Global Robot Localization with Random Finite Set Statistics*

Adrian N. Bishop
National ICT Australia (NICTA)
Australian National University (ANU)
adrian.bishop@anu.edu.au

Patric Jensfelt
Center for Autonomous Systems
Royal Institute of Technology (KTH)
patric@kth.se

Abstract – *We re-examine the problem of global localization of a robot using a rigorous Bayesian framework based on the idea of random finite sets. Random sets allow us to naturally develop a complete model of the underlying problem accounting for the statistics of missed detections and of spurious/erroneously detected (potentially unmodeled) features along with the statistical models of robot hypothesis disappearance and appearance. In addition, no explicit data association is required which alleviates one of the more difficult sub-problems. Following the derivation of the Bayesian solution, we outline its first-order statistical moment approximation, the so called probability hypothesis density filter. We present a statistical estimation algorithm for the number of potential robot hypotheses consistent with the accumulated evidence and we show how such an estimate can be used to aid in re-localization of kidnapped robots. We discuss the advantages of the random set approach and examine a number of illustrative simulations.*

Keywords: Robot localization; multiple-hypothesis localization; PHD filtering; random-set-based localization.

1 Introduction

The general approach to global localization (when not using GPS or artificial beacons such as bar codes and transponders) is to compare information (or features) extracted from sensor readings with an a priori map associated with the global reference frame. Each comparison carries some evidence about where the robot may be, and the challenge is then, as efficiently as possible, to find the correct pose, or a number of poses, that are in some statistical sense the most consistent with the accumulated evidence.

The approach proposed in this work most closely resembles the multiple hypothesis localization algorithms such as [1–3]. For brevity, we must point to the literature [4] for information on the other common approaches; most notably the Monte-Carlo (particle-filter type) algorithms [5, 6].

In the multiple-hypothesis technique [1–3], a set of Gaussians is used to represent individual pose hypotheses. The advantage of this is that a standard Kalman based pose tracking module can be used to independently update each hypothesis in a simple scheme commonly known as multiple hypothesis tracking (MHT). A further advantage of using a set of Gaussians is that it enables us to explicitly reason about each hypothesis whereas a particle filter, or a position probability grid, in principle requires you to process a much larger number of particles/cells (and their weights/probabilities). This is computationally challenging and is often circumvented by performing thresholding on the probabilities or clustering to form fewer hypotheses. Further advantages of the MHT-based approaches is found in [2].

A disadvantage of the multi-hypothesis techniques is that they inherently don't solve some of the most difficult problems related to localization (albeit they do with external algorithm components). In particular, the data-association problem and the problem of estimating (in an integrated/optimal fashion) a meaningful statistic regarding the number of poses consistent with the accumulated evidence and system models are not inherent. The second problem is often not studied explicitly in most localization algorithms[1] but plays an important role in the localization performance when false positive feature detections occur and/or more general measurement/system models are considered (such as in this paper). In such cases, the former data-association problem is further complicated and so-called clutter-rejection must be incorporated.

Furthermore, the standard vector-based stochastic differential (or difference) equation framework is arguably not a natural formulation for the multi-hypothesis tracking problem. Instead, in this paper we will explicitly exploit the framework of random finite sets (RFS - a term made precise later) and stochastic point processes [7]. A theoretically op-

*A.N. Bishop was with the Royal Institute of Technology (KTH) in Stockholm, Sweden at the time of the original submission. The authors were supported by the Swedish Foundation for Strategic Research (SSF) through the Centre for Autonomous Systems (CAS) and by the EU project 'CogX'. A.N. Bishop was also supported by NICTA. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

[1]Some algorithms can potentially provide an ad-hoc estimate on the number of consistent robot poses, e.g. by counting particle clusters or hypotheses above some weight. However, the framework discussed in this paper provides such an estimate inherently and, in particular, it provides an estimate of the mathematically *expected* number of hypotheses consistent with the data.

timal, Bayesian filtering, framework can then be formulated using the concept of finite set statistics (FISST) [8].

For computational reasons, Mahler [9] proposed a first-oder moment approximation to the full Bayesian solution and termed the first-moment, the probability hypothesis density (PHD). A generic sequential Monte Carlo implementation [10–12] has been proposed and accompanied by convergence results [12–14]. Alternatively, an analytic solution to the PHD recursion was presented in [15] for problems involving linear Gaussian target dynamics, a Gaussian birth model and linear Gaussian (partial) observations. It is shown in [15] that when the initial prior intensity is a Gaussian mixture, the posterior intensity at any subsequent time step is also a Gaussian mixture. Furthermore, the Gaussian-mixture PHD recursions can approximate the true posterior intensity to any desired degree of accuracy [16].

See [7, 8, 11] for a comprehensive background on random finite set-based estimation and the PHD filter.

1.1 Contribution

The PHD filter has primarily been examined in the context of target tracking (with various sensors etc) [7]. However, a recent paper [17] has examined the performance of the PHD filter in solving the simultaneous localization and mapping (SLAM) problem. The PHD filter-based SLAM implementation [17] was shown to outperform the base implementation of FastSLAM in a number of adverse environments. In this paper, we re-examine the problem of global robot localization using a rigorous Bayesian framework based on the concept of random finite sets and the PHD filter.

Random sets allow us to naturally develop a complete model of the underlying problem accounting for the statistics of missed detections and the statistics of spurious/erroneously detected (potentially unmodeled) features. In addition we incorporate the statistical models of robot hypothesis disappearance (death) and appearance (typically after kidnapping or error). No explicit data association is required which alleviates one of the more difficult subproblems. Following the derivation of a complete and integrated Bayesian solution, we outline its first-order statistical moment approximation, i.e. the probability hypothesis density filter. We present a statistical estimation algorithm for the expected number of potential robot hypotheses consistent with the accumulated evidence and we show how such an estimate can be used to aid in re-localization of kidnapped robots. We then discuss the advantages of the random set approach and examine a number of illustrative examples.

Our technique's ability to handle missed detections, false alarms (false feature detections) and to associate an entire set of measurement hypotheses to a set of robot pose hypotheses within an integrated framework is significant. The stochastic vector realizations of robot localization (including the traditional multiple hypothesis techniques [1–4]) have to deal with such problems explicitly and outside any Bayes optimal (or approximated) filter. To the best of the authors' knowledge, this paper presents the first complete Bayesian localization solution involving the concept of random finite sets and details the first implementation of the PHD filter for global robot localization.

2 Conceptual Model

The idea behind the set-based, multiple hypothesis generation and tracking technique presented in this paper is illustrated in Figure 1.

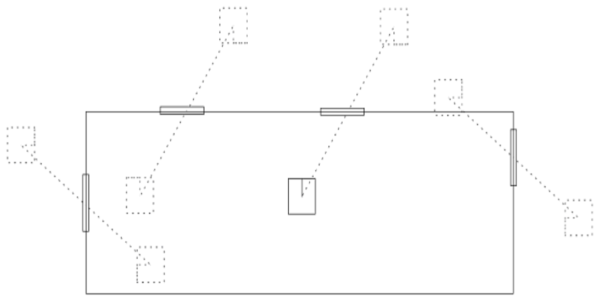

Figure 1: Multiple robot hypotheses are generated by a single measurement of some features in the environment given a priori knowledge that such features appear in a number of places in the global map.

In this illustration we see a situation where the true position of the robot is given by the solid square in the middle of the room in the figure. A door is detected in front and slightly to the right of the robot. Matching this feature to the map, consisting of four doors in one room, results in eight potential robot poses. These eight poses give rise to eight hypotheses regarding the pose of the robot. In the formulation outlined in this work, each hypothesis generated by a single feature detection is input to the estimation algorithm as an additional measurement. The grouping of such *measurement hypotheses* are modeled as random sets. In addition, the state of the robot's knowledge is modeled as a random set of *robot pose hypotheses*. The idea is that by making more observations of features in the environment and matching these to the existing robot pose hypotheses, those pose hypotheses which are not supported by the measurement set, i.e. the measurement hypotheses, can be eliminated (in a manner to be made precise) from the robot pose state set.

2.1 The Robot Set State Model

The true pose of a single robot is represented by the random variable $\mathbf{R}_t$ measured on the space $\mathcal{E} \subseteq \mathbb{R}^{n_r}$ with realization $\mathbf{r}_t$. The number of robot pose hypotheses is time-varying and given by N_t at time t. We denote the individual pose hypotheses by $\mathbf{X}_t^i$. The set of state pose hypotheses at time t is denoted by $\mathfrak{X}_t$.

The true state of a single robot is assumed to obey

$$\mathbf{R}_t = \psi_t \left(\mathbf{r}_{t-1} \right) + \mathbf{W}_t \qquad (1)$$

where the input $\{\mathbf{W}_t\}$ is a sequence of independent Gaussian random variables that account for control input errors and unmodeled dynamics etc. More general, non-Gaussian, error inputs can be accommodated in the general framework outlined in this paper.

Note that the transition density for the individual robot hypotheses $\mathbf{X}_t^i \in \mathfrak{X}_t$ is now given by

$$f_{t|t-1}^i(\mathbf{x}_t^i|\mathbf{x}_{t-1}^i) = \mathcal{N}(\mathbf{x}_t^i;\ \psi_t\left(\mathbf{x}_{t-1}^i\right), \mathbf{\Sigma}_t) \tag{2}$$

where $\mathcal{N}(\cdot;\ \mathbf{m}, \mathbf{P})$ denotes a Gaussian density function with mean $\mathbf{m}$ and covariance $\mathbf{P}$ and $\mathbf{\Sigma}_t$ is the covariance of $\mathbf{W}_t$.

The state set transition model in this paper incorporates statistical models of hypothesis appearance (birth) and hypothesis disappearance (death). New hypotheses might appear when the robot is kidnapped or in the case of localization error.

The probability that any hypothesis i continues to exist at time t given that it exists at $t-1$ is given by the survival probability p_S^i. Now it follows that

$$\mathfrak{X}_t \quad = \quad \left[\bigcup_{\mathbf{X}_{t-1}^i \in \mathfrak{X}_{t-1}} \mathfrak{S}_{t|t-1}(\mathbf{X}_{t-1}^i) \right] \cup b_t \mathfrak{B}_t \tag{3}$$

where $b_t \in \{0, 1\}$ and $b_t\mathfrak{B}_t$ is defined to be $\mathfrak{B}_t$ when $b_t = 1$ or $\emptyset$ when $b_t = 0$. Also,

$$\mathfrak{S}_{t|t-1}(\mathbf{X}_{t-1}^i) = \begin{cases} \mathbf{X}_t^i \cap \mathcal{E} & \text{with probability } p_S^i \\ \emptyset & \text{with probability } 1 - p_S^i \end{cases} \tag{4}$$

where the evolution of $\mathbf{X}_{t-1}^i$ follows (2). If we neglect $\mathfrak{B}_t$, then it is clear that we are modeling the motion and death of a number of robot hypotheses in (3). That is, if $b_t = 0$ then our set transition model accounts only for random hypothesis disappearance.

The input $b_t \in \{0, 1\}$ acts as a kidnapped robot switch which is set to 1 based on the outcome of kidnapped robot test outlined later in the paper. If $b_t = 1$ then we permit a statistical model of the new hypotheses, i.e. the possible locations of the kidnapped robot. The new hypotheses born at time t are characterized by a Poisson random finite set $\mathfrak{B}_t$ with intensity

$$\beta_t = \sum_{i=1}^{J_t^\beta} w_t^{(\beta,i)} \mathcal{N}(\mathbf{x};\ \mathbf{m}_t^{(\beta,i)}, \mathbf{P}_t^{(\beta,i)}) \tag{5}$$

which can approximate any arbitrary intensity function as closely as desired in the sense of the L^1 error [18]. The test for switching $b_t \in \{0, 1\}$ will be detailed later in the paper.

Using finite set statistics (FISST), we can find an explicit expression for the random set state transition density[2]. Now the multiple-hypothesis transition density $f_{t|t-1}(\mathfrak{X}_t|\mathfrak{X}_{t-1})$ under the model (3) and with $b_t = 1$ is given by

$$f_{t|t-1}(\mathfrak{X}_t|\mathfrak{X}_{t-1}) = \prod_{\mathbf{x} \in \mathfrak{X}_t} \beta_t \cdot \prod_{\mathbf{x}_t^i \in \mathfrak{X}_{t-1}} (1 - p_S^i) \cdot$$
$$e^{-\left(\sum_{i=1}^{J_t^\beta} w_t^{(\beta,i)}\right)} \cdot \sum_\theta \prod_{i:\theta(i)>0} \frac{p_S^i f_{t|t-1}(\mathbf{x}_t^{\theta(i)}|\mathbf{x}_{t-1}^i)}{\beta_t(1 - p_S^i)} \tag{6}$$

[2]In the case of random finite sets we make no distinction between the random sets and their realizations.

and with $b_t = 0$ it is given by

$$f_{t|t-1}(\mathfrak{X}_t|\mathfrak{X}_{t-1}) = \prod_{\mathbf{x}_t^i \in \mathfrak{X}_{t-1}} (1 - p_S^i) \cdot$$
$$\sum_\theta \prod_{i:\theta(i)>0} \frac{p_S^i f_{t|t-1}(\mathbf{x}_t^{\theta(i)}|\mathbf{x}_{t-1}^i)}{\beta_t(1 - p_S^i)} \tag{7}$$

where the summation is taken over all associations $\theta : \{1, \ldots, N_{t-1}\} \rightarrow \{1, \ldots, N_t\}$; see [7]. To the best of our knowledge, this is the first complete set-based transition density function proposed for global robot localization.

2.2 The Measurement Set Model

We consider n_s information sources, e.g. sensors on the robot. Each information source j generates an output

$$\mathbf{Z}_t^{(j,i)} = \zeta_t^j\left(\mathbf{r}_t, \mathcal{F}_{\phi^j(i)}\right) + \mathbf{V}_t^j, \ \text{ for } i = 1, \ldots, M_t^j \tag{8}$$

in the observation space $\mathcal{M} \subseteq \mathbb{R}^{n_{z_j}}$ where typically $n_{z_j} \leq n_r$. Note that certain measurement spaces, such as bearing measurements, can be approximated by subspaces of the real line. The input $\mathcal{F}_{\phi(i)} \in \mathcal{G}$ is some feature in the global environment model $\mathcal{G}$ and M_t^j is the number of measurement hypotheses generated by the j^{th} source given the true robot pose $\mathbf{r}_t$ and $\mathcal{G}$. The function $\phi^j : \{1, \ldots, M_t^j\} \rightarrow \{1, \ldots, \text{number of features}\}$ relates the index of the generated measurement hypotheses to the set of features in the global model $\mathcal{G}$[3]. The input $\{\mathbf{V}_t^j\}$ is a sequence of independent Gaussian random variables. Of course, more general noise models can also be considered in this framework.

The measurement likelihood function is

$$g_t^{(j,i)}(\mathbf{z}_t^{(j,i)}|\mathbf{r}_t, \mathcal{F}_{\phi^j(i)}) = \mathcal{N}(\mathbf{z}_t^{(j,i)};\ \zeta_t^j\left(\mathbf{r}_t, \mathcal{F}_{\phi^j(i)}\right), \mathbf{\Lambda}_t^j) \tag{9}$$

where $\mathbf{\Lambda}_t^j$ is the covariance of $\mathbf{V}_t^j$.

The considered measurement model incorporates measurements of the true robot pose (or nonlinear functions of such) and false measurement hypotheses generated by ambiguities (e.g. multiple occurrences of particular features) in the environment. In addition, we account for spurious false (clutter) measurements which are caused by false detections (e.g. of unmodeled features in the environment or simply detection/recognition errors). Finally, we also consider the possibility of missed measurements.

The probability that some modeled feature $\mathcal{F}_{\phi^j(i)}$ is actually detected by sensor j is given by the detection probability p_D^j, i.e. the probability of missing a measurement is $1 - p_D^j$. Spurious false (clutter) measurements at sensor j are approximated by a Poisson random finite set $\mathfrak{C}_t^j$ with intensity

$$\kappa_t^j = \gamma_t^j \mathcal{U}(\mathcal{G}) \tag{10}$$

[3]For example, the robot might, using sensor j, measure the bearing to some detected feature, e.g. a door, in the environment. This single measurement constrains the robot position to a number of rays associated with this and similar features, e.g. other doors, in the environment. Each constraint is treated as a measurement hypothesis and the total number of such hypotheses is M_t^j.

where $\mathcal{U}(\mathcal{G})$ denotes a uniform density function over the environment. The clutter corresponds to the spurious set of measurement hypotheses generated by erroneous detections or the detection of features not in the environment model. The detection probability p_D^j can be a function of the environment model $\mathcal{G}$ and the true robot pose. Thus, we can use p_D^j to model the sensor geometry etc [19].

Now it follows that the set of *measurement hypotheses* at sensor j is given by

$$\mathfrak{Z}_t^j \quad = \quad \left[\bigcup_{i=1,\ldots,h_j} \mathfrak{D}_t(\mathbf{R}_t, \mathcal{F}_{\phi^j(i)}) \right] \cup \mathfrak{C}_t^j \qquad (11)$$

where

$$\mathfrak{D}_t(\mathbf{R}_t, \mathcal{F}_{\phi^j(i)}) = \begin{cases} \{\mathbf{Z}_t^{(j,i)}\}_{i=1}^{h_j} & \text{with prob. } p_D^j \\ \emptyset & \text{with prob. } 1 - p_D^j \end{cases} \qquad (12)$$

and where $\mathbf{Z}_t^{(j,i)}$ is modeled by (8) and (9) with $g_t^{(j,i)}$. Now the entire set of evidence at time t is given by

$$\mathfrak{Z}_t = \bigcup_{i=1}^{n_s} \{\mathfrak{Z}_t^i, i\} \qquad (13)$$

where the union is disjoint and $M_t = \sum_{i=1}^{n_s} M_t^i$. The measurement likelihood function corresponding to the single-sensor model (11) is given by

$$\begin{aligned} g_t^j(\mathfrak{Z}_t^j|\mathfrak{X}_t) \quad &= \quad e^{-\gamma_t} \cdot \prod_{\mathbf{z}\in\mathfrak{Z}_t^j} \kappa_t \cdot \prod_{\mathbf{z}\in\mathfrak{Z}_t^j} (1-p_D) \cdot \\ & \sum_{\theta^j} \prod_{i:\theta^j(i)>0} \frac{p_D g_t(\mathbf{z}_t(\theta^j(i))|\mathbf{x}_t^i)}{\kappa_t(1-p_D)} \end{aligned} \qquad (14)$$

where the summation is taken over all associations θ^j : $\{1,\ldots,N_t\} \to \{1,\ldots,M_t^j\}$ and where $\mathbf{z}_t(\theta^j(i))$ is an element in $\mathfrak{Z}_t^j$ marked by the function θ^j; see [7]. The multi-sensor likelihood function $g_t(\mathfrak{Z}_t|\mathfrak{X}_t)$ is then given by

$$g_t(\mathfrak{Z}_t|\mathfrak{X}_t) = \prod_{j=1}^{n_s} g_t^j(\mathfrak{Z}_t^j|\mathfrak{X}_t) \qquad (15)$$

under the assumptions adopted in (11). To the best of the author's knowledge, this measurement model is the most general considered in the literature on global robot localization.

3 A General Bayesian Localization Algorithm

The aim of global localization is to use the measured data $\mathfrak{Z}_t$ and some dynamical constraint on the random robot pose $\mathbf{R}_t$ to estimate the set $\mathfrak{X}_t$ of potential robot positions in the environment. If $|\mathfrak{X}_t| = 1$ then the robot is said to be uniquely localized and it is of course the hope that in such cases $\mathbf{X}_t^1 \approx \mathbf{R}_t$ for the single estimate $\mathbf{X}_t^1 \in \mathfrak{X}_t$. The notion of a "set" of information points $\mathfrak{Z}_t$ and a "set" of hypotheses $\mathfrak{X}_t$ is critical as it allows us to side-step the problem of associating measurement points to a priori hypotheses in addition to other bookkeeping localization tasks.

Let $p_t(\mathfrak{X}_t|\mathfrak{Z}_{1:t})$ denote the multiple hypothesis posterior density. Then, the optimal Bayes localization filter propagates the posterior in time via the recursion

$$\begin{aligned} & p_{t|t-1}(\mathfrak{X}_t|\mathfrak{Z}_{1:t-1}) = \\ & \int f_{t|t-1}(\mathfrak{X}_t|\mathfrak{X}) p_{t-1}(\mathfrak{X}|\mathfrak{Z}_{1:t-1}) \, \mu_S(\mathrm{d}\mathfrak{X}) \end{aligned} \qquad (16)$$

$$\begin{aligned} & p_t(\mathfrak{X}_t|\mathfrak{Z}_{1:t}) = \\ & \frac{g_t(\mathfrak{Z}_t|\mathfrak{X}_t) p_{t|t-1}(\mathfrak{X}_t|\mathfrak{Z}_{1:t-1})}{\int g_t(\mathfrak{Z}_t|\mathfrak{X}_t) p_{t|t-1}(\mathfrak{X}_t|\mathfrak{Z}_{1:t-1}) \, \mu_S(\mathrm{d}\mathfrak{X})} \end{aligned} \qquad (17)$$

where μ_S is an appropriate reference measure on the collection of finite sets of $\mathcal{E}$. FISST is the first systematic approach to multi-object filtering that uses random finite sets in the Bayesian framework presented above [7, 15]. The general recursive Bayesian filter based on density functions defined for random finite set models suffers from a severe computational requirement and only a few implementations have been studied (using Monte-Carlo methods and for the problem of multi-sensor/multi-target tracking) [7, 11, 20, 21].

4 A First-Order Moment Approximation: The PHD Filter

The probability hypothesis density filter is an approximation developed to alleviate the computational intractability of the general Bayes filter. The PHD filter propagates the posterior intensity, a first-order statistical moment of the posterior state.

Assumption 1. *The predicted multi-target random finite set governed by $p_{t|t-1}$ is Poisson.*

For a random finite set $\mathfrak{X}$ on $\mathcal{E}$ with probability distribution P, its first-order moment is a non-negative function v on $\mathcal{E}$, called the intensity, such that for each region $\mathcal{A} \subseteq \mathcal{E}$

$$\int |\mathfrak{X} \cap \mathcal{A}| \; \mathrm{P}(\mathrm{d}\mathfrak{X}) = \int_{\mathcal{A}} v(\mathbf{x}) \, \mathrm{d}\mathbf{x} \qquad (18)$$

where

$$\mathrm{E}[N] = \int_{\mathcal{E}} v(\mathbf{x}) \, \mathrm{d}\mathbf{x} \qquad (19)$$

and $\mathrm{E}[N]$ denotes the expected number of elements in $\mathfrak{X}$. The local maxima of v are points in $\mathfrak{X}$ with the highest local concentration of expected number of elements, and hence can be used to generate estimates for the elements of $\mathfrak{X}$.

Let v_t and $v_{t|t-1}$ denote the intensities associated with the multiple target posterior density p_t and the multiple target predicted density $p_{t|t-1}$. The posterior intensity is $v_t(\mathbf{x}) = v_t^{n_s}(\mathbf{x})$ where

$$\begin{aligned} & v_t^k(\mathbf{x}) = (1 - p_D^k) v_t^{k-1}(\mathbf{x}) + \\ & \sum_{\mathbf{z}\in\mathfrak{Z}_t^k} \frac{p_D^k g_t^{(k,')}(\mathbf{z}|\mathbf{x}) v_t^{k-1}(\mathbf{x})}{\kappa_t^k + p_D^k \int g_t^{(k,')}(\mathbf{z}|\mathbf{x}') v_t^{k-1}(\mathbf{x}') \, \mathrm{d}\mathbf{x}'} \end{aligned} \qquad (20)$$

and $v_t^0(\mathbf{x}) = v_{t|t-1}(\mathbf{x})$. The PHD predictor is given by

$$v_{t|t-1}(\mathbf{x}) = b_t\beta_t + p'_S \int f_{t|t-1}(\mathbf{x}|\mathbf{x}')v_{t-1}(\mathbf{x}')\,\mathrm{d}\mathbf{x}' \quad (21)$$

Following [7] we note that the PHD filter (similarly to the full recursive multi-target Bayesian estimator), admits explicit statistical models for missed detections, false alarms and the geometry of the sensor's field of view. In addition, the PHD filter admits explicit statistical models of robot hypothesis disappearance (death) and appearance (due to, for example, kidnapping). In addition, at every step the PHD filter computes an estimate of the number of robot hypotheses consistent with the data up until this step.

The last property aids in clutter-rejection and will be used to derive a probabilistically justified test for kidnapping.

5 Multi-Sensor and Multi-Hypothesis Gaussian-Sum PHD Filter

In this section we present an implementation of the PHD filer based on a mixture of Gaussians algorithm.

We firstly suppose that each hypothesis is constrained by a linear model of the form

$$\mathbf{X}_t^i = \mathbf{\Phi}_t\mathbf{x}_{t-1}^i + \mathbf{W}_t \quad (22)$$

Also, the output of information source (sensor) j obeys

$$\mathbf{Z}_t^{(j,i)} = \mathbf{\Gamma}_t^j\mathbf{r}_t + \mathbf{V}_t^j \quad (23)$$

In addition, we make a reasonable assumption that the survival probability p_S^i is independent of the individual state hypothesis, i.e. $p_S^i = p_S$. Under the above assumptions the following Gaussian-Mixture PHD filter (GM-PHD) is an exact implementation of the conceptual PHD filter.

Proposition 1 ([15]). *Suppose the modeling assumptions presented hold and that the posterior intensity at time $t-1$ is a Gaussian mixture of the form*

$$v_{t-1}(\mathbf{x}) = \sum_{i=1}^{J_{t-1}} w_{t-1}^i \mathcal{N}(\mathbf{x};\ \mathbf{m}_{t-1}^i, \mathbf{P}_{t-1}^i) \quad (24)$$

Then, the predicted intensity at time t is given by

$$v_{t|t-1}(\mathbf{x}) = \beta_t + p_S \sum_{i=1}^{J_{t-1}} w_{t-1}^i \mathcal{N}(\mathbf{x};\ \mathbf{m}_{t|t-1}^i, \mathbf{P}_{t|t-1}^i) \quad (25)$$

where

$$\mathbf{m}_{t|t-1}^i = \mathbf{\Phi}_t\mathbf{m}_{t-1}^i \quad (26)$$

$$\mathbf{P}_{t|t-1}^i = \mathbf{\Sigma}_t + \mathbf{\Phi}_t\mathbf{P}_{t-1}^i\mathbf{\Phi}_t^\top \quad (27)$$

and is also a Gaussian mixture.

Proposition 2 (Adapted from [15]). *Suppose the modeling assumptions presented hold and that the predicted intensity at time t is a Gaussian mixture of the form*

$$v_{t|t-1}(\mathbf{x}) = \sum_{i=1}^{J_{t|t-1}} w_{t|t-1}^i \mathcal{N}(\mathbf{x};\ \mathbf{m}_{t|t-1}^i, \mathbf{P}_{t|t-1}^i) \quad (28)$$

Then, the posterior intensity at t is $v_t(\mathbf{x}) = v_t^{n_s}(\mathbf{x})$ where

$$v_t^k(\mathbf{x}) = (1-p_D)v_t^{k-1}(\mathbf{x}) + \sum_{\mathbf{z}\in\mathfrak{Z}_t^k}\sum_{i=1}^{J_{t|t-1}} w_{t|t}^{(i,k)}(\mathbf{z})\mathcal{N}(\mathbf{x};\ \mathbf{m}_{t|t}^{(i,k)}(\mathbf{z}), \mathbf{P}_{t|t}^{(i,k)}) \quad (29)$$

$$= \sum_{i=1}^{J_t^k} w_t^{(i,k)}\mathcal{N}(\mathbf{x};\ \mathbf{m}_t^{(i,k)}, \mathbf{P}_t^{(i,k)}) \quad (30)$$

and $v_t^0(\mathbf{x}) = v_{t|t-1}(\mathbf{x})$ and where

$$w_{t|t}^{(i,k)}(\mathbf{z}) = \frac{p_D w_t^{(i,k-1)} q_t^{(i,k)}(\mathbf{z})}{\kappa_t + p_D\sum_{\ell=1}^{J_{t|t-1}^k} w_t^{(\ell,k-1)} q_t^{(\ell,k)}(\mathbf{z})} \quad (31)$$

$$q_t^{(i,k)}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{\Gamma}_t\mathbf{m}_t^{(i,k-1)}, \mathbf{\Lambda}_t + \mathbf{\Gamma}_t\mathbf{P}_t^{(i,k-1)}\mathbf{\Gamma}_t^\top)$$

$$\mathbf{m}_{t|t}^i(\mathbf{z}) = \mathbf{m}_t^{(i,k-1)} + \mathbf{K}_t^{(i,k)}(\mathbf{z} - \mathbf{\Gamma}_t\mathbf{m}_t^{(i,k-1)}) \quad (32)$$

$$\mathbf{K}_t^{(i,k)} = \mathbf{P}_t^{(i,k-1)}\mathbf{\Gamma}_t^\top(\mathbf{\Lambda}_t + \mathbf{\Gamma}_t\mathbf{P}_t^{(i,k-1)}\mathbf{\Gamma}_t^\top)^{-1}$$

$$\mathbf{P}_{t|t}^i = (\mathbf{I} - \mathbf{K}_t^{(i,k)}\mathbf{\Gamma}_t)\mathbf{P}_t^{(i,k-1)} \quad (33)$$

and $v_t(\mathbf{x})$ is also a Gaussian mixture.

The preceding propositions show how the Gaussian components of the posterior intensity are analytically propagated in time (for the linear Gaussian measurement and hypothesis dynamic model[4]) [15].

5.1 Accounting for Nonlinear Models

Instead of (22) and (23) we know each state pose hypothesis $\mathbf{X}_t^i \in \mathfrak{X}_t$ is constrained by the transition density

$$f_{t|t-1}^i(\mathbf{x}_t^i|\mathbf{x}_{t-1}^i) = \mathcal{N}(\mathbf{x}_t^i;\ \psi_t\left(\mathbf{x}_{t-1}^i\right), \mathbf{\Sigma}_t) \quad (34)$$

and each measurement hypothesis likelihood function is

$$g_t^{(j,i)}(\mathbf{z}_t^{(j,i)}|\mathbf{r}_t, \mathcal{F}_{\phi^j(i)}) = \mathcal{N}(\mathbf{z}_t^{(j,i)};\ \zeta_t^j\left(\mathbf{r}_t, \mathcal{F}_{\phi^j(i)}\right), \mathbf{\Lambda}_t^j) \quad (35)$$

It then follows that the transition density for the individual hypotheses is now approximately given by

$$f_{t|t-1}^i(\mathbf{x}_t^i|\mathbf{x}_{t-1}^i) = \mathcal{N}(\mathbf{x}_t^i;\ \mathbf{\Phi}_t\mathbf{x}_{t-1}^i, \mathbf{\Sigma}_t) \quad (36)$$

where $\mathbf{\Sigma}_t$ is the covariance $\mathbf{W}_t$ and now

$$\mathbf{\Phi}_t = \left.\frac{\partial\psi_t\left(\mathbf{x}_{t-1}\right)}{\partial\mathbf{x}_{t-1}}\right|_{\mathbf{x}_{t-1}=\mathbf{m}_{t-1}^i} \quad (37)$$

In addition, the measurement likelihood functions can be approximated by

$$g_t^j(\mathbf{z}_t^j|\mathbf{x}_t) = \mathcal{N}(\mathbf{z}_t^j;\ \mathbf{\Gamma}_t^j\mathbf{x}_t, \mathbf{\Lambda}_t^j) \quad (38)$$

where

$$\mathbf{\Gamma}_t^j = \left.\frac{\partial\zeta_t^j\left(\mathbf{x}_t, \mathbf{0}\right)}{\partial\mathbf{x}_t}\right|_{\mathbf{x}_t=\mathbf{m}_{t|t-1}^i} \quad (39)$$

Now plugging the Jacobians $\mathbf{\Phi}_t^j$ and $\mathbf{\Gamma}_t^j$ into the previously outlined GM-PHD filter (covariance formulas) leads to an approximation in the spirit of the extended Kalman filter[5].

[4] In order to be computationally feasible the authors of [15] proposed a simple pruning and merging strategy.

[5] In [15] an unscented extension of the GM-PHD filter is also outlined in a similar manner.

5.2 The Expected Number of Hypotheses

The predicted number of hypotheses is given by

$$\mathrm{E}[N_{t|t-1}] = p_S \mathrm{E}[N_t] + b_t \sum_{i=1}^{J_t^\beta} w_t^{(\beta,i)} \quad (40)$$

while the updated, expected, number of hypotheses is given by $\mathrm{E}[N_t] = \mathrm{E}[N_t^{n_s}]$ where

$$\mathrm{E}[N_t^k] = (1 - p_D^k)\mathrm{E}[N_t^{k-1}] + \sum_{\mathbf{z}\in \mathfrak{Z}_t^k} \sum_{i=1}^{J_{t|t-1}} w_{t|t}^{(i,k)}(\mathbf{z}) \quad (41)$$

and $\mathrm{E}[N_t^0] = \mathrm{E}[N_{t|t-1}]$. Note that while some other approaches can potentially provide an estimate on the number of consistent robot poses, e.g. by counting particle clusters or hypotheses above some weight, the PHD framework inherently provides a statistically relevant estimate of the *expected* number of hypotheses consistent with the data.

6 Recovery from Localization Error and Kidnapping

The kidnapped robot switch $b_t \in \{0, 1\}$ allows us to switch a statistical model of new hypothesis appearance into the state hypothesis set transition model when we suspect the need to permit new hypotheses. The switch is given by

$$b_t = \begin{cases} 0 & \text{if } \mathrm{E}[N_t] \geq \tau_k \\ 1 & \text{if } \mathrm{E}[N_t] < \tau_k \end{cases} \quad (42)$$

where $\tau_k << 1$ is a kidnapping threshold on the expected number of state pose hypotheses. The idea behind the kidnapping switch is that if the expected number of hypotheses falls below some threshold (typically $<< 1$) then the robot has either been kidnapped or the localization algorithm has encountered an error. In either case, we are justified in suspecting that the new hypotheses are required (assuming we know that the true robot is still located in the environment).

The spatial distribution of the Poisson random finite birth set $\mathfrak{B}_t$ is a sum of Gaussian components β_t and can be tailored to account for any a priori information available or can be used to approximate a uniform density over the environment model (as closely as desired [18]).

7 Experiments

We evaluated the localization algorithm using both simulated and real feature measurements[6]. In both cases, we use a real robot trajectory[7]. The floor plan with corner and door features is shown in Figure 2.

[6]The experiments conducted were designed to highlight a number advantages of the PHD filter framework such as providing an estimate on the expected number of robot poses and kidnap detection using real data.

[7]The robot employed in the experiments is a Pioneer P3X.

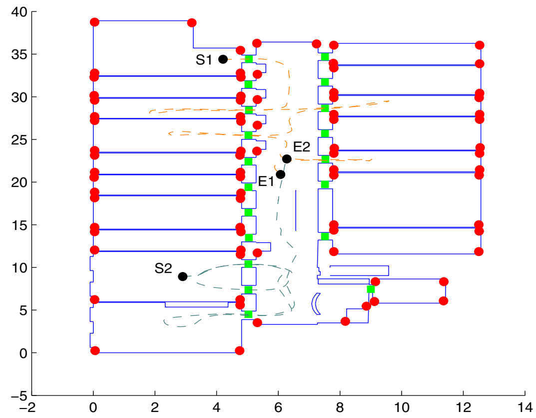

Figure 2: The floorplan at the Center for Autonomous Systems (CAS) at KTH with corner and door features highlighted. Two real robot trajectories are shown with start positions (S1 and S2) and end positions (E1 and E2).

7.1 Example 1

As previously stated, one particular advantage of the PHD framework is the inherent estimation of the number of robot poses consistent with the accumulated evidence and the system models. Firstly, we illustrate this behavior for a robot traveling along the first trajectory (S1→E1) shown in Figure 2 with real door features extracted as in [2]. In this case, we would expect that the expected number of hypotheses should not approach one for some time due to the symmetry in the features. The sequence of localization is illustrated in Figure 3 and verifies our expectations.

The estimated expected number of hypotheses for an individual run is depicted in Figure 4 along with the maximum Gaussian component weight scaled by this expected value.

From Figure 3 and 4 we see that the robot is uniquely localized just prior to E1 (where the expected number of hypotheses goes to ≈ 1). We purposefully used only sparse door features to create a symmetry along trajectory 1.

7.2 Example 2

In this example we highlight the re-localization capabilities of the proposed algorithm. The robot initially travels along trajectory 2 and detects only door features. At position E2 the robot is then kidnapped and placed at position S1. We expect to see the expected number of hypotheses drop to zero and the re-localization sequence should be initiated. The estimation of the expected number of hypotheses is shown in Figure 5 and verifies our expectations.

The robot is uniquely localized relatively early during the transversal of trajectory 2 (as clearly seen in Figure 5). This is because of the relatively little symmetry in the lower half of the environment model. Then following the kidnapping, we see the re-localization sequence proceed similarly to the initial localization sequence examined in Example 1. The robot is then uniquely re-localized just prior to E1.

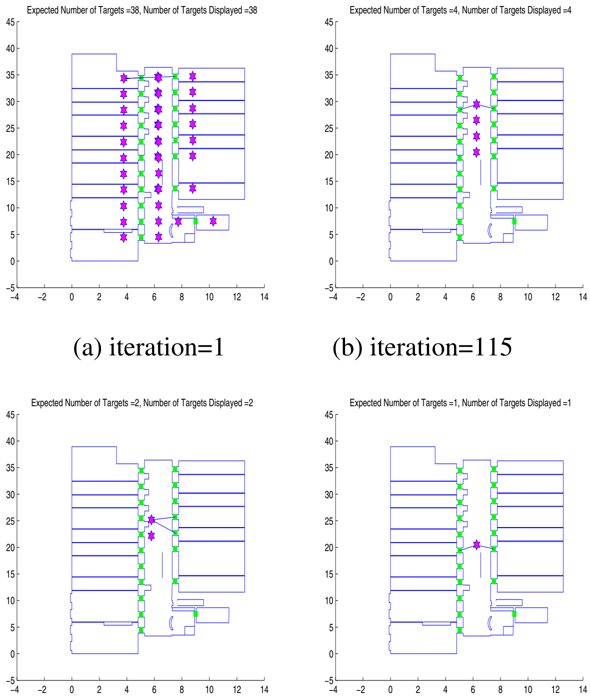

(a) iteration=1 (b) iteration=115

(c) iteration=800 (d) iteration=1100

Figure 3: The expected number of hypotheses displayed in the title is rounded to the nearest integer value. The number of displayed hypotheses is the number of Gaussian components with a weight above 0.5. The time step between iterations is 100ms. The bearings to the measured features in the map are drawn at the true robot pose.

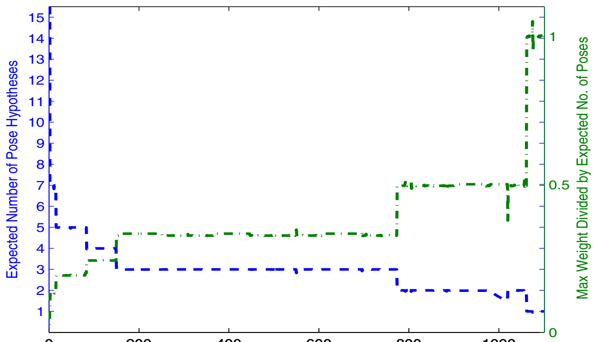

Figure 4: The expected number of hypotheses is displayed on the left for the duration of the experiment. The maximum Gaussian component weight scaled by the expected number of hypotheses is shown on the scale on the right.

We used only doors again in this example to highlight that the algorithm works on real data and to highlight again the estimation of the expected number of hypotheses (which should be greater than one initially and immediately following the first feature detection after the kidnapping detection).

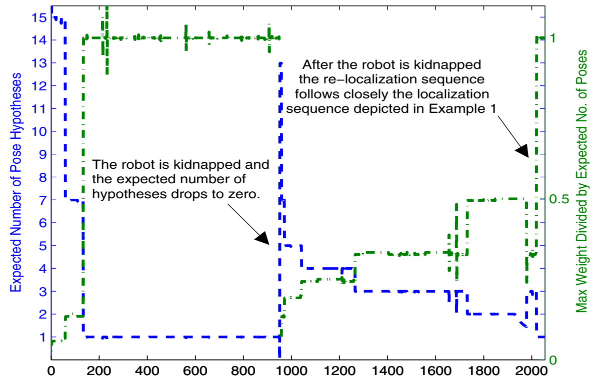

Figure 5: The expected number of hypotheses is displayed on the scale on the left. The maximum Gaussian component weight scaled by the expected number of hypotheses is shown on the scale on the right. The robot is kidnapped around iteration 940 from position E2 and placed at S1.

7.3 Example 3

This example is similar to Example 2 except we now use simulated door and corner features. We expect the robot to be localized much quicker (both initially and following kidnapping) as a result of the reduced symmetry.

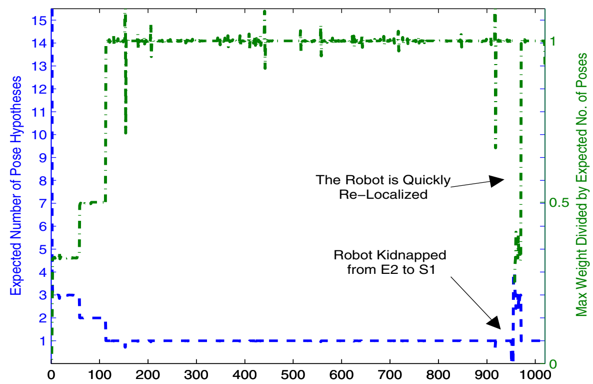

Figure 6: The expected number of hypotheses is displayed on the scale on the left. The maximum Gaussian component weight scaled by the expected number of hypotheses is shown on the scale on the right. The robot is kidnapped at some point from position E2 and placed at position S1.

Note that following re-localization, the robot follows the remainder of trajectory 1 as in the previous example (except the number of hypotheses is fewer (indeed the robot is uniquely localized) as a result of the extra corner features).

The speed at which we can detect a localization error (or robot kidnapping) depends on the chosen model parameters. We can trade robustness due to missed/false detections for an increased speed of error recovery (kidnapped detection). This trade off is inherent since in a single iteration, for example, a kidnapped robot and a robot missing some true detections while simultaneously detecting some false positives are indistinguishable.

8 Conclusion

The problem of robot localization is certainly not new. However, a true recursive Bayesian solution which eliminates the data-association problem and incorporates missed/false positive detections and hypothesis birth and death etc has not been examined previously in the context of random finite sets. We have outlined the first-order moment approximation of the integrated Bayesian solution (i.e. the PHD filter) for a general class of localization problems. The algorithm provided in this paper is based on rigorous statistical analysis of random finite sets and accommodates a very general measurement and robot motion model in an integrated framework. We have presented experimental results using both real and simulated data to illustrate some of the fundamental advantages of the integrated approach.

References

[1] S.I. Roumeliotis and G.A. Bekey. Bayesian estimation and kalman filtering: a unified framework for mobile robot localization. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation (ICRA'00)*, pages 2985–2992, 2000.

[2] Patric Jensfelt and Steen Kristensen. Active global localisation for a mobile robot using multiple hypothesis tracking. *IEEE Transactions on Robotics and Automation*, 17(5):748–760, October 2001.

[3] K.O. Arras, J.A. Castellanos, and R. Siegwart. Feature-based multi-hypothesis localization and tracking for mobile robots using geometric constraints. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'02)*, pages 1371–1377, 2002.

[4] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press, Cambridge, Massachusetts, 2005. ISBN 0-262-20162-3.

[5] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proceedings of the 16th AAAI Conference on Artificial Intelligence (AAAI'99).*, July 1999.

[6] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, May 2001.

[7] R.P.S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Boston, M.A., 2007.

[8] I.R. Goodman, R.P.S. Mahler, and H.T. Nguyen. *Mathematics of Data Fusion*. Kluwer Academic Publishers, London, U.K., 1997.

[9] R.P.S. Mahler. Multi-target Bayes filtering via first-order multi-target moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1152–1178, 2003.

[10] H. Sidenbladh. Multi-target particle filtering for the Probability Hypothesis Density. In *Proceedings of the 2003 International Conference on Information Fusion*, pages 800–806, Cairns, Australia, 2003.

[11] M. Vihola. Random sets for multitarget tracking and data fusion. Technical report, Department of Information Technology, Tampere University of Technology, Licentiate Thesis, 2004.

[12] B.-N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo methods for multi-target filtering with random finite sets. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4):12241245, 2005.

[13] A. Johansen, S. Singh, A. Doucet, and B.-N. Vo. Convergence of the sequential Monte Carlo implementation of the PHD filter. *Methodology and Computing in Applied Probability*, 8(2):265–291, 2006.

[14] D. Clark and J. Bell. Convergence results for the particle PHD filter. *IEEE Transactions on Signal Processing*, 54(7):2652–2661, July 2006.

[15] B.-N. Vo and W.K. Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, November 2006.

[16] D. Clark and B.-N. Vo. Convergence analysis of the Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 55(4):1204–1212, April 2007.

[17] B. Kalyan, K. W. Lee, and W. S. Wijesoma. FISST-SLAM: Finite Set Statistical Approach to Simultaneous Localization and Mapping. *The International Journal of Robotics Research (In Press)*, 2009.

[18] B.D.O. Anderson and J.B. Moore. *Optimal Filtering*. Prentice Hall, Englewood Cliffs, N.J., 1979.

[19] A.N. Bishop, B. Fidan, B.D.O. Anderson, K. Dogancay, and P.N. Pathirana. Optimality analysis of sensor-target localization geometries. *Automatica*, 46(3):479–492, March 2010.

[20] H. Sidenbladh and S.-L. Wirkander. Tracking random sets of vehicles in terrain. In *Proceedings of the 2003 IEEE Workshop on Multi-Object Tracking*, Madison, Wisconsin, June 2003.

[21] W.K. Ma, B.-N. Vo, S. Singh, and A. Baddeley. Tracking an unknown time-varying number of speakers using tdoa measurements: A random finite set approach. *IEEE Transactions on Signal Processing*, 54(9):3291–3304, September 2006.