



**CHALMERS**  
UNIVERSITY OF TECHNOLOGY

---

# **A Study of Poisson Multi-Bernoulli Mixture Conjugate Prior in Multiple Target Estimation**

Sampling-based Data Association, Multi-Bernoulli Mixture Approximation and Performance Evaluation

Master's thesis in Communication Engineering

YUXUAN XIA



MASTER'S THESIS 2017:Ex091

# A Study of Poisson Multi-Bernoulli Mixture Conjugate Prior in Multiple Target Estimation

Sampling-based Data Association, Multi-Bernoulli Mixture  
Approximation and Performance Evaluation

YUXUAN XIA



**CHALMERS**  
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering  
*Division of Signals Processing and Biomedical Engineering*  
CHALMERS UNIVERSITY OF TECHNOLOGY  
Gothenburg, Sweden 2017

A Study of Poisson Multi-Bernoulli Mixture Conjugate Prior in Multiple Target Estimation  
YUXUAN XIA

© YUXUAN XIA, 2017.

Supervisor: Karl Granström, Lennart Svensson, Department of Electrical Engineering  
Examiner: Karl Granström, Department of Electrical Engineering

Master's Thesis 2017:EX091  
A Study of Poisson Multi-Bernoulli Mixture Conjugate Prior in Multiple Target Estimation  
Chalmers University of Technology  
SE-412 96 Gothenburg  
Telephone +46 31 772 1000

Typeset in L<sup>A</sup>T<sub>E</sub>X  
Printed by [Chalmers University of Technology]  
Gothenburg, Sweden 2017

A Study of Poisson Multi-Bernoulli Mixture Conjugate Prior in Multiple Target Estimation

YUXUAN XIA

Department of Electrical Engineering  
Chalmers University of Technology

## Abstract

Multiple target tracking (MTT) denotes the process of estimating the set of target trajectories based on a sequence of noise-corrupted measurements including missed detections and false alarms. Nowadays, MTT has found applications in numerous areas, such as air traffic control, autonomous vehicles and robotics, computer vision and biomedical research. In recent years, a significant research trend in MTT is the development of conjugate distributions in Bayesian probability theory based on random finite set. Two popular frameworks have been studied so far, one is based on labelled multi-Bernoulli conjugate prior, and the other is based on unlabelled multi-Bernoulli conjugate prior. The first contribution of this thesis is a performance comparison study of filters based on multi-Bernoulli conjugate prior. In this part of work, we focus on point target tracking that each target is assumed to give rise to at most one measurement per time scan. The simulation results show that the Poisson multi-Bernoulli filters arguably provide the best overall performance.

Due to the rapid development of high-resolution sensors equipped on autonomous vehicles, e.g., near-field radar and lidar, a target may occupy multiple sensor cells on any given scan, leading to the so-called extended target. Solving the multiple extended target tracking problem is mainly complicated by the unknown correspondence between targets and measurements that a huge number of data association events need to be considered. Methods of how to solve the data association in a single step that maximises the desired likelihood function using sampling methods are presented. The second contribution of this thesis is the performance evaluation of different sampling algorithms, which are integrated into the Poisson multi-Bernoulli mixture (PMBM) filter.

As an approximation of the PMBM filter, the Poisson multi-Bernoulli (PMB) filter has shown superior performance in point target tracking, but it is not yet clear how to adapt this algorithm to extended target tracking. The third contribution of this thesis is that we present an extended target PMB filter, along with its gamma Gaussian inverse Wishart implementation. The simulation results show that the PMB filter can retain most of the advantages of the PMBM filter.

Keywords: multiple target tracking, extended targets, random finite sets, random matrix model, Bayesian estimation, conjugate prior, sampling methods, variational inference.



## Acknowledgements

My deepest thanks goes to my supervisor Karl Granström for his timely support, valuable academic discussions and elaborate feedback. Many thanks also to my supervisor Lennart Svensson for his insightful comments, motivation, and immense knowledge. Their guidance helped me in all the time of research and writing of this thesis. I could not have imagined having a better advisor and mentor for my Master thesis.

Besides my supervisors, I would like to thank Ángel F. García-Fernández at Aalto University for his valuable discussion and thoroughly feedback on the Fusion paper. The paper would have not been a success without his collaboration.

Last but not the least, I would like to thank my family: my parents, my uncle and aunt, my grandparents for supporting me spiritually throughout writing this thesis and my my life in general.

Yuxuan Xia, Gothenburg, September 2017



# List of Publications

This thesis is based on the work contained in the following papers:

## **Paper I**

Yuxuan Xia, Karl Granström, Lennart Svensson, and Ángel F. García-Fernández. “Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering”, in 20th International Conference on Information Fusion, Xi’an, July 2017.

**Paper II** Karl Granström, Lennart Svensson, Stephan Reuter, Yuxuan Xia and Maryam Fatemi, “Likelihood-based data association for extended object tracking using sampling methods”, Submitted to IEEE Transactions on Intelligent Vehicles.

**Paper III** Yuxuan Xia, Karl Granström, Lennart Svensson and Maryam Fatemi, “Poisson multi-Bernoulli filter for extended object estimation”, to be submitted.



# Contents

<b>Introduction</b>	<b>1</b>
1.1 Prospect . . . . .	1
1.2 Motivation and Main Focus . . . . .	2
1.3 Contributions of the Appended Papers . . . . .	4
1.4 Conclusions and Future Work . . . . .	5
<b>Paper I: Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering</b>	<b>7</b>
2.1 Introduction . . . . .	7
2.2 Background . . . . .	8
2.2.1 Bayesian Filtering Recursion . . . . .	10
2.2.2 Standard Point Target Transition and Measurement Model . . . . .	10
2.3 Labelled Multi-Bernoulli Conjugate Prior . . . . .	11
2.3.1 Delta Generalised Labelled Multi-Bernoulli . . . . .	11
2.3.2 Labelled Multi-Bernoulli . . . . .	12
2.3.3 Poisson Multi-Bernoulli Mixture . . . . .	12
2.3.4 Poisson Multi-Bernoulli . . . . .	13
2.4 Approximations for Computational Tractability . . . . .	14
2.5 Simulations and Results . . . . .	14
2.5.1 State Space Models . . . . .	15
2.5.2 Performance Evaluation . . . . .	15
2.5.3 Scenario with Well-spaced Targets . . . . .	16
2.5.4 Scenario with Coalescence . . . . .	18
2.6 Conclusions and Future Work . . . . .	20
<b>Paper II: Likelihood-based data association for extended object tracking using sampling methods</b>	<b>23</b>
3.1 Introduction . . . . .	23
3.2 Problem formulation, previous work, and proposed approach . . . . .	25
3.2.1 Problem formulation . . . . .	25
3.2.2 Previous two-step approach: clustering and assignment . . . . .	26
3.2.3 Proposed single-step approach . . . . .	27
3.3 Multiple extended object tracking . . . . .	28
3.3.1 Random set modelling . . . . .	28
3.3.2 Standard extended object measurement model . . . . .	28
3.3.3 The PMBM conjugate prior . . . . .	29

3.3.4	PMBM data association . . . . .	30
3.4	Multi-object density approximation . . . . .	31
3.4.1	Multi-object density approximation . . . . .	32
3.4.2	Approximation error upper bound . . . . .	32
3.4.3	Weight approximations . . . . .	33
3.5	Sampling methods . . . . .	37
3.5.1	Assignment variable . . . . .	37
3.5.2	MCMC sampling methods . . . . .	38
3.5.2.1	Gibbs Sampling . . . . .	38
3.5.2.2	Split/Merge Metropolis-Hastings . . . . .	38
3.5.3	Stochastic optimisation . . . . .	41
3.5.4	Implementational aspects . . . . .	42
3.5.4.1	Reducing the cost of the likelihood computation . . . . .	42
3.5.4.2	Initialisation . . . . .	42
3.5.4.3	Finding a subset of associations . . . . .	43
3.6	Simulation results . . . . .	43
3.7	Experimental results . . . . .	44
3.8	Concluding remarks . . . . .	48

**Paper III: Poisson multi-Bernoulli filter for extended object estimation** **53**

4.1	Introduction . . . . .	53
4.2	Background . . . . .	55
4.2.1	Bayesian multi-object filtering . . . . .	55
4.2.2	Random set modelling . . . . .	56
4.2.3	Standard extended target measurement model . . . . .	57
4.2.4	Standard target transition model . . . . .	57
4.2.5	PMBM conjugate prior and data association . . . . .	57
4.3	Poisson Multi-Bernoulli Filter . . . . .	58
4.3.1	PMB filter recursion . . . . .	58
4.3.2	Pre-existing track formation . . . . .	60
4.3.2.1	Optimal assignment . . . . .	60
4.3.2.2	Efficient approximation of feasible set . . . . .	61
4.3.3	New track formation . . . . .	62
4.3.4	Recycling . . . . .	63
4.4	GGIW Implementation . . . . .	64
4.4.1	Single target models . . . . .	64
4.4.2	MBM merging . . . . .	65
4.4.2.1	E-step . . . . .	65
4.4.2.2	M-step . . . . .	66
4.5	Simulations and Results . . . . .	67
4.5.1	State space model . . . . .	67
4.5.2	Performance evaluation . . . . .	67
4.5.3	Simulation study . . . . .	68
4.6	Conclusions . . . . .	69

**Bibliography** **71**

# Introduction

This introduction section is organised as follows. Some future prospects are introduced in Section 1.1. Motivation and the main focus of the thesis are discussed in Section 1.2, and main contributions of each appended paper are described in Section 1.3. We summarise the conclusions and discuss research ideas for future extensions in Section 1.4, and the outline of the thesis is presented in Section 1.5.

## 1.1 Prospect

The Gothenburg region is the heart of the Swedish automotive industry. Many well-known companies in the region, such as Volvo Cars, Autoliv and Zenuity, have put a lot of efforts in the research and development of autonomous vehicles. For example, there will be 100 self-driving cars operating on public roads in Gothenburg in 2017, and further research towards more advanced and mature autonomous driving technologies will continue.

The emergence of the concept of self-driving vehicles has been a driving force for the rapid development of autonomous driving technologies. In order for a car to navigate without human input, the advanced control system should make correct decisions based on the environment as they perceive it. An important part of the environment perception technology is multiple target tracking (MTT) techniques, which are vital for collision avoidance and traffic safety. They enable the moving autonomous vehicle to reliably and efficiently model static as well as dynamical obstacles, such as landmarks, pedestrians and other nearby moving vehicles. The MTT problem is the process of estimating the set of target trajectories based on a sequence of noise-corrupted measurements including missed detections and false alarms. In the context of autonomous driving, the targets being tracked are the moving objects around the self-driving vehicle. Typical sensors used to collect measurements in self-driving vehicles include the near-field radar and lidar.

Nowadays, MTT techniques have attracted intensive attentions and found numerous applications in various fields, for example, land, sea, air, and space surveillance for military use; and collision avoidance, navigation, and image processing for civilian use. During the last decade, advances in MTT techniques, along with modern sensing and computing technologies, has opened up a remarkable research venue in the automotive industry. We believe that this thesis includes important contributions to the MTT research field.

## 1.2 Motivation and Main Focus

In many MTT applications, a recursive filtering approach is employed as it allows estimates to be made each time new observations are received without having to reprocess existing observations. Popular solutions to MTT include the Joint Probabilistic Data Association filter [1], Multiple Hypothesis Tracker [2] and approaches based on random finite sets (RFS) [3]. An RFS is a set of random variables whose cardinality is also a discrete random variable. By representing the targets and measurements in the form of finite sets, RFS-based methods can provide a convenient mechanism for solving two of the major challenges involved in MTT: one is the unknown number of targets presented in the surveillance area, and the other is the data association problem due to the unknown measurement-to-target correspondence.

In this work, we focus on the estimation of the current set of targets, which refers to multiple target filtering. The random set theory was first developed by Mahler [3] to cast the multiple target filtering problem in the Bayesian framework. The Bayes filter is used to predict and update the multitarget set density recursively, so that the uncertainty of target number and states can be propagated in time. The multitarget set posterior density, i.e., the density of the set of targets at time  $k$  given all the observations received up to time  $k$ , is generally intractable. In the literature, there are two main filter types that implement the Bayes recursion approximately for the multitarget set density. The first is based on moment approximations of the posterior distributions, such as the Probability Hypothesis Density filter (PHD) [4, 5] and the Cardinalised PHD (CPHD) filter [6, 7], which avoid handling the data association uncertainty explicitly via approximation. The second is based on parameterised density representations. Filters of this type are based on conjugate priors, which can provide accurate approximations to the exact posterior distributions.

In the context of MTT, conjugacy means that *if we begin with a multitarget density of a conjugate prior form, then all subsequent predicted and updated multitarget densities will also be of the conjugate prior form* [8]. Different from the concept of conjugate prior defined in the exponential family, the data association problem in MTT causes the parameters needed to represent the posterior density increase very fast over time, thus harming the tractability of conjugacy. Despite this, different approximation techniques can be employed to reduce the number of parameters. The advantage of using conjugate prior in MTT algorithms is that the true posterior can be approximated arbitrary well, as long as a sufficient number of parameters is used. Filters based on parameterised density representations have been developed for labelled and unlabelled multi-Bernoulli mixture (MBM) conjugate priors respectively. A labelled RFS can be considered as an RFS, where each variable is augmented with a unique label.

Filters based on labelled multi-Bernoulli (MB) representation include the Generalised Labelled MB filter (GLMB) [8, 9, 10] and its approximation the Labelled MB (LMB) filter [11], whereas filters based on unlabelled MB representation include the Poisson multi-Bernoulli Mixture (PMBM) filter [12, 13] and its approximation

the Poisson multi-Bernoulli (PMB) filter [14]. Compared to the PMBM and PMB filters, the GLMB and LMB filters are able to produce the estimation of targets trajectories. These algorithms have shown superior performance over algorithms based on moment approximations, but it is not yet clear how they compare to each other regarding computational time and filtering performance. One contribution of this thesis is the performance evaluation of MB conjugate priors for multiple point target filtering using the Generalised Optimal Sub-pattern Assignment (GOSPA) metric [15]. In this part of work, we assume that each target gives rise to at most one measurement and that targets do not have any shared measurements. Simulation results demonstrate that the PMB filters arguably provide the best overall performance concerning GOSPA error and computational time.

Due to the development of high-resolution radar and lidar sensors equipped on autonomous vehicles, it is common that a target may occupy several sensor resolution cells. The tracking of such a target leads to the so-called extended target tracking problem, and the objective is to recursively determine the shape and kinematic parameters of the target over time. In extended target tracking, a non-standard measurement model is needed to model the number and the spatial distribution of generated measurements for each target. A common choice for modelling the number of measurements is the inhomogeneous Poisson Point Process (PPP), proposed in [16]. As for the modelling of the spatial distribution, two popular models are the Random Hypersurface Models [17] and the Gaussian inverse Wishart (GIW) approach [18, 19]. The former is designed for general star-convex shape; the latter relies on the elliptic shape that it models the spatial distribution of target-generated measurements as Gaussian with unknown mean and covariance. The Gamma GIW (GGIW) model [20, 21] is an extension of the GIW model by incorporating the estimation of target measurement rates. Many of the existing multiple extended target tracking algorithms are based on RFS with the distribution of target-generated measurements modelled as GIW or GGIW, e.g., the PHD filter [22, 23], the CPHD filter [21], the GLMB filter and the LMB filter [24]; as well as the PMBM filter [25, 26]. It was shown in [26] that the extended target PMBM filter has improved performance compared to the extended target GLMB and LMB filters.

Since each extended target can generate multiple measurements per time scan and it is unknown which group of measurements are generated by the same object (target or clutter), the problem of data association is very challenging in multiple extended target tracking. In previous work [20, 25, 26], the data association can be split into two parts: first, clustering methods are used to find multiple partitions of the set of measurements; second, assignment methods are used to assign each measurement cell to a target or a clutter source. In the recently presented work [27], a sampling based method called stochastic optimisation (SO) is proposed, which directly maximises the multitarget likelihood function and solves the data association problem in a single step. It has been demonstrated in [27] that the SO method has improved performance in comparison to methods which involve clustering and assignment. In this thesis, an extension of the work [27] is presented, in which the benefits of solving the data association problem using sampling methods are highlighted and

more sampling algorithms are investigated. Specifically, we discuss different aspects of four different sampling methods, namely, the Gibbs sampling, the merge/split Metropolis Hastings (MH), the combined Gibbs and merge/split MH and the SO, and compare their performance after integrating them into the PMBM filter in a simulation study.

For the extended PMBM filter, several approximation techniques can be used to reduce the number of parameters in the posterior density representation. First, gating [3] can be used to separate targets and measurements into independent subgroups. Second, for each gating group, the multiple updated MB components obtained after solving the data association problem are pruned to only contain the components whose cumulative sum of the relative weights exceeds a threshold. Third, the remaining MB components with equal number of Bernoulli components and small symmetrised KL divergence can be merged [26]. In order to further reduce the computational complexity of the PMBM filter, it would be interesting to investigate the possibility of approximating the MBM as a single MB. A variational method to find the best-fitting MB that minimises the set Kullback-Leibler divergence from the MBM was presented in [14] for point target tracking. The proposed PMB filter is a computationally efficient approximation of the PMBM filter, but it is not yet clear how the variational method of [14] can be used on extended target tracking. In this thesis, a PMB filter for extended target estimation is presented, along with its GGIW implementation. We have discussed two different implementations of the variational MB algorithm to approximate the MBM representing pre-existing tracks as a single MB. In addition, a method to create new tracks in a reasonable manner is proposed. The performance of the PMB filter is compared to the PMBM filter and the PHD filter in a simulation study.

## 1.3 Contributions of the Appended Papers

### **Paper I: Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering**

In this paper, we evaluate the performance of labelled and unlabelled MB conjugate priors for multi-target filtering. We analyse various aspects of six different filters: 1)  $\delta$ -GLMB with separate prediction and update steps, 2)  $\delta$ -GLMB with joint prediction and update steps, 3) LMB, 4) PMBM, 5) PMB using Murty's algorithm, and 6) PMB using loopy belief propagation. Two different scenarios are considered: 1) targets are well-spaced and 2) targets are in close proximity. The benefit of recycling for the PMBM filter is also studied. Filtering performance is assessed using the GOSPA metric.

I carried out the comparison study and wrote the paper under the supervision of the co-authors.

## **Paper II: Likelihood-based data association for extended object tracking using sampling methods**

In this paper, we show that it is possible to handle the data association in a single step that maximises the desired likelihood function. For single step data association, we use algorithms based on stochastic sampling, and integrate them into the PMBM filter. Four sampling algorithms are compared in a simulation study: 1) Gibbs sampling, 2) merge/split MH, 3) the combined Gibbs and merge/split MH and 4) SO. The best one is compared to clustering and assignment in an experiment, where Velodyne data is used to track pedestrians.

I carried out the simulation study comparing four different sampling methods and participated in the discussion with the co-authors.

## **Paper III: Poisson multi-Bernoulli filter for extended object estimation**

In this paper, we present a PMB filter for multiple extended targets estimation. Different methods to merge the MBM are presented, along with their GGIW implementation. The performance of the PMB filter is compared to the PMBM filter and the PHD filter in three simulated scenarios: 1) high target number and clutter density, 2) dense birth and 3) parallel manoeuvre.

I carried out the implementations and wrote the paper under the supervision of the co-authors.

## **1.4 Conclusions and Future Work**

### **Paper I: Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering**

PMBM is a more efficient filter structure than  $\delta$ -GLMB. The Poisson birth model provides the unlabelled filter with the ability to quickly detect several targets that appear at the same time at very short distance from each other, and it enables the use of recycling methods to improve the estimation performance of unlabelled filters further. In the scenario with coalescence, unlabelled filters have much better performance than labelled filters, and the PMB filter presents the best overall performance.

MTT algorithms based on conjugate priors are developing very fast in recent years. New implementations and new algorithms continue to be developed. It would be interesting to compare the performance of labelled and unlabelled filters with more variants. Another future work direction is to evaluate the performance of unlabelled MB filters using a MB birth model.

## **Paper II: Likelihood-based data association for extended object tracking using sampling methods**

The data association problem in extended object tracking can be solved in a single, likelihood-based, step using sampling algorithms, which are shown to outperform previous work that has been based on the combination of clustering algorithms and optimal assignment algorithms. The simulation study shows that both the combined Gibbs and merge/split MH and the SO are good sampling algorithms to solve the data association problem.

An important topic for future work is investigating how the sampling algorithms can be implemented as efficiently as possible, for example, by utilising parallelisation, by storing temporary variables and by testing different algorithms to split a measurement cell into two parts.

## **Paper III: Poisson multi-Bernoulli filter for extended object estimation**

The proposed PMB filter is an efficient approximation of the PMBM filter for extended target tracking. But the simulation results show that the PMB filter can produce distorted target extent estimation especially when targets move in close proximity.

A future work direction is to investigate how to improve the estimation of target extent in the PMB filter, for example, by adjusting the weights of different cost functions in the E-step. Since the approximation of the MBM representing pre-existing tracks and new tracks are considered separately in this paper, it would also be interesting to investigate how to integrate the formation of new tracks in the variational MB algorithm.

# Paper I: Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering

## Abstract

In this paper, we evaluate the performance of labelled and unlabelled multi-Bernoulli conjugate priors for multi-target filtering. Filters are compared in two different scenarios with performance assessed using the generalised optimal sub-pattern assignment (GOSPA) metric. The first scenario under consideration is tracking of well-spaced targets. The second scenario is more challenging and considers targets in close proximity, for which filters may suffer from coalescence. We analyse various aspects of the filters in these two scenarios. Though all filters have pros and cons, the Poisson multi-Bernoulli filters arguably provide the best overall performance concerning GOSPA and computational time.

## 2.1 Introduction

Multiple target tracking (MTT) involves the processing of sets of measurements obtained from multiple targets in order to estimate their current states. Solving the MTT problem is mainly complicated by the unknown correspondence between targets and measurements, known as data association. Popular solutions to MTT are the joint probabilistic data association (JPDA) filter [1], multiple hypothesis tracker (MHT) [2] and algorithms based on random finite sets (RFS) [3]. Developments using RFS have yielded a variety of tracking methods that avoid handling the data association uncertainty explicitly, such as probability hypothesis density (PHD) filter [4, 5], cardinalised PHD (CPHD) filter [6, 7], multiple target multi-Bernoulli (MeMber) filter [3], and cardinality-balanced MeMber (CB-MeMber) filter [28].

A significant trend in RFS-based MTT is the development of conjugate distributions in Bayesian probability theory, which means that the posterior distribution has the same functional form as the prior. MTT algorithms based on conjugate priors for labelled RFS [8, 9, 10, 11] and unlabelled RFS [12, 13, 14] can provide accurate approximations to the exact posterior distributions, and appealing performance for MTT has been demonstrated. Filters based on labelled multi-Bernoulli (MB) conju-

---

gate prior using MB birth model include Delta generalised labelled multi-Bernoulli ( $\delta$ -GLMB) [8, 9] and its approximation labelled multi-Bernoulli (LMB) [11]. Filters based on unlabelled MB conjugate prior using Poisson birth model include Poisson multi-Bernoulli mixture (PMBM) [12, 13] and its variational approximation Poisson multi-Bernoulli (PMB) [14].

These algorithms are all important contributions to the MTT literature, but it is not yet clear how they compare to each other regarding computational time and performance. The main contribution of this paper is the performance evaluation of MB conjugate priors for multi-target filtering using the generalised optimal sub-pattern assignment (GOSPA) metric [15]. Compared with the unnormalised OSPA metric [29], the GOSPA metric allows for further breaking down the cardinality mismatch into errors due to missed and false targets. In this work, we assume that each target gives rise to at most one measurement and that targets do not have any shared measurements. Simulation results demonstrate that PMBM is a more computationally efficient filter structure than  $\delta$ -GLMB since it tends to yield better estimation performance than  $\delta$ -GLMB with less computational time in a given test scenario. The advantage of a Poisson birth model over an MB birth model in MTT is also demonstrated in specific examples. In addition, the unlabelled RFS filter shows the ability to resolve the coalescence phenomenon, and its performance under this scenario can be further improved by applying the variational approximation method [14].

The paper is organised as follows. Background on RFS and Bayesian filtering is provided in Section II. Section III and IV summarise the labelled and unlabelled MB conjugate priors with their corresponding filter implementations. Section V discusses different methods that can be used to reduce the computational cost. Simulation results are presented in Section VI, and conclusions are drawn in Section VII.

## 2.2 Background

In RFS-based methods [3], target states and observations are represented in the form of finite sets. The system state at time  $k$  is modelled as a set  $X_k$ , and the set of measurements obtained at time  $k$  is denoted as  $Z_k$ , including clutter and target-generated measurements with unknown origin. The sequence of all the measurement sets received so far up to time  $k$  is denoted as  $Z^k$ . More notations are given in Table 2.1.

A labelled RFS can be formed from its unlabelled version by incorporating labels into target states such that each state  $x \in \mathbb{X}$  is augmented with a unique discrete label  $l \in \mathbb{L}$ . Let  $\mathcal{L}(\mathbf{X})$  be the set of unique labels in  $\mathbf{X}$ , then a finite subset  $\mathbf{X}$  on space  $\mathbb{X} \times \mathbb{L}$  has distinct labels if and only if  $\delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|) = 1$  [9].

A non-homogeneous Poisson point process (PPP) with intensity function  $\lambda(x)$  has

---

**Table 2.1:** Notations

---

- Single target states are represented by lower-case letter, e.g.,  $x$ ; multi-target states are represented by upper-case letters, e.g.,  $X$ ; labelled states are represented by bold letters, e.g.,  $\mathbf{x}$ ,  $\mathbf{X}$ ; spaces are represented by blackboard bold letters, e.g.,  $\mathbb{X}$ .
- $|X|$ : set cardinality, i.e., number of elements in set  $X$ .
- Kronecker delta function

$$\delta_Y(X) = \begin{cases} 1 & \text{if } Y = X, \\ 0 & \text{otherwise.} \end{cases}$$

- $h^X = \prod_{x \in X} h(x)$ .
- $\Pi_N$ : set of permutation functions on  $I_N \triangleq \{1, \dots, N\}$

$$\Pi_N = \{\pi : I_N \rightarrow I_N | i \neq j \Rightarrow \pi(i) \neq \pi(j)\}.$$

- $\uplus$ : disjoint set union, i.e.,  $Y \uplus U = X$  means that  $Y \cup U = X$  and  $Y \cap U = \emptyset$ .
  - $I_m$ : identity matrix of size  $m \times m$ .
- 

RFS density

$$f^{ppp}(X) = \exp\left(-\int \lambda(x)dx\right) \cdot \lambda^X, \quad (2.1)$$

where  $|X|$  is Poisson distributed, and elements  $x \in X$  are independently and identically distributed (i.i.d.). A Poisson process is often used to model clutter and target birth in unlabelled RFS filters.

A Bernoulli process with probability of existence  $r$  and existence-conditioned probability density function (PDF)  $f(x)$  has RFS density

$$f(X) = \begin{cases} 1 - r & X = \emptyset \\ r \cdot f(x) & X = \{x\} \\ 0 & \text{otherwise,} \end{cases} \quad (2.2)$$

where  $|X|$  is Bernoulli distributed with parameter  $r$ . A labelled Bernoulli RFS  $\mathbf{X}$  is a Bernoulli RFS  $X$  augmented with label  $l$  corresponding to the non-empty Bernoulli component  $x$ , i.e.,  $\mathbf{X} = \{(x, l)\}$  (or  $\mathbf{X} = \emptyset$ ). A Bernoulli process can capture the target uncertainty of both existence and state, and it is also used in labelled RFS filters to model target birth.

Multiple targets can be naturally represented through an MB RFS. An MB RFS  $X$  is a union of independent Bernoulli RFSs  $X_i$ , i.e.,  $X = \bigcup_{i=1}^N X_i$ ,

$$f^{mb}(X) = \sum_{\uplus_i^N X_i = X} \prod_{i=1}^N f_i(X_i). \quad (2.3)$$

An MB RFS can either be labelled or unlabelled.

---

## 2.2.1 Bayesian Filtering Recursion

An RFS density can be expressed using a cardinality distribution  $c(n)$ , and joint conditional state distributions  $f_n(x_1, \dots, x_n|n)$ , yielding [12]

$$f(\{x_1, \dots, x_n\}) = c(n) \sum_{\pi \in \Pi_n} f_n(x_{\pi(1)}, \dots, x_{\pi(n)}|n), \quad (2.4)$$

where  $\pi(i)$  is the  $i$ -th component of vector  $\pi$ . A labelled RFS and its unlabelled version have the same cardinality distribution [9].

The multi-target distribution at time  $k$ , given all measurement sets up to and including time  $k'$  is denoted as  $f_{k|k'}(X_k|Z^{k'})$ , and  $f_k(Z_k|X_k)$  is the multi-target measurement likelihood. The multi-target Bayes filter propagates the target set PDF  $f_{k|k-1}(X_k|Z^{k-1})$  in time using the Bayes update [3, p. 484]

$$f_{k|k}(X_k|Z^k) \propto f_k(Z_k|X_k) f_{k|k-1}(X_k|Z^{k-1}), \quad (2.5)$$

and the Chapman-Kolmogorov prediction [3, p. 484]

$$f_{k+1|k}(X_{k+1}|Z^k) = \int f_{k+1|k}(X_{k+1}|X_k) f_{k|k}(X_k|Z^k) \delta X_k, \quad (2.6)$$

where the set integral is defined as [3, p. 361]:

$$\int f(X) \delta X \triangleq f(\emptyset) + \sum_{n=1}^{\infty} \frac{1}{n!} \int f(\{x_1, \dots, x_n\}) d(x_1, \dots, x_n). \quad (2.7)$$

## 2.2.2 Standard Point Target Transition and Measurement Model

The dynamic point target transition model utilised to solve the RFS-based MTT problem is based on the following assumptions

- *Targets arrive according to a PPP or an MB process, independently of existing targets.*
- *At each time step, targets remain with a probability of survival  $P^s(x)$ . Targets depart according to i.i.d. Markov processes with probability  $1 - P^s(x)$ .*
- *Targets' next states only depend on their current states. Target motion follows i.i.d. Markov processes with transition density  $f_{k|k-1}(x_k|x_{k-1})$ .*

The multi-target transition kernel can be written as the convolution of transition density for survival targets  $f_S^{mb}(\cdot|X)$  and transition density for new born targets  $f_B(\cdot)$

$$f(X_{k+1}|X_k) = \sum_{S_k \uplus B_k = X_{k+1}} f_S^{mb}(S_k|X_k) f_B(B_k), \quad (2.8)$$

where  $S_k$  is the set of survival targets, and  $B_k$  is the set of new born targets.

In the measurement model, the following assumptions are made.

- *Each target-generated measurement is only conditioned on its corresponding target. The single target measurement likelihood is  $f_k(z_k|x_k)$ .*

- At each time step, existing targets may or may not be detected. The detection probability is  $P^d(x)$ .
- The sensor may receive measurements not originating from any target, known as false alarm or clutter. At each time step, the clutter arrives according to a PPP with intensity  $\lambda^a(z)$ , independently of targets and target-generated measurements.

The multi-target likelihood can be presented in a convolution form

$$f_k(Z_k|X_k) = \sum_{T_k \uplus C_k = Z_k} f_k^{mb}(T_k|X_k) f_k^{PPP}(C_k), \quad (2.9)$$

where  $T_k$  is the set of target-generated measurements, and  $C_k$  is the set of clutter measurements.

## 2.3 Labelled Multi-Bernoulli Conjugate Prior

This section presents a general description of  $\delta$ -GLMB and its efficient approximation LMB. The reader is referred to [8, 9, 10, 11] for detailed analytic derivations and mathematical expressions.

### 2.3.1 Delta Generalised Labelled Multi-Bernoulli

A  $\delta$ -GLMB RFS is a labelled RFS with state space  $\mathbb{X}$  and label space  $\mathbb{L}$  distributed according to [9]

$$f(\mathbf{X}) = \delta_{|\mathbf{X}|}(|\mathcal{L}(\mathbf{X})|) \sum_{(I, \xi) \in \mathcal{F}(\mathbb{L}) \times \Xi} w^{(I, \xi)} \delta_I(\mathcal{L}(\mathbf{X})) (f^{(\xi)})^{\mathbf{X}}, \quad (2.10)$$

where  $I \in \mathcal{F}(\mathbb{L})$  is the set of target labels, and  $\xi \in \Xi$  represents a history of association maps. Each pair  $(I, \xi)$  is a hypothesis with probability  $w^{(I, \xi)}$ , which satisfies

$$\sum_{(I, \xi) \in \mathcal{F}(\mathbb{L}) \times \Xi} w^{(I, \xi)} = 1. \quad (2.11)$$

If the prior distribution is a  $\delta$ -GLMB of the form (2.10), then under the multi-target likelihood function (2.9), the posterior and predicted distribution are both  $\delta$ -GLMB [9]. According to this conjugacy, the  $\delta$ -GLMB filter can recursively propagate a  $\delta$ -GLMB RFS density in time via the Bayes update and prediction equation (2.5) and (2.6).

A labelled MB model is used for target birth, and we denote the label space for new born targets as  $\mathbb{B}$ . The likelihood for a birth hypothesis yields [8]

$$w_B(L) = [1 - r(\cdot)]^{(\mathbb{B}-L)} [r(\cdot)]^L, \quad (2.12)$$

where  $r(l)$  denotes the existence probability of a target with label  $l \in I$ . For the surviving targets, their labels are kept from last time step, while the labels of birth targets are newly assigned. In the  $\delta$ -GLMB prediction step, each component in the prior generates a new set of predicted components. In the  $\delta$ -GLMB update step, each

---

component in the predicted density generates a set of updated components for each possible measurement-to-label mapping. Instead of computing the filtering density in two separate steps, a fast implementation has been proposed in [10], in which the prediction and update steps are combined into a single step by formulating a new mapping between the components of the current and previous filtering density. Compared with the original mapping used in update step, the associations for non-survival and unconfirmed birth targets are included in this new mapping.

### 2.3.2 Labelled Multi-Bernoulli

The LMB filter [11] is based on approximating the posterior and predicted densities (2.5) and (2.6) as labelled MB RFS densities. In contrast to  $\delta$ -GLMB, the number of components maintained in the recursive Bayesian filtering only grows linearly with time.

The multi-target prediction in the LMB filter is identical to the prediction for MeM-Ber [28] with target labels interpreted as component indices. The predicted existence probability and density distribution are re-weighted by the survival probability and transition density. In the update step, the posterior multi-target density is approximated by exact marginalisation; thus the LMB filter propagates only one component. The predicted LMB is first represented as  $\delta$ -GLMB, then a full  $\delta$ -GLMB update is applied directly before the posterior collapsing back to a matching LMB approximation.

### 2.3.3 Poisson Multi-Bernoulli Mixture

The PMBM conjugate prior for point target MTT was developed in [12]. It is a linear combination of independent PPP and multi-Bernoulli mixture (MBM) components with the following form

$$f(X) = \sum_{Y \uplus W = X} f^{ppp}(Y) f^{mbm}(W), \quad (2.13)$$

where  $Y$  stands for targets that have not yet been detected, and  $W$  stands for targets that have been detected at least once. For targets that have already been detected, their distributions can be described as an MBM of the form [12]:

$$f^{mbm}(X) = \sum_{a=\{h_1, \dots, h_N\} \in \mathcal{A}} w_a \sum_{\uplus_i^N X_i = X} \prod_{i=1}^N f_{h_i}(X_i), \quad (2.14)$$

where each of the MB components  $a \in \mathcal{A}$  corresponds to a particular data association with weight  $w_a$  satisfying

$$\sum_{a=\{h_1, \dots, h_N\} \in \mathcal{A}} w_a = 1. \quad (2.15)$$

Each data association hypothesis is made up of single target hypotheses  $\{h_1, \dots, h_N\}$  on each target [12]. One single target hypothesis can incorporate events, including that the target never existed, that the target existed before and that the target

continues to exist, represented via a Bernoulli process. Missed detections may occur on some proportions of targets that are hypothesised to be born. A target that is hypothesised to exist but has never been detected is treated as an unknown target [12]. The distribution of unknown targets is represented by a PPP.

In the prediction step, the MB describing pre-existing tracks and the PPP describing unknown targets are predicted individually. By having a Poisson birth model, PPP for new born targets can be easily incorporated into the predicted PPP. In the update step, PPP and MBM are updated independently. Two single target hypotheses are created for each measurement, and then the PPP intensity is updated by the miss detection probability. The first single target hypothesis covers the case that the measurement is associated with a previous target so that this hypothesis has zero existence probability and the corresponding PDF of the Bernoulli distribution has no effect. The second single target hypothesis covers the case that the measurement corresponds to a false alarm or a new target. For targets surviving from previous time steps, new single target hypotheses are included from missed detections or updates of previous hypotheses using one of the new measurements.

### 2.3.4 Poisson Multi-Bernoulli

A PMB is a union of a PPP describing unknown targets and an MB process describing already detected targets. In the PMB recursion, the PMB density is preserved in prediction step, whereas the MB component becomes an MBM due to data association. A variational approximation method was presented in [14] to obtain the best-fitting MB  $g(X)$  that minimises the Kullback-Leibler (KL) divergence from the MBM distribution  $f(X)$ . In [14], it is shown that this optimisation problem can be solved approximately as:

$$\min_{q(h,j) \in \mathcal{M}} - \sum_{j=1}^N \int \left( \sum_{h \in \mathcal{H}} q(h,j) f_h(X) \right) \cdot \log g_j(X) \delta X, \quad (2.16)$$

where  $q(h,j)$  is the probability that a Bernoulli component in  $f(X)$  that is utilising hypothesis  $h$  is assigned to the  $j$ -th Bernoulli component of  $g(X)$ , and  $\mathcal{M}$  is an approximated polytope needed for tractability

$$\mathcal{M} = \left\{ q(h,j) \geq 0 \left| \sum_{h \in \mathcal{H}} q(h,j) = p_j \forall j \in \{1, \dots, N\}, \sum_{j=1}^N q(h,j) = p_h \forall h \in \mathcal{H} \right. \right\}. \quad (2.17)$$

The algorithm is initialized with

$$p_j(h) = \sum_{a=(h_1, \dots, h_N) \in \mathcal{A} | h_j=h} w_a, \quad (2.18)$$

which can either be calculated using the best data association hypotheses and their corresponding weights obtained from Murty's algorithm [30] or be approximated based on loopy belief propagation (LBP) [31].

---

## 2.4 Approximations for Computational Tractability

In previous sections, we have summarised the  $\delta$ -GLMB, LMB, PMBM and PMB. The common bottleneck for all the filters is the large number of hypotheses generated in the update step. Since it is not tractable to compute all the possible components, efficient truncation methods should be implemented by only maintaining components with most significant weights<sup>1</sup>.

If hypothesis weights are generated in non-decreasing order, the  $M$ -best mappings (i.e. highest weights) can be selected without computing the weights of all possible mappings exhaustively. The  $M$ -best mappings can be found, e.g., by solving the ranked assignment problem using Murty's algorithm. Implementation details, e.g., how to construct the cost function for an association map, can be found in [9], [10] and [13]. In the implementation of  $\delta$ -GLMB with separate prediction and update steps, additional approximation is needed to truncate the independent surviving and birth hypotheses separately. The hypotheses with highest weights can be determined by solving the K-shortest path problem [32].

To further reduce the computational complexity of data association, gating can be implemented in all the proposed filters to remove unlikely target-to-measurement pairs before the update step. Moreover, clustering can be used in LMB, PMBM and PMB to partition well-spaced targets and the corresponding measurements falling into their gates into independent groups, which allows for parallel computing.

For unlabelled RFS filters, in each MB component, the recycling method of [33] can be applied to Bernoulli components with low existence probability. The recycled components are approximated as being Poisson and are incorporated into the PPP representing unknown targets for generating possible new targets in subsequent steps. Thus, the number of single target hypotheses in each data association hypothesis is reduced while maintaining the significant information of already detected targets. In addition, recycling can help recover performance loss due to the pruning of Bernoulli components with small existence probability [31]. However, having a PPP for each MB in PMBM would result in multiple PPP updates, which may instead increase the computational cost. To address this problem, we first find the best-fitting MB of the MBM to be recycled, and then the approximated MB is further approximated as being PPP and recycled.

## 2.5 Simulations and Results

In this section we show simulation results that compare six different filters: 1)  $\delta$ -GLMB with separate prediction and update steps [8, 9], 2)  $\delta$ -GLMB with joint prediction and update steps [10], 3) LMB [11], 4) PMBM [13], 5) PMB [14] using

---

<sup>1</sup>For PMB using LBP, marginal association probabilities are used to merge hypotheses.

---

Murty’s algorithm, and 6) PMB using LBP [31]. All the codes are written in MATLAB, except Murty’s algorithm, which is written in C++<sup>2</sup>. The benefit of recycling for PMBM is also studied. The estimation performances of the filters are evaluated in two different scenarios. In the first scenario, targets are well-separated most of the time. The second scenario is more complicated, involving targets that get in close proximity first and then separate.

### 2.5.1 State Space Models

The kinematic target state is a vector of position and velocity  $x_k = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$ . A single measurement is a vector of position  $z_k = [z_{x,k}, z_{y,k}]^T$ . Targets follow a linear-Gaussian constant velocity model,  $f_{k|k-1}(x_k|x_{k-1}) = \mathcal{N}(x_k; F_k x_{k-1}, Q_k)$  with parameters

$$F_k = I_2 \otimes \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}, \quad Q_k = \sigma_v^2 I_2 \otimes \begin{bmatrix} T^4/4 & T^3/2 \\ T^3/2 & T^2 \end{bmatrix}$$

where  $T = 1s$  is the sampling period, and  $\sigma_v$  is the standard deviation of motion noise. The linear-Gaussian measurement likelihood model has density  $g(z_k|x_k) = \mathcal{N}(z_k; H_k x_k, R_k)$  with parameters  $H_k = I_2 \otimes [1 \ 0]$  and  $R_k = \sigma_\epsilon^2 I_2$ , where  $\sigma_\epsilon$  is the standard deviation of measurement noise. The target survival probability  $P_k^s$  and the detection probability  $P_k^d$  are assumed to be constant throughout the simulation. The clutter follows a Poisson RFS with uniform density, giving an average of  $\lambda^{fa}$  per time step. We consider cases with  $P_k^s = 0.99$ ,  $P_k^d \in \{0.75, 0.98\}$ , and  $\lambda^{fa} \in \{10, 30\}$ .

For unlabelled filters, the Poisson birth intensity is a Gaussian mixture  $\lambda^b(x) = \sum_{n=1}^{n_b} \lambda^{b,n} \mathcal{N}(x; \mu^{b,n}, \Sigma^{b,n})$  with all the Gaussian components sharing the same weight  $\lambda^{b,n} = 0.1$  and covariance. The PPP intensity of unknown targets is represented as  $\lambda^u(x) = \sum_{n=1}^{n_u} \lambda^{u,n} \mathcal{N}(x; \mu^{u,n}, \Sigma^{u,n})$ , with the initial value assumed to be  $\lambda_{|0}^u(\mathbf{x}) = \lambda^b(\mathbf{x})$ . For labelled filters, the birth density is a labelled MB RFS with all the Bernoulli components sharing the same existence probability  $r^{b,n} = 0.1$ . The existence-conditioned PDF of each Bernoulli is Gaussian distributed. A labelled Poisson RFS can also be used to model target birth in  $\delta$ -GLMB [9], but this is not a suitable option [13]. As a Poisson RFS has a cardinality distribution that expands from zero to infinity, it is not efficient to consider each possible birth hypothesis separately via a  $\delta$ -GLMB filter. Parameter settings for different filters are stated in Table 2.2.

### 2.5.2 Performance Evaluation

Given a multi-target posterior density, several state estimators are available [9]. In this work, we choose to extract the target states by finding the maximum a posteriori (MAP) cardinality estimate, following the methods suggested in [9] and [13].

Filtering performance is assessed using GOSPA metric. It was mentioned in [15] that the GOSPA metric with parameter  $\alpha = 2$  can be rewritten using the 2D assignment

---

<sup>2</sup>For the implementations of  $\delta$ -GLMB with separate prediction and update steps and LMB, we use the codes that Profs Ba-Ngu Vo and Ba-Tuong Vo share online. The authors thank them for providing the codes.

---

**Table 2.2:** Filter Parameters

---

- Gating size in percentage is 0.999.
  - $\delta$ -GLMB: capping threshold on the number of hypotheses is 3000; pruning threshold on the weight of hypothesis is  $10^{-5}$ ; requested number of components used in K-shortest path and Murty’s algorithm is  $\lceil 3000 \cdot w \rceil$ , where  $w$  is the weight of hypothesis; requested number of birth components used in K-shortest path algorithm is 10.
  - LMB: capping threshold on the number of targets is 100; pruning threshold is  $10^{-4}$ ; requested number of birth and survival components used in K-shortest path algorithm are 10 and 30 respectively; requested number of components used in Murty’s algorithm is  $\lceil 100 \cdot w \rceil$  in the first testing scenario and  $\lceil 300 \cdot w \rceil$  in the second; the number of Gaussian components in each Gaussian mixture is limited to 10 and components are merged with Mahalanobis distance smaller than 1.
  - PMBM: capping threshold on the number of hypotheses is 100; pruning threshold on the weight of hypotheses and weight of Gaussian components in PPP is  $10^{-4}$ ; requested number of components used in Murty’s algorithm is  $\lceil 100 \cdot w \rceil$ ; pruning threshold on target existence probability is  $10^{-4}$  without recycling, and  $10^{-1}$  with recycling.
  - PMB: corresponding pruning threshold is set to the same as PMBM; requested number of components used in Murty’s algorithm is 20; convergence threshold used in LBP is  $10^{-6}$ .
- 

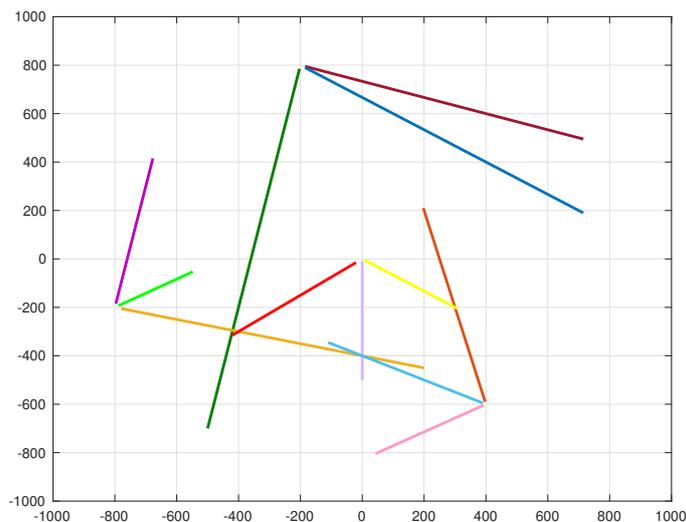
set as

$$d_p^{(c,2)}(X, Y) \triangleq \left[ \min_{\gamma \in \Gamma(|X|, |Y|)} \left( \sum_{(i,j) \in \gamma} d(x_i, y_j)^p + \frac{c^p}{2} (|X| + |Y| - 2|\gamma|) \right) \right]^{\frac{1}{p}}, \quad (2.19)$$

where  $\Gamma(|X|, |Y|)$  is the set of all possible 2D assignment sets,  $c$  denotes the cut-off value at base distance and  $p$  determines the severity of penalising the outliers in the localisation component. GOSPA allows for decomposing the multi-target metric into localisation error  $\sum_{(i,j) \in \gamma} d(x_i, y_j)^p$ , missed detection error  $c^p(|X| - |\gamma|)/2$ , and false detection error  $c^p(|Y| - |\gamma|)/2$ , considering  $X$  as the ground truth and  $Y$  as the estimates. In the simulations, we choose  $p = 1, c = 100$ , and results are shown over 200 Monte Carlo trials.

### 2.5.3 Scenario with Well-spaced Targets

In this scenario, targets are born from four localised positions with ground truth shown in figure 2.1. The standard deviations of motion and measurement noises are  $\sigma_v = 5m/s^2$  and  $\sigma_\epsilon = 10m$ , respectively. In the birth intensity, each Gaussian component has the same covariance  $\Sigma^{b,n} = 100 \times I_4$  and the mean values are  $\mu^{b,1} = [0, 0, 0, 0]^T$ ,  $\mu^{b,2} = [400, -600, 0, 0]^T$ ,  $\mu^{b,3} = [-200, 800, 0, 0]^T$  and  $\mu^{b,4} = [-800, -200, 0, 0]^T$ . The initial target positions are randomly sampled from these Gaussian densities.

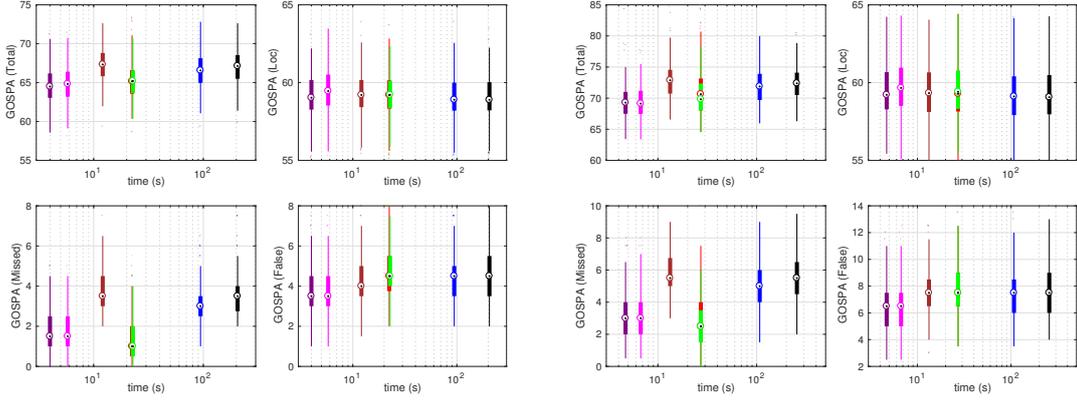


**Figure 2.1:** Target trajectories on the region  $[-1000, 1000]m \times [-1000, 1000]m$ . Up to twelve targets travel along straight lines with constant speed in the duration time of 100s with three targets born at time 1 and 20, two targets born at time 40, 60 and 80 and two targets dead at time 50 and 70. Specifically, the three targets born at time 20 are from roughly the same position, and the two targets born at time 40 are from roughly the same position.

Figure 2.2 presents the GOSPA errors and the average computation time required for a complete Monte Carlo simulation. The GOSPA errors were averaged over time, and are shown as box plots. In cases with high detection probability, the median of the GOSPA error due to localisation errors is very similar for all filters, while in cases of low detection probability, filters using variational approximations have slightly higher median value compared with the rest. It is clearly visible that PMBM exhibits a similar or better estimation performance than  $\delta$ -GLMB while saving computational time notably. Errors due to missed and false targets can be reduced by increasing the requested number of components used in K-shortest path or Murty’s algorithm, but this would further increase the computational burden.

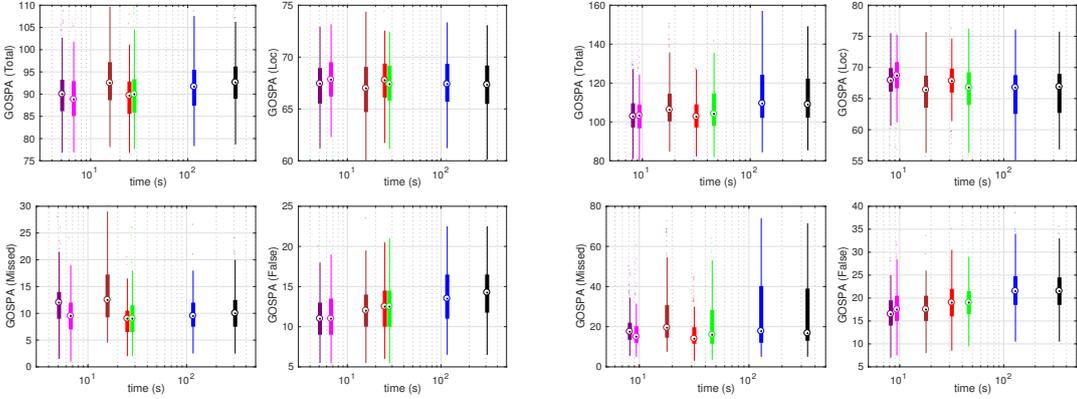
The benefit of recycling is most clear when the probability of detection is low. In figure 2.2 we see that PMBM with recycling requires less computation time, and has lower GOSPA error, compared to PMBM without recycling. As for the two variants of PMB, LBP is a faster implementation than Murty’s algorithm, and it tends to yield fewer false targets, but more missed targets.

Figure 2.3 shows the mean GOSPA error cost for missed targets against time. Here, we can see that labelled RFS filters experience higher peaks than unlabelled filters around initialization period, and at time 20 and 40, when there are two or more new targets born from roughly the same position.



(a)  $\lambda^{fa} = 10, P_k^d = 0.98$

(b)  $\lambda^{fa} = 30, P_k^d = 0.98$



(c)  $\lambda^{fa} = 10, P_k^d = 0.75$

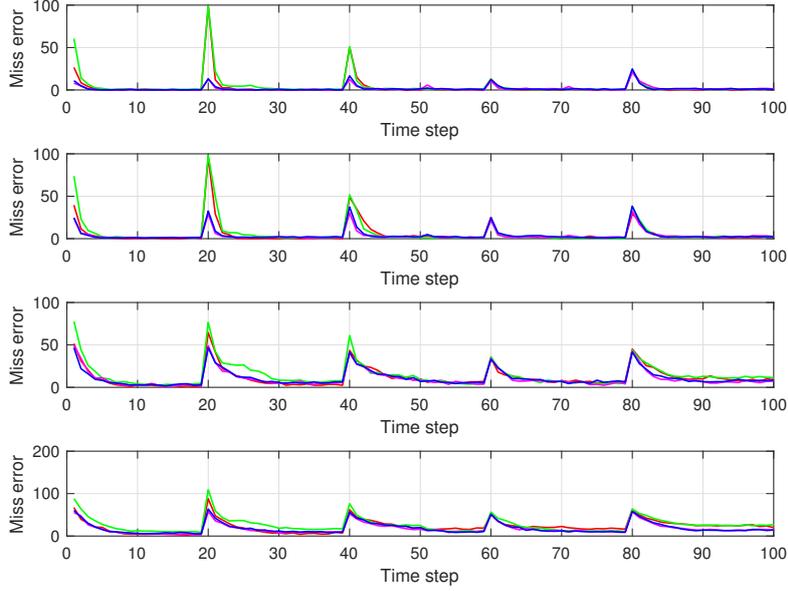
(d)  $\lambda^{fa} = 30, P_k^d = 0.75$

**Figure 2.2:** Filtering performance evaluation of PMB (LBP) with recycling (purple), PMB (Murty) with recycling (magenta), LMB (brown), PMBM with recycling (red), PMBM without recycling (green),  $\delta$ -GLMB with joint prediction and update (blue) and  $\delta$ -GLMB (black) using GOSPA in a scenario where targets are well-spaced. Results are presented in box plots. On each box, ‘ $\odot$ ’ indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The end of ‘Whisker’ corresponds to approximately  $\pm 2.7$  units of standard deviation.

## 2.5.4 Scenario with Coalescence

In the following simulation, filters are evaluated in a more challenging scenario involving six targets which are in close proximity at the midpoint of the simulation, achieved by initializing at the midpoint and running forward and backwards dynamically. When multiple targets are in close proximity, multiple estimates may be placed on the same target that leads to missed detection [12].

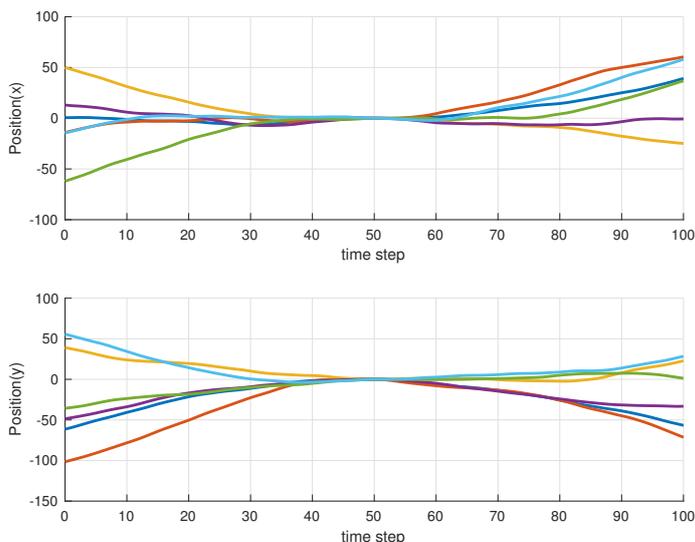
The midpoint is initialized as  $x_{50} = \mathcal{N}(x; 0, 10^{-6} \times I_4)$ . One possible realisation of target trajectories is shown in figure 2.4. Standard deviation of motion and mea-



**Figure 2.3:** Mean GOSPA error cost for missed targets of  $\delta$ -GLMB with joint prediction and update (red), LMB (green), PMBM with recycling (magenta) and PMB using Murty’s algorithm with recycling (blue). Testing cases from top to bottom are (a), (b), (c) and (d) respectively.

surement noises are  $\sigma_v = 0.2m/s^2$  and  $\sigma_\epsilon = 1m$  respectively. For unlabelled filters, the Poisson birth intensity only contains a single Gaussian,  $\lambda^b(x) = 0.1\mathcal{N}(x; 0, \Sigma^b)$ , where  $\Sigma^b = \text{diag}[100^2, 1, 100^2, 1]$  covers the position and velocity region of interest. For labelled filters, the birth intensity is a single Bernoulli with  $r^b = 0.1$  and  $f^b(x) = \mathcal{N}(x; 0, \Sigma^b)$ . To generate clear plots, for each type of filter we choose to evaluate only the one that yields fewer missed targets: PMB using Murty’s algorithm with recycling, PMBM with recycling,  $\delta$ -GLMB with joint prediction and update and LMB. The results of compared filters are shown in figure 2.5.

One significant difference between the unlabelled and labelled filters is that unlabelled filters initiate the tracker faster at initialization period. In cases with high detection probability, none of these filters suffers from the coalescence phenomenon. However, when detection probability decreases to  $P_k^d = 0.75$ , we can observe a substantial increase in missed targets for all the filters after targets become closely spaced. In this case, the labelled filters present fewer missed targets around midpoint that correspond to coalescence than unlabelled filters, though it is difficult for LMB to detect targets when the false alarm rate increases to  $\lambda^{fa} = 30$ . Compared with labelled filters, unlabelled filters exhibit good robustness to coalescence effects, and the performance is restored after targets become separated. As for the error due to false targets, we can observe that labelled filters suffer less from false detections than unlabelled filters. Regarding the localisation error, it is increased abnormally shortly after the midpoint for labelled filters, though  $\delta$ -GLMB is an exception to



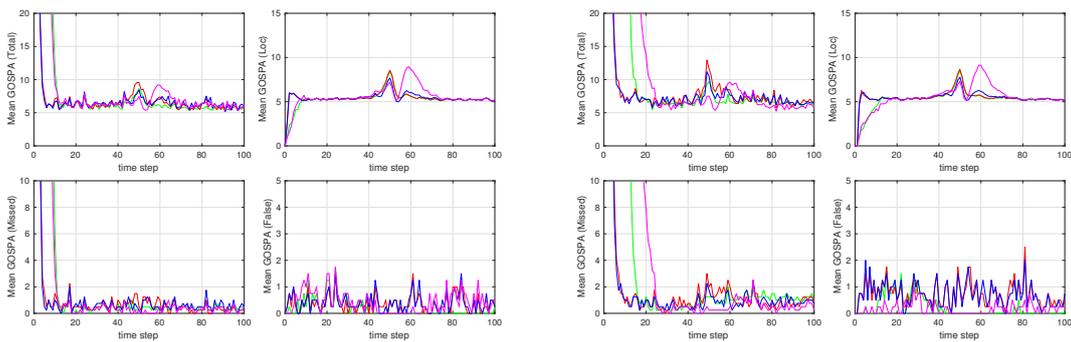
**Figure 2.4:** One dimension of target trajectories in a single Monte Carlo trial on the region  $[-100, 100]m \times [-100, 100]m$ . Six targets are all born at time step 1 and exist throughout the duration time of 100s.

this in cases of high detection probability.

Furthermore, the results demonstrate the success of PMB in resolving the coalescence phenomenon, which presents even less deterioration than PMBM. The advantage with PMB is that it merges tracks using a (variational) technique that reduces coalescence [14, 34]. In the simulation, we found that the data association hypotheses have multiple dominant weights when coalescence happens. In this case, information is lost if we only extract the best states from a single MB component in PMBM. However, approximating MBM as a single MB by minimising the KL divergence yields an MB RFS density which tends to be closer to the full RFS density when targets are in close proximity.

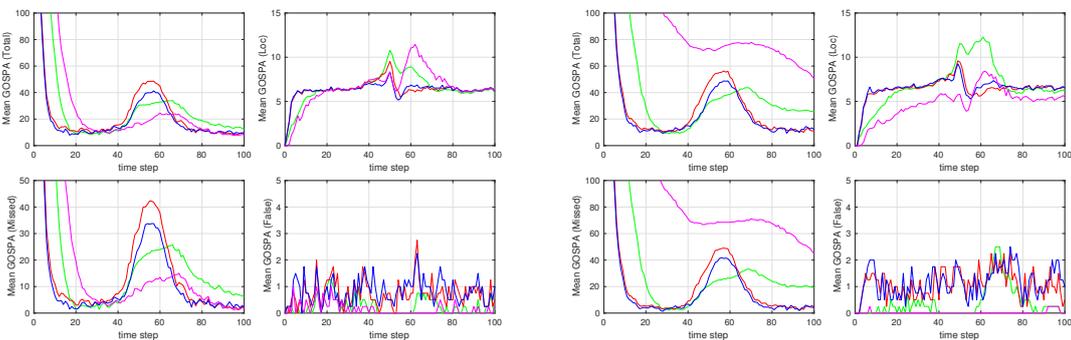
## 2.6 Conclusions and Future Work

In this paper, we have evaluated the  $\delta$ -GLMB, LMB, PMBM and PMB filters in two different scenarios. Estimation error due to localisation, false targets and missed targets have been assessed using the GOSPA metric. Results from the scenario with well-spaced targets show that PMBM is a more efficient filter structure than  $\delta$ -GLMB since it yields similar or better performance with less computational time. The Poisson birth model provides the unlabelled filter with the ability to detect multiple targets born from roughly the same position quickly. Also, approximating the distribution of unknown targets as being Poisson enables the use of recycling methods to improve estimation performance further. Under the simulation with coalescence, the labelled filters have fewer missed targets when targets are in close proximity and fewer false targets than unlabelled filters. The unlabelled filter can



(a)  $\lambda^{fa} = 10, P_k^d = 0.98$

(b)  $\lambda^{fa} = 30, P_k^d = 0.98$



(c)  $\lambda^{fa} = 10, P_k^d = 0.75$

(d)  $\lambda^{fa} = 30, P_k^d = 0.75$

**Figure 2.5:** Filtering performance evaluation of  $\delta$ -GLMB with joint prediction and update (green), LMB (magenta), PMBM with recycling (red), and PMB (Murty) with recycling (blue) using GOSPA in a challenging scenario with coalescence.

resolve the coalescence phenomenon, whereas  $\delta$ -GLMB and LMB are affected at various degrees in cases with low detection probability. By applying variational approximation, PMB presents the best overall performance in this challenging scenario.

In the simulation study, filters are evaluated using the MAP cardinality estimator with fixed parameter settings. It would be interesting to evaluate and compare filters with different estimators and with different complexities by tuning their parameters, e.g., pruning threshold and requested number of components used in Murty’s algorithm. A technique to solve the data association problem based on Gibbs sampling was proposed in [35], and it has been applied in  $\delta$ -GLMB with joint prediction and update [10] to replace Murty’s algorithm to truncate hypotheses. Thus, a potential direction for future work is to evaluate the performance and computational time of different filters using Gibbs sampling.



# Paper II: Likelihood-based data association for extended object tracking using sampling methods

## Abstract

Environment perception is a key enabling technology in autonomous vehicles, and multiple object tracking is an important part of this. The use of high resolution sensors, such as automotive radar and lidar, leads to the extended object tracking problem, with multiple detections per tracked object. For computationally feasible multiple extended object tracking, the data association problem must be handled. Previous work has relied on a two step approach, using clustering algorithms, together with assignment algorithms, to achieve this. In this paper, we show that it is possible to handle the data association in a single step that maximises the desired likelihood function. Single step data association is beneficial, because it enables better use of the measurement model and the predicted multi-object density. For single step data association, we use algorithms based on stochastic sampling, and integrate them into a Poisson Multi-Bernoulli Mixture filter. Four sampling algorithms are compared in a simulation study; the best one is compared to clustering and assignment in an experiment, where Velodyne data is used to track pedestrians. The results from the experiment clearly show that the proposed single-step likelihood-based method achieves the best performance, especially in challenging scenarios where the pedestrians are very close.

Keywords: Target tracking, extended target, data association, Markov Chain Monte Carlo, Metropolis-Hastings, Gibbs sampling, pedestrian tracking, autonomous vehicles, lidar, Velodyne.

## 3.1 Introduction

The realization of highly automated driving functions, as well as advanced driver assistance systems, requires a precise and consistent environment model that incorporates dynamic as well as static obstacles. Object-based representations are typically used for dynamic obstacles whose state is estimated over time using multiple target tracking (MTT) approaches. The main challenge of MTT is the simultaneous

---

estimation of the number of objects and their individual states, based on a sequence of noisy measurements including missed detections and false alarms.

Modern lidar and radar sensors typically deliver multiple detections per object per time step, since the object’s extent is not negligible in comparison to the sensor resolution. During the last decade, extended object tracking approaches, capable of handling several measurements per object, have been proposed. A single extended object measurement model has to model the number of measurements per object, as well as the distribution of each measurement. The Poisson Point Process (PPP) model [16] is a very popular model, where the individual measurements are modelled as spatially distributed around the object. Two popular spatial distribution models are the Gaussian inverse Wishart (GIW), or random matrix, model [18, 19], and the Random Hypersurface Models [17]; a longer discussion of spatial distribution models is given in [36]. The GIW model describes the spatial distribution as Gaussian with unknown mean and covariance; this implies elliptically shaped objects. In [20, 21], the GIW approach was extended to additionally estimate the Poisson measurement rate of each object, resulting in the Gamma GIW (GGIW) model.

The most common approach to multiple extended object tracking is to model the problem using random finite sets (RFS) [3]. The GIW and GGIW models have been integrated into RFS based filters, e.g., the Probability Hypothesis Density (PHD) filter [22, 23]; the Cardinalized PHD (CPHD) filter [21]; the  $\delta$ -Generalized Labeled Multi-Bernoulli filter and the Labeled Multi-Bernoulli filter ( $\delta$ -GLMB and LMB) [24]; as well as the Poisson Multi-Bernoulli Mixture (PMBM) filter [25, 26]. The PHD and CPHD filters are examples of moment approximations of the multi-object density, while the  $\delta$ -GLMB and PMBM are examples of multi-object conjugate priors. Common to all mentioned multiple extended object tracking filters is that they have to deal with the unknown measurement origin—it is unknown if a measurement’s source is clutter or an object—something that leads to the combinatorially complex data association problem.

In previous work [21, 22, 23, 24, 25, 26] the data association problem is handled in two stages: first, clustering methods are used to find a set of different ways in which the measurements can be clustered; second, assignment methods are used to assign measurement clusters to objects. While these approaches have produced good tracking results, they have difficulty handling dense scenarios with close objects, and rely on heuristics rather than directly maximising the multi-object likelihood function.

In this paper, we show that it is possible to handle the data association problem in a single step, i.e., partitioning and assignment (C&A) simultaneously, while focusing on the multi-object likelihood. To achieve this we use different stochastic sampling methods, which are then integrated into the PMBM filter for multiple extended object tracking. A preliminary version of this work was presented in [27]. In this paper, we present an extended version of that work, whose contributions can be summarised as follows:

- 
- The data association problem inevitably leads to approximations of the filtered multi-object density. In Section IV, we discuss different methods for density approximation and analyse the resulting approximation error.
  - In Section V, we present four different sampling algorithms that can be used to find data associations. The first was presented in [35], the second and third are extended object tracking adaptations of Bayesian non-parametric sampling algorithms presented in [37, 38], and the fourth is entirely novel work.
  - Results from an extensive simulation study, in which we compare the four different sampling algorithms, are presented in Section VI. Based on the simulation results we can identify which of the sampling algorithms has more promising performance.
  - An experiment where Velodyne data is used to track pedestrians is presented in Section VII. Here, sampling methods for single-step likelihood based data association are compared to two-step C&A based data association. The results clearly show the benefit of the proposed approach.

In Section II, we give a problem formulation, discuss previous work and explain how the proposed approach improves upon it. A review of multiple extended object tracking is given in Section III. The paper is concluded in Section VIII.

## 3.2 Problem formulation, previous work, and proposed approach

In this section, we clarify our contributions, and we explain why they are important. To do so, we give some background to the problem and describe the previous approaches that can be found in the literature.

### 3.2.1 Problem formulation

Let  $\mathbf{x}_k^i$  denote the state of the  $i$ th object at discrete time step  $k$ , and let the object set be denoted  $\mathbf{X}_k$ . The object set is modeled as a random finite set (RFS) [3], meaning that the number of objects, i.e., the object set cardinality is a time-varying discrete random variable, and each object state  $\mathbf{x}_k \in \mathbf{X}_k$  is a random variable. The set of measurements obtained at time step  $k$  is also modeled as an RFS, and is denoted  $\mathbf{Z}_k$ . Further,  $\mathbf{Z}^k$  denotes all measurement sets  $\mathbf{Z}_m$  from  $m = 0$  up to, and including,  $m = k$ . The MTT problem consists of using the information in the sequence of measurement sets  $\mathbf{Z}^k$  to estimate the set of objects  $\mathbf{X}_k$ , i.e., we are interested in both the number of objects, as well as the states of the individual objects.

The multi-object density is denoted  $f_{k|k}(\mathbf{X}_k|\mathbf{Z}^k)$ . This distribution captures what is known about the set of objects, given the sets of measurements. In this work we use the PMBM conjugate prior [12, 25, 26] to represent the multi-object distribution. The extended object PMBM filter [25, 26] propagates in time the object set PMBM density using a Bayes update and a Chapman-Kolmogorov prediction. The focus in

---

this work is on the Bayes update, defined as

$$f_{k|k}(\mathbf{X}_k|\mathbf{Z}^k) = \frac{f_k(\mathbf{Z}_k|\mathbf{X}_k)f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}^{k-1})}{\int f_k(\mathbf{Z}_k|\mathbf{X}_k)f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}^{k-1})\delta\mathbf{X}_k}, \quad (3.20)$$

where  $f_k(\mathbf{Z}_k|\mathbf{X}_k)$  is the multi-object measurement set density, and  $\int g(\mathbf{X})\delta\mathbf{X}$ , for some function  $g(\mathbf{X})$ , denotes the set integral, defined in [3].

Because of the unknown measurement origin, the update includes consideration of all possible data associations  $A$ , i.e., all possible ways to assign measurements to measurement sources. This has combinatorial complexity, and approximations are necessary to enable computationally tractable tracking algorithms. The approximation consists of finding a subset of associations  $A$  that are highly probable, i.e., that have high likelihoods  $\mathcal{L}_{k|k-1}^A$ , and then truncating the associations that are unlikely. Finding a subset of highly probable associations in multiple extended object tracking is the problem considered in this paper.

### 3.2.2 Previous two-step approach: clustering and assignment

Data association in extended object tracking can be understood by separating it into two parts, called partitioning and assignment, see, e.g., [36]. The partitioning divides  $\mathbf{Z}_k$  into non-empty disjoint subsets called cells—this specifies which measurements are from the same object. An illustrative example of partitioning is given in [36, Sec. IV.C]. Given a partitioning, an assignment then assigns the cells to measurement sources, either clutter or objects. Previous work, see, e.g., [22, 23, 24, 25], has dealt with the data association problem by first using combinations of different clustering algorithms to compute a set of partitions, and then using, for each partition, assignment algorithms, e.g., Murty’s algorithm [30], to assign the cells to different objects. This efficiently reduces the number of data associations, and empirical results are good.

In early multiple extended object tracking, a clustering scheme called distance partitioning (DP) was proposed [22]. DP builds upon a single linkage hierarchical clustering algorithm, and given an upper and a lower threshold for the spatial proximity of measurements, it finds a subset of partitions. DP has been shown to work well when the objects are spatially separated, however it has also been shown to result in missed objects when two or more objects are spatially close [22]. The reason is that the detections from multiple spatially close objects are also spatially close, and clustering algorithms based on spatial proximity, which DP and also Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [39] are two examples of, will then return a single measurement cluster, incorrectly indicating that there is only one object.

To improve the tracking results for close objects, two different extensions based on k-means clustering and expectation maximisation for Gaussian mixtures, called sub-partition (SP) and EM-partition (EMP), respectively, were proposed in [22, 23].

---

Both simulations, as well as experiments where 2D lidar is used to track pedestrians, have shown that handling of close objects is possible when DP is used together with SP and EMP, see [23]. Nevertheless, a few drawbacks can be identified:

- The clustering does not consider the multi-object likelihood, which is the quantity that we ultimately would like to maximise with the computed association.
- To handle close objects multiple clustering algorithms are necessary, see, e.g., [23].
- Performance depends on the clustering parameters. In the worst case, incorrect parameter values lead to poor clustering, which leads to poor tracking results.

Even though intuition about the tracking problems can be used as guidance when selecting and parameterising the clustering algorithms, there is no guarantee that the computed clusterings will actually yield highly probable associations. Instead, one has to rely on empirical evidence that confirms that this is the case.

### 3.2.3 Proposed single-step approach

In this paper, we show that it is possible to solve the extended object data association problem in a single step, while focusing on the multi-object likelihood. We achieve this by using sampling based methods, and show that these methods can be integrated in existing tracking filters, such as the GGIW-PMBM filter.

The main idea in the sampling methods is to start with an initial data association for the measurements, and then randomly take actions that change the association. By basing the probability of selecting an action on its impact on the multi-object likelihood—the higher the resulting likelihood, the higher the probability of selecting the action—the sampling approaches efficiently find highly probable associations. From an obtained sequence of associations, a subset of associations with high likelihoods is then taken. The use of sampling algorithms for data association in extended object tracking is inspired by works in Bayesian nonparametrics, see, e.g., [37, 38], and work on extended object mapping [35].

The benefits of a single-step approach based on sampling methods, in comparison to previous two-step approaches based on C&A, can be summarised as follows:

- It works directly on the likelihood, which is directly related to the quantity which we wish to maximise.
- It works equally well for both close and distant objects.
- It only has one parameter: the number of iterations.

In order to evaluate the tracking performance and the computational cost of single-step data association, sampling methods are integrated into the PMBM filter. In a simulation study, different sampling algorithms are compared, and the best one is used in an experiment where Velodyne lidar data is used for pedestrian tracking. In the experiment, single step data association using sampling is compared to the most common variants of two step C&A data associations, namely DP, and DP+SP/EMP.

Note that the sampling methods are not limited to the PMBM filter; it is straight-

---

forward to integrate them into other types of multiple extended object tracking filters. For example, when integrated into the PHD filter the tracking performance is improved; this is similar to the PMBM results presented in this paper. Due to page length constraints, the PHD filter results were not included in this paper.

### 3.3 Multiple extended object tracking

In this section, we review random set modelling of multiple extended object tracking, and review the standard extended object tracking model. Next we briefly introduce the PMBM conjugate prior for multiple extended object tracking, and then give an overview of the data association within the PMBM measurement update, under the assumed standard measurement model.

#### 3.3.1 Random set modelling

Two types of RFSs are important in this work: the PPP and the Bernoulli process. A PPP is an RFS whose cardinality is Poisson distributed, and each object is independent and identically distributed (i.i.d.). A PPP has one parameter, the intensity  $D(\mathbf{x})$ . The intensity can be broken down into two parts  $D(\mathbf{x}) = \mu f(\mathbf{x})$ : the Poisson rate  $\mu$  and the spatial distribution  $f(\mathbf{x})$ .

A Bernoulli RFS  $\mathbf{X}$  is an RFS that is empty with probability  $1-r$  or, with probability  $r$ , contains a single element with distribution  $f(\mathbf{x})$ . In other words, the cardinality is Bernoulli distributed with parameter  $r$ . A typical assumption in MTT is that the objects are independent, see, e.g., [40]. A multi-Bernoulli (MB) RFS  $\mathbf{X}$  is the union of independent Bernoulli RFSs  $\mathbf{X}^i$ ,  $\mathbf{X} = \uplus_{i \in \mathbb{I}} \mathbf{X}^i$ , where  $\mathbb{I}$  is an index set. The MB distribution is defined entirely by the parameters  $\{r^i, f^i(\cdot)\}_{i \in \mathbb{I}}$  of the involved Bernoulli RFSs. Here  $|\mathbb{I}|$  is the maximum number of objects that the MB RFS can represent.

Lastly, an MB mixture (MBM) is an RFS whose multi-object density is a normalized weighted sum of MB densities, where the components correspond to, e.g., different data association sequences. An MBM density is defined entirely by the set of parameters  $\{(\mathcal{W}^j, \{r^{j,i}, f^{j,i}(\cdot)\}_{i \in \mathbb{I}^j})\}_{j \in \mathbb{J}}$ , where  $\mathbb{J}$  is an index set for the MBs in the MBM (also called components of the MBM),  $\mathcal{W}^j$  is the weight of the  $j$ th MB, and  $\mathbb{I}^j$  is an index set for the Bernoullis in the  $j$ th MB.

#### 3.3.2 Standard extended object measurement model

The set of measurements obtained at time step  $k$  is the union of object generated measurements and clutter measurements,  $\mathbf{Z}_k = (\cup_i \mathbf{W}_k^i) \cup \mathbf{K}_k$ , where  $\mathbf{K}_k$  denotes the set of clutter measurements, and  $\mathbf{W}_k^i$  denotes the set of measurements from the  $i$ th object. Standard MTT assumptions are that the sets are all independent, and that the measurement origin is unknown, i.e., for the measurement set  $\mathbf{Z}_k$  it is not known which measurements are clutter, nor is it known which measurements originated from which object.

---

The set of clutter measurements  $\mathbf{Z}_k$  is modelled as a PPP with rate  $\lambda$  and spatial distribution  $c(\mathbf{z})$ , and the clutter PPP intensity is  $\kappa(\mathbf{z}) = \lambda c(\mathbf{z})$ . An extended object with state  $\mathbf{x}_k^i$  is detected with state dependent probability of detection  $p_D(\mathbf{x}_k^i)$ , and, if it is detected, the object measurement set  $\mathbf{W}_k^i$  is modelled as a PPP with state dependent Poisson rate  $\gamma(\mathbf{x}_k^i)$  and spatial distribution  $\phi(\cdot|\mathbf{x}_k^i)$ . For a non-empty set of measurements ( $|\mathbf{W}_k^i| > 0$ ) the conditional extended object measurement set likelihood is denoted

$$\ell_{\mathbf{W}_k^i}(\mathbf{x}_k^i) = p_D(\mathbf{x}_k^i) e^{-\gamma(\mathbf{x}_k^i)} \prod_{\mathbf{z} \in \mathbf{W}_k^i} \gamma(\mathbf{x}_k^i) \phi(\mathbf{z}|\mathbf{x}_k^i). \quad (3.21)$$

Note that this is the product of the probability of detection and the PPP density. The effective probability of detection for an extended object with state  $\mathbf{x}_k^i$  is  $p_D(\mathbf{x}_k^i)(1 - e^{-\gamma(\mathbf{x}_k^i)})$ , where  $1 - e^{-\gamma(\mathbf{x}_k^i)}$  is the Poisson probability of generating at least one detection. Accordingly, the effective probability of missed detection, i.e., the probability that the object is not detected, is

$$q_D(\mathbf{x}_k^i) = 1 - p_D(\mathbf{x}_k^i) + p_D(\mathbf{x}_k^i) e^{-\gamma(\mathbf{x}_k^i)}. \quad (3.22)$$

Note that  $q_D(\mathbf{x}_k^i)$  is the conditional likelihood for an empty set of measurements, i.e.,  $\ell_\emptyset(\mathbf{x}_k^i) = q_D(\mathbf{x}_k^i)$  (cf. (3.21)).

### 3.3.3 The PMBM conjugate prior

The PMBM conjugate prior was developed for extended objects in [25, 26], and for point objects in [12]. The PPP describes the distribution of the objects that are thus far undetected, while the MBM describes the distribution of the objects that have been detected at least once. Thus, the set of objects can be divided into two disjoint subsets,  $\mathbf{X} = \mathbf{X}^u \uplus \mathbf{X}^d$ , corresponding to undetected objects  $\mathbf{X}^u$  and detected objects  $\mathbf{X}^d$ . The PMBM density is defined entirely by the parameters

$$D^u, \{(\mathcal{W}^j, \{(r^{j,i}, f^{j,i})\}_{i \in \mathbb{I}^j})\}_{j \in \mathbb{J}}, \quad (3.23)$$

where

- $D^u(\cdot)$  is the PPP intensity for the set of undetected objects  $\mathbf{X}^u$ . Higher/lower intensity in a location implies a higher/lower probability that a yet undetected object is located there. This can be used to model object occlusions, see [26] for an example.
- $j$  is an index for the MB components, and  $\mathbb{J}$  is an index set. There are  $|\mathbb{J}|$  MB components. Each MB component corresponds to a unique global hypothesis for the detected objects, i.e., a particular history of data associations for all detected objects.
- The probability, or weight, of the  $j$ th MB component is  $\mathcal{W}^j$ .
- For the  $j$ th MB component,  $i$  is an index for the Bernoulli estimates, and  $\mathbb{I}^j$  is an index set. The  $j$ th component has  $|\mathbb{I}^j|$  Bernoulli estimates. In the  $j$ th global association hypothesis, the  $i$ th Bernoulli has probability of existence  $r^{j,i}$  and object state pdf  $f^{j,i}(\cdot)$ .

### 3.3.4 PMBM data association

Let the predicted multi-object density be a PMBM density, with parameters

$$D^u, \{(\mathcal{W}^j, \{(r^{j,i}, f^{j,i})\}_{i \in \mathbb{I}})\}_{j \in \mathbb{J}}, \quad (3.24)$$

let the measurements in the set  $\mathbf{Z}$  be indexed by  $m \in \mathbb{M}$ ,

$$\mathbf{Z} = \{\mathbf{z}^m\}_{m \in \mathbb{M}} \quad (3.25)$$

and let  $\mathcal{A}^j$  be the space of all data associations  $A$  for the  $j$ th predicted hypothesis. A data association  $A \in \mathcal{A}^j$  is an assignment of each measurement in  $\mathbf{Z}$  to a source, either to the background (clutter or new object) or to one of the existing objects indexed by  $\mathbb{I}^j$ . Note that  $\mathbb{M} \cap \mathbb{I}^j = \emptyset$  for all  $j$ .

Formally a data association  $A \in \mathcal{A}^j$  consists of a partition of  $\mathbb{M} \cup \mathbb{I}^j$  into non-empty disjoint subsets called index cells, denoted  $C \in A$ . The meaning of a cell is that all elements in the cell are associated together<sup>3</sup>. Due to the standard MTT assumption that the objects generate measurements independent of each other, an index cell contains either no object index, or one object index, i.e.,  $C \cap \mathbb{I}^j \neq \emptyset \Rightarrow |C \cap \mathbb{I}^j| = 1$  for all  $C \in A$ . If the index cell  $C$  contains a object index, then let  $i_C$  denote the corresponding object index. Further, let  $\mathbf{C}_C$  denote the measurement cell that corresponds to the index cell  $C$ , i.e., the set of measurements

$$\mathbf{C}_C = \cup_{m \in C} \mathbf{z}^m. \quad (3.26)$$

By [26, Theorem 1], the updated multi-object density is also a PMBM. Given a predicted set of PMBM parameters, for each predicted global hypothesis there are multiple possible data associations, each of which will result in a MB component in the updated MB mixture. The weight for the component in the updated PMBM, that resulted from updating predicted global hypothesis  $j \in \mathbb{J}$  with association  $A \in \mathcal{A}^j$ , is [26]

$$\mathcal{W}_A^j = \frac{\mathcal{W}^j \mathcal{L}_A^j}{\sum_{j' \in \mathbb{J}} \sum_{A' \in \mathcal{A}^{j'}} \mathcal{W}^{j'} \mathcal{L}_{A'}^{j'}} \quad (3.27)$$

where  $\mathcal{W}^j$  is the predicted weight, and  $\mathcal{L}_A^j$  is the likelihood of predicted hypothesis  $j$  and association  $A \in \mathcal{A}^j$ .

---

<sup>3</sup>For example, let  $\mathbb{M} = (m_1, m_2, m_3)$  and  $\mathbb{I} = (i_1, i_2)$ , i.e., three measurements and two objects. One valid partition of  $\mathbb{M} \cup \mathbb{I}$ , i.e., one of the possible associations, is  $\{i_1, m_1, m_2\}, \{m_3\}, \{i_2\}$ . The meaning of this is that measurements  $m_1, m_2$  are associated to object  $i_1$ , object  $i_2$  is not detected, and measurement  $m_3$  is not associated to any previously detected object, i.e., measurement  $m_3$  is either clutter or from a new object.

The likelihood  $\mathcal{L}_A^j$  can be expressed as [26]

$$\mathcal{L}_A^j = \prod_{\substack{C \in A: \\ C \cap \mathbb{I}^j = \emptyset \\ C \cap \mathbb{M} \neq \emptyset}} \mathcal{L}_{\mathbf{C}_C}^b \prod_{\substack{C \in A: \\ C \cap \mathbb{I}^j \neq \emptyset \\ C \cap \mathbb{M} \neq \emptyset}} \mathcal{L}_{\mathbf{C}_C}^{j,i_C} \prod_{\substack{C \in A: \\ C \cap \mathbb{I}^j \neq \emptyset \\ C \cap \mathbb{M} = \emptyset}} \mathcal{L}_{\emptyset}^{j,i_C} \quad (3.28a)$$

$$\mathcal{L}_{\mathbf{C}}^b = \begin{cases} \kappa^{\mathbf{C}} + \langle D^u; \ell_{\mathbf{C}} \rangle & \text{if } |\mathbf{C}| = 1 \\ \langle D^u; \ell_{\mathbf{C}} \rangle & \text{if } |\mathbf{C}| \geq 1 \end{cases} \quad (3.28b)$$

$$\mathcal{L}_{\mathbf{C}}^{j,i} = r^{j,i} \langle f^{j,i}; \ell_{\mathbf{C}} \rangle \quad (3.28c)$$

$$\mathcal{L}_{\emptyset}^{j,i} = 1 - r^{j,i} + r^{j,i} \langle f^{j,i}; q_D \rangle \quad (3.28d)$$

where  $\langle a; b \rangle = \int a(x)b(x)dx$ . The three products in (3.28a) correspond to

- cells that are associated to the background,
- cells that are associated to previously detected objects,
- previously detected objects that are miss-detected.

By multiplication by 1 in (3.28a) we get

$$\mathcal{L}_A^j = \prod_{\substack{C \in A: \\ C \cap \mathbb{I}^j = \emptyset \\ C \cap \mathbb{M}_k \neq \emptyset}} \mathcal{L}_{\mathbf{C}_C}^b \prod_{\substack{C \in A: \\ C \cap \mathbb{I}^j \neq \emptyset \\ C \cap \mathbb{M}_k \neq \emptyset}} \frac{\mathcal{L}_{\mathbf{C}_C}^{j,i_C}}{\mathcal{L}_{\emptyset}^{j,i_C}} \prod_{i \in \mathbb{I}^j} \mathcal{L}_{\emptyset}^{j,i} \quad (3.29)$$

$$= \mathcal{L}_{\emptyset}^j \prod_{\substack{C \in A: \\ C \cap \mathbb{M}_k \neq \emptyset}} \mathcal{L}_{\mathbf{C}_C}^j \quad (3.30)$$

where  $\mathcal{L}_{\emptyset}^j = \prod_{i \in \mathbb{I}^j} \mathcal{L}_{\emptyset}^{j,i}$  is the likelihood that none of the previously detected objects are detected in this time step, and

$$\mathcal{L}_{\mathbf{C}_C}^j = \begin{cases} \mathcal{L}_{\mathbf{C}_C}^b & \text{if } C \cap \mathbb{I}^j = \emptyset \\ \mathcal{L}_{\mathbf{C}_C}^{j,i_C} (\mathcal{L}_{\emptyset}^{j,i_C})^{-1} & \text{if } C \cap \mathbb{I}^j = \{i_C\} \end{cases} \quad (3.31)$$

is the likelihood for measurement cell  $\mathbf{C}_C$  under association  $A \in \mathcal{A}^j$ .

The updated weight resulting from updating predicted hypothesis  $j$  with association  $A \in \mathcal{A}^j$ , see (3.27), can thus be written as follows

$$\mathcal{W}_A^j \propto \mathcal{W}^j \mathcal{L}_{\emptyset}^j \prod_{\substack{C \in A: \\ C \cap \mathbb{M} \neq \emptyset}} \mathcal{L}_{\mathbf{C}_C}^j \quad (3.32)$$

Notice that, for predicted hypothesis  $j$ , the predicted weight  $\mathcal{W}^j$  and the likelihood of all objects being missed  $\mathcal{L}_{\emptyset}^j$  are constant and independent of the data association  $A \in \mathcal{A}^j$ . In other words, the difference lies in the product over the cell likelihoods  $\mathcal{L}_{\mathbf{C}_C}^j$ .

### 3.4 Multi-object density approximation

As noted above, the size of the association space  $\mathcal{A}^j$  is extremely large, it is not possible to consider all associations in the tracking update, and approximations

are necessary. In this section, we discuss some important theoretical aspects of multi-object density approximation, with a focus on the approximation errors that follow from considering only a subset of data associations. It will be shown that the approximation errors are smallest when the data association is approximated by only considering data associations with large likelihoods (3.28). This highlights the benefit of using the proposed single-step likelihood based data association.

### 3.4.1 Multi-object density approximation

It is shown in [26] that, following the Bayes update (3.20), the density is a mixture density, where each component in the mixture corresponds to a particular association sequence. In this section, for the sake of brevity, we express this compactly as

$$f(\mathbf{X}) = \sum_{A \in \mathcal{A}} \mathcal{W}^A f(\mathbf{X}|A), \quad (3.33a)$$

$$\mathcal{W}^A = \frac{\mathcal{L}^A}{\sum_{A' \in \mathcal{A}} \mathcal{L}^{A'}}. \quad (3.33b)$$

Here,  $\mathcal{A}$  is the finite discrete association space,  $\mathcal{L}^A$  is the likelihood of the association  $A \in \mathcal{A}$ , the weight  $\mathcal{W}^A$  is the probability of the association  $A \in \mathcal{A}$ , and  $\sum_{A \in \mathcal{A}} \mathcal{W}^A = 1$ .

The number of possible associations, i.e., the size  $|\mathcal{A}|$  of the association space, is combinatorial, which leads to (at least) two challenges:

1. The number of mixture components, i.e., number of summands in (3.33a), is computationally intractable.
2. The association weights (3.33b) are in general difficult to compute, because the normalisation requires a summation over  $A \in \mathcal{A}$ . However, note that the association likelihoods  $\mathcal{L}^A$  can be computed straightforwardly.

The prevailing approach to these challenges is to compute an approximate mixture density that has a computationally tractable number of components, and to approximate the component weights. The approximate density

$$f_{\hat{\mathcal{A}}}(\mathbf{X}) = \sum_{A \in \hat{\mathcal{A}}} \hat{\mathcal{W}}^A f(\mathbf{X}|A) \quad (3.34)$$

considers a subset of associations,  $\hat{\mathcal{A}} \subseteq \mathcal{A}$ , and approximate association weights  $\hat{\mathcal{W}}^A$ , where  $\sum_{A \in \hat{\mathcal{A}}} \hat{\mathcal{W}}^A = 1$  and  $\hat{\mathcal{W}}^{A'} = 0$  for all  $A' \notin \hat{\mathcal{A}}$ .

### 3.4.2 Approximation error upper bound

The  $L_1$ -norm of a function  $g(\mathbf{X})$  is defined as

$$\|g\|_1 = \int \text{abs}(g(\mathbf{X})) \delta \mathbf{X}, \quad (3.35)$$

and it is a useful performance measure for density approximation, see, e.g., [9]. Using the triangle inequality, it is straightforward to show that the  $L_1$ -error of the

density approximation is upper bounded,

$$\|f - f_{\hat{\mathcal{A}}}\|_1 \leq \sum_{A \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{W}^A + \sum_{A \in \hat{\mathcal{A}}} \text{abs}(\mathcal{W}^A - \hat{\mathcal{W}}^A). \quad (3.36)$$

Equality holds, e.g., in the special case when there is no overlapping support between the conditional densities, i.e.,

$$\text{supp}f(\cdot|A) \cap \text{supp}f(\cdot|A') = \emptyset \quad (3.37)$$

for all  $A \in \mathcal{A}$  and  $A' \in \mathcal{A}$  such that  $A \neq A'$ .

The  $L_1$ -error upper bound (3.36) shows that the approximation error is not only determined by which associations are included in the subset  $\hat{\mathcal{A}}$ , but also by how the approximate association weights  $\hat{\mathcal{W}}^A$  are chosen.

### 3.4.3 Weight approximations

Given a subset of associations, one possibility is to compute the approximate mixture weights as

$$\hat{\mathcal{W}}_L^A = \frac{\mathcal{L}^A}{\sum_{A' \in \hat{\mathcal{A}}} \mathcal{L}^{A'}}. \quad (3.38)$$

We call this *likelihood based weight approximation* (L). For approximate weights computed as in (3.38), the upper bound on the  $L_1$ -error is

$$\|f - f_{\hat{\mathcal{A}}}\|_1 \leq 2 \sum_{A \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{W}^A. \quad (3.39)$$

In other words, the approximation error is less than two times the sum of the weights of the mixture components that are pruned, i.e., the components whose weights are approximated as zero. An interpretation of (3.39) is that, by pruning components with small weights, the sum of pruned weights will be small, and the approximation error will be small.

Consider a subset of associations  $\hat{\mathcal{A}}$ , such that

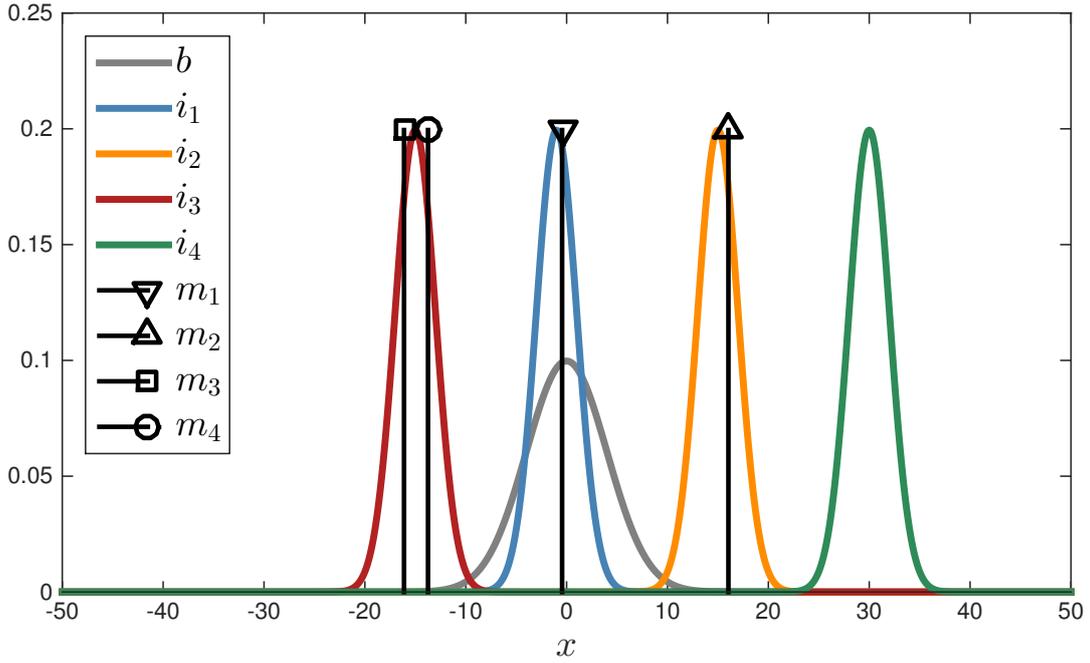
$$\sum_{A \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{W}^A = \varepsilon \Leftrightarrow \sum_{A \in \hat{\mathcal{A}}} \mathcal{W}^A = 1 - \varepsilon, \quad (3.40)$$

for  $\varepsilon > 0$ . We can rewrite this in terms of the likelihoods as follows,

$$\varepsilon = \frac{\sum_{A \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{L}^A}{\sum_{A' \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{L}^{A'} + \sum_{A'' \in \hat{\mathcal{A}}} \mathcal{L}^{A''}}. \quad (3.41)$$

From this we see that if the following holds for the sums of likelihoods,

$$\sum_{A' \in \hat{\mathcal{A}}} \mathcal{L}^{A'} \gg \sum_{A'' \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{L}^{A''}, \quad (3.42)$$



**Figure 3.6:** Illustrative example with four object densities  $i_1$  to  $i_4$ , a PPP intensity  $b$  for the background (new object or clutter), and four measurements  $m_1$  to  $m_4$ . Note that the only association that is highly ambiguous is measurement  $m_1$ : is it from object  $i_1$  or from the background  $b$ ? The corresponding significant associations are given in Table 3.3.

**Table 3.3:** The two most significant associations, and their corresponding weights, for the scenario in Figure 3.6

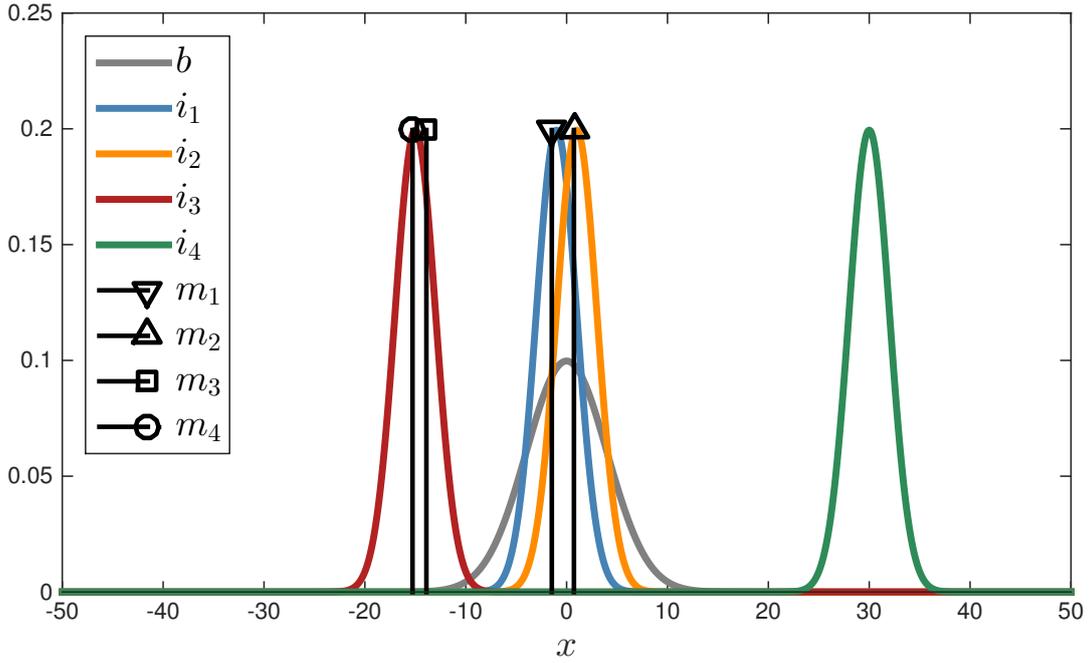
---

$A_{39} :$	$\{i_1, m_1\}, \{i_2, m_2\}, \{i_3, m_3, m_4\}, \{i_4\}$	$\mathcal{W}^A = 0.87$
$A_{616} :$	$\{i_1\}, \{i_2, m_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_1\}$	$\mathcal{W}^A = 0.13$

---

then  $\varepsilon$  in (3.41) will be a small number. In other words, for the density approximation (3.34) with likelihood based weights (3.38), we can make the  $L_1$ -error arbitrarily small by including in  $\hat{\mathcal{A}}$  all associations  $A$  with large likelihoods  $\mathcal{L}^A$ . This is equivalent to including in  $\hat{\mathcal{A}}$  all associations  $A$  that have large weights  $\mathcal{W}^A$ . Note that we can make the upper bound on the  $L_1$ -error arbitrarily small, regardless of the sizes of  $\hat{\mathcal{A}}$  and  $\mathcal{A}$ ; as long as (3.42) holds,  $|\hat{\mathcal{A}}|$  can be several orders of magnitude smaller than  $|\mathcal{A}|$ . We illustrate this with the following example.

**Example 3.4.1** Consider a one-dimensional scenario where there are four existing objects denoted  $i_1$  to  $i_4$ , a PPP representing object birth denoted  $b$ , and four measurements denoted  $m_1$  to  $m_4$ . In this case there are 799 possible data associations  $A_i$ , for the sake of simplicity indexed  $i$  from 1 to 799. Note that the order of the associations does not have a specific meaning; one was simply chosen for the illustration.



**Figure 3.7:** Illustrative example with four object densities  $i_1$  to  $i_4$ , a PPP intensity  $b$  for the background (new object or clutter), and four measurements  $m_1$  to  $m_4$ . Compared to Figure 3.6, the scenario is more complex; the associations of both  $m_1$  and  $m_2$  are ambiguous. For  $m_1$  and  $m_2$  there are three different highly probable measurement sources,  $i_1$ ,  $i_2$  and  $b$ . Further, it is unknown if both measurements are from the same source, or if they are from two different sources. The corresponding significant associations are given in Table 3.4.

One realisation of such a scenario is shown in Figure 3.6, where the object densities and the measurements are shown. In this case only two associations have significant weights; these two associations are given in Table 3.3. The two associations correspond to the only highly ambiguous assignment in this scenario: is measurement  $m_1$  from existing object  $i_1$ , or from a new object? If  $\hat{A} = \{A_{39}, A_{616}\}$ , then the density approximation (3.34) with likelihood based weights (3.38) has an  $L_1$ -error with upper bound (3.39) equal to 0.008. Thus, despite truncating over 99.7% of the possible associations, the approximation error has a very low upper bound.

If object  $i_2$  and measurement  $m_2$  are closer to object  $i_1$  and the birth PPP intensity, the data association becomes much more ambiguous. This is illustrated in Figure 3.7. In this case, if  $\hat{A}$  includes only the two most probable associations, then the density approximation (3.34) with likelihood based weights (3.38) has an  $L_1$ -error with upper bound of 0.9.

For the scenario in Figure 3.7, the ten associations with highest weights are given in Table 3.4. Note that, as expected, in all of the ten significant associations, measurements  $m_3$  and  $m_4$  are associated to object  $i_3$ , and object  $i_4$  is not associated to any measurement. If  $\hat{A}$  is set to include the ten associations listed in Table 3.4, then the  $L_1$ -error has an upper bound of 0.0005, while truncating 98.7% of the possible

---

**Table 3.4:** The ten most significant associations, and their corresponding weights, for the scenario in Figure 3.7

---

$A_{39}$	$\{i_1, m_1\}, \{i_2, m_2\}, \{i_3, m_3, m_4\}, \{i_4\}$	$\mathcal{W}^A = 0.29$
$A_{180}$	$\{i_1\}, \{i_2, m_1, m_2\}, \{i_3, m_3, m_4\}, \{i_4\}$	$\mathcal{W}^A = 0.27$
$A_{154}$	$\{i_1, m_2\}, \{i_2, m_1\}, \{i_3, m_3, m_4\}, \{i_4\}$	$\mathcal{W}^A = 0.14$
$A_{13}$	$\{i_1, m_1, m_2\}, \{i_2\}, \{i_3, m_3, m_4\}, \{i_4\}$	$\mathcal{W}^A = 0.13$
$A_{616}$	$\{i_1\}, \{i_2, m_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_1\}$	$\mathcal{W}^A = 0.04$
$A_{119}$	$\{i_1, m_1\}, \{i_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_2\}$	$\mathcal{W}^A = 0.04$
$A_{260}$	$\{i_1\}, \{i_2, m_1\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_2\}$	$\mathcal{W}^A = 0.04$
$A_{727}$	$\{i_1\}, \{i_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_1, m_2\}$	$\mathcal{W}^A = 0.02$
$A_{579}$	$\{i_1, m_2\}, \{i_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_1\}$	$\mathcal{W}^A = 0.02$
$A_{766}$	$\{i_1\}, \{i_2\}, \{i_3, m_3, m_4\}, \{i_4\}, \{m_1\}, \{m_2\}$	$\mathcal{W}^A = 0.01$

---

associations. □

The example illustrates that to ensure a tight upper bound on the  $L_1$ -error, one must ensure that the subset of associations  $\hat{\mathbf{A}}$  includes all associations with non-negligible likelihoods. The number of associations that have non-negligible likelihoods depends on the scenarios; in the examples in Figure 3.6 and Figure 3.7, this number is 2 and 10, respectively. However, in a typical object tracking scenario, the number of non-negligible associations is not known a priori. Further, this number may be larger than what the computational resources allow, meaning that associations with non-negligible likelihoods must be truncated.

For such cases, the likelihood based weight approximation (3.38) is not guaranteed to yield small errors, and we may possibly obtain better approximations using other weights. When a sampling method has been used to obtain a set of association samples  $\mathbf{A} = \{A^{(i)}\}_{i=1}^N$ , an obvious alternative is to approximate the density (3.33) using the Monte Carlo average,

$$f_{\mathbf{A}}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N f(\mathbf{X}|A^{(i)}) = \sum_{A \in \mathcal{A}} \frac{N^A}{N} f(\mathbf{X}|A), \quad (3.43)$$

where  $N^A$  is the number of samples  $A^{(i)} \in \mathbf{A}$  that are equal to  $A \in \mathcal{A}$ . Not all  $A \in \mathcal{A}$  will be sampled at least once, meaning that  $N^A = 0$  for  $A \notin \mathbf{A}$ . Thus, the Monte Carlo density approximation (3.43) is of the same form as the density approximation (3.34), with  $\hat{\mathcal{A}} = \cup_{i=1}^N \{A^{(i)}\}$  and approximate mixture weights

$$\hat{\mathcal{W}}_S^A = \frac{N^A}{N}. \quad (3.44)$$

We call this *relative frequency based weight approximation* ( $S$ ). For approximate weights computed as in (3.44), the upper bound on the  $L_1$ -error becomes

$$\|f - f_{\mathbf{A}}\|_1 \leq \sum_{A \in \mathcal{A} \setminus \hat{\mathcal{A}}} \mathcal{W}^A + \sum_{A \in \hat{\mathcal{A}}} \text{abs}(\mathcal{W}^A - \hat{\mathcal{W}}_S^A). \quad (3.45)$$

---

Empirical results from simulated tracking scenarios show that weight approximation based on relative frequency (3.44) can give the same tracking performance as weight approximation based on the likelihood (3.38), but relative frequency requires a considerably high number of iterations  $N$ . There are two main reasons for this:

1. Relative frequency weights should (preferably) be based on samples from the stationary distribution. In general, sampling methods, e.g., Markov Chain Monte Carlo (MCMC) algorithms, require a period of “burn-in” before one can reasonably assume that the samples are from the stationary distribution. This means that the batch of samples from the “burn-in” cannot be used, as they are not from the stationary distribution.
2. Relative frequency weights are more accurate, the more samples there are. This means that, after the “burn-in” period, the sampling algorithm has to run for a sufficient number of iterations.

Because, holding all else equal, more iterations lead to higher computational cost, we favour likelihood based weight approximation.

## 3.5 Sampling methods

In this section, we discuss sampling based methods for obtaining a sequence of samples from the discrete distribution over the associations,  $P\{A\}$ . Due to the complex nature of the extended object tracking data associations problem, direct sampling from  $P\{A\}$  is difficult, and approximate sampling methods are therefore used. Specifically, we discuss two existing MCMC methods, called Gibbs sampling and split/merge Metropolis Hastings, and we propose a new algorithm called Stochastic Optimisation (SO).

### 3.5.1 Assignment variable

We start by defining an assignment variable  $\varphi$ ; this variable facilitates the description, as well as the implementation, of the sampling algorithms. Let  $\varphi$  be a vector of length  $|\mathbf{Z}_k|$ , with entries  $\varphi_m$  that indicate which sources the measurements are associated to. If  $\varphi_m \in \mathbb{I}^j$ , then the measurement with index  $m$  is associated to Bernoulli estimate with index  $\varphi_m$ , and if  $\varphi_m \notin \mathbb{I}^j$ , then the measurement with index  $m$  is associated to a background object (clutter or new object). Further, if  $\varphi_{m_1} = \varphi_{m_2}$ , then measurements  $\mathbf{z}_k^{m_1}$  and  $\mathbf{z}_k^{m_2}$  are associated to the same source, and if  $\varphi_{m_1} \neq \varphi_{m_2}$ , then measurements are associated to different sources. Note that for  $\varphi_m \notin \mathbb{I}^j$ , it does not matter what value  $\varphi_m$  takes, it only matters which measurements are associated together (i.e., for which  $m$  we have equal  $\varphi_m$ ). A set of measurements associated to the same source is, analogously to (3.26), denoted

$$\mathbf{C}_c = \bigcup_{m:\varphi_m=c} \mathbf{z}_k^m. \quad (3.46)$$

Given an assignment vector  $\varphi$  we can obtain the equivalent association  $A(\varphi)$  by simply forming subsets of the measurement indices that have equal assignments, and including the object index if  $\varphi_m \in \mathbb{I}^j$ . Note that the associations  $A$  are unique,

---

however, the assignments are not; i.e., two assignments  $\varphi \neq \varphi'$  can have the same equivalent association,  $A(\varphi) = A(\varphi')$ .

Given  $\varphi^{(t)}$ , where superscript  $t$  denotes the iteration, the next assignment  $\varphi^{(t+1)}$  is obtained by the sampling methods as follows. First, a measurement index  $r_1$  is randomly sampled; let  $\mathbf{C}_{c_1}^{(t)}$  denote the corresponding cell, i.e.,  $\varphi_{r_1}^{(t)} = c_1$ . By performing an *action*  $\alpha$  that involves the measurement we can obtain a different assignment vector, i.e., a different association. Thus, starting from the assignment  $\varphi^{(t)}$  and performing a randomly sampled action  $\alpha$ , we get the next assignment  $\varphi^{(t+1)} = \varphi_\alpha^{(t)}$ , i.e., a randomly sampled association  $A_\alpha^{(t)} = A(\varphi_\alpha^{(t)})$ .

In the implementations the number of iterations  $N$  is set adaptively to  $N = N_I(|\mathbb{M}| + |\mathbb{I}^j|)$ , where  $N_I > 1$  is a scaling factor. The following sub-sections presents the different sampling methods in detail.

### 3.5.2 MCMC sampling methods

MCMC methods are a class of algorithms that construct a Markov chain whose stationary distribution is equal to the probability distribution that is of interest, which in this context is  $P\{A\}$ . MCMC methods have been used to handle data association in point object tracking, see, e.g., [10, 41, 42, 43, 44]. However, those results do not apply here, because of the significant differences between point object tracking and extended object tracking, especially in terms of how the data association is defined.

#### 3.5.2.1 Gibbs Sampling

Gibbs sampling is an iterative MCMC method that can be used when direct sampling from a multivariate distribution is difficult. Applied to extended object tracking data association, the basic idea is that given the randomly chosen measurement index  $r_1$ , the corresponding association is randomly changed while holding all other associations fixed. For the assignment vector, this means that  $\varphi_r$  is randomly changed, while all other elements of  $\varphi$  are constant.

A Gibbs sampling algorithm was developed for the PMBM conjugate prior in [35]; the Gibbs algorithm in [35] was in turn inspired by work by Neal [45], where the ideas were applied to Dirichlet Process mixture models. For conjugate models, the Gibbs algorithm uses the following conditional distribution, cf. [45, Eqn. 3.2] and [35, Eqn. 45],

$$P\{\varphi^{(t+1)} = \varphi_\alpha^{(t)} \mid \varphi^{(t)}, r_1\} \propto \mathcal{L}_{k|k-1}^{j, A_\alpha^{(t)}}. \quad (3.47)$$

The Gibbs actions, and the updates of the association vector, are given in Table 3.5, where  $\mathbb{S}^{(t)} = \mathbb{I}^j \cup \{\varphi_m^{(t)}\}$  is used for notational brevity.

#### 3.5.2.2 Split/Merge Metropolis-Hastings

In [38] it was noted that the Gibbs sampling [45] can be slow, because it changes at most the assignment of a single measurement in each iteration. To speed things

---

**Table 3.5:** Single measurement actions

---

Update assignments as

$$\varphi_m^{(t+1)} = \begin{cases} c_\star & \text{if } m = r_1 \\ \varphi_m^{(t)} & \text{otherwise} \end{cases} \quad (3.48)$$

where there are three different possibilities for  $c_\star$ ,

1.  $c_\star = c_1$ : Do not change the assignment vector.
  2.  $c_\star \in \mathbb{S}^{(t)} \setminus c_1$ : Move selected index  $r$  to one of the existing cells.
  3.  $c_\star \notin \mathbb{S}^{(t)}$ : Move selected index  $r$  to a new cell.
- 

up, a Metropolis-Hastings (MH) algorithm was proposed in [38], in which so called splits and merges were used. The splits and merges change the assignments for entire cells in each iteration, which leads to a faster algorithm. MH algorithms are suitable when we wish to sample from a distribution which we can evaluate up to a normalising constant; this is the case for MTT data associations where we can evaluate the likelihood  $\mathcal{L}^A$  of an association  $A$ , but not the weight  $\mathcal{W}^A$ .

In the MH algorithm, one additional measurement index  $r_2$  is randomly chosen. If the two indices belong to the same cell, i.e., if  $\varphi_{r_1} = \varphi_{r_2} = c_1$ , then corresponding cell is a candidate for a cell split. If  $\varphi_{r_1} = c_1 \neq \varphi_{r_2} = c_2$ , then the two corresponding cells are considered for a cell merge. For a merge, there is one way to merge the two involved cells. However, for a split, there are multiple ways in which the two sub-cells can be assigned, i.e., there are multiple possible split-actions. When the MH algorithm is integrated into a tracking filter, we select the split-action that has highest likelihood. The split/merge actions, and the updates of the association vector, are given in Table 3.6.

There are many ways to split a set of measurement indices into two subsets. We have evaluated two alternatives in a simulation study, namely restricted Gibbs sampling, suggested in [38], and the  $k$ -means++ algorithm [46]. Empirically we found that both can yield equally good results, however, restricted Gibbs requires a larger number of iterations, and therefore has a significantly higher computational cost. Further, in a simulation scenario with very dense birth, restricted Gibbs gave worse results than  $k$ -means++, even with a very high number of iterations. From these tests we conclude that  $k$ -means++ is more suitable for use in an extended object tracking context.

The Metropolis-Hastings (MH) algorithm with merge and split actions in [38] has acceptance probability of the form [38, Eqn. 3.1]

$$\begin{aligned} P\{\alpha\} &= P\left\{\varphi^{(t+1)} = \varphi^{(t)}(\alpha) \mid \varphi^{(t)}\right\} \\ &= \min \left[ 1, \frac{q\left(\varphi^{(t)} \mid \varphi^{(t)}(\alpha)\right) P\{A_\alpha^{(t)}\}}{q\left(\varphi^{(t)}(\alpha) \mid \varphi^{(t)}\right) P\{A^{(t)}\}} \right] \end{aligned} \quad (3.51)$$

where the action  $\alpha$  is either split or merge;  $q(\cdot|\cdot)$  is a proposal density;  $P\{A\}$  is the

---

**Table 3.6:** Cell merge and cell split actions

---

1. Cell merge:

$$\varphi_m^{(t+1)} = \begin{cases} c_2 & \text{if } \varphi_m^{(t)} = c_1 \\ \varphi_m^{(t)} & \text{otherwise} \end{cases} \quad (3.49)$$

2. Cell split: First, split set of indices  $\{m\}_{\varphi_m^{(t)}=c_1}$  into two subsets denoted  $R_1$  and  $R_2$ . Next, given the split, either the indices  $m \in R_1$  are assigned to another cell,

$$\varphi_m^{(t+1)} = \begin{cases} c_\star & \text{if } m \in R_1 \\ \varphi_m^{(t)} & \text{otherwise} \end{cases} \quad (3.50a)$$

or the indices  $m \in R_2$  are assigned to another cell,

$$\varphi_m^{(t+1)} = \begin{cases} c_\star & \text{if } m \in R_2 \\ \varphi_m^{(t)} & \text{otherwise} \end{cases} \quad (3.50b)$$

For both alternatives, we may have  $c_\star \in \mathbb{S}^{(t)}$  and the split cell is assigned to an existing cell, or  $c_\star \notin \mathbb{S}^{(t)}$  and the split cell is assigned to a new cell.

---

density that we want to sample from; and, for the PMBM filter, we have that

$$\frac{P\{A_\alpha^{(t)}\}}{P\{A^{(t)}\}} = \frac{\mathcal{L}_{k|k-1}^{j, A_\alpha^{(t)}}}{\mathcal{L}_{k|k-1}^{j, A^{(t)}}}. \quad (3.52)$$

Choosing the proposal density is in general not simple, and it is especially difficult here, considering the very large dimension of the MTT data association problem. A general discussion of different methods for cell split, and the corresponding proposal densities, is given in [38]. Empirically, we have found that for object tracking with the PMBM filter, best results are achieved if the proposal density is chosen to be symmetric, i.e.,  $\frac{q(\varphi^{(t)}|\varphi^{(t)}(\alpha))}{q(\varphi^{(t)}(\alpha)|\varphi^{(t)})} = 1$ . In this case the acceptance probabilities become

$$P\{\text{merge}\} = \min \left[ 1, \mathcal{L}_{k|k-1}^{j, A_{\text{merge}}^{(t)}} \left( \mathcal{L}_{k|k-1}^{j, A^{(t)}} \right)^{-1} \right], \quad (3.53a)$$

$$P\{\text{split}\} = \min \left[ 1, \mathcal{L}_{k|k-1}^{j, A_{\text{split}}^{(t)}} \left( \mathcal{L}_{k|k-1}^{j, A^{(t)}} \right)^{-1} \right]. \quad (3.53b)$$

The acceptance probabilities (3.53) give a split/merge MH algorithm that is greedy: if the likelihood of the association resulting from the action is larger than the likelihood of the current association, the action is performed with probability one.

In [38] it is noted that improved performance can be obtained if the merge/split MH algorithm is followed by refinement using Gibbs sampling. The combination of split/merge MH and Gibbs has also been implemented and evaluated for extended object tracking.

---

### 3.5.3 Stochastic optimisation

In this section, we present a novel approach to handling the extended object data association problem, that is inspired by the previous work [35, 38]. The presented method is called stochastic optimisation (SO) because it takes stochastic steps towards maximising the likelihood (3.28a). This sampling method combines the Gibbs actions and the merge/split actions and considers them all simultaneously. For the split actions, there is one important difference compared to the MH algorithm: here, given a cell split, we consider all possible ways to assign the two subcells. In the MH, only the split with highest likelihood was considered. The main intuition behind mixing the Gibbs actions, the merge action, and all possible split actions, is that one should obtain an algorithm that requires fewer iterations, because in each iteration there is a broader variety of different actions to choose from. Indeed, results from a simulation study confirms that this is the case.

The downside to mixing the Gibbs actions and the split/merge actions is that the resulting algorithm does not satisfy all the theoretical properties of an MCMC method. This means that we cannot guarantee that in the limit we are sampling from the distribution of interest  $P\{A\}$ . However, in the context of multiple extended object tracking, we are interested in finding probable data associations, not in generating samples from the posterior distribution.

In the SO algorithm, the possible actions affect either the selected measurement index, or all measurement indices with the same assignment. The selected measurement index can be moved to another cell, either an existing or a new cell. Alternatively, the selected cell can be merged with an existing cell, or split into two parts. The updates to the assignment vector are as presented in Table 3.5 and Table 3.6. More details are found in [26]; due to page length constraints they could not be included.

Given an assignment  $\varphi^{(t)}$ , the probability for the next assignment resulting from an action  $\alpha$  is given by the relative likelihood (cf. (3.27)),

$$P\{\varphi^{(t+1)} = \varphi_\alpha^{(t)} \mid \varphi^{(t)}\} = \frac{\mathcal{L}_{k|k-1}^{j, A_\alpha^{(t)}}}{\sum_{\alpha'} \mathcal{L}_{k|k-1}^{j, A(\varphi_{\alpha'}^{(t)})}} \quad (3.54)$$

where  $A_\alpha^{(t)} = A(\varphi_\alpha^{(t)})$  is the equivalent association that follows from applying action  $\alpha$  to assignment  $\varphi^{(t)}$ , and  $\sum_{\alpha} P\{\varphi^{(t+1)} = \varphi_\alpha^{(t)} \mid \varphi^{(t)}\} = 1$ . As noted above, the algorithm is not a proper MCMC algorithm. However, by using (3.54), data associations with high posterior probabilities will be sampled more often, which is a key property for our purposes. We are explicitly looking for data associations with high posterior probabilities, i.e., large  $\mathcal{L}_{k|k-1}^{j, A_\alpha^{(t)}}$ .

---

## 3.5.4 Implementational aspects

### 3.5.4.1 Reducing the cost of the likelihood computation

As shown in (3.30), each likelihood  $\mathcal{L}_{k|k-1}^{j,A_\alpha^{(t)}}$  is proportional to a product over the cells resulting from the association. However, for each action  $\alpha$ , only a subset of the cells given by  $A$  will be changed. This fact can be utilized to simplify the computations of the probabilities (3.54). For example, given the assignment  $\varphi^{(t)}$  with equivalent association  $A_\alpha^{(t)} = \alpha(\varphi_\alpha^{(t)})$ , and an action  $\alpha$  that affects cells  $c_1$  and  $c_2$ , we obtain a new assignment  $\varphi_\alpha^{(t)}$  with equivalent association  $A_\alpha^{(t)}$  and likelihood

$$\mathcal{L}_{k|k-1}^{j,A_\alpha^{(t)}} = \mathcal{L}_{k|k-1}^{j,A^{(t)}} \mathcal{L}_{C_{c_1}}^{A_\alpha^{(t)}} \left( \mathcal{L}_{C_{c_1}}^{A^{(t)}} \right)^{-1} \mathcal{L}_{C_{c_2}}^{A_\alpha^{(t)}} \left( \mathcal{L}_{C_{c_2}}^{A^{(t)}} \right)^{-1} \quad (3.55)$$

For the general case, we express this as  $\mathcal{L}_{k|k-1}^{j,A_\alpha^{(t)}} = \mathcal{L}_{k|k-1}^{j,A^{(t)}} \mathcal{L}_\alpha^{(t)}$ , where  $\mathcal{L}_\alpha^{(t)}$  represents the product of the likelihoods of the cells that are affected by the action. It is straightforward to calculate what  $\mathcal{L}_\alpha^{(t)}$  is for the different types of actions. Inserting this into (3.54) and simplifying gives

$$P \left\{ \varphi^{(t+1)} = \varphi_\alpha^{(t)} \mid \varphi^{(t)} \right\} = \frac{\mathcal{L}_\alpha^{(t)}}{\sum_{\alpha'} \mathcal{L}_{\alpha'}^{(t)}} \quad (3.56)$$

which is cheaper to compute than (3.54) is. Naturally, this can also be used to lower the computational cost of computing the Gibbs conditional distribution (3.47) and the split/merge MH acceptance probabilities (3.53).

### 3.5.4.2 Initialisation

The sampling algorithms above—Gibbs sampling, split/merge MH, and SO—can in theory be initialised with any valid partition; given enough iterations the sampling algorithm should yield good results. Two simple initialisations are starting with a partition with all the measurement in individual cells, or starting with a partition with all measurements in a single cell. However, empirically we have found that the different algorithms show varying sensitivity to the initialisation.

- Gibbs sampling requires an extremely large number of iterations if all measurements are initialised in a single cell. In general, Gibbs sampling can fail to move a group of measurements from one cell to another (this corresponds to the cell split in MH and SO). The reason for this is that the Gibbs sampling has to go via intermediate associations that have relatively low likelihood.
- Split/merge MH has shown poor performance when all measurements are initialised in individual cells. In this case several cell merges are required in the first iterations. However, before larger cells can be formed, the algorithm must pass through intermediate associations with relatively low likelihood.

In comparison, the SO algorithm has not shown any sensitivity to the initialisation. We attribute, at least in part, this robustness to the initialisation to the fact that the SO algorithm in each iteration considers more actions, compared to what Gibbs and MH do.

By using the predicted PMBM parameters, it is possible to find an initialisation that

---

**Table 3.7:** Simulation scenario parameters

Scen	$\lambda_c$	$p_D$	$p_S$	$N_I^G$	$N_I^{MH}$	$N_I^{MHG}$	$N_I^{SO}$
1	60	0.90	0.99	10	10	5 + 3	3
2	20	0.90	0.99	20	40	20 + 10	10
3	10	0.98	0.99	20	50	20 + 10	20

allows the sampling algorithm to run in fewer iterations. The assignment variables are initialised as

$$\varphi_m^{(0)} = \begin{cases} \hat{i}_m & \text{if } \hat{i}_m \in \mathbb{I}^j \\ m + |\mathbb{I}^j| & \text{otherwise} \end{cases}, \quad (3.57a)$$

$$\hat{i}_m = \arg \max_{i \in \mathbb{I}^j \cup \{0\}} \ell_{j,i}^{m,\text{init}}, \quad (3.57b)$$

$$\ell_{j,0}^{m,\text{init}} = \lambda c(\mathbf{z}^m) + p_D \langle D^u; \ell_{\{\mathbf{z}^m\}} \rangle, \quad (3.57c)$$

$$\ell_{j,i}^{m,\text{init}} = p_D r^{j,i} \gamma(\hat{\mathbf{x}}^{j,i}) \phi(\mathbf{z}^m | \hat{\mathbf{x}}^{j,i}), \quad (3.57d)$$

where  $\hat{\mathbf{x}}^{j,i} = E_{f_{j,i}}[\mathbf{x}]$ . In other words, if the maximum PPP likelihood corresponds to a object, then the measurement is associated to that object. Otherwise the measurements initial association is a background object.

### 3.5.4.3 Finding a subset of associations

After the algorithm is run for  $T$  iterations, a set of highly probable unique associations is obtained by sorting the unique equivalent associations  $A^{(t)}$  in order of descending likelihood  $\mathcal{L}_{k|k-1}^{A^{(t)}}$ , and then taking the associations whose cumulative sum of the relative likelihood exceeds a threshold. In the simulation and experiment presented below, the threshold was set to 99.99%.

## 3.6 Simulation results

Four different sampling method were compared in a simulation study: Gibbs sampling (G), split/merge Metropolis-Hastings (MH), split/merge Metropolis-Hastings with Gibbs refinement (MHG), and stochastic optimisation (SO). The sampling algorithms were integrated into the GGIW-PMBM filter [25, 26], and tested in three scenarios:

1. The first scenario contains 27 objects, and it tests the tracking filters capabilities of handling objects that appear/disappear in dense clutter.
2. The second scenario has five objects and tests the tracking filters capabilities of handling dense birth, i.e., when several objects appear at the same time step at very short distance from each other. Close objects has been shown in previous work, see, e.g., [22, 23], to lead to challenging data association, because the resulting measurements form a single cluster, rather than several distinct clusters. Dense birth is especially difficult, because there is no predicted information that can be utilised.

---

**Table 3.8:** Average time [s] to run one full Monte Carlo simulation

Scen	G	MH	MHG	SO
1	62.7	69.0	57.6	63.7
2	119.5	44.5	45.8	56.6
3	48.4	98.1	88.6	147.3

3. The third scenario has two objects that manoeuvre while staying next to each other. Here, in addition to the challenge of close objects, the non-linear manoeuvre adds additional complexity because the object motion model is linear, i.e., there is a model mismatch in the prediction step.

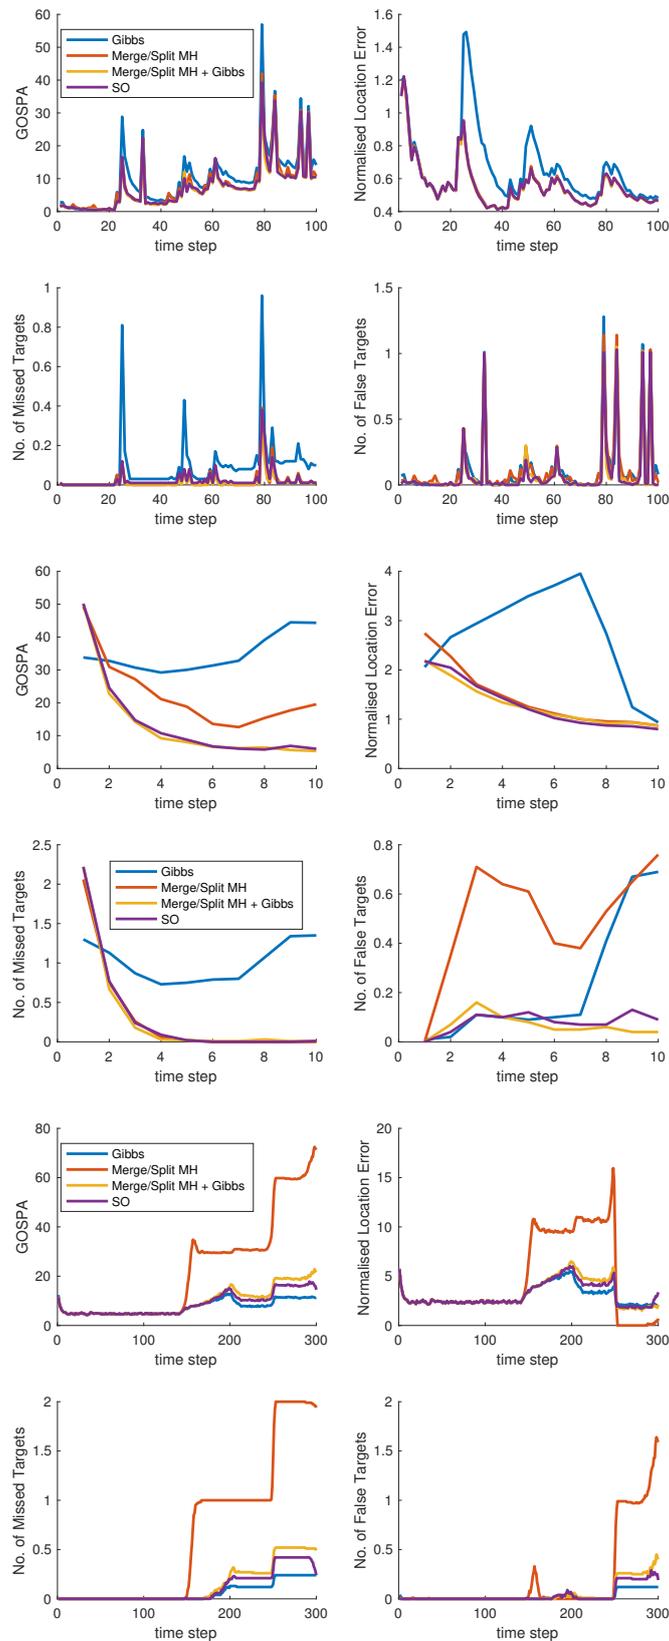
The scenario parameters clutter rate  $\lambda_c$ , detection probability  $p_D$ , and survival probability  $p_S$ , as well as the iteration number scaling factors  $N_I$  used for the different sampling methods, are summarized in Table 3.7. The scaling factors  $N_I$  were tuned such that the resulting tracking algorithms would yield good tracking performance at a reasonable computational cost, as measured by the average cycle time.

For each scenario, 100 Monte Carlo runs were simulated. The tracking performance is evaluated using the Gaussian Wasserstein distance (GWD) [47] integrated into the GOSPA metric [15]. The results are shown in Figure 3.8. The results show that in general G and MH lead to worst tracking performance, and require a larger number of iterations than the other two methods. In general, MHG and SO give approximately equal performance, however, SO achieves this in fewer iterations. The simulation study shows that both Gibbs actions and merge/split MH actions are necessary for finding a good set of highly probable data associations in as few iterations as possible.

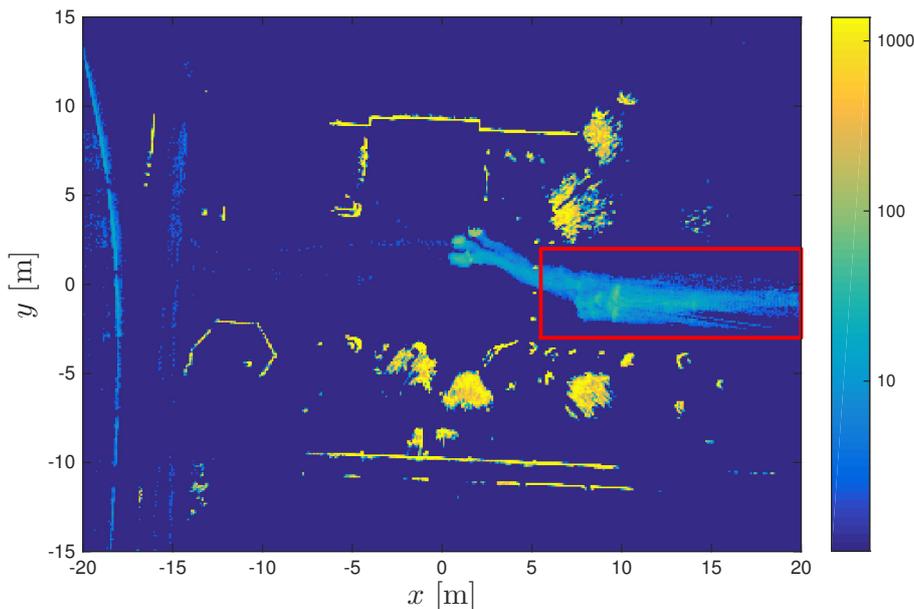
Note that, in the implementations that we have, fewer iterations did not automatically translate to lower average cycle time, because the cycle time is greatly affected by how many actions we have to compute the likelihood of the resulting association for. The average times to run a single Monte Carlo simulation of the three scenarios are given in Table 3.8. The average cycle times can be lowered by utilising parallelisation, as well as storing information to avoid computing the same things twice.

## 3.7 Experimental results

The data for the experimental results is recorded with the test vehicle of Ulm University. For the evaluation, the data of four Velodyne Puck 16 is used. The sensors are mounted on the roof, one in the front, one over each of the B-pillars and one at the back. The raw measurements are projected to the ground plane and discretized in the  $x$ - $y$  plane using grid cells of 0.1 m in order to reduce the total number of measurements. The data for the evaluation was recorded with the car parked on the campus of Ulm University. A histogram of the discretized measurements over



**Figure 3.8:** Results from simulation of three scenarios, from top to bottom: 1) 27 objects, 2) dense birth, and 3) parallel manoeuvre.



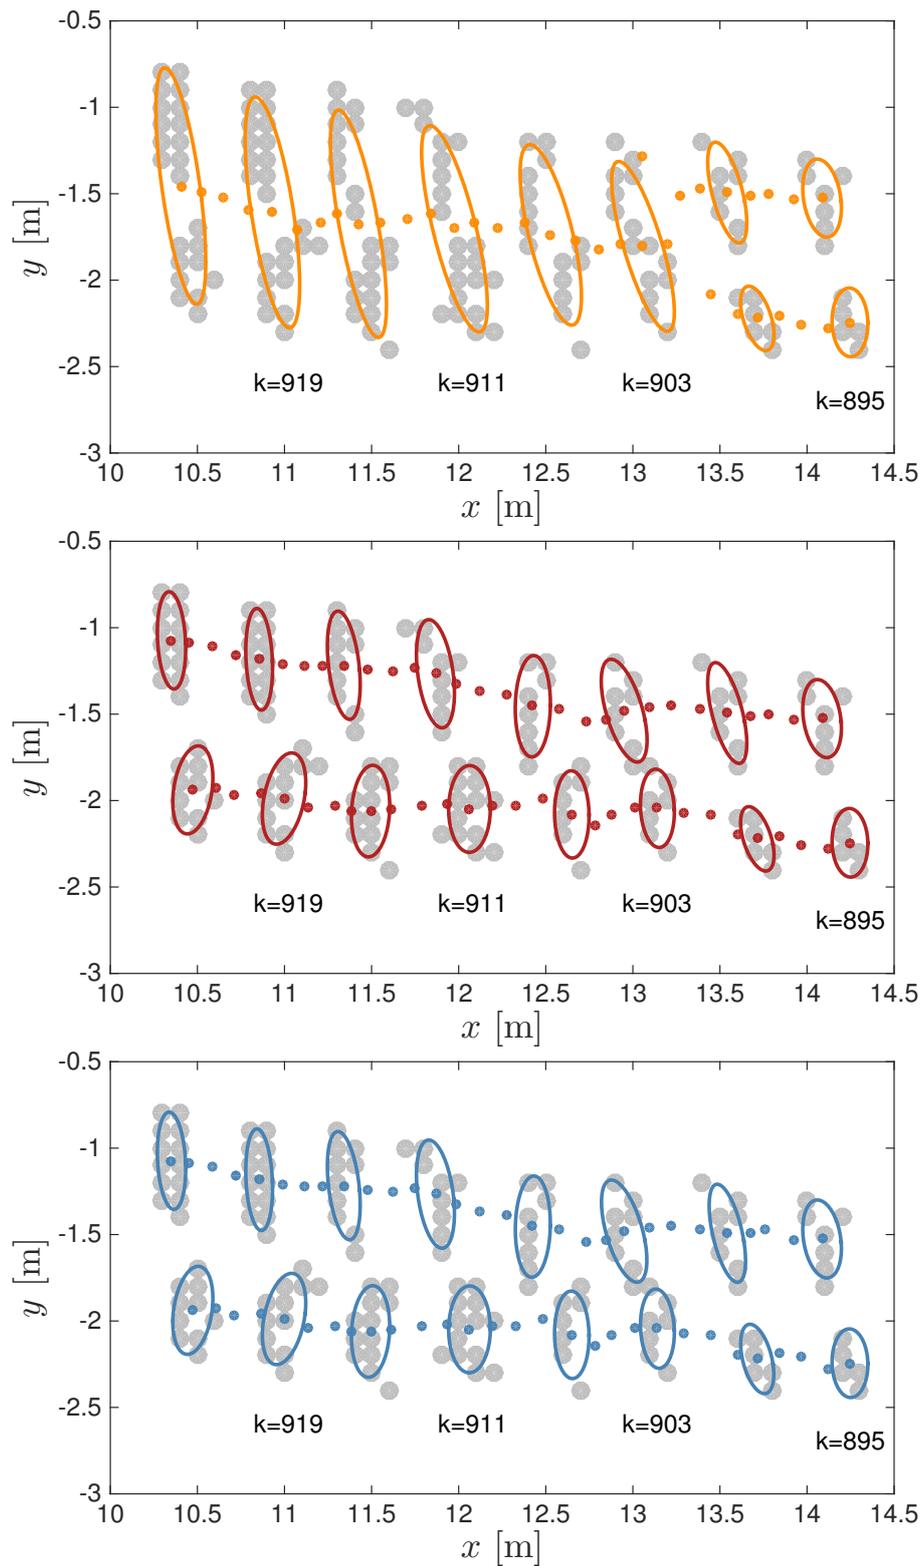
**Figure 3.9:** Visualisation of Velodyne data. The 2D world was gridded, and for each grid the number of times a detection fell inside was recorded. The logarithmic colorscale corresponds to the number of times the grid cell was detected. Stationary structures can be seen as yellow, the traces left by moving objects can be seen as green/light-blue. The data that was used to evaluate pedestrian tracking was taken from the area indicated by the red rectangle.

10 minutes is visualized in Figure 3.9.

We have compared three different variants of the GGIW-PMBM extended object tracking filter. Due to page length constraints the GGIW-PMBM filter details cannot be given; the reader is referred to [25, 26] for details. All three filters are identical except for the different approximations used to handle the data association in the update: PMBM-DP uses distance partition (DP); PMBM-SP/EMP uses DP, subpartition (SP) and expectation maximisation partition (EMP); and PMBM-SO uses stochastic optimisation (SO).

For most parts of the data the three compared tracking filters have identical performance; this is expected when the objects are well separated because all three filters are based on the same single object models. Therefore we highlight results from parts of the data that proved to be challenging due to spatially close objects. In this data the pedestrians move such that their respective measurements change between forming one measurement cluster per pedestrian, and forming joint measurement clusters (i.e., two or more pedestrians per cluster).

Figure 3.10 shows two pedestrians moving in the same direction. As soon as the objects become close, PMBM-DP merges the objects, while PMBM-SP/EMP and PMBM-SO can maintain both objects. Note how PMBM-DP maintains a single object, even after the data clearly forms two clusters; because the separation happens



**Figure 3.10:** Measurements (gray), PMBM-DP (orange), PMBM-SP/EMP (red), and PMBM-SO (blue). The object position is shown for each time step. For clarity, measurements and ellipses are shown for a subset of time steps. As indicated by the time steps ( $k$ ), the three objects move left to right.

---

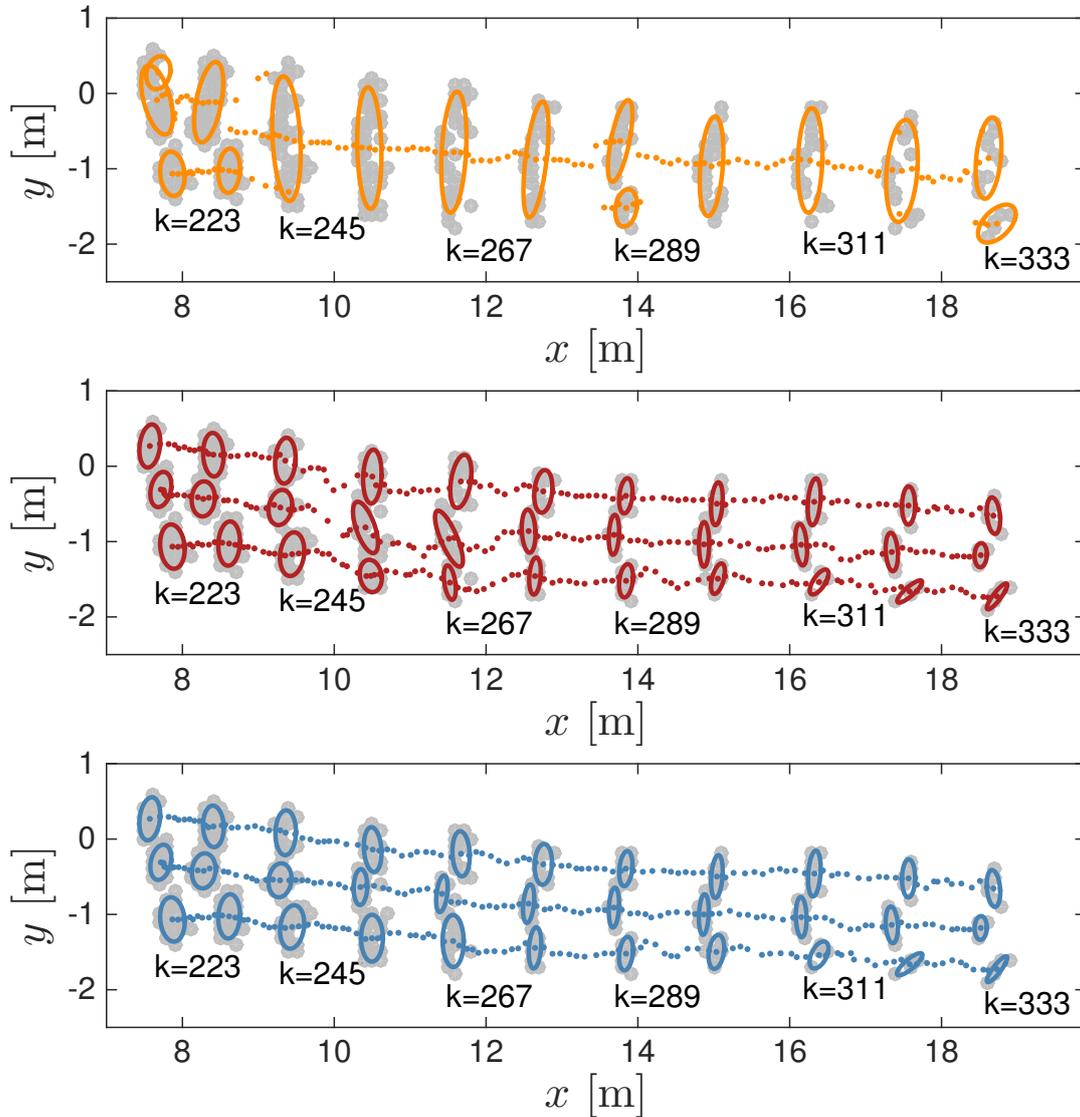
slowly, the object estimate is able to slowly adapt. A similar scenario, but with three objects, is shown in Figure 3.11. In this case PMBM-DP merges the close objects again. PMBM-SP/EMP can maintain three objects, however, the estimated tracks are not as smooth as the results from PMBM-SO. In these two scenarios, because the objects are moving in the same direction and with the same speed, the merging can be seen as a benign error—the tracked object represents a group of pedestrians. However, merging of objects moving in different direction is a worse problem, as it causes estimates that are significantly erroneous.

Figures 3.12 and 3.13 show scenarios in which objects pass each other at close distance, while moving in opposite directions. For increased clarity, we have plotted each time step on its own, rather than all time steps in one plot as in Figures 3.10 and 3.11. Velocity vectors indicate the direction the object is heading; for increased clarity the length of the vectors are tripled in the plots. Note that in the first figures, time steps  $k = 1111$  and  $k = 748$ , because the objects have been well separated all three filters have the same estimates. When the objects become close, PMBM-SO manages to keep track of all three objects, while both PMBM-DP and PMBM-SP/EMP produce significantly worse estimates of the objects. As the objects separate again, PMBM-DP and PMBM-SP/EMP can recover all three objects, however, the estimated extents are too large compared to the results from PMBM-SO, see time steps  $k = 1122$  and  $k = 757$ . Note also that PMBM-SO produces better estimates of the velocity vector than the other two filters do.

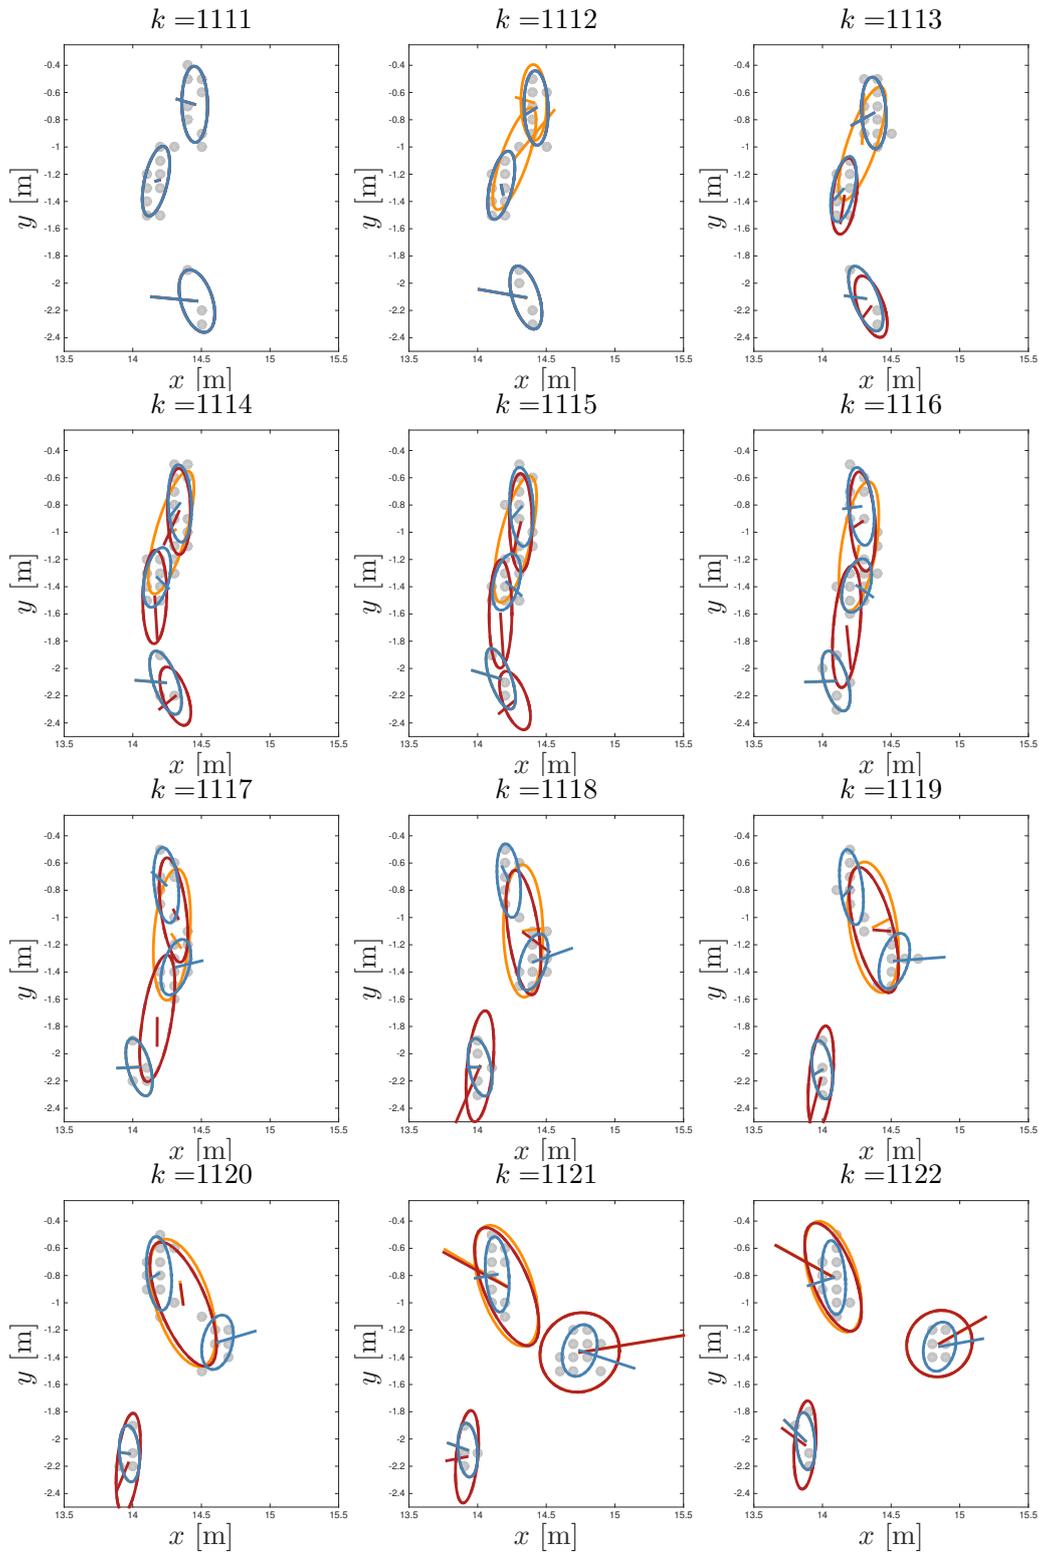
### 3.8 Concluding remarks

In this paper, we show that it is possible to solve the data association problem in extended object tracking in a single, likelihood-based, step. This is achieved using sampling algorithms, which are shown to outperform previous work that has been based on the combination of clustering algorithms and optimal assignment algorithms. An important topic for future work is investigating how the sampling algorithms can be implemented as efficiently as possible, for example by utilising parallelisation.

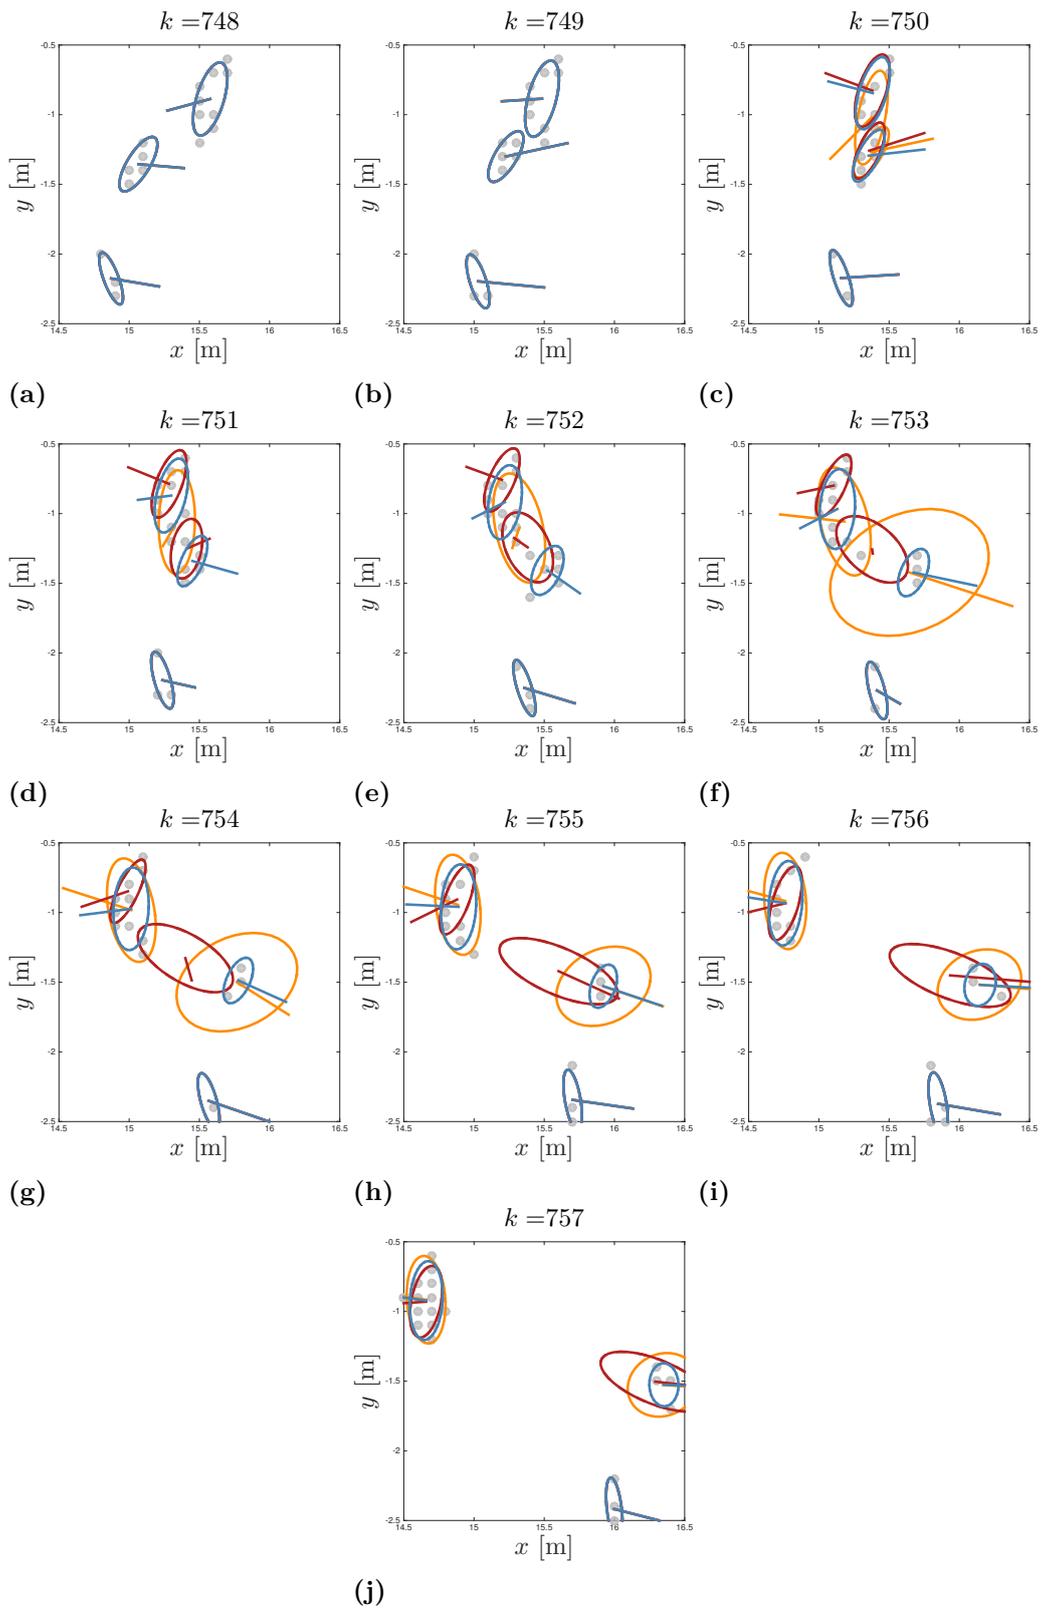
The sampling methods were evaluated in both simulations and in experiments where Velodyne data was used to track pedestrians. The experiments clearly showed that improved tracking performance is obtained when the sampling methods are used, especially when the pedestrians pass close to each other when moving in opposite directions. This result is important in an autonomous driving context: for safe and robust operation, the decision making requires not only accurate position estimates for all surrounding moving objects, but also accurate estimates of their motion parameters, such as velocity.



**Figure 3.11:** Measurements (gray), PMBM-DP (orange), PMBM-SP/EMP (red), and PMBM-SO (blue). The object position is shown for each time step. For clarity, measurements and ellipses are shown for a subset of time steps. As indicated by the time steps ( $k$ ), the three objects move left to right.



**Figure 3.12:** Measurements (gray), PMBM-DP (orange), PMBM-SP/EMP (red), and PMBM-SO (blue).



**Figure 3.13:** Measurements (gray), PMBM-DP (orange), PMBM-SP/EMP (red), and PMBM-SO (blue).



# Paper III: Poisson multi-Bernoulli filter for extended object estimation

## Abstract

In this paper, a Poisson multi-Bernoulli (PMB) filter for multiple extended targets estimation is presented. The PMB filter is based on the Poisson multi-Bernoulli mixture (PMBM) conjugate prior and approximates the multi-Bernoulli mixture (MBM) in the posterior as a single multi-Bernoulli. Different methods to merge the MBM are presented, along with their gamma Gaussian inverse Wishart implementation. The performance of the PMB filter is compared to the PMBM filter and the Probability Hypothesis Density filter in different simulated scenarios.

Keywords: Multiple target tracking, extended target, random matrix model, random finite sets, Bayesian estimation, variational inference

## 4.1 Introduction

Multiple target tracking (MTT) denotes the process of estimating the set of target trajectories based on a sequence of noise-corrupted measurements, including missed detections and false alarms. Traditionally, MTT algorithms have been tailored to the “point target” assumption: each target is modelled as a point without spatial extent, and that each target gives rise to at most one measurement per time scan. However, the high-resolution modern radar and lidar sensors makes the “point target” assumption unrealistic, since it is common that a target may occupy several sensor resolution cells. The tracking of such a target leads to the so-called extended target tracking problem, and the objective is to recursively determine the extent and kinematic parameters of the target over time. In this paper, we focus on the estimation of the current set of targets, which refers to multiple target filtering.

In extended target tracking, a non-standard measurement model is needed to model the number and the spatial distribution of generated measurements for each target. A common choice for modelling the number of measurements is the inhomogeneous Poisson Point Process (PPP), proposed in [16]. As for the modelling of the spatial

---

distribution, two popular models are the Random Hyper-surface Models [17] and the Gaussian inverse Wishart (GIW) approach [18, 19]. The former is designed for general star-convex shape; the latter relies on the elliptic shape and it models the spatial distribution of target-generated measurements as Gaussian with unknown mean and covariance. The Gamma GIW (GGIW) model [20, 21] is an extension of the GIW model that incorporates the estimation of target measurement rates.

Solving the MTT problem is complicated by the unknown correspondence between targets and measurements, known as data association. Because each target can generate multiple measurements per time scan, the problem of data association is even more challenging in multiple extended target tracking, compared to multiple point target tracking. Many of the existing multiple extended target tracking algorithms are based on random finite sets (RFS) [3], with the distribution of target-generated measurements modelled as GIW or GGIW, e.g., the Probability Hypothesis Density (PHD) filter [22, 23]; the Cardinalised PHD (CPHD) filter [21]; the  $\delta$ -Generalised Labelled Multi-Bernoulli filter and its approximation the Labelled Multi-Bernoulli filter ( $\delta$ -GLMB and LMB) [24]; as well as the Poisson Multi-Bernoulli Mixture (PMBM) filter [25, 26]. The PHD and CPHD filters are based on moment approximations of posterior densities, while the  $\delta$ -GLMB filter and the PMBM filter are based on conjugate priors that can provide accurate approximations to the exact posterior densities.

Due to the unknown number of data associations, the number of MB components in the posterior density of filters based on conjugate priors grows rapidly as more data is observed. Different methods to keep the number of MB components at a tractable level have been discussed in [24, 25, 26, 27]. The extended target LMB filter is an efficient approximation of the extended target  $\delta$ -GLMB filter, in which the MBM is approximated as a single MB [24]. To our best knowledge, this type of approximation, which simplifies the MBM in the posterior as a single MB, has not been presented for the extended target PMBM filter yet. Thus, it would be interesting to seek an approximation of the extended target PMBM filter, such that the approximated posterior density is a Poisson MB (PMB).

An algorithm based on variational Bayesian approximation was presented in [14]. It is designed for point target tracking, and finds the best-fitting PMB by approximating the true posterior MBM with the MB that minimises the Kullback-Leibler (KL) divergence. A performance evaluation of filters based on MB conjugate prior for point target estimation given in [48] has shown that the PMB filter has the best overall performance regarding estimation error and computational time, but it is not yet clear how the variational MB algorithm [14] can be used on extended target tracking.

The main contribution of this paper is presenting a PMB filter for multiple extended target estimation, along with its GGIW implementation. Two different implementations of the variational MB algorithm for merging MBM describing already detected targets are studied, one is based on the efficient approximation of feasible set pro-

---

**Table 4.9:** Notations

- 
- $\langle a, b \rangle = \int a(x)b(x)dx$ : inner product of  $a(x)$  and  $b(x)$ .
  - $|V|$ : determinant of matrix  $V$ .
  - $|\mathbf{X}|$ : cardinality of set  $\mathbf{X}$ .
  - $\Pi_N$ : set of permutation functions on  $I_N \triangleq \{1, \dots, N\}$

$$\Pi_N = \{\pi : I_N \rightarrow I_N | i \neq j \Rightarrow \pi(i) \neq \pi(j)\}.$$

- $\uplus$ : disjoint set union, i.e.,  $Y \uplus U = X$  means that  $Y \cup U = X$  and  $Y \cap U = \emptyset$ .
  - $D_{\text{KL}}(p||q) = \int p(x) \log \left( \frac{p(x)}{q(x)} \right) dx$ : KL divergence between  $p(x)$  and  $q(x)$ .
  - $\Gamma_d(\cdot)$ : multivariate gamma function.
  - $\varphi_0(\cdot)$ : digamma function.
  - $\text{Tr}(X)$ : trace of matrix  $X$ .
  - $\mathbf{I}_m$ : identity matrix of size  $m \times m$ .
- 

posed in [14], and the other is based on the optimal assignment following a similar optimisation procedure used by the set joint probabilistic data association (SJPDA) filter [34]. In addition, we propose a method to merge the MBM describing newly detected target in a reasonable way.

The paper is organised as follows. Background on multiple extended target tracking is given in Section II. The PMB filter and its GGIW implementation are presented in Section III and IV. Simulation results are presented in Section V.

## 4.2 Background

In this section, we first give some background on Bayesian filtering and RFS modelling. Next, the standard target transition model and extended target measurement model used in this work are outlined. At last, we review the PMBM conjugate prior and data association.

### 4.2.1 Bayesian multi-object filtering

In RFS-based MTT methods, target states and measurements are represented in the form of finite sets. Let  $\mathbf{x}_k$  denote the single target state at discrete time step  $k$ , and let  $\mathbf{X}_k$  denote the target set. The target set cardinality  $|\mathbf{X}_k|$  is a time-varying discrete random variable, and each target state  $\mathbf{x}_k \in \mathbf{X}_k$  is also a random variable. The set of measurements obtained at time step  $k$  is denoted as  $\mathbf{Z}_k$ , including clutter and target-generated measurements with unknown origin. The sequence of all the measurement sets received so far up to time step  $k$  is denoted as  $\mathbf{Z}^k$ . More notations are given in Table 4.9.

---

The objective of multiple target filtering is to estimate the set of targets  $\mathbf{X}_k$ , including the number of targets and individual target state, using the information contained in the measurement set sequence  $\mathbf{Z}^k$ . In extended target tracking, the target state models both kinematic properties and target extent. Let  $f_{k|k}(\mathbf{X}_k|\mathbf{Z}^k)$ ,  $f_{k,k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})$  and  $f_k(\mathbf{Z}_k|\mathbf{X}_k)$  denote the multi-target set density, the multi-target transition density and the multi-target measurement likelihood respectively. The multi-target Bayes filter propagates in time the multi-target set density  $f_{k-1|k-1}(\mathbf{X}_{k-1}|\mathbf{Z}^{k-1})$  using the Chapman-Kolmogorov prediction

$$f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}^{k-1}) = \int f_{k,k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})f_{k-1|k-1}(\mathbf{X}_{k-1}|\mathbf{Z}^{k-1})\delta\mathbf{X}_{k-1}, \quad (4.58)$$

and then updates the density using the Bayes update

$$f_{k|k}(\mathbf{X}_k|\mathbf{Z}^k) = \frac{f_k(\mathbf{Z}_k|\mathbf{X}_k)f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}^{k-1})}{\int f_k(\mathbf{Z}_k|\mathbf{X}_k)f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}^{k-1})\delta\mathbf{X}_k}, \quad (4.59)$$

where the set integral,  $\int f(\mathbf{X})\delta\mathbf{X}$ , is defined in [3, Sec 11.3.3].

## 4.2.2 Random set modelling

Two basic forms of RFS distribution used in this work are the PPP and the Bernoulli process. A PPP is an RFS whose cardinality is Poisson distributed, and each target is independent and identically distributed (i.i.d.). The PPP intensity  $D(\mathbf{x}) = \mu f(\mathbf{x})$  is determined by the scalar Poisson rate  $\mu$  and the spatial distribution  $f(\mathbf{x})$ . The PPP density is given by

$$f(\mathbf{X}) = e^{-\mu} \prod_{\mathbf{x} \in \mathbf{X}} \mu f(\mathbf{x}). \quad (4.60)$$

A Bernoulli process with probability of existence  $r$  and existence-conditioned probability density function (PDF)  $f(\mathbf{x})$  has RFS density

$$f(\mathbf{X}) = \begin{cases} 1 - r & \mathbf{X} = \emptyset \\ r \cdot f(\mathbf{x}) & \mathbf{X} = \{\mathbf{x}\} \\ 0 & \text{otherwise,} \end{cases} \quad (4.61)$$

where the cardinality  $|\mathbf{X}|$  is Bernoulli distributed with parameter  $r \in [0, 1]$ . The Bernoulli process offers a convenient way to capture both the uncertainty regarding target existence and state. In MTT, targets are typically assumed to be independent. Thus, multiple targets can be represented as a multi-Bernoulli RFS  $\mathbf{X}$ , which is a disjoint union of independent Bernoulli RFSs  $\mathbf{X}^i$ , i.e.,  $\mathbf{X} = \uplus_{i \in \mathbb{I}} \mathbf{X}^i$ , where  $\mathbb{I}$  is an index set. The RFS density of an MB process can be represented as

$$f(\mathbf{X}) = \begin{cases} \sum_{\sum_{i \in \mathbb{I}} \mathbf{X}^i = \mathbf{X}} \prod_{i \in \mathbb{I}} f^i(\mathbf{X}^i) & |\mathbf{X}| \leq |\mathbb{I}| \\ 0 & |\mathbf{X}| > |\mathbb{I}|. \end{cases} \quad (4.62)$$

The MB distribution can be defined entirely by the parameters  $\{r^i, f^i(\cdot)\}_{i \in \mathbb{I}}$  of the involved Bernoulli RFSs. Lastly, an MBM density is an RFS density that is a normalised, weighted sum of MB densities. In MTT, the weights typically correspond

---

to the probability of different data association sequences. An MBM is defined entirely by the parameters  $\{(\mathcal{W}^j, \{r^{j,i}, f^{j,i}(\cdot)\}_{i \in \mathbb{I}^j})\}_{j \in \mathbb{J}}$ , where  $\mathbb{I}^j$  is the index set of Bernoulli components in the  $j$ th MB,  $r^{j,i}$  and  $f^{j,i}(\cdot)$  are the existence probability and existence-conditioned PDF of the  $i$ th Bernoulli process in the  $j$ th MB, and  $\mathcal{W}^j$  is the weight of the  $j$ th MB.

### 4.2.3 Standard extended target measurement model

The set of measurements  $\mathbf{Z}_k$  is a union of a set of clutter measurements and sets of target-generated measurements. The clutter is modelled as a PPP with Poisson rate  $\lambda$  and spatial distribution  $c(\mathbf{z})$ , independent of targets and any target-generated measurements. Each extended target may give rise to multiple measurements with a state dependent detection probability  $p^D(\mathbf{x})$ . If the extended target is detected, the target-generated measurements are modelled as a PPP with Poisson rate  $\gamma(\mathbf{x})$  and spatial distribution  $\phi(\mathbf{z}|\mathbf{x})$ , independent of all other targets and their corresponding generated measurements. The extended target set measurement likelihood for a nonempty set of measurements  $\mathbf{Z}$  is the product of target detection probability and the PPP density of target-generated measurements [26]

$$\ell_{\mathbf{Z}}(\mathbf{x}) = p^D(\mathbf{x})e^{-\gamma(\mathbf{x})} \prod_{\mathbf{z} \in \mathbf{Z}} \gamma(\mathbf{x})\phi(\mathbf{z}|\mathbf{x}). \quad (4.63)$$

For an extended target state  $\mathbf{x}$ , the effective detection probability is the product of target detection probability and the probability that target generates at least one measurement, which is  $1 - e^{-\gamma(\mathbf{x})}$ . Then the effective probability of missed detection can be calculated accordingly as

$$q^D(\mathbf{x}) = 1 - p^D(\mathbf{x}) + p^D(\mathbf{x})e^{-\gamma(\mathbf{x})}. \quad (4.64)$$

This is also the measurement likelihood for an empty measurement set,  $\ell_{\emptyset}(\mathbf{x})$ .

### 4.2.4 Standard target transition model

In this work, we focus on the update, hence the prediction and the transition model is only briefly discussed. In the standard target transition model, it is assumed that targets arrive according to a PPP, independently of any pre-existing targets. At each time step, a single target remains with a probability of survival  $p_k^S(\mathbf{x})$ . The targets depart according to i.i.d. Markov processes with probability  $1 - p_k^S(\mathbf{x})$ . The target state at next time step only depends on its current state. Targets evolve independently according to i.i.d. Markov process with transition density  $f_{k+1,k}(\mathbf{x}_{k+1}|\mathbf{x}_k)$ .

### 4.2.5 PMBM conjugate prior and data association

The PMBM conjugate prior for multiple extended target filtering was developed in [25, 26]. In the PMBM filter, the target set is a union of two disjoint sets of undetected targets  $\mathbf{X}^u$  and detected objects  $\mathbf{X}^d$ , i.e.,  $\mathbf{X} = \mathbf{X}^u \uplus \mathbf{X}^d$ . The distribution of targets that have not been detected yet  $\mathbf{X}^u$  is described by a PPP, while the distribution of targets that have been detected at least once  $\mathbf{X}^d$  is described by an

---

MBM, which is a normalised weighted sum of MB densities. The PMBM density is defined entirely by the parameters,

$$D^u, \{(\mathcal{W}^j, \{r^{j,i}, f^{j,i}(\cdot)\}_{i \in \mathbb{I}^j})\}_{j \in \mathbb{J}}, \quad (4.65)$$

where  $D^u(\cdot)$  is the PPP intensity of the set of undetected targets.

The PMBM form is preserved by prediction and update [12]. Each MB corresponds to a unique *global hypothesis* for the detected targets, i.e., a particular history of data associations for all detected targets. Given a predicted PMBM density, for each predicted global hypothesis there are multiple possible data associations, each of which will result in an MB component in the updated MBM. Let  $\mathbb{M}$  be the set of indices of measurement set  $\mathbf{Z}$  and let  $\mathcal{A}^j$  be the space of all data associations  $A$  for the  $j$ th predicted hypothesis. Formally a data association  $A \in \mathcal{A}^j$  consists of a partition of  $\mathbb{M} \cup \mathbb{I}^j$  into non-empty disjoint subsets called index cells, denoted  $C \in \mathcal{A}$  [26]. An index cell can contain at most one target index. If the index cell  $C$  contains a target index, then let  $i_C$  denote the corresponding target index. Further, let  $\mathbf{C}_C$  denote the measurement cell that corresponds to the index cell  $C$ .

Global hypotheses are made up of *single target hypotheses*, each of which can represent the event that target  $i_C$  is updated with  $\mathbf{C}_C$ , the event that target  $i_C$  is not detected and the event that it is the first detection of the background (clutter or new target), which means  $C \cap \mathbb{I}^j = \emptyset$ , via a Bernoulli process. A *track* is defined as a collection of single target hypotheses corresponding to the same target that was first detected. As a consequence, a new track is created for each measurement cell. If the measurement cell is associated to a pre-existing track, the new track would have existence probability equal to zero. A complexity analysis of the data association problem given in [26] has shown that the number of MB distributions grows very rapidly over time. Approximations of the data association problem have been presented in [26, 27], which can keep the number of global hypotheses at a tractable level. In this work, we focus on approximating the MBM as a single MB after solving the data association problem.

## 4.3 Poisson Multi-Bernoulli Filter

In the PMB filter, the PMB density is preserved in the prediction step, whereas in the update step the posterior first becomes a PMBM density due to data association, then it is approximated as a PMB afterwards. In this section, the prediction and update steps of the PMB filter are presented. We then discuss approximation methods that result in a PMB posterior. The time index is omitted for brevity.

### 4.3.1 PMB filter recursion

Given a posterior PMB density with parameters  $D^u, (r^i, f^i(\cdot))_{i \in \mathbb{I}}$  and the standard dynamical model, the predicted density is a PMB density with parameters

$$D_+^u, (r_+^i, f_+^i(\cdot))_{i \in \mathbb{I}}, \quad (4.66)$$

where

$$D_+^u(\mathbf{x}) = D_+^b(\mathbf{x}) + \langle D^u, p^S f_{k+1,k} \rangle, \quad (4.67a)$$

$$r_+^i = \langle f^i, p^S \rangle r^i, \quad (4.67b)$$

$$f_+^i(\mathbf{x}) = \frac{\langle f^i, p^S f_{k+1,k} \rangle}{\langle f^i, p^S \rangle}, \quad (4.67c)$$

and  $D^b(\cdot)$  denotes the PPP birth intensity. Detailed derivations can be found in [12].

Given a prior PMB density with parameters (4.66), a set of measurements  $\mathbf{Z}$ , and the standard measurement model, the updated density is a PMBM (4.65). The updated PPP intensity is  $D^u(\mathbf{x}) = q^D(\mathbf{x})D_+^u(\mathbf{x})$ . For tracks continuing from previous time steps, a hypothesis can be included as a missed detection or an update using a measurement cell  $\mathbf{C}_C$ . For missed detection hypotheses, the Bernoulli densities have parameters ( $i \in \mathbb{I}$ )

$$r^{j,i} = \frac{r_+^i \langle f_+^i, q^D \rangle}{1 - r_+^i + r_+^i \langle f_+^i, q^D \rangle}, \quad (4.68a)$$

$$f^{j,i}(\mathbf{x}) = \frac{q^D(\mathbf{x})f_+^i(\mathbf{x})}{\langle f_+^i, q^D \rangle}. \quad (4.68b)$$

For hypotheses updating pre-existing tracks, the Bernoulli densities have parameters ( $i \in \mathbb{I}$ )

$$r^{j,i} = 1, \quad (4.69a)$$

$$f^{j,i}(\mathbf{x}) = \frac{\ell_{\mathbf{C}_C}(\mathbf{x})f_+^i(\mathbf{x})}{\langle f_+^i, \ell_{\mathbf{C}_C} \rangle}. \quad (4.69b)$$

For new tracks with corresponding measurement cell associated to the background, the updated densities are Bernoulli densities with parameters ( $i \in \mathbb{I}^j \setminus \mathbb{I}$ )

$$r^{j,i} = \frac{\langle D_+^u, \ell_{\mathbf{C}_C} \rangle}{\lambda c(\mathbf{C}_C) + \langle D_+^u, \ell_{\mathbf{C}_C} \rangle}, \quad (4.70a)$$

$$f^{j,i} = \frac{\ell_{\mathbf{C}_C}(\mathbf{x})D_+^u(\mathbf{x})}{\langle D_+^u, \ell_{\mathbf{C}_C} \rangle}. \quad (4.70b)$$

After updating, the number of Bernoulli components representing detected targets in each MB remains unchanged, while the number of Bernoulli components representing newly detected targets becomes  $|\mathbb{I}^j \setminus \mathbb{I}|$  in the  $j$ th MB. Because how to solve the data association problem is out of the scope of this paper, mathematical expressions regarding association likelihood and global hypotheses weights are not given. The reader is referred to [26] for more details. In the MBM approximation, we seek to find the MB distribution  $\{\hat{r}^i, \hat{f}^i(\cdot)\}_{i \in \hat{\mathbb{I}}}$  that best matches the MBM  $\{(\mathcal{W}^j, \{r^{j,i}, f^{j,i}(\cdot)\}_{i \in \mathbb{I}^j})\}_{j \in \mathbb{J}}$ . In this work, the MBM approximations for pre-existing tracks and new tracks are considered separately.

---

### 4.3.2 Pre-existing track formation

In this subsection, the MBM describing new tracks is disregarded, and we denote the MBM  $f(\mathbf{X})$  describing pre-existing tracks as  $\{(\mathcal{W}^j, \{r^{j,i}, f^{j,i}(\cdot)\}_{i=1}^N})\}_{j \in \mathbb{J}}$ , where  $N$  is the number of pre-existing tracks. A variational method was presented in [14] to obtain the best-fitting MB  $g(\mathbf{X})$  that minimises the set KL divergence from the MBM distribution  $f(\mathbf{X})$ :

$$\arg \min_g \int f(\mathbf{X}) \log \frac{f(\mathbf{X})}{g(\mathbf{X})} = \arg \max_g \int f(\mathbf{X}) \log g(\mathbf{X}) d\mathbf{X}. \quad (4.71)$$

An approximate solution is based on minimising the upper bound of the true objective (4.71), following a similar process to expectation-maximisation [49]. The correspondence between the underlying Bernoulli distribution in  $f(\mathbf{X})$  and the Bernoulli component in the best-fitting distribution  $g(\mathbf{X})$  is treated as missing data  $q(\pi)$ . An approximated upper bound to the objective of (4.71) is given by [14]

$$D_{\text{UB}}(f(\mathbf{X})||g(\mathbf{X})) = - \sum_{j \in \mathbb{J}, \pi \in \Pi_N^j} \mathcal{W}^j q_j(\pi^j) \sum_{i=1}^N \int f^{j, \pi^j(i)}(\mathbf{X}) \log g^i(\mathbf{X}) \delta \mathbf{X}, \quad (4.72)$$

where  $\Pi_N^j$  is the set of all ways to assign the Bernoulli components in the  $j$ th MB  $f^j(\mathbf{X})$  to the Bernoulli components in  $g(\mathbf{X})$ ; the missing data  $q_j(\pi^j)$  is constrained to vary only with the  $j$ th MB, and it satisfies  $q_j(\pi^j) \geq 0$  and  $\sum_{\pi^j \in \Pi_N^j} q_j(\pi^j) = 1$ . The standard method for solving the form of (4.72) is by block coordinate descent, which alternates between minimisation with respect to  $g^i(\mathbf{X})$  (M-step) and  $q_j(\pi^j)$  (E-step) [14].

Because solving the above problem suffers from combinatorial complexity, approximation is needed to obtain tractable solution. In the remainder of this subsection, two different approaches to solving this minimisation problem are studied. The first one is based on finding the best assignment for each MB in the MBM; the second one is based on an approximation of the set of missing data distribution  $q_j(\pi^j)$ .

#### 4.3.2.1 Optimal assignment

The MB density is invariant to the indexing of the Bernoulli components, and therefore the selection of the assignment mapping  $\pi$  in each MB will not change the MBM process but only determine which Bernoulli components are going to be merged. The minimisation problem can be interpreted as the Bernoulli components in each MB are permuted in such a manner that the upper bound (4.72) is minimised but the density of the reordered  $f(\mathbf{X})$  remains unchanged, following a similar approach to KLSJDPA [34]. The similarities between the best-fitting MB and KLSJPDA were explored in [14]. Empirically, we found that finding a set of most likely assignments for each MB is computationally heavy. Hence, we choose to find the single optimal assignment  $\hat{\pi}^j$  for each MB  $f^j(\mathbf{X})$ . In this case, the missing data becomes  $q_j(\pi^j) = 1 \forall j$ , and the E-step can be expressed as

$$\hat{\pi}^j = \underset{\pi^j}{\operatorname{argmin}} - \sum_{i=1}^N \int f^{j, \pi^j(i)}(\mathbf{X}) \log g^i(\mathbf{X}) \delta \mathbf{X}, \quad (4.73)$$

where the optimal assignment  $\hat{\pi}^j$  can be obtained using methods such as the auction algorithm. The minimisation of (4.72) with respect to the approximated MB  $g(\mathbf{X})$  simplifies to

$$g^i(\mathbf{X}) = \sum_{j \in \mathbb{J}} \mathcal{W}^j f^{j, \hat{\pi}^j(i)}(\mathbf{X}). \quad (4.74)$$

#### 4.3.2.2 Efficient approximation of feasible set

It was revealed in [14] that the minimisation of the upper bound (4.72) can be solved approximately as

$$\operatorname{argmin}_{q(h,i) \in \mathcal{M}} - \sum_{i=1}^N \int \left( \sum_{h \in \mathcal{H}} q(h,i) f^h(\mathbf{X}) \right) \log g^i(\mathbf{X}) \delta \mathbf{X}, \quad (4.75)$$

where  $\mathcal{H}$  is the set of single target hypotheses,  $q(h,i)$  is a simplified representation of  $q_j(\pi^j)$ , which specifies the weight of single target hypothesis density  $f^h(\mathbf{X})$  in the new Bernoulli component  $g^i(\mathbf{X})$ , and the feasible set  $\mathcal{M}$  is an approximation needed for tractability

$$\mathcal{M} = \left\{ q(h,i) \geq 0 \mid \sum_{h \in \mathcal{H}} q(h,i) = 1 \forall i \in \{1, \dots, N\}, \sum_{i=1}^N q(h,i) = p_h \forall h \in \mathcal{H} \right\}. \quad (4.76)$$

The constraint  $p_h$  satisfies  $p_h = \sum_{i=1}^N p_i(h)$ , where

$$p_i(h) = \sum_{f^j(\mathbf{X}) | f^{j,i}(\mathbf{X}) = f^h(\mathbf{X})} \mathcal{W}^j. \quad (4.77)$$

The derivation of (4.75) can be found in [14]. Note that here the missing data distribution is no longer constrained to vary only with the global hypotheses. In this case, each approximated Bernoulli component can be expressed as the weighted sum of different single target hypothesis densities, and the M-step becomes

$$g^i(\mathbf{X}) = \sum_{h \in \mathcal{H}} q(h,i) f^h(\mathbf{X}), \quad (4.78)$$

while the E-step reverts to a LP:

$$\operatorname{argmin}_{q(h,i)} \sum_{h \in \mathcal{H}} \sum_{i=1}^N -q(h,i) \int f^h(\mathbf{X}) \log g^i(\mathbf{X}) \delta \mathbf{X}, \quad (4.79)$$

$$\begin{aligned} \text{subject to } & \sum_{i=1}^N q(h,i) = p_h \forall h, \\ & \sum_{h \in \mathcal{H}} q(h,i) = 1 \forall i, \\ & q(h,i) \geq 0 \forall h, i. \end{aligned}$$

Problem of this type can be solved using methods such as the simplex algorithm. The resulting algorithm can be initialised with  $q(h,i) = p_i(h) \forall h, i$ .

---

### 4.3.3 New track formation

A big barrier in the extended target PMB filter is the formation of new tracks. In point target tracking, a new track is created for each measurement. Given  $|\mathbf{Z}|$  measurements, there are  $|\mathbf{Z}|$  mutually independent possible new tracks. The new tracks can be formed out of the marginal association probabilities (cf. (4.77)) [12]. In extended target tracking, measurements are partitioned into measurement cells and a new track is created for each measurement cell. Given  $|\mathbf{Z}|$  measurements, there are  $2^{|\mathbf{Z}|} - 1$  possible ways in which we can form a subset of measurements. Each such subset can, in an association, be associated to a new target. This means that in extended target tracking, the updated MBM can at most contain  $2^{|\mathbf{Z}|} - 1$  different possible new targets. Let  $N_+^j$  denote the number of measurement cells in the  $j$ th updated global hypothesis. According to the definition of track, there are  $N_+^j$  new tracks created, but only  $N_+^j - N$  of them, which are associated to newly detected targets, have valid existence probability in the  $j$ th global hypothesis.

One strategy to form new tracks, which is similar to [12], is only to merge single target hypotheses across different global hypotheses that are created by the same measurement cell and treat each different single target hypothesis as a candidate for a potential new target. In other words, the existence probability of new track created by measurement cell  $\mathbf{C}$  is determined by summing over the data association hypotheses where  $\mathbf{C}$  is associated to the background. This approach is simple to implement; nevertheless, it has several drawbacks. First, a large number of new tracks with low existence probability can be created using this approach. Second, this approach approximates new tracks generated by different measurement cells as independent. But if some of the measurement cells have shared measurements, which is the typical case, there might be a high dependency between the new tracks that are created by these measurement cells. This is not consistent with the assumption that tracks are independent.

In this subsection, a method to form new tracks in the PMB filter is presented. The intuition behind this proposed method is that we want to merge highly correlated Bernoulli components across different MB distributions so that similar new tracks will not be formed in the same local region and new tracks with significant different Bernoulli densities will not be merged. Let  $\mathbf{C}^{j,i}$  denote the  $i$ th measurement cell that is associated with the background in the  $j$ th global hypothesis, and let  $I^{j,i}$  denote the label of the corresponding Bernoulli component which represents this created new track. The objective is to give the same label to Bernoulli components representing new tracks that are considered similar enough.

Bernoulli components across different global hypotheses may have the same label, but the Bernoulli components in the same global hypothesis must be assigned to different labels due to the independence assumption. For any pair of Bernoulli densities, the symmetrised KL divergence is used to measure the similarity, which is defined as

$$D_{\text{SKL}}(p||q) = D_{\text{KL}}(p||q) + D_{\text{KL}}(q||p). \quad (4.80)$$

---

Because reordering the MB distributions will not change the MBM process, we assume that the MB distributions are sorted in the descending order of their weights so that  $f^1(\mathbf{X})$  has the highest weight. Starting from  $I^{j,i} = 1$ , we calculate the symmetrised KL divergence of Bernoulli densities between the new track created by  $\mathbf{C}^{1,1}$  and the new tracks in other global hypotheses. For each  $f^j(\mathbf{X})$  with  $j \geq 2$ , pick the Bernoulli component which represents the most similar track to  $f^{1,N+1}(\mathbf{X})$  in  $f^j(\mathbf{X})$ . If the minimum symmetrised KL divergence is below a certain threshold  $\tau$ , let the label of the selected Bernoulli component be the same with  $I^{1,1}$ . This procedure is repeated for every Bernoulli density of new track in every  $f^j(\mathbf{X})$  until all the Bernoulli components are assigned a label. At last, the number of new tracks formed is equal to the number of unique labels, which is also the largest label index  $I_{max}$ .

New tracks can be formed by only merging Bernoulli components across different MB distributions with the same label. For the  $l$ th ( $l \in \{1, \dots, I_{max}\}$ ) new track in the approximated MB  $g(\mathbf{X})$ , its Bernoulli density can be written as

$$g^{N+l}(\mathbf{X}) = \sum_{j \in \mathbb{J}} \mathcal{W}^j \sum_{i=1}^{N_+^j - N} f^{j,N+i}(\mathbf{X}) \delta(I^{j,i}(\mathbf{X}) - l). \quad (4.81)$$

The pseudo code of the resulting algorithm is given in Table 4.10. Empirical result shows that the number of tracks formed using this approach can be kept to a relatively small number.

#### 4.3.4 Recycling

For the approximated MB distribution  $g(\mathbf{X})$ , the recycling method of [33] can be applied to Bernoulli components with low existence probability. The recycled components are approximated as being Poisson and are incorporated into the PPP representing undetected targets for generating possible new targets in subsequent steps.

Suppose that a Bernoulli component  $f(\mathbf{X})$  in the form of (2.2) is approximated as a PPP

$$f(\mathbf{X}) \approx \tilde{f}(\mathbf{X}) = e^{-u} \prod_{\mathbf{x} \in \mathbf{X}} D(\mathbf{x}). \quad (4.82)$$

As shown in [3, p. 359], it is optimal to set  $D(\mathbf{x}) = \mu f(\mathbf{x})$  and  $\mu = r$ . The value of the KL divergence at this optimal choice

$$D_{\text{KL}}(f(\mathbf{X}) || \tilde{f}(\mathbf{X})) = r + (1 - r) \log(1 - r), \quad (4.83)$$

is very small for existence probability  $r$  less than 0.1 [33]. Denote  $\mathbb{R}$  as the set of indices of Bernoulli components to be recycled. After recycling, the total PPP intensity of undetected targets can be expressed as

$$\hat{D}^u(\mathbf{x}) = D^u(\mathbf{x}) + \sum_{i \in \mathbb{R}} r^i f^i(\mathbf{x}). \quad (4.84)$$

---

**Table 4.10:** Pseudo Code of New Tracks Forming

---

```

Input:  $f^{j,i}(\mathbf{X}), \mathcal{W}^j, \tau$ 
Output:  $g(\mathbf{X})$ 
Sort  $f^j(\mathbf{X})$  in the descending order of  $\mathcal{W}^j$ ;
 $I \leftarrow 1$ ;
for all  $j = \{1, \dots, |\mathbb{J}|\}$  do
  for all  $i = \{1, \dots, N_+^j - N\}$  do
    if  $I^{j,i}$  is empty then
       $I^{j,i} \leftarrow I$ ;
       $I \leftarrow I + 1$ ;
    end if
  for all  $j^+ = \{j + 1, \dots, |\mathbb{J}|\}$  do
    for all  $i^+ = \{1, \dots, N_+^{j^+} - N\}$  do
      if  $I^{j^+,i^+}$  is empty then
         $d^{i^+} \leftarrow D_{\text{SKL}}(f^{j,N+i}(\mathbf{X}) || f^{j^+,N+i^+}(\mathbf{X}))$ ;
      end if
    end for
     $[d^*, i^*] \leftarrow \min(d)$ ;
    if  $d^* < \tau$  then
       $I^{j^+,i^*} \leftarrow I$ ;
    end if
  end for
end for
for all  $i = \{1, \dots, I\}$  do
  Calculate  $g^{N+i}(\mathbf{X})$  using (4.81).
end for

```

---

The benefits of recycling in the point target PMBM and PMB filters have been evaluated in [48]. An additional benefit of recycling in the extended target PMB filter is that we can ensure that only new tracks with relatively high existence probability are formed.

## 4.4 GGIW Implementation

In this section, some implementation details of the PMB filter are presented. A GGIW implementation of the PMBM filter has been given in [26]. To make comparison easy, we choose to use the random matrix model [18, 19], in which the target shape is approximated as an ellipse.

### 4.4.1 Single target models

In the random matrix model, it is assumed that the measurements are Gaussian distributed around the centroid of the target. The extended target state  $\mathbf{x}_k$  is the combination of a random Poisson rate  $\gamma_k$  modelling the average number of measurements generated by the target, a random vector  $\xi_k$  describing the target kinematic state, and a random covariance matrix  $\chi_k$  describing the target size and shape, i.e.,  $\mathbf{x}_k = \{\xi_{k+1}, \chi_{k+1}, \gamma_{k+1}\}$ .

---

The motions models are

$$\xi_{k+1} = F(\xi_k) + w_k, \quad (4.85a)$$

$$\chi_{k+1} = M(\chi_k)\chi_k M(\chi_k)^T, \quad (4.85b)$$

$$\gamma_{k+1} = \gamma_k, \quad (4.85c)$$

where  $F(\cdot)$  is a motion model,  $w_k$  is zero mean Gaussian noise and  $M(\cdot)$  is a transformation matrix. The measurement likelihood for a single measurement  $\mathbf{z}$  is

$$\phi(\mathbf{z}_k | \mathbf{x}_k) = \mathcal{N}(\mathbf{z}_k; H_k \xi_k, \chi_k), \quad (4.86)$$

where  $H_k$  is the measurement model. The single target conjugate prior for the PPP model (4.60) with single measurement likelihood (4.86) can be expressed as

$$f_{k|k}(\mathbf{x}_k) = \mathcal{GAM}(\gamma_k; a_{k|k}, b_{k|k}) \mathcal{N}(\xi_k; m_{k|k}, P_{k|k}) \mathcal{IW}_d(\chi_k; v_{k|k}, V_{k|k}) \triangleq \mathcal{GGIW}(\xi_k; \zeta_{k|k}), \quad (4.87)$$

where  $d$  is the dimension of the random matrix  $V_{k|k}$  and

$\zeta_{k|k} = \{a_{k|k}, b_{k|k}, m_{k|k}, P_{k|k}, v_{k|k}, V_{k|k}\}$  is the set of GGIW density parameters.

## 4.4.2 MBM merging

The GGIW implementations regarding the prediction and update of PPP and Bernoulli components are not presented due to page limits. The reader is referred to [25, 26] for more details. In this subsection, we choose to present the GGIW implementations regarding the block coordinate descent used to merge the MBM representing pre-existing tracks. The time index is omitted for simplicity.

### 4.4.2.1 E-step

In order to solve the optimisation problems (4.73) and (4.79) arisen in the E-step, the cross entropy between two Bernoulli-GGIW distribution must be calculated. Since the gamma distribution, the Gaussian distribution and the inverse Wishart distributions are mutually independent, a tractable solution can be derived.

Suppose  $f^h(\mathbf{X})$  and  $g^i(\mathbf{X})$  are two Bernoulli process with the following form

$$f^h(\mathbf{X}) = \begin{cases} 1 - r^h & \mathbf{X} = \emptyset, \\ r^h \mathcal{GGIW}(\mathbf{x}^h; \zeta^h) & \mathbf{X} = \{\mathbf{x}\}, \end{cases} \quad (4.88a)$$

$$g^i(\mathbf{X}) = \begin{cases} 1 - r^i & \mathbf{X} = \emptyset, \\ r^i \mathcal{GGIW}(\mathbf{x}^i; \zeta^i) & \mathbf{X} = \{\mathbf{x}\}. \end{cases} \quad (4.88b)$$

---

The cross entropy between  $f^h(\mathbf{X})$  and  $g^i(\mathbf{X})$  can be decomposed into:

$$\begin{aligned}
- \int f^h(\mathbf{X}) \log g^i(\mathbf{X}) d\mathbf{X} &= -(1 - r^h) \log(1 - r^i) \\
&\quad - r^h \log r^i - r^h \left( \int \mathcal{N}(\xi^h; m^h, P^h) \log \mathcal{N}(\xi^i; m^i, \hat{P}^i) d\mathbf{X} \right. \\
&\quad \quad + \int \mathcal{GAM}(\gamma^h; a^h, b^h) \log \mathcal{GAM}(\gamma^i; a^i, b^i) d\mathbf{X} \\
&\quad \quad \left. + \int \mathcal{IW}(\chi^h; v^h, V^h) \log \mathcal{IW}(\chi^i; v^i, V^i) d\mathbf{X} \right), \quad (4.89)
\end{aligned}$$

where

$$\begin{aligned}
\int \mathcal{N}(\xi^h; m^h, P^h) \log \mathcal{N}(\xi^i; m^i, P^i) d\mathbf{X} &= \\
- \frac{d}{2} \log(2\pi) - \frac{1}{2} \log |P^i| - \frac{1}{2} \text{Tr} \left( (P^h + (m^h - m^i)(m^h - m^i)^T) (P^i)^{-1} \right), \quad (4.90a)
\end{aligned}$$

$$\begin{aligned}
\int \mathcal{GAM}(\gamma^h; a^h, b^h) \log \mathcal{GAM}(\gamma^i; a^i, b^i) d\mathbf{X} &= \\
a^i \log b^i - \log \Gamma(a^i) + (a^i - 1)(\psi_0(a^h) - \log b^h) - b^i \frac{a^h}{b^h}, \quad (4.90b)
\end{aligned}$$

$$\begin{aligned}
\int \mathcal{IW}(\chi^h; v^h, V^h) \log \mathcal{IW}(\chi^i; v^i, V^i) d\mathbf{X} &= \\
- \frac{(v^i - d - 1)d}{2} \log 2 + \frac{v^i - d - 1}{2} \log |V^i| - \log \Gamma_d \left( \frac{v^i - d - 1}{2} \right) \\
- \frac{v^i}{2} \left( \log |V^h| - d \log 2 - \sum_{j=1}^d \psi_0 \left( \frac{v^h - d - j}{2} \right) \right) - \frac{1}{2} \text{Tr} \left( (v^h - d - 1)(V^h)^{-1} V^i \right). \quad (4.90c)
\end{aligned}$$

#### 4.4.2.2 M-step

Given a Bernoulli-GGIW mixture, the existence probability of the approximated Bernoulli is a weighted sum of the existence probability of each Bernoulli component. The mixture reduction for multivariate Gaussian distribution is straightforward. Theorems describing how a sum of an arbitrary number of Gamma components or inverse Wishart component can be merged into a single Gamma or inverse Wishart component were presented in [20] and [50] respectively. They are both performed via analytically minimising of the KL divergence. The same merging techniques also apply to merging the MBM representing new tracks (4.81). Corresponding mathematical expressions are not presented here due to page limits.

Empirically, we have found that in extended target filtering with GGIW implementation it is generally not advisable to merge the whole GGIW components. The

---

main reason is the extent state: merging two densities with significantly different extent estimates will result in an approximate density in which the extent estimates are distorted. This problem is exacerbated in the extended PMB filter, since the distorted extent states contained by the approximated single MB can easily lead to poor target state estimations in subsequent time steps. A simple strategy to handle this problem is to use a criterion for deciding which components should be merged. In this work, the KL divergence is used as the similarity measure between any pair of GGIW distributions. The component with the highest weight  $\mathcal{GGIW}(\xi^{h^*}; \zeta^{h^*})$  is chosen as the comparison baseline, which is merged with all other components  $\mathcal{GGIW}(\xi^h; \zeta^h)$  for which it holds

$$D_{\text{KL}}(\mathcal{GGIW}(\xi^{h^*}; \zeta^{h^*}) || \mathcal{GGIW}(\xi^h; \zeta^h)) < \tau, \quad (4.91)$$

where threshold  $\tau$  determines how aggressively the GGIW components are going to be merged.

## 4.5 Simulations and Results

In this section, we show simulation results that compare the PHD filter [22, 23], the PMBM filter [25, 26], and two variants of the PMB filter presented in this paper, namely one with the optimal assignment (OA) and one with the efficient approximation of feasible set (EAFS). In the PMBM filter, MB distributions with an equal number of Bernoulli components and small symmetrised KL divergence are merged [26].

### 4.5.1 State space model

The kinematic target state is a vector of position and velocity  $x_k = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$ . The single measurement is a vector of position  $z_k = [z_{x,k}, z_{y,k}]^T$ . The random matrix  $V_k$  is two-dimensional. The motion model  $F(\cdot)$  and process noise  $Q_k$  are

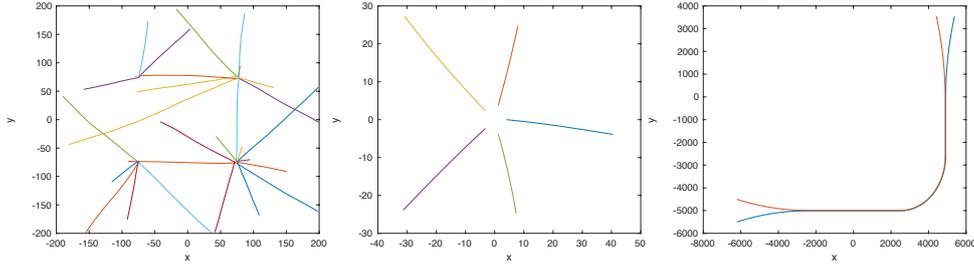
$$F(\xi_k) = \mathbf{I}_2 \otimes \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \xi_k, \quad Q_k = \sigma_v^2 \mathbf{I}_2 \otimes \begin{bmatrix} T^4/4 & T^3/2 \\ T^3/2 & T^2 \end{bmatrix},$$

where  $T = 1s$  is the sampling period, and  $\sigma_v$  is the standard deviation of motion noise. Because the kinematic state motion model is constant velocity, the extent transformation function  $M$  is an identity matrix, i.e.,  $M(\xi_k) = \mathbf{I}_2$ .

### 4.5.2 Performance evaluation

For GGIW-PMB, the target states are extracted by taking the mean vector of all Bernoulli components with existence probability larger than 0.5. For GGIW-PMBM, target state extraction is performed analogously, but only from the MB distribution with the highest weight. For GGIW-PHD, target states are obtained from all the GGIW components with weight larger than 0.5.

The estimation performance is evaluated using the Gaussian Wasserstein distance



**Figure 4.14:** True target trajectories of three scenarios, from left to right: 1) 27 objects, 2) dense birth, and 3) parallel manoeuvre.

**Table 4.11:** Average time [s] to run one full Monte Carlo simulation

Scenario	PHD	PMBM	PMB w/ OA	PMB w/ EAFS
1	34.1	57.6	37.9	41.1
2	21.4	45.8	7.7	8.5
3	92.0	41.4	25.5	25.1

(GWD) [51] integrated into the generalised optimal sub-pattern assignment (GOSPA) metric [15]. The GOSPA metric allows for decomposing the estimation error into localisation error, missed detection error and false detection error.

### 4.5.3 Simulation study

Filters are evaluated in three different scenarios. True target trajectories are shown in Fig. 4.14. For each scenario, the result is averaged over 100 Monte Carlo trials. In the first scenario, 27 randomly generated targets are born from four localised positions, and they appear and disappear the surveillance area at different time steps. The parameters were set to  $p^D = 0.90$ ,  $p^S = 0.99$  and  $\lambda = 60$ . This scenario illustrates how the different filters behave with a high target number and high clutter density scenario. In the second scenario, five targets are born at a very short distance from each other at the same time step. The parameters were set to  $p^D = 0.90$ ,  $p^S = 0.99$  and  $\lambda = 20$ . This scenario tests different filters capabilities of handling dense birth. In the third scenario, two targets first get close, and then they manoeuvre in close proximity before splitting. The parameters were set to  $p^D = 0.98$ ,  $p^S = 0.99$  and  $\lambda = 10$ . Here, the data association is very challenging due to coalescence and motion model mismatch.

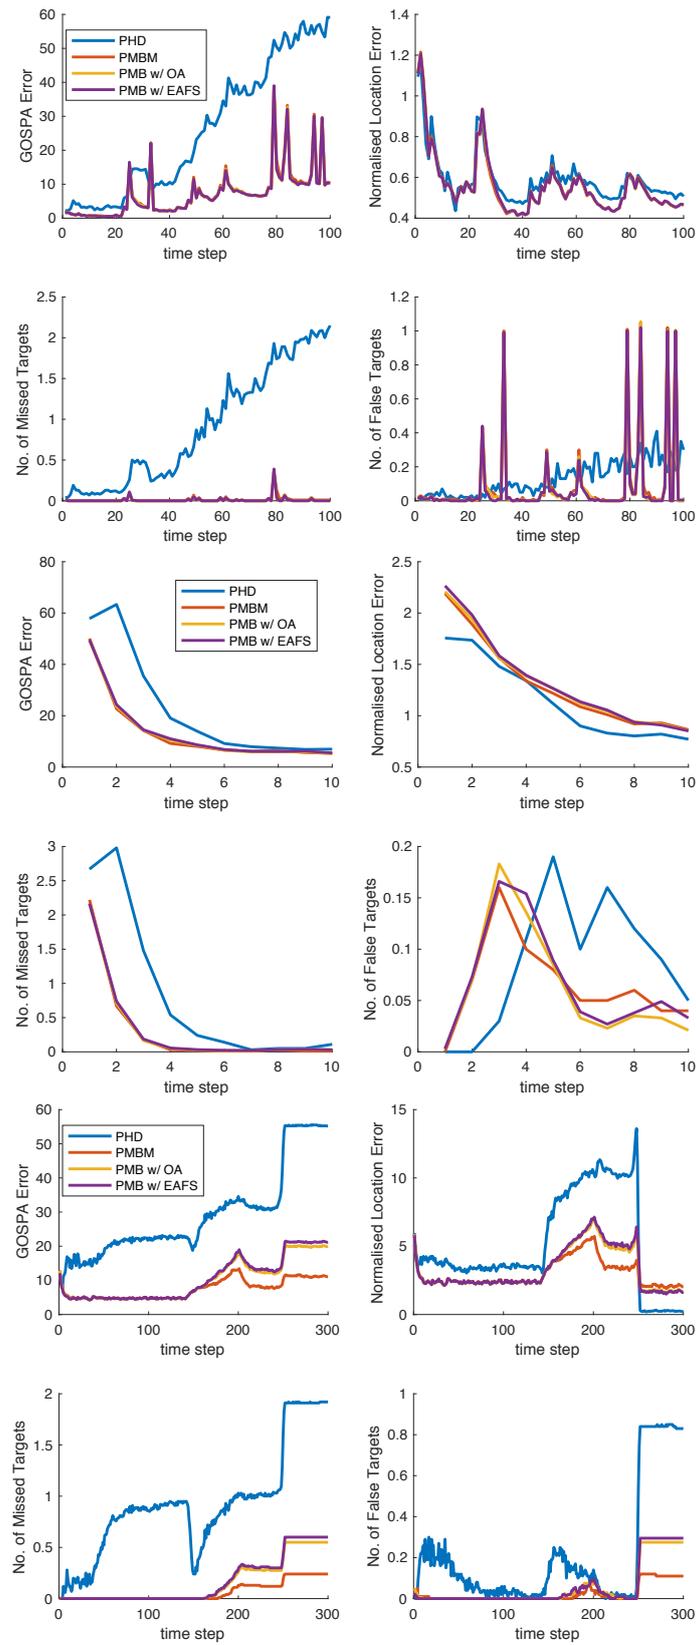
The GOSPA performance of different filters is shown in Fig. 4.15, and the average times to run a single Monte Carlo trial of the three scenarios are given in Table 4.11. From the results of the simulation study, we see that filters based on MB conjugate priors have better performance than the PHD filter. The two variants of the PMB filter showed very similar performance in these three scenarios. The PMB filter has lower computational complexity than the PMBM filter, and the difference of average running time is most distinct in the scenario with dense birth. But the PMBM filter

---

is able to produce better target extent estimations than the PMB filter especially in the scenario with parallel manoeuvre.

## 4.6 Conclusions

This paper has proposed a tractable and efficient extended target filtering algorithm based on a PMBM conjugate prior approximation to the posterior density. A simulation study shows that the proposed PMB filter retains the advantage of the PMBM filter but with lower computational complexity. Possible future work includes how to improve the estimation of target extent and how to incorporate the formation of new tracks in the variational MB algorithm.



**Figure 4.15:** Results from simulation of three scenarios, from left to right: 1) 27 objects, 2) dense birth, and 3) parallel manoeuvre.

# Bibliography

- [1] Thomas Fortmann, Yaakov Bar-Shalom, and Molly Scheffe. “Sonar tracking of multiple targets using joint probabilistic data association”. In: *IEEE journal of Oceanic Engineering* 8.3 (1983), pp. 173–184.
- [2] Samuel Blackman and Robert Popoli. “Design and analysis of modern tracking systems(Book)”. In: *Norwood, MA: Artech House* (1999).
- [3] Ronald PS Mahler. *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.
- [4] Ronald PS Mahler. “Multitarget Bayes filtering via first-order multitarget moments”. In: *IEEE Transactions on Aerospace and Electronic systems* 39.4 (2003), pp. 1152–1178.
- [5] B-N Vo and W-K Ma. “The Gaussian mixture probability hypothesis density filter”. In: *IEEE Transactions on signal processing* 54.11 (2006), pp. 4091–4104.
- [6] Ronald Mahler. “PHD filters of higher order in target number”. In: *IEEE Transactions on Aerospace and Electronic Systems* 43.4 (2007).
- [7] Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni. “Analytic implementations of the cardinalized probability hypothesis density filter”. In: *IEEE Transactions on Signal Processing* 55.7 (2007), pp. 3553–3567.
- [8] Ba-Tuong Vo and Ba-Ngu Vo. “Labeled random finite sets and multi-object conjugate priors”. In: *IEEE Transactions on Signal Processing* 61.13 (2013), pp. 3460–3475.
- [9] Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung. “Labeled random finite sets and the Bayes multi-target tracking filter”. In: *IEEE Transactions on Signal Processing* 62.24 (2014), pp. 6554–6567.
- [10] Ba Ngu Vo, Ba-Tuong Vo, and Hung Hoang. “An Efficient Implementation of the Generalized Labeled Multi-Bernoulli Filter”. In: *IEEE Transactions on Signal Processing* (2016).
- [11] Stephan Reuter et al. “The labeled multi-Bernoulli filter”. In: *IEEE Transactions on Signal Processing* 62.12 (2014), pp. 3246–3260.
- [12] Jason L Williams. “Marginal multi-bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member”. In: *IEEE Transactions on Aerospace and Electronic Systems* 51.3 (2015), pp. 1664–1687.
- [13] Ángel F. García-Fernández et al. “Poisson multi-Bernoulli mixture filter: direct derivation and implementation”. In: *arXiv preprint arXiv:1703.04264* (2017).
- [14] Jason L Williams. “An efficient, variational approximation of the best fitting multi-Bernoulli filter”. In: *IEEE Transactions on Signal Processing* 63.1 (2015), pp. 258–273.

- [15] Abu Sajana Rahmathullah, Ángel F García-Fernández, and Lennart Svensson. “Generalized optimal sub-pattern assignment metric”. In: *arXiv preprint arXiv:1601.05585* (2016).
- [16] Kevin Gilholm and David Salmond. “Spatial distribution model for tracking extended objects”. In: *IEEE Proceedings-Radar, Sonar and Navigation* 152.5 (2005), pp. 364–371.
- [17] Marcus Baum and Uwe D Hanebeck. “Extended object tracking with random hypersurface models”. In: *IEEE Transactions on Aerospace and Electronic Systems* 50.1 (2014), pp. 149–159.
- [18] Johann Wolfgang Koch. “Bayesian approach to extended object and cluster tracking using random matrices”. In: *IEEE Transactions on Aerospace and Electronic Systems* 44.3 (2008).
- [19] Michael Feldmann, Dietrich Franken, and Wolfgang Koch. “Tracking of extended objects and group targets using random matrices”. In: *IEEE Transactions on Signal Processing* 59.4 (2011), pp. 1409–1420.
- [20] Karl Granström and Umut Orguner. “Estimation and maintenance of measurement rates for multiple extended target tracking”. In: *Information Fusion (FUSION), 2012 15th International Conference on*. IEEE. 2012, pp. 2170–2176.
- [21] Christian Lundquist, Karl Granström, and Umut Orguner. “An extended target CPHD filter and a gamma Gaussian inverse Wishart implementation”. In: *IEEE Journal of Selected Topics in Signal Processing* 7.3 (2013), pp. 472–483.
- [22] Karl Granström, Christian Lundquist, and Umut Orguner. “Extended target tracking using a Gaussian-mixture PHD filter”. In: *IEEE Transactions on Aerospace and Electronic Systems* 48.4 (2012), pp. 3268–3286.
- [23] Karl Granström and Umut Orguner. “A PHD filter for tracking multiple extended targets using random matrices”. In: *IEEE Transactions on Signal Processing* 60.11 (2012), pp. 5657–5671.
- [24] Michael Beard et al. “Multiple extended target tracking with labeled random finite sets”. In: *IEEE Transactions on Signal Processing* 64.7 (2016), pp. 1638–1653.
- [25] Karl Granström, Maryam Fatemi, and Lennart Svensson. “Gamma Gaussian inverse-Wishart Poisson multi-Bernoulli filter for extended target tracking”. In: (2016), pp. 893–900.
- [26] Karl Granström, Maryam Fatemi, and Lennart Svensson. “Poisson multi-Bernoulli conjugate prior for multiple extended object estimation”. In: *arXiv preprint arXiv:1605.06311* (2016).
- [27] M. Fatemi K. Granström S. Reuter and L. svensson. “Pedestrian tracking using Velodyne data – stochastic optimization for extended object tracking”. In: (2017).
- [28] Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni. “The cardinality balanced multi-target multi-Bernoulli filter and its implementations”. In: *IEEE Transactions on Signal Processing* 57.2 (2009), pp. 409–423.
- [29] Dominic Schuhmacher, Ba-Tuong Vo, and Ba-Ngu Vo. “A consistent metric for performance evaluation of multi-object filters”. In: *IEEE Transactions on Signal Processing* 56.8 (2008), pp. 3447–3457.

- 
- [30] KG Murthy. “An algorithm for ranking all the assignments in order of increasing costs”. In: *Operations Research* 16.3 (1968), pp. 682–687.
- [31] Jason Williams and Roslyn Lau. “Approximate evaluation of marginal association probabilities with belief propagation”. In: *IEEE Transactions on Aerospace and Electronic Systems* 50.4 (2014), pp. 2942–2959.
- [32] David Eppstein. “Finding the k shortest paths”. In: *SIAM Journal on computing* 28.2 (1998), pp. 652–673.
- [33] Jason L Williams. “Hybrid Poisson and multi-Bernoulli filters”. In: *Information Fusion (FUSION), 15th International Conference on*. IEEE. 2012, pp. 1103–1110.
- [34] Lennart Svensson et al. “Set JPDA filter for multitarget tracking”. In: *IEEE Transactions on Signal Processing* 59.10 (2011), pp. 4677–4691.
- [35] Maryam Fatemi et al. “Poisson Multi-Bernoulli Mapping Using Gibbs Sampling”. In: *IEEE Transactions on Signal Processing* 65.11 (2017), pp. 2814–2827.
- [36] Karl Granström and Marcus Baum. “Extended object tracking: Introduction, overview and applications”. In: *arXiv preprint arXiv:1604.00970* (2016).
- [37] Radford M Neal. “Markov chain sampling methods for Dirichlet process mixture models”. In: *Journal of computational and graphical statistics* 9.2 (2000), pp. 249–265.
- [38] Sonia Jain and Radford M Neal. “A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model”. In: *Journal of Computational and Graphical Statistics* 13.1 (2004), pp. 158–182.
- [39] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In: *Kdd*. Vol. 96. 34. 1996, pp. 226–231.
- [40] Yaakov Bar-Shalom, Peter K Willett, and Xin Tian. *Tracking and data fusion*. YBS publishing, 2011.
- [41] Hanna Pasula et al. “Tracking many objects with many sensors”. In: *IJCAI*. Vol. 99. 1999, pp. 1160–1171.
- [42] Shan Cong, Lang Hong, and Devert Wicker. “Markov-chain Monte-Carlo approach for association probability evaluation”. In: *IEE Proceedings-Control Theory and Applications* 151.2 (2004), pp. 185–193.
- [43] Songhwai Oh, Stuart Russell, and Shankar Sastry. “Markov chain Monte Carlo data association for multi-target tracking”. In: *IEEE Transactions on Automatic Control* 54.3 (2009), pp. 481–497.
- [44] Stephan Reuter et al. “A fast implementation of the Labeled Multi-Bernoulli filter using gibbs sampling”. In: *Intelligent Vehicles Symposium (IV), 2017 IEEE*. IEEE. 2017, pp. 765–772.
- [45] Radford M Neal. “Markov chain sampling methods for Dirichlet process mixture models”. In: *Journal of computational and graphical statistics* 9.2 (2000), pp. 249–265.
- [46] David Arthur and Sergei Vassilvitskii. “k-means++: The advantages of careful seeding”. In: *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics. 2007, pp. 1027–1035.

- [47] Shishan Yang, Marcus Baum, and Karl Granström. “Metrics for performance evaluation of elliptic extended object tracking methods”. In: (2016), pp. 523–528.
- [48] Yuxuan Xia et al. “Performance Evaluation of Multi-Bernoulli Conjugate Priors for Multi-Target Filtering”. In: *20th International Conference on Information Fusion*. Xi’an, P.R. China, 2017.
- [49] Arthur P Dempster, Nan M Laird, and Donald B Rubin. “Maximum likelihood from incomplete data via the EM algorithm”. In: *Journal of the royal statistical society. Series B (methodological)* (1977), pp. 1–38.
- [50] Karl Granström and Umut Orguner. “On the reduction of Gaussian inverse Wishart mixtures”. In: (2012), pp. 2162–2169.
- [51] Clark R Givens, Rae Michael Shortt, et al. “A class of Wasserstein metrics for probability distributions.” In: *The Michigan Mathematical Journal* 31.2 (1984), pp. 231–240.