



Poisson Multi-Bernoulli Mapping Using Gibbs Sampling

Downloaded from: <https://research.chalmers.se>, 2025-06-18 04:26 UTC

Citation for the original published paper (version of record):

Fatemi, M., Granström, K., Svensson, L. et al (2017). Poisson Multi-Bernoulli Mapping Using Gibbs Sampling. IEEE Transactions on Signal Processing, 65(11): 2814-2827.
<http://dx.doi.org/10.1109/tsp.2017.2675866>

N.B. When citing this work, cite the original published paper.

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

Poisson Multi-Bernoulli Mapping Using Gibbs Sampling

Maryam Fatemi, Karl Granström, Lennart Svensson, Francisco J. R. Ruiz, and Lars Hammarstrand

Abstract—This paper addresses the mapping problem. Using a conjugate prior form, we derive the exact theoretical batch multi-object posterior density of the map given a set of measurements. The landmarks in the map are modeled as extended objects, and the measurements are described as a Poisson process, conditioned on the map. We use a Poisson process prior on the map and prove that the posterior distribution is a hybrid Poisson, multi-Bernoulli mixture distribution. We devise a Gibbs sampling algorithm to sample from the batch multi-object posterior. The proposed method can handle uncertainties in the data associations and the cardinality of the set of landmarks, and is parallelizable, making it suitable for large-scale problems. The performance of the proposed method is evaluated on synthetic data and is shown to outperform a state-of-the-art method.

Index Terms—Statistical mapping, extended object, Monte Carlo methods, inference algorithms, sampling methods, Gibbs sampling.

I. INTRODUCTION

During recent years, self-driving cars have been the subject of extensive research. Many different functionalities expected from a self-driving vehicle are facilitated by having accurate localization and mapping capabilities. These capabilities require information of the environment which is gained by on-board sensors such as radars, cameras and internal sensors. The problem of mapping an unknown environment and estimating the unknown trajectory of the vehicle concurrently is referred to as simultaneous localization and mapping (SLAM) [1] [2]. SLAM has attracted a lot of attention, especially in the robotics community. Examples of SLAM algorithms are EKF-SLAM [3] [4], FastSLAM [5], LMB-SLAM [6] and Graph-SLAM [7]. The map produced by these methods is a statistical map of the environment as viewed by a particular sensor. To achieve higher positioning accuracy, it is common to separate the localization and mapping problems. In this case, the statistical mapping problem is solved offline using data from a vehicle equipped with a reference positioning system [8] [9] [10]. In such a system, if the position uncertainties are negligible, the focus can be on the statistical mapping.

Mapping has been addressed in, e.g., [10], [11] and [12]. Two aspects that make the problem challenging are the unknown data association between the measurements and the

landmarks, and the unknown number of landmarks. In [11] and [12], these two aspects are addressed through representing the map and the measurements as random finite sets (RFS), thus incorporating the uncertainties in the data association and the number of landmarks into the model. They use the probability hypothesis density (PHD) filter [13] to recursively approximate the posterior density of the map. In addition, [11] and [12] model the sensor detections as point objects, i.e., each object can generate at most one measurement at each time step. For some automotive sensors, the distance to landmarks is often such that a landmark is covered by more than one resolution cell of the sensor. It is therefore more reasonable to use an extended object model, where each object can generate more than one measurement at each time step.

A commonly used measurement model for extended objects is the inhomogeneous Poisson point process (PPP) proposed in [14]. In this model, a detected object generates a Poisson distributed number of measurements, and these measurements are spatially distributed around the object. This extended object model has been used in [10], [15], [16], [17]. In [10], the authors present a vector-based extended object mapping, where a variational Bayesian expectation maximization (VBEM) algorithm [18] is used to approximate the posterior. In the algorithm presented in [10], the data associations and the map are jointly estimated; however, the uncertainty of the number of landmarks is not taken into account.

In the mapping problem, where the landmarks are assumed to be extended objects, we need to cluster together the measurements that were generated by the same landmark in order to infer the properties of that landmark. However, the number of clusters is unknown because it corresponds to the number of landmarks. This resembles the problem addressed in the Bayesian nonparametric (BNP) literature [19] [20]. The BNP approach consists in fitting a model that can adapt its complexity to the observed data [20]. In particular, the Dirichlet process (DP) [21] defines a prior distribution over an infinite-dimensional parameter space, ensuring that only a finite number of parameters are needed to generate any finite dataset. The posterior distribution reveals the number of parameters (clusters) that best fit the data. Both Markov Chain Monte Carlo (MCMC) methods [22] [23] and variational methods [24] have been proposed to perform inference in this context, where the former is the most widely used method of inference [20].

In this paper, we develop an MCMC method to estimate sensor maps that contain extended objects. The proposed method can be applied to any sensor that is subject to false alarms, missed detections, and data association uncertainty. The prior

Maryam Fatemi was with the Department of Signals and Systems, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden. She is now with Autoliv Sverige AB (email: maryam.fatemi@autoliv.com). Karl Granström, Lennart Svensson and Lars Hammarstrand are with the Department of Signals and Systems, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden (emails: {karl.granstrom, lennart.svensson, lars.hammarstrand}@chalmers.se). Francisco J. R. Ruiz is with the Department of Engineering, University of Cambridge, United Kingdom, and with the Computer Science Department, Columbia University, NY, USA (email: f.ruiz@columbia.edu).

of the map is modeled by a Poisson process which can incorporate the uncertainty regarding the number of landmarks in a systematic manner, which gives important advantages compared to, e.g., the method of [10], where the prior of the map does not incorporate uncertainties in the number of landmarks. Therefore, our method is capable of handling the uncertainties in both data associations and unknown number of landmarks. We derive the exact batch multi-object posterior density of the map using the conjugate prior form described in [25].¹ To estimate the posterior, we derive a collapsed Gibbs sampling [22] [23] algorithm which is analogous to widely used Gibbs sampling methods in BNP. We show that our method has the advantage of being parallelizable, making it suitable for large scale problems.

The paper is organized as follows. The problem formulation is presented in Section II. Because in our derivations we make extensive use of the results given in [25], Section III presents a background on this paper. The batch posterior density is derived in Section IV followed by the proposed sampling method in Section V, where the similarities and differences between an alternative BNP approach and the proposed method are discussed. The inference algorithm is evaluated on a simulated scenario in Section VII and conclusions are made in Section VIII.

II. PROBLEM FORMULATION

We are interested in mapping an area of interest (AOI). A vehicle equipped with a sensor navigates the AOI and collects data at discrete time steps, and our task is to construct a statistical model of those measurements as a function of the sensor pose. At each time instance the sensor only measures within its limited field of view (FOV) and the union of all these FOVs outlines the complete observed area (OA). In some scenarios the vehicle does not visit all parts of the AOI, in which case the OA is a subset of the AOI. Because of this, it is necessary to reason about both the parts of the AOI that have been observed, and the parts of the AOI that have not been observed.

The AOI is assumed to contain objects, here called landmarks. The map is the set of landmarks described by an RFS Θ , for which the cardinality and the properties (distribution) of its members are unknown. Each landmark is modelled as an elliptical extended object described by the landmark state θ containing its 2D position μ , extent modelled as a random matrix Σ , and expected number of measurements ω .

Mapping of the AOI is to estimate the batch posterior density of the set of landmarks given the set of collected measurements from K time steps, i.e., $f(\Theta|\mathbf{Z})$, where $\mathbf{Z} = \mathbf{Z}_1 \cup \mathbf{Z}_2 \dots \cup \mathbf{Z}_K$ is the union of K sets of measurements. These measurements are modelled by the standard PPP extended object measurement model. Since the OA is a subset of the AOI, it is suitable to model the set of landmarks as the union of two disjoint subsets: a set of detected landmarks Θ^d and a set of undetected landmarks Θ^u . The detected landmarks are within the OA, while the undetected ones

are mostly outside the OA. However, because of detection uncertainty, there is a non-zero probability that there is one or more undetected landmarks inside the OA. The batch posterior density, denoted by $f(\Theta|\mathbf{Z})$, is a Poisson MBM (PMBM) density. For the standard PPP extended object measurement model [14], this is a conjugate prior [25].

In what follows, we describe the models and assumptions based on which we derive $f(\Theta|\mathbf{Z})$ in Section IV. The models and derivations are based on finite set statistics [27]. RFS based methods [27] facilitate statistical inference in problems where the variables of interest and/or the measurements are modelled as finite sets. These methods are specially attractive because they can accommodate uncertainties in both the number of variables and their states. These two aspects match the statistical mapping problem very well since we are interested in estimating the number of landmarks as well as their properties. A brief overview of some key concepts of the random set theory required for following the presented models and the derivations, such as the distribution of a set-valued variable and its probability generating functional (p.g.fl), is given in Appendix A.

A. Standard extended object measurement model

In the standard extended object measurement model [27, pp. 431-432] [14], it is assumed that the measurements form an RFS that is the union of two independent RFSs: one for landmark generated measurements, and one for clutter. The clutter is modelled by a PPP with rate λ_c and spatial density $c(\mathbf{z})$. The p.g.fl of the clutter measurement model is

$$\begin{aligned} G_k^c[g] &= \exp(\lambda_c \langle c; g \rangle - \lambda_c) \\ &= \exp(\langle \kappa; g \rangle - \langle \kappa; 1 \rangle), \end{aligned} \quad (1)$$

where $\kappa(\mathbf{z}) = \lambda_c c(\mathbf{z})$ is the clutter PPP intensity.

A landmark is detected at time k with state dependent probability of detection $p_D^k(\theta)$. Measurements generated from landmarks are modelled by a PPP with rate $\gamma_k(\theta)$ and spatial distribution $\phi(\mathbf{z}_k|\theta)$. The conditional set likelihood of extended landmark measurements is expressed as [25]

$$\begin{aligned} l_{\mathbf{z}_k}(\theta) &= p_D^k(\theta) p(\mathbf{Z}_k|\theta) \\ &= p_D^k(\theta) e^{-\gamma_k(\theta)} \prod_{\mathbf{z}_k \in \mathbf{Z}_k} \gamma_k(\theta) \phi(\mathbf{z}_k|\theta). \end{aligned} \quad (2)$$

The Poisson rate $\gamma_k(\theta)$ models the expected number of detections from a landmark with state θ .

The p.g.fl of the measurement model of a single extended landmark is

$$G_k^l[g|\theta] = 1 - p_D^k(\theta) + p_D^k(\theta) \exp(\gamma \langle \phi; g \rangle - \gamma) \quad (3)$$

where $\langle \phi; g \rangle = \int \phi(\mathbf{z}|\theta) g(\mathbf{z}) d\mathbf{z}$ is a function of θ . It is assumed that the measurements originated from different landmarks are independent, therefore, the p.g.fl of the measurement model of multiple landmarks can be described by

$$G_k^l[g|\Theta] = (1 - p_D^k(\theta) + p_D^k(\theta) \exp(\gamma \langle \phi; g \rangle - \gamma))^{\Theta} \quad (4)$$

Using (58), (1) and (4) the p.g.fl for the complete set of measurements can be expressed by

$$G_k[g|\Theta] = G_k^c[g] G_k^l[g|\Theta]. \quad (5)$$

¹ An extended version of this paper was presented in [26], where the proof of conjugacy is included.

B. Modelling assumptions

The prior RFS density of the map is described by a PPP. The map is assumed to be static, i.e., the appearance of landmarks on the map or their disappearance from the map is not accounted for in this work.

It is further assumed that the measurement batch contains measurements from K time steps, and is denoted by the set \mathbf{Z} . A data association hypothesis over \mathbf{Z} describes which measurements belong to which landmark, and partitions this set into non-empty subsets, called cells, whose members belong to the same landmark. The i th cell in partition j is denoted by $\mathbf{C}^{j,i}$, further,

$$\biguplus_i \{\mathbf{C}^{j,i}\} = \mathbf{Z}, \forall j. \quad (6)$$

In each cell we group the measurements by their time tags. Measurements with time tag k form a set denoted by $\mathbf{V}_k^{j,i}$, therefore,

$$\biguplus_k \{\mathbf{V}_k^{j,i}\} = \mathbf{C}^{j,i}, \forall j, i. \quad (7)$$

Note that $\mathbf{V}_k^{j,i}$ could possibly be empty for some k , as each landmark is within the FOV for a limited number of time steps.

In the clutter model described by (1), we assume that clutter measurements are uniformly distributed over the field of view of the sensor, i.e.,

$$c(\mathbf{z}) = V^{-1}, \quad (8)$$

where V is the observation volume of the sensor. In (2), we model

$$\gamma_k(\theta) = \omega f(\mu, \mathbf{x}_k) \quad (9)$$

where ω is the expected number of detections generated by a landmark, μ is the 2D position of the landmark, \mathbf{x}_k is the pose of the sensor at time k , and $f(\mu, \mathbf{x}_k)$ is the FOV function. This function is equal to one if μ is in the FOV at time k , and zero otherwise. The connection between the sensor's field of view and the extent of a landmark is not taken into consideration for simplicity. This implies that a landmark is assumed to be outside the FOV if its centre is outside the FOV, even if a part of its extent is inside. The spatial distribution is modelled as

$$\phi(\mathbf{z}_k|\theta) = \mathcal{N}(\mathbf{z}_k; \mu, \Sigma), \quad (10)$$

where Σ is a random matrix that describes the landmark extent. In addition, this model assumes that landmark measurements are generated by the centroid of the landmark, and since in reality the measurements could have been generated from any scattering point on the landmark, the measurement error corresponds to the extent of the landmark [28]. In this model it is assumed that the sensor measurement noise is negligible compared to a landmark's extent. That is, if Σ_0 is the true extent and \mathbf{R} the measurement noise covariance, it is assumed that \mathbf{R} is small compared to Σ_0 . More specifically, the measurement covariance in (10) is the added effect of Σ_0 and \mathbf{R} , and is approximated as $\Sigma = \Sigma_0 + \mathbf{R} \approx \Sigma_0$.

III. BACKGROUND

In this section we present a summary of the sequential posterior RFS density for multiple extended objects. The reader is referred to [25] and [26] for a comprehensive description.

In [26], the authors prove that a Poisson multi-Bernoulli mixture (PMBM) RFS density is a conjugate prior for the Poisson extended object measurement model (presented in Section II-A). In addition, it is shown that updating a Poisson prior with this measurement model results in a PMBM posterior. Therefore, for a batch problem, if the prior is a Poisson process, the posterior resulting from the batch update would be PMBM.

In this section the results of [25] are summarized. Note that only the results for the update step are presented here, this is because in this paper we are dealing with a static state (map) and there is no need to perform a prediction step. That is, it can be said that the predicted density is equal to the posterior density of the previous time step.

Similar to Section II, the set of landmarks Θ is divided into two disjoint subsets of detected and undetected landmarks denoted by Θ^d and Θ^u , respectively. It is assumed that the p.g.fl of the prior density at time $k+1$ is

$$G_k[h] = G_k^u[h]G_k^d[h], \quad (11)$$

where

$$G_k^u[h] = \exp(\lambda_k^u \langle f_k^u; h \rangle - \lambda_k^u) \quad (12)$$

$$G_k^d[h] = \sum_{j=1}^{N_k} W_k^j \prod_{i=1}^{I_k^j} (1 - r_k^{j,i} + r_k^{j,i} \langle f_k^{j,i}; h \rangle), \quad (13)$$

the Poisson rate of the undetected landmarks at time k is denoted by λ_k^u , and f_k^u is the spatial density of undetected landmarks at time k . The p.g.fl of detected landmarks is an MBM where the weight of the j th multi-Bernoulli component is denoted by W_k^j . This component corresponds to a partitioning of the set \mathbf{Z}_k that consists of I_k^j cells each denoted by $\mathbf{C}^{j,i}$. A cell refers to a subset of \mathbf{Z}_k that cannot be empty. Accordingly, N_k is the number of ways that the set \mathbf{Z}_k can be partitioned into different cells, and it is equal to the Bell number of order $|\mathbf{Z}_k|$ [29]. In each partition, measurements assigned to the same cell are assumed to have originated from the same source, i.e., the same landmark or clutter. The probability of existence of a landmark corresponding to cell $\mathbf{C}^{j,i}$ is $r_k^{j,i}$ and its spatial density is denoted by $f_k^{j,i}$.

The update step has two parts, updating the undetected landmarks, and updating the previously detected ones. Upon receiving the measurements at time $k+1$, the set \mathbf{Z}_{k+1} is split into two subsets, namely, measurements not generated by previously detected landmarks \mathbf{Y}_{k+1} , and measurements generated by previously detected landmarks $\mathbf{Z}_{k+1} \setminus \mathbf{Y}_{k+1}$. The newly detected landmarks will form a MB RFS, conditioned on a partitioning \mathcal{P} of the set \mathbf{Y}_{k+1} ,

$$G_{k+1}^{\mathcal{P}}[h] = \prod_{\Upsilon \in \mathcal{P}} (1 - r_{\Upsilon} + r_{\Upsilon} \langle f_{\Upsilon}; h \rangle), \quad (14)$$

where

$$r_{\Upsilon} = \begin{cases} 1 & \text{if } |\Upsilon| > 1 \\ \frac{1}{\kappa_{\Upsilon} + \mathcal{L}_{\Upsilon}} & \text{if } |\Upsilon| = 1 \end{cases} \quad (15)$$

$$f_{\Upsilon}(\theta) = \frac{l_{\Upsilon}(\theta) \lambda_k^u f_k^u(\theta)}{\lambda_k^u \langle f_k; l_{\Upsilon} \rangle} \quad (16)$$

$$\mathcal{L}_{\Upsilon} = \lambda_k^u \langle f_k; l_{\Upsilon} \rangle. \quad (17)$$

Here, a cell of the partition \mathcal{P} is denoted by Υ , and $|\Upsilon|$ denotes the cardinality of the set Υ . In addition, l_{Υ} is the conditional set likelihood described in (2). The rate and spatial density of the PPP process for undetected landmarks are expressed by

$$\lambda_{k+1}^u = \lambda_k^u \langle f_k^u; q_D^{k+1} \rangle \quad (18)$$

$$f_{k+1}^u(\theta) = \frac{q_D^{k+1}(\theta) f_k^u(\theta)}{\langle f_k^u; q_D^{k+1} \rangle} \quad (19)$$

$$q_D^{k+1}(\theta) = 1 - p_D^{k+1}(\theta) + p_D^{k+1}(\theta) e^{-\gamma_{k+1}(\theta)}, \quad (20)$$

where q_D^{k+1} is the effective probability of missed detection at time $k+1$, p_D^{k+1} is the probability of detection and $\gamma_{k+1}(\theta)$ is the intensity of the PPP process of the measurements.

The set $\mathbf{Z}_{k+1} \setminus \mathbf{Y}_{k+1}$ contains measurements originated from the previously detected landmarks, and is partitioned into disjoint, possibly empty, subsets $\mathbf{V}_{k+1}^{j,i}$. The multi-Bernoulli component corresponding to this set is described by

$$G_{k+1|\mathbf{Y}_{k+1}}^{\{\mathbf{V}_{k+1}^{j,i}\}_i}[h] = \prod_i (1 - r_{k+1}^{j,i} + r_{k+1}^{j,i} \langle f_{k+1}^{j,i}; h \rangle) \quad (21)$$

where $\biguplus_i \{\mathbf{V}_{k+1}^{j,i}\} = \mathbf{Z}_{k+1} \setminus \mathbf{Y}_{k+1}, \forall j$. In addition, the probabilities of existence, the spatial distributions and the predicted likelihoods depend on the cardinality of the disjoint subsets. That is, if $\mathbf{V}_{k+1}^{j,i} \neq \emptyset$,

$$r_{k+1}^{j,i} = 1 \quad (22)$$

$$f_{k+1}^{j,i}(\theta) = \frac{l_{k+1}^{j,i}(\theta) f_k^{j,i}(\theta)}{\langle f_k^{j,i}; l_{k+1}^{j,i} \rangle} \quad (23)$$

$$\mathcal{L}_{k+1}^{\mathbf{V}_{k+1}^{j,i}} = r_{k+1}^{j,i} \langle f_k^{j,i}; l_{k+1}^{j,i} \rangle. \quad (24)$$

If $\mathbf{V}_{k+1}^{j,i} = \emptyset$,

$$r_{k+1}^{j,i} = \frac{r_k^{j,i} \langle f_k^{j,i}; q_D^{k+1} \rangle}{1 - r_k^{j,i} + r_k^{j,i} \langle f_k^{j,i}; q_D^{k+1} \rangle} \quad (25)$$

$$f_{k+1}^{j,i}(\theta) = \frac{q_D^{k+1}(\theta) f_k^{j,i}(\theta)}{\langle f_k^{j,i}; q_D^{k+1} \rangle} \quad (26)$$

$$\mathcal{L}_{k+1}^{\mathbf{V}_{k+1}^{j,i}} = 1 - r_k^{j,i} + r_k^{j,i} \langle f_k^{j,i}; q_D^{k+1} \rangle. \quad (27)$$

In (22) the existence probability of object i in global hypothesis j is equal to one. However, note that global hypothesis j is not certain, but has a probability $W_k^j \in [0, 1]$. Consequently, (22) merely implies that object i exists with probability 1 conditioned on hypothesis j .

The weight of each multi-Bernoulli component is computed by the predicted partition likelihood

$$\mathcal{L}_P = \prod_{\Upsilon \in \mathcal{P} | |\Upsilon| > 1} \mathcal{L}_{\Upsilon} \times \prod_{\Upsilon \in \mathcal{P} | |\Upsilon| = 1} (\kappa_{\Upsilon} + \mathcal{L}_{\Upsilon}), \quad (28)$$

and the predicted MB likelihood

$$\mathcal{L}_{k+1}^{\{\mathbf{V}_{k+1}^{j,i}\}_i} = \prod_i \mathcal{L}_{k+1}^{\mathbf{V}_{k+1}^{j,i}}. \quad (29)$$

IV. BATCH MULTI-OBJECT POSTERIOR DENSITY

In this section, we present the exact theoretical multi-object batch posterior density. This distribution is calculated by performing a batch update using all the collected measurements. Naturally, we expect the batch posterior to be the same as the sequential posterior after having updated the latter with all collected measurements.

A. The batch update

The resulting batch multi-object posterior density has a PMBM form where the PPP part describes the posterior of the undetected landmarks and the MBM part describes the posterior of the detected landmarks and the clutter. The multi-object batch density and the update for parameters of the PPP and the MBM are described in Theorem 1, the proof of which is presented in Appendix B.

Theorem 1. *The multi-object batch posterior density of the map Θ given the set of measurements \mathbf{Z} has a PMBM form and is expressed as*

$$f(\Theta) = \sum_{\Theta^d \subset \Theta} \left(\sum_{j=1}^{N_K} W_K^j \sum_{\alpha^j \in P_{I_K^j}^{|\Theta^d|}} \prod_{i=1}^{I_K^j} f_K^{j,i}(\Theta_{\alpha^j(i)}) \right) \times \left(e^{-\lambda_K^u} \prod_{\theta \in (\Theta \setminus \Theta^d)} \lambda_K^u f_K^u(\theta) \right), \quad (30)$$

where the first factor describes an MBM, the second factor is a PPP and the entire set Θ is the union of these disjoint sets. In the MBM, the index j describes a partitioning of an entire set of the measurements and represents an MB in the mixture, whereas, the index i corresponds to different Bernoulli components in an MB. The mapping $P_{I_K^j}^{|\Theta^d|}$, defined in (78), accounts for all different ways of assigning each Bernoulli of each multi-Bernoulli component to a landmark for different number of landmarks, and α^j is a function that maps each Bernoulli component of the j -th multi-Bernoulli into either a landmark $l \in \{1, 2, \dots, |\Theta^d|\}$ or clutter. In addition, $\Theta_{\alpha^j(i)}$ is described by (79) and the index K denotes the final time step.

The parameters of the (j, i) -th Bernoulli component of the MBM density corresponding to the set \mathbf{Z} , the cell $\mathbf{C}^{j,i}$ and subsets $\mathbf{V}_k^{j,i}$ are given by

$$f_K^{j,i}(\theta) = \frac{f_0^u(\theta) \prod_{k=1}^K l_k^{j,i}}{\langle f_0^u; \prod_{k=1}^K l_k^{j,i} \rangle} \quad (31)$$

$$r_K^{j,i} = \begin{cases} \frac{L_K^{j,i}}{\lambda_{c(\mathbf{z})} + L_K^{j,i}} & \text{if } |\mathbf{C}^{j,i}| = 1 \\ 1 & \text{if } |\mathbf{C}^{j,i}| > 1 \end{cases}, \quad (32)$$

where

$$l_k^{j,i} = \begin{cases} p_D^k(\theta) e^{-\gamma_k(\theta)} \prod_{\mathbf{z} \in V_k^{j,i}} \gamma_k(\theta) \phi(\mathbf{z}|\theta) & \mathbf{V}_k^{j,i} \neq \emptyset \\ q_D^k(\theta) & \mathbf{V}_k^{j,i} = \emptyset \end{cases} \quad (33)$$

$$L_K^{j,i} = \lambda_0^u \langle f_0^u; \prod_{k=1}^K l_k^{j,i} \rangle \quad (34)$$

and $q_D^k(\theta) = 1 - p_D^k(\theta) + p_D^k(\theta) e^{-\gamma_k(\theta)}$, is the effective probability of missed detection at time k . The weight of each multi-Bernoulli component, W_K^j is expressed as

$$W_K^j \propto \prod_i \mathcal{L}_K^{j,i} \quad (35)$$

where $\mathcal{L}_K^{j,i}$ is the likelihood of the MB which corresponds to cell i in partition j . This likelihood depends on $|\mathbf{C}^{j,i}|$ and is described by

$$\mathcal{L}_K^{j,i} = \begin{cases} \lambda_c c(\mathbf{z}) + L_K^{j,i} & |\mathbf{C}^{j,i}| = 1 \\ L_K^{j,i} & |\mathbf{C}^{j,i}| > 1 \end{cases} \quad (36)$$

The PPP, which corresponds to the density of undetected landmarks, does not depend on the measurements. The Poisson rate and the spatial density of the PPP are expressed as

$$\lambda_K^u = \lambda_0^u \langle f_0^u; \prod_{k=1}^K q_D^k \rangle \quad (37)$$

$$f_K^u(\theta) = \frac{f_0^u(\theta) \prod_{k=1}^K q_D^k(\theta)}{\langle f_0^u; \prod_{k=1}^K q_D^k \rangle}. \quad (38)$$

In the MBM, each partition of the measurements indexed by j represents an MB. Each cell in the j th partition (MB) is denoted by $\mathbf{C}^{j,i}$, corresponds to a Bernoulli component, and is assumed to contain measurements belonging to the same source (landmark). The probability of existence, expressed in (32), is one for the Bernoulli components that have been assigned more than one measurement, i.e., the Bernoulli components which correspond to $|\mathbf{C}^{j,i}| > 1$. This is because, the clutter process generates independent point measurements which implies that two clutter measurements cannot be assigned to the same cell. Accordingly, the probability of existence for Bernoulli components associated with a single measurement is less than one, since it is not certain whether this measurement is generated by clutter or a landmark. According to (31), the spatial density of each Bernoulli component of the MBM originates from the spatial density of an undetected landmark, and is calculated by updating the spatial density of the undetected landmark with all the measurements assigned to the corresponding Bernoulli component. Additionally, the weight of each MBM, described in (35), is proportional to a product of the likelihood of each cell.

In (37), the term $\langle f_0^u; \prod_{k=1}^K q_D^k \rangle$ is the probability that a landmark with density $f_0^u(\theta)$ is not detected from time 1 to K . The effective probability of missed detection at time k is q_D^k , which according to (33) could also be regarded as the likelihood of an empty measurement set at time k . This probability is 1 outside the FOV and lower inside. The rate of the PPP describing undetected landmarks is decreased inside

the union of the FOVs that have been surveyed by the sensor. This rate should be lowest over the areas which have been surveyed the most. Outside the union of the surveyed FOVs, λ_K^u does not change. The spatial density of the undetected landmarks given by (38) expresses that this density does not change outside the union of the FOVs. Furthermore, the changes to the spatial density inside the FOVs indicate that the expected number of detections, parametrized in this density through ω , should be decreased inside the FOVs.

B. Approximations

A major difficulty in computing the PMBM density in (30) is the second summation which has N_K terms, i.e., the Bell number of order $|\mathbf{Z}|$. Since the sequence of Bell numbers grows very fast and set \mathbf{Z} corresponds to the measurements collected over many time steps, N_K is possibly a very large number. For example, for only 10 measurements, $N_K = 115975$.

Given the level of complexity, computing the posterior density exactly is intractable and there is a need to incorporate appropriate approximations into the computations. Approximating the posterior distribution is the central algorithmic problem of Bayesian inference. Most Bayesian inference algorithms can be grouped into variational methods and MCMC methods. Variational inference approximates the posterior with a simpler family of distributions, whose parameters are chosen to minimize the Kullback-Leibler divergence between the approximation and the true posterior [24] [30]. In contrast, in MCMC methods, a Markov chain is defined over the hidden variables such that the equilibrium distribution of this chain is the true posterior [31] [22] [32]. Drawing samples from the chain eventually results in obtaining samples from the posterior. The advantage of MCMC methods is that the sample distribution is generally guaranteed to converge to the posterior provided that we draw enough samples. Variational inference does not have such property [30].

In Section V, we present an inference algorithm based on Gibbs sampling. Gibbs sampling is a simple form of MCMC sampling where the Markov chain is constructed by sampling from conditional distributions of each hidden variable given the rest of the variables and the measurements.

V. SAMPLING THE PARTITIONS

In this section, we present an MCMC based inference method. As the parameter space of the map can be quite large, sampling from those parameters could result in a slow and inefficient inference algorithm. However, using certain assumptions, the model structure enables us to marginalize all the continuous variables analytically, which has been demonstrated to substantially improve the mixing properties of the Gibbs sampler [23]. Therefore, inspired by the same situation in Dirichlet process inference methods which deal with the situation where the number of classes is unknown [23], we devise an efficient algorithm to sample from the distribution of the partitions, i.e., the distribution of the data association hypotheses over the complete batch of measurements.

Each sample is a partition of the measurements which corresponds to an MB in the MBM described in (30). That is, each sample corresponds to a distribution of the detected landmarks. To calculate an estimated map, we perform an averaging over the landmarks present in the substantial number of the final samples of the chain (avoiding the burn-in samples).

In this section we focus on inferring the posterior density of the detected landmarks, i.e.,

$$f_{\theta^d}(\Theta^d) = \sum_{j=1}^{N_K} W_K^j \sum_{\alpha^j \in P_{I_K^j}^n} \prod_{i=1}^{I_K^j} f_K^{j,i}(\Theta_{\alpha^j(i)}), \quad (39)$$

where the index j describes the partitioning \mathcal{P} of the entire set of data. We can now give j a probabilistic interpretation, and observe that the joint distribution over the partition and the detected landmarks, described by

$$f(\Theta^d, j) = W_K^j \sum_{\alpha^j \in P_{I_K^j}^n} \prod_{i=1}^{I_K^j} f_K^{j,i}(\Theta_{\alpha^j(i)}), \quad (40)$$

matches the marginal distribution in (39).

Accordingly, the probability mass function of the partitions can be obtained by marginalizing out the detected landmarks,

$$\Pr\{j\} = W_K^j. \quad (41)$$

This distribution describes the probabilities of partitions of the measurements, i.e., the probabilities of different data association hypotheses over the batch of the measurements, which are equal to the weights of each multi-Bernoulli component W_K^j . The ability to compute the weights, at least up to a proportionality constant, is key to a sampling method that samples from the partitions.

In this section a closed form solution for the computation of the weights (up to a proportionality constant) is described and a Gibbs sampler which can sample from the partitions is proposed.

A. Computing the weights

Computing the weights in (41) is possible if the integral in (34) has a closed form solution. To find a closed form solution the following assumptions have been made: the probability of detection is constant inside the field of view, and the parameters of the map are a priori distributed according to a conjugate prior form expressed by

$$\begin{aligned} p(\mu) &= \mathcal{U}(\mu), \\ p(\Sigma) &= \mathcal{IW}(\Sigma; \mathbf{S}_0, \nu_0) \\ p(\omega) &= \mathcal{GAM}(\omega; \alpha_0, \beta_0), \end{aligned} \quad (42)$$

where $\mathcal{U}(\mu)$ is a uniform distribution over the AOI. If we denote this region by \mathcal{A} , the uniform distribution can be expressed as $\mathcal{U}(\mu) = \frac{\mathcal{I}(\mathcal{A})}{V_{\mathcal{A}}}$, where $\mathcal{I}(\mathcal{A})$ is an indicator function over \mathcal{A} and $V_{\mathcal{A}}$ is the volume of this region. In addition, $\mathcal{IW}(\cdot)$ denotes an inverse Wishart density and $\mathcal{GAM}(\cdot)$ denotes a gamma distribution. Using the above assumptions, the calculation of the weights is tractable.

The integral in (34) corresponds to having updated parameters of a landmark with the measurements in cell $\mathbf{C}^{j,i}$, that is with multiple empty and non-empty subsets $\mathbf{V}_k^{j,i}$. The result of an update with multiple non-empty $\mathbf{V}_k^{j,i}$ is a normal inverse Wishart gamma distribution whose parameters have been updated by the measurements in these subsets.

Empty $\mathbf{V}_k^{j,i}$ s correspond to missed detection events. Such events only affect the parameters of the resulting gamma distribution. Each missed detection creates two hypotheses regarding the distribution of the expected number of detections from a landmark. The two hypotheses describe two different ways by which our measurement model accounts for an empty set of detections. One of the hypotheses corresponds to events such as occlusion in the FOV, in which case we do not need to change our prior belief regarding the expected number of detections from a landmark. The other hypothesis states that the landmark in the FOV has generated zero measurements thus we should expect fewer detections from this landmark, and this information should be used to update our prior belief.

Accordingly, the result of an update by multiple empty and non-empty sets of measurements is a posterior with unimodal normal-inverse Wishart part, and multi modal gamma part. The gamma has 2^{N_\emptyset} modes, where N_\emptyset is the number of empty sets in $\mathbf{C}^{j,i}$ for which $f(\mu, \mathbf{x}_k) = 1$. The multimodal gamma distribution should be reduced to a unimodal gamma distribution, to avoid computational complexity issues. In this paper we have used a simple merging technique which works well for high p_D and low true landmark Poisson rate. In scenarios with low p_D and/or high landmark Poisson rate a more accurate gamma mixture reduction technique, such as the merging technique based on Kullback-Leibler divergence [33] presented in [34], can be used.

Using the above assumptions, a landmark updated by the measurements in the cell $\mathbf{C}^{j,i}$ will have a normal gamma inverse Wishart distribution expressed as,

$$\mathcal{N}(\mu; \mu_K^{j,i}, (c_K^{j,i})^{-1} \Sigma) \mathcal{GAM}(\omega; \alpha_K^{j,i}, \beta_K^{j,i}) \mathcal{IW}(\Sigma; \mathbf{S}_K^{j,i}, \nu_K^{j,i}),$$

whose parameters are given by

$$\begin{aligned} \nu_K^{j,i} &= \nu_0 + |\mathbf{C}^{j,i}| - 1 \\ \mathbf{S}_K^{j,i} &= \mathbf{S}_0 + \sum_{\mathbf{z} \in \mathbf{C}^{j,i}} (\mathbf{z} - \bar{\mathbf{z}})(\mathbf{z} - \bar{\mathbf{z}})^T \\ \mu_K^{j,i} &= \bar{\mathbf{z}} \\ c_K^{j,i} &= |\mathbf{C}^{j,i}| \\ \alpha_K^{j,i} &= \alpha_0 + |\mathbf{C}^{j,i}| \\ \beta_K^{j,i} &= \beta_0 + N_1 + \sum_{n=0}^{N_\emptyset} \binom{N_\emptyset}{n} (p_D)^n (1 - p_D)^{N_\emptyset - n} \\ &= \beta_0 + N_1 + p_D N_\emptyset \end{aligned} \quad (43)$$

where $\bar{\mathbf{z}}$ is the mean of the measurements in cell $\mathbf{C}^{j,i}$, and $\beta_K^{j,i}$ is calculated by a weighted sum over all the hypotheses due to missed detections under the assumption of having a constant probability of detection for all landmarks in the FOV, i.e., $p_D^k = p_D$ for all k . In addition, N_1 is the cardinality of a set of time stamps k for which $\mathbf{V}_k^{j,i} \neq \emptyset$. Accordingly the integral

in (34) which corresponds to the weight of the j th partition of the measurements is expressed as

$$W_K^j \propto \prod_i \frac{\prod_{k: \mathbf{v}_k^{j,i} \neq \emptyset} p_D^k}{\mathbf{V}_A} \frac{\beta_0^{\alpha_0} \Gamma(\alpha_K^{j,i})}{(\beta_K^{j,i})^{\alpha_K^{j,i}} \Gamma(\alpha_0)} \times \frac{|\mathbf{S}_0|^{\nu_0/2} \Gamma_2(\nu_K^{j,i}/2)}{\pi^{|\mathbf{C}^{j,i}|-1} (c_K^{j,i})^{0.5} \Gamma_2(\nu_0/2) |\mathbf{S}_K^{j,i}|^{\nu_K^{j,i}/2}}. \quad (44)$$

The results in (43) and (44) are derived using similar mathematical derivations as [28] and [35].

In this section the probability mass function of the partition function was derived. We showed that the probability of each partitioning of the measurements is equal to W_K^j . In addition, it was shown that with certain assumptions regarding the prior densities and the likelihood of the measurements, these weights can be calculated up to a proportionality constant. In the remaining of the section, we propose a Gibbs sampler to sample from the partition function.

B. A Gibbs sampler

The goal is to sample from the partitions. Each partitioning of the measurements corresponds to a different multi-Bernoulli component in the MBM. To sample from the partitions we need to devise a method that facilitates moving from one partition to the next. In this section we present one such method. Let $\mathcal{P}^{(t)}$ be the partition at the t -th iteration of the Gibbs sampler. Assume that $\mathcal{P}^{(t)} = \mathcal{P}^{(j)}$, where $\mathcal{P}^{(j)}$ is a valid partition with cells $\mathbf{C}^{j,i}$. The cells are indexed by $i \in \{1, 2, \dots, n_{\mathcal{P}^{(j)}}\}$, where $n_{\mathcal{P}^{(j)}} = |\mathcal{P}^{(j)}|$ is the number of cells in partition $\mathcal{P}^{(j)}$. Note that a partition is invariant to the ordering of the cells, i.e., if we permute the cell indices we have the same partition.

The $(t+1)$ -th partition is obtained as follows. First a measurement \mathbf{z} is randomly selected from the measurement set. Assume without loss of generality that the selected measurement belongs to the ℓ -th cell, $\mathbf{z} \in \mathbf{C}^{j,\ell}$. New partitions can be obtained by performing “actions” that involve the selected measurement \mathbf{z} . Specifically, we consider the following actions:

- Move \mathbf{z} from $\mathbf{C}^{j,\ell}$ to the m -th cell $\mathbf{C}^{j,m}$, where $m \in \{1, 2, \dots, n_{\mathcal{P}^{(j)}}\}$. We denote the resulting partition as $\mathcal{P}_m^{(j')}$.
- Move \mathbf{z} from $\mathbf{C}^{j,\ell}$ to a new cell. We denote the resulting partition as $\mathcal{P}_0^{(j')}$.

The probability of moving to a partition is

$$\Pr \left\{ \mathcal{P}^{(t+1)} = \mathcal{P}_m^{(j')} \mid \mathcal{P}^{(t)} = \mathcal{P}^{(j)}, \mathbf{z} \in \mathbf{C}^{j,\ell} \right\} = \frac{W_K^{\mathcal{P}_m^{(j')}}}{\sum_{m'=0}^{n_{\mathcal{P}^{(j)}}} W_K^{\mathcal{P}_{m'}^{(j')}}} \quad (45)$$

for $m = 0, 1, \dots, n_{\mathcal{P}^{(j)}}$, where $W_K^{\mathcal{P}_m^{(j')}}$ is the weight of the partition (44). Note that $m = 0$ means that we get the same partition, i.e., this action corresponds to $\mathcal{P}^{(t+1)} = \mathcal{P}^{(t)}$. Further, for the special case $|\mathbf{C}^{j,\ell}| = 1$, moving \mathbf{z} to a new cell results in an equivalent partition, i.e., if $|\mathbf{C}^{j,\ell}| = 1$ and $m = \ell$ then $\mathcal{P}^{(t+1)} = \mathcal{P}^{(t)}$. Because of this, to avoid double

counting the probability of $\mathcal{P}^{(t+1)} = \mathcal{P}^{(t)}$, if $|\mathbf{C}^{j,\ell}| = 1$ we set

$$\Pr \left\{ \mathcal{P}^{(t+1)} = \mathcal{P}_\ell^{(j')} \mid \mathcal{P}^{(t)} = \mathcal{P}^{(j)}, \mathbf{z} \in \mathbf{C}^{j,\ell} \right\} = 0. \quad (46)$$

The algorithm can be initialized with any valid partition. Two simple examples are starting with a partition with all the measurements in individual cells, or starting with a partition with all measurements in a single cell. In the simulation section, we use the former alternative.

The sum in the denominator of (45) might pose a challenge as it could include a large number of cells. There are two relieving aspects for the calculation of this sum. First, we expect many of the terms to be zero as they could correspond to infeasible data associations. For example, if the time tag of \mathbf{z} is k and the landmark formed by the measurements in $\mathbf{C}^{j,m}$ is out of the field of view at time k , then it can be concluded that the corresponding $W_K^{\mathcal{P}_m^{(j')}} = 0$. Such cases can be removed from the possible partitions. Second, since each move only changes two cells, we can reuse the parameters calculated for the other cells when calculating the weight of the newly formed partition, i.e.,

$$\begin{aligned} W_K^{\mathcal{P}_m^{(j')}} &\propto \prod_{\tau} \mathcal{L}_K^{\mathcal{P}_m^{(j')}, \tau} \\ &= \mathcal{L}_K^{\mathcal{P}_m^{(j')}, \ell} \mathcal{L}_K^{\mathcal{P}_m^{(j')}, m} \prod_{\tau: \tau \neq \ell, m} \mathcal{L}_K^{\mathcal{P}_m^{(j')}, \tau} \\ &= \mathcal{L}_K^{\mathcal{P}_m^{(j')}, \ell} \mathcal{L}_K^{\mathcal{P}_m^{(j')}, m} \prod_{\tau: \tau \neq \ell, m} \mathcal{L}_K^{\mathcal{P}^{(j)}, \tau} \end{aligned} \quad (47)$$

where $\mathcal{L}_K^{\mathcal{P}_m^{(j')}, \tau}$ denotes the likelihood of cell τ of partition $\mathcal{P}_m^{(j')}$. Note that when a measurement is moved out of a cell, the cell might become empty and we define, $\mathcal{L}_K^{\mathcal{P}_m^{(j')}, \emptyset} \triangleq 1$.

C. Parallelization

The proposed sampling method can be parallelized. This can be viewed as a result of the product in (35), limited FOV of the sensor and the size of the landmarks. For example, measurements that have been detected in two non-overlapping field of views have zero probability of belonging to the same landmark. This implies that we can process two such groups of measurements in parallel without having to introduce any approximations. Moreover, measurements that are far away (relative to our current belief of the size of the landmarks) will have very low probability of belonging to the same landmark. While this probability could be non-zero in practice, we can use this information to obtain further parallelization by gating those measurements. This results in an approximated but accelerated sampling procedure.

D. Comparison to BNP

One commonly used prior model in BNPs is the Dirichlet process [21], which is typically applied as a building block in the Dirichlet process mixture model (DPMM). This model assumes that there is an infinite number of components, of which a finite number have generated the finite dataset at hand.

This type of modelling results in a posterior distribution that assigns higher probability mass to the number of components that best explains the observations. This avoids the need to pre-specify the number of components before observing the data [19] [20].

In this paper, the map has been assigned a Poisson process prior. Similarly to the DPMM, our model allows us to infer the number of components (landmarks) given the data. The difference is that the Poisson prior assumes that there exists a finite and unknown number of landmarks a priori.

The Poisson prior model of the map is motivated by its favorable conjugate properties to the standard measurement model of extended landmarks. In [26] it is shown that the PMBM prior is conjugate to the PPP extended target measurement model. In a sequential setting, at the first time step, our Poisson prior is updated by this measurement model, resulting in a PMBM posterior. Consequently, the posterior will maintain its PMBM form in the subsequent time steps. Similarly, in a batch setting the Poisson prior and the PPP measurement model result in a PMBM posterior. Therefore, while the Poisson model for the map ensures a closed form posterior, it is not clear how a Dirichlet process prior would shape the posterior given the PPP measurement model.

The Gibbs sampling method used in this paper is similar to the methods discussed in [23] [36], where the associations of measurements to components are updated sequentially one at a time, conditioned on the data and the rest of the associations. This results in an MCMC chain whose equilibrium distribution is the posterior density of the component associations. This conditional probability can be viewed as the probability of a particular partitioning of the measurements, since one can move from one partitioning of the measurements to another by changing a component (landmark) assignment of one measurement. From this perspective, the equilibrium distribution of the MCMC chain is the posterior density of the measurements' partitions. In the proposed method we have derived the probability distribution of partitions of the measurements formed by changing the landmark assignment of one measurement.

VI. ESTIMATION OF THE MAP

In this section we describe how the result of the sampling method is used for estimating the map of the environment.

We assume that the chain has run long enough for convergence, as a result the final samples are samples of the posterior density of the partitioning of the measurements. That is, the number of sampling iterations should be chosen such that the burn-in period of the sampler is avoided. Given a partition of the measurements, the association of the measurements to landmarks is known. Therefore, each sample of the posterior corresponds to a map of the environment. The estimated map is computed by averaging over a number of final samples of the chain. The averaging is performed over the positions, extents and the weights of the landmarks of each sample (map).

In each partition of the measurements, every cell that contains more than one measurement is considered to have been generated by one landmark. This landmark has the prior distribution described in (42) and is updated by the measurements

assigned to it according to (43). Among the cells which contain a single measurement, some are generated by landmarks. To decide whether a single measurement is generated by a landmark, a threshold is set on the probability of existence of the cells with cardinality one. Cells whose probability of existence are above the threshold are considered to have been generated by a landmark. Consequently, the parameters for these landmarks are updated by the single measurement assigned to them. According to (43), the parameters of the extent of such landmarks are not updated, since for a uniform position prior a single measurement only contains information about the position and not the extent of the landmark.

To perform the averaging over the samples, we need to solve an association problem between the cells of the samples. That is, we need to establish which cells across the samples belong to the same landmark. This is done by specifying an area around the position of each cell. The cells whose distance is less than the specified area are assumed to belong to the same landmark. To estimate the properties of a landmark, we average over the properties of the cells that belong to that landmark. We do not use spurious cells in this calculation. Spurious cells are those that exist in a very small percentage of the chosen samples. The resulting estimation method is presented in Algorithm 1.

Gibbs Sampling

Choose an initial partitioning \mathcal{P}_I of the measurements,

See Section VI.

for $n=1$: #iterations **do**

 Randomly choose a measurement \mathbf{z} from a cell.

 Evaluate the weight of possible partitions formed by moving \mathbf{z} to a different cell according to Eq. (44).

 Choose a partition \mathcal{P} according to Eq. (45).

end

Map Estimation

Choose a number of final samples of the chain, see Section VI.

for $j=1$: #Samples **do**

 sample j corresponds to partition \mathcal{P} .

for $i=1$: #Cells in \mathcal{P} **do**

if $r^{\mathcal{P},i} > \text{threshold}$ **then**

 update the parameters of the corresponding landmark according to Eq. (43).

end

end

end

Average over the landmarks of the samples, see Section VI.

Algorithm 1: Pseudo-code of the mapping algorithm.

The initial partitioning of the measurements can be chosen arbitrarily so long as the chosen partition is feasible. An infeasible partition corresponds to a partition with at least one cell containing measurements from spatially far apart areas of the map. Because the measurements in the same cell are hypothesized to belong to the same landmark, such a partition is infeasible and has zero weight. A simple choice of an initial partitioning of the measurements is a partition with single-

member cells where each measurement belongs to its own cell.

The map generated by each sample of the MCMC chain can be viewed as the intensity function of the PPP that has generated the measurements. This per sample map is later used for performance evaluation in Section VII. The intensity function is an unnormalized Gaussian mixture. The mixture corresponding to the estimated map is given by

$$\zeta(\hat{\theta}) = \sum_{l=1}^{N_E} \hat{\omega}_l \mathcal{N}(\mu_l; \hat{\mu}_l, \hat{\Sigma}_l), \quad (48)$$

where $\hat{\theta}$ is formed by choosing an arbitrary order of the estimated landmarks, and N_E is the estimated number of landmarks. Additionally, $\hat{\omega}_l$, $\hat{\mu}_l$ and $\hat{\Sigma}_l$ are the mean of the updated gamma, normal and inverse Wishart distribution of landmark l , respectively.

VII. SIMULATIONS AND RESULTS

In this section the proposed sampling algorithm is evaluated on a synthetic mapping scenario where a sensor travels over a track and collects measurements. The scenario is depicted in Figure 1. The proposed algorithm is compared to the VBEM mapping method presented in [10].

The VBEM method approximates the posterior density of the map by simpler densities and is very computationally efficient; however, this algorithm does not take into account the uncertainty in the number of landmarks. A challenging part of the VBEM is the initialization of the landmarks. The number of landmarks should be set when initializing the algorithm. For mapping a large area one needs to generate a large number of landmarks which cover the AOI reasonably. This is very important as VBEM is sensitive to initialization. The VBEM is guaranteed to find a local optimum which may be far from the global optimum. The Gibbs sampling method samples from the posterior and considers the uncertainties in the number of landmarks at a higher computational price. This algorithm is expected to converge to the true posterior given enough number of MCMC samples.

We use the integrated squared error (ISE) [37] as performance measure, which is defined as the squared L2 norm of the difference of two functions. In our problem, the first function is $\zeta(\theta)$, the non-normalized Gaussian mixture formed by the true map, and the second is $\zeta(\hat{\theta})$. The ISE is given by

$$J(\theta, \hat{\theta}) = \|\zeta(\theta) - \zeta(\hat{\theta})\|_2^2 = \int (\zeta(\theta) - \zeta(\hat{\theta}))^2 d\theta. \quad (49)$$

Here, θ is formed by an arbitrary order of the true landmarks. This measure of performance accounts for all properties of the map, i.e., positions, extents and weights of landmarks.




Figure 1. The scenario on which the algorithm is evaluated. The vehicle travels one lap over the track and collects 543 measurements in total over 190 time steps.

Furthermore, in [38] it is shown that (49) has a simple and exact form for Gaussian mixtures, which is described as

$$J(\theta, \hat{\theta}) = J_{TT} - 2J_{TE} + J_{EE} \quad (50)$$

$$J_{TT} = \sum_{l_1=1}^{N_T} \sum_{l_2=1}^{N_T} \omega_{l_1} \omega_{l_2} \mathcal{N}(\mu_{l_1}; \mu_{l_2}, \Sigma_{l_1} + \Sigma_{l_2}) \quad (51)$$

$$J_{TE} = \sum_{l_1=1}^{N_T} \sum_{l_2=1}^{N_E} \omega_{l_1} \hat{\omega}_{l_2} \mathcal{N}(\mu_{l_1}; \hat{\mu}_{l_2}, \Sigma_{l_1} + \hat{\Sigma}_{l_2}) \quad (52)$$

$$J_{EE} = \sum_{l_1=1}^{N_E} \sum_{l_2=1}^{N_E} \hat{\omega}_{l_1} \hat{\omega}_{l_2} \mathcal{N}(\hat{\mu}_{l_1}; \hat{\mu}_{l_2}, \hat{\Sigma}_{l_1} + \hat{\Sigma}_{l_2}), \quad (53)$$

where N_T is the true number of landmarks.

The following parameter values have been used in the simulations. The prior values for different parameters of the map are set to $\mathbf{S}_j^0 = 5\mathbf{I}$, $\nu_j^0 = 5$ and $\alpha_0 = 0.1$, $\beta_0 = 0.2$. The covariance of the measurement noise is set to $\mathbf{R} = 0.01^2\mathbf{I}$ and $\lambda_c = 1$. The sensor's field of view has a range of 60m and an angle of $\pm 30^\circ$. The probability of detection is set to one inside the field of view and zero outside.

The true map consists of 20 landmarks. These landmarks, together with clutter, have generated 543 measurements over 190 time scans, therefore, there are $10^{931.111}$ ways to partition this set into non-empty cells. The sampling algorithm is run for 120000 iterations, each taking on average 0.2 seconds on a desktop computer with Windows 7 professional, Intel(R) core(TM) i5 CPU 650@ 3.20 GHZ, and 8 GB of RAM. We have chosen a high number of iterations to ensure that we avoid the burn-in period of the sampler. Our results indicate that the algorithm converges earlier (see below).

We initialize the algorithm with one cell per measurement. This corresponds to each detection being an individual landmark or clutter, depending on the probability of existence. Since this setting corresponds in general to a very unlikely configuration (under the posterior), we considered a burn-in period to remove the contribution of the initial samples of the




Figure 2. Comparison between the true and the estimated positions of the landmarks for the Gibbs sampling method.




Figure 3. A magnified area of the map estimated by the Gibbs sampling method, depicting the measurements, the true and the estimated landmarks.

algorithm. We chose this initialization scheme to show that our algorithm does converge even from a poor starting point.

The map is estimated by averaging over the final 40000 samples of the chain. In each sample of the map, the cells with existence probability higher than 0.5 are considered as landmarks. The estimated position of the landmarks is compared to the true positions in Figure 2. It can be seen that most of the estimated positions are very close to the true ones. In Figure (3), an area from the lower left corner of the map has been magnified to illustrate the true map, the measurements and the estimated map of the area. Figure 4 depicts the posterior Poisson rate of the undetected landmarks calculated according to (37), the sensor trajectory and the detected landmarks.

The estimated expected number of clutter measurements is calculated by averaging over the final 40000 samples of the MCMC chain. In each sample, the single-member cells with probability of existence lower than 0.5 are assumed to be clutter. In addition, for each sample, the total number of clutter cells divided by the number of time steps is equal to the expected number of clutter measurements. Following this averaging method, the estimated expected number of clutter measurements is 0.7626.

Figure 6 depicts a comparison of the estimated number of landmarks per iteration between the VBEM method described in [10] and our method. The number of landmarks in the VBEM is considered to be those which have weights larger than 0.01 and in the Gibbs sampling the cells with probability of existence larger than 0.5. It should be noted that the chosen weight threshold for the VBEM is the same as the one used in [10] to estimate the number of landmarks. We can see that the Gibbs sampling method provides a more accurate estimate of the number of landmarks. The parameters of the VBEM are set to the same values as [10]. The VBEM is initialized with 300 landmarks uniformly distributed over the AOI (recall that the Gibbs sampling method is initialized with 543 single-member cells).

The ISE of the map corresponding to each sample of the




Figure 4. The posterior Poisson rate of the undetected landmarks and the estimated detected landmarks illustrated in one figure. The background is the posterior Poisson rate of the undetected landmarks, where the brighter the color the more frequently the area has been in the FOV and, consequently, the smaller the expected number of undetected landmarks. Accordingly, the darker the background, the less frequently the area has been observed by the sensor (covered by the FOV over multiple scans) and the higher the Poisson rate. On top of the background, we have depicted the sensor trajectory and the detected landmarks, marked by ellipses corresponding to their extent.

MCMC chain is compared to the ISE resulting from each iteration of the VBEM in Figure 5. It can be seen that the Gibbs sampling method has a lower ISE compared to the VBEM. We can see that the Gibbs sampling algorithm provides a more accurate map both in terms of the estimated number of landmarks and the ISE.

VIII. CONCLUSIONS

In this paper, we derive the exact theoretical batch multi-object posterior density of the map using a conjugate prior form described in [25]. We propose a solution that uses an extended object model and is capable of handling the uncertainties in the cardinality of the set of landmarks. We




Figure 5. Comparison between the ISE of the VBEM and the ISE of the proposed algorithm.




Figure 6. The estimated number of landmarks per iteration for VBEM and the proposed algorithm.

use Gibbs sampling to estimate the batch multi-object posterior density of the map. Our method is analogous to well-known sampling approaches in BNPs. The results show that the proposed algorithm can estimate the number of landmarks as well as their properties and outperforms the VBEM method. In addition, we discuss possible parallelization of the proposed algorithm, which makes it suitable for large scale problems. A parallel implementation of the proposed method is a topic for future investigations. Preliminary results indicate that it is possible to devise a sequential mapping algorithm based on the presented batch mapping algorithm. An interesting future research line is to continue this work and compare a complete and well-designed sequential mapping algorithm with the batch algorithm.

ACKNOWLEDGMENTS

Francisco J. R. Ruiz acknowledges the EU H2020 program (Marie Skłodowska-Curie grant agreement 706760) which is supported by ONR N00014-11-1-0651.

APPENDIX A RFS BACKGROUND

If the RFS of interest is denoted by $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$, the RFS density is expressed as

$$f(\Theta) = p(n) \sum_{\eta} f_n(\theta_{\eta(1)}, \dots, \theta_{\eta(n)}) \quad (54)$$

where $p(n)$ is a cardinality distribution, $f_n(\theta_{\eta(1)}, \dots, \theta_{\eta(n)})$ is a cardinality-conditioned joint distribution and η accounts for all $n!$ possible permutations ensuring that $f(\Theta)$ is permutation invariant.

P.g.fl.s provide an alternative representation of a RFS density, one that is often simpler to deal with. The p.g.fl of a RFS density is a transform of the density defined through the set integral as [27, p. 371]

$$G[h] = \int h^{\Theta} f(\Theta) \delta\Theta \quad (55)$$

where $h(\theta)$ is a test function, $h^{\Theta} \triangleq \prod_{\theta \in \Theta} h(\theta)$ and a set integral is

$$\int \tau(\Theta) \delta\Theta \triangleq \tau(\emptyset) + \sum_{n=1}^{\infty} \frac{1}{n!} \int \tau(\{\theta_1, \theta_2, \dots, \theta_n\}) d\theta_1 \dots d\theta_n. \quad (56)$$

The RFS density of a union of two independent RFS $\Theta = \Theta_1 \cup \Theta_2$ is described by

$$f(\Theta) = \sum_{\Theta_1 \subseteq \Theta} f_{\Theta_1}(\Theta_1) f_{\Theta_2}(\Theta - \Theta_1) \quad (57)$$

and the corresponding p.g.fl is [27, p. 372]

$$G_{\Theta}[h] = G_{\Theta_1}[h] G_{\Theta_2}[h]. \quad (58)$$

Additionally, a RFS density can be derived from its p.g.fl by [27, pp. 375-376, 384]

$$\begin{aligned} f(\Theta) &= \frac{\delta}{\delta\Theta} G[h]|_{h=0} \\ &= \frac{\delta^{|\Theta|}}{\prod_{\theta \in \Theta} \delta\theta} G[h]|_{h=0}, \end{aligned} \quad (59)$$

where

$$\frac{\delta}{\delta\theta} G[h] \triangleq \lim_{\epsilon \downarrow 0} \frac{G[h + \epsilon\delta_{\theta}] - G[h]}{\epsilon},$$

the cardinality of the set Θ is denoted by $|\Theta|$ and $\delta_{\theta}(\theta')$ is a Dirac delta function centred at $\theta' = \theta$.

Our model makes use of two random processes, namely, the PPP and the MBM. The RFS density of these processes together with their p.g.fl.s have been previously derived in [27] and discussed in [39] [25].

The RFS density and the p.g.fl of a PPP RFS are described by [27]

$$\begin{aligned} f(\Theta) &= \exp(-\lambda) \prod_{\theta \in \Theta} \lambda f(\theta) \\ G[h] &= \exp(\lambda \langle f; h \rangle - \lambda) \end{aligned} \quad (60)$$

where $\langle f; h \rangle = \int f(\theta) h(\theta) d\theta$ and the cardinality of Θ is distributed according to a Poisson distribution with rate λ .

Further, the members of Θ are independent and identically distributed (iid) according to $f(\theta)$.

A Bernoulli RFS has the following RFS density and p.g.fl,

$$f(\Theta) = \begin{cases} 1 - r, & \Theta = \emptyset \\ rf(\theta), & \Theta = \theta \\ 0, & \text{otherwise} \end{cases} \quad (61)$$

$$G[h] = 1 - r + r\langle f; h \rangle$$

where r is the probability of existence and $f(\theta)$ is the existence-conditioned distribution.

A multi-Bernoulli process is formed by a union of N independent Bernoulli processes, $\Theta = \bigcup_{i=1}^N \Theta_i$. The p.g.fl of this process resulting from (58) and (61) is expressed as

$$G[h] = \prod_i (1 - r_i + r_i \langle f_i; h \rangle). \quad (62)$$

The RFS density of this process can be described as [39]

$$f(\{\theta_1, \theta_2, \dots, \theta_n\}) = \sum_{\alpha \in P_N^n} \prod_{i=1}^n f_i(\Theta_{\alpha(i)}), \quad (63)$$

where

$$P_N^n = \{\alpha : \{1, \dots, N\} \rightarrow \{0, \dots, n\} \mid \{1, \dots, n\} \subset \alpha(\{1, \dots, N\}), \text{ if } \alpha(i) > 0, i \neq l \text{ then } \alpha(i) \neq \alpha(l)\} \quad (64)$$

and

$$\Theta_{\alpha(i)} = \begin{cases} \emptyset & \alpha(i) = 0 \\ \{\theta_{\alpha(i)}\} & \alpha(i) > 0 \end{cases}. \quad (65)$$

That is, P_N^n accounts for all different ways of assigning the N component of the multi-Bernoulli to the n landmarks. In addition, each assignment α is a function that maps each Bernoulli component into either a landmark $l \in \{1, 2, \dots, n\}$ or clutter, which is indexed by zero.

The RFS density of a MBM process is formed by a normalized sum over multi-Bernoulli (MB) RFSs, that is,

$$f(\Theta) = \sum_j W^j f^j(\Theta) \quad (66)$$

where each weight W^j is related to one data association hypothesis, $\sum_j W^j = 1$, and each $f^j(\Theta)$ has the form presented in (63). In addition, the p.g.fl of the MBM process in (66) is formed by a weighted sum of the p.g.fl of MB RFSs described by (62),

$$G[h] = \sum_j W^j \prod_i (1 - r^{j,i} + r^{j,i} \langle f^{j,i}; h \rangle). \quad (67)$$

A. P.g.fl form of multi-object posterior density

In this section, the p.g.fl of the multi-object posterior is presented for both sequential and batch processing. The p.g.fl of the sequential posterior density can be calculated using the p.g.fl of the corresponding measurement model. Accordingly,

this p.g.fl at time k given the RFS of the measurements up to and including time k , $\mathbf{Z}_{1:k}$, is calculated by [27]

$$G_{k|k}[h] = \frac{\delta F[g, h] |_{g=0}}{\delta \mathbf{Z}_k |_{\dot{g}=0, \dot{h}=0}} \quad (68)$$

$$F[g, h] \triangleq \int h^{\Theta_k} G_k[g | \Theta_k] f_{k|k-1}(\Theta_k | \mathbf{Z}_{1:k-1}) \delta \Theta_k \quad (69)$$

$$G_k[g | \Theta_k] \triangleq \int g^{\mathbf{Z}_k} f_k(\mathbf{Z}_k | \Theta_k) \delta \mathbf{Z}_k \quad (70)$$

where $f_{k|k-1}(\Theta | \mathbf{Z}_{1:k-1})$ is the predicted density of the state. Using this p.g.fl, the posterior density of the RFS of landmarks can be obtained. The complete derivation of this p.g.fl under certain assumptions is given in [26] and summarized in Section III.

The batch posterior p.g.fl can be expressed as

$$G_{K|K}[h] = \frac{\delta F[g_1, g_2, \dots, g_K, h] |_{g_1=0, g_2=0, \dots, g_K=0}}{\delta \mathbf{Z}_1, \mathbf{Z}_2, \dots, \delta \mathbf{Z}_K |_{\dot{g}_1, \dot{g}_2, \dots, \dot{g}_K, \dot{h}=0}} \quad (71)$$

where

$$F[g_1, g_2, \dots, g_K, h] \triangleq \int h^{\Theta} G_K[g | \Theta] f_{0|0}(\Theta) \delta \Theta \quad (72)$$

and

$$G_K[g | \Theta] = \int g_1^{\mathbf{Z}_1} g_2^{\mathbf{Z}_2} \dots g_K^{\mathbf{Z}_K} \prod_{k=1}^K p(\mathbf{Z}_k | \Theta) \delta \mathbf{Z}_1 \dots \delta \mathbf{Z}_K \\ = \prod_{k=1}^K G_k[g_k | \Theta]. \quad (73)$$

Additionally, $p(\mathbf{Z}_k | \Theta)$ is the likelihood function at time k and $f_{0|0}(\Theta)$ is the prior density. These equations are the same as those presented in [40] in a multi-sensor set-up. Both (68) and (71) can be used to derive the batch posterior density.

APPENDIX B

THE BATCH MULTI-OBJECT POSTERIOR DENSITY

In this appendix we present the detailed derivation of the batch multi-object posterior for a static map.

A. The PMBM density

Since the map consists of two disjoint subsets of detected and undetected landmarks, the RFS density of the map can be written as [27]

$$f(\Theta) = \sum_{\Theta^d \subset \Theta} f_d(\Theta^d) f_u(\Theta \setminus \Theta^d). \quad (74)$$

Additionally, following the results regarding the relation between probability density and p.g.fl in [27] we can write

$$f_d(\Theta^d) = \frac{\delta}{\delta \Theta} G^d[h] |_{h=0} \\ = \frac{\delta}{\delta \Theta} \left(\sum_j W_K^j G^{d,j}[h] \right) |_{h=0} \\ = \sum_j W_K^j \frac{\delta}{\delta \Theta} G^{d,j}[h] |_{h=0} \quad (75)$$

where

$$G^{d,j}[h] = \prod_{i=1}^{I_K^j} (1 - r_K^{j,i} + r_K^{j,i} \langle f_K^{j,i}; h \rangle), \quad (76)$$

and

$$\frac{\delta}{\delta \Theta} G^{d,j}[h]|_{h=0} = \sum_{\alpha^j \in P_{I_K^j}^{|\Theta^d|}} \prod_{i=1}^{I_K^j} f_K^{j,i}(\Theta_{\alpha^j(i)}). \quad (77)$$

Similar to (64) and (65), the mapping $P_{I_K^j}^{|\Theta^d|}$ and the components $\Theta_{\alpha^j(i)}$ are defined as

$$\begin{aligned} P_{I_K^j}^{|\Theta^d|} &= \{\alpha^j : \{1, \dots, I_K^j\} \rightarrow \{0, \dots, |\Theta^d|\}\} \\ &\quad \{1, \dots, |\Theta^d|\} \subset \alpha(\{1, \dots, I_K^j\}), \\ \text{if } \alpha^j(i) > 0, i \neq l \quad \text{then } \alpha^j(i) &\neq \alpha^j(l) \end{aligned} \quad (78)$$

and

$$\Theta_{\alpha^j(i)} = \begin{cases} \emptyset & \alpha^j(i) = 0 \\ \{\theta_{\alpha^j(i)}\} & \alpha^j(i) > 0 \end{cases}, \quad (79)$$

where $P_{I_K^j}^{|\Theta^d|}$ accounts for all different ways of assigning each Bernoulli of each multi-Bernoulli component to a landmark for different number of landmarks. For each multi-Bernoulli, $|\Theta^d|$ cannot exceed I_K^j . In addition, α^j is a function that maps each Bernoulli component of the j -th multi-Bernoulli into either a landmark $l \in \{1, 2, \dots, |\Theta^d|\}$ or clutter. Similar to (61), the spatial density of each Bernoulli component is defined as

$$f_K^{j,i}(\Theta) = \begin{cases} 1 - r_K^{j,i} & \Theta = \emptyset \\ r_K^{j,i} f_K^{j,i}(\theta) & \Theta = \{\theta\} \\ 0 & \text{otherwise} \end{cases}. \quad (80)$$

The Poisson process accounting for the undetected landmarks is given by

$$f_u(\Theta \setminus \Theta^d) = e^{-\lambda_K^u} \prod_{\theta \in \Theta^u} \lambda_K^u f_K^u(\theta) \quad (81)$$

where λ_K^u and $f_K^u(\theta)$ are described in (37) and (38), respectively.

B. Components of the PMBM density

The probability of existence of each Bernoulli component in (32) depends on the cardinality of $\mathbf{C}^{j,i}$. By applying (15), (22) and (25) sequentially we can see that if $|\mathbf{C}^{j,i}| > 1$ then $r_K^{j,i} = 1$. Similarly, the probability of existence for $|\mathbf{C}^{j,i}| = 1$ in (32) is derived by applying (15) and (25), sequentially.

The spatial density of each Bernoulli component in (31) is derived using (16), (23) and (26).

The Poisson rate and the spatial density of undetected landmarks described in (37) and (38) are derived by sequentially applying (18) and (19), respectively.

REFERENCES

- [1] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *IEEE Robot. Autom. Mag.*, vol. 13, no. 2, 2006.
- [2] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robot. Autom. Mag.*, vol. 13, no. 3, 2006.
- [3] R. C. Smith and P. Cheeseman, "On the representation and estimation of spatial uncertainty," *The international journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986.
- [4] M. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, 2001.
- [5] M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit *et al.*, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *AAAI/IAAI*, 2002, pp. 593–598.
- [6] H. Deusch, S. Reuter, and K. Dietmayer, "The labeled multi-Bernoulli SLAM filter," *Signal Processing Letters, IEEE*, vol. 22, no. 10, pp. 1561–1565, 2015.
- [7] G. Grisetti, R. Kummerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based SLAM," *IEEE Intell. Transp. Syst. Mag.*, vol. 2, no. 4, pp. 31–43, 2010.
- [8] J. Ziegler *et al.*, "Making Bertha drive? An autonomous journey on a historic route," *IEEE Intell. Transp. Syst. Mag.*, vol. 6, no. 2, 2014.
- [9] M. Lundgren, E. Stenborg, L. Svensson, and L. Hammarstrand, "Vehicle self-localization using off-the-shelf sensors and a detailed map," in *2014 Proc. Intelligent Vehicles Symp. Proceedings*, pp. 522–528.
- [10] M. Lundgren, L. Svensson, and L. Hammarstrand, "Variational bayesian expectation maximization for radar map estimation," *IEEE Trans. Signal Process.*, vol. 64, no. 6, pp. 1391–1404, 2016.
- [11] C. Lundquist, L. Hammarstrand, and F. Gustafsson, "Road intensity based mapping using radar measurements with a probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 59, no. 4, 2011.
- [12] M. Adams, B.-N. Vo, R. Mahler, and J. Mullane, "SLAM gets a PHD: New concepts in map estimation," *IEEE Robot. Autom. Mag.*, vol. 21, no. 2, pp. 26–37, 2014.
- [13] R. P. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [14] K. Gilholm, S. Godsill, S. Maskell, and D. Salmond, "Poisson models for extended target and group tracking," in *Optics & Photonics 2005*. International Society for Optics and Photonics, 2005, pp. 59 130R–59 130R.
- [15] R. P. Mahler, "PHD filters for nonstandard targets, I: Extended targets," in *12th International Conference on Information Fusion*. IEEE, 2009, pp. 915–921.
- [16] U. Orguner, C. Lundquist, and K. Granström, "Extended target tracking with a cardinalized probability hypothesis density filter," in *14th International Conference on Information Fusion*. IEEE, 2011, pp. 1–8.
- [17] K. Granström, C. Lundquist, and O. Orguner, "Extended target tracking using a Gaussian-mixture PHD filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 48, no. 4, pp. 3268–3286, 2012.
- [18] M. J. Beal, "Variational algorithms for approximate Bayesian inference," Ph.D. dissertation, University of London, 2003.
- [19] N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker, *Bayesian nonparametrics*. Cambridge University Press, 2010, vol. 28.
- [20] S. J. Gershman and D. M. Blei, "A tutorial on Bayesian nonparametric models," *Journal of Mathematical Psychology*, vol. 56, no. 1, pp. 1–12, 2012.
- [21] T. S. Ferguson, "A Bayesian analysis of some nonparametric problems," *The annals of statistics*, pp. 209–230, 1973.
- [22] C. Robert and G. Casella, *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [23] R. M. Neal, "Markov chain sampling methods for Dirichlet process mixture models," *Journal of Computational and Graphical Statistics*, vol. 9, no. 2, pp. 249–265, 2000.
- [24] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [25] K. Granström, M. Fatemi, and L. Svensson, "Gamma Gaussian inverse-Wishart Poisson multi-Bernoulli Filter for Extended Target Tracking," in *19th International Conference on Information Fusion*. IEEE, 2016, pp. 893–900.
- [26] —. (2016, August) Poisson multi-Bernoulli conjugate prior for multiple extended object estimation. ArXiv:1605.06311. [Online]. Available: <https://arxiv.org/pdf/1605.06311>

- [27] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.
- [28] J. W. Koch, "Bayesian approach to extended object and cluster tracking using random matrices," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 44, no. 3, pp. 1042–1059, 2008.
- [29] G.-C. Rota, "The number of partitions of a set," *The American Mathematical Monthly*, vol. 71, no. 5, pp. 498–504, May 1964.
- [30] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Foundations and Trends in Machine Learning*, vol. 1, no. 1–2, pp. 1–305, Jan. 2008.
- [31] J. S. Liu, *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag New York, 2004.
- [32] A. B. Owen, *Monte Carlo theory, methods and examples*, 2013.
- [33] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Mathematical Statistics*, pp. 79–86, 1951.
- [34] K. Granström and U. Orguner, "Estimation and maintenance of measurement rates for multiple extended target tracking," in *15th International Conference on Information Fusion*. IEEE, 2012, pp. 2170–2176.
- [35] —, "A PHD filter for tracking multiple extended targets using random matrices," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5657–5671, 2012.
- [36] S. N. MacEachern, "Estimating normal means with a conjugate style Dirichlet process prior," *Communications in Statistics-Simulation and Computation*, vol. 23, no. 3, pp. 727–741, 1994.
- [37] J. L. Williams, "Gaussian mixture reduction for tracking multiple maneuvering targets in clutter," DTIC Document, Tech. Rep., 2003.
- [38] D. F. Crouse, P. Willett, K. Pattipati, and L. Svensson, "A look at Gaussian mixture reduction algorithms," in *14th International Conference on Information Fusion*. IEEE, 2011, pp. 1–8.
- [39] J. L. Williams, "Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA and association-based MeMBer," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 3, pp. 1664–1687, 2015.
- [40] R. Mahler, "The multisensor PHD filter, I: General solution via multitarget calculus," in *Proc. Signal Processing, Sensor Fusion, and Target Recognition XVIII*, vol. 7336. SPIE, 2009.



Maryam Fatemi received the M.Sc. degree in communication systems engineering from AmirKabir University of Technology, Tehran, Iran in 2008, and received the Ph.D. degree in electrical engineering from Chalmers University of Technology, Gothenburg, Sweden in 2016. Since November 2016 she has joined Autoliv Sverige AB. Her research interests include Bayesian Inference and nonlinear filtering with applications to sensor data fusion, autonomous driving and active safety systems.



Karl Granström (M'08) is a postdoctoral research fellow at the Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden. He received the MSc degree in Applied Physics and Electrical Engineering in May 2008, and the PhD degree in Automatic Control in November 2012, both from Linköping University, Sweden. He previously held postdoctoral positions at the Department of Electrical and Computer Engineering at University of Connecticut, USA, from September 2014 to August 2015, and at the Department of

Electrical Engineering of Linköping University from December 2012 to August 2014. His research interests include estimation theory, multiple model estimation, sensor fusion and target tracking, especially for extended targets. He has received paper awards at the Fusion 2011 and Fusion 2012 conferences.



Lennart Svensson was born in Älvängen, Sweden in 1976. He received the M.S. degree in electrical engineering in 1999 and the Ph.D. degree in 2004, both from Chalmers University of Technology, Gothenburg, Sweden.

He is currently Professor of Signal Processing, again at Chalmers University of Technology. His main research interests include machine learning and Bayesian inference in general, and nonlinear filtering and tracking in particular.



Francisco J. R. Ruiz is a Postdoctoral Research Scientist who works with David Blei at the Data Science Institute and the Department of Computer Science at Columbia University, and with Zoubin Ghahramani at the Engineering Department of the University of Cambridge. Francisco holds a Marie Skłodowska-Curie fellowship in the context of the E.U. Horizon 2020 program. He completed his Ph.D. and M.Sc. from the University Carlos III in Madrid. His research is focused on statistical machine learning, specially Bayesian modeling and inference.



Lars Hammarstrand was born in Landvetter, Sweden in 1979. He received his M.Sc. and Ph.D. degree in electrical engineering from Chalmers University of Technology, Gothenburg, Sweden, in 2004 and 2010, respectively.

Between 2004 and 2011, he was with the Active Safety and Chassis Department at Volvo Car Corporation, Gothenburg, conducting research on tracking and sensor fusion methods for active safety systems. Currently, Lars is an Assistant Professor at the Signal Processing group at Chalmers University

of Technology where his main research interests are in the fields of estimation, sensor fusion, self-localization and mapping, especially with application to self-driving vehicles.