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

Yuanqi Du [*1] **Michael Plainer** [*2] **Rob Brekelmans** [*3] **Chenru Duan** [4] **Frank Noe** [5 6 7] **Carla P. Gomes** [1]
Alan Apsuru-Guzik [3 8] **Kirill Neklyudov** [9 10]

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 simulation of this transform is infeasible, as it requires sufficiently many forward trajectories to estimate rare event probabilities. In this work, we propose a variational formulation of Doob's h-transform — 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., 2024), drug discovery (Kirmizialtin et al., 2012; 2015; Dickson, 2018), and materials science (Xi et al., 2013; Selli 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). 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 draw explicit connections to the Doob's h-transform 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; Holdijk et al., 2024) 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 as an inner loop 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

[*]Equal contribution [1]Cornell University [2]Technische Universität Berlin [3]Vector Institute [4]Massachusetts Institute of Technology [5]Freie Universität Berlin [6]Rice University [7]Microsoft Research AI4Science [8]University of Toronto [9]Université de Montréal [10]MILA Quebec AI Institute. Correspondence to: <yuanqidu@cs.cornell.edu>, <michael.plainer@outlook.com>, <k.necludov@gmail.com>.

Accepted by the Structured Probabilistic Inference & Generative Modeling workshop of ICML 2024, Vienna, Austria. Copyright 2024 by the author(s).

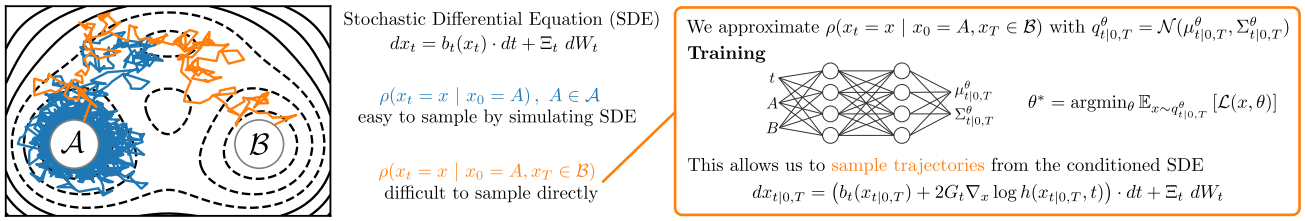

Figure 1: Given reference dynamics, transition path sampling seeks to capture the conditional or posterior distribution over paths which reach a terminal set $x_T \in \mathcal{B}$. However, simulating the reference dynamics (blue) can be wasteful since we rarely obtain paths (orange) which reach (the 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.

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 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 transition probability $\rho(x_{t+dt} = y \,|\, x_t = x)$. Starting from an initial $x_0 = A$, the density of the path is

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

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*, i.e. the posterior distribution over intermediate states

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

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,

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

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

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

where W_t denotes the Wiener process. Note, U is a potential energy, $k_B \mathcal{T}$ 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 (1)–(2), we consider the forward or reference stochastic differential equation (SDE),

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

Proposition 1. *The following PDEs are obeyed by (a) the marginal density of the conditioned process $\rho_{t|0,T}(x) = \rho(x_t = x \,|\, x_0 = A, x_T \in \mathcal{B})$ and (b) the h-function $h(x,t)$ (which implicitly depends on $\mathcal{B}$),*

$$\frac{\partial \rho_{t|0,T}(x)}{\partial t} + \big\langle \nabla_x, \rho_{t|0,T}(x)\big(b_t(x) + 2G_t \nabla_x \log h(x,t)\big)\big\rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \rho_{t|0,T}(x) = 0\,, \tag{3a}$$

$$\frac{\partial h(x,t)}{\partial t} + \big\langle \nabla_x h(x,t), b_t(x)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} h(x,t) = 0\,. \tag{3b}$$

Reparameterizing (3b) in terms of $s(x,t) \coloneqq \log h(x,t)$, we can also write

$$\frac{\partial s(x,t)}{\partial t} + \big\langle \nabla s(x,t), G_t \nabla s(x,t)\big\rangle + \big\langle \nabla s(x,t), b_t(x_t)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s(x,t) = 0. \tag{3c}$$

with drift $b_t : \mathbb{R}^N \to \mathbb{R}^N$ and diffusion matrix $\Xi_t \in \mathbb{R}^{N\times N}$ such that $G_t \coloneqq \frac{1}{2}\Xi_t \Xi_t^T$ is positive definite.[1] We denote the induced path measure as $\mathbb{P}^{\text{ref}}_{0:T} \in \mathcal{P}(\mathcal{C}([0,T] \to \mathbb{R}^N))$.

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

Proposition 2. [Jamison (1975, Thm. 2)] *Let $h(x,t) \coloneqq \rho(x_T \in \mathcal{B} \,|\, x_t = x)$ denote the conditional transition density with respect to the reference process in (5). Letting $G_t \coloneqq \frac{1}{2}\Xi_t\Xi_t^T$, the SDE*

$$dx_{t|0,T} = \Big(b_t(x_{t|0,T}) + 2G_t \nabla_x \log h(x_{t|0,T}, t)\Big)dt + \Xi_t\, dW_t \tag{6}$$

is associated with the following transition probabilities

$$\rho(x_t = y \,|\, x_s = x, x_T \in \mathcal{B}) = \frac{h(y,s)}{h(x,t)} \rho(x_t = y \,|\, x_s = x), \tag{7}$$

for $s < t < T$, where we omit the dependence of $h(x,t)$ on $\mathcal{B}$ for simplicity of notation.

See App. A.1 for proof, and note that (7) is simply an application of Bayes rule $\rho(x_t = y \,|\, x_s = x, x_T \in \mathcal{B}) = \big(\rho(x_T \in \mathcal{B} | x_s = y)/\rho(x_T \in \mathcal{B} | x_t = x)\big) \cdot \rho(x_t = y \,|\, x_s = x)$, with the unconditioned reference transition density as the prior. Furthermore, the conditioned transition probabilities in (7) allow us to directly construct the transition path in Sec. 2.1. Using Bayes rule, we have

$$\begin{aligned}\rho(x_{T-dt}, \ldots, x_{dt} \,|\, x_0 = A, x_T \in \mathcal{B}) \\ = \frac{h(x_{T-dt}, T-dt)}{h(A,0)} \rho(x_{T-dt} \ldots, x_{dt} \,|\, x_0 = A)\end{aligned} \tag{8}$$

after telescoping cancellation of h-functions and rewriting the denominator as $h(A,0) = \rho(x_T \in \mathcal{B} \,|\, x_0 = A)$. Thus, we can solve the TPS problem by exactly solving for the h-function and simulating the SDE in (6).

Finally, the h-function and marginal density of the conditioned process satisfy the forward and backward Kolmogorov equations in Prop. 1, which will be useful when deriving our variational objectives in the next section. Note that we use $\langle \nabla_x, \bullet \rangle = \text{div}(\bullet)$ for the divergence operator. See App. A.1 for proof of Prop. 1.

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.

[1] See (16) in Sec. 3.2.2 to write (2) in the form of (5)

Theorem 1. *The following Lagrangian action minimization has a unique solution which matches Doob's h-transform in Prop. 2, where the optimal $q^*_{t|0,T}(x)$ and $v^*_{t|0,T}(x) = \nabla_x \log h(x,t)$ satisfy the PDEs in Prop. 1,*

$$\mathcal{S} = \min_{q,v} \int_0^T dt \int dx\; q_{t|0,T}(x) \big\langle v_{t|0,T}(x), G_t\, v_{t|0,T}(x)\big\rangle\,, \tag{4a}$$

$$\text{s.t.}\quad \frac{\partial q_{t|0,T}(x)}{\partial t} = -\big\langle \nabla_x, q_{t|0,T}(x)\big(b_t(x) + 2G_t\, v_{t|0,T}(x)\big)\big\rangle + \sum\nolimits_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x), \tag{4b}$$

$$q_0(x) = \delta(x - A), \qquad q_T(x) = \delta(x - B)\,. \tag{4c}$$

3.1. Doob's Lagrangian

Consider reference dynamics given in the form of either (1) or (2), 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 Thm. 1 below, we define a Lagrangian action whose minimization yields the optimal $q^*_{t|0,T}(x) = \rho_{t|0,T}(x)$ and $v^*_{t|0,T}(x) = \nabla_x \log h(x,t)$ from Prop. 2 and A.1.

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 parameterization in Sec. 3.2.

Unconstrained Dual Objective. Introducing Lagrange multipliers to enforce the constraints (4b)-(4c) and eliminating $v_{t|0,T}$, we obtain an alternative unconstrained version of (4).

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

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

if $q_{t|0,T}$ satisfies (4c). Note $v_{t|0,T}(x) = \nabla_x s(x,t)$, with $s^*(x,t) = \log h(x,t)$ at optimality.[2]

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 (3c) of App. A.1.

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

$$\mathbb{Q}^v_{0:T} : \; dx_t = \big(b_t(x_t) + 2G_t \, v_{t|0,T}(x_{t|0,T})\big) \cdot dt + \Xi_t \, dW_t \,, \tag{9}$$

with $x_0 = A$. Note that the density $q_{t|0,T}$ of $\mathbb{Q}^v_{0:T}$ satisfies the Fokker-Planck equation in (4b) (Särkkä and Solin, 2019, Sec. 5.2) Using the Girsanov Theorem, the objective in (4a) can then be viewed as a KL divergence minimization over path measures $\mathbb{Q}^v_{0:T}$ which satisfy the boundary constraints.

[2] Again, note that we omit the dependence of $s(x,t)$ and $h(x,t)$ on the conditioning information B.

Corollary 2. *The Lagrangian objective in Thm. 1 is equivalent to the following optimization of $\mathbb{Q}^v_{0:T}$*

$$\mathcal{S} := \min_{\mathbb{Q}^v_{0:T} \text{ s.t. } \mathbb{Q}^v_0 = \delta_A, \mathbb{Q}^v_T = \delta_B} D_{KL}[\mathbb{Q}^v_{0:T} : \mathbb{P}^{ref}_{0:T}] \tag{10}$$

where the minimizing argument recovers the path measure $\mathbb{P}^*_{0:T}$ associated with the SDE in (6).

Our Lagrangian action minimization thus solves a Schrödinger Bridge (SB) problem (Schrödinger, 1932; Léonard, 2014) with Dirac delta functions as the endpoint measures. Our objective in (4a) particularly resembles optimal control formulations of SB (Chen et al., 2016; 2021b). While it is well-known that the Doob h-transforms (and large deviations more generally) play a role in the solution to SB problems (e.g. 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 (4a). 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 (10), indicating that we need to be able to efficiently sample from the conditioned process in (9) or $q_{t|0,T}$ in (4). 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 (4b) 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 $\mathbb{Q}^v_{0:T}$ 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 (9) (Challenge 1), provide *analytic* solutions for $v_{t|0,T}$ satisfying (4b) 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 (1) in

Sec. 3.2.1, before presenting extensions to mixture paths and the second-order setting (2) in Sec. 3.2.3–3.2.2.

3.2.1. First-Order Dynamics and General Approach

Tractable Drift $v_{t|0,T}$ for Variational Doob Objective. Consider modifying the Fokker-Planck constraint in (4b), to absorb all drift terms into a single vector field $u_{t|0,T}$,

$$\frac{\partial q_{t|0,T}(x)}{\partial t} = -\langle \nabla_x, q_{t|0,T}(x)\, u_{t|0,T}(x)\rangle + \sum_{ij}(G_t)_{ij}\frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x). \tag{11}$$

For arbitrary $q_{t|0,T}$, solving for *any* $u_{t|0,T}(x)$ satisfying (11) 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 $u^{(q)}_{t|0,T}$ which satisfies (11). 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.

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

$$u^{(q)}_{t|0,T}(x) := \frac{\partial \mu_{t|0,T}}{\partial t} + \left[\frac{1}{2}\frac{\partial \Sigma_{t|0,T}}{\partial t}\Sigma^{-1}_{t|0,T} - G_t\, \Sigma^{-1}_{t|0,T}\right](x - \mu_{t|0,T}) \tag{12}$$

satisfies the Fokker-Planck equation (11) for $q_{t|0,T}$ and diffusion coefficients $G_t = \frac{1}{2}\Xi_t\Xi_t^T$.

Given $u^{(q)}_{t|0,T}$ corresponding to $q_{t|0,T}$, we can simply solve for the $v_{t|0,T}$ satisfying the Fokker-Planck equation in (4b) in our variational Doob objective (Thm. 1). Since G_t was assumed to be invertible and the base drift b_t is known,

$$v^{(q)}_{t|0,T}(x) = \frac{1}{2}(G_t)^{-1}\Big(u^{(q)}_{t|0,T}(x) - b_t(x)\Big), \tag{13}$$

We may now evaluate terms involving $v_{t|0,T}$ in our Lagrangian objective in (4) using (13) 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 $v^{(q)}_{t|0,T}$ for a given $q_{t|0,T} \in \mathcal{Q}_G$ as above, our variational Doob objective in (4a) reduces to a single optimization over the marginals $q_{t|0,T}$ of a conditioned process which satisfies the boundary conditions (4c).

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 $\Sigma_{t|0,T} = \texttt{diag}(\{\sigma^2_{t|0,T,d}\}_{d=1}^D)$. We parameterize a neural network $\text{NNET}_\theta : [0,T]\times\mathbb{R}^D\times\mathbb{R}^D \to \mathbb{R}^D\times\mathbb{R}^D$ 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 for $t = \frac{t}{T}$

$$x_{t|0,T} = \mu^{(\theta)}_{t|0,T} + \Sigma^{(\theta)}_{t|0,T}\,\epsilon, \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \mathbb{I}_D). \tag{14}$$
$$\mu^{(\theta)}_{t|0,T} = \left(1 - \frac{t}{T}\right)A + \frac{t}{T}\,B + \frac{t}{T}\left(1 - \frac{t}{T}\right)\text{NNET}_\theta(t, A, B)_{[:D]}$$
$$\Sigma^{(\theta)}_{t|0,T} = \frac{t}{T}\left(1 - \frac{t}{T}\right)\texttt{diag}\big(\text{NNET}_\theta(t, A, B)_{[D:]}\big) + \sigma^2_{\text{min}}\mathbb{I}.$$

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 (4b) *without simulating* the corresponding SDE with nonlinear drift (9). Further, the $t(1-t)$ coefficients in (14) ensure that our parameterization satisfies the (smoothed) boundary conditions by design (Challenge 3). Although we add σ^2_{min} to ensure invertibilty of $\Sigma_{t|0,T}$ (see (12)) as $t \to 0$ or $t \to T$, we preserve $q_0(x_0) = \mathcal{N}(x_0 \,|\, A, \sigma^2_{\text{min}}\mathbb{I}_D) \approx \delta(x_0 - A)$ and $q_T(x_T) = \mathcal{N}(x_T \,|\, B, \sigma^2_{\text{min}}\mathbb{I}_D) \approx \delta(x_T - B)$.

Reparameterization Gradients. Since we have now shown that our parameterization satisfies the constraints (4b)-(4c) 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 (4a),

$$\nabla_\theta \mathbb{E}_{q^{(\theta)}_{t|0,T}(x)}\left[\left\langle v^{(q,\theta)}_{t|0,T}(x), G_t\, v^{(q,\theta)}_{t|0,T}(x)\right\rangle\right] \tag{15}$$
$$= \mathbb{E}_{\mathcal{N}(\epsilon|0,\mathbb{I}_D)}\left[\nabla_\theta\left\langle v^{(q,\theta)}_{t|0,T}\big((g(t,\epsilon;\theta)\big), G_t\, v^{(q,\theta)}_{t|0,T}\big((g(t,\epsilon;\theta)\big)\right\rangle\right]$$

where $x = g(t, \epsilon; \theta)$ is the parameterized map in (14) and $v^{(q,\theta)}_{t|0,T}$ depends on θ via $\mu^{(\theta)}_{t|0,T}$, $\Sigma^{(\theta)}_{t|0,T}$ in (12)–(13). In practice, we approximate gradients using a single sample of ϵ at uniformly sampled discrete time points $0 \le t \le T$ which represent physical time (e.g., femtoseconds).

Refer to App. C for a complete algorithm of this procedure.

3.2.2. Second-Order Dynamics

To handle the case of the second-order dynamics in (2), 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 (2) 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

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

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 $\text{NNET}_\theta : [0,T] \times \mathbb{R}^{2D} \times \mathbb{R}^{2D} \to \mathbb{R}^{2D} \times \mathbb{R}^{2D}$, which outputs mean perturbations and per-dimension variances to calculate $\mu^{\bar{x}}_{t|0,T}, \mu^{\bar{v}}_{t|0,T}$ and $\Sigma^{\bar{x}}_{t|0,T}, \Sigma^{\bar{v}}_{t|0,T}$ and sample $(\bar{x}, \bar{v})$, as in (14). Note that we parameterize $\Sigma^{\bar{x}}_{t|0,T}, \Sigma^{\bar{v}}_{t|0,T}$ separately, matching the block diagonal form of (16). We calculate $v^{(q)}_{t|0,T}(\bar{x}, \bar{v}) := [v^{\bar{x}(q)}_{t|0,T}, v^{\bar{v}(q)}_{t|0,T}]$ from $u^{(q)}_{t|0,T}(\bar{x}, \bar{v}) = [u^{\bar{x}(q)}_{t|0,T}, u^{\bar{v}(q)}_{t|0,T}]$ as in (12)–(13), with $G_t^{-1} = (\frac{1}{2}\Xi_t \Xi_t^T)^{-1}$ given by (16). The Lagrangian objective in (4) minimizes the norm of the concatenated vector $v^{(q)}_{t|0,T}(\bar{x}, \bar{v})$, 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 marginals which follow the unimodal Gaussian distributions in the previous sections. To increase the expressivity of our variational family of conditioned processes, we consider extending our parameterization to mixtures of Gaussians, $q_{t|0,T} \in \mathcal{Q}^K_{\text{MoG}}$. For a given set of K mixture weights w^k and component Gaussian paths $q^k_{t|0,T}$, the following identity allows us to recover the drift $u^{(q)}_{t|0,T}$ of the corresponding mixture distribution $q_{t|0,T}$.

Proposition 4. *Given a set of processes $q^k_{t|0,T}(x)$ and mixtures weights w^k, the vector field satisfying the Fokker-Planck equation in (11) for the mixture $q_{t|0,T}(x) = \sum_k w^k q^k_{t|0,T}(x)$ is given by*

$$u^{(q)}_{t|0,T}(x) = \sum_{k=1}^{K} \frac{w^k q^k_{t|0,T}(x)}{\sum_{j=1}^{K} w^j q^j_{t|0,T}(x)} u^{(q,k)}_{t|0,T}(x)\,, \qquad (17)$$

where $u^{(q,k)}_{t|0,T}(x)$ satisfies the Fokker-Planck equation in (11) for $q^k_{t|0,T}(x)$. 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 $v^{(q)}_{t|0,T}(x)$ by comparing $u^{(q)}_{t|0,T}(x)$ for the mixture of Gaussian path $q_{t|0,T} \in \mathcal{Q}^K_{\text{MoG}}$ to the reference drift $b_t(x)$ as in (13), and proceed to minimize its norm as in (4). We use Gumbel softmax reparamerization 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 $\{\mu^{(\theta)}_{t|0,T}, \Sigma^{(\theta)}_{t|0,T}\}_{k=1}^K$ and 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 $x_0, x_T \sim q^{\text{data}}_{0,T}$ 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 $q^{\text{data}}_{0,T}$ 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 $x_0, x_T \sim q^{\text{data}}_{0,T}$ 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 $v^{(q)}_{t|0,T}$ 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 Vanden-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 in which 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 guar-

Method	MCMC (variable length)	MCMC	Ours
# Evaluations (↓)	3.53M	1.03B	1.28M
Max Energy (↓)	-13.77± 16.43	-17.80 ± 14.77	-14.81 ± 13.73
MinMax Energy (↓)	-40.75	-40.21	-40.56
Log-Likelihood (↑)	-	866.56 ± 17.00	907.15
Max Log-Likelihood (↑)	-	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). MinMax energy reports the lowest maximum energy of all paths.

antees, 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 against in Sec. 5, are more expensive but typically sample diverse paths faster. Recent machine learning approaches such as Plainer et al. (2023); Lelièvre et al. (2023) aim to reduce the need for MD. Holdijk et al. (2024) propose a stochastic optimal control method that simulates (9) with a learned drift, but can be inefficient unless the terminal state is sampled frequently.

5. Experiments

We investigate the capabilities of our approach across a variety of different settings. We first illustrate features of our method on toy examples 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 the following link. 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. E 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 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 learning shooting point selection. This baseline gives theoretical guarantees about the ensemble and thus can be considered as a proxy for the ground truth.

Method	States	# Evals (↓)	Max Energy (↓)	MinMax Energy (↓)
MCMC (variable length)	CV	25.82M	1,212.81 ± 19,444.46	28.67
MCMC*	CV	1.29B	288.46 ± 128.31	60.52
MCMC (variable length)	relaxed	80.23M	269.16 ± 248.51	39.11
MCMC	relaxed	N/A	N/A	N/A
MCMC (variable length)	exact	N/A	N/A	N/A
MCMC	exact	N/A	N/A	N/A
Ours (Cartesian)	exact	38.40M	804.24 ± 0.20	803.62
Ours (Cartesian, 2 Mixtures)	exact	51.20M	828.77 ± 27.34	803.44
Ours (Internal)	exact	51.20M	352.20 ± 0.04	352.08
Ours (Internal, 2 Mixtures)	exact	51.20M	371.16 ± 82.88	239.66

Table 2: Transition path sampling for alanine dipeptide. For MCMC methods, we compare different state definitions of $\mathcal{A}, \mathcal{B}$: '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 they require significantly more than 1 billion 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.

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 (2): *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

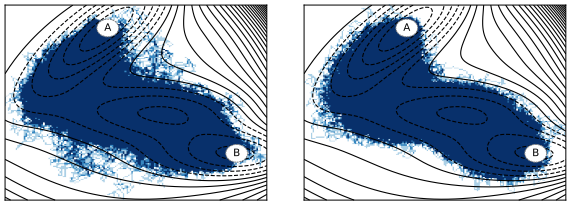

(a) MCMC (b) Ours

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

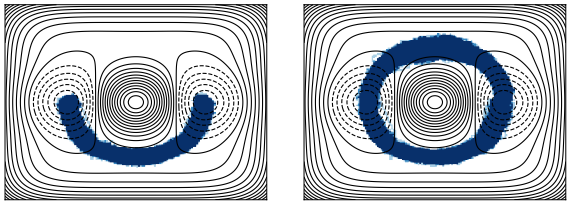

(a) Single Gaussian (b) Mixture of Gaussians

Figure 3: Expressivity of Gaussian vs. mixture of Gaussian paths on a symmetric potential with two transition path modes.

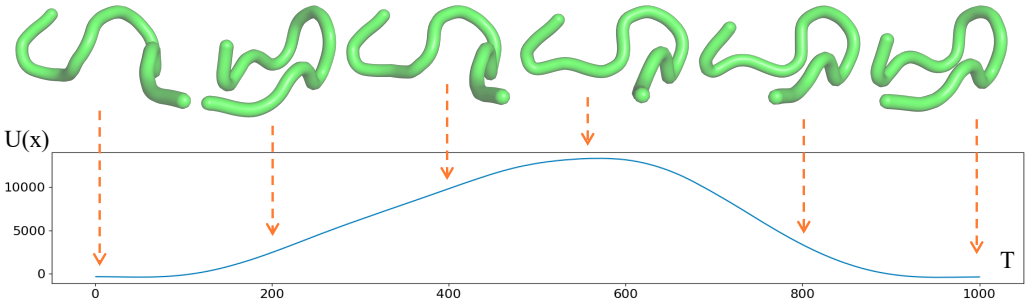

Figure 4: Transition path for the protein Chignolin. The energy plot demonstrates that the conformation goes through a high energy barrier in a total of $T = 1,000fs$, with the highest energy state reached at $567fs$.

freedom) that cannot be summarized as easily. We use an AMBER14 forcefield (Maier et al., 2015) implemented in OpenMM (Eastman et al., 2017) but use DMFF (Wang et al., 2023) to backpropagate through the energy evaluations.

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. E), either with or without Gaussian mixture. For our method, operating in internal coordinates takes more iterations to converge but generates better results compared to Cartesian coordinates, where internal coordinates have nicely distributed input and our network does need not learn equivariances (Du et al., 2022). Similarly, Gaussian mixture paths perform slightly better than a single Gaussian path due to the extra 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 for reaching precise states A, B ('exact') instead of larger regions ('relaxed'), or for computing fixed-length trajectories. Variable length MCMC with relaxed endpoint conditions and fixed-length MCMC with CV perform well on this task, but our method is competitive using fewer evaluations and more strict boundary conditions.

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 12.8M potential energy evaluation calls and a transition with a duration of $T = 1ps$, which is faster compared to $T = 0.6\mu s$ in Lindorff-Larsen et al. (2011).

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 methods with MCMC-based baselines and show comparable accuracy with lower computational costs on both synthetic datasets and real-world molecular systems. Finally, our method might be improved or extended by (1) accounting for conditioning on a set of terminal events, (2) amortizing over many state pairs or systems and finally learning an unconditioned process, and (3) accommodating variable length paths.

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: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.

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*.

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*.

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*.

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. (2024). 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)*.

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).

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.

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).

Lindorff-Larsen, K., Piana, S., Dror, R. O., and Shaw, D. E. (2011). How fast-folding proteins fold. *Science*, 334(6055):517–520.

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 *International Conference on Learning Representations*.

Liu, G.-H., Vahdat, A., Huang, D.-A., Theodorou, E., Nie, W., and Anandkumar, A. (2023b). I^2sb: 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).

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*.

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.

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*. 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.

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.

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 schr\" odinger 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).

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.

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. 2.[Jamison (1975, Thm. 2)] *Let $h(x,t) \coloneqq \rho(x_T \in \mathcal{B} \,|\, x_t = x)$ denote the conditional transition density with respect to the reference process in (5). Letting $G_t \coloneqq \frac{1}{2}\Xi_t \Xi_t^T$, the SDE*

$$dx_{t|0,T} = \Big(b_t(x_{t|0,T}) + 2G_t \nabla_x \log h(x_{t|0,T}, t)\Big) \cdot dt + \Xi_t \, dW_t \tag{18}$$

is associated with the following transition probabilities

$$\rho(x_t = y \,|\, x_s = x, x_T \in \mathcal{B}) = \frac{h(y,s)}{h(x,t)} \rho(x_t = y \,|\, x_s = x), \tag{19}$$

for $s < t < T$, where we omit the dependence of $h(x,t)$ on $\mathcal{B}$ for simplicity of notation.

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

Proposition. 1. *The following PDEs are obeyed by (a) the marginal density of the conditioned process $\rho_{t|0,T}(x) = \rho(x_t = x \,|\, x_0 = A, x_T \in \mathcal{B})$ and (b) the h-function $h(x,t)$ (which implicitly depends on $\mathcal{B}$),*

$$\frac{\partial \rho_{t|0,T}(x)}{\partial t} + \big\langle \nabla_x, \rho_{t|0,T}(x)\big(b_t(x) + 2G_t \nabla_x \log h(x,t)\big)\big\rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \rho_{t|0,T}(x) = 0\,, \tag{20a}$$

$$\frac{\partial h(x,t)}{\partial t} + \big\langle \nabla_x h(x,t), b_t(x)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} h(x,t) = 0\,. \tag{20b}$$

Reparameterizing (3b) in terms of $s(x,t) \coloneqq \log h(x,t)$, we can also write

$$\frac{\partial s(x,t)}{\partial t} + \big\langle \nabla s(x,t), G_t \nabla s(x,t)\big\rangle + \big\langle \nabla s(x,t), b_t(x_t)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} s(x,t) = 0. \tag{20c}$$

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

$$\frac{\partial}{\partial s} p(x_{t+s} = y \,|\, x_t = x) = -\big\langle \nabla_y, p(x_{t+s} = y \,|\, x_t = x) b_{t+s}(y)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} p(x_{t+s} = y \,|\, x_t = x), \tag{21}$$

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

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

$$p(x_{t+s} = y \,|\, x_t = x, x_T \in \mathcal{B}) = \frac{p(x_T \in \mathcal{B} \,|\, x_{t+s} = y)}{p(x_T \in \mathcal{B} \,|\, x_t = x)} p(x_{t+s} = y \,|\, x_t = x)\,, \tag{22}$$

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

$$\underbrace{p(x_T \in \mathcal{B} \,|\, x_t = x)}_{h(x,t)} = \int dy \; \underbrace{p(x_T \in \mathcal{B} \,|\, x_{t+s} = y)}_{h(y,t+s)} p(x_{t+s} = y \,|\, x_t = x)\,, \tag{23}$$

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

$$0 = \int dy \left[p(x_{t+s} = y \,|\, x_t = x) \frac{\partial h(y, t+s)}{\partial s} + \frac{\partial p(x_{t+s} = y \,|\, x_t = x)}{\partial s} h(y, t+s) \right]. \tag{24}$$

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

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

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

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

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

$$\frac{\partial h(y, s)}{\partial s} + \big\langle \nabla_y h(y, s), b_s(y) \big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} h(y, s) = 0\,. \tag{25}$$

as desired to prove the optimality condition in (3b).

To prove (3a), 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

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

Note that $h(x, 0)$ 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

$$\frac{\partial}{\partial s} p(y, s \,|\, \mathcal{B}) = -\big\langle \nabla_y, p(y, s \,|\, \mathcal{B}) (b_s(y) + 2D \nabla_y \log h(y, s)) \big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial y_i \partial y_j} p(y, s \,|\, \mathcal{B})\,. \tag{26}$$

which matches the desired PDE in (3a) thereby proving the first two parts of App. A.1.

Finally, to show (20c), we index time using t in Eq. (25) and change variables $h(x, t) = e^{s(x,t)}$,

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

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

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

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

which demonstrates (20c) since the inner term must be zero. □

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 $q_t(x)s_t(x) \to 0$ vanishes at the boundary $x \to \pm\infty$, which is used when integrating by parts in x.

Corollary 3. (Cor 1. in main text). *The Lagrangian objective in Thm. 1 which solves Doob's h-transform is equivalent to*

$$\mathcal{S} = \min_q \max_s \; s(B,1) - s(A,0) - \int_0^1 dt \int dx \; q_{t|0,T}\left(\frac{\partial s}{\partial t} + \left\langle \nabla s, G_t \nabla s \right\rangle + \left\langle \nabla s, b_t \right\rangle + \left\langle \nabla, G_t \nabla s \right\rangle\right)$$

if $q_{t|0,T}$ satisfies (4c). Note $v_{t|0,T}(x) = \nabla_x s(x,t)$, with $s^*(x,t) = \log h(x,t)$ at optimality

Proof. Consider the following action functional

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

The Lagrangian of this optimization problem is

$$\mathcal{L} = \int_0^1 dt \int dx \left[q_t \left\langle v_t, G_t v_t \right\rangle + s_t \left(\frac{\partial q_t}{\partial t} + \left\langle \nabla, q_t(b_t + 2G_t v_t) \right\rangle - \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_t \right) \right] ,$$

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

$$\begin{aligned}
\frac{\delta \mathcal{L}}{\delta v_t}[\mathfrak{h}_t] = \; & q_t \left\langle (G_t + G_t^T) v_t, \mathfrak{h}_t \right\rangle - q_t \left\langle 2 G_t^T \nabla s_t, \mathfrak{h}_t \right\rangle = 0 \\
& \implies v_t = \nabla s_t ,
\end{aligned} \tag{27}$$

Substituting into the above, we have

$$\mathcal{L} = \int_0^1 dt \int dx \left[s_t \frac{\partial q_t}{\partial t} - q_t \left\langle \nabla s_t, G_t \nabla s_t \right\rangle + s_t \left\langle \nabla, q_t b_t \right\rangle - s_t \left\langle \nabla, G_t \nabla q_t \right\rangle \right] . \tag{28}$$

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

$$\begin{aligned}
\mathcal{L} &= \int dx \; q_1 s_1 - \int dx \; q_0 s_0 + \int_0^1 dt \int dx \left[-q_t \frac{\partial s_t}{\partial t} - q_t \left\langle \nabla s_t, G_t \nabla s_t \right\rangle - q_t \left\langle \nabla s_t, b_t \right\rangle + \left\langle \nabla s_t, G_t \nabla q_t \right\rangle \right] \\
&= \int dx \; q_1 s_1 - \int dx \; q_0 s_0 + \int_0^1 dt \int dx \left[-q_t \frac{\partial s_t}{\partial t} - q_t \left\langle \nabla s_t, G_t \nabla s_t \right\rangle - q_t \left\langle \nabla s_t, b_t \right\rangle - q_t \left\langle \nabla, G_t \nabla s_t \right\rangle \right] \\
&= \int dx \; q_1 s_1 - \int dx \; q_0 s_0 - \int_0^1 dt \int dx \; q_t \left[\frac{\partial s_t}{\partial t} + \left\langle \nabla s_t, G_t \nabla s_t \right\rangle + \left\langle \nabla s_t, b_t \right\rangle + \left\langle \nabla, G_t \nabla s_t \right\rangle \right]
\end{aligned} \tag{29}$$

where in the second line, we integrate by parts in x again. Enforcing $q_1(x) = \delta(x - B)$ 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. □

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

App. A.1,

$$\mathcal{S} = \min_{q,v} \int_0^T dt \int dx\; q_{t|0,T}(x) \big\langle v_{t|0,T}(x), G_t\, v_{t|0,T}(x) \big\rangle\,, \tag{30a}$$

$$\text{s.t.}\quad \frac{\partial q_{t|0,T}(x)}{\partial t} = -\big\langle \nabla_x, q_{t|0,T}(x)\big(b_t(x) + 2G_t\, v_{t|0,T}(x)\big)\big\rangle + \sum_{ij} (G_t)_{ij} \frac{\partial^2}{\partial x_i \partial x_j} q_{t|0,T}(x), \tag{30b}$$

$$q_0(x) = \delta(x - A), \qquad q_T(x) = \delta(x - B)\,. \tag{30c}$$

Namely, the optimal $q^*_{t|0,T}(x)$ obeys (3a) and the optimal $v^*_{t|0,T}(x) = \nabla_x \log h(x,t) = \nabla_x s(x,t)$ follows (3b) or (20c).

Proof. The proof proceeds from (28) above,

$$\mathcal{S} = \min_q \max_s \mathcal{L} = \min_q \max_s \int_0^1 dt \int dx \left[s_t \frac{\partial q_t}{\partial t} - q_t \big\langle \nabla s_t, G_t \nabla s_t \big\rangle + s_t \big\langle \nabla, q_t b_t \big\rangle - s_t \big\langle \nabla, G_t \nabla q_t \big\rangle \right]. \tag{31}$$

We first show that the optimality condition with respect to s_t yields the Fokker-Planck equation for q_t in App. A.1 (3a), before deriving the PDE in (3b) as the optimality condition with respect to q_t.

Optimality Condition for (30) recovers App. A.1 (3a): The variation with respect to s_t of (31) is simple, apart from the intermediate term. For a perturbation direction $\mathfrak{h}_t$, we seek

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

where $\bullet$ 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

$$\int dx\; \frac{\delta(\bullet)}{\delta s_t}\, \mathfrak{h}_t = \left[-2 \int dx\; q_t \big\langle \nabla \mathfrak{h}_t, G_t \nabla s_t \big\rangle \right] \overset{(i)}{=} \left[\int dx\; \mathfrak{h}_t \Big(2 \big\langle \nabla, q_t G_t \nabla s_t \big\rangle \Big) \right] \tag{32}$$

where in (i) we integrate by parts x.

We are now ready to set the variation of (31) with respect to s_t (in an arbitrary direction $\mathfrak{h}_t$) equal to zero. Using (32), we have

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta s_t}[\mathfrak{h}_t] = 0 &= \frac{\partial q_t}{\partial t} + 2 \big\langle \nabla, q_t G_t \nabla s_t \big\rangle + \big\langle \nabla, q_t b_t \big\rangle - \big\langle \nabla, G_t \nabla q_t \big\rangle \\ \implies \quad 0 &= \frac{\partial q_t}{\partial t} + \Big\langle \nabla, q_t \Big(b_t + 2 G_t \nabla s_t \Big) \Big\rangle - \big\langle \nabla, G_t \nabla q_t \big\rangle \end{aligned} \tag{33}$$

which matches the desired optimality condition for the conditioned process in App. A.1 (3a).

Optimality Condition for (30) recovers App. A.1 (3b): Starting again from (31), we take the variation with respect to q_t. First, we repeat identical steps (integrate by parts in both x and t) to reach (29),

$$\mathcal{L} = \int dx\; q_1 s_1 - \int dx\; q_0 s_0 - \int_0^1 dt \int dx\; q_t \left[\frac{\partial s_t}{\partial t} + \big\langle \nabla s_t, G_t \nabla s_t \big\rangle + \big\langle \nabla s_t, b_t \big\rangle + \big\langle \nabla, G_t \nabla s_t \big\rangle \right]$$

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

$$\frac{\delta \mathcal{L}}{\delta q_t}[\mathfrak{h}_t] = 0 = \frac{\partial s_t}{\partial t} + \big\langle \nabla s_t, G_t \nabla s_t \big\rangle + \big\langle \nabla s_t, b_t \big\rangle + \big\langle \nabla, G_t \nabla s_t \big\rangle \tag{34}$$

which is the desired PDE for $s(x,t) = \log h(x,t)$ in (20c). To obtain (3b), we note an identity used to simplify the last term

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

Now, substituting $s(x,t) = \log h(x,t)$ into Eq. (34), we obtain

$$\frac{1}{h_t}\frac{\partial h_t}{\partial t} + \frac{1}{h_t^2}\langle \nabla h_t, G_t \nabla h_t\rangle + \frac{1}{h_t}\langle \nabla h_t, b_t\rangle - \frac{1}{h_t^2}\langle \nabla h_t, G_t \nabla h_t\rangle + \frac{1}{h_t}\langle \nabla, G_t \nabla h_t\rangle = 0\,,$$
$$\implies \quad \frac{\partial h_t(x)}{\partial t} + \langle \nabla h_t(x), b_t(x)\rangle + \langle \nabla, G_t \nabla h_t\rangle = 0, \tag{35}$$

which matches (3b) 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 $h_t(x) = p(x_T \in B' \,|\, x_t = x)$ for some set 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 $h_t(x) = p(x_T = B \,|\, x_t = x)$. □

Corollary 2. *The Lagrangian objective in Thm. 1 is equivalent to the following optimization of $\mathbb{Q}^v_{0:T}$*

$$\mathcal{S} := \min_{\mathbb{Q} \text{ s.t. } \mathbb{Q}=\delta, \mathbb{Q}=\delta} D_{KL}[\mathbb{Q}^v_{0:T} : \mathbb{P}^{ref}_{0:T}] \tag{10}$$

where the minimizing argument recovers the path measure $\mathbb{P}^*_{0:T}$ associated with the SDE in (6).

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 $x_0 = A$,

$$\mathbb{P}^{\text{ref}}_{0:T}: \qquad dx_t = b_t(x_t)\cdot dt + \Xi_t\, dW_t\,, \tag{36}$$
$$\mathbb{Q}^v_{0:T}: \qquad dx_t = \big(b_t(x_t) + 2G_t\, v_{t|0,T}(x_{t|0,T})\big)\cdot dt + \Xi_t\, dW_t\,, \tag{37}$$

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

$$\frac{d\mathbb{Q}^v_{0:T}}{d\mathbb{P}^{\text{ref}}_{0:T}} = \frac{q_{t|0,T}(x_0, ...x_T)}{\rho(x_0, ...x_T)} = \exp\Big\{ -\frac{1}{2}\int_0^T \langle 2G_t\, v_{t|0,T}(x_t), (G_t)^{-1}\, 2G_t\, v_{t|0,T}(x_t)\rangle dt$$
$$- \int 2\big(G_t\, v_{t|0,T}(x_t)\big)^T G_t^{-1} dW_t\Big\} \tag{38}$$

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

$$D_{KL}[\mathbb{Q}^v_{0:T} : \mathbb{P}^{\text{ref}}_{0:T}] = 2\int_0^1 dt \int dx_t\; q_{t|0,T}(x_t)\,\langle v_{t|0,T}(x_t), G_t\, v_{t|0,T}(x_t)\rangle, \tag{39}$$

which matches (4a) 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 (37) satisfies the Fokker-Planck equation in (4b) (Särkkä and Solin, 2019, Sec. 5.2). We respect (4c) by optimizing over endpoint-constrained path measures, which yields

$$\mathcal{S} = \min_{\mathbb{Q} \text{ s.t. } \mathbb{Q}=\delta, \mathbb{Q}=\delta} D_{KL}[\mathbb{Q}^v_{0:T} : \mathbb{P}^{\text{ref}}_{0:T}] \tag{40}$$

as desired. □

B. Gaussian Path Parameterizations

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

$$u^{(q)}_{t|0,T}(x) := \frac{\partial \mu_{t|0,T}}{\partial t} + \left[\frac{1}{2}\frac{\partial \Sigma_{t|0,T}}{\partial t}\Sigma^{-1}_{t|0,T} - G_t\,\Sigma^{-1}_{t|0,T}\right](x - \mu_{t|0,T}) \tag{41}$$

satisfies the Fokker-Planck equation (11) 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,

$$\log q_t(x) = -\frac{1}{2}(x-\mu_t)^T\Sigma_t^{-1}(x-\mu_t) - \frac{d}{2}\log(2\pi) - \frac{1}{2}\log\det\Sigma_t\,, \tag{42a}$$

$$\nabla_x \log q_t(x) = -\Sigma_t^{-1}(x-\mu_t)\,, \tag{42b}$$

$$\frac{\partial}{\partial t}\log q_t(x) = (x-\mu_t)^T\Sigma_t^{-1}\frac{\partial\mu_t}{\partial t} + \frac{1}{2}(x-\mu_t)^T\Sigma_t^{-1}\frac{\partial\Sigma_t}{\partial t}\Sigma_t^{-1}(x-\mu_t) - \frac{1}{2}\mathrm{tr}\left(\Sigma_t^{-1}\frac{\partial\Sigma_t}{\partial t}\right) \tag{42c}$$

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

$$\begin{aligned}\frac{\partial q_t}{\partial t} &= -\langle\nabla_x, q_t u_t^{\mathrm{o}}\rangle = -q_t\langle\nabla_x, u_t^{\mathrm{o}}\rangle + \langle\nabla_x q_t, \nabla_x u_t^{\mathrm{o}}\rangle \\ \implies \frac{\partial}{\partial t}\log q_t &= -\langle\nabla_x, u_t^{\mathrm{o}}\rangle - \langle\nabla_x \log q_t, u_t^{\mathrm{o}}\rangle\end{aligned} \tag{43}$$

The vector field satisfying this equation is

$$u_t^{\mathrm{o}}(x) = \frac{\partial\mu_t}{\partial t} + \frac{1}{2}\frac{\partial\Sigma_t}{\partial t}\Sigma_t^{-1}(x-\mu_t) \tag{44}$$

which we can confirm using the identities in (42). Indeed, for the terms on the RHS of Eq. (43),

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

Putting these terms and the time derivative from (42c) into Eq. (43) 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 (11). 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

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

Finally, we use the identities in (42) to obtain

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

□

Proposition 4. *Given a set of processes $q^k_{t|0,T}(x)$ and mixtures weights w^k, the vector field satisfying the Fokker-Planck equation in (11) for the mixture $q_{t|0,T}(x) = \sum_k w^k q^k_{t|0,T}(x)$ is given by*

$$u^{(q)}_{t|0,T}(x) = \sum_{k=1}^{K}\frac{w^k q^k_{t|0,T}(x)}{\sum_{j=1}^{K} w^j q^j_{t|0,T}(x)} u^{(q,k)}_{t|0,T}(x)\,, \tag{17}$$

where $u^{(q,k)}_{t|0,T}(x)$ satisfies the Fokker-Planck equation in (11) for $q^k_{t|0,T}(x)$. 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. Algorithm

The pseudocode presented in Algorithm 1 shows the training procedure for a single sample t. In practice, we average over a batch of independent samples and predict multiple μ, Σ for k Gaussian mixtures.

Algorithm 1 Variational Doob h-Transform for a Single Sample with Gaussian Paths

Input: Reference drift b_t, diffusion matrix $G_t = \Xi_t \Xi_t^T$, conditioning endpoints $x_0 = A, x_T = B$
while not converged **do**
Sample x_t for $t \sim \mathcal{U}(0, T)$ from current $q^{(\theta)}_{t|0,T}$ using (14)

$$x_t = \mu^{(\theta)}_{t|0,T} + \Sigma^{(\theta)}_{t|0,T}\,\epsilon, \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \mathbb{I}_D),$$
$$\mu^{(\theta)}_{t|0,T} = (1 - \frac{t}{T})A + \frac{t}{T}\,B + \frac{t}{T}(1 - \frac{t}{T})\textsc{nnet}_\theta(t, A, B)_{[:D]},$$
$$\Sigma^{(\theta)}_{t|0,T} = \frac{t}{T}(1 - \frac{t}{T})\texttt{diag}\big(\textsc{nnet}_\theta(t, A, B)_{[D:]}\big) + \sigma^2_{\text{min}}\mathbb{I}.$$

Calculate $u_{t|0,T}(x_t)$ using the output of the neural network and (12)

$$u^{(q)}_{t|0,T}(x_t) = \frac{\partial \mu^{(\theta)}_{t|0,T}}{\partial t} + \left[\frac{1}{2}\frac{\partial \Sigma^{(\theta)}_{t|0,T}}{\partial t}(\Sigma^{(\theta)}_{t|0,T})^{-1} - G_t\,(\Sigma^{(\theta)}_{t|0,T})^{-1}\right]\big(x_t - \mu^{(\theta)}_{t|0,T}\big)$$

Calculate $v_{t|0,T}(x_t)$ using (13) and the underlying drift $b_t(x_t)$

$$v^{(q)}_{t|0,T}(x_t) = \frac{1}{2}(G_t)^{-1}\Big(u^{(q)}_{t|0,T}(x_t) - b_t(x_t)\Big)$$

Calculate $\mathcal{L}$ from Thm. 1

$$\mathcal{L} = \langle v^{(q)}_{t|0,T}(x_t), G_t v^{(q)}_{t|0,T}(x_t)\rangle$$

Update $\theta \leftarrow \text{optimizer}(\theta, \nabla_\theta \mathcal{L})$
end while
return

D. Extended Related Work

D.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).

E. Further Experimental Details

E.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

$$L(x_0, x_1, \ldots, x_{N-1}) = \rho(x_0) \cdot \prod_{i=0}^{N-2} \pi(x_{i+1} \,|\, x_i)\,. \tag{45}$$

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)

$$\rho(z, v) \propto \exp\left(-\frac{U(z)}{k_B T}\right) \cdot \mathcal{N}\big(v \,|\, 0, k_B T \cdot M^{-1}\big)\,, \tag{46}$$

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.

E.2. Toy Potentials

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

$$x_{t+1} = x_t - dt \cdot \nabla_x U(x_t) + \sqrt{dt} \cdot \texttt{diag}(\xi) \cdot \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1)\,, \tag{47}$$

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

$$\begin{aligned} U(x, y) = &- 200 \cdot \exp\big(-(x-1)^2 - 10y^2\big) \\ &- 100 \cdot \exp\big(-x^2 - 10 \cdot (y - 0.5)^2\big) \\ &- 170 \cdot \exp\big(-6.5 \cdot (0.5 + x)^2 + 11 \cdot (x + 0.5) \cdot (y - 1.5) - 6.5 \cdot (y - 1.5)^2\big) \\ &+ 15 \cdot \exp\big(0.7 \cdot (1 + x)^2 + 0.6 \cdot (x + 1) \cdot (y - 1) + 0.7 \cdot (y - 1)^2\big)\,. \end{aligned} \tag{48}$$

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.

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 maximum energy on the trajectory in Fig. 5b reveals that all methods have a similar quality of paths.

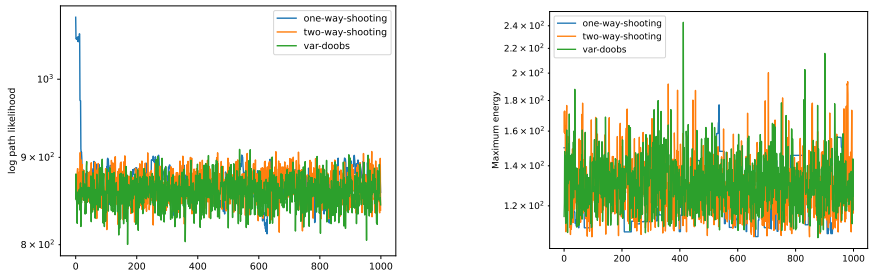

(a) Log Path Likelihood (b) Maximum Energy

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.

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

$$\begin{aligned} U(x, y) = &+ 2 \cdot \exp\big(-(12x^2 + 12y^2)\big) \\ &- 1 \cdot \exp\big(-(12 \cdot (x + 0.5)^2 + 12y^2)\big) \\ &- 1 \cdot \exp\big(-(12 \cdot (x - 0.5)^2 + 12y^2)\big) + x^6 + y^6 \,. \end{aligned} \tag{49}$$

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 $w \notin \theta$.

E.3. 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.

E.4. Molecular Systems

To simulate molecular dynamics, we rely on the AMBER14 forcefield (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 $\nabla_\theta U\Big(x_{t|0,T} \sim \mathcal{N}(\mu^{(\theta)}_{t|0,T}, \Sigma^{(\theta)}_{t|0,T})\Big)$, 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.

Visualization of transition for alanine dipeptide. In Fig. 6, we show a transition sampled without any noise from the model with internal coordinates and 2 Gaussian mixtures.

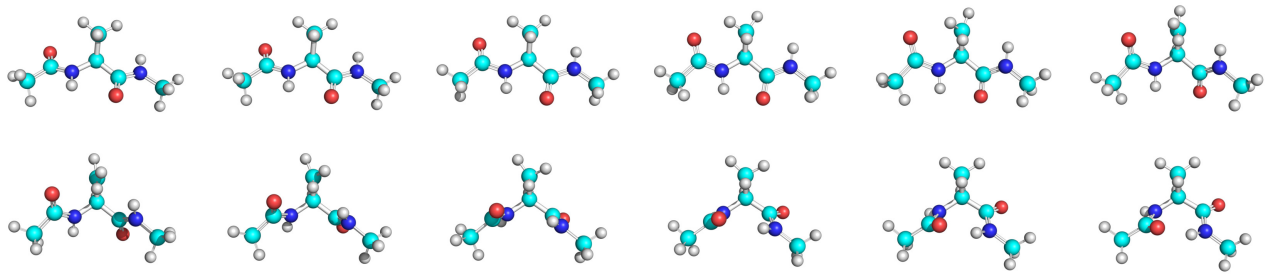

Figure 6: Transition path for the alanine dipeptide.

E.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.

F. 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.