
Doob’s Lagrangian: A Sample-Efficient Variational Approach to Transition Path Sampling

Yuanqi Du*¹ Michael Plainer*^{2,3,4,5} Rob Brekelmans*⁶ Chenru Duan^{7,8}
Frank Noé^{4,9,10} Carla P. Gomes¹ Alán Aspuru-Guzik^{6,11} Kirill Neklyudov^{12,13}

¹Cornell University ²Zuse School ELIZA ³Technische Universität Berlin
⁴Freie Universität Berlin ⁵Berlin Institute for the Foundations of Learning and Data
⁶Vector Institute ⁷Massachusetts Institute of Technology ⁸Deep Principle, Inc.
⁹Rice University ¹⁰Microsoft Research AI4Science ¹¹University of Toronto
¹²Université de Montréal ¹³Mila Quebec AI Institute

Abstract

Rare event sampling in dynamical systems is a fundamental problem arising in the natural sciences, which poses significant computational challenges due to an exponentially large space of trajectories. For settings where the dynamical system of interest follows a Brownian motion with known drift, the question of conditioning the process to reach a given endpoint or desired rare event is definitively answered by Doob’s h -transform. However, the naive estimation of this transform is infeasible, as it requires simulating sufficiently many forward trajectories to estimate rare event probabilities. In this work, we propose a variational formulation of Doob’s h -transform as an optimization problem over trajectories between a given initial point and the desired ending point. To solve this optimization, we propose a simulation-free training objective with a model parameterization that imposes the desired boundary conditions by design. Our approach significantly reduces the search space over trajectories and avoids expensive trajectory simulation and inefficient importance sampling estimators which are required in existing methods. We demonstrate the ability of our method to find feasible transition paths on real-world molecular simulation and protein folding tasks.

1 Introduction

Conditioning a stochastic process to obey a particular endpoint distribution, satisfy desired terminal conditions, or observe a rare event is a problem with a long history (Schrödinger, 1932; Doob, 1957) and wide-ranging applications from generative modeling (De Bortoli et al., 2021; Chen et al., 2021a; Liu et al., 2022, 2023c; Somnath et al., 2023) to molecular simulation (Anderson, 2007; Wu et al., 2022; Plainer et al., 2023; Holdijk et al., 2023), drug discovery (Kirmizialtin et al., 2012, 2015; Dickson, 2018), and materials science (Xi et al., 2013; Selli et al., 2016; Sharma et al., 2016).

Transition Path Sampling. In this work, we take a particular interest in the problem of *transition path sampling* (TPS) in computational chemistry (Dellago et al., 2002; Weinan and Vanden-Eijnden, 2010), which attempts to describe how molecules transition between local energy minima or metastable states under random fluctuations or the influence of external forces. Understanding such transitions has numerous applications for combustion, catalysis, battery, material design, and protein folding (Zeng et al., 2020; Klucznik et al., 2024; Blau et al., 2021; Noé et al., 2009; Escobedo et al., 2009).

*Equal contribution. Correspondence: k.nekludov@gmail.com (Kirill Neklyudov)




Figure 1: Given reference dynamics, transition path sampling seeks to capture the conditional or posterior distribution over paths which reach a terminal set $x_T \in \mathcal{B}$. However, simulating the reference dynamics (blue) can be wasteful since we rarely obtain paths (orange) which reach (the vicinity of) the terminal set \mathcal{B} . This is a major challenge for techniques based on importance sampling or Monte Carlo estimation, even when adding a control term to the reference dynamics. By contrast, our approach optimizes a tractable variational distribution over transition paths with a parameterization which satisfies the initial and terminal conditions by design.

While the TPS problem is often framed as finding the ‘most probable path’ transitioning between states (Dürr and Bach, 1978; Vanden-Eijnden and Heymann, 2008), we build upon connections between TPS and Doob’s h -transform (Das et al., 2019, 2021, 2022; Koehl and Orland, 2022; Singh and Limmer, 2023) and seek to match the *full* posterior distribution over conditioned processes.

Doob’s h -Transform. For Brownian motion diffusion processes, conditioning is known to be achieved by Doob’s h -transform (Doob, 1957; Särkkä and Solin, 2019). However, solving this problem amounts to estimating rare event probabilities or matching a complex target distribution. Approaches which involve simulation of trajectories to construct Monte Carlo expectations or importance sampling estimators (Papaspiliopoulos and Roberts, 2012; Schauer et al., 2017; Yan et al., 2022; Holdijk et al., 2023) can be extremely inefficient if the target event is rare or endpoint distribution is difficult to match. Recent methods based on score matching (Heng et al., 2021) or nonlinear Feynman-Kac formula (Chopin et al., 2023) still require simulation during optimization.

Variational Formulation of Doob’s h -Transform. In this work, we propose a variational formulation of Doob’s h -transform as the solution to an optimization on the space of paths of probability distributions. We focus on solving for the Doob transform conditioning on a particular terminal point, which is natural in the TPS setting (see Fig. 1). Taking inspiration from recent bridge matching methods (Peluchetti, 2021, 2023; Liu et al., 2022; Lipman et al., 2022; Shi et al., 2023; Liu et al., 2023a), we propose a parameterization with the following attractive features.

1. **Every Sample Matters.** In contrast to most existing approaches, our training method is *simulation-free*, thereby avoiding computationally wasteful simulation methods to estimate rare-event probabilities and inefficient importance or rejection sampling. We thus refer to our approach as being *sample-efficient*.
2. **Optimization over Sampling.** We propose an expressive variational family of approximations to the conditioned process, which are tractable to sample and can be optimized using neural networks with end-to-end backpropagation.
3. **Problem-Informed Parameterization.** Our parameterization enforces the boundary conditions *by design*, thereby reducing the search space for optimization and efficiently making use of the conditioning information.

We begin by linking the problem of transition path sampling to the Doob’s h -transform and recalling background results in Sec. 2. We present our variational formulation in Sec. 3.1 and detail our optimization algorithm throughout Sec. 3.2. We demonstrate the ability of our approach to achieve comparable performance to Markov Chain Monte Carlo (MCMC) methods with notably improved efficiency on synthetic, and real-world molecular simulation tasks in Sec. 5.

2 Background

2.1 Transition Path Sampling

Consider a forward or reference stochastic process with states x_t and the density of transition probability $\rho_{t+dt}(y|x_t = x) := \rho(x_{t+dt} = y | x_t = x)$. Starting from an initial point $x_0 = A$, the probability density of a discrete-time path is given as

$$\rho(x_T, \dots, x_{dt} | x_0 = A) = \prod_{t=dt}^{T-dt} \rho(x_{t+dt} | x_t) \cdot \rho(x_{dt} | x_0 = A). \quad (1)$$

The problem of rare event sampling aims to condition this reference stochastic process on some event at time T , for example, that the final state belongs to a particular set $x_T \in \mathcal{B}$. We are interested in

sampling from the entire *transition path*, namely the posterior distribution over intermediate states

$$\rho(x_{T-dt}, \dots, x_{dt} | x_0 = A, x_T \in \mathcal{B}) = \frac{\rho(x_T \in \mathcal{B}, x_{T-dt}, \dots, x_{dt} | x_0 = A)}{\rho(x_T \in \mathcal{B} | x_0 = A)}. \quad (2)$$

Moving to continuous time, we focus on the transition path sampling problem in the case where the reference process is given by a Brownian motion. In particular, we are motivated by applications in computational chemistry (Dellago et al., 2002; Weinan and Vanden-Eijnden, 2010), where the reference process is given by molecular dynamics following either overdamped Langevin dynamics,

$$dx_t = -(\gamma M)^{-1} \nabla_x U(x_t) \cdot dt + (\gamma M)^{-1/2} \sqrt{2k_B \mathcal{T}} \cdot dW_t, \quad (3)$$

or the second-order Langevin dynamics with spatial coordinates \bar{x}_t and velocities \bar{v}_t ,

$$d\bar{x}_t = \bar{v}_t \cdot dt, \quad d\bar{v}_t = \left(-M^{-1} \nabla_x U(\bar{x}_t) - \gamma \bar{v}_t \right) \cdot dt + M^{-1/2} \sqrt{2\gamma k_B \mathcal{T}} \cdot dW_t. \quad (4)$$

for a potential energy function U , where W_t denotes the Wiener process. Note that $k_B \mathcal{T}$ is the Boltzman constant times temperature, M is the mass matrix, and γ is the friction coefficient.

2.2 Doob's h -transform

Doob's h -transform addresses the question of conditioning a reference Brownian motion to satisfy a terminal condition such as $x_T \in \mathcal{B}$, thereby providing an avenue to solve the transition path sampling problem described above. Without loss of generality, and to provide a unified treatment of the dynamics in (3)–(4), we consider the forward or reference stochastic differential equation (SDE),

$$\mathbb{P}_{0:T}^{\text{ref}} : \quad dx_t = b_t(x_t) \cdot dt + \Xi_t dW_t, \quad x_0 \sim \rho_0, \quad (5)$$

with drift vector field $b_t : \mathbb{R}^N \rightarrow \mathbb{R}^N$ and diffusion coefficient matrix $\Xi_t \in \mathbb{R}^{N \times N}$ such that $G_t := \frac{1}{2} \Xi_t \Xi_t^T$ is positive definite.² We denote the induced path measure as $\mathbb{P}_{0:T}^{\text{ref}} \in \mathcal{P}(\mathcal{C}([0, T] \rightarrow \mathbb{R}^N))$, i.e. a measure over continuous functions from time to \mathbb{R}^N .

Remarkably, Doob's h -transform (Doob, 1957; Särkkä and Solin, 2019, Sec. 7.5) shows that conditioning the reference process (5) on $x_T \in \mathcal{B}$ results in another Brownian motion process.

Proposition 1. [Jamison (1975, Thm. 2)] *Let $h_{\mathcal{B}}(x, t) := \rho_T(x_T \in \mathcal{B} | x_t = x)$ denote the conditional transition probability of the reference process in (5). Then,*

$$\mathbb{P}_{0:T}^* : \quad dx_{t|T} = \left(b_t(x_{t|T}) + 2G_t \nabla_x \log h_{\mathcal{B}}(x_{t|T}, t) \right) \cdot dt + \Xi_t dW_t \quad x_0 \sim \rho_0 \quad (6)$$

where we use $x_{t|T}$ to denote a conditional process. The SDE in (6) is associated with the following transition probabilities for $s < t < T$,

$$\rho_t(y | x_s = x, x_T \in \mathcal{B}) = \frac{h_{\mathcal{B}}(y, t)}{h_{\mathcal{B}}(x, s)} \rho_t(y | x_s = x), \quad (7)$$

Note that all of our subsequent results hold for the case when \mathcal{B} is a point-mass, with the only change being that the h -function becomes a density, $h_{\mathcal{B}}(x, t) = \rho_T(\mathcal{B} | x_t = x)$.

See App. A.1 for proof, and note that (7) is simply an application of Bayes rule $\rho_t(y | x_s = x, x_T \in \mathcal{B}) = \rho_T(x_T \in \mathcal{B} | x_t = y) \rho_t(y | x_s = x) / \rho_T(x_T \in \mathcal{B} | x_s = x)$ with the unconditioned or reference transition probability as the prior. Furthermore, the conditioned transition probabilities in (7) allow us to directly construct the transition path (2). Using Bayes rule, we have

$$\rho(x_{T-dt}, \dots, x_{dt} | x_0 = A, x_T \in \mathcal{B}) = \frac{h_{\mathcal{B}}(x_{T-dt}, T-dt)}{h_{\mathcal{B}}(A, 0)} \rho(x_{T-dt}, \dots, x_{dt} | x_0 = A)$$

after telescoping cancellation of h -functions and rewriting the denominator in (2) as $h_{\mathcal{B}}(A, 0)$. Thus, we can solve the TPS problem by exactly solving for the h -function and simulating the SDE in (6).

Finally, the h -process and temporal marginals $\rho_t(x | x_0 = A, x_T \in \mathcal{B})$ of the conditioned process satisfy the following forward and backward Kolmogorov equations, which will be useful in deriving our variational objectives in the next section. Note, we use $\langle \nabla_x, \cdot \rangle = \text{div}(\cdot)$ for the divergence operator, and we use $\rho_{t|0,T}$ to indicate the dependence on both $x_0 = A$ (via the initial condition of (8a)) and $x_T \in \mathcal{B}$ (via the h -transform $h_{\mathcal{B}}$). See App. A.1 for the proof.

²Note that the second-order dynamics in (4) can be represented using

$$x_t = \begin{bmatrix} \bar{x}_t \\ \bar{v}_t \end{bmatrix}, \quad b_t(x_t) = \begin{bmatrix} \bar{v}_t \\ -M^{-1} \nabla_x U(\bar{x}_t) - \gamma \bar{v}_t \end{bmatrix}, \quad G_t = \begin{bmatrix} 0 & 0 \\ 0 & M^{-1/2} \sqrt{2\gamma k_B \mathcal{T}} \end{bmatrix}.$$

Proposition 2. *The following PDEs are obeyed by (a) the density of the conditioned process $\rho_{t|0,T}(x) := \rho_t(x | x_0 = A, x_T \in \mathcal{B})$ and (b) the h -function $h_{\mathcal{B}}(x, t)$,*

$$\frac{\partial \rho_{t|0,T}(x)}{\partial t} + \langle \nabla_x, \rho_{t|0,T}(x) (b_t(x) + 2G_t \nabla_x \log h_{\mathcal{B}}(x, t)) \rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \rho_{t|0,T}(x) = 0, \quad (8a)$$

$$\frac{\partial h_{\mathcal{B}}(x, t)}{\partial t} + \langle \nabla_x h_{\mathcal{B}}(x, t), b_t(x) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} h_{\mathcal{B}}(x, t) = 0. \quad (8b)$$

Reparameterizing (8b) in terms of $s_B(x, t) := \log h_{\mathcal{B}}(x, t)$, we can also write

$$\frac{\partial s_B(x, t)}{\partial t} + \langle \nabla_{s_B}(x, t), G_t \nabla s_B(x, t) \rangle + \langle \nabla_{s_B}(x, t), b_t(x) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s_B(x, t) = 0. \quad (8c)$$

3 Method

We first present a novel variational objective whose minimum corresponds to the Doob h -transform in Sec. 3.1, and then propose an efficient parameterization to solve for the h -transform in Sec. 3.2.

3.1 Doob's Lagrangian

Consider reference dynamics given in the form of either (3) or (4), with known drift b_t or energy U . We will restrict our attention to conditioning on a terminal rare event of reaching a given endpoint $x_T = B$, along with an initial point $x_0 = A$. We approach solving for Doob's h -transform via a *least action principle* where, in the following theorem, we define a Lagrangian action whose minimization yields the optimal $q_{t|0,T}^*(x) = \rho_{t|0,T}(x)$ and $v_{t|0,T}^*(x) = \nabla_x \log h_{\mathcal{B}}(x, t)$ from Prop. 1 and 2.

Theorem 1. *The following Lagrangian action functional has a unique solution which matches the Doob h -transform in Prop. 2,*

$$\mathcal{S} = \min_{q,v} \int_0^T dt \int dx q_{t|0,T}(x) \langle v_{t|0,T}(x), G_t v_{t|0,T}(x) \rangle, \quad (9a)$$

$$\text{s.t. } \frac{\partial q_{t|0,T}(x)}{\partial t} = -\langle \nabla_x, q_{t|0,T}(x) (b_t(x) + 2G_t v_{t|0,T}(x)) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x), \quad (9b)$$

$$q_0(x) = \delta(x - A), \quad q_T(x) = \delta(x - B). \quad (9c)$$

The optimal $q_{t|0,T}^(x)$ obeys (8a), and $v_{t|0,T}^*(x) = \nabla_x \log h_{\mathcal{B}}(x, t) = \nabla_x s_B(x, t)$ obeys (8b)-(8c).*

This objective will form the basis for our computational approach, with proof of Thm. 1 deferred to App. A.2. We proceed briefly to contextualize our variational objective and highlight several optimization challenges which will be solved by our proposed parameterization in Sec. 3.2.

Unconstrained Dual Objective. Introducing Lagrange multipliers to enforce the constraints in (9b)–(9c) and eliminating $v_{t|0,T}$, we obtain an alternative, unconstrained version of (9a).

Corollary 1. *The Lagrangian objective in Thm. 1 which solves Doob's h -transform is equivalent to*

$$\mathcal{S} = \min_{q_{t|0,T}} \max_s s_B(B, T) - s_B(A, 0) - \int_0^T dt \int dx q_{t|0,T} \left(\frac{\partial s_B}{\partial t} + \langle \nabla_{s_B}, G_t \nabla s_B \rangle + \langle \nabla_{s_B}, b_t \rangle + \langle \nabla, G_t \nabla s_B \rangle \right)$$

if $q_{t|0,T}$ satisfies (9c). Note $v_{t|0,T}(x) = \nabla_x s_B(x, t)$, with $s_B^(x, t) = \log h_{\mathcal{B}}(x, t)$ at optimality.*³

This objective is similar to the objectives optimized by Action Matching methods (Neklyudov et al., 2023, 2024). Notably, the objective in Cor. 1 is expressed *directly* in terms of the (log) of the h -function for fixed conditioning information $x_T = B$. We also note that the Hamilton Jacobi-style quantity, whose expectation appears in the final term, is zero at optimality in (8c) of Prop. 2.

³Compared to (8c), we write $\sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s_B(x, t) = \langle \nabla, G_t \nabla s_B(x, t) \rangle$ for simplicity of notation.

Path Measure Perspective. We next relate our variational objective in [Thm. 1](#) to a KL divergence optimization over path measures. Let $\mathbb{P}_{0:T}^{\text{ref}}$ denote the law of the reference SDE in [\(5\)](#) with fixed $\mathbb{P}_0^{\text{ref}} = \delta(x_0 - A)$. Let $\mathbb{Q}_{0:T}^v$ denote the law of a controlled process similar to [\(6\)](#), but with a variational $v_{t|0,T}$ in place of $\nabla_x \log h_B$,

$$\mathbb{Q}_{0:T}^v : \quad dx_t = (b_t(x_{t|T}) + 2G_t v_{t|0,T}(x_{t|T})) \cdot dt + \Xi_t dW_t, \quad x_0 = A. \quad (10)$$

Note that the density $q_{t|0,T}$ of $\mathbb{Q}_{0:T}^v$ evolve according to the Fokker-Planck equation in [\(9b\)](#) ([Särkkä and Solin, 2019, Sec. 5.2](#)). Using the Girsanov Theorem, the objective in [\(9a\)](#) can then be viewed as a KL divergence minimization over path measures $\mathbb{Q}_{0:T}^v$ which satisfy the boundary constraints.

Corollary 2. *The following Schrödinger Bridge (SB) problem*

$$\mathcal{S} := \min_{\mathbb{Q}_{0:T}^v \text{ s.t. } \mathbb{Q}_0^v = \delta_A, \mathbb{Q}_T^v = \delta_B} D_{KL}[\mathbb{Q}_{0:T}^v : \mathbb{P}_{0:T}^{\text{ref}}] \quad (11)$$

yields the path measure $\mathbb{P}_{0:T}^$ associated with the SDE in [\(6\)](#) as its unique minimizing argument. The temporal marginals of $\mathbb{P}_{0:T}^*$ are equal to those which optimize the Lagrangian objective in [Thm. 1](#).*

Our Lagrangian action minimization thus corresponds to the solution of an SB problem ([Schrödinger, 1932; Léonard, 2014](#)) with Dirac delta functions as the endpoint measures. Our objective in [\(9a\)](#) particularly resembles optimal control formulations of SB ([Chen et al., 2016, 2021b, Prob. 4.4, 5.3](#)). While it is well-known that the Doob h -transform (and large deviation theory more generally) plays a role in the solution to SB problems ([Jamison, 1975; Léonard, 2014](#)), our interest in the transition path sampling problem leads to specific computational decisions below. See [Sec. 4](#) for further discussion.

Challenges of Optimizing [\(9a\)](#). We highlight several distinctive features of our problem which inform the development of new computational methods in [Sec. 3.2](#).

1. First, we perform optimization over the *first* argument of the KL divergence in [\(11\)](#), indicating that we need to be able to efficiently sample from the conditioned process in [\(10\)](#) or $q_{t|0,T}$ in [\(9\)](#). This appears challenging due to the nonlinearity of both the reference and variational drifts, b_t and $v_{t|0,T}$.
2. For a given $q_{t|0,T}$, it can be difficult to solve for $v_{t|0,T}$ which satisfies the Fokker-Planck equation in [\(9b\)](#) or ∇s which solves the inner optimization in [Cor. 1](#).
3. Finally, we would like to strictly enforce the boundary constraints on $q_{t|0,T}$ or $\mathbb{Q}_{0:T}^v$ to avoid simulating or wasting computation on trajectories for which $x_T \neq B$.

In fact, our parameterization of $q_{t|0,T}$ in [Sec. 3.2](#) will *completely avoid* simulation of the SDE in [\(10\)](#) during training ([Challenge 1](#)), provide *analytic* solutions for $v_{t|0,T}$ satisfying [\(9b\)](#) with a given $q_{t|0,T}$ ([Challenge 2](#)), and *exactly* enforce the boundary constraints ([Challenge 3](#)).

3.2 Computational Approach

We now propose a family of Gaussian (mixture) path parameterizations $q_{t|0,T}$ which overcome the computational challenges posed in the previous section, while still maintaining expressivity. We present all aspects of our proposed method in the context of the first-order dynamics [\(3\)](#) in [Sec. 3.2.1](#), before presenting extensions to mixture paths and the second-order setting [\(4\)](#) in [Sec. 3.2.2–3.2.3](#).

3.2.1 First-Order Dynamics and General Approach

Tractable Drift $v_{t|0,T}$ for Variational Doob Objective. We begin by considering a modification of the Fokker-Planck constraint in [\(9b\)](#), with all drift terms absorbed into a single vector field $u_{t|0,T}$,

$$\frac{\partial q_{t|0,T}(x)}{\partial t} = -\left\langle \nabla_x, q_{t|0,T}(x) u_{t|0,T}(x) \right\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x). \quad (12)$$

For arbitrary $q_{t|0,T}$, solving for *any* $u_{t|0,T}(x)$ satisfying [\(12\)](#) can be a difficult optimization problem, whose solution is not unique without some cost-minimizing assumption ([Neklyudov et al., 2023](#)).

To sidestep this optimization, and address [Challenge 2](#), we restrict attention to variational families of $q_{t|0,T} \in \mathcal{Q}$ where it is *analytically tractable* to calculate a vector field $u_{t|0,T}^{(q,\theta)}$ which satisfies [\(12\)](#). We first consider the family of Gaussian paths \mathcal{Q}_G , in similar fashion to (conditional) flow matching methods ([Lipman et al., 2022; Tong et al., 2023; Liu et al., 2023a](#)), with proof in [App. B](#).

Algorithm 1: Training (Single Gaussian)

Input: Reference drift b_t , diffusion matrix G_t
Conditioning endpoints
while not converged **do**
 Sample $t \sim \mathcal{U}(0, T)$
 Sample $x_t \sim q_{t|0,T}^{(\theta)}$ using (15)
 Calculate $u_{t|0,T}^{(q,\theta)}(x_t)$ using (13)
 Calculate $v_{t|0,T}^{(q,\theta)}(x_t)$ using $u_{t|0,T}^{(q,\theta)}(x_t), b_t(x_t)$, (14)
 Calculate $\mathcal{L} = \langle v_{t|0,T}^{(q,\theta)}(x_t), G_t v_{t|0,T}^{(q,\theta)}(x_t) \rangle$ (Thm. 1)
 Update $\theta \leftarrow \text{optimizer}(\theta, \nabla_{\theta} \mathcal{L})$
end while
return θ

Algorithm 2: Sampling Trajectories

def get_drift(x_t, t):
 Evaluate $\mu_{t|0,T}^{(\theta)}, \Sigma_{t|0,T}^{(\theta)}$ at t
 return drift $u_{t|0,T}^{(q,\theta)}(x_t)$ using (13)

 Sample initial state $x_0 \sim \mathcal{N}(A, \sigma_{\min}^2)$
 return SDESolve($x_0, \text{get_drift}, T$)

Algorithms for training with a single Gaussian path (Alg. 1) and sampling or generating transition paths at test time (Alg. 2). Note that we sample from the marginals $q_{t|0,T}$ during training, but generate paths by simulating the SDE (10).

Proposition 3. For the family of endpoint-conditioned marginals $q_{t|0,T}(x) = \mathcal{N}(x | \mu_{t|0,T}, \Sigma_{t|0,T})$,

$$u_{t|0,T}^{(q,\theta)}(x) := \frac{\partial \mu_{t|0,T}}{\partial t} + \left[\frac{1}{2} \frac{\partial \Sigma_{t|0,T}}{\partial t} \Sigma_{t|0,T}^{-1} - G_t \Sigma_{t|0,T}^{-1} \right] (x - \mu_{t|0,T}) \quad (13)$$

satisfies the Fokker-Planck equation (12) for $q_{t|0,T}$ and diffusion coefficients $G_t = \frac{1}{2} \Xi_t \Xi_t^T$.

Given $u_{t|0,T}^{(q,\theta)}$ corresponding to $q_{t|0,T}$, we can simply solve for the $v_{t|0,T}$ satisfying the Fokker-Planck equation in (9b) in our variational Doob objective (Thm. 1). Since G_t was assumed to be invertible and the base drift b_t is known, we have

$$v_{t|0,T}^{(q,\theta)}(x) = \frac{1}{2} (G_t)^{-1} \left(u_{t|0,T}^{(q,\theta)}(x) - b_t(x) \right). \quad (14)$$

We may now evaluate terms involving $v_{t|0,T}$ in our Lagrangian objective in (9) using (14) directly, without spending effort to solve an inner minimization over $v_{t|0,T}$ (thus addressing Challenge 2).

Optimization over $q_{t|0,T}$ satisfying Boundary Constraints. Given the ability to evaluate $v_{t|0,T}^{(q,\theta)}$ for a given $q_{t|0,T} \in \mathcal{Q}_G$ as above, our variational Doob objective in (9a) reduces to a single optimization over the marginals $q_{t|0,T}$ of a conditioned process which satisfies the boundary conditions (9c).

We consider parameterizing the mean $\mu_{t|0,T}$ and covariance $\Sigma_{t|0,T}$ of our Gaussian path $q_{t|0,T}$ using a neural network. For simplicity, we consider a diagonal parameterization $\Sigma_{t|0,T} = \text{diag}(\{\sigma_{t|0,T,d}^2\}_{d=1}^D)$. We parameterize a neural network $\text{NNET}_{\theta} : [0, T] \times \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}^D \times \mathbb{R}^D$ which inputs time t and boundary conditions $x_0 = A, x_T = B$, and outputs vectors of mean perturbations and per-dimension variances. Finally, using index notation to separate the output, we construct

$$x_{t|0,T} = \mu_{t|0,T}^{(\theta)} + \Sigma_{t|0,T}^{(\theta)} \epsilon, \quad \text{where } \epsilon \sim \mathcal{N}(0, \mathbb{I}_D). \quad (15a)$$

$$\mu_{t|0,T}^{(\theta)} = \left(1 - \frac{t}{T}\right) A + \frac{t}{T} B + \frac{t}{T} \left(1 - \frac{t}{T}\right) \text{NNET}_{\theta}(t, A, B)_{[:D]} \quad (15b)$$

$$\Sigma_{t|0,T}^{(\theta)} = \frac{t}{T} \left(1 - \frac{t}{T}\right) \text{diag}(\text{NNET}_{\theta}(t, A, B)_{[D:]}) + \sigma_{\min}^2 \mathbb{I}_D. \quad (15c)$$

Crucially, our Gaussian parameterization addresses Challenge 1, in that we can easily draw samples $x_{t|0,T} \sim q_{t|0,T}$ from our variational conditioned process (9b) *without simulating* the corresponding SDE with nonlinear drift (10). Further, the coefficients in (15b) and (15c) ensure that, as $t \rightarrow 0$ or $t \rightarrow T$, our parameterization satisfies the (smoothed) boundary conditions by design (Challenge 3). Although we add σ_{\min}^2 to ensure invertibility of $\Sigma_{t|0,T}$ (see (13)), we preserve $q_0(x_0) = \mathcal{N}(x_0 | A, \sigma_{\min}^2 \mathbb{I}_D) \approx \delta(x_0 - A)$ and $q_T(x_T) = \mathcal{N}(x_T | B, \sigma_{\min}^2 \mathbb{I}_D) \approx \delta(x_T - B)$.

Reparameterization Gradients. Having shown that our parameterization satisfies the constraints (9b)-(9c) by design, we can finally optimize our variational Doob objective with respect to $q_{t|0,T} \in \mathcal{Q}_G$ using the reparameterization trick (Kingma and Welling, 2013; Rezende et al., 2014). In particular, for the expectation at each t in (9a), we rewrite

$$\nabla_{\theta} \mathbb{E}_{q_{t|0,T}^{(\theta)}}(x) \left[\langle v_{t|0,T}^{(q,\theta)}(x), G_t v_{t|0,T}^{(q,\theta)}(x) \rangle \right] = \mathbb{E}_{\mathcal{N}(\epsilon|0, \mathbb{I}_D)} \left[\nabla_{\theta} \langle v_{t|0,T}^{(q,\theta)}(g(t, \epsilon; \theta)), G_t v_{t|0,T}^{(q,\theta)}(g(t, \epsilon; \theta)) \rangle \right],$$

where $x = g(t, \epsilon; \theta)$ is the mapping in (15) and $v_{t|0,T}^{(q,\theta)}$ depends on θ via $\mu_{t|0,T}^{(\theta)}, \Sigma_{t|0,T}^{(\theta)}$ in (13)-(14).

Full Training Algorithm. In practice, we sample a batch of times $\{t_i\}_{i=1}^M$ uniformly from the interval $t \in [0, T]$. For each time point, we approximate the gradient using a single-sample estimate of the expectation above (or (9)), which yields a simulation-free training procedure. The full training algorithm is outlined in Alg. 1.

Sampling of Trajectories. While we sample directly from $q_{t|0,T}^{(\theta)}$ during training, we can sample full trajectories which obey this sequence of marginals at test time (Alg. 2). In particular, we simulate SDE trajectories with drift $u_{t|0,T}^{(q,\theta)}(x)$ and diffusion coefficient G_t using an appropriate solver. Note that this generation scheme sidesteps computationally expensive evaluation of the force field or base drift $b_t(x_t)$. We visualize example sampling trajectories in Fig. 2.

3.2.2 Second-Order Dynamics

To handle the case of the second-order dynamics in (4), we can adapt our recipe from the previous section with minimal modifications by extending the state space $x \in \mathbb{R}^D$ to include velocities \bar{v} , with $x = (\bar{x}, \bar{v}) \in \mathbb{R}^{2D}$. However, note that the dynamics in (4) are no longer stochastic in the spatial coordinates \bar{x} . To ensure invertibility of G_t and existence of the h -transform, we add a small nonzero diffusion coefficient in the coordinate space \bar{x} , so that the reference process in Eq. (5) is given by

$$x_t = \begin{bmatrix} \bar{x}_t \\ \bar{v}_t \end{bmatrix}, \quad b_t(x_t) = \begin{bmatrix} \bar{v}_t \\ -M^{-1} \nabla_x U(\bar{x}_t) - \gamma \bar{v}_t \end{bmatrix}, \quad \Xi_t = \begin{bmatrix} \xi_{\min} \mathbb{I}_D & 0 \\ 0 & M^{-1/2} \sqrt{2\gamma k_B \mathcal{T}} \end{bmatrix}. \quad (16)$$

All steps in our algorithm proceed in similar fashion to Sec. 3.2.1. We now parameterize $q_{t|0,T}(\bar{x}, \bar{v})$ using $\text{NNET}_\theta : [0, T] \times \mathbb{R}^{2D} \times \mathbb{R}^{2D} \rightarrow \mathbb{R}^{2D} \times \mathbb{R}^{2D}$, which outputs mean perturbations and per-dimension variances to calculate $\mu_{t|0,T}^{\bar{x}}, \mu_{t|0,T}^{\bar{v}}$ and $\Sigma_{t|0,T}^{\bar{x}}, \Sigma_{t|0,T}^{\bar{v}}$ and sample (\bar{x}, \bar{v}) , as in (15). Note that we parameterize $\Sigma_{t|0,T}^{\bar{x}}, \Sigma_{t|0,T}^{\bar{v}}$ separately, matching the block diagonal form of (16). We calculate $v_{t|0,T}^{(q)}(\bar{x}, \bar{v}) := [v_{t|0,T}^{\bar{x}(q)}, v_{t|0,T}^{\bar{v}(q)}]$ from $u_{t|0,T}^{(q)}(\bar{x}, \bar{v}) = [u_{t|0,T}^{\bar{x}(q)}, u_{t|0,T}^{\bar{v}(q)}]$ as in (13)–(14), with $G_t^{-1} = (\frac{1}{2} \Xi_t \Xi_t^T)^{-1}$ given by (16). The Lagrangian objective in (9) minimizes the norm of the concatenated vector $v_{t|0,T}^{(q)}(\bar{x}, \bar{v})$, which depends on the reference drift $b_t(\bar{x}, \bar{v})$ in (16).

3.2.3 Gaussian Mixture Paths

Note that the true Doob h -transform may not yield marginal distributions which are unimodal Gaussians as in the previous section. To increase the expressivity of our variational family of conditioned processes, we now extend our parameterization to mixtures of Gaussians, $q_{t|0,T} \in \mathcal{Q}_{\text{MoG}}^K$. Given a set of K mixture weights $w := \{w^k\}_{k=1}^K$ and component Gaussian paths $\{q_{t|0,T}^k\}_{k=1}^K$, the following identity allows us to obtain the drift $u_{t|0,T}^{(q,\theta)}$ of the corresponding mixture distribution $q_{t|0,T}$.

Proposition 4. *Given a set of processes $q_{t|0,T}^k(x)$ and mixtures weights w^k , the vector field satisfying the Fokker-Planck equation in (12) for the mixture $q_{t|0,T}(x) = \sum_k w^k q_{t|0,T}^k(x)$ is given by*

$$u_{t|0,T}^{(q,\theta)}(x) = \sum_{k=1}^K \frac{w^k q_{t|0,T}^k(x)}{\sum_{j=1}^K w^j q_{t|0,T}^j(x)} u_{t|0,T}^{(q,k)}(x), \quad (17)$$

where $u_{t|0,T}^{(q,k)}(x)$ satisfies the Fokker-Planck equation in (12) for $q_{t|0,T}^k(x)$. This identity holds for both first-order dynamics in spatial coordinates only or second-order dynamics in $x = (\bar{x}, \bar{v})$.

Finally, we can calculate $v_{t|0,T}^{(q,\theta)}(x)$ by comparing $u_{t|0,T}^{(q,\theta)}(x)$ for the mixture of Gaussian path $q_{t|0,T} \in \mathcal{Q}_{\text{MoG}}^K$ to the reference drift $b_t(x)$ as in (14), and proceed to minimize its norm as in (9). In practice, we use Gumbel softmax reparameterization gradients (Maddison et al., 2016; Jang et al., 2017) to optimize the mixture weights $\{w^k\}_{k=1}^K$ alongside the neural network parameters $\{\theta^k\}_{k=1}^K$ for each Gaussian component $\{\mu_{t|0,T}^{(\theta)}, \Sigma_{t|0,T}^{(\theta)}\}_{k=1}^K$ and either first- or second-order dynamics.

4 Related Work

(Aligned) Schrödinger Bridge Matching Methods. Many existing ‘bridge matching’ approaches (Shi et al., 2023; Peluchetti, 2021, 2023; Liu et al., 2022; Lipman et al., 2022; Liu et al., 2023b) for SB and generative modeling rely on convenient properties of Brownian bridges and would require calculating h -transforms to simulate bridges for general reference processes. Our conditional

Gaussian path parameterization is similar to Liu et al. (2023a); Neklyudov et al. (2024), where analytic bridges are not available for SB problems with nonlinear reference drift or general costs.

Somnath et al. (2023); Liu et al. (2023b) attempt to solve the SB problem given access to aligned data $x_0, x_T \sim q_{0,T}^{\text{data}}$ assumed to be drawn from an optimal coupling. While the method in Somnath et al. (2023) involves approximating an h -transform, their goal is to obtain an unconditioned vector field v_t to simulate a Markov process. However, De Bortoli et al. (2023) use Doob’s h -transform to argue the learned Markov process will not preserve the empirical coupling unless $q_{0,T}^{\text{data}}$ is the optimal coupling for the SB problem, and show that an ‘augmented’ $v_{0,t}$ which conditions on x_0 can correct this issue.

After training on a dataset of $x_0, x_T \sim q_{0,T}^{\text{data}}$ pairs using our method, we could consider using an (augmented) bridge matching objective (Shi et al., 2023; De Bortoli et al., 2023) to distill our learned $v_{t|0,T}^{(q)}$ into a vector field v_t or $v_{0,t}$ which does not condition on the endpoint. Our use of a Gaussian path parameterization with samples from a fixed endpoint coupling and no Markovization step corresponds to a simplified version of the conditional optimal control step in Liu et al. (2023a).

Transition Path Sampling. We refer to the surveys of Dellago et al. (2002); Weinan and Vandeen-Eijnden (2010); Bolhuis and Swenson (2021) for an overview of the TPS problem. Least action principles for TPS have a long history, building upon the Freidlin-Wentzell (Freidlin and Wentzell, 1998) and Onsager-Machlup (Onsager and Machlup, 1953; Dürr and Bach, 1978) Lagrangian functionals in the zero-noise limit and finite-noise cases. In particular, the Onsager-Machlup functional relates maximum a posteriori estimators or ‘most probable (conditioned) paths’ to the minimizers of an action functional similar to Thm. 1, where example algorithms include (Vanden-Eijnden and Heymann, 2008; Sheppard et al., 2008). By contrast, our approach targets the *entire* posterior over transition paths using an expressive variational family. While Lu et al. (2017) provide analysis for the Gaussian family, we draw connections with Doob’s h -transform and extend to mixtures of Gaussians.

Shooting methods are among the most popular for sampling the posterior of transition paths. From a path that satisfies the boundary conditions (obtained, e.g., using high-temperature simulations), shooting picks points and directions to propose alterations, then simulates new trajectories and accepts or rejects using Metropolis-Hastings (MH) (Juraszek and Bolhuis, 2008; Borrero and Dellago, 2016; Jung et al., 2017; Falkner et al., 2023; Jung et al., 2023). While the MCMC corrections yield theoretical guarantees, shooting methods involve expensive molecular dynamics (MD) simulations and need to balance high rejection rates with large changes in trajectories. One-way shooting methods sample paths efficiently but yield highly correlated samples. Two-way shooting methods, which we compare to in Sec. 5, are more expensive but typically sample diverse paths faster. Recent machine learning approaches (e.g. Plainer et al. (2023); Lelièvre et al. (2023)) aim to reduce the need for MD.

Finally, various related methods rely on iterative simulation of SDE in (10) during training to learn the control drift term. Yan et al. (2022); Holdijk et al. (2023) are motivated from the perspective of stochastic optimal control, while Das et al. (2021); Rose et al. (2021) develop actor-critic methods using closely-related ideas from soft reinforcement learning. The variational method in Das et al. (2019) optimizes the rate function quantifying the probability of the rare events, while Singh and Limmer (2023) solves the Kolmogorov backward equation to learn the Doob’s h -transform. However, all of these methods may be inefficient if the desired terminal state is sampled infrequently.

5 Experiments

We investigate the capabilities of our approach across a variety of different settings. We first illustrate features of our method on toy potentials before continuing to real-world molecular systems, including a commonly-used benchmark system, alanine dipeptide, and a small protein, Chignolin. The code behind our method is available at <https://github.com/plainerman/variational-doob>. Before diving into the experiments, we introduce the evaluation procedure and baseline methods.

Evaluation metrics. In our evaluation, we emphasize two key quantities: accuracy and efficiency. Efficiency is evaluated by the number of calls to the potential energy function, which requires extensive computation and dominates the runtime of larger molecules. For accuracy, we evaluate the log-likelihood of each sampled path and the maximum energy point (saddle point/transition state) along each sampled path. A good method samples many probable paths (i.e., high log-likelihood) and an accurate transition state (i.e., small maximum energy). See App. D for further details.

Baselines. We compare our approach against the MCMC-based two-way shooting method with uniform point selection with variable or fixed length trajectories. We found that two-way shooting




Figure 2: Comparing path histograms and trajectories of TPS using fixed-length two-way shooting and comparing it with our variational approach.




Figure 3: Illustration of the expressivity of unimodal Gaussian versus mixture of Gaussian paths on a symmetric potential with two transition path modes.

| Method | # Evaluations (\downarrow) | Max Energy (\downarrow) | MinMax Energy (\downarrow) | Log-Likelihood (\uparrow) | Max Log-Likelihood (\uparrow) |
|-----------------|--------------------------------|-----------------------------|--------------------------------|-------------------------------|-----------------------------------|
| MCMC (variable) | 3.53M | -13.77 ± 16.43 | -40.75 | - | - |
| MCMC | 1.03B | -17.80 ± 14.77 | -40.21 | 866.56 ± 17.00 | 907.15 |
| Ours | 1.28M | -14.81 ± 13.73 | -40.56 | 858.50 ± 17.61 | 909.74 |

Table 1: Transition path sampling experiment for Müller-Brown potential. We report the number of potential evaluations needed to sample 1,000 paths, as well as the maximum energy and the likelihood of each path (including mean and standard deviation). The methods marked with ‘variable’ use a variable length setting. MinMax energy reports the lowest maximum energy across all paths (i.e., energy of lowest transition state).

produced the most diverse path ensembles among possible baselines, although the acceptance probability can be relatively low for systems dominated by diffusive dynamics (Brotzakis and Bolhuis, 2016) and might be improved by better shooting point selection. This baseline gives theoretical guarantees about the ensemble and thus can be considered as a proxy for the ground truth. In that sense, our goal is not to beat two-way shooting but to approximate it with fewer potential evaluations.

5.1 Synthetic Datasets

Müller-Brown Potential. The Müller-Brown potential is a popular benchmark to study transition path sampling between metastable states. It consists of three local minima, and we aim to sample transition paths connecting state A and state B with a circular state definition. In Fig. 2, we visualize the potential and the sampled paths and can see that the same ensemble is sampled for both our method and two-way shooting. Our method exhibits a slightly reduced variance for unlikely transitions. In Table 1, we can observe that MCMC-based methods require many potential evaluations to achieve a good result, which comes from the low acceptance rate (especially when fixing the lengths of trajectories). Our method requires fewer energy evaluations (1 million vs. 1 billion) while finding paths with similar energy and likelihood. Note that the likelihood for variable approaches has been omitted, as it is governed by the number of steps in the trajectory and cannot be compared directly.

Gaussian Mixture. We further consider a potential in which the states are separated by a symmetric high-energy barrier that allows for two distinct reaction channels. In Fig. 3, we observe that a single Gaussian path cannot model a system with multiple modes of transition paths. Nevertheless, this issue can be resolved using a mixture of Gaussian paths, with slightly increased computational cost.

The Case for Neural Networks. According to our empirical study, the neural network parameterization of the Gaussian distribution statistics $\mu_{t|0,T}, \Sigma_{t|0,T}$ is an invaluable part of our framework. As an ablation, we consider parameterizing $\mu_{t|0,T}, \Sigma_{t|0,T}$ as piecewise linear splines whose intermediate points are updated using the same gradient-based optimizer as used for neural network training. In App. D.3, we report results comparing the W1 distance of learned marginals using neural network versus spline parameterizations, observing that splines yield inferior results even after an order of magnitude more potential function evaluations. We thus conclude that spline parameterizations are not competitive for learning transition paths and continue to focus on our neural-network approach.

5.2 Second-order Dynamics and Molecular Systems

Experiment Setup. We evaluate our methods on real-world high-dimensional molecular systems governed by the second-order dynamics (4): *alanine dipeptide* and *Chignolin*. Alanine dipeptide is a well-studied system of 22 atoms (66 total degrees of freedom), where the molecule can be described by two collective variables (CV): the dihedral angles ϕ, ψ . Chignolin is a larger system consisting of 10 residues with 138 atoms (414 total degrees of freedom) that cannot be summarized as easily. We use an AMBER14 force field (Maier et al., 2015) implemented in OpenMM (Eastman et al., 2017) but use DMFF (Wang et al., 2023) to backpropagate through the energy evaluations.

| Method | States | # Evaluations (\downarrow) | Max Energy (\downarrow) | MinMax Energy (\downarrow) |
|------------------------------|---------|--------------------------------|-----------------------------|--------------------------------|
| MCMC (variable length) | CV | 21.02M | 740.70 ± 695.79 | 52.37 |
| MCMC* | CV | 1.29B* | 288.46 ± 128.31 | 60.52 |
| MCMC (variable length) | relaxed | 187.54M | 412.65 ± 334.70 | 26.97 |
| MCMC | relaxed | > 10B | N/A | N/A |
| MCMC (variable length) | exact | > 10B | N/A | N/A |
| MCMC | exact | > 10B | N/A | N/A |
| Ours (Cartesian) | exact | 38.40M | 304.16 ± 0.02 | 304.10 |
| Ours (Cartesian, 2 Mixtures) | exact | 51.20M | 838.23 ± 79.93 | 328.46 |
| Ours (Cartesian, 5 Mixtures) | exact | 51.20M | 570.26 ± 273.93 | 293.70 |
| Ours (Internal) | exact | 38.40M | 69.63 ± 0.01 | 69.59 |
| Ours (Internal, 2 Mixtures) | exact | 51.20M | 69.26 ± 0.21 | 68.79 |
| Ours (Internal, 5 Mixtures) | exact | 51.20M | 69.49 ± 0.28 | 69.06 |

Table 2: Transition path sampling for alanine dipeptide. For MCMC methods, we compare different state definitions of A, B : ‘CV’ uses ϕ, ψ angles. ‘Exact’ uses a very small threshold of aligned root-mean-square deviation (RMSD) around reference states A, B (as in Ours). ‘Relaxed’ uses a larger threshold of RMSD around A, B . The method marked with a * only samples 100 paths due to computational limitations, while others sample 1,000. Fields with N/A are intractable as a single trajectory requires more than 1 billion potential evaluations.




Figure 4: Transition path for the protein Chignolin. The energy plot a transition path in which the protein folds in $T = 1,000 fs$, and passes a high energy barrier at $460 fs$ with about $3,000 \frac{kJ}{mol}$.

Alanine Dipeptide. In Table 2, we report results for four variants of our models, which either predict Cartesian coordinates or internal coordinates in the form of bond lengths and dihedral angles (compare App. D.4), either with or without Gaussian mixture. For our method, operating in internal coordinates takes more iterations to converge but yields better results compared to Cartesian coordinates, where the internal coordinates are distributed similarly to Gaussians and our network does not need to learn equivariances (Du et al., 2022). Similarly, Gaussian mixture paths perform slightly better than a single Gaussian path due to the additional expressiveness. We note that paths sampled with Gaussian mixture exhibit a larger variance in max energy as they represent multiple reaction channels.

We find that prior-informed definitions of the desired initial and target states (i.e., CV) are necessary for MCMC to work efficiently with fixed-length trajectories. Finding these CVs in practice is challenging and only possible in this instance because the molecule is small and well-studied. For the larger system size in Table 2, it becomes intractable to use MCMC to connect precise states A, B (‘exact’) instead of larger regions (‘relaxed’), even with a single trajectory. Variable length MCMC with relaxed endpoint conditions with CV perform well on this task, but our method is competitive using fewer evaluations and more strict boundary conditions. Fixed-length MCMC, even with prior-informed knowledge, can only find 100 trajectories while needing 50 times more potential evaluations compared to variable length.

Chignolin. The folding dynamics of Chignolin already pose a challenge and have not yet been well-studied compared to alanine dipeptide. We illustrate the qualitative experimental results for this system in Fig. 4. Operating in Cartesian space, our model samples a feasible transition within 25.6M potential energy evaluation calls and a transition with a duration of $T = 1 ps$.

6 Conclusion, Limitations and Future Work

In this paper, we propose an efficient computational framework for transition path sampling with Brownian dynamics. We formulate the transition path sampling problem by using Doob’s h -transform to condition a reference stochastic process, and propose a variational formulation for efficient optimization. Specifically, we propose a simulation-free training objective and model parameterization that imposes boundary conditions as hard constraints. We compare our method with MCMC-based baselines and show comparable accuracy with lower computational costs on both synthetic datasets and real-world molecular systems. Our method is currently limited by rigidly defining states A and B to be a point mass with Gaussian noise instead of any arbitrary set. Finally, our method might be improved by accommodating variable length paths.

Acknowledgments and Disclosure of Funding

The authors would like to thank Juno Nam and Soojung Yang for spotting the unphysical steric barrier in the original pair of initial and target states and Guan-Horng Liu, Maurice Weiler, Hannes Stärk, and Yanze Wang for helpful discussions. The work of Yuanqi Du and Carla P. Gomes was supported by the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Futures program; the National Science Foundation (NSF), the Air Force Office of Scientific Research (AFOSR); the Department of Energy; and the Toyota Research Institute (TRI). The work of Michael Plainer was supported by the Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA) through the DAAD program Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the German Ministry of Education and Research and by Max Planck Society. The work of Kirill Neklyudov was supported by IVADO.

References

- Anderson, J. B. (2007). Predicting rare events in molecular dynamics. *Advances in Chemical Physics*, 91:381–431.
- Batatia, I., Benner, P., Chiang, Y., Elena, A. M., Kovács, D. P., Riebesell, J., Advincula, X. R., Asta, M., Baldwin, W. J., Bernstein, N., et al. (2023). A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096*.
- Blau, S. M., Patel, H. D., Spotte-Smith, E. W. C., Xie, X., Dwaraknath, S., and Persson, K. A. (2021). A chemically consistent graph architecture for massive reaction networks applied to solid-electrolyte interphase formation. *Chemical science*, 12(13):4931–4939.
- Bolhuis, P. G. and Swenson, D. W. H. (2021). Transition path sampling as markov chain monte carlo of trajectories: Recent algorithms, software, applications, and future outlook. *Advanced Theory and Simulations*, 4(4).
- Borrero, E. and Dellago, C. (2016). Avoiding traps in trajectory space: Metadynamics enhanced transition path sampling. *The European Physical Journal Special Topics*, 225(8-9):1609–1620.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs.
- Brotzakis, Z. F. and Bolhuis, P. G. (2016). A one-way shooting algorithm for transition path sampling of asymmetric barriers. *The Journal of Chemical Physics*, 145(16):164112.
- Castellan, G. W. (1983). *Physical Chemistry*. Addison-Wesley, Reading, Mass, 3rd ed edition.
- Chen, T., Liu, G.-H., and Theodorou, E. (2021a). Likelihood training of schrödinger bridge using forward-backward sdes theory. In *International Conference on Learning Representations*.
- Chen, Y., Georgiou, T. T., and Pavon, M. (2016). On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169:671–691.
- Chen, Y., Georgiou, T. T., and Pavon, M. (2021b). Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *Siam Review*, 63(2):249–313.
- Chopin, N., Fulop, A., Heng, J., and Thiery, A. H. (2023). Computational doob h-transforms for online filtering of discretely observed diffusions. In *International Conference on Machine Learning*, pages 5904–5923. PMLR.
- Das, A., Kuznets-Speck, B., and Limmer, D. T. (2022). Direct evaluation of rare events in active matter from variational path sampling. *Physical Review Letters*, 128(2):028005.
- Das, A., Limmer, D. T., and Limmer, D. T. (2019). Variational control forces for enhanced sampling of nonequilibrium molecular dynamics simulations. *The Journal of chemical physics*, 151(24).
- Das, A., Rose, D. C., Garrahan, J. P., and Limmer, D. T. (2021). Reinforcement learning of rare diffusive dynamics. *The Journal of Chemical Physics*, 155(13).

- De Bortoli, V., Liu, G.-H., Chen, T., Theodorou, E. A., and Nie, W. (2023). Augmented bridge matching. *arXiv preprint arXiv:2311.06978*.
- De Bortoli, V., Thornton, J., Heng, J., and Doucet, A. (2021). Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709.
- Dellago, C., Bolhuis, P. G., and Geissler, P. L. (2002). Transition path sampling. *Advances in chemical physics*, 123:1–78.
- Dickson, A. (2018). Mapping the ligand binding landscape. *Biophysical Journal*, 115(9):1707–1719.
- Doob, J. L. (1957). Conditional brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85:431–458.
- Du, W., Zhang, H., Du, Y., Meng, Q., Chen, W., Zheng, N., Shao, B., and Liu, T.-Y. (2022). Se (3) equivariant graph neural networks with complete local frames. In *International Conference on Machine Learning*, pages 5583–5608. PMLR.
- Dürr, D. and Bach, A. (1978). The onsager-machlup function as lagrangian for the most probable path of a diffusion process. *Communications in Mathematical Physics*, 60:153–170.
- Eastman, P., Swails, J., Chodera, J. D., McGibbon, R. T., Zhao, Y., Beauchamp, K. A., Wang, L.-P., Simmonett, A. C., Harrigan, M. P., Stern, C. D., Wiewiora, R. P., Brooks, B. R., and Pande, V. S. (2017). OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology*, 13(7):e1005659.
- Escobedo, F. A., Borrero, E. E., and Araque, J. C. (2009). Transition path sampling and forward flux sampling. applications to biological systems. *Journal of Physics: Condensed Matter*, 21(33):333101.
- Falkner, S., Coretti, A., Romano, S., Geissler, P., and Dellago, C. (2023). Conditioning normalizing flows for rare event sampling. *arXiv preprint arXiv:2207.14530*.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. (2021). Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8.
- Freidlin, M. and Wentzell, A. (1998). *Random perturbations of Dynamical Systems*. Springer.
- Gabrié, M., Rotskoff, G. M., and Vanden-Eijnden, E. (2022). Adaptive monte carlo augmented with normalizing flows. *Proceedings of the National Academy of Sciences*, 119(10):e2109420119.
- Heng, J., De Bortoli, V., Doucet, A., and Thornton, J. (2021). Simulating diffusion bridges with score matching. *arXiv preprint arXiv:2111.07243*.
- Holdijk, L., Du, Y., Hooft, F., Jaini, P., Ensing, B., and Welling, M. (2023). Stochastic optimal control for collective variable free sampling of molecular transition paths. *Advances in Neural Information Processing Systems*, 36.
- Jamison, B. (1975). The markov processes of schrödinger. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 32(4):323–331.
- Jang, E., Gu, S., and Poole, B. (2017). Categorical reparametrization with gumble-softmax. In *International Conference on Learning Representations (ICLR 2017)*. OpenReview. net.
- Jung, H., Covino, R., Arjun, A., Leitold, C., Dellago, C., Bolhuis, P. G., and Hummer, G. (2023). Machine-guided path sampling to discover mechanisms of molecular self-organization. *Nature Computational Science*, 3(4):334–345.
- Jung, H., Ichi Okazaki, K., and Hummer, G. (2017). Transition path sampling of rare events by shooting from the top. *The Journal of Chemical Physics*, 147(15).

- Juraszek, J. and Bolhuis, P. G. (2008). Rate constant and reaction coordinate of trp-cage folding in explicit water. *Biophysical Journal*, 95(9):4246–4257.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *International Conference on Learning Representations*.
- Kirmizialtin, S., Johnson, K. A., and Elber, R. (2015). Enzyme selectivity of HIV reverse transcriptase: Conformations, ligands, and free energy partition. *The Journal of Physical Chemistry B*, 119(35):11513–11526.
- Kirmizialtin, S., Nguyen, V., Johnson, K. A., and Elber, R. (2012). How conformational dynamics of DNA polymerase select correct substrates: Experiments and simulations. *Structure*, 20(4):618–627.
- Klucznik, T., Syntrivanis, L.-D., Baś, S., Mikulak-Klucznik, B., Moskal, M., Szymkuć, S., Mlynarski, J., Gadina, L., Beker, W., Burke, M. D., et al. (2024). Computational prediction of complex cationic rearrangement outcomes. *Nature*, 625(7995):508–515.
- Koehl, P. and Orland, H. (2022). Sampling constrained stochastic trajectories using brownian bridges. *The Journal of Chemical Physics*, 157(5).
- Lelièvre, T., Robin, G., Sekkat, I., Stoltz, G., and Cardoso, G. V. (2023). Generative methods for sampling transition paths in molecular dynamics. *ESAIM: Proceedings and Surveys*, 73:238–256.
- Léonard, C. (2014). A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems-A*, 34(4):1533–1574.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. (2022). Flow matching for generative modeling. *International Conference on Learning Representations*.
- Liu, G.-H., Lipman, Y., Nickel, M., Karrer, B., Theodorou, E., and Chen, R. T. (2023a). Generalized schrödinger bridge matching. In *The Twelfth International Conference on Learning Representations*.
- Liu, G.-H., Vahdat, A., Huang, D.-A., Theodorou, E., Nie, W., and Anandkumar, A. (2023b). I²sb: Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872*.
- Liu, X., Wu, L., Ye, M., and Liu, Q. (2022). Let us build bridges: Understanding and extending diffusion generative models. *arXiv preprint arXiv:2208.14699*.
- Liu, X., Wu, L., Ye, M., and qiang liu (2023c). Learning diffusion bridges on constrained domains. In *The Eleventh International Conference on Learning Representations*.
- Lu, Y., Stuart, A., and Weber, H. (2017). Gaussian approximations for transition paths in brownian dynamics. *SIAM Journal on Mathematical Analysis*, 49(4):3005–3047.
- Maddison, C. J., Mnih, A., and Teh, Y. W. (2016). The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*.
- Maier, J. A., Martinez, C., Kasavajhala, K., Wickstrom, L., Hauser, K. E., and Simmerling, C. (2015). ff14sb: improving the accuracy of protein side chain and backbone parameters from ff99sb. *Journal of chemical theory and computation*, 11(8):3696–3713.
- Mehdi, S., Smith, Z., Herron, L., Zou, Z., and Tiwary, P. (2024). Enhanced sampling with machine learning. *Annual Review of Physical Chemistry*, 75.
- Neklyudov, K., Brekelmans, R., Severo, D., and Makhzani, A. (2023). Action matching: Learning stochastic dynamics from samples. In *International Conference on Machine Learning*.
- Neklyudov, K., Brekelmans, R., Tong, A., Atanackovic, L., Liu, Q., and Makhzani, A. (2024). A computational framework for solving wasserstein lagrangian flows. *International Conference on Machine Learning*.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. (2019). Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147.

- Noé, F., Schütte, C., Vanden-Eijnden, E., Reich, L., and Weikl, T. R. (2009). Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences*, 106(45):19011–19016.
- Onsager, L. and Machlup, S. (1953). Fluctuations and irreversible processes. *Physical Review*, 91(6):1505.
- Papaspiliopoulos, O. and Roberts, G. (2012). Importance sampling techniques for estimation of diffusion models. *Statistical methods for stochastic differential equations*, 124:311–340.
- Peluchetti, S. (2021). Non-denoising forward-time diffusions.
- Peluchetti, S. (2023). Diffusion bridge mixture transports, Schrödinger bridge problems and generative modeling. *arXiv preprint arXiv:2304.00917*.
- Plainer, M., Stärk, H., Bunne, C., and Günnemann, S. (2023). Transition path sampling with boltzmann generator-based MCMC moves. In *Generative AI and Biology Workshop*.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR.
- Rose, D. C., Mair, J. F., and Garrahan, J. P. (2021). A reinforcement learning approach to rare trajectory sampling. *New Journal of Physics*, 23(1):013013.
- Rotskoff, G. M. (2024). Sampling thermodynamic ensembles of molecular systems with generative neural networks: Will integrating physics-based models close the generalization gap? *Current Opinion in Solid State and Materials Science*, 30:101158.
- Särkkä, S. and Solin, A. (2019). *Applied stochastic differential equations*, volume 10. Cambridge University Press.
- Schauer, M., van der Meulen, F., and van Zanten, H. (2017). Guided proposals for simulating multi-dimensional diffusion bridges. *Bernoulli*, 23(4A).
- Schrödinger, E. (1932). Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pages 269–310.
- Selli, D., Boulfelfel, S. E., Schapotschnikow, P., Donadio, D., and Leoni, S. (2016). Hierarchical thermoelectrics: crystal grain boundaries as scalable phonon scatterers. *Nanoscale*, 8(6):3729–3738.
- Sharma, N. D., Singh, J., Vijay, A., Samanta, K., Dogra, S., and Bandyopadhyay, A. K. (2016). Pressure-induced structural transition trends in nanocrystalline rare-earth sesquioxides: A raman investigation. *The Journal of Physical Chemistry C*, 120(21):11679–11689.
- Sheppard, D., Terrell, R., and Henkelman, G. (2008). Optimization methods for finding minimum energy paths. *The Journal of chemical physics*, 128(13).
- Shi, Y., De Bortoli, V., Campbell, A., and Doucet, A. (2023). Diffusion schrödinger bridge matching. *arXiv preprint arXiv:2303.16852*.
- Shoghi, N., Kolluru, A., Kitchin, J. R., Ulissi, Z. W., Zitnick, C. L., and Wood, B. M. (2023). From molecules to materials: Pre-training large generalizable models for atomic property prediction. In *The Twelfth International Conference on Learning Representations*.
- Sidky, H., Chen, W., and Ferguson, A. L. (2020). Molecular latent space simulators. *Chemical Science*, 11(35):9459–9467.
- Singh, A. N. and Limmer, D. T. (2023). Variational deep learning of equilibrium transition path ensembles. *The Journal of Chemical Physics*, 159(2).
- Smith, J. S., Isayev, O., and Roitberg, A. E. (2017). Ani-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chemical science*, 8(4):3192–3203.

- Somnath, V. R., Pariset, M., Hsieh, Y.-P., Martinez, M. R., Krause, A., and Bunne, C. (2023). Aligned diffusion schrödinger bridges. In *Uncertainty in Artificial Intelligence*, pages 1985–1995. PMLR.
- Tong, A., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Fatras, K., Wolf, G., and Bengio, Y. (2023). Conditional flow matching: Simulation-free dynamic optimal transport. *arXiv preprint arXiv:2302.00482*, 2(3).
- Vanden-Eijnden, E. and Heymann, M. (2008). The geometric minimum action method for computing minimum energy paths. *The Journal of chemical physics*, 128(6).
- Wang, H., Zhang, L., Han, J., and Weinan, E. (2018). Deepmd-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Computer Physics Communications*, 228:178–184.
- Wang, X., Li, J., Yang, L., Chen, F., Wang, Y., Chang, J., Chen, J., Feng, W., Zhang, L., and Yu, K. (2023). DMFF: An open-source automatic differentiable platform for molecular force field development and molecular dynamics simulation. *Journal of Chemical Theory and Computation*, 19(17):5897–5909.
- Weinan, E. and Vanden-Eijnden, E. (2010). Transition-path theory and path-finding algorithms for the study of rare events. *Annual review of physical chemistry*, 61(2010):391–420.
- Wu, L., Gong, C., Liu, X., Ye, M., and Liu, Q. (2022). Diffusion-based molecule generation with informative prior bridges. *Advances in Neural Information Processing Systems*, 35:36533–36545.
- Xi, L., Shah, M., and Trout, B. L. (2013). Hopping of water in a glassy polymer studied via transition path sampling and likelihood maximization. *The Journal of Physical Chemistry B*, 117(13):3634–3647.
- Yan, J., Touchette, H., and Rotskoff, G. M. (2022). Learning nonequilibrium control forces to characterize dynamical phase transitions. *Physical Review E*, 105(2):024115.
- Zeng, J., Cao, L., Xu, M., Zhu, T., and Zhang, J. Z. (2020). Complex reaction processes in combustion unraveled by neural network-based molecular dynamics simulation. *Nature communications*, 11(1):5713.
- Zhang, D., Bi, H., Dai, F.-Z., Jiang, W., Zhang, L., and Wang, H. (2022). Dpa-1: Pretraining of attention-based deep potential model for molecular simulation. *arXiv preprint arXiv:2208.08236*.

A Proofs

A.1 Proofs from Sec. 2.2 (Doob's h -Transform Background)

Proposition 1. [Jamison (1975, Thm. 2)] Let $h_{\mathcal{B}}(x, t) := \rho_T(x_T \in \mathcal{B} | x_t = x)$ denote the conditional transition probability of the reference process in (5). Then,

$$\mathbb{P}_{0:T}^* : \quad dx_{t|T} = \left(b_t(x_{t|T}) + 2G_t \nabla_x \log h_{\mathcal{B}}(x_{t|T}, t) \right) \cdot dt + \Xi_t dW_t \quad x_0 \sim \rho_0 \quad (6)$$

where we use $x_{t|T}$ to denote a conditional process. The SDE in (6) is associated with the following transition probabilities for $s < t < T$,

$$\rho_t(y | x_s = x, x_T \in \mathcal{B}) = \frac{h_{\mathcal{B}}(y, t)}{h_{\mathcal{B}}(x, s)} \rho_t(y | x_s = x), \quad (7)$$

Note that all of our subsequent results hold for the case when \mathcal{B} is a point-mass, with the only change being that the h -function becomes a density, $h_{\mathcal{B}}(x, t) = \rho_T(\mathcal{B} | x_t = x)$.

Proof. See Jamison (1975) for a simple proof based on Ito's Lemma, assuming smoothness and strict positivity of h . \square

Proposition 2. The following PDEs are obeyed by (a) the density of the conditioned process $\rho_{t|0,T}(x) := \rho_t(x | x_0 = A, x_T \in \mathcal{B})$ and (b) the h -function $h_{\mathcal{B}}(x, t)$,

$$\frac{\partial \rho_{t|0,T}(x)}{\partial t} + \langle \nabla_x, \rho_{t|0,T}(x) (b_t(x) + 2G_t \nabla_x \log h_{\mathcal{B}}(x, t)) \rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \rho_{t|0,T}(x) = 0, \quad (8a)$$

$$\frac{\partial h_{\mathcal{B}}(x, t)}{\partial t} + \langle \nabla_x h_{\mathcal{B}}(x, t), b_t(x) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} h_{\mathcal{B}}(x, t) = 0. \quad (8b)$$

Reparameterizing (8b) in terms of $s_{\mathcal{B}}(x, t) := \log h_{\mathcal{B}}(x, t)$, we can also write

$$\frac{\partial s_{\mathcal{B}}(x, t)}{\partial t} + \langle \nabla_{s_{\mathcal{B}}}(x, t), G_t \nabla_{s_{\mathcal{B}}}(x, t) \rangle + \langle \nabla_{s_{\mathcal{B}}}(x, t), b_t(x) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s_{\mathcal{B}}(x, t) = 0. \quad (8c)$$

Proof. Let $p(x_{t+s} = y | x_t = x)$ denote the transition probability of a reference diffusion process

$$\frac{\partial}{\partial s} p(x_{t+s} = y | x_t = x) = -\langle \nabla_y, p(x_{t+s} = y | x_t = x) b_{t+s}(y) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} p(x_{t+s} = y | x_t = x), \quad (18)$$

where $(G_t)_{ij} = \frac{1}{2} \Xi_{t+s}^T \Xi_{t+s}$.

Now we condition the process on the end-point value $x_T \in \mathcal{B}$, and we get another kernel, i.e.

$$p(x_{t+s} = y | x_t = x, x_T \in \mathcal{B}) = \frac{p(x_T \in \mathcal{B} | x_{t+s} = y)}{p(x_T \in \mathcal{B} | x_t = x)} p(x_{t+s} = y | x_t = x). \quad (19)$$

We let $h_{\mathcal{B}}(x, t) = p(x_T \in \mathcal{B} | x_t = x)$ denote the conditional probability over the desired endpoint condition given $x_t = x$. According to laws of conditional probability, we can describe how $h_{\mathcal{B}}(x, t)$ changes in time using the unconditioned transition probability

$$\underbrace{p(x_T \in \mathcal{B} | x_t = x)}_{h(x, t)} = \int dy \underbrace{p(x_T \in \mathcal{B} | x_{t+s} = y)}_{h(y, t+s)} p(x_{t+s} = y | x_t = x), \quad (20)$$

we take the derivative $\frac{\partial}{\partial s}$ on both sides, and we get

$$0 = \int dy \left[p(x_{t+s} = y | x_t = x) \frac{\partial h_{\mathcal{B}}(y, t+s)}{\partial s} + \frac{\partial p(x_{t+s} = y | x_t = x)}{\partial s} h_{\mathcal{B}}(y, t+s) \right]. \quad (21)$$

Using the FP equation for the transition probability and integrating by parts, we have

$$0 = \int dy p(x_{t+s} = y | x_t = x) \left[\frac{\partial h_{\mathcal{B}}(y, t+s)}{\partial s} + \langle \nabla_y h_{\mathcal{B}}(y, t+s), b_t(y) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} h_{\mathcal{B}}(y, t+s) \right].$$

Note that this holds $\forall x$, hence, we have

$$\frac{\partial h_{\mathcal{B}}(y, t+s)}{\partial s} + \langle \nabla_y h_{\mathcal{B}}(y, t+s), b_{t+s}(y) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} h_{\mathcal{B}}(y, t+s) = 0,$$

without any loss of generality we can set $t = 0$

$$\frac{\partial h_{\mathcal{B}}(y, s)}{\partial s} + \langle \nabla_y h_{\mathcal{B}}(y, s), b_s(y) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} h_{\mathcal{B}}(y, s) = 0. \quad (22)$$

as desired to prove the optimality condition in (8b).

To prove (8a), denote $p(y, s) = p(x_s = y | x_0 = x)$ and differentiate $p(x_s = y | x_0 = x, x_T \in \mathcal{B}) = \frac{h(y, s)}{h(x, 0)} p(y, s)$ as

$$\begin{aligned} & \frac{\partial}{\partial s} p(x_s = y | x_0 = x, x_T \in \mathcal{B}) \\ &= \frac{1}{h_{\mathcal{B}}(x, 0)} \left[p(y, s) \frac{\partial h_{\mathcal{B}}(y, s)}{\partial s} + h_{\mathcal{B}}(y, s) \frac{\partial p(y, s)}{\partial s} \right] \\ &= \frac{1}{h_{\mathcal{B}}(x, 0)} \left[- \langle \nabla_y h_{\mathcal{B}}(y, s), p(y, s) b_s(y) \rangle - p(y, s) \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} h_{\mathcal{B}}(y, s) \right. \\ &\quad \left. - h_{\mathcal{B}}(y, s) \langle \nabla_y, p(y, s) b_s(y) \rangle + h_{\mathcal{B}}(y, s) \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} p(y, s) \right] \\ &= - \left\langle \nabla_y, \frac{h_{\mathcal{B}}(y, s)}{h_{\mathcal{B}}(x, 0)} p(y, s) b_s(y) \right\rangle - p(y, s) \left\langle \nabla_y, 2D \nabla_y \frac{h_{\mathcal{B}}(y, s)}{h_{\mathcal{B}}(x, 0)} \right\rangle \\ &\quad \pm \left\langle \nabla_y p(y, s), 2D \nabla_y \frac{h_{\mathcal{B}}(y, s)}{h_{\mathcal{B}}(x, 0)} \right\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} \left(\frac{h_{\mathcal{B}}(y, s)}{h_{\mathcal{B}}(x, 0)} p(y, s) \right), \end{aligned}$$

Note that $h_{\mathcal{B}}(x, 0)$ can be pulled outside the differential operator because it is a function of x . The PDE for the new kernel $p(y, s | \mathcal{B}) = p(x_s = y | x_0 = x, x_T \in \mathcal{B})$ (conditioned on the end-point) becomes

$$\frac{\partial}{\partial s} p(y, s | \mathcal{B}) = - \langle \nabla_y, p(y, s | \mathcal{B}) (b_s(y) + 2D \nabla_y \log h_{\mathcal{B}}(y, s)) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} p(y, s | \mathcal{B}). \quad (23)$$

which matches the desired PDE in (8a) thereby proving the first two parts of Prop. 2.

Finally, to show (8c), we index time using t in Eq. (22) and change variables $h_{\mathcal{B}}(x, t) = e^{s(x, t)}$,

$$\frac{\partial e^{s(x, t)}}{\partial t} + \langle \nabla_x e^{s(x, t)}, b_t(x) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} e^{s(x, t)} = 0.$$

$$e^{s(x, t)} \frac{\partial s(x, t)}{\partial t} + e^{s(x, t)} \langle \nabla_x s(x, t), b_t(x) \rangle + \langle \nabla, G_t \nabla e^{s(x, t)} \rangle = 0$$

Next, we simplify $\langle \nabla, G_t \nabla e^{s(x, t)} \rangle = \langle \nabla, G_t e^{s(x, t)} \nabla s(x, t) \rangle = \langle \nabla e^{s(x, t)}, G_t \nabla s(x, t) \rangle + e^{s(x, t)} \langle \nabla, G_t \nabla s(x, t) \rangle = e^{s(x, t)} \langle \nabla s(x, t), G_t \nabla s(x, t) \rangle + e^{s(x, t)} \langle \nabla, G_t \nabla s(x, t) \rangle$ to finally write

$$e^{s(x, t)} \left(\frac{\partial s(x, t)}{\partial t} + \langle \nabla_x s(x, t), b_t(x) \rangle + \langle \nabla s(x, t), G_t \nabla s(x, t) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s(x, t) \right) = 0$$

which demonstrates (8c) since the inner term must be zero. \square

A.2 Proofs from Sec. 3.1 (Lagrangian Action Minimization for Doob's h -Transform)

We begin by proving [Cor. 1](#), whose proof actually contains the initial steps needed to prove our main theorem [Thm. 1](#). In both proofs, we omit conditioning notation $q_t \leftarrow q_{t|0,T}$ for simplicity and assume $q_t(x)s_t(x) \rightarrow 0$ vanishes at the boundary $x \rightarrow \pm\infty$, which is used when integrating by parts in x .

Corollary 1. *The Lagrangian objective in [Thm. 1](#) which solves Doob's h -transform is equivalent to*

$$\mathcal{S} = \min_q \max_s s_B(B, T) - s_B(A, 0) - \int_0^T dt \int dx q_{t|0,T} \left(\frac{\partial s_B}{\partial t} + \langle \nabla s_B, G_t \nabla s_B \rangle + \langle \nabla s_B, b_t \rangle + \langle \nabla, G_t \nabla s_B \rangle \right)$$

if $q_{t|0,T}$ satisfies [\(9c\)](#). Note $v_{t|0,T}(x) = \nabla_x s_B(x, t)$, with $s_B^*(x, t) = \log h_B(x, t)$ at optimality. ⁴

Proof. Consider the following action functional

$$\begin{aligned} \mathcal{S} &= \min_{q,v} \int dt \int dx q_t(x) \langle v_t(x), G_t v_t(x) \rangle, \\ \text{s.t. } \frac{\partial q_t(x)}{\partial t} &= -\langle \nabla_x, q_t(x)(b_t(x) + 2G_t v_t(x)) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_t(x), \\ q_0(x) &= \delta(x - A), \quad q_1(x) = \delta(x - B). \end{aligned}$$

The Lagrangian of this optimization problem is

$$\mathcal{L} = \int_0^T dt \int dx \left[q_t \langle v_t, G_t v_t \rangle + s_t \left(\frac{\partial q_t}{\partial t} + \langle \nabla, q_t(b_t + 2G_t v_t) \rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_t \right) \right],$$

where s_t is the dual variable and we omit the optimization arguments, with $\mathcal{S} = \min_{q,v} \max_s \mathcal{L}$. Swapping the order of optimizations under strong duality, we take the variation with respect to v_t in an arbitrary direction \mathfrak{h}_t . Using $G_t = G_t^T$, we obtain

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta v_t}[\mathfrak{h}_t] &= q_t \langle (G_t + G_t^T) v_t, \mathfrak{h}_t \rangle - q_t \langle 2G_t^T \nabla s_t, \mathfrak{h}_t \rangle = 0 \\ \implies v_t &= \nabla s_t, \end{aligned} \tag{24}$$

Substituting into the above, we have

$$\mathcal{L} = \int_0^T dt \int dx \left[s_t \frac{\partial q_t}{\partial t} - q_t \langle \nabla s_t, G_t \nabla s_t \rangle + s_t \langle \nabla, q_t b_t \rangle - s_t \langle \nabla, G_t \nabla q_t \rangle \right]. \tag{25}$$

Integrating by parts in t and in x , assuming that $q_t(x)s_t(x) \rightarrow 0$ as $x \rightarrow \pm\infty$, yields

$$\begin{aligned} \mathcal{L} &= \int dx q_T s_T - \int dx q_0 s_0 + \int_0^T dt \int dx \left[-q_t \frac{\partial s_t}{\partial t} - q_t \langle \nabla s_t, G_t \nabla s_t \rangle - q_t \langle \nabla s_t, b_t \rangle + \langle \nabla s_t, G_t \nabla q_t \rangle \right] \\ &= \int dx q_T s_T - \int dx q_0 s_0 + \int_0^T dt \int dx \left[-q_t \frac{\partial s_t}{\partial t} - q_t \langle \nabla s_t, G_t \nabla s_t \rangle - q_t \langle \nabla s_t, b_t \rangle - q_t \langle \nabla, G_t \nabla s_t \rangle \right] \\ &= \int dx q_T s_T - \int dx q_0 s_0 - \int_0^T dt \int dx q_t \left[\frac{\partial s_t}{\partial t} + \langle \nabla s_t, G_t \nabla s_t \rangle + \langle \nabla s_t, b_t \rangle + \langle \nabla, G_t \nabla s_t \rangle \right] \end{aligned} \tag{26}$$

where in the second line, we integrate by parts in x again. Enforcing $q_T(x) = \delta(x - B)$ and $q_0(x) = \delta(x - A)$ and recalling $\mathcal{S} = \min_q \max_s \mathcal{L}$ after eliminating v_t , we recover the optimization in the statement of the corollary. \square

Theorem. 1. *The following Lagrangian action functional has a unique solution which matches the Doob h -transform in [Prop. 2](#),*

$$\mathcal{S} = \min_{q,v} \int_0^T dt \int dx q_{t|0,T}(x) \langle v_{t|0,T}(x), G_t v_{t|0,T}(x) \rangle, \tag{27a}$$

$$\text{s.t. } \frac{\partial q_{t|0,T}(x)}{\partial t} = -\langle \nabla_x, q_{t|0,T}(x)(b_t(x) + 2G_t v_{t|0,T}(x)) \rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x), \tag{27b}$$

$$q_0(x) = \delta(x - A), \quad q_T(x) = \delta(x - B). \tag{27c}$$

⁴Compared to [\(8c\)](#), we write $\sum_{ij} (G_t)_{ij} \frac{\partial}{\partial x_i \partial x_j} s_B(x, t) = \langle \nabla, G_t \nabla s_B(x, t) \rangle$ for simplicity of notation.

Namely, the optimal $q_{t|0,T}^*(x)$ obeys (8a) and the optimal $v_{t|0,T}^*(x) = \nabla_x \log h_B(x, t) = \nabla_x s(x, t)$ follows (8b) or (8c).

Proof. The proof proceeds from (25) above,

$$S = \min_q \max_s \mathcal{L} = \min_q \max_s \int_0^T dt \int dx \left[s_t \frac{\partial q}{\partial t} - q_t \langle \nabla s_t, G_t \nabla s_t \rangle + s_t \langle \nabla, q_t b_t \rangle - s_t \langle \nabla, G_t \nabla q_t \rangle \right]. \quad (28)$$

We first show that the optimality condition with respect to s_t yields the Fokker-Planck equation for q_t in Prop. 2 (8a), before deriving the PDE in (8b) as the optimality condition with respect to q_t .

Optimality Condition for (27) recovers Prop. 2 (8a): The variation with respect to s_t of (28) is simple, apart from the intermediate term. For a perturbation direction \mathfrak{h}_t , we seek

$$\int dx \frac{\delta(\cdot)}{\delta s_t} \mathfrak{h}_t = \frac{d}{d\varepsilon} \left[- \int dx q_t \langle \nabla(s_t + \varepsilon \mathfrak{h}_t), G_t \nabla(s_t + \varepsilon \mathfrak{h}_t) \rangle \right] \Big|_{\varepsilon=0},$$

where \cdot indicates the functional on the right hand side. Proceeding to differentiate with respect to ε , we use linearity to pull $\frac{d}{d\varepsilon}$ inside the integral and apply it first to obtain $\frac{d}{d\varepsilon}(s_t + \varepsilon \mathfrak{h}_t) = \mathfrak{h}_t$. Using the product rule, recognizing the symmetry of terms, and evaluating at $\varepsilon = 0$, we are left with

$$\int dx \frac{\delta(\cdot)}{\delta s_t} \mathfrak{h}_t = \left[-2 \int dx q_t \langle \nabla \mathfrak{h}_t, G_t \nabla s_t \rangle \right] \stackrel{(i)}{=} \left[\int dx \mathfrak{h}_t \left(2 \langle \nabla, q_t G_t \nabla s_t \rangle \right) \right] \quad (29)$$

where in (i) we integrate by parts x .

We are now ready to set the variation of (28) with respect to s_t (in an arbitrary direction \mathfrak{h}_t) equal to zero. Using (29), we have

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta s_t} [\mathfrak{h}_t] = 0 &= \frac{\partial q_t}{\partial t} + 2 \langle \nabla, q_t G_t \nabla s_t \rangle + \langle \nabla, q_t b_t \rangle - \langle \nabla, G_t \nabla q_t \rangle \\ \implies 0 &= \frac{\partial q_t}{\partial t} + \langle \nabla, q_t (b_t + 2G_t \nabla s_t) \rangle - \langle \nabla, G_t \nabla q_t \rangle \end{aligned} \quad (30)$$

which matches the desired optimality condition for the conditioned process in Prop. 2 (8a).

Optimality Condition for (27) recovers Prop. 2 (8b): Starting again from (28), we take the variation with respect to q_t . First, we repeat identical steps (integrate by parts in both x and t) to reach (26),

$$\mathcal{L} = \int dx q_T s_T - \int dx q_0 s_0 - \int_0^T dt \int dx q_t \left[\frac{\partial s_t}{\partial t} + \langle \nabla s_t, G_t \nabla s_t \rangle + \langle \nabla s_t, b_t \rangle + \langle \nabla, G_t \nabla s_t \rangle \right]$$

where it is now clear that taking the variation with respect to q_t and setting equal to zero yields

$$\frac{\delta \mathcal{L}}{\delta q_t} [\mathfrak{h}_t] = 0 = \frac{\partial s_t}{\partial t} + \langle \nabla s_t, G_t \nabla s_t \rangle + \langle \nabla s_t, b_t \rangle + \langle \nabla, G_t \nabla s_t \rangle \quad (31)$$

which is the desired PDE for $s(x, t) = \log h_B(x, t)$ in (8c). To obtain (8b), we note an identity used to simplify the last term

$$\sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \log h_t = \langle \nabla, G_t \nabla \log h_t \rangle = \left\langle \nabla, \frac{1}{h_t} G_t \nabla h_t \right\rangle = -\frac{1}{h_t^2} \langle \nabla h_t, G_t \nabla h_t \rangle + \frac{1}{h_t} \langle \nabla, G_t \nabla h_t \rangle.$$

Now, substituting $s(x, t) = \log h_B(x, t)$ into Eq. (31) and abbreviating $\log h_B(\cdot, t) = \log h_t(\cdot)$, we obtain

$$\begin{aligned} \frac{1}{h_t} \frac{\partial h_t}{\partial t} + \frac{1}{h_t^2} \langle \nabla h_t, G_t \nabla h_t \rangle + \frac{1}{h_t} \langle \nabla h_t, b_t \rangle - \frac{1}{h_t^2} \langle \nabla h_t, G_t \nabla h_t \rangle + \frac{1}{h_t} \langle \nabla, G_t \nabla h_t \rangle &= 0, \\ \implies \frac{\partial h_t(x)}{\partial t} + \langle \nabla h_t(x), b_t(x) \rangle + \langle \nabla, G_t \nabla h_t \rangle &= 0, \end{aligned} \quad (32)$$

which matches (8b) as desired.

The last equation defines the backward Kolmogorov equation for the diffusion process with the drift $b_t(x)$ and covariance matrix G_t , i.e. the function $h_t(x)$ defines the conditional density $h_t(x) = p(x_T \in \mathcal{B}' | x_t = x)$ for some set \mathcal{B}' , which agrees with the forward process with the same drift and covariance. The boundary condition $q_T(x) = \delta(x - B)$ together with the backward equation define the unique solution to this PDE. Since the PDEs and the boundary conditions are the same as in Doob's h -transform, we have $h_t(x) = p(x_T = B | x_t = x)$. \square

Corollary 2. *The following SB problem*

$$\mathcal{S} := \min_{\mathbb{Q} \text{ s.t. } \mathbb{Q}=\delta, \mathbb{Q}=\delta} D_{KL}[\mathbb{Q}_{0:T}^v : \mathbb{P}_{0:T}^{\text{ref}}] \quad (11)$$

yields the path measure $\mathbb{P}_{0:T}^*$ associated with the SDE in (6) as its unique minimizing argument. The temporal marginals of $\mathbb{P}_{0:T}^*$ are equal to those which optimize the Lagrangian objective in [Thm. 1](#).

Proof. We use the Girsanov theorem ([Särkkä and Solin, 2019](#), Sec. 7.3) to calculate the KL divergence between the following two Brownian diffusions with fixed initial condition $x_0 = A$,

$$\mathbb{P}_{0:T}^{\text{ref}} : \quad dx_t = b_t(x_t) \cdot dt + \Xi_t dW_t, \quad (33)$$

$$\mathbb{Q}_{0:T}^v : \quad dx_t = (b_t(x_t) + 2G_t v_{t|0,T}(x_{t|0,T})) \cdot dt + \Xi_t dW_t, \quad (34)$$

In particular, noting the difference of drifts is $b_t(x_t) + 2G_t v_{t|0,T}(x_t) - b_t(x_t) = 2G_t v_{t|0,T}(x_t)$, the likelihood ratio is given by

$$\frac{d\mathbb{Q}_{0:T}^v}{d\mathbb{P}_{0:T}^{\text{ref}}} = \frac{q_{t|0,T}(x_0, \dots, x_T)}{\rho(x_0, \dots, x_T)} = \exp \left\{ -\frac{1}{2} \int_0^T \langle 2G_t v_{t|0,T}(x_t), (G_t)^{-1} 2G_t v_{t|0,T}(x_t) \rangle dt \right. \\ \left. - \int 2(G_t v_{t|0,T}(x_t))^T G_t^{-1} dW_t \right\} \quad (35)$$

We finally calculate the KL divergence, noting that, after taking the log, the expectation of the integral $\int (\cdot) dW_t$ in the final term vanishes,

$$D_{KL}[\mathbb{Q}_{0:T}^v : \mathbb{P}_{0:T}^{\text{ref}}] = 2 \int_0^T dt \int dx_t q_{t|0,T}(x_t) \langle v_{t|0,T}(x_t), G_t v_{t|0,T}(x_t) \rangle, \quad (36)$$

which matches (9a) up to a constant factor of 2 does not change the optimum. We finally compare to the constraints in [Thm. 1](#). First, it is clear that the diffusion in (34) satisfies the Fokker-Planck equation in (9b) ([Särkkä and Solin, 2019](#), Sec. 5.2). We respect (9c) by optimizing over endpoint-constrained path measures, which yields

$$\mathcal{S} = \min_{\mathbb{Q} \text{ s.t. } \mathbb{Q}=\delta, \mathbb{Q}=\delta} D_{KL}[\mathbb{Q}_{0:T}^v : \mathbb{P}_{0:T}^{\text{ref}}] \quad (37)$$

as desired. \square

B Gaussian Path Parameterizations

Proposition 3. *For the family of endpoint-conditioned marginals $q_{t|0,T}(x) = \mathcal{N}(x | \mu_{t|0,T}, \Sigma_{t|0,T})$,*

$$u_{t|0,T}^{(q,\theta)}(x) := \frac{\partial \mu_{t|0,T}}{\partial t} + \left[\frac{1}{2} \frac{\partial \Sigma_{t|0,T}}{\partial t} \Sigma_{t|0,T}^{-1} - G_t \Sigma_{t|0,T}^{-1} \right] (x - \mu_{t|0,T}) \quad (13)$$

satisfies the Fokker-Planck equation (12) for $q_{t|0,T}$ and diffusion coefficients $G_t = \frac{1}{2} \Xi_t \Xi_t^T$.

Proof. Consider the following identities for the Gaussian family of marginals $q_t(x) = \mathcal{N}(x | \mu_t, \Sigma_t)$, where we omit conditioning $q_t \leftarrow q_{t|0,T}$ for simplicity of notation,

$$\log q_t(x) = -\frac{1}{2} (x - \mu_t)^T \Sigma_t^{-1} (x - \mu_t) - \frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det \Sigma_t, \quad (38a)$$

$$\nabla_x \log q_t(x) = -\Sigma_t^{-1} (x - \mu_t), \quad (38b)$$

$$\frac{\partial}{\partial t} \log q_t(x) = (x - \mu_t)^T \Sigma_t^{-1} \frac{\partial \mu_t}{\partial t} + \frac{1}{2} (x - \mu_t)^T \Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial t} \Sigma_t^{-1} (x - \mu_t) - \frac{1}{2} \text{tr} \left(\Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial t} \right) \quad (38c)$$

We begin by solving for a vector field $u_t^o(x)$ that satisfies the continuity equation (where u_t^o denotes the drift of an ODE)

$$\frac{\partial q_t}{\partial t} = -\langle \nabla_x, q_t u_t^o \rangle = -q_t \langle \nabla_x, u_t^o \rangle + \langle \nabla_x q_t, \nabla_x u_t^o \rangle \\ \implies \frac{\partial}{\partial t} \log q_t = -\langle \nabla_x, u_t^o \rangle - \langle \nabla_x \log q_t, u_t^o \rangle \quad (39)$$

The vector field satisfying this equation is

$$u_t^o(x) = \frac{\partial \mu_t}{\partial t} + \frac{1}{2} \frac{\partial \Sigma_t}{\partial t} \Sigma_t^{-1} (x - \mu_t) \quad (40)$$

which we can confirm using the identities in (38). Indeed, for the terms on the RHS of Eq. (39),

$$\begin{aligned} -\langle \nabla_x, u_t^o \rangle &= -\frac{1}{2} \text{tr} \left(\Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial t} \right), \\ -\langle \nabla_x \log q_t, u_t^o \rangle &= \left\langle \Sigma_t^{-1} (x - \mu_t), \frac{\partial \mu_t}{\partial t} \right\rangle + \frac{1}{2} (x - \mu_t)^T \Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial t} \Sigma_t^{-1} (x - \mu_t). \end{aligned}$$

Putting these terms and the time derivative from (38c) into Eq. (39) we conclude the proof.

However, we are eventually interested in finding the formula for the drift u_t that satisfies the Fokker-Planck equation in (12). That is, to describe the same evolution of density $\frac{\partial q(x)}{\partial t}$, the relationship between u_t and u_t^o is as follows

$$\begin{aligned} \frac{\partial q_t(x)}{\partial t} &= -\langle \nabla_x, q_t u_t^o \rangle = -\langle \nabla_x, q_t u_t \rangle + \langle \nabla_x, G_t \nabla_x q_t \rangle \\ &= -\langle \nabla_x, q_t u_t \rangle + \langle \nabla_x, G_t q_t \nabla_x \log q_t \rangle \\ &= -\left\langle \nabla_x, q_t \underbrace{(u_t - G_t \nabla_x \log q_t)}_{u^o} \right\rangle \end{aligned}$$

Finally, we use the identities in (38) to obtain

$$\begin{aligned} u_t &= u_t^o + G_t \nabla_x \log q_t = \frac{\partial \mu_t}{\partial t} + \frac{1}{2} \frac{\partial \Sigma_t}{\partial t} \Sigma_t^{-1} (x - \mu_t) - G_t \Sigma_t^{-1} (x - \mu_t) \\ \implies u_t &= \frac{\partial \mu_t}{\partial t} + \left[\frac{1}{2} \frac{\partial \Sigma_t}{\partial t} \Sigma_t^{-1} - G_t \Sigma_t^{-1} \right] (x - \mu_t) \end{aligned}$$

□

Proposition 4. Given a set of processes $q_{t|0,T}^k(x)$ and mixtures weights w^k , the vector field satisfying the Fokker-Planck equation in (12) for the mixture $q_{t|0,T}(x) = \sum_k w^k q_{t|0,T}^k(x)$ is given by

$$u_{t|0,T}^{(q,\theta)}(x) = \sum_{k=1}^K \frac{w^k q_{t|0,T}^k(x)}{\sum_{j=1}^K w^j q_{t|0,T}^j(x)} u_{t|0,T}^{(q,k)}(x), \quad (17)$$

where $u_{t|0,T}^{(q,k)}(x)$ satisfies the Fokker-Planck equation in (12) for $q_{t|0,T}^k(x)$. This identity holds for both first-order dynamics in spatial coordinates only or second-order dynamics in $x = (\bar{x}, \bar{v})$.

Proof. See Peluchetti (2023) Theorem 1 and its proof in their App. A. □

C Extended Related Work

C.1 Machine Learning for Molecular Simulation

The main dilemma of molecular dynamics comes from the accuracy and efficiency trade-off—accurate simulation requires solving the Schrödinger equation which is computationally intractable for large systems, while efficient simulation relies on empirical force fields which is inaccurate. Recently, there has been a surge of work in applying machine learning approaches to accelerate molecular simulation. One successful paradigm is machine learning force field (MLFF) which leverages the transferability and efficiency of machine learning methods to fit force/energy prediction models on quantum mechanical data and transfer across different atomic systems Smith et al. (2017); Wang et al. (2018). More recently, increasing attention has been focused on building atomic foundation models to encompass all types of molecular structures Batatia et al. (2023); Shoghi et al. (2023); Zhang et al. (2022).

Sampling is a classical problem in molecular dynamics to draw samples from the Boltzmann distribution of molecular systems. Classical methods mainly rely on Markov chain Monte Carlo (MCMC) or MD which requires long mixing time for multimodal distributions with high energy barriers [Rotskoff \(2024\)](#). Generative models in machine learning demonstrate promises in alleviating this problem by learning to draw independent samples from the Boltzmann distribution of molecular systems (known as Boltzmann generator) [Noé et al. \(2019\)](#). Numerous methods have been developed to utilize generative models as a proposal distribution for escaping local minima in running MCMC methods [Gabrié et al. \(2022\)](#). However, one critical issue is that generative models rely on training from samples. Although recent advances have been developed to learn from unnormalized density (i.e., energy) function, the training inefficiency limits their applicability to solve high-dimensional molecular dynamics problems. To circumvent the curse of dimensionality for the sampling problem, another branch of work study to learn coarse-grained representation with neural networks [Sidky et al. \(2020\)](#). For broader literature of applying machine learning to enhanced sampling, we refer the reader to [Mehdi et al. \(2024\)](#).

D Further Experimental Details

D.1 Evaluation Metrics

To assess the quality of our approach in terms of performance and physicalness of paths, we compare them under different metrics to well-established TPS techniques. One important describing factor of a trajectory is the molecule’s highest energy during the transition. These high-energy states are often referred as transition states and less likely to occur but they determine importance factors during chemical reaction such as reaction rate. As such, we will look at the maximum energy along the transition path and use it to compare the ensemble of trajectories more efficiently. The main goal is to show that lower energy of the transition states can be sampled by the methods.

However, the maximum energy does not account for the fact that the transition path needs to be sequential, and each step needs to be coherent based on the previous position and momentum. For this, we also compare the likelihood of the paths (i.e., unnormalized density) by computing the probably of being in the start state $\rho(x_0)$ and multiplying it with the step probability such that

$$L(x_0, x_1, \dots, x_{N-1}) = \rho(x_0) \cdot \prod_{i=0}^{N-2} \pi(x_{i+1} | x_i). \quad (41)$$

For the step probability π , we solve the Langevin leap-frog implementation as implemented in OpenMM to solve $\mathcal{N}(x_{i+1} | x_i + dt \cdot b_t(x), dt\sigma_i^2)$. As for the starting probability, we compute the unnormalized density of the Boltzmann distribution for our start state z and assume that the velocity v can be sampled independently ([Castellan, 1983](#), Sec. 4.6)

$$\rho(z, v) \propto \exp\left(-\frac{U(z)}{k_B T}\right) \cdot \mathcal{N}(v | 0, k_B T \cdot M^{-1}), \quad (42)$$

with the Boltzmann constant k_B and the diagonal matrix M containing the mass of each atom.

As for the performance, the number of energy evaluations will be the main determining factor of the runtime for larger molecular systems, especially for proteins. We hence compare the use of the number of energy computations as a proxy for hardware-independent relative measurements. In our tests, this number aligned with the relative runtime of these approaches.

D.2 Toy Potentials

The toy systems move according to the following integration scheme (first-order Euler)

$$x_{t+1} = x_t - dt \cdot \nabla_x U(x_t) + \sqrt{dt} \cdot \text{diag}(\xi) \cdot \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1), \quad (43)$$

following the definition of our stochastic system in [Sec. 2.2](#) with a time-independent Wiener process, where ξ is a constant time-independent standard deviation for all dimensions.

Müller-Brown. The underlying Müller-Brown potential that has been used for our experiments can be written as

$$\begin{aligned}
 U(x, y) = & -200 \cdot \exp(-(x-1)^2 - 10y^2) \\
 & -100 \cdot \exp(-x^2 - 10 \cdot (y-0.5)^2) \\
 & -170 \cdot \exp(-6.5 \cdot (0.5+x)^2 + 11 \cdot (x+0.5) \cdot (y-1.5) - 6.5 \cdot (y-1.5)^2) \\
 & +15 \cdot \exp(0.7 \cdot (1+x)^2 + 0.6 \cdot (x+1) \cdot (y-1) + 0.7 \cdot (y-1)^2).
 \end{aligned} \tag{44}$$

We used a first-order Euler integration scheme to simulate transition paths with 275 steps and a dt of $10^{-4}s$. ξ was chosen to be 5 and 1,000 transition paths were simulated. We have used an MLP with four layers and a hidden dimension of 128 each, with swish activations. It has been trained for 2,500 steps with a batch size of 512 and a single Gaussian.

In Fig. 5a, we compare the likelihood of the sampled paths. We can see that one-way shooting takes time until the path is decorrelated from the initial trajectory, which is shorter and thus has a higher likelihood. All MCMC methods exhibit this behavior, which is typically alleviated by using a warmup period in which all paths are discarded. After that, all methods exhibit similar likelihood, with our method having a slightly lower likelihood. Looking at the transition state (i.e., maximum energy on the trajectory) in Fig. 5b reveals that all methods have a similar quality of paths.

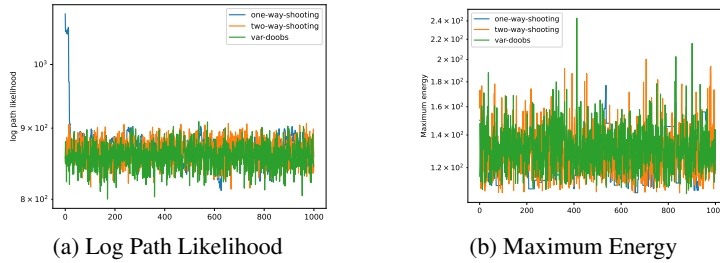


Figure 5: In a, we compare the log likelihood of sampled trajectories, where a higher likelihood is generally more favorable. The plot in b shows the maximum energy of each individual trajectory. A high maximum energy means that the molecule needs to be in an excited state during the transition, making it less likely to occur under lower temperatures.

We can further analyze the quality of our method by investigating the difference between the “ground truth” marginal $\rho_{t|0,T}(x)$ and the learned marginal $q_{t|0,T}(x)$. For this, we compute the Wasserstein W1 distance (Flamary et al., 2021) between the marginal observed by fixed-length two-way shooting (which we assume to be close to the ground truth) and our variational approach. We observe a mean W1 distance of 0.130 ± 0.026 and visualize it along the time coordinate t (in steps) in Fig. 6.




Figure 6: In this figure, we compare the Wasserstein W1 distance between the marginals. The densities are almost identical at the beginning and end state. Note that at the third local minimum of the Müller-Brown potential along the trajectory (i.e., reached at about 200 steps), the marginals align more closely as well.

Dual-Channel Double-Well. To demonstrate the advantage of mixtures, we have used the two-dimensional potential

$$\begin{aligned}
 U(x, y) = & +2 \cdot \exp(-(12x^2 + 12y^2)) \\
 & -1 \cdot \exp(-(12 \cdot (x+0.5)^2 + 12y^2)) \\
 & -1 \cdot \exp(-(12 \cdot (x-0.5)^2 + 12y^2)) + x^6 + y^6.
 \end{aligned} \tag{45}$$

In this case, we have used $dt = 5 * 10^{-4}s$ with a transition time of $T = 1s$ and $\xi = 0.1$. As for the MLP, we have used the same structure as in the Müller-Brown example but trained it for 20,000 iterations. The corresponding weights to Prop. 4 are $w = [\frac{1}{2}, \frac{1}{2}]$ and are fixed for this experiment and hence $w \notin \theta$.

D.3 Neural Network Ablation Study

In Fig. 7 we compare how different parameterizations of $\mu_{t|0,T}^{(\theta)}$, and $\Sigma_{t|0,T}^{(\theta)}$ impact the quality of trajectories on the Müller-Brown potential. For this, we compare linear and cubic splines (with 20 knots) with neural networks. As a metric to estimate the quality, we compare the Wasserstein W1 distance between the learned marginal $q_{t|0,T}(x)$ and the marginal observed by fixed-length two-way shooting (i.e., baseline). We notice that using linear splines results in the highest W1 distance, while cubic splines improve the quality. Using neural networks, however, yields the best approximation.

We have fixed the computational budget for all systems, which means that we have trained splines for more epochs than the neural network (since they are slower to train). For high-dimensional systems, the runtime is mostly determined by the number of potential evaluations and not the complexity of the architecture. We thus conclude that the additional expressivity provided by neural networks is necessary for more complicated (molecular) systems and does not introduce much computational overhead.




Figure 7: We compare the Wasserstein W1 distance between the marginals with different parameterization techniques of $\mu_{t|0,T}^{(\theta)}$, and $\Sigma_{t|0,T}^{(\theta)}$.

D.4 Molecular Systems

To simulate molecular dynamics, we rely on the AMBER14 force field (amber14/protein.ff14SB Maier et al. (2015)) without a solvent, as implemented in OpenMM (Eastman et al., 2017). As OpenMM does not support auto-differentiation, we do not use OpenMM for the simulations themselves, but utilize DMFF (Wang et al., 2023) which is a differentiable framework implemented in JAX (Bradbury et al., 2018) for molecular simulation. This is needed because during training we compute $\nabla_{\theta} U(x_{t|0,T} \sim \mathcal{N}(\mu_{t|0,T}^{(\theta)}, \Sigma_{t|0,T}^{(\theta)}))$, where the concrete $x_{t|0,T}$ is sampled based on the parameters of the neural network.

For the concrete simulations, we ran them with the timestep $dt = 1fs$, with $T = 1ps$, $\gamma = 1ps$, and Temp = 300K. To compute the MCMC two-way shooting baselines, we use the same settings and consider trajectories as failed, if they exceed 2,000 steps without reaching the target.

Neural Network Parameterization. We parameterize our model with neural networks, a 5-layer MLP with ReLU activation function and 256/512 hidden units for alanine dipeptide and Chignolin, respectively. The neural networks are trained using an Adam optimizer with learning rate 10^{-4} .

We represent the molecular system in two ways: (1) in Cartesian coordinates, which are the 3D coordinates of each atoms, and with (2) internal coordinate which instead uses bond length, angle and dihedral angle along the molecule, where we use the same parameterization as in (Noé et al., 2019).

Our state definition includes a variance parameter for the initial and target marginal distributions at $t = 0$ and $t = T$, we choose the variance to be 10^{-8} which almost does not change the energy of the perturbed system.

Visualization of Transition for Alanine Dipeptide. In Fig. 8, we show a transition sampled without any noise from the model with internal coordinates and 2 Gaussian mixtures.




Figure 8: Transition path for the alanine dipeptide.

Comparison of Sampled Paths During Training. In our training procedure, the marginal starts with a linear interpolation between A and B , which produces very unlikely paths with potentially high energy states. In Fig. 9, we compare how the quality of paths changes depending on the number of training iterations (i.e., the number of potential evaluations). We show the curve for a single Gaussian mixture with Cartesian coordinates. Similar trends can be observed in other settings.




Figure 9: In this figure, we compare the quality of paths based on the current training step (i.e., potential evaluations). We observe that with increasing training time, paths with higher likelihood are sampled.

Loss Curves. In this section, we would like to investigate the training losses of different configurations. For this, we plot the exponential moving average of the loss ($\alpha = 0.001$) to better highlight the trends of the noisy variational loss. Fig. 10 compares the results of different training settings. We can observe that mixtures can decrease the overall loss, but all model variations converge to a similar loss value.




Figure 10: Visualization of the loss for different training setups. These setups are identical to what has been reported in Table 2.

D.5 Computational Resources

All our experiments involving training were conducted on a single NVIDIA A100 80GB. The baselines themselves were computed on a M3 Pro 12-core CPU.

E Societal Impact

Our research concerns the efficient sampling of transition paths which are crucial for a variety of tasks in biology, chemistry, materials science and engineering. Our research could potentially benefit research areas from combustion, catalysis, protein design to battery design. Nevertheless, we do not foresee special potential negative impacts to be discussed here.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The major claims made in the abstract and introduction are backed up by the theoretical and empirical results in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: In [Sec. 6](#), we discuss potential limitations and future work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: A detailed set of assumptions is presented in [Sec. 3](#) and further proofs for theoretical results are included in [App. A](#).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We disclose all the information needed to reproduce the main experimental results and make our code available at <https://github.com/plainerman/variational-doob>.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The data we use are synthetic datasets that can be generated by publicly available codes and real-world molecular systems taken from previous literature cited in the paper. Our code is publicly available at <https://github.com/plainerman/variational-doob>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We report the experimental details in [Sec. 5](#) and further specify the parameter and settings in [App. D](#).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We appropriately report the mean and standard deviation in the experiments where suitable (compare [Table 1](#) and [Table 2](#)).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We include computational resources used to run the model in [App. D.5](#).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Yes, we believe our research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the potential societal impact of our method in [App. E](#).

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not believe that any safe guards needs to be in place, as our approach does not pose any direct risks. As for the data, only publicly available and commonly used resources have been used and, hence, does not pose a risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use public datasets and libraries in this research and have cited the sources.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines: No new assets are introduced in this research.

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing or human subjects are involved in this research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No human subjects are involved in this research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.