



Guided smoothing and control for diffusion processes

Downloaded from: <https://research.chalmers.se>, 2025-11-24 12:24 UTC

Citation for the original published paper (version of record):

Eklund, O., Lang, A., Schauer, M. (2026). Guided smoothing and control for diffusion processes. Stochastic Processes and their Applications, 192. <http://dx.doi.org/10.1016/j.spa.2025.104806>

N.B. When citing this work, cite the original published paper.



Guided smoothing and control for diffusion processes

Oskar Eklund, Annika Lang^{*,} Moritz Schauer^{*}

Department of Mathematical Sciences, Chalmers University of Technology and the University of Gothenburg, SE-412 96 Gothenburg, Sweden

ARTICLE INFO

Keywords:

Stochastic differential equation
Smoothing
Guided process
Enlargement of filtration
Reaction–diffusion
Hamilton–Jacobi–Bellman equation

ABSTRACT

The smoothing distribution is the conditional distribution of the diffusion process in the space of trajectories given noisy observations made continuously in time. It is generally difficult to sample from this distribution. We use the theory of enlargement of filtrations to show that the conditional process has an additional drift term derived from the backward filtering distribution that is moving or guiding the process towards the observations. This term is intractable, but its effect can be equally introduced by replacing it with a heuristic, where importance weights correct for the discrepancy. From this Markov Chain Monte Carlo and sequential Monte Carlo algorithms are derived to sample from the smoothing distribution. The choice of the guiding heuristic is discussed from an optimal control perspective and evaluated. The results are tested numerically on a stochastic differential equation for reaction–diffusion.

1. Introduction

Diffusion processes are continuous-time probabilistic models with applications across various fields, including finance, physics, and engineering, see, e.g., [1–3]. They naturally emerge by introducing Gaussian random perturbations (white noise) to deterministic systems. We consider the problem of smoothing, that is, determining the conditional distribution of the process' trajectory given observational data by a variant of Bayes' rule. Intuitively, this conditional distribution represents the knowledge about the stochastic trajectory given the observation data. We start with introducing our most central stochastic processes. Let $X = (X_t, t \in [0, T])$ be a d -dimensional diffusion process of interest, which is indirectly observed as detailed below. Its dynamics are governed by the stochastic differential equation (SDE)

$$dX_t = b(t, X_t) dt + \sigma(t, X_t) dW_t, \quad (1.1)$$

on the time interval $t \in [0, T]$ with $T > 0$ some final time, initial value $X_0 = x_0$, and where W is a d' -dimensional Brownian motion. Here, the drift coefficient $b : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and the dispersion coefficient $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d'}$ are sufficiently smooth functions. We denote $a = \sigma \sigma'$.

Let $Y = (Y_t, t \in [0, T])$ be the D -dimensional observation process with initial value $Y_0 = 0$ and dynamics governed by the SDE

$$dY_t = H_t X_t dt + d\beta_t. \quad (1.2)$$

Here, for each $t \in [0, T]$, $H_t \in \mathbb{R}^{D \times d}$ is a linear observation operator and β is a D -dimensional Brownian motion on our probability space, independent of W and X .

* Corresponding author.

E-mail address: annika.lang@chalmers.se (A. Lang).

Assume that Y was observed (up to time T) and the conditional distribution (or smoothing distribution) of X given Y is of interest. In particular, we are interested in sampling paths of X given Y and computing conditional expectations with respect to the smoothing distribution,

$$\mathbb{E}[f(X) | Y],$$

correctly up to Monte Carlo error. For suitable functions f this includes for example quantities of interests such as the mean and variance of the unobserved process X given the data.

The problem is intractable except in a few easy cases. For example, when b is linear in space and σ is constant in space, we can use the Rauch–Tung–Striebel smoother, see [4], to compute it. Naively, smoothing an approximation of the process gives an approximation to the smoothing distribution. But such direct approximations themselves are not suitable for Monte Carlo methods, as in general the conditional law and the approximation are (as measures) mutually singular, precluding the use of the approximation as proposal distribution in the Monte Carlo context when samples of the smoothing distribution are sought.

As solution for this, we propose to use “guided processes” (see [5]) for sampling from the actual smoothing distribution. They are also derived from tractable approximations to the smoothing distribution, for example via the Rauch–Tung–Striebel smoother of a linearized auxiliary system, albeit in a more indirect manner: The auxiliary process is only used to change the drift component of the process X by a linear term, preventing the measures to become mutually singular.

The process with changed drift approximates the conditional process X given observations well and its samples can be used to sample the *actual* conditional process by Monte Carlo, up to Monte Carlo error.

High frequency observations are for example relevant in automated stock trading, but also observations at 50-year intervals such as glacier data can constitute high frequency data as geological processes happen on time scales of million years. Due to their large dimension, observation vectors in these applications are better modelled as functional observation data, and likewise the latent process as continuous time processes, see, e.g., [6].

Method and results

Here, we describe our method of computing expectations with respect to the smoothing distribution. We work on the complete filtered probability space $(\Omega, \mathcal{A}, (\mathcal{G}_t), \mathbb{P})$ supporting the independent Brownian motions W and β and consider the processes X and Y defined on the fixed time interval $[0, T]$.

In Section 2, using the theory of enlargement of filtrations, we introduce a filtration $\mathcal{H} = (\mathcal{H}_t)$ containing information about Y . Under the new filtration, the meaning of “adapted Brownian motion” changes, which is reflected by a change in the drift of the process X . In our first main theorem, we show that X under $\mathbb{P}$ solves the stochastic differential equation

$$dX_t = b(t, X_t)dt + \sigma(t, X_t)(U_t dt + dW_t^*), \quad (\star)$$

for a data dependent, $\mathcal{H}$ -adapted control process $U = (U_t, t \in [0, T])$ taking values in $\mathbb{R}^{d'}$, where W^* is again a d' -dimensional Brownian motion, but now also adapted to $\mathcal{H}$ under $\mathbb{P}$. The theorem includes a characterization of the control process U which induces the dependence on Y , see (2.2).

Note that X in $(\star)$ is still the same process under the same measure $\mathbb{P}$, just expressed differently to make it possible to sample X conditional having seen Y already. For this, note that the new Brownian motion W^* can be simulated in the filtration $\mathcal{H}$ like any other Brownian motion and numerical solutions to $(\star)$ are obtained using standard integration procedures.

The new representation makes it therefore straight forward to sample X given Y if one can get hold on the control process U by solving (2.2). However, this is not always practical, and we may settle for a suboptimal, but tractable data dependent control term U° to substitute U . In Section 3, we define the guided process, which solves the SDE

$$dX_t = b(t, X_t)dt + \sigma(t, X_t)(U_t^\circ dt + dW_t^\circ), \quad (\circ)$$

where W° is again a d' -dimensional Brownian motion adapted to $\mathcal{H} = (\mathcal{H}_t)$ under a different probability measure we denote $\mathbb{P}^\circ$.

We choose U° such that X given Y can be easily sampled under $\mathbb{P}^\circ$. For that, we approximate the SDE (1.1) by a linear, and therefore tractable, process, for which the corresponding control process is known in closed form. This generalizes the backward filtering forward guiding paradigm of [7] to continuous time observation models.

Under $\mathbb{P}^\circ$, X solves the smoothing problem only approximately. But we then derive the Girsanov change of measure which changes the control back to U , expressed as function which gives each sample X of $(\circ)$ a weight $\mathcal{P}(X)$. This is our second main theorem.

In line with the theory of sequential Monte Carlo, the change of measure is tractable up to an unknown constant. This allows us to use Monte Carlo methods to sample the actual conditional process. We devise a Markov Chain Monte Carlo sampler using solutions to $(\circ)$ as proposals which produces dependent samples of $X | Y$ when it has reached stationary.

Ideally, we choose U° such that $\mathbb{P}^\circ$ is close to $\mathbb{P}$. To do so in a principled manner, we parametrize U_θ° and $\mathbb{P}_\theta^\circ$ and vary θ as to minimize the Kullback–Leibler divergence between $\mathbb{P}$ and $\mathbb{P}_\theta^\circ$. This variational problem can be formulated as optimal control problem similar to those considered in [8].

Notation

We use P_X with a subscript to denote the law of a random process X (i.e. the push forward of $\mathbb{P}$ under X). We denote the conditional distribution of X given Y as $P_{X|Y}$. Expectations under a decorated measure, e.g. under $\mathbb{P}^\circ$, are indicated by a decorated expectation $\mathbb{E}^\circ$, likewise the law of X is denoted P_X° , etc. We also denote the integral $\int_0^T Y_t dX_t$ by $Y \cdot X$. The processes $[X, Y]$ and $[X]$ denote the quadratic (co-)variation process with coordinate processes $[X, Y]^{(i,j)} = [X^{(i)}, X^{(j)}]$ and $\mathcal{E}(X) = \exp\left(X - \frac{1}{2}[X]\right)$ the Doléans–Dade exponential process, up to a scaling factor the unique solution Z of $dZ_t = Z_t dX_t$, $Z_0 = 1$. By $\sigma(Y)$ respective $\sigma(C)$, we denote the σ -field generated by some random variable Y or by the sets in a collection C , and we let $\mathcal{F} \vee \mathcal{G} := \sigma(\mathcal{F} \cup \mathcal{G})$.

Related work

The problem of smoothing and filtering for the model (1.1) and (1.2) is well-studied. The exact smoothing distribution as well as the filtering distribution are in general unknown, although it can be found in the Gaussian affine case, see [4,9,10]. In general, indirect observation of diffusion processes causes changes in the underlying probability measure which are described by Girsanov’s theorem. These changes, as done in this work, can be alternatively, but to large extend equivalently, seen as changes in the information flow that can be modelled via enlargement of filtrations, see [11].

Alternatively to both approaches, the Girsanov framework or the framework of enlargement of filtrations, conditional processes can be obtained as solutions to optimal control problems under Kullback–Leibler loss. In this direction, the conditional process can be derived as optimal controlled process, see [8], with

$$U_t = \sigma'(t, X_t) \nabla \log v(t, X_t),$$

being the optimal control defined in terms of the solution v of a backward stochastic partial differential equation derived in [12]. The use of stochastic partial differential equations brings additional complexity which we can avoid in the main line of our work. Nevertheless, we will explore the connection in more detail.

Earlier, similar approaches based on guided processes have been used in the case of discrete-time observation processes, see [7,13]. These approaches are based on work on bridge processes derived in [14,15], which generalize the Brownian bridge. Here, we extend the theory to continuous time observations. Also in this line of work, one can choose between the perspectives of changes of measures and of enlargement of filtrations, see [16, p. 125] and [17].

As reference literature for the smoothing and filtering problems, we use [1,6,18,19]. For filtering techniques using machine learning or deep learning methods, see [20,21]. For related work on optimal control, see e.g. [22].

Outline

This paper is organized as follows: In Section 2, we derive a stochastic differential equation for the smoothed process and relate it to a backward stochastic partial differential equation for which we give a solution in the case of linear processes. In Section 3, we define the guided proposal process, derive the change of measure between the conditional process and the proposal process, and illustrate how to use this in sampling and in variational inference. The last Section 4 is devoted to a numerical experiment. The proofs of theorems are included in the main text, remaining proofs are gathered in the Appendix.

2. Smoothing for continuously observed diffusions

In this section, we derive a stochastic differential equation for the smoothed process using the approach of *initial enlargement of filtration* (grossissement initial de filtration), see [11]. Let $\mathcal{F} = (\mathcal{F}_t, t \in [0, T])$ (recall that $\mathcal{G}$ contains information about both Brownian motions W and β) be the complete and right continuous filtration created by W . The enlarged filtration $\mathcal{H} = (\mathcal{H}_s)$, also complete and right continuous,

$$\mathcal{H}_s = \bigcap_{t>s} (\mathcal{F}_t \vee \sigma(Y)),$$

models the flow of information when the entire realization Y is made known from the beginning of time. Now the process seen as semimartingale under the new filtration will have a new canonical decomposition which we derive in this section.

For the simpler case where $Y = X_T$, the process X under the filtration $\mathcal{H}$ (with $\sigma(Y) = \sigma(X_T)$ accordingly), fulfils Eq. (★) with

$$U_t = \sigma'(t, X_t) \nabla_x \log v(t, X_t; T, X_T), \quad (2.1)$$

where $v(t, x; T, y) = \mathbb{P}(X_T \in dy \mid X_t = x)/dy$ is the transition density, that is the likelihood to still observe $Y = X_T$ when the process is in x at time t .

In the case of continuous observations considered here we can express U similarly in terms of a density of the conditional distribution of Y of an event A ,

$$P_{Y|\mathcal{F}_t}(A) = \mathbb{E}[\mathbf{1}_A(Y) \mid \mathcal{F}_t],$$

which is derived in the next section using a decoupling change of measure.

2.1. Decoupling change of measure for the joint process

We introduce the progressive log-likelihood process $\varphi = (\varphi_t, t \in [0, T])$ taking values in $\mathbb{R}$ with dynamics governed by the SDE

$$d\varphi_t = (H_t X_t)' dY_t - \frac{1}{2} \|H_t X_t\|^2 dt$$

and initial value $\varphi_0 = 0$. Exponentiated it gives the likelihood of seeing a sample $(Y_s, s \in [0, t])$ given $(X_s, s \in [0, t])$ so far relative to the Wiener law, so $\varphi = \log \mathcal{E}((HX)' \cdot Y)$.

Assumption 2.1. Assume that the drift vector b and the dispersion matrix σ are measurably defined on $[0, T] \times \mathbb{R}^d$, with σ bounded and that the observation matrix H is bounded on $[0, T]$ and measurable. Assume that the law of the solution to (1.1) is unique.

Assumption 2.2 (Novikov). There is $\epsilon > 0$ such that $\mathbb{E} \exp(\epsilon \sup_{t \in [0, T]} \|X_t\|^2) < \infty$.

Under this local Novikov-type assumption, compare [23, Exercise 1.40, p. 338], using φ as Radon–Nikodym density, we can define a new, equivalent measure

$$d\tilde{\mathbb{P}} = e^{-\varphi_T} d\mathbb{P},$$

and the process $\mathcal{E}((HX)' \cdot Y)$ is an $\mathcal{G}$ -adapted martingale under $\tilde{\mathbb{P}}$.

Under this new measure, the process X will be decoupled from the observation process Y as this is the change of measure which removes the drift term $H_t X_t dt$ from Y . But the marginal law P_X of X remains the same, for example $\mathbb{E}[f(X)] = \tilde{\mathbb{E}}[f(X)]$ for path functionals f .

Then, under the new measure, by Girsanov's theorem, $(W_t, \tilde{\beta}_t) := (W_t, \beta_t + \int_0^t H_s X_s ds) = (W_t, Y_t)$ is an $\mathbb{R}^{d'+D}$ -dimensional Brownian motion and X is independent of Y . Hence,

$$\tilde{P}_{X,Y} = P_X \otimes \tilde{P}_Y,$$

where $\tilde{P}_Y$ is the Wiener measure restricted to $[0, T]$ and a tilde indicates that we consider the push forward under $\tilde{\mathbb{P}}$. As in [12, p. 150, (1.6)], we have the following characterization of the conditional expectation given Y in form of a Kallianpur–Striebel type formula.

Proposition 2.3. For bounded measurable real test functions f on the path space $C([0, T], \mathbb{R}^d)$, it holds

$$\mathbb{E}[f(X) | Y] = \tilde{\mathbb{E}}[f(X) Z_T | Y], \quad \text{where } Z_t = \frac{e^{\varphi_t}}{\mathbb{E}[e^{\varphi_t} | Y]}.$$

2.2. The dynamics of the conditional process

The process $V = (V_t, t \in [0, T])$, given by

$$V_t = \tilde{\mathbb{E}}[\exp(\varphi_T - \varphi_t) | \mathcal{H}_t],$$

can be understood as likelihood of the future part $(Y_r, t \leq r \leq T)$ of the observation Y given X_t and plays a similar role as the transition density $q(t, X_t, T, y)$ in the case of observing $X_T = y$. The dynamics of the process X under the new filtration are governed by the score process, an $\mathcal{H}$ -adapted process $U = (U_t, t \in [0, T])$ such that

$$\int_0^t U_s V_s ds = [W, V]_t. \quad (2.2)$$

Theorem 2.4. Under Assumptions 2.1 and 2.2, there exists a predictable integrable process U such that (2.2) holds and

$$W_t^* = W_t - \int_0^t U_s ds$$

is an $\mathcal{H}$ -Brownian motion under $\mathbb{P}$. The process X from (1.1) solves $(\star)$ with U and W^* .

Proof. Define $p_t = \Phi_t V_t = \tilde{\mathbb{E}}[\Phi_T | \mathcal{H}_t]$ with $\Phi_t = \exp(\varphi_t)$. The process p is an $\mathcal{H}$ -martingale under $\tilde{\mathbb{P}}$. From the martingale representation theorem, which holds under initial enlargement of $\mathcal{F}$ by the $\tilde{\mathbb{P}}$ -independent $\sigma(Y)$, see [24, theorem III.4.33], there is a square integrable stochastic process Δ with

$$p_t = p_0 + \int_0^t \Delta'_s dW_s.$$

Also $[W, p]$ is absolutely continuous. Therefore, as $p > 0$,

$$U_t = p_t^{-1} \Delta_t = p_t^{-1} \frac{\partial}{\partial t} [W, p]_t$$

solves $p = p_0 \mathcal{E}(U' \cdot W)$ by $p^{-1} \cdot p = (p^{-1} \Delta)' \cdot W = U' \cdot W$ and

$$\int_0^t U_s p_s ds = [W, p]_t. \quad (2.3)$$

Integration by parts yields

$$Wp = W \cdot p + p \cdot W + [W, p],$$

and with $A_t = \int_0^t U_s ds$ by (2.3)

$$Ap = A \cdot p + p \cdot A = A \cdot p + [W, p].$$

Also W is an $\mathcal{H}$ -martingale under $\tilde{\mathbb{P}}$. So,

$$W^* p = (W - A)p = (Wp - [W, p]) - (Ap - [W, p])$$

is an $\mathcal{H}$ -martingale under $\tilde{\mathbb{P}}$ as difference of two $\mathcal{H}$ -martingales. Then by (abstract) Bayes' formula

$$\mathbb{E}[W_t^* | \mathcal{H}_s] = \tilde{\mathbb{E}}[W_t^* p_t | \mathcal{H}_s] / p_s = W_s^*.$$

As $[W^*]_t = tI$, by Lévy's characterization, W^* is an $\mathcal{H}$ -Brownian motion under $\tilde{\mathbb{P}}$. The SDE (★) follows by substituting $dW_t^* + U_t dt$ for dW_t in (1.1).

Finally, by $[W, \Phi] = 0$ we have from integration by parts $[W, p] = [W, \Phi V] = [W, \Phi \cdot V + V \cdot \Phi + [V, \Phi]] = \Phi \cdot [W, V]$. We see that U also solves (2.2) by $\int_0^t U_s V_s ds = \Phi^{-1} \cdot [W, p] = (\Phi^{-1} \Phi) \cdot [W, V] = [W, V]$. $\square$

The term U guides the process closer to the observations Y and plays a similar role as the corresponding term in case of a single discrete observation in (2.1). We derive a corollary that further helps to gain intuitive understanding of U and will help to compute U in the linear Gaussian case later.

Corollary 2.5. *Let V be parametrized by an r -dimensional Itô-process $\Theta = (\Theta_t, t \in [0, T])$, independent of X , in the sense that there is a function $v : [0, T] \times \mathbb{R}^d \times \mathbb{R}^r \rightarrow [0, \infty)$, continuously differentiable in space, such that*

$$v(t, X_t, \Theta_t) = V_t.$$

Then

$$U_t = \sigma'(t, X_t) \nabla_x \log v(t, X_t, \Theta_t).$$

So $(\sigma')^{-1}U$ is the direction where the likelihood V increases, σU acts as a drift in the corresponding direction in coordinates of X , see [25].

2.3. Connection to stochastic PDEs

Without the assumption that V can be parametrized by a finite dimensional process, U is correspondingly determined by the following backward stochastic partial differential equation (PDE), due to [12,26,27], in particular [26, Sec. 5] for the case of unbounded coefficients. Define the formal differential operator

$$\begin{aligned} \mathfrak{L}f(t, x) &= \langle b(t, x), \nabla f(t, x) \rangle + \frac{1}{2} \text{Tr}(a(t, x) \text{Hess}_f(t, x)) \\ &= \sum_{i=1}^d b_i(t, x) \frac{\partial}{\partial x_i} f(t, x) + \frac{1}{2} \sum_{i,j=1}^d a_{ij}(t, x) \frac{\partial^2}{\partial x_i \partial x_j} f(t, x) \end{aligned}$$

associated with X .

One possible way to make practical use of Theorem 2.4 is using these results to compute U .

Proposition 2.6. *Under the additional conditions that b is once and σ and $\sigma\sigma'$ are twice bounded differentiable in space, $\sigma\sigma'$ is uniformly elliptic, and H is bounded differentiable in time,*

$$v(t, x) = \tilde{\mathbb{E}} \left[\frac{\Phi_T}{\Phi_t} \mid X_t = x, Y \right]$$

satisfies the stochastic PDE

$$dv(t, \cdot) + \mathfrak{L}v(t, \cdot) dt + v(t, \cdot) (H_t \cdot)' \diamond dY_t = 0 \quad (2.4)$$

with terminal condition $v(T, \cdot) \equiv 1$, where $\diamond$ indicates a backward Itô integral.

We refer to [27] for further motivation, for the notion of a backward Itô integral and for what entails to be a solution in this context. The stochastic term reads in integral form

$$\int_0^t G(v(s, \cdot), (H_s \cdot)') \diamond dY_s,$$

where for $c : \mathbb{R}^d \rightarrow \mathbb{R}$, $g : \mathbb{R}^d \rightarrow \mathbb{R}^D$, we have $G(c, g)(x) = c(x)g(x)$.

Remark 2.7. When additional to Y also a random variable ζ that depends (only) on X_T is noisily observed, then the terminal condition changes to $v(T, \cdot) = L(x; \zeta)$, where $L(x; y)$ is the conditional probability density of ζ given $X_t = x$ (the likelihood of x). Then $V_T = L(X_T; \zeta)$ and

$$v(t, x) = \mathbb{E} \left[\frac{\Phi_T}{\Phi_t} L(X_T; \zeta) \mid X_t = x, Y, \zeta \right].$$

Thus one can easily extend our results to the setting where both continuous and discrete observations are available, extending [5,7] by this setting.

Remark 2.8. From (2.4) by Itô's formula [12, p. 133],

$$d \log v(t, \cdot) + \left[\mathcal{L} \log v(t, \cdot) + \frac{1}{2} \|\sigma' \nabla \log v(t, \cdot)\|^2 \right] dt + (H_t \cdot)' dY_t - \frac{1}{2} \|(H_t \cdot)\|^2 dt = 0.$$

Details on the derivation can be found in [Appendix](#).

2.4. Backward filter for the likelihood process in the linear case

Consider a linear process (1.1) with $b(t, x) = B_t x + m_t$ and $\sigma(t, x) = \sigma_t$ constant in space, where $\sigma_t \in \mathbb{R}^{d \times d'}$, $B_t \in \mathbb{R}^{d \times d}$, and $m_t \in \mathbb{R}^d$. In line with [28], the solution of the backward stochastic PDE (2.4) can be alternatively obtained by a backward filter. Introduce

$$\mathcal{L}^\dagger f(t, x) = - \sum_{i=1}^d b_i(t, x) \frac{\partial}{\partial x_i} f(t, x) + \frac{1}{2} \sum_{i,j=1}^d a_{ij}(t, x) \frac{\partial^2}{\partial x_i \partial x_j} f(t, x)$$

and let

$$c(t) = -\text{Tr}(B_t).$$

The operator $(-\frac{\partial}{\partial t} + \mathcal{L}^\dagger)f + cf$ is the formal adjoint of $\frac{\partial}{\partial t} + \mathcal{L}$. Now one can think of $-\frac{\partial}{\partial t} + \mathcal{L}^\dagger$, where we split off the bounded multiplicative operator c , as the space-time generator of a linear process $X_t^\dagger$ running backwards in time starting from a Gaussian density $\pi^\dagger(T, \cdot)$.

Using classical filtering theory for the process $X_{T-r}^\dagger$, we can derive the Kushner–Stratonovich equation for the proper filtering density

$$\pi^\dagger(s, x) = \mathbb{P}(X_s^\dagger \in dx \mid Y_t, t \in [s, T]) / dx,$$

where $Y_T - Y_t = \int_t^T H_s X_s^\dagger ds + \beta_T - \beta_t$. Relating $\pi^\dagger(s, x)$ and $v(t, x)$, we obtain a guiding term from the filtering density of the $X^\dagger$ process. More specifically, there is a process $C = (C_t, t \in [0, T])$ such that $v(t, x) = \mathbb{E}[\Phi_T / \Phi_t L(X_T) \mid X_t = x, Y, \zeta]$ in the setting of [Remark 2.7](#) with the conditional distribution of ζ given X_T given by $N(B_\zeta X_T, \Sigma_\zeta)$, a normal distribution with parameters $B_\zeta X_T, \Sigma_\zeta$, and

$$v(s, x) = C_t \pi^\dagger(s, x),$$

with terminal condition

$$v(T, x) = C_T \pi^\dagger(T, x), \tag{2.5}$$

where $v(T, x) \equiv 1$ is obtained as limiting case.

In a statistical context, C can depend on unknown parameters and therefore cannot be neglected. For example, maximum likelihood estimation of B and m can be implemented by gradient descent on $v(0, x_0)$, in which case knowledge of C is needed. We have the following proposition proven in [Appendix](#) and giving directly the backward filtering equation for the likelihood process v making C explicit.

Proposition 2.9. The process v is given by

$$v(t, x) = \frac{C_t}{\sqrt{(2\pi)^d |P_t|}} \exp \left(-\frac{1}{2} (x - v_t)' P_t^{-1} (x - v_t) \right), \tag{2.6}$$

where the parameters v_t, P_t, C_t are obtained from solving a system of backward stochastic differential equations

$$\begin{aligned} dv_t &= (B_t v_t + m_t) dt - P_t H_t' (dY_t - H_t v_t dt), \\ dP_t &= (B_t P_t + P_t B_t' - a_t + P_t H_t' H_t P_t) dt, \\ d \log C_t &= \text{Tr}(B_t) dt - (H_t v_t)' \diamond dY_t + \frac{1}{2} \|H_t v_t\|^2 ds \end{aligned} \tag{2.7}$$

with $a = \sigma \sigma'$ and with terminal condition (2.5), assuming that P_t has a well-defined inverse for $t \in [0, T)$ with bounded limit for $t \rightarrow T$. Consequently, the smoothing process X solves

$$dX_t = (B_t X_t + m_t) dt + a_t P_t^{-1} (v_t - X_t) dt + \sigma_t dW_t^*.$$

Note that result contains Theorem 2.6 in [7] as special case for $H_t \equiv 0$.

3. Guided processes

The process U in Theorem 2.4 is hardly ever tractable outside the linear Gaussian case. Hence, direct forward simulation of the conditional process by solving $(\star)$,

$$dX_t = b(t, X_t) dt + \sigma(t, X_t) U_t dt + \sigma(t, X_t) dW_t^\star \quad (\star)$$

is in general not possible. In this section, we propose an approximation based on proposition 2.9. For a tractable, $\mathcal{H}$ -adapted approximation U° , the process X solves Eq. $(\circ)$,

$$dX_t = b(t, X_t) dt + \sigma(t, X_t) U_t^\circ dt + \sigma(t, X_t) dW_t^\circ, \quad (\circ)$$

with $dW_t^\circ = dW_t^\star + (U_t - U_t^\circ) dt = dW_t - U_t^\circ dt$.

Suppose (i) $\mathbb{E}[\mathcal{E}((U^\circ - U)' \cdot W^\star)_T] = 1$. Then by Girsanov's theorem, $d\mathbb{P}^\circ = \mathcal{E}((U^\circ - U)' \cdot W^\star)_T d\mathbb{P}$ defines a law under which W° is a Brownian motion, so weak solutions of $(\circ)$ can be obtained numerically under $\mathbb{P}^\circ$. If also (ii) $\mathbb{E}^\circ[\mathcal{E}((U - U^\circ)' \cdot W^\circ)_T] = 1$, (weighted) Monte Carlo samples from $X | Y$ can be obtained from those solutions by weighting them with Girsanov's likelihood $Z = \mathcal{E}((U - U^\circ) \cdot W^\circ)_T = \mathcal{E}((U^\circ - U) \cdot W^\star)_T^{-1}$. We come back to this after the next section.

3.1. Guiding based on a linear approximation

It remains to choose U° that approximates U , express Z containing the intractable U only as proportionality constant, and show equivalence between $\mathbb{P}$ and $\mathbb{P}^\circ$ assumed in (i) and (ii).

Let us split $b(t, x) = B_t x + m_t + g(t, x)$ into an (affine) linear part and a non-linear g , so that ideally $\tilde{b}(t, x) = B_t x + m_t$ can be thought of as approximation of b , likewise we approximate the dispersion coefficient $\sigma(t, x)$ by a dispersion coefficient $\tilde{\sigma}_t$ constant in space. Again, let $\tilde{a}_t = \tilde{\sigma}_t \tilde{\sigma}_t'$. Finally, if ζ is observed, we approximate the conditional distribution of ζ by $N(B_\zeta X_T, \Sigma_\zeta)$ for some matrix B_ζ and noise covariance Σ_ζ . If ζ was not observed, we chose $B_\zeta \equiv 0$ and Σ_ζ large, compare with [7].

We leverage that we can solve the smoothing problem for the process $d\tilde{X}_t = \tilde{b}(t, \tilde{X}_t) dt + \tilde{\sigma}_t dW_t$ with the same observation model $d\tilde{Y}_t = H_t \tilde{X}_t dt + d\beta_t$.

Under regularity assumptions, Proposition 2.9 applies to the system $(\tilde{X}, \tilde{Y})$ and the corresponding solution (C, ν, P) defines, with Y substituted for $\tilde{Y}$, an backward filter

$$\tilde{v}(t, x) = \left[\frac{C_t}{\sqrt{(2\pi)^d |P_t|}} \exp\left(-\frac{1}{2}(x - \nu_t)' P_t^{-1} (x - \nu_t)\right) \right]$$

with $\tilde{V}_t = \tilde{v}(t, X_t)$ and

$$U_t^\circ = u^\circ(t, X_t) \quad (3.1)$$

where

$$\begin{aligned} u^\circ(t, x) &:= \sigma'(t, x) \nabla_x \log \tilde{v}(t, x) \\ &= \sigma'(t, x) P_t^{-1} (\nu_t - x). \end{aligned}$$

In fact, it is enough if the conclusion of Proposition 2.9 holds true.

Theorem 3.1. Under Assumptions 2.1 and 2.2 and assuming that ν and P solve (2.7) from Proposition 2.9 for B, m and a with $P_T = \Sigma_\zeta$ and $\nu_T = B_\zeta \zeta$, the measures $\mathbb{P}$ and $\mathbb{P}^\circ$ are equivalent with Radon–Nikodym derivative

$$\frac{d\mathbb{P}}{d\mathbb{P}^\circ} = \frac{\tilde{V}_0 V_T}{V_0 \tilde{V}_T} \exp\left(\int_0^T \psi(t, X_t^\circ, \nu_t) dt\right), \quad (3.2)$$

where

$$\begin{aligned} \psi(t, x, \nu) &= (b(t, x) - \tilde{b}(t, x))' P_t^{-1} (\nu - x) \\ &\quad - \frac{1}{2} \text{Tr}((a(t, x) - \tilde{a}_t) P_t^{-1}) + \frac{1}{2} (\nu - x)' P_t^{-1} (a(t, x) - \tilde{a}_t) P_t^{-1} (\nu - x). \end{aligned}$$

For bounded, measurable, real-valued test functions f defined on the path space $C([0, T], \mathbb{R}^d)$,

$$\mathbb{E}[f(X) | Y] = \frac{\tilde{V}_0 V_T}{V_0 \tilde{V}_T} \mathbb{E}\left[f(X^\circ) \exp\left(\int_0^T \psi(t, X_t^\circ, \nu_t) dt\right) \middle| Y\right].$$

When ζ was observed, we have

$$\mathbb{E}[f(X) | Y] = \frac{\tilde{V}_0 V_T}{V_0 \tilde{V}_T} \mathbb{E}\left[f(X^\circ) \exp\left(\int_0^T \psi(t, X_t^\circ, \nu_t) dt\right) \middle| Y, \zeta\right].$$

Proof. Note that $\mathbb{P}$ and $\mathbb{P}^\circ$ coincide on $\mathcal{H}_0$.

By Yor's formula $\mathcal{E}(X)\mathcal{E}(Y - [X, Y]) = \mathcal{E}(X + Y)$,

$$\begin{aligned}\mathcal{E}((-U + U^\circ)' \cdot W^*)_t &= \mathcal{E}(-U' \cdot W^*)_t \mathcal{E}((U^\circ)' \cdot W^* + [U' \cdot W^*, (U^\circ)' \cdot W^*])_t \\ &= \frac{p_0}{p_t} \mathcal{E}(-(U^\circ)' \cdot W^\circ)_t^{-1},\end{aligned}$$

where we used $p = p_0 \mathcal{E}(U' \cdot W) = p_0 \mathcal{E}(-U' \cdot W^*)^{-1}$ by (2.3) and

$$\begin{aligned}\mathcal{E}((U^\circ)' \cdot W^* + [U' \cdot W^*, (U^\circ)' \cdot W^*]) \\ &= \exp\left((U^\circ)' \cdot W^\circ + \int_0^\cdot (U_s^\circ - U_s)' U_s^\circ ds + [U' \cdot W^*, (U^\circ)' \cdot W^*] - \frac{1}{2}[(U^\circ)' \cdot W^*]\right) \\ &= \exp\left((U^\circ)' \cdot W^\circ + \frac{1}{2}[(U^\circ)' \cdot W^\circ]\right) = \mathcal{E}(-(U^\circ)' \cdot W^\circ)^{-1}.\end{aligned}$$

By Lemma 3.2,

$$\mathcal{E}(-(U^\circ)' \cdot W^\circ)_0 = \frac{\Phi_0}{\Phi_T} \frac{\tilde{v}(0, X_0)}{\tilde{v}(T, X_T)} \exp\left(\int_0^T \psi(t, X_t, v_t) dt\right).$$

Combining the preceding displays gives (3.2). It remains to show equivalence between $\mathbb{P}$ and $\mathbb{P}^\circ$. By equivalence of $\mathbb{P}$ and $\tilde{\mathbb{P}}$, it suffices to show equivalence between $\tilde{\mathbb{P}}$ and $\mathbb{P}^\circ$. Under $\tilde{\mathbb{P}}$, $\hat{v}_t = \tilde{\mathbb{E}} v_t$ is bounded and $v - \hat{v}$ is a centred Gaussian process on the separable Banach space $C([0, T], \mathbb{R}^d)$ with supremum norm $\|\cdot\|_\infty$, so by Fernique's theorem since P_T is finite, see [29, theorem 2.7], there is $\epsilon > 0$ such that

$$\tilde{\mathbb{E}} \exp(\epsilon \|v - \hat{v}\|_\infty / 2) < \infty \quad (3.3)$$

and therefore

$$\tilde{\mathbb{E}} \exp(\epsilon \|v\|_\infty / 2) \leq \tilde{\mathbb{E}} \exp(\epsilon \|v - \hat{v}\|_\infty / 2 + \epsilon \|\hat{v}\|_\infty / 2) < \infty.$$

Then

$$\mathcal{E}((U^\circ)' \cdot W)_t = \mathcal{E}\left(\int_0^t (\sigma'(s, X_s) P_s^{-1} (v_s - X_s))' dW_s\right)$$

and

$$\begin{aligned}\tilde{\mathbb{E}} \exp\left(\frac{1}{2} \int_t^{\min(t+\delta, T)} \|\sigma'(s, X_s) P_s^{-1} (v_s - X_s)\|^2 ds\right) \\ \leq \tilde{\mathbb{E}} \exp\left(\frac{1}{2} \bar{a} \bar{p} \delta (\|v - \hat{v}\|_\infty^2 + \|\hat{v}\|_\infty^2 + \|X\|_\infty^2)\right) \\ \leq \exp\left(\frac{1}{2} \bar{a} \bar{p} \delta \|\hat{v}\|_\infty^2\right) \cdot \tilde{\mathbb{E}} \exp\left(\frac{1}{2} \bar{a} \bar{p} \delta \|v - \hat{v}\|_\infty^2\right) \cdot \tilde{\mathbb{E}} \exp\left(\frac{1}{2} \bar{a} \bar{p} \delta \|X\|_\infty^2\right),\end{aligned}$$

where $\bar{a} = \sup_{t \in [0, T], x \in \mathbb{R}^d} \|\sigma(t, x) \sigma'(t, x)'\|$ and $\bar{p} = \sup_{t \in [0, T]} \|P_t^{-2}\|$ where $\|\cdot\|$ denotes the spectral norm. In the last step, we have used that $\hat{v}$ is deterministic and v and X are independent under $\tilde{\mathbb{P}}$.

Therefore, we may choose $\delta > 0$ small enough, such that by (3.3) and Assumption 2.2, the preceding display is finite and bounded in $t \in [0, T]$. This establishes that for each $t \in [0, T]$ Novikov's condition holds for $(U^\circ)' \cdot W$ for a short time δ independent of t , showing equivalence between $\tilde{\mathbb{P}}$ and $\mathbb{P}^\circ$. $\square$

Note the parallels between (3.2) and formula (1.11) in [7] for the discretely observed case.

In the proof we used the following lemma, which is proven in Appendix directly from Proposition 2.9 using Itô's formula. Heuristically, the result follows from Remark 2.8.

Lemma 3.2. In the setting of Theorem 3.1,

$$\begin{aligned}d \log \tilde{v}(t, X_t) &= -(H_t X_t)' dY_t + \frac{1}{2} \|H_t X_t\|^2 dt \\ &\quad + (\sigma'(t, X_t) P_t^{-1} (v_t - X_t))' dW_t^\circ + \frac{1}{2} \|\sigma(s, X_t) P_t^{-1} (X_t - v_t)\|^2 dt \\ &\quad + (b(t, X_t) - (B_t X_t + m_t))' P_t^{-1} (v_t - X_t) dt \\ &\quad - \frac{1}{2} \text{Tr}((a(t, X_t) - \bar{a}_t) P_t^{-1}) dt + \frac{1}{2} (X_t - v_t)' P_t^{-1} (a(t, X_t) - \bar{a}_t) P_t^{-1} (X_t - v_t) dt.\end{aligned}$$

3.2. Monte Carlo method

The last Theorem 3.1 can be used directly to estimate the expectation of functionals of the conditional process. It would be natural to derive estimates by Monte Carlo sampling, but if the Monte Carlo weights $\frac{dP_{X|Y=y}}{dP^\circ}(X^\circ)$ have too high variance, it can be better to use instead Markov Chain Monte Carlo methods for example. Here we give the Metropolis–Hastings sampler with an autoregressive

random walk on the driving Brownian motion as example. In line with the spirit of the paper, we formulate the Metropolis–Hastings algorithm in the space of sample paths, see [30], and do not discuss approximation errors. In our implementation in Section 4, we use a fine time grid to make sure that the approximation error of the SDEs is small compared to the Monte Carlo sampling error.

In order to apply the Metropolis–Hastings algorithm, we need to define the proposal kernel Q , a measurable function α for the acceptance probabilities and the target distribution π . On Wiener space equipped with the Wiener measure P_W , we define the proposal kernel

$$Q(w, dw')$$

indirectly, by conditional on the current Brownian path w proposing a new Brownian path

$$w' = \rho w + \sqrt{1 - \rho^2} \bar{W},$$

where $\bar{W}$ is an independent Brownian motion and $\rho \in (0, 1)$ is an autocorrelation parameter governing the size of random walk steps in Wiener space. We denote the strong solution of (5), whose existence we assume, by F° such that

$$X = F^\circ(W^\circ).$$

The target distribution $P_{X|Y=y}$ has density $\frac{dP_{X|Y=y}}{dP^\circ}$, known up to a constant, with respect to the “reference measure” P° . Thus, to sample from $P_{X|Y=y}$, we can sample from the measure π defined via its P_W density,

$$\frac{d\pi}{dP_W} = \frac{dP_{X|Y=y}}{dP^\circ} \circ F^\circ,$$

which has $P_{X|Y=y}$ as push forward under F° , such that $F^\circ(Z) \sim P_{X|Y=y}$, where $Z \sim \pi$. To sample from π we are using Metropolis–Hastings with proposal kernel Q . To be precise, given z , we propose $z' = \rho z + \sqrt{1 - \rho^2} \bar{W}$ and accept the proposed z' with probability $\alpha(z, z')$ as new sample z , otherwise we reject the proposal and retain the old sample z . Here,

$$\alpha(z, z') = \min \left(\frac{\frac{d\pi}{dP_W}(z')}{\frac{d\pi}{dP_W}(z)}, 1 \right)$$

is the Metropolis–Hastings acceptance rate.

In Appendix we prove the following proposition that this procedure leads to the correct stationary distribution.

Proposition 3.3. *The measure π is the stationary distribution of the Metropolis–Hastings chain with proposal kernel Q and acceptance probability α .*

3.3. Bellman principle and optimal control for the smoothing distribution

Finding the smoothing distribution can be thought of as a stochastic control problem, see e.g. [22]. Hence, in this section, we revisit heuristically the process defined in (5) from a stochastic control perspective, this also elucidates (2.4). Just as under additional regularity assumptions $V_t = v(t, X_t)$, we can make the ansatz $U_t = u(t, X_t)$ for a choice of (sufficiently smooth) $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$. Here we consider to choose u which minimizes a Kullback–Leibler type cost function: Let $\mathbb{P}^u$ defined by the change of measure $\frac{d\mathbb{P}^u}{d\mathbb{P}} = \mathcal{E}(U' \cdot W) = \mathcal{E}(-U' \cdot W^u)^{-1} = (\frac{d\mathbb{P}}{d\mathbb{P}^u})^{-1}$.

We define the reward function (the negative of a cost function)

$$J_s^t = \int_s^t (H X_r)' dY_r - \frac{1}{2} \int_s^t \|H X_r\|^2 dr - \int_s^t u(r, X_r)' dW_r^u - \int_s^t \frac{1}{2} \|u(r, X_r)\|^2 dr,$$

which is the sum of the log-likelihood process and the negative log-likelihood between $\mathbb{P}^u$ under which W^u is a Brownian motion and $\mathbb{P}$. This objective can be made light of in terms of the variational formula of [31], which relates a conditional distribution to a Kullback–Leibler optimization problem: As example, let $\mathbb{P}^*$ be the conditional or posterior measure induced having observed a datum with likelihood Z , using $\mathbb{P}$ as prior. That is $\mathbb{P}^*$ is the measure given by $d\mathbb{P}^* = Z d\mathbb{P}$ using the likelihood Z as Radon–Nikodym derivative with respect to the prior. By the variational formula, a consequence of Jensen’s inequality, $\mathbb{P}^*$ is the maximizer of $\mathbb{E}^Q \log Z - \text{KL}(Q \parallel \mathbb{P})$ over measures $Q \ll \mathbb{P}$. Here we essentially restrict our search to random measures of the form $Q = \mathbb{P}^u$ with

$$\log \frac{d\mathbb{P}^u}{d\mathbb{P}} = \int_0^T u(t, X_t)' dW_t - \frac{1}{2} \int_0^T \|u(t, X_t)\|^2 dt$$

for some process $u(t, X_t)$, observe $\mathbb{E}^u[\log \frac{d\mathbb{P}^u}{d\mathbb{P}}] = \text{KL}(\mathbb{P}^u \parallel \mathbb{P})$, and search for an optimal control u^* attaining the maximum of the conditional expectation. ($\mathbb{E}^u$ denoting the expectation under $\mathbb{P}^u$.) To solve this, we consider the optimal expected reward-to-go conditional on Y , that is

$$\log v(s, x) := \sup_u \mathbb{E}^u \left[J_s^T \mid X_s = x, Y \right],$$

and use the Bellman principle: “An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision”. (See [32], Chap. III.3.) So

$$\sup_u \mathbb{E}^u \left[J_s^{s+\epsilon} + \log v(s + \epsilon, X_{s+\epsilon}) \mid X_s = x, Y \right] - \log v(s, x) = 0.$$

Taking the formal derivative,

$$\begin{aligned} & (d \log v)(s, \cdot) + (\mathfrak{L} \log v)(s, \cdot) ds + (H_t \cdot)' dY_s - \frac{1}{2} \|H_t \cdot\|^2 ds \\ & + \sup_u \left(u'(s, \cdot) \sigma'(s, x) \frac{\partial}{\partial x} \log v(s, \cdot) - \frac{1}{2} \|u(s, \cdot)\|^2 \right) ds = 0. \end{aligned}$$

The supremum $\frac{1}{2} \|u(s, \cdot)\|^2$ of the right-most term is attained in

$$u(s, x) = \sigma'(s, x) \frac{\partial}{\partial x} \log v(s, x)$$

(see [22]), with $v(s, x)$ solving the parabolic stochastic differential equation

$$d \log v(t, \cdot) + \mathfrak{L} \log v(t, \cdot) dt + (H_t \cdot)' \diamond dY_t + \frac{1}{2} \|\sigma(s, \cdot)' \nabla \log v(t, \cdot)\|^2 dt - \frac{1}{2} \|(H_t \cdot)'\|^2 dt = 0. \quad (3.4)$$

We now have the Hamilton–Jacobi–Bellman stochastic parabolic differential equation for $\log v$ and v . We can relate (3.4) to (2.4) by Remark 2.8. By the verification theorem (c.f. [22]), the reward-to-go $\log v$ characterizes the optimal control u^* : If v is the solution with boundary condition $\log v(T, \cdot) \equiv 0$, then

$$u^*(t, x) = \sigma(t, x)' \nabla \log v(t, x)$$

is the optimal control and the drift of X under $\mathbb{P}^*$ is $b + a \nabla \log v$.

We can therefore think of (★) as solution to a Kullback–Leibler control problem.

3.4. Optimizing the guide

The performance of the method depends on the choice of linear auxiliary process, see the discussion in van der Meulen and Schauer [33] (Section 4.4) for general ways to approach this. Let $\theta = \text{vec}(B, m, \sigma)$ be the tuneable parameters, determining X° using a time-homogeneous linear process as guide. We write P_θ° to emphasize the dependence of the law of X° on the choice of θ . In line with Section 3.3 we propose to simulate the process X° and the Jacobian process DX° , see [34,35],

$$(DX_t^\circ)^{(i, \cdot)} = \partial_{\theta_i} X_t^\circ, t \in [0, T], \quad i \in \{1, \dots, d\},$$

using a differentiable numerical solver and to maximize the reward

$$J = \int_0^T (H_t X_t^\circ)' dY_t - \frac{1}{2} \int_0^T \|H_t X_t^\circ\|^2 dt + \log L(X_T^\circ; \zeta) - \int_s^t u(r, X_r)' dW_r - \int_s^t \frac{1}{2} \|u(r, X_r)\|^2 dr, \quad (3.5)$$

using gradient based optimization methods with gradients $\nabla_\theta J$.

4. Application: Reaction–diffusion

We consider the following discretization of a stochastic reaction diffusion system as a $d = 100$ dimensional stochastic differential equation

$$dX_t = -5\Lambda X_t dt + F(X_t) dt + dW_t, \quad X_0 = 0, \quad (4.1)$$

for $t \in [0, 1]$, where $W = (W_t)$ is an $\mathbb{R}^d$ -valued Brownian motion. The tridiagonal matrix $\Lambda \in \mathbb{R}^{d \times d}$ is given as

$$\Lambda^{(1,1)} = \Lambda^{(d,d)} = 1, \quad \Lambda^{(i,i)} = 2, \quad 1 < i < d, \quad \Lambda^{(i,i+1)} = \Lambda^{(i+1,i)} = -1, \quad 1 \leq i < d$$

and a non-linear force

$$F(x)_i = 2x_i - 2x_i^3, \quad x \in \mathbb{R}^d,$$

that causes spontaneous organization.

This system corresponds to a semidiscrete finite difference approximation in space of the stochastic reaction–diffusion equation

$$d\mathbb{X}(t, x) = \left(5\Delta \mathbb{X}(t, x) + 2\mathbb{X}(t, x) - 2\mathbb{X}^3(t, x) \right) dt + d\mathbb{W}(t, x)$$

with initial condition $\mathbb{X}(0, x) = 0$ and truncated to a finite space interval of $\mathbb{R}$ with grid size 1 and Neumann boundary conditions. Here $\Delta = \partial_{xx}$ denotes the Laplace operator, and $\mathbb{W}$ is a cylindrical Wiener process.

We observe the process

$$Y_t = \int_0^t H X_s ds + \beta_t, \quad Y_t = 0,$$

with $H = 5$ as well as the noisy final value $\zeta = X_1 + Z$, $Z \sim N(0, 0.1I_d)$.

We implemented our procedure for this model in Julia, [36]. All code is available from [37] and <https://github.com/mschauer/GuiSDE.jl>. The Euler–Maruyama scheme with step size 0.001 is employed to generate a sample of X as ground truth and corresponding observation Y . Fig. 1 shows the process X on the left and coordinate processes $X^{(50)}$, $Y^{(50)}$ on the right.




Fig. 1. Forward simulation of system (4.1). Left: Heatmap of $(X_t^{(j)})$. (Horizontal axis: time. Vertical: Coordinates. Colour: Value.) Right: Coordinate process $X^{(50)}$ and, in bold, the corresponding observation process $Y^{(50)}$.




Fig. 2. Left: Estimate $\mathbb{E}[X_t | Y, \zeta]$. Right: A sample of $X | Y, \zeta$.

We approximate X by the linear process $\tilde{X} = (\tilde{X}_t, t \in [0, 1])$ given by

$$d\tilde{X}_t = -5\Lambda\tilde{X}_t dt + dW_t, \quad \tilde{X}_0 = 0.$$

Given our observation Y , we solve the backward system from Proposition 2.9 to obtain U° , see (3.1).

We use the autoregressive random walk Metropolis–Hastings from Section 3.2 with 5000 iterations to compute an estimate of the process $\mu_t^* = \mathbb{E}[X_t | Y, \zeta]$. Fig. 2 shows the obtained estimate as well as a sample of $X | Y, \zeta$.

Finally, we fit a variational approximation to the conditional law $X | Y, \zeta$ using the approach of Section 3.4 on the system. Here we take the same parameters, except the dimension is $d = 20$, to allow the use of (simpler) forward mode automatic differentiation. The results are shown in Fig. 3.

For B we consider all symmetric, tridiagonal matrices, parametrized as

$$B_t^{(i,i)} = \theta_i, \quad 1 \leq i \leq d, \quad B_t^{(i,i+1)} = B_t^{(i+1,i)} = \theta_{d+i}, \quad 1 \leq i < d$$

and m_t we parametrize as $m_t^{(i)} = \theta_{2d+i-1}$. We denote by θ° the ad hoc choice $B(\theta^\circ) = -5\Lambda$ from earlier. We use the Julia package ForwardDiff [38] to compute gradients and the Julia package Optimizers from Julia's Flux ecosystem [39] for the optimization.

We run Adam with learning rate parameter $\eta = 0.01$ and starting value $\theta = 0$ for 2000 iterations to find θ^* maximizing the expected reward (3.5). The training took 187 min on a single Intel Xeon Platinum 8180 CPU at 2.50 GHz.

Fig. 4 shows the training loss curve and the entries of the matrix $B(\theta^*)$, and for comparison, the matrix -5Λ , as heatmap. Our procedure finds a decent approximation to the posterior. To compare, we compare both $\hat{\mu}_t = \mathbb{E}_{\theta^*}^\circ[X_t | Y, \zeta]$ and $\mu_t^\circ = \mathbb{E}_{\theta^\circ}^\circ[X_t | Y, \zeta]$ with μ^* , the posterior mean. Here $\mathbb{P}_\theta^\circ$ is the $\mathbb{P}^\circ$ measure corresponding to the choice θ of the variational parameter. While μ° is already quite close to μ^* , $\hat{\mu}$ is much closer, μ^* see Fig. 5.




Fig. 3. Left: Forward simulation of the system (4.1), $d = 20$. Centre: Posterior mean computed with Metropolis-Hastings for a noisy observation of the left sample. Right: Mean of variational approximation of the same posterior.




Fig. 4. Left: Loss curve (negative reward versus iteration). Centre: $B(\theta^*)$ versus $-5A$. Right: $m(\theta^*)$.




Fig. 5. Left: Error $|\mu_t^{*,(i)} - \mu_t^{\circ,(i)}|$ corresponding to the ad hoc choice. Right: Error $|\mu_t^{*,(i)} - \hat{\mu}_t^{(i)}|$ corresponding to the variational optimum.

Declaration of competing interest

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

Acknowledgements

The authors like to thank Kasper Bågmark for helpful discussions. The work presented in this article was supported by the Chalmers AI Research Centre (CHAIR) through the project “Stochastic Continuous-Depth Neural Network” by the Chalmers foundation, by the European Union (ERC) through grant no. 101088589, “StochMan”, and by the Swedish Research Council (VR) through grant no. 2020–04170. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Appendix. Proofs

Proof of Corollary 2.5. By Clark's formula, [34, IV.41], cited after [40], and the assumed independence,

$$\begin{aligned} [p, W]_t &= [p(\cdot, X_\cdot, \Theta_\cdot), W]_t = \int_0^t (\Phi_t \nabla_x v(s, X_s, \Theta_s))' d[X]_s \\ &= \int_0^t (\sigma'(s, X_s) \nabla_x \log v(s, X_s, \Theta_s))' v(s, X_s, \Theta_s) ds. \quad \square \end{aligned}$$

Derivation of Remark 2.8. Rewrite for clarity the stochastic PDE (2.4)

$$dv(t, \cdot) = -\mathfrak{L}v(t, \cdot) dt - v(t, \cdot)(H_t \cdot)' \diamond dY_t.$$

By Itô's formula for stochastic backward integrals

$$d \log v(t, \cdot) = \frac{1}{v(t, \cdot)} dv(t, \cdot) + \frac{1}{2} \|(H_t \cdot)'\|^2 dt.$$

Plugging in dv yields

$$d \log v(t, \cdot) = -\frac{1}{v(t, \cdot)} \mathfrak{L}v(t, \cdot) dt - (H_t \cdot)' \diamond dY_t + \frac{1}{2} \|H(\cdot)\|^2 dt$$

and

$$d \log v(t, \cdot) + \mathfrak{L} \log v(t, \cdot) dt + \frac{1}{2} \|\sigma' \nabla \log v(t, \cdot)\|^2 dt + (H_t \cdot)' \diamond dY_t - \frac{1}{2} \|(H_t \cdot)'\|^2 dt = 0$$

Lastly, $(H_t \cdot)' \diamond dY_t = (H_t \cdot)' dY_t$. $\square$

Proof of Proposition 2.9. This can be proven via (2.4) by choosing the ansatz (2.6), so that

$$C_t = \int v(t, x) dx, \quad v_t = \frac{1}{C_t} \int x v(t, x) dx, \quad \text{and} \quad P_t = \frac{1}{C_t} \int (x - v_t)(x - v_t)' v(t, x) dx.$$

We only derive the stochastic differential equation for the process C , the equations for v and P are derived in [28].

Note now that

$$\nabla_x v(t, x) = P_t^{-1}(v_t - x)v(t, x)$$

and

$$\frac{1}{2} \sum_{i,j=1}^n a_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} v(t, x) = \frac{1}{2} (\|\sigma'_t P_t^{-1}(v_t - x)\|^2 - \text{Tr}(a'_t P_t^{-1})) v(t, x).$$

Furthermore, we observe that

$$\int \frac{1}{2} (\|\sigma'_t P_t^{-1}(v_t - x)\|^2 - \text{Tr}(a'_t P_t^{-1})) v(t, x) dx = 0$$

and

$$\int (B_t x + m_t)' P_t^{-1}(v_t - x)v(t, x) dx = -\text{Tr}(B_t)C_t.$$

So, plugging everything into (2.4) and using stochastic Fubini to integrate over x , we obtain as equation for C

$$dC_t = \text{Tr}(B_t)C_t dt - C_t(H_t v_t)' \diamond dY_t$$

and therefore

$$d \log C_t = \text{Tr}(B_t) dt - (H_t v_t)' \diamond dY_t + \frac{1}{2} \|H_t v_t\|^2 dt. \quad \square$$

Proof of Lemma 3.2. The result follows from applying Itô's formula and collecting terms.

$$\begin{aligned} d \log \tilde{v}(t, X_t) &= + d \log C_t - \frac{1}{2} d \log |P_t| + d \left(-\frac{1}{2} (X_t - v_t)' P_t^{-1} (X_t - v_t) \right) \\ &= \text{Tr}(B_t) dt - (H_t v_t)' \diamond dY_t + \frac{1}{2} \|H_t v_t\|^2 dt - \frac{1}{2} \text{Tr} \left(\left(\frac{\partial}{\partial t} P_t \right) P_t^{-1} \right) dt \\ &\quad - (dX_t - dv_t)' P_t^{-1} (X_t - v_t) - (dX_t - dv_t)' P_t^{-1} (dX_t - dv_t)/2 \\ &\quad - \frac{1}{2} (X_t - v_t)' \left(\frac{\partial}{\partial t} P_t^{-1} \right) (X_t - v_t) dt \\ &= - (H_t v_t)' dY_t + \frac{1}{2} \|H_t v_t\|^2 dt - (H_t dv_t)' dY_t - \frac{1}{2} \text{Tr}(P_t H_t' H_t) dt + \frac{1}{2} \text{Tr}(\tilde{a}_t P_t^{-1}) dt \\ &\quad + (dX_t - dv_t)' P_t^{-1} (v_t - X_t) - (dX_t - dv_t)' P_t^{-1} (dX_t - dv_t)/2 \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{2} (X_t - v_t)' P_t^{-1} (B_t P_t + P_t B_t') P_t^{-1} (X_t - v_t) dt - \frac{1}{2} (X_t - v_t)' P_t^{-1} \tilde{a}_t P_t^{-1} (X_t - v_t) dt \\
& + \frac{1}{2} ((X_t - v_t)' H_t' H_t (X_t - v_t)) dt \\
& = - (H_t v_t)' dY_t + \frac{1}{2} \|H_t v_t\|^2 dt \\
& + \text{Tr}(H_t P_t H_t') dt - \frac{1}{2} \text{Tr}(P_t H_t' H_t) dt + \frac{1}{2} \text{Tr}(\tilde{a}_t P_t^{-1}) dt \\
& + (b(t, X_t) dt - (X v_t + m) dt)' P_t^{-1} (v_t - X_t) \\
& + (H_t (v_t - X_t))' dY_t - (H_t' H_t v_t)' (v_t - X_t) dt \\
& + (\sigma(t, X_t) dW_t)' P_t^{-1} (v_t - X_t) \\
& - (dX_t)' P_t^{-1} (dX_t)/2 - (dv_t)' P_t^{-1} (dv_t)/2 \\
& - (X (X_t - v_t))' P_t^{-1} (v_t - X_t) dt - \frac{1}{2} (X_t - v_t)' P_t^{-1} \tilde{a}_t P_t^{-1} (X_t - v_t) dt \\
& + \frac{1}{2} X_t' H_t' H_t X_t dt - v_t' H_t' H_t X_t dt + \frac{1}{2} v_t' H_t' H_t v_t dt \\
& = - (H_t X_t)' dY_t + \frac{1}{2} \|H_t X_t\|^2 dt \\
& + (b(t, X_t) dt - (X X_t + m) dt)' P_t^{-1} (v_t - X_t) \\
& - \frac{1}{2} \text{Tr}((a(t, X_t) - \tilde{a}) P_t^{-1}) dt + \frac{1}{2} (X_t - v_t)' P_t^{-1} (-a - \tilde{a}) P_t^{-1} (X_t - v_t) dt \\
& + (\sigma'(t, X_t) P_t^{-1} (v_t - X_t))' dW_t + \frac{1}{2} \|\sigma(s, X_t) P_t^{-1} (X_t - v_t)\|^2 dt \quad \square
\end{aligned}$$

Proof of Proposition 3.3. To establish detailed balance, and therefore prove the proposition, by [30], page 2, we need to show

$$\alpha(z, z') \pi(dz) Q(z, dz') = \alpha(z', z) \pi(dz') Q(z', dz).$$

In our case,

$$\begin{aligned}
\alpha(z, z') \pi(dz) Q(z, dz') &= \alpha(z, z') \frac{d\pi}{dP_W}(z) P_W(dz) Q(z, dz') \\
&= \alpha(z', z) \frac{d\pi}{dP_W}(z') P_W(dz) Q(z, dz') \\
&= \alpha(z', z) \frac{d\pi}{dP_W}(z') P_W(dz') Q(z', dz) \\
&= \alpha(z', z) \pi(dz') Q(z', dz),
\end{aligned}$$

follows from $\min(a/b, 1) \cdot b = \min(b/a, 1) \cdot a$ and the symmetry of the measure

$$\lambda(dw, dw') = P_W(dw) Q(w, dw'). \quad \square$$

References

- [1] B. Øksendal, Stochastic Differential Equations, Springer, Berlin, Heidelberg, 2003, <http://dx.doi.org/10.1007/978-3-642-14394-6>.
- [2] G.N. Milstein, M.V. Tretyakov, Stochastic Numerics for Mathematical Physics, in: Scientific Computation, Springer, Berlin, ISBN: 3-540-21110-1, 2004, <http://dx.doi.org/10.1007/978-3-662-10063-9>.
- [3] S. Särkkä, A. Solin, Applied Stochastic Differential Equations, in: Institute of Mathematical Statistics Textbooks, vol. 10, Cambridge University Press, 2019, <http://dx.doi.org/10.1017/9781108186735>.
- [4] H.E. Rauch, F. Tung, C.T. Striebel, Maximum likelihood estimates of linear dynamic systems, AIAA J. 3 (8) (1965) 1445–1450.
- [5] F. van der Meulen, M. Schauer, Automatic backward filtering forward guiding for Markov processes and graphical models, 2020, <http://dx.doi.org/10.48550/ARXIV.2010.03509>.
- [6] S. Särkkä, Bayesian Filtering and Smoothing, in: Institute of Mathematical Statistics Textbooks, Cambridge University Press, 2013, <http://dx.doi.org/10.1017/CBO9781139344203>.
- [7] M. Mider, M. Schauer, F. van der Meulen, Continuous-discrete smoothing of diffusions, Electron. J. Stat. 15 (2) (2021) <http://dx.doi.org/10.1214/21-ejs1894>.
- [8] J. Kim, P. Mehta, An optimal control derivation of nonlinear smoothing equations, in: Studies in Systems, Decision and Control, in: Studies in Systems, Decision and Control, Springer, Cham, 2020, pp. 295–311, http://dx.doi.org/10.1007/978-3-030-51264-4_12.
- [9] R.E. Kalman, A New Approach to Linear Filtering and Prediction Problems, J. Basic Eng. (ISSN: 0021-9223) 82 (1) (1960) 35–45, <http://dx.doi.org/10.1115/1.3662552>.
- [10] R.E. Kalman, R.S. Bucy, New results in linear filtering and prediction theory, J. Basic Eng. 83 (1961) 95–108.
- [11] T. Jeulin, M. Yor, Grossissements de filtrations: exemples et applications, université de Paris VI. Séminaire de calcul stochastique, in: Lecture notes in mathematics, Springer, Berlin, Heidelberg, ISBN: 9783540152101, 1985, <http://dx.doi.org/10.1007/BFb0075765>.
- [12] E. Pardoux, Stochastic partial differential equations and filtering of diffusion processes, Stochastics 3 (1–4) (1980) 127–167, <http://dx.doi.org/10.1080/17442507908833142>.
- [13] O. Papaspiliopoulos, G. Roberts, Importance sampling techniques for estimation of diffusion models, in: M. Kessler, A. Lindner, M. Sørensen (Eds.), Statistical Methods for Stochastic Differential Equations, Chapman and Hall/CRC, ISBN: 9780429063251, 2012, <http://dx.doi.org/10.1201/b12126>.
- [14] B. Delyon, Y. Hu, Simulation of conditioned diffusion and application to parameter estimation, Stochastic Process. Appl. (ISSN: 0304-4149) 116 (11) (2006) 1660–1675, <http://dx.doi.org/10.1016/j.spa.2006.04.004>.
- [15] M. Schauer, F. van der Meulen, H. van Zanten, Guided proposals for simulating multi-dimensional diffusion bridges, Bernoulli 23 (4A) (2017) 2917–2950, <http://dx.doi.org/10.3150/16-BEJ833>.
- [16] M. Schauer, Bayesian inference for discretely observed diffusion processes (Ph.D. thesis), Technical University Delft, 2015.
- [17] J.-L. Marchand, Conditionnement de Processus Markoviens (Ph.D. thesis), Université Rennes 1, 2012.

- [18] A. Bain, D. Crisan, *Fundamentals of Stochastic Filtering*, Springer, New York, 2009, <http://dx.doi.org/10.1007/978-0-387-76896-0>.
- [19] R. Handel, *Filtering, Stability, and Robustness* (Ph.D. thesis), California Institute of Technology, 2007, <http://dx.doi.org/10.7907/4p53-1h42>.
- [20] J. Han, A. Jentzen, W. E, Solving high-dimensional partial differential equations using deep learning, *Proc. Natl. Acad. Sci.* (ISSN: 1091-6490) 115 (34) (2018) 8505–8510, <http://dx.doi.org/10.1073/pnas.1718942115>.
- [21] K. Bågmark, A. Andersson, S. Larsson, An energy-based deep splitting method for the nonlinear filtering problem, *Partial. Differ. Equations Appl.* (ISSN: 2662-2971) 4 (2) (2023) <http://dx.doi.org/10.1007/s42985-023-00231-5>.
- [22] B. Tzen, M. Raginsky, Theoretical guarantees for sampling and inference in generative models with latent diffusions, in: A. Beygelzimer, D. Hsu (Eds.), *Proceedings of the Thirty-Second Conference on Learning Theory*, in: *Proceedings of Machine Learning Research*, vol. 99, PMLR, 2019, pp. 3084–3114.
- [23] D. Revuz, M. Yor, *Continuous Martingales and Brownian Motion*, third ed., in: *Grundlehren der mathematischen Wissenschaften*, Springer, Berlin, 1998.
- [24] J. Jacod, A.N. Shiryaev, *Limit Theorems for Stochastic Processes*, second ed., in: *Grundlehren der mathematischen Wissenschaften*, Springer, Berlin, 2002.
- [25] M. Corstanje, F. van der Meulen, M. Schauer, S. Sommer, Simulating conditioned diffusions on manifolds, 2024, [arXiv:2403.05409](https://arxiv.org/abs/2403.05409). [math.PR].
- [26] E. Pardoux, Équations du filtrage non linéaire de la prédiction et du lissage, *Stochastics* 6 (3–4) (1982) 193–231, <http://dx.doi.org/10.1080/17442508208833204>.
- [27] A. Heunis, On the stochastic differential equations of filtering theory, *Appl. Math. Comput.* (ISSN: 0096-3003) 37 (3) (1990) 185–218, [http://dx.doi.org/10.1016/0096-3003\(90\)90055-8](http://dx.doi.org/10.1016/0096-3003(90)90055-8).
- [28] D. Fraser, J. Potter, The optimum linear smoother as a combination of two optimum linear filters, *IEEE Trans. Autom. Control* 14 (4) (1969) 387–390, <http://dx.doi.org/10.1109/tac.1969.1099196>.
- [29] G. Da Prato, J. Zabczyk, *Stochastic Equations in Infinite Dimensions*, second ed., in: *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, 2014.
- [30] L. Tierney, A note on Metropolis-Hastings kernels for general state spaces, *Ann. Appl. Probab.* 8 (1) (1998) 1–9, <http://dx.doi.org/10.1214/aoap/1027961031>.
- [31] M.D. Donsker, S.R.S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time. IV, *Comm. Pure Appl. Math.* 36 (2) (1983) 183–212, <http://dx.doi.org/10.1002/cpa.3160360204>.
- [32] R. Bellman, *Dynamic Programming*, in: *Rand Corporation Research Study*, Princeton University Press, 1957.
- [33] F. van der Meulen, M. Schauer, Bayesian estimation of discretely observed multi-dimensional diffusion processes using guided proposals, *Electron. J. Stat.* 11 (2017) <http://dx.doi.org/10.1214/17-EJS1290>.
- [34] L.C.G. Rogers, D. Williams, *Diffusions, Markov Processes, and Martingales. Vol. 2*, in: *Cambridge Mathematical Library*, Cambridge University Press, Cambridge, ISBN: 0-521-77593-0, 2000, Itô calculus, Reprint of the second (1994) edition.
- [35] X. Li, T.-K.L. Wong, R.T.Q. Chen, D.K. Duvenaud, Scalable gradients and variational inference for stochastic differential equations, in: C. Zhang, F. Ruiz, T. Bui, A.B. Dieng, D. Liang (Eds.), *Proceedings of the 2nd Symposium on Advances in Approximate Bayesian Inference*, in: *Proceedings of Machine Learning Research*, vol. 118, PMLR, 2020, pp. 1–28.
- [36] J. Bezanson, A. Edelman, S. Karpinski, V.B. Shah, Julia: A fresh approach to numerical computing, *SIAM REview* 59 (1) (2017) 65–98, <http://dx.doi.org/10.1137/141000671>.
- [37] M. Schauer, A. Lang, Julia package and numerical experiments for “Guided smoothing and control for diffusion processes”, 2025, <http://dx.doi.org/10.5281/zenodo.14984662>.
- [38] J. Revels, M. Lubin, T. Papamarkou, *Forward-mode automatic differentiation in Julia*, 2016, [ArXiv:1607.07892](https://arxiv.org/abs/1607.07892) [Cs.MS].
- [39] M. Innes, et al., *Fashionable modelling with Flux*, 2018, CoRR, [abs/1811.01457](https://arxiv.org/abs/1811.01457). [arXiv:1811.01457](https://arxiv.org/abs/1811.01457).
- [40] J. Dawkins, *Applying Itô’s formula on a C^1 only differentiable function yielding a martingale*, 2018, (Version: 2018-02-05). *Mathematics Stack Exchange*.