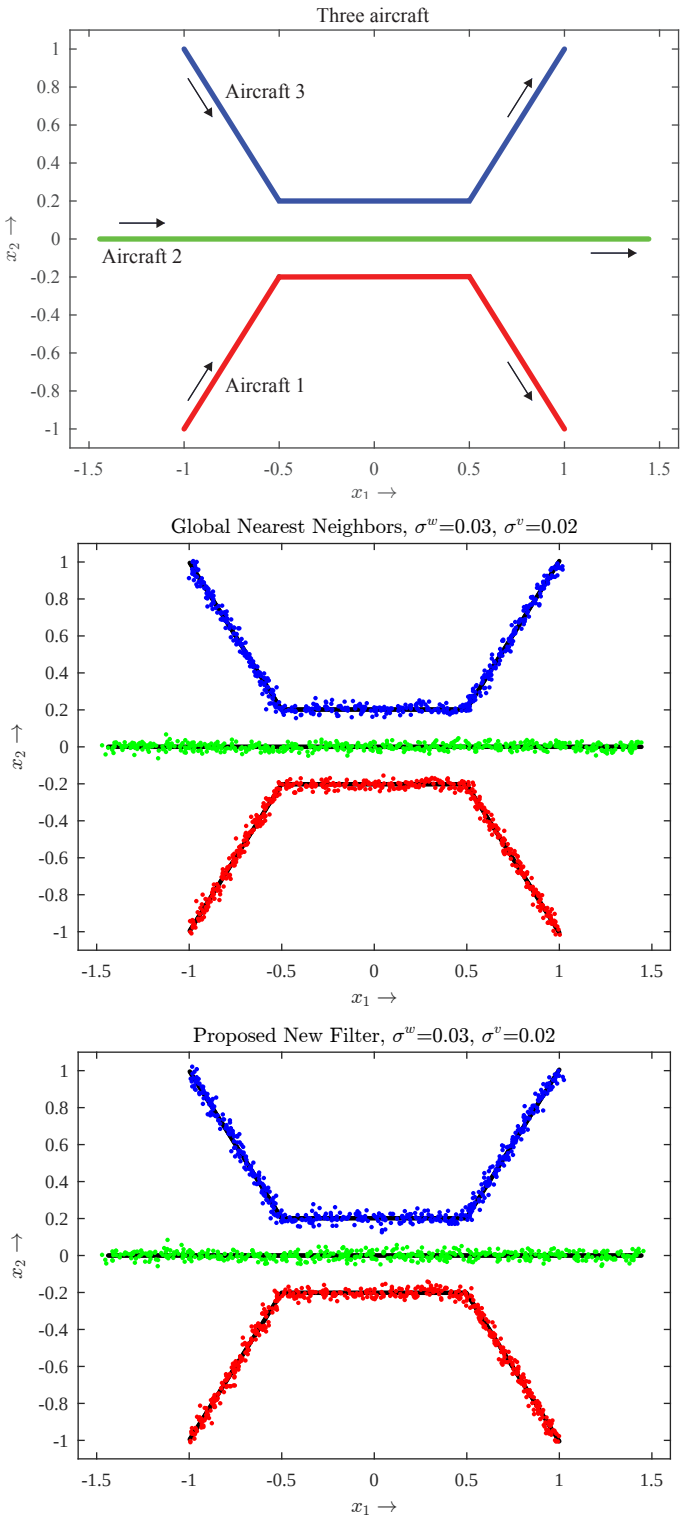

Figure 1. Three aircraft example, low measurement noise level. (top) Setup: Three aircraft flying along straight line segments. (middle) Tracking results of a tracker with a two-step measurement update comprising a global association step and a Kalman filter update. (bottom) Results of the proposed association-free tracker.

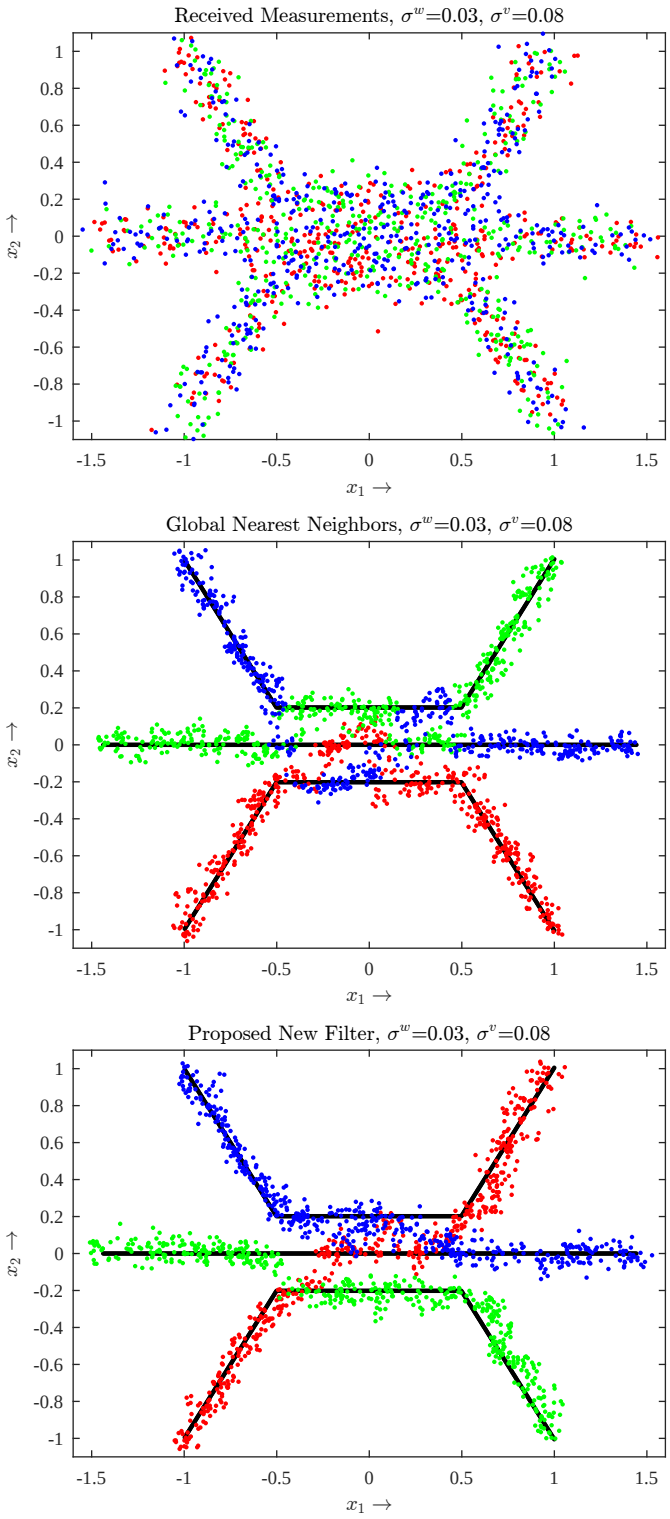

Figure 2. Three aircraft example, high measurement noise level. (top) Received measurements with colors indicating their random permutation. (middle) Tracking results of a tracker with a two-step measurement update comprising a global association step and a Kalman filter update. (bottom) Results of the proposed association-free tracker.

7. CONCLUSIONS

The proposed multi-target tracking approach for unknown target-measurement associations has a simple, elegant, and intuitive derivation, is straightforward to implement, and does not need tuning factors. The filter is very fast as the minimization of the association-free set distance measure is done simultaneously and implicitly by the filter. No explicit optimization is required. Hence, it shows a complexity in the number of targets and measurements determined by the complexity of the filter used (e.g., unscented Kalman filter). This is in contrast to association-based multi-target filters as these first perform an association step and then a filter step based on the association. The association step can be costly, e.g., for global nearest neighbor methods with many targets and measurements. As a result, the proposed multi-target filter is scalable and can be used for a large number of targets and measurements.

The simulations demonstrated that in the considered simple scenario, the proposed new association-free multi-target tracker yields results comparable to a Kalman filter based tracker with global nearest neighbor (GNN) association. However, the new tracker can handle complex scenarios that cannot simply be treated with a GNN-based tracker. More complex scenarios will be discussed in the next section.

8. FUTURE WORK

Enhancement of filter: The association-free multi-target filter proposed in Sec. 4 of this paper is a sample-based nonlinear Kalman filter that copes with the implicit nonlinear multi-target measurement equation (3) in Sec. 3. Although providing satisfactory results even for scenarios with relatively high noise, this type of filter is a strong approximation of a hypothetical optimal filter. Future work will be concerned with devising filters that provide a better approximation to the optimal filtering results. The challenge here is that no explicit likelihood can be derived from the multi-target measurement equation (3), so we are looking for likelihood-free filters.

An obvious approach to enhancing the filter quality by just slightly modifying the filter proposed in this paper is to employ smoothing, especially fixed-lag smoothing. The filter estimates the current multi-target state $\underline{\boldsymbol{x}}_k$ at time step k based on measurement vectors $\underline{\boldsymbol{y}}_0, \underline{\boldsymbol{y}}_1, \ldots, \underline{\boldsymbol{y}}_k, \ldots, \underline{\boldsymbol{y}}_{k+n}$, which comes at the expense of having a lag time n.

More complex scenarios: This paper considered the most basic multi-target tracking scenario, where the number of targets is known and at one time step exactly one measurement is obtained per target (detection probability of one) but the measurements are unlabeled. It was selected to simplify the explanation and forms the basis for approaching more complex scenarios:

- A slightly more complex scenario involves a time-varying number of unlabeled measurements per target with the numbers of measurements per target being equal. This scenario can be directly attacked with the method given in this paper. The same holds for nonlinear relations between targets and measurements occurring, e.g., in distance or angle measurements.
- A far more complex scenario arises when there are different (and possibly random) numbers of measurements per target at a given time step. This scenario will be approached next. It requires a modified set distance measure as targets now have to be updated based on regions of different measurement density.
- More realistic scenarios include false measurements, outliers, clutter, and missing measurements. This extension will be handled by adapting the approach for robustifying[28] the Kernel-SME filter[11] to the filter proposed in this paper.
- So far, we assumed the number of targets to be known and constant. In practice, the number of targets usually varies with time and is unknown. This problem will also be tackled by using set distance measures.

Batch estimation: Most of the work done in multi-target tracking focuses on recursive estimation. The multi-target state is propagated through time via a motion model and is recursively updated based on given measurements so that a new estimate is provided at every time step.

When real-time constraints are not an issue, measurements can be collected over a certain time and processed as a batch. Batch processing[29,30] has the advantage that all measurements are considered at once, e.g., later measurements are already used for estimating early target states. In addition, the entire tracks are optimized so that spatial constraints can more easily be considered. Closely related to batch processing is moving horizon estimation,[31,32] where a sliding horizon is considered for local batch processing.

The approach proposed in this paper and in Association-Free Direct Filtering[13] will be adapted to association-free batch tracking as the used point set distance measure is well suited for performing a global optimization of the tracks.

Acknowledgment

Marcus Baum was supported by the German Research Foundation (DFG) under grant BA 5160/1-1.

REFERENCES

[1] Ristic, B., Arulampalam, S., and Gordon, N., [*Beyond the Kalman Filter: Particle Filters for Tracking Applications*], Artech House Publishers (2004).

[2] Bar-Shalom, Y., Willett, P. K., and Tian, X., [*Tracking and Data Fusion: A Handbook of Algorithms*], YBS Publishing (2011).

[3] Blackman, S., "Multiple Hypothesis Tracking for Multiple Target Tracking," *IEEE Aerospace and Electronic Systems Magazine* **19**, 5–18 (Jan. 2004).

[4] Bar-Shalom, Y., Daum, F., and Huang, J., "The Probabilistic Data Association Filter," *IEEE Control Systems Magazine* **29**, 82–100 (December 2009).

[5] Svensson, L., Svensson, D., Guerriero, M., and Willett, P., "Set JPDA Filter for Multi-Target Tracking," *IEEE Transactions on Signal Processing* **59**(10) (2011).

[6] Mahler, R. P. S., [*Statistical Multisource-Multitarget Information Fusion*], Artech House, Boston, Mass. (2007).

[7] Vo, B.-T., Vo, B.-N., and Cantoni, A., "The Cardinality Balanced Multi-Target Multi-Bernoulli Filter and Its Implementations," *IEEE Transactions on Signal Processing* **57**, 409–423 (Feb. 2009).

[8] Reuter, S., Vo, B.-T., Vo, B.-N., and Dietmayer, K., "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing* **62**, 3246–3260 (July 2014).

[9] Kamen, E., "Multiple Target Tracking based on Symmetric Measurement Equations," *IEEE Transactions on Automatic Control* **37**, 371 –374 (March 1992).

[10] Kamen, E. and Sastry, C., "Multiple Target Tracking using Products of Position Measurements," *IEEE Transactions on Aerospace and Electronic Systems* **29**, 476 –493 (April 1993).

[11] Baum, M. and Hanebeck, U. D., "The Kernel-SME Filter for Multiple Target Tracking," in [*Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*], (July 2013).

[12] Kurz, G., Baum, M., and Hanebeck, U. D., "Real-time Kernel-based Multiple Target Tracking for Robotic Beating Heart Surgery," in [*Proceedings of the Eighth IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM 2014)*], (June 2014).

[13] Hanebeck, U. D. and Baum, M., "Association-Free Direct Filtering of Multi-Target Random Finite Sets with Set Distance Measures," in [*Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)*], (July 2015).

[14] Evensen, G., "The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation," *Ocean Dynamics* **53**(4), 343–367 (2003).

[15] Evensen, G., [*Data Assimilation: The Ensemble Kalman Filter*], Springer Science & Business Media (2009).

[16] Hanebeck, U. D., "Optimal Reduction of Multivariate Dirac Mixture Densities," *at - Automatisierungstechnik, Oldenbourg Verlag* **63**, 265–278 (Apr. 2015).

[17] Hanebeck, U. D., "Optimal Reduction of Multivariate Dirac Mixture Densities," *arXiv preprint: Systems and Control (cs.SY)* (Nov. 2014).

[18] Lefebvre, T., Bruyninckx, H., and De Schutter, J., "The Linear Regression Kalman Filter," in [*Nonlinear Kalman Filtering for Force-Controlled Robot Tasks*], *Springer Tracts in Advanced Robotics* **19** (2005).

[19] Wan, E. A. and Van Der Merwe, R., "The Unscented Kalman Filter for Nonlinear Estimation," in [*Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*], 153–158, IEEE (2000).

[20] Julier, S. J. and Uhlmann, J. K., "Unscented Filtering and Nonlinear Estimation," *Proceedings of the IEEE* **92**, 401– 422 (Mar. 2004).

[21] Steinbring, J. and Hanebeck, U. D., "S2KF: The Smart Sampling Kalman Filter," in [*Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*], (July 2013).

[22] Steinbring, J. and Hanebeck, U. D., "LRKF Revisited: The Smart Sampling Kalman Filter (S2KF)," *Journal of Advances in Information Fusion* **9**, 106–123 (Dec. 2014).

[23] Steinbring, J., Pander, M., and Hanebeck, U. D., "The Smart Sampling Kalman Filter with Symmetric Samples," *Journal of Advances in Information Fusion* **11**, 71–90 (June 2016).

[24] Julier, S., Uhlmann, J., and Durrant-Whyte, H. F., "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators," *IEEE Transactions on Automatic Control* **45**, 477–482 (Mar. 2000).

[25] Julier, S. and Uhlmann, J., "Reduced Sigma Point Filters for the Propagation of Means and Covariances Through Nonlinear Transformations," in [*Proceedings of the 2002 American Control Conference*], **2**, 887–892 vol.2 (2002).

[26] Golub, G. H. and Van Loan, C. F., [*Matrix Computations (3rd Ed.)*], Johns Hopkins University Press, Baltimore, MD, USA (1996).

[27] Hanebeck, U. D., Huber, M. F., and Klumpp, V., "Dirac Mixture Approximation of Multivariate Gaussian Densities," in [*Proceedings of the 2009 IEEE Conference on Decision and Control (CDC 2009)*], (Dec. 2009).

[28] Baum, M., Yang, S., and Hanebeck, U. D., "The Kernel-SME Filter with False and Missing Measurements," in [*Proceedings of the 19th International Conference on Information Fusion (Fusion 2016)*], (July 2016).

[29] Smith, K., Carleton, A., and Lepetit, V., "General Constraints for Batch Multiple-Target Tracking Applied to Large-Scale Videomicroscopy," in [*2008 IEEE Conference on Computer Vision and Pattern Recognition*], 1–8 (June 2008).

[30] Vu, T., Vo, B. N., and Evans, R., "A Particle Marginal Metropolis-Hastings Multi-Target Tracker," *IEEE Transactions on Signal Processing* **62**, 3953–3964 (Aug. 2014).

[31] Rawlings, J. B. and Bakshi, B. R., "Particle Filtering and Moving Horizon Estimation," *Computers & Chemical Engineering* **30**, 1529–1541 (Sept. 2006).

[32] Rawlings, J. B., "Moving Horizon Estimation," in [*Encyclopedia of Systems and Control*], Baillieul, J. and Samad, T., eds., 1–7, Springer London (2013). DOI: 10.1007/978-1-4471-5102-9_4-1.