

Feedback Particle Filter with Mean-field Coupling

Tao Yang, Prashant G. Mehta, Sean P. Meyn

Abstract—Motivated by several recent success stories in the area of mean-field games, in recent work we have introduced a simulation-based approximation to the nonlinear filter for the continuous-discrete time filtering problem [14]. The filter has a feedback control structure: The control is chosen as a solution of a certain variational problem, so that the posterior distribution of a particle matches as closely as possible the posterior distribution of the true state given the observations.

In this paper, we provide an extension, referred to as the *feedback particle filter*, for solution to the continuous time filtering problem. The feedback particle filter admits an innovation error-based feedback structure. The filter is shown to be consistent with the nonlinear filter. An algorithm is introduced and implemented in two numerical examples. A numerical comparison of the feedback particle filter with the bootstrap particle filter is provided.

I. INTRODUCTION

We consider a scalar filtering problem:

$$dX_t = a(X_t) dt + \sigma_B dB_t, \quad (1a)$$

$$dZ_t = h(X_t) dt + \sigma_W dW_t, \quad (1b)$$

where $X_t \in \mathbb{R}$ is the state at time t , $Z_t \in \mathbb{R}$ is the observation process, $a(\cdot)$, $h(\cdot)$ are C^1 functions, and $\{B_t\}$, $\{W_t\}$ are mutually independent standard Wiener processes.

The objective of the filtering problem is to estimate the posterior distribution p^* of X_t given the history $\mathcal{Z}_t := \sigma(Z_s : s \leq t)$. If $a(\cdot)$, $h(\cdot)$ are linear functions, the solution is given by the finite-dimensional Kalman filter. The theory of nonlinear filtering is described in the classic monographs [11], [10]. The filter is infinite dimensional since it defines the evolution, in the space of probability measures, of $\{p^*(\cdot, t) : t \geq 0\}$.

The article [3] surveys numerical methods to approximate the nonlinear filter. One approach described in this survey is particle filtering (see also [1], [5], [12]).

In this paper, we introduce the *feedback particle filter* to approximate the nonlinear filter. As in the usual particle filter, it is formulated through the construction of a large number of coupled stochastic processes $\{X_t^i : 1 \leq i \leq N\}$: The value $X_t^i \in \mathbb{R}$ is the state for the i^{th} particle at time t .

We assume the initial conditions $\{X_0^i\}_{i=1}^N$ are i.i.d., and drawn from initial distribution $p^*(x, 0)$ of X_0 . The dynamics of the i^{th} particle are defined by a controlled system,

$$dX_t^i = a(X_t^i) dt + \sigma_B dB_t^i + v(X_t^i, t) dI_t^i + \frac{1}{2} \sigma_W^2 v(X_t^i, t) v'(X_t^i, t) dt \quad (2)$$

in which $v'(x, t) = \frac{\partial v}{\partial x}(x, t)$, where $\{B_t^i\}$ are mutually independent standard Wiener processes, and I^i is intended to mirror the *innovations process* that appears in the nonlinear filter,

$$dI_t^i := dZ_t - \frac{1}{2} (h(X_t^i) + \hat{h}) dt \quad (3)$$

where $\hat{h} = \frac{1}{N} \sum_{i=1}^N h(X_t^i)$. Then, for any t and any set $A \in \mathcal{B}(\mathbb{R})$, we define the approximation of p^* as the empirical distribution, $p^{(N)}(A, t) := \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{X_t^i \in A\}$.

This filter is an extension of the filter introduced in our earlier paper [14]. As in this previous work, we arrive at a control problem: How to choose the *gain function* $v(x, t)$ so that the empirical distribution $p^{(N)}$ approximates p^* , in the sense that

$$\lim_{N \rightarrow \infty} p^{(N)}(A, t) = \int_A p^*(x, t) dx$$

In this paper we show that the gain function is the solution to the Euler-Lagrange boundary value problem (E-L BVP):

$$-\frac{\partial}{\partial x} \left(\frac{1}{p(x, t)} \frac{\partial}{\partial x} \{p(x, t) v(x, t)\} \right) = \frac{1}{\sigma_W^2} h'(x), \quad (4)$$

with boundary conditions $\lim_{x \rightarrow \pm\infty} p(x, t) v(x, t) = 0$, where $h'(x) = \frac{d}{dx} h(x)$, and $p(x, t)$ denotes the conditional distribution of X_t^i given \mathcal{Z}_t .

Note that the gain function needs to be obtained for each value of time t . If $h'(x) \geq 0$ for each x , it follows from the minimum principle for elliptic BVPs that the gain function v is non-negative valued [7].

The implementation of the particle filter appears to suffer from the same drawback as importance sampling: We must compute the object $p(x, t)$ that we wish to simulate. In practice, one considers appropriate parametrizations, where the parameters are approximated empirically in terms of the particles $\{X_t^i\}_{i=1}^N$ alone.

The contributions of this paper beyond the prior work [14] are as follows:

- **Consistency.** We show that the particle filter model (2) is consistent with nonlinear filter in the following sense: Suppose at time 0, $p(x, 0) = p^*(x, 0)$ and the gain function $v(x, t)$ is obtained as the solution to (4). Then, for all $t > 0$,

$$p(x, t) = p^*(x, t).$$

- **Algorithms.** We propose algorithms for synthesis of the gain function $v(x, t)$. If $a(\cdot)$ and $h(\cdot)$ are linear and the density p is Gaussian, then the gain function admits a closed-form solution in terms of variance alone. The variance is approximated empirically as a sample covariance.

In the nonlinear case, we approximate the density as a sum of Gaussian and provide exact and approximate formulae

Financial support from the NSF grant EECs-0925534 and the AFOSR grant FA9550-09-1-0190 is gratefully acknowledged.

T. Yang, P. Mehta and S. Meyn are with Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign taoyang1@illinois.edu; mehtapg@illinois.edu; meyn@illinois.edu

for the solution of the gain function. With the sum of Gaussian approximation, each Gaussian models a sub-cluster of particles $\{X_t^i\}$.

In recent decades, there have been many important advances in importance sampling based approaches for particle filtering; cf., [1], [5], [3]. The crucial distinction here is that there is no resampling of particles. We believe that the introduction of control in the feedback particle filter has several advantages:

Innovation error. The innovation error-based feedback structure is a key feature of the feedback particle filter (2).

The innovation error is now based on the average value of the prediction $h(X_t^i)$ of the i^{th} -particle and the prediction \hat{h} due to the entire population.

Variance reduction. We believe that the feedback can help reduce the high variance that is sometimes observed in the usual particle filter. Numerical results in Sec V support this claim — See Fig. 2 for a comparison of the feedback particle filter and the bootstrap filter for the linear filtering problem.

Applications. Bayesian inference is an important paradigm used to model functions of certain neural circuits in brain [6]. Compared to techniques that rely on importance sampling, a feedback particle filter may provide a more neurobiologically plausible model to implement filtering and inference functions [14].

The biggest drawback is the need to solve the BVP at each time-step, that additionally requires one to approximate the density. We are encouraged however by the extensive set of tools in feedback control: After all, one rarely needs to solve the HJB equations in closed-form to obtain a reasonable feedback control law [9]. Moreover, there are many approaches in nonlinear and adaptive control to both approximate control laws as well as learn/adapt these in online fashion [2]. For feedback particle filter, this is a subject of continuing research.

The outline of this paper is as follows. In Sec. II, we begin by reviewing the results of our earlier paper [14] that introduced the continuous-discrete time filtering problem. The nonlinear filter is introduced in III, and algorithms discussed in IV. We conclude with a discussion of some numerical results in Sec. V. The proofs all appear in the Appendix.

II. PRELIMINARIES

In this section we briefly summarize the main results of our earlier paper [14] that dealt with the continuous-discrete time filtering problem.

For the continuous-discrete time filtering problem, the equation for dynamics is given by (1a), and the observations are made only at discrete times $\{t_n\}$:

$$Z_{t_n} = h(X_{t_n}) + W_{t_n}, \quad (5)$$

where $\{W_{t_n}\}$ is i.i.d and drawn from $N(0, \sigma_W^2)$. We denote $\mathcal{Z}_n^\square := \sigma\{Z_k : k \leq n\}$.

The particle model in this case is a hybrid dynamical system: At time t_n , assuming that $X_{t_{n-1}}^i$ is given, the i^{th}

particle evolves on $[t_{n-1}, t_n)$ according to the stochastic differential equation

$$dX_t^i = a(X_t^i) dt + \sigma_B dB_t^i, \quad t_{n-1} \leq t < t_n. \quad (6)$$

At the end of this time-horizon, there is a potential jump that is determined by the control input $U_{t_n}^i$:

$$X_{t_n}^i = X_{t_n^-}^i + U_{t_n}^i, \quad (7)$$

where $X_{t_n^-}^i$ denotes the right limit of $\{X_t^i : t_{n-1} \leq t < t_n\}$. The specification (7) defines the initial condition for the process on the next interval $[t_n, t_{n+1})$.

A. Belief maps

For each n we denote:

- 1) p_n^* and p_n^{*-} : The conditional distribution of X_{t_n} given \mathcal{Z}_n^\square and $\mathcal{Z}_{n-1}^\square$, respectively.
- 2) p_n and p_n^- : The conditional distribution of $X_{t_n}^i$ given \mathcal{Z}_n^\square and $\mathcal{Z}_{n-1}^\square$, respectively.

These densities evolve according to recursions of the form,

$$p_n^* = \mathcal{P}^*(p_{n-1}^*, Z_{t_n}), \quad p_n = \mathcal{P}(p_{n-1}, Z_{t_n}). \quad (8)$$

The mappings \mathcal{P}^* and \mathcal{P} can be decomposed into two parts. The first part is identical for each of these mappings: The transformation that takes p_{n-1} to p_n^- coincides with the mapping from p_{n-1}^* to p_n^{*-} . In each case it is defined by the Kolmogorov forward equation associated with the diffusion on $[t_{n-1}, t_n)$.

The second part of the mapping is different: The transformation that takes p_n^{*-} to p_n^* is obtained from Bayes' rule. The transformation that takes p_n^- to p_n depends upon the choice of control $U_{t_n}^i$ in (7).

In [14], the input $U_{t_n}^i$ is restricted to be an *admissible function* of $X_{t_n}^i$ and the observations $\{Z_{t_k} : k \leq n\}$.

Definition 1 (Admissible function): A function $u : \mathbb{R} \rightarrow \mathbb{R}$ is said to be *admissible* if u is twice continuously differentiable and

$$\lim_{x \rightarrow \pm\infty} u(x) p_n^-(x) = 0.$$

The space of admissible functions is denoted as C_b^2 . ■ We suppress the dependency on the observations, writing $U_{t_n}^i = u(x)$ when $X_{t_n}^i = x$.

B. Variational Problem

Our goal is to choose an admissible function u so that the mapping \mathcal{P} approximates the mapping \mathcal{P}^* in (8). In particular, we assume $p_n^- = p_n^{*-}$ and seek

$$u_n(x) = \arg \min_{u \in C_b^2} \text{KL}(p_n \| p_n^*), \quad (9)$$

where $\text{KL}(\cdot)$ denote the Kullback-Leibler (KL) divergence between the two distributions. The KL divergence can be expressed as

$$\begin{aligned} \text{KL}(p_n \| p_n^*) = C - \int_{\mathbb{R}} p_n^-(x) \left\{ \ln |1 + u'(x)| \right. \\ \left. + \ln(p_n^-(x + u(x)) p_{Z_{t_n}}(Z_{t_n} | x + u(x))) \right\} dx \end{aligned} \quad (10)$$

where $p_{Z_t|X_t}(\cdot) = \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp\left(-\frac{(Z_t - h(\cdot))^2}{2\sigma_W^2}\right)$, and $C = \int_{\mathbb{R}} p_n^-(x) \ln(p_n^-(x) p_{Z_t|X_t}(Z_t)) dx$ is a constant that does not depend on u ; see [14] for the calculation.

The solution to (9) is described in the following:

Proposition 1 (Proposition 1 in [14]): Suppose that the admissible function u is a minimizer for the optimization problem (9). Then it is a solution of the following Euler-Lagrange (E-L) BVP:

$$\frac{d}{dx} \left(\frac{p_n^-(x)}{|1+u'(x)|} \right) = p_n^-(x) \frac{\partial}{\partial u} (\ln(p_n^-(x+u) p_{Z_t|X_t}(Z_t|x+u))), \quad (11)$$

with boundary conditions $\lim_{x \rightarrow \pm\infty} u(x) p_n^-(x) = 0$. ■

We refer to the minimizer as the *optimal control function*.

III. CONTINUOUS-TIME FILTERING

Consider now the continuous time filtering problem (1a, 1b) introduced in Section I.

We denote as $p^*(x, t)$ the conditional distribution of X_t given $\mathcal{Z}_t = \sigma(Z_s : s \leq t)$. The evolution of $p^*(x, t)$ is described by the Kushner-Stratonovich (K-S) equation:

$$dp^* = \mathcal{L}^\dagger p^* dt + \frac{1}{\sigma_W^2} (h - \hat{h})(dZ_t - \hat{h} dt) p^*, \quad (12)$$

where $\hat{h} = \int h(x) p^*(x, t) dx$.

A. Belief state dynamics & control architecture

The model for the particle filter is given by,

$$dX_t^i = (a(X_t^i) + u(X_t^i, t)) dt + v(X_t^i, t) dZ_t + \sigma_B dB_t^i, \quad (13)$$

where $X_t^i \in \mathbb{R}$ is the state for the i^{th} particle at time t , and $\{B_t^i\}$ are mutually independent standard Wiener processes. We assume the initial conditions $\{X_0^i\}_{i=1}^N$ are i.i.d., independent of $\{B_t^i\}$, and drawn from the initial distribution $p^*(x, 0)$ of X_0 . Both $\{B_t^i\}$ and $\{X_0^i\}$ are also assumed to be independent of X_t, Z_t .

There are two types of conditional distributions of interest in our analysis:

- 1) $p(x, t)$: Defines the conditional dist. of X_t^i given \mathcal{Z}_t .
- 2) $p^*(x, t)$: Defines the conditional dist. of X_t given \mathcal{Z}_t .

The functions $\{u(x, t), v(x, t)\}$ represent the continuous time-counterparts of the optimal control function $u_n(x)$ (see (9)). We say that these functions are *optimal* if $p \equiv p^*$. That is, given $p^*(\cdot, 0) = p(\cdot, 0)$, our goal is to choose $\{u, v\}$ in the feedback particle filter so that the evolution equations of these conditional distributions coincide (see (12) and (15)).

For purposes of control, we regard $\{\pi_t\}$ as the *belief state*: For any $A \in \mathcal{B}(\mathbb{R})$ we write $\pi_t(A) := \int_A p(x, t) dx = \mathbb{P}(X_t^i \in A | \mathcal{Z}_t)$. In the theory of optimal control for HMMs it is known that the belief state may be regarded as a state process, and in this way partially observed optimal control is translated to a fully observed control problem. We adopt a similar formalism here: The state process for purposes of control is $\Phi_t^i = (X_t^i, \pi_t)$, evolving on $\mathbb{R} \times \mathcal{P}(\mathbb{R})$, where $\mathcal{P}(\mathbb{R})$ denotes the space of probability measures on \mathbb{R} . The control for the particle filter is *static state feedback* in the following sense: We obtain an optimal pair $\{u, v\}$ of the

form $u(x, t) = \mathcal{U}(x, \pi_t)$, $v(x, t) = \mathcal{V}(x, \pi_t)$, where \mathcal{U} and \mathcal{V} are real-valued mappings on $\mathbb{R} \times \mathcal{P}(\mathbb{R})$.

As in Sec II, we require these functions to be admissible:

Definition 2 (Admissible functions): The pair of functions $\{u, v\}$ is said to be *admissible* if, for each x, t , the random variables $u(x, t)$ and $v(x, t)$ are $\mathcal{Z}_t = \sigma(Z_s : s \leq t)$ measurable. Moreover, with probability one, $u : \mathbb{R}^2 \rightarrow \mathbb{R}$, $v : \mathbb{R}^2 \rightarrow \mathbb{R}$ are twice continuously differentiable in their first argument, and satisfy for each t ,

$$\lim_{x \rightarrow \pm\infty} u(x, t) p(x, t) = 0, \quad (14a)$$

$$\lim_{x \rightarrow \pm\infty} v(x, t) p(x, t) = 0. \quad (14b)$$

where p is the posterior distribution of X_t^i given \mathcal{Z}_t . The space of admissible functions $\{u, v\}$ is denoted C_b^2 . ■

The evolution equation for the belief state is described in the next result. Its proof appears in Appendix A.

Proposition 2: Consider the process X_t^i that evolves according to the particle filter model (13). The conditional distribution of X_t^i given the filtration \mathcal{Z}_t , $p(x, t)$, satisfies the forward equation

$$dp = \mathcal{L}^\dagger p dt - \frac{\partial}{\partial x} (vp) dZ_t - \frac{\partial}{\partial x} (up) dt + \sigma_W^2 \frac{1}{2} \frac{\partial^2}{\partial x^2} (pv^2) dt, \quad (15)$$

where $\mathcal{L}^\dagger p = -\frac{\partial(pa)}{\partial x} + \frac{\sigma_B^2}{2} \frac{\partial^2 p}{\partial x^2}$. ■

B. Consistency with the nonlinear filter

The main result of this section is the construction of an optimal pair $\{u, v\}$ under the following general assumption:

Assumption A1 The conditional distributions (p^*, p) are C^2 , with $p^*(x, t) > 0$ and $p(x, t) > 0$, for all $x \in \mathbb{R}$, $t > 0$. ■

We henceforth choose $\{u, v\}$ as the solution to a certain E-L BVP based on p : The function v as the solution to

$$-\frac{\partial}{\partial x} \left(\frac{1}{p(x, t)} \frac{\partial}{\partial x} \{p(x, t)v(x, t)\} \right) = \frac{1}{\sigma_W^2} h'(x), \quad (16)$$

with boundary condition (14b). The function $u(\cdot, t) : \mathbb{R} \rightarrow \mathbb{R}$ is obtained as:

$$u(x, t) = v(x, t) \left(-\frac{1}{2}(h(x) + \hat{h}) + \frac{1}{2}\sigma_W^2 v'(x, t) \right), \quad (17)$$

where $\hat{h} = \int h(x) p(x, t) dx$. We assume moreover that the boundary conditions given in (14a) also hold, so that $\{u, v\}$ is admissible. The BVP is motivated by the continuous-time limit of (11), obtained on letting $t_{n+1} - t_n$ go to zero; the calculations appear in Appendix B.

Existence and uniqueness of $\{u, v\}$ is obtained in the following proposition — Its proof is given in Appendix B.

Proposition 3: Consider the BVP (16), subject to Assumption A1. Then,

- 1) There exists a unique solution v , subject to the boundary condition (14b).
- 2) The solution satisfies $v(x, t) \geq 0$ for all x, t , provided $h'(x) \geq 0$ for all x . ■

The following theorem shows that the two evolution equations (12) and (15) are identical. The proof appears in Appendix C.

Theorem 1: Consider the two evolution equations for p and p^* , defined according to the solution of the forward equation (15) and the K-S equation (12), respectively. Suppose that the control functions $u(x,t)$ and $v(x,t)$ are obtained according to (16) and (17), respectively. Then, provided $p(x,0) = p^*(x,0)$, we have for all $t \geq 0$,

$$p(x,t) = p^*(x,t)$$

■

C. Linear Gaussian case

We provide here a special case for linear system:

$$dX_t = a' X_t dt + \sigma_B dB_t, \quad (18a)$$

$$dZ_t = h' X_t dt + \sigma_W dW_t, \quad (18b)$$

where a', h' are real numbers. We assume the initial distribution $p^*(x,0)$ is Gaussian with mean μ_0 and variance Σ_0 .

The following lemma provides the solution of the optimal control functions $u(x,t), v(x,t)$ in the linear Gaussian case.

Lemma 1: Consider the linear observation equation (18b). Suppose $p(x,t) = \frac{1}{\sqrt{2\pi\Sigma_t}} \exp(-\frac{(x-\mu_t)^2}{2\Sigma_t})$ is assumed to be Gaussian with mean μ_t and variance Σ_t . Then the solution of E-L BVP (4) is given by:

$$v(x,t) = \frac{h'\Sigma_t}{\sigma_W^2}, \quad u(x,t) = -\frac{(h')^2\Sigma_t}{2\sigma_W^2}(x + \mu_t) \quad (19)$$

■

The formulae (19) are verified by direct substitution in the ODE (4) where the distribution p is Gaussian.

The optimal control yields the following form for the particle filter in this linear Gaussian model:

$$dX_t^i = a' X_t^i dt + \sigma_B dB_t^i + \frac{h'\Sigma_t}{\sigma_W^2} \left(dZ_t - h' \frac{X_t^i + \mu_t}{2} dt \right). \quad (20)$$

Now we show that $p = p^*$ in this case. That is, the conditional distributions of X and X^i coincide, and are defined by the well-known dynamic equations that characterize the mean and the variance of the continuous-time Kalman filter.

Theorem 2: Consider the linear Gaussian filtering problem defined by the state-observation equations (18a,18b). In this case the posterior distributions of X and X^i are Gaussian, whose conditional mean and covariance are given by the respective SDE and the ODE,

$$d\mu_t = a' \mu_t dt + \frac{h'\Sigma_t}{\sigma_W} \left(dZ_t - h' \mu_t dt \right)$$

$$\frac{d}{dt} \Sigma_t = 2a' \Sigma_t + \sigma_B^2 - \frac{(h')^2 \Sigma_t^2}{\sigma_W^2}$$

■

The result is verified by substituting $p(x,t) = \frac{1}{\sqrt{2\pi\Sigma_t}} \exp(-\frac{(x-\mu_t)^2}{2\Sigma_t})$ in the forward equation (15). The details are omitted on account of space, and because the result is a special case of Theorem 1.

Notice that particle system (20) is not practical since it requires computation of the conditional mean and variance $\{\mu_t, \Sigma_t\}$. If we are to compute these quantities, then there is no reason to run a particle filter!

In practice $\{\mu_t, \Sigma_t\}$ are approximated as sample means and sample covariances from the ensemble $\{X_t^i\}_{i=1}^N$:

$$\mu_t \approx \mu_t^{(N)} := \frac{1}{N} \sum_{i=1}^N X_t^i,$$

$$\Sigma_t \approx \Sigma_t^{(N)} := \frac{1}{N-1} \sum_{i=1}^N (X_t^i - \mu_t^{(N)})^2.$$

The resulting equation (20) for the i^{th} particle is given by

$$dX_t^i = a' X_t^i dt + \sigma_B dB_t^i + \frac{h'\Sigma_t^{(N)}}{\sigma_W^2} \left(dZ_t - h' \frac{X_t^i + \mu_t^{(N)}}{2} dt \right). \quad (21)$$

It is very similar to the mean-field ‘‘synchronization-type’’ control laws and oblivious equilibria constructions as in [8], [15], [13]. As $N \rightarrow \infty$, the empirical distribution of the particle system approximates the posterior distribution $p^*(x,t)$ (by Theorem 2).

IV. SYNTHESIS OF THE GAIN FUNCTION $v(x,t)$

Implementation of the nonlinear filter (2) requires solution of the E-L BVP (4) to obtain the gain function $v(x,t)$ for each fixed t . If $p(x,t)$ is known, the linear BVP admits a closed-form solution (43) – the main issue thus is the approximation of the distribution $p(x,t)$.

In this section, we consider the following approximation:

Assumption A2 For each fixed t , the distribution $p(x,t)$ is a sum of Gaussian:

$$p(x,t) = \sum_{j=1}^m \lambda_t^j q_t^j(x), \quad (22)$$

where $q_t^j(x) = q(x; \mu_t^j, \Sigma_t^j) = \frac{1}{\sqrt{2\pi\Sigma_t^j}} \exp(-\frac{(x-\mu_t^j)^2}{2\Sigma_t^j})$,

$\lambda_t^j > 0, \sum \lambda_t^j = 1$. We assume an ordering so $\mu_t^j < \mu_t^k$ for $j < k$. ■

The approximation is motivated by the numerical algorithm: At each discrete time-step t , we have particle states $\{X_t^i\}_{i=1}^N$. We identify m -clusters each of which is assumed to be localized in \mathbb{R} . We approximate the j^{th} -cluster with a Gaussian pdf with weight $\lambda_t^j \in (0,1)$, empirical mean μ_t^j and variance Σ_t^j .


For the ease of presentation, we also assume $h(x) = x$ in the observation model (1b):

Assumption A3 We assume the observation model (1b) with $h(x) \equiv x$.

This assumption is not critical. Other modalities can also be considered as discussed in Remark 1 below.

The following proposition provides a closed-form solution of the E-L BVP with $p(x,t)$ of the form (22):

Proposition 4: Consider the BVP (4) with $h(x) = x$. Suppose $p(x,t)$ is of the form (22). Then the solution is given



j	λ^j	μ^j
1	0.2	-1
2	0.5	0
3	0.3	1

TABLE I

TABLE II: Taylor series coefficients	
v_0	$\Sigma^j + \sqrt{2\pi\Sigma^j} \left(\frac{1}{2}(\hat{h} - \mu^j) + \sum_{k < j} \frac{\lambda^k}{\lambda^j} (\hat{h} - \mu^k) \right)$
v_1	$(\hat{h} - \mu^j)$
v_2	$\sqrt{\frac{2\pi}{\Sigma^j}} \left(\frac{1}{2}(\hat{h} - \mu^j) + \sum_{k < j} \frac{\lambda^k}{\lambda^j} (\hat{h} - \mu^k) \right)$

(c)

Fig. 1. (a) Comparison of the solutions (23) and (25). (b) Comparison of the exact (solid) and approximate (dashed) solutions. In the background, the three Gaussians are also shown. (c) For both parts (a) and (b) $p(x, t)$ is a sum of ($m = 3$) Gaussian densities with weights λ^j and means μ^j tabulated in Table I. Table II tabulates the formulae for the Taylor series coefficients in (26).

by

$$pv(x, t) = \frac{1}{\sigma_W^2} \left(\sum_{j=1}^m \lambda_t^j (\hat{h} - h(\mu_t^j)) Q_t^j(x) + \sum_{j=1}^m h'(\mu_t^j) \lambda_t^j \Sigma_t^j q_t^j(x) \right), \quad (23)$$

where $Q_t^j(x) = \int_{-\infty}^x q_t^j(y) dy$. ■

The proof is omitted – it is a straightforward verification by direct substitution of the solution in the ODE (4).

Remark 1: For general h , the expression in (23) represents an approximate solution of the E-L BVP in the asymptotic limit that $\Sigma_t^j \rightarrow 0$ for $j = 1, \dots, m$.

This is seen by considering a BVP

$$-\frac{\partial}{\partial x} \left(\frac{1}{p(x, t)} \frac{\partial w}{\partial x} \right) = \frac{1}{\sigma_W^2} h'(x), \quad (24)$$

with boundary condition $\lim_{x \rightarrow \pm\infty} w(x) = 0$. Consider also a limiting distribution $p^0(x, t) = \sum_{j=1}^m \lambda_t^j \delta(x - \mu_t^j)$. With $p(x, t) = p^0(x, t)$ in (24), the weak solution is given by the staircase function:

$$w(x) = w^0(x) := \frac{1}{\sigma_W^2} \sum_{j=1}^m \lambda_t^j (\hat{h} - h(\mu_t^j)) H(x - \mu_t^j), \quad (25)$$

where $H(\cdot)$ is the Heaviside function.

For small values of Σ_t^j , the solution given by (23) is a small perturbation of $w^0(x)$ in (25); see Fig. 1(a) which depicts the two solutions for $h(x) = x$ and (λ_t^j, μ_t^j) tabulated in Table I and $\Sigma_t^j = 0.01$.

In the case of general h , the approximate nature of solution (23) follows by considering a perturbation argument in the asymptotic limit $\Sigma_t^j \rightarrow 0$. ■

Now given $w(x, t) = pv(x, t)$, the gain function is obtained as $v(x, t) = w(x, t)/p(x, t)$. Since dividing by p may be a problem, we also provide asymptotic formulae for values of x near the j^{th} empirical mean ($x \approx \mu_t^j$):

1) For $2 \leq j \leq m-1$, the gain function may be approximated by using a Taylor series approximation: For $x \approx \mu_t^j$,

$$v(x, t) = \frac{1}{\sigma_W^2} \left(v_0 + v_1(x - \mu_t^j) + v_2(x - \mu_t^j)^2 \right), \quad (26)$$

where formulae for the coefficients appear in Table II.

2) For $i = 1$ or m , the Taylor series does not yield a good approximation and one can use the following formula:

$$\text{For } x \approx \mu_t^1 : v(x, t) = \frac{1}{\sigma_W^2} \left(\Sigma_t^1 + (\hat{h} - \mu_t^1) \frac{Q_t^1(x, t)}{q_t^1(x, t)} \right), \quad (27a)$$

$$\text{For } x \approx \mu_t^m : v(x, t) = \frac{1}{\sigma_W^2} \left(\Sigma_t^m + (\hat{h} - \mu_t^m) \frac{Q_t^m - 1}{q_t^m}(x, t) \right). \quad (27b)$$

Using (27b), $\lim_{x \rightarrow -\infty} v(x, t) = \frac{\Sigma_t^1}{\sigma_W^2}$, $\lim_{x \rightarrow \infty} v(x, t) = \frac{\Sigma_t^m}{\sigma_W^2}$. So, the asymptotic value of the gain is consistent with the formula for the gain obtained in the linear case (see (19)).

Figure 1 (b) depicts a comparison between the exact numerical solution v and the approximate formulae; parameter values are the same as in Table I. We obtain the approximate solution for $x \in [\mu_t^j - 4\sqrt{\Sigma_t^j}, \mu_t^j + 4\sqrt{\Sigma_t^j}]$ – for $j = 2$ using (26), and for $j = 1$ and 3 using (27a) and (27b), respectively.

Remark 2: Depending upon the problem as well as the available computational resources, one may choose different approximate structures for the gain function. For example, using the Taylor series approximation, one possibility is to choose $v(x, t) = \frac{1}{\sigma_W^2} v_0$, a constant for the j^{th} cluster. A better choice, however, may be to pick the constant to be the mean value taken over ± 1 standard deviation, i.e., $v(x, t) = \frac{1}{\sigma_W^2} (v_0 + \frac{1}{3} v_2 \Sigma_t^j)$.

A. Algorithm

At time t , we assume m clusters with empirical mean μ_t^j and variance Σ_t^j , ordered such that $\mu_t^j < \mu_t^k$ for $j < k$.

We assign the i^{th} -particle to the j^{th} -cluster based on proximity ($j = \arg \min_k |X_t^i - \mu_t^k|$) and evaluate the gain function $v(X_t^i, t)$ according to (27a, 27b) or (26).

Now, one problem with directly implementing the particle filter model (2) is that there possibly are multiple time-scales in the problem: $v(x, t)$ may become very large for particles that lie between two clusters (see Fig. 1(b)). So, dt would need to be chosen extremely small to avoid numerical instabilities. This is impractical.

To help deal with this issue, we propose to break up the trajectory into the following two phases:

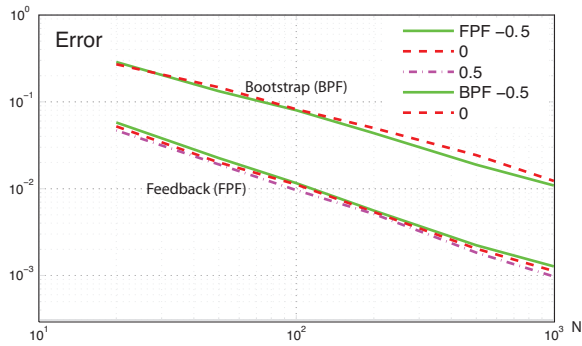


Fig. 2. Comparison of the mse using feedback particle filter and the bootstrap particle filter.

- (i) *Control phase.* If $v(X_t^i, t) < Cv(\mu_t^i, t)$ then we use the particle filter model (2) to obtain dX_t^i . The constant $C > 1$ is selected based on the available dt .
- (ii) *Flight phase.* If $v(X_t^i, t) \geq Cv(\mu_t^i, t)$ then we set $X_t^i = \mu_t^j$ or $X_t^i = \mu_t^{j+1}$ depending upon the sign of the innovation term dI_t^i .

Note that the flight phase implements motion during the fast time-scale: It allows the particle to escape the cluster without the need to make dt arbitrarily small.

V. NUMERICS

A. Linear case

We first provide a comparison between the feedback particle filter and the bootstrap particle filter for the linear problem (18a, 18b).

For the linear filtering problem, the optimal solution is given by the Kalman filter. We use this solution to define the relative mean-squared error:

$$mse = \frac{1}{T} \int_0^T \left(\frac{\Sigma_t^{(N)} - \Sigma_t}{\Sigma_t} \right)^2 dt, \quad (28)$$

where Σ_t is the error covariance using the Kalman filter, and $\Sigma_t^{(N)}$ is its approximation using the particle filter.

Figure 2 depicts a comparison between mse obtained using the feedback particle filter (21) and the bootstrap filter. The latter implementation is based on an algorithm taken from Ch. 9 of [1]. For simulation purposes, we used a range of values of $a' \in [-1, 1]$, $h = 3$, $\sigma_B = 1$, $\sigma_W = 0.5$, $dt = 0.01$, and $T = 50$. The plot in Fig. 2 is generated using simulations with $N = 20, 50, 100, 200, 500, 1000$ particles.

We refer the reader to our earlier paper [14] for additional simulation plots for the pdf and estimates obtained using the linear feedback particle filter (21).

B. Nonlinear example

We consider

$$dX_t = X_t(1 - X_t^2) dt + \sigma_B dB_t, \quad (29a)$$

$$dZ_t = X_t dt + \sigma_W dW_t, \quad (29b)$$

where $\sigma_B = 0.4$, $\sigma_W = 0.2$. Without noise, the ODE (29a) has two stable equilibria at ± 1 . With noise, the state of the SDE “transitions” between these two “equilibria” (see Fig. 3).




Fig. 3. Comparison of the true state $X(t)$ and the conditional mean $\bar{X}(t)$ by using feedback particle filter. The error $\mathcal{E}(t) = X(t) - \bar{X}(t)$ remains small even during a transition of the state.

Figure 3 depicts the simulation results obtained using the nonlinear feedback particle filter (2). The implementation is based on the algorithm presented in Sec IV-A with $dt = 0.01$. The control gain $v(x, t)$ is chosen via the constant approximation discussed in Remark 2. We initialize the simulation with $m = 2$ clusters at ± 1 . After a brief period of transients, these clusters merge into a single cluster, which adequately tracks the true state including the transition events.

We refer the reader to our earlier paper [14] for another numerical example that uses the nonlinear feedback particle filter (2) for filtering of a nonlinear oscillator system.

APPENDIX

A. Derivation of the Forward Equation

We denote the filtration $\mathcal{B}_t = \sigma(X_0^i, B_s^i : s \leq t)$, and we recall that $\mathcal{Z}_t = \sigma(Z_s : s \leq t)$ for $t \geq 0$. These two filtrations are *independent* by construction.

On denoting $\tilde{a}(x, t) = a(x) + u(x, t)$, the particle evolution (13) is expressed,

$$X_t^i = X_0^i + \int_0^t \tilde{a}(X_s^i, s) ds + \int_0^t v(X_s^i, s) dZ(s) + \sigma_B B_t^i. \quad (30)$$

By assumption on Lipschitz continuity of \tilde{a} and v , there exists a unique solution that is adapted to the larger filtration $\mathcal{B}_t \vee \mathcal{Z}_t = \sigma(X_0^i, B_s^i, Z_s : s \leq t)$. In fact, there is a functional F_t such that

$$X_t^i = F_t(X_0^i, B_t^i, Z^t), \quad (31)$$

where $Z^t := \{Z_s : 0 \leq s \leq t\}$ denotes the trajectory.

The conditional distribution of X_t^i given $\mathcal{Z}_t = \sigma(Z_s : s \leq t)$ was introduced in Sec. II-A: Its density is denoted $p(x, t)$, defined by any bounded and measurable function $f: \mathbb{R} \rightarrow \mathbb{R}$ via,

$$\mathbb{E}(f(X_t^i) | \mathcal{Z}_t) = \int_{\mathbb{R}} p(x, t) f(x) dx =: \langle p_t, f \rangle.$$

We begin by stating a lemma that is the key to proving Proposition 2.

Lemma 2: Suppose that f is an $\mathcal{B}_t \vee \mathcal{Z}_t$ -adapted process satisfying $\mathbb{E} \int_0^t |f(s)|^2 ds < \infty$. Then,

$$\mathbb{E} \left[\int_0^t f(s) ds | \mathcal{Z}_t \right] = \int_0^t \mathbb{E}[f(s) | \mathcal{Z}_s] ds, \quad (32)$$

$$\mathbb{E} \left[\int_0^t f(s) dZ_s | \mathcal{Z}_t \right] = \int_0^t \mathbb{E}[f(s) | \mathcal{Z}_s] dZ_s. \quad (33)$$

We next provide a proof of the Proposition 2 and follow it up with the proof of the Lemma 2. ■

Proof of Proposition 2 Applying Itô's formula to equation (13) gives, for any smooth and bounded function f ,

$$df(X_t^i) = \mathcal{L}f(X_t^i) dt + v(X_t^i, t) \frac{\partial f}{\partial x}(X_t^i) dZ_t + \sigma_B \frac{\partial f}{\partial x}(X_t^i) dB_t^i,$$

where $\mathcal{L}f := (a+u) \frac{\partial f}{\partial x} + \frac{1}{2}(\sigma_W^2 v^2 + \sigma_B^2) \frac{\partial^2 f}{\partial x^2}$. Therefore,

$$f(X_t^i) = f(X_0^i) + \int_0^t \mathcal{L}f(X_s^i) ds + \int_0^t v(X_s^i, s) \frac{\partial f}{\partial x}(X_s^i) dZ_s + \sigma_B \int_0^t \frac{\partial f}{\partial x}(X_s^i) dB_s^i.$$

Taking conditional expectations on both sides,

$$\begin{aligned} \langle p_t, f \rangle &= E(f(X_0^i) | \mathcal{Z}_t) + E \left[\int_0^t \mathcal{L}f(X_s^i) ds | \mathcal{Z}_t \right] \\ &\quad + E \left[\int_0^t v(X_s^i, s) \frac{\partial f}{\partial x}(X_s^i) dZ_s | \mathcal{Z}_t \right] \\ &\quad + \sigma_B E \left[\int_0^t \frac{\partial f}{\partial x}(X_s^i) dB_s^i | \mathcal{Z}_t \right] \end{aligned}$$

On applying Lemma 2, and the fact that B_t^i is a Wiener process, we conclude that

$$\langle p_t, f \rangle = \langle p_0, f \rangle + \int_0^t \langle p_s, \mathcal{L}f \rangle ds + \int_0^t \langle p_s, v \frac{\partial f}{\partial x} \rangle dZ_s.$$

The forward equation (15) follows using integration by parts. ■

We now provide a proof of Lemma 2.

Proof of Lemma 2 The key is the functional form (31) of the solution X_t^i : It says that apart from the past values of Z , the solution depends only upon initial condition X_0^i and Wiener process B_t^i that are both independent of Z .

First we suppose that f is simple, i.e.,

$$f(s) = \sum_{i=1}^k F_i 1_{(a_i, b_i]}(s),$$

where $(a_i, b_i]$ are disjoint intervals of $[0, t]$ and F_i is measurable with respect to $\mathcal{B}_{a_i} \vee \mathcal{L}_{a_i}$. For general f satisfying the assumptions of the lemma, the result will then follow via an application of the dominated convergence theorem [4].

Once we restrict to simple functions, the essence of the proof is to establish the identity,

$$E(F_i | \mathcal{Z}_t) = E(F_i | \mathcal{L}_{a_i}). \quad (34)$$

Under the measurability assumption we can write $F_i = \phi(\zeta, \xi)$, where $\zeta \in \mathcal{L}_{a_i}$, $\xi \in \mathcal{B}_{a_i}$ are random variables, and ϕ is a real-valued function. The random variable ξ is independent of \mathcal{Z}_t , so that

$$E(F_i | \mathcal{Z}_t) = E(\bar{\phi}(\zeta) | \mathcal{Z}_t),$$

with $\bar{\phi}(\cdot) = E[\phi(\cdot, \xi)]$. Using the fact that $\zeta \in \mathcal{L}_{a_i} \subset \mathcal{Z}_t$ we obtain (34):

$$E(F_i | \mathcal{Z}_t) = \bar{\phi}(\zeta) |_{\mathcal{Z}_t} = E(F_i | \mathcal{L}_{a_i})$$

The desired results follow easily from (34): To obtain (32) we write,

$$\begin{aligned} E \left(\int_0^t f(s) ds | \mathcal{Z}_t \right) &= \sum_{i=1}^k E(F_i (b_i - a_i) | \mathcal{Z}_t) \\ &= \sum_{i=1}^k E(F_i | \mathcal{L}_{a_i}) (b_i - a_i) \\ &= \int_0^t E(f(s) | \mathcal{L}_s) ds. \end{aligned}$$

The proof of (33) is similar:

$$\begin{aligned} E \left[\int_0^t f(s) dZ_s | \mathcal{Z}_t \right] &= \sum_{i=1}^k E[F_i (Z_{b_i} - Z_{a_i}) | \mathcal{Z}_t] \\ &= \sum_{i=1}^k E[F_i | \mathcal{L}_{a_i}] (Z_{b_i} - Z_{a_i}) \\ &= \int_0^t E[f(s) | \mathcal{Z}_s] dZ_s, \end{aligned}$$

where the second equality uses the fact that Z is adapted to \mathcal{Z}_t , and $a_i < b_i \leq t$ for each i . ■

B. Euler-Lagrange BVP

In this section we describe, formally, the continuous-time limit of the discrete-time E-L BVP (11).

In the continuous-time case, the control is of the form:

$$U_t^i = u(X_t^i, t) dt + v(X_t^i, t) dZ_t. \quad (35)$$

Substituting this in the E-L BVP (11) for the continuous-discrete time case, we arrive at the formal equation:

$$\begin{aligned} \frac{\partial}{\partial x} \left(\frac{p(x, t)}{1 + u' dt + v' dZ_t} \right) &= p(x, t) \frac{\partial}{\partial \tilde{u}} \left(\ln p(x + u dt + v dZ_t, t) \right. \\ &\quad \left. + \ln p_{\text{az}|x}(dZ_t | x + u dt + v dZ_t) \right), \end{aligned} \quad (36)$$

where $p_{\text{az}|x}(dZ_t | \cdot) = \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp\left(-\frac{(dZ_t - h(\cdot))^2}{2\sigma_W^2}\right)$ and $\tilde{u} = u dt + v dZ_t$.

For notational ease, we use primes to denote partial derivatives with respect to x : p is used to denote $p(x, t)$, $p' := \frac{\partial p}{\partial x}(x, t)$, $p'' := \frac{\partial^2 p}{\partial x^2}(x, t)$, $u' := \frac{\partial u}{\partial x}(x, t)$, $v' := \frac{\partial v}{\partial x}(x, t)$ etc. Note that the time t is fixed.

A sketch of calculations to obtain (16) and (17) starting from (36) appears in the following three steps:

Step 1: The three terms in (36) are simplified as:

$$\begin{aligned} \frac{\partial}{\partial x} \left(\frac{p}{1 + u' dt + v' dZ_t} \right) &= p' - f_1 dt - (p'v' + pv'') dZ_t \\ p \frac{\partial}{\partial \tilde{u}} \ln p(x + u dt + v dZ_t) &= p' + f_2 dt + (p''v - \frac{p'^2 v}{p}) dZ_t \\ p \frac{\partial}{\partial \tilde{u}} \ln p_{\text{az}|x}(dZ_t | x + u dt + v dZ_t) &= \frac{p}{\sigma_W^2} [h' dZ_t + (h''v - hh') dt] \end{aligned}$$

where we have used Itô's rules $dZ_t^2 = \sigma_W^2 dt$, $dZ_t dt = 0$ etc., and the functions

$$\begin{aligned} f_1 &= (p'u' + pu'') - \sigma_W^2 (p'v'^2 + 2pv'v''), \\ f_2 &= (p''u - \frac{p'^2 u}{p}) + \sigma_W^2 v^2 \left(\frac{1}{2} p''' - \frac{3p'p''}{2p} + \frac{p'^3}{p^2} \right). \end{aligned}$$

Collecting terms in $O(dZ_t)$ and $O(dt)$, after some simplification, leads to the following ODEs:

$$\mathcal{E}(v) = \frac{1}{\sigma_W^2} h'(x) \quad (37)$$

$$\mathcal{E}(u) = -\frac{1}{\sigma_W^2} h(x)h'(x) + h''(x)v + \sigma_W^2 G(x,t) \quad (38)$$

where $\mathcal{E}(v) = -\frac{\partial}{\partial x} \left(\frac{1}{p(x,t)} \frac{\partial}{\partial x} \{p(x,t)v(x,t)\} \right)$, and $G = -2v'v'' - (v')^2(\ln p)' + \frac{1}{2}v^2(\ln p)'''$.

Step 2. Suppose (u, v) are admissible solutions of the E-L BVP (37)-(38). Then it is claimed that

$$-(pv)' = \frac{h - \hat{h}}{\sigma_W^2} p \quad (39)$$

$$-(pu)' = -\frac{(h - \hat{h})\hat{h}}{\sigma_W^2} p - \frac{1}{2}\sigma_W^2 (pv^2)'' \quad (40)$$

Recall that admissible here means

$$\lim_{x \rightarrow \pm\infty} p(x,t)u(x,t) = 0, \quad \lim_{x \rightarrow \pm\infty} p(x,t)v(x,t) = 0. \quad (41)$$

To show (39), integrate (37) once to obtain

$$-(pv)' = \frac{1}{\sigma_W^2} hp + Cp,$$

where the constant of integration $C = -\frac{\hat{h}}{\sigma_W^2}$ is obtained by integrating once again between $-\infty$ to ∞ and using the boundary conditions for v (41). This gives (39).

To show (40), we denote its right hand side as \mathcal{R} and claim

$$\left(\frac{\mathcal{R}}{p} \right)' = -\frac{hh'}{\sigma_W^2} + h''v + \sigma_W^2 G. \quad (42)$$

The equation (40) then follows by using the ODE (38) together with the boundary conditions for u (41). The verification of the claim involves a straightforward calculation, where we use (37) to obtain expressions for h' and v'' . The details of this calculation are omitted on account of space.

Step 3. The E-L equation for v is given by (37) which is the same as (16). The proof of (17) involves a short calculation starting from (40), which is simplified to the form (17) by using (39).

Proof of Proposition 3. Consider the ODE (16). It is a linear ODE whose unique solution is given by

$$v(x,t) = \frac{1}{p(x,t)} \left(C_1 + C_2 \int_{-\infty}^x p(y,t) dy - \frac{1}{\sigma_W^2} \int_{-\infty}^x h(y)p(y,t) dy \right), \quad (43)$$

where the constant of integrations $C_1 = 0$ and $C_2 = \frac{\hat{h}}{\sigma_W^2}$ because of the boundary conditions for v . Part 2 is an easy consequence of the minimum principle for elliptic PDEs [7].

C. Consistency with the nonlinear filter

Proof of Theorem 1 It is only necessary to show that with this choice of $\{u, v\}$, we have $dp(x,t) = dp^*(x,t)$, for all x and t , in the sense that they are defined by identical stochastic differential equations. Recall dp^* is defined according to

the K-S equation (12), and dp according to the forward equation (15).

If v solves the E-L BVP (16) then using (43),

$$\frac{\partial}{\partial x}(pv) = -\frac{1}{\sigma_W^2}(h - \hat{h})p. \quad (44)$$

On multiplying both sides of (17) by $-p$, we have

$$\begin{aligned} -up &= \frac{1}{2}(h - \hat{h})pv - \frac{1}{2}\sigma_W^2(pv) \frac{\partial v}{\partial x} + \hat{h}pv \\ &= -\frac{1}{2}\sigma_W^2 \frac{\partial(pv)}{\partial x} v - \frac{1}{2}\sigma_W^2(pv) \frac{\partial v}{\partial x} + \hat{h}pv \\ &= -\frac{1}{2}\sigma_W^2 \frac{\partial}{\partial x}(pv^2) + \hat{h}pv, \end{aligned}$$

where we have used (44) to obtain the second equality. Differentiating once with respect to x and using (44) once again,

$$-\frac{\partial}{\partial x}(up) + \frac{1}{2}\sigma_W^2 \frac{\partial^2}{\partial x^2}(pv^2) = -\frac{\hat{h}}{\sigma_W^2}(h - \hat{h})p. \quad (45)$$

Using (44)-(45) in the forward equation (15), we have

$$dp = \mathcal{L}^\dagger p + \frac{1}{\sigma_W^2}(h - \hat{h})(dZ_t - \hat{h}dt)p$$

This is precisely the SDE (12), as desired.

REFERENCES

- [1] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*. Springer, Cambridge, Mass, 2010.
- [2] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Atena Scientific, Cambridge, Mass, 1996.
- [3] A. Budhiraja, L. Chen, and C. Lee. A survey of numerical methods for nonlinear filtering problems. *Physica D: Nonlinear Phenomena*, 230(1-2):27 – 36, 2007.
- [4] J. L. Doob. *Stochastic Processes*. John Wiley & Sons, New York, 1953.
- [5] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte-Carlo Methods in Practice*. Springer-Verlag, April 2001.
- [6] K. Doya, S. Ishii, A. Pouget, and R. P. N. Rao. *Bayesian Brain*. Comput. Neurosci. MIT Press, Cambridge, MA, 2007.
- [7] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 1998.
- [8] M. Huang, P. E. Caines, and R. P. Malhame. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε -Nash equilibria. *IEEE Trans. Automat. Control*, 52(9):1560–1571, 2007.
- [9] D. H. Jacobson and D. Q. Mayne. *Differential dynamic programming*. American Elsevier Pub. Co., New York, NY, 1970.
- [10] Gopinath Kallianpur. *Stochastic filtering theory*, volume 13 of *Applications of Mathematics*. Springer-Verlag, New York, 1980.
- [11] Harold J. Kushner. *Probability methods for approximations in stochastic control and for elliptic equations*. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1977. Mathematics in Science and Engineering, Vol. 129.
- [12] R. Van Handel. Uniform time average consistency of Monte Carlo particle filters. *Stoch. Proc. Applns.*, 119(11):3835–3861, 2009.
- [13] G. Y. Weintraub, L. Benkard, and B. Van Roy. Oblivious equilibrium: A mean field approximation for large-scale dynamic games. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2006.
- [14] T. Yang, P. G. Mehta, and S. P. Meyn. A mean-field control-oriented approach to particle filtering. In *Proc. of American Control Conference*, June 2011.
- [15] H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag. Synchronization of coupled oscillators is a game. In *American Control Conference*, pages 1783–1790, June 2010.