Signal Processing 130 (2017) 64–73

Contents lists available at ScienceDirect

Signal Processing

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

Particle-gating SMC-PHD filter

Yiyue Gao [a,*], Defu Jiang [b], Ming Liu [b]

[a] College of Energy and Electrical Engineering, Hohai University, Nanjing 210098, China
[b] College of Computer and Information, Hohai University, Nanjing 210098, China

ARTICLE INFO

Article history:
Received 4 February 2016
Received in revised form
13 June 2016
Accepted 19 June 2016
Available online 21 June 2016

Keywords:
Multi-target tracking
Probability hypothesis density filter
Particle filter
Gating
Clutter density

ABSTRACT

The Sequential Monte Carlo (SMC) implementation for the probability hypothesis density (PHD) filter, referred to as the SMC-PHD filter, is a good candidate for multi-target tracking (MTT) problems. It recursively propagates the weighted particle set that approximates the multi-target posterior density. In this paper, we propose an improved SMC-PHD filter for MTT called the particle-gating SMC-PHD filter. First, a robust gating based on particles propagated from a previous time period is proposed to select the observations of survival targets. Second, a sigma-nearest-gating is proposed to accurately select the observations of new targets. By employing only the observations obtained by the above algorithms to update the state estimations, the overall processing speed of the filter is significantly improved. In addition, a softening factor is suggested to lower the average number of clutters in the updater. This provides more accurate estimation compared with the basic SMC-PHD filter. Finally, the respective real-time and tracking performances of the proposed SMC-PHD filter are verified by the simulation results.

1. Introduction

Multi-target tracking (MTT) jointly determines the time-varying number of targets and their states from a sequence of noisy observations that originate from true targets and clutters. The finite set statistics (FISST) proposed by Mahler [1,2] provides a systematic treatment of these unordered targets and observations —i.e., the states of targets and observations are modeled by a random finite set (RFS) —in which both the number of elements and their values are random. Extending the RFS and the Bayesian inference to MTT problems, the probability hypothesis density (PHD) filter has been proposed [3–6] as a more tractable alternative to the optimal Bayesian filter, which recursively propagates the intensity function of the multi-target posterior density.

The PHD filter has been approximately implemented via the Sequential Monte Carlo (SMC) method [7], the finite Gaussian mixtures (GM) technique [8], or their hybrids [9–11]. Owing to the requirement of the Gaussian mixture assumption, the applications of the GM-PHD filter are constrained in linear and mildly nonlinear system. On the other hand, the SMC-PHD filter is applicable for most general situations and has been receiving considerable attention [12] on account of its advantage of the freedom of linear and Gaussian assumptions.

The computational complexity of the SMC-PHD filter is strongly dependent on the number of targets and observations and the total number of particles [7]. It is thus highly time intensive, especially in dense clutter environments. When the number of particles is fixed, the degradation of the filtering speed is primarily caused by the observations originating from clutters participating in updating particles. If real observations are selected and most clutters are discarded in the updater, the filtering speed can be greatly improved.

An intuitive idea is to select real observations by the state estimates from the previous time. Various gating techniques have been applied in the PHD filter and have achieved good real-time performance [13,14] for the GM-PHD filter and [15–17] SMC-PHD filter. Li et al. [15] proposed a Sigma-gating that updates particles using only the local nearby measurements inside a specified sigma-gate. Zheng et al. [16] used a data-driven mechanism with a gating technique to improve the real-time performance by selecting the measurements nearest to the expected positions of state estimates. The approach of Shi et al. [17], which was similar to that of Zheng et al. [16], opportunistically selected a fixed number of observations from a varying number of observations for filtering. According to our knowledge, these gating techniques are primarily based on state estimates and supply no solution for miss detection.

The multi-target density filter is only based on observation; therefore, disturbed by clutters or detection uncertainty, the state estimates are likely to be erroneous. They may drift away from real positions, or their number may be more or less than the real number of targets. Thus, the output of gating techniques based on state estimates may provide false observations originating from

* Corresponding author.
E-mail addresses: njgyiyue@163.com (Y. Gao), surfer_jiangdf0801@163.com (D. Jiang), 140407080002@hhu.edu.cn (M. Liu).

http://dx.doi.org/10.1016/j.sigpro.2016.06.017
0165-1684/

clutter or it may miss some real observations. In other words, a small amount of survival observations will be discarded and not employed to update the state estimations; accordingly, the track of the small amount of survival targets will be lost.

It is worth noting that if the gating method is based on the state estimation extracted from the multi-sensor multi-Bernoulli filter [18,19], the output will be more accurate than that from the PHD filter. The reason is multi-sensor target tracking supplies a better detection performance compared with single-sensor target tracking, as well as multi-Bernoulli filter [20] provides more accurate state estimation. However, the two imperfections mentioned above will just be improved, not be solved completely. In this case, we still need to continue the investigation on gating method for the single-sensor PHD filter.

A gating method [21] that is free of state estimates was proposed based on the propagated particles from a previous time. This method is based on the fact that, in the SMC-PHD filter, although particles interfered by nearby clutters may produce erroneous state estimates, they inherently contain more real information about survival targets than state estimates. However, all the particles were used to select true observations for designing importance functions of the particle PHD filter in [21], which leads to the overall processing speed not improving or even being somewhat degraded.

In this paper, we employ the propagated particles to select true observations from a novel perspective. By randomly extracting a small amount of particles from the propagated particles, the selection operation does not affect the overall filtering speed; moreover, the posterior density information is sufficiently used. Since the missed target's posterior is only approximated by very few particles, we propose the definition of the ownerless particles and the claimed particles, which supply the gist at the next time to select the observation of a re-detected target.

The benefits of our approach are twofold. First, more accurate tracking is obtained because our particle gating supplies more available real observations than gating methods based on state estimates for the SMC-PHD filter. Moreover, the softening factor cooperates with the gating behavior of eliminating most clutters by lowering the clutter density. Second, high-speed processing is achieved because only the observation nearest to the center in the specified gate participates in updating state estimations. Meanwhile, without a birth observation, no newborn particles are predicted in the sigma-nearest-gating.

The remainder of this paper is organized as follows. The technical background, including the PHD filter, as well as its SMC implementation and influence of state estimates on gating methods, is provided in Section 2. The method of designing the proposed particle-gating SMC-PHD filter is presented in Section 3, where its advantages over current gating methods are analyzed. The performance evaluations and results are described in Section 4. Our concluding remarks are presented in Section 5.

2. Background

2.1. The PHD filter

In a multi-target system, the multi-target state and multi-target observation are two collections of individual targets and observations, in which the state and the observation at time k are two vectors of possibly different dimensions. As targets may survive or die, or a new target appears in time, the dimensions of the two collections may also change in time. Moreover, there exists no ordering for the elements of the multi-target state and observation. Applying the random set theory [3], the multi-target state and observation are naturally represented as finite subsets, X_k and Z_k, which are defined as follows:

The defined $N(k)$ targets are located at $x_{k,1}, \ldots, x_{k,N(k)}$ in the single-target state space E_S (e.g. $\mathbb{R}^{n_x}$) at time k. Then, the multi-target state is represented as:

$$X_k = \left\{ x_{k,1}, \ldots, x_{k,N(k)} \right\} \in \mathcal{F}(E_S) \tag{1}$$

where $\mathcal{F}(E_S)$ is the collection of all finite subsets of the space E_S. These states may be of survive targets, spawned targets, or birth targets. Similarly, there are $M(k)$ observations located at $z_{k,1}, \ldots, z_{k,M(k)}$ in the single-target observation space E_O (e.g. $\mathbb{R}^{n_z}$) at time k, then the multi-target observation is represented as:

$$Z_k = \left\{ z_{k,1}, \ldots, z_{k,M(k)} \right\} \in \mathcal{F}(E_O) \tag{2}$$

in which $\mathcal{F}(E_O)$ is the collection of all finite subsets of the space E_O, and the observation may be derived from a real target, clutter, or false alarm.

Let $Z^{(k)}$: $Z_1, \ldots, Z_k$ denote the time sequence of the observation set. Assuming that each target evolves and generates observation independently of each other, and one target generates no more than one observation at each scan. The clutter points are independent from the observations; moreover, their average number is Poisson-distributed [3]. Based on the above assumptions, the multi-target Bayesian filter [1] is proposed as the theoretically optimal approach to multi-target tracking by recursively propagating the multi-target posterior $p_{k|k}(X_k|Z^{(k)})$.

Moreover, assuming $p_{k|k}(X_k|Z^{(k)})$ and the multi-target prior $p_{k|k-1}(X_k|Z^{(k-1)})$ to be Poisson-distributed, by respectively compressing them into the first moment $D_{k|k}(x_k|Z^{(k)})$ and $D_{k|k-1}(x_k|Z^{(k-1)})$ known as PHD, we have the PHD filter [3] as the approximation of the multi-target Bayesian filter. Their relationship is given as follows:

$$\begin{array}{ccccccc}
\ldots \to p_{k-1|k-1}(X_{k-1}|Z^{(k-1)}) & \to & p_{k|k-1}(X_k|Z^{(k-1)}) & \to & p_{k|k}(X_k|Z^{(k)}) & \to & \ldots \\
\downarrow & & \downarrow & & \downarrow & & \\
\ldots \to D_{k-1|k-1}(x_{k-1}|Z^{(k-1)}) & \to & D_{k|k-1}(x_k|Z^{(k-1)}) & \to & D_{k|k}(x_k|Z^{(k)}) & \to & \ldots
\end{array}$$

in which the top line represents the multi-target Bayesian filter, while the bottom line represents the PHD filter. Additionally, the down arrows represent that the multi-target density is compressed into its first-order moment. The PHD filter is completed by recursively propagating the PHD of the multi-target posterior. Using the abbreviation $D_{k|k}(x_k|Z^{(k)}) \overset{\text{abbr}}{=} D_{k|k}(x_k)$, the PHD predictor is

$$D_{k|k-1}(x_k) = \gamma_k(x_k) + \int \phi_{k|k-1}(x_k, x_{k-1}).\ D_{k-1|k-1}(x_{k-1})dx_{k-1} \tag{3}$$

with one abbreviation used

$$\phi_{k|k-1}(x_k, x_{k-1}) = p_{S,k}(x_{k-1}).\ f_{k|k-1}(x_k|x_{k-1}) + b_{k|k-1}(x_k|x_{k-1})$$

where $\gamma_k(x_k)$ is the PHD of the RFS of the new targets at time k. In addition, $b_{k|k-1}(x_k|x_{k-1})$ is the PHD of the RFS of targets spawned from the previous state x_{k-1}, $f_{k|k-1}(x_k|x_{k-1})$ is the transition density of individual targets, and $p_{S,k}(x_{k-1})$ is the probability that the target still survives at time k.

The PHD updater is

$$D_{k|k}(x_k) = \left(1 - p_{D,k}(x_k) + \sum_{z \in Z_k} \frac{p_{D,k}(x_k) g_k(z|x_k)}{\lambda c(z) + D_{k|k-1}[p_{D,k} g_k]} \right) D_{k|k-1}(x_k) \tag{4}$$

where $p_{D,k}(x_k)$ is the detection probability of a target with state x_k, $g_k(z|x_k)$ is the likelihood of individual targets, and $D_{k|k-1}[h] = \int h(x_{k-1}).\ D_{k|k-1}(x_{k-1})dx_{k-1}$. $\lambda c(z)$ is the PHD of clutter at time k, where λ denotes the average number of clutter points per scan, and $c(.)$ denotes the spatial distribution of each clutter