
Sampling via Controlled Stochastic Dynamical Systems

Benjamin J. Zhang

Massachusetts Institute of Technology
Cambridge, MA 02139
bjz@mit.edu

Tuhin Sahai

Raytheon Technologies Research Center
Berkeley, CA 94705
tuhin.sahai@rtx.com

Youssef M. Marzouk

Massachusetts Institute of Technology
Cambridge, MA 02139
ymarz@mit.edu

Abstract

We present a general framework for constructing controlled stochastic dynamical systems that exactly sample from a class of probability distributions with Gaussian tails. Given a target distribution and a reference stochastic differential equation (SDE), the Doob h -transform produces a controlled stochastic process whose marginal at a finite time T will be equal to the target distribution. Our method constructs a reference linear SDE and uses the eigenfunctions of its associated Markov operator to approximate the Doob h -transform. The control is approximated by projecting the ratio between the target density and the reference system's time T marginal onto the span of a finite set of eigenfunctions. This projection is performed by minimizing the Kullback-Leibler (KL) divergence from the marginal produced by the approximate control to the true target distribution. In practice, the method lacks robustness due to the high sensitivity to the algorithm's parameters.

1 Introduction

A common problem in statistics and machine learning is that of computing expectations with respect to complex probability distributions. These distributions frequently arise as posterior distributions in Bayesian statistics. Estimating these quantities efficiently via Monte Carlo requires computationally efficient schemes for producing samples that approximate the distribution. We present a framework for constructing a family of *controlled* stochastic dynamical systems that can *exactly* sample from a class of probability distributions with Gaussian tails on \mathbb{R}^d .

Recently the theory of controlled diffusion processes has been gaining attention in statistics and machine learning. For example, it provides a unifying approach to many problem in statistics, including sampling [1], data assimilation [2], optimal transport [3], stochastic optimal control [4, 5], and rare event simulation [6, 7, 8]. Given a reference stochastic differential equation (SDE), an initial distribution, and a target distribution, one aims to find a feedback control such that the marginal at some finite future time T is equal to the target distribution. Finding the optimal control enables *exact sampling* of the target. This problem is also known as the Schrödinger bridge problem [9]. The optimal control is known by different names in different communities, including the *Doob h -transform* [6, 10] and the Föllmer drift [11, 12]. The primary challenge is that to find the optimal control one must choose between two computationally expensive options: solving a series of stochastic optimal control problems or solving a high-dimensional partial differential equation (PDE).

We present a special case of the approach when the reference system is a linear SDE, and show that the optimal control can be expressed in terms of the eigenfunctions of the reference system's Markov generator. Relating these eigenfunctions to the target distribution only requires solving a static optimization problem. The resulting controlled (nonlinear) SDE can produce samples from a broad class of target distributions.

2 Background on controlled diffusion processes

In this section we review some relevant notions from the theories of SDEs and controlled diffusion processes [13, 14, 15, 16]. Let $\{X_t\}_{t \in [0, T]}$ be a time-homogeneous d -dimensional diffusion process on $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\})$. The evolution of the diffusion is described by the SDE

$$dX_t = \mathbf{A}(X_t) dt + \mathbf{B}(X_t) dW_t, \quad X_0 = x_0 \quad (1)$$

where the drift term $\mathbf{A}(x)$ maps \mathbb{R}^d to itself, the diffusion term $\mathbf{B}(x)$ maps \mathbb{R}^d to the space of $d \times r$ matrices, and W_t is a standard r -dimensional Brownian motion. To guarantee existence and uniqueness of the solution, we assume that \mathbf{A} and \mathbf{B} are Lipschitz continuous and have linear growth. A standard tool that describes and that is used to analyze SDEs is the Markov generator defined as $\mathcal{A}\psi = \langle \mathbf{A}(x), \nabla\psi \rangle + \frac{1}{2} \text{Tr} [\mathbf{B}(x)\mathbf{B}(x)^* \nabla^2\psi]$, where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product, ∇ is the gradient, ∇^2 is the Hessian, and Tr is the trace. This is a *linear* operator that acts on the space of twice-continuously differentiable functions and describes the evolution of expectations of the SDE through the Kolmogorov backward equation (KBE). For $f \in \mathcal{C}^2(\mathbb{R}^d)$, define $\Phi(t, x) = \mathbb{E}[f(X_T) | X_t = x]$. The KBE is $\partial_t \Phi + \mathcal{A}\Phi = 0$ with terminal condition $\Phi(T, x) = f(x)$. Now, given an unnormalized *target density* $\pi(x)$, we wish to find the optimal feedback control $u(t, x)$ such that the controlled diffusion process

$$dY_t = [\mathbf{A}(Y_t) + \mathbf{B}u(t, Y_t)] dt + \mathbf{B}(Y_t) dW_t, \quad Y_0 = x_0 \quad (2)$$

has its time T marginal equal to the target distribution. The control that achieves this goal is called the Doob h -transform, which we describe in the following proposition. Let $p(t, t', x, x')$ denote the Markov transition kernel of (1); that is, for a set $A \in \mathcal{F}$, $\mathbb{P}[X_{t'} \in A | X_t = x] = \int_A p(t, t', x, x') dx'$. The transition kernel of the controlled process relates to that of the reference process for a certain class of controls. Define $\eta_{t, x_0}(x)$ to be the probability density of X_t with initial condition x_0 .

Proposition 2.1 (Doob h -transform [10, 12, 17]). *Let $f \in \mathcal{C}^2(\mathbb{R}^d)$ be strictly positive over \mathbb{R}^d and $\Phi(t, x) = \mathbb{E}[f(X_T) | X_t = x]$ be the solution to the KBE. Let $p(t, t', x, x')$ be the Markov transition kernel of the process in (1). If $u(t, x) = \mathbf{B}(x)^* \nabla \log \Phi(t, x)$ is the feedback control in (2), then the transition kernel of the controlled process is $p^u(t, t', x, x') = \Phi(t', x') p(t, t', x, x') \Phi(t, x)^{-1}$. Moreover, observe that by letting $t = 0$, $t' = T$, and $x = x_0$, we have*

$$p^u(0, T, x_0, x') = f(x') \eta_{T, x_0}(x') \Phi(0, x_0)^{-1} := \eta_{T, x_0}^u(x'), \quad (3)$$

which is the marginal density of the controlled process at time T .

Now, assuming that the target distribution is absolutely continuous with respect to the marginal distribution of the reference SDE at time T , notice that if we can choose $f(x) = \pi(x)/\eta(x)$ and solve the KBE, then the corresponding Doob h -transform will lead to a controlled SDE that *samples from the target distribution*. We do not need access to the normalized target density, as the denominator in (3) is itself the normalization constant. In the next section, we show that there exists a family of SDEs whose Doob h -transforms can be found by finding cheap solutions to the KBE.

Previous work on controlled SDEs for importance sampling and rare event simulation of SDEs typically do not take this approach, citing the difficulty of solving the KBE for high-dimensional systems [6]. Instead, they consider a stochastic optimal control perspective. If we introduce the variable transform $U(t, x) = -\log \Phi(t, x)$, we obtain a stochastic Hamilton-Jacobi-Bellman equation [6], which has the following variational formulation

$$U(t, x) = \inf_u \mathbb{E} \left[\frac{1}{2} \int_t^T \|u(t, Y_t)\|^2 dt - \log f(x) \Big| Y_t = x \right]. \quad (4)$$

This approach is common in the rare event simulation literature [6, 7, 8]. A clear proof of Proposition 2.1 using the stochastic optimal control formulation is provided in [12]. The stochastic control perspective has the added benefit of also describing an information-theoretic approach to the problem [7]. In the next section, we show that there exists a family of SDEs whose Doob h -transforms can be found *without* directly solving these computationally challenging problems.

3 Construction of the controlled SDE sampler

We construct a method for sampling a target density π on \mathbb{R}^d , based on finding the Doob h -transform using the KBE. In the following sections, we describe 1) the reference process that admits a closed-form solution to its KBE, 2) how we project the likelihood function described in Proposition 2.1 onto the basis of eigenfunctions, and 3) the role of the terminal marginal.

3.1 Choosing a reference process

To make sampling via controlled SDEs tractable, we must find a way to approximate the Doob h -transform without solving a high-dimensional PDE. Since we only care about sampling from some target distribution, we have freedom to choose the reference dynamical system. Our desiderata for a reference system are as follows: we need to be able to compute the marginal density of the uncontrolled system at time T that contains the support of the target density, and we need a way to solve the KBE without expensive computations. We argue that the reference system should be a *linear* SDE, also known as an *Ornstein-Uhlenbeck* (OU) process. The drift term is $A(x) = -x$ and the diffusion matrix $B(x) = B$ is constant and user-designed depending on the target density.

If the initial condition is deterministic, the density can be derived exactly for all time. In this case, $X_t \sim \mathcal{N}(x_0 e^{-t}, \Sigma_t)$ where $\Sigma_t = \frac{1}{2}(1 - e^{-2t})BB^*$ [15]. Furthermore, the corresponding Markov generator, called the OU operator, has a discrete spectrum and the eigenfunctions of the system can be derived exactly. The eigenfunctions of the OU operator are products of Hermite polynomials. In particular, let $B^*e_i = \mu_i e_i$, where $\|e_i\| = 1$, and let $\mathbf{n} \in \mathbb{N}_0^d$ be multi-indices. The eigenfunctions are $\phi_{\mathbf{n}}(x) = \prod_{i=1}^d \text{He}_{n_i} \left(\frac{\sqrt{2}}{\mu_i} \langle x, e_i \rangle \right)$ with eigenvalues $\lambda_{\mathbf{n}} = -\sum_{i=1}^d n_i$, where $\text{He}_{n_i}(x)$ denotes the Hermite polynomial of degree n_i [18]. This spectral decomposition lets us find an *exact* solution to the KBE as long as $f(x)$ can be expressed in terms of the eigenfunctions. Notice that if $f(x) = \sum_{\mathbf{n}} c_{\mathbf{n}} \phi_{\mathbf{n}}(x)$, then $\Phi(t, x) = \sum_{\mathbf{n}} c_{\mathbf{n}} e^{\lambda_{\mathbf{n}}(T-t)} \phi_{\mathbf{n}}(x)$.

3.2 Projecting onto eigenfunctions

The OU process gives us eigenfunctions that yield efficient solutions to the KBE. We now need to find the expansion coefficients c_i for a given $f(x)$ and set of eigenfunctions $\{\phi_{\mathbf{n}}\}_{\mathbf{n} \in \mathcal{I}}$. We find this “projection” by minimizing the KL divergence from the approximate density to the target.

Define $\tilde{f}(x, c) = \sum_{\mathbf{n} \in \mathcal{I}} c_{\mathbf{n}} \phi_{\mathbf{n}}(x)$, where $\mathcal{I} \subset \mathbb{N}_0^d$ is some set of multi-indices. Let $\pi(x)$ and $\pi_0(x)$ be the unnormalized and normalized target densities, respectively, and let $\tilde{\pi}$ be the approximate density. Let η be the density of the uncontrolled system at time T . Define $f(x) = \pi(x)/\eta(x)$ and let $\tilde{f}(x, c)$ be its approximation. Then we may write the exact and approximate densities as $\pi_0(x) = \gamma^{-1} f(x)\eta(x)$, and $\tilde{\pi}(x) = \tilde{\gamma}^{-1} \tilde{f}(x, c)\eta(x)$, where γ and $\tilde{\gamma}$ are the normalizing constants of $\pi(x)$ and $\tilde{\pi}(x)$. Here, γ is not known, but $\tilde{\gamma}$ can be computed exactly: $\tilde{\gamma} = \mathbb{E}[\tilde{f}(X_T, c) | X_0 = x_0] = \sum_{\mathbf{n} \in \mathcal{I}} c_{\mathbf{n}} e^{\lambda_{\mathbf{n}} T} \phi_{\mathbf{n}}(x_0)$. The KL divergence from $\tilde{\pi}$ to π_0 is $\mathcal{D}_{\text{KL}}(\pi_0(x) \| \tilde{\pi}(x)) = \mathbb{E}_{\pi_0} [\log \pi_0(x) - \log \tilde{\pi}(x)]$. Minimizing this divergence amounts to maximizing $\mathbb{E}_{\pi_0} [\log \tilde{\pi}]$ over the space of admissible probability densities. This objective can be directly approximated, and yields the following optimization problem.

$$\max_{c \in \mathbb{R}^{|\mathcal{I}|}} \mathbb{E}_{\eta} \left[f(x) \log \tilde{\gamma}(c)^{-1} \tilde{f}(x, c) \right]. \quad (5)$$

Solving this optimization problem should guarantee that $\tilde{f}(x, c)$ is positive since if it were otherwise, $\tilde{\pi}$ would no longer be a density function. The objective function is approximated using Monte Carlo samples from η . Doing so enforces positivity at the sample points, since $\log \tilde{f}$ diverges otherwise; it is possible, however, that an approximation is negative elsewhere in the domain. Indeed, another option for performing the projection include using the other direction of KL divergence, $\mathcal{D}_{\text{KL}}(\pi(x) \| \tilde{\pi}_0(x))$, which yields a different optimization problem. In practice we have found that the direction we used in this formulation produced better results for the target distributions we have explored.

3.3 Choosing the terminal marginal η_T and the initial condition

Next we must choose the marginal at time T , η_{T, x_0} , and design the reference OU process accordingly. Since the only closed-form solutions to the OU process are normal distributions, we restrict ourselves

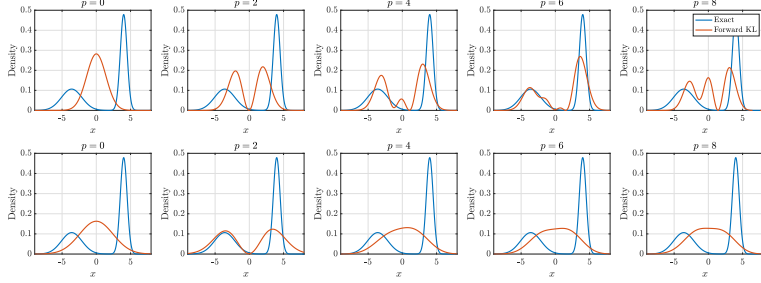


Figure 1: 1-D Gaussian mixture model target. Red curve is the density of the controlled SDE at $T = 1$, and p denotes the maximum polynomial order. Time $T = 1$ marginal of top row has variance 4 while bottom row has variance 36. It is difficult to tune the parameters to match this distribution.

this class. We also assume that the initial condition is deterministic. Choosing otherwise complicates the problem, and is closer to the full Schrödinger bridge [9]. In the optimization problem in (5), we are in effect evaluating the expectation $\mathbb{E}_{\pi_0}[\log \tilde{\pi}]$ using $\eta_{T,x_0} \equiv \eta$ as a biasing distribution. This means that we should choose η such that the objective function can be estimated with low variance. One way to choose η is to apply another method that would give some (crude) normal approximation to π , such as the Laplace approximation, or mean-field variational Bayes. These heuristics are not sufficient when the target distribution is multi-modal. Doing so allows us to approximate the scale of the target and shift it closer to the origin. Therefore, we choose the initial condition to be $x_0 = 0$.

Suppose we have a normal approximation $\mathcal{N}(0, \Sigma)$ to π . To find the OU process that has this marginal at time T , we first find the eigenvalue decomposition of $\Sigma = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^*$. The eigenvectors identify the principal directions for the diffusion term. The diffusion matrix is chosen to be $\mathbf{B} = \sqrt{\frac{2\mathbf{\Lambda}}{1-e^{-2T}}}\mathbf{V}$.

After obtaining the Doob h -transform, the resulting controlled SDE (2) can be simulated independently multiple times, we can use the samples directly for approximate inference or importance sampling. We summarize the algorithm in the Appendix. We also discuss some remarks on the expressiveness of this class of functions for approximating densities.

4 Numerical examples and discussion

We first demonstrate the algorithm on a 1-D Gaussian mixture model. The target distribution is $\pi(x) = 0.6\mathcal{N}(x; 4, 0.5^2) + 0.4\mathcal{N}(x; -3.6, 1.5^2)$. We choose the reference process to be such that the marginal at time $T = 1$ is $\mathcal{N}(x; 0, 4)$ and $\mathcal{N}(x; 0, 36)$ for the top and bottom rows, respectively. The optimization problem in (5) is solved with 10000 samples. In Figure 1, we show how the approximate density approaches the exact density as the basis is enriched. Notice that when the marginal at time T is chosen poorly, it is difficult to adequately capture the target distribution. This is not due to the difficulty of evaluating the objective function or the nature of the optimization problem. Rather, the class of approximating distributions associated with the eigenfunctions and the marginal at time T is not expressive enough to describe the target.

Meanwhile, in Figure 2, when the algorithm parameters are tuned properly, the method can indeed capture a different Gaussian mixture model. Here, the target distribution is $\pi(x) = 0.6\mathcal{N}(x; 1.8, 0.7^2) + 0.4\mathcal{N}(x; -2.6, 0.9^2)$. We choose the reference process to be such that the marginal at time $T = 1$ is $\mathcal{N}(x; 0, 1.4^2)$. The optimization problem in (5) is still solved with 10000 samples. In Figure 1, we show how the approximate density approaches the exact density as the basis is enriched as well as histograms that show samples from the approximate density.

Next, we evaluate the method for a highly non-Gaussian 2-D distribution. The target distribution is a slight modification of an example in [19]; an exact formula for the density is in the appendix. The reference process is chosen such that the distribution at time $T = 1$ is a normal with mean 0 and covariance $\Sigma = \text{diag}(0.6, 1)$. The optimization objective is evaluated with $M = 20000$ samples from η . We use a total order basis of degree up to p for $p \in \{0, 2, 4\}$ as shown in top row of Figure 3, and see that the density is well approximated. In the bottom row of Figure 3, we see that other reasonable choices for the density η lead to poor approximations of the target.

Our formulation of sampling via controlled SDEs is elegant in theory. In practice, however, it only works as long as the parameters of the algorithm are chosen wisely. Approaches to finding parameters

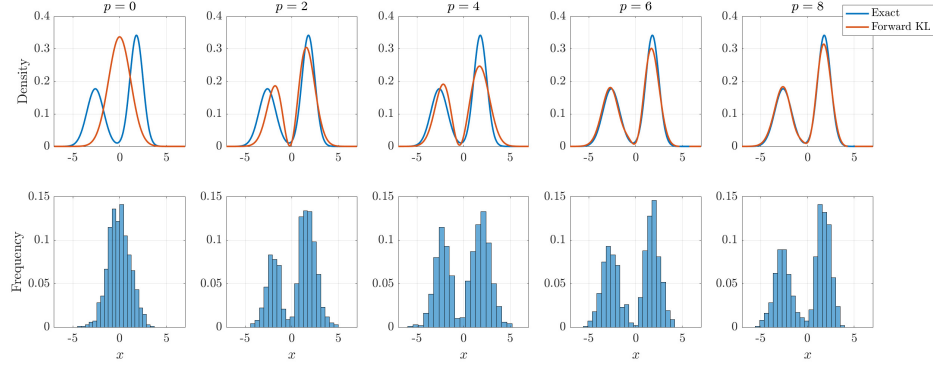


Figure 2: A different 1-D Gaussian mixture model. Red curve is the density of the controlled SDE at $T = 1$, and p denotes the maximum polynomial order. Time $T = 1$ marginal has variance 1.4. In this case, there are parameters such that the approximating class of distributions captures the target.




Figure 3: Top left three figures show the approximate density produced by the controlled SDE for total degree up to p . top rightmost figure shows the exact target density. Red dots show the simulated points of the controlled SDE. Bottom left figure has reference density with $\Sigma = 1.5\mathbf{I}$. Bottom figure has reference density with $\Sigma = 0.3\mathbf{I}$.

such that the method samples from the target distribution well enough for importance sampling or approximate inference are, as yet, elusive. Furthermore, based on the sensitivity of the method to the design parameters, this approach is not robust to poorly chosen parameters. Therefore, it is difficult to fathom that the sampling method will be scalable to high dimensional, non-Gaussian distributions.

The 1-D numerical example shows that certain distributions that are multi-modal are not well described by the approximating class of distributions no matter the choice of the reference density. The reference density with smaller variance does not have enough support where the bulk of the target lies, while the density with larger variance has polynomials that are not expressive enough. The 2-D example shows that even when there is a reference density such that the target can be well approximated, a poor choice of the reference can lead to inaccurate approximations of the target.

Most of the algorithm's issues can be traced to the choice of the terminal marginal. We have found that the successful implementation of the algorithm is highly sensitive to this choice. This makes finding a good terminal marginal through some other means difficult as being slightly off from an "optimal" choice can lead to a useless optimal control in the SDE. Moreover, while the Hermite polynomials are complete in $L^2(\nu)$, in practice, the resulting class of functions is not as expressive as one would like, which leads the nonconvex objective function to be even more difficult to evaluate and optimize.

Other approaches to sampling via controlled SDEs based on the SBP demands a fair amount of computational effort. For example Schrödinger bridge samplers require solutions to many stochastic optimal control problems [1]. Our formulation only requires the solution of a single, static optimization problem, which critically relies on the use of linear systems for the reference process since their eigenfunctions can be computed analytically. Moreover the terminal marginal for linear SDEs can also be computed exactly. If we change to a different class of reference processes, then we can no longer appeal to these virtues.

References

- [1] Espen Bernton, Jeremy Heng, Arnaud Doucet, and Pierre E Jacob. Schrödinger bridge samplers. *arXiv preprint arXiv:1912.13170*, 2019.
- [2] Sebastian Reich. Data assimilation: The Schrödinger perspective. *Acta Numerica*, 28:635–711, 2019.
- [3] Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- [4] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, part i. *IEEE Transactions on Automatic Control*, 61(5):1158–1169, 2015.
- [5] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution—part iii. *IEEE Transactions on Automatic Control*, 63(9):3112–3118, 2018.
- [6] Eric Vanden-Eijnden and Jonathan Weare. Rare event simulation of small noise diffusions. *Communications on Pure and Applied Mathematics*, 65(12):1770–1803, 2012.
- [7] Carsten Hartmann, Lorenz Richter, Christof Schütte, and Wei Zhang. Variational characterization of free energy: Theory and algorithms. *Entropy*, 19(11):626, 2017.
- [8] Wei Zhang, Han Wang, Carsten Hartmann, Marcus Weber, and Christof Schütte. Applications of the cross-entropy method to importance sampling and optimal control of diffusions. *SIAM Journal on Scientific Computing*, 36(6):A2654–A2672, 2014.
- [9] Erwin Schrödinger. *Über die umkehrung der naturgesetze*. Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyter u. Company, 1931.
- [10] Joseph L Doob. *Classical potential theory and its probabilistic counterpart: Advanced problems*, volume 262. Springer Science & Business Media, 2012.
- [11] Hans Föllmer. An entropy approach to the time reversal of diffusion processes. In *Stochastic Differential Systems Filtering and Control*, pages 156–163. Springer, 1985.
- [12] Belinda Tzen and Maxim Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. *arXiv preprint arXiv:1903.01608*, 2019.
- [13] Grigorios A Pavliotis. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer, 2014.
- [14] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [15] Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 2012.
- [16] Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.
- [17] Paolo Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991.
- [18] Giorgio Metafune, Diego Pallara, and Enrico Priola. Spectrum of Ornstein-Uhlenbeck operators in l_p spaces with respect to invariant measures. *Journal of Functional Analysis*, 196(1):40–60, 2002.
- [19] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- [20] Michele Pavon, Esteban G Tabak, and Giulio Trigila. The data-driven Schrödinger bridge. *arXiv preprint arXiv:1806.01364*, 2018.

- [21] Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
- [22] Jie Shen and Li-Lian Wang. Sparse spectral approximations of high-dimensional problems based on hyperbolic cross. *SIAM Journal on Numerical Analysis*, 48(3):1087–1109, 2010.

A Appendix

A.1 Algorithm summary

Algorithm 1: Computing the Doob h -transform

- Input:** Unnormalized target density $\pi(x)$, set of multi-indices $\mathcal{I} \subset \mathbb{N}_0^d$
- Output:** Optimal control $u(t, x)$
- 1: Find an approximation $\eta(x) = \mathcal{N}(0, \Sigma)$ to $\pi(x)$, define $f(x) = \pi(x)/\eta(x)$
 - 2: Compute $\Sigma = \mathbf{V}\Lambda\mathbf{V}^*$
 - 3: Set $\mathbf{B} = \sqrt{\frac{2\Lambda}{1-e^{-2T}}}\mathbf{V}$
 - 4: Construct eigenfunctions $\{\phi_{\mathbf{n}}(x)\}_{\mathbf{n} \in \mathcal{I}}$
 - 5: Draw M independent $X^{(i)} \sim \mathcal{N}(0, \Sigma)$
 - 6: Solve $c^* = \arg \max_{c \in \mathbb{R}^{|\mathcal{I}|}} \frac{1}{M} \sum_{i=1}^M f(X^{(i)}) \log \frac{\tilde{f}(X^{(i)}, c)}{\bar{\gamma}(c)}$ where $\tilde{f}(X^{(i)}, c) = \sum_{\mathbf{n} \in \mathcal{I}} c_{\mathbf{n}} \phi_{\mathbf{n}}(X^{(i)})$, and $\bar{\gamma}(c) = \sum_{\mathbf{n} \in \mathcal{I}} c_{\mathbf{n}} e^{\lambda_{\mathbf{n}}(T-t)} \phi_{\mathbf{n}}(x_0)$
 - 7: Doob h -transform is $u(t, x) = \mathbf{B}^* \nabla \log \sum_{\mathbf{n} \in \mathcal{I}} c_{\mathbf{n}}^* e^{\lambda_{\mathbf{n}}(T-t)} \phi_{\mathbf{n}}(x)$.
-

A.2 Expressiveness of the Hermite polynomials

The OU eigenfunctions form a complete set in the weighted L^2 space with weight function equal to the invariant density η_{∞} of the SDE. For example, in 1-D we need

$$\int_{\mathbb{R}} |f(x)|^2 \eta_{\infty}(x) dx = \int_{\mathbb{R}} \frac{\pi(x)^2}{\eta_T(x)^2} \eta_{\infty}(x) dx < \infty. \quad (6)$$

One difficulty with having a deterministic initial condition is that the distribution of the OU process at finite time T will always be a Gaussian that is *narrower* than the invariant distribution. This means the ratio between the invariant density and the reference marginal η_T grows as $\exp(x^2)$, which implies $\pi(x)$ must have Gaussian tails for the integral to be finite. It is possible to extend this method to include distributions with heavier tails based on the Schrödinger bridge literature [1, 20, 2], though the formulation will be more complicated.

Lastly we address the question of how to choose the set of multi-indices $\mathcal{I} \subset \mathbb{N}_0^d$. In low dimensions (approximately $d \leq 5$), it is practically feasible to use a total order basis $\|\mathbf{n}\|_1 \leq p$ for $p \in \mathbb{N}$. For higher dimensional problems, we will need to incorporate additional structure of the target density into the problem to reduce the number of basis functions. One option is to take advantage of the target density’s Markov structure [21]. If we know that the density factors into a product of potential functions that are only dependent on a subset of the variables, we can construct a reference process whose marginal matches the structure of the target. This allows us to decouple the problem according to the Markov structure, gives us information on how to choose the basis functions, and makes the approach more scalable.

Another generally applicable choice is the sparse truncation that corresponds to choosing $\|\mathbf{n}\|_q \leq p$ for $q \in [0, 1)$. Doing so assumes that $f(x)$ is well approximated by polynomials that depend primarily on certain eigenvector directions, i.e., with only lower-order cross terms. This choice originates from the high-dimensional approximation literature [22].

A.3 Additional numerical details

In general, we have not yet prescribed a definitive way of choosing a “good” terminal reference distribution η_T . We do note, however, that the controlled SDEs seem to perform better when the ratio between the target and reference densities $f(x) = \pi(x)/\eta(x)$ grow to infinity as $|x| \rightarrow \infty$. This is due to the fact that polynomials diverge away from the origin, so they are poor at approximating

functions that decay away from the origin. This implies that η should decay faster than π and is why for the 1-D and 2-D examples, the reference distribution is chosen so that the bulk of the target distribution contains the bulk of the reference.

The density of the two-dimensional example in Section 4 is $\pi(x) = \exp(-U(x))$ where

$$U(x) = 0.5 \left(\frac{\|x\|_2 - 1.5}{0.7} \right)^2 - \log \left[\exp \left(- \left(\frac{x_1 - 2}{0.8\sqrt{2}} \right)^2 \right) + \exp \left(- \left(\frac{x_1 + 1.5}{0.8\sqrt{2}} \right)^2 \right) \right]. \quad (7)$$