

Tutorial article

A survey of feedback particle filter and related controlled interacting particle systems (CIPS)[☆]Amirhossein Taghvaei^a, Prashant G. Mehta^{b,*}^a William E. Boeing Department of Aeronautics & Astronautics, University of Washington, Seattle, 98195, WA, USA^b Coordinated Science Laboratory and the Department of Mechanical Science and Engineering, University of Illinois, Urbana-Champaign, 61801, IL, USA

ARTICLE INFO

Keywords:

Nonlinear filtering
Mean-field optimal control
Optimal transportation theory

ABSTRACT

In this survey, we describe controlled interacting particle systems (CIPS) to approximate the solution of the optimal filtering and the optimal control problems. Part I of the survey is focussed on the feedback particle filter (FPF) algorithm, its derivation based on optimal transportation theory, and its relationship to the ensemble Kalman filter (EnKF) and the conventional sequential importance sampling–resampling (SIR) particle filters. The central numerical problem of FPF—to approximate the solution of the Poisson equation—is described together with the main solution approaches. An analytical and numerical comparison with the SIR particle filter is given to illustrate the advantages of the CIPS approach. Part II of the survey is focussed on adapting these algorithms for the problem of reinforcement learning. The survey includes several remarks that describe extensions as well as open problems in this subject.

1. Introduction

In many applications, dynamic models exist only in the form of a simulator. Our aim is to provide a survey of a class of algorithms, that use *only* a model simulator, to solve the two canonical problems of Control Theory:

- Design of optimal filter (in the sense of estimation);
- Design of optimal control law.

In this survey, such simulation-based algorithms are broadly referred to as *controlled interacting particle systems (CIPS)*. Our research group's most well known contribution to CIPS is the feedback particle filter (FPF), which is also the main focus of this survey. The FPF algorithm is useful to approximate the optimal (nonlinear) filter. By making use of the duality between optimal control and filtering, the FPF algorithm is extended to approximate the solution of an optimal control problem.

We begin by describing the high-level idea for the two problems of optimal filtering and optimal control.

1.1. CIPS in optimal filtering

Mathematical problem: In continuous-time and continuous-space settings of the problem, the standard model of nonlinear (or stochastic) filtering is the following Itô stochastic differential equations (SDEs):

$$\text{State: } dX_t = a(X_t)dt + \sigma_B(X_t)dB_t, \quad X_0 \sim p_0, \quad (1a)$$

$$\text{Observation: } dZ_t = h(X_t)dt + dW_t, \quad (1b)$$

where $X_t \in \mathbb{R}^d$ and $Z_t \in \mathbb{R}^m$ are the state and observation, respectively, at time t , p_0 is the probability density function (PDF) at the initial time $t = 0$ (p_0 is referred to as the prior density), and $\{B_t\}_{t \geq 0}$, $\{W_t\}_{t \geq 0}$ are mutually independent standard Wiener processes (W.P.) taking values in $\mathbb{R}^q$ and $\mathbb{R}^m$, respectively. The mappings $a(\cdot)$, $h(\cdot)$, $\sigma_B(\cdot)$, and the density $p_0(\cdot)$ are smooth (continuously differentiable) functions. The linear Gaussian model is obtained when the drift terms $a(\cdot)$, and $h(\cdot)$ are linear functions, $\sigma_B(\cdot)$ is a constant matrix, and p_0 is a Gaussian density.

[☆] The authors gratefully acknowledge the continued support of the National Science Foundation through the current grant 1761622, and the past grants 1435555 and 1334987.

The authors' prior research—reported here—is based on contributions from a number of graduate students and collaborators at the University of Illinois at Urbana-Champaign. These collaborations are gratefully acknowledged.

* Corresponding author.

E-mail addresses: amirtag@uw.edu (A. Taghvaei), mehtapg@illinois.edu (P.G. Mehta).

The filtering problem is to compute the conditional PDF of the state X_t given the time-history (filtration) of observations up to time t . The conditional PDF is denoted by p_t and is referred to as the posterior density.

CIPS algorithm: involves construction of N stochastic processes $\{X_t^i \in \mathbb{R}^d : t \geq 0, 1 \leq i \leq N\}$ where the i th process (particle) evolves according to the SDE:

$$dX_t^i = \underbrace{a(X_t^i)dt + \sigma_B(X_t^i)dB_t^i}_{\text{ith copy of model (1a)}} + dU_t^i, \quad X_0^i \stackrel{\text{i.i.d.}}{\sim} p_0, \quad (2)$$

where $U := \{U_t^i : t \geq 0, 1 \leq i \leq N\}$ is referred to as the *coupling* (with $U = 0$, the N processes are un-coupled). The goal is to design the coupling U so that the empirical distribution of the N particles at any time t approximates the posterior p_t :

$$\frac{1}{N} \sum_{i=1}^N f(X_t^i) \approx \int_{\mathbb{R}^d} f(x) p_t(x) dx, \quad \forall f \in C_b(\mathbb{R}^d), \quad (3)$$

where “ $\approx$ ” means that the approximation error goes to zero (in a suitable sense) as $N \rightarrow \infty$ ($C_b(\mathbb{R}^d)$ is the space of continuous and bounded functions on $\mathbb{R}^d$).

A key breakthrough, that appeared around 2010, is that U can be realized as a mean-field type feedback control law (“mean-field type” means that the control law depends also on the statistics of the stochastic process). Feedback particle filter (FPF) is one such example of a mean-field type control law. In this paper, we describe the FPF, relate it to its historical precursor, the ensemble Kalman filter (EnKF) algorithm, and summarize the important developments in this area.

For the filtering model (1), the idea of controlling the particles to approximate the posterior appears in the work of three groups working independently: the first example of such a control law appears in Crisan and Xiong (2010) using a certain smoothed form of observations. The FPF formula appears in Yang, Mehta, and Meyn (2011a, 2011b) and its special case for the linear Gaussian model is described in Bergemann and Reich (2012) and Reich (2011). A comparison of these three early works can be found in Pathiraja, Reich, and Stannat (2021). For the discrete-time filtering models, closely related ideas and algorithms were proposed, also around the same time-frame, by Daum and Huang (2008), El Moselhy and Marzouk (2012), Reich (2013) and Yang, Blom, and Mehta (2014) (see Spantini, Baptista, & Marzouk, 2022 for a recent review of this literature).

Our early work on FPF was closely inspired by the pioneering developments in mean-field games (Huang, Caines, & Malhame, 2007; Huang, Malhame, & Caines, 2006). The topic of mean-field games and mean-field type optimal control is concerned with control and decision problems arising in interacting particle systems. Over the past decade, this topic has grown in significance with theory and applications described in several monographs (Bensoussan et al., 2013; Carmona et al., 2018; Gomes, Pimentel, & Voskanyan, 2016). In the Physics literature, the study of interacting particle systems is a classical subject (Liggett, 1985). A canonical example of an interacting particle system is the coupled oscillators model of Kuramoto (Dörfler & Bullo, 2014; Kuramoto, 1975; Strogatz, 2000). Extensions of the classical Kuramoto model to mean-field games appears in Carmona and Graves (2020) and Yin, Mehta, Meyn, and Shanbhag (2011) and to FPF is given in Tilton, Hsiao-Wecksler, and Mehta (2012).

Design of CIPS to approximate the optimal control law is a more recent development. The idea is described next.

1.2. CIPS in optimal control

Mathematical problem: Consider a finite-horizon deterministic optimal control problem:

$$\min_u J(u) = \int_0^T \left(\frac{1}{2} |c(x_t)|^2 + \frac{1}{2} u_t^T R u_t \right) dt + g(x_T), \quad (4a)$$

$$\text{subject to: } \dot{x}_t = a(x_t) + b(x_t)u_t, \quad x_0 = x. \quad (4b)$$

where $x_t \in \mathbb{R}^d$ is the state at time t and $u := \{u_t \in \mathbb{R}^m : 0 \leq t \leq T\}$ is the control input. The mappings $a(\cdot)$, $b(\cdot)$, $c(\cdot)$, $g(\cdot)$ are smooth functions and R is a strictly positive-definite matrix (henceforth denoted as $R > 0$). The linear quadratic (LQ) model is obtained when $a(x) = Ax$, $b(x) = B$, $c(x) = Cx$, and $g(x) = x^T P_T x$. The infinite-time horizon ($T = \infty$) case is referred to as the linear quadratic regulator (LQR) problem.

CIPS algorithm: involves construction of N stochastic processes $\{Y_t^i \in \mathbb{R}^d : 0 \leq t \leq T, 1 \leq i \leq N\}$ where the i th particle evolves according to an SDE

$$dY_t^i = \underbrace{a(Y_t^i)dt + b(Y_t^i)dv_t^i}_{\text{ith copy of model (4b)}} + U_t^i dt, \quad 0 \leq t \leq T, \quad (5a)$$

where the input $v := \{v_t^i \in \mathbb{R}^m : 0 \leq t \leq T\}$ and the coupling $U := \{U_t^i \in \mathbb{R}^d : 0 \leq t \leq T\}$ are obtained as part of the design. The goal is to design v and U so that the empirical distribution of the N particles at time t approximates a smooth density p_t encoding the optimal control law $u_t = \phi_t^*(x_t)$ where

$$\phi_t^*(x) = R^{-1} b^T(x) \nabla \log p_t(x), \quad 0 \leq t \leq T, \quad (5b)$$

and ∇ denotes the gradient operator. In the infinite-horizon case, a stationary policy is obtained by letting $T \rightarrow \infty$.

The righthand-side of the formula (5b) is a consequence of the log transformation. The transformation relates the value function of an optimal control problem to the posterior density of the dual optimal filtering problem (Fleming & Mitter, 1982; Mitter & Newton, 2003). This manner of converting an optimal control problem into an optimal filtering problem (and vice-versa) is referred to as the minimum energy duality (Hijab, 1980; Mortensen, 1968). The use of this duality to express and solve an estimation problem as an optimal control problem is a standard approach in model predictive control (Rawlings, Mayne, & Diehl, 2017, Ch. 4). The CIPS (5a) comes about from the use of duality in the opposite direction whereby an optimal control problem (4) is solved using a filtering-type algorithm. Related constructions, based on somewhat different algorithmic approaches, is an important theme in the Robotics literature (Hoffmann & Rostalski, 2017; Kappen, 2005a, 2005b; Todorov, 2007; Toussaint, 2009; Vijayakumar, Rawlik, & Toussaint, 2013) (see Levine, 2018 for a recent review).

Both (2) and (5) are examples of a “simulation-based” algorithm because multiple copies—of the model (1a) and (4b), respectively—are run in a Monte-Carlo manner. *The main message of our paper is that through a suitable design of interactions between simulations—referred to as coupling—yields powerful algorithms for solving optimal filtering and optimal control problems.*

1.3. Relationship to other simulation-based algorithms

For the two problems of filtering and control, related simulation-based solution approaches are considered in the data assimilation (DA) and reinforcement learning (RL) communities, respectively. These relationships are discussed next.

Data assimilation (DA). The term “Data Assimilation” means assimilating real-time observations (“data”) into models—which typically exist only as a software code. The term is used by a community of researchers working in geophysical and atmospheric sciences (Evensen, 2006; Houtekamer & Mitchell, 2001; Reich & Cotter, 2015; Van Leeuwen & Evensen, 1996). The most celebrated application is weather prediction and forecast. For the abstract mathematical model, the nonlinear filter gives the optimal solution. In practice, the filter must be approximated in a computationally tractable form. For this purpose, the EnKF algorithm was first introduced in Evensen (1994) as an alternative to the extended Kalman filter (EKF). In geophysical applications, there are two issues that adversely affect the implementation of an extended Kalman filter:

1. In high-dimensions, it is a challenge to compute the Kalman gain. This is because the formula for the Kalman gain is based on the solution of a certain differential Riccati equation (DRE). The matrix-valued nature of the DRE means that any algorithm is $\mathcal{O}(d^2)$ in the dimension d of the state-space.
2. The model parameters are not explicitly available to write down the DRE let alone solve it. This is a concern whenever the model exists only in the form of a black-box numerical simulator.

In an EnKF implementation, N processes are simulated (same as (2)). In order to compute the Kalman gain, the solution of the DRE at time t is approximated by the empirical covariance of the ensemble $\{X_t^i\}_{i=1}^N$. Because an explicit solution of the DRE is avoided, an EnKF can be implemented using only a model simulator. This property has historically proved to be an important factor in applications. Notably, the EnKF algorithm is a workhorse for the weather prediction application (Evensen, 2003; Houtekamer & Zhang, 2016). The computational complexity of the EnKF is $\mathcal{O}(Nd)$ and in high-dimensions, N is chosen to be much smaller than d .

The historical significance of the FPF is that it represents a simulation-based solution of the nonlinear filtering problem (1), for arbitrary types of non-Gaussian posterior density p_t (under some mild technical conditions). Moreover, the EnKF was shown to arise as a special case in the linear Gaussian setting of the problem. Like the Kalman filter, the FPF formula has a “gain times error” feedback structure which is useful in several ways, e.g., to handle additional uncertainty in signal and measurement models. For these reasons, FPF can be viewed as a modern extension to the Kalman filter, a viewpoint stressed in a prior review paper (Taghvaei, De Wiljes, Mehta, & Reich, 2018).

For the nonlinear filtering problem (1), the FPF represents an alternative solution approach to the sequential importance sampling–resampling (SIR) particle filters and its many variants (Bain & Crisan, 2009; Del Moral, 2004; Doucet & Johansen, 2009; Gordon, Salmond, & Smith, 1993). In an SIR filter, the posterior is approximated as (compare with (3))

$$\int_{\mathbb{R}^d} f(x) p_t(x) dx \approx \sum_{i=1}^N W_t^i f(X_t^i), \quad \forall f \in C_b(\mathbb{R}^d),$$

where X_t^i is a copy of the hidden state X_t and $\{W_t^i\}_{i=1}^N$ are the importance weights obtained from the Bayes’ formula. In practice, all but a few weights can become very small—an issue known as particle degeneracy. This issue is ameliorated using a re-sampling procedure. The salient feature of the FPF, compared to the conventional particle filters, is that the weights are uniform ($= \frac{1}{N}$) by construction. Because of this difference, FPF does not suffer from the particle degeneracy issue and does not require re-sampling. In several independent numerical evaluations and comparisons, it has been observed that FPF exhibits smaller simulation variance (Berntorp, 2015; Stano, Tilton, & Babuska, 2014; Tilton, Ghiotto, & Mehta, 2013; Yang, Mehta, & Meyn, 2013) and better scaling properties with the problem dimension compared to particle filters (Surace, Kutschireiter, & Pfister, 2019; Yang, Lugeisen, Mehta, & Meyn, 2016). Some of these analytical and numerical comparisons are highlighted in the paper.

Reinforcement learning (RL). RL is concerned with solving optimal control problems, such as (4) and its extensions. All of the standard choices are treated in the literature: continuous and discrete state-space and time, deterministic and stochastic dynamics, discounted and average cost structures, and finite and infinite time-horizon (Bertsekas & Tsitsiklis, 1996; Meyn, 2022). What makes the RL paradigm so different from optimal control as formalized by Bellman and Pontryagin in the 1950s is that in RL the system identification step is usually avoided. Instead, the optimal policy is approximated (“learned”) based on input–output measurements.

In popular media, RL is described as an “agent” that learns an approximately optimal policy based on interactions with the environment. Important examples of this idea include advertising, where there is no scarcity of real-time data. In the vast majority of applications we are not so fortunate, which is why successful implementation usually requires simulation of the physical system for the purposes of training. For example, DeepMind’s success story with Go and Chess required weeks of simulation for training on a massive collection of super-computers (Schrittwieser et al., 2020).

These success stories are largely empirical. In order to better understand the theoretical foundations of RL, there has been a concerted recent interest, in the Control community, to revisit the classical linear quadratic (LQ) optimal control problem (Dean, Mania, Matni, Recht, & Tu, 2020; Fazel, Ge, Kakade, & Mesbahi, 2018; Malik et al., 2020; Mohammadi, Zare, Soltanolkotabi, & Jovanović, 2022; Tu & Recht, 2019). The two issues discussed as part of DA are relevant also to this problem: In high-dimensions, it is a challenge to solve the Riccati equation, and typically the model parameters are not explicitly available in RL settings of the problem.

An outgrowth of this recent work is a class of simulation-based algorithms where multiple copies of the simulator are run in parallel to learn and iteratively improve the solution of the DRE. The CIPS algorithm (5a) has the same structure where the important distinction is that the simulations are now coupled with a coupling term. We include comparisons on a benchmark problem to show how coupling helps improve performance over state-of-the-art.

1.4. Structure of the paper and outline

This paper is divided into two parts as follows:

- Part I on CIPS for the optimal filtering problem (1). It comprises Section 2 - Section 5.
- Part II on CIPS for the optimal control problem (4). It comprises Section 6.

The paper is written so that the key ideas are easily accessible together with an understanding of the main computational problems and algorithms for the same. For example, a reader should be able to implement the FPF and EnKF algorithms after reading Section 3 and Section 4. The more theoretical aspects related to optimal transportation theory appear in a self-contained manner in Section 5. The other significant aspect of this survey is analytical and numerical comparison against competing approaches. These appear in Section 3.4 for part I where a comparison with the SIR filter is discussed; and in Section 6.6 for part II where a comparison with RL algorithms for the LQR problem is described.

In writing any survey or review paper, one must make a choice of not only the topics to include but also the ones to leave out. Our choice is guided by our own area of expertise and by the intended audience in the Control community (where most of our own prior work has been published). We have stressed the interpretation of coupling as a mean-field feedback control law and highlighted its connection to optimal transportation. Both of these are important research themes in the community with related work on mean-field optimal control. The mathematics is most elegant in the continuous-time settings of the problem which is also the setting of this paper. A number of important aspects have not been covered in detail: On the theoretical side, the well-posedness of the mean-field model and justification of the mean-field limit are both hard mathematical topics. For a reader interested in some of these topics, we have included some high level remarks with references where additional details can be found. On the practical side, important issues arise on account of numerical discretization of the SDEs. Such numerical aspects have been entirely left out of this paper.

We make note of two final points: (i) While the paper presents some relatively novel ideas that are closely inspired by and connected to

the work in mean-field modeling and control, and therefore of interest to the Control community, these algorithms have older roots (EnKF) in the DA community. Along with the discussion in the Introduction, several remarks are included to highlight these roots and connections. (ii) While the CIPS algorithms solve some problems (such as particle degeneracy), they also create new ones. This informs the structure of the paper with a dedicated Section 4 on the central numerical problem of FPF. In particular, the discussion of the bias–variance trade-off in Section 4.3 is helpful to understand some of the key limitations in high dimensions.

PART I

2. Background on optimal filtering

Consider the filtering problem for the model (1). The sigma-algebra (on the time–history) of observations up to time t is denoted by $\mathcal{Z}_t := \sigma(\mathcal{Z}_s : 0 \leq s \leq t)$. The posterior density p_t is defined as follows:

$$\int_{\mathbb{R}^d} f(x) p_t(x) dx := \mathbb{E}[f(X_t) | \mathcal{Z}_t], \quad \forall f \in C_b(\mathbb{R}^d),$$

where the conditional expectation on the righthand-side is referred to as the nonlinear filter. The integral on the lefthand-side is denoted by $\langle p_t, f \rangle$.

The posterior p_t is optimal in the sense that, among all $\mathcal{Z}_t$ -measurable random variables, $\langle p_t, f \rangle$ represents the best mean-squared error (MSE) estimate of the random variable $f(X_t)$:

$$\langle p_t, f \rangle = \arg \min_{S \in \mathcal{Z}_t} \mathbb{E}[|f(X_t) - S|^2], \quad (6)$$

where the notation “ $S \in \mathcal{Z}_t$ ” means S is allowed to be $\mathcal{Z}_t$ -measurable, i.e., an arbitrary measurable function of observations up to time t .

For the model (1), the evolution of the posterior p_t is given by the Kushner–Stratonovich stochastic partial differential equation (Xiong, 2008, Ch. 5). In the special linear Gaussian setting of the problem, the equation admits a finite-dimensional representation given by the Kalman–Bucy filter.

2.1. Linear Gaussian model and the Kalman–Bucy filter

The linear Gaussian model is a special case of (1a)–(1b) and takes the following form:

$$dX_t = AX_t + \sigma_B dB_t, \quad X_0 \sim \mathcal{N}(m_0, \Sigma_0), \quad (7a)$$

$$dZ_t = HX_t dt + dW_t, \quad (7b)$$

where A, H, σ_B are matrices of appropriate dimensions and the prior is a Gaussian density with mean m_0 and variance Σ_0 . It is denoted by $\mathcal{N}(m_0, \Sigma_0)$.

For the linear Gaussian model (7), it can be shown that the posterior p_t is a Gaussian density. It is denoted by $\mathcal{N}(m_t, \Sigma_t)$, where m_t and Σ_t are conditional mean and covariance, respectively. Their evolution is described by the Kalman–Bucy filter (Kalman & Bucy, 1961):

$$dm_t = Am_t + K_t(dZ_t - Hm_t dt), \quad m_0 \text{ (given)} \quad (8a)$$

$$\frac{d}{dt} \Sigma_t = \text{Ric}(\Sigma_t), \quad \Sigma_0 \text{ (given)} \quad (8b)$$

where $K_t := \Sigma_t H^T$ is referred to as the Kalman gain, and the Riccati function

$$\text{Ric}(\Sigma) := A\Sigma + \Sigma A^T + \Sigma_B - \Sigma H^T H \Sigma$$

with $\Sigma_B := \sigma_B \sigma_B^T$.

Apart from the linear Gaussian model, there are very few examples where the equation for the posterior p_t admits a finite-dimensional representation (Beneš, 1981). In the general setting of the nonlinear model (1) with a non-Gaussian posterior, p_t is numerically approximated.

3. Feedback particle filter

Feedback particle filter (FPF) is a numerical algorithm to approximate the posterior p_t for the filtering model (1). Before describing the FPF, it is helpful to consider a simpler static problem.

3.1. Intuitive explanation with a simpler example

Suppose the state X and the observation Y are vector-valued random variables of dimension d and m , respectively. The probability distribution (prior) of X is denoted by P_X and the joint distribution of (X, Y) is denoted by P_{XY} . For any given function $f \in C_b(\mathbb{R}^d)$, the problem is to obtain an MSE estimate of the unknown $f(X)$ from a single observation of Y . Adapting (6) to the simple case,

$$S_f^*(Y) = \arg \min_{S_f(\cdot)} \mathbb{E}[|f(X) - S_f(Y)|^2], \quad (9)$$

where on the righthand-side $S_f : \mathbb{R}^m \rightarrow \mathbb{R}$ is allowed to be an arbitrary function of the $\mathbb{R}^m$ -valued observation (the sub-script means that the function may depend also upon f). The optimal estimator gives the conditional expectation, i.e., $\mathbb{E}[f(X)|Y] = S_f^*(Y)$.

Example 3.1 (Linear Estimation and the Update Formula for Kalman Filter). Consider the case where f is linear, $f(x) = a^T x$, and $S_f(\cdot)$ is restricted to be an affine function of its argument:

$$S_f(y) = u^T y + b,$$

where $u \in \mathbb{R}^m$ and $b \in \mathbb{R}$ parametrize the estimator. With such a choice, the optimization problem (9) is finite-dimensional whose solution is readily obtained as

$$S_f^*(Y) = a^T (\mathbb{E}[X] + K(Y - \mathbb{E}[Y])),$$

where $K = \Sigma_{XY} \Sigma_Y^{-1}$, $\Sigma_{XY} = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])^T]$, $\Sigma_Y = \mathbb{E}[(Y - \mathbb{E}[Y])(Y - \mathbb{E}[Y])^T]$, and it is assumed that Σ_Y is invertible with inverse Σ_Y^{-1} . Because the vector a is arbitrary, this also shows that the optimal linear estimate of X is $\mathbb{E}[X] + K(Y - \mathbb{E}[Y])$. Under the stronger assumption that X and Y are jointly Gaussian, it can be shown that this is in fact the optimal estimate of X among all functions $S_f(\cdot)$ (not necessarily affine) (Hajek, 2015, Prop. 3.9). Therefore, in the Gaussian case

$$\mathbb{E}[X|Y] = \mathbb{E}[X] + K(Y - \mathbb{E}[Y]).$$

The righthand-side is the update formula for the discrete-time Kalman filter. Note that the interpretation of the formula as the conditional expectation works only in the Gaussian case. In general, the formula gives only the best linear estimator.

The example above illustrates the special and important case of obtaining optimal linear estimators. The question is how to extend the procedure to the nonlinear setting, i.e., the setting where both the function $f(\cdot)$ and the estimator $S_f(\cdot)$ are allowed to be nonlinear functions of their arguments. This is achieved through the concept of CIPS whose construction proceeds in two steps:

Step 1: Let $\bar{X}_0$ be an independent copy of X . Design a control U such that, upon setting $\bar{X}_1 = \bar{X}_0 + U$,

$$S_f^*(Y) = \mathbb{E}[f(\bar{X}_1)|Y], \quad \forall f \in C_b(\mathbb{R}^d),$$

Note that the control is not allowed to depend on the function f . It is designed to give the best estimate for any choice of function f . It is not yet clear that such a control exists. But for now, let us assume that it exists and moreover takes the form $U = u(\bar{X}_0, Y)$. (Typically, the mapping $u(\cdot, \cdot)$ is designed to be a deterministic function but may in general also be random.)

Step 2: Generate N independent samples (particles) $\{X_0^1, \dots, X_0^N\}$ from P_X , update each particle according to

$$X_1^i = X_0^i + u(X_0^i, Y), \quad i = 1, 2, \dots, N,$$

and form a Monte-Carlo approximation of the estimate:

$$S_f^*(Y) \approx \frac{1}{N} \sum_{i=1}^N f(X_i^i).$$

Example 3.2 (CIPS and the Update Formula for EnKF). Continuing with Example 3.1 where P_{XY} is assumed to be Gaussian, two formulae are described for the transformation $\bar{X}_0 \mapsto \bar{X}_1$. The first of these formulae is based on optimal transportation theory. The second formula is based on the perturbed form of the discrete-time EnKF algorithm.

- Optimal transport formula is given by a deterministic affine mapping

$$\bar{X}_1 = A(\bar{X}_0 - \mathbb{E}[\bar{X}_0]) + K(Y - \mathbb{E}[Y]) + \mathbb{E}[\bar{X}_0],$$

where A is the unique such symmetric positive-definite solution to a Lyapunov equation

$$A\Sigma_X A = \Sigma_X - \Sigma_{XY} \Sigma_Y^{-1} \Sigma_{YX}.$$

- Perturbed EnKF formula. Let $(\bar{X}_0, \bar{Y}_0)$ be an independent copy of (X, Y) then

$$\bar{X}_1 = \bar{X}_0 + K(Y - \bar{Y}_0),$$

where the formula for K is same as in Example 3.1.

It is readily verified that, in either case, $\bar{X}_1$ is a Gaussian random variable whose conditional mean and variance equals the conditional mean and variance of X .

We defer the details on how these formulae came about to Section 5.3 instead remarking here on several features which apply also to more general settings:

1. The transformation $\bar{X}_0 \mapsto \bar{X}_1$ is not unique.
2. Both the transformations are of “mean-field type” whereby the transformation depends also on statistics, e.g., $\mathbb{E}[X]$ and $\mathbb{E}[Y]$, of (X, Y) .
3. In the optimal transport formula, $u(\cdot, \cdot)$ is a deterministic function. In the EnKF formula, $u(x, y) = K(y - \bar{Y}_0)$ is a random map because $\bar{Y}_0$ is a random variable.

These considerations provide the background for the feedback particle filter algorithm which is described next.

3.2. Feedback particle filter

Just like the static example, the construction of FPF proceeds in two steps.

Step 1: Construct a stochastic process, denoted by $\bar{X} = \{\bar{X}_t\}_{t \geq 0}$, according to a controlled SDE:

$$d\bar{X}_t = a(\bar{X}_t)dt + \sigma_B(\bar{X}_t)d\bar{B}_t + u_t dt + K_t dZ_t, \quad \bar{X}_0 \sim p_0, \quad (10)$$

where the controls u_t and K_t are designed so that the conditional density of $\bar{X}_t$ equals the posterior density p_t .

Step 2: Simulate N stochastic processes, denoted by $X^i = \{X_t^i\}_{t \geq 0}$ for $i = 1, 2, \dots, N$, according to (10).

The two steps are summarized below:

$$\underbrace{\langle p_t, f \rangle}_{\text{exactness condition}} \stackrel{\text{Step 1}}{=} \mathbb{E}[f(\bar{X}_t) | \mathcal{Z}_t] \stackrel{\text{Step 2}}{\approx} \frac{1}{N} \sum_{i=1}^N f(X_t^i).$$

The exactness condition refers to the fact that $\bar{X}_t$ has the same conditional density as X_t . The N processes $\{X_t^i\}_{i=1}^N$ are referred to as particles.

At this point, the first of these two steps appears to be aspirational. Even in the case of the static example, it is not at all clear that the function $u(\cdot, \cdot)$ exists in the general non-Gaussian case, and even if

it does, it can be computed in a tractable manner. The case of the stochastic process where u_t and K_t are allowed to be measurable with respect to the past values of observations Z and state $\bar{X}$ appears, at the first glance, to be entirely hopeless.

The surprising (at least at the time of its discovery) breakthrough of the FPF is that the control terms u_t and K_t are given by a simple feedback control law where the computation reduces to solving a linear Poisson equation at each time-step.

FPF: The process $\bar{X}$ is defined according to the SDE

$$d\bar{X}_t = \underbrace{a(\bar{X}_t)dt + \sigma_B(\bar{X}_t)d\bar{B}_t}_{\text{copy of model (1a)}} + \underbrace{K_t(\bar{X}_t) \circ (dZ_t - \frac{h(\bar{X}_t) + \bar{h}_t}{2} dt)}_{\text{FPF feedback control law}}, \quad \bar{X}_0 \sim p_0 \quad (11)$$

where $\{\bar{B}_t\}_{t \geq 0}$ is a copy of the process noise $\{B_t\}_{t \geq 0}$, and $\bar{h}_t := \mathbb{E}[h(\bar{X}_t) | \mathcal{Z}_t]$. The $\circ$ indicates that the SDE is expressed in its Stratonovich form. At any fixed time t , the gain $K_t(\cdot)$ is a $d \times m$ matrix-valued function obtained by solving m partial differential equations: for $j = 1, 2, \dots, m$, the j th column $K_t^{(j)} := \nabla \phi^{(j)}$ where $\phi^{(j)}$ is the solution of the Poisson equation:

$$-\frac{1}{\rho(x)} \nabla \cdot (\rho(x) \nabla \phi^{(j)}(x)) = (h^{(j)}(x) - \bar{h}^{(j)}), \quad x \in \mathbb{R}^d \quad (12)$$

where the density $\rho = \bar{p}_t$ (the conditional density of $\bar{X}_t$ at time t), $h^{(j)}$ is the j th component of the observation function h , $\bar{h}^{(j)} = \int h^{(j)}(x) \rho(x) dx$, and ∇ and $\nabla \cdot$ denote the gradient and the divergence operators, respectively. For a succinct presentation, the functions $\{\phi^{(j)}\}_{j=1}^m$ are collected to form the vector-valued function $\phi = [\phi^{(1)}, \dots, \phi^{(m)}]$. With such a notation, the gain function K_t is the Jacobian $\nabla \phi = [\nabla \phi^{(1)}, \dots, \nabla \phi^{(m)}]$.

The process $\bar{X}$ is an example of a mean-field process because its evolution depends upon its own statistics. An SDE of this type is called a McKean–Vlasov SDE or a mean-field SDE. Accordingly, (11) is referred to as the mean-field FPF.

The main result, first proved in Yang, Mehta, and Meyn (2013), is that the mean-field process thus defined is exact.

Theorem 3.3 (Thm 3.3, Yang, Mehta, & Meyn, 2013). Consider the filtering model (1). Suppose $\{p_t\}_{t \geq 0}$ denotes the conditional density of the process $\{X_t\}_{t \geq 0}$. Suppose the mean-field process $\{\bar{X}_t\}_{t \geq 0}$ defined by (11)–(12) is well-posed with conditional density denoted by $\{\bar{p}_t\}_{t \geq 0}$. Then, provided $\bar{p}_0 = p_0$,

$$\bar{p}_t = p_t, \quad \forall t > 0.$$

Remark 3.4 (Well-Posedness and Poincaré Inequality). The well-posedness of (11)–(12) means that a strong solution $\bar{X}$ exists with a well-defined density $\{\bar{p}_t\}_{t \geq 0}$. To show well-posedness, apart from the standard Lipschitz condition on the drift terms $a(\cdot)$ and $\sigma_B(\cdot)$, the main technical condition is that the posterior density p_t (of X_t) satisfies the Poincaré inequality (PI), and $\int |h(x)|^2 p_t(x) dx < \infty$ (Laugesen, Mehta, Meyn, & Raginsky, 2015, Theorem 2.2). (A probability density $\rho = e^{-V}$ satisfies the PI if $x^T \nabla V(x) \geq \alpha |x|$ for $|x| \geq R$ where α and R are positive constants (Bakry, Barthe, Cattiaux, & Guillin, 2008, Cor. 1.6). This condition is true, e.g., whenever ρ has a Gaussian tail.) An explanation of the relevance of the PI for the well-posedness (existence, uniqueness, and regularity) of the solution ϕ of the Poisson equation (12) is deferred to Section 4, where algorithms for its approximation are also described. Once a solution ϕ of the Poisson equation is obtained together with necessary apriori estimates, well posedness of $\bar{X}$ follows from the standard theory of mean-field SDEs (Carmona et al., 2018). Although the general case remains open, it has been possible to prove the PI under certain additional conditions on the filtering model (1) (Pathiraja et al., 2021, Lemma 5.1), (Laugesen et al., 2015, Prop 2.1).

We next describe the finite- N algorithm which is how the FPF is implemented in practice.

CIPS: The particles $\{X_t^i : t \geq 0, 1 \leq i \leq N\}$ evolve according to:

$$dX_t^i = a(X_t^i)dt + \sigma(X_t^i)dB_t^i + K_t^{(N)}(X_t^i) \circ (dZ_t - \frac{h(X_t^i) + h_t^{(N)}}{2} dt), \quad (13)$$

$$X_0^i \stackrel{\text{i.i.d.}}{\sim} p_0, \quad i = 1, \dots, N,$$

where $\{B_t^i\}_{i \geq 1}$, for $i = 1, 2, \dots, N$, are mutually independent W.P., $h_t^{(N)} := N^{-1} \sum_{i=1}^N h(X_t^i)$, and $K_t^{(N)}$ is the output of an algorithm that is used to approximate the solution to the Poisson equation (12):

$$K_t^{(N)} := \text{Algorithm}(\{X_t^i\}_{i=1}^N; h).$$

The notation is suggestive of the fact that algorithm is adapted to the ensemble $\{X_t^i\}_{i=1}^N$ and the function h ; the density $\bar{p}_t$ is not known in an explicit form. Before describing the algorithms for gain function approximation in (the following) Section 4, we discuss the linear Gaussian case.

The main computational challenge to simulate the finite- N FPF (13) is the computation of the gain function. The difficulty arises because, for a general nonlinear observation function h and a non-Gaussian density ρ , there are no known closed-form solutions of the Poisson equation (12). In the linear Gaussian special case, with a linear observation function $h(x) = Hx$ and a Gaussian density, the Poisson equation admits an explicit solution whereby the gain function is given by the Kalman gain:

Proposition 3.5 (Lem. 3.4, Yang, Mehta, & Meyn, 2013). Consider the Poisson equation (12). Suppose ρ is a Gaussian density $\mathcal{N}(m, \Sigma)$ and $h(x) = Hx$. Then its unique solution is given by:

$$\phi(x) = (H\Sigma)(x - m), \quad x \in \mathbb{R}^d.$$

Consequently, the gain function $\nabla \phi(x) = \Sigma H^T$ is the Kalman gain.

Using the Kalman gain, the FPF algorithm simplifies to a square-root form of the ensemble Kalman filter (EnKF) algorithm. This is described next.

3.3. Ensemble Kalman filter

In the linear Gaussian case, upon replacing the gain function with the Kalman gain, the mean-field FPF (11) is the Itô-SDE

$$d\bar{X}_t = A\bar{X}_t dt + \sigma_B dB_t + \bar{\Sigma}_t H^T (dZ_t - \frac{H\bar{X}_t + H\bar{m}_t}{2} dt), \quad (14)$$

where

$$\bar{m}_t = \mathbb{E}[\bar{X}_t | \mathcal{Z}_t],$$

$$\bar{\Sigma}_t = \mathbb{E}[(\bar{X}_t - \bar{m}_t)(\bar{X}_t - \bar{m}_t)^T | \mathcal{Z}_t].$$

As a corollary of Theorem 3.3, the mean-field process $\bar{X}$ is exact which, in the linear Gaussian case, means that the conditional density of $\bar{X}_t$ is Gaussian whose mean $\bar{m}_t$ and the covariance matrix $\bar{\Sigma}_t$ evolve according to the Kalman filter (8). A direct proof showing (14) is exact appears in Section 5.1.

The finite- N FPF is obtained as follows:

$$dX_t^i = AX_t^i dt + \sigma_B dB_t^i + \Sigma_t^{(N)} H^T (dZ_t - \frac{HX_t^i + Hm_t^{(N)}}{2} dt), \quad (15a)$$

where the mean-field terms in (14) are approximated empirically as follows:

$$m_t^{(N)} := \frac{1}{N} \sum_{j=1}^N X_t^j, \quad (15b)$$

$$\Sigma_t^{(N)} := \frac{1}{N-1} \sum_{j=1}^N (X_t^j - m_t^{(N)})(X_t^j - m_t^{(N)})^T. \quad (15c)$$

The linear Gaussian FPF (15) is identical to the square-root form of the ensemble Kalman filter (Bergemann & Reich, 2012, Eq. 3.3).

Remark 3.6 (Historical Context for EnKF). The EnKF algorithm was first introduced in Evensen (1994), in the discrete-time setting of the filtering problem. At the time, the algorithm was introduced as an alternative to the extended Kalman filter (EKF). As already mentioned in Section 1, a major reason for using an EnKF is that, unlike EKF, it does not require an explicit solution of the DRE (Burgers, Van Leeuwen, & Evensen, 1998; Houckamer & Mitchell, 1998; Van Leeuwen & Evensen, 1996). Since its introduction, a number of distinct types of EnKF algorithms have appeared in the literature. Amongst these, the most well-known types are as follows: (i) EnKF based on perturbed observation (Evensen, 2003); and (ii) The square root EnKF (Anderson, 2001; Bishop, Etherton, & Majumdar, 2001; Whitaker & Hamill, 2002). The details for these algorithms can be found in Reich and Cotter (2015, Ch. 6-7). The two aforementioned types of the EnKF algorithm have also been extended to the continuous-time setting (Bergemann & Reich, 2012). In these settings, the EnKF is usually referred to as the ensemble Kalman–Bucy filter (EnKBF). A review of the EnKBF algorithm and its connection to the FPF algorithm can be found in Taghvaei et al. (2018). The EnKBF algorithm and the linear FPF admits several extensions: (i) EnKBF with perturbed observation (Bergemann & Reich, 2012; Del Moral & Tugaut, 2018); (ii) Stochastic linear FPF (Yang et al., 2016, Eq. (26)) which is same as the square root EnKBF (Bergemann & Reich, 2012); (iii) Deterministic linear FPF (Taghvaei & Mehta, 2016, Eq. (15)) (de Wiljes, Reich, & Stannat, 2018). EnKF was recently extended to the case with correlated observation noise (Ertel & Stannat, 2022). An excellent recent survey on this topic appears in Calvello, Reich, and Stuart (2022).

Remark 3.7 (Current Research on EnKF). Error analysis of the EnKF algorithm remains an active area of research. For the discrete-time EnKF algorithm, these results appear in Kelly, Law, and Stuart (2014), Kwiatkowski and Mandel (2015), Le Gland, Monbet, and Tran (2009), Mandel, Cobb, and Beezley (2011) and Tong, Majda, and Kelly (2016). The analysis for continuous-time EnKF is more recent (Bishop & Del Moral, 2018, 2020; Chen, Luo, Shi, & Yau, 2021; de Wiljes et al., 2018; Del Moral, Kurtzmann, & Tugaut, 2017; Del Moral & Tugaut, 2018; Taghvaei & Mehta, 2018). Typically, one is interested in obtaining a uniform error bound as follows:

$$\mathbb{E}[\|m_t^{(N)} - m_t\|^2] + \mathbb{E}[\|\Sigma_t^{(N)} - \Sigma_t\|^2] \leq \frac{C}{\sqrt{N}}, \quad (16)$$

where (m_t, Σ_t) are the solutions of the Kalman filter (8) and $(m_t^{(N)}, \Sigma_t^{(N)})$ are obtained from simulating an EnKF; and $C > 0$ is a time-independent constant. In the most recent result (Bishop & Del Moral, 2020), (16) is shown under the assumption that $H^T H$ is a positive-definite matrix. It is expected that (16) also holds under the weaker condition of the pair (A, H) being detectable, which is the condition for the stability of the Kalman filter. However, a complete resolution is still open. A comprehensive review of recent developments in this area can be found in Bishop and Del Moral (2020).

3.4. Comparison with importance sampling

In this section, we provide an analytical comparison of the FPF with the importance sampling-based particle filter. For this purpose, consider a parameter estimation example with a fully observed model as follows:

$$dX_t = 0, \quad X_0 \sim \mathcal{N}(0, \sigma_0^2 I_d) = p_0, \quad (17)$$

$$dZ_t = X_t dt + \sigma_w dW_t,$$

where the time $t \in [0, 1]$, $\sigma_w, \sigma_0 > 0$, and I_d is the $d \times d$ identity matrix. The posterior p_1 at time $t = 1$ is a Gaussian $\mathcal{N}(m_1, \Sigma_1)$ with

$$m_1 = \frac{\sigma_0^2}{\sigma_0^2 + \sigma_w^2} Z_1 \text{ and } \Sigma_1 = \frac{\sigma_0^2 \sigma_w^2}{\sigma_0^2 + \sigma_w^2} I_d.$$




Fig. 1. Numerical comparison for the filtering model (17). Level sets of the MSE, using: (a) importance sampling-based algorithm (18) and (b) the FPF (19). As the state dimension d grows, in order to have same performance (MSE), the number of particles N must increase as 2^d for (18) while they increase as $d^{1/2}$ for (19).

Let $\{X_0^i\}_{i=1}^N$ be N i.i.d samples from the prior p_0 . The importance sampling-based particle filter yields an empirical approximation of the posterior p_1 as follows:

$$\pi_{\text{PF}}^{(N)}(f) := \sum_{i=1}^N W_1^i f(X_0^i), \quad W_1^i = \frac{e^{-\frac{|Z_1 - X_0^i|^2}{2\sigma_w^2}}}{\sum_{i=1}^N e^{-\frac{|Z_1 - X_0^i|^2}{2\sigma_w^2}}}. \quad (18)$$

In contrast, given the initial samples $\{X_0^i\}_{i=1}^N$, the FPF approximates the posterior by implementing a feedback control law as follows:

$$\pi_{\text{FPF}}^{(N)}(f) := \frac{1}{N} \sum_{i=1}^N f(X_1^i), \quad dX_t^i = \frac{\Sigma_t^{(N)}}{\sigma_w^2} (dZ_t - \frac{X_t^i + m_t^{(N)}}{2} dt), \quad (19)$$

where the mean $m_t^{(N)}$ and covariance $\Sigma_t^{(N)}$ are empirically approximated using (15b) and (15c), respectively.

The MSE in estimating the conditional expectation of a given function f is defined as follows:

$$\text{MSE}_*(f) := \mathbb{E}[|\pi_*^{(N)}(f) - \langle p_1, f \rangle|^2],$$

where the subscript $*$ is either the PF or the FPF.

For $f(x) = \frac{1}{\sqrt{d}} 1^T x$, a numerically computed plot of the level-sets of MSE, as a function of N and d , is depicted in Fig. 1-(a)-(b). The expectation is approximated by averaging over $M = 1000$ independent simulations. It is observed that, in order to have the same error, the importance sampling-based approach requires the number of samples N to grow exponentially with the dimension d , whereas the growth using the FPF for this numerical example is $O(d^{1/2})$. This conclusion is consistent with other numerical studies reported in the literature (Berntorp, 2015; Stano et al., 2014; Surace et al., 2019).

For the purposes of the analysis, a modified form of the particle filter is considered whereby the denominator is replaced by its exact form:

$$\pi_{\text{PF}}^{(N)}(f) := \sum_{i=1}^N \bar{W}_1^i f(X_0^i), \quad \bar{W}_1^i = \frac{e^{-\frac{|Z_1 - X_0^i|^2}{2\sigma_w^2}}}{N \mathbb{E}[e^{-\frac{|Z_1 - X_0|^2}{2\sigma_w^2}} | Z_1]}. \quad (20)$$

Proposition 3.8 (Prop. 4 in Taghvaei & Mehta, 2020). Consider the filtering problem (17) with state dimension d . Suppose $\sigma_0 = \sigma_w = \sigma > 0$ and $f(x) = a^T x$ where $a \in \mathbb{R}^d$ with $|a| = 1$. Then:

1. The MSE, for the modified importance sampling estimator (20) is given by

$$\text{MSE}_{\text{PF}}(f) = \frac{\sigma^2}{N} \left(3(2^d) - \frac{1}{2} \right) \geq \frac{\sigma^2}{N} 2^{d+1}.$$

2. The MSE for the FPF estimator (19) is bounded as

$$\text{MSE}_{\text{FPF}}(f) \leq \frac{\sigma^2}{N} (3d^2 + 2d). \quad (21)$$

Remark 3.9 (Curse of Dimensionality (CoD)). In the limit as $d \rightarrow \infty$, the performance of the importance sampling-based particle filters is studied in the literature (Bengtsson, Bickel, & Li, 2008; Bickel et al., 2008; Rebeschini et al., 2015; Snyder, Bengtsson, Bickel, & Anderson, 2008). The main focus of these studies is on the particle degeneracy (or the weight collapse) issue: it is shown that if $\frac{\log N \log d}{d} \rightarrow 0$ then the largest weight $\max_{1 \leq i \leq N} W_1^i \rightarrow 1$ in probability. Consequently, in order to prevent the weight collapse, the number of particles must grow exponentially with the dimension. This phenomenon is referred to as the curse of dimensionality for the particle filters. In contrast, the weights in an FPF are uniform by design (see (19)). Therefore, the FPF does not suffer from the weight collapse issue and, in particular, does not require resampling. A complete comparison of the two types of particle filters remains open (see Abedi, Surace, & Pfister, 2022 for recent progress on this topic).

Remark 3.10 (Scaling with the Dimension). The scaling with dimension depicted in Fig. 1 (b) suggests that the $O(d^2)$ bound in (21) is loose. This is the case because, in deriving the bound, the inequality $\|\cdot\|_2 \leq \|\cdot\|_F$ is used, where $\|\cdot\|$ and $\|\cdot\|_F$ denote the induced and Frobenius norms, respectively (Taghvaei & Mehta, 2020, Appendix E). The inequality is loose particularly so as the dimension grows. Also, it is observed that the MSE for the particle filter grows slightly slower than the lower-bound 2^d . This is because the lower-bound is obtained for the modified particle filter (20), while the MSE is numerically evaluated for the standard particle filter (18). The correlation between the numerator and denominator in (18) reduces the MSE.

3.5. Extensions of FPF

In deriving the FPF, the main modeling assumption is the nature of observation model (1b). (Such a model is referred to as the white noise observation model.) In several follow on works, the basic FPF is extended to handle more general types of models for the state process. These extensions are briefly described next.

FPF on Riemannian manifolds. The feedback control form of the FPF formula (11) holds not only for the Euclidean state-space but also for the cases where the state $\{X_t\}_{t \geq 0}$ evolves on a Riemannian manifold, such as the matrix Lie groups. These extensions are described in Zhang, Taghvaei, and Mehta (2016a, 2016b, 2017a, 2017b). In these papers,

the FPF is shown to provide an intrinsic description of the filter that automatically satisfies the geometric constraints of the manifold. The gain is expressed as $\text{grad } \phi$ and obtained as a solution of the Poisson equation. It is shown that the gain is also intrinsic that furthermore does not depend upon the choice of the Riemannian metric. For the special case when the manifold is a matrix Lie group, explicit formulae for the filter are derived, using the matrix coordinates. Filters for two example problems are presented: the attitude estimation problem on $SO(3)$ and the robot localization problem in $SE(3)$. Comparisons are also provided between the FPF and popular algorithms for attitude estimation, namely the multiplicative EKF, the invariant EKF, the unscented quaternion estimator, the invariant ensemble Kalman filter, and the bootstrap particle filter. Specifically, under a certain assumption of a “concentrated distribution”, the evolution equations for the mean and the covariance are shown to be identical to the left invariant EKF algorithm.

FPF on discrete state-space. In Yang, Mehta, and Meyn (2015), FPF is extended to the filtering problem where the hidden state $\{X_t\}_{t \geq 0}$ is a continuous-time Markov process that evolves on a finite state-space. (For this model, the optimal nonlinear filter is called the Wonham filter.) A standard algorithm to simulate a Markov process is based on the use of Poisson counters to simulate transitions between discrete states. In order to define the process $\bar{X}$, a control process U is introduced that serves to modulate the rates of these counters based on causal observations of data Z . An explicit formula for the FPF feedback control law is derived and shown to be exact. Similar to (11), the formula is in the form of “gain times error” where the gain is now obtained by solving a certain linear matrix problem. The linear matrix problem is the finite state-space counterpart of the Poisson equation (12).

FPF with data association and model uncertainty. In applications such as multiple target tracking, the filtering problem often involves additional uncertainties in the state model (1a) and the observation model (1b). In the classical linear Gaussian settings, algorithms based on the Kalman filter have been developed to provide a solution to these problems. These algorithms are referred to as the interacting multiple model (IMM) filter (Blom, 2013) and the probabilistic data association (PDA) filter (Bar-Shalom, Daum, & Huang, 2009). In the PDA filter, the Kalman gain is allowed to vary based on an estimate of the instantaneous uncertainty in the observations. In the IMM filter, multiple Kalman filters are run in parallel and their outputs combined to form an estimate.

Like the Kalman filter, the FPF is easily extended to handle additional uncertainties in the observation and signal models: These extensions, namely, the probabilistic data association (PDA)-FPF and the interacting multiple model (IMM)-FPF are derived in our prior works (Yang, Blom, & Mehta, 2013; Yang, Huang, & Mehta, 2012; Yang & Mehta, 2018). Structurally, the FPF based implementations are similar to the classical algorithms based on the Kalman filter. In the linear Gaussian settings, the equations for the mean and the variance of the FPF-based filters evolve according the classical PDA and IMM filters.

Collective inference FPF. The term “collective inference” is used to describe filtering problems with a large number of aggregate and anonymized data (Sheldon & Dietterich, 2011; Singh, Haasler, Zhang, Karlsson, & Chen, 2020). Some of these problems have gained in importance recently because of COVID-19. Indeed, the spread of COVID-19 involves dynamically evolving hidden processes (e.g., number of infected, number of asymptomatic etc.) that must be deduced from noisy and partially observed data (e.g., number of tested positive, number of deaths, number of hospitalized etc.). In carrying out data assimilation for such problems, one typically only has aggregate observations. For example, while the number of daily tested positives is available, the information on the disease status of any particular agent in the population is not known.

In Kim, Taghvaei, Chen, and Mehta (2021), the FPF algorithm is extended for a model with M agents and M observations. The M observations are non-agent specific. Therefore, in its basic form, the problem is characterized by data association uncertainty whereby the association between the observations and agents must be deduced in addition to the agent state. In Kim, Taghvaei, et al. (2021), the large- M limit is interpreted as a problem of collective inference. This viewpoint is used to derive the equation for the empirical distribution of the hidden agent states. An FPF algorithm for this problem is presented and illustrated via numerical simulations. Formulae are described for both the Euclidean and the finite state-space case. The classical FPF algorithm is shown to be the special case (with $M = 1$) of these more general results. The simulations help show that the algorithm well approximates the empirical distribution of the hidden states for large M .

Before closing this section, we remark on the Stratonovich form of the mean-field FPF SDE (11). The FPF is expressed in this form because of two reasons:

1. The feedback control law is “gain times error” which is appealing to control engineers, and structurally similar to the update formula in a Kalman filter. Moreover, for the linear Gaussian model, the gain is the Kalman gain.
2. Expressed in its Stratonovich form, the gain times error formula carries over to the Riemannian manifolds settings. This is because of the intrinsic nature of the Stratonovich form (Zhang et al., 2017b, Remark 1).

Notably, for the linear Gaussian model, the gain function is a constant (i.e., does not depend upon x) and therefore the Stratonovich form and the Itô form are the same. For the general case, the Itô form involves a Wong-Zakai correction term as described in the following remark.

Remark 3.11 (Itô Form of FPF). In its Itô form, the mean-field FPF (11) is expressed as

$$d\bar{X}_t = a(\bar{X}_t)dt + \sigma(\bar{X}_t)d\bar{B}_t + K_t(\bar{X}_t)(dZ_t - \frac{h(\bar{X}_t) + \tilde{h}_t}{2}dt) + \frac{1}{4} \sum_{j=1}^m \nabla |K_t^{(j)}(\bar{X}_t)|^2 dt,$$

where $\frac{1}{4} \sum_{j=1}^m \nabla |K_t^{(j)}(\bar{X}_t)|^2$ is the Wong-Zakai correction term. The Itô-Stratonovich relationship discussed here is based on interpreting $K_t(x)$ as a function of space x and time t , and interpreting the $\circ$ in the Stratonovich form only with respect to the space x . In a recent paper (Pathiraja et al., 2021, Sec. 3), the gain function is defined and interpreted as a function of space x and the density. This is natural because the dependence upon time t comes because of the changes in density ($\bar{\rho}_t$) as the time evolves. Because the density is a stochastic process, it is argued that the appropriate interpretation of $\circ$ in the Stratonovich form should involve both space x and the density. Using such an interpretation, the Stratonovich form involves extra-terms that are solutions to accompanying Poisson equations.

4. Algorithms for gain function approximation

The exact gain K is a $d \times m$ matrix-valued function, where the j th column of K is the solution of the Poisson equation (12) for $j = 1, \dots, m$. For the ease of presentation, the exposition in this section is restricted to the scalar-valued observation setting, i.e. $m = 1$, so that K becomes a d -dimensional vector-valued function and the superscript j is dropped from the Poisson equation (12).

In practice, the Poisson equation must be solved numerically. The numerical gain function approximation problem is as follows:

input: samples $\{X^i : 1 \leq i \leq N\} \stackrel{\text{i.i.d.}}{\sim} \rho, h(\cdot)$
output: gain function $\{K^i : 1 \leq i \leq N\}$




Fig. 2. Gain function approximation problem in the feedback particle filter. The exact gain function $K(x) = \nabla\phi(x)$ where ϕ solves the Poisson equation (12). The numerical problem is to approximate $K^i = \nabla\phi(x)|_{x=X_t^i}$ using only the particles $\{X^i : 1 \leq i \leq N\}$ sampled from density ρ (depicted as shaded region). The dashed line indicates the constant gain approximation, where the gain function is approximated by its expected value according to (26).

where ρ is the (posterior) density and $K^i := K(X_t^i)$. The explicit dependence on time t is suppressed in this section. An illustration of the gain function approximation problem appears in Fig. 2.

4.1. Motivation and overview of approaches

The Poisson equation is a linear PDE. In order to motivate the various solution approaches, it is useful to first consider a finite-dimensional counterpart

$$Ax = b, \quad (22)$$

where A is a $n \times n$ (strictly) positive-definite symmetric matrix and the righthand-side b is a given $n \times 1$ vector. The problem is to compute the unknown $n \times 1$ vector x . For this purpose, the following equivalent formulations of the finite-dimensional problem are first introduced:

1. x is the solution of the weak form

$$y^T Ax = y^T b, \quad \forall y \in \mathbb{R}^n.$$

2. For some chosen positive ϵ , x is the solution to the fixed-point equation

$$x = e^{-\epsilon A} x + \int_0^\epsilon e^{-sA} b \, ds.$$

3. x is the solution of an optimization problem

$$x = \arg \min_{z \in \mathbb{R}^n} \frac{1}{2} z^T A z - z^T b.$$

When n is large, these formulations are useful to numerically approximate the solution of (22):

1. For each fixed $y \in \mathbb{R}^n$, the weak form is a single equation. By restricting y to a suitable low-dimensional subspace $S \subset \mathbb{R}^n$, the number of linear equations is reduced for the purposes of obtaining an approximate solution (possibly also in S).
2. The fixed-point equation is useful because $e^{-\epsilon A}$ is a strict contraction for $\epsilon > 0$ (because A is strictly positive-definite). So, a good initial guess for x can readily be improved by using the Banach iteration.
3. The optimization form is useful to develop alternate (e.g., search type) algorithms to obtain the solution.

With this background, we turn our attention to the Poisson equation (12) expressed succinctly as

$$-\Delta_\rho \phi = (h - \bar{h}),$$

where $\bar{h} := \int h(x)\rho(x)dx$ and $\Delta_\rho := \frac{1}{\rho} \nabla \cdot (\rho \nabla)$. The linear operator Δ_ρ is referred to as the probability weighted Laplacian. Functional analytic considerations require introduction of the function spaces:

$L^2(\rho)$ is the space of square integrable functions with respect to ρ with inner product $\langle f, g \rangle := \int f(x)g(x)\rho(x)dx$; $H^1(\rho)$ is the Hilbert space of functions in $L^2(\rho)$ whose first derivative, defined in the weak sense, is also in $L^2(\rho)$; and $H_0^1(\rho) = \{\psi \in H^1(\rho) \mid \int \psi(x)\rho(x)dx = 0\}$.

These definitions are important because $H_0^1(\rho)$ is the natural space for the solution ϕ of the Poisson equation (12). The operator $-\Delta_\rho$ is symmetric (self-adjoint) and positive definite because

$$-\langle f, \Delta_\rho g \rangle = \langle \nabla f, \nabla g \rangle = -\langle \Delta_\rho f, g \rangle, \quad \forall f, g \in H_0^1(\rho).$$

In the infinite-dimensional settings, one requires an additional technical condition—the Poincaré inequality (PI)—to conclude that the operator is in fact strictly positive-definite (Taghvaei, Mehta, & Meyn, 2020, Sec. 2.2). Assuming the PI holds, it is also readily shown that Δ_ρ^{-1} is well defined, i.e., a unique solution $\phi \in H_0^1(\rho)$ exists for any given $h \in L^2(\rho)$ (Yang et al., 2016, Thm. 2).

For the purposes of numerical approximation, entirely analogous to the finite-dimensional case, the following equivalent formulations of the Poisson equation are introduced:

1. ϕ is a solution of the weak form

$$\langle \nabla \psi, \nabla \phi \rangle = \langle \psi, h - \bar{h} \rangle \quad \forall \psi \in H_0^1(\rho). \quad (23)$$

2. For some chosen positive ϵ , ϕ is a solution of the fixed-point equation

$$\phi = e^{\epsilon \Delta_\rho} \phi + \int_0^\epsilon e^{s \Delta_\rho} (h - \bar{h}) ds. \quad (24)$$

The notation $e^{\epsilon \Delta_\rho}$ is used to denote the semigroup associated with Δ_ρ (Bakry, Gentil, & Ledoux, 2013). The semigroup is readily shown to be a Markov operator.

3. ϕ is the solution of an optimization problem

$$\phi = \arg \min_{f \in H_0^1(\rho)} \frac{1}{2} \langle \nabla f, \nabla f \rangle + \langle f, h - \bar{h} \rangle. \quad (25)$$

Each of the three formulations has been used to develop numerical algorithms for gain function approximation. A review of the resulting constructions appears in the following three subsections:

4.2. Galerkin and constant gain approximation

The starting point is the weak form (23). A relaxation is considered whereby $\psi \in S = \text{span}\{\psi_1, \dots, \psi_M\}$, a finite-dimensional subspace of $H_0^1(\rho)$. The functions $\psi_1, \dots, \psi_M$ need to be picked and are referred to as the basis functions. The resulting algorithm is referred to as the Galerkin algorithm (Yang et al., 2016, Sec 3.3). The algorithm is given in Algorithm 4.2.

Algorithm 1 Synthesis of the gain function: Galerkin approximation

Input: $\{X^i\}_{i=1}^N$, $\{h(X^i)\}_{i=1}^N$, basis functions $\{\psi_l(x)\}_{l=1}^L$.

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

- 1: Calculate $h^{(N)} = \frac{1}{N} \sum_{i=1}^N h(X_t^i)$.
 - 2: Calculate $b_k = \frac{1}{N} \sum_{i=1}^N (h(X_t^i) - h^{(N)}) \psi_k(X_t^i)$.
 - 3: Calculate $A_{kl} = \frac{1}{N} \sum_{i=1}^N \nabla \psi_l(X_t^i)^T \nabla \psi_k(X_t^i)$.
 - 4: Solve the linear matrix equation $A\kappa = b$ for κ , where $A = [A_{kl}]$ and $b = [b_k]$.
 - 5: $K^i = \sum_{l=1}^L \kappa_l \nabla \psi_l(X_t^i)$.
-

The most important special case of the Galerkin algorithm is obtained upon picking S to be the subspace spanned by the d coordinate functions $\{x_1, x_2, \dots, x_d\}$. The special case yields the *constant gain approximation* of the gain K as its expected value. Remarkably, the

Algorithm 2 Synthesis of the gain function: constant gain approximation

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

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

- 1: Calculate $\hat{h}^{(N)} = \frac{1}{N} \sum_{i=1}^N h(X^i)$.
- 2: $K^i = \frac{1}{N} \sum_{j=1}^N X_t^j (h(X_t^j) - \hat{h}^{(N)})$

expected value admits a closed-form expression which is then readily approximated empirically using the particles:

$$\begin{aligned} K^{(\text{const. appr.})} &:= \int \nabla \phi(x) \rho(x) dx = \int (h(x) - \bar{h}) x \rho(x) dx \\ &\approx \frac{1}{N} \sum_{i=1}^N (h(X^i) - \hat{h}^{(N)}) X^i, \end{aligned} \quad (26)$$

where $\hat{h}^{(N)} := N^{-1} \sum_i h(X^i)$. (See Fig. 2 for an illustration of the constant gain approximation.) With the constant gain approximation, the FPF algorithm is a nonlinear EnKF algorithm (Taghvaei et al., 2018). While its derivation starting from an FPF is novel, the formula (26) has been used as a heuristic in the EnKF literature (Bergemann & Reich, 2012; Evensen, 2006).

The main issue with the Galerkin approximation is that it is in general very difficult to pick the basis functions. There have been a number of studies to refine and improve upon this formula (Berntorp, 2018; Berntorp & Grover, 2016; Matsuura, Ohata, Nakakuki, & Hirokawa, 2016; Radhakrishnan, Devraj, & Meyn, 2016; Radhakrishnan & Meyn, 2018; Yang et al., 2016; Yang, Mehta, & Meyn, 2013). In the following two subsections, we describe two approximations which appear to be more promising approaches in general settings.

4.3. Diffusion map-based algorithm

The starting point is the fixed-point equation (24) based on the Markov semigroup $e^{\epsilon \Delta_\rho}$. For small values of ϵ , there is a well known approximation of $e^{\epsilon \Delta_\rho}$ in terms of the so-called diffusion map (which too is a Markov operator):

$$(T_\epsilon f)(x) := \frac{1}{n_\epsilon(x)} \int_{\mathbb{R}^d} \frac{g_\epsilon(|x-y|)}{\sqrt{\int g_\epsilon(|y-z|) \rho(z) dz}} f(y) \rho(y) dy, \quad (27)$$

where $g_\epsilon(z) := e^{-\frac{z^2}{4\epsilon}}$ is the Gaussian kernel in $\mathbb{R}$ and $n_\epsilon(x)$ is the normalization factor chosen so that $\int (T_\epsilon 1)(x) dx = 1$ (Coifman & Lafon, 2006). A representative approximation result is as follows:

Proposition 4.1 (Prop. 3.4 in Taghvaei et al., 2020). Let $n \in \mathbb{N}$, $t_0 < \infty$, and $t \in (0, t_0)$ with $\epsilon = \frac{t}{n}$. Then, for all functions f such that $f, \nabla f \in L^4(\rho)$:

$$\|(T_\epsilon^n - e^{t\Delta_\rho})f\|_{L^2(\rho)} \leq \frac{t^{\frac{3}{2}}}{n} C(\|f\|_{L^4(\rho)} + \|\nabla f\|_{L^4(\rho)}),$$

where the constant C depends only on t_0 and ρ .

Because the diffusion map (27) is defined using Gaussian kernels, its empirical approximation is straightforward:

$$(T_\epsilon^{(N)} f)(x) = \frac{1}{n_\epsilon^{(N)}(x)} \sum_{i=1}^N \frac{g_\epsilon(|x-X^i|)}{\sqrt{\sum_{j=1}^N g_\epsilon(|X^i-X^j|)}} f(X^i),$$

where $n_\epsilon^{(N)}(x)$ is the normalization factor. The nature of the approximation is as follows:

Proposition 4.2 (Prop. 3.5 in Taghvaei et al., 2020). Consider the diffusion map kernel T_ϵ and its empirical approximation $\{T_\epsilon^{(N)}\}_{N \in \mathbb{N}}$. Then for any bounded continuous function $f \in C_b(\mathbb{R}^d)$:

Algorithm 3 Synthesis of the gain function: diffusion map-based algorithm

Input: $\{X^i\}_{i=1}^N, \{h(X^i)\}_{i=1}^N, \Phi_{\text{prev}}, \epsilon, L$.

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

- 1: Calculate $g_{ij} := e^{-\frac{|X^i-X^j|^2}{4\epsilon}}$ for $i, j = 1$ to N .
- 2: Calculate $k_{ij} := \frac{g_{ij}}{\sqrt{\sum_l g_{il}} \sqrt{\sum_l g_{jl}}}$ for $i, j = 1$ to N .
- 3: Calculate $d_i = \sum_j k_{ij}$ for $i = 1$ to N .
- 4: Calculate $\tau_{ij} := \frac{k_{ij}}{d_i}$ for $i, j = 1$ to N .
- 5: Calculate $\pi_i = \frac{d_i}{\sum_j d_j}$ for $i = 1$ to N .
- 6: Calculate $\hat{h} = \sum_{i=1}^N \pi_i h(X^i)$.
- 7: Initialize $\Phi = \Phi_{\text{prev}}$.
- 8: **for** $t = 1$ to L **do**
- 9: $\Phi_i = \sum_{j=1}^N \tau_{ij} \Phi_j + \epsilon(h - \hat{h})$ for $i = 1$ to N .
- 10: **end for**
- 11: Calculate $r_i = \Phi_i + \epsilon h_i$ for $i = 1$ to N .
- 12: Calculate $s_{ij} = \frac{1}{\epsilon} \tau_{ij} (r_j - \sum_{k=1}^N \tau_{ik} r_k)$ for $i, j = 1$ to N .
- 13: Calculate $K^i = \sum_j s_{ij} X^j$ for $i = 1$ to N .

1. (Almost sure convergence) For all $x \in \mathbb{R}^d$

$$\lim_{N \rightarrow \infty} (T_\epsilon^{(N)} f)(x) = (T_\epsilon f)(x), \quad \text{a.s.}$$

2. (Convergence rate) For any $\delta \in (0, 1)$, in the asymptotic limit as $N \rightarrow \infty$,

$$\int |(T_\epsilon^{(N)} f)(x) - (T_\epsilon f)(x)|^2 \rho(x) dx \leq O\left(\frac{\log(\frac{N}{\delta})}{N \epsilon^d}\right),$$

with probability higher than $1 - \delta$.

With these approximations, the fixed-point equation (24) is approximated in two steps:

1. The semigroup $e^{\epsilon \Delta_\rho}$ is approximated by the diffusion map T_ϵ :

$$(\text{step 1}) \quad \phi_\epsilon = T_\epsilon \phi_\epsilon + \epsilon(h - \bar{h}_\epsilon), \quad (28a)$$

$$\text{where } \bar{h}_\epsilon = \int h(x) \rho^{(\epsilon)}(x) dx \text{ with } \rho^{(\epsilon)}(x) = \frac{n_\epsilon(x) \rho(x)}{\int n_\epsilon(x) \rho(x) dx}.$$

2. T_ϵ is approximated by its empirical approximation $T_\epsilon^{(N)}$:

$$(\text{step 2}) \quad \phi_\epsilon^{(N)} = T_\epsilon^{(N)} \phi_\epsilon^{(N)} + \epsilon(h - \bar{h}_\epsilon^{(N)}), \quad (28b)$$

$$\text{where } \bar{h}_\epsilon^{(N)} = \int h(x) \rho_\epsilon^{(N)}(x) dx \text{ with } \rho_\epsilon^{(N)}(x) = \frac{\sum_{i=1}^N n_\epsilon(x) \delta_{X^i}}{\sum_{i=1}^N n_\epsilon(X^i)}.$$

Based on the finite-dimensional fixed-point equation (28b), an algorithm for gain function approximation is given in Algorithm 3. In the algorithm, the gain function is approximated by the formula

$$K_\epsilon^{(N)} = \nabla [T_\epsilon^{(N)} \phi_\epsilon^{(N)} + \epsilon T_\epsilon^{(N)} (h - \bar{h}_\epsilon^{(N)})].$$

There are alternative ways to approximate the gain function in terms of $\phi_\epsilon^{(N)}$. While these solutions have the same asymptotic in the limit as $\epsilon \rightarrow 0$, they behave differently when ϵ is large. The specific approximation selected here does not require derivative of the observation function and converges to the constant gain approximation as ϵ becomes large (Taghvaei et al., 2020, Remark 4.8).

Error analysis. The error in diffusion map approximation comes from two sources:

1. The bias error due to the diffusion map approximation of the semigroup (step 1); and
2. The variance error due to empirical approximation in terms of particles (step 2).




Fig. 3. Bias variance trade-off in the diffusion map-based gain function approximation algorithm: (a) Gain function computed for different values of ϵ with $N = 200$ particles. The dashed line is the constant gain solution (26). As ϵ gets larger, the diffusion map gain converges to the constant gain. (b) Plot of the MSE as a function of ϵ . The shaded area in the background of part (a) is the density ρ which is taken as sum of two Gaussians $\mathcal{N}(-1, \sigma^2)$ and $\mathcal{N}(+1, \sigma^2)$ with $\sigma^2 = 0.2$. The exact gain function $K(x)$ is computed for $h(x) = x$ by using an (exact) integral formula for the solution (Taghvaei et al., 2020, Eq. 4.6). In part (b), the MSE is computed as an empirical approximation of the lefthand-side of (29) by averaging over 1000 simulation runs.




Fig. 4. Bias-variance trade-off as a function of (a) the state dimension $d \in \{1, 2, 5, 10\}$ (for a fixed $N = 1000$); and (b) the number of particles $N \in \{100, 200, 500, 1000\}$ (for a fixed $d = 1$). In the vector case, $\rho(x) = \rho_b(x_1) \prod_{n=2}^d \rho_g(x_n)$ where ρ_b is the bimodal density (same as in Fig. 3) and ρ_g is the Gaussian density.

The error is analyzed in Taghvaei et al. (2020) where the following result is proved:

Proposition 4.3 (Thm. 4.3 and 4.4 in Taghvaei et al., 2020). Consider the fixed-point formulation of the Poisson equation (24), its diffusion-map approximation (28a), and its empirical approximation (28b).

- For each fixed $\epsilon > 0$, there exists a unique solution to (28a) with a uniform bound $\|\phi_\epsilon\|_{L^2(\rho_\epsilon)} \leq C\|h\|_{L^2(\rho_\epsilon)}$. In the asymptotic limit as $\epsilon \rightarrow 0$

$$\|\phi_\epsilon - \phi\|_{L^2(\rho_\epsilon)} \leq O(\epsilon).$$

- The operator $T_\epsilon^{(N)}$ is a strict contraction on $L_0^2(\rho_\epsilon^{(N)})$ and the fixed-point equation (28b) admits a unique solution. The approximate solution $\phi_\epsilon^{(N)}$ converges to the kernel solution ϕ_ϵ

$$\lim_{N \rightarrow \infty} \|\phi_\epsilon^{(N)} - \phi_\epsilon\|_{L^\infty(\Omega)} = 0, \quad \text{a.s.}$$

The following diagram illustrates the convergence and the respective types of errors:

$$\phi_\epsilon^{(N)} \xrightarrow[N \uparrow \infty]{\text{(variance)}} \phi_\epsilon \xrightarrow[\epsilon \downarrow 0]{\text{(bias)}} \phi.$$

A quantitative bound on the mean-squared error (MSE) is obtained in the asymptotic limit as $\epsilon \downarrow 0$ and $N \rightarrow \infty$ as follows:

$$\underbrace{\left(\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N |K^i - \nabla \phi(X^i)|^2 \right] \right)}_{\text{MSE}} \leq \underbrace{O(\epsilon^2)}_{\text{bias}} + \underbrace{O\left(\frac{1}{\epsilon^{(2+d)} N}\right)}_{\text{variance}}. \quad (29)$$

where $\{K^i\}_{i=1}^N$ is computed from the Algorithm 3 and $\nabla \phi$ is the exact gain function from solving the Poisson equation (12). The error due to bias converges to zero as $\epsilon \rightarrow 0$ and the error due to variance converges to zero as $N \rightarrow \infty$. There is trade-off between the two errors: To reduce bias, one must reduce ϵ . However, for any fixed value of N , one can reduce ϵ only up to a point where the variance starts increasing. The bias-variance trade-off is illustrated with the aid of a scalar ($d = 1$) example in Fig. 3: If ϵ is large, the error due to bias dominates, while if ϵ is small, the error due to variance dominates. A numerical illustration of scalings with N and d appears in Fig. 4. Additional details on both these examples can be found in Taghvaei et al. (2020, Sec. 5).

Remark 4.4 (Relationship to the Constant Gain Formula (26)). There is a remarkable and somewhat unexpected relationship between the diffusion map and the constant gain approximation (Taghvaei et al., 2020, Prop. 4.7). In particular, in the limit as $\epsilon \rightarrow \infty$, the diffusion map gain converges to the constant gain (26). This suggests a systematic procedure to improve upon the constant gain by de-tuning the value of ϵ away from the $[\epsilon = \infty]$ limit. For any fixed N , a finite value of ϵ is chosen to minimize the MSE according to the bias variance trade-off. Based on this, a rule of thumb for choosing the ϵ value appears in Taghvaei et al. (2020, Remark 5.1).

Remark 4.5 (Analysis of FPF with Diffusion Map Approximation). An analysis of the finite- N FPF using the diffusion map approximation appears in Pathiraja and Stannat (2021). Under mild technical conditions on the drift $a(\cdot)$, $\sigma(\cdot)$, $h(\cdot)$, it is shown that the finite- N FPF is

Table 1
Applications and evaluation of the feedback particle filter.

Authors	Applications of FPF	Reference	Year
del Moral and Horton	Quantum harmonic oscillators	del Moral and Horton (2021)	2021
Wang et. al.	Unmanned aerial vehicle tracking	Wang, Wang, and Cui (2021)	2021
Su et. al.	Soil estimation	Su, Zhou, Yu, Yuan, and Fu (2019)	2021
Kumar and Mishra	Marine applications	Zheng, Yu, and Yuan (2019)	2019
Berntorp and Grover	Satellite tracking and re-entry	Berntorp (2015)	2015
Surace et. al.	Evaluation and comparison of FPF	Surace et al. (2019)	2017
Stano	Hopper-dredger model	Stano (2018), Stano et al. (2014)	2014
Matsuura et. al.	Target state estimation	Matsuura et al. (2016)	2016
Kutschireiter et. al.	Neuronal dynamics	Kutschireiter, Surace, Sprekeler, and Pfister (2017)	2016
Tilton et. al.	Coupled oscillators	Tilton et al. (2012)	2013
Tilton et. al.	Marine estimation	Tilton et al. (2013)	2013

well-posed, i.e., a strong solution exists for all time t (Pathiraja & Stannat, 2021, Thm. 1.1). Based on a propagation of chaos type analysis, convergence estimates are derived to relate the finite- N system to its mean-field limit (Pathiraja & Stannat, 2021, Thm. 1.2). These estimates are shown to hold up to a certain stopping time. For arbitrary time t , well-posedness and convergence remains an open problem.

4.4. Variational approximation

The starting point is the variational form (25). The objective function is denoted by $J(f)$ with its empirical approximation is obtained as

$$J^{(N)}(f) := \frac{1}{N} \sum_{i=1}^N \frac{1}{2} |\nabla f(X^i)|^2 - f(X^i)(h(X^i) - h^{(N)})$$

The problem of minimizing the empirical approximation over all functions is ill-posed: the minimum is unbounded and minimizer does not exist. (Abstractly, this is because the empirical probability distribution does not satisfy the Poincaré inequality.) Therefore, we consider

$$\min_{f_\theta \in \mathcal{F}_\theta} J^{(N)}(f_\theta)$$

where $\mathcal{F}_\theta$ is a parameterized class of functions. A function in the class $\mathcal{F}_\theta$ is denoted by $f_\theta(x)$ or $f(x; \theta)$ where $\theta \in \Theta$ is the parameter, and Θ is the parameter set. The two main examples are as follows:

1. $\mathcal{F}_\theta = \{\sum_{j=1}^M \theta_j \psi_j; \psi_j \in H_0^1, \theta_j \in \mathbb{R} \text{ for } j = 1, \dots, M\}$ is a linear combination of selected basis functions. With a linear parametrization, the solution of the empirical optimization problem is given by the Galerkin algorithm (Yang et al., 2016, Remark 5).
2. $\mathcal{F}_\theta$ is a neural network where the parameters θ are the weights in the network.

In practice, it is not possible to solve the optimization problem exactly, but up to some optimization gap. In particular, let $\phi_\theta^{(N)}$ be the output of an optimization algorithm that solves the problem up to ϵ error, i.e.,

$$J(\phi_\theta^{(N)}) \leq \min_{f \in H_0^1} J(f) + \epsilon.$$

The good news is that it is possible to upper-bound the error in approximating the gain function in terms of this optimization gap.

Proposition 4.6 (Prop. 1 in Olmez, Taghvaei, & Mehta, 2020). Let $K_\theta^{(N)} = \nabla \phi_\theta^{(N)}$ where $\phi_\theta^{(N)}$ is the output of an optimization algorithm that solves the minimization objective $J(f)$ with ϵ optimality gap. Then

$$\|K_\theta^{(N)} - K\|_{L^2}^2 \leq 2\epsilon,$$

where $K = \nabla \phi$ is the exact gain function.

The optimization gap ϵ depends on the selected parametrization $\mathcal{F}_\theta$, number of particles N , and the iteration number of the employed optimization algorithm. Its characterization and analysis is open and

the subject of ongoing work. In general, such analysis falls under the framework of statistical learning theory (Anthony et al., 1999; Shalev-Shwartz & Ben-David, 2014).

The numerical results using this approach are depicted in Fig. 5. These results are for the bimodal example introduced in Fig. 3. The gain function is parameterized using a two-layer residual NN with 32 neurons per layer. The Adam algorithm is used to learn the parameters of the NN. Additional details on the numerics can be found in Olmez et al. (2020).

4.5. Numerical evaluation of FPF

Numerical evaluations of the FPF algorithm, and comparisons with the nonlinear extensions of the Kalman filter and conventional particle filters, have been subject of several publications (some of these studies are tabulated in the Table 1). Notable amongst these is the early work of Berntorp (2015) who both extended the algorithm and applied it to two highly nonlinear applications in aerospace, namely, the *re-entry* and *two-body* problems. Another notable early work is Stano et al. (2014) on the application of estimating soil-dependent time-varying parameters of the hopper sedimentation model. The study includes extensive comparisons with the conventional particle filters. While these studies report favorable accuracy and computational cost, the application of FPF to truly high-dimensional and nonlinear problems remains still open. In particular, beyond the toy examples, we do not know of any application where the diffusion map approximation has been applied.

Remark 4.7. Curse of dimensionality (CoD) is one of the main issues motivating the recent work on particle filters (Bengtsson et al., 2008; Beskos, Crisan, Jasra, & Whiteley, 2014; Bickel et al., 2008; Rebeschini et al., 2015). The analysis presented in Section 3.4 helps show that, at least in the linear Gaussian settings of the problem, the FPF/EnKF algorithm does not suffer from CoD. This is because the Poisson equation admits an explicit solution in this case. Because FPF is an exact algorithm, if/when CoD can be avoided in the nonlinear case really depends on the quality of the gain function approximation. The bias-variance analysis of the diffusion map algorithm, presented in Section 4.3, is helpful to see some of the tradeoffs. The analysis suggests that to avoid CoD one must take advantage of (i) the underlying regularity of the gain function (e.g., constant in the linear Gaussian case), and/or (ii) inherent low-dimensional structure in the problem (e.g., approximation of posterior density in a low-dimensional manifold where a good diffusion map approximation can be obtained). One promising avenue is the variational gain function approximation using neural networks, as described in Section 4.4. It remains to be seen whether some of the outstanding successes of neural networks in other fields can be replicated to avoid CoD in the particle filters.





Fig. 5. Results of the variational gain function approximation using a neural network parametrization: Plot of (a) the gain function; and (b) the optimization gap as the number of iterations of the Adam algorithm. The problem setup is the same as Fig. 3.

5. Optimal transport theory

In this section, we describe a systematic procedure to construct the exact mean-field process $\bar{X}$ introduced as step 1 in (10). The first aspect to note is that while the FPF (11) provides an explicit formula for u and K , the formula is not unique: One can interpret (10) as transporting the prior density p_0 at time $t = 0$ to the posterior density p_t at time t . Clearly, there are infinitely many maps that transport one density into another. This suggests that there are infinitely many choices of control laws that all lead to exact filters. This is not surprising: The exactness condition specifies only the marginal density at times t , which is not enough to uniquely identify a stochastic process, e.g., the joint density at two time instants has not been specified.

In the following, we first discuss the non-uniqueness issue for the simpler linear Gaussian model. The non-uniqueness naturally motivates optimal transport ideas to uniquely solve for u and K . This is the subject of the remainder of this section to derive the feedback control law for the FPF (11).

5.1. Non-uniqueness issue in linear-Gaussian setting

Consider the linear Gaussian FPF (14) for the mean-field process $\{\bar{X}_t\}_{t \geq 0}$. The conditional mean and variance of $\bar{X}_t$ are denoted by $\bar{m}_t$ and $\bar{\Sigma}_t$, respectively. The conditional mean evolves according to

$$d\bar{m}_t = A\bar{m}_t dt + \bar{K}_t(dZ_t - H\bar{m}_t dt),$$

where $\bar{K}_t := \bar{\Sigma}_t H^T$. Define an error process $\xi_t := \bar{X}_t - \bar{m}_t$. Its equation is given by

$$d\xi_t = (A - \frac{1}{2}\bar{\Sigma}_t H^T H)\xi_t + \sigma_B d\bar{B}_t.$$

This is a linear system and therefore the variance of ξ_t , which equals $\bar{\Sigma}_t$ (by definition), evolves according to the Lyapunov equation

$$\begin{aligned} \frac{d}{dt} \bar{\Sigma}_t &= (A - \frac{1}{2}\bar{\Sigma}_t H^T H)\bar{\Sigma}_t + \bar{\Sigma}_t (A - \frac{1}{2}\bar{\Sigma}_t H^T H)^T + \Sigma_B \\ &= \text{Ricc}(\bar{\Sigma}_t). \end{aligned}$$

The derivation helps show that the equations for the mean and variance are identical to the Kalman filter equations, (8a) and (8b), respectively, and thus proves the exactness property of the linear FPF (14).

These arguments suggest the following general procedure to construct an exact $\bar{X}$ process: Express $\bar{X}_t$ as a sum of two terms:

$$\bar{X}_t = \bar{m}_t + \xi_t, \quad t \geq 0,$$

where $\bar{m}_t$ evolves according to (8a) and the evolution of ξ_t is defined by the SDE:

$$d\xi_t = G_t \xi_t dt + \sigma_t d\bar{B}_t + \sigma'_t d\bar{W}_t,$$

where $\{\bar{W}_t\}_{t \geq 0}$ and $\{\bar{B}_t\}_{t \geq 0}$ are independent copies of the measurement noise $\{W_t\}_{t \geq 0}$ and the process noise $\{B_t\}_{t \geq 0}$, respectively, and G_t , σ_t , and σ'_t satisfy the matrix equation (for each time)

$$G_t \bar{\Sigma}_t + \bar{\Sigma}_t G_t^T + \sigma_t \sigma_t^T + \sigma'_t (\sigma'_t)^T = \text{Ricc}(\bar{\Sigma}_t), \quad t \geq 0. \quad (30)$$

By construction, the equation for the variance is given by the Riccati equation (8b). The result is summarized in the following Proposition:

Proposition 5.1 (Prop. 1 in Taghvaei, Mehta, & Georgiou, 2022). Consider the linear-Gaussian filtering problem (7) and the following family of the mean-field processes

$$\begin{aligned} d\bar{X}_t &= A\bar{m}_t dt + \bar{K}_t(dZ_t - H\bar{m}_t dt) \\ &\quad + G_t(\bar{X}_t - \bar{m}_t)dt + \sigma_t d\bar{B}_t + \sigma'_t d\bar{W}_t, \quad \bar{X}_0 \sim N(m_0, \Sigma_0), \end{aligned}$$

where G_t , σ_t , and σ'_t satisfy the consistency condition (30). Then, $\bar{X}_t$ is exact, i.e. the density of $\bar{X}_t$ is Gaussian $N(\bar{m}_t, \bar{\Sigma}_t)$ where $\bar{m}_t$ and $\bar{\Sigma}_t$ solve the Kalman filter equations, (8a) and (8b), respectively.

In general, with different choices of σ_t and σ'_t , there are infinitely many solutions for (30). Below, we describe three solutions that lead to three established form of EnKF and linear FPF:

1. EnKF with perturbed observation (Reich, 2011, Eq. (27)):

$$G_t = A - \bar{\Sigma}_t H^T H, \quad \sigma_t = \sigma_B, \quad \sigma'_t = \bar{\Sigma}_t H^T.$$

2. Stochastic linear FPF (Yang et al., 2016, Eq. (26)) or square-root form of the EnKF (Bergemann & Reich, 2012, Eq. (3.3)) :

$$G_t = A - \frac{1}{2}\bar{\Sigma}_t H^T H, \quad \sigma_t = \sigma_B, \quad \sigma'_t = 0.$$

3. Deterministic linear FPF (Taghvaei & Mehta, 2016, Eq. (15)) (de Wiljes et al., 2018, Eq. (82)):

$$G_t = A - \frac{1}{2}\bar{\Sigma}_t H^T H + \frac{1}{2}\Sigma_B \bar{\Sigma}_t^{-1}, \quad \sigma_t = 0, \quad \sigma'_t = 0.$$

Fix σ_t, σ'_t . Then given any particular solution G_t of (30), one can construct a family of solutions $G_t + \bar{\Sigma}_t^{-1} \Omega_t$, where Ω_t is any arbitrary skew-symmetric matrix (Taghvaei & Mehta, 2020, Sec. III-B). For the linear Gaussian problem, the non-uniqueness issue is well known in literature: The two forms of EnKF, the perturbed observation form Reich (2011) and the square-root form Bergemann and Reich (2012) are standard. A homotopy of exact deterministic and stochastic EnKFs is given in Kim, Taghvaei, and Mehta (2018). An explanation for the non-uniqueness in terms of the Gauge transformation appears in Abedi and Surace (2019). An extension to the case with correlated noise appears in Kang, Chen, Tao, and Yau (2022).

Given the non-uniqueness issue, a natural question is how to identify a unique $\bar{X}$ process? For this purpose, optimal transport theory is

described in the following Section 5.2. For the linear Gaussian case, the theory is used to derive the following optimal transport form of the linear FPF (see Taghvaei & Mehta, 2016, 2020 for details):

$$\begin{aligned} d\bar{X}_t = & A\bar{X}_t dt + \frac{1}{2} \Sigma_B \bar{\Sigma}_t^{-1} (\bar{X}_t - \bar{m}_t) dt \\ & + \frac{1}{2} \bar{K}_t (dZ_t - \frac{H\bar{X}_t + H\bar{m}_t}{2} dt) + \Omega_t \bar{\Sigma}_t^{-1} (\bar{X}_t - \bar{m}_t) dt, \end{aligned} \quad (31)$$

where $\Omega_t = \Omega_t^{\text{OPT}}$ is a specific skew-symmetric matrix. The optimal transport FPF (31) is exact and has two differences compared to the linear FPF (14):

1. The stochastic term $\sigma_B d\bar{B}_t$ is replaced with the deterministic term $\frac{1}{2} \Sigma_B \bar{\Sigma}_t^{-1} (\bar{X}_t - \bar{m}_t) dt$. Given a Gaussian prior, the two terms yield the same posterior. However, in a finite- N implementation, the stochastic term serves to introduce an additional error of order $O(\frac{1}{\sqrt{N}})$ (Taghvaei & Mehta, 2018, Prop. 4).
2. The SDE (31) has an extra term involving the skew-symmetric matrix Ω_t . The extra term does not effect the posterior, i.e., $\bar{X}$ is exact for all skew-symmetric choices of Ω_t . The specific optimal choice $\Omega_t = \Omega_t^{\text{OPT}}$ serves to pick the symmetric solution G_t of the consistency equation (30). For the scalar ($d = 1$) case, the skew-symmetric term is zero. Therefore, in the scalar case, the update formula in the linear FPF (14) is optimal. In the vector case, it is optimal iff $\Omega_t^{\text{OPT}} \equiv 0$.

5.2. FPF formula

In this section, we provide a justification for the feedback control formula in the FPF (11). It is helpful to begin with the simpler deterministic case.

5.2.1. Deterministic path

Let $\mathcal{P}_2(\mathbb{R}^d)$ be the space of everywhere positive probability densities on $\mathbb{R}^d$ with finite second moment. Given a smooth path $\{p_t \in \mathcal{P}_2(\mathbb{R}^d) : t \geq 0\}$ the problem is to construct a stochastic process $\{\bar{X}_t\}_{t \geq 0}$ such that the probability density of $\bar{X}_t$, denoted as $\bar{p}_t$, equals p_t for all $t \geq 0$. The exactness condition is expressed as

$$\bar{p}_t = p_t, \quad \forall t \geq 0. \quad (32)$$

As has already been noted, there are infinitely many stochastic processes that satisfy the exactness condition. A unique choice is made by prescribing an additional optimality criterion based on the optimal transport theory.

To make these considerations concrete, assume that the given path $\{p_t\}_{t \geq 0}$ evolves according to the PDE

$$\frac{\partial p_t}{\partial t} = \mathcal{V}(p_t), \quad t > 0,$$

where $\mathcal{V}(\cdot)$ is an operator (e.g., the Laplacian) that acts on probability densities. (This necessarily restricts the operator $\mathcal{V}$, e.g., $\int \mathcal{V}(\rho)(x) dx = 0$ for all $\rho \in \mathcal{P}_2(\mathbb{R}^d)$.) The following model is assumed for the process $\{\bar{X}_t\}_{t \geq 0}$:

$$\frac{d}{dt} \bar{X}_t = u_t(\bar{X}_t), \quad \bar{X}_0 \sim p_0, \quad (33)$$

where $u_t(\cdot)$ is a control law that needs to be designed. From the continuity equation, the exactness condition (32) is satisfied if

$$-\nabla \cdot (\bar{p}_t u_t) = \mathcal{V}(\bar{p}_t), \quad \forall t > 0. \quad (34)$$

The non-uniqueness issue is now readily seen: The first-order PDE (34) admits infinitely many solutions. A unique solution $u_t(\cdot)$ is picked by minimizing the transportation cost from $\bar{X}_t$ to $\bar{X}_{t+\Delta t}$ in the limit as $\Delta t \rightarrow 0$. The L^2 -Wasserstein cost is particularly convenient because

$$\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t^2} \mathbb{E}[|X_{t+\Delta t} - X_t|^2] = \int_{\mathbb{R}^d} |u_t(x)|^2 \bar{p}_t(x) dx.$$

Therefore, for each fixed t , the control law $u_t(\cdot)$ is obtained by solving the constrained optimization problem

$$\min_{u_t(\cdot)} \int_{\mathbb{R}^d} |u_t(x)|^2 \bar{p}_t(x) dx, \quad \text{s.t.} \quad -\nabla \cdot (\bar{p}_t u_t) = \mathcal{V}(\bar{p}_t).$$

By a standard calculus of variation argument, the optimal solution is obtained as $u_t^* = \nabla \phi_t$ where ϕ_t solves the Poisson equation $-\nabla \cdot (\bar{p}_t \nabla \phi_t) = \mathcal{V}(\bar{p}_t)$. The resulting stochastic process $\bar{X}$ is defined by

$$\frac{d\bar{X}_t}{dt} = \nabla \phi_t(\bar{X}_t), \quad \bar{X}_0 \sim p_0,$$

ϕ_t solves the PDE $-\nabla \cdot (\bar{p}_t \nabla \phi_t) = \mathcal{V}(\bar{p}_t)$.

The process is exact by construction.

Example 5.2. Suppose the given path is a solution of the heat equation $\frac{\partial p_t}{\partial t} = \Delta p_t$ ($\mathcal{V}(\cdot)$ is the Laplacian). The solution of the Poisson equation is easily obtained as $\phi_t = \log(\bar{p}_t)$. The optimal transport process then evolves according to

$$\frac{d}{dt} \bar{X}_t = -\nabla \log(\bar{p}_t(\bar{X}_t)), \quad \bar{X}_0 \sim p_0. \quad (35a)$$

This process should be compared to the well known example

$$dX_t = dB_t, \quad X_0 \sim p_0, \quad (35b)$$

where $\{B_t\}_{t \geq 0}$ is a W.P. The density for X_t also solves the heat equation. In the language of optimal transportation theory, the coupling defining (35a) is deterministic while it is stochastic in (35b).

5.2.2. Stochastic path

In the filtering problem, the path of the posterior probability density is stochastic (because it depends upon the random observations $\{Z_t\}_{t \geq 0}$). Therefore, the preceding discussion is not directly applicable. Suppose the stochastic path $\{p_t\}_{t \geq 0}$ is governed by a stochastic PDE

$$dp_t = \mathcal{H}(p_t) dt,$$

where $\mathcal{H}(\cdot)$ is an operator that acts on probability densities and $\{I_t : t \geq 0\}$ is a W.P.

Consider the following SDE model:

$$d\bar{X}_t = u_t(\bar{X}_t) dt + K_t(\bar{X}_t) dI_t, \quad \bar{X}_0 \sim p_0$$

where, compared to the deterministic model (33), an additional stochastic term is now included. The problem is to identify control laws $u_t(\cdot)$ and $K_t(\cdot)$ such that the conditional density of $\bar{X}_t$ equals p_t . Upon writing the evolution equation for the conditional density of $\bar{X}_t$ (Yang et al., 2016, Prop. 1), the exactness condition is formally satisfied by all such $u_t(\cdot)$ and $K_t(\cdot)$ that solve

$$-\nabla \cdot (\bar{p}_t K_t) = \mathcal{H}(\bar{p}_t), \quad (36a)$$

$$-\nabla \cdot (\bar{p}_t u_t) + \frac{1}{2} (\nabla \cdot (\bar{p}_t K_t) K_t + \bar{p}_t K_t \nabla K_t) = 0. \quad (36b)$$

These equations are the stochastic counterpart of (34), and as with (34), their solution is not unique.

The unique solution is obtained by requiring that the coupling from $\bar{X}_t$ and $\bar{X}_{t+\Delta t}$ is optimal in the limit as $\Delta t \rightarrow 0$. In contrast to the deterministic setting, the leading term in the transportation cost $\mathbb{E}[|\bar{X}_{t+\Delta t} - \bar{X}_t|^2]$ is $O(\Delta t)$ whereby

$$\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \mathbb{E}[|\bar{X}_{t+\Delta t} - \bar{X}_t|^2] = \int_{\mathbb{R}^d} |K_t(x)|^2 \bar{p}_t(x) dx.$$

Therefore, for each fixed t , the control law $K_t(\cdot)$ is obtained by solving the constrained optimization problem

$$\min_{K_t(\cdot)} \int_{\mathbb{R}^d} |K_t(x)|^2 \bar{p}_t(x) dx, \quad \text{s.t.} \quad -\nabla \cdot (\bar{p}_t K_t) = \mathcal{H}_t(\bar{p}_t).$$

As before, the optimal solution is given by $K_t^* = \nabla \phi_t$ where ϕ_t solves the second-order PDE $-\nabla \cdot (\bar{p}_t \nabla \phi_t) = \mathcal{H}(\bar{p}_t)$.

It remains to identify the control law $u_t(\cdot)$. For this purpose, the second-order term in the infinitesimal Wasserstein cost is used:

$$\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t^2} \left(\mathbb{E}[|\bar{X}_{t+\Delta t} - \bar{X}_t|^2] - \Delta t \int_{\mathbb{R}^d} |K_t^*(x)|^2 \bar{p}_t(x) dx \right) = \int_{\mathbb{R}^d} |u_t(x)|^2 \bar{p}_t(x) dx.$$

The righthand-side is minimized subject to the constraint (36b). Remarkably, the optimal solution is obtained in closed form as

$$u_t^* = -\frac{1}{2\bar{p}_t} \mathcal{H}(\bar{p}_t) \nabla \phi_t + \frac{1}{2} \nabla^2 \phi_t \nabla \phi_t + \xi_t,$$

where ξ_t is the (unique such) divergence free vector field (i.e., $\nabla \cdot (p_t \xi_t) = 0$) such that u_t^* is of a gradient form. That (36b) can be solved in an explicit manner was a major surprise at the time of its discovery (Yang et al., 2011b; Yang, Mehta, & Meyn, 2013). The resulting optimal transport process is

$$d\bar{X}_t = \nabla \phi_t(\bar{X}_t) \circ (dI_t - \frac{1}{2\bar{p}_t} \mathcal{H}(\bar{p}_t) dt) + \xi_t(\bar{X}_t) dt, \quad \bar{X}_0 \sim p_0. \quad (37)$$

It is also readily shown that the process $\{\bar{X}_t\}_{t \geq 0}$ is in fact exact for any choice of divergence free vector field $\{\xi_t\}_{t \geq 0}$. The most convenient such choice is to simply set $\xi_t \equiv 0$. The resulting filter is exact and furthermore also (infinitesimally) optimal to the first-order.

For the special case of the nonlinear filtering problem, $\mathcal{H}(\rho) = (h - \bar{h})\rho$ where $\bar{h} = \int h(x)\rho(x)dx$ and $dI_t = (dZ_t - \bar{h}_t dt)$ is the increment of the innovation process. For these choices, the optimal transport stochastic process (37) becomes

$$d\bar{X}_t = \nabla \phi_t(\bar{X}_t) \circ (dZ_t - \frac{1}{2} (h(\bar{X}_t) + \bar{h}_t) dt) + \xi_t(\bar{X}_t) dt.$$

The feedback control law in the FPF algorithm (11) represents the particular sub-optimal choice $\xi_t \equiv 0$. The choice is optimal for $d = 1$.

Remark 5.3. The sub-optimality of FPF is not a problem because the filter is exact. A case for FPF may be made on computational grounds. Because it requires a solution of a single Poisson equation, the FPF control law is the simplest possible control law leading to an exact filter. A natural question then is whether there is any advantage to be had by using the optimal transport control law? As of yet, the answer to this question is not clear. The same question arises in the optimal transport map estimation problem (Makkuva, Taghvaei, Oh, & Lee, 2020): why aim for the optimal transport map as opposed to say Knothe–Rosenblatt rearrangement (Villani, 2009, Ch. 1)? As an additional point, there is also a freedom in replacing the quadratic cost function in the optimal transport problem. An argument for the optimal transport map with quadratic cost function can be made on the account of its special geometrical structure: the optimal map is the gradient of a convex function. This may lead to nice computational and stability properties when the map is approximated with particles/samples.

5.3. Optimal transport formula for the static example

We now revisit the static example introduced in Section 3.1 with the aim of deriving an explicit form of the control U and relating it to the FPF. As explained in Section 3.1, the problem is to find a control U such that $\mathbb{E}[f(X)|Y] = \mathbb{E}[f(\bar{X}_1)|Y]$ for all functions $f \in C_b(\mathbb{R}^d)$, where $\bar{X}_1 = \bar{X}_0 + U$ and $\bar{X}_0$ is an independent copy of X . This condition is equivalently expressed as $(\bar{X}_1, Y) \sim P_{XY}$, and the problem of finding U is formulated as the following optimal transportation problem:

$$\min_{U \in \sigma(\bar{X}_0, Y)} \mathbb{E}[|U|^2], \quad (38)$$

s.t. $\bar{X}_1 = \bar{X}_0 + U, \quad (\bar{X}_1, Y) \sim P_{XY},$

where the notation $U \in \sigma(\bar{X}_0, Y)$ means that U is allowed to be measurable with respect to $\bar{X}_0$ and Y . This is an optimal transportation problem between $(\bar{X}_0, Y) \sim P_X \otimes P_Y$ and $(X, Y) \sim P_{XY}$ where the transportation is constrained to be of the form $(\bar{X}_0, Y) \rightarrow (\bar{X}_0 + U, Y)$, i.e., the second argument Y remains fixed. Its solution is obtained as an extension of the celebrated Brenier's result (Brenier, 1991) as follows:

Theorem 5.4 (Thm. 1 in Taghvaei & Hosseini, 2022). Consider the optimal transportation problem (38). Suppose P_X admits a density with respect to the Lebesgue measure. Then the optimal control is

$$U = \nabla \bar{\Phi}(\bar{X}_0; Y) - \bar{X}_0,$$

where $\bar{\Phi}$ is the minimizer of the dual Kantorovich problem

$$\min_{\Phi \in CVX_x} \mathbb{E}[\Phi(\bar{X}_0; Y) + \Phi^*(X; Y)], \quad (39)$$

where $\Phi \in CVX_x$ means $x \mapsto \Phi(x; y)$ is convex in x for all y and $\Phi^*(x; y) := \sup_z z^T x - \Phi(z; y)$ is the convex conjugate of Φ with respect to x .

Remark 5.5 (Relationship to the Update Formula for FPF). In the continuous-time limit, the dual Kantorovich problem (39) is related to the variational form (25) of the Poisson equation (12). In particular, with $\Delta Z_t = h(X_t)\Delta t + \Delta W_t$, the solution to the problem (39) is as follows (Taghvaei & Hosseini, 2022, Prop. 2):

$$\bar{\Phi}(\bar{X}_t; \Delta Z_t) = \frac{1}{2} |\bar{X}_t|^2 + \phi(\bar{X}_t) \Delta Z_t + \psi(\bar{X}_t) \Delta t + O(\Delta t^2)$$

where ϕ is the solution to the Poisson equation (12) with ρ taken as the density of P_X , and ψ is the unique such function such that $\nabla \psi = -\frac{h+\bar{h}}{2} \nabla \phi + \frac{1}{4} \nabla |\nabla \phi|^2 + \xi$ where ξ is divergence free. Therefore, the optimal transformation $\bar{X}_t \mapsto \bar{X}_{t+\Delta t}$ is given by,

$$\begin{aligned} \bar{X}_{t+\Delta t} &= \nabla_x \bar{\Phi}(\bar{X}_t; \Delta Z_t) \\ &= \bar{X}_t + \nabla \phi(\bar{X}_t) (\Delta Z_t - \frac{h(\bar{X}_t) + \bar{h}_t}{2} \Delta t) \\ &\quad + \frac{1}{4} \nabla |\nabla \phi(\bar{X}_t)|^2 \Delta t + \xi(\bar{X}_t) \Delta t + O(\Delta t^2) \end{aligned}$$

which in the limit as $\Delta t \rightarrow 0$ is the SDE for the optimal transport FPF (37).

Remark 5.6 (Stochastic Optimization and DNNs). The variational problem (39) is a stochastic optimization problem which allows for application of machine learning tools to approximate its solution. In particular, deep neural networks (DNNs) can be used to parametrize the function Φ and stochastic optimization algorithms employed to learn the parameters. Preliminary results in this direction are presented in Taghvaei and Hosseini (2022) with a comprehensive development the subject of ongoing work.

PART II

6. CIPS for optimal control

In order to elucidate the ideas as clearly as possible, our focus in this paper is entirely on the linear quadratic (LQ) problem. Its extension to the nonlinear optimal control problem (4) can be found in Joshi, Taghvaei, Mehta, and Meyn (2022).

6.1. Problem statement and background

The finite-horizon linear quadratic (LQ) optimal control problem is a special case of (4) as follows:

$$\min_u J(u) = \int_0^T \frac{1}{2} (|Cx_t|^2 + u_t^T R u_t) dt + x_T^T P_T x_T \quad (40a)$$

$$\text{subject to: } \dot{x}_t = Ax_t + Bu_t, \quad x_0 = x \quad (40b)$$

It is assumed that (A, B) is controllable, (A, C) is observable, and matrices $P_T, R > 0$. The $[T = \infty]$ limit is referred to as the linear quadratic regulator (LQR) problem.

It is well known that the optimal control $u_t = \phi_t(x_t)$ where the optimal policy is linear

$$\phi_t(x) = K_t x \quad \text{where} \quad K_t = -R^{-1} B^T P_t, \quad 0 \leq t \leq T$$

is the optimal gain matrix and $\{P_t : 0 \leq t \leq T\}$ is a solution of the backward (in time) differential Riccati equation (DRE)

$$-\frac{d}{dt}P_t = A^T P_t + P_t A + C^T C - P_t B R^{-1} B^T P_t, \quad P_T \text{ (given)} \quad (41)$$

The algebraic Riccati equation (ARE) is obtained by setting the left-hand side to 0. As $T \rightarrow \infty$, for each fixed time t , $P_t \rightarrow P^\infty$, exponentially fast (Kwakernaak & Sivan, 1972, Thm. 3.7), where $P^\infty > 0$ is the unique such positive-definite solution of the ARE, and therefore the optimal gain converges, $K_t \rightarrow K^\infty := -R^{-1} B^T P^\infty$. Approximation of the gain K^∞ is a goal in recent work on model-based RL for the LQR problem (Fazel et al., 2018; Mohammadi et al., 2022).

6.2. Objectives and assumptions

For the reasons noted in Section 1, we are interested in a simulation-based solution that does not rely on an explicit solution of the DRE (41). To clarify what is meant by a simulation-based solution in the context of model-based RL, we make a formal assumption as follows:

Assumption 1.

1. Functions $f(x, \alpha) = Ax + B\alpha$ and $c(x) = Cx$ are available in the form of an oracle (which allows function evaluation at any state action pair $(x, \alpha) \in \mathbb{R}^d \times \mathbb{R}^m$).
2. Matrices R and P_T are available. Both of these matrices are strictly positive-definite.
3. Simulator is available to simulate (40b).
4. Simulator provides for an ability to add additional inputs outside the control channel (e.g., see (5a)).

This assumption is motivated from the data assimilation literature where it is entirely standard and widely used in applications, such as weather prediction, involving EnKF. Part 1 of the assumption means that the matrices A, B, C are not available explicitly. Rather, for any given $(x, \alpha) \in \mathbb{R}^d \times \mathbb{R}^m$, the vectors $f(x, \alpha)$ and $c(x)$ can be evaluated. Function evaluation forms for the dynamics and the cost function is also a standard assumption for any model-based RL algorithm. Part 2 of the assumption is not too restrictive for the following two reasons:

1. In physical systems, one is typically able to assess relative costs for different control inputs (actuators). This knowledge can be used to select R .
2. For the LQR problem, under mild technical conditions, the optimal policy is stationary and does not depend upon the choice of P_T .

If these matrices are not available, one possibility is to take R and P_T to be identity matrices of appropriate dimensions. The main restriction comes from part 3 of the assumption. However, as the widespread use of EnKF amply demonstrates, it is not un-realistic to assume it for a simulation-based solution. Of course, it will not be possible with a physical experiment.

6.3. Dual EnKF

The dual EnKF algorithm is obtained from making use of duality between optimal control and filtering. For this purpose, we need to first dualize the DRE (41). Under the assumptions of this paper, $P_t > 0$ for $0 \leq t \leq T$ whenever $P_T > 0$ (Brockett, 2015, Sec. 24). Set $S_t = P_t^{-1}$. It is readily verified that $\{S_t : 0 \leq t \leq T\}$ also solves a DRE (which represents the dual of (41))

$$\frac{d}{dt}S_t = AS_t + S_t A^T - BR^{-1}B^T + S_t C^T C S_t, \quad S_T = P_T^{-1} \quad (42)$$

The strategy is to approximate $\{S_t : 0 \leq t \leq T\}$ using simulations. As before, the construction proceeds in two steps: (i) definition of an exact mean-field process; and (ii) its finite- N approximation.

Step 1. Mean-field process: Define a stochastic process $\bar{Y} = \{\bar{Y}_t \in \mathbb{R}^d : 0 \leq t \leq T\}$ as a solution of the following backward (in time) SDE:

$$d\bar{Y}_t = A\bar{Y}_t dt + B d\bar{\eta}_t + \frac{1}{2} \bar{S}_t C^T (C\bar{Y}_t + C\bar{n}_t) dt, \quad 0 \leq t < T$$

$$\bar{Y}_T \sim \mathcal{N}(0, S_T) \quad (43)$$

where $\eta = \{\eta_t \in \mathbb{R}^m : 0 \leq t \leq T\}$ is a W.P. with covariance matrix R^{-1} , and

$$\bar{n}_t := \mathbb{E}[\bar{Y}_t], \quad \bar{S}_t := \mathbb{E}[(\bar{Y}_t - \bar{n}_t)(\bar{Y}_t - \bar{n}_t)^T], \quad 0 \leq t < T \quad (44)$$

The meaning of the backward arrow on $d\bar{\eta}$ in (43) is that the SDE is simulated backward in time starting from the terminal condition specified at time $t = T$. The reader is referred to Nualart and Pardoux (1988, Sec. 4.2) for the definition of the backward Itô-integral. The mean-field process is useful because of the following proposition.

Proposition 6.1 (Prop. 1 in Joshi et al., 2022). The solution to the SDE (43) is a Gaussian stochastic process, in which the mean and covariance of $\bar{Y}_t$ are given by

$$\bar{n}_t = 0, \quad \bar{S}_t = S_t, \quad 0 \leq t \leq T$$

Consequently, $\bar{X}_t := \bar{S}_t^{-1}(\bar{Y}_t - \bar{n}_t)$ is also a Gaussian random variable with

$$\mathbb{E}[\bar{X}_t] = 0, \quad \mathbb{E}[\bar{X}_t \bar{X}_t^T] = P_t, \quad 0 \leq t \leq T$$

The significance of Proposition 6.1 is that the optimal control policy $\phi_t(\cdot)$ can now be obtained in terms of the statistics of the random variable $\bar{X}_t$. Specifically, we have the following two cases:

1. If the matrix B is explicitly known then the optimal gain matrix $K_t = -R^{-1} B^T \mathbb{E}[\bar{X}_t \bar{X}_t^T]$
2. If B is unknown, define the Hamiltonian (the continuous-time counterpart of the Q-function (Mehta & Meyn, 2009)):

$$H(x, \alpha, t) := \underbrace{\frac{1}{2} |Cx|^2 + \frac{1}{2} \alpha^T R \alpha + x^T \mathbb{E}[\bar{X}_t \bar{X}_t^T] (Ax + B\alpha)}_{\text{cost function}} \quad \underbrace{\quad}_{\text{model (40b)}}$$

from which the optimal control law is obtained as

$$\phi_t(x) = \arg \min_{\alpha \in \mathbb{R}^m} H(x, \alpha, t)$$

by recalling the minimum principle, which states that the optimal control is the unique minimizer of the Hamiltonian. It is noted that the Hamiltonian $H(x, \alpha, t)$ is in the form of an oracle because $(Ax + B\alpha)$ is the right-hand side of the simulation model (40b).

Step 2. Finite- N approximation: The particles $\{Y_t^i \in \mathbb{R}^d : 0 \leq t \leq T, i = 1, \dots, N\}$ evolve according to the backward SDE:

$$dY_t^i = \underbrace{AY_t^i dt + B d\eta_t^i}_{\text{ith copy of model (40b)}} + \underbrace{S_t^{(N)} C^T \left(\frac{CY_t^i + Cn_t^{(N)}}{2} \right) dt}_{\text{coupling}} \quad (45)$$

$$Y_T^i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, P_T^{-1}), \quad 1 \leq i \leq N$$

$\eta^i := \{\eta_t^i : 0 \leq t \leq T\}$ is an i.i.d. copy of η and

$$n_t^{(N)} = \frac{1}{N} \sum_{i=1}^N Y_t^i$$

$$S_t^{(N)} = \frac{1}{N-1} \sum_{i=1}^N (Y_t^i - n_t^{(N)})(Y_t^i - n_t^{(N)})^T$$




Fig. 6. Comparison of the numerical solution obtained from the EnKF, the DRE, and the ARE. Note the x -axis for these plots is $T - t$ for $0 \leq t \leq T$. d is the state-dimension.

The CIPS (45) is referred to as the *dual EnKF*.

Optimal control: Set $X_t^i = (S_t^{(N)})^{-1}(Y_t^i - n_t^{(N)})$. There are two cases as before:

1. If the matrix B is explicitly known then

$$K_t^{(N)} = -\frac{1}{N-1} \sum_{i=1}^N R^{-1} (B^T X_t^i) (X_t^i)^T \quad (46)$$

2. If B is unknown, define the Hamiltonian

$$H^{(N)}(x, \alpha, t) := \underbrace{\frac{1}{2} |Cx|^2 + \frac{1}{2} \alpha^T R \alpha}_{\text{cost function}} + \frac{1}{N-1} \sum_{i=1}^N (x^T X_t^i) (X_t^i)^T \underbrace{(Ax + B\alpha)}_{\text{model (40b)}}$$

from which the optimal control policy is approximated as

$$\phi_t^{(N)}(x) = \arg \min_{a \in \mathbb{R}^m} H^{(N)}(x, a, t)$$

There are several zeroth-order approaches to solve the minimization problem, e.g., by constructing 2-point estimators for the gradient. Since the objective function is quadratic and the matrix R is known, m queries of $H^{(N)}(x, \cdot, t)$ are sufficient to compute $\phi_t^{(N)}(x)$.

The overall dual EnKF algorithm can be found in Joshi et al. (2022, Algorithm 1 and 2).

6.4. Relating dual EnKF to model-based RL

The following remarks are included to help provide an intuitive explanation of the various aspects of the dual EnKF and relate these to the model-based RL:

1. **Representation.** In designing any RL algorithm, the first issue is the representation of the unknown value function (P_t in the linear case). Our novel idea is to represent P_t in terms of statistics (variance) of the particles. Such a representation is distinct from representing the value function, or its proxies, such as the Q function, within a parameterized class of functions.

2. **Value iteration.** The algorithm is entirely simulation based: N copies of the model (40b) are simulated in parallel where the terms on the right hand-side of (45) have the following intuitive interpretations:

- (a) **Dynamics:** The first term “ $AY_t^i dt$ ” on the right-hand side of (45) is simply a copy of uncontrolled dynamics in the model (40b).
- (b) **Control:** The second term “ $Bd\eta_t^i$ ” is the control input for the i th particle. It is specified as a W.P. with covariance R^{-1} . One may interpret this as an approach to exploration whereby cheaper control directions are explored more.
- (c) **Coupling:** The third term, referred to as the coupling, effectively implements the value iteration step. Coupling has a “gain times error” structure where $S_t^{(N)} C^T$ is the gain and $\frac{1}{2}(CY_t^i + Cn_t^{(N)})$ is the counterpart of the error in the linear FPF (14).

3. **Arrow of time.** The particles are simulated backward—from terminal time $t = T$ to initial time $t = 0$. This is different from most model-based RL but consistent with the dynamic programming (DP) equation which also proceeds backward in time.

6.5. Convergence and error analysis

In Joshi et al. (2022, Prop. 3), under certain additional assumptions on system matrices, the following error bound is derived:

$$\mathbb{E}[\|S_t^{(N)} - \tilde{S}_t\|_F] \leq \frac{C_1}{\sqrt{N}} + C_2 e^{-2\lambda(T-t)} \mathbb{E}[\|S_T^{(N)} - \tilde{S}_T\|_F], \quad (47)$$

where C_1, C_2, λ are positive constants and $\|\cdot\|_F$ denotes Frobenius norm for matrices. The significance of the bound (47) is as follows: The constant λ is same as the rate that governs the convergence of the solution of the DRE (41) to the stationary solution (of the infinite-horizon LQR problem). *This means that the dual EnKF learns the optimal LQR gain exponentially fast with a rate that is as good as one would obtain from directly solving the DRE.*




Fig. 7. Open and closed-loop poles for the two plots (parts (a) and (b)) depicted in Fig. 6.

Convergence is numerically illustrated for a d -dimensional system expressed in its controllable canonical form

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ a_1 & a_2 & a_3 & a_4 & \dots & a_d \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

where the entries $(a_1, \dots, a_d) \in \mathbb{R}^d$ are i.i.d. samples from $\mathcal{N}(0, 1)$. The matrices C, R, P_T are identity matrices of appropriate dimension. For numerics, we fix $T = 10$, chose the time-discretization step as 0.02, and use $N = 1000$ particles to simulate the dual EnKF.

Fig. 6(a) depicts the convergence of the four entries of the matrix $P_t^{(N)}$ for the case where $d = 2$. Fig. 6(b) depicts the analogous results for $d = 10$. Fig. 7(a) and Fig. 7(b) depict the open-loop poles (eigenvalues of the matrix A) and the closed-loop poles (eigenvalues of the matrix $(A + BK_0^{(N)})$), for $d = 2$ and $d = 10$, respectively. Note that the closed-loop poles are stable, whereas some open-loop poles have positive real parts.

6.6. Comparison to literature

We present a comparison of the dual EnKF with policy gradient algorithms in Mohammadi et al. (2022) (denoted as [M21]) and Fazel et al. (2018) (denoted as [F18]). In these prior works, by restricting the control policies to the linear form $u_t = Kx_t$, the LQR problem reduces to the finite-dimensional static optimization problem:

$$K^* = \arg \min_K J(K) = \mathbb{E} \left(\int_0^\infty x_t^T Q x_t + u_t^T R u_t dt \right) \quad (48)$$

where the expectation is over the initial condition. The authors apply a pure-actor method using “zeroth order” methods to approximate gradient descent, much like the early REINFORCE algorithm for RL (Sutton & Barto, 2018).

A qualitative comparison of the dual EnKF with these prior algorithms is given in Table 2. Choosing $t = 0$ in (47), the error is smaller than ϵ if the number of particles $N > O(\frac{1}{\epsilon^2})$ and the simulation time $T > O(\log(\frac{1}{\epsilon}))$, while the iteration number is one. This is compared with policy optimization approach in Fazel et al. (2018) where the number of particles and the simulation time scales polynomially with ϵ , while the number of iterations scale as $O(\log(\frac{1}{\epsilon}))$. This result is later refined in Mohammadi et al. (2022) where the required number of particles and the simulation time are shown to be $O(1)$ and $O(\log(\frac{1}{\epsilon}))$ respectively (although this result is valid with probability that approaches zero as the number of iterations grow (Mohammadi et al., 2022, Thm. 3)).

A numerical comparison is made on the benchmark spring mass damper example borrowed from Mohammadi, Zare, Soltanolkotabi, and

Table 2

Computational complexity comparison of the algorithms to achieve ϵ error in approximating the infinite-horizon LQR optimal gain.

Algorithm	Particles/samples	Simulation time	Iterations
dual EnKF	$O(\frac{1}{\epsilon^2})$	$O(\log(\frac{1}{\epsilon}))$	1
Fazel et al. (2018)	$\text{poly}(\frac{1}{\epsilon})$	$\text{poly}(\frac{1}{\epsilon})$	$O(\log(\frac{1}{\epsilon}))$
Mohammadi et al. (2022)	$O(1)$	$O(\log(\frac{1}{\epsilon}))$	$O(\log(\frac{1}{\epsilon}))$

Jovanovic (2019, Sec. VI). Fig. 8 depicts the relative mean-squared error, defined as

$$\text{MSE} := \frac{1}{T} \mathbb{E} \left(\int_0^T \frac{\|P_t - P_t^{(N)}\|_F^2}{\|P_t\|_F^2} dt \right)$$

Two trends are depicted in the figure: the $O(\frac{1}{N})$ decay of the MSE as N increases (for d fixed), which is a numerical illustration of the error bound (47), and a plot of the MSE as a function of dimension d (for N fixed).

A side-by-side comparison with [F18] and [M21] is depicted in Fig. 9. The comparison is for the following metrics (taken from Mohammadi et al., 2022):

$$\text{error}^{\text{gain}} = \frac{\|K^{\text{est}} - K^\infty\|_F}{\|K^\infty\|_F}, \quad \text{error}^{\text{value}} = \frac{c^{\text{est}} - c^\infty}{c_{\text{init}}^{(N)} - c^\infty}$$

where the LQR optimal gain K^∞ and the optimal value c^∞ are computed from solving the ARE. The value $c_{\text{init}}^{(N)}$ is approximated using the initial gain $K = 0$ (Note such a gain is not necessary for EnKF). Because [F18] is for discrete-time system, an Euler approximation is used to obtain a discrete-time model.

In the numerical experiments, the dual EnKF is found to be significantly more computationally efficient—by two orders of magnitude or more. The main reason for the order of magnitude improvement in computational time is as follows: An EnKF requires only a single iteration over a fixed time-horizon. In contrast, [F18] and [M21] require several steps of gradient descent, with each step requiring an evaluation of the LQR cost, and because these operations must be done serially, these computations are slower.

In carrying out these comparisons, the same time-horizon $[0, T]$ and discretization time-step Δt was used for all the algorithms. It is certainly possible that some of these parameters can be optimized to improve the performance of the other algorithms. In particular, one may consider shorter or longer time-horizon T or use parallelization to speed up the




Fig. 8. Performance of the dual EnKF algorithm: MSE as a function of the number of particles N and system dimension d .




Fig. 9. Comparison with algorithms in Fazel et al. (2018) (labeled [F18]) and Mohammadi et al. (2022) (labeled [M21]). The comparisons depict the computation time (in Python) as a function of the relative error in approximating the LQR gain and cost.

gradient calculation. Codes are made available on Github for interested parties to independently verify these comparisons.¹

6.7. Extension to the nonlinear problem (4)

An extension of the dual EnKF algorithm for the nonlinear optimal control problem (4) appears in Joshi et al. (2022, Sec. 3). In the general nonlinear setting, the empirical distribution of the N particles approximates the minus log of the value function, leading to the optimal control law (5b). The algorithm involves the solution of a Poisson equation, similar to the Poisson equation that appears in the FPF algorithm. The dual EnKF algorithm for the LQ problem arises as a special case when the Poisson equation admits an analytical solution. An interested reader can find additional details in Joshi et al. (2022) where some numerical results for the problem of stabilizing an inverted pendulum on the cart are also described.

7. Discussion and conclusion

In this survey, we described CIPS to approximate the solution of the optimal filtering and the optimal control problems (in parts I and II, respectively). As explained in Section 1, there are close parallels

with DA and RL. In this section, we expand on some of these parallels with the goal of highlighting some important points and directions for future work.

Data assimilation, sampling, optimal transportation. CIPS may be viewed as a sampling algorithm. The FPF control law (coupling) is designed to sample from the posterior. Compared to the conventional particle filters, coupling is beneficial because the issue of particle degeneracy is avoided (as discussed in Section 3.4). To design the coupling, optimal transportation theory provides a useful framework (as described in Section 5). Variations of the basic approach described here have been used in construction of a class of filtering algorithms (Garbuno-Inigo, Hoffmann, Li, & Stuart, 2020; Halder & Georgiou, 2017, 2018, 2019; Luo, 2019). The optimal transport formulation has also been extended to the Schrödinger bridge setting by considering a cost with respect to the (prior) dynamics, or considering an entropic regularization (Chen, Georgiou, & Pavon, 2016; Reich, 2019). In related works, the coupling viewpoint along with geometric notions from optimal transportation theory, have enabled application of optimization algorithms to design sampling schemes (Chen, Zhang, Wang, Li, & Chen, 2018; Chizat & Bach, 2018; Frogner & Poggio, 2018; Liu & Wang, 2016; Liu et al., 2018; Richemond & Maginnis, 2017; Taghvaei & Mehta, 2019; Zhang, Chen, Li, & Carin, 2018; Zhang, Taghvaei, & Mehta, 2019).

Part II of this paper is motivated by the enormous success of the CIPS (EnKF) in DA.

¹ <https://github.com/anantjoshi97/EnKF-RL>

Reinforcement learning and optimal control. Compared to typical RL approaches, there are two key innovations/differences:

1. Representation of the unknown value function in terms of the statistics (variance) of a suitably designed process; and
2. Design of interactions (coupling) between simulations for the purposes of policy optimization.

We fully believe that the two key innovations may be useful for many other types of models including MDPs and partially observed problems. In the LQ setting of the problem, doing so is beneficial because of the learning rate: Since the $[N = \infty]$ limit is exact for the LQ problem, the dual EnKF algorithm yields a learning rate that closely approximates the exponential rate of convergence of the solution of the DRE. This is rigorously established with the aid of error bound (47). In numerical examples, this property is shown to lead to an order of magnitude better performance than the state-of-the-art algorithms.

Apart from RL, model predictive control (MPC) is another area where a model in the form of a simulator is assumed to design optimal control for problems such as (4) (Rawlings et al., 2017). Using duality, MPC methods have been adapted to design the moving horizon estimator (MHE). A big selling point of MPC is its ability to handle constraints which has not been a major theme in the DA literature. Another notable distinction is that while MPC aims to find a single (optimal) trajectory, CIPS simulate multiple stochastic trajectories in a Monte Carlo manner. Notably, the solution of the deterministic optimal control problem (4) is based on simulating (5a) which is an SDE. For the stochastic MPC problems, multiple simulations have been considered in the scenario-based approach (Campi & Garatti, 2018).

Some perspectives on future research. In basic sciences, there are a number of important examples of interacting particle systems. This paper presents results on the theme of “CIPS as an algorithm”. The most historical of such algorithms is the EnKF which is used to solve the problem of data assimilation. It is hoped that this survey convinces the reader that the paradigm is also useful for solving other problems in estimation and control. A major selling point of CIPS, and also the reason for widespread use of the EnKF, is that it is able to work directly with a simulator. Therefore, it is amenable as a solution method for complex systems where models typically exist only in the form of a simulator. Apart from the open problems described in the main body of the paper, a few themes for future research are as follows:

- MPC offers a useful benchmark for CIPS. With the exception of the geometric approaches, e.g., FPF on Riemannian manifolds (Zhang et al., 2017b), constraints has not been an important theme in design of CIPS. It is an important problem to extend the design of mean-field process to handle general types of constraints in inputs and states. One possible next step is to extend the dual EnKF to the inequality-constrained LQR problems.
- RL could be an important application for CIPS. A key difference is that CIPS-based solution does not rely on function approximation. Instead, the value function is approximated in terms of the distribution of the particles. This has some advantages, e.g., avoids the need to select basis functions, and some disadvantages, e.g., availability of computational resources. It will be useful to understand some of these trade-offs.
- Relationship to mean-field games and optimal control should be further developed. CIPS represent simple examples of mean-field type control laws. However, derivation of these control laws is, more often than not, rooted in methods from optimal transportation theory (Section 5). It remains an open problem to derive the FPF control law starting from a mean-field optimal control type objective (some partial results in this direction appear in Zhang et al., 2019).

- The lack of progress to obtain FPF as a solution of an optimal control problem is symptomatic of a satisfactory duality theory between optimal filtering and optimal control (Todorov, 2008). Recent progress in this direction has been made in some work originating in our group (Kim, 2022; Kim & Mehta, 2022b; Kim, Mehta, & Meyn, 2019). While the focus of this new work has thus far been on dual characterization of stochastic observability (Kim & Mehta, 2021a, 2022a) and its use in filter stability analysis (Kim & Mehta, 2021b; Kim, Mehta, & Meyn, 2021), it will be interesting to explore connections both to FPF and to mean-field control. Duality-based derivation of the EnKF has previously been considered in Kim et al. (2018).
- Extensions to partially observed optimal control problems. For the linear Gaussian model, algorithms described in parts I and II are easily combined to obtain a CIPS for the partially observed problem. The solution is based on the separation principle: A forward (in time) EnKF is run to solve the optimal filtering problem; and a completely independent backward (in time) dual EnKF is run to solve the optimal control problem. For the nonlinear problem, there may be benefit to couple the forward and backward CIPS.
- Distributionally robust FPF. In order to handle uncertainty in signal and observation models, it may be useful to explore methods from distributionally robust optimization framework (Rahimian & Mehrotra, 2019). The framework has been used to develop the Wasserstein robust Kalman filter for the linear Gaussian model (Shafieezadeh Abadeh, Nguyen, Kuhn, & Mohajerin Esfahani, 2018). Its extension to the nonlinear filtering model (1) is open and may be possible based on the optimal transport formulation of the FPF.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

- Abedi, E., & Surace, S. C. (2019). Gauge freedom within the class of linear feedback particle filters. In *2019 IEEE 58th conference on decision and control* (pp. 666–671). <http://dx.doi.org/10.1109/CDC40024.2019.9029897>.
- Abedi, E., Surace, S. C., & Pfister, J. P. (2022). A unification of weighted and unweighted particle filters. *SIAM Journal on Control and Optimization*, 60(2), 597–619.
- Anderson, J. L. (2001). An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129(12), 2884–2903.
- Anthony, M., Bartlett, P. L., Bartlett, P. L., et al. (1999). *Neural network learning: Theoretical foundations*, vol. 9. Cambridge University Press Cambridge.
- Bain, A., & Crisan, D. (2009). *Fundamentals of stochastic filtering*, vol. 3. Springer, <http://dx.doi.org/10.1007/978-0-387-76896-0>.
- Bakry, D., Barthe, F., Cattiaux, P., & Guillin, A. (2008). A simple proof of the Poincaré inequality for a large class of probability measures including the log-concave case. *Electronic Communications in Probability*, 13, 60–66. <http://dx.doi.org/10.1214/ECP.v13-1352>.
- Bakry, D., Gentil, I., & Ledoux, M. (2013). *Analysis and geometry of Markov diffusion operators*, vol. 348. Springer Science & Business Media.
- Bar-Shalom, Y., Daum, F., & Huang, J. (2009). The probabilistic data association filter. *IEEE Control Systems Magazine*, 29(6), 82–100.
- Beneš, V. (1981). Exact finite-dimensional filters for certain diffusions with nonlinear drift. *Stochastics. An International Journal of Probability and Stochastic Processes*, 5(1–2), 65–92.
- Bengtsson, T., Bickel, P., & Li, B. (2008). *IMS lecture notes - Monograph series in probability and statistics: Essays in honor of David F. Freedman: vol. 2, Curse of dimensionality revisited: Collapse of the particle filter in very large scale systems* (pp. 316–334). Institute of Mathematical Sciences.
- Bensoussan, A., Frehse, J., Yam, P., et al. (2013). *Mean field games and mean field type control theory*, vol. 101. Springer.

- Bergemann, K., & Reich, S. (2012). An ensemble Kalman–Bucy filter for continuous data assimilation. *Meteorologische Zeitschrift*, 21(3), 213.
- Berntorp, K. (2015). Feedback particle filter: Application and evaluation. In *18th int. conf. information fusion*. URL: <https://ieeexplore.ieee.org/document/7266752>.
- Berntorp, K. (2018). Comparison of gain function approximation methods in the feedback particle filter. In *2018 21st international conference on information fusion* (pp. 123–130). IEEE, <http://dx.doi.org/10.23919/ICIF.2018.8455574>.
- Berntorp, K., & Grover, P. (2016). Data-driven gain computation in the feedback particle filter. In *2016 American control conference* (pp. 2711–2716). <http://dx.doi.org/10.1109/ACC.2016.7525328>.
- Bertsekas, D., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Cambridge, Mass: Athena Scientific.
- Beskos, A., Crisan, D., Jasra, A., & Whiteley, N. (2014). Error bounds and normalising constants for sequential Monte Carlo samplers in high dimensions. *Advances in Applied Probability*, 46(1), 279–306. <http://dx.doi.org/10.1017/s0001867800007047>.
- Bickel, P., Li, B., Bengtsson, T., et al. (2008). Sharp failure rates for the bootstrap particle filter in high dimensions. In *Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh* (pp. 318–329). Institute of Mathematical Statistics.
- Bishop, A. N., & Del Moral, P. (2018). On the stability of matrix-valued Riccati diffusions. arXiv preprint [arXiv:1808.00235](https://arxiv.org/abs/1808.00235).
- Bishop, A. N., & Del Moral, P. (2020). On the mathematical theory of ensemble (linear-Gaussian) Kalman–Bucy filtering. arXiv preprint [arXiv:2006.08843](https://arxiv.org/abs/2006.08843).
- Bishop, C. H., Etherton, B. J., & Majumdar, S. J. (2001). Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Monthly Weather Review*, 129(3), 420–436.
- Blom, H. A. (2013). The continuous time roots of the interacting multiple model filter. *Integrated Tracking, Classification, and Sensor Management*, 127–162.
- Brenier, Y. (1991). Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4), 375–417.
- Brockett, R. W. (2015). *Finite dimensional linear systems*. SIAM.
- Burgers, G., Van Leeuwen, P. J., & Evensen, G. (1998). Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126(6), 1719–1724.
- Calvelli, E., Reich, S., & Stuart, A. M. (2022). Ensemble Kalman methods: A mean field perspective. arXiv preprint [arXiv:2209.11371](https://arxiv.org/abs/2209.11371).
- Campi, M. C., & Garatti, S. (2018). *Introduction to the scenario approach*. SIAM.
- Carmona, R., Delarue, F., et al. (2018). *Probabilistic theory of mean field games with applications I-II*. Springer.
- Carmona, R., & Graves, C. V. (2020). Jet lag recovery: Synchronization of circadian oscillators as a mean field game. *Dynamic Games and Applications*, 10(1), 79–99.
- Chen, Y., Georgiou, T. T., & Pavon, M. (2016). On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2), 671–691.
- Chen, X., Luo, X., Shi, J., & Yau, S. S. T. (2021). General convergence result for continuous-discrete feedback particle filter. *International Journal of Control*, 1–15.
- Chen, C., Zhang, R., Wang, W., Li, B., & Chen, L. (2018). A unified particle-optimization framework for scalable Bayesian sampling. In *Conference on uncertainty in artificial intelligence*.
- Chizat, L., & Bach, F. R. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Neural information processing systems*.
- Coifman, R. R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1), 5–30. <http://dx.doi.org/10.1016/j.acha.2006.04.006>.
- Crisan, D., & Xiong, J. (2010). Approximate McKean–Vlasov representations for a class of SPDEs. *Stochastics*, 82(1), 53–68. <http://dx.doi.org/10.1080/17442500902723575>.
- Daum, F., & Huang, J. (2008). Particle flow for nonlinear filters with log-homotopy. In *Proc. SPIE*, vol. 6969 (p. 696918). <http://dx.doi.org/10.1117/12.764909>.
- de Wiljes, J., Reich, S., & Stannat, W. (2018). Long-time stability and accuracy of the ensemble Kalman–Bucy filter for fully observed processes and small measurement noise. *SIAM Journal on Applied Dynamical Systems*, 17(2), 1152–1181. <http://dx.doi.org/10.1137/17m1119056>.
- Dean, S., Mania, H., Matni, N., Recht, B., & Tu, S. (2020). On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4), 633–679. <http://dx.doi.org/10.1007/s10208-019-09426-y>.
- Del Moral, P. (2004). *Feynman–Kac formulae* (pp. 47–93). Springer.
- del Moral, P., & Horton, E. (2021). Quantum harmonic oscillators and Feynman–Kac path integrals for linear diffusive particles. arXiv preprint [arXiv:2106.14592](https://arxiv.org/abs/2106.14592).
- Del Moral, P., Kurtzmann, A., & Tugaut, J. (2017). On the stability and the uniform propagation of chaos of a class of extended ensemble Kalman–Bucy filters. *SIAM Journal on Control and Optimization*, 55(1), 119–155. <http://dx.doi.org/10.1137/16M1087497>.
- Del Moral, P., & Tugaut, J. (2018). On the stability and the uniform propagation of chaos properties of ensemble Kalman–Bucy filters. *Annals of Applied Probability*, 28(2), 790–850. <http://dx.doi.org/10.1214/17-AAP1317>.
- Dörfler, F., & Bullo, F. (2014). Synchronization in complex networks of phase oscillators: A survey. *Automatica*, 50(6), 1539–1564.
- Doucet, A., & Johansen, A. M. (2009). A tutorial on particle filtering and smoothing: Fifteen years later. In *Handbook of nonlinear filtering*, vol. 12 (pp. 656–704).
- El Moselhy, T. A., & Marzouk, Y. M. (2012). Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23), 7815–7850. <http://dx.doi.org/10.1016/j.jcp.2012.07.022>.
- Ertel, S., & Stannat, W. (2022). Analysis of the ensemble Kalman–Bucy filter for correlated observation noise. arXiv preprint [arXiv:2205.14253](https://arxiv.org/abs/2205.14253).
- Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, 99(C5), 10143–10162. <http://dx.doi.org/10.1029/94JC00572>.
- Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4), 343–367.
- Evensen, G. (2006). *Data assimilation. the ensemble Kalman filter*. New York: Springer-Verlag.
- Fazel, M., Ge, R., Kakade, S., & Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. In *International conference on machine learning* (pp. 1467–1476). PMLR, [ISSN: 2640-3498] URL: <http://proceedings.mlr.press/v80/fazel18a.html>.
- Fleming, W. H., & Mitter, S. K. (1982). Optimal control and nonlinear filtering for nondegenerate diffusion processes. *Stochastics*, 8(1), 63–77. <http://dx.doi.org/10.1080/17442508208833228>.
- Frogner, C., & Poggio, T. (2018). Approximate inference with Wasserstein gradient flows. arXiv preprint [arXiv:1806.04542](https://arxiv.org/abs/1806.04542).
- Garbuno-Inigo, A., Hoffmann, F., Li, W., & Stuart, A. M. (2020). Interacting langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM Journal on Applied Dynamical Systems*, 19(1), 412–441.
- Gomes, D. A., Pimentel, E. A., & Voskanyan, V. (2016). *Regularity theory for mean-field game systems*. Springer.
- Gordon, N. J., Salmond, D. J., & Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEEE proceedings F (Radar and signal processing)*, vol. 140 (pp. 107–113). <http://dx.doi.org/10.1049/ip-f-2.1993.0015>.
- Hajek, B. (2015). *Random processes for engineers*. Cambridge University Press.
- Halder, A., & Georgiou, T. T. (2017). Gradient flows in uncertainty propagation and filtering of linear Gaussian systems. In *2017 IEEE 56th annual conference on decision and control* (pp. 3081–3088). IEEE.
- Halder, A., & Georgiou, T. T. (2018). Gradient flows in filtering and Fisher-Rao geometry. In *2018 annual American control conference* (pp. 4281–4286). IEEE.
- Halder, A., & Georgiou, T. T. (2019). Proximal recursion for the Wonham filter. In *2019 IEEE 56th annual conference on decision and control*. IEEE, URL: http://georgiou.eng.uci.edu/papers/CDC2019_Wonham-5.pdf.
- Hijab, O. B. (1980). *Minimum energy estimation* (Ph.D. thesis), Berkeley: University of California.
- Hoffmann, C., & Rostalski, P. (2017). Linear optimal control on factor graphs — A message passing perspective —. *IFAC-PapersOnLine*, 50(1), 6314–6319. <http://dx.doi.org/10.1016/j.ifacol.2017.08.914>, URL: <https://www.sciencedirect.com/science/article/pii/S2405896317313800>. 20th IFAC world congress.
- Houtekamer, P. L., & Mitchell, H. L. (1998). Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126(3), 796–811.
- Houtekamer, P., & Mitchell, H. (2001). A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129, 123–136.
- Houtekamer, P. L., & Zhang, F. (2016). Review of the ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144(12), 4489–4532.
- Huang, M., Caines, P. E., & Malhame, R. P. (2007). Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ -Nash equilibria. *IEEE Transactions on Automatic Control*, 52(9), 1560–1571.
- Huang, M., Malhame, R. P., & Caines, P. E. (2006). Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Communications in Information and Systems*, 6(3), 221–251.
- Joshi, A. A., Taghvaei, A., Mehta, P. G., & Meyn, S. P. (2022). Controlled interacting particle algorithms for simulation-based reinforcement learning. *Systems & Control Letters*, 170, Article 105392.
- Kalman, R. E., & Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1), 95–108. <http://dx.doi.org/10.1115/1.3658902>.
- Kang, J., Chen, X., Tao, Y., & Yau, S. S. T. (2022). Optimal transportation particle filter for linear filtering systems with correlated noises. *IEEE Transactions on Aerospace and Electronic Systems*, 58(6), 5190–5203.
- Kappen, H. J. (2005a). Linear theory for control of nonlinear stochastic systems. *Physical Review Letters*, 95, Article 200201. <http://dx.doi.org/10.1103/PhysRevLett.95.200201>, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.95.200201>.
- Kappen, H. J. (2005b). Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11), P11011. <http://dx.doi.org/10.1088/1742-5468/2005/11/p11011>.
- Kelly, D., Law, K. J., & Stuart, A. M. (2014). Well-posedness and accuracy of the ensemble Kalman filter in discrete and continuous time. *Nonlinearity*, 27(10), 2579. <http://dx.doi.org/10.1088/0951-7715/27/10/2579>.
- Kim, J. W. (2022). *Duality for nonlinear filtering* (Ph.D. thesis), Urbana: University of Illinois at Urbana-Champaign.
- Kim, J. W., & Mehta, P. G. (2021a). A dual characterization of observability for stochastic systems. *IFAC-PapersOnLine*, 54(9), 659–664.
- Kim, J. W., & Mehta, P. G. (2021b). A dual characterization of the stability of the Wonham filter. In *2021 IEEE 60th conference on decision and control* (pp. 1621–1628).

- Kim, J. W., & Mehta, P. G. (2022a). Duality for nonlinear filtering I: Observability. arXiv preprint [arXiv:2208.06586](https://arxiv.org/abs/2208.06586).
- Kim, J. W., & Mehta, P. G. (2022b). Duality for nonlinear filtering II: Optimal control. arXiv preprint [arXiv:2208.06587](https://arxiv.org/abs/2208.06587).
- Kim, J. W., Mehta, P. G., & Meyn, S. (2019). What is the Lagrangian for nonlinear filtering? In *2019 IEEE 58th conference on decision and control* (pp. 1607–1614). Nice, France: IEEE.
- Kim, J. W., Mehta, P. G., & Meyn, S. (2021). The conditional Poincaré inequality for filter stability. In *2021 IEEE 60th conference on decision and control* (pp. 1629–1636).
- Kim, J. W., Taghvaei, A., Chen, Y., & Mehta, P. G. (2021). Feedback particle filter for collective inference. *Foundations of Data Science*, 3(3), 543–561. <http://dx.doi.org/10.3934/fods.2021018>, URL: [/article/id/97f98e95-60ec-4ef7-be71-93835f30b466](https://arxiv.org/abs/2010.11337).
- Kim, J. W., Taghvaei, A., & Mehta, P. G. (2018). Derivation and extensions of the linear feedback particle filter based on duality formalisms. In *2018 IEEE conference on decision and control* (pp. 7188–7193). IEEE.
- Kuramoto, Y. (1975). Self-entrainment of a population of coupled non-linear oscillators. In *International symposium on mathematical problems in theoretical physics* (pp. 420–422). Springer.
- Kutschireiter, A., Surace, S. C., Sprekeler, H., & Pfister, J. P. (2017). Nonlinear Bayesian filtering and learning: a neuronal dynamics for perception. *Scientific Reports*, 7.
- Kwakernaak, H., & Sivan, R. (1972). *Linear optimal control systems*, vol. 1. Wiley-Interscience New York.
- Kwiatkowski, E., & Mandel, J. (2015). Convergence of the square root ensemble Kalman filter in the large ensemble limit. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1), 1–17.
- Laugesen, R. S., Mehta, P. G., Meyn, S. P., & Raginsky, M. (2015). Poisson's equation in nonlinear filtering. *SIAM Journal on Control and Optimization*, 53(1), 501–525. <http://dx.doi.org/10.1137/13094743X>.
- Le Gland, F., Monbet, V., & Tran, V. (2009). *Large sample asymptotics for the ensemble Kalman filter* (Ph.D. thesis), INRIA.
- Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. arXiv:1805.00909.
- Liggett, T. M. (1985). *Interacting particle systems*, vol. 2. Springer.
- Liu, Q., & Wang, D. (2016). Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in neural information processing systems* (pp. 2378–2386).
- Liu, C., Zhuo, J., Cheng, P., Zhang, R., Zhu, J., & Carin, L. (2018). Accelerated first-order methods on the wasserstein space for Bayesian inference. arXiv preprint [arXiv:1807.01750](https://arxiv.org/abs/1807.01750).
- Luo, X. (2019). Multivariate feedback particle filter via F-divergence and the well-posedness of its admissible control input. arXiv preprint [arXiv:1902.08745](https://arxiv.org/abs/1902.08745).
- Makkuva, A., Taghvaei, A., Oh, S., & Lee, J. (2020). Optimal transport mapping via input convex neural networks. In *International conference on machine learning* (pp. 6672–6681). PMLR.
- Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P. L., & Wainwright, M. J. (2020). Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. *Journal of Machine Learning Research*, 21(21), 1–51, URL: <http://jmlr.org/papers/v21/19-198.html>.
- Mandel, J., Cobb, L., & Beezley, J. D. (2011). On the convergence of the ensemble Kalman filter. *Applications of Mathematics*, 56(6), 533–541.
- Matsuura, Y., Ohata, R., Nakakuki, K., & Hirokawa, R. (2016). Suboptimal gain functions of feedback particle filter derived from continuation method. In *AIAA guidance, navigation, and control conference* (p. 1620). <http://dx.doi.org/10.2514/6.2016-1620>.
- Mehta, P. G., & Meyn, S. P. (2009). Q-learning and Pontryagin's minimum principle. In *Proceedings of the 48th IEEE conference on decision and control (CDC) held jointly with 2009 28th Chinese control conference* (pp. 3598–3605). IEEE.
- Meyn, S. (2022). *Control systems and reinforcement learning*. Cambridge University Press.
- Mitter, S. K., & Newton, N. J. (2003). A variational approach to nonlinear estimation. *SIAM Journal on Control and Optimization*, 42(5), 1813–1833.
- Mohammadi, H., Zare, A., Soltanolkotabi, M., & Jovanovic, M. R. (2019). Global exponential convergence of gradient methods over the nonconvex landscape of the linear quadratic regulator. In *2019 IEEE 58th conference on decision and control* (pp. 7474–7479). [ISSN: 2576-2370] <http://dx.doi.org/10.1109/CDC40024.2019.9029985>.
- Mohammadi, H., Zare, A., Soltanolkotabi, M., & Jovanović, M. R. (2022). Convergence and sample complexity of gradient methods for the model-free linear-quadratic regulator problem. *IEEE Transactions on Automatic Control*, 67(5), 2435–2450. <http://dx.doi.org/10.1109/TAC.2021.3087455>.
- Mortensen, R. E. (1968). Maximum-likelihood recursive nonlinear filtering. *Journal of Optimization Theory and Applications*, 2(6), 386–394. <http://dx.doi.org/10.1007/BF00925744>.
- Nualart, D., & Pardoux, É. (1988). Stochastic calculus with anticipating integrands. *Probability Theory and Related Fields*, 78(4), 535–581.
- Olmez, S. Y., Taghvaei, A., & Mehta, P. G. (2020). Deep FPF: Gain function approximation in high-dimensional setting. In *2020 59th IEEE conference on decision and control* (pp. 4790–4795). IEEE.
- Pathiraja, S., Reich, S., & Stannat, W. (2021). McKean-Vlasov SDEs in nonlinear filtering. *SIAM Journal on Control and Optimization*, 59(6), 4188–4215. <http://dx.doi.org/10.1137/20M1355197>.
- Pathiraja, S., & Stannat, W. (2021). Analysis of the feedback particle filter with diffusion map based approximation of the gain. *Foundations of Data Science*, 3(3), 615–645. <http://dx.doi.org/10.3934/fods.2021023>, URL: [/article/id/adaea096-8149-4632-8129-918680dd5871](https://arxiv.org/abs/2010.11337).
- Radhakrishnan, A., Devraj, A., & Meyn, S. (2016). Learning techniques for feedback particle filter design. In *Conference on decision and control* (pp. 648–653). IEEE, <http://dx.doi.org/10.1109/CDC.2016.7799106>.
- Radhakrishnan, A., & Meyn, S. (2018). Feedback particle filter design using a differential-loss reproducing kernel Hilbert space. In *2018 annual American control conference* (pp. 329–336). IEEE, <http://dx.doi.org/10.23919/ACC.2018.8431689>.
- Rahimian, H., & Mehrotra, S. (2019). Distributionally robust optimization: A review. arXiv preprint [arXiv:1908.05659](https://arxiv.org/abs/1908.05659).
- Rawlings, J. B., Mayne, D. Q., & Diehl, M. (2017). *Model predictive control: Theory, computation, and design*, vol. 2. WI: Nob Hill Publishing Madison.
- Rebeschini, P., Van Handel, R., et al. (2015). Can local particle filters beat the curse of dimensionality? *Annals of Applied Probability*, 25(5), 2809–2866.
- Reich, S. (2011). A dynamical systems framework for intermittent data assimilation. *BIT Numerical Analysis*, 51, 235–249. <http://dx.doi.org/10.1007/s10543-010-0302-4>.
- Reich, S. (2013). A nonparametric ensemble transform method for Bayesian inference. *SIAM Journal on Scientific Computing*, 35(4), A2013–A2024.
- Reich, S. (2019). Data assimilation: The Schrödinger perspective. *Acta Numerica*, 28, 635–711.
- Reich, S., & Cotter, C. (2015). *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press.
- Richemond, P. H., & Maginnis, B. (2017). On Wasserstein reinforcement learning and the Fokker–Planck equation. arXiv preprint [arXiv:1712.07185](https://arxiv.org/abs/1712.07185).
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., et al. (2020). Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839), 604–609. <http://dx.doi.org/10.1038/s41586-020-03051-4>.
- Shafieezadeh Abadeh, S., Nguyen, V. A., Kuhn, D., & Mohajerin Esfahani, P. M. (2018). Wasserstein distributionally robust Kalman filtering. *Advances in Neural Information Processing Systems*, 31.
- Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
- Sheldon, D. R., & Dietterich, T. (2011). Collective graphical models. *Advances in Neural Information Processing Systems*, 24.
- Singh, R., Haasler, I., Zhang, Q., Karlsson, J., & Chen, Y. (2020). Incremental inference of collective graphical models. *IEEE Control Systems Letters*, 5(2), 421–426.
- Snyder, C., Bengtsson, T., Bickel, P., & Anderson, J. (2008). Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12), 4629–4640.
- Spantini, A., Baptista, R., & Marzouk, Y. (2022). Coupling techniques for nonlinear ensemble filtering. *SIAM Review*, 64(4), 921–953.
- Stano, P. M. (2018). The collapse of sequential Bayesian estimator in two-target tracking problem. *Studia Informatica: Systems and Information Technology*, 1.
- Stano, P. M., Tilton, A. K., & Babuska, R. (2014). Estimation of the soil-dependent time-varying parameters of the hopper sedimentation model: The FPF versus the BPF. *Control Engineering Practice*, 24, 67–78. <http://dx.doi.org/10.1016/j.conengprac.2013.11.005>.
- Strogatz, S. H. (2000). From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1–4), 1–20.
- Su, Z., Zhou, Z. X., Yu, M. H., Yuan, W., & Fu, J. Q. (2019). Online estimation of soil grain diameter during dredging of Hopper dredger using continuous-discrete feedback particle filter. *Sensors and Materials*, 31(3), 953–968.
- Surace, S. C., Kutschireiter, A., & Pfister, J. P. (2019). How to avoid the curse of dimensionality: scalability of particle filters with and without importance weights. *SIAM Review*, 61(1), 79–91. <http://dx.doi.org/10.1137/17M1125340>.
- Sutton, R. S., & Barto, A. G. (2018). *Adaptive computation and machine learning, Reinforcement learning: An introduction* (2nd ed.). (p. xxii+526). Cambridge, MA: MIT Press, 3889951.
- Taghvaei, A., De Wiljes, J., Mehta, P. G., & Reich, S. (2018). Kalman filter and its modern extensions for the continuous-time nonlinear filtering problem. *Journal of Dynamic Systems, Measurement, and Control*, 140(3), Article 030904. <http://dx.doi.org/10.1115/1.4037780>.
- Taghvaei, A., & Hosseini, B. (2022). An optimal transport formulation of Bayes' law for nonlinear filtering algorithms. arXiv preprint [arXiv:2203.11869](https://arxiv.org/abs/2203.11869).
- Taghvaei, A., & Mehta, P. G. (2016). An optimal transport formulation of the linear feedback particle filter. In *American control conference* (pp. 3614–3619). IEEE, <http://dx.doi.org/10.1109/acc.2016.7525474>.
- Taghvaei, A., & Mehta, P. G. (2018). Error analysis of the stochastic linear feedback particle filter. In *2018 IEEE conference on decision and control* (pp. 7194–7199). IEEE.
- Taghvaei, A., & Mehta, P. (2019). Accelerated flow for probability distributions. In *International conference on machine learning* (pp. 6076–6085). PMLR.
- Taghvaei, A., & Mehta, P. G. (2020). An optimal transport formulation of the ensemble Kalman filter. *IEEE Transactions on Automatic Control*, 66(7), 3052–3067.
- Taghvaei, A., Mehta, P. G., & Georgiou, T. T. (2022). Optimality vs stability trade-off in ensemble Kalman filters. *IFAC-PapersOnLine*, 55(30), 335–340.
- Taghvaei, A., Mehta, P. G., & Meyn, S. P. (2020). Diffusion map-based algorithm for gain function approximation in the feedback particle filter. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3), 1090–1117.

- Tilton, A. K., Ghiotto, S., & Mehta, P. G. (2013). A comparative study of nonlinear filtering techniques. In *Proc. 16th int. conf. on inf. fusion* (pp. 1827–1834).
- Tilton, A. K., Hsiao-Wecksler, E. T., & Mehta, P. G. (2012). Filtering with rhythms: Application to estimation of gait cycle. In *2012 American control conference* (pp. 3433–3438). <http://dx.doi.org/10.1109/ACC.2012.6315665>.
- Todorov, E. (2007). Linearly-solvable Markov decision problems. In B. Schölkopf, J. Platt, & T. Hoffman (Eds.), *Advances in neural information processing systems*, vol. 19. MIT Press, URL: <https://proceedings.neurips.cc/paper/2006/file/d806ca13ca3449af72a1ea5aedbed26a-Paper.pdf>.
- Todorov, E. (2008). General duality between optimal control and estimation. In *2008 47th IEEE conference on decision and control* (pp. 4286–4292).
- Tong, X. T., Majda, A. J., & Kelly, D. (2016). Nonlinear stability and ergodicity of ensemble based Kalman filters. *Nonlinearity*, 29(2), 657.
- Toussaint, M. (2009). Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning* (pp. 1049–1056). New York, NY, USA: Association for Computing Machinery, <http://dx.doi.org/10.1145/1553374.1553508>.
- Tu, S., & Recht, B. (2019). The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. In *Conference on learning theory* (pp. 3036–3083). PMLR, [ISSN: 2640-3498] URL: <http://proceedings.mlr.press/v99/tu19a.html>.
- Van Leeuwen, P. J., & Evensen, G. (1996). Data assimilation and inverse methods in terms of a probabilistic formulation. *Monthly Weather Review*, 124(12), 2898–2913.
- Vijayakumar, S., Rawlik, K., & Toussaint, M. (2013). On stochastic optimal control and reinforcement learning by approximate inference. In N. Roy, P. Newman, & S. Srinivasa (Eds.), *Robotics: Science and systems VIII* (pp. 353–360).
- Villani, C. (2009). *Optimal transport: Old and new*, vol. 338. Springer.
- Wang, Y., Wang, X., & Cui, N. (2021). Quantized feedback particle filter for unmanned aerial vehicles tracking with quantized measurements. *Proceedings of the Institution of Mechanical Engineers, Part G (Journal of Aerospace Engineering)*, 235(2), 257–270.
- Whitaker, J., & Hamill, T. M. (2002). Ensemble data assimilation without perturbed observations. *Monthly Weather Review*, 130(7), 1913–1924. [http://dx.doi.org/10.1175/1520-0493\(2002\)130<1913:edawpo>2.0.co;2](http://dx.doi.org/10.1175/1520-0493(2002)130<1913:edawpo>2.0.co;2).
- Xiong, J. (2008). *Oxford graduate texts in mathematics: vol. 18, An introduction to stochastic filtering theory*. Oxford University Press.
- Yang, T., Blom, H. A., & Mehta, P. G. (2013). Interacting multiple model-feedback particle filter for stochastic hybrid systems. In *52nd IEEE conference on decision and control* (pp. 7065–7070). IEEE.
- Yang, T., Blom, H., & Mehta, P. G. (2014). The continuous-discrete time feedback particle filter. In *2014 American control conference* (pp. 648–653). IEEE.
- Yang, T., Huang, G., & Mehta, P. G. (2012). Joint probabilistic data association-feedback particle filter for multiple target tracking applications. In *2012 American control conference* (pp. 820–826). IEEE.
- Yang, T., Laugesen, R. S., Mehta, P. G., & Meyn, S. P. (2016). Multivariable feedback particle filter. *Automatica*, 71, 10–23. <http://dx.doi.org/10.1016/j.automatica.2016.04.019>.
- Yang, T., & Mehta, P. G. (2018). Probabilistic data association-feedback particle filter for multiple target tracking applications. *Journal of Dynamic Systems, Measurement, and Control*, 140(3).
- Yang, T., Mehta, P. G., & Meyn, S. P. (2011a). Feedback particle filter with mean-field coupling. In *2011 50th IEEE conference on decision and control and European control conference* (pp. 7909–7916). IEEE.
- Yang, T., Mehta, P. G., & Meyn, S. P. (2011b). A mean-field control-oriented approach to particle filtering. In *Proceedings of the 2011 American control conference* (pp. 2037–2043). IEEE.
- Yang, T., Mehta, P. G., & Meyn, S. P. (2013). Feedback particle filter. *IEEE Transactions on Automatic Control*, 58(10), 2465–2480.
- Yang, T., Mehta, P. G., & Meyn, S. P. (2015). Feedback particle filter for a continuous-time Markov chain. *IEEE Transactions on Automatic Control*, 61(2), 556–561.
- Yin, H., Mehta, P. G., Meyn, S. P., & Shanbhag, U. V. (2011). Synchronization of coupled oscillators is a game. *IEEE Transactions on Automatic Control*, 57(4), 920–935.
- Zhang, R., Chen, C., Li, C., & Carin, L. (2018). Policy optimization as Wasserstein gradient flows. *arXiv preprint arXiv:1808.03030*.
- Zhang, C., Taghvaei, A., & Mehta, P. G. (2016a). Attitude estimation with feedback particle filter. In *2016 IEEE 55th conference on decision and control* (pp. 5440–5445). IEEE.
- Zhang, C., Taghvaei, A., & Mehta, P. G. (2016b). Feedback particle filter on matrix Lie groups. In *2016 American control conference* (pp. 2723–2728). IEEE.
- Zhang, C., Taghvaei, A., & Mehta, P. G. (2017a). Attitude estimation of a wearable motion sensor. In *2017 American control conference* (pp. 4570–4575). IEEE.
- Zhang, C., Taghvaei, A., & Mehta, P. G. (2017b). Feedback particle filter on Riemannian manifolds and matrix Lie groups. *IEEE Transactions on Automatic Control*, 63(8), 2465–2480.
- Zhang, C., Taghvaei, A., & Mehta, P. G. (2019). A mean-field optimal control formulation for global optimization. *IEEE Transactions on Automatic Control*, 64(1), 279–286.
- Zheng, H., Yu, M., & Yuan, W. (2019). Parameter identification of ship model based on feedback particle filter. *Chinese Journal of Ship Research*, 14(3), 158–162.