Tracking in Dense Clutter With The PHD Filter

By Trevor M Wood

Maths Institute, 24-29 St. Giles, University of Oxford, OX1 3LB [woodtm@maths.ox.ac.uk]

Abstract

The probability hypothesis density (PHD) filter is an alternative to currently used methods for multitarget tracking. In this paper, we present a performance comparison between the Gaussian mixture approximation to the PHD filter and the conventional nearest neighbour (NN) and probabilistic data association (PDA) filters. Results show a performance advantage comparable to the reported advantage for the multiple hypothesis tracker (MHT).

1. Introduction

When tracking using either active sonar or radar, a transmitter sends out a pulse. The pulse propagates outwards, bouncing off anything it encounters, and returns to a detector. The analysis of the returning signal will yield a set of detections of potential objects. Some of these detections will originate from targets that are of interest and must be tracked but there may also be many false alarms, known as clutter. In the sonar case, for example, clutter can be caused by a number of things such as reverberations of the pulse from the surface of the sea and the seabed. In dense clutter, where there are many more false alarms than detections of true targets, it is the difficult job of multitarget tracking to determine from this noisy data how many targets are present and where they are located.

In the case of a single target and a single set of measurements, the methods used for tracking, such as the Kalman filter, are well known. The solution when confronted with the multitarget tracking problem has typically been to use a heuristic method for reducing the situation back to the single target case. This usually means using a data association algorithm, as well as a track management algorithm, to decide when to initialise and delete tracks (see, for example, Blackman & Popoli 1999) .

The finite set statistics (FISST) framework, largely due to Mahler, provides the tools to perform a Bayesian inference on the full set of target states, conditioned on the set of measurements. The Bayesian framework is natural due to the prevalence of the Kalman and particle filters, which are both linked to the Bayes filter, in existing tracking methods. FISST allows a unified tracking process, with no need for separate heuristics for data association or track management (see Mahler 2007). This unified tracker is known as the multitarget Bayes filter.

The multitarget Bayes filter from FISST is computationally intractable and hence must be approximated. Propagating the 1st moment of the multitarget posterior, called the probability hypothesis density (PHD), gives the PHD filter. The PHD filter itself must also be approximated for tractability, and two known methods use sequential Monte Carlo methods (SMC-PHD) and Gaussian mixtures (GM-PHD), as described in Vo et. al 2005 and Vo & Ma 2006 respectively.

There are several examples in the literature of successful implementations of the PHD filter. This study will focus on a practical analysis of the performance of the PHD filter in dense clutter as compared with competing existing tracking methods, namely those based on nearest neighbour (NN) and probabilistic data association (PDA). It is hoped that this will allow a prediction of expected performance particularly for active sonar and radar systems.

2. A Brief Background of Random Finite Sets and the PHD Filter

In a single target system, the state and measurement of a target at a given time are two vectors, possibly with different lengths. In this case, it is normal to use a statistical dynamical model for the evolution of the state and another statistical model for how the measurement relates to the state. From this starting point, it is possible to make an inference about the state, conditioned on the measurements received using Bayes rule, and from this an estimate can be extracted. Often, the full computation of the posterior on the state space is computationally intractable, and so well known approximations such as the Kalman filter or particle filter are used.

In a multitarget system, it is necessary to estimate *how many* targets are present, as well as *where* they are. In FISST, the full set of target states is treated as a single entity called the multitarget state. The multitarget state is then a random finite set, which means that the number of members is random, and the members themselves (the single target state vectors) are also random. Similarly, while a single measurement is a random vector, the full set of measurements is also a random finite set.

In the single target methods, there is a prediction step and an update step; the same is true in the multitarget FISST approach. The prediction of the multitarget state is analogous in the expected way to the single target case. Each individual target is predicted according to the single target dynamical model, but the multitarget prediction also incorporates statistical models for target appearance and disappearance, target spawning and, in some cases, co-ordinated target motion. This is the multitarget equivalent to the Markov transition density in the Bayes filter. Similarly, the update step is analogous in the expected way to the single target case, but the multitarget update also incorporates statistical models for missed detections, clutter and, in some cases, state dependent effects on sensor field of view. This is the multitarget equivalent to the likelihood function of the Bayes filter.

The single target Bayes filter is computationally intractable and so it is unsurprising that the multitarget Bayes filter is also. Just as the Kalman filter approximates the single target Bayes filter by using only two moments (mean and covariance) a similar moment matching approximation for the multitarget Bayes filter has become prevalent. This is the PHD filter which propagates the PHD; this is the first moment of the multitarget Bayes filter. The PHD has a useful physical interpretation as a density of expected targets. This means that integrating the PHD over a region of the state space will give the expected number of targets in that region. The PHD filter requires that the following assumptions are made for the multitarget state X_k and measurement Z_k at time k.

$$X_k = \left(\bigcup_{x \in X_{k-1}} S_{k|k-1}(x) \right) \cup \left(\bigcup_{x \in X_{k-1}} B_{k|k-1}(x) \right) \cup \Gamma_k \tag{2.1}$$

$$Z_k = K_k \cup \left(\bigcup_{x \in X_{k-1}} \Theta_k(x) \right) \tag{2.2}$$

where $S_{k|k-1}$ represents targets surviving from the previous time step and includes the single target dynamical model, $B_{k|k-1}$ represents targets spawned from existing targets and includes the motion model for spawned targets, Γ_k represents new targets, K_k represents false alarms and Θ_k represents target detections and includes models for missed detections as well as the single target measurement model.

A clear exposition of the details of the PHD filter is available in the PhD thesis by Clark, while a more in depth look at the details and derivation may be found in Mahler 2007.

3. Implementation

Broadly speaking, there are two types of scheme for implementing the PHD filter. These are i) Sequential Monte Carlo methods (SMC-PHD) ii) Gaussian Mixtures (GM-PHD). In order to be more succinct This paper will restrict focus to GM-PHD, which we presently consider to be the most promising for widespread implementation. It is worth pointing out that, despite perhaps confounding references to set theory, the GM-PHD filter is effectively a bank of Kalman filters, internally managed via the probability hypothesis density. Thus, the departure from current methods is more in the theory than the detail.

The details of the GM-PHD algorithm may be found in Vo & Ma 2006, which includes many implementation details such as the process of pruning, merging and limiting Gaussian components to prevent a combinatorial explosion. A scheme for data association within the GM-PHD algorithm is available in Panta et. al 2006, who also suggest some more complex schemes similar to the N-scan pruning used in multiple hypothesis tracking. However, it is our opinion, formed during extensive testing, that the simpler schemes suggested are sufficient in most tracking situations. It is shown in Mahler 2007 that the GM-PHD filter has the desirable computational characteristic of being linearly proportional to the number of targets and the number of measurements.

4. Simulation Results

For the purpose of comparing the performance of the FISST methods with conventional methods, we considered a two-dimensional scenario with one target observed in clutter over the region [-40,40]x[0,40] (which can be thought of as km). The simulated point target moves in a straight line for 50 time steps. As there is exactly one target at each time step, equation (2.1) for the multitarget state at time step k simplifies to $X_k = \{x_k\} = \{[p_x, \dot{p}_x, p_y, \dot{p}_y]_k\}$, where (p_x, p_y) and $(\dot{p}_x, \dot{p}_y)$ are the positions and velocities of the target respectively. Similarly, equation (2.2) simplifies to $Z_k = K_k \cup \hat{Z}_k$ where K_k is uniformly distributed clutter over the region and the number of clutter points is Poisson distributed with mean (refered to as the clutter level) ranging from 10 to 2000 and:

$$\hat{Z}_k = \begin{cases} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} x_k + w_k & \textit{with probability } P_D \\ \emptyset & \textit{else} \end{cases} \tag{4.1}$$

where w_k is a multivariate Gaussian random variable with covariance matrix $\sigma^2 \mathbf{I}_2$. For all of the tests here $\sigma = 0.3$.

Previous evaluations of PHD filter implementations have tended to focus on the Wasserstein distance (see, for example, Vo et. al 2005). We will instead adopt a more practical, but less mathematically rigorous, method following the approach taken in comparing the performance of data association schemes in Chapter 6 of Blackman & Popoli 1999. Better performance in the tests conducted here would be consistent with a lower average

Wasserstein distance, as described in Hoffman & Mahler 2004. However, it is hoped that the results presented here will allow an easier comparison and prediction of expected performance gain using the GM-PHD filter.

4.1. *Test 1: Track Maintenance*

1000 Monte Carlo tests were run for several levels of clutter and $P_D = 1$ and 0.75 . The track on the target was initialised semi-automatically by identifying the measurement which originated from the target on the first time step. A track was deemed to have been maintained if a track was reported within 1 of the true target location for 80% of the time steps. Gaussian components in the GM-PHD filter come with an associated weight, which can be interpreted as the probability of representing a real target. For this test, a track is considered to be reported in GM-PHD if the associated Gaussian component has weight > 0.2. The results are presented in Table 1.

In Blackman & Popoli 1999, the results for a comparison of NN, PDA and multiple hypothesis tracking (MHT) lead to the conclusion that PDA gives comparable performance at "several" times the clutter level as NN while MHT gives comparable performance for 10-100 times the clutter level. Seeking similar generalities within the results in Table 1, it seems that the comparison between PDA and NN would be the same, whereas GM-PHD is able to give comparable performance to NN at 10-20 times the clutter level, a little worse than the reported results for MHT.

There are, however, some mitigating factors for this worse performance. The first is the increased complexity of implementing an MHT filter. Blackman's MHT implementation which provided the '10-100' result took "6 person-months to develop". Panta et. al 2004 state that MHT is more complicated to implement than PHD and that "the performance of the MHT filter depends heavily upon the particular implementation... techniques that are ad-hoc in general". The second factor is the choice of criteria for a maintained track. In Bar-Shalom & Tse 1975, a comparative study of NN and PDA using a similar approach to the one used here, the criteria for a maintained track is to have one update in the last 20 time steps. This seems a little weak to claim a track has been maintained for the whole run, but this criteria would suit GM-PHD well. This is because lost tracks in conventional methods tend to be lost entirely and often early on (which explains Bar-Shalom and Tse's choice of criteria), whereas track losses in GM-PHD are usually due to the weight falling below the threshold too many times. There would be similar issues when deciding what constitutes a maintained track with MHT, but Blackman and Popoli do not spell out their criteria. It is hoped that the simple criteria presented above will be considered a reasonable starting point. Incidentally, had we chosen Bar-Shalom's criteria, a smaller set of tests indicate that the GM-PHD's performance improvement over GNN would have exceeded that reported for MHT by Blackman and Popoli.

4.2. *Test 2: Initialisation*

100 Monte Carlo tests were run for several levels of clutter. The target must be fully detected and initialised by the trackers. Two initialisation schemes were used for the conventional trackers. The first scheme was 'M/N' whereby a target is initialised if it is updated M_1 out of N time steps and deleted if there is no update on M_2 out of N time steps. Here $M_1 = 3$, $M_2 = 5$ and $N = 5$. The second scheme was the sequential probability ratio test (SPRT) scheme described in Chapter 6 of Blackman & Popoli 1999 with the track confirmation threshold set so that the false track confirmation probability is 10^{-3}. As in Test 1, a track is considered reported by the PHD filter if the associated Gaussian has weight > 0.2. We consider the average number of false tracks reported per time step and the time taken to detect and report the true target. Results are only

	$P_D = 1$			$P_D = 0.75$		
clutter	NN	PDA	PHD	NN	PDA	PHD
10	1	1	1	0.96	0.97	1
20	0.99	1	1	0.94	0.97	1
50	0.97	0.99	1	0.92	0.96	1
100	0.93	0.99	1	0.84	0.95	0.99
200	0.88	0.99	1	0.73	0.93	0.99
500	0.65	0.91	1	0.47	0.47	0.92
1000	0.38	0.37	0.90	0.23	0.02	0.71
2000	0.14	0.09	0.41	0.07	0	0.14

TABLE 1. Proportion of tracks maintained from 1000 Monte Carlo Tests.

clutter	NN - M/N		NN - SPRT		PDA - M/N		PDA - SPRT		GM - PHD	
	FT	Decl.	FT	Decl.	FT	Decl.	FT	Decl.	FT	Decl.
10	0.5	3.1	0	3.2	1.9	3	0.2	3.1	0.2	2
20	2.7	3.1	0	3.5	7.2	3	0.3	3.5	0.4	2.2
50	23.1	3.1	0	4.6			0.9	4.4	0.3	3.1
100							1.3	5.2	0.3	4
200							1.2	6.6	0.3	5.4
500									0.1	10.4

TABLE 2. Number of false tracks (FT) declared per time step and time taken to declare real target in full tracking scenario (Decl.) averaged over 100 Monte Carlo tests.

reported in those cases where the target was detected in more than 75% of the Monte Carlo tests and the level of false tracks declared is less than 50 per time step, as any tracker that cannot meet both of these criteria would be considered too unreliable for use. The results are presented in Table 2.

The results in Table 2 show a performance improvement for the GM-PHD filter in initialisation which is comparable to its improvement in track maintenance. It should also be noted that the GM-PHD filter is able in all tests to detect and declare the real target a full time step before the SPRT scheme while also having a lower level of reported false targets.

5. Conclusions

The paper has undertaken a performance comparison of the GM-PHD filter with NN and PDA. It was argued in Section 2 that the general approach of the PHD filter is more satisfactory than that of the conventional methods. The PHD filter is a natural extension of the Bayes filter based methods used for single target tracking rather than being something tacked on to the single target methods. The incorporation of explicit statistical models of the effects making the tracking problem difficult are also an advantage and can give the FISST methods more flexibility.

It was found that substantial performance improvements were possible, with the biggest improvements found in high levels of clutter. This would suggest a direct application to active sonar in shallow water highly reverberant environments. Given the results presented in Section 4, the PHD filter might be considered an alternative to MHT. Attempts to extrapolate the relative performances of GM-PHD and MHT by looking at their respective comparisons with the other filters cannot be fully satisfactory and further work comparing the MHT and PHD filters directly would be valuable. The PHD filter has some preferable characteristics; it is less complex to implement. It is also expected that due to the computational complexity being linearly proportional to both the number of targets and measurements, that the GM-PHD will also have a computational advantage over MHT, but again, a direct comparison between GM-PHD and MHT would be useful.

This paper has focused only on a linear Gaussian case with a single non-manoeuvring target for simplicity. However, extensions to nonlinear measurement models (Vo & Ma 2006) and manoeuvring targets (Punithakumar et. al 2008) are well known. The filter implemented here has been adapted for use on sea trial data using a towed array in a highly cluttered shallow water environment tracking manoeuvring targets with measurements in range and bearing. Similar performance was achieved.

Acknowledgements

This work is generously supported by EPRSC, Thales Underwater Systems and Thales Aerospace. I would also like to give thanks to my academic supervisors David Allwright, Philip Bond and Irene Moroz as well as my industrial supervisors Glen Davidson and Stephen Long (both from Thales) for guidance and useful discussions.

REFERENCES

Y. Bar-Shalom and E. Tse Tracking in a Cluttered Environment with Probabilistic Data Association *Automatica* **2** 451-460, 1975

S. Blackman and R. Popoli Design and Analysis of Modern Tracking Systems *Artech House*, 1999.

D.E. Clark Target Tracking with the Probability Hypothesis Density Filter *PhD Thesis, Heriot-Watt University*, 2005

J.R. Hoffman and R.P.S. Mahler Multitarget miss distance via optimal assignment *IEEE Transactions on Systems, Man and Cybernetics - Part A* **34** 327–336, 2004

R.P.S. Mahler Statistical Multisource-Multitarget Information Fusion *Artech House*, 2007

K. Panta, B. -N. Vo, S. Singh and A. Doucet Probability hypothesis density filter versus multiple hypothesis tracking *Signal Processing, Sensor Fusion and Target Recognition XIII* 284-295, 2004

K. Panta, B. Vo, and D.E Clark An efficient track management scheme for the gaussian-mixture probability hypothesis density tracker *Fourth International Conference on Intelligent Sensing and Information Processing* 230–235, 2006

K. Punithakumar, T. Kirubarajan, A. Sinha Multiple-Model Probability Hypothesis Density Filter for Tracking Manoeuvring Targets *IEEE Transactions on Aerospace and Electronic Systems* **44** 87–98, 2008

B.-N. Vo, S. Singh, A. Doucet Sequential Monte Carlo Methods for Multitarget Filtering with Random Finite Sets *IEEE Transactions on Aerospace and Electronic Systems* **41** 1224–1245, 2005

B.-N. Vo and W. Ma The Gaussian Mixture Probability Hypothesis Density Filter *IEEE Transactions on Signal Processing* **54** 4091–4104, 2006