
Computational Doob’s h -transforms for Online Filtering of Discretely Observed Diffusions

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 This paper is concerned with online filtering of discretely observed nonlinear diffu-
2 sion processes. Our approach is based on the fully adapted auxiliary particle filter
3 which involves Doob’s h -transforms that are typically intractable. We propose a
4 computational framework to approximate these h -transforms by solving the under-
5 lying backward Kolmogorov equations using nonlinear Feynman-Kac formulas.
6 The methodology allows one to train a locally optimal particle filter prior to the
7 data-assimilation procedure. Numerical experiments illustrate that the proposed
8 approach can be orders of magnitude more efficient than the bootstrap particle fil-
9 ter in the regime of highly informative observations and when the observations are
10 extreme under the model.

11 1 Introduction

12 Diffusion processes are fundamental tools in applied mathematics, statistics and machine learning.
13 This rich class of continuous-time models have been used to model real-world phenomena in disci-
14 plines as varied as life-sciences, engineering, economics and finance. However, working with dif-
15 fusions can be challenging as its transition densities are only tractable in simple and specific cases
16 such as (geometric) Brownian motions, Ornstein-Uhlenbeck (OU) processes and Cox-Ingersoll-Ross
17 processes. This difficulty has hindered the use of standard methodologies for inference and data-
18 assimilation of models driven by diffusion processes. Specialized methodologies have been devel-
19 oped to circumvent or mitigate these issues [35, 4, 3, 14, 13, 5, 37].

20 Consider a time-homogeneous multivariate diffusion process $d\mathbf{X}_t = \mu(\mathbf{X}_t) dt + \sigma(\mathbf{X}_t) dB_t$ that
21 is discretely observed at regular intervals. Noisy observations \mathbf{y}_k of the latent process \mathbf{X}_{t_k} is
22 collected at time $t_k \equiv kT$ for $k \geq 1$. We consider the online filtering problem which consists in
23 estimating the conditional laws $\pi_k(d\mathbf{x}) = \mathbb{P}(\mathbf{X}_{t_k} \in d\mathbf{x} | \mathbf{y}_1, \dots, \mathbf{y}_k)$, i.e. the filtering distributions,
24 as observations are collected. We focus on the use of Particle Filters (PF) that approximate the
25 filtering distributions with a system of weighted particles. Although many previous works have
26 relied on the Bootstrap Particle Filter (BPF), which simulates particles from the diffusion process,
27 it can perform poorly in challenging scenarios as it fails to take the incoming observation \mathbf{y}_k into
28 account. This issue can be partially tackled by relying on resampling at intermediate times between
29 observations [10, 31]. The goal of this article is to show that the (locally) optimal approach given by
30 the Fully Adapted Auxiliary Particle Filter (FA-APF) [33] can be implemented. This necessitates
31 simulating a conditioned diffusion process, which can be formulated as a control problem involving
32 an intractable Doob’s h -transform [36, 8]. We propose the *Computational Doob’s h -Transform*
33 (CDT) framework for efficiently approximating these quantities. The method relies on nonlinear
34 Feynman-Kac formulas for solving backward Kolmogorov equations simultaneously for all possible
35 observations. Importantly, this preprocessing step only needs to be performed once before starting the
36 online filtering procedure. Numerical experiments illustrate that the proposed approach can be orders

37 of magnitude more efficient than the BPF in the regime of highly informative observations or when
 38 the observations are extreme under the model. A PyTorch implementation to reproduce our numerical
 39 experiments is available at <https://anonymous.4open.science/r/CompDoobTransform/>.

40 **Notations.** For two matrices $A, B \in \mathbb{R}^{d,d}$, their Frobenius inner product is defined as $\langle A, B \rangle_F =$
 41 $\sum_{i,j=1}^d A_{i,j} B_{i,j}$. The Euclidean inner product for $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$ is denoted as $\langle \mathbf{u}, \mathbf{v} \rangle = \sum_{i=1}^d u_i v_i$.

42 2 Background

43 2.1 Filtering of discretely observed diffusions

44 Consider an homogeneous diffusion process $\{\mathbf{X}_t\}_{t \geq 0}$ in $\mathcal{X} = \mathbb{R}^d$ with initial distribution $\rho_0(dx)$ and
 45 dynamics

$$d\mathbf{X}_t = \mu(\mathbf{X}_t) dt + \sigma(\mathbf{X}_t) dB_t, \quad (1)$$

46 described by the drift and volatility functions $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d,d}$. We assume standard
 47 smoothness and growth conditions [27] for a unique strong solution of (1) to exist for all times. The
 48 associated semi-group of transition probabilities $p_s(d\hat{\mathbf{x}} | \mathbf{x})$ satisfies $\mathbb{P}(\mathbf{X}_{t+s} \in A | \mathbf{X}_t = \mathbf{x}) =$
 49 $\int_A p_s(d\hat{\mathbf{x}} | \mathbf{x})$ for any $s, t > 0$ and measurable $A \subset \mathcal{X}$. The process $\{\mathbf{B}_t\}_{t \geq 0}$ is a standard \mathbb{R}^d -
 50 valued Brownian motion. The diffusion process $\{\mathbf{X}_t\}_{t \geq 0}$ is discretely observed at time $t_k = kT$,
 51 for $k \geq 1$, for some inter-observation time $T > 0$. The \mathcal{Y} -valued observation $\mathbf{Y}_k \in \mathcal{Y}$ at time t_k is
 52 modelled by the likelihood function $g : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ in the sense that for any measurable $A \subset \mathcal{Y}$,
 53 we have $\mathbb{P}(\mathbf{Y}_k \in A | \mathbf{X}_{t_k} = \mathbf{x}_k) = \int_A g(\mathbf{x}_k, \mathbf{y}) d\mathbf{y}$ for some dominating measure $d\mathbf{y}$ on \mathcal{Y} . For a
 54 test function $\varphi : \mathcal{X} \rightarrow \mathbb{R}$, the generator of the diffusion process $\{\mathbf{X}_t\}_{t \geq 0}$ is given by

$$\mathcal{L}\varphi = \langle \mu, \nabla \varphi \rangle + \frac{1}{2} \langle \sigma \sigma^\top, \nabla^2 \varphi \rangle_F. \quad (2)$$

55 This article is concerned with approximating the filtering distributions $\pi_k(dx) = \mathbb{P}(\mathbf{X}_{t_k} \in dx |$
 56 $\mathbf{y}_1, \dots, \mathbf{y}_k)$. For notational convenience, we set $\pi_0(dx) \equiv \rho_0(dx)$.

57 2.2 Particle filtering

58 Particle Filters (PF), also known as Sequential Monte Carlo methods, are a set of Monte Carlo
 59 algorithms that can be used to solve filtering problems (see [7] for a recent textbook on the topic). PFs
 60 evolve a set of $N \geq 1$ particles $\mathbf{x}_t^{1:N} = (\mathbf{x}_t^1, \dots, \mathbf{x}_t^N) \in \mathcal{X}^N$ forward in time using a combination of
 61 *propagation* and *resampling* operations.

62 To initialize the PF, each initial particle $\mathbf{x}_0^j \in \mathcal{X}$ for $1 \leq j \leq N$ is sampled independently from
 63 the distribution $\rho_0(dx)$ so that $\pi_0(dx) \approx N^{-1} \sum_{j=1}^N \delta(dx; \mathbf{x}_0^j)$. Approximations of the filtering
 64 distribution π_k for $k \geq 1$ are built recursively as follows. Given the Monte Carlo approximation
 65 of the filtering distribution at time t_k , $\pi_k(dx) \approx N^{-1} \sum_{j=1}^N \delta(dx; \mathbf{x}_{t_k}^j)$, the particles $\mathbf{x}_{t_k}^{1:N}$ are
 66 propagated independently forward in time by $\hat{\mathbf{x}}_{t_{k+1}}^j \sim q_{k+1}(d\hat{\mathbf{x}} | \mathbf{x}_{t_k}^j)$, using a Markov kernel
 67 $q_{k+1}(d\hat{\mathbf{x}} | \mathbf{x})$ specified by the user. The BPF corresponds to the choice of Markov kernel $q_{k+1}^{\text{BPF}}(d\hat{\mathbf{x}} |$
 68 $\mathbf{x}) = \mathbb{P}(\mathbf{X}_{t_{k+1}} \in d\hat{\mathbf{x}} | \mathbf{X}_{t_k} = \mathbf{x})$ while the FA-APF [33] corresponds to the choice

$$q_{k+1}^{\text{FA-APF}}(d\hat{\mathbf{x}} | \mathbf{x}) = \mathbb{P}(\mathbf{X}_{t_{k+1}} \in d\hat{\mathbf{x}} | \mathbf{X}_{t_k} = \mathbf{x}, \mathbf{Y}_{k+1} = \mathbf{y}_{k+1}). \quad (3)$$

69 Each particle $\hat{\mathbf{x}}_{t_{k+1}}^j$ is associated with a normalized weight $\bar{W}_{k+1}^j = W_{k+1}^j / \sum_{i=1}^N W_{k+1}^i$, where the
 70 unnormalized weights $W_{k+1}^j > 0$ are defined as

$$W_{k+1}^j = \frac{p_T(d\hat{\mathbf{x}}_{t_{k+1}}^j | \mathbf{x}_{t_k}^j)}{q_{k+1}(d\hat{\mathbf{x}}_{t_{k+1}}^j | \mathbf{x}_{t_k}^j)} g(\hat{\mathbf{x}}_{t_{k+1}}^j, \mathbf{y}_{k+1}). \quad (4)$$

71 The BPF and FA-APF correspond respectively to having

$$W_{k+1}^{j,\text{BPF}} = g(\hat{\mathbf{x}}_{t_{k+1}}^j, \mathbf{y}_{k+1}) \quad \text{and} \quad W_{k+1}^{j,\text{FA-APF}} = \mathbb{E}[g(\mathbf{X}_{t_{k+1}}, \mathbf{y}_{k+1}) | \mathbf{X}_{t_k} = \mathbf{x}_{t_k}^j]. \quad (5)$$

72 The weights are such that $\pi_{k+1}(d\mathbf{x}) \approx \sum_{j=1}^N \bar{W}_{k+1}^j \delta(d\mathbf{x}; \mathbf{x}_{t_{k+1}}^j)$. The *resampling* step consists in
73 defining a new set of particles $\mathbf{x}_{t_{k+1}}^{1:N}$ with $\mathbb{P}(\mathbf{x}_{t_{k+1}}^j = \hat{\mathbf{x}}_{t_{k+1}}^i) = \bar{W}_{k+1}^i$. This resampling scheme
74 ensures that the equally weighted set of particles $\mathbf{x}_{t_{k+1}}^{1:N}$ provides a Monte Carlo approximation of the
75 filtering distribution at time t_{k+1} in the sense that $\pi_{k+1}(d\mathbf{x}) \approx N^{-1} \sum_{j=1}^N \delta(d\mathbf{x}; \mathbf{x}_{t_{k+1}}^j)$. Note that
76 the particles $\mathbf{x}_{t_{k+1}}^{1:N}$ do not need to be resampled independently given the set of propagated particles
77 $\hat{\mathbf{x}}_{t_{k+1}}^{1:N}$. We refer the reader to [15] for a recent discussion of resampling schemes within PFs and to
78 [9] for a book-length treatment of the convergence properties of this class of Monte Carlo methods.

79 In most settings, the FA-APF [33] that minimizes a local variance criterion [12] leads to better
80 performance when compared to the BPF. This gain in efficiency can be very substantial when the
81 signal-to-noise ratio is high or when observations contain outliers under the model specification.
82 Nevertheless, implementing FA-APF requires sampling from the conditioned transition probability
83 in (3), which is typically not feasible in practice. We will show in the following that this can be
84 achieved in our setting by simulating a conditioned diffusion. We note also that standard strategies to
85 approximate the FA-APF do not apply to our setup as the latent state process evolves on a higher
86 frequency relative to the observations.

87 2.3 Conditioned and controlled diffusions

88 As the diffusion process (1) is assumed to be time-homogeneous, it suffices to focus on the initial
89 interval $[0, T]$ and study the dynamics of the diffusion $\mathbf{X}_{[0,T]} = \{\mathbf{X}_t\}_{t \in [0,T]}$ conditioned upon the
90 first observation $\mathbf{Y}_T = \mathbf{y}$. The conditioned dynamics can also be described by a diffusion process.
91 Contrarily to the original diffusion, the conditioned process is not time-homogeneous in general. The
92 conditioned process is described by the same volatility function but with a different drift term that
93 takes the future observation $\mathbf{Y}_T = \mathbf{y}$ into account.

94 Before deriving the exact form of the conditioned diffusion in Section 2.4, this section describes a
95 more general setting that will be of crucial importance in our proposed numerical scheme. For a
96 *control* function $\mathbf{c} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ and a given observation $\mathbf{y} \in \mathcal{Y}$, consider the controlled
97 diffusion process $\{\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}\}_{t \in [0, T]}$ satisfying

$$d\mathbf{X}_t^{\mathbf{c}, \mathbf{y}} = \underbrace{\mu(\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}) dt + \sigma(\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}) d\mathbf{B}_t}_{\text{(original dynamics)}} + \underbrace{[\sigma \mathbf{c}](\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}, \mathbf{y}, t) dt}_{\text{(control drift term)}}. \quad (6)$$

98 The dynamics of the controlled diffusion is identical to the original diffusion, except for the additional
99 drift term $[\sigma \mathbf{c}](\mathbf{x}, \mathbf{y}, t) \in \mathbb{R}^d$ described by the control function \mathbf{c} . For $\mathbf{y} \in \mathcal{Y}$ and a test function
100 $\varphi : \mathcal{X} \rightarrow \mathbb{R}$, the generator of the controlled diffusion is given by

$$\mathcal{L}^{\mathbf{c}, \mathbf{y}, t} \varphi(\mathbf{x}) = \mathcal{L} \varphi(\mathbf{x}) + \langle [\sigma \mathbf{c}](\mathbf{x}, \mathbf{y}, t), \nabla \varphi(\mathbf{x}) \rangle. \quad (7)$$

101 Let $\mathbb{P}_{[0,T]}$ and $\mathbb{P}_{[0,T]}^{\mathbf{c}, \mathbf{y}}$ denote the probability measures on the space of continuous functions
102 $C([0, T], \mathbb{R}^d)$ generated by the original and controlled diffusions respectively. Under mild growth as-
103 sumptions on the control \mathbf{c} , the two measures are equivalent and Girsanov's theorem [16] shows that

$$\frac{d\mathbb{P}_{[0,T]}}{d\mathbb{P}_{[0,T]}^{\mathbf{c}, \mathbf{y}}}(\mathbf{X}_{[0,T]}) = \exp \left\{ -\frac{1}{2} \int_0^T \|\mathbf{c}(\mathbf{X}_t, \mathbf{y}, t)\|^2 dt - \int_0^T \langle \mathbf{c}(\mathbf{X}_t, \mathbf{y}, t), d\mathbf{B}_t \rangle \right\}. \quad (8)$$

104 Our main objective is to construct a control function $\mathbf{c}_* : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ so that, for any
105 observation value $\mathbf{y} \in \mathcal{Y}$, the controlled diffusion $\mathbf{X}_{[0,T]}^{\mathbf{c}_*, \mathbf{y}}$ has the same dynamics as the original
106 diffusion $\mathbf{X}_{[0,T]}$ conditioned upon the observation $\mathbf{Y}_T = \mathbf{y}$, i.e. for any measurable set $A \subset$
107 $C([0, T], \mathbb{R}^d)$, we have

$$\mathbb{P}(\mathbf{X}_{[0,T]}^{\mathbf{c}_*, \mathbf{y}} \in A) = \mathbb{E}[\mathbf{1}(\mathbf{X}_{[0,T]} \in A) g(\mathbf{X}_T, \mathbf{y})] / \mathbb{E}[g(\mathbf{X}_T, \mathbf{y})]. \quad (9)$$

108 We will give an exact expression of this control in Section 2.4 and propose a numerical scheme to
109 approximate it in Section 3.1.

110 **2.4 Doob's h -transform**

111 To simplify notation, we shall denote the conditioned process $\mathbf{X}_{[0,T]} \mid (\mathbf{Y}_T = \mathbf{y})$ as $\widehat{\mathbf{X}}_{[0,T]}$. To
 112 describe its dynamics, we introduce the function

$$h(\mathbf{x}, \mathbf{y}, t) = \mathbb{E}[g(\mathbf{X}_T, \mathbf{y}) \mid \mathbf{X}_t = \mathbf{x}] = \int_{\mathcal{X}} g(\mathbf{x}_T, \mathbf{y}) p_{T-t}(d\mathbf{x}_T \mid \mathbf{x}) \quad (10)$$

113 which gives the probability of observing $\mathbf{Y}_T = \mathbf{y}$ when the diffusion has state $\mathbf{x} \in \mathcal{X}$ at time
 114 $t \in [0, T]$. We recall that the likelihood function $g : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ was defined in Section 2.1.
 115 The definition in (10) implies that the function $h : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}_+$ satisfies the backward
 116 Kolmogorov equation [27],

$$(\partial_t + \mathcal{L})h = 0, \quad (11)$$

117 with terminal condition $h(\mathbf{x}, \mathbf{y}, T) = g(\mathbf{x}, \mathbf{y})$ for all $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$. For $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ and an
 118 infinitesimal increment $\delta > 0$, we have

$$\begin{aligned} \mathbb{E}[\varphi(\widehat{\mathbf{X}}_{t+\delta}) \mid \widehat{\mathbf{X}}_t = \mathbf{x}] &= \mathbb{E}[\varphi(\mathbf{X}_{t+\delta}) g(\mathbf{X}_T, \mathbf{y}) \mid \mathbf{X}_t = \mathbf{x}] / \mathbb{E}[g(\mathbf{X}_T, \mathbf{y}) \mid \mathbf{X}_t = \mathbf{x}] \\ &= \mathbb{E}[\varphi(\mathbf{X}_{t+\delta}) h(\mathbf{X}_{t+\delta}, \mathbf{y}, t + \delta) \mid \mathbf{X}_t = \mathbf{x}] / h(\mathbf{x}, \mathbf{y}, t) \\ &= \varphi(\mathbf{x}) + \delta \left\{ \frac{\mathcal{L}[\varphi h]}{h} \right\}(\mathbf{x}, \mathbf{y}, t) + O(\delta^2). \end{aligned} \quad (12)$$

119 Furthermore, since the function h satisfies (11), some algebra shows that $\mathcal{L}[\varphi h]/h = \mathcal{L}\varphi +$
 120 $\langle \sigma \sigma^\top \nabla \log h, \nabla \varphi \rangle$. By taking $\delta \rightarrow 0$, this heuristic derivation shows that the generator of the condi-
 121 tioned diffusion equals $\mathcal{L}\varphi + \langle \sigma \sigma^\top \nabla \log h, \nabla \varphi \rangle$. Hence $\widehat{\mathbf{X}}_{[0,T]}$ satisfies the dynamics of a controlled
 122 diffusion (6) with control function $\mathbf{c}_*(\mathbf{x}, \mathbf{y}, t) = [\sigma^\top \nabla \log h](\mathbf{x}, \mathbf{y}, t)$. We refer readers to [36, 8] for
 123 a formal treatment of Doob's h -transform.

124 **2.5 Nonlinear Feynman-Kac formula**

125 Obtaining the control function $\mathbf{c}_*(\mathbf{x}, \mathbf{y}, t) = [\sigma^\top \nabla \log h](\mathbf{x}, \mathbf{y}, t)$ by solving the backward Kol-
 126 mogorov equation in (11) for each possible observation $\mathbf{y} \in \mathcal{Y}$ is computationally not feasible. Fur-
 127 thermore, when the dimensionality of the state-space \mathcal{X} becomes larger, standard numerical methods
 128 for solving Partial Differential Equations (PDEs) such as Finite Differences or the Finite Element
 129 Method become impractical. For these reasons, we propose instead to approximate the control func-
 130 tion \mathbf{c}_* with neural networks, and employ methods based on automatic differentiation and the nonlin-
 131 ear Feynman-Kac approach to solve semilinear PDEs [19, 20, 24, 17, 6, 22, 23, 1, 18].

132 As the non-negative function h typically decays exponentially for large $\|\mathbf{x}\|$, it is computationally
 133 more stable to work on the logarithmic scale and approximate the *value* function $v(\mathbf{x}, \mathbf{y}, t) =$
 134 $-\log[h(\mathbf{x}, \mathbf{y}, t)]$. Using the fact that h satisfies the PDE (11), the value function satisfies

$$(\partial_t + \mathcal{L})v = \frac{1}{2} \|\sigma^\top \nabla v\|^2, \quad v(\mathbf{x}, \mathbf{y}, T) = -\log[g(\mathbf{x}, \mathbf{y})] \quad \text{for all } (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}. \quad (13)$$

135 Let $\{\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}\}_{t \in [0, T]}$ be a controlled diffusion defined in Equation (6) with a given control function
 136 $\mathbf{c} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ and define the diffusion process $\{V_t\}_{t \in [0, T]}$ as $V_t = v(\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}, \mathbf{y}, t)$. Itô
 137 Lemma shows that for any observation $\mathbf{Y}_T = \mathbf{y}$ and $0 \leq s \leq T$, we have

$$V_T = V_s + \int_s^T \left(\frac{1}{2} \|\mathbf{Z}_t\|^2 + \langle \mathbf{c}, \mathbf{Z}_t \rangle \right) dt + \int_s^T \langle \mathbf{Z}_t, d\mathbf{B}_t \rangle$$

138 with $\mathbf{Z}_t = [\sigma^\top \nabla v](\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}, \mathbf{y}, t)$ and $V_T = -\log[g(\mathbf{X}_T^{\mathbf{c}, \mathbf{y}}, \mathbf{y})]$. In summary, the pair of processes
 139 (V_t, \mathbf{Z}_t) are such that the following equation holds,

$$-\log[g(\mathbf{X}_T^{\mathbf{c}, \mathbf{y}}, \mathbf{y})] = V_s + \int_s^T \left\{ \frac{1}{2} \|\mathbf{Z}_t\|^2 + \langle \mathbf{c}, \mathbf{Z}_t \rangle \right\} dt + \int_s^T \langle \mathbf{Z}_t, d\mathbf{B}_t \rangle. \quad (14)$$

140 Crucially, under mild growth and regularity assumptions on the drift and volatility function $\mu :$
 141 $\mathcal{X} \rightarrow \mathbb{R}^d$ and $\sigma : \mathcal{X} \rightarrow \mathbb{R}^{d,d}$, the pair of processes (V_t, \mathbf{Z}_t) is the unique solution to Equation (14)
 142 [28, 29, 30, 40]. This result can be used as a building block for designing Monte Carlo approximations
 143 of the solution to semilinear and fully nonlinear PDEs [18, 34, 21]

144 **3 Method**

145 **3.1 Computational Doob's h -transform**

146 As before, consider a diffusion $\{\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}\}_{t \in [0, T]}$ controlled by a function $\mathbf{c} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ and
 147 driven by the standard Brownian motion $\{\mathbf{B}_t\}_{t \geq 0}$. Furthermore, for two functions $N_0 : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$
 148 and $N : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$, consider the diffusion process $\{V_t\}_{t \in [0, T]}$ defined as

$$V_t = V_0 + \int_0^t \left\{ \frac{1}{2} \|\mathbf{Z}_t\|^2 + \langle \mathbf{c}(\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}, \mathbf{y}, t), \mathbf{Z}_t \rangle \right\} dt + \int_0^t \langle \mathbf{Z}_t, d\mathbf{B}_t \rangle, \quad (15)$$

149 where the initial condition V_0 and the process $\{\mathbf{Z}_t\}_{t \in [0, T]}$ are defined as

$$V_0 = N_0(\mathbf{X}_0^{\mathbf{c}, \mathbf{y}}, \mathbf{y}) \quad \text{and} \quad \mathbf{Z}_t = N(\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}, \mathbf{y}, t). \quad (16)$$

150 Importantly, we remind the reader that the two diffusion processes $\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}$ and V_t are driven by the
 151 same Brownian motion \mathbf{B}_t . The uniqueness result mentioned at the end of Section 2.5 implies
 152 that, if for any choice of initial condition $\mathbf{X}_0^{\mathbf{c}, \mathbf{y}} \in \mathcal{X}$ and terminal observation $\mathbf{y} \in \mathcal{Y}$ the condition
 153 $V_T = -\log[g(\mathbf{X}_T^{\mathbf{c}, \mathbf{y}}, \mathbf{y})]$ is satisfied, then we have that for all $(\mathbf{x}, \mathbf{y}, t) \in \mathcal{X} \times \mathcal{Y} \times [0, T]$

$$N_0(\mathbf{x}, \mathbf{y}) = -\log h(\mathbf{x}, \mathbf{y}, 0) \quad \text{and} \quad N(\mathbf{x}, \mathbf{y}, t) = -[\sigma^\top \nabla \log h](\mathbf{x}, \mathbf{y}, t). \quad (17)$$

154 In particular, the optimal control is given by $\mathbf{c}_*(\mathbf{x}, \mathbf{y}, t) = -N(\mathbf{x}, \mathbf{y}, t)$.

155 These remarks suggest parametrizing the functions $N_0(\cdot, \cdot)$ and $N(\cdot, \cdot, \cdot)$ by two neural networks with
 156 respective parameters $\theta_0 \in \Theta_0$ and $\theta \in \Theta$ while minimizing the loss function

$$L(\theta_0, \theta; \mathbf{c}) = \mathbb{E} \left[\left(V_T + \log[g(\mathbf{X}_T^{\mathbf{c}, \mathbf{Y}}, \mathbf{Y})] \right)^2 \right]. \quad (18)$$

157 The above expectation is with respect to the distribution of the Brownian motion $\{\mathbf{B}_t\}_{t \geq 0}$, the initial
 158 condition $\mathbf{X}_0^{\mathbf{c}, \mathbf{Y}} \sim \eta_{\mathbf{X}}(d\mathbf{x})$ of the controlled diffusion, and the observation $\mathbf{Y} \sim \eta_{\mathbf{Y}}(d\mathbf{y})$ at time T .
 159 In practice, we will let the three sources of randomness be independent of each other. The spread
 160 of the distributions $\eta_{\mathbf{X}}$ and $\eta_{\mathbf{Y}}$ should be large enough to cover typical states under the filtering
 161 distributions $\pi_k, k \geq 1$ and future observations to be filtered respectively. Specific choices will be
 162 detailed for each application in Section 4. For offline problems, one could learn in a data-driven
 163 manner by selecting $\eta_{\mathbf{Y}}$ as the empirical distribution of actual observations. Furthermore, any control
 164 function $\mathbf{c} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ with mild growth and regularity assumptions can be employed
 165 within our methodology: specific choices are discussed at the end of this section.

166 **CDT algorithm.** The following outlines our training procedure to learn neural networks N_0 and N
 167 that satisfy (17). To minimize the loss function (18), any stochastic gradient algorithm can be used
 168 with a user-specified *mini-batch* size of $J \geq 1$. The following steps are iterated until convergence.

- 169 1. Choose a control $\mathbf{c} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$, possibly based on the current neural network
 170 parameters $(\theta_0, \theta) \in \Theta_0 \times \Theta$.
- 171 2. Simulate independent Brownian paths $\mathbf{B}_{[0, T]}^j$, initial conditions $\mathbf{X}_0^j \sim \eta_{\mathbf{X}}(d\mathbf{x})$, and observa-
 172 tions $\mathbf{Y}^j \sim \eta_{\mathbf{Y}}(d\mathbf{y})$ for $1 \leq j \leq J$.
- 173 3. Generate the controlled trajectories: the j -th sample path $\mathbf{X}_{[0, T]}^j$ is obtained by forward
 174 integration of the controlled dynamics in Equation (6) with initial condition \mathbf{X}_0^j , control
 175 $\mathbf{c}(\cdot, \mathbf{Y}^j, \cdot)$, and the Brownian path $\mathbf{B}_{[0, T]}^j$.
- 176 4. Generate the value trajectories: the j -th sample path $V_{[0, T]}^j$ is obtained by forward integration
 177 of the dynamics in Equation (15)–(16) with the Brownian path $\mathbf{B}_{[0, T]}^j$ and the current neural
 178 network parameters $(\theta_0, \theta) \in \Theta_0 \times \Theta$.
- 179 5. Construct a Monte Carlo estimate of the loss function (18):

$$\widehat{L} = J^{-1} \sum_{j=1}^J (V_T^j + \log[g(\mathbf{X}_T^j, \mathbf{Y}^j)])^2 \quad (19)$$

180 6. Use automatic differentiation to compute $\partial_{\theta_0} \widehat{L}$ and $\partial_{\theta} \widehat{L}$ and update the parameters (θ_0, θ) .

181 Importantly, if the control function \mathbf{c} in *Step:1* does depend on the current parameters (θ_0, θ) , the
 182 gradient operations executed in *Step:6* should not be propagated through the control function \mathbf{c} . A
 183 standard `stop-gradient` operation available in most popular automatic differentiation frameworks
 184 can be used for this purpose.

185 **Time-discretization of diffusions.** For clarity of exposition, we have described our algorithm in
 186 continuous-time. In practice, one would have to discretize these diffusion processes, which is entirely
 187 straightforward. Although any numerical integrator could potentially be considered, the experiments
 188 in Section 4 employed the standard Euler-Maruyama scheme [25].

189 **Parametrizations of functions N_0 and N .** In all numerical experiments presented in Section 4, the
 190 functions N_0 and N are parametrized with fully-connected neural networks with two hidden layers
 191 and the Leaky ReLU activation function except in the last layer. Future work could explore other
 192 neural network architectures for our setting.

193 **Choice of controlled dynamics.** In challenging scenarios where observations are highly informative
 194 and/or extreme under the model, choosing a good control function to implement *Step:1* of the
 195 proposed algorithm can be crucial. We focus on two possible implementations:

- 196 • **CDT static scheme:** a simple (and naive) choice is not using any control, i.e. $\mathbf{c}(\mathbf{x}, \mathbf{y}, t) \equiv$
 197 $0 \in \mathbb{R}^d$ for all $(\mathbf{x}, \mathbf{y}, t) \in \mathcal{X} \times \mathcal{Y} \times [0, T]$.
- 198 • **CDT iterative scheme:** use the current approximation of the optimal control \mathbf{c}_* described
 199 by the parameters $(\theta_0, \theta) \in \Theta_0 \times \Theta$. This corresponds to setting $\mathbf{c}(\mathbf{x}, \mathbf{y}, t) = -N(\mathbf{x}, \mathbf{y}, t)$.

200 While using a *static control* approach can perform reasonably well in some situations, our results in
 201 Section 4 suggest that the *iterative control* procedure is a more reliable strategy. This is consistent
 202 with findings in the stochastic optimal control literature [38, 32]. This choice of control function
 203 drives the forward process $\mathbf{X}_t^{\mathbf{c}, \mathbf{y}}$ to regions of the state-space where the likelihood function is large and
 204 helps mitigate convergence and stability issues. Furthermore, Section 4 reports that (at convergence),
 205 the solutions N_0 and N can be significantly different. The *iterative control* procedure leads to more
 206 accurate solutions and, ultimately, better performance when used for online filtering.

207 3.2 Online filtering

208 Before performing online filtering, we first run the CDT algorithm described in Section 3.1 to construct
 209 an approximation of the optimal control $\mathbf{c}_*(\mathbf{x}, \mathbf{y}, t) = [\sigma^\top \nabla \log h](\mathbf{x}, \mathbf{y}, t)$. For concreteness, denote
 210 by $\widehat{\mathbf{c}} : \mathcal{X} \times \mathcal{Y} \times [0, T] \rightarrow \mathbb{R}^d$ the resulting approximate control, i.e. $\widehat{\mathbf{c}}(\mathbf{x}, \mathbf{y}, t) = -N(\mathbf{x}, \mathbf{y}, t)$ where
 211 $N(\cdot, \cdot, \cdot)$ is parametrized by the final parameter $\theta \in \Theta$. Similarly, denote by $\widehat{V}_0 : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ the
 212 approximation of the initial value function $v(\mathbf{x}, \mathbf{y}, 0) = -\log h(\mathbf{x}, \mathbf{y}, 0)$, i.e. $\widehat{V}_0(\mathbf{x}, \mathbf{y}) = N_0(\mathbf{x}, \mathbf{y})$
 213 where $N_0(\cdot, \cdot)$ is parametrized by the final parameter $\theta_0 \in \Theta_0$.

214 For implementing online filtering with $N \geq 1$ particles, consider a current approximation $\pi_k(dx) =$
 215 $N^{-1} \sum_{j=1}^N \delta(dx; \mathbf{x}_{t_k}^j)$ of the filtering distributions at time $t_k \geq 0$. Given the future observation
 216 $\mathbf{Y}_{k+1} = \mathbf{y}_{k+1}$, the particles $\mathbf{x}_{t_k}^{1:N}$ are then propagated forward by exploiting the approximately
 217 optimal control $(\mathbf{x}, t) \mapsto \widehat{\mathbf{c}}(\mathbf{x}, \mathbf{y}_{k+1}, t - t_k)$. In particular, $\widehat{\mathbf{x}}_{t_{k+1}}^j$ is obtained by setting $\widehat{\mathbf{x}}_{t_{k+1}}^j = \widehat{\mathbf{X}}_{t_{k+1}}^j$
 218 where $\{\widehat{\mathbf{X}}_t^j\}_{t \in [t_k, t_{k+1}]}$ follows the controlled diffusion

$$d\widehat{\mathbf{X}}_t^j = \underbrace{\mu(\widehat{\mathbf{X}}_t^j) dt + \sigma(\widehat{\mathbf{X}}_t^j) dB_t^j}_{\text{(original dynamics)}} + \underbrace{[\sigma \widehat{\mathbf{c}}](\widehat{\mathbf{X}}_t^j, \mathbf{y}_{k+1}, t - t_k) dt}_{\text{(approximately optimal control)}} \quad (20)$$

219 initialized at $\widehat{\mathbf{X}}_{t_k}^j = \mathbf{x}_{t_k}^j$. Each propagated particle $\widehat{\mathbf{x}}_{t_{k+1}}^j$ is associated with a normalized
 220 weight $\overline{W}_{k+1}^j = W_{k+1}^j / \sum_{i=1}^N W_{k+1}^i$ where $W_{k+1}^j = (d\mathbb{P}_{[t_k, t_{k+1}]}^{\widehat{\mathbf{c}}, \mathbf{y}_{k+1}} / d\mathbb{P}_{[t_k, t_{k+1}]}^{\widehat{\mathbf{c}}, \mathbf{y}_{k+1}})(\widehat{\mathbf{X}}_{[t_k, t_{k+1}]}^j) \times$
 221 $g(\widehat{\mathbf{x}}_{t_{k+1}}^j, \mathbf{y}_{k+1})$. We recall that the probability measures $\mathbb{P}_{[t_k, t_{k+1}]}$ and $\mathbb{P}_{[t_k, t_{k+1}]}^{\widehat{\mathbf{c}}, \mathbf{y}_{k+1}}$ correspond to the
 222 original and controlled diffusions on the interval $[t_k, t_{k+1}]$. Girsanov's theorem, as described in Equa-
 223 tion (8), implies that

$$W_{k+1}^j = \exp \left\{ -\frac{1}{2} \int_{t_k}^{t_{k+1}} \|\mathbf{Z}_t^j\|^2 dt + \int_{t_k}^{t_{k+1}} \langle \mathbf{Z}_t^j, d\mathbf{B}_t^j \rangle + \log g(\mathbf{x}_{t_{k+1}}^j, \mathbf{y}_{k+1}) \right\}$$

224 where $\mathbf{Z}_t^j = -\widehat{\mathbf{c}}(\widehat{\mathbf{X}}_t^j, \mathbf{y}_{k+1}, t - t_k)$. Similarly to Equation (15), consider the diffusion process
 225 $\{V_t^j\}_{t \in [t_k, t_{k+1}]}$ defined by the dynamics $dV_t^j = -\frac{1}{2}\|\mathbf{Z}_t^j\|^2 dt + \langle \mathbf{Z}_t^j, d\mathbf{B}_t^j \rangle$ with initialization at
 226 $V_{t_k}^j = \widehat{V}_0(\mathbf{x}_{t_k}^j, \mathbf{y}_{k+1})$. Therefore the weight can be re-written as

$$W_{k+1}^j = \exp \left\{ \underbrace{V_{t_{k+1}}^j + \log g(\mathbf{x}_{t_{k+1}}^j, \mathbf{y}_{k+1})}_{\approx 0} \right\} \exp \left\{ -\widehat{V}_0(\mathbf{x}_{t_k}^j, \mathbf{y}_{k+1}) \right\}, \quad (21)$$

227 and computed by numerically integrating the process $\{V_t^j\}_{t \in [t_k, t_{k+1}]}$. Given the definition of the loss
 228 function in (18), we can expect the term within the first exponential to be close to zero. In the ideal
 229 case where $\widehat{\mathbf{c}}(\mathbf{x}, \mathbf{y}, t) \equiv \mathbf{c}_*(\mathbf{x}, \mathbf{y}, t)$ and $\widehat{V}_0(\mathbf{x}, \mathbf{y}) \equiv -\log h(\mathbf{x}, \mathbf{y}, 0)$, one recovers the exact AF-APF
 230 weights in (5). Once the unnormalized weights (21) are computed, the resampling steps are identical
 231 to those described in Section 2.2 for a standard PF. For practical implementations, all the processes
 232 involved in the proposed methodology can be straightforwardly time-discretized. To distinguish
 233 between CDT learning with static or iterative control, we shall refer to the resulting approximation of
 234 FA-APF as Static-APF and Iterative-APF respectively.

235 4 Experiments

236 This section presents numerical results obtained on three models. All experiments employed 2000
 237 iterations of the Adam optimizer with a learning rate of 0.01 and a mini-batch size of 1000 sample
 238 paths with 10 different observations. Training times took around one to two minutes on a standard
 239 CPU. We note that this compute time is marginal compared to the cost of running filters with many
 240 particles and/or to assimilate large number of observations. Moreover, we can also benefit from
 241 the use of hardware accelerators. We set the inter-observation time as $T = 1$ and employed the
 242 Euler-Maruyama integrator with a stepsize of 0.02 for all examples. Our results are not sensitive to
 243 the choice of T and discretization stepsize if it is sufficiently small. We examined the performance
 244 of each particle filter by computing its effective sample size (ESS) averaged over observation
 245 times and independent repetitions, the evidence lower bound (ELBO) $\mathbb{E}[\log \widehat{p}(\mathbf{y}_1, \dots, \mathbf{y}_K)]$, and
 246 the variance $\text{Var}[\log \widehat{p}(\mathbf{y}_1, \dots, \mathbf{y}_K)]$, where $\widehat{p}(\mathbf{y}_1, \dots, \mathbf{y}_K)$ denotes its unbiased estimator of the
 247 marginal likelihood of the time-discretized filter $p(\mathbf{y}_1, \dots, \mathbf{y}_K)$. When testing particle filters with
 248 varying number of observations K , we increased the number of particles linearly with K to keep
 249 marginal likelihood estimators stable [2].

250 4.1 Ornstein-Uhlenbeck model

251 We considered an Ornstein-Uhlenbeck process given by (1) with $\mu(\mathbf{x}) = -\mathbf{x}$, $\sigma(\mathbf{x}) = 1$ and the Gaus-
 252 sian observation model $g(\mathbf{x}, \mathbf{y}) = \mathcal{N}(\mathbf{y}; \mathbf{x}, \sigma_{\mathbf{Y}}^2)$. We chose $\eta_{\mathbf{X}} = \mathcal{N}(0, 1/2)$ as the stationary distri-
 253 bution and $\eta_{\mathbf{Y}} = \mathcal{N}(0, 1/2 + \sigma_{\mathbf{Y}}^2)$ as the implied distribution of the observation when training neural
 254 networks with the CDT iterative scheme. We took different values of $\sigma_{\mathbf{Y}} \in \{0.125, 0.25, 0.5, 1.0\}$
 255 to vary the informativeness of observations. Analytically tractability in this example allows us to
 256 visualize the quality of our neural network approximations in Figure 1 and consider two idealized
 257 particle filters, namely an APF with exact networks (Exact-APF) and the FA-APF. Comparing our
 258 proposed Iterative-APF to Exact-APF and FA-APF enables us to distinguish between neural network
 259 approximation errors and time-discretization errors. We note that all PFs except the FA-APF involve
 260 time-discretization.

261 Columns 1 to 4 of Figure 2 summarize our numerical findings when filtering simulated observations
 262 from the model. We see that the performance of BPF deteriorates as the observations become more
 263 informative, which is to be expected. Furthermore, when $\sigma_{\mathbf{Y}}$ is small, the impact of our neural
 264 network approximation and time-discretization becomes more noticeable. For the values of $\sigma_{\mathbf{Y}}$ and
 265 the number of observations K considered, we obtained around an order of magnitude gain in efficiency
 266 over BPF. From Column 5, we note that these gains become very substantial when we filter $K = 100$
 267 observations that are simulated with observation noise that are several standard deviations larger than
 268 $\sigma_{\mathbf{Y}} = 0.25$ under the model specification. In particular, while the ELBO of BPF diverges as we
 269 increase the degree of noise in the simulated observations, the ELBO of Iterative-APF remains stable.

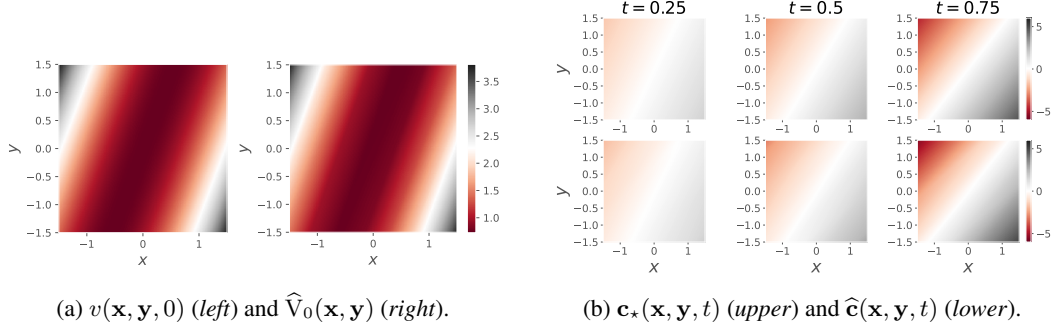


Figure 1: Neural network approximations for Ornstein-Uhlenbeck model with $\sigma_Y = 0.5$.

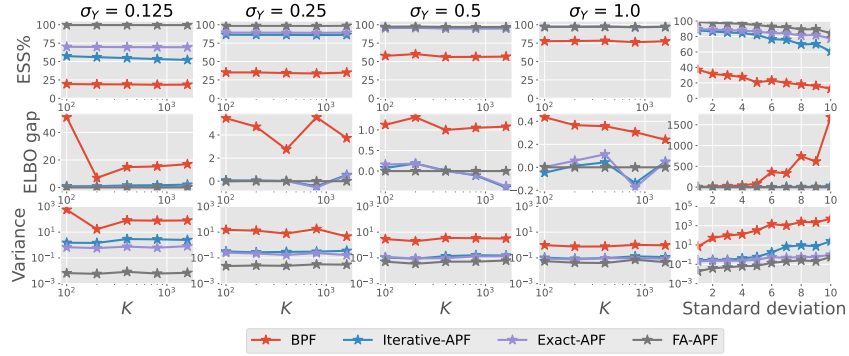


Figure 2: Results for Ornstein-Uhlenbeck model based on 100 independent repetitions of each PF. The ELBO gap in the second row is relative to FA-APF.

270 4.2 Logistic diffusion model

271 Next we consider a logistic diffusion process [11, 26] to model the dynamics of a population size
 272 $\{\mathbf{P}_t\}_{t \geq 0}$, defined by

$$d\mathbf{P}_t = (\theta_3^2/2 + \theta_1 - \theta_2\mathbf{P}_t)\mathbf{P}_t dt + \theta_3\mathbf{P}_t d\mathbf{B}_t. \quad (22)$$

273 We apply the Lamperti transformation $\mathbf{X}_t = \log(\mathbf{P}_t)/\theta_3$ and work with the process $\{\mathbf{X}_t\}_{t \geq 0}$ that
 274 satisfies (1) with $\mu(\mathbf{x}) = \theta_1/\theta_3 - (\theta_2/\theta_3)\exp(\theta_3\mathbf{x})$ and $\sigma(\mathbf{x}) = 1$. Following [26], we adopt
 275 a negative binomial observation model $g(\mathbf{x}, \mathbf{y}) = \mathcal{NB}(\mathbf{y}; \theta_4, \exp(\theta_3\mathbf{x}))$ for counts $\mathbf{y} \in \mathbb{N}_0$ with
 276 dispersion $\theta_4 > 0$ and mean $\exp(\theta_3\mathbf{x})$. We set $(\theta_1, \theta_2, \theta_3, \theta_4)$ as the parameter estimates obtained
 277 in [26]. Noting that (22) admits a Gamma distribution with shape parameter $2(\theta_3^2/2 + \theta_1)/\theta_3^2 - 1$
 278 and rate parameter $2\theta_2/\theta_3^2$ as stationary distribution [11], we select $\eta_{\mathbf{X}}$ as the push-forward under
 279 the Lamperti transformation and $\eta_{\mathbf{Y}}$ as the implied distribution of the observation when training
 280 neural networks under both static and iterative CDT schemes. To induce varying levels of informative
 281 observations, we considered $\theta_4 \in \{1.069, 4.303, 17.631, 78.161\}$.

282 Figure 3 displays our filtering results for various number of simulated observations from the model
 283 (Columns 1 to 4) and for $K = 100$ observations that are simulated with an observation model with
 284 several standard deviations larger than $\theta_4 = 17.631$ under the model specification (Column 5). In the
 285 latter setup, we solved for different values of θ_4 in the negative binomial observation model to induce
 286 larger standard deviations. The behaviour of BPF and Iterative-APF is similar to the previous example
 287 as the observations become more informative with larger values of θ_4 . Iterative-APF outperforms
 288 both BPF and Static-APF over all combinations of θ_4 and K considered, and also when filtering
 289 observations that are increasingly extreme under the model. We note also that the APFs trained using
 290 the CDT static scheme can sometimes give unstable results, particularly in challenging scenarios.

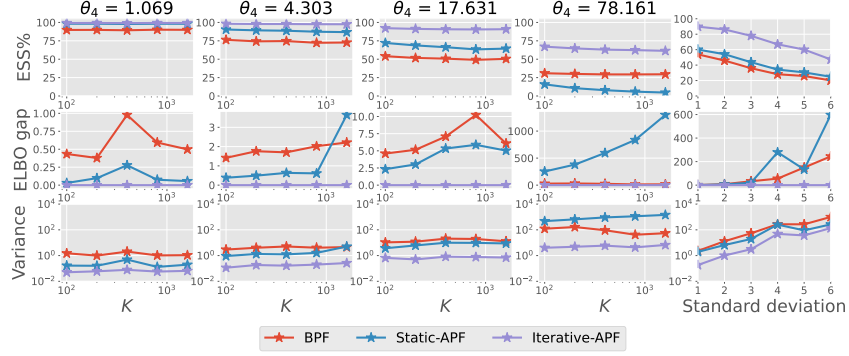


Figure 3: Results for logistic diffusion model based on 100 independent repetitions of each PF. The ELBO gap in the second row is relative to Iterative-APF.

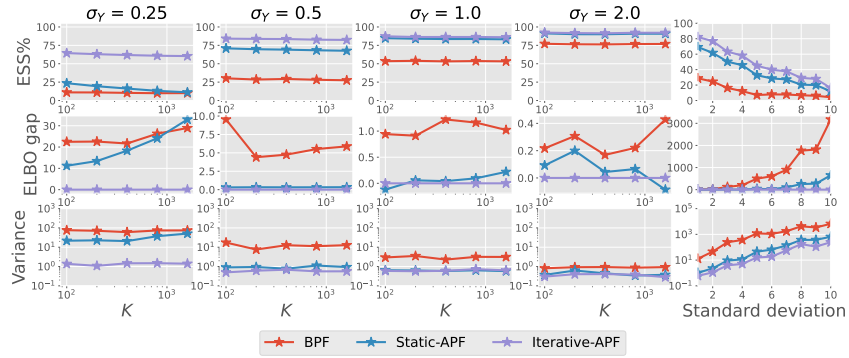


Figure 4: Results for cell model based on 100 independent repetitions of each PF. The ELBO gap in the second row is relative to Iterative-APF.

291 4.3 Cell model

292 Lastly, we examine a cell differentiation and development model from [39]. Cellular expression
 293 levels $\mathbf{X}_t = (\mathbf{X}_{t,1}, \mathbf{X}_{t,2})$ of two genes are modelled by (1) with

$$\mu(\mathbf{x}) = \begin{pmatrix} \mathbf{x}_1^4 / (2^{-4} + \mathbf{x}_1^4) + 2^{-4} / (2^{-4} + \mathbf{x}_2^4) - \mathbf{x}_1 \\ \mathbf{x}_2^4 / (2^{-4} + \mathbf{x}_2^4) + 2^{-4} / (2^{-4} + \mathbf{x}_1^4) - \mathbf{x}_2 \end{pmatrix} \quad (23)$$

294 and $\sigma(\mathbf{x}) = \sqrt{0.1} I_d$. The terms in (23) describe self-activation, mutual inhibition and inactivation
 295 respectively, and the volatility captures intrinsic and external fluctuations. We initialize the diffusion
 296 process from the undifferentiated state of $\mathbf{X}_0 = (1, 1)$ and consider the Gaussian observation model
 297 $g(\mathbf{x}, \mathbf{y}) = \mathcal{N}(\mathbf{y}; \mathbf{x}, \sigma_Y^2 I_2)$. To train neural networks under both static and iterative CDT schemes,
 298 we selected $\eta_{\mathbf{X}}$ and $\eta_{\mathbf{Y}}$ as the empirical distributions obtained by simulating states and observations
 299 from the model for 2000 time units.

300 Figure 4 illustrates our numerical results for various number of observations K and $\sigma_Y \in$
 301 $\{0.25, 0.5, 1.0, 2.0\}$. It shows that Iterative-APF offers significant gains over BPF and Static-APF
 302 when filtering observations that are informative (see Columns 1 to 4) and highly extreme under the
 303 model specification of $\sigma_Y = 0.5$ (see Column 5). In this example, Static-APF did not exhibit any
 304 unstable behaviour and its performance lies somewhere in between BPF and Iterative-APF.

305 References

306 [1] C. Beck, A. Jentzen, et al. Machine learning approximation algorithms for high-dimensional
 307 fully nonlinear partial differential equations and second-order backward stochastic differential
 308 equations. *Journal of Nonlinear Science*, 29(4):1563–1619, 2019.

309 [2] J. Bérard, P. Del Moral, and A. Doucet. A lognormal central limit theorem for particle
 310 approximations of normalizing constants. *Electronic Journal of Probability*, 19:1–28, 2014.

- 311 [3] A. Beskos, O. Papaspiliopoulos, and G. O. Roberts. Retrospective exact simulation of diffusion
312 sample paths with applications. *Bernoulli*, 12(6):1077–1098, 2006.
- 313 [4] A. Beskos, O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead. Exact and computationally ef-
314 ficient likelihood-based estimation for discretely observed diffusion processes (with discussion).
315 *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):333–382,
316 2006.
- 317 [5] J. Blanchet and F. Zhang. Exact simulation for multivariate Itô diffusions. *Advances in Applied*
318 *Probability*, 52(4):1003–1034, 2020.
- 319 [6] Q. Chan-Wai-Nam, J. Mikael, and X. Warin. Machine learning for semi linear PDEs. *Journal*
320 *of Scientific Computing*, 79(3):1667–1712, 2019.
- 321 [7] N. Chopin, O. Papaspiliopoulos, et al. *An introduction to sequential Monte Carlo*. Springer,
322 2020.
- 323 [8] K. L. Chung and J. B. Walsh. *Markov processes, Brownian motion, and time symmetry*, volume
324 249. Springer Science & Business Media, 2006.
- 325 [9] P. Del Moral. *Feynman-Kac formulae: genealogical and interacting particle systems with*
326 *applications*, volume 88. Springer, 2004.
- 327 [10] P. Del Moral and L. M. Murray. Sequential Monte Carlo with highly informative observations.
328 *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):969–997, 2015.
- 329 [11] B. Dennis and R. F. Costantino. Analysis of steady-state populations with the gamma abundance
330 model: application to Tribolium. *Ecology*, 69(4):1200–1213, 1988.
- 331 [12] A. Doucet, A. M. Johansen, et al. A tutorial on particle filtering and smoothing: Fifteen years
332 later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.
- 333 [13] P. Fearnhead, O. Papaspiliopoulos, and G. O. Roberts. Particle filters for partially observed
334 diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):
335 755–777, 2008.
- 336 [14] P. Fearnhead, O. Papaspiliopoulos, G. O. Roberts, and A. Stuart. Random-weight particle
337 filtering of continuous time processes. *Journal of the Royal Statistical Society: Series B*
338 *(Statistical Methodology)*, 72(4):497–512, 2010.
- 339 [15] M. Gerber, N. Chopin, and N. Whiteley. Negative association, ordering and convergence of
340 resampling methods. *The Annals of Statistics*, 47(4):2236–2260, 2019.
- 341 [16] I. V. Girsanov. On transforming a certain class of stochastic processes by absolutely continuous
342 substitution of measures. *Theory of Probability & Its Applications*, 5(3):285–301, 1960.
- 343 [17] J. Han, A. Jentzen, et al. Deep learning-based numerical methods for high-dimensional parabolic
344 partial differential equations and backward stochastic differential equations. *Communications*
345 *in Mathematics and Statistics*, 5(4):349–380, 2017.
- 346 [18] J. Han, A. Jentzen, and E. Weinan. Solving high-dimensional partial differential equations using
347 deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- 348 [19] C. Hartmann, L. Richter, C. Schütte, and W. Zhang. Variational characterization of free energy:
349 Theory and algorithms. *Entropy*, 19(11):626, 2017.
- 350 [20] C. Hartmann, O. Kebiri, L. Neureither, and L. Richter. Variational approach to rare event
351 simulation using least-squares regression. *Chaos: An Interdisciplinary Journal of Nonlinear*
352 *Science*, 29(6):063107, 2019.
- 353 [21] C. Huré, H. Pham, and X. Warin. Deep backward schemes for high-dimensional nonlinear
354 PDEs. *Mathematics of Computation*, 89(324):1547–1579, 2020.
- 355 [22] M. Hutzenthaler and T. Kruse. Multilevel Picard approximations of high-dimensional semilinear
356 parabolic differential equations with gradient-dependent nonlinearities. *SIAM Journal on*
357 *Numerical Analysis*, 58(2):929–961, 2020.
- 358 [23] M. Hutzenthaler, A. Jentzen, T. Kruse, T. Anh Nguyen, and P. von Wurstemberger. Overcoming
359 the curse of dimensionality in the numerical approximation of semilinear parabolic partial
360 differential equations. *Proceedings of the Royal Society A*, 476(2244):20190630, 2020.

- 361 [24] O. Kebiri, L. Neureither, and C. Hartmann. Adaptive importance sampling with forward-
362 backward stochastic differential equations. In *International workshop on Stochastic Dynamics*
363 *out of Equilibrium*, pages 265–281. Springer, 2017.
- 364 [25] P. E. Kloeden and E. Platen. Stochastic differential equations. In *Numerical Solution of*
365 *Stochastic Differential Equations*, pages 103–160. Springer, 1992.
- 366 [26] J. Knappe and P. De Valpine. Fitting complex population models by combining particle filters
367 with Markov chain Monte Carlo. *Ecology*, 93(2):256–263, 2012.
- 368 [27] B. Oksendal. *Stochastic differential equations: an introduction with applications*. Springer
369 Science & Business Media, 2013.
- 370 [28] E. Pardoux and S. Peng. Adapted solution of a backward stochastic differential equation.
371 *Systems & Control Letters*, 14(1):55–61, 1990.
- 372 [29] E. Pardoux and S. Peng. Backward stochastic differential equations and quasilinear parabolic
373 partial differential equations. In *Stochastic partial differential equations and their applications*,
374 pages 200–217. Springer, 1992.
- 375 [30] E. Pardoux and S. Tang. Forward-backward stochastic differential equations and quasilinear
376 parabolic PDEs. *Probability Theory and Related Fields*, 114(2):123–150, 1999.
- 377 [31] J. Park and E. L. Ionides. Inference on high-dimensional implicit dynamic models using a
378 guided intermediate resampling filter. *Statistics and Computing*, 30(5):1497–1522, 2020.
- 379 [32] M. Pereira, Z. Wang, I. Exarchos, and E. A. Theodorou. Learning deep stochastic optimal
380 control policies using forward-backward SDEs. *arXiv preprint arXiv:1902.03986*, 2019.
- 381 [33] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the*
382 *American statistical association*, 94(446):590–599, 1999.
- 383 [34] M. Raissi. Forward-backward stochastic neural networks: Deep learning of high-dimensional
384 partial differential equations. *arXiv preprint arXiv:1804.07010*, 2018.
- 385 [35] G. O. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models
386 using the Metropolis–Hastings algorithm. *Biometrika*, 88(3):603–621, 2001.
- 387 [36] L. C. G. Rogers and D. Williams. *Diffusions, Markov processes and Martingales: Volume 2:*
388 *Itô Calculus*, volume 2. Cambridge university press, 2000.
- 389 [37] H. Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a
390 survey. *International Statistical Review*, 72(3):337–354, 2004.
- 391 [38] S. Thijssen and H. Kappen. Path integral control and state-dependent feedback. *Physical Review*
392 *E*, 91(3):032104, 2015.
- 393 [39] J. Wang, K. Zhang, L. Xu, and E. Wang. Quantifying the Waddington landscape and biological
394 paths for development and differentiation. *Proceedings of the National Academy of Sciences*,
395 108(20):8257–8262, 2011.
- 396 [40] J. Yong and X. Y. Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*, vol-
397 ume 43. Springer Science & Business Media, 1999.

398 Checklist

- 399 1. For all authors...
- 400 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
401 contributions and scope? [Yes] See Introduction Section
- 402 (b) Did you describe the limitations of your work? [Yes] See Introduction Section
- 403 (c) Did you discuss any potential negative societal impacts of your work? [No]
- 404 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
405 them? [Yes]
- 406 2. If you are including theoretical results...
- 407 (a) Did you state the full set of assumptions of all theoretical results? [N/A]
- 408 (b) Did you include complete proofs of all theoretical results? [N/A]
- 409 3. If you ran experiments...

- 410 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
411 perimental results (either in the supplemental material or as a URL)? [Yes] A Py-
412 Torch implementation to reproduce our numerical experiments is available at <https://anonymous.4open.science/r/CompDoobTransform/>
413
- 414 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
415 were chosen)? [Yes] See Experiment Section
- 416 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
417 ments multiple times)? [N/A]
- 418 (d) Did you include the total amount of compute and the type of resources used (e.g., type
419 of GPUs, internal cluster, or cloud provider)? [Yes] CPU and compute time reported
- 420 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 421 (a) If your work uses existing assets, did you cite the creators? [N/A]
- 422 (b) Did you mention the license of the assets? [N/A]
- 423 (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
- 424 (d) Did you discuss whether and how consent was obtained from people whose data you're
425 using/curating? [N/A]
- 426 (e) Did you discuss whether the data you are using/curating contains personally identifiable
427 information or offensive content? [N/A]
- 428 5. If you used crowdsourcing or conducted research with human subjects...
- 429 (a) Did you include the full text of instructions given to participants and screenshots, if
430 applicable? [N/A]
- 431 (b) Did you describe any potential participant risks, with links to Institutional Review
432 Board (IRB) approvals, if applicable? [N/A]
- 433 (c) Did you include the estimated hourly wage paid to participants and the total amount
434 spent on participant compensation? [N/A]