Automatica 71 (2016) 10–23

Contents lists available at ScienceDirect

Automatica

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

Multivariable feedback particle filter☆

Tao Yang [a], Richard S. Laugesen [b], Prashant G. Mehta [b], Sean P. Meyn [c]

[a] GE Global Research, Niskayuna, NY, USA
[b] University of Illinois, Urbana–Champaign, IL, USA
[c] University of Florida, Gainesville, FL, USA

ARTICLE INFO

Article history:
Received 5 August 2014
Received in revised form
21 December 2015
Accepted 12 April 2016
Available online 31 May 2016

Keywords:
Nonlinear filtering
Particle filtering
Estimation theory

ABSTRACT

This paper presents the multivariable extension of the *feedback particle filter* (FPF) algorithm for the nonlinear filtering problem in continuous-time. The FPF is a control-oriented approach to particle filtering. The approach does not require importance sampling or resampling and offers significant variance improvements; in particular, the algorithm can be applied to systems that are not stable. This paper describes new representations and algorithms for the FPF in the general multivariable nonlinear non-Gaussian setting. Theory surrounding the FPF is improved: Exactness of the FPF is established in the general setting, as well as well-posedness of the associated boundary value problem to obtain the filter gain. A Galerkin finite-element algorithm is proposed for approximation of the gain. Its performance is illustrated in numerical experiments.

1. Introduction

In a recent work, we introduced a new feedback control-based formulation of the particle filter for the nonlinear filtering problem (Yang et al., 2012; Yang, Mehta, & Meyn, 2011a,b, 2013). The resulting filter is referred to as the *feedback particle filter*. In our prior journal article (Yang et al., 2013), the filter was described for the scalar case, where the signal and observation processes are both real-valued. The aim of this paper is to generalize the scalar results of our earlier paper to the multivariable filtering problem:

$$\mathrm{d}X_t = a(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}B_t, \tag{1a}$$

$$\mathrm{d}Z_t = h(X_t)\,\mathrm{d}t + \mathrm{d}W_t, \tag{1b}$$

where $X_t \in \mathbb{R}^d$ is the state at time t, $Z_t \in \mathbb{R}^m$ is the observation vector, and $\{B_t\}, \{W_t\}$ are two mutually independent Wiener processes taking values in $\mathbb{R}^d$ and $\mathbb{R}^m$. The mappings $a(\cdot) : \mathbb{R}^d \to \mathbb{R}^d$, $h(\cdot) : \mathbb{R}^d \to \mathbb{R}^m$ and $\sigma(\cdot) : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ are C^1 functions. The covariance matrix of the observation noise $\{W_t\}$ is assumed to be positive definite. The function h is a column vector whose jth coordinate is denoted as h_j (i.e., $h = (h_1, h_2, \ldots, h_m)^T$). By scaling, we assume without loss of generality that the covariance matrices associated with $\{B_t\}, \{W_t\}$ are identity matrices. Unless otherwise noted, the stochastic differential equations (SDEs) are expressed in Itô form.

The objective of filtering is to estimate the posterior distribution of X_t given the time history of observations $\mathcal{Z}_t := \sigma(Z_s : 0 \le s \le t)$. The density of the posterior distribution is denoted by p^*, so that for any measurable set $A \subset \mathbb{R}^d$,

$$\int_{x\in A} p^*(x,t)\;\mathrm{d}x = \mathsf{P}\{X_t \in A \mid \mathcal{Z}_t\}.$$

The filter is infinite-dimensional since it defines the evolution, in the space of probability measures, of $\{p^*(\,\cdot\,,t) \;:\; t \ge 0\}$. If $a(\,\cdot\,)$, $h(\,\cdot\,)$ are linear functions, the solution is given by the finite-dimensional Kalman–Bucy filter. The article (Budhiraja, Chen, & Lee, 2007) surveys numerical methods to approximate the nonlinear filter. One approach described in this survey is particle filtering.

The particle filter is a simulation-based algorithm to approximate the filtering task (Doucet, de Freitas, & Gordon, 2001). The key step is the construction of N stochastic processes $\{X_t^i : 1 \le i \le N\}$: The value $X_t^i \in \mathbb{R}^d$ is the state for the ith particle at time t. For each time t, the empirical distribution formed by, the particle population is used to approximate the posterior distribution. Recall that this is defined for any measurable set $A \subset \mathbb{R}^d$ by,

$$p^{(N)}(A,t) = \frac{1}{N}\sum_{i=1}^{N} 1\{X_t^i \in A\}.$$

☆ Financial support from the NSF grants 1334987 and 1462773, and the Simons foundation grant 204296 is gratefully acknowledged. The material in this paper was presented at the 51st IEEE Conference on Decision and Control, December 10–13, 2012, Maui, HI, USA. This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen. The conference version of this paper appeared in Yang, Laugesen, Mehta, and Meyn (2012).

E-mail addresses: t.yang@ge.com (T. Yang), laugesen@illinois.edu (R.S. Laugesen), mehtapg@illinois.edu (P.G. Mehta), meyn@ufl.edu (S.P. Meyn).

http://dx.doi.org/10.1016/j.automatica.2016.04.019
0005-1098/

A common approach in particle filtering is called *sequential importance sampling*, where particles are generated according to their importance weight at every time step (Bain & Crisan, 2010; Doucet et al., 2001).

In our earlier papers (Yang et al., 2011a,b, 2013), an alternative feedback control-based approach to the construction of a particle filter was introduced. The resulting particle filter, referred to as the feedback particle filter (FPF), was described for the scalar filtering problem (where $d = m = 1$). The main result of this paper is to present the FPF for the multivariable filtering problem (1a)–(1b). In the following, this algorithm is described followed by a statement of the original contributions of this paper and comparison to relevant literature.

The feedback particle filter is a controlled system. The dynamics of the ith particle have the following gain feedback form,

$$\mathrm{d}X_t^i = a(X_t^i)\,\mathrm{d}t + \sigma(X_t^i)\,\mathrm{d}B_t^i + \underbrace{\mathsf{K}(X_t^i,t)\,\mathrm{d}I_t^i + \Omega(X_t^i,t)\,\mathrm{d}t}_{\mathrm{d}U_t^i}, \tag{2}$$

where $\{B_t^i\}$ are mutually independent standard Wiener processes, $\mathbf{I}_t^i$ is similar to the *innovation process* that appears in the nonlinear filter,

$$\mathrm{d}I_t^i := \mathrm{d}Z_t - \frac{1}{2}(h(X_t^i) + \hat{h})\,\mathrm{d}t, \tag{3}$$

where $\hat{h} := \mathsf{E}[h(X_t^i)|\mathcal{Z}_t]$. In a numerical implementation, we approximate $\hat{h} \approx \frac{1}{N}\sum_{i=1}^{N} h(X_t^i) =: \hat{h}^{(N)}$.

The gain function K is obtained by solving a weighted Poisson equation: For $j = 1, 2, \ldots, m$, the function ϕ_j is a solution to the second-order boundary value problem (BVP),

$$\text{BVP}\qquad \begin{aligned} \nabla\cdot(p(x,t)\nabla\phi_j(x,t)) &= -(h_j(x) - \hat{h}_j)p(x,t),\\ \int \phi_j(x,t)p(x,t)\,\mathrm{d}x &= 0 \quad \text{(normalization)}, \end{aligned} \tag{4}$$

for all $x \in \mathbb{R}^d$ where ∇ and $\nabla\cdot$ denote the gradient and the divergence operators, respectively, and p denotes the conditional density of X_t^i given $\mathcal{Z}_t$, and $\hat{h}_j := \mathsf{E}[h_j(X_t^i)|\mathcal{Z}_t]$. Although this paper is limited to $\mathbb{R}^d$, for domains with boundary, the BVP is accompanied by a Neumann boundary condition,

$$\nabla\phi(x,t)\cdot\hat{n}(x) = 0,$$

for all x on the boundary of the domain where $\hat{n}(x)$ is a unit normal vector at the boundary point x.

In terms of BVP solution, the gain function is given by

$$[\mathsf{K}]_{lj} = \frac{\partial\phi_j}{\partial x_l}. \tag{5}$$

Note that the gain function K is matrix-valued (with dimension $d \times m$) and it needs to be obtained for each value of time t. Also recall that h is a column vector with m entries.

Finally, $\Omega = (\Omega_1, \Omega_2, \ldots, \Omega_d)^T$ is the Wong–Zakai correction term:

$$\Omega_l(x,t) := \frac{1}{2}\sum_{k=1}^{d}\sum_{s=1}^{m}\mathsf{K}_{ks}(x,t)\frac{\partial\mathsf{K}_{ls}}{\partial x_k}(x,t). \tag{6}$$

The controlled system (2)–(6) is called the multivariable feedback particle filter.

The inspiration for controlling a single particle – via the control input U_t^i in (2) – comes from the mean-field game formalism; cf., Huang, Caines, and Malhame (2007) and Yin, Mehta, Meyn, and Shanbhag (2010). With no control input ($U_t^i = 0$), the particle system (2) implements a Monte Carlo propagation of the (unconditioned) probability distribution for (1a). One interpretation of the control input U_t^i is that it implements the 'Bayesian update step' to account for conditioning due to observations (1b). The gain times error structure is reminiscent of the Bayesian update formula in the Kalman filter (see also Remark 1). Structurally, such an update procedure is very different from the importance sampling based implementation of the Bayes rule in conventional particle filters. While the FPF is naturally a continuous-time algorithm, an importance sampling-based procedure typically requires discretization of time; cf., Bain and Crisan (2010). In discrete-time, approximations of the posterior distribution are typically used as importance densities.

The contributions of this paper are as follows:

- **Exactness.** The feedback particle filter (2) is shown to be exact, given an exact initialization $p(\cdot,0) = p^*(\cdot,0)$. Consequently, if the initial conditions $\{X_0^i\}_{i=1}^N$ are drawn from the initial density $p^*(\cdot,0)$ of X_0, then, as $N \to \infty$, the empirical distribution of the particle system approximates the posterior density $p^*(\cdot,t)$ for each t.
- **Well-posedness.** A weak formulation of the BVP (4) is introduced, and used to prove an existence–uniqueness result for ϕ_j in a suitable function space. Certain a priori bounds are derived for the gain function to show that the resulting control input in (2) is admissible. (That is, the filter (2) is well-posed in the Itô sense.)
- **Numerical algorithms.** Based on the weak formulation, a Galerkin finite-element algorithm is proposed for approximation of the gain function $\mathsf{K}(x,t)$. The algorithm is completely adapted to data (that is, it does not require an explicit approximation of $p(x,t)$ or computation of derivatives). Certain closed-form expressions for the gain function are derived in a few special cases. The conclusions are illustrated with numerical examples.

In recent years, there has been a burgeoning interest in application of ideas and techniques from statistical mechanics to nonlinear estimation and control theory. Although some of these applications are classical (see e.g. Del Moral, 2013, 2004; Del Moral, Patras, & Rubenthaler, 2011; Del Moral & Rio, 2011), the recent impetus comes from explosive interest in mean-field games, starting with the two papers from 2007: Lasry and Lions paper titled "Mean-field games" (Lasry & Lions, 2007) and a paper in IEEE TAC by Huang et al. (2007). These papers spurred interest in analysis and synthesis of *controlled interacting particle systems*.

For the continuous-time filtering problem, an approximate particle filtering algorithm appears in the 2009 paper of Crisan and Xiong (2009). In the 2003 paper of Mitter, an optimal control problem for particle filtering is formulated based on duality (Mitter & Newton, 2003). A comparison between the algorithms proposed in these papers and the feedback particle filter appears in Yang et al. (2013). Certain mean-field game inspired approximate algorithms for nonlinear estimation appear in Fallah, Malhamé, and Martinelli (2013a,b), Pequito, Aguiar, Sinopoli, and Gomes (2011). In discrete-time settings, Daum and Huang have introduced the *particle flow filter* algorithm (Daum & Huang, 2010). A detailed comparison of the feedback particle filter to Daum's particle flow filter appears in Yang, Blom, and Mehta (2014). There are by now a growing list of papers on application of such controlled algorithms to: physical activity recognition (Tilton, Hsiao-Wecksler, & Mehta, 2012; Tilton, Mehta, & Meyn, 2013), estimation of soil parameters in dredging applications (Stano, Tilton, & Babuska, 2014), estimation and control in the presence of communication channels (Ma & Coleman, 2011), target state estimation (Daum & Huang, 2010; Tilton, Ghiotto, & Mehta, 2013), satellite tracking (Berntorp, 2015) and weather forecasting (Reich, 2011).

The outline of the remaining part of this paper is as follows. The nonlinear filter is introduced and shown to be exact in Section 2. The weak formulation of the BVP appears in Section 3 where well-posedness results are derived and the numerical Galerkin algorithm is described. A self-contained summary of the finite-N FPF