Information Fusion 31 (2016) 110–126

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

An overview of particle methods for random finite set models

Branko Ristic[a,*], Michael Beard[b], Claudio Fantacci[c]

[a] School of Engineering, RMIT University, Rm 10.8.14, 376-392 Swanston Street, Melbourne, VIC 3000, Australia
[b] Curtin University, Bentley, WA, Australia
[c] Università di Firenze, DINFO, Florence, Italy

ARTICLE INFO

Article history:
Received 16 November 2015
Revised 4 February 2016
Accepted 7 February 2016
Available online 15 February 2016

Keywords:
Stochastic nonlinear filtering
Sequential Monte Carlo estimation
Random set models
Target tracking
Bearings-only measurements

ABSTRACT

This overview paper describes the particle methods developed for the implementation of the class of Bayes filters formulated using the random finite set formalism. It is primarily intended for the readership already familiar with the particle methods in the context of the standard Bayes filter. The focus in on the Bernoulli particle filter, the probability hypothesis density (PHD) particle filter and the generalised labelled multi-Bernoulli (GLMB) particle filter. The performance of the described filters is demonstrated in the context of bearings-only target tracking application.

1. Introduction

In many areas of science and engineering there is a need to infer the behaviour of a stochastic dynamic system, using its partial and indirect observations. By combining (typically nonlinear) mathematical models of system evolution and sensor measurements, one can formulate the optimal sequential estimator in the Bayesian framework. This estimator, commonly referred to as the Bayes-optimal (or simply Bayes) filter, provides a recursive formula for the complete probabilistic characterisation of the dynamic system in the form of a time-varying posterior probability density of its state [1].

For most nonlinear/non-Gaussian formulations, analytic closed-from solutions of the Bayes filter are intractable. Practical solutions, therefore, need to be based on approximations. Particle filters are a class of Monte Carlo simulation based methods which can provide very accurate approximations of the Bayes filter. Despite being computationally expensive, particle filters have become universally popular, primarily due to their accuracy, relatively simple implementation and the ever increasing speed of computers. As a result of their widespread application, a few good tutorials and books have been published on the subject of particle filters [2–7].

Particle filters have been introduced and traditionally applied as the approximate solutions of the *standard* Bayes filter, formulated during the 1960s [1] under the following assumptions: (i) the stochastic dynamic system (object, phenomenon) is permanently active (or present); (ii) observations are noisy, but collected with perfect detection (i.e. there are no false or missed detections). All the aforementioned tutorials and books discuss the particle filters in this context only. However, in many practical applications, one may have to deal with multiple stochastic dynamic systems (objects), which can be simultaneously active (present), and which can randomly switch *on* and *off* (appear/disappear). In addition, perfect detection using surveillance sensors (e.g. radar, sonar, video cameras) is rarely possible [8]. Until recently, particle filters have been applied to this class of problems using a clever combination of Bayesian estimation theory with ad-hoc logic. However, the recent advances in Bayesian estimation using *random finite set* (RFS) models [9] resulted in elegant and rigorous mathematical formulations of the Bayes-optimal and principled Bayes-suboptimal filters, applicable to multiple interacting on/off switching systems with possibly imperfect detection of measurements.

This overview paper describes the particle methods developed for the implementation of the new class of RFS-Bayes filters. It is primarily intended for the readership already familiar with the particle methods in the context of the standard Bayes filter. One of the most popular and convincing applications of particle filters, versus standard approximation methods, such as the Extended Kalman filter (EKF) [10] and unscented Kalman filter (UKF) [11], has been for bearings-only tracking problems [5, chap. 6]. Hence, this application has been chosen to demonstrate throughout the paper different RFS-Bayes particle filters and their performance. The paper is organised as follows. Section 2 reviews the particle method for the standard Bayes filter. The elements of mathematics

* Corresponding author. Tel.: +61 3 9925 3768.
E-mail address: branko.ristic@rmit.edu.au (B. Ristic).

http://dx.doi.org/10.1016/j.inffus.2016.02.004
1566-2535/

Notation

$\mathcal{X}$	the single object state space
$\mathcal{F}(\mathcal{X})$	the multiple object state space
$\mathcal{L}$	the space of labels
$\mathbf{x}$	the state of a single object (a random vector)
$\mathbf{X}$	the state of multiple-objects (a random finite set, RFS)
$\mathbb{X}$	the state of multiple labelled objects (a labelled RFS)
$\mathcal{Z}$	measurement space
$\mathbf{z}$	a measurement of a single object (a random vector)
$\mathbf{Z}$	a detector output measurement (a RFS)
k	discrete-time index
$p(\mathbf{x})$	a probability density function (PDF) of $\mathbf{x} \in \mathcal{X}$
$\tilde{p}(\mathbf{x}, \ell)$	a PDF of a labelled random vector $(\mathbf{x}, \ell) \in \mathcal{X} \times \mathcal{L}$
$f(\mathbf{X})$	a PDF of an RFS variable $\mathbf{X} \in \mathcal{F}(\mathcal{X})$
$\tilde{f}(\mathbb{X})$	a PDF of a labelled RFS variable $\mathbb{X} \in (\mathcal{F}(\mathcal{X}) \times \mathcal{L})$

for random finite set models are presented in Section 3. The particle method for the RFS Bayes-optimal filter and its special case, the Bernoulli filter, are discussed in Section 4. A multi-target particle filter, referred to as the PHD particle filter, is presented in Section 5. The labelled RFS Bayes tracking filters implemented using the particle method are discussed in Section 6. The summary and pointers to advanced research topics are given in Section 7.

2. Particle methods for the standard Bayes filter

In order to familiarise with the notation, let us start with a quick review of the standard Bayes filter and the corresponding particle methods. The problem is by no means simple and it is still an active and highly relevant research topic.

2.1. Problem formulation and the standard Bayes filter

Suppose the state vector $\mathbf{x}_k \in \mathcal{X}$ provides the complete specification of the state of a dynamic system (object, phenomenon) at time t_k. Here $\mathcal{X} \subseteq \mathcal{R}^{n_x}$ is the state space, while k is the discrete-time index corresponding to t_k. Let us adopt the discrete-time additive-noise formulation, specified by two equations:

$$\mathbf{x}_k = \mathbf{f}_{k-1}(\mathbf{x}_{k-1}) + \mathbf{v}_{k-1}, \tag{1}$$

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k) + \mathbf{w}_k, \tag{2}$$

referred to as the *dynamics equation* and the *measurement equation*, respectively. Function $\mathbf{f}_{k-1} : \mathcal{R}^{n_x} \to \mathcal{R}^{n_x}$ in (1) is a nonlinear transition function defining the temporal evolution of the state vector as a first-order Markov process. Random disturbances $\mathbf{v}_k \in \mathcal{R}^{n_x}$, also known as process noise, are assumed to be independent identically distributed (IID) according to the probability density function (PDF) $p_{\mathbf{v}}$. Function $\mathbf{h}_k : \mathcal{R}^{n_x} \to \mathcal{R}^{n_z}$ in (2) defines the relationship between the state $\mathbf{x}_k$ and the measurement $\mathbf{z}_k \in \mathcal{Z}$, where $\mathcal{Z} \subseteq \mathcal{R}^{n_z}$ is the measurement space. Random disturbances $\mathbf{w}_k \in \mathcal{R}^{n_z}$, also known as measurement noise, are assumed independent of $\mathbf{v}_k$, and modelled as an IID process with the PDF $p_{\mathbf{w}}$. Typically $n_z < n_x$, giving rise to the term *partial observations* of the system.

In the formulation specified by (1) and (2), the functions $\mathbf{f}_k$ and $\mathbf{h}_k$, the probability distributions $p_{\mathbf{v}}$ and $p_{\mathbf{w}}$, and the PDF of the state vector at initial time $k = 0$, (i.e. $p_0(\mathbf{x}_0)$), are all assumed known. Eqs. (1) and (2) effectively define two probability functions: the *transitional density* $\pi_{k|k-1}(\mathbf{x}_k|\mathbf{x}_{k-1}) = p_{\mathbf{v}}(\mathbf{x}_k - \mathbf{f}_{k-1}(\mathbf{x}_{k-1}))$ and the *likelihood functiong*$_k(\mathbf{z}_k|\mathbf{x}_k) = p_{\mathbf{w}}(\mathbf{z}_k - \mathbf{h}_k(\mathbf{x}_k))$. The problem is to compute recursively the posterior PDF of the state, denoted as $p_{k|k}(\mathbf{x}_k|\mathbf{z}_{1:k})$ at discrete-time k, where the notation $\mathbf{z}_{1:k}$ stands for the sequence $\mathbf{z}_1, \mathbf{z}_2, \cdots, \mathbf{z}_k$.

The solution is usually presented as a two step procedure. Let $p_{k-1|k-1}(\mathbf{x}_{k-1}|\mathbf{z}_{1:k-1})$ denote the posterior PDF at $k-1$. The first step *predicts* the density of the state to time k via the Chapman–Kolmogorov equation [1]:

$$p_{k|k-1}(\mathbf{x}_k|\mathbf{z}_{1:k-1}) = \int \pi_{k|k-1}(\mathbf{x}_k|\mathbf{x}')\, p(\mathbf{x}'|\mathbf{z}_{1:k-1})\, d\mathbf{x}'. \tag{3}$$

The second step applies Bayes rule to *update*$p(\mathbf{x}_k|\mathbf{z}_{1:k-1})$ using measurement $\mathbf{z}_k$:

$$p_{k|k}(\mathbf{x}_k|\mathbf{z}_{1:k}) = \frac{g_k(\mathbf{z}_k|\mathbf{x}_k)\; p_{k|k-1}(\mathbf{x}_k|\mathbf{z}_{1:k-1})}{\int g_k(\mathbf{z}_k|\mathbf{x})\, p_{k|k-1}(\mathbf{x}|\mathbf{z}_{1:k-1})\, d\mathbf{x}}. \tag{4}$$

Knowing the posterior $p_{k|k}(\mathbf{x}_k|\mathbf{z}_{1:k})$, one can compute a point estimate of the state $\hat{\mathbf{x}}_k$ (e.g. as the mean or the mode of the posterior) and a confidence (or credible) interval.

The closed-form analytic solution to (3) and (4) can be found only in some special cases. One important case is when $\mathbf{f}_k$ and $\mathbf{h}_k$ are linear functions and PDFs $p_{\mathbf{v}}$, $p_{\mathbf{w}}$ and p_0 are Gaussian; the solution in this case is the Kalman filter. In general, however, stochastic filtering via (3) and (4) can be solved only numerically. Many algorithms have been proposed for this purpose, including analytic approximations (e.g. Extended Kalman filter and its variants), grid-based methods (where the posterior PDF is evaluated at a finite and fixed set of points), Gaussian sum filters (where the posterior PDF is approximated by a Gaussian mixture), unscented transforms [11] and particle filters [5,12].

2.2. A primer on the particle method

Suppose the posterior density at discrete-time $k-1$ is approximated by a set of random samples (particles) $\{w_{k-1}^{(i)}, \mathbf{x}_{k-1}^{(i)}\}_{1 \le i \le N}$, where $\mathbf{x}_{k-1}^{(i)}$ is the state of particle i and $w_{k-1}^{(i)}$ is its weight. The weights are normalised, that is $\sum_{i=1}^{N} w_{k-1}^{(i)} = 1$. This approximation of the posterior improves as $N \to \infty$. Given $\{w_{k-1}^{(i)}, \mathbf{x}_{k-1}^{(i)}\}_{1 \le i \le N}$ and using the measurement $\mathbf{z}_k$ at time k, the key question is how to form the particle approximation of the posterior at k, i.e. $p_{k|k}(\mathbf{x}_k|\mathbf{z}_{1:k})$, denoted $\{w_k^{(i)}, \mathbf{x}_k^{(i)}\}_{1 \le i \le N}$.

The computation of the weights and particles at time k is based on the concept of importance sampling [13]. Let us introduce a proposal or importance density $q_k(\mathbf{x}_k|\mathbf{x}_{k-1}, \mathbf{z}_k)$, whose support contains the support of the posterior PDF at time k. Then the (preliminary) particles at time k are drawn from the importance density:

$$\tilde{\mathbf{x}}_k^{(i)} \sim q_k(\mathbf{x}_k|\mathbf{x}_{k-1}^{(i)}, \mathbf{z}_k), \tag{5}$$

whose weights are computed as follows:

$$\tilde{w}_k^{(i)} = w_{k-1}^{(i)} \frac{g_k(\mathbf{z}_k|\tilde{\mathbf{x}}_k^{(i)})\, \pi_{k|k-1}(\tilde{\mathbf{x}}_k^{(i)}|\mathbf{x}_{k-1}^{(i)})}{q_k(\tilde{\mathbf{x}}_k^{(i)}|\mathbf{x}_{k-1}^{(i)}, \mathbf{z}_k)} \tag{6}$$

$$w_k^{(i)} = \frac{\tilde{w}_k^{(i)}}{\sum_{j=1}^{N} \tilde{w}_k^{(j)}} \tag{7}$$

for $i = 1, \ldots, N$. This recursive procedure starts at time $k = 0$ by sampling N times from the initial PDF p_0.

The described particle method, also known as sequential importance sampling (SIS), inevitably fails after many iterations, because all particle weights, except a few, become zero (a poor approximation of the posterior PDF due to particle degeneracy). The collapse of the SIS scheme can be prevented by resampling the particles. The resampling step chooses N particles from $\{w_k^{(i)}, \tilde{\mathbf{x}}_k^{(i)}\}_{1 \le i \le N}$, where the selection of particles is based on their weights: the