



CHALMERS

Chalmers Publication Library

Set JPDA Filter for Multitarget Tracking

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

IEEE Transactions on Signal Processing (ISSN: 1053-587X)

Citation for the published paper:

Svensson, L. ; Svensson, D. ; Guerriero, M. (2011) "Set JPDA Filter for Multitarget Tracking". IEEE Transactions on Signal Processing, vol. 59(10), pp. 4677-4691.

<http://dx.doi.org/10.1109/TSP.2011.2161294>

Downloaded from: <http://publications.lib.chalmers.se/publication/146503>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Set JPDA Filter for Multi-Target Tracking

Lennart Svensson, Daniel Svensson*, Marco Guerriero[†]
and Peter Willett[‡], *Fellow, IEEE*

Abstract

In this article we show that when targets are closely spaced, traditional tracking algorithms can be adjusted to perform better under a performance measure that disregards identity. More specifically, we propose an adjusted version of the Joint Probabilistic Data Association (JPDA) filter, which we call the Set JPDA (SJPDA) filter. Through examples and theory we motivate the new approach, and show its possibilities. To decrease the computational requirements, we further show that the SJPDA filter can be formulated as a continuous optimization problem which is fairly easy to handle. Optimal approximations are also discussed, and an algorithm, KLSJPDA, which provides optimal Gaussian approximations in the Kullback-Leibler sense is derived. Finally, we evaluate the SJPDA filter on two scenarios with closely spaced targets, and compare the performance in terms of the mean Optimal Subpattern Assignment (MOSPA) measure with the JPDA filter, and also with the Gaussian-mixture CPHD filter. The results show that the SJPDA filter performs substantially better than the JPDA filter, and almost as well as the more complex GM-CPHD filter.

Index Terms

Target tracking, filtering theory, random finite set theory, recursive estimation, Bayes methods.

*Lennart Svensson and Daniel Svensson are with the Department of Signals and Systems, Chalmers University of Technology, SE-412 96, Gothenburg, Sweden (E-mail: {lennart.svensson,daniel.svensson}@chalmers.se).

[†]Marco Guerriero is with Research and Advanced System Design, Elt. Elettronica, S.p.A., 00 131 Rome, Italy (E-mail: marco.guerriero@elt.it).

[‡]Peter Willett is with the Electrical and Computer Engineering Department, University of Connecticut, Storrs, CT 06269-2157, USA (E-mail: willett@engr.uconn.edu).

1 Introduction

Traditional target tracking algorithms are designed to track targets over time, and discriminate between them by labeling them at each time instant. Examples of such algorithms are the Probabilistic Data Association (PDA) and Joint PDA (JPDA) filters [1,2], and the Multiple Hypothesis Tracking (MHT) algorithm [3–6]. In many applications, such as those in which evidence of target type is being accumulated, the identity (labeling) of the targets is of great importance; but in other cases where one might say that "a threat is a threat", it is not. In this article we show that by not considering target identity, the traditional algorithms can be significantly improved both in terms of density approximations and of estimation accuracy, when evaluated with a metric that disregards target identity.

For the traditional algorithms, the aim is to minimize the Mean Square Error (MSE) between target states and corresponding track estimates. When target identity is not of interest, minimization under such a measure subjects itself to an unnecessary constraint. Instead we need a measure which only describes how good an estimate is in determining where the targets are, regardless of which is which. In this article, we propose the use of the Optimal Subpattern Assignment (OSPA) metric [7]. More specifically, we study the Mean OSPA (MOSPA). The tracking problem can thus be formulated as the minimization of the MOSPA which, as we will see, is significantly different from the problem of minimizing the MSE.

A framework suitable for the description of tracking without target identity is Finite Set Statistics (FISST) [8]. Within that framework, the aim is to track an unordered, or unlabeled, set of targets, described as a Random Finite Set (RFS). Two popular algorithms have been derived within FISST, namely the Probability Hypothesis Density (PHD) and the Cardinalized PHD (CPHD) filters [9–12]. We notice two drawbacks with FISST-based techniques. First, it is difficult to find an analytical expression for the posterior density of the RFS. Therefore, the standard technique of calculating and approximating the posterior density is not applicable. Instead, the filters operate on a first-order approximation of the density of unordered targets, referred to as the intensity function. Second, the first-order approximation has the effect that all targets be assumed independent and identically distributed. Due to these drawbacks, we believe that it is relevant to search for alternative approaches. In the current article, one such alternative is suggested, where the advantages of FISST methods and those of classic recursive filtering are combined, and where it is shown how traditional algorithms can be improved for the problem of minimizing MOSPA.

The approach that we propose relies on the fact that there is a relation

between an ordered density and the density of a random finite set, here called the unordered density, or set density. In fact, there is an infinite number of ordered densities that correspond to the same set density. We refer to a set of such ordered densities as an RFS family. Since optimal MOSPA estimates can be derived from the set density, any ordered density in the corresponding RFS family can be used without losing optimality. We thus have a previously unrecognized possibility of switching ordered densities within traditional tracking filters, to obtain better approximations and performance, when target identity is irrelevant. This is useful since the different densities in the RFS family are more or less suitable to use in practice.

Most tracking algorithms rely on Gaussian approximations. That is, the posterior density is often described as a Gaussian mixture, where the number of mixture components is controlled. In the JPDA filter, which is the filter that we specifically consider in this article, the posterior density is approximated by a single Gaussian in each iteration. By combining traditional methods and FISST techniques, we show that this Gaussian approximation can be made more accurate, when target identity is not of interest. The approach is to utilize the possibility of switching from the original Gaussian mixture, to another density in the same RFS family that can be better approximated as Gaussian. We further show that there is an optimal way, in Kullback-Leibler sense, of making that density switch. From this, we develop a refined JPDA algorithm for optimal approximations, which we call Kullback-Leibler Set JPDA (KLSJPDA). The drawback with the algorithm is its computational demand. We thus also develop a more computationally efficient algorithm, called the Set JPDA (SJPDA). The SJPDA algorithm, which is also based on density switches, optimizes MOSPA performance under the constraint of remaining within the family of Gaussian mixtures, and this generally also leads to a density which is better approximated as Gaussian. Through examples and theory we discuss the reason for that. The details about the optimization criterion for the SJPDA algorithm, and its relation to optimal MOSPA estimates, are found in [13]. Throughout the article, we make the assumption that the number of targets is known.

The article is outlined as follows. In Section II, the problem formulation is stated. Section III concerns two conceptual solutions to the formulated problem, based on traditional approaches and FISST, respectively. In Section IV, a conceptual solution based on the new approach is presented. We also discuss how and why traditional tracking methods can be improved when MOSPA is the target cost function. Section V considers optimal Gaussian approximations in the Kullback-Leibler sense, and the derivation of the KLSJPDA algorithm. In Section VI, the SJPDA filter is derived, and the optimization step is formulated as a continuous optimization problem. Section

VII regards evaluations of the proposed KLSJPDA and SJPDA algorithms. For the SJDPA algorithm, we see that the tracking performance in terms of MOSPA is better than that of the JPDA filter, and almost as good as that of the more complex Gaussian mixture CPHD filter [12]. We also see that the track loss probability is dramatically decreased, which leads to a much longer track life. For the KLSJPDA algorithm we observe similar performance as the SJPDA filter, but the computational complexity is much larger. In Section VIII the article is concluded.

2 Problem formulation

2.1 Target modeling assumptions

In this article we study the problem of tracking an arbitrary, but known and fixed, number of targets. Further, we are not interested in the identities of the targets. To reflect this, we are using a different cost function from the standard formulation, in which MSE is used.

To formulate the problem, we introduce the vector of ordered (labeled) target states

$$\mathbf{X}_k = \left[(\mathbf{x}_k^{(1)})^T \ (\mathbf{x}_k^{(2)})^T \ \dots \ (\mathbf{x}_k^{(n)})^T \right]^T, \quad (1)$$

where $\mathbf{x}_k^{(i)}$ is the state vector of target number i at time k , and where n is the number of targets in the scene. We further introduce the collection of measurements \mathbf{Z}^k ,

$$\mathbf{Z}^k = \{\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_k\}, \quad (2)$$

up to the current time step k , where \mathbf{Z}_k is a matrix of measurement vectors at time k .

For the general case, the process model is governed by

$$\mathbf{x}_k = f_{k-1}(\mathbf{x}_{k-1}, \mathbf{v}_{k-1}), \quad (3)$$

where f_{k-1} is a nonlinear (prediction) function and \mathbf{v}_{k-1} is a realization of the process noise, which has some assumed distribution. As seen, we sometimes omit the superscript on the state vectors, and write \mathbf{x}_k instead of $\mathbf{x}_k^{(i)}$. The corresponding linear-Gaussian motion model is written as

$$\mathbf{x}_k = \mathbf{F}_{k-1}\mathbf{x}_{k-1} + \mathbf{v}_{k-1}, \quad (4)$$

where $\mathbf{v}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{k-1})$.

The general measurement model is given by

$$\mathbf{z}_k = h_k(\mathbf{x}_k, \mathbf{w}_k) \quad (5)$$

and the corresponding linear-Gaussian model by

$$\mathbf{z}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{w}_k, \quad (6)$$

where $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$ is measurement noise. The measurement set \mathbf{Z}_k does not only include target-generated measurements, but also spurious measurements due to false alarms and clutter. If we let m_k be the number of target-generated measurements (smaller than, or equal to, n), and c_k be the number of clutter measurements at time k , the measurement set is given by

$$\mathbf{Z}_k = \{\mathbf{z}_k^{(1)}, \dots, \mathbf{z}_k^{(m_k)}, \mathbf{z}_k^{c,(1)}, \dots, \mathbf{z}_k^{c,(c_k)}\}. \quad (7)$$

The target-generated measurements, $\mathbf{z}_k^{(1)}, \dots, \mathbf{z}_k^{(m_k)}$, are governed by (5) or (6), and the clutter measurements, $\mathbf{z}_k^{c,(1)}, \dots, \mathbf{z}_k^{c,(c_k)}$, follow a density function which is target-state independent.

2.2 MOSPA measure and optimal estimates

The ubiquitous measure in the literature is the squared error (SE)

$$\begin{aligned} \text{SE}(\hat{\mathbf{X}}_k, [\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n_k)}]) \\ = \left((\hat{\mathbf{x}}_k^{(1)} - \mathbf{x}_k^{(1)})^2 + \dots + (\hat{\mathbf{x}}_k^{(n_k)} - \mathbf{x}_k^{(n_k)})^2 \right). \end{aligned} \quad (8)$$

To evaluate the measure, we need estimates of the states of target 1, 2, and so on, i.e., of the labeled targets. Since this paper is about describing where there are targets, rather than where a target with a certain label is, the squared error is not a good measure. We note for example in the two-target case, that if the identities of the targets have been mixed up, the squared error can be very large, even though there might be two accurate tracks available, i.e., even though $\hat{\mathbf{x}}_k^{(1)} \approx \mathbf{x}_k^{(2)}$ and $\hat{\mathbf{x}}_k^{(2)} \approx \mathbf{x}_k^{(1)}$. Therefore, we seek a measure that can capture the quality of an algorithm to estimate the set of targets.

Several multi-target performance measures have been proposed in the literature. An ad-hoc optimal assignment-based approach, with arbitrary cost function, was given in [14], while the first rigorous theory of multi-object distances was given in [15]. The measure, which is based on the optimal assignment approach, is called the Optimal Mass Transfer (OMAT)

metric. When the number of targets is known, the OMAT metric is identical to the optimal assignment procedure. OMAT has a number of weaknesses, discussed in [7]. Apart from not being a metric, the major weaknesses appear when the number of targets is not known and the estimated number is not always equal to the true number. As a remedy, a new measure called the Optimal Subpattern Assignment (OSPA) metric, was proposed in [7]. Since OSPA both is a true metric and an intuitively appealing measure which has received much attention of late, we use it as the basis measure for the problem.

Let \mathbf{X} be the set of true target states and $\widehat{\mathbf{X}}$ be the set of target estimates, in our case both with n elements. The OSPA measure $\tilde{d}_p^{(c)}$ is then defined as

$$\tilde{d}_p^{(c)}(\widehat{\mathbf{X}}, \mathbf{X}) = \left(\frac{1}{n} \left(\min_{\pi \in \Pi_n} \sum_{i=1}^n d^{(c)}(\mathbf{x}^{(i)}, \hat{\mathbf{x}}^{\pi(i)})^p \right) \right)^{1/p}. \quad (9)$$

Here, $d^{(c)}(\mathbf{x}, \hat{\mathbf{x}}) \triangleq \min(c, d(\mathbf{x}, \hat{\mathbf{x}}))$ is the distance d between \mathbf{x} and $\hat{\mathbf{x}}$, cut-off at c . Further, Π_n is the set of all possible permutations of $\widehat{\mathbf{X}}$. The notation $\hat{\mathbf{x}}^{\pi(i)}$ describes the i^{th} permutation (re-ordering) of the vector $\hat{\mathbf{x}}$. In this article, we let d be the Euclidean distance, and we use a quadratic measure ($p = 2$). In practice, the measure performs an optimal assignment of target estimates to true target states, possibly clamped at c .¹

To describe the performance of an estimator, and to have a measure for which we can define an optimal algorithm, we average over all state vectors, which gives us a definition of the *mean* OSPA (MOSPA) measure

$$\text{MOSPA}_p^{(c)}(\widehat{\mathbf{X}}) \triangleq \mathbb{E}_{p(\mathbf{x}|\mathbf{z}^k)} \{ \tilde{d}_p^{(c)} \}. \quad (10)$$

An optimal estimator, in the MOSPA sense, is an estimator which minimizes the MOSPA measure. Such an estimator is referred to as a minimum MOSPA (MMOSPA) estimator².

2.3 Motivation of the problem and the MOSPA measure

The problem that we study in this paper is the problem of estimating the unordered set of targets, for which the MMOSPA estimator is optimal. We

¹For $c = \infty$, the above measure $\tilde{d}_p^{(c)}$ is equal to the OMAT metric. A consequence of this relation, using results from [15], is that the OSPA measure for known target numbers, and $c = \infty$, reduces to the optimal assignment approach, presented earlier.

²Note that the relation between OSPA, MOSPA and MMOSPA is analogous to the relation between the common acronyms SE, MSE and MMSE.

here give three examples of when minimizing MOSPA provides a solution which is more reasonable than what is obtained when minimizing the Mean Square Error (MSE), using an ordered density.

The first application is in radar cueing, i.e., for the problem of steering a radar sensor to areas of high target existence probability. In such cases, there is no interest in which target is which, but the question of where there are targets is very important. The risk with using an ordered density for this application is that if there is an uncertainty in the labeling of the targets, the posterior density will be multimodal. Then, there is a high probability that the mean value is in an area of low target existence probability. Steering the main lobe of the sensor to such an area is thus prone to low probability of true-target returns. By disregarding ordering, the posterior density can be made less multimodal, and the steering of the main lobe can then with higher probability be directed to an area where targets are likely to be.

A second example is in the automotive industry: In collision avoidance systems, it is not of interest to know which car is which—the only interest is to avoid all cars. Finally, when tracking extended objects using radar measurements, those objects are often described by a set of reflectors. The problem of interest is then to track that set of reflectors, and not try to distinguish which reflector is which.

3 Conceptual solutions and the JPDA approximation

For the case of target tracking, there are two well-studied optimal, or conceptual, solutions, and these will be discussed in this section. By presenting these conceptual solutions, we believe that it is easier to understand the new approach of this paper, which is based on a third conceptual solution introduced in Section 4.

The basis of both conventional solutions is first to calculate an optimal description of the joint target density, and then to derive MMOSPA estimates from the optimal description. In the description of the first conceptual solution, we also describe the data association problem, and how it is solved in the conventional frameworks. At the end, we describe the JPDA approximation to the first conceptual solution.

3.1 Conceptual solution I – ordered densities

The first conceptual solution is to use the traditional approach of first calculating the ordered posterior density $p(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n_k)} | \mathbf{Z}^k)$, and then to derive

MMOSPA state estimates from it. For the conventional problem of minimizing MSE, the traditional approach is the foundation for methods such as JPDA and MHT. The first conceptual solution to the considered problem is presented in Table 1.

Table 1: Conceptual solution I – ordered densities

1)	Compute the ordered density $p(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)} \mathbf{Z}^k)$ recursively.
2)	Derive MMOSPA state estimates from the ordered density.

For an ordered density with known number of targets, n , the state vectors of the targets can be stacked in a long vector \mathbf{X}_k as shown in (1). To compute MMOSPA estimates, we seek to calculate the posterior density $p(\mathbf{X}_k | \mathbf{Z}^k)$. By marginalizing over all global data association hypotheses, \mathcal{H} , the density is given by

$$p(\mathbf{X}_k | \mathbf{Z}^k) = \sum_{h \in \mathcal{H}} p(\mathbf{X}_k | h, \mathbf{Z}^k) \Pr \{h | \mathbf{Z}^k\} \quad (11)$$

where the densities $p(\mathbf{X}_k | h, \mathbf{Z}^k)$ are easy to express. A global data association hypothesis, h , determines which measurements are clutter and which are target-generated, and in the latter case their target of origin. If the process and measurement models are linear and Gaussian, the posterior density is a Gaussian mixture, with increasingly many components over time. Thus, the optimal data association solution is practically infeasible, and sub-optimal solutions are required. One such solution is given by the JPDA filter, which we describe in Section 3.3.

Apart from being computationally infeasible, there is also a second difficulty with the conceptual solution, namely that it is not obvious how estimates with low MOSPA should be derived from the ordered density. As we show in Example 1 in Section 4.2, the posterior mean, which is normally the estimate used in the traditional problem of minimizing MSE, might not at all be a suitable estimate in MOSPA sense.

3.2 Conceptual solution II – unordered densities

The second conceptual solution relies on the fact that to compute MMOSPA estimates, it is sufficient to know the unordered posterior density $p(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\} | \mathbf{Z}^k)$, since the labeling of the targets has no influence on the MOSPA measure. An optimal solution utilizing the sufficiency of the unordered density is given in Table 2.

Table 2: Conceptual solution II – unordered densities

1)	Compute the unordered density $p(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\} \mathbf{Z}^k)$ recursively.
2)	Derive MMOSPA state estimates from the unordered density.

The second conceptual solution has the same two difficulties as the first one, viz. that approximations are necessary to keep complexity at a constant level and that derivation of MMOSPA estimates from an unordered density is a research topic in itself.

Step one of the second conceptual solution (cf. Table 2) is the foundation of the family of PHD/CPHD algorithms [9–12] with the Gaussian-mixture CPHD filter [12] as perhaps the most prominent one.

3.3 JPDA approximation

The JPDA filter is an approximative solution to the data association problem, described in the spirit of the first conceptual solution. The approach of the filter is to recursively approximate the multi-modal posterior density in (11) by a single Gaussian. The posterior density can be rewritten as

$$p(\mathbf{X}_k | \mathbf{Z}^k) = \sum_{h \in \mathcal{H}} \frac{p(\mathbf{Z}_k | \mathbf{X}_k, h, \mathbf{Z}^{k-1}) p(\mathbf{X}_k | \mathbf{Z}^{k-1})}{p(\mathbf{Z}_k | \mathbf{Z}^{k-1})} \times \Pr \{h | \mathbf{Z}^{k-1}\}. \quad (12)$$

Assuming linear and Gaussian models, and a Gaussian prior density $p(\mathbf{X}_k | \mathbf{Z}^{k-1})$ (given by the JPDA approximation at time index $k - 1$), the above density is a Gaussian mixture.

The JPDA filter performs the following steps:

1. Formulate all global data association hypotheses, \mathcal{H} , which describe possible origins of \mathbf{Z}_k .
2. For each data association hypothesis $h \in \mathcal{H}$, update the predicted density for each target i with the assigned measurement j as a Kalman filter update. The output of the filter is the mean value $\mathbf{x}_{k|k}^{i,h}$ and covariance matrix $\mathbf{P}_{k|k}^{i,h}$ (cf. (18)–(19)).
3. Calculate the weight β_h of each mixture component h .

4. Use moment matching to approximate the Gaussian mixture by a single Gaussian, see (17).

The mixture weights β_h are given by

$$\beta_h = \frac{\bar{\beta}_h}{\sum_{h=1}^{N_{\mathcal{H}}} \bar{\beta}_h} \quad (13)$$

$$\bar{\beta}_h = \prod_{\mathcal{S}_0^h} \lambda \cdot \prod_{\mathcal{S}_u^h} (1 - P_D) \cdot \prod_{\{i,j\} \in \mathcal{S}_a^h} P_D g_{ij}, \quad (14)$$

where a constant detection probability is assumed, \mathcal{S}_0^h is the set of unassigned measurements, \mathcal{S}_u^h is the set of unassigned targets, and \mathcal{S}_a^h is a set including the pairs of detected targets, i , and their assigned measurements, j . Further,

$$g_{ij} = \mathcal{N}(\boldsymbol{\nu}_{ij}; \mathbf{0}, \mathbf{S}_k) = \frac{1}{|2\pi\mathbf{S}_k|^{1/2}} e^{-\frac{d_{ij}^2}{2}} \quad (15)$$

$$d_{ij}^2 = \boldsymbol{\nu}_{ij}^T \mathbf{S}_k^{-1} \boldsymbol{\nu}_{ij}, \quad \boldsymbol{\nu}_{ij} = \mathbf{z}_k^j - \mathbf{H}_k \mathbf{x}_{k|k-1}^{(i)}. \quad (16)$$

As described above, the first three steps of the JPDA algorithm describe the calculation of the components in the Gaussian mixture, while the final step is to approximate that Gaussian mixture density for each target with a single Gaussian. This is done by moment matching, i.e.,

$$p(\mathbf{x}_k^{(i)} | \mathbf{Z}^k) \cong \mathcal{N}(\mathbf{x}_k^{(i)}; \mathbf{x}_{k|k}^i, \mathbf{P}_{k|k}^i) \quad (17)$$

where³

$$\mathbf{x}_{k|k}^i = \sum_{h=1}^{N_{\mathcal{H}}} \beta_h \mathbf{x}_{k|k}^{i,h} \quad (18)$$

$$\mathbf{P}_{k|k}^i = \sum_{h=1}^{N_{\mathcal{H}}} \beta_h \left\{ \mathbf{P}_{k|k}^{i,h} + (\mathbf{x}_{k|k}^i - \mathbf{x}_{k|k}^{i,h}) (\mathbf{x}_{k|k}^i - \mathbf{x}_{k|k}^{i,h})^T \right\} \quad (19)$$

and where $N_{\mathcal{H}}$ is the total number of hypotheses.

The JPDA algorithm is often described in an alternative, but equivalent, fashion [5], where the computation of the state estimates includes the calculation of a weighted measurement residual which is used in an ordinary Kalman filter update.

³Although the sums in (18) and (19) can be done as written, in practice there would be a step of marginalization over the single-target association events.

4 Conceptual solution III – ordered densities with switching

4.1 The RFS family

In this section we discuss the relationship between ordered and unordered densities, and the effects and possibilities of that relationship. Further, we introduce a new, third, conceptual solution to the problem of interest, and illustrate its properties using simple examples.

The first two conceptual solutions included the calculation of the ordered, and unordered, posterior densities of the joint target states, respectively. A key insight, upon which we capitalize, is that there is a relation between these densities. For n targets, the relation is the following

$$\begin{aligned} p(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\} = \{\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_n\}) \\ = \sum_{i=1}^{n!} p(\mathbf{x}_k^{(1)} = \boldsymbol{\alpha}_{m_1^i}, \dots, \mathbf{x}_k^{(n)} = \boldsymbol{\alpha}_{m_n^i}), \end{aligned} \quad (20)$$

where $[\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_n]$ is a point in the joint target state space, and m_j^i , for $j = 1, \dots, n$, is index j in permutation number i . To go from an ordered density to a set density, we thus sum over all possible permutations of the state vector. For instance, when $n = 2$ a natural choice is to set $m_1^1 = 1$, $m_2^1 = 2$ and $m_1^2 = 2$ and $m_2^2 = 1$. One important consequence of this relation is described in the following proposition.

Proposition 1 *For $n > 1$, the mapping from densities of ordered state vectors, $p_i(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)})$, to RFS densities, $p(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\})$, is many-to-one.*

Since many ordered densities correspond to the same unordered density, and since the RFS density is sufficient to derive optimal estimates, it is fair to say that the ordered density contains more information than necessary. For all densities which correspond to the same unordered density, we make the following definition:

Definition 1 *When two labeled densities, $p_1(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)})$ and $p_2(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)})$, correspond to the same RFS density, we say that they belong to the same RFS family.*

So, using (20), we obtain the same RFS density regardless if $p_1(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)})$ or $p_2(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)})$ is used. In Fig. 1, two labeled densities in the same RFS

family are shown, together with the RFS density. Obviously, even though the ordered densities belong to the same RFS family, their shape, expected values and covariance matrices can be very different.

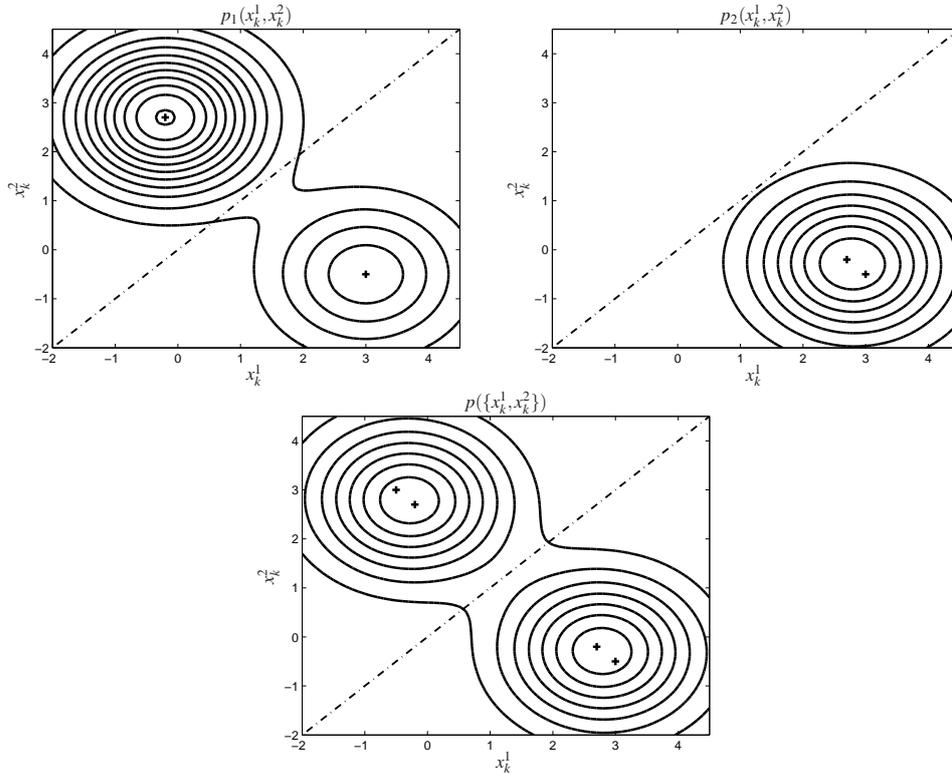


Figure 1: Marginalized posterior densities before [top left] and after [top right] switching the indexes under \mathcal{H}_2 , and the RFS density [bottom]. In our opinion, a Gaussian approximation is much more appropriate in the middle figure than in the left one. The symmetry line is dash-dotted and the pluses indicate expected values of the Gaussian mixture components. The RFS density is inherently symmetric and is not affected by label switches in the ordered densities.

The MMOSPA estimate can be calculated from the RFS density, but also from the ordered densities. Since Proposition 1 tells us that several ordered densities correspond to the same RFS density, all these densities should result in the same MMOSPA estimates. This assumes, of course, that the estimates are computed using optimal algorithms. To find the optimal estimates, we can first calculate the RFS density and then find the MMOSPA estimate from it using optimal algorithms. All ordered densities which correspond to the same RFS density must hence yield the same MMOSPA estimate. We

summarize this in a proposition.

Proposition 2 *All ordered densities within an RFS family yield the same MMOSPA estimate. Further, that MMOSPA estimate is equal to the MMO-SPA estimate obtained from the corresponding unordered density.*

So if we have the RFS family at a certain time, we can use any of the densities in the family to find the MMOSPA estimate, or we can use the corresponding RFS density. But what happens when we update the ordered densities with the same data? Will they still correspond to the same RFS family? With the following proposition, and its proof, we show that it is actually so.

Proposition 3 *Suppose $p_1(\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(n)})$ and $p_2(\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(n)})$ are two labeled densities within the same RFS family⁴. When these densities are updated using the same set of measurements, \mathbf{Z}_k , the updated densities still belong to the same RFS family.*

Proof of proposition 3 *The key to this result is the fact that given the state vector, the likelihood is the same for all terms in (20), i.e., for all permutations of the target positions. Intuitively, this means that once the information about the target identities is lost, it can not be recovered from future data. Consequently, after the Bayesian update we get $p_1(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\} = \{\alpha_1, \dots, \alpha_n\} | \mathbf{Z}_k)$*

$$= \sum_{i=1}^{n!} \left[\frac{p(\mathbf{Z}_k | \mathbf{x}_k^{(1)} = \alpha_{m_1^i}, \dots, \mathbf{x}_k^{(n)} = \alpha_{m_n^i})}{p(\mathbf{Z}_k)} \times p_1(\mathbf{x}_k^{(1)} = \alpha_{m_1^i}, \dots, \mathbf{x}_k^{(n)} = \alpha_{m_n^i}) \right] \quad (21)$$

$$= \frac{p(\mathbf{Z}_k | \mathbf{x}_k^{(1)} = \alpha_1, \dots, \mathbf{x}_k^{(n)} = \alpha_n)}{p(\mathbf{Z}_k)} \times \sum_{i=1}^{n!} p_1(\mathbf{x}_k^{(1)} = \alpha_{m_1^i}, \dots, \mathbf{x}_k^{(n)} = \alpha_{m_n^i}) \quad (22)$$

$$= p_2(\{\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n)}\} = \{\alpha_1, \dots, \alpha_n\} | \mathbf{Z}_k), \quad (23)$$

⁴We also assume that all targets have identical properties conditioned on the state vector, i.e., that the measurement equations are the same, independent of the target number. If this is not so, we have no business using the unordered densities in the first place.

where the last equality is due to $p_1(\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(n)})$ and $p_2(\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(n)})$ being part of the same RFS family.

As we can see, all labeled densities that correspond to the same RFS density, $\frac{1}{n!} \sum_{i=1}^{n!} p(\mathbf{x}_k^1 = \boldsymbol{\alpha}_{m_1^i}, \dots, \mathbf{x}_k^n = \boldsymbol{\alpha}_{m_n^i})$ before measurement update, still belong to the same (only updated) RFS density, $p(\{\mathbf{x}_k^1 = \boldsymbol{\alpha}_1, \dots, \mathbf{x}_k^n = \boldsymbol{\alpha}_n\} | \mathbf{Z}_k)$, after the measurement update.

To phrase this differently, all densities that belong to the same RFS family should result in the same MMOSPA estimates, both now and for all future times, i.e., also when new data is available. We can therefore replace our density at hand with any other density within the RFS family, without influencing the estimates⁵. We conclude our findings in a theorem.

Theorem 1 *In a recursive filtering framework, at each time instant we have the possibility of switching between ordered densities in the RFS family without affecting current or future MMOSPA estimates.*

From Theorem 1, we see that there is a previously unrecognized possibility of switching densities at any time in a filtering framework, without affecting the optimal performance. Based on this, we can describe a third conceptual solution (see Table 3) to the problem of this paper—a conceptual solution that utilizes the possibility of switching densities within the RFS family.

Table 3: Conceptual solution III – ordered densities with switching

1)	Compute the ordered density $p(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n_k)} \mathbf{Z}^k)$ recursively.
2)	Replace $p(\mathbf{x}_k^{(1)}, \dots, \mathbf{x}_k^{(n_k)} \mathbf{Z}^k)$ with another ordered density which corresponds to the same RFS density.
3)	Derive MMOSPA state estimates from the ordered density.

This conceptual solution provides us with new possibilities and tools for designing novel suboptimal algorithms with better performance. As a switch within the RFS family does not affect the estimates, we could for instance make a change to a density which is more accurately approximated with a Gaussian density. The relevance of that is obvious, considering for example the JPDA filter, which relies heavily on Gaussian approximations. Since

⁵It is important, to go between (21) and (22), as with other steps in our development, that the targets be truly exchangeable and indistinguishable. For example, if target 1 has a different RCS than target 2, then this cannot be claimed.

computational complexity must be kept at a constant level, all practical algorithms need approximations, and Gaussian approximations are common. In the RFS family, there may be a density which much more resembles a Gaussian density. An example of that is seen in Fig. 1. By switching to that density, a more accurate Gaussian approximation is possible.

In practice, the posterior mean is often used as estimate also when the cost function is not the squared error. With the same reasoning as above, there may thus also exist densities within the RFS family whose expected value is closer to the optimal estimate, than the expected value of the original density. In fact, in [13] it is shown that there is a density in the RFS family whose mean value is the MMOSPA estimate. In Section 4.2, we study an example which shows that a density switch can lead to both improved approximations and to improved state estimates when the posterior mean is the estimator.

4.2 An example of the use of an RFS family

By considering an example, we here illustrate and discuss the potential benefits of density switches, both in terms of accuracy of the Gaussian approximations and the accuracy of the state estimates, when the posterior mean is used as estimate.

Example 1: Consider a scenario where we have two Gaussian distributed targets. The example is illustrated in Fig. 1. The probability of detection is one for both targets, and we have received two detections. We represent the two possible data associations by the hypotheses \mathcal{H}_1 and \mathcal{H}_2 . Now, suppose that calculations yield the numbers⁶,

$$\Pr\{\mathcal{H}_1\} = 0.3, \quad \mathbf{x}_k|\mathcal{H}_1 \sim \mathcal{N}\left(\begin{bmatrix} 3 \\ -0.5 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right) \quad (24)$$

$$\Pr\{\mathcal{H}_2\} = 0.7, \quad \mathbf{x}_k|\mathcal{H}_2 \sim \mathcal{N}\left(\begin{bmatrix} -0.2 \\ 2.7 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right). \quad (25)$$

We realize that the two sets $\{x_k^1 = \beta_1, x_k^2 = \beta_2\}$ and $\{x_k^1 = \beta_2, x_k^2 = \beta_1\}$ represent the same set of targets. We may therefore move density from one such labeled point to the other, without changing the RFS density. One way to do this is by switching the indexes under \mathcal{H}_2 ,

$$\Pr\{\mathcal{H}_1\} = 0.3, \quad \mathbf{x}_k|\mathcal{H}_1 \sim \mathcal{N}\left(\begin{bmatrix} 3 \\ -0.5 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right) \quad (26)$$

$$\Pr\{\mathcal{H}_2\} = 0.7, \quad \mathbf{x}_k|\mathcal{H}_2 \sim \mathcal{N}\left(\begin{bmatrix} 2.7 \\ -0.2 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right). \quad (27)$$

⁶All probabilities and densities are conditioned on data, but this is omitted for notational convenience.

The densities $p_1(\mathbf{x}_k)$, given by (24)–(25), and $p_2(\mathbf{x}_k)$, given by (26)–(27), thus correspond to the same RFS density

$$p(\{x_k^1, x_k^2\}) = p_1(\mathbf{x}_k) + p_1(\pi\mathbf{x}_k) = p_2(\mathbf{x}_k) + p_2(\pi\mathbf{x}_k), \quad (28)$$

where the matrix π is a permutation matrix, defined as

$$\pi = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (29)$$

Since the ordered densities p_1 and p_2 belong to the same RFS family, they should, ideally, render the same estimates. As illustrated in Fig. 1, the suggested switch leads to a simpler problem, since *the marginalized densities of $x_k^{(1)}$ and $x_k^{(2)}$ can be approximated by a Gaussian density more accurately.*

The dash-dotted line in Fig. 1 is a symmetry line. When probability mass is moved from one labeled point to another, the movement is through this line to the mirror point. Of course, the example is selected to highlight the advantages with a switch, and one can easily construct situations when it is better not to switch indices. Still, the example illustrates a general technique that can be employed by most tracking algorithms that use merging.

To further improve the understanding of the concepts in Example 1, we stress the relation to the RFS densities. Let $\tilde{p}_1(x_k^{(1)}, x_k^{(2)})$ and $\tilde{p}_2(x_k^{(1)}, x_k^{(2)})$ denote the Gaussian approximations of $p_1(x_k^{(1)}, x_k^{(2)})$ and $p_2(x_k^{(1)}, x_k^{(2)})$, respectively. Both $p_1(x_k^{(1)}, x_k^{(2)})$ and $p_2(x_k^{(1)}, x_k^{(2)})$ correspond to the same RFS density, i.e., $p_1(\{x_k^{(1)}, x_k^{(2)}\}) = p_2(\{x_k^{(1)}, x_k^{(2)}\})$ (cf. (28)). As the approximation $\tilde{p}_2(x_k^{(1)}, x_k^{(2)}) \approx p_2(x_k^{(1)}, x_k^{(2)})$ is fairly accurate, it follows that $\tilde{p}_2(\{x_k^{(1)}, x_k^{(2)}\}) \approx p_2(\{x_k^{(1)}, x_k^{(2)}\}) = p_1(\{x_k^{(1)}, x_k^{(2)}\})$. However, it is likely that the approximation $\tilde{p}_1(\{x_k^{(1)}, x_k^{(2)}\}) \approx p_1(\{x_k^{(1)}, x_k^{(2)}\})$ is less accurate. Hence, by switching densities we will make approximations that better preserve the information about the desired RFS density, $p_1(\{x_k^{(1)}, x_k^{(2)}\})$.

It is not obvious how to compute MMOSPA estimates from a given RFS density. Instead, the proposed algorithms in Sections 5 and 6 use the MMSE estimates, i.e., the posterior means, of a density of ordered targets. The idea is to select a density within the RFS family such that the MMSE estimates are close to the MMOSPA estimates. For the considered example, we illustrate the importance that the choice of density has on the posterior means and the MOSPA performance, by studying the MMSE state vector estimates before and after the index switch. In the original indexation, the posterior means are

$$\hat{x}_k^{(1)} = 0.76, \quad \hat{x}_k^{(2)} = 1.74 \quad (30)$$

whereas the posterior means after the switch are

$$\hat{x}_k^{(1)} = 2.79, \quad \hat{x}_k^{(2)} = -0.29. \quad (31)$$

The latter posterior means are probably close to the optimal estimates since, under both hypotheses, one target is fairly close to 2.79 whereas the other target is reasonably close to -0.29 . Hence, although the initial objective was to improve the Gaussian approximation, the density switch also seems to yield MMSE estimates which are closer to the MMOSPA estimates.

For more results on the connection between accurate Gaussian approximations and improved estimates, we refer to Section 6.

5 Optimal approximations in the Kullback-Leibler sense

In this section, we present a way of switching densities within the RFS family such that the new density can be most accurately approximated with a Gaussian density, in the Kullback-Leibler sense. Note that the description is only made for a two-target case, but the results can be generalized.

In what follows, we assume that the posterior density is a Gaussian mixture, denoted

$$p(\mathbf{x}) = \sum_{h=1}^{N_{\mathcal{H}}} \beta_h \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_h, \mathbf{P}_h). \quad (32)$$

For the two-target case, we make the following definition of the RFS family, \mathcal{A}_p , of $p(\mathbf{x})$ (cf. (28))

$$\mathcal{A}_p \triangleq \{f : f(\mathbf{x}) + f(\pi\mathbf{x}) = p(\mathbf{x}) + p(\pi\mathbf{x})\}. \quad (33)$$

Here, \mathbf{x} is the stacked vector of the state vectors $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ of the two targets, and π is a permutation matrix (cf. (29)).

Depending on the choice of density $f(\mathbf{x})$, the Gaussian approximation may be more or less accurate. Our objective here is to

1. find a density $f(\mathbf{x})$ that enables the most accurate Gaussian approximations; and then to
2. find the Gaussian density $\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})$ that best approximates $f(\mathbf{x})$.

As a measure for how accurate a density approximation is, we will use the Kullback-Leibler divergence [16]. Of course, our interest in $f(\mathbf{x})$ is only as a means to obtain $\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})$, which means that one could imagine a technique that does not involve $f(\mathbf{x})$. The mathematical definition of the problem that we would like to solve is

$$\{\bar{\mathbf{x}}, \mathbf{R}\} = \arg \min_{\bar{\mathbf{x}}, \mathbf{R}} \left\{ \min_{f(\mathbf{x}) \in \mathcal{A}_p} \text{KL} [f(\mathbf{x}) || \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})] \right\}, \quad (34)$$

where

$$\text{KL} [f(\mathbf{x}) || \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})] \triangleq \int f(\mathbf{x}) \log \frac{f(\mathbf{x})}{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})} d\mathbf{x}, \quad (35)$$

and where KL denotes Kullback-Leibler divergence. In the standard procedure (used in for instance PDA and JPDA) one would use $f(\mathbf{x}) = p(\mathbf{x})$ and only optimize over $\bar{\mathbf{x}}$ and \mathbf{R} , for which the optimal solution is given by moment matching. By also optimizing over $f(\mathbf{x})$ we believe that the approximation errors will decrease significantly compared to JPDA, and slightly compared to the algorithm that we present in Section 6.

To design an algorithm that can search for $\bar{\mathbf{x}}$ and \mathbf{R} , the following results are very useful.

Theorem 2 *The solution to*

$$\{\bar{\mathbf{x}}, \mathbf{R}\} = \arg \min_{\bar{\mathbf{x}}, \mathbf{R}} \text{KL} [f(\mathbf{x}) || \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})] \quad (36)$$

is given by moment matching,

$$\bar{\mathbf{x}} = \mathbb{E}_{f(\mathbf{x})} \{\mathbf{x}\} \quad (37)$$

$$\mathbf{R} = \text{Cov}_{f(\mathbf{x})} \{\mathbf{x}\}. \quad (38)$$

Furthermore, the density $f(\mathbf{x}) \in \mathcal{A}_p$ that minimizes $\text{KL} [f(\mathbf{x}) || \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})]$ is

$$f(\mathbf{x}) = (p(\mathbf{x}) + p(\pi\mathbf{x})) \cdot \frac{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})}{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R}) + \mathcal{N}(\pi\mathbf{x}; \bar{\mathbf{x}}, \mathbf{R})}. \quad (39)$$

Proof of Theorem 2 *The results in Eq. (37) and (38) are well known, see e.g. [17]. For (39), see [18].*

Based on the above theorem, we propose an iterative optimization algorithm. The algorithm is illustrated in Fig. 2, where the initial density is $p(\mathbf{x})$.

1. Initiate with $i = 1$ and let $\bar{\mathbf{x}}_0$ and \mathbf{R}_0 be the first two moments of $p(\mathbf{x})$.

2. Set

$$f_i(\mathbf{x}) = \frac{p(\mathbf{x}) + p(\pi\mathbf{x})}{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_{i-1}, \mathbf{R}_{i-1}) + \mathcal{N}(\pi\mathbf{x}; \bar{\mathbf{x}}_{i-1}, \mathbf{R}_{i-1})}. \quad (40)$$

3. Compute

$$\bar{\mathbf{x}}_i = \mathbb{E}_{f_i(\mathbf{x})} \{\mathbf{x}\} \quad (41)$$

$$\mathbf{R}_i = \text{Cov}_{f_i(\mathbf{x})} \{\mathbf{x}\}. \quad (42)$$

4. If $\bar{\mathbf{x}}_i \approx \bar{\mathbf{x}}_{i-1}$ and $\mathbf{R}_i \approx \mathbf{R}_{i-1}$ we stop. Otherwise, set $i = i + 1$ and go back to 2.

The algorithm is essentially straightforward. However, to implement it we need the ability to calculate expected values with respect to the densities $f_i(\mathbf{x})$, $i = 1, 2, \dots$

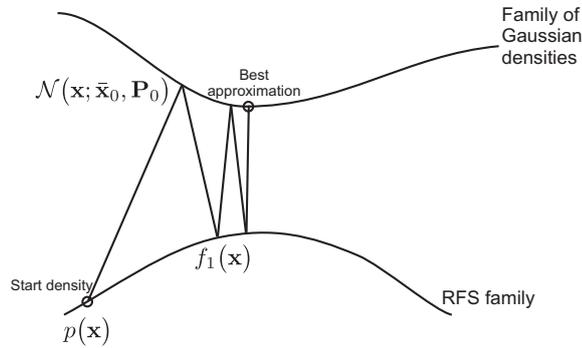


Figure 2: Illustration of the iterative optimization algorithm which finds the best Gaussian approximation of the RFS family, in the Kullback-Leibler sense. The curves represent the functional spaces of densities within the RFS family (below), and the Gaussian densities (above).

It appears complicated to find analytical expressions for the expected values in (41) and (42) and we therefore suggest a numerical method based on importance sampling. As importance density we use

$$q(\mathbf{x}) = \frac{1}{2}(p(\mathbf{x}) + p(\pi\mathbf{x})). \quad (43)$$

The expected values can be reformulated as

$$\bar{\mathbf{x}}_i = \mathbb{E}_{f_i(\mathbf{x})} \{ \mathbf{x} \} = \mathbb{E}_{q(\mathbf{x})} \left\{ \mathbf{x} \frac{f_i(\mathbf{x})}{q(\mathbf{x})} \right\} \quad (44)$$

$$\mathbf{R}_i = \mathbb{E}_{f_i(\mathbf{x})} \{ (\mathbf{x} - \bar{\mathbf{x}}_i)(\mathbf{x} - \bar{\mathbf{x}}_i)^T \} \quad (45)$$

$$= \mathbb{E}_{q(\mathbf{x})} \left\{ (\mathbf{x} - \bar{\mathbf{x}}_i)(\mathbf{x} - \bar{\mathbf{x}}_i)^T \frac{f_i(\mathbf{x})}{q(\mathbf{x})} \right\}, \quad (46)$$

and for the selected importance density it is easy to evaluate the ratio

$$\frac{f_i(\mathbf{x})}{q(\mathbf{x})} = 2 \frac{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_{i-1}, \mathbf{R}_{i-1})}{\mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}_{i-1}, \mathbf{R}_{i-1}) + \mathcal{N}(\pi\mathbf{x}; \bar{\mathbf{x}}_{i-1}, \mathbf{R}_{i-1})}. \quad (47)$$

Based on i.i.d. samples $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \sim q(\mathbf{x})$, we can approximate the desired entities as

$$\bar{\mathbf{x}}_i \approx \sum_{n=1}^N \mathbf{x}_n w_n \quad (48)$$

$$\mathbf{R}_i \approx \sum_{n=1}^N (\mathbf{x}_n - \bar{\mathbf{x}}_i)(\mathbf{x}_n - \bar{\mathbf{x}}_i)^T w_n, \quad (49)$$

where

$$w_n = \frac{f_i(\mathbf{x}_n)/q(\mathbf{x}_n)}{\sum_{r=1}^N f_i(\mathbf{x}_r)/q(\mathbf{x}_r)}. \quad (50)$$

In (49), we replace $\bar{\mathbf{x}}_i$ with the approximated value from (48). A beneficial property with the suggested importance function is that since the RFS family is preserved we can use the same samples $\mathbf{x}_1, \dots, \mathbf{x}_N$ for all iterations $i = 1, 2, \dots$.

In this section, we have presented the optimal approach of performing density switches within the RFS family, such that the final density can most accurately be approximated with a Gaussian density. We have also presented a numerical approach of computing the required expected values. Since the presented algorithm is computationally demanding, we are interested in a sub-optimal approach, which still has the properties of enabling better Gaussian approximations than the JPDA algorithm, and which also presents better estimates than JPDA in MOSPA sense. The development of such an approach is the topic of the following section.

6 Set JPDA algorithm

In this section, we derive a sub-optimal approach to the third conceptual solution, called the Set JPDA (SJPDA) algorithm. The filter utilizes the fact that several ordered densities correspond to the same set density, which provides the possibility to switch between those ordered densities. The criterion that we try to minimize with our switch of densities is optimal in terms of the MOSPA estimates. Still, we argue that it in most cases also gives better density approximations than the JPDA algorithm.

The Set JPDA algorithm is a modification of the classic JPDA filter. The difference is that once the posterior density is described as a weighted sum of Gaussian densities, we allow ourself to switch that density for another ordered density in the same RFS family. On a high level, the SJDA algorithm works as follows:

- 1) Formulate the set of global measurement hypotheses, \mathcal{H} , and calculate conditional densities of all targets as well as the probabilities of all hypotheses. Approximate the targets as independent.
- 2) Reorder the target indexes under the different hypotheses with the objective to make the marginalized densities resemble Gaussian densities (by minimization in one's favorite sense, see Sections 6.1 and 6.2 for details).
- 3) Approximate the marginalized posterior densities of all targets as independent Gaussian. Then go back to 1.

A block diagram description of the SJPDA filter is shown in Fig. 3. As seen in the figure, the difference between SJPDA and JPDA lies in the switching block. If that block is removed, we obtain the JPDA filter.

The key aspect of the SJPDA filter is the switching of densities. To obtain a filter with good performance, and which enables accurate Gaussian approximations, the switching criterion is very important. In the following two sections, we propose a goal function, and formulate the problem of minimizing that function (while remaining within the family of Gaussian mixtures) as a continuous optimization problem.

6.1 Goal function proposal and motivation

For the considered problem, the labeled density is a Gaussian mixture. In this section, we propose and motivate a goal function for the optimization problem of finding the best switching of such densities in the SJPDA filter.

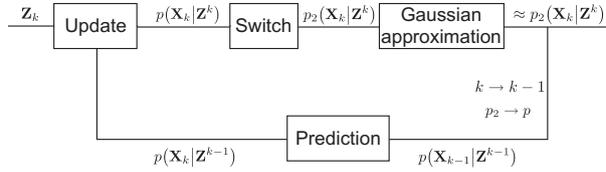


Figure 3: SJPDA block diagram.

Generally speaking, a Gaussian mixture can be accurately approximated by a single Gaussian density as long as the Gaussian mixture is not too distinctly multimodal. In the SJPDA filter, we wish to adjust the indexation within each global hypothesis, in order to find a labeled density which is less multimodal. The goal function that we propose is hence a function which measures the multi-modality of the density. The proposed goal function is

$$\sum_{i=1}^n \text{tr}\{\mathbf{P}_k^i\}, \quad (51)$$

for which \mathbf{P}_k^i is given in (19). In the following, we motivate that goal function, basically by arguing that through minimizing of the function, we will obtain both better approximations and better estimates than what we would obtain using the JPDA filter.

In [13], the problem of minimizing MOSPA is considered in more detail. An important result from the paper is that *the problem of minimizing the sum of the trace of the posterior covariances is equal to the problem of minimizing MOSPA*, given that the expected value is used as estimate and that we search in the entire RFS family (not restricting the search to Gaussian mixtures in the RFS family). Thus, the switching criterion that we use is optimal for the estimation problem.

The second advantage with the cost function is that it enables better Gaussian approximations. The reason for this is that by minimizing the trace of the covariance matrices, we make the posterior density less multimodal. We therefore argue that the Gaussian approximations in the SJPDA filter are at least as accurate as the JPDA filter approximations, where equality holds if no switches are made. However, for the density approximations, there are no optimality results. That is, one cannot show that the approximations are optimal, for example in a Kullback-Leibler sense, by using the proposed goal function. Instead, for optimal approximations, we refer to the algorithm given in Section 5.

We wish to minimize

$$\begin{aligned} & \text{tr} \{ \text{Cov}_{q(\mathbf{x})} \{ \mathbf{x} \} \} \\ &= \sum_{i=1}^{n!} \sum_{h=1}^{N_{\mathcal{H}}} \phi_{i,h} \left\{ (\pi_i \boldsymbol{\mu}_h - \bar{\mathbf{x}})^T (\pi_i \boldsymbol{\mu}_h - \bar{\mathbf{x}}) + \text{tr} \{ \pi_i \mathbf{P}_h \pi_i^T \} \right\} \end{aligned} \quad (58)$$

$$= \left\{ \sum_{i=1}^{n!} \sum_{h=1}^{N_{\mathcal{H}}} \phi_{i,h} \boldsymbol{\mu}_h^T \boldsymbol{\mu}_h \right\} - \bar{\mathbf{x}}^T \bar{\mathbf{x}} + \sum_{h=1}^{N_{\mathcal{H}}} \text{tr} \{ \mathbf{P}_h \} \sum_{i=1}^{n!} \phi_{i,h} \quad (59)$$

$$= \left\{ \sum_{h=1}^{N_{\mathcal{H}}} \beta_h \boldsymbol{\mu}_h^T \boldsymbol{\mu}_h \right\} - \boldsymbol{\phi}^T \mathbf{V}^T \mathbf{V} \boldsymbol{\phi} + \sum_{h=1}^{N_{\mathcal{H}}} \text{tr} \{ \mathbf{P}_h \} \beta_h, \quad (60)$$

where $-\boldsymbol{\phi}^T \mathbf{V}^T \mathbf{V} \boldsymbol{\phi}$ is the only part that depends on $\boldsymbol{\phi}$, subject to the constraints presented in (53) and (54). To summarize, we have the optimization problem

$$\begin{aligned} & \min_{\boldsymbol{\phi}} \quad -\boldsymbol{\phi}^T \mathbf{V}^T \mathbf{V} \boldsymbol{\phi}. \quad (61) \\ & \text{s.t.} \quad \begin{cases} \phi_{i,h} \geq 0 & \forall i, h \\ \sum_{i=1}^{n!} \phi_{i,h} = \beta_h & \forall h \end{cases} \end{aligned}$$

The problem is non-convex. The goal function is concave, and it is to be minimized over a convex region, which has the implication that the optimal point is along the border of the constraints. We also note that since the optimal solution is on the border of the constraints, the permutations either occur or not, i.e., there will be no partial permutations in the optimal point. A benefit with the continuous formulation is that it enables the use of commercial optimization solvers, which scale well with increasing number of targets and hypotheses. To use an optimization solver, we need knowledge regarding existence of local minima, and how they can be avoided.

It can be analytically shown that there are no local minima for the case of scalar state vectors and two targets. But when the dimensionality of the state vector increases, there will in some situations exist local minimum points in which an optimization algorithm can be trapped. It is thus important to select a suitable starting point for the algorithm. Through empirical studies, we have seen that it is only in a small region around the minimum point that the negative gradient points in the direction of the local minimum. Therefore, only starting points in a small region around potential local minima should be avoided. For the problems considered in this article, we have found a suitable set of starting points which almost always lead the selected optimization algorithm to the global minimum.

For more details on the characteristics of the optimization problem, see [18].

6.3 Algorithm description

The SJPDA algorithm has five steps which are executed for each time index k . In the following, we describe each of those five steps in more detail, and summarize the algorithm in Table 4.

The algorithm description starts at time k , at which we have available a set of measurements, \mathbf{Z}_k , and predicted states $\mathbf{x}_{k|k-1}^i$ and covariance matrices $\mathbf{P}_{k|k-1}^i$ for each target i . Note that steps I, II and V are identical to the JPDA algorithm (one version of it). Note also that we assume linear and Gaussian process and measurement models in the description, although the SJPDA algorithm can be easily extended to handle nonlinear models using, for example, an extended Kalman filter (EKF).

Step I: Data hypothesis extraction

The first step of the SJPDA algorithm is to formulate all possible global data association hypotheses, \mathcal{H} . A hypothesis is possible if it describes the origin of each measurement in \mathbf{Z}_k (target-generated or false), and if the total number of target-originated measurements is at most n .

Step II: Data update

The second step of the algorithm is to update the predicted state vectors and covariance matrices, and to calculate the mixture weights, β_h . For a certain target i , the updated state and covariance matrix under hypothesis h are given by

$$\mathbf{x}_{k|k}^{i,h} = \mathbf{x}_{k|k-1}^i + \mathbf{K}_k^i (\mathbf{H}_k \mathbf{x}_{k|k-1}^i - \mathbf{z}_k^{i,h}) \quad (62)$$

$$\mathbf{P}_{k|k}^{i,h} = (\mathbf{I} - \mathbf{K}_k^i \mathbf{H}_k) \mathbf{P}_{k|k-1}^i, \quad (63)$$

where $\mathbf{z}_k^{i,h}$ is the measurement associated to target i under hypothesis h . If no measurement is associated, we use $\mathbf{z}_k^{i,h} = \mathbf{0}$. An expression for the mixture weights is given in (13).

Step III: Optimization

Step three of the SJPDA algorithm is the main step, namely to find the optimal permutation of state vectors under the data association hypotheses. We find the optimum solution, ϕ^* , to (61) by applying an optimization solver

Table 4: SJPDA algorithm description

I: Data hypothesis extraction

formulate all $N_{\mathcal{H}}$ global data association hypotheses

II: Data update

for each data association hypothesis $h = 1, 2, \dots, N_{\mathcal{H}}$, do

for each target $i = 1, \dots, n$, do

update the expected value and covariance matrix of the predicted state vector $\hat{\mathbf{x}}_k^i$ according to (62)–(63)

calculate the mixture weight β_h according to (13)

III: Optimization

solve the problem defined in (61) using an optimization solver

a suitable pair of starting points is given in (64)–(65), and a useful third initiation point is to perform a 70% permutation under the N (e.g. 10) largest hypotheses

IV: Permutation and update

compute the expected value and covariance matrix of the joint state vector \mathbf{X}_k after data update, optimal permutation and Gaussian approximation, according to (68)–(69)

V: Prediction

for each target $i = 1, 2, \dots, N_t$, do

predict the state vector and covariance matrix of target i , according to (70)–(71).

with different initiation points. A suitable pair of starting points for the two-target problem is

$$\phi_{\text{start},1} = [0.7\beta_1 \quad 0.3\beta_1 \quad \dots \quad 0.7\beta_{N_h} \quad 0.3\beta_{N_h}]^T \quad (64)$$

$$\begin{aligned} \phi_{\text{start},2} = [0.7\beta_1 \quad 0.3\beta_1 \quad 0.3\beta_2 \quad 0.7\beta_2 \\ \dots \quad 0.7\beta_{N_h} \quad 0.3\beta_{N_h}]^T. \end{aligned} \quad (65)$$

Step IV: Permutation and update

After the optimal solution ϕ^* has been found, we permute the state vectors according to that solution. The permutation does not have to be done explicitly, since we are only interested in the Gaussian approximation of the permuted posterior density. To express the mean and covariance after permutation, we describe the calculation in the joint state vector \mathbf{X}_k . Before permutation, the expected value and covariance matrix of that vector is

$$\mathbf{X}_{k|k}^h = \left[(\mathbf{x}_{k|k}^{1,h})^T (\mathbf{x}_{k|k}^{2,h})^T \cdots (\mathbf{x}_{k|k}^{n,h})^T \right]^T \quad (66)$$

$$\mathbf{P}_{k|k}^h = \text{diag}\{\mathbf{P}_{k|k}^{1,h}, \mathbf{P}_{k|k}^{2,h}, \dots, \mathbf{P}_{k|k}^{n,h}\}. \quad (67)$$

After permutation, the expected value and covariance matrix are given by

$$\mathbf{X}_{k|k} = \mathbf{V}\phi^* \quad (68)$$

$$\begin{aligned} \mathbf{P}_{k|k} = & \sum_{h=1}^{N_h} \sum_{j=1}^{n!} \phi^*(i(h, j)) \left[\mathbf{\Pi}_j \mathbf{P}_{k|k}^h \mathbf{\Pi}_j \right. \\ & \left. + (\mathbf{\Pi}_j \mathbf{X}_{k|k}^h - \mathbf{X}_{k|k}) (\mathbf{\Pi}_j \mathbf{X}_{k|k}^h - \mathbf{X}_{k|k})^T \right], \end{aligned} \quad (69)$$

where the function $i(h, j)$ gives the index in ϕ^* which corresponds to permutation i and hypothesis h . We approximate the targets as independent, which implies that the posterior covariance matrix \mathbf{P}_k^i for target i (the indices are irrelevant) is the corresponding block in the matrix $\mathbf{P}_{k|k}$.

Step V: Prediction

The final step of the SJPDA algorithm is to let $k \rightarrow k - 1$, and predict the state and covariance matrices of the set of targets to the next time step k . For linear models, that prediction is done using the ordinary Kalman prediction equations

$$\mathbf{x}_{k|k-1}^i = \mathbf{F}_k \mathbf{x}_{k-1|k-1}^i \quad (70)$$

$$\mathbf{P}_{k|k-1}^i = \mathbf{F}_k \mathbf{P}_{k-1|k-1}^i \mathbf{F}_k^T + \mathbf{Q}_k. \quad (71)$$

In Fig. 4, the densities of the joint state vector (with two targets and scalar states) is illustrated for different steps in the JPDA and SJPDA algorithms. In the example, the predicted density is a Gaussian, given by the top figure. The second row of densities show the density after measurement update and optimal permutations (for SJPDA). From the figures, we see that the SJPDA density can be described much more accurately by a Gaussian density. The last row shows the Gaussian approximations of the JPDA

and SJPDA posterior densities. It is clear that the covariance matrix of the SJPDA approximation is smaller.

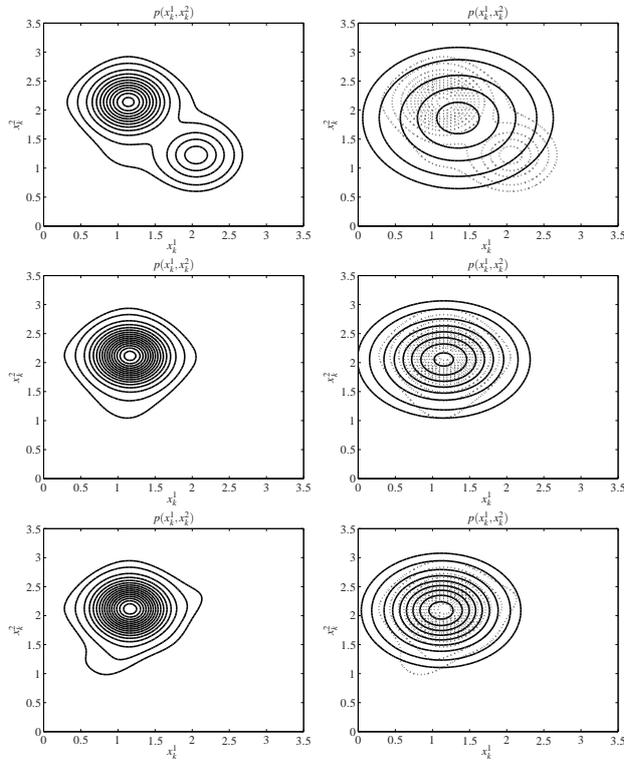


Figure 4: One-step iteration of the JPDA, SJPDA and KLSJPDA filters. The prior density is a Gaussian, there are two targets, the detection probability is 0.7, and two target-originated measurements are received. Posterior density [top left], posterior density after SJPDA optimization [middle left], posterior density after KLSJPDA optimization [bottom left], JPDA approximation [top right], SJPDA approximation [middle right], KLSJPDA approximation [bottom right]. The dashed contours in the right-column figures represent the densities that are approximated. Clearly, both the SJPDA and KLSJPDA posterior densities are better approximated by a Gaussian density than the original posterior density.

7 Evaluations

Two different tracking scenarios are considered for evaluation of the SJPDA algorithm, and comparison with JPDA and a Gaussian-mixture CPHD filter with known target number. The first scenario, for which also the KLSJPDA

algorithm is evaluated, is tracking of two targets, whose trajectories are as illustrated in Fig. 5. The parameters used in the evaluation are $l = 10$ m, $d_1 = 0.5$ m, $d_2 = 30$ m, $\varphi = \pi/3$, and $v = 1$ m/s. The second scenario, illustrated in Fig. 8, is also a two-target example, but in that scenario the two targets move around a hexagonal shape. The aim of that scenario is to evaluate how fast the JPDA and SJPDA algorithms lose track in a challenging scenario. In both scenarios, a sensor collects measurements at even time intervals of $T_0 = 1$ second.

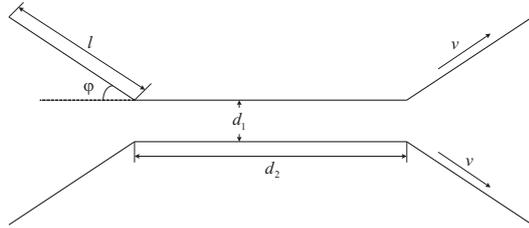


Figure 5: Illustration of the scenario for the first evaluation.

In the filtering algorithms, a nearly constant velocity model is assumed (see (4) for a general description), which is governed by the system matrix

$$\mathbf{F}_{k-1} = \begin{bmatrix} \mathbf{I}_{2 \times 2} & T_0 \mathbf{I}_{2 \times 2} \\ \mathbf{0}_{2 \times 2} & \mathbf{I}_{2 \times 2} \end{bmatrix}, \quad (72)$$

where $\mathbf{I}_{2 \times 2}$ is a 2×2 identity matrix, and $\mathbf{0}_{2 \times 2}$ is a 2×2 matrix of zeroes, and by the process noise \mathbf{v}_{k-1} which is zero-mean Gaussian with covariance matrix

$$\mathbf{Q} = q_0 \begin{bmatrix} T_0^3 \mathbf{I}_{2 \times 2} / 3 & T_0^2 \mathbf{I}_{2 \times 2} / 2 \\ T_0^2 \mathbf{I}_{2 \times 2} / 2 & T_0 \mathbf{I}_{2 \times 2} \end{bmatrix}, \quad (73)$$

where q_0 is a tuning parameter. In the simulations, the best tuning parameter is selected for each filter, where the parameter which yields the lowest average MOSPA is used. Further, the measurement model is assumed linear and Gaussian (cf. (6)) with observation matrix

$$\mathbf{H}_k = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}, \quad (74)$$

which means that the sensor delivers position measurements. The Gaussian measurement noise \mathbf{w}_k is zero-mean, with covariance matrix

$$\mathbf{R}_k = \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix}, \quad (75)$$

where we have used $\sigma_x = \sigma_y = 0.2\text{m}$.

Fig. 6 shows an example of the JPDA and SJPDA output for $d_1 = 0.5\text{m}$, $\lambda = 0.01$, and $P_D = 1$. In the figure we see the track coalescence tendency of the JPDA filter, which makes it hard for the filter to detect the separation of the tracks. This leads to high MOSPA after the track separation, and it also leads to a high risk of losing tracks. In Fig. 7, the MOSPA performance over 100 Monte Carlo runs is shown for detection probabilities of 1 and 0.85, respectively. The clutter intensity in the simulations is 0.02 m^{-2} .

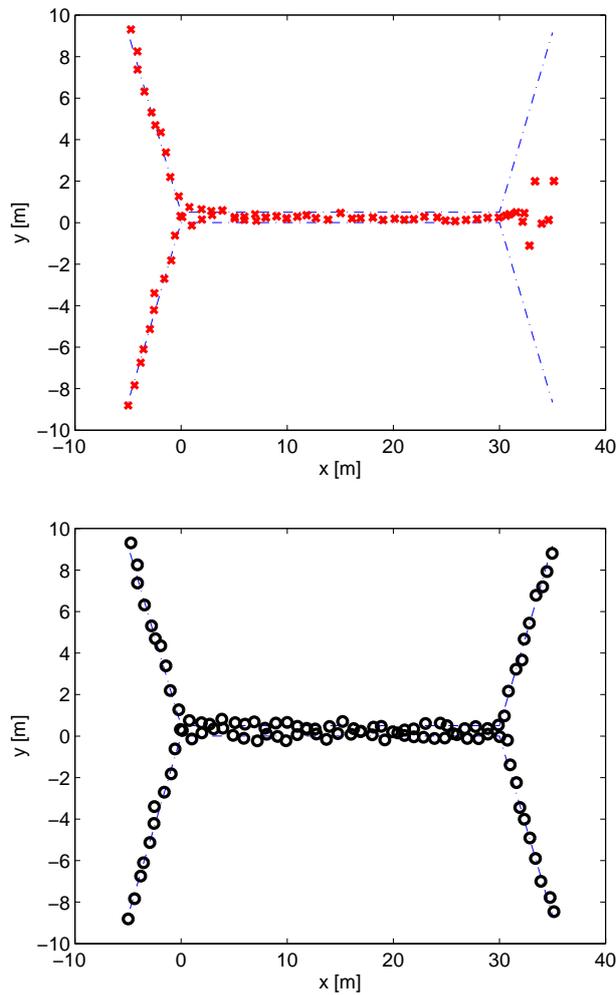


Figure 6: Output of JPDA [top] and SJPDA [bottom] for one sequence of measurements.

The figures show that the average OSPA performance of the SJPDA filter is better than for the JPDA filter for almost the entire scenario, and the

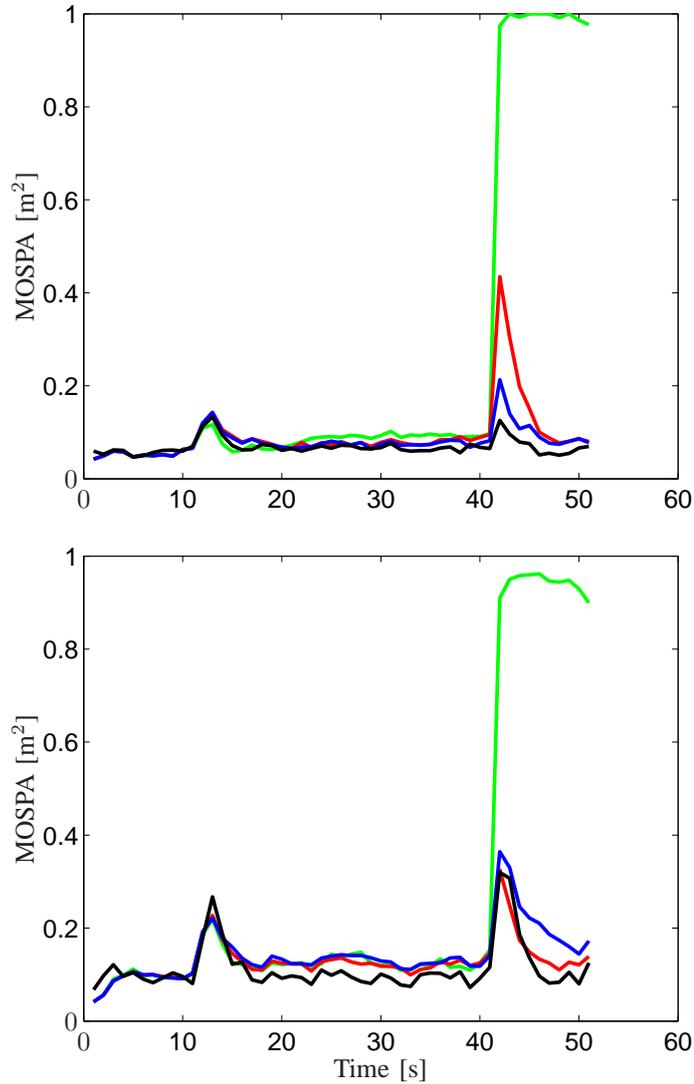


Figure 7: Evaluation of JPDA (green), SJPDA (red), KLSJPDA with 10 000 samples (blue), and GM-CPHD (black) on the two-target scenario for detection probabilities $P_d = 1$ [top] and $P_d = 0.85$ [bottom].

figures also show that the difference between the filters are very large at the time when the tracks separate. This is due to the track coalescence of the JPDA filter. Note that the filters are run with different process noises, where the noise level has been selected to yield the best average MOSPA. When the filters are run with the same process noise, the SJPDA filter always gives an average OSPA which is lower than, or equal to, that of the JPDA filter. For the KLSJPDA algorithm, the performance is similar to that of the SJPDA

filter, whereas its computational complexity is much higher. We thus suggest that the SJPDA filter be used.

When compared to GM-CPHD, we see that both the SJPDA and the KLSJPDA filters for large detection probabilities performs on average equally well as GM-CPHD (for unity detection probability actually slightly better), and almost as well for lower detection probabilities. The results are rather surprising since the GM-CPHD is a more complex algorithm. For the GM-CPHD algorithm, the best process noise parameter q_0 in the aforementioned set is used, and the filter also uses pruning (threshold 0.0001) and merging (threshold $U = 1$, cf. [12]). We note that the somewhat similar JPDA* [20] was compared to the SJPDA in [21]; in fact the JPDA* performs well, but the SJPDA is notably superior.

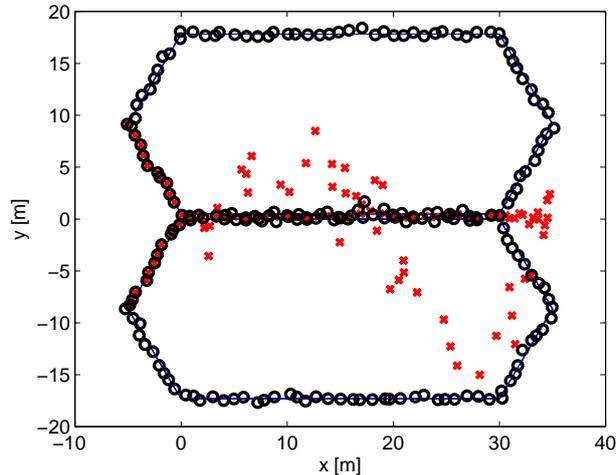


Figure 8: Example of filter output of the JPDA (crosses) and SJPDA (circles) algorithms for the track life evaluation scenario (for $P_D = 1$). In the plot, the output for the first 101 time steps is presented, which is the time it takes for the two targets to make one loop around the path. In the evaluations, the full trajectories make ten loops around the hexagonal path.

For the second scenario, the average track life and average track loss for JPDA and SJPDA are summarized in Table 5, for four different values of detection probability. In the evaluation, the best process noise parameter in the set $q_0 = \{0.05, 0.1, 0.2, 0.3, 0.4\}$ has been used, and otherwise the same parameter values as before. A track is considered lost if the covariance matrix elements corresponding to the uncertainty in the x or y dimensions have surpassed 225m^2 , the estimate of a track is further away than 50 meters from the true value, or if there has been only two measurements in the gate

of a track for the last five scans. The selected process noise is such that the average track life is maximized. The trajectories that the two targets travel along are illustrated in Fig. 8, where example outputs of the JPDA and SJPDA filters are included. The angle φ (cf. Fig. 5) is $\pi/3$, and the targets travel at a maximum 10 laps around the circuits, which correspond to 1010 seconds, or almost 17 minutes. From the table we see that the SJPDA filter has much longer average track life than the JPDA filter, due to its avoidance of track coalescence and better description of where there are targets. The lower track-loss probabilities of the SJPDA filter are evident for all four detection probabilities, but the improvement compared to the JPDA filter is better the higher the detection probability.

Table 5: Average track lives and average track loss probabilities for the JPDA and SJPDA algorithms, for different detection probabilities, P_D .

		JPDA	SJPDA	Improvement SJPDA vs. JPDA
$P_D = 1$	Length	97s	993s	924%
	Loss	100%	3%	97%
$P_D = 0.95$	Length	128s	734s	473%
	Loss	100%	53%	47%
$P_D = 0.9$	Length	152s	456s	200%
	Loss	100%	88%	12%
$P_D = 0.85$	Length	142s	251s	77%
	Loss	100%	100%	0%

8 Conclusions

In this article we have shown how traditional tracking algorithms can be improved when target identity is not of interest. The presented approach uses the relation between the density of ordered (labeled) targets, and the density of unordered (the set of) targets. More specifically, there is an infinite number of ordered densities which correspond to the same unordered density, and by switching between those densities we can obtain densities which have better characteristics. In order to reflect the fact that target identity is irrelevant, we use the Mean Optimal Subpattern Assignment (MOSPA) metric instead of the Mean Squared Error (MSE).

In the first part of the article, we discuss how the new problem of minimizing MOSPA differs from the classic tracking problem, and show that there is substantial room for improvements of the traditional algorithms when evaluated according to MOSPA.

In the second part of the article we study the Joint Probabilistic Data Association (JPDA) filter and how it can be adjusted to the considered problem. First, we describe how optimal Gaussian approximations in Kullback-Leibler sense can be found. From this, we develop a new filter called the Kullback-Leibler Set JPDA (KLSJPDA) filter. Since the filter is computationally intensive, we propose another approach which operates on a goal function which leads to both good Gaussian approximations and low MOSPA. The minimization of that goal function can be formulated as a continuous optimization problem. From this, we develop another adjustment of the JPDA filter called the Set JPDA (SJPDA) filter.

In the final part of the paper, the SJPDA and KLSJPDA filters are evaluated on two simulation examples and compared to the JPDA and the Gaussian-mixture Cardinalized Probability Hypothesis Density (GM-CPHD) filters. The results show that the SJPDA and KLSJPDA filters have similar performance and that they perform substantially better than JPDA in terms of MOSPA. However, due to its lower complexity, we prefer the SJPDA filter. The results also show that the SJPDA filter has a much longer average track length than JPDA for the considered scenario. Further, it is seen that the SJPDA filter performs almost as good as the more complex GM-CPHD filter, and that it even performs slightly better for a unity detection probability.

References

- [1] Y. Bar-Shalom and E. Tse, "Tracking in a cluttered environment with probabilistic data association," *Automatica*, vol. 11, no. 5, pp. 451–460, September 1975.
- [2] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE Journal of Oceanic Engineering*, vol. 8, no. 3, pp. 173–183, July 1983.
- [3] D. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, December 1979.
- [4] S. Deb, M. Yeddanapudi, K. Pattipati, and Y. Bar-Shalom, "A generalized s-d assignment algorithm for multisensor-multitarget state estima-

- tion,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 2, pp. 523–538, April 1997.
- [5] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, 1999.
- [6] S. Blackman, “Multiple hypothesis tracking for multiple target tracking,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, January 2004.
- [7] 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, August 2008.
- [8] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
- [9] —, “Multitarget Bayes filtering via first-order multitarget moments,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, October 2003.
- [10] B.-N. Vo and W.-K. Ma, “The Gaussian mixture probability hypothesis density filter,” *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, November 2006.
- [11] R. Mahler, “PHD filters of higher order in target number,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, October 2007.
- [12] B.-T. Vo, B.-N. Vo, and A. Cantoni, “Analytic implementations of the cardinalized probability hypothesis density filter,” *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3553–3567, July 2007.
- [13] M. Guerriero, L. Svensson, D. Svensson, and P. Willett, “Shooting two birds with two bullets: How to find minimum mean OSPA estimates,” in *Proceedings of the 13th International Conference on Information Fusion*, 2010.
- [14] O. Drummond and B. Fridling, “Ambiguities in evaluating performance of multiple target tracking algorithms,” in *Signal and Data Processing of Small Targets*, ser. Proc. SPIE, vol. 1698, 1992, pp. 326–337.
- [15] J. Hoffman and R. Mahler, “Multitarget miss distance via optimal assignment,” *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, vol. 34, no. 3, pp. 327–336, May 2004.

- [16] S. Kullback, *Information Theory and Statistics*. NY: Wiley, 1959.
- [17] A. Runnalls, “Kullback-Leibler approach to Gaussian mixture reduction,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 3, pp. 989–999, July 2007.
- [18] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, “Set JPDA filter for multi-target tracking,” Department of Signals and Systems, Chalmers University of Technology, Tech. Rep. R012/2010, 2010.
- [19] L. Svensson, D. Svensson, and P. Willett, “Set JPDA algorithm for tracking unordered sets of targets,” in *Proceedings of the 12th International Conference on Information Fusion*, 2009.
- [20] H. Blom and E. Bloem, “Probabilistic data association avoiding track coalescence,” *IEEE Transactions on Automatic Control*, vol. 45, no. 2, pp. 247–259, February 2000.
- [21] D. Crouse, Y. Bar-Shalom, P. Willett, and L. Svensson, “The JPDAF in practical systems: Computation and snake oil,” in *Signal and Data Processing of Small Targets*, ser. Proc. SPIE, 2010.