Using a Probabilistic Hypothesis Density filter to confirm tracks in a multi-target environment

Paul Horridge and Simon Maskell
QinetiQ Ltd
Malvern Technology Centre
St. Andrews Road
Malvern, Worcs, UK.
p.horridge@signal.qinetiq.com, s.maskell@signal.qinetiq.com

Abstract: In this paper, we aim to perform scalable multi-target particle filter tracking. Previously, the authors presented an approach to track initiation and deletion which maintains an existence probability on each track, including a “search track” which represents the existence probability and state distribution of an unconfirmed track. This approach was seen to perform well even in cases of low detection probability and high clutter levels, but modelling all unconfirmed tracks by a single-target search track can be problematic if more than one target appears in a sensor’s field of view at the same time. To address this, we replace the search track with a Probabilistic Hypothesis Density (PHD) filter which can maintain a density over several unconfirmed tracks. A method is proposed to derive probabilities of measurements originating from targets, allowing us to confirm tracks when these probabilities reach a threshold. We observe that in so doing, we implicitly solve the track-labelling challenge that otherwise exists with PHD filters. This is shown to maintain good tracking performance for high-clutter, low-detection scenarios while addressing the shortcomings of the single-target search track approach. We also show results from a scenario with obscured regions where the target cannot be detected, and show that targets can be tracked through the obscurations.

1 Introduction

In [4], we considered the problem of track initiation and deletion using existence probabilities on each track, including a “search track” which searches the target space for new targets. When the existence probability of the search track reaches a threshold, we create a confirmed track with the probability distribution of the search track, and create a new search track to search for additional targets. Confirmed tracks are deleted when their existence probabilities drop below a deletion threshold.

This was shown to perform well in a tracking scenario with high clutter and a low probability of detecting the target, but there are problems with the approach which we address in the current paper:

- Since more than one unconfirmed target may potentially be in the search region, the distribution of a newly confirmed track may be multi-modal. This can create artefacts where the track mean is initially between two or more targets and shoots across the target space when the filter collapses onto a single mode.
- Since the search track implicitly assumes that there is at most one undiscovered

erschienen im Tagungsband der INFORMATIK 2011
Lecture Notes in Informatics, Band P192
ISBN 978-3-88579-286-4

target in the target space at a time, it can only confirm one track at a time, leading to a delay in confirmation and targets being missed.

The PHD filter proposed by Mahler [7] performs multi-target tracking by maintaining a measure of the mean number of targets over the state space. It does not store explicit track labels, but its ability to maintain a density of several tracks with a single set of samples makes it an attractive replacement for the search track of [4].

In this paper, we show how to use the PHD filter to derive a probability, for each measurement in the current scan, that the measurement originated from an unconfirmed target. When these probabilities reach a specified threshold, we show how to obtain the distributions of newly promoted tracks and remove them from the PHD density to allow it to search for additional targets. The distributions of newly promoted tracks will only depend on one of the measurements of the current scan, avoiding the problem of initially multi-modal tracks. Furthermore, a new track may potentially be promoted for each measurement of the current scan, removing the restriction that only one track can be promoted at a time.

Track labelling using a PHD filter has been performed before — for example, in [6], peaks are extracted by using a finite resolution grid, approximating the particles in a resolution cell by a single weight and the weighted average of particles in that cell. Alternatively, in [11], the PHD filter estimate is represented as a mixture of Gaussian components, and track estimates are taken to be the components whose weight exceeds a threshold. Our approach allows labelled tracking without the need for a grid approximation, and without constraining us to a Gaussian mixture representation.

2 Problem formulation

The problem of interest is essentially the same as that in [4], which we recap here. Suppose that at time step $k \in \mathbb{N}$ we have N_k confirmed tracks. Each track $i = 1, \ldots, N_k$ has a binary existence variable e_k^i representing whether or not it corresponds to a real target. Conditional on the existence event $E_k^i \triangleq \{e_k^i = 1\}$, the track has a state $\mathbf{x}_k^i$.

At each time step k we receive a scan of M_k measurements

$$Y^k = \{\mathbf{y}_k^1, \ldots, \mathbf{y}_k^{M_k}\} \tag{1}$$

of the targets. These measurements may include spurious measurements (clutter) and targets may not necessarily be detected. Also, which measurements originated from which targets is not known.

Given the measurements $Y^{1:k} = \{Y^1, \ldots, Y^k\}$ received so far, we wish to maintain for each track i at time step k a probability of existence

$$p(E_k^i | Y^{1:k}) \tag{2}$$

and a representation of the target state distribution conditional on its existence

$$p(\mathbf{x}_k^i | Y^{1:k}, E_k^i). \tag{3}$$

When the existence probability of a confirmed track falls below a deletion threshold, the track is deleted.

As well as maintaining a number of confirmed tracks, we also wish to have a means of searching for and confirming new tracks. This is achieved by using a PHD filter to propagate a density of the number of unconfirmed tracks over the target state space, as described in Section 4.

2.1 Target births, deaths and motion models

We assume that an existing target at time step $k-1$ will cease to exist by time step k with some known probability P^k_{death}, independently of the other targets, i.e.

$$P(E^i_k|E^i_{k-1}) \quad = \quad 1 - P^k_{death}. \tag{4}$$

Furthermore, the expected density of new targets appearing in the state space between time steps $k-1$ and k is given by $\gamma_k(\cdot)$. Targets which continue to exist are assumed to move independently of each other with a given Markov transition

$$f(\mathbf{x}^i_k|\mathbf{x}^i_{k-1}) = p(\mathbf{x}^i_k|\mathbf{x}^i_{k-1}, E^i_{k-1}, E^i_k). \tag{5}$$

2.2 Measurement model

At time step k, we receive a scan of measurements from a sensor. Note that we can have different sensors reporting at different time steps so we are not restricted to having a single sensor. However, since we receive measurements from a single known sensor at each time step, we can omit the sensor index from the notation.

Multiple sensors reporting simultaneously can be accommodated by considering them as sequential with a zero-time time step. We accept that there is some debate in the literature [8], [10] as to whether this approximation is optimal; we acknowledge our approach may be suboptimal and perceive that the approximation we adopt can be approved upon. However, we do not regard the multi-sensor PHD component of our algorithm as critical to its performance and so have simply opted for an approximation strategy that eases our software engineering effort at this stage.

We assume that each target is detected independently with some probability $P_d(\mathbf{x})$ depending on its state $\mathbf{x}$. Given that the target is detected, let

$$g(\mathbf{y}|\mathbf{x}) \tag{6}$$

be the density function of the probability distribution of the resulting measurement $\mathbf{y}$.

In addition to measurements generated by targets, a number of clutter measurements are generated. The number of clutter measurements is Poisson-distributed with mean λV, where V is the volume of the measurement space and λ is a known parameter representing the clutter density. Each clutter measurement is uniformly distributed over the measurement space.

The M_k measurements are assumed to be shuffled in some random order, with each of the $M_k!$ permutations equally likely. Hence which measurements correspond to which targets is not directly known and association probabilities must be inferred.

3 Tracking the confirmed tracks

In tracking the confirmed tracks, we follow the standard approach of predicting forward the target states and then updating based on the current scan of received measurements. This is almost identical to the approach used in [4] except for the omission of the search track. The PHD filter used for initiation is processed separately, as described in Section 4.

This part of our approach also similar to others' previous work [9], though we consider particle filters for each target, rather than a more restrictive Gaussian or Gaussian mixture approximation.

3.1 Prediction

As in [4], we see by (4) and (5) that

$$\begin{aligned} p(E_k^i|Y^{1:k-1}) &= (1-P_{death}^k)p(E_{k-1}^i|Y^{1:k-1}) && (7)\\ p(\mathbf{x}_k^i|E_k^i,Y^{1:k-1}) &= \int f(\mathbf{x}_k^i|\mathbf{x}_{k-1}^i)p(\mathbf{x}_{k-1}^i|E_k^i,Y^{1:k-1})\,d\mathbf{x}_{k-1}^i. && (8)\end{aligned}$$

3.2 Update

For each confirmed track $i = 1, \dots, N_k$, let a_k^i be the index of the measurement generated by track i, or 0 if the track is not detected. Then as in [4], we can write the updated track distribution as a mixture over the measurement hypotheses:

$$p(\mathbf{x}_k^i|E_k^i,Y^{1:k}) = \sum_{a_k^i=0}^{M_k} p(\mathbf{x}_k^i|a_k^i,\mathbf{y}_k^{a_k^i},E_k^i,Y^{1:k-1})p(a_k^i|E_k^i,Y^{1:k}). \qquad (9)$$

The updated existence probabilities can be calculated by normalizing out the association variables as follows:

$$p(E_k^i|Y^{1:k}) = \sum_{a_k^i=0}^{M_k} p(a_k^i,E_k^i|Y^{1:k}). \qquad (10)$$

Hence to calculate the updated existence probability and the track state conditional on existence, it is sufficient to calculate the track state conditional on the measurement hypothesis

$$p(\mathbf{x}_k^i|a_k^i,\mathbf{y}_k^{a_k^i},E_k^i,Y^{1:k-1}) \qquad (11)$$

and the measurement hypotheses probabilities

$$p(a_k^i,E_k^i|Y^{1:k}). \qquad (12)$$

The necessary calculations are given in [4]. We also use the EHM2 algorithm [3] to enforce the mutual exclusion constraint that no more than one of the confirmed tracks can use each measurement.

4 Performing track initiation using the PHD filter

To maintain a belief on the number and positions of unconfirmed tracks, we use a PHD filter [7]. This maintains a measure of the expected number of targets across the state space. We perceive the novelty of our approach stems from the method we propose for extricating a track from the PHD filter (as described in Section 4.2.2).

Let α_k be the a density function of the expected number of unconfirmed targets at measurement epoch k, i.e. if A is a measurable subset of the target space, then the expected number of targets in A is

$$\int_A \alpha_k(\mathbf{x})\,d\mathbf{x}. \qquad (13)$$

Note that α_k is not generally a probability density since it integrates to the expected number of unconfirmed targets present. A natural way to approximate this density is by a set of L_k weighted particles $\{w_k^i, \mathbf{x}_k^i\}$ [12]

$$\alpha_k(\mathbf{x}) \quad = \quad \sum_{i=1}^{L_k} w_k^i \delta_{\mathbf{x}_k^i}(\mathbf{x}). \tag{14}$$

The PHD filter propagates this density. This is done using the usual tracking approach of predicting the density from time epoch $k-1$ to epoch k, then updating based on the current scan of measurements Y^k.

4.1 Predicting the PHD filter

From [12], the prediction operator $\Phi_{k|k-1}$ is defined as

$$(\Phi_{k|k-1}\alpha)(\mathbf{x}) \quad = \quad \int (1 - P_{death}^k) f(\mathbf{x}|\xi)\alpha(\xi)\, d\xi + \gamma_k(\xi) \tag{15}$$

for an integrable function α on the state space. [12] also provides a method of obtaining a particle approximation to $\alpha_{k|k-1} = \Phi_{k|k-1}\alpha_k$ using importance sampling. Generally, it is desirable to sample from a proposal distribution which is conditional on the current scan of measurements Y^k, especially when sampling birth particles, to avoid wasting computational effort on large numbers of particles in areas of low measurement likelihood.

4.2 Updating the PHD filter and confirming new tracks

4.2.1 Accounting for measurements being used by the confirmed tracks

Unlike the PHD implementation in [12], our PHD filter is being used in conjunction with a separate tracker for the confirmed tracks. This means that it is necessary to account for the fact that some measurements are used by confirmed tracks and we do not want these measurements to start up additional tracks.

One possible solution is to only pass measurements which are not in the gates of any of the confirmed tracks to the initiator. This is undesirable since if a track's gate is large or targets are closely spaced, measurements from unconfirmed targets may appear in the gates of confirmed tracks and these targets may be missed by the tracker. Another approach is to weight the impact of measurements by the probability that they are unused. The probability that measurement $\mathbf{y}_k^i$ is unused by any of the confirmed tracks can be approximated by

$$p(\mathbf{y}_k^i \text{ unused}) \quad \approx \quad \prod_{t=1}^{N_k} (1 - p(a_k^t = 0|Y^k)). \tag{16}$$

This can be easily calculated from quantities calculated during the confirmed tracking stage which we perform prior to the update of new tracks. We set this to be ρ_i, a weighting factor for the measurement which we use in the PHD update step.

Another potential application of estimating the probability of a measurement being used is in online estimation of clutter density, with the impact of a measurement on the clutter density estimate being related to its probability of being unused. This is a possible avenue of future work.

4.2.2 Confirming new tracks

Adapting the PHD update operator given in [12] to incorporate the weighting factors specified in Subsection 4.2.1 gives the update operator Ψ_k:

$$(\Psi_k \alpha_{k|k-1})(\mathbf{x}) = \left[(1 - P_d(\mathbf{x})) + \sum_{i=1}^{M_k} \frac{\rho_i P_d(x) g(\mathbf{y}_k^i|\mathbf{x})}{\lambda + \int P_d(\mathbf{x}')g(\mathbf{y}_k^i|\mathbf{x}')\alpha_{k|k-1}(\mathbf{x}')\,d\mathbf{x}'}\right]\alpha_{k|k-1}(\mathbf{x}). \tag{17}$$

We decompose (17) according to the measurement hypotheses as follows:

$$(\Psi_k \alpha_{k|k-1})(\mathbf{x}) = \sum_{i=0}^{M_k} (\Psi_k^i \alpha_{k|k-1})(\mathbf{x}) \tag{18}$$

where

$$(\Psi_k^0 \alpha_{k|k-1})(x) = (1 - P_d(\mathbf{x}))\alpha_{k|k-1}(\mathbf{x}) \tag{19}$$

and

$$(\Psi_k^i \alpha_{k|k-1})(x) = \frac{\rho_i P_d(\mathbf{x}) g(\mathbf{y}_k^i|\mathbf{x})\alpha_{k|k-1}(\mathbf{x})}{\lambda + \int P_d(\mathbf{x}')g(\mathbf{y}_k^i|\mathbf{x}')\alpha_{k|k-1}(\mathbf{x}')\,d\mathbf{x}'} \tag{20}$$

for $i = 1, \ldots, M_k$. We consider $\Psi_k^0 \alpha_{k|k-1}$ to be the density of the expected number of targets which were not detected and $\Psi_k^i \alpha_{k|k-1}$ for $i = 1, \ldots, M_k$ to be the density of the expected number of targets producing measurement $\mathbf{y}_k^i$.

Note that $\int (\Psi_k^i \alpha_{k|k-1})(x)\,dx \leq 1$ for $i = 1, \ldots, M_k$. This is expected since a measurement can have come from only one target. Furthermore, the expected number of targets producing measurement $\mathbf{y}_k^i$ is equal to the probability that $\mathbf{y}_k^i$ originated from a target. Hence we can threshold on this probability to determine whether to confirm tracks. From this, we obtain a probability p_i that each measurement derives from an unconfirmed track:

$$p_i = \int \left(\Psi_k^i \alpha_{k|k-1}\right)(\mathbf{x})\,d\mathbf{x}. \tag{21}$$

Our overall tracking algorithm for each scan of measurements is therefore as follows, where P_c is a specified confirmation probability threshold:

- Predict the confirmed tracks (as described in Subsection 3.1).
- Predict the PHD filter (as described in Subsection 4.1).
- Update the confirmed tracks (Subsection 3.2) and calculate the values of ρ_i (Subsubsection 4.2.1).
- Update the PHD filter and confirm new tracks:
  - Predict α_{k-1} forward according to the prediction operator: $\alpha_{k|k-1} = \Phi_{k|k-1}\alpha_{k-1}$.
  - For each measurement hypothesis $i = 1, \ldots, M_k$ with $p_i \geq P_c$, start a new confirmed track with existence probability p_i and probability density $\Psi_k^i \alpha_{k|k-1}/p_i$.
  - Update $\alpha_{k|k-1}$ using the target densities of the remaining measurement hypotheses: $\alpha_k = \Psi_k^0 \alpha_{k|k-1} + \sum_{i:p_i<P_c} \Psi_k^i \alpha_{k|k-1}$.
- Delete the confirmed tracks whose existence probability is less than the deletion threshold.

5 Results

Here we demonstrate the tracking performance in two scenarios with simulated target trajectories and measurements, the first with a high level of clutter and low detection probability, and the second with obscured regions in the target space to demonstrate the utility of a scalable multi-target particle filter in the context of tracking through the obscurations.

In both scenarios the target position space is taken to be the region $[0, 1] \times [0, 1]$. 100 target trajectories are simulated as follows: the initial position of a target is taken to be uniform over $[0, 1] \times [0, 1]$ and the initial velocity in each of the x and y coordinates is sampled independently from a Gaussian distribution with zero mean and standard deviation 4×10^{-4}. The trajectory is propagated using a constant velocity model with noise intensity 10^{-8} in each of the x and y dimensions.

Measurements of detected targets are of x and y position, with independent Gaussian noise in each coordinate of standard deviation 7×10^{-4}. Measurement epochs occur at intervals of 1 second for 180 seconds. Each scan contains a number of clutter measurements which is different for each of the two scenarios, and the clutter measurements are uniformly distributed over the position space.

The tracker in each scenario will promote new tracks from the PHD filter initiator if the existence probability calculated in (21) reaches 0.9, and delete confirmed tracks if their existence probabilities drop below 0.1. We also use importance sampling to draw birth particles for the PHD filter from a mixture proposal [2] which includes a component which draws samples from “under” a (randomly sampled for each particle) measurement [1].

5.1 Tracking through dense clutter

In this scenario, we set the number of clutter measurements per scan to be Poisson-distributed with mean 500, and the probability of detection to be 0.5. This is to show the tracker’s performance in a high-clutter, low detection probability environment. The measurements and true target positions at the first time step are shown in Figure 1(a). 100000 particles are used for the PHD filter and 1000 particles for each of the confirmed tracks. The tracker takes 3349 seconds to run on a 3GHz PC.

We assert that particle filters are well suited to tracking of targets in dense clutter since, conditional on the path through the state space, the association hypotheses are independent over time. Hence, for each particle, no association tree needs to be explored. This is reminiscent of the ML-PDA [5] though our approach has the advantage that the trajectories are inherently stochastic, rather than being parametric curves. Note that both our approach (and the ML-PDA) completely circumvent the computational explosion that is encountered by Multi-Hypothesis Trackers (MHTs). MHTs (whether explicitly or implicitly) consider association histories. When the clutter rate increases, the breadth of the association tree grows and so the depth that can be accommodated with a given number of hypotheses falls. It is the authors belief that this phenomenon limits the clutter levels that can be accommodated by MHTs since such an MHT must hurry when the association becomes more ambiguous. This is highly undesirable since ambiguity should introduce an implicit lag into the filter while the ambiguity is resolved. The authors perceive that the particle filter approach described herein does not suffer from this problem, can delay (implicit) association decisions in response to ambiguity and so is surprisingly effective relative to the state-of-the-art in high clutter environments

Figure 1(b) shows the confirmed tracks from the tracker and the true target trajectories. We see that in spite of the high clutter density, the tracker manages to track the vast majority of the targets.

5.2 Tracking through obscurations using a map

Using a particle filter-based tracker allows us to deal with nonlinear constraints such as a state-dependent probability of detection. We demonstrate this with the following tracking scenario.

The target position space is split into a 20×20 grid of squares, each of which independently has a 0.2 probability of being obscured. A target in an unobscured region is assumed to have a probability of 0.9 of being detected in each scan, whereas targets in obscured regions are not detected. The number of clutter measurements for each scan is Poisson-distributed with mean 10.

We track the data with two variations — in the first case, we assume that a map of the obscured regions is available and so we can use the state-dependent probability of detection in the algorithm. In the second case, for comparison, we assume no knowledge of the obscuration regions and use a fixed probability of detection of 0.72 (accounting for the fact that 0.8 of the region is obscured and an unobscured target is detected with probability 0.9). In both cases, we use 10000 particles for the PHD filter and 1000 particles for each of the confirmed tracks.

The tracker takes 844 seconds to run for the obscuration map scheme, and 756 seconds for the fixed probability of detection scheme, on the same 3GHz PC as in Subsection 5.1 above. The difference in timings here is likely due mostly to the increased number of tracks confirmed by the obscuration map tracker.

Figure 2 shows the confirmed tracks output under each scheme. We see that using an obscuration map to accurately model the state-dependent probability of detection allows better tracking of targets through the obscurations. To further illustrate the point, Figure 3(a) shows the number of confirmed tracks at each time step for each scheme compared to the true number of targets (note that the true number of targets decreases from 100 since some of them leave the target region). We see that using the obscuration map does a better job of estimating the number of targets at each time step. Furthermore, Figure 2(b) shows a histogram of the lengths of tracks, showing that using a fixed probability of detection in this case produces a large number of fragmented tracks.

6 Conclusions

We have adapted a previous tracking algorithm which uses existence probabilities to perform track confirmation and deletion in order to deal with shortcomings arising from trying to model multiple unconfirmed tracks with a single-target “search” track. The new approach has been tested on scenarios with many targets present, and found to perform well even in the presence of high levels of clutter. Provided that the obscurations are known, the tracker can also perform well in the presence of obscured regions where targets cannot be detected.

7 Acknowledgements

This work has been supported by the UK Ministry of Defence.

References

[1] A. Doucet, S. Godsill and C. Andrieu, “On sequential Monte Carlo sampling methods for Bayesian filtering”, *Statistics and Computing*, Vol. 10, 2000.

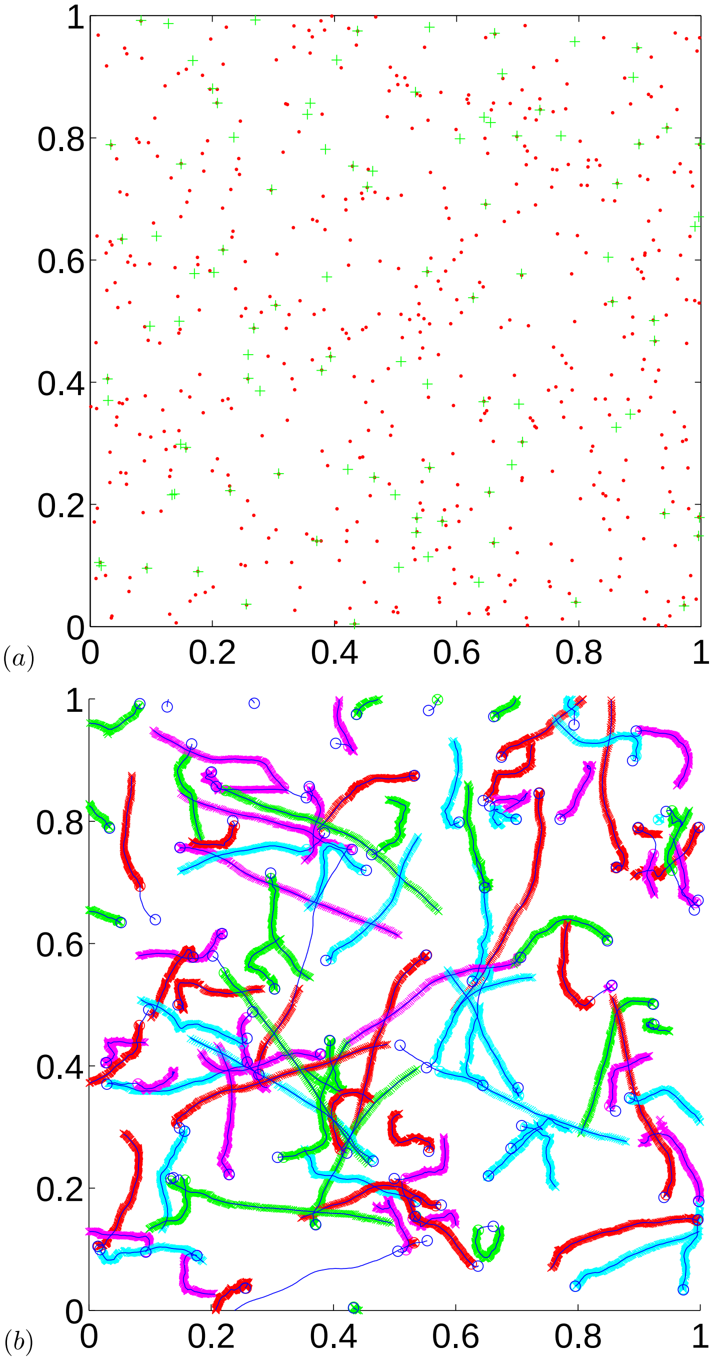

Figure 1: In the high clutter scenario, (a) measurements (red) and true target positions (green) for the first scan, and (b) tracks (multiple colours) and ground truth (dark blue).

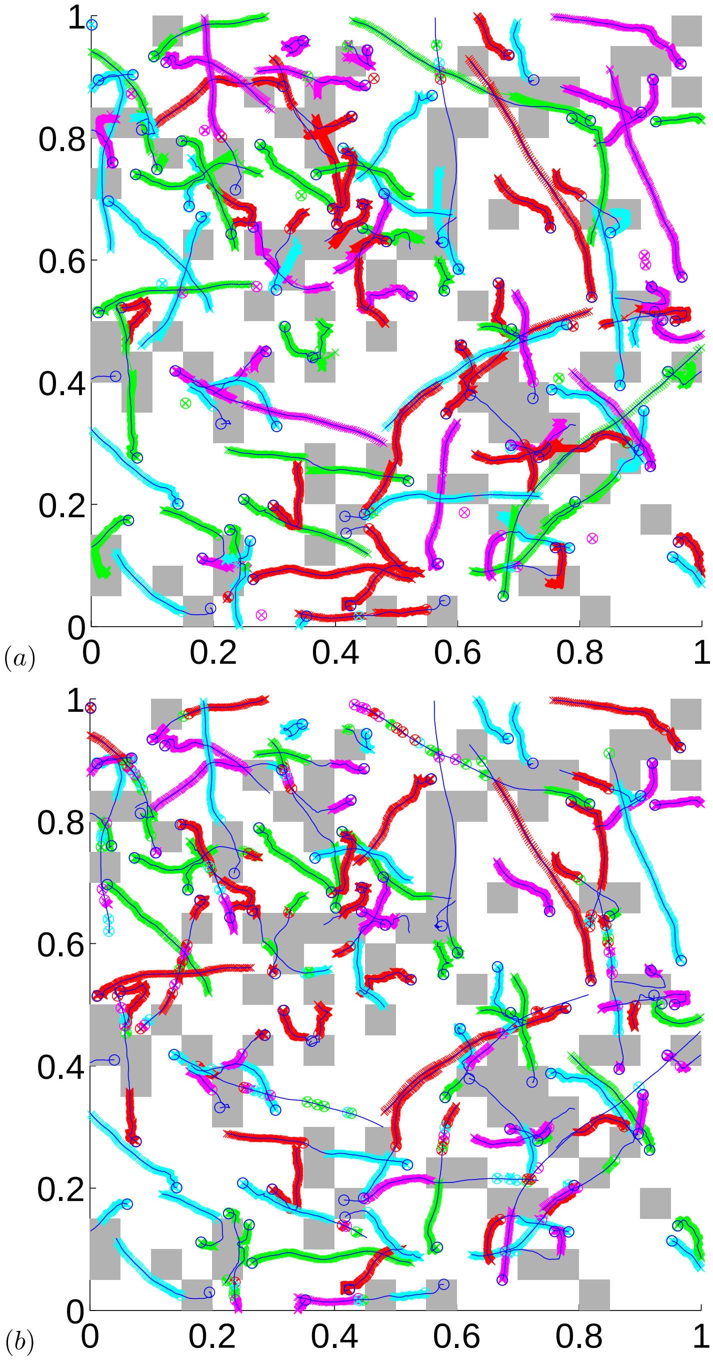

Figure 2: Tracks (multiple colours) and ground truth (dark blue) in the obscuration scenario, when (a) using an obscuration map, and (b) using a fixed probability of detection.

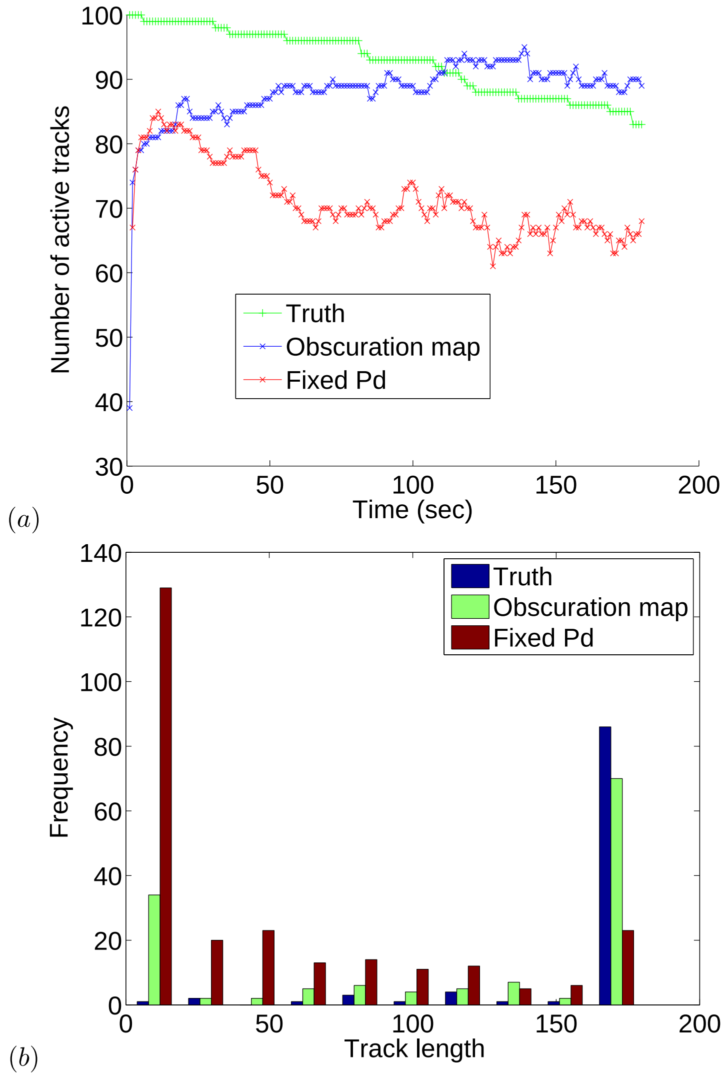

Figure 3: (a) Number of confirmed tracks and true number of targets present over time for the obscuration scenario. (b) Histogram of track lengths and true target lifetimes for the obscuration scenario.

[2] T. Hesterberg, “Weighted Average Importance Sampling and Defensive Mixture Distributions”, *Technometrics*, Vol. 37, No. 2, 1995.

[3] P. Horridge and S. Maskell, “Real-Time tracking of hundreds of targets with efficient exact JPDAF implementation”, *International Conference on Information Fusion*, 2006.

[4] P. Horridge and S. Maskell, “Searching for, initiating and tracking multiple targets using existence probabilities”, *International Conference on Information Fusion*, 2009.

[5] C. Jauffret and Y. Bar-Shalom, “Track formation with bearing and frequency measurements in clutter”, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 26, 1990.

[6] L. Lin, Y. Bar-Shalom and T. Kirubarajan, “Track labeling and PHD filter for multitarget tracking”, *IEE Transactions on Aerospace and Electronic Systems*, Vol. 42, No. 3, 2006.

[7] R. Mahler, “A theoretical foundation for the Stein-Winter Probabilistic Hypothesis Density (PHD) multi-target tracking approach”, *MSS National Symposium on Sensor and Data Fusion*, Vol. 1, San Antonio, Texas, 2000.

[8] R. Mahler, “The multisensor PHD filter: I. General solution via multitarget calculus”, *Proceedings of SPIE 7336*, 73360E, 2009.

[9] D. Musicki and R. Evans, “Joint integrated probabilistic data association: JIPDA”, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 40, No. 3, 2004.

[10] R. Streit, “Multisensor multitarget intensity filter”. *Proceedings of Fusion*, 2008.

[11] B. Vo and W. Ma, “The Gaussian mixture probability hypothesis density filter”, *IEEE Transactions on Signal Processing*, Vol. 54, No. 11, 2006.

[12] B. Vo, S. Singh and A. Doucet, “Sequential Monte Carlo implementation of the PHD filter for sequential tracking”, *International Conference on Information Fusion*, 2003.