Kalman Filter and its Modern Extensions for the Continuous-time Nonlinear Filtering Problem

Amirhossein Taghvaei, Jana de Wiljes, Prashant G. Mehta, Sebastian Reich

Abstract

This paper is concerned with the filtering problem in continuous-time. Three algorithmic solution approaches for this problem are reviewed: (i) the classical Kalman-Bucy filter which provides an exact solution for the linear Gaussian problem, (ii) the ensemble Kalman-Bucy filter (EnKBF) which is an approximate filter and represents an extension of the Kalman-Bucy filter to nonlinear problems, and (iii) the feedback particle filter (FPF) which represents an extension of the EnKBF and furthermore provides for a consistent solution in the general nonlinear, non-Gaussian case. The common feature of the three algorithms is the gain times error formula to implement the update step (to account for conditioning due to observations) in the filter. In contrast to the commonly used sequential Monte Carlo methods, the EnKBF and FPF avoid the resampling of the particles in the importance sampling update step. Moreover, the gain times innovation feedback structure provides for error correction potentially leading to smaller simulation variance and improved stability properties. The paper also describes numerical algorithms for gain function approximation in the FPF as well as their relationship to optimal transport and coupling of measures. This paper is concerned with the filtering problem in continuous-time. Three algorithmic solution approaches for this problem are reviewed: (i) the classical Kalman-Bucy filter which provides an exact solution for the linear Gaussian problem, (ii) the ensemble Kalman-Bucy filter (EnKBF) which is an approximate filter and represents an extension of the Kalman-Bucy filter to nonlinear problems, and (iii) the feedback particle filter (FPF) which represents an extension of the EnKBF and furthermore provides for a consistent solution in the general nonlinear, non-Gaussian case. The common feature of the

A. Taghvaei and P. G. Mehta are with the Coordinated Science Laboratory and the Department of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign (UIUC) taghvae2@illinois.edu; mehtapg@illinois.edu

J. de Wiljes is with Institut für Mathematik at Universität Potsdam wiljes@uni-potsdam.de

S. Reich is with Institut für Mathematik at Universität Potsdam and Department of Mathematics and Statistics at University of Reading sereich@uni-potsdam.de

Financial support from the NSF CMMI grants 1334987 and 1462773 is gratefully acknowledged.

three algorithms is the gain times error formula to implement the update step (to account for conditioning due to observations) in the filter. In contrast to the commonly used sequential Monte Carlo methods, the EnKBF and FPF avoid the resampling of the particles in the importance sampling update step. Moreover, the gain times innovation feedback structure provides for error correction potentially leading to smaller simulation variance and improved stability properties. The paper also describes numerical algorithms for gain function approximation in the FPF as well as their relationship to optimal transport and coupling of measures. T

I. Introduction

Since the pioneering work of Kalman from the 1960s, sequential state estimation has been extended to application areas far beyond its original aims such as numerical weather prediction [1] and oil reservoir exploration (history matching) [2]. These developments have been made possible by clever combination of Monte Carlo techniques with Kalman-like techniques for assimilating observations into the underlying dynamical models. The most prominent of these algorithms are the ensemble Kalman filter (EnKF), the randomized maximum likelihood (RML) method and the unscented Kalman filter (UKF) invented independently by several research groups [3]–[6] in the 1990s. The EnKF in particular can be viewed as a cleverly designed random dynamical system of interacting particles which is able to track a reference solution with a relatively small number of particles. The EnKF naturally extends to nonlinear dynamical systems and has become very popular in recent years with applications to, for example, atmosphere-ocean dynamics and oil reservoir exploration. This interacting particle perspective has also led to many new filter algorithms in recent years which go beyond the inherent Gaussian approximation of an EnKF during the data assimilation step [7].

In this paper, we review the interacting particle perspective in the context of the continuous-time filtering problems and demonstrate its close relation to Kalman's and Bucy's original feedback control structure of the data assimilation step. More specifically, we highlight the feedback control structure of three classes of algorithms for approximating the posterior distribution: (i) the classical Kalman-Bucy filter which provides an exact solution for the linear Gaussian problem, (ii) the ensemble Kalman-Bucy filter (EnKBF) which is an approximate filter and represents an extension of the Kalman-Bucy filter to nonlinear problems, and (iii) the feedback particle filter (FPF) which represents an extension of the EnKBF and furthermore provides for a consistent solution in the general nonlinear, non-Gaussian case.

A closely related goal is to provide comparison between these algorithms. The common feature of the three algorithms is the gain times error formula to implement the update step (to account for conditioning due to observations) in the filter. It is shown that both the EnKBF and the FPF algorithm are consistent with the Kalman-Bucy filter in the linear Gaussian setting. The difference is that while the Kalman-Bucy filter is an exact algorithm, the two particle-based algorithms are approximate with error decreasing to zero as the number of particles increases to infinity.

In the class of interacting particle algorithms discussed, the FPF represents the most general solution to the nonlinear non-Gaussian filtering problem. The challenge with implementing the FPF arises due to the gain function computation. The gain function equals the Kalman gain in the linear Gaussian setting and must be numerically approximated in the general setting. One particular closed-form approximation of the gain function is the constant gain approximation. In this case, the FPF is shown to reduce to the EnKBF algorithm.

The outline of the remainder of this paper is as follows: The continuous-time filtering problem and the classic Kalman-Bucy filter are summarized in Sections II and III, respectively. The Kalman-Bucy filter is then put into the context of interacting particle systems in the form of the EnKBF in Section IV. Section V together with Appendices A and A provides a consistent definition of the FPF and a discussion of alternative approximation techniques which lead to consistent approximations to the filtering problem as the number of particles, N, goes to infinity. It is shown that the EnKBF can be viewed as an approximation to the FPF. Four algorithmic approaches to gain function approximation are described and their relationship discussed. The paper concludes with discussion of some future research directions in Section VI.

II. Problem Statement

In continuous time, the model for nonlinear filtering problem is described by the nonlinear stochastic differential equations (sdes):

$$\text{Signal:} \qquad \mathrm{d}X_t = a(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}B_t, \quad X_0 \sim p_0^* \tag{1a}$$

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

where $X_t \in \mathbb{R}^d$ is the (hidden) state at time t, the initial condition X_0 is sampled from a given prior density p_0^*, $Z_t \in \mathbb{R}^m$ is the observation or the measurement 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 known 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 j-th coordinate is denoted as h_j (i.e., $h = (h_1, h_2, \ldots, h_m)^{\mathrm{T}}$). By scaling, it is assumed without loss of generality that the covariance matrices associated with $\{B_t\}$, $\{W_t\}$ are identity matrices. Unless otherwise noted, the stochastic differential equations (sde) are expressed in Itô form.

In applications, the continuous time filtering models are often expressed as:

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = a(X_t) + \sigma(X_t)\dot{B}_t \tag{2a}$$

$$Y_t := \frac{\mathrm{d}Z_t}{\mathrm{d}t} = h(X_t) + \dot{W}_t \tag{2b}$$

where $\dot{B}_t$ and $\dot{W}_t$ are mutually independent white noise processes (Gaussian noise) and $Y_t \in \mathbb{R}^m$ is the vector valued observation at time t. The sde-based model is preferred here because of its mathematical rigor. Any sde involving Z_t is converted into an ODE involving Y_t by formally dividing the sde by $\mathrm{d}t$ and replacing $\frac{\mathrm{d}Z_t}{\mathrm{d}t}$ by Y_t (See also Remark 1).

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

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

One example of particular interest is when the mappings $a(x)$ and $h(x)$ are linear, $\sigma(x)$ is a constant matrix that does not depend upon x, and the prior density p_0^* is Gaussian. The associated problem is referred to as the linear Gaussian filtering problem. For this problem, the posterior density is known to be Gaussian. The resulting filter is said to be finite-dimensional because the posterior is completely described by finitely many statistics – conditional mean and variance in the linear Gaussian case.

For the general nonlinear non-Gaussian case, however, the filter is infinite-dimensional because it defines the evolution, in the space of probability measures, of $\{p^*(\cdot, t) : t \ge 0\}$. The particle filter is a simulation-based algorithm to approximate the posterior: The key step is the construction of N interacting stochastic processes $\{X_t^i : 1 \le i \le N\}$: The value $X_t^i \in \mathbb{R}^d$ is the

state for the i-th 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,

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

The first controlled interacting particle representation of the continuous-time filtering problem can be found in [8], [9]. The close connection of such interacting particle formulations to the gain factor and innovation structure of the classic Kalman filter has been made explicit starting with [10] and has led to the FPF formulation considered in this paper.

III. Kalman-Bucy Filter

Consider the linear Gaussian problem: The mappings $a(x) = Ax$ and $h(x) = Hx$ where A and H are $d \times d$ and $m \times d$ matrices; the process noise covariance $\sigma(x) = \sigma$, a constant $d \times d$ matrix; and the prior density is Gaussian, denoted as $\mathcal{N}(\hat{X}_0, \Sigma_0)$.

For this problem, the posterior density is known to be Gaussian, denoted as $\mathcal{N}(\hat{X}_t, \Sigma_t)$, where $\hat{X}_t$ and Σ_t are the conditional mean and variance. Their evolution is described by the finite-dimensional Kalman-Bucy filter:

$$\mathrm{d}\hat{X}_t = A\hat{X}_t\,\mathrm{d}t + \mathsf{K}_t\Big(\mathrm{d}Z_t - H_t\hat{X}_t\,\mathrm{d}t\Big) \tag{3a}$$

$$\frac{\mathrm{d}\Sigma_t}{\mathrm{d}t} = A\Sigma_t + \Sigma_t A^{\mathrm{T}} + \sigma\sigma^{\mathrm{T}} - \Sigma_t H^{\mathrm{T}} H\Sigma_t \tag{3b}$$

where $\mathsf{K}_t := \Sigma_t H_t^{\mathrm{T}}$ is referred to as the Kalman gain and the filter is initialized with the initial conditions $\hat{X}_0$ and Σ_0 of the prior density[1].

The evolution equation for the mean is a sde because of the presence of stochastic forcing term Z_t on the right-hand side. The evolution equation for the variance Σ_t is an ode that does not depend upon the observation process.

The Kalman filter is one of the most widely used algorithm in engineering. Although the filter describes the posterior *only* in linear Gaussian settings, it is often used as an approximate algorithm even in more general settings, e.g., by defining the matrices A and H according to the

[1] A^{T} denotes the transpose of the matrix A; the hat notation is used throughout the paper to denote a mean (e.g., $\hat{X}_t = \mathsf{E}[X_t|\mathscr{Z}_t]$).

Jacobians of the mappings a and h:

$$A := \frac{\partial a}{\partial x}(\hat{X}_t), \quad H := \frac{\partial h}{\partial x}(\hat{X}_t)$$

The resulting algorithm is referred to as the extended Kalman filter:

$$\mathrm{d}\hat{X}_t = a(\hat{X}_t)\,\mathrm{d}t + \mathsf{K}_t\Big(\mathrm{d}Z_t - h(\hat{X}_t)\,\mathrm{d}t\Big)$$

$$\frac{\mathrm{d}\Sigma_t}{\mathrm{d}t} = A\Sigma_t + \Sigma_t A^{\mathrm{T}} + \sigma(\hat{X}_t)\sigma^{\mathrm{T}}(\hat{X}_t) - \Sigma_t H^{\mathrm{T}} H \Sigma_t$$

where $\mathsf{K}_t = \Sigma_t H^{\mathrm{T}}$ is used as the formula for the gain.

The Kalman filter and its extensions are recursive algorithms that process measurements in a sequential (online) fashion. At each time t, the filter computes an error $\mathrm{d}Z_t - H\hat{X}_t\,\mathrm{d}t$ (called the *innovation error*) which reflects the new information contained in the most recent measurement. The filter state $\hat{X}_t$ is corrected at each time step via a $(\text{gain} \times \text{error})$ update formula.

The error correction feedback structure (see Fig. 1) is important on account of robustness. A filter is based on an idealized model of an underlying stochastic dynamic process. The self-correcting property of the feedback provides robustness, allowing one to tolerate a degree of uncertainty inherent in any model.

The simple intuitive nature of the update formula is invaluable in design, testing and operation of the filter. For example, the Kalman gain is proportional to H which scales with the signal-to-noise ratio of the measurement model. In practice, the gain may be 'tuned' to optimize the filter performance. To minimize online computations, an offline solution of the algebraic Ricatti equation (obtained after equating the right-hand side of the variance ode (3b) to zero) may be used to obtain a constant value for the gain.

The basic Kalman filter has also been extended to handle filtering problems involving additional uncertainties in the signal model and the observation model. The resulting (approximate) algorithms are referred to as the interacting multiple model (IMM) filter and the probabilistic data association (PDA) filter, respectively. In the PDA filter, the gain changes based on an estimate of the instantaneous uncertainty in the measurements. In the IMM filter, multiple Kalman filters are run in parallel and their outputs combined to obtain an estimate.

One explanation of the feedback control structure of the Kalman filter is based on duality between estimation and control [11]. Although limited to linear Gaussian problems, these considerations also help explain the differential Ricatti equation structure for the variance ode (3b).

Although widely used, the extended Kalman filter can suffer from stability issues because of the sometimes very crude approximation of the nonlinear model. The observed divergence arises on account of two inter-related reasons: (i) Even with Gaussian process and measurement noise, the nonlinearity of the mappings $a(\cdot)$, $\sigma(\cdot)$ and $h(\cdot)$ can lead to non-Gaussian forms of the posterior density p^*; and (ii) the Jacobians A and H used in propagating the covariance can lead to large errors in approximation of the gain particularly if the Hessian of these mappings is large. These issues have necessitated development of particle based algorithms described in the following sections.

IV. Ensemble Kalman-Bucy Filter

For pedagogical reasons, the ensemble Kalman-Bucy filter (EnKBF) is best described for the linear Gaussian problem – also the approach taken in this section. The extension to the nonlinear non-Gaussian problem is then immediate, similar to the extension from the Kalman filter to the extended Kalman filter.

Even in linear Gaussian settings, a particle filter may be a computationally efficient option for problems with very large state dimension d (e.g., weather models in meteorology). For large d, the computational bottleneck in simulating a Kalman filter arises due to propagation of the covariance matrix according to the differential Riccati equation (3b). This computation scales as $O(d^2)$ in memory. In an EnKBF implementation, one replaces the exact propagation of the covariance matrix by an empirical approximation with N particles

$$\Sigma_t^{(N)} = \frac{1}{N-1}\sum_{i=1}^{N}(X_t^i - \hat{X}_t^{(N)})(X_t^i - \hat{X}_t^{(N)})^{\mathrm{T}} \tag{4}$$

This computation scales as $O(Nd)$. The same reduction in computational cost can be achieved by a reduced rank Kalman filter. However, the connection to empirical measures (3) is crucial to the application of the EnKBF to nonlinear dynamical systems.

The EnKF algorithm was first developed in a discrete-time setting [4]. Since then various formulations of the EnKF have been proposed [7], [12]. Below we state two time-continuous formulations of the EnKBF.

A. Stochastic EnKBF

The conceptual idea of the stochastic EnKBF algorithm is to introduce a zero mean perturbation (noise term) in the innovation error to achieve consistency for the variance update. In the

time-continuous stochastic EnKBF algorithm, the particles evolve according to

$$\mathrm{d}X_t^i = AX_t^i\,\mathrm{d}t + \sigma\,\mathrm{d}B_t^i + \Sigma_t^{(N)}H^{\mathrm{T}}\Big(\mathrm{d}Z_t^i - HX_t^i\,\mathrm{d}t + \mathrm{d}V_t^i\Big) \tag{5}$$

for $i = 1,\ldots,N$, where $X_t^i \in \mathbb{R}^d$ is the state of the i^{th} particle at time t, the initial condition $X_0^i \sim p_0^*$, B_t^i is a standard Wiener process, and V_t^i is a standard Wiener process assumed to be independent of X_0^i, B_t^i, X_t, Z_t [12].. The variance $\Sigma_t^{(N)}$ is obtained empirically using (4). Note that the N particles only interact through the common covariance matrix $\Sigma_t^{(N)}$.

The idea of introducing a noise process first appeared for the discrete-time EnKF. The derivation of the continuous-time stochastic EnKBF can be found in [12] or [13]. It is based on a limiting argument whereby the discrete-time update step is formally viewed as an Euler-Maruyama discretization of a stochastic SDE. For the linear Gaussian problem, the stochastic EnKBF algorithm is consistent in the limit as $N \to \infty$. This means that the conditional distribution of X_t^i is Gaussian whose mean and variance evolve according the Kalman filter, equations (3a) and (3b), respectively.

The update formula used in (5) is not unique. A deterministic analogue is described next.

B. Deterministic EnKBF

A deterministic alternative update variant of the EnKBF first proposed in [14] is given by:

$$\mathrm{d}X_t^i = AX_t^i\,\mathrm{d}t + \sigma\,\mathrm{d}B_t^i + \Sigma_t^{(N)}H^{\mathrm{T}}\left(\mathrm{d}Z_t - \frac{HX_t^i + H\hat{X}_t^{(N)}}{2}\mathrm{d}t\right) \tag{6}$$

for $i = 1,\ldots,N$. A proof of the consistency of this deterministic variant of the EnKBF for linear systems can be found in [15]. There are close parallels between the deterministic EnKBF and the FPF which is explored further in Section V. In the deterministic formulation of the EnKBF the interaction between the N particles arises through the covariance matrix $\Sigma_t^{(N)}$ and the mean $\hat{X}_t^{(N)}$.

Although the stochastic and the deterministic EnKBF algorithms are consistent for the linear Gaussian problem, they can be easily extended to the nonlinear non-Gaussian settings. However, the resulting algorithm will in general not be consistent.

C. Well-posedness and accuracy of the EnKBF

Recent research on EnKF has focussed on the long term behavior and accuracy of filters applicable for nonlinear data assimilation [15]–[18]. In particular the mathematical justification

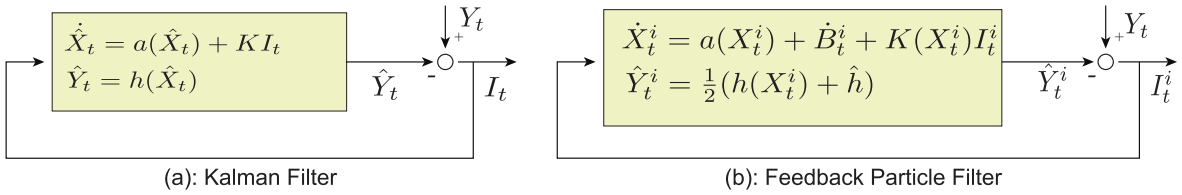

(a): Kalman Filter (b): Feedback Particle Filter

Fig. 1: Innovation error-based feedback structure for the (a) Kalman filter and (b) nonlinear feedback particle filter.

for the feasibility of the EnKF and its continuous counterpart in the small ensemble limit are of interest. These studies of the accuracy for a finite ensemble are of exceptional importance due to the fact that a large number of ensemble members is not an option from a computational point of view in many applicational areas. The authors of [17] were able to show mean-squared asymptotic accuracy in the large-time limit for a variant of a stochastic EnKBF algorithm. Well-posedness of the discrete and continuous formulation of the EnKF is also proven. Similar results concerning the well-posedness and accuracy for the deterministic variant (6) of the EnKBF are derived in [15].

Note that a fully observed system is assumed for the discussed accuracy and well-posedness results. Yet an investigation of the well-posedness for partially observed systems is particularly relevant as the analysis update in such cases can cause a divergence of the approximated signal in the sense that the signal is lost or the values reach machine infinity [19].

V. Feedback Particle Filter

The FPF is a controlled sde:

$$\mathrm{d}X^i_t = a(X^i_t)\,\mathrm{d}t + \sigma(X^i_t)\,\mathrm{d}B^i_t + \underbrace{\mathsf{K}_t(X^i_t)\circ(\mathrm{d}Z_t - \frac{h(X^i_t)+\hat{h}_t}{2}\,\mathrm{d}t)}_{\text{update}}, \quad X^i_0 \sim p^*_0 \tag{7}$$

for $i = 1,\ldots,N$, where (similar to EnKBF) $X^i_t \in \mathbb{R}^d$ is the state of the i^{th} particle at time t, the initial condition $X^i_0 \sim p^*_0$, B^i_t is a standard Wiener process, and $\hat{h}_t := \mathsf{E}[h(X^i_t)|\mathscr{Z}_t]$. Both B^i_t and X^i_0 are mutually independent and also independent of X_t, Z_t. The $\circ$ in the update term indicates that the sde is expressed in its Stratonovich form.

The gain function K_t is vector-valued (with dimension $d \times m$) and it needs to be obtained for each fixed time t. The gain function is defined as a solution of a pde introduced in the following subsection. For the linear Gaussian problem, K_t is the Kalman gain. For the general nonlinear non-Gaussian, the gain function needs to be numerically approximated. Algorithms for this are also summarized in the following subsection.

Remark 1: Given that the Stratonovich form provides a mathematical interpretation of the (formal) ode model [20, see Section 3.3 of the sde textbook by Øksendal], we also obtain the (formal) ode model of the filter. Denoting $Y_t \doteq \frac{\mathrm{d}Z_t}{\mathrm{d}t}$ and white noise process $\dot{B}^i_t \doteq \frac{\mathrm{d}B^i_t}{\mathrm{d}t}$, the ODE model of the filter is given by,

$$\frac{\mathrm{d}X^i_t}{\mathrm{d}t} = a(X^i_t) + \sigma(X^i_t)\dot{B}^i_t + \mathsf{K}(X^i,t)\cdot\left(Y_t - \frac{1}{2}(h(X^i_t)+\hat{h})\right) \tag{8}$$

for $i = 1,\ldots,N$. The feedback particle filter thus provides a generalization of the Kalman filter to nonlinear systems, where the innovation error-based feedback structure of the control is preserved (see Fig. 1). For the linear Gaussian case, the gain function is the Kalman gain. For the nonlinear case, the Kalman gain is replaced by a nonlinear function of the state (See Fig. 3).

Remark 2: It is shown in Appendix A that, under the condition that the gain function can be computed exactly, FPF is an exact algorithm. That is, if the initial condition X^i_0 is sampled from the prior p^*_0 then

$$\mathsf{P}[X_t \in A \mid \mathscr{Z}_t] = \mathsf{P}[X^i_t \in A \mid \mathscr{Z}_t], \quad \forall\, A \subset \mathbb{R}^d,\ t > 0$$

In a numerical implementation, a finite number, N, of particles is simulated and $\mathsf{P}[X^i_t \in A \mid \mathscr{Z}_t] \approx \frac{1}{N}\sum_{i=1}^N 1[X^i_t \in A]$ by the Law of Large Numbers (LLN).

The considerations in the Appendix are described in a more general setting, e.g., applicable to stochastic processes X_t and X^i_t evolving on manifolds. This also explains why the update formula has a Stratonovich form. For sdes on a manifold, it is well known that the Stratonovich form is invariant to coordinate transformations (i.e., intrinsic) while the Ito form is not. A more in-depth discussion of the FPF for Lie groups appears in [21].

A. Gain function

For pedagogical reasons primarily to do with notational convenience, the gain function is defined here for the case of scalar-valued observation[2]. In this case, the gain function K_t is

[2]The extension to multi-valued observation is straightforward and appears in [22].

defined in terms of the solution of the weighted Poisson equation:

$$\begin{aligned} -\nabla\cdot(\rho(x)\nabla\phi(x)) &= (h(x)-\hat{h})\rho(x), \quad x\in\mathbb{R}^d \\ \int \phi(x)\rho(x)\,\mathrm{d}x &= 0 \end{aligned} \tag{9}$$

where $\hat{h} := \int h(x)\rho(x)\,\mathrm{d}x$, ∇ and $\nabla\cdot$ denote the gradient and the divergence operators, respectively, and at time t, $\rho(x) = p(x,t)$ denotes the density of X_t^i[3]. In terms of the solution $\phi(x)$ of (9), the gain function at time t is given by

$$\mathsf{K}_t(x) = \nabla\phi(x) \tag{10}$$

Remark 3: The gain function $\mathsf{K}_t(x)$ is not uniquely defined through the filtering problem. Formula (10) represents one choice of the gain function. More generally, it is sufficient to require that $\mathsf{K}_t = \mathsf{K}$ is a solution of

$$-\nabla\cdot(\rho(x)\mathsf{K}(x)) = (h(x)-\hat{h})\rho(x), \quad x\in\mathbb{R}^d \tag{11}$$

with $\rho(x) = p(x,t)$ at time t.

One justification for choosing the gradient form solution, as in (10), is its L^2 optimality. The general solution of (11) is given by

$$\mathsf{K} = \nabla\phi + v$$

where ϕ is the solution of (9) and v solves $\nabla\cdot(\rho v) = 0$. It is easy to see that

$$\mathsf{E}[|\mathsf{K}|^2] = \mathsf{E}[|\nabla\phi|^2] + \mathsf{E}[|v|^2]$$

Therefore, $\mathsf{K} = \nabla\phi$ is the minimum L^2-norm solution of (11). By interpreting the L^2 norm as the kinetic energy, the gain function $\mathsf{K}_t = \nabla\phi$, defined through (9), is seen to be optimal in the sense of optimal transportation [23], [24].

An alternative solution of (11) is provided through the definition

$$\mathsf{K}_t(x) = \frac{1}{\rho(x)}\nabla\tilde{\phi}(x)$$

[3]Although this paper is limited to $\mathbb{R}^d$, the proposed algorithm is applicable to nonlinear filtering problems on differential manifolds, e.g., matrix Lie groups (For an intrinsic form of the Poisson equation, see [21]). For domains with boundary, the pde is accompanied by a Neumann boundary condition:

$$\nabla\phi(x)\cdot n(x) = 0$$

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

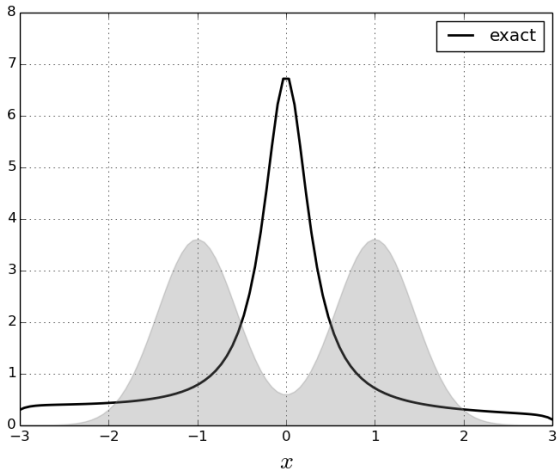

Fig. 2: The exact solution to the Poisson equation using the formula (12). The density ρ is the sum of two Gaussians $N(-1,\sigma^2)$ and $N(+1,\sigma^2)$, and $h(x)=x$. The density is depicted as the shaded curve in the background.

which leads to a standard Poisson equation in the unknown potential $\tilde{\phi}$ for which the fundamental solution is explicitly known. This fact is exploited in the interacting particle filter representations of [8], [9].

There are two special cases of (9) – summarized as part of the following two examples – where the exact solution can be found.

Example 1: In the scalar case (where $d=1$), the Poisson equation is:

$$-\frac{1}{\rho(x)}\frac{\mathrm{d}}{\mathrm{d}x}(\rho(x)\frac{\mathrm{d}\phi}{\mathrm{d}x}(x)) = h-\hat{h}$$

Integrating once yields the solution explicitly,

$$\mathsf{K}(x) = \frac{\mathrm{d}\phi}{\mathrm{d}x}(x) = -\frac{1}{\rho(x)}\int_{-\infty}^{x}\rho(z)(h(z)-\hat{h})\,\mathrm{d}z \tag{12}$$

For the particular choice of ρ as the sum of two Gaussians $\mathscr{N}(-1,\sigma^2)$ and $\mathscr{N}(+1,\sigma^2)$ with $\sigma^2=0.2$ and $h(x)=x$, the solution obtained using (12) is depicted in Fig. 2.

Example 2: Suppose the density ρ is a Gaussian $\mathscr{N}(\mu,\Sigma)$. The observation function $h(x)=Hx$, where $H\in\mathbb{R}^{1\times d}$. Then, $\phi = x^{\mathrm{T}}\Sigma H^{\mathrm{T}}$ and the gain function $\mathsf{K}=\Sigma H^{\mathrm{T}}$ is the Kalman gain.

In the general non-Gaussian case, the solution is not known in an explicit form and must be numerically approximated. Note that even in the two exact cases, one would need to numerically approximate the solution because the density ρ is not available in an explicit form.

The problem statement for numerical approximation is as follows:

Problem statement: Given N samples $\{X^1,\dots,X^i,\dots,X^N\}$ drawn i.i.d. from ρ, approximate the vector-valued gain function$\{\mathsf{K}^1,\dots,\mathsf{K}^i,\dots,\mathsf{K}^N\}$, where $\mathsf{K}^i := \mathsf{K}(X^i) = \nabla\phi(X^i)$. The density ρ is not explicitly known.

Four numerical algorithms for approximation of the gain function appear in the following four subsections[4].

B. Constant Gain Approximation

The constant gain approximation is the best – in the least-square sense – constant approximation of the gain function (see Figure 3). Mathematically, it is obtained by considering the following least-square optimization problem:

$$\kappa^* = \arg\min_{\kappa\in\mathbb{R}^d} \mathsf{E}\left[|\mathsf{K}-\kappa|^2\right]$$

By using a standard sum of the squares argument, $\kappa^* = \mathsf{E}[\mathsf{K}]$. By multiplying both sides of the pde (9) by x and integrating by parts, the expected value is computed explicitly as

$$\kappa^* = \mathsf{E}[\mathsf{K}] = \int_{\mathbb{R}^d} (h(x)-\hat{h})\, x\, \rho(x)\, \mathrm{d}x$$

The integral is evaluated empirically to obtain the following approximate formula for the gain:

$$\mathsf{K}^i \equiv \frac{1}{N}\sum_{j=1}^{N} (h(X^j)-\hat{h}^{(N)})\, X^j \tag{13}$$

where $\hat{h}^{(N)} = N^{-1}\sum_{j=1}^N h(X^j)$. The formula (13) is referred to as the *constant gain approximation* of the gain function; cf., [22]. It is a popular choice in applications [26]–[28].

Example 3: Consider the linear case where $h(x) = Hx$. The constant gain approximation formula equals

$$\mathsf{K}^i = \frac{1}{N}\sum_{j=1}^{N} (HX^j - H\hat{X}^{(N)})\, X^i = \Sigma^{(N)}H$$

where $\hat{X}^{(N)} := \frac{1}{N}\sum_{i=1}^N X^i$ and $\Sigma^N = \frac{1}{N}\sum_{i=1}^N (X^i-\hat{X}^{(N)})(X^i-\hat{X}^{(N)})^{\mathrm{T}}$ are the empirical mean and the empirical variance, respectively. That is, for the linear Gaussian case, the FPF algorithm with the constant gain approximation gives the deterministic EnKBF algorithm (6).

[4]These algorithms are based on the existence-uniqueness theory for solution ϕ of the Poisson equation pde (9), as described in [25].

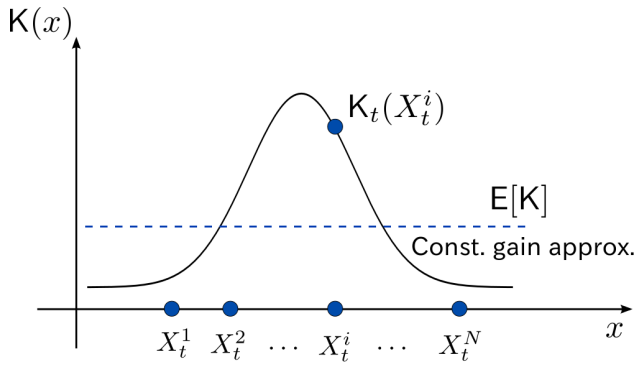

Fig. 3: Constant gain approximation in the feedback particle filter

Algorithm 1 Constant gain approximation

Input: $\{X^i\}_{i=1}^N$, $\{h(X^i)\}_{i=1}^N$,

Output: $\{\mathsf{K}^i\}_{i=1}^N$,

1: Calculate $\hat{h}^{(N)} = \frac{1}{N}\sum_{i=1}^N h(X^i)$

2: Calculate $\mathsf{K}_c = \frac{1}{N}\sum_{i=1}^N (h(X^i) - \hat{h}^{(N)})\, X^i$

3: $K^i = \mathsf{K}_c$ for $i = 1,\ldots,N$

C. Galerkin Approximation

The Galerkin approximation is a generalization of the constant gain approximation where the gain function $\mathsf{K} = \nabla\phi$ is now approximated in a finite-dimensional subspace $S := \text{span}\{\psi_1,\ldots,\psi_M\}$[5]. Mathematically, the Galerkin solution $\nabla\phi^{(M)}$ is defined as the optimal least-square approximation of $\nabla\phi$ in S, i.e,

$$\phi^{(M)} = \arg\min_{\psi\in S} \ \mathsf{E}[|\nabla\phi - \nabla\psi|^2]$$

The least-square solution is easily obtained by applying the projection theorem which gives

$$\mathsf{E}[\nabla\phi^{(M)}\cdot\nabla\psi] = \mathsf{E}[(h-\hat{h})\,\psi], \quad \forall\ \psi\in S \tag{14}$$

[5] S is a finite-dimensional subspace in the Sobolev space $H_0^1(\mathbb{R}^d;\rho)$ – defined as the space of functions f that are square-integrable with respect to density ρ and whose (weak) derivatives are also square-integrable with respect to density ρ. H_0^1 is the appropriate space for the solution ϕ of the Poisson equation (9).

By denoting $(\psi_1(x), \psi_2(x), \ldots, \psi_M(x)) =: \psi(x)$ and expressing $\phi^{(M)}(x) = c \cdot \psi(x)$, the finite dimensional system (14) is expressed as a linear matrix equation

$$Ac = b \tag{15}$$

where A is a $M \times M$ matrix and b is a $M \times 1$ vector whose entries are given by the respective formulae:

$$[A]_{lk} = \mathsf{E}[\nabla\psi_l \cdot \nabla\psi_k]$$

$$[b]_l = \mathsf{E}[(h - \hat{h})\,\psi_l]$$

Example 4: Two types of approximations follow from consideration of two types of basis functions:

(i) The constant gain approximation is obtained by taking basis functions as $\psi_l(x) = x_l$ for $l = 1, \ldots, d$. With this choice, A is the identity matrix and the Galerkin gain function is a constant vector:

$$\mathsf{K}(x) = \int x\,(h(x) - \hat{h})\,\rho(x)\,\mathrm{d}x$$

It's empirical approximation is

$$\mathsf{K}^i \equiv \frac{1}{N}\sum_{j=1}^{N} (h(X^j) - \hat{h}^{(N)})\,X^j$$

(ii) With a single basis function $\psi(x) = h(x)$, the Galerkin solution is

$$\mathsf{K}(x) = \frac{\int (h(x) - \hat{h})^2 \rho(x)\,\mathrm{d}x}{\int |\nabla h(x)|^2 \rho(x)\,\mathrm{d}x}\,\nabla h(x)$$

It's empirical approximation is obtained as

$$\mathsf{K}^i = \frac{\sum_{j=1}^{N} (h(X_t^i) - \hat{h}^{(N)})^2}{\sum_{j=1}^{N} |\nabla h(X_t^i)|^2}\,\nabla h(X_t^i)$$

In practice, the matrix A and the vector b are approximated empirically, and the equation (15) solved to obtain the empirical approximation of c, denoted as $c^{(N)}$ (see Table 2 for the Galerkin algorithm). In terms of this empirical approximation, the gain function is approximated as

$$\mathsf{K}^i = \nabla\phi^{(M,N)}(X^i) := c^{(N)} \cdot \nabla\psi(X^i)$$

The main problem with the Galerkin approximation is how to chose (a small number of) basis functions, particularly when the dimension d is large. This motivates the following non-parametric approaches which do not require a selection of basis functions.

Algorithm 2 Galerkin approximation of the gain function

Input: $\{X^i\}_{i=1}^N$, $\{h(X^i)\}_{i=1}^N$, $\{\psi_1,\dots,\psi_M\}$,

Output: $\{\mathsf{K}^i\}_{i=1}^N$,

1: $\hat{h}^{(N)} = \frac{1}{N}\sum_{i=1}^N h(X^i)$

2: $[A^{(N)}]_{lk} := \frac{1}{N}\sum_{i=1}^N \nabla\psi_l(X^i)\cdot\nabla\psi_k(X^i)$, for $l,k = 1,\dots,M$

3: $[b^{(N)}]_k := \frac{1}{N}\sum_{i=1}^N \psi_k(X^i)(h(X^i)-\hat{h}^{(N)})$, for $k = 1,\dots,M$

4: Calculate $c^{(N)}$ by solving $A^{(N)}c^{(N)} = b^{(N)}$

5: $\mathsf{K}^i = \sum_{k=1}^M c_k^{(N)}\nabla\psi_k(X^i)$

D. Kernel-based Approximation

The linear operator $-\frac{1}{\rho}\nabla\cdot(\rho\nabla) =: \Delta_\rho$ for the pde (9) is a generator of a Markov semigroup, denoted as $e^{\varepsilon\Delta_\rho}$ for $\varepsilon > 0$. It follows that the solution ϕ of (9) is equivalently expressed as, for any fixed $\varepsilon > 0$,

$$\phi = e^{\varepsilon\Delta_\rho}\phi + \int_0^\varepsilon e^{s\Delta_\rho}(h-\hat{h})\,\mathrm{d}s \tag{16}$$

The fixed-point representation is useful because $e^{\varepsilon\Delta_\rho}$ can be approximated by a finite-rank operator

$$T_\varepsilon^{(N)} f(x) := \sum_{i=1}^N k_\varepsilon^{(N)}(x,X^i) f(X^i),$$

where the kernel

$$k_\varepsilon^{(N)}(x,y) = \frac{1}{n_\varepsilon^{(N)}(x)} \frac{g_\varepsilon(x-y)}{\sqrt{\frac{1}{N}\sum_{i=1}^N g_\varepsilon(x-X^i)}\sqrt{\frac{1}{N}\sum_{i=1}^N g_\varepsilon(y-X^i)}}$$

is expressed in terms of the Gaussian kernel $g_\varepsilon(z) := (4\pi\varepsilon)^{-\frac{d}{2}}\exp\left(-\frac{|z|^2}{4\varepsilon}\right)$ for $z\in\mathbb{R}^d$, and $n_\varepsilon^{(N)}(x)$ is a normalization factor chosen such that $T_\varepsilon^{(N)}1 = 1$. It is shown in [29], [30] that $e^{\varepsilon\Delta\rho} \approx T_\varepsilon^{(N)}$ as $\varepsilon\downarrow 0$ and $N\to\infty$.

The approximation of the fixed-point problem (16) is obtained as

$$\phi_\varepsilon^{(N)} = T_\varepsilon^{(N)}\phi_\varepsilon^{(N)} + \varepsilon(h-\hat{h}), \tag{17}$$

where $\int_0^\varepsilon e^{s\Delta_\rho}(h-\hat{h})\,\mathrm{d}s \approx \varepsilon(h-\hat{h})$ for small $\varepsilon > 0$. The method of successive approximation is used to solve the fixed-point equation for $\phi_\varepsilon^{(N)}$. In a recursive simulation, the algorithm is initialized with the solution from the previous time-step.

The gain function is obtained by taking the gradient of the two sides of (17). For this purpose, it is useful to first define a finite-rank operator:

$$\begin{aligned}\nabla T_\varepsilon^{(N)} f(x) &:= \sum_{i=1}^N \nabla k_\varepsilon^{(N)}(x,X^i) f(X^i) \\ &= \frac{1}{2\varepsilon}\left[\sum_{i=1}^N k_\varepsilon^{(N)}(x,X^i) f(X^i)\left(X^i - \sum_{j=1}^N k_\varepsilon^{(N)}(x,X^j)X^j\right)\right]\end{aligned} \tag{18}$$

In terms of this operator, the gain function is approximated as

$$\mathsf{K}^i = \nabla T_\varepsilon^{(N)} \phi_\varepsilon^{(N)}(X^i) + \varepsilon \nabla T_\varepsilon^{(N)}(h - \hat{h}^{(N)})(X^i) \tag{19}$$

where $\phi_\varepsilon^{(N)}$ on the righthand-side is the solution of (17).

For $i, j \in \{1,2,\ldots,N\}$, denote

$$a_{ij} := \frac{1}{2\varepsilon} k_\varepsilon^{(N)}(X^i,X^j)\left(r_j - \sum_{l=1}^N k_\varepsilon^{(N)}(X^i,X^l) r_l\right)$$

where $r_i := \phi_\varepsilon^{(N)}(X^i) + \varepsilon h(X^i) - \varepsilon \hat{h}^{(N)}$. Then, the formula (19) is succinctly expressed as

$$\mathsf{K}^i = \sum_{j=1}^N a_{ij} X^j \tag{20}$$

It is easy to verify that $\sum_{j=1}^N a_{ij} = 0$ and as $\varepsilon \to \infty$, $a_{ij} = N^{-1}(h(X^j) - \hat{h}^{(N)})$. Therefore, as $\varepsilon \to \infty$, K^i equals the constant gain approximation formula (13).

E. Optimal Coupling-based Approximation

Optimal coupling-based approximation is another non-parametric approach to directly compute the gain function K^i from the ensemble $\{X^i\}_{i=1}^N$. It is based upon the ensemble transform for optimally transporting (coupling) measures [7].

The relationship to the gain function approximation is as follows: Define an ε-parametrized family of densities by $\rho_\varepsilon(x) := \frac{\rho(x)e^{\varepsilon h(x)}}{\int \rho(y)e^{\varepsilon h(y)}\,\mathrm{d}y}$ and consider the optimal transport problem

$$\begin{aligned}\text{Objective:} \qquad & \min_{S_\varepsilon} \mathsf{E}\left[|S_\varepsilon(X) - X|^2\right] \\ \text{Constraints:} \qquad & X \sim \rho, \quad S_\varepsilon(X) \sim \rho_\varepsilon\end{aligned} \tag{21}$$

The solution to this problem, denoted as S_ε, is referred to as the optimal transport map. It is formally shown in the Appendix Sec. A that $\frac{\mathrm{d}S_\varepsilon}{\mathrm{d}\varepsilon}\big|_{\varepsilon=0} = \mathsf{K}$, the gain function, solution of the Poisson equation (9).

Algorithm 3 Kernel-based approximation of the gain function

Input: $\{X^i\}_{i=1}^N$, $\{h(X^i)\}_{i=1}^N$, Φ_{prev}, L

Output: $\{\mathsf{K}^i\}_{i=1}^N$

Calculate $g_{ij} := \exp(-|X^i - X^j|^2/4\varepsilon)$ for $i,j = 1$ to N

Calculate $k_{ij} := \frac{g_{ij}}{\sqrt{\sum_l g_{il}}\sqrt{\sum_l g_{jl}}}$ for $i,j = 1$ to N

Calculate $T_{ij} := \frac{k_{ij}}{\sum_l k_{il}}$ for $i,j = 1$ to N

Calculate $\hat{h}^{(N)} = \frac{1}{N}\sum_{i=1}^N h(X^i)$

Initialize $\Phi_i = \Phi_{\text{prev},i}$ for $i = 1$ to N

for $l = 1$ to L **do**

  Calculate $\Phi_i = \sum_{j=1}^N T_{ij}\Phi_j + \varepsilon(h(X^i) - \hat{h}^{(N)})$

  Calculate $\Phi_i = \Phi_i - \frac{1}{N}\sum_{j=1}^N \Phi_j$

end for

Calculate $r_i = \Phi_i + \varepsilon(h(X^i) - \hat{h}^{(N)})$

Calculate $a_{ij} = \frac{1}{2\varepsilon}T_{ij}\left(r_j - \sum_{l=1}^N T_{il}r_l\right)$

Calculate $\mathsf{K}^i = \sum_{j=1}^N a_{ij}X^j$

The ensemble transform is a non-parametric algorithm to approximate the solution of the optimal transportation problems, of the type (21) and more generally, given *only* samples X^i drawn from ρ. For the problem of gain function approximation, the algorithm involves first solving the linear program:

$$\begin{aligned} \text{Objective:} \quad & \min_{\{a_{ij}\}} \sum_{i=1}^N \sum_{j=1}^N a_{ij}|X^i - X^j|^2 \\ \text{Constraints:} \quad & \sum_{j=1}^N a_{ij} = 0, \quad \sum_{i=1}^N a_{ij} = h(X^j) - \hat{h}^{(N)}, \\ & a_{ij} \geq 0 \text{ for } i \neq j \end{aligned} \tag{22}$$

The objective in (22) is motivated through an empirical approximation of the objective in (21) where the coupling constants a_{ij} have the interpretation of transition rates. The two equality constraints arise due to the specification of the two marginals ρ and ρ_ε in (21) where it is noted

that the particles X^i are sampled i.i.d. from ρ. A derivation of the linear program (22) based on the ensemble transform method appears in the Appendix Sec. A.

In terms of a solution, denoted as $\{a_{ij}\}_{i,j=1}^N$, of the linear program (22), the gain function is obtained as

$$\mathsf{K}^i = \sum_{j=1}^N a_{ij} X^j \tag{23}$$

The approximation (23) is structurally similar to the constant gain approximation formula (13) and also the kernel-gain approximation formula (20). In all three cases, the gain K^i at the i^{th} particle is approximated as a linear combination of the particle states $\{X^j\}_{j=1}^N$. Such approximations are computationally attractive whenever $N \ll d$, i.e., when the dimension of state space is high but the dynamics is confined to a low-dimensional subset which, however, is not a priori known.

VI. Conclusion

We have summarized an interacting particle representation of the classic Kalman-Bucy filter and its extension to nonlinear systems and non-Gaussian distributions under the general framework of FPFs. This framework is attractive since it maintains key structural elements of the classic Kalman-Bucy filter, namely, a gain factor and the innovation. In particular, the EnKF has become widely used in data assimilation for atmosphere-ocean dynamics and oil reservoir exploration. Robust extensions of the EnKF to non-Gaussian distributions are urgently needed and FPFs provide a systematic approach for such extensions in the spirit of Kalman's original work. However, interacting particle representations come at a price; they require approximate solutions of an elliptic PDE or a coupling problem, when viewed from a probabilistic perspective. Hence, robust and efficient numerical techniques for FPF-based interacting particle systems and study of their long-time behavior will be a primary focus of future research.

References

[1] E. Kalnay, *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2002.

[2] D. Oliver, A. Reynolds, and N. Liu, *Inverse theory for petroleum reservoir characterization and history matching*. Cambridge: Cambridge University Press, 2008.

[3] P. Kitanidis, "Quasi-linear geostatistical theory for inversion," *Water Resources Research*, vol. 31, pp. 2411–2419, 1995.

[4] G. Burgers, P. van Leeuwen, and G. Evensen, "On the analysis scheme in the ensemble Kalman filter," *Mon. Wea. Rev.*, vol. 126, pp. 1719–1724, 1998.

[5] P. Houtekamer and H. Mitchell, "A sequential ensemble Kalman filter for atmospheric data assimilation," *Mon. Wea. Rev.*, vol. 129, pp. 123–136, 2001.

[6] S. J. Julier and J. K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Signal processing, sensor fusion, and target recognition. Conference No. 6*, vol. 3068, Orlando FL, 1997, pp. 182–193.

[7] S. Reich and C. Cotter, *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge, UK: Cambridge University Press, 2015.

[8] D. Crisan and J. Xiong, "Numerical solutions for a class of SPDEs over bounded domains," *ESAIM: Proc.*, vol. 19, pp. 121–125, 2007.

[9] ——, "Approximate McKean-Vlasov representations for a class of SPDEs," *Stochastics*, vol. 82, no. 1, pp. 53–68, 2010.

[10] T. Yang, P. G. Mehta, and S. P. Meyn, "Feedback particle filter with mean-field coupling," in *Decision and Control and European Control Conference (CDC-ECC), 50th IEEE Conference on*, 2011, pp. 7909–7916.

[11] R. E. Kalman *et al.*, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[12] K. Law, A. Stuart, and K. Zygalakis, *Data Assimilation: A Mathematical Introduction*. New York: Springer-Verlag, 2015.

[13] S. Reich, "A dynamical systems framework for intermittent data assimilation," *BIT Numerical Analysis*, vol. 51, pp. 235–249, 2011.

[14] K. Bergemann and S. Reich, "An ensemble Kalman-Bucy filter for continuous data assimilation," *Meteorolog. Zeitschrift*, vol. 21, pp. 213–219, 2012.

[15] J. de Wiljes, S. Reich, and W. Stannat, "Long-time stability and accuracy of the ensemble Kalman-Bucy filter for fully observed processes and small measurement noise," University Potsdam, Tech. Rep. https://arxiv.org/abs/1612.06065, 2016.

[16] P. Del Moral, A. Kurtzmann, and J. Tugaut, "On the stability and the uniform propagation of chaos of extended ensemble Kalman-Bucy filters," INRIA Bordeaux Research Center, Tech. Rep. arXiv:1606.08256v1, 2016.

[17] K. L. D. Kelly and and A. Stuart, "Well-posedness and accuracy of the ensemble Kalman filter in discrete and continuous time," *Nonlinearity*, vol. 27, no. 10, p. 2579, 2014.

[18] D. Kelly and A. M. Stuart, "Ergodicity and accuracy of optimal particle filters for Bayesian data assimilation," Courant Institute of Mathematical Sciences, Tech. Rep. https://arxiv.org/abs/1611.08761, 2016.

[19] D. Kelly, A. J. Majda, and X. T. Tong, "Concrete ensemble Kalman filters with rigorous catastrophic filter divergence," *PNAS*, 2015.

[20] B. Oksendal, *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

[21] C. Zhang, A. Taghvaei, and P. G. Mehta, "Feedback particle filter on matrix Lie groups," in *American Control Conference (ACC), 2016*. IEEE, 2016, pp. 2723–2728.

[22] T. Yang, R. S. Laugesen, P. G. Mehta, and S. P. Meyn, "Multivariable feedback particle filter," *Automatica*, vol. 71, pp. 10–23, 2016.

[23] C. Villani, *Topics in Optimal Transportation*. American Mathematical Soc., 2003, no. 58.

[24] L. C. Evans, "Partial differential equations and Monge-Kantorovich mass transfer," *Current developments in mathematics*, pp. 65–126, 1997.

[25] R. S. Laugesen, P. G. Mehta, S. P. Meyn, and M. Raginsky, "Poisson's equation in nonlinear filtering," *SIAM J. Control Optim.*, vol. 53, no. 1, pp. 501–525, 2015.

[26] P. M. Stano, A. K. Tilton, and R. Babuška, "Estimation of the soil-dependent time-varying parameters of the hopper sedimentation model: The FPF versus the BPF," *Control Engineering Practice*, vol. 24, pp. 67–78, 2014.

[27] A. K. Tilton, S. Ghiotto, and P. G. Mehta, "A comparative study of nonlinear filtering techniques," in *Information Fusion (FUSION), 2013 16th International Conference on*, 2013, pp. 1827–1834.

[28] K. Berntorp, "Feedback particle filter: Application and evaluation," in *18th Int. Conf. Information Fusion, Washington, DC*, 2015.

[29] R. R. Coifman and S. Lafon, "Diffusion maps," *Appl. Comput. Harmon. A.*, vol. 21, no. 1, pp. 5–30, 2006.

[30] M. Hein, J. Audibert, and U. Luxburg, "Graph Laplacians and their convergence on random neighborhood graphs," *J. Mach. Learn. Res.*, vol. 8, no. Jun, pp. 1325–1368, 2007.

[31] J. Xiong, *An introduction to stochastic filtering theory*, ser. Oxford Graduate Texts in Mathematics. Oxford University Press, 2008, vol. 18.

[32] S. Reich, "A nonparametric ensemble transform method for Bayesian inference," *SIAM. J. Sci. Comput.*, vol. 35, pp. A2013–A2014, 2013.

APPENDIX

The objective of this section is to describe the consistency result for the feedback particle algorithm (7), in the sense that the posterior distribution of the particle exactly matches the posterior distribution in the mean-field limit as $N \to \infty$. To put this in a mathematical framework, let π_t^* denote the conditional distribution of $X_t \in \mathscr{X}$ given history (filtration) of observations $\mathscr{Z}_t := \sigma\{Z_s, s \le t\}$[6]. For a real-valued function $f : \mathscr{X} \to \mathbb{R}$ define the action of π_t^* on f according to:

$$\pi_t^*(f) := \mathsf{E}[f(X_t)|\mathscr{Z}_t].$$

The time evolution of $\pi_t^*(f)$ is described by Kushner-Stratonovich pde (see Theorem 5.7 in [31]),

$$\pi_t^*(f) = \pi_0^*(f) + \int_0^t \pi_s^*(\mathscr{L}f)\,\mathrm{d}s + \int_0^t (\pi_s^*(fh) - \pi_s^*(f)\pi_s^*(f))\,(\mathrm{d}Z_s - \pi_s^*(h)\,\mathrm{d}s)\,, \tag{24}$$

where $f \in C_c^\infty(\mathscr{X})$ (smooth functions with compact support), and

$$\mathscr{L}f := a(x)\cdot\nabla f(x) + \frac{1}{2}\sum_{k,l=1}^{d} \sigma_k(x)\sigma_l(x)\frac{\partial^2 f}{\partial x_k \partial x_l}(x).$$

Next define π_t to be the conditional distribution of $X_t^i \in \mathscr{X}$ given $\mathscr{Z}_t$. It's action on real-valued function is defined according to:

$$\pi_t(f) := \mathsf{E}[f(X_t^i)|\mathscr{Z}_t].$$

[6]The state space $\mathscr{X}$ is $\mathbb{R}^d$ in this paper but the considerations of this section also apply when $\mathscr{X}$ is a differential manifold.

The time evolution of $\pi_t(f)$ is described by the Fokker-Planck equation (see Proposition 1 in [22]):

$$\pi_t(f) = \pi_0(f) + \int_0^t \pi_s(\mathscr{L}f)\,\mathrm{d}s + \int_0^t \pi_s(\mathsf{K}\cdot\nabla f)\,\mathrm{d}Z_s + \int_0^t \pi_s(u\cdot\nabla f)\,\mathrm{d}s + \frac{1}{2}\int_0^t \pi_s(\mathsf{K}\cdot\nabla(\mathsf{K}\cdot\nabla f)\,\mathrm{d}s, \tag{25}$$

where K is the gain function and u is the control function as follows:

1) Gain function: Let $\phi \in H^1(\mathscr{X};\pi_t)$[7] be the solution of a (weak form of the) Poisson equation:

$$\begin{aligned} \pi_t\big(\nabla\phi\cdot\nabla\psi\big) &= \pi_t\big((h-\pi_t(h))\psi\big), \\ \pi_t(\phi) &= 0 \qquad \text{(mean-zero)}, \end{aligned} \tag{26}$$

for all $\psi \in H^1(\mathscr{X};\pi_t)$. The gain function $\mathsf{K}(x,t) = \nabla\phi(x)$.

2) Control function: The function $u(x,t) = -\frac{1}{2}\mathsf{K}(x,t)\big(h(x)+\pi_t(h)\big)$.

The existence-uniqueness of the gain function K requires additional assumptions on the distribution π_t and the function h.

(i) **Assumption A1:** The probability distribution π_t admits a spectral gap. That is , $\exists\lambda > 0$ such that for all functions $f \in H^1_0(\mathscr{X},\pi_t)$,

$$\pi_t(|f|^2) \le \frac{1}{\lambda}\pi_t(|\nabla f|^2),$$

for $t \in [0,T]$.

(ii) **Assumption A2:** The function $h \in L^2(\mathscr{X};\pi_t)$, the space of square-integrable functions with respect to π_t.

Under the Assumptions (A1) and (A2) the Poisson equation (26) has a unique solution $\phi \in H^1_0(\mathscr{X},\pi_t)$ and the resulting control and gain function will be admissible [22]. The consistency of the feedback particle filter is stated in the following Theorem.

Theorem 1: Let π^*_t and π_t satisfy the forward equations (24) and (25), respectively. Then assuming $\pi^*_0 = \pi_0$, we have:

$$\pi^*_t(f) = \pi_t(f), \tag{27}$$

for all $t \in [0,T]$ and all functions $f \in C^\infty_c(\mathbb{R}^d)$.

[7] $H^1(\mathscr{X};\rho)$ is the Sobolev space of functions on $\mathscr{X}$ that are square-integrable with respect to density ρ and whose (weak) derivatives are also square-integrable with respect to density ρ.

Proof: Using (24) and (25) it is sufficient to show the following two identities:

$$\pi_s(\mathsf{K}\cdot\nabla f) = \pi_s(fh) - \pi_s(f)\pi_s(h),$$
$$\pi_s(u\cdot\nabla f) + \frac{1}{2}\pi_s(\mathsf{K}\cdot\nabla(\mathsf{K}\cdot\nabla f) = -\left(\pi_s(fh) - \pi_s(f)\pi_s(h)\right)\pi_s(h).$$

The first identity is obtained by using $\mathsf{K} = \nabla\phi$ and the weak form of the Poisson equation (26) with $\psi = f$. The second identity is obtained similarly. Use the expression for the control function u to obtain,

$$\begin{aligned}\pi_s(u\cdot\nabla f) &= -\pi_s(\frac{h+\pi_s(h)}{2}\mathsf{K}\cdot\nabla f)\\ &= -\frac{1}{2}\pi_s((h-\pi_s(h))\mathsf{K}\cdot\nabla f) - \pi_s(h)\pi_s(\mathsf{K}\cdot\nabla f)\\ &= -\frac{1}{2}(\mathsf{K}\cdot\nabla(\mathsf{K}\cdot\nabla f)) - \pi_s(h)\pi_s((h-\pi_s(h))f),\end{aligned}$$

where in the last step the weak form of the Poisson equation (26) is used for the $\psi = \mathsf{K}\cdot\nabla f$ and $\psi = f$. This concludes the second identity and hence the Theorem. ■

Let S_ε be the optimal transport map, solution of the problem (21). It is known that this map is of a gradient form [23]. In the following we furthermore assume that the map $S_\varepsilon =: \nabla\Phi_\varepsilon$ is C^1 in the parameter ε. For any test function ψ

$$C_\varepsilon = \int \psi(\nabla\Phi_\varepsilon(x))\rho(x)\,\mathrm{d}x - \int \psi(x)\rho_\varepsilon(x)\,\mathrm{d}x \equiv 0$$

Therefore,

$$\left.\frac{\mathrm{d}C_\varepsilon}{\mathrm{d}\varepsilon}\right|_{\varepsilon=0} = \int \nabla\phi(x)\cdot\nabla\psi(x)\rho(x)\,\mathrm{d}x - \int (h(x)-\hat{h})\psi(x)\rho(x)\,\mathrm{d}x = 0$$

where $\phi := \left.\frac{\mathrm{d}\Phi_\varepsilon}{\mathrm{d}\varepsilon}\right|_{\varepsilon=0}$. This is the weak form of the Poisson equation (9). Therefore,

$$\mathsf{K} = \nabla\phi = \left.\frac{\mathrm{d}S_\varepsilon}{\mathrm{d}\varepsilon}\right|_{\varepsilon=0}$$

A. *Derivation of the Linear Program* (22)

The ensemble transform method, introduced in [32], is an algorithm to sample N particles from ρ_ε, given N samples $\{X^i\}_{i=1}^N$ from ρ. The main step in this algorithm is to solve the following linear program:

$$\begin{aligned}&\text{Objective:} \quad \min_{\{t_{ij}\}} \sum_{i=1}^N\sum_{j=1}^N t_{ij}|X^i - X^j|^2\\ &\text{Constraints:} \quad \sum_{j=1}^N t_{ij} = 1, \quad \sum_{i=1}^N t_{ij} = w_j N, \quad t_{ij}\ge 0\end{aligned} \tag{28}$$

where $w_j := \frac{e^{\varepsilon h(X^j)}}{\sum_i e^{\varepsilon h(X^i)}}$ are the importance weights related to ρ_ε. The solution, denoted as t^*_{ij}, is referred to as the optimal coupling. The associated optimal value is an approximation of the optimal value of the objective in (21). The latter is the celebrated Wasserstein distance between ρ and ρ_ε. In terms of the optimal coupling, the N samples from ρ_ε are obtained as

$$S^i_\varepsilon = \sum_{j=1}^{N} t^*_{ij} X^j$$

Now, in the limit as $\varepsilon \downarrow 0$, samples S^i_ε can be used to obtain the samples K^i for the optimal vector-field K (see (21)). More precisely,

$$\mathsf{K}^i = \left.\frac{\mathrm{d}S^i_\varepsilon}{\mathrm{d}\varepsilon}\right|_{\varepsilon=0} = \lim_{\varepsilon\downarrow 0} \frac{S^i_\varepsilon - X^i}{\varepsilon} = \lim_{\varepsilon\downarrow 0} \frac{\sum_{j=1}^N (t^*_{ij} - \delta_{ij}) X^j}{\varepsilon}$$

where δ_{ij} is the Dirac delta tensor ($\delta_{ij} = 1$ if $i = j$ and 0 otherwise).

This also suggests a more direct approach whereby the linear program is solved with transition rates $a_{ij} = \frac{t_{ij} - \delta_{ij}}{\varepsilon}$. The linear program (22) is expressed with a_{ij} where the constraint is obtained by taking the following limit

$$\lim_{\varepsilon\downarrow 0} \frac{w^j N - 1}{\varepsilon} = h(X^j) - \hat{h}^{(N)}$$

The solution of the linear program (22) directly yields a formula for K^i in (23).