Dedicated to the memory of Ari Arapostathis

# Controlled Interacting Particle Algorithms for Simulation-based Reinforcement Learning

Anant A. Joshi<sup>a</sup>, Amirhossein Taghvaei<sup>b</sup>, Prashant G. Mehta<sup>a</sup>, Sean P. Meyn<sup>c</sup>

 <sup>a</sup>Coordinated Science Laboratory and the Department of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign
 <sup>b</sup>William E. Boeing Department of Aeronautics and Astronautics at University of Washington Seattle
 <sup>c</sup>Department of Electrical and Computer Engineering at University of Florida Gainesville

## Abstract

This paper is concerned with optimal control problems for control systems in continuous time, and interacting particle system methods designed to construct approximate control solutions. Particular attention is given to the linear quadratic (LQ) control problem. There is a growing interest in revisiting this classical problem, in part due to the successes of reinforcement learning (RL). The main question of this body of research (and also of our paper) is to approximate the optimal control law *without* explicitly solving the Riccati equation. A novel simulation-based algorithm, namely a dual ensemble Kalman filter (EnKF), is introduced. The algorithm is used to obtain formulae for optimal control, expressed entirely in terms of the EnKF particles. An extension to the nonlinear case is also presented. The theoretical results and algorithms are illustrated with numerical experiments.

#### 1. Introduction

The field of reinforcement learning (RL) is concerned with optimal control, to design a policy for a dynamical system that minimizes some performance criterion. All of the standard choices are treated in the literature: discounted cost, finite time-horizon, and average cost. 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 based on input-output measurements.

Preprint submitted to Systems and Control Letters

July 11, 2022

There are two standard approaches to obtain an algorithm for this purpose: (i) critic methods, in which a value function is approximated within a parameterized family, and the policy is obtained as a functional of the approximation, and (ii) actor methods in which a parameterized family of policies is given, and the algorithm is designed to obtain the best policy within this family.

In popular media, RL is often described as an "agent" that learns an approximately optimal policy based on interactions with the environment. Important examples of this ideal 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 [1].

This paper focuses on model-based RL in which the model is available only in the form of a simulator. The proposed approach is novel, drawing on mean-field techniques similar to those appearing in state estimation (data assimilation) in high dimension. It is likely that the concepts will lead to new approaches for online RL—see directions for future research in the conclusions.

We consider the finite-horizon optimal control problem,

$$\min_{u} \quad J(u) = \int_{0}^{T} \left( \frac{1}{2} |c(x_t)|^2 + \frac{1}{2} u_t^{\top} R u_t \right) \, \mathrm{d}t + g(x_T) \tag{1a}$$

subject to: 
$$\dot{x}_t = a(x_t) + b(x_t)u_t, \ x_0 = x$$
 (1b)

where  $x_t \in \mathbb{R}^d$  is the state at time t, and  $u = \{u_t \in \mathbb{R}^m : 0 \le t \le T\}$ is the control input. The functions  $a(\cdot), b(\cdot), c(\cdot), g(\cdot)$  are continuously differentiable  $(C^1)$ , and the control penalty matrix positive definite,  $R \succ 0$ .

In the linear quadratic (LQ) setting the model is linear (a(x) = Ax and b(x) = B) and the cost function is quadratic  $(c(x) = Cx \text{ and } g(x) = x^{\intercal}P_Tx)$ . The infinite-time horizon  $(T = \infty)$  case is referred to as the linear quadratic regulator (LQR) problem. Although it is a historical problem, LQR has been the subject of recent research interest in the control community. The goals of this research are much like ours: design simulations for the purposes of learning the optimal control policy.

The proposed solution involves construction of N stochastic processes  $\{Y_t^i \in \mathbb{R}^d : 0 \le t \le T, 1 \le i \le N\}$  where the *i*-th particle evolves according

to a stochastic differential equation (SDE) of the form,

$$dY_t^i = \underbrace{a(Y_t^i) dt + b(Y_t^i) dv_t^i}_{i-\text{th copy of model (1b)}} + \mathcal{U}_t^i dt, \quad 0 \le t \le T$$
(2)

where the input  $v^i = \{v_t^i \in \mathbb{R}^m : 0 \leq t \leq T\}$  and the *data assimilation* process  $\mathcal{U}^i = \{\mathcal{U}_t^i \in \mathbb{R}^d : 0 \leq t \leq T\}$  is obtained as part of the RL design. The goal is to design these processes so that the empirical distribution of the N particles at time t approximates a smooth density  $p_t$  (for the  $N = \infty$ mean-field limit), encoding the optimal policy as follows:

$$\Phi_t^*(x) = R^{-1} b^\top(x) \nabla \log p_t(x) \tag{3}$$

where  $\nabla$  denotes the gradient operator. In the infinite-horizon case, a stationary policy is obtained by letting  $T \to \infty$ .

We make the following assumption:

- **Assumption 1.** 1. Functions  $f(x, \alpha) = a(x) + b(x)\alpha$  and c(x) 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 (2). In particular, this requires an ability to add additional inputs outside the control channel.

Part 1 of the assumption is standard for any RL algorithm. Part 2 is not too restrictive for the following reasons: In physical systems, one typically is able to assess relative costs for different control inputs (actuators). For the LQR problem, under certain technical conditions, the optimal policy is stationary and does not depend upon the choice of  $P_T$ . A 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.

A motivation comes from data assimilation applications such as weather prediction and geosciences where Assumption 1 is standard. The ensemble Kalman Filter (EnKF) is a particle system method which serves as a workhorse in these applications [2, 3]. The computational complexity of the EnKF is  $\mathcal{O}(Nd)$  where d is dimension, and N is the number of particles, with  $N \ll d$  typical in these applications. Part of the tremendous success of the EnKF in these domains is that it works directly with the simulator. Multiple copies are run in a Monte-Carlo manner where the data assimilation process is used to assimilate the most recent measurement.

The goal of the research summarized here is to create approximation techniques with similar success for applications in control.

## 1.1. Contributions of this paper

In order to elucidate these new ideas as clearly as possible, the main focus of this paper is on the linear quadratic (LQ) problem. This also allows us to highlight and contrast our work with recent developments. The algorithm (2) for the LQ problem is presented first in Sec. 2 before describing its nonlinear extension in Sec. 3. The details of the original contributions of our work are as follows:

1. For the LQ problem, the proposed algorithm is an ensemble Kalman filter (EnKF) referred to here as the *dual EnKF*. The mean-field limit  $(N = \infty)$  of the dual EnKF is shown to be exact (Prop. 1). For the finite-N algorithm, an error bound on the approximation is obtained under additional assumptions on the model (see (11)). An extensive discussion is included in Sec. 2.4 to situate the algorithm in the RL landscape. In particular, it is shown that (i) the process  $v^i$  implements the exploration step of RL whereby the cheap control directions are explored more; and (ii) the process  $\mathcal{U}^i$  implements the value iteration step of RL.

2. For the nonlinear problem (1), a dual algorithm is presented to approximate the Hamilton-Jacobi-Bellman (HJB) equation. The algorithm requires a solution of a (linear) Poisson equation that is far more easily approximated. It is shown that the dual EnKF algorithm for the LQ problem is a special case in which the Poisson equation admits an analytical solution.

**3.** A numerical comparison of the dual EnKF algorithm against the state-ofthe-art is described for benchmark examples. It is shown that the proposed algorithm can be up to two orders of magnitude more computationally efficient (Fig. 4). Scalings with respect to both the number of particles Nand the problem dimension d are numerically illustrated and compared with analytical bounds (Fig. 3).

#### 1.2. Literature review

There are three areas of prior work that are related to the subject of this paper: (i) RL algorithms for the LQR problem; (ii) EnKF and related control-type algorithms for data assimilation; and (iii) duality theory between optimal control and estimation.

(i) **RL for LQR:** The LQ problem has a rich and storied history in modern control theory going back to the original work on the subject [4]. Its solution requires solving a Riccati equation – the differential Riccati equation (DRE) in finite-horizon settings or the algebraic Riccati equation (ARE) in the infinite-horizon setting. There is a large body of literature devoted to the analytical study of the Riccati equations [5] and specialized numerical techniques have been developed to efficiently compute the solution [6].

There are two issues which makes the LQ and related problems a topic of recent research interest: (i) In high-dimensions, the matrix-valued nature of the DRE or ARE means that any algorithm is  $\mathcal{O}(d^2)$  in the dimension d of the state-space, and (ii) the model parameters may not be explicitly available to write down the DRE (or the ARE) let alone solve it. The latter is a concern, e.g., when the model exists only in the form of a black-box numerical simulator.

These two issues have motivated the recent research on the infinitehorizon linear quadratic regulator (LQR) problem [7, 8, 9, 10, 11]. Of particular interest are policy gradient type algorithms that seek to bypass solving an ARE by directly searching over the space of stabilizing gain matrices. The algorithms are of iterative type where each iteration requires a policy evaluation step (using N simulations much like (2)). This step is used to estimate a gradient which is then used to obtain a improved policy based on a gradient descent procedure. Global convergence rate estimates are established for both discrete-time [7, 12] and continuous-time [11] settings of the LQR problem. Extensions to the  $H_{\infty}$  regularized LQR [13] and Markov jump linear systems [14] have also been carried out. In the recent thesis [15, Chapter 4], finite horizon extensions are considered under additional assumptions.

Additional comparison with this prior work appears in Sec. 2.4 and numerical comparison is in Sec. 4.

(ii) EnKF for data assimilation: Although novel for RL, the proposed algorithms are inspired by the data assimilation (nonlinear filtering) literature [3]. During the past decade, a key breakthrough in the data assimilation theory and its applications is the design of controlled interactions between particles (such as  $\mathcal{U}_t^i$  in (2)) to approximate the solution of the nonlinear filtering problem; c.f., [16] and references therein. Such an approach is in contrast to the traditional importance sampling type approaches which suf-

fer from issues such as particle degeneracy [17]. Two well known examples of the controlled interacting particle systems are the ensemble Kalman filter (EnKF) and the feedback particle filter (FPF). The EnKF is an exact algorithm for the linear Gaussian filtering problem [18, 19] while the FPF is an exact algorithm for the nonlinear non-Gaussian case [20]. The two major algorithmic contributions of this work represent the optimal control (dual) counterparts of the EnKF and the FPF.

Notably, EnKF is a workhorse in data assimilation applications such as weather prediction where models are simulation-based and high-dimensional [2, 3]. As noted above, these two issues have also motivated much of recent work on the LQ problem in the control community.

(iii) Duality: The formula (3) for the optimal policy is a consequence of the so-called log transformation. The transformation relates the value function of an optimal control problem to the posterior distribution of the dual optimal estimation problem [21]. Duality is an old subject [22, Chapter 7.5],[23, Chapter 15]. In recent years, there has been renewed interest in duality for algorithm design. In much of the classical literature on the subject, duality was used to obtain optimal control algorithms for solving estimation problems [24]. Although it remains an important theme [25], some of the more recent work has been in the opposite direction: to solve optimal control problems by using sampling techniques [26]. Our work fits within this latter body of work.

A salient aspect of this paper is a detailed comparison with literature appearing in each of the three main sections after technical details have been presented.

#### 1.3. Paper outline

The outline of the remainder of this paper is as follows. The LQ optimal control problem and its dual EnKF solution is described in Sec. 2. The nonlinear extension and its connection to duality appears in Sec. 3. The algorithms are illustrated with numerical examples in Sec. 4. The proofs appear in the Appendix.

**Notation:**  $\mathcal{N}(m, \Sigma)$  is a Gaussian probability distribution with mean mand covariance  $\Sigma$ . The notation  $\Sigma \succ 0$  is used when the matrix  $\Sigma$  is positive definite. The  $n \times n$  identity matrix is denoted  $\mathcal{I}_n$ . The trace of a matrix is denoted by  $\text{Tr}(\cdot)$ . For a smooth function  $f : \mathbb{R}^d \to \mathbb{R}, \nabla f(x) = \begin{bmatrix} \frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_d} \end{bmatrix}^{\top}$ denotes the gradient of f, and  $\nabla^2 f(x) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_m \partial x_n}(x) \end{bmatrix}_{n,m=1}^d$  denotes the Hessian matrix. For a smooth vector-field  $v : \mathbb{R}^d \to \mathbb{R}^d$ ,  $\nabla \cdot v(x) = \sum_{n=1}^d \frac{\partial v_n}{\partial x_n}(x)$ denotes the divergence. And for a smooth tensor  $D : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ ,  $\nabla \cdot D(x)$ is a vector field whose *m*-th component is  $\sum_{n=1}^d \frac{\partial D_{mn}}{\partial x_n}(x)$ , and  $\nabla^2 \cdot D(x) = \sum_{n,m=1}^d \frac{\partial^2 D_{mn}}{\partial x_n \partial x_m}(x)$ .

# 2. LQ problem

The finite-horizon linear quadratic (LQ) optimal control problem is a special case of (1) as follows:

$$\min_{u} \quad J(u) = \int_{0}^{T} \frac{1}{2} |Cx_{t}|^{2} + \frac{1}{2} u_{t}^{\top} R u_{t} \, \mathrm{d}t + \frac{1}{2} x_{T}^{\top} P_{T} x_{T}$$
(4a)

subject to: 
$$\dot{x}_t = Ax_t + Bu_t, \quad x_0 = x$$
 (4b)

It is assumed that (A, B) is controllable, (A, C) is observable, and the matrices  $P_T$ ,  $R \succ 0$ . 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

$$\phi_t(x) = K_t x$$
 where  $K_t = -R^{-1}B^\top P_t$ 

is the optimal gain matrix and  $P_t$  is a solution of the backward (in time) DRE

$$-\frac{\mathrm{d}}{\mathrm{d}t}P_t = A^{\mathsf{T}}P_t + P_t A + C^{\mathsf{T}}C - P_t B R^{-1} B^{\mathsf{T}}P_t, \quad P_T \text{ (given)}$$
(5)

The ARE is obtained by setting the left-hand side to 0. As  $T \to \infty$ , for each fixed time  $t, P_t \to P^{\infty}$ , exponentially fast [27, Remark 2.1], where  $P^{\infty} \succ 0$  is the unique positive-definite solution of the ARE, and therefore the optimal gain converges,  $K_t \to K^{\infty} := -R^{-1}B^{\top}P^{\infty}$ . Approximation of the LQR gain  $K^{\infty}$  is a goal in recent RL research [7, 11].

#### 2.1. Main contribution: Dual EnKF algorithm

Under the assumptions of this paper,  $P_t \succ 0$  for  $0 \le t \le T$  whenever  $P_T \succ 0$  [28, Sec. 24]. Set  $S_t = P_t^{-1}$ . It is readily verified that  $S_t$  also solves a backward DRE

$$\frac{\mathrm{d}}{\mathrm{d}t}S_t = AS_t + S_t A^{\top} - BR^{-1}B^{\top} + S_t C^{\top}CS_t, \quad S_T = P_T^{-1} \tag{6}$$

Our objective is to approximate  $S_t$  using simulations. The proposed 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  $Y = \{Y_t \in \mathbb{R}^d : 0 \le t \le T\}$  as a solution of the following backward (in time) stochastic differential equation (SDE):

$$dY_t = AY_t dt + B d\dot{\eta}_t + \frac{1}{2}\bar{S}_t C^\top (CY_t + C\bar{n}_t) dt, \ 0 \le t < T$$
$$Y_T \stackrel{d}{=} \mathcal{N}(0, S_T)$$
(7)

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

$$\bar{n}_t := \mathsf{E}[Y_t], \quad \bar{S}_t := \mathsf{E}[(Y_t - \bar{n}_t)(Y_t - \bar{n}_t)^{\top}]$$
 (8)

The process Y is an example of a mean-field process because its evolution depends upon the statistics  $(\bar{n}_t, \bar{S}_t)$  of the process. An SDE of this type is called a McKean-Vlasov SDE. The meaning of the backward arrow on  $d\bar{\eta}$  in (7) is that the SDE is simulated backward in time starting from the terminal condition specified at time t = T. The reader is referred to [29, Sec. 4.2] for the definition of the backward Itô-integral.

The mean-field process is useful because of the following proposition whose proof is included in Appendix A.

**Proposition 1.** The solution to the SDE (7) is a Gaussian stochastic process, in which the mean and covariance of  $Y_t$  are given by

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

Consequently,  $X_t := \bar{S}_t^{-1}(Y_t - \bar{n}_t)$  is also a Gaussian satisfying

$$\mathsf{E}(X_t) = 0, \quad \mathsf{E}(X_t X_t^{\top}) = P_t, \quad 0 \le t \le T$$

The significance of Prop. 1 is that the optimal control policy  $\phi_t(\cdot)$  can now be obtained in terms of the statistics of the random variable  $X_t$ . Specifically, we have the following two cases:

(i) If the matrix B is explicitly known then the optimal gain matrix

$$K_t = -R^{-1}B^{\top}\mathsf{E}(X_tX_t^{\top})$$

(ii) If B is unknown, define the Hamiltonian (the continuous-time counterpart of the Q-function [30]):

$$H(x,\alpha,t) := \underbrace{\frac{1}{2}|Cx|^2 + \frac{1}{2}\alpha^{\top}R\alpha}_{\text{cost function}} + x^{\top}\mathsf{E}(X_tX_t^{\top})\underbrace{(Ax + B\alpha)}_{\text{model (4b)}}$$

from which the optimal control law is obtained as

$$\phi_t(x) = \operatorname*{arg\,min}_{\alpha \in \mathbb{R}^m} H(x, \alpha, t)$$

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 (4b).

Step 2. Finite-N approximation: The mean-field process is empirically approximated by simulating a system of controlled interacting particles according to

$$dY_t^i = \underbrace{AY_t^i dt + B d\tilde{\eta}_t^{-i}}_{i-\text{th copy of model (4b)}} + \underbrace{S_t^{(N)} C^{\top} \left(\frac{CY_t^i + Cn_t^{(N)}}{2}\right)}_{\text{data assimilation process}} dt, \qquad (9)$$

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

 $\eta^i$  is an i.i.d copy of  $\eta,\, n_t^{(N)} = N^{-1} \sum_{i=1}^N Y_t^i,$  and

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

The data assimilation process has a linear feedback control structure where  $S_t^{(N)}C^{\top}$  is the Kalman gain matrix and  $\frac{1}{2}(CY_t^i + Cn_t^{(N)})$  is the state feedback term similar to the error in the FPF [31]. The process serves to couple the particles. Without it, the particles are independent of each other.

The finite-N system (9) 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:

(i) If the matrix B is explicitly known then

$$K_t^{(N)} = -\frac{1}{N-1} \sum_{i=1}^N R^{-1} (B^\top X_t^i) (X_t^i)^\top$$
(10)

(ii) If B is unknown, define the Hamiltonian

$$H^{(N)}(x,\alpha,t) := \underbrace{\frac{1}{2}|Cx|^2 + \frac{1}{2}\alpha^\top R\alpha}_{\text{cost function}} + \frac{1}{N-1}\sum_{i=1}^N (x^\top X^i_t)(X^i_t)^\top \underbrace{(Ax + B\alpha)}_{\text{model (4b)}}$$

from which the optimal control policy is approximated as

$$\Phi_t^{(N)}(x) = \operatorname*{arg\,min}_{a \in \mathbb{R}^m} H^{(N)}(x, a, t)$$

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, mqueries of  $H^{(N)}(x, \cdot, t)$  are sufficient to compute  $\Phi_t^{(N)}(x)$ .

The overall algorithm including its numerical approximation appears in Appendix E.

#### 2.2. Remarks

The following remarks are included to help provide an intuitive explanation to various aspects of the dual EnKF.

1. Representation. In designing any RL algorithm, the first issue is representation of the unknown value function ( $P_t$  in the linear case). Our novel idea is to represent  $P_t$  is in terms of statistics (variance) of the particles. Such a representation is fundamentally 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 (4b) are simulated in parallel where the terms on the right hand-side of (9) have the following intuitive interpretations:

1. Dynamics: The first term " $AY_t^i dt$ " on the right-hand side of (9) is simply a copy of uncontrolled dynamics in the model (4b).

2. Control: The second term is the control input U for the *i*-th particle, specified as a white noise process with covariance  $R^{-1}$ . One may interpret this as an approach to exploration whereby cheaper control directions are explored more.

While there are similarities with traditional approaches to RL, the novelty comes from the data assimilation process that represents an original contribution.

**3.** Arrow of time. The particles are simulated backward – from terminal time t = T to initial time t = 0. This is consistent with the dynamic programming (DP) equation which also proceeds backward in time.

#### 2.3. Convergence and error analysis

The mean-field process (7) represents the mean-field limit of the finite-N system (9), as the number of particles  $N \to \infty$ . The convergence analysis is a challenging problem but impressive progress has been made in some groundbreaking work appearing in recent years [32, 33]. In Appendix B, under certain additional assumptions on system matrices, the following error bound is derived:

$$\mathsf{E}[\|S_t^{(N)} - \bar{S}_t\|_F] \le \frac{C_1}{\sqrt{N}} + C_2 e^{-2\lambda(T-t)} \mathsf{E}[\|S_T^{(N)} - \bar{S}_T\|_F],$$
(11)

where  $C_1, C_2$  are time-independent positive constants, and  $|| \cdot ||_F$  denotes Frobenius norm for matrices. The proof largely follows the techniques developed in [33].

#### 2.4. Comparison to literature

Function approximation: Classical RL algorithms for the LQR problem are based on a linear function approximation, using quadratic basis functions, of the value function or the Q-function [34, 35, 36]. The basic idea is to run the system for a time horizon T, and successively update an estimate of the parameters based on new data collected, using a least-squares approximation. Convergence guarantees typically require (i) a persistence of excitation condition, see e.g. [35, Equation (9)], [37, Remark 3, Page 173] and (ii) use of the on-policy methods whereby the parameters are learned for a given fixed policy (which is subsequently improved), see e.g. [34, Page 299]. For the deterministic LQR problem, the persistence of excitation condition is difficult to justify using on-policy RL methods. These limitations have spurred recent research on the LQR problem. **Policy gradient algorithms:** An inspiration for our work comes from the pioneering contributions of [11] and [7] who consider the infinite-horizon LQR objective ((4) with  $T = \infty$ ). With  $x_0$  drawn from a given initial distribution  $\mathcal{D}$ , and control policies restricted to the linear form  $u_t = Kx_t$ , the optimal control problem reduces to the finite-dimensional static optimization problem:

$$K^{\star} = \underset{K}{\operatorname{arg\,min}} J(K) = \mathsf{E}\left(\int_{0}^{\infty} x_{t}^{\top} Q x_{t} + u_{t}^{\top} R u_{t} \, dt\right) \tag{12}$$

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 [38].

In a technical tour de force, a Lyapunov function is obtained to carry out convergence analysis of the approximate gradient descent algorithm. The result is surprising because the problem is non-convex in K. Error bounds are obtained to quantify the effect of finite T and the finite number of iterations of the gradient descent algorithm. The number of particles  $N_g$  is of the order of the dimension of the system [12, Section VIII].

The trade-off between our algorithm and this prior work is as follows: While policy optimization methods require multiple iterations with a small number  $N_g$  of particles, the EnKF requires *only* a single iteration with relatively larger number N of particles. Using the EnKF, it is not necessary to have a stabilizing initial gain.

For a quantitative comparison, consider using the EnKF algorithm to approximate the infinite-horizon optimal gain (or equivalently the solution to the algebraic Ricatti equation). Choosing t = 0 in (11), the error is smaller than  $\varepsilon$  if the number of particles  $N > O(\frac{1}{\varepsilon^2})$  and the simulation time  $T > O(\log(\frac{1}{\varepsilon}))$ , while the iteration number is one. This is compared with policy optimization approach in [7] where the number of particles and the simulation time scales polynomially with  $\varepsilon$ , while the number of iterations scale as  $O(\log(\frac{1}{\varepsilon}))$ . This result is later refined in [11] where the required number of particles and the simulation time are shown to be O(1) and  $O(\log(\frac{1}{\varepsilon}))$ respectively (although this result is valid with probability that approaches zero as the number of iterations grow [11, Thm. 3].).

The overall comparison between the three algorithms appears in Sec. 4.

Algorithm	particles/samples	simulation time	iterations
EnKF	$O(\frac{1}{\varepsilon^2})$	$O(\log(\frac{1}{\varepsilon}))$	1
[7]	poly $\left(\frac{1}{\varepsilon}\right)$	poly $\left(\frac{1}{\varepsilon}\right)$	$O(\log(\frac{1}{\varepsilon}))$
[11]	O(1)	$O(\log(\frac{1}{\varepsilon}))$	$O(\log(\frac{1}{\varepsilon}))$

Table 1: Computational complexity comparison of the algorithms to achieve  $\varepsilon$  error in approximating the infinite-horizon LQR optimal gain.

#### 3. Nonlinear extensions

We return to the nonlinear optimal control problem (1) in Sec. 1. Its solution is obtained using a standard DP argument.

**Dynamic programming:** For  $t \in (0, T)$ , the value function

$$v_t(x) := \min_{\{u_s: t \le s \le T\}} \int_t^T \left(\frac{1}{2} |c(x_s)|^2 + \frac{1}{2} u_s^\top R u_s\right) \, \mathrm{d}s + g(x_T) \tag{13}$$

From the DP optimality principle, the value function solves the HJB equation

$$\frac{\partial v_t}{\partial t} + \frac{1}{2}c^2 - \frac{1}{2}\nabla v_t^\top D\nabla v_t + a^\top \nabla v_t = 0, \quad v_T(x) = g(x), \quad x \in \mathbb{R}^d$$
(14)

where  $D(x) := b(x)R^{-1}b^{\top}(x)$ , and the optimal control input is of the state feedback form  $u_t = \phi_t(x_t)$  where

$$\Phi_t(x) = -R^{-1}b^{\top}(x)\nabla v_t(x), \quad x \in \mathbb{R}^d$$
(15)

is the optimal control policy. For the LQ special case, the value function  $v_t(x) = \frac{1}{2}x^{\top}P_tx$  is quadratic and the HJB equation (14) reduces to the DRE (5) for the matrix  $P_t$ .

In the following, a mean-field process is introduced to solve the HJB equation based on the use of a log transformation.

Log transformation: Define a probability density as

$$p_t(x) := \frac{\exp(-v_t(x))}{\int \exp(-v_t(x)) \, \mathrm{d}x}, \quad x \in \mathbb{R}^d$$

In Appendix C, it is shown that the density solves a backward nonlinear

PDE:

$$\frac{\partial p_t}{\partial t} = p_t(h_t - \hat{h}_t) - \nabla \cdot (p_t(a - \nabla \cdot D)) - \frac{1}{2} \nabla^2 \cdot (p_t D)$$
$$p_T(x) = \frac{\exp(-g(x))}{\int \exp(-g(x)) \, \mathrm{d}x}, \quad x \in \mathbb{R}^d$$
(16)

where

$$h_t(x) := \frac{1}{2} |c(x)|^2 + (\nabla \cdot a)(x) - \frac{1}{2} \nabla^2 \cdot D(x) + \frac{1}{2} \operatorname{Tr} \left( (D(x)) \nabla^2 \log(p_t(x)) \right)$$

and  $\hat{h}_t := \int h_t(x) p_t(x) \, \mathrm{d}x.$ 

Our objective is to design simulations to sample from  $p_t$ . As in the LQ case, the construction proceeds in two steps: (i) definition of an exact mean-field process and (ii) its finite-N approximation.

**Mean-field process:** A mean-field process  $Y = \{Y_t \in \mathbb{R}^d : 0 \le t \le T\}$  is defined as follows:

$$dY_t = a(Y_t) dt + b(Y_t) d\bar{\eta}_t + \nabla \cdot D(Y_t) dt + \mathcal{V}_t(Y_t) dt,$$
  

$$Y_T \stackrel{d}{=} p_T$$
(17)

where  $\eta := \{\eta_t \in \mathbb{R}^m : 0 \le t \le T\}$  is a W.P. with covariance  $R^{-1}$ , and  $\mathcal{V}_t(\cdot)$  is a vector-field that solves the first order linear PDE

$$-\frac{1}{\bar{p}_t(x)}\nabla \cdot (\bar{p}_t(x)\mathcal{V}_t(x)) = (h_t(x) - \bar{h}_t), \quad \forall \ x \in \mathbb{R}^d$$
(18)

where  $\bar{h}_t := \int h_t(x)\bar{p}_t(x) \, dx$  and  $\bar{p}_t$  is the density of  $Y_t$  at time t.

The following proposition relates the density  $\bar{p}_t(x)$  of the mean-field process and the value function  $v_t(x)$  of the optimal control problem. Its proof appears in Appendix D.

**Proposition 2.** Suppose  $p_T = \bar{p}_T$ . Then

$$p_t(x) = \bar{p}_t(x), \quad \forall \ x \in \mathbb{R}^d, \ 0 \le t < T$$

Consequently, the optimal control law is given by

$$\Phi_t(x) = R^{-1} b^\top(x) \nabla \log \bar{p}_t(x)$$

Consistency with the LQ setting: With a(x) = Ax, c(x) = Cx, b(x) = B, and  $p_t = \mathcal{N}(0, P_t^{-1})$ . Then  $\nabla \cdot D = 0$  and the function  $h_t(x)$  simplifies considerably because

$$\nabla \cdot a(x) = (\text{constant}), \quad \nabla^2 \cdot D = 0,$$
  
 $\operatorname{Ir} \left( D \nabla^2 \log(p_t(x)) \right) = (\text{constant})$ 

Therefore, the right-hand side of the PDE (18) is given by

$$h_t(x) - \bar{h}_t = \frac{1}{2} |Cx|^2 - \frac{1}{2} \operatorname{Tr}(C^{\top} C(\bar{S}_t + \bar{n}_t \bar{n}_t^{\top}))$$

It is straightforward to verify that

$$\mathcal{V}_t(x) = \frac{1}{2} S_t C^\top C(x + \bar{n}_t)$$

solves the PDE (18), from which it follows that the equation for Y reduces to the form described in (7).

The first order PDE (18) is well known to arise in the nonlinear data assimilation literature [39, 40, 41]. One of the issues with the PDE is that its solution is not unique. For this reason, it is useful to consider the gradient form solution such that  $\mathcal{V}_t(x) = \nabla \phi_t(x)$ . The resulting PDE

$$-\frac{1}{\bar{p}_t(x)}\nabla\cdot(\bar{p}_t(x)\nabla\phi_t(x)) = h_t(x) - \bar{h}_t$$

is referred to as the Poisson equation, where the operator on the left-hand side is the weighted Laplacian. Based on assuming a suitable Poincare inequality, there is a well developed theory for existence and uniqueness of the solution of the Poisson equation [42, Theorem 1]. Given its importance in nonlinear filtering, numerical algorithms for solving the PDE is an area of ongoing research [39, 43]. Approximate formulae for the solution are also available, e.g., the constant gain approximation formula [44, Example 2].

## 3.1. Dual EnKF for nonlinear systems

Although one may numerically approximate the solution of the Poisson equation, one difficulty is that such approximations will require explicit forms of the vector-fields a(x) and b(x), and will violate Assumption 1. It is noted that the terms simplify in the following case:

- 1. If a(x) is conservative then  $\nabla \cdot a(x) = 0$ .
- 2. If b(x) = B then  $\nabla^2 \cdot D(x) = 0$  and  $\nabla \cdot D = 0$ .

Upon these simplifications, the mean-field process becomes

$$dY_t = a(Y_t) dt + b(Y_t) d\bar{\eta}_t + \mathcal{V}_t(Y_t) dt$$

where  $\mathcal{V}_t$  is obtained from solving the PDE (18) with

$$h_t = \frac{1}{2}|c|^2 + \frac{1}{2}\operatorname{Tr}\left(B^{\mathsf{T}}R^{-1}B\nabla^2\log(p_t)\right)$$

Now, it is natural to consider a Gaussian approximation of the density  $p_t$  whereupon  $h_t(x) = \frac{1}{2}|c(x)|^2 + (\text{constant})$ . This is useful to obtain a dual EnKF algorithm:

$$dY_t^i = \underbrace{a(Y_t^i) dt + b(Y_t^i) d\mathring{\eta}_t^i}_{i-\text{th copy of model (1b)}} + \mathsf{K}_t^{(N)}(\frac{c(Y_t^i) + \hat{c}_t^{(N)}}{2}) dt,$$
$$Y_T^i \overset{\text{i.i.d}}{\sim} \exp(-g_T), \quad 1 \le i \le N$$

where (as before)  $\eta^i := \{\eta^i_t \in \mathbb{R}^m : i : 0 \le t \le T\}$  is an independent copy of  $\eta$ ,  $\hat{c}_t^{(N)} := N^{-1} \sum_{i=1}^N c(Y_t^i)$ , and the gain is a constant matrix:

$$\mathsf{K}_t^{(N)} = \frac{1}{N-1} \sum_{i=1}^N (Y_t^i - n_t^{(N)}) (c(Y_t^i) - \hat{c}_t^{(N)})^\top$$

One may interpret the above as the dual counterpart of the FPF algorithm with a constant gain approximation [31, Example 2].

The optimal control is approximated as in the foregoing via the Hamiltonian,

$$H^{(N)}(x,\alpha,t) := \frac{1}{2}|c(x)|^2 + \frac{1}{2}\alpha^{\top}R\alpha + \frac{1}{N}\sum_{i=1}^{N}(x^{\top}X_t^i)(X_t^i)^{\top}\underbrace{(a(x) + b(x)\alpha)}_{i-\text{th copy of model (1b)}}$$

where as before  $X_t^i \coloneqq (S_t^{(N)})^{-1}(Y_t^i - n_t^{(N)})$ . Pseudo-code for the dual EnKF appears in Appendix E.

#### 3.2. Comparison with literature

In the introduction of [45], the authors write "Transformations based on an exponential change of measures have a rich tradition ... and are regularly re-discovered". Indeed, the pathwise (robust) representation of the nonlinear filter is based on a log transformation and its link to the HJB equation is at least as old as the works of [21, 46]. In the early 2000s, these classical ideas were re-purposed and extended for the purposes of algorithm design. There were two sets of ground-breaking contributions:

1. Inference as control. In [47], Mitter and Newton proposed a dual optimal control formulation of the nonlinear smoothing equations (see [25] for a recent review including a discussion of log transformation).

2. Control as inference. In [48, 49], Kappen described the so called path integral formulation of optimal control, where the log transformation is used to convert the HJB equation into a linear equation. In a closely related but independent work, Todorov used duality to express a class of optimal control problems as graphical inference problems [50]. Both these works continue to impact RL for robotics (a recent review is in [26]).

A key idea in these works is the classical connection between Kullback-Leibler (KL) divergence and Bayes' formula: Let P denote the law for a stochastic process X and  $Q^z$  denote the conditional law for X given an observation path z (this is given for inference problems). Let us construct a controlled process  $X^u$  and denote its law as  $P^u$  (this is given for control problems). Assuming  $P^u$  is absolutely continuous with respect to P (denoted  $P^u \ll P$ ), let us define the objective function as the KL divergence between  $P^u$  and  $Q^z$  as follows:

$$\min_{\mathsf{P}^{u}} \quad \mathsf{E}_{\mathsf{P}^{u}} \Big( \log \frac{\mathrm{d}\mathsf{P}^{u}}{\mathrm{d}\mathsf{P}} \Big) - \mathsf{E}_{\mathsf{P}^{u}} \Big( \log \frac{\mathrm{d}\mathsf{Q}^{z}}{\mathrm{d}\mathsf{P}} \Big)$$

In going from inference to control, a model for the controlled process  $X^u$  is prescribed. In going from control to inference, the integral state cost is interpreted as the conditional law  $Q^z$  (the second expectation). Of course, this places restriction on both the structure of the control system and the structure of the running cost. In both Mitter-Newton and in Kappen, the model structure is as follows:

$$dX_t^u = b(X_t^u, t) dt + \sigma(X_t^u)(U_t dt + dB_t)$$

where  $b(\cdot), \sigma(\cdot)$  are  $C^1$  vector fields and  $\tilde{B}$  is a W.P. For such a model,  $\mathsf{P}^u \ll \mathsf{P}$ and divergence (the first expectation) equals the quadratic control cost based on the use of the Girsanov transformation [51, Eq. (35)]. Extension of these concepts to discrete Markov decision processes (MDP) can be found in [52, Chapter 3] and is referred to as linearly solvable MDPs in [50].

In [53], Kappen and Ruiz write "Despite these elegant theoretical results, this idea has not been used much in practice. The essential problem is the representation of the controller as a parametrized model and how to adapt the parameters such as to optimize the importance sampler". Indeed, the design of algorithms based on these ideas remains an important area of research.

Since our focus is on inference algorithms for solving optimal control problems, we mention some salient points: The most direct approach is based on exact or approximate inference to compute the posterior. Computationally efficient message passing algorithms for the same are attractive in the linear Gaussian settings or if the state and action space is finite [54, 55]. The optimal control formulation of the smoothing equations in the linear Gaussian case is completely classical [23, Chapter 15], as are the message passing algorithms for the these cases. In a discrete MDP setting, a relevant example is the posterior policy iteration algorithm [56, Section II-C].

For nonlinear SDEs, the link is again classical – based on log transformation relating the pathwise filter and the HJB equation [25, Section 3.5]. The optimal policy is expressed as a certain Feyman-Kac type expectation which is approximated using importance sampling. For MDPs as well, the use of importance sampling for policy evaluation while sampling from another (simpler) policy is a standard approach in RL [38, Chapter 5.5, 5.7]. It allows the user to explore the state space using an exploratory policy while updating the optimal policy.

In practice, approximations are necessary. Based on the KL divergence, a natural approximation is to parametrize the control policy as  $\theta$  and denote the law of the controlled process as  $P^{\theta}$ . Then policy improvement is obtained using

$$\theta \longleftarrow \underset{\theta}{\operatorname{arg\,min}} \mathsf{E}_{P^{\theta}} \Big( \log \frac{\mathrm{d}P^{\theta}}{\mathrm{d}P} \Big) - \mathsf{E}_{\tilde{P}^{\theta}} \Big( \log \frac{\mathrm{d}Q^{z}}{\mathrm{d}P} \Big)$$

The resulting algorithm is referred to as the cross-entropy method in [53] where formulae for the gradient are also obtained and approximated using importance sampling. Related concepts and algorithms appear in a somewhat more general form in [56] for discrete state-space MDPs.

Given this history, we make the following points to distinguish our work from this earlier literature: 1. Log transformation. While our use of the log transformation is same as the path integral approach of Kappen [48, 49], an important difference is that for us  $p_t$  is a (normalized) probability density. The governing equation (16) is nonlinear because of the terms involving  $\bar{h}_t$ . In contrast, the path integral method works with the un-normalized density whose equation is linear. The linearity is crucial for the Feynan Kac formula and its empirical approximation using importance sampling. For us, the equation for the normalized density is necessary because our aim is to construct a McKean-Vlasov SDE.

2. Algorithm. The controlled interacting particle system via a finite-N approximation of the McKean-Vlasov SDE is original. It is conceptually and structurally distinct from earlier work, same as the distinction between important sampling and control-type algorithms in the filtering context; the latter class of algorithms is of much recent origin [18, 20]. In particular, we are not aware of any work using EnKF (or similar constructions) to solve an optimal control problem.

## 4. Numerics

The performance of the dual EnKF algorithm is numerically evaluated for three benchmark examples. In each of the three examples, the optimal control problem is formulated as an infinite-horizon LQR problem. This allows also for a comparison with the state-of-the-art methods that have focussed on this problem.

In a numerical implementation, the terminal time T is fixed and EnKF is simulated to obtain an empirical approximation  $\{P_t^{(N)} \in \mathbb{R}^{d \times d} : 0 \le t < T\}$ , typically using  $P_T = I$ , the identity matrix. For the sake of comparison, the exact  $\{P_t \in \mathbb{R}^{d \times d} : 0 \le t \le T\}$  is obtained by numerically integrating the DRE (5). The stationary solution  $P^{\infty}$  is obtained as a solution of the ARE using scipy package in Python. Pseudo-code is contained in Algorithm 2 of Appendix E. All the code is available on Github [57].

4.1. Linear system with randomly chosen entries

A d-dimensional system is in its controllable canonical form

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

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

Figure 1a depicts the convergence of the four entries of the matrix  $P_t^{(N)}$  for the case where the state dimension d = 2. Figure 1b depicts the analogous results for d = 10. Figures 2a and 2b 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.



Figure 1: 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  $0 \le t \le T$ .

#### 4.2. Mass spring damper system

We present numerical comparison of EnKF with policy gradient algorithms in [11] (denoted as [M21]) and [7] (denoted as [F18]). Comparison is made on the benchmark spring mass damper example [58, Section VI]. Additional details on modeling along with the numerical values of various simulation parameters can be found in Appendix F.1.

Figure 3 depicts the variation of the relative mean-squared error, defined as

MSE := 
$$\frac{1}{T} \mathsf{E} \left( \int_0^T \frac{\|P_t - P_t^{(N)}\|_F^2}{\|P_t\|_F^2} \, \mathrm{d}t \right)$$



Figure 2: Open and closed-loop poles.

The figure depicts two trends: the  $O(\frac{1}{N})$  decay of the MSE as N increases (for d fixed), which is an illustration of the error bound (11), 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. 4. The comparison is for the following metrics (taken from [11]):

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

where the LQR optimal gain  $K^{\infty}$  and the optimal value  $c^{\infty}$  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, we use the Euler approximation to obtain a discretetime model. Such an approximation is consistent with our choice of numerical integration in Algorithm 2.

To obtain the relationship between the error and computational time, the number of particles N is varied in the EnKF algorithm while the number of gradient descent steps is changed in [M21] and [F18].

In the numerical experiments, the dual EnKF is found to be significantly more computationally efficient—by two orders of magnitude or more. Comparison was carried out for a range of d and is qualitatively similar, see Appendix F.2. 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 [0, T]. We found that the number of particles (N)



Figure 3: Mean-squared error (MSE) as a function of the number of particles N and system dimension d

for the EnKF algorithm is typically one or two orders of magnitude larger than  $N_g$ . Since our algorithm is designed to be written as a matrix vector multiplication, vectorization features of the **numpy** package in Python yield significant gains in computational time. 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 our comparisons, the same time-horizon [0, T] and discretization time-step  $\Delta t$  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 (over the  $N_g$  copies) to speed up the gradient calculation. Codes are made available on Github for interested parties to independently verify these comparisons [57].

#### 4.3. Nonlinear cart-pole system

Figure 5 depicts the closed-loop trajectories of a four-dimensional nonlinear cart pole model. The control acts as external force applied to the cart. The four-dimensional state for the system is  $(\theta, x, \dot{\theta}, \dot{x})$ , where  $\theta \in S^1$ (the circle) is the angle of the pole (pendulum) as measured from the stable equilibrium,  $x \in \mathbb{R}$  is the displacement of cart along the horizontal. The control objective is to balance the pole – stabilize the system at the inverted equilibrium ( $\pi, 0, 0, 0$ ), assuming full state feedback. (See Appendix G for details on the model parameters and their numerical values).



Figure 4: Comparison with algorithms in [7] (labeled [F18]) and [11] (labeled [M21]). The comparisons depict the computation time (in Python) as a function of the relative error in approximating the LQR gain and cost.

For the purposes of control design, the nonlinear system is first linearized at the desired equilibrium and an LQR problem is formulated based on [59, Chapter 3.2]. The (nonlinear) dual EnKF is used to approximate the optimal control law which is numerically evaluated on the fully nonlinear model. Figure 5 depicts the numerically obtained results. It was found that reasonable levels of performance is obtained with as few as N = 10 particles. With N = 1000 particles, the closed-loop trajectories are virtually indistinguishable from the exact optimal control solution.

# 5. Conclusions

In this paper, we present a new class of algorithms for learning optimal policies using simulations. A key message is that log transforms combined with mean field techniques can lead to simulation based methods for optimal policy approximation. We have demonstrated this for LQ in full detail, and shown how the techniques generalize to nonlinear systems.

There are two key innovations: (i) the representation of the unknown value function in terms of the statistics (variance) of a suitably designed process; and (ii) design of interactions 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. For policy



Figure 5: Trajectories of the closed-loop nonlinear cart pole.

evaluation, use of Monte Carlo techniques is already standard. It is shown in this paper is that by designing careful interaction amongst simulations, one can also solve the policy optimization problem.

Another notable aspect is the learning rate. Because the  $N = \infty$  limit is exact tor the LQR problem, the proposed algorithms 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 (11) (although such an analysis is conservative). In numerical examples, this property is shown to lead to an order of magnitude better performance than the state-of-the-art algorithms.

Given the enormous success of EnKF in data assimilation [3, 2], the contributions of this paper potentially open up new opportunities for RL. It is our hope that the paper will engender new synergies between the data assimilation and the RL communities.

# Appendix A. Proof of Prop. 1

The equation for the mean  $\bar{n}_t$  is obtained by taking the expectation of SDE (7),

$$\mathrm{d}\bar{n}_t = (A + \bar{S}_t C^\top C) \bar{n}_t \,\mathrm{d}t$$

Because  $\bar{n}_T = 0$ , we have  $\bar{n}_t = 0$  for all  $t \in [0, T]$ .

The equation for the covariance  $\bar{S}_t$  is obtained by writing the SDE for the error  $e_t := \bar{\mathcal{Y}}_t - \bar{n}_t$ :

$$\mathrm{d}e_t = (A + \frac{1}{2}\bar{S}_t C^\top C)e_t\,\mathrm{d}t + B\,\mathrm{d}\dot{\bar{\eta}}_t,$$

Using the Itô rule for  $e_t e_t^{\top}$ ,

$$d(e_t e_t^{\mathsf{T}}) = B \, \mathrm{d} \dot{\eta}_t e_t^{\mathsf{T}} + e_t (B \, \mathrm{d} \dot{\eta}_t)^{\mathsf{T}} - B R^{-1} B^{\mathsf{T}} + (A + \frac{1}{2} \bar{S}_t C^{\mathsf{T}} C) (e_t e_t^{\mathsf{T}}) \, \mathrm{d} t + (e_t e_t^{\mathsf{T}}) (A + \frac{1}{2} \bar{S}_t C^{\mathsf{T}} C)^{\mathsf{T}}$$

The Itô correction term appears with a negative sign because the SDE involves a backward Wiener process  $\dot{\eta}_t$  [29, Sec. 4.2]. Taking an expectation yields the following equation for  $\bar{S}_t$ :

$$\frac{\mathrm{d}}{\mathrm{d}t}\bar{S}_t = (A + \frac{1}{2}\bar{S}_tC^{\top}C)\bar{S}_t + \bar{S}_t(A + \frac{1}{2}\bar{S}_tC^{\top}C)^{\top} - BR^{-1}B^{\top}$$

The SDE is identical to the SDE for  $S_t$ . Because  $\bar{S}_T = S_T$ , we have  $\bar{S}_t = S_t$  for all  $t \in [0, T]$ . The conclusion that  $Y_t$  is Gaussian follows from the fact that with  $\bar{n}_t = n_t$  and  $\bar{S}_t = S_t$ , the SDE for  $Y_t$  is an Ornstein-Uhlenbeck SDE with a Gaussian terminal condition.

The proof for the rest of proposition is straightforward. By definition,

$$\mathsf{E}[X_t] = \mathsf{E}[\bar{S}_t^{-1}(Y_t - \bar{n}_t)] = \bar{S}_t^{-1}(\bar{n}_t - \bar{n}_t) = 0 \mathsf{E}[X_t X_t^{\top}] = \mathsf{E}[\bar{S}_t^{-1}(Y_t - \bar{n}_t)(Y_t - \bar{n}_t)^{\top}\bar{S}_t^{-1}] = \bar{S}_t^{-1} = S_t^{-1} = P_t$$

## Appendix B. Error analysis

**Notation:** Let  $S^d_+ \subset S^d \subset \mathbb{R}^{d \times d}$  denote the set of symmetric positive definite matrices and symmetric matrices respectively. Let  $\langle Q_1, Q_2 \rangle \coloneqq \operatorname{Tr}(Q_1 Q_2^{\top})$  denote the Frobenius inner product, and  $|| \cdot ||_F \coloneqq \sqrt{\langle Q_1, Q_1 \rangle}$  denote the Frobenius inner product for  $Q_1, Q_2 \in \mathbb{R}^{d \times d}$ .

The objective is to study the error between the empirical covariance of the particles  $S_t^{(N)}$  and its mean-field limit  $S_t$ . To simplify the presentation, we use the time-reversed quantitative  $\Omega_t^{(N)} := S_{T-t}^{(N)}$  and  $\Omega_t := \bar{S}_{T-t}$ . According to the Proposition 1,  $\Omega_t$  satisfies the Riccati equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\Omega_t = \operatorname{Ricc}(\Omega_t) := -A\Omega_t - \Omega_t A^{\mathsf{T}} - \Omega_t C^{\mathsf{T}} C \Omega_t + \Sigma_B, \qquad (B.1)$$

where  $\Sigma_B := BR^{-1}B^{\top}$ . The time-evolution for  $\Omega_t^{(N)}$  is obtained by the application of the Itô rule to its definition [60, Prop. 4.2]

$$d\Omega_t^{(N)} = \operatorname{Ricc}(\Omega_t^{(N)}) \, \mathrm{d}t + \frac{1}{\sqrt{N}} \, \mathrm{d}M_t, \tag{B.2}$$

where  $\{M_t : t \ge 0\}$  is a martingale given by

$$dM_t = \frac{1}{\sqrt{N}} \sum_{i=1}^N F_t^i (B \operatorname{d} \eta_t^i)^{\mathsf{T}} + B \operatorname{d} \eta_t^i (F_t^i)^{\mathsf{T}}, \quad F_{T-t}^i \coloneqq X_t^i - n_t^{(N)}$$

with quadratic variation

$$d\langle M \rangle_t = \text{Tr}(\Sigma_B)\Omega_t^{(N)} + \Sigma_B\text{Tr}(\Omega_t^{(N)}) + \Sigma_B\Omega_t^{(N)} + \Omega_t^{(N)}\Sigma_B$$

The error analysis is based on a sensitivity analysis of the Riccati equation. Let  $\phi(t, Q)$  denote the semigroup associated with the Riccati equation such that for any positive definite matrix  $Q \in S^d_+$ ,

$$\frac{\partial \phi}{\partial t}(t,Q) = \operatorname{Ricc}(\phi(t,Q)), \quad \phi(0,Q) = Q.$$

We define the first-order and the second-order derivatives which are the linear and bilinear operators  $\frac{\partial \phi}{\partial Q}(t,Q)$ :  $S^d \to S^d$  and  $\frac{\partial^2 \phi}{\partial Q^2}(t,Q)$ :  $S^d \times S^d \to S^d$  respectively that satisfy

$$\frac{\partial \phi}{\partial Q}(t,Q)(Q_1) = \left. \frac{\mathrm{d}}{\mathrm{d}\epsilon} \right|_{\epsilon=0} \phi(t,Q+\epsilon Q_1)$$
$$\frac{\partial^2 \phi}{\partial Q^2}(t,Q)(Q_1,Q_1) = \left. \frac{\mathrm{d}^2}{\mathrm{d}\epsilon^2} \right|_{\epsilon=0} \phi(t,Q+\epsilon Q_1).$$

We also let  $\|\frac{\partial \phi}{\partial Q}(t,Q)\|_{F,F}$  and  $\|\frac{\partial^2 \phi}{\partial Q^2}(t,Q)\|_{F,F}$  denote the induced-norm of these operators with respect to the Frobenius norm. The following lemma expresses the error as a stochastic integral that involves the semigroup.

**Lemma 1.** Consider  $\Omega_t$  and  $\Omega_t^{(N)}$  defined in (B.1) and (B.2) respectively. Then

$$\Omega_t^{(N)} - \Omega_t = \frac{1}{\sqrt{N}} \int_0^t \frac{\partial \phi}{\partial Q} (t - s, \Omega_s^{(N)}) (dM_s) + \frac{1}{2N} \int_0^t \frac{\partial^2 \phi}{\partial Q^2} (t - s, \Omega_s^{(N)}) (dM_s, dM_s) + \phi(t, \Omega_0^{(N)}) - \phi(t, \Omega_0)$$
(B.3)

**PROOF.** The proof follows by expressing the difference

$$\Omega_t^{(N)} - \Omega_t = \phi(0, \Omega_t^{(N)}) - \phi(t, \Omega_0)$$
  
=  $\phi(0, \Omega_t^{(N)}) - \phi(t, \Omega_0^{(N)}) + \phi(t, \Omega_0^{(N)}) - \phi(0, \Omega_0)$   
=  $\int_0^t d_s \phi(t - s, \Omega_s^{(N)}) + \phi(t, \Omega_0^{(N)}) - \phi(t, \Omega_0),$ 

and evaluating the differential

$$d_s\phi(t-s,\Omega_s^{(N)}) = -\frac{\partial\phi}{\partial t}(t-s,\Omega_s^{(N)}) ds + \frac{\partial\phi}{\partial Q}(t-s,\Omega_s^{(N)})(d\Omega_s^{(N)}) + \frac{1}{2}\frac{\partial^2\phi}{\partial Q^2}(t-s,\Omega_s^{(N)})(d\Omega_s^{(N)}, d\Omega_s^{(N)}),$$

and using the identity  $\frac{\partial \phi}{\partial t}(t,Q) = \frac{\partial \phi}{\partial Q}(t,Q)(\operatorname{Ricc}(Q)).$ 

The preceding lemma can be viewed as the extension of the Alekseev-Gröbner formula to matrix-valued stochastic differential equations [61]. The explicit form of this expression appears in [62, Sec. 5.3].

The error bound follows from uniform bounds on the terms involved in the integral (B.3). Such uniform bounds are available if the Riccati equation enjoys the following stability property.

**Assumption 2.** Consider the semigroup corresponding to the Riccati equation (B.1). There are positive constants  $c_1$ ,  $c_2$ , and  $\lambda$  such that  $\forall Q \in S^d_+$ :

$$\|\frac{\partial\phi}{\partial Q}(t,Q)\|_{F,F} \le c_1 e^{-2\lambda t}, \quad \|\frac{\partial^2\phi}{\partial Q^2}(t,Q)\|_{F,F} \le c_2 e^{-2\lambda t}.$$

These bounds are directly related to the exponential stability of the closedloop linear system under optimal feedback control [60, Sec. 2]. The exponential decay holds when the linear system is controllable and observable. However, the fact that the constants  $c_1$  and  $c_2$  are uniform among all initial matrices Q is still open. See [60, 63] for detailed analysis of the Riccati equation where these uniform bounds are shown to hold under the additional assumption that the matrix C is full-rank. **Proposition 3.** Let  $\bar{S}_t$  be the mean-field covariance defined in (8) and  $S_t^N$  be the empirical covariance of the particles defined in (B.2). Then, under Assumption 2, the error between  $S_t^{(N)}$  and  $\bar{S}_t$  satisfies the upper-bound

$$\mathsf{E}[\|S_t^{(N)} - \bar{S}_t\|_F] \le \frac{C_1}{\sqrt{N}} + C_2 e^{-2\lambda(T-t)} \mathsf{E}[\|S_T^{(N)} - \bar{S}_T\|_F], \qquad (B.4)$$

where  $C_1, C_2$  are time-independent positive constants.

**PROOF.** Using (B.3) and the triangle inequality, the expected norm of the difference satisfies

$$\mathsf{E}[\|\Omega_t^{(N)} - \Omega_t\|_F] \le \frac{R_1}{\sqrt{N}} + \frac{R_2}{2N} + R_3$$

where

$$\begin{aligned} R_1 &= \mathsf{E}\left[\left\|\int_0^t \frac{\partial \phi}{\partial Q}(t-s,\Omega_s)(\,\mathrm{d}M_s)\right\|_F\right]\\ R_2 &= \mathsf{E}\left[\int_0^t \left\|\frac{\partial^2 \phi}{\partial Q^2}(t-s,\Omega_s)(\,\mathrm{d}M_s,\,\mathrm{d}M_s)\right\|_F\right]\\ R_3 &= \mathsf{E}\left[\left\|\phi(t,\Omega_0^{(N)}) - \phi(t,\Omega_0)\right\|_F\right] \end{aligned}$$

The first term

$$R_{1} \leq \left[ \mathsf{E} \left[ \left\| \int_{0}^{t} \frac{\partial \phi}{\partial Q} (t - s, \Omega_{s}) (\,\mathrm{d}M_{s}) \right\|_{F}^{2} \right] \right]^{\frac{1}{2}}$$
$$= \left[ \int_{0}^{t} \mathsf{E} \left[ \left\| \frac{\partial \phi}{\partial Q} (t - s, \Omega_{s}) (\,\mathrm{d}M_{s}) \right\|_{F}^{2} \right] \right]^{\frac{1}{2}}$$
$$\leq \left[ \int_{0}^{t} \mathsf{E} \left[ \left\| \frac{\partial \phi}{\partial Q} (t - s, \Omega_{s}) \right\|_{F,F}^{2} \left\| \mathrm{d}M_{s} \right\|_{F}^{2} \right] \right]^{\frac{1}{2}}$$
$$\leq \left[ \int_{0}^{t} 4c_{1}^{2} e^{-4\lambda(t-s)} \mathrm{Tr}(\Sigma_{B}) \mathsf{E}[\mathrm{Tr}(\Omega_{s}^{(N)})] \,\mathrm{d}s \right]^{\frac{1}{2}}$$

where we used Jensen's inequality in the first step, Itö isometry in the second step, and Assumption 2 in the last step. The second term,

$$R_{2} \leq \mathsf{E}\left[\int_{0}^{t} \|\frac{\partial^{2}\phi}{\partial Q^{2}}(t-s,\Omega_{s})\|_{F} \|dM_{s}\|_{F}^{2}\right]$$
$$\leq \int_{0}^{t} 4c_{2}e^{-2\lambda(t-s)}\mathrm{Tr}(\Sigma_{B})\mathsf{E}[\mathrm{Tr}(\Omega_{s}^{(N)})]ds$$

where we used Assumption 2. The third term,

$$R_3 \le c_1 e^{-2\lambda t} \mathsf{E}[\|\Omega_0^{(N)} - \Omega_0\|_F]$$

because of the bound on the first derivative in Assumption 2. Upon using the bound  $\mathsf{E}[\operatorname{Tr}(\Omega_t^{(N)})] \leq \operatorname{Tr}(\Sigma_t) \leq \sup_{t\geq 0} \operatorname{Tr}(\Sigma_t) =: \sigma^2$  from [60, Thm. 5.2], we conclude

$$\mathsf{E}[\|\Omega_t^{(N)} - \Omega_t\|_F] \le (c_1 + c_2\sqrt{\epsilon})\sqrt{\epsilon} + c_1 e^{-2\lambda t} \mathsf{E}[\|\Omega_0^{(N)} - \Omega_0\|_F]$$

where  $\epsilon := \frac{\sigma^2 \operatorname{Tr}(\Sigma_B)}{\lambda N}$ . Changing t to T - t concludes the proof.

# Appendix C. Evolution of density in (16)

By definition, the probability density

$$p_T(x) = \frac{\exp(-g(x))}{\int \exp(-g(x)) \, \mathrm{d}x}, \quad x \in \mathbb{R}^d$$

Write  $v_t = -\log(p_t) + \beta_t$  where  $\beta_t := \log(\int p_t(x)dx)$  is a time-dependent constant to ensure  $p_t$  is normalized. In terms of  $p_t$  and  $\beta_t$ , the HJB equation (14) for  $v_t$  is written as

$$-\frac{1}{p_t}\frac{\partial p_t}{\partial t} + \dot{\beta}_t + \frac{1}{2}|c|^2 - \frac{1}{p}a^T \nabla p - \frac{1}{2p_t} \operatorname{Tr}(D\nabla^2 p_t) - \frac{1}{2} \operatorname{Tr}((Q-D)\nabla^2 \log(p_t)) = 0$$

where we used  $\nabla^2 \log(p_t) = \frac{1}{p_t} \nabla^2 p_t - \frac{1}{p_t^2} \nabla p_t \nabla p_t^{\top}$ . Multiplying by  $p_t$  yields

$$\frac{\partial p_t}{\partial t} = (h_t + \dot{\beta}_t)p_t - \nabla \cdot (p_t a) + \nabla \cdot (p_t \nabla \cdot D) - \frac{1}{2}\nabla^2 \cdot (p_t D)$$

where we used

$$h_t := \frac{1}{2} |c|^2 + \nabla \cdot a - \frac{1}{2} \nabla^2 \cdot D + \frac{1}{2} \operatorname{Tr}((D - Q) \nabla^2 \log(p_t))$$
$$a^T \nabla p_t = \nabla \cdot (p_t a) - p_t \nabla \cdot a$$
$$\operatorname{Tr}(D \nabla^2 p_t) = \nabla^2 \cdot (p_t D) - 2 \nabla \cdot (p_t \nabla \cdot D) + p_t \nabla^2 \cdot D$$

Noting  $\int \frac{\partial p_t}{\partial t} dx = 0$ , we obtain

$$\dot{\beta} = -\int h_t(x)p_t(x)\,\mathrm{d}x = -\hat{h}_t$$

which in turn gives the PDE (16) for  $p_t$ .

## Appendix D. Proof of Prop. 2

The proof for  $\bar{p}_t = p_t$  follows from showing that the evolution equation for  $\bar{p}_t$  and  $p_t$  are identical. Consider the SDE (17). The evolution equation for the density  $\bar{p}_t$  is the Fokker-Planck equation:

$$\frac{\partial \bar{p}_t}{\partial t} = -\nabla \cdot (\bar{p}_t a) - \nabla \cdot (\bar{p}_t \nabla \cdot D) - \nabla \cdot (\bar{p}_t \mathcal{V}_t) - \frac{1}{2} \nabla^2 \cdot (\bar{p}_t D)$$

where the diffusion term  $\frac{1}{2}\nabla^2 \cdot (\bar{p}_t D)$  appears with a negative sign because  $\dot{\bar{\eta}}_t$  is a backward Wiener process.

It is easily see that if the vector-field  $\mathcal{V}_t(\cdot)$  solves the PDE (18) then the evolution equations for  $p_t$  and  $\bar{p}_t$  are identical.

# Appendix E. Algorithm for implementing nonlinear dual EnKF

The algorithm to approximate the optimal control policy for (1) is divided into an online and offline component.

**Offline algorithm.** (Algorithm 1) to compute  $\{P_t^{(N)} : 0 \le t \le T\}$ . It is based on the finite-N approximation of the dual EnKF (9). For a numerical solution of the SDE, we use the simplest Euler scheme which can be swapped with a higher order scheme.

**Online algorithm.** (Algorithm 2) to compute the optimal control for a given state  $X_t = x$  at time t. In addition to the simulator, this algorithm also requires  $P_t^{(N)}$  computed from the offline algorithm. It is based on minimizing the Hamiltonian function.

The algorithm is described for the general nonlinear case. The LQ is the special case when f(x, u) = Ax + Bu and c(x) = Cx.

In a numerical implementation of the offline algorithm, there are two sources of error: (i) because of finite-N approximation; and (ii) because of time-discretization step size  $\Delta t$ . The first type of error scales as  $O(\frac{1}{\sqrt{N}})$  as shown in the bound (11). For SDEs, the second type of error scales as  $O(\Delta t)$ using the Euler scheme [64].

#### Appendix F. Details of Example 4.2

Appendix F.1. Coupled mass spring damper system

This system is taken from [58]. The matrices A and B are as follows:

$$A = \begin{bmatrix} 0_{d_s \times d_s} & \mathcal{I}_{d_s} \\ -\mathbb{T} & -\mathbb{T} \end{bmatrix}, \quad B = \begin{bmatrix} 0_{d_s \times d_s} \\ \mathcal{I}_{d_s} \end{bmatrix}$$

Algorithm 1 [offline] EnKF algorithm to approximate  $\{P_t : 0 \le t \le T\}$ 

**Require:** Simulation time T, simulation step-size  $\Delta t$ , number of particles N, simulator f(x, u) = a(x) + b(x)u, terminal cost  $g_T$ , cost function c(x), and control cost matrix R. 1: **return**  $\{P_k^{(N)}(\cdot) : k = 0, 1, 2, \dots, \frac{T}{\Delta t} - 1\}$ 2:  $T_F = \frac{T}{\Delta t}$ 3: Initialize  $\{Y_{T_F}^i\}_{i=1}^N \stackrel{\text{i.i.d}}{\sim} \exp(-g_T)$ 4: calculate  $n_{T_F}^{(N)} = N^{-1} \sum_{i=1}^N Y_{T_F}^i$ 5: for  $k = T_F$  to 1 do 6: Calculate  $\hat{c}_k^{(N)} = N^{-1} \sum_{i=1}^N c(Y_k^i)$ 7: Calculate  $M_k^{(N)} = (N-1)^{-1} \sum_{i=1}^N (Y_k^i - n_k^{(N)}) (c(Y_k^i) - \hat{c}_k^{(N)})^\top$ for i = 1 to N do  $\Delta \eta_k^i \stackrel{\text{i.i.d}}{\sim} \mathcal{N}(0, \frac{1}{\Delta t} R^{-1})$ 8: 9:  $\begin{array}{l} \Delta Y_k^i = f(Y_k^i, \overset{\Delta}{\Delta} \eta_k^i) \Delta t + \frac{1}{2} M_k^{(N)} (c(Y_k^i) + \hat{c}_k^{(N)}) \Delta t \\ Y_{k-1}^i = Y_k^i - \Delta Y_k^i \end{array}$ 10: 11: Calculate  $n_{k-1}^{(N)} = N^{-1} \sum_{i=1}^{N} Y_{k-1}^{i}$ Calculate  $S_{k-1}^{(N)} = (N-1)^{-1} \sum_{i=1}^{N} (Y_{k-1}^{i} - n_{k-1}^{(N)}) (Y_{k-1}^{i} - n_{k-1}^{(N)})^{\top}$  $P_{k-1}^{(N)} = (S_{k-1}^{(N)})^{-1}$ 12:13:14: 15:16: end for

Algorithm 2 [online] EnKF algorithm to calculate optimal control for (1)

**Require:** Simulation time *T*, simulation step-size  $\Delta t$ , number of particles  $N, \{P_k^{(N)} : k = 0, 1, 2, \dots, \frac{T}{\Delta t}\}$  from the offline algorithm 1, Hamiltonian function  $\mathcal{H}(x, y, \alpha) = y^T(a(x) + b(x)\alpha) + \frac{1}{2}|c(x)|^2 + \frac{1}{2}\alpha^\top R\alpha, \{e_i\}_{i=1}^m$  the standard basis of  $\mathbb{R}^m$ 

- 1: **return** optimal control input  $\{u_k^{(N)} \in \mathbb{R}^m : k = 0, 1, 2, \dots, \frac{T}{\Delta t} 1\}.$
- 2: Define  $T_F \coloneqq \frac{T}{\Delta t}$
- 3: for k = 0 to  $\vec{T}_F 1$  do
- 4: Observe state of the system, denoted  $x_k$
- 5: Define  $y_k = P_k x_k$
- 6: for i = 1 to m do

7: 
$$\langle u_k^{(N)}, e_i \rangle = \mathcal{H}(x_k, y_k, R^{-1}e_i) - \mathcal{H}(x_k, y_k, 0) - \frac{1}{2}(R^{-1})_{ii}$$

- 8: end for
- 9: Apply control  $u_k^{(N)}$  to the true system

10: **end for** 

where  $d_s = \frac{d}{2}$  is the number of masses and  $\mathbb{T} \in \mathbb{R}^{d_s \times d_s}$  is a Toeplitz matrix with 2 on the main diagonal and -1 on the first sub-diagonal and first superdiagonal. Numerical values of parameters used in simulations are listed in Table F.2.

Parameter Name	Symbol	Numerical value				
Model Parameters						
	C for $d = 2$	$\sqrt{5}\mathcal{I}_d$				
I.O. parametera	C for $d > 2$	$\mathcal{I}_d$				
LQ parameters	R	$\mathcal{I}_{d_s}$				
	$P_T$	$\mathcal{I}_d$				
Simulation Parameters						
Simulation time	T	10				
Step size	$\Delta t$	0.02				

Table F.2: Model parameters for the coupled mass spring damper system

#### Appendix F.2. Comparison between EnKF and policy-gradient methods

The hyper-parameters required to implement the algorithms of [M21], and [F18] algorithms are as follows. The simulation time horizon T = 10, and the step-size  $\Delta t = 0.01$  is the same for all of EnKF, [F18] and [M21]. The initial guess  $K^0 = 0$ , initial distribution  $\mathcal{D}^0 = \mathcal{N}(0, \mathcal{I}_d)$ , and gradient descent step  $\alpha = 0.0001$  for both [M21] and [F18]. The values of the other hyper parameters, namely the smoothing parameter r and number of particles in gradient calculation  $N_g$  are in Table F.3. The numerical results for d = 10are depicted in Figure 4 and for d = 2, 4 in Figure F.6. Additionally, Figure F.7 shows comparison for error in cost. While calculating cost, the system is initialised with a  $\mathcal{N}(0, 0.1\mathcal{I}_d)$  distribution to keep the simulation setup as close to the setting of [M21] and [F18] as possible.

The simulations are implemented in Python 3 on a Intel Xeon E3-1240 V2 3.40 Ghz CPU, and the process\_time() function from the time module is used to evaluate the execution time.



Figure F.7: Comparison of relative in error in cost

Hyper-param.		[M21]			[F18]	
d	2	4	10	2	4	10
r	$10^{-1}$	$10^{-1}$	$10^{-3}$	$10^{-1}$	$10^{-1}$	$10^{-1}$
$N_g$	2	4	10	2	4	10

Table F.3: Hyper-parameter values for policy gradient

## Appendix G. Cart-pole system

The nonlinear model is taken from [59, Chapter 3.2.1]:

$$\dot{\theta} = \omega$$
$$\dot{\omega} = \frac{-F\cos(\theta) - ml\omega^2\cos(\theta)\sin(\theta) - (m+M)g\sin(\theta)}{l(M+m\sin^2(\theta))}$$
$$\dot{x} = v$$
$$\dot{v} = \frac{F + m\sin(\theta)(l\omega^2 + g\cos(\theta))}{M + m\sin^2(\theta)}$$

For the specification of the LQ cost, we first linearize the system about the desired inverted equilibrium  $(\pi, 0, 0, 0)$ . The associated A and B matrices are as follows:

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \frac{(M+m)g}{Ml} & 0 & 0 & 0 \\ \frac{mg}{M} & 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ \frac{1}{Ml} \\ 0 \\ \frac{1}{M} \end{bmatrix}$$

Note these are used only to obtain the LQR solution (for comparison) but not needed to implement the dual EnKF. The model parameters and the simulation parameters are are listed in Table G.4.

## References

- J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, T. Lillicrap, D. Silver, Mastering atari, go, chess and shogi by planning with a learned model, Nature 588 (7839) (2020) 604–609. doi:10.1038/ s41586-020-03051-4. URL https://doi.org/10.1038/s41586-020-03051-4
- [2] G. Evensen, Data Assimilation. The Ensemble Kalman Filter, Springer-Verlag, New York, 2006.
- [3] S. Reich, C. Cotter, Probabilistic forecasting and Bayesian data assimilation, Cambridge University Press, 2015.

Parameter name	Symbol	Numerical value					
Model parameters							
Mass of ball	m	0.08					
Mass of cart	M	1					
Length of rod	l	0.7					
Gravity	g	9.81					
Unstable equilibrium	$(ar{ heta},ar{x},ar{\omega},ar{v})$	$(\pi,0,0,0)$					
Initial condition	$(\theta(0), x(0), \omega(0), v(0))$	$(1.25\pi, -0.1, 0, 0)$					
	C	diag([10, 10, 1, 1])					
LQ parameters	R	10					
	$P_T$	$\mathcal{I}_4$					
Simulation parameters							
Simulation time	T	10					
Step size	$\Delta t$	0.0002					

Table G.4: Parameters for the cart-pole system

- [4] R. E. Kalman, Contributions to the theory of optimal control, Boletin de la Sociedad Matematica Mexicana (2) 5 (1960) 102–109.
- [5] P. Lancaster, L. Rodman, Algebraic Riccati Equations, Clarendon press, 1995.
- [6] A. J. Laub, Invariant subspace methods for the numerical solution of Riccati equations, in: The Riccati Equation, Springer, 1991, pp. 163– 196.
- M. Fazel, R. Ge, S. Kakade, M. Mesbahi, Global Convergence of Policy Gradient Methods for the Linear Quadratic Regulator, in: International Conference on Machine Learning, PMLR, 2018, pp. 1467–1476, iSSN: 2640-3498.
   URL http://proceedings.mlm.proce/u80/fazel182.html

 ${\rm URL}\ {\tt http://proceedings.mlr.press/v80/fazel18a.html}$ 

- [8] S. Tu, B. Recht, The Gap Between Model-Based and Model-Free Methods on the Linear Quadratic Regulator: An Asymptotic Viewpoint, in: Conference on Learning Theory, PMLR, 2019, pp. 3036-3083, iSSN: 2640-3498. URL http://proceedings.mlr.press/v99/tu19a.html
- [9] S. Dean, H. Mania, N. Matni, B. Recht, S. Tu, On the Sample Complexity of the Linear Quadratic Regulator, Found Comput Math 20 (4) (2020) 633-679. doi:10.1007/s10208-019-09426-y. URL https://doi.org/10.1007/s10208-019-09426-y
- [10] D. Malik, A. Pananjady, K. Bhatia, K. Khamaru, P. L. Bartlett, M. J. Wainwright, Derivative-Free Methods for Policy Optimization: Guarantees for Linear Quadratic Systems, Journal of Machine Learning Research 21 (21) (2020) 1–51. URL http://jmlr.org/papers/v21/19-198.html
- [11] H. Mohammadi, A. Zare, M. Soltanolkotabi, M. R. Jovanović, Convergence and sample complexity of gradient methods for the model-free linear-quadratic regulator problem, IEEE Transactions on Automatic Control 67 (5) (2022) 2435–2450. doi:10.1109/TAC.2021.3087455.
- [12] H. Mohammadi, M. Soltanolkotabi, M. R. Jovanovic, On the Linear Convergence of Random Search for Discrete-Time LQR, IEEE Control Systems Letters 5 (3) (2021) 989–994, conference Name: IEEE Control Systems Letters. doi:10.1109/LCSYS.2020.3006256.
- K. Zhang, B. Hu, T. Basar, On the Stability and Convergence of Robust Adversarial Reinforcement Learning: A Case Study on Linear Quadratic Systems, Advances in Neural Information Processing Systems 33 (2020) 22056-22068.
   URL https://proceedings.neurips.cc//paper/2020/hash/ fb2e203234df6dee15934e448ee88971-Abstract.html
- [14] J. P. Jansch-Porto, B. Hu, G. E. Dullerud, Convergence guarantees of policy optimization methods for Markovian jump linear systems, in: 2020 American Control Conference (ACC), IEEE, 2020, pp. 2882–2887.
- [15] K. Zhang, Reinforcement learning for multi-agent and robust control systems, Ph.D. thesis, University of Illinois at Urbana-Champaign (2021).

- [16] A. Taghvaei, P. G. Mehta, Optimal transportation methods in nonlinear filtering, IEEE Control Systems Magazine 41 (4) (2021) 34–49. doi: 10.1109/MCS.2021.3076391.
- [17] S. C. Surace, A. Kutschireiter, J.-P. Pfister, How to avoid the curse of dimensionality: scalability of particle filters with and without importance weights, Siam Review 61 (1) (2019) 79–91. doi:10.1137/17M1125340.
- [18] S. Reich, A dynamical systems framework for intermittent data assimilation, BIT Numerical Analysis 51 (2011) 235-249. doi:10.1007/ s10543-010-0302-4.
- [19] K. Bergemann, S. Reich, An ensemble Kalman-Bucy filter for continuous data assimilation, Meteorologische Zeitschrift 21 (3) (2012) 213.
- [20] T. Yang, P. G. Mehta, S. P. Meyn, Feedback particle filter, IEEE Transactions on Automatic Control 58 (10) (2013) 2465-2480. doi: 10.1109/TAC.2013.2258825.
- W. H. Fleming, S. K. Mitter, Optimal Control and Nonlinear Filtering for Nondegenerate Diffusion Processes, Stochastics 8 (1) (1982) 63-77. doi:10.1080/17442508208833228. URL https://doi.org/10.1080/17442508208833228
- [22] K. J. K. J. Astrom, Introduction to stochastic control theory, Mathematics in science and engineering; v. 70, Academic Press, New York, 1970.
- [23] T. Kailath, A. Sayed, B. Hassibi, Linear Estimation, Prentice-Hall information and system sciences series, Prentice Hall, 2000.
- [24] R. E. Mortensen, Maximum-likelihood recursive nonlinear filtering, Journal of Optimization Theory and Applications 2 (6) (1968) 386-394. doi:10.1007/BF00925744. URL https://doi.org/10.1007/BF00925744
- [25] J. W. Kim, P. G. Mehta, An optimal control derivation of nonlinear smoothing equations, in: Proceedings of the Workshop on Dynamics, Optimization and Computation held in honor of the 60th birthday of Michael Dellnitz, Springer, 2020, pp. 295–311.

- [26] S. Levine, Reinforcement learning and control as probabilistic inference: Tutorial and review (2018). arXiv:1805.00909.
- [27] D. Ocone, E. Pardoux, Asymptotic stability of the optimal filter with respect to its initial condition, SIAM Journal on Control and Optimization 34 (1) (1996) 226-243. doi:10.1137/s0363012993256617.
- [28] R. W. Brockett, Finite dimensional linear systems, SIAM, 2015.
- [29] D. Nualart, E. Pardoux, Stochastic calculus with anticipating integrands, Probability Theory and Related Fields 78 (4) (1988) 535–581.
- [30] P. G. Mehta, S. P. Meyn, Q-learning and Pontryagin's minimum principle, in: Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference, IEEE, 2009, pp. 3598–3605.
- [31] T. Yang, R. Laugesen, P. Mehta, S. Meyn, Multivariable feedback particle filter, Automatica 71 (2016) 10-23. doi:https: //doi.org/10.1016/j.automatica.2016.04.019. URL https://www.sciencedirect.com/science/article/pii/ S000510981630142X
- [32] A. N. Bishop, P. Del Moral, On the stability of matrix-valued Riccati diffusions, arXiv preprint arXiv:1808.00235 (2018).
   URL https://arxiv.org/abs/1808.00235
- [33] A. N. Bishop, P. Del Moral, On the mathematical theory of ensemble (linear-gaussian) kalman-bucy filtering (2020). doi:10.48550/ARXIV. 2006.08843. URL https://arxiv.org/abs/2006.08843
- [34] S. J. Bradtke, Reinforcement learning applied to linear quadratic regulation, in: NIPS, 1992, pp. 295-302. URL http://papers.nips.cc/paper/712-reinforcement-learning-applied-to-linear-
- [35] S. Bradtke, B. Ydstie, A. Barto, Adaptive linear quadratic control using policy iteration, in: Proceedings of 1994 American Control Conference ACC '94, Vol. 3, 1994, pp. 3475–3479 vol.3. doi:10.1109/ACC.1994. 735224.

- [36] K. G. Vamvoudakis, D. Vrabie, F. L. Lewis, Online adaptive learning of optimal control solutions using integral reinforcement learning, in: 2011 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL), 2011, pp. 250–257. doi:10.1109/ADPRL. 2011.5967359.
- [37] M. Palanisamy, H. Modares, F. L. Lewis, M. Aurangzeb, Continuoustime q-learning for infinite-horizon discounted cost linear quadratic regulator problems, IEEE Transactions on Cybernetics 45 (2) (2015) 165– 176. doi:10.1109/TCYB.2014.2322116.
- [38] R. S. Sutton, A. G. Barto, Reinforcement learning: an introduction, 2nd Edition, Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA, 2018.
- [39] S. Pathiraja, S. Reich, W. Stannat, Mckean-vlasov sdes in nonlinear filtering, arXiv preprint arXiv:2007.12658 (2020).
- [40] F. Daum, J. Huang, A. Noushin, Generalized Gromov method for stochastic particle flow filters, in: SPIE Defense+ Security, International Society for Optics and Photonics, 2017, pp. 102000I-102000I. doi:10.1117/12.2248723.
- [41] D. Crisan, J. Xiong, Approximate McKean-Vlasov representations for a class of SPDEs, Stochastics 82 (1) (2010) 53-68. doi:10.1080/ 17442500902723575.
- [42] R. S. Laugesen, P. G. Mehta, S. P. Meyn, M. Raginsky, Poisson's equation in nonlinear filtering, SIAM Journal on Control and Optimization 53 (1) (2015) 501–525. URL 10.1137/13094743X
- [43] A. Taghvaei, P. G. Mehta, S. P. Meyn, Diffusion map-based algorithm for gain function approximation in the feedback particle filter, SIAM/ASA Journal on Uncertainty Quantification 8 (3) (2020) 1090–1117.
- [44] T. Yang, R. S. Laugesen, P. G. Mehta, S. P. Meyn, Multivariable feedback particle filter, Automatica 71 (2016) 10-23. doi:10.1016/ j.automatica.2016.04.019.

- [45] C. Hartmann, C. Schütte, Efficient rare event simulation by optimal nonequilibrium forcing, Journal of Statistical Mechanics: Theory and Experiment 2012 (11) (2012) P11004.
- [46] W. H. Fleming, Exit probabilities and optimal stochastic control, Applied Mathematics and Optimization 4 (1) (1977) 329–346.
- [47] S. K. Mitter, N. J. Newton, A variational approach to nonlinear estimation, SIAM journal on control and optimization 42 (5) (2003) 1813–1833.
- [48] H. J. Kappen, Linear theory for control of nonlinear stochastic systems, Phys. Rev. Lett. 95 (2005) 200201. doi:10.1103/PhysRevLett.95. 200201. URL https://link.aps.org/doi/10.1103/PhysRevLett.95.200201
- [49] H. J. Kappen, Path integrals and symmetry breaking for optimal control theory, Journal of Statistical Mechanics: Theory and Experiment 2005 (11) (2005) P11011–P11011. doi:10.1088/1742-5468/2005/11/p11011.
  URL https://doi.org/10.1088/1742-5468/2005/11/p11011
- [50] E. Todorov, Linearly-solvable markov decision problems, in:
   B. Schölkopf, J. Platt, T. Hoffman (Eds.), Advances in Neural Information Processing Systems, Vol. 19, MIT Press, 2007.
   URL https://proceedings.neurips.cc/paper/2006/file/ d806ca13ca3449af72a1ea5aedbed26a-Paper.pdf
- [51] L. C. G. Rogers, D. Williams, Diffusions, Markov Processes and Martingales, 2nd Edition, Vol. 2 of Cambridge Mathematical Library, Cambridge University Press, 2000. doi:10.1017/CB09780511805141.
- [52] R. van Handel, Filtering, stability, and robustness, Ph.D. thesis, California Institute of Technology (2019). doi:10.7907/4p53-1h42. URL https://resolver.caltech.edu/CaltechETD: etd-12122006-164640
- [53] H. J. Kappen, H. C. Ruiz, Adaptive importance sampling for control and inference, Journal of Statistical Physics 162 (5) (2016) 1244–1266. doi:10.1007/s10955-016-1446-7. URL https://doi.org/10.1007/s10955-016-1446-7

[54] M. Toussaint, Robot trajectory optimization using approximate inference, in: Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09, Association for Computing Machinery, New York, NY, USA, 2009, p. 1049–1056. doi:10.1145/1553374. 1553508.

URL https://doi.org/10.1145/1553374.1553508

- [55] C. Hoffmann, P. Rostalski, Linear optimal control on factor graphs a message passing perspective —, IFAC-PapersOnLine 50 (1) (2017) 6314-6319, 20th IFAC World Congress. doi:https://doi.org/10.1016/j.ifacol.2017.08.914.
  URL https://www.sciencedirect.com/science/article/pii/S2405896317313800
- [56] S. Vijayakumar, K. Rawlik, M. Toussaint, On stochastic optimal control and reinforcement learning by approximate inference, in: N. Roy, P. Newman, S. Srinivasa (Eds.), Robotics: Science and Systems VIII, 2013, pp. 353–360.
- [57] [link]. URL https://github.com/anantjoshi97/EnKF-RL
- [58] H. Mohammadi, A. Zare, M. Soltanolkotabi, M. R. Jovanovic, Global exponential convergence of gradient methods over the nonconvex landscape of the linear quadratic regulator, in: 2019 IEEE 58th Conference on Decision and Control (CDC), 2019, pp. 7474–7479, iSSN: 2576-2370. doi:10.1109/CDC40024.2019.9029985.
- [59] R. Tedrake, Underactuated Robotics: Algorithms for Walking, Running, Swimming, Flying, and Manipulation (course Notes for MIT 6.832), last accessed on 16 May 2021. URL http://underactuated.mit.edu/
- [60] A. N. Bishop, P. Del Moral, On the mathematical theory of ensemble (linear-gaussian) kalman-bucy filtering, arXiv preprint arXiv:2006.08843 (2020).
- [61] P. del Moral, S. S. Singh, A forward-backward stochastic analysis of diffusion flows, arXiv preprint arXiv:1906.09145 (2019).

- [62] A. N. Bishop, P. Del Moral, On the stability of matrix-valued riccati diffusions, Electronic Journal of Probability 24 (2019) 1–40.
- [63] A. N. Bishop, P. Del Moral, On the stability of Kalman-Bucy diffusion processes, SIAM Journal on Control and Optimization 55 (6) (2017) 4015-4047. doi:10.1137/16m1102707.
- [64] P. E. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Springer, Applications of Mathematics, 1999.