

Functional Adjoint Sampler: Scalable Sampling on Infinite Dimensional Spaces

Byoungwoo Park¹ Juho Lee¹ Guan-Horng Liu²

Abstract

Learning-based methods for sampling from the Gibbs distribution in finite-dimensional spaces have progressed quickly, yet theory and algorithmic design for infinite-dimensional function spaces remain limited. This gap persists despite their strong potential for sampling the paths of conditional diffusion processes, enabling efficient simulation of trajectories of diffusion processes that respect rare events or boundary constraints. In this work, we present the adjoint sampler for infinite-dimensional function spaces, a stochastic optimal control-based diffusion sampler that operates in function space and targets Gibbs-type distributions on infinite-dimensional Hilbert spaces. Our Functional Adjoint Sampler (FAS) generalizes Adjoint Sampling (Havens et al., 2025) to Hilbert spaces based on a SOC theory called stochastic maximum principle, yielding a simple and scalable matching-type objective for a functional representation. We show that FAS achieves superior performance on transition path sampling across synthetic potential and real molecular systems, including Alanine Dipeptide and Chignolin. Codes are available at <https://github.com/bw-park/FAS>.

1. Introduction

The study of constructing Markov diffusion processes under certain conditions—such as rare events and prescribed terminal distributions—has long been a central theme in probability and statistical physics (Doob, 2001; Chetrite and Touchette, 2015). With the recent explosive interests in generative modeling (Song et al., 2021; Lipman et al., 2023), the framework has emerged in many other machine learning applications, including data assimilation (Chopin et al., 2023; Park et al., 2025), population modeling (Liu

et al., 2022; Chen et al., 2023), molecule dynamics simulation (Plainer et al., 2023; Holdijk et al., 2022), and—of particular focus of this work—sampling from the Gibbs distribution (Tzen and Raginsky, 2019).

Specifically, we are interested in generating samples from a Gibbs distribution $\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})}$ given its unnormalized energy function $U(\mathbf{x})$. Modern computational methods—known as *diffusion samplers*—achieve this by learning diffusion processes whose terminal distribution matches the law of π . Among these, adjoint-based diffusion samplers (Havens et al., 2025; Liu et al., 2026; Choi et al., 2026; Guo et al., 2026) form a distinctive family by reformulating the sampling problem as a stochastic optimal control problem (SOC) (Fleming and Soner, 1993; Kappen, 2005) and, from which, employing Adjoint Matching (AM) (Domingo-Enrich et al., 2025). For instance, Adjoint Sampling (AS) (Havens et al., 2025) learns the drift of a *memoryless*¹ stochastic differential equation with an adjoint-matching objective that depends solely on on-policy samples and a simulation-free target. This yields a highly scalable method compared to vanilla SOC-based methods that rely on simulation-based targets (Zhang and Chen, 2022), as well as those that require importance-weighted samples (Phillips et al., 2024; Akhound-Sadegh et al., 2024).

Despite their practical successes, all of the aforementioned adjoint samplers generate samples only in finite-dimensional state spaces, *e.g.*, $\mathbf{x} \in \mathbb{R}^d$. This essentially precludes their applications to sampling Gibbs distributions defined on more general *trajectory* spaces—a family of distributions that appears prevalently in computational chemistry, with a representative example known as the *transition path sampling* (TPS) (Bolhuis et al., 2002; Vanden-Eijnden et al., 2010). At its core, TPS aims to identify rare reactive paths by sampling ensembles of dynamical trajectories between metastable states driven by stochastic fluctuations or external forcing (Dürr and Bach, 1978). Studying these transitions enables accurate estimations of reaction rates and mechanisms that explain how molecules transition between metastable states, benefits various applications such as protein folding (Dobson, 2003; Noé et al., 2009), enzyme catalysis (Basner and Schwartz, 2005), drug design (Tiwarly et al.,

¹KAIST ²FAIR at Meta. Correspondence to: Byoungwoo Park <bw.park@kaist.ac.kr>, Guan-Horng Liu <ghliu@meta.com>.

¹Practically, AS (Havens et al., 2025) fulfills the memoryless condition with a fixed initial condition.

2015) and nucleation (Sosso et al., 2016).

As TPS aims to generate a Gibbs distribution of trajectories following specific energy structures, it is nature to consider samplers formulated directly in these *infinite-dimensional function spaces*. Indeed, trajectories of $\mathbb{R}^d$ -valued diffusion processes are commonly viewed as random functions (Hairer et al., 2009) where the associated probability laws typically reside on function spaces such as $\mathcal{H} := L^2(\mathbf{L}, \mathbb{R}^d)$, where $\mathbf{L} \subseteq \mathbb{R}$ denotes a time interval. In this case, classical sampling methods based on Langevin-type infinite-dimensional stochastic differential equations have been investigated (Stuart et al., 2004; Beskos et al., 2008), despite remaining primarily of theoretical interest rather than practical utility. On the other hand, existing diffusion samplers for TPS (Holdijk et al., 2022; Seong et al., 2025) remains confined to finite dimensional state spaces because they rely on reference dynamics driven by local energy gradients, which preclude a scalable AM objective.

In this paper, we introduce **Functional Adjoint Sampler** (FAS), a scalable matching algorithm that extends AS to infinite-dimensional function spaces. Formally, we recast sampling from Gibbs-type distributions in their infinite-dimensional formulation and build diffusion samplers in function spaces. Leveraging infinite-dimensional SOC, we cast the non-equilibrium sampling (NES) as a function-space generalization of (Tzen and Raginsky, 2019; Zhang and Chen, 2022). Furthermore, under standard regularity, we demonstrate that Adjoint Matching (AM) can also be extended to function spaces by applying the *Stochastic Maximum Principle* (Bensoussan, 1983; Du and Meng, 2013), providing a simple and scalable training objective.

In practice, we take a particular interest in the problem of TPS, our method generates conditional paths across diverse settings, from synthetic potential to molecular systems such as Alanine Dipeptide and Chignolin. Specifically, FAS operates directly in infinite dimensional function spaces by choosing a basis that enforces the prescribed boundary condition (*i.e.*, Dirichlet boundary) of the simulated path, thereby yielding perfect transition between metastable states without extra projection or penalty terms. This principled formulation ensures well posedness of the conditional diffusion and supports stable discretizations, while outperforming prior works that rely on finite-dimensional dynamics.

We summarize our contributions as follows:

- We introduces a function-space diffusion sampler that draws samples from Gibbs-type distribution on infinite-dimensional path spaces. Taking an infinite-dimensional Ornstein–Uhlenbeck semigroup as the reference dynamics, We casts NES as a infinite-dimensional SOC problem.
- Based on the SOC theory called *Stochastic Maximum*

Principle, we propose FAS, a generalization of the matching-based training objective of AS (Havens et al., 2025) to Hilbert spaces and yielding a scalable training algorithm suitable for functional representations.

- We demonstrate superior performance of FAS on TPS, outperforming prior finite-dimensional methods across synthetic potential and challenging real molecular systems including Alanine Dipeptide and Chignolin, while maintaining perfect transition-hit rates to a target states.

Notation. Throughout the paper, we consider a real and separable Hilbert space $\mathcal{H}$, equipped with the norm $\|\cdot\|_{\mathcal{H}}$ and inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$. We denote $\mathbb{P}^{(\cdot)}$ the path measure on the space of all continuous mappings $\Omega = C([0, T], \mathcal{H})$ and $\mathbf{X}^{(\cdot)} \in \mathcal{H}$ the stochastic processes associated with this path measure. We denote the evaluated value of a (random) function $\mathbf{X}[\cdot] : \mathbb{R}^n \rightarrow \mathbb{R}^d$ at a point $u \in \mathbb{R}^n$ from domain space $\mathbb{R}^n$ as $\mathbf{X}[u] \in \mathbb{R}^d$, which will be omitted unless necessary for clarity. For a function $\mathcal{V} : [0, T] \times \mathcal{H} \rightarrow \mathbb{R}$, we define $D_{\mathbf{x}}\mathcal{V}, D_{\mathbf{xx}}\mathcal{V}$ as the first and second order Fréchet derivatives with respect to the variable $\mathbf{x} \in \mathcal{H}$, respectively, and $\partial_t \mathcal{V}$ as the derivative with respect to the time variable $t \in [0, T]$. The Hermitian adjoint will be denoted as $\dagger$.

2. Preliminaries

Here, we provide a concise overview of the necessary probability concepts for studying diffusion samplers formulated in infinite-dimensional state spaces. Additional review can be found in Appendix A.

Gaussian measure in Hilbert spaces Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space and $(\mathcal{H}, \mathcal{B}(\mathcal{H}))$ be a measurable state space, and let $\mathbf{X} : \Omega \rightarrow \mathcal{H}$ be an $\mathcal{H}$ -valued random variable inducing the push-forward measure $\nu := \mathbf{X}_{\#}\mathbb{Q}$. Then, the measure ν is called Gaussian when every $\mathbb{R}$ -valued random variable $\langle u, \mathbf{X} \rangle_{\mathcal{H}}$ is Gaussian for all $u \in \mathcal{H}$. In this case, there is a unique mean function $\mathbf{m} \in \mathcal{H}$ satisfying $\langle \mathbf{m}, u \rangle_{\mathcal{H}} = \mathbb{E}_{\nu}[\langle \mathbf{X}, u \rangle_{\mathcal{H}}]$ and unique non-negative self-adjoint covariance operator $Q : \mathcal{H} \rightarrow \mathcal{H}$ defined by $\langle u, Qv \rangle_{\mathcal{H}} = \mathbb{E}_{\nu}[\langle \mathbf{X} - \mathbf{m}, u \rangle_{\mathcal{H}} \langle \mathbf{X} - \mathbf{m}, v \rangle_{\mathcal{H}}]$ for all $u, v \in \mathcal{H}$, where it is trace class operator *i.e.*, $\text{Tr}(Q) < \infty$. We write $\nu = \mathcal{N}(\mathbf{m}, Q)$ and say that $\mathbf{X}$ is centered when $\mathbf{m} = 0$. Assume now that $\mathbf{X}$ is centred with law $\mathcal{N}(0, Q)$. Then, there exists an eigen-system $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R} \times \mathcal{H} : k \in \mathbb{N}\}$ such that $Q\phi^{(k)} = \lambda^{(k)}\phi^{(k)}$ and $\text{Tr}(Q) = \sum_{k=1}^{\infty} \langle Q\phi^{(k)}, \phi^{(k)} \rangle_{\mathcal{H}} = \sum_{k=1}^{\infty} \lambda^{(k)} < \infty$.

Q-Wiener process and Cameron-Martin space In infinite-dimensional spaces, we distinguish two Gaussian driving noises. Let $\mathcal{U}$ be a separable Hilbert space with orthonormal basis $\{\phi^{(k)}\}_{k \geq 1}$. A *cylindrical Wiener process* is the natural extension of $\mathbb{R}$ -valued Wiener process, formally

represented by the series $\mathbf{W}_t = \sum_{k=1}^{\infty} \mathbf{W}_t^{(k)} \phi^{(k)}$, where $\{\mathbf{W}_t^{(k)}\}_{k \geq 1}$ are independent Brownian motions. It satisfies independent increments $\mathbf{W}_{t+\Delta_t} - \mathbf{W}_t \sim \mathcal{N}(0, \Delta_t \mathbf{I}_{\mathcal{U}})$ but the covariance $\mathbf{I}_{\mathcal{U}}$ is *not* trace class when $\dim \mathcal{U} = \infty$. In other words, $\mathcal{N}(0, \Delta_t \mathbf{I}_{\mathcal{U}})$ is not supported in $\mathcal{U}$ itself, thereby $\mathbf{W}_t$ is *not* an $\mathcal{U}$ -valued random variable. Here, we embed $\mathcal{U}$ into another Hilbert space $\mathcal{H}$ and define an Hilbert–Schmidt operator $Q^{1/2} : \mathcal{U} \rightarrow \mathcal{H}$ such that $Q = Q^{1/2}(Q^{1/2})^\dagger$ is trace class on $\mathcal{H}$. Then, we can introduce the following $\mathcal{H}$ -valued covariated noise process:

$$\mathbf{W}_t^Q := Q^{1/2} \mathbf{W}_t = \sum_{k=1}^{\infty} \sqrt{\lambda^{(k)}} \mathbf{W}_t^{(k)} \phi^{(k)} \in \mathcal{H}, \quad (1)$$

where $\lambda^{(k)}$ is an eigen-value of Q . Consequently, we obtain, $\mathbf{W}_{t+\Delta_t}^Q - \mathbf{W}_t^Q \sim \mathcal{N}(0, \Delta_t Q)$ with trace class operator Q . We call $\mathbf{W}_t^Q$ as Q -Wiener process, which is now $\mathcal{H}$ -valued random variable. To formalize the set of admissible control directions associated with $\mathbf{W}_t^Q$, we introduce the *Cameron–Martin Space* $\mathcal{H}_0 \subseteq \mathcal{H}$ equipped with inner product:

$$\langle u, v \rangle_{\mathcal{H}_0} = \langle Q^{-1/2} u, Q^{-1/2} v \rangle_{\mathcal{H}}, \quad \forall u, v \in \mathcal{H}_0. \quad (2)$$

Infinite dimensional stochastic differential equations

With the infinite dimensional Wiener processes defined in (1), we introduce the infinite-dimensional stochastic differential equations as follows (Da Prato and Zabczyk, 2014):

$$d\mathbf{X}_t = \mathcal{A}\mathbf{X}_t dt + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{X}_0 = \mathbf{x}_0 \in \mathcal{H}, \quad (3)$$

where $\mathcal{A} : \mathcal{H} \rightarrow \mathcal{H}$ is a linear operator and $\sigma_t : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ is a noise schedule. Under the mild conditions in Theorem B.1, the dynamics in (3) generates an Ornstein–Uhlenbeck C_0 -semigroup (Da Prato and Zabczyk, 2002) and admit the unique *mild* solution. The resulting stochastic process $\{\mathbf{X}_t\}_{t \geq 0}$ is then an $\mathcal{H}$ -valued Gaussian process with trace-class covariance operator:

$$\mathbf{X}_t \sim \nu_t^{\mathbf{x}_0} := \mathcal{N}(e^{t\mathcal{A}} \mathbf{x}_0, Q_t), \quad (4)$$

where $Q_t = \int_0^t e^{(t-s)\mathcal{A}} \sigma_s^2 Q e^{(t-s)\mathcal{A}^\dagger} ds$. Finally, as $t \rightarrow \infty$, the law of $\mathbf{X}_\infty$ converges to the Gaussian invariant law

$$\nu_\infty = \mathcal{N}(0, Q_\infty), \quad \text{with } Q_\infty = -\frac{1}{2} \sigma_\infty^2 \mathcal{A}^{-1} Q. \quad (5)$$

3. Functional Adjoint Sampler

In this section, we present a scalable diffusion sampler that operates directly in infinite-dimensional Hilbert spaces, targeting a *distribution of functions*. Full proofs and necessary derivations are provided in Section B.

3.1. Non-Equilibrium Sampling in Function Spaces

Sampling from a Gibbs distribution is a fundamental task. The goal is to draw samples from a target distribution π that

is known only through an unnormalized energy $U : \mathcal{H} \rightarrow \mathbb{R}$:

$$d\pi(\mathbf{x}) = \frac{1}{\mathcal{Z}} e^{-U(\mathbf{x})} d\nu(\mathbf{x}), \quad \mathcal{Z} = \int_{\mathcal{H}} e^{-U(\mathbf{x})} d\nu(\mathbf{x}), \quad (6)$$

where ν is the measure defined on some space $\mathcal{H}$. In finite dimensions, *e.g.*, $\mathcal{H} = \mathbb{R}^d$, the Gibbs density can be expressed with respect to Lebesgue measure *i.e.*, $d\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})} d\mathbf{x}$. However, due to the absence of an equivalent formulation of Lebesgue measure for Hilbert spaces, the expression in (6) is well-defined only after specifying an appropriate reference measure ν . The following lemma states a condition for the target law to be well defined for arbitrary Hilbert spaces.

Lemma 1 (Finite Normalization Constant). *Let us assume $-U(\mathbf{x}) \leq \beta \|\mathbf{x}\|_{\mathcal{H}}^2 - C$ with constants $\beta, C > 0$. Then, for any $\text{Tr}(Q) < \infty$, choosing $\nu = \mathcal{N}(0, Q)$ in (6) results in a finite normalization constant $\mathcal{Z} < \infty$.*

From Lemma 1, we have a simple yet effective choice of reference measure ν , the centered Gaussian $\mathcal{N}(0, Q)$ with trace-class covariance Q . This ensures that $\mathcal{Z} < \infty$, and, consequently, the target π in (6) is well-defined on $\mathcal{H}$.

To sample from (6) on $\mathcal{H}$, it is natural to consider an $\mathcal{H}$ -valued stochastic process that follows the gradient flow of U (Da Prato et al., 1996; Suzuki, 2020)—in the same spirit of overdamped Langevin dynamics in $\mathbb{R}^d$. Concretely, one may define the infinite-dimensional dynamics as follows:

$$d\mathbf{X}_t = [\mathcal{A}\mathbf{X}_t - Q D_{\mathbf{x}} U(\mathbf{x})] dt + \sigma_t d\mathbf{W}_t^Q, \quad (7)$$

Convergence of the dynamics (7) into the target measure (6) is guaranteed when $\sigma_t := \sigma$, and direct simulation via space-time discretization of the $\mathcal{H}$ -valued dynamics is standard. However, the resulting dynamics are of limited practical use as the mixing is only effective near equilibrium *i.e.*, $t \rightarrow \infty$, because they rely on asymptotic convergence to the invariant reference measure ν_∞ in (5) required for the well-defined normalizing constant $\mathcal{Z} < \infty$ by following Lemma 1.

Non equilibrium sampling (NES) Hence, to reach the target distribution (6) within a finite time horizon $T < \infty$, we recast our sampling problem as a NES problem. Unlike the Langevin-type processes in (7), which approach the target only asymptotically at equilibrium with reference ν_∞ , we employ an importance-sampling (IS) via change of measure to sample the target efficiently in finite time:

$$\begin{aligned} d\pi(\mathbf{x}) &\propto e^{-U(\mathbf{x})} \frac{d\nu_\infty(\mathbf{x})}{d\nu_T^{\mathbf{x}_0}(\mathbf{x})} d\nu_T^{\mathbf{x}_0}(\mathbf{x}) \\ &= e^{-U(\mathbf{x}) - \log \frac{d\nu_T^{\mathbf{x}_0}(\mathbf{x})}{d\nu_\infty(\mathbf{x})}} d\nu_T^{\mathbf{x}_0}(\mathbf{x}), \end{aligned} \quad (8)$$

where $\nu_T^{\mathbf{x}_0}$ is marginal law of the reference dynamics (3) at time T . Hence, for the NES in infinite dimensions we need

an extra correction term in the importance weight, given by the Radon–Nikodým derivative (RND). It corrects the mismatch between the two Gaussian measures so that after reweighting, samples drawn from $\nu_T^{\mathbf{x}_0}$ have the desired target π , while guarantees the π to be well-defined. However, compared to finite dimensional cases, the infinite dimensional RND exists only when the two measures are exactly compatible, because of absence of Lebesgue measure.

When this compatibility condition fails, the two measures are singular. Hence, no density can be written explicitly. Fortunately, under Ornstein–Uhlenbeck semigroup generated by (3), $\nu_T^{\mathbf{x}_0}$ and ν_∞ are mutually absolutely continuous Gaussian measures (Da Prato and Zabczyk, 2002). Consequently, we can derive the explicit RND, as provided below.

Theorem 2 (Explicit RND). *Suppose Theorem B.1 holds. Then, for any $0 < t < \infty$ and initial condition $\mathbf{x}_0 \in \mathcal{H}$, $\nu_t^{\mathbf{x}_0} = \mathcal{N}(e^{tA}\mathbf{x}_0, Q_t)$ and $\nu_\infty = \mathcal{N}(0, Q_\infty)$ are mutually absolutely continuous. Moreover, for any $\mathbf{x} \in \mathcal{H}$, we obtain the explicit RND $q_t(\mathbf{x}_0, \mathbf{x}) := \frac{d\nu_t^{\mathbf{x}_0}}{d\nu_\infty}(\mathbf{x})$ defined as follows:*

$$q_t(\mathbf{x}_0, \mathbf{x}) = \det(\Theta_t)^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0} \rangle_\infty + \langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x} \rangle_\infty - \frac{1}{2} \langle e^{2tA} \Theta_t^{-1} \mathbf{x}, \mathbf{x} \rangle_\infty \right], \quad (9)$$

where we define $\mathbf{m}_t^{\mathbf{x}_0} = e^{tA}\mathbf{x}_0$, $\Theta_t = 1 - e^{2tA}$ and Q_∞ norm $\langle u, v \rangle_\infty := \langle Q_\infty^{-\frac{1}{2}} u, Q_\infty^{-\frac{1}{2}} v \rangle_{\mathcal{H}}$.

Theorem 2 provides an explicit RND, so IS can in principle generate samples from π in (6). In practice, unfortunately, the importance weights are computed on a finite domain discretization, and when that discretization is refined the number of samples required grows exponentially proportional to discretization. Therefore, IS might become impractical.

3.2. Infinite-Dimensional Stochastic Optimal Control

Inspired by recent advances in diffusion samplers (Zhang and Chen, 2022; Vargas et al., 2023; Berner et al., 2024; Liu et al., 2026; Choi et al., 2026), we reformulate the NES problem as a variational inference problem on path space $\Omega = C([0, T], \mathcal{H})$ to build an efficient sampler by introducing a learnable control to steer the process (3) to the π in finite horizon $T < \infty$. Let us introduce a target path measure $\mathbb{P}^*$, which is an extension of target π in (6) into Ω :

$$d\mathbb{P}^*(\mathbf{X}) \propto e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{x}_0, \mathbf{X}_T)} d\mathbb{P}(\mathbf{X}), \quad (10)$$

where $\mathbb{P}$ is path measure induced by the process in (3). Now, we consider variational path measures $\mathbb{P}^\alpha$ induced by infinite-dimensional SDE with control (Fabbri et al., 2017):

$$d\mathbf{X}_t^\alpha = [\mathbf{A}\mathbf{X}_t^\alpha + \sigma_t Q^{1/2} \alpha(\mathbf{X}_t^\alpha, t)] dt + \sigma_t d\mathbf{W}_t^Q \quad (11)$$

where $\alpha : \mathcal{H} \times [0, T] \rightarrow \mathcal{U}$ such that $Q^{1/2} \alpha \in \mathcal{H}_0 \subseteq \mathcal{H}$ is a Markov control functional. The primary reason for studying this path measure formulation is that the Gibbs-type target distribution π in (6) can be effectively induced by choosing optimal control α_t^* in (11). By invoking generalized Girsanov theorem (Da Prato and Zabczyk, 2014, Theorem 10.14), such an optimal control can be found by solving the following infinite-dimensional optimization problem (Fuhrman, 2003; Theodorou et al., 2018):

$$\min_{\alpha} \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \frac{1}{2} \|\alpha_s\|_{\mathcal{H}}^2 ds + g(\mathbf{X}_T^\alpha) \right] \text{ s.t. (11),} \quad (12)$$

where $g(\mathbf{x}) = U(\mathbf{x}) + \log q_T(\mathbf{x}_0, \mathbf{x})$. Accordingly, the optimal control α_t^* which minimize the objective in (11) is equivalent to minimizer of Kullback–Liebler divergence i.e., $\alpha^* = \min_{\alpha} D_{\text{KL}}(\mathbb{P}^\alpha | \mathbb{P}^*)$ with target path measure $\mathbb{P}^*$ in (10). In other word, this formulation enables us to the exact sampling from target π , namely $\mathbf{X}_T^{\alpha^*} \sim \pi$ in finite time $T < \infty$. This successfully extends the finite-dimensional control-based NES found in (Tzen and Raginsky, 2019; Zhang and Chen, 2022) into the Hilbert spaces $\mathcal{H}$.

3.3. Infinite Dimensional Adjoint Matching

The SOC problem in (12) follows the *least-action* principle in path space, where the goal is to find the *global* optimal control minimizes the terminal cost $g(\mathbf{x})$ while minimizing the action. In practice, solving this global problem by direct path-wise minimization is expensive because every update requires gradients through the full trajectory (Havens et al., 2025). A practical workaround is to move to the dual formulation. *Adjoint Matching* (AM; Domingo-Enrich et al., 2025) offers a scalable route by reformulating the global SOC problem into the local matching problem.

Extending AM to infinite dimensions requires a rigorous definition of the *first-order adjoint process* that appears in the corresponding infinite-dimensional SOC formulation. Below, we will demonstrate that a specific mathematical concept in SOC called *Stochastic Maximum Principle* (SMP; Bensoussan, 1983) provides existence and uniqueness of the adjoint process via coupled *Forward-Backward SDEs* (FBSDEs; Du and Meng, 2013) system and necessary condition.

Lemma 3 (Stochastic Maximum Principle). *Consider the general infinite-dimensional SOC problem:*

$$\min_{\alpha \in \mathcal{A}} \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T l(t, \mathbf{X}_t^\alpha, \alpha_t) dt + g(\mathbf{X}_T^\alpha) \right], \quad (13)$$

$$\text{s.t. } d\mathbf{X}_t^\alpha = [\mathbf{A}\mathbf{X}_t^\alpha + f(t, \mathbf{X}_t^\alpha, \alpha_t)] dt + \sigma(t, \mathbf{X}_t^\alpha) d\mathbf{W}_t^Q.$$

With fixed initial condition $\mathbf{X}_0^\alpha = \mathbf{x}_0$. Now, let Assumption B.1–B.2 holds and define the Hamiltonian functional $H : [0, T] \times \mathcal{H} \times \mathcal{U} \times \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$:

$$H(t, \mathbf{X}, \alpha, \mathbf{Y}, \mathbf{Z}) := l(t, \mathbf{X}, \alpha) + \langle \mathbf{Y}, f(t, \mathbf{X}, \alpha) \rangle_{\mathcal{H}} + \langle \mathbf{Z}, \sigma(t, \mathbf{X}) \rangle_{\mathcal{H}} \quad (14)$$

Suppose $\mathbf{X}^*$ is the optimally controlled process driven by optimal control α^* . Then, the following first-order adjoint infinite-dimensional backward SDEs has a unique solution:

$$\begin{aligned} d\mathbf{Y}_t^* &= -[\mathcal{A}^\dagger \mathbf{Y}_t^* + D_{\mathbf{x}}H(t, \mathbf{X}_t^*, \alpha_t^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*)] dt \\ &\quad + \mathbf{Z}_t^* d\mathbf{W}_t^Q, \quad \mathbf{Y}_T^* = D_{\mathbf{x}}g(\mathbf{X}_T^*) \end{aligned} \quad (15)$$

Moreover, we get necessary condition of optimal control $\alpha_t^* = \arg \min_{\alpha \in \mathbb{A}} H(t, \mathbf{X}_t^*, \alpha, \mathbf{Y}_t^*, \mathbf{Z}_t^*)$.

Lemma 3 states that, for the general SOC problem (13), there exists a first-order adjoint pairs $(\mathbf{Y}_t, \mathbf{Z}_t)$ solving a backward SDE in (15). This result provides a *necessary condition* for optimality: any control that violates the backward BSDE in (15) or the point-wise minimization of Hamiltonian for all $t \in [0, T]$ cannot be optimal control.

In other words, SMP provides a *point-wise* optimality condition determined by the corresponding Hamiltonian H , allowing optimal control α_t^* to be computed *locally* rather than *path-wise* computation. This principle can be effective, since solving path-wise minimization in (12) is often computationally prohibitive. With this local characterization in hand, we now derive a proper training objective for the infinite-dimensional SOC problem (12) based on the SMP. By applying Lemma 3 into our SOC problem in (12), the corresponding Hamiltonian H is defined by:

$$\begin{aligned} H(t, \mathbf{X}_t^\alpha, \bar{\alpha}_t, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha) &:= \frac{1}{2} \|\bar{\alpha}_t\|_{\mathcal{H}}^2 + \langle \mathbf{Y}_t^\alpha, \sigma_t Q^{1/2} \bar{\alpha}_t \rangle_{\mathcal{H}} \\ &\quad + \langle \mathbf{Z}_t^\alpha, \sigma_t Q^{1/2} \rangle_{\mathcal{H}}, \end{aligned} \quad (16)$$

with a fixed state control pair $(\mathbf{X}_t^\alpha, \bar{\alpha})$. Note that the Hamiltonian H in (16) is convex with respect to the control $\bar{\alpha}_t$, hence any critical point is global minimum. It implies that if $(\mathbf{X}_t^\alpha, \bar{\alpha})$ solves the backward SDEs in (15), then the optimal control is characterized by the first-order stationarity of H :

$$\begin{aligned} \nabla_{\alpha} H(t, \mathbf{X}_t^\alpha, \alpha_t, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha) &= \alpha_t + \langle \mathbf{Y}_t^\alpha, \sigma_t Q^{1/2} \rangle_{\mathcal{H}} \\ &= \nabla_{\alpha} \frac{1}{2} \left\| \alpha_t + \sigma_t Q^{1/2} \mathbf{Y}_t^\alpha \right\|_{\mathcal{H}}^2. \end{aligned} \quad (17)$$

It follows that, with the state variables $(\mathbf{X}_t^\alpha, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha)$ remain fixed for the chosen control $\bar{\alpha}_t$, the minimization over α_t reduces to the convex quadratic in (17), whose unique minimizer is $\alpha_t = -\sigma_t Q^{1/2} \mathbf{Y}_t^\alpha$. Leveraging this observation, we now state our main result:

Proposition 4 (Adjoint Matching in $\mathcal{H}$). *Consider the following infinite-dimensional matching objective:*

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{\mathbb{P}^\alpha} \left[\frac{1}{2} \left\| \alpha^\theta(\mathbf{X}_t^\alpha, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^\alpha \right\|_{\mathcal{H}}^2 \right] dt \quad (18)$$

$$d\mathbf{Y}_t^\alpha = -\mathcal{A}^\dagger \mathbf{Y}_t^\alpha dt + \mathbf{Z}_t^\alpha d\mathbf{W}_t^Q, \quad \mathbf{Y}_T^\alpha = D_{\mathbf{x}}g(\mathbf{X}_T^\alpha). \quad (19)$$

where $\bar{\alpha} := \text{stopgrad}(\alpha^\theta)$. Then, the critical point of $\mathcal{L}$ is optimal control α^* for the SOC problem in (12).

Compared to the original SOC problem (12), SMP-based matching objective in (18) is inherently scalable since it can avoid path-wise backpropagation and reduces training to point-wise convex minimizations with $\mathbf{X}_t^\alpha$ held fixed. This local structure enables straightforward parallelization over time. Note that when $\mathcal{H} := \mathbb{R}^d$ with $Q = \mathbf{I}_d$, then the objective function in (18) recovers the *lean* AM objective in (Domingo-Enrich et al., 2025). In other words, our formulation strictly generalizes AM into the Hilbert spaces.

Yet, for $\mathcal{L}$ in (18), the stochastic term $\mathbf{Z}_t^\alpha$ can incur a computational burden because this term requires simulating the BSDE in (19) to obtain the solution $\mathbf{Y}_t^\alpha$, which markedly increases the cost of each optimization step. Since this BSDE is adapted to the filtration generated by the same Q -Wiener process $\mathbf{W}_t^Q$, we substitute the adjoint process $\mathbf{Y}_t^\alpha$ with its conditional expectation $\mathbb{E}_{\mathbb{P}^\alpha}[\mathbf{Y}_t^\alpha | \mathbf{X}_t^\alpha]$.

Proposition 5 (Unbiased Estimator). *Define the sample-wise adjoint matching objective with conditional expectation $\tilde{\mathbf{Y}}_t^\alpha := \mathbb{E}_{\mathbb{P}^\alpha}[\mathbf{Y}_t^\alpha | \mathbf{X}_t^\alpha]$ for any sample trajectory $\mathbf{X}_t^\alpha \sim \mathbb{P}^\alpha$:*

$$\tilde{\mathcal{L}}(\theta) := \int_0^T \frac{1}{2} \left\| \alpha^\theta(\mathbf{X}_t^\alpha, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^\alpha \right\|_{\mathcal{H}}^2 dt \quad (20)$$

Then, we get unbiased gradient estimator $\frac{d}{d\theta} \mathbb{E}_{\mathbb{P}^\alpha}[\tilde{\mathcal{L}}(\theta)] = \frac{d}{d\theta} \mathcal{L}(\theta)$ with same critical point in (18).

In practice, we can estimate $\tilde{\mathbf{Y}}_t^\alpha$ from the trajectory saved during simulation of $\mathbf{X}_T^\alpha$ for efficient computation. This preserves the same critical point α^* for the SOC problem in (12) while reducing the computational burden because the substitution yields a closed form solution of the BSDE:

$$\begin{aligned} \tilde{\mathbf{Y}}_t^\alpha &= \mathbb{E}_{\mathbb{P}^\alpha}[\mathbf{Y}_T^\alpha | \mathbf{X}_t^\alpha] = e^{-(T-t)\mathcal{A}^\dagger} \mathbb{E}_{\mathbb{P}^\alpha}[D_{\mathbf{x}}g(\mathbf{X}_T^\alpha) | \mathbf{X}_t^\alpha] \\ &\approx e^{-(T-t)\mathcal{A}^\dagger} D_{\mathbf{x}}g(\mathbf{X}_T^\alpha). \end{aligned} \quad (21)$$

Finally, since the optimal path measure $\mathbb{P}^*$ in (6) admits the endpoint disintegration $\mathbb{P}^*(\mathbf{X}_t) = \mathbb{P}^*(\mathbf{X}_T) \mathbb{P}_{t|T}(\mathbf{X}_t | \mathbf{X}_T)$; see Section B.8 for details. Substituting these relations into (18) yields our proposed training objective for the efficient and scalable sampling in function spaces:

$$\mathcal{L}_{\text{FAS}} = \int_0^T \mathbb{E}_{\mathbb{P}_{t|T} \mathbb{P}_T^\alpha} \left[\frac{1}{2} \left\| \alpha^\theta(\mathbf{X}_t, t) + \mathcal{C}_t D_{\mathbf{x}}g(\mathbf{X}_T^\alpha) \right\|_{\mathcal{H}}^2 \right] dt, \quad (22)$$

where $\mathcal{C}_t := \sigma_t Q^{1/2} e^{-(T-t)\mathcal{A}^\dagger}$. Note that the path measure $\mathbb{P}^\alpha$ induced by current control $\bar{\alpha}$ may not satisfy the endpoint disintegration. In this case, we can project the path measure to a conditional path measure that does satisfy the disintegration (Havens et al., 2025).

3.4. Numerical Computation

Because FAS operates on infinite-dimensional path spaces, both training and sampling require appropriate discretization and simulation schemes. This requires specifying the

Algorithm 1 Galerkin SDE simulation

Input: Initial condition $\mathbf{x}_0$, basis functions $\{\phi^{(k)}\}_{k=1}^K$.
 Set initial condition $\mathbf{X}_0^\alpha = \mathbf{x}_0$
for time t **in range** $[0, T]$ **do**
 Estimate control $\alpha_t := \alpha(\mathbf{X}_t^\alpha, t)$
 Sample Gaussian noise $\epsilon \sim \mathcal{N}(0, \mathbf{I}_K)$
 $\{\mathbf{X}_t^{(k)}\}_{k \in K} \leftarrow \text{Galerkin expansion}(\mathbf{X}_t^\alpha)$
 $\{\alpha_t^{(k)}\}_{k \in K} \leftarrow \text{Galerkin expansion}(\alpha_t)$
 for $k = 1, \dots, K$ **do in parallel**
 Euler-Step for each $\mathbf{X}_t^{(k)}$ using (24)
 end for
 $\mathbf{X}_t^\alpha \leftarrow \text{Galerkin truncation}(\{\mathbf{X}_{t+dt}^{(k)}\}_{k \in K})$
end for
Output: $\mathbf{X}_T^\alpha \sim \mathbb{P}_T^\alpha$

operators $\mathcal{A}$ and Q , which are also key modeling choices reflecting the domain structure of a given task. For NES, we adopt Theorem B.1, ensuring well-posed reference dynamics and yielding an explicit RND in Theorem 2. Hence, we restrict to a specific class of operator pairs $(\mathcal{A}, Q)$.

Proposition 6 (Laplacian Operator). *For the negative Laplacian $\mathcal{A} = -\Delta$ on a bounded domain with Dirichlet boundary condition, take $Q = \mathcal{A}^{-s}$ with $s > \frac{d}{2}$. Then, $\mathcal{A}$ and Q satisfy the conditions in Theorem B.1.*

Boundary condition Since the Laplacian Δ admits different realizations depending on the boundary conditions (BCs) (Evans, 2022), we restrict to Dirichlet BCs in this study. While Neumann BCs are common in other domains (Rissanen et al., 2023; Lim et al., 2023b; Park et al., 2024), our goal is conditional path sampling, where Dirichlet BCs are natural because they keep boundary values fixed. For example, in TPS, a transition path between metastable states **A** and **B** is meaningful when these states are fixed as boundary conditions *i.e.*, $\mathbf{X}_t[0] = \mathbf{A}$, $\mathbf{X}_t[L] = \mathbf{B}$. To do so, we use the Dirichlet sine basis for $k = [1, \dots, K]$:

$$\phi^{(k)}[u] = \left(\frac{2}{L}\right)^{1/2} \sin\left(\frac{\pi k u}{L}\right), \quad u \in [0, L]. \quad (23)$$

This basis satisfies $\phi^{(k)}[0] = \phi^{(k)}[L] = 0$, and therefore naturally enforces zero boundary values for any path. In our TPS implementation, we write $\mathbf{X}_t = \mathbf{x}_0 + \mathbf{R}_t$, where $\mathbf{x}_0$ is a fixed reference path connecting **A** to **B** and $\mathbf{R}_t$ is expanded in the above sine-basis in (23). Hence, $\mathbf{R}_t[0] = \mathbf{R}_t[L] = 0$ for all t , so the reconstructed path always satisfies $\mathbf{X}_t[0] = \mathbf{A}$ and $\mathbf{X}_t[L] = \mathbf{B}$ for a path $\mathbf{X}_t[\cdot] : \mathbb{R}_{0>} \rightarrow \mathbb{R}^d$ and for any $t \in [0, T]$. We defer the details of enforcing BCs for TPS under our formulation to Appendix C.2.

Numerical simulation Simulating infinite-dimensional SDEs is a well-studied problem (Lord et al., 2014). Since it

Algorithm 2 Functional Adjoint Sampler

Input: Differentiable cost $g(\mathbf{x})$, parametrized $\alpha_\theta(\mathbf{x}, t)$
 Initialize $\alpha_\theta^{(0)} := 0$
for step m **in** $1, 2, \dots$ **do**
 Sample $t \in \mathcal{U}[0, T]$
 $\bar{\theta} \rightarrow \text{Stopgrad}(\theta)$
 $\mathbf{X}_T^{\bar{\alpha}} \leftarrow \text{Galerkin SDE}(\mathbf{x}_0, \alpha_{\bar{\theta}}, \{\phi^{(k)}\}_{k=1}^K)$
 ▷ Algorithm 1 for Galerkin SDE
 Compute adjoint $\mathbf{Y}_t^{\bar{\alpha}}$ using (21)
 Sample $\mathbf{X}_{t|T}^{\bar{\alpha}} \sim \mathbb{P}_{t|T}(\cdot | \mathbf{X}_T^{\bar{\alpha}})$
 ▷ cf. Section C.1 for bridge sampling
 Update control $\alpha_\theta^{(m)}$ by minimizing (22)
end for
Output: $\alpha_\theta^{(M)} \approx \alpha^*$

often requires a finite-dimensional approximation, we use a Galerkin truncation scheme (Shardlow, 1999) by approximating the function with K basis functions $\{\phi^{(k)}\}_{k=1}^K$:

$$\begin{aligned}
 \mathbf{X}_t &\approx \sum_{k=1}^K \mathbf{X}_t^{(k)} \phi^{(k)}, & (\text{Galerkin truncation}) \\
 \mathbf{X}_t^{(k)} &:= \langle \mathbf{X}_t, \phi^{(k)} \rangle_{\mathcal{H}} = \int_0^L \mathbf{X}_t[u] \phi^{(k)}[u] du. & (\text{Galerkin expansion})
 \end{aligned}$$

Hence, with the finite eigen-system $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R} \times \mathcal{H} : k \in [1, \dots, K]\}$ of Dirichlet Laplacian Δ , the $\mathcal{H}$ -valued controlled SDE (11) reduces to the K -dimensional coefficient SDEs after Galerkin expansion (GE):

$$\begin{aligned}
 d\mathbf{X}_t^{(k)} &= \left[-\lambda^{(k)} \mathbf{X}_t^{(k)} + \sigma_t(\lambda^{(k)})^{-s/2} \langle \alpha(\mathbf{X}_t, t), \phi^{(k)} \rangle_{\mathcal{H}} \right] dt \\
 &\quad + \sigma_t(\lambda^{(k)})^{-s/2} d\mathbf{W}_t^{(k)},
 \end{aligned} \quad (24)$$

where $\{\mathbf{W}^{(k)}\}_{k=1}^K$ are independent standard Brownian motions. In particular, since (24) is a finite-dimensional mode-wise diagonal system, we can simulate each Galerkin mode $\mathbf{X}_t^{(k)}$ independently using a standard numerical SDE integrator, as summarized in Algorithm 1. At each time step, we first expand the current path $\mathbf{X}_t^\alpha$ and control α onto the first $|K|$ basis functions, evolve the resulting coefficient SDEs in parallel, and then reconstruct the path by the corresponding $|K|$ -mode Galerkin truncation. This Galerkin SDE simulation is used in the FAS training procedure, summarized in Algorithm 2. Appendix C.3 provides additional details for the TPS instantiation and corresponding DST-based Galerkin SDE simulation².

²Other parameterizations can be applied to different input domains with different BCs, such as image and PDE domains (a rectangular domain $\mathbb{R}^2$) or molecular conformations (a graph domain $\mathcal{G}$), as discussed in Appendix C.5.

4. Related Works

Infinite-dimensional generative models Many recent works extend finite-dimensional generative models to infinite-dimensional function spaces. Score-based diffusion models (Song et al., 2021) have been extended to function space (Lim et al., 2023a;b; Pidstrigach et al., 2023; Franzese et al., 2023). Later, matching-based generative models follow the same direction. For example, flow matching (Lipman et al., 2023) has been generalized to function space (Kerrigan et al., 2024), and Schrödinger bridge matching (Shi et al., 2024) has been adapted as well (Park et al., 2024). See (Franzese and Michiardi, 2025) for a survey on this topic. However, these algorithms operate in a *sample-to-sample* setting where paired samples from both distributions are available for training. To the best of our knowledge, scalable function-space generative modeling for *sample-to-energy* such as targeting sampling from a Gibbs distribution $\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})}$ in (6), has not been investigated.

Diffusion sampler Diffusion samplers address the *sample-to-energy* setting, where the goal is to draw samples from the Gibbs distribution $\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})}$. Based on SOC theory, this line of work reformulates sampling as an optimization problem and methods for learning such samplers have advanced rapidly (Zhang and Chen, 2022; Vargas et al., 2023; Berner et al., 2024; Zhu et al., 2026). Due to the target is specified only through a complex energy functional U , optimization tends to be computationally demanding. This challenge motivates adjoint-based samplers (Havens et al., 2025; Liu et al., 2026; Guo et al., 2026), which avoid simulation-based targets and yield a method that scales better. However, existing methods are limited to finite-dimensional settings, and extensions to function space remain unexplored, despite the appeal and practical importance.

Transition path sampling TPS aims to sample reactive trajectories connecting metastable states in molecular systems (Dellago et al., 1998; Vanden-Eijnden et al., 2010). Classical TPS and enhanced sampling methods construct Markov chains over trajectories or bias molecular dynamics to increase the probability of rare transitions (Izrailev et al., 1999), but they often require long simulations, carefully chosen collective variables, or expensive accept-reject procedures. Recent learning-based methods mitigate these issues by learning controlled dynamics or variational path distributions for TPS (Holdijk et al., 2022; Seong et al., 2025; Blessing et al., 2026). A related line of work instead optimizes discretized path functionals to find minimum-energy or high-probability transition paths (Petersen and Covino, 2025; Ramakrishnan et al., 2025; Raja et al., 2025). We leave introduction of transition path sampling and additional related work in Appendix A.3. Compared to prior works that operate after fixing a finite-discretization, FAS formulates TPS as sampling from a path-space Gibbs mea-

sure in (6) in Hilbert space $\mathcal{H}$. This Hilbert-space formulation separates the target path measure from its numerical discretization, allowing the learned sampler to be reused across different path resolutions.

5. Experiments

5.1. Experimental Setup

We evaluate FAS on three benchmarks for TPS. We start with a synthetic potential, then move to two molecular systems. For the synthetic case, following (Du et al., 2024), we use the standard TPS benchmark, the Müller–Brown potential. For molecular systems, we consider alanine dipeptide and chignolin, with 22 and 138 atoms in 3D, respectively. All molecular simulations use Cartesian coordinates, and we compute potential energy with OpenMM (Eastman et al., 2023). Further experimental details appear in Section C.

Baselines and evaluation We compare FAS with: unbiased MD (UMD) at elevated temperature, steered MD (SMD) with forces on selected collective variables, and deep learning methods PIPS (Holdijk et al., 2022), TPS-DPS (Seong et al., 2025), TR-LV (Blessing et al., 2026), and DL (Du et al., 2024). Note that DL uses a different training environments, so we report it only on the synthetic benchmark. We evaluate with metrics that capture reliability and probabilistic fidelity:

- **RMSD**: We first align the heavy atoms of the endpoint $\mathbf{X}[L]$ to the target state $\mathbf{B}$ with the Kabsch algorithm (Kabsch, 1976). Then calculate the root mean square deviation (**RMSD**) between the heavy atom positions of $\mathbf{X}[L]$ and $\mathbf{B}$.
- **THP**: It is the indicator-based success rate of trajectories reaching the target metastable state $\mathbf{B}$. Given endpoints $\{\mathbf{X}^{(i)}[L]\}_{i=1}^N$ from N paths, we define **THP** as:

$$\mathbf{THP} = 100 \times \frac{|\{i : \mathbf{X}^{(i)}[L] \in \mathcal{B}\}|}{N}. \quad (25)$$

where $\mathcal{B} = \{\mathbf{X}[L] \mid \|\mathbf{X}[L] - \mathbf{B}\| < \epsilon\}$ and choose ϵ following (Seong et al., 2025).

- **ETS**: It measure how well the method find probable transition states when crossing the energy barrier. Given a transition path $\mathbf{X}$ of length L that reaches $\mathcal{B}$, we define

$$\mathbf{ETS}(\mathbf{X}) = \max_{u \in [0, L]} V(\mathbf{X}[u]). \quad (26)$$

On the synthetic potential, we follow (Du et al., 2024) and additionally report the log likelihood (**LLK**) of a transition path distribution (TPD) defined in (27).

Table 1. Comparison on synthetic potential. We report the metrics averaged over 64 paths. Best results are highlighted.

Methods	THP (%) $\uparrow$	ETS (J) $\downarrow$	LLK (J) $\uparrow$
MCMC	-	-13.77 $\pm$ 16.43	-
DL (Du et al., 2024)	82	-14.93 $\pm$ 12.61	10.42 $\pm$ 0.42
FAS (Ours)	100	-36.70 $\pm$ 3.09	11.48 $\pm$ 0.05




Figure 1. TPS for synthetic potential. Sampled transition path on energy landscape.

Energies We adopt the TPD described in (Holdijk et al., 2022; Du et al., 2024) for the Müller-Brown potential:

$$\mathbb{P}(\mathbf{X}) = \mathbf{1}_A(\mathbf{X}[0])\mathbf{1}_B(\mathbf{X}[L]) \quad (27)$$

$$\prod_{u=0}^{L-1} \mathcal{N}(\mathbf{X}[u+1]|\mathbf{X}[u] - \frac{1}{\gamma_m} \nabla_{\mathbf{X}} V(\mathbf{X}[u])\delta_u, \frac{2k_B T}{\gamma_m} \mathbf{I}_{3N} \delta_u),$$

where both $\mathbf{1}_A$ and $\mathbf{1}_B$ are indicator functions and V is potential function. For molecular systems, we observed that TPS in (27) led to unstable training because it relies on local potential gradients. We therefore derive a TPD based on the Feynman-Kac formula (Del Moral, 2011), which removes the reliance on local gradients:

$$\mathbb{P}(\mathbf{X}) = \mathbf{1}_A(\mathbf{X}[0])\mathbf{1}_B(\mathbf{X}[L]) \quad (28)$$

$$\prod_{u=0}^{L-1} e^{-\frac{1}{k_B T} V(\mathbf{X}[u])} \mathcal{N}(\mathbf{X}[u+1]|\mathbf{X}[u], \frac{2k_B T}{\gamma_m} \mathbf{I}_{3N} \delta_u).$$

We provide further discussions on the choice of TPD in Appendix A.6.

5.2. Experimental Results

Müller Brown potential The Müller Brown potential is a widely adopted benchmark for TPS and features three local minima. Following the setup in (Du et al., 2024), we aim to sample transition paths connecting the top left state to the bottom right state. Table 1 reports the results on the the Müller Brown potential. FAS attains lower ETS and higher LLK than baselines, showing it yields reliable paths that remain in low energy valleys as shown in Figure 1.

Alanine dipeptide We evaluate FAS on alanine dipeptide, a peptide with two alanine residues, and target transitions




Figure 2. TPS on alanine dipeptide. (a) Projected transition paths sampled by FAS on the conformational landscape as a function of CVs. The saddle point that separates the two meta-stable states is depicted as $\star$ (orange star) and the highest energy along paths sampled from FAS depicted as $\bullet$ (blue circle). (b) Potential energy plot. Grey box highlights the transition-state region where paths reach the ETS. (c) Visualization of molecule transition trajectory for given meta-stable states C5 (far left) and C7ax (far right).

between meta-stable states C5 and C7ax. We use torsion angles ϕ and ψ as collective variables (CVs) from atomic coordinates. Table 2 summarizes alanine dipeptide results. Notably, our function space approaches achieve perfect THP with zero RMSD due to the Dirichlet boundary formulation enforces these endpoints by construction. Perfect THP implies that all sampled paths are reactive, which maximizes compute use and prevents bias in transition rate and mechanism estimates hence lowers variance of path statistics. Moreover, FAS also attain significantly lower ETS than baselines, implying the path generated by FAS is reliable.

Moreover, as shown in Figure 2a, the free energy surface over these CVs shows a high barrier with a saddle point. Reliable reactive paths are expected to cross near this saddle because it is the lowest barrier route. The highest energy along paths sampled from FAS occurs close to the saddle, which supports that the sampled paths are physically meaningful transition paths. Figure 2c provides a visualization of configuration for the molecule transition.

Chignolin Finally, we evaluate sampling of the folding dynamics of the small protein Chignolin. As shown in Table 2, FAS achieves the lowest ETS. Because of chignolin has a higher dimensional and more rugged landscape with a narrow folding funnel and multiple intermediates basin, baselines show larger RMSD and lower THP than on alanine dipeptide. Baselines that rely on gradient driven or steered updates often traverse strained conformations. This inflates RMSD even after Kabsch alignment and lowers the chance of reaching the folded basin within the time horizon. In contrast, FAS enforces the endpoints throughout and searches

Table 2. Comparison between baselines on the molecular transition path sampling of the alanine dipeptide and chignolin. We report the metrics averaged over 64 paths. Best results are highlighted.

Method	Alanine Dipeptide			Chignolin		
	RMSD (Å ↓)	THP (% ↑)	ETS (kJ/mol ↓)	RMSD (Å ↓)	THP (% ↑)	ETS (kJ/mol ↓)
UMD	1.19±0.03	6.25	812.47±148.80	7.23±0.93	1.56	388.17
SMD (Izrailev et al., 1999)	0.56±0.27	54.69	78.40±12.76	0.85±0.24	34.38	179.52±138.87
PIPS (Holdijk et al., 2022)	0.66±0.15	43.75	28.17±10.86	4.66±0.17	0.00	-
TPS-DPS (Seong et al., 2025)	0.47±0.18	39.58±28.13	46.34±10.16	1.06±0.08	25.00±10.69	-189.81±23.01
TR-LV (Blessing et al., 2026)	0.29±0.03	61.25±4.05	49.11±5.84	0.90±0.01	43.95±5.64	-303.98±28.65
FAS (Ours)	0.00±0.00	100.00±0.00	-29.46±3.87	0.00±0.00	100.00±0.00	-360.85±53.71




Figure 3. **TPS on Chignolin.** (a) Potential energy plot. (b) Visualization of protein folding trajectory from **unfolded** (far left) to **folded** (far right). Grey box highlights the transition-state region where trajectories reach the **ETS**.

only within admissible paths. This boundary enforcement removes endpoint drift and avoids late corrective pulls. As a result, transitions are smoother and lead to lower **ETS**. We illustrate visualization of the folding dynamics in Figure 3b.

5.3. Discussions

Discretization invariance sampling Infinite dimensional formulation enables inference on paths at arbitrary time discretizations. Table 3 reports **ETS** across varying resolutions. Specifically, we train α^θ with a discretization of implicit time interval $[0, 1]$ into 100 uniform time steps (denoted $\times 1$) and evaluate on unseen, more finer time discretization (denoted $\times 2-100$). Across both potentials, FAS maintains nearly unchanged **ETS** on unseen resolutions, which indicate discretization-agnostic behavior. In particular, for the synthetic Müller Brown potential, **ETS** and **LLK** remain stable even when the resolution increases by $\times 100$, where $[0, 1]$ is discretized into 10^4 steps. Compared to synthetic potential, alanine dipeptide exhibits higher **ETS** variability on unseen discretizations because molecular potentials are highly sensitive to configuration. We leave additional discussions on discretization invariance sampling and further ablation studies in Section C.8.

6. Conclusion

We introduced the Functional Adjoint Sampler (FAS), a scalable diffusion sampler for Gibbs-type distributions on Hilbert spaces. FAS generalizes adjoint matching algorithms (Domingo-Enrich et al., 2025; Havens et al., 2025)

Table 3. **ETS** across various unseen discretization levels. **Disc** represents how fine the generated path is, where $\times r$ multiplies the base number of discretization by r .

Müller Brown			Alanine Dipeptide	
Disc	ETS(kJ/mol ↓)	LLK (↑)	Disc	ETS(kJ/mol ↓)
$\times 1$	-36.70±3.09	11.48±0.05	$\times 1$	-29.46±3.87
$\times 10$	-38.55±1.41	11.65±0.02	$\times 2$	-28.72±41.52
$\times 100$	-37.83±3.17	11.72±0.09	$\times 4$	-33.58±35.59

to infinite-dimensional Hilbert spaces, enabling efficient sampling over functional representations such as paths evaluated on a one-dimensional time grid, and it demonstrates superior performance on path sampling tasks. Despite these promising results, we focused on a one dimensional domain. However, FAS readily extends to other tasks such as PDE generation and graph based domains, and we leave a comprehensive study for future work.

Acknowledgments

The authors would like to thank Benjamin Kurt Miller, Juno Nam, Yuanqi Du, and Ricky T. Q. Chen for the helpful discussions and comments. This work was partly supported by Institute for Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No.RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST), No.RS-2024-00509279, Fundamental Research in Artificial Intelligence, No.RS-2022-II220713, Meta-learning Applicable to Real-world Problems).

Impact Statement

This paper aims to advance the field of machine learning. Our work may have societal implications, but we do not identify any that warrant specific emphasis here.

References

- Tara Akhound-Sadegh, Jarrid Rector-Brooks, Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated denoising energy matching for sampling from boltzmann densities. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=gVjMwLDFoQ>.
- Jodi E Basner and Steven D Schwartz. How enzyme dynamics helps catalyze a reaction in atomic detail: a transition path sampling study. *Journal of the American Chemical Society*, 127(40):13822–13831, 2005.
- Alain Bensoussan. Stochastic maximum principle for distributed parameter systems. *Journal of the Franklin Institute*, 315(5-6):387–406, 1983.
- Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=oYIjw37pTP>.
- Alexandros Beskos, Gareth Roberts, Andrew Stuart, and Jochen Voss. Mcmc methods for diffusion bridges. *Stochastics and Dynamics*, 8(03):319–350, 2008.
- Denis Blessing, Julius Berner, Lorenz Richter, Carles Domingo-Enrich, Yuanqi Du, Arash Vahdat, and Gerhard Neumann. Trust region constrained measure transport in path space for stochastic optimal control and inference. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026. URL <https://openreview.net/forum?id=6RlbOEcos4>.
- Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annual review of physical chemistry*, 53(1):291–318, 2002.
- René Carmona. *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*. SIAM, 2016.
- Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos Theodorou. Deep momentum multi-marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, 36:57058–57086, 2023.
- Raphaël Chetrite and Hugo Touchette. Nonequilibrium markov processes conditioned on large deviations. In *Annales Henri Poincaré*, volume 16, pages 2005–2057. Springer, 2015.
- Jaemoo Choi, Yongxin Chen, Molei Tao, and Guan-Horng Liu. Non-equilibrium annealed adjoint sampler. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026. URL <https://openreview.net/forum?id=ay7WDSq0Kb>.
- Nicolas Chopin, Andras Fulop, Jeremy Heng, and Alexandre H. Thiery. Computational doob’s h-transforms for online filtering of discretely observed diffusions, 2023.
- G. Da Prato and J. Zabczyk. *Second Order Partial Differential Equations in Hilbert Spaces*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2002. ISBN 9780521777292. URL <https://books.google.co.kr/books?id=TUt-SAYofVMC>.
- G. Da Prato and J. Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2014. ISBN 9781107055841. URL <https://books.google.co.kr/books?id=bxkmAwAAQBAJ>.
- G. Da Prato, J. Zabczyk, P.J. Zabczyk, J.W.S. Cassels, and S.P.G.N.J. Hitchin. *Ergodicity for Infinite Dimensional Systems*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 1996. ISBN 9781107362499. URL <https://books.google.co.kr/books?id=gEe2oQEACAAJ>.
- Giuseppe Da Prato and Jerzy Zabczyk. Differentiability of the feynman-kac semigroup and a control application. *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti Lincei. Matematica e Applicazioni*, 8(3):183–188, 1997.
- P. Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Probability and Its Applications. Springer New York, 2011. ISBN 9781441919021. URL <https://books.google.co.kr/books?id=JN9NcgAACAAJ>.
- Christoph Dellago, Peter G Bolhuis, Félix S Csajka, and David Chandler. Transition path sampling and the calculation of rate constants. *The Journal of chemical physics*, 108(5):1964–1977, 1998.
- Christopher M Dobson. Protein folding and misfolding. *Nature*, 426(6968):884–890, 2003.
- Carles Domingo-Enrich, Michal Drozdal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning

- flow and diffusion generative models with memory-less stochastic optimal control. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=xQBRrtQM8u>.
- J.L. Doob. *Classical Potential Theory and Its Probabilistic Counterpart*. Classics in Mathematics. Springer Berlin Heidelberg, 2001. ISBN 9783540412069. URL <https://books.google.com/books?id=vXdVqv8HHKMC>.
- Kai Du and Qingxin Meng. A maximum principle for optimal control of stochastic evolution equations. *SIAM Journal on Control and Optimization*, 51(6):4343–4362, 2013.
- Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Aspuru-Guzik, and Kirill Neklyudov. Doob’s lagrangian: A sample-efficient variational approach to transition path sampling. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=ShJWT0n7kX>.
- Detlef Dürr and Alexander Bach. The onsager-machlup function as lagrangian for the most probable path of a diffusion process. *Communications in Mathematical Physics*, 60(2):153–170, 1978.
- Peter Eastman, Raimondas Galvelis, Raúl P. Peláez, Charles R. A. Abreu, Stephen E. Farr, Emilio Gallicchio, Anton Gorenko, Michael M. Henry, Frank Hu, Jing Huang, Andreas Krämer, Julien Michel, Joshua A. Mitchell, Vijay S. Pande, João PGLM Rodrigues, Jaime Rodriguez-Guerra, Andrew C. Simmonett, Sukrit Singh, Jason Swails, Philip Turner, Yuanqing Wang, Ivy Zhang, John D. Chodera, Gianni De Fabritiis, and Thomas E. Markland. Openmm 8: Molecular dynamics simulation with machine learning potentials. *arXiv preprint arXiv:2310.03121*, 2023.
- Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2022.
- Giorgio Fabbri, Fausto Gozzi, and Andrzej Swiech. Stochastic optimal control in infinite dimension. *Probability and Stochastic Modelling*. Springer, 2017.
- W.H. Fleming and H.M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Applications of mathematics. Springer-Verlag, 1993. ISBN 9780387979274. URL https://books.google.co.kr/books?id=5k4_AQAAIAAJ.
- Giulio Franzese and Pietro Michiardi. Generative diffusion models in infinite dimensions: a survey. *Philosophical Transactions A*, 383(2299):20240322, 2025.
- Giulio Franzese, Giulio Corallo, Simone Rossi, Markus Heinonen, Maurizio Filippone, and Pietro Michiardi. Continuous-time functional diffusion processes. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=VPPrir0p5b6>.
- Marco Fuhrman. A class of stochastic optimal control problems in hilbert spaces: B-sdes and optimal control laws, state constraints, conditioned processes. *Stochastic processes and their applications*, 108(2):263–298, 2003.
- Wei Guo, Yuchen Zhu, Xiaochen Du, Juno Nam, Yongxin Chen, Rafael Gómez-Bombarelli, Guan-Hong Liu, Molei Tao, and Jaemoo Choi. Discrete adjoint schrödinger bridge sampler. *arXiv preprint arXiv:2602.08243*, 2026.
- Martin Hairer, Andrew M Stuart, and Jochen Voss. Sampling conditioned diffusions. *Trends in stochastic analysis*, 353:159–186, 2009.
- Aaron J Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Daniel S. Levine, Brandon M Wood, Bin Hu, Brandon Amos, Brian Karner, Xiang Fu, Guan-Hong Liu, and Ricky T. Q. Chen. Adjoint sampling: Highly scalable diffusion samplers via adjoint matching. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=6EglOrHmg2>.
- Lars Holdijk, Yuanqi Du, Priyank Jaini, Ferry Hooft, Bernd Ensing, and Max Welling. Path integral stochastic optimal control for sampling transition paths. In *ICML 2022 2nd AI for Science Workshop*, 2022.
- Sergei Izrailev, Sergey Stepaniants, Barry Isralewitz, Dorina Kosztin, Hui Lu, Ferenc Molnar, Willy Wriggers, and Klaus Schulten. Steered molecular dynamics. In *Computational Molecular Dynamics: Challenges, Methods, Ideas: Proceedings of the 2nd International Symposium on Algorithms for Macromolecular Modelling, Berlin, May 21–24, 1997*, pages 39–65. Springer, 1999.
- Hannes Jónsson, Greg Mills, and Karsten W Jacobsen. Nudged elastic band method for finding minimum energy paths of transitions. In *Classical and quantum dynamics in condensed phase simulations*, pages 385–404. World Scientific, 1998.
- Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of Crystallography*, 32(5):922–923, 1976.
- Hilbert J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.

- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=k7FuTOWMOc7>.
- Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Functional flow matching. In *International Conference on Artificial Intelligence and Statistics*, pages 3934–3942. PMLR, 2024.
- Jae Hyun Lim, Nikola B Kovachki, Ricardo Baptista, Christopher Beckham, Kamyar Azizzadenesheli, Jean Kossaifi, Vikram Voleti, Jiaming Song, Karsten Kreis, Jan Kautz, et al. Score-based diffusion models in function space. *arXiv preprint arXiv:2302.07400*, 2023a.
- Sungbin Lim, Eunbi Yoon, Taehyun Byun, Taewon Kang, Seungwoo Kim, Kyungjae Lee, and Sungjoon Choi. Score-based generative modeling through stochastic evolution equations in hilbert spaces. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b. URL <https://openreview.net/forum?id=GrElRvXnEj>.
- Kresten Lindorff-Larsen, Stefano Piana, Kim Palmo, Paul Maragakis, John L Klepeis, Ron O Dror, and David E Shaw. Improved side-chain torsion potentials for the amber ff99sb protein force field. *Proteins: Structure, Function, and Bioinformatics*, 78(8):1950–1958, 2010.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized schrödinger bridge. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=fp33Nsh005>.
- Guan-Horng Liu, Jaemoo Choi, Yongxin Chen, Benjamin Kurt Miller, and Ricky T. Q. Chen. Adjoint schrödinger bridge sampler. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026. URL <https://openreview.net/forum?id=rMhQB1hh4c>.
- Gabriel J Lord, Catherine E Powell, and Tony Shardlow. *An introduction to computational stochastic PDEs*, volume 50. Cambridge University Press, 2014.
- James A Maier, Carmenza Martinez, Koushik Kasavajhala, Lauren Wickstrom, Kevin E Hauser, and Carlos Simmerling. ff14sb: improving the accuracy of protein side chain and backbone parameters from ff99sb. *Journal of chemical theory and computation*, 11(8):3696–3713, 2015.
- Hai Nguyen, Daniel R Roe, and Carlos Simmerling. Improved generalized born solvent model parameters for protein simulations. *Journal of chemical theory and computation*, 9(4):2020–2034, 2013.
- Frank Noé, Christof Schütte, Eric Vanden-Eijnden, Lothar Reich, and Thomas R Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences*, 106(45):19011–19016, 2009.
- Byoungwoo Park, Jungwon Choi, Sungbin Lim, and Juho Lee. Stochastic optimal control for diffusion bridges in function spaces. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=WyQW4G57Zd>.
- Byoungwoo Park, Hyungi Lee, and Juho Lee. Amortized control of continuous state space feynman-kac model for irregular time series. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=8zJR0n6k5v>.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Magnus Petersen and Roberto Covino. PINN-MEP: Continuous neural representations for minimum energy path discovery in molecular systems. In *Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025. URL <https://openreview.net/forum?id=5DbwuFklkO>.
- Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Particle denoising diffusion sampler. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=vMUnnS4OWC>.
- Jakiw Pidstrigach, Youssef Marzouk, Sebastian Reich, and Sven Wang. Infinite-dimensional diffusion models for function spaces. *arXiv e-prints*, pages arXiv–2302, 2023.

- Michael Plainer, Hannes Stärk, Charlotte Bunne, and Stephan Günnemann. Transition path sampling with boltzmann generator-based mcmc moves. *arXiv preprint arXiv:2312.05340*, 2023.
- Md Ashiqur Rahman, Zachary E Ross, and Kamyar Aziz-zadenesheli. U-NO: U-shaped neural operators. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=j3oQF9coJd>.
- Sanjeev Raja, Martin Sipka, Michael Psenka, Tobias Kreiman, Michal Pavelka, and Aditi S. Krishnapriyan. Action-minimization meets generative modeling: Efficient transition path sampling with the onsager-machlup functional. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=QwoGfQzuMa>.
- Kalyan Ramakrishnan, Lars L. Schaaf, Chen Lin, Guan-grun Wang, and Philip Torr. Implicit neural representations for chemical reaction paths. *The Journal of Chemical Physics*, 163(3):034109, 07 2025. ISSN 0021-9606. doi: 10.1063/5.0267023. URL <https://doi.org/10.1063/5.0267023>.
- Severi Rissanen, Markus Heinonen, and Arno Solin. Generative modelling with inverse heat dissipation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=4PJUBT9f201>.
- Kiyoung Seong, Seonghyun Park, Seonghwan Kim, Woo Youn Kim, and Sungsoo Ahn. Transition path sampling with improved off-policy training of diffusion path samplers. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=WQV9kB1qSU>.
- Tony Shardlow. Numerical methods for stochastic parabolic pdes. *Numerical functional analysis and optimization*, 20(1-2):121–145, 1999.
- Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.
- Søren Smidstrup, Andreas Pedersen, Kurt Stokbro, and Hannes Jónsson. Improved initial guess for minimum energy path calculations. *The Journal of chemical physics*, 140(21), 2014.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PXTIG12RRHS>.
- Gabriele C Sosso, Ji Chen, Stephen J Cox, Martin Fitzner, Philipp Pedevilla, Andrea Zen, and Angelos Michaelides. Crystal nucleation in liquids: Open questions and future challenges in molecular dynamics simulations. *Chemical reviews*, 116(12):7078–7116, 2016.
- Andrew M Stuart, Jochen Voss, and Petter Wilberg. Conditional path sampling of sdes and the langevin mcmc method. *Communications in Mathematical Sciences*, 2004.
- Taiji Suzuki. Generalization bound of globally optimal non-convex neural network training: Transportation map estimation by infinite dimensional langevin dynamics. *Advances in Neural Information Processing Systems*, 33: 19224–19237, 2020.
- Evangelos A. Theodorou, George I. Boutselis, and Kaivalya Bakshi. Linearly solvable stochastic optimal control for infinite-dimensional systems. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4110–4116, 2018. doi: 10.1109/CDC.2018.8618948.
- Pratyush Tiwary, Vittorio Limongelli, Matteo Salvalaglio, and Michele Parrinello. Kinetics of protein–ligand unbinding: Predicting pathways, rates, and rate-limiting steps. *Proceedings of the National Academy of Sciences*, 112(5):E386–E391, 2015.
- Belinda Tzen and Maxim Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 3084–3114. PMLR, 25–28 Jun 2019. URL <https://proceedings.mlr.press/v99/tzen19a.html>.
- Eric Vanden-Eijnden et al. Transition-path theory and path-finding algorithms for the study of rare events. *Annual review of physical chemistry*, 61:391–420, 2010.
- Francisco Vargas, Andrius Ovsianas, David Fernandes, Mark Girolami, Neil D Lawrence, and Nikolas Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):3, 2023.
- Yuyang Wang, Ahmed AA Elhag, Navdeep Jaitly, Joshua M Susskind, and Miguel Ángel Bautista. Swallowing the bitter pill: Simplified scalable conformer generation. In *International Conference on Machine Learning*, pages 50400–50418. PMLR, 2024.

Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=_uCb2ynRu7Y.

Yuchen Zhu, Wei Guo, Jaemoo Choi, Guan-Horng Liu, Yongxin Chen, and Molei Tao. MDNS: Masked diffusion neural sampler via stochastic optimal control. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026. URL <https://openreview.net/forum?id=xIH95kXNR2>.

A. Overview of Relevant Concepts

We first give a concise overview of core probabilistic notions in infinite dimensional spaces, then describe stochastic differential equations on Hilbert spaces (Da Prato and Zabczyk, 2002; 2014).

A.1. Gaussian Measure and Generalized Wiener Processes.

Gaussian measures are central because the infinite dimensional setting lacks a Lebesgue reference measure, and a Gaussian with a trace class covariance provides a natural baseline that keeps stochastic integrals square integrable and provides explicit RND, both essential for our framework.

Gaussian Measure. Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space and $(\mathcal{H}, B(\mathcal{H}))$ be a measurable state spaces where $\mathcal{H}$ be a separable Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norm $\|\cdot\|_{\mathcal{H}}$ and $B(\mathcal{H})$ is Borel measure on $\mathcal{H}$. We consider the $\mathcal{H}$ -valued random variable $\mathbf{X} : \Omega \rightarrow \mathcal{H}$. Now, define the push-forward measure $\nu := \mathbf{X}_{\#}\mathbb{Q}$ by $\nu(\mathbf{B}) = \mathbb{Q}(\mathbf{X}^{-1}(\mathbf{B}))$ for every Borel set $\mathbf{B} \subset B(\mathcal{H})$. We call ν Gaussian if every $\mathbb{R}$ -valued projection $\langle u, \mathbf{X} \rangle_{\mathcal{H}}$ is a Gaussian random variable for all $u \in \mathcal{H}$. When $\mathcal{H} = \mathbb{R}^d$ with the standard inner product, this definition reproduces the usual Gaussian distribution.

For a Gaussian measure ν , there exists a unique mean function $\mathbf{m} \in \mathcal{H}$ that satisfies:

$$\mathbf{m} = \int_{\mathcal{H}} \mathbf{X} d\nu, \quad \langle \mathbf{m}, u \rangle_{\mathcal{H}} = \mathbb{E}_{\nu} [\langle \mathbf{X}, u \rangle_{\mathcal{H}}], \quad \forall u \in \mathcal{H}. \quad (29)$$

Moreover, there exists a unique bounded linear operator $Q : \mathcal{H} \rightarrow \mathcal{H}$ such that

$$Qu = \int_{\mathcal{H}} \langle \mathbf{X}, u \rangle_{\mathcal{H}} \mathbf{X} d\nu - \langle \mathbf{m}, u \rangle_{\mathcal{H}} \mathbf{m}, \quad \forall u \in \mathcal{H}. \quad (30)$$

A Gaussian measure is fully specified by its mean and covariance operator. If the mean is zero, so that $\nu = \mathcal{N}(0, Q)$, we say that ν is *centred*. The Fernique's integrability theorem (Da Prato and Zabczyk, 2014, Theorem 2.7) states that for any centred Gaussian measure ν on a separable Banach space $\mathcal{B}$ (in particular on $\mathcal{H}$) there exists $r > 0$ such that $\int_{\mathcal{H}} \exp(r \|\mathbf{X}\|_{\mathcal{H}}^2) d\nu < \infty$. Consequently, it implies that every polynomial moment is finite *i.e.*, $\mathbb{E}_{\nu} [\|\mathbf{X}\|_{\mathcal{H}}^2] < \infty$. Using definition (30) and an arbitrary orthonormal basis of $\mathcal{H}$, $\{\phi^{(k)}\}_{k \geq 1}$ for each $\phi^{(k)} \in \mathcal{H}$ for every $k \in \mathbb{N}$:

$$\text{Tr}(Q) = \sum_{k \geq 1} \langle Q\phi^{(k)}, \phi^{(k)} \rangle_{\mathcal{H}} = \sum_{k \geq 1} \mathbb{E}_{\nu} [\langle \mathbf{X}, \phi^{(k)} \rangle_{\mathcal{H}}^2] = \mathbb{E}_{\nu} [\|\mathbf{X}\|_{\mathcal{H}}^2] < \infty. \quad (31)$$

This implies that the symmetric and non-negative covariance operator Q is trace class, *i.e.*, $\text{Tr}(Q) < \infty$, hence implies that Q is compact. Hence, there exists the eigen-system $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R} \times \mathcal{H} : k \in \mathbb{N}\}$, where $\{\phi^{(k)}\}_{k \geq 1}$ is orthonormal eigen-basis and $\{\lambda^{(k)}\}_{k \geq 1}$ is positive eigenvalues such that

$$Q\phi^{(k)} = \lambda^{(k)}\phi^{(k)}, \quad \text{and} \quad \text{Tr}(Q) = \sum_{k \geq 1} \langle Q\phi^{(k)}, \phi^{(k)} \rangle_{\mathcal{H}} = \sum_{k \geq 1} \lambda^{(k)} < \infty. \quad (32)$$

We now extend the notion of a Wiener process to an infinite dimensional SDE.

Wiener Processes on Hilbert Spaces. A natural way to generalize the Wiener process to a Hilbert space $\mathcal{H}$ is through the *cylindrical Wiener process*. It is introduced as the formal series of independent $\mathbb{R}$ -valued Wiener processes $\{\mathbf{W}^{(k)}\}_{k \geq 1}$ along an orthonormal basis $\{\phi^{(k)}\}_{k \geq 1}$:

$$\mathbf{W}_t = \sum_{k \geq 1} \mathbf{W}_t^{(k)} \phi^{(k)}, \quad t \geq 0. \quad (33)$$

However, the series $\mathbf{W}_t$ dose not converge in $\mathcal{H}$ because

$$\mathbb{E} [\|\mathbf{W}_t\|_{\mathcal{H}}^2] = \sum_{k \geq 1} \mathbb{E} [(\mathbf{W}_t^{(k)})^2] = t \sum_{k \geq 1} 1 = \infty, \quad (34)$$

thereby $\mathbf{W}_t$ fails to be an $\mathcal{H}$ -valued random variable. Hence, it can be used only under additional smoothing to obtain $\mathcal{H}$ -valued noise process. This smoothing is done with trace class operator Q . Define $\mathbf{W}_t^Q = Q^{1/2} \mathbf{W}_t =$

$\sum_{k \geq 1} \sqrt{\lambda^{(k)}} \mathbf{W}_t^{(k)} \phi^{(k)}$. Then, we have

$$\mathbb{E} \left[\left\| \mathbf{W}_t^Q \right\|_{\mathcal{H}}^2 \right] = \sum_{k \geq 1} \mathbb{E} \left[\lambda^{(k)} (\mathbf{W}_t^{(k)})^2 \right] = t \sum_{k \geq 1} \lambda^{(k)} < \infty, \quad (35)$$

where $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R} \times \mathcal{H} : k \in \mathbb{N}\}$ is eigen-system of Q . Hence $\mathbf{W}_t^Q$ is $\mathcal{H}$ -valued process. This smoothed process is called a Q -Wiener process or *coloured noise* because the noise covariance operator Q introduces a spatial correlation.

A.2. Stochastic Differential Equations in Hilbert Spaces

With the infinite dimensional Wiener process $\mathbf{W}_t^Q$ well defined $\mathcal{H}$, we then define the infinite-dimensional SDEs as follows:

$$d\mathbf{X}_t = \mathcal{A}\mathbf{X}_t dt + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{X}_0 = \mathbf{x}_0 \in \mathcal{H}, \quad (36)$$

where $\mathcal{A} : \mathcal{H} \rightarrow \mathcal{H}$ is a linear operator and $\sigma_t > 0$. We recall the classical mild formulation driven by a general Q -Wiener process on $\mathcal{H}$. We then restrict attention to a cylindrical Wiener process on $\mathcal{H}_0$, introduce an appropriate Hilbert–Schmidt embedding that can be naturally aligned with our framework, and outline the minimal assumptions that still guarantee existence of a solution for the infinite-dimensional SDE under the choice of cylindrical Wiener process.

Mild solution. Let a linear operator $\mathcal{A} : \mathcal{H} \rightarrow \mathcal{H}$ generate a strongly continuous C_0 semi-group $\{e^{t\mathcal{A}}\}_{t \geq 0}$ on $\mathcal{H}$. Assume $\mathbf{W}^Q$ is a centred Q -Wiener process with trace-class covariance operator Q and that $e^{t\mathcal{A}}Q^{1/2}$ is Hilbert-Schmidt for any $t > 0$. Under these conditions, the infinite-dimensional SDE (or stochastic PDE) in (3) admit the unique mild solution:

$$\mathbf{X}_t = e^{t\mathcal{A}}\mathbf{x}_0 + \int_0^t e^{(t-s)\mathcal{A}} \sigma_s d\mathbf{W}_s^Q, \quad t \geq 0. \quad (37)$$

The resulting stochastic process $\{\mathbf{X}_t\}_{t \geq 0}$ is then $\mathcal{H}$ -valued random variable since the stochastic convolution $\int_0^t e^{(t-s)\mathcal{A}} d\mathbf{W}_s^Q = \int_0^t e^{(t-s)\mathcal{A}} Q^{1/2} d\mathbf{W}_s$ is lies on $\mathcal{H}$ because $e^{(t-s)\mathcal{A}} Q^{1/2}$ is Hilbert-Schmidt. We can compute its mean $\mathbf{m}_t$ and covariance operator Q_t for any $t > 0$ as follows:

$$\mathbf{m}_t = e^{t\mathcal{A}}\mathbf{x}_0, \quad Q_t = \int_0^t \sigma_s^2 e^{(t-s)\mathcal{A}} Q e^{(t-s)\mathcal{A}^\dagger} ds. \quad (38)$$

Due to the Hille-Yosida Theorem (Da Prato and Zabczyk, 2014, Theorem A.3), $\mathcal{A}$ satisfies $\|e^{t\mathcal{A}}\| \leq M e^{-wt}$ for some $M \geq 1, w \geq 0$. Hence, for every $0 \leq s \leq t, T_{st} := e^{(t-s)\mathcal{A}} Q e^{(t-s)\mathcal{A}^\dagger}$ satisfies:

$$\text{Tr}(T_{st}) \leq \|e^{(t-s)\mathcal{A}}\| \text{Tr}(Q) \|e^{(t-s)\mathcal{A}^\dagger}\| \leq M^2 e^{-2w(t-s)} \text{Tr}(Q) < \infty. \quad (39)$$

Implies that T_{st} is a trace class operator. Therefore, since the bound is integrable on $[0, t]$,

$$\text{Tr}(Q_t) = \text{Tr}(\sigma_t^2 \int_0^t T_{st} ds) \leq \sigma_t^2 \int_0^t \text{Tr}(T_{st}) ds < \infty. \quad (40)$$

Hence Q_t is a trace-class operator for every $t < \infty$. Hence, we can conform that the $\mathcal{H}$ -valued stochastic process $\mathbf{X}_t$ has a Gaussian law $\nu_t = \mathcal{N}(\mathbf{m}_t, Q_t)$.

Invariant Measure. Let $\mathcal{A}$ is dissipative (i.e., $\langle \mathcal{A}\mathbf{x}, \mathbf{x} \rangle_{\mathcal{H}} \leq -\lambda \|\mathbf{x}\|_{\mathcal{H}}^2$ for some $\lambda > 0$), the integral converges as $t \rightarrow \infty$ and we define the noise schedule such that $\sup_t \sigma_t < \infty, \sigma_t \rightarrow \sigma_\infty$. Then, the operator Q_t solves the following differential Lyapunov equation:

$$\dot{Q}_t = \mathcal{A}Q_t + Q_t\mathcal{A}^\dagger + \sigma_t^2 Q, \quad Q_t = e^{t\mathcal{A}}Q_0e^{t\mathcal{A}^\dagger} + \int_0^t \sigma_{t-r}^2 e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} dr. \quad (41)$$

Because $Q_0 = 0$ since we assume Dirac delta initial condition $\mathbf{x}_0$ in (3), we obtain

$$Q_t \xrightarrow{t \rightarrow \infty} Q_\infty := \sigma_\infty^2 \int_0^\infty e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} dt \quad (42)$$

Now, let $F(t) = \sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger}$. Then, we get:

$$\frac{d}{dt} F(t) = \mathcal{A}F(t) + F(t)\mathcal{A}^\dagger = \mathcal{A}\sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} + \sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} \mathcal{A}^\dagger. \quad (43)$$

Since $F(t) \xrightarrow{t \rightarrow \infty} 0$ and $F(0) = \sigma_\infty^2 Q$, by integrating (43) from 0 to ∞ , we get:

$$\mathcal{A}Q_\infty + Q_\infty \mathcal{A}^\dagger = \int_0^\infty \frac{d}{dt} F(t) dt = \lim_{t \rightarrow \infty} F(t) - F(0) = -\sigma_\infty^2 Q \quad (44)$$

Hence, we get the operator Lyapunov equation (Da Prato and Zabczyk, 2002, Proposition 10.1.4):

$$\mathcal{A}Q_\infty + Q_\infty \mathcal{A}^\dagger + \sigma_\infty^2 Q = 0, \quad (45)$$

where Q_∞ is trace class operator. so $\mathcal{N}(0, \infty)$ is the unique invariant Gaussian measure of the process in (3). Finally, if $\mathcal{A}$ is trace-class, self-adjoint operator (i.e., $\mathcal{A} = \mathcal{A}^\dagger$) and $\mathcal{A}$ and Q are commute (i.e., $Q\mathcal{A} = \mathcal{A}Q$), we get $Q_\infty = -\frac{1}{2}\sigma_\infty^2 \mathcal{A}^{-1}Q$ (Da Prato and Zabczyk, 2002, Proposition 10.1.6).

A.3. Brief Introduction of the Transition Path Theory

In this section, we introduce the classic Traditional *transitional path theory* (TPT; Vanden-Eijnden et al., 2010), a method to sample and quantify reactive pathways meta stable states between **A** and **B** in complex system. Here, we briefly introduce two representative methods for TPT, *transition path sampling* (TPS) and *chain of states* (CoS) methods (e.g., NEB, string). Then, we will make explicit connection between these two viewpoints and show how our infinite-dimensional formulation provides a principled bridge between them by using the Feynman-Kac formula (Del Moral, 2011).

Following the notion of (Dellago et al., 1998), let us denote a discrete space-time path $\mathbf{P} := \{\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_L\}$, where $\mathbf{p}_u = (\mathbf{r}_u, \mathbf{v}_u) \in \mathbb{R}^{6N}$ is configuration with $\mathbf{r}_u \in \mathbb{R}^{3N}$ the atom-wise positions and $\mathbf{v}_u \in \mathbb{R}^{3N}$ the atom-wise velocities with N number of atom on time interval $\mathbf{L} = [0, L]$. We assume that consecutive configuration states follow a Markov transition $p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1})$, then the path probability is given by product of the transition probabilities as shown in below:

$$\mathbb{P}(\mathbf{P}) = e^{-\frac{1}{k_B T} V(\mathbf{p}_0)} \prod_{u=0}^{L-1} p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}), \quad (46)$$

where $V : \mathbb{R}^{6N} \rightarrow \mathbb{R}$ is the energy function. In this setting, our goal is to sample the transition path bridging boundary constraints $\pi_{\mathbf{A}}(\mathbf{p}_0)$ and $\pi_{\mathbf{B}}(\mathbf{p}_L)$ from the transition path distribution:

$$\mathbb{P}^*(\mathbf{P}) := e^{-S(\mathbf{P})} = \pi_{\mathbf{A}}(\mathbf{p}_0)\pi_{\mathbf{B}}(\mathbf{p}_L)\mathbb{P}(\mathbf{P}), \quad (47)$$

where $\pi_{\mathbf{A}}(\mathbf{p}_0)$ and $\pi_{\mathbf{B}}(\mathbf{p}_L)$ enforces the path to be started and terminated with prescribed **A** and **B**, respectively. Note that we can choose Markov transition $p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1})$ freely, if it conserves the Gibbs distribution normalized (Dellago et al., 1998). In following, we show how the choice of Markov transition differentiating the TPS from CoS-based *maximum energy path* (MEP) optimization.

A.4. Transition Path Sampling

Classic TPS is usually formulated for Langevin dynamics in the under-damped regime. However we focus on the over-damped case for its simple form, and the extension to the under-damped case is straightforward. The molecule system governed by over-damped Langevin dynamics is given by:

$$d\mathbf{p}_t = -\frac{1}{\gamma \mathbf{m}} \nabla_{\mathbf{p}} V(\mathbf{p}_t) dt + \sqrt{\frac{2k_B T}{\gamma \mathbf{m}}} d\mathbf{W}_t, \quad (48)$$

where we denote γ as friction and $\mathbf{m}$ as a mass. In this case, the discrete Markov transition kernel of overdamped dynamics in (48) follows from an Euler-Maruyama discretization, since the increment over δ_t is Gaussian with mean $\mu_t = \mathbf{p}_t - \frac{1}{\gamma \mathbf{m}} \nabla_{\mathbf{p}} V(\mathbf{p}_t) \delta_t$ and covariance $\Sigma_t = \frac{2k_B T}{\gamma \mathbf{m}} \delta_t \mathbf{I}_{3N}$:

$$\mathbf{p}_{u+1} = \mathbf{p}_u - \frac{1}{\gamma \mathbf{m}} \nabla_{\mathbf{p}} V(\mathbf{p}_t) \delta_t + \sqrt{\frac{2k_B T}{\gamma \mathbf{m}}} \delta_t \zeta, \quad \zeta \sim \mathcal{N}(0, \mathbf{I}_{3N}). \quad (49)$$

$$\text{Hence, } p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}) = \mathcal{N}(\mathbf{p}_{u+1} | \mu_t, \Sigma_t). \quad (50)$$

Hence, substituting (49) into (46) yields the TPS transition path distribution, which is widely used in recent literature (Holdijk et al., 2022; Seong et al., 2025; Blessing et al., 2026). Concretely, starting from the overdamped dynamics in (48), this line of work sets up a finite-dimensional SOC problem by biasing the potential V with a neural-network based control α^θ :

$$d\mathbf{p}_t^\alpha = -\frac{1}{\gamma\mathbf{m}}\nabla_{\mathbf{p}} [V(\mathbf{p}_t^\alpha) + \alpha^\theta(t, \mathbf{p}_t^\alpha)] dt + \sqrt{\frac{2k_B T}{\gamma\mathbf{m}}} d\mathbf{W}_t, \quad \mathbf{p}_0^\alpha = \mathbf{A}. \quad (51)$$

Now, the goal of the optimization is to learn α^θ that steers the prior dynamics start from a given metastable state $\mathbf{A}$ into another target metastable state $\mathbf{B}$ by maximizing the log-likelihood with respect to the TPD in (47).

A.5. Chain of States Methods

If the given system has smooth energy landscape, TPS provides an adequate description of the transition process. In this case, the main target is the *transition state*, which is a saddle point on the potential energy landscape (Vanden-Eijnden et al., 2010). However, for a complex energy landscape, when the initial and target states are separated by intermediates basins, the MEP that bridges the relevant minima through the sequence of transition states is preferred, which motivates CoS methods. The objective of CoS is to parametrize the geometric path as discrete chain $\mathbf{P}^\theta := \{\mathbf{p}_0 = \mathbf{A}, \mathbf{p}_1^\theta, \dots, \mathbf{p}_{L-1}^\theta, \mathbf{p}_L = \mathbf{B}\}$ over the grid with fixed boundary and to optimize a geometric functional. For example, the objective of elastic-band method (Jónsson et al., 1998) is given by:

$$S(\mathbf{P}^\theta) = \sum_{u=0}^{L-1} \left[V(\mathbf{p}_u^\theta) + \frac{\kappa}{2} \|\mathbf{p}_{u+1}^\theta - \mathbf{p}_u^\theta\|^2 \right], \quad (52)$$

where $\kappa > 0$ is constant. Recent studies (Petersen and Covino, 2025; Ramakrishnan et al., 2025) optimize the neural-network parameterization of continuous path $\mathbf{P}^\theta$ by minimizing objectives of the similar form in (52) to find the MEP. However, CoS formulations lack a probabilistic foundation, so they cannot target the reactive path ensemble or uncertainty without additional modeling.

A.6. Choice of the Terminal Cost in (12)

Previously, the transition path distribution of TPS is often derived by plugging the Langevin transition kernel (49) into the factorized form (46) as follows:

$$\mathbb{P}(\mathbf{P}) = e^{-\frac{1}{k_B T} V(\mathbf{p}_0)} \prod_{u=0}^{L-1} p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}), \quad (53a)$$

$$\text{where } p(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}) = \mathcal{N}(\mathbf{p}_{u+1} | \mathbf{p}_u - \frac{1}{\gamma\mathbf{m}} \nabla_{\mathbf{p}} V(\mathbf{p}_u) \delta_u, \frac{2k_B T}{\gamma\mathbf{m}} \mathbf{I}_{3N} \delta_u). \quad (53b)$$

Intuitively, (53a) gives the path log-likelihood of a discrete trajectory under the overdamped Langevin dynamics (48) at temperature T . Then, to enforce the endpoint condition $\pi_{\mathbf{B}}(\mathbf{p}_L)$ for TPS:

$$\mathbb{P}_{\text{TPS}}^*(\mathbf{P}) = e^{-S(\mathbf{P})} = \pi_{\mathbf{A}}(\mathbf{p}_0) \pi_{\mathbf{B}}(\mathbf{p}_L) \mathbb{P}(\mathbf{P}), \quad (54)$$

For the terminal cost U in (12), one can take the TPS energy in (54) as a negative log-likelihood (NLL) *i.e.*, $U(\mathbf{X}) = -\log \mathbb{P}_{\text{TPS}}^*(\mathbf{X})$, and then optimize the control α^θ to sample transition paths. In practice, however, we observed that this choice led to instability during FAS training. Indeed, the simulation of Langevin dynamics in (48) moves according to the local gradient $-\nabla V$, which helps stabilize trajectories, because this force steers paths toward descending the potential and suppress large uphill moves. In contrast, our infinite-dimensional SDE in (11) injects noise and evolves on a fixed training grid without explicitly aligning the path with geometric of $-\nabla V$, so small perturbations of the path can cause large oscillations in V along the grid, thereby U fluctuates across gradient steps.

Hence, we instead derive a principled path-objective that is well-suited for our path parameterization. Compared to TPS, by adopting well-known reweighting principle of Markov process known as the Feynman-Kac formula (Del Moral, 2011), we can rewrite the path distribution as *tilted* Markov chain:

$$\tilde{\mathbb{P}}(\mathbf{P}) = \prod_{u=0}^{L-1} e^{-\frac{1}{k_B T} V(\mathbf{p}_u)} \tilde{p}(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}), \quad (55a)$$

$$\text{where } \tilde{p}(\mathbf{p}_u \rightarrow \mathbf{p}_{u+1}) = \mathcal{N}(\mathbf{p}_{u+1} | \mathbf{p}_u, \frac{2k_B T}{\gamma\mathbf{m}} \mathbf{I}_{3N} \delta_u). \quad (55b)$$

Note that, the potential $V(\mathbf{p}_u)$ is carried by the path distribution rather than by the dynamics. Although the goal of (55) is same as (53), which is to concentrate the resulting path measure in high-potential region, this construction differs from (53), the high-potential can be achieved by direct reweighting, not from the local gradient $-\nabla V$. This viewpoint enable us to derive the CoS type objective (52) that is suited to our path-parametrization, leads to the transition path distribution:

$$\mathbb{P}_{\text{CoS}}^*(\mathbf{P}) \propto e^{-\tilde{S}(\mathbf{P})} = \pi_{\mathbf{A}}(\mathbf{P}_0)\pi_{\mathbf{B}}(\mathbf{P}_L)\tilde{\mathbb{P}}(\mathbf{P}), \quad (56)$$

With this transition path distribution, the NLL $U(\mathbf{X}) = -\log \mathbb{P}_{\text{CoS}}^*(\mathbf{X})$ can be written as follows:

$$U(\mathbf{X}) = -[\log \pi_{\mathbf{A}}(\mathbf{X}[0]) + \pi_{\mathbf{B}}(\mathbf{X}[L])] + \frac{1}{\kappa_B T} \sum_{u=0}^{L-1} \left[V(\mathbf{X}[u]) + \frac{\kappa}{2} \|\mathbf{X}[u+1] - \mathbf{X}[u]\|^2 \right] \quad (57)$$

$$\stackrel{(i)}{=} C + \underbrace{\frac{1}{\kappa_B T} \sum_{u=0}^{L-1} \left[V(\mathbf{X}[u]) + \frac{\kappa}{2} \|\mathbf{X}[u+1] - \mathbf{X}[u]\|^2 \right]}_{\text{CoS objective described in (52)}} \quad (58)$$

where $\kappa = \frac{1}{2} \frac{\gamma m}{\delta_u}$ and (i) follows from a fixed endpoints of our methods. Note that the second term on the right-hand side equals the CoS objective in (52). Hence (56) can be viewed as a probabilistic interpretation of classic elastic band methods. Compared to the NLL based on (54), (57) is better aligned with a path parameterization. It pushes every grid point $\mathbf{X}[u]$ for all $u \in \mathbf{L}$ toward low values of V and penalizes rough increments, which damps oscillations of V and yields smoother gradients, thereby training is more stable particular on sharp potentials such as molecular systems.

B. Proofs and Derivations

Assumption B.1. We assume a linear operator $\mathcal{A} : \mathcal{H} \rightarrow \mathcal{H}$ that generates a C_0 -semigroup and is exponentially stable, dissipative i.e., $\langle \mathcal{A}\mathbf{x}, \mathbf{x} \rangle_{\mathcal{H}} \leq -\lambda \|\mathbf{x}\|_{\mathcal{H}}^2$ for some $\lambda > 0$ and self-adjoint i.e., $\mathcal{A} = \mathcal{A}^\dagger$. We also take a Q -Wiener process with a self-adjoint trace-class operator Q that commutes with $\mathcal{A}$ i.e., $Q\mathcal{A} = \mathcal{A}Q$.

Assumption B.2. For each $(\mathbf{x}, \alpha) \in \mathcal{H} \times \mathcal{U}$, the maps $t \mapsto f(t, \mathbf{x}, \alpha)$, $t \mapsto g(t, \mathbf{x}, \alpha)$, and $t \mapsto l(t, \mathbf{x}, \alpha)$ are predictable. The function $h(\mathbf{x})$ is $\mathcal{F}_T$ -measurable. For each $(t, \alpha, \omega) \in [0, T] \times \mathcal{U} \times \Omega$, the functions f , g , l , and h are globally twice Fréchet differentiable with respect to $\mathbf{x}$. The derivatives $D_{\mathbf{x}}f$, $D_{\mathbf{x}}g$, $D_{\mathbf{x}\mathbf{x}}f$, $D_{\mathbf{x}\mathbf{x}}g$, $D_{\mathbf{x}\mathbf{x}}l$, and $D_{\mathbf{x}\mathbf{x}}h$ are continuous in $\mathbf{x}$ and uniformly bounded by a constant K . Moreover, we assume the growth bounds:

$$|f(t, \mathbf{x}, \alpha)| + \|g(t, \mathbf{x}, \alpha)\| + \|D_{\mathbf{x}}l(t, \mathbf{x}, \alpha)\| + \|D_{\mathbf{x}}h(\mathbf{x})\| \leq K(1 + \|\mathbf{x}\|_{\mathcal{H}} + |\alpha|_{\mathcal{U}}), \quad (59)$$

$$|l(t, \mathbf{x}, \alpha)| + |h(\mathbf{x})| \leq K(1 + \|\mathbf{x}\|_{\mathcal{H}} + |\alpha|_{\mathcal{U}}^2). \quad (60)$$

Assumption B.3. The function $\mathcal{V} : [0, T] \times \mathcal{H} \rightarrow \mathbb{R}$ and its derivatives $D_{\mathbf{x}}\mathcal{V}$, $D_{\mathbf{x}\mathbf{x}}\mathcal{V}$, $\partial_t \mathcal{V}$ are uniformly continuous on bounded subsets of $[0, T] \times \mathcal{H}$ and $(0, T) \times \mathcal{H}$, respectively. Moreover, for all $(t, \mathbf{x}) \in (0, T) \times \mathcal{H}$, there exists $C_1, C_2 > 0$ such that

$$|\mathcal{V}(t, \mathbf{x})| + |D_{\mathbf{x}}\mathcal{V}(t, \mathbf{x})| + |\partial_t \mathcal{V}(t, \mathbf{x})| + \|D_{\mathbf{x}\mathbf{x}}\mathcal{V}(t, \mathbf{x})\| + |\mathcal{A}^\dagger D_{\mathbf{x}}\mathcal{V}(t, \mathbf{x})| \leq C_1(1 + |\mathbf{x}|)^{C_2}. \quad (61)$$

B.1. Preliminary

We recall two key ingredients. The first is the product rule for Hilbert space semimartingales, also called stochastic integration by parts. We use this rule to derive the adjoint matching objective. The second is a verification theorem that we use to establish the optimality.

Lemma B.4 (Stochastic integration by parts in $\mathcal{H}$). Assume A, C are $\mathcal{H}$ -valued predictable processes with $\int_0^T \mathbb{E}\|A_s\|^2 ds < \infty$ and $\int_0^T \mathbb{E}\|C_s\|^2 ds < \infty$ for every $T > 0$ and B, D are predictable processes with values in the Hilbert-Schmidt space $\mathcal{L}_2(U_Q, \mathcal{H})$ and $\int_0^T \mathbb{E}\|B_s\|_{\mathcal{L}_2}^2 ds < \infty$, $\int_0^T \mathbb{E}\|D_s\|_{\mathcal{L}_2}^2 ds < \infty$ for every $T > 0$. Then, let us define two stochastic convolutions:

$$\mathbf{X}_t = \mathbf{X}_0 + \int_0^t A_s ds + \int_0^t B_s d\mathbf{W}_s^Q, \quad \mathbf{Y}_t = \mathbf{Y}_0 + \int_0^t C_s ds + \int_0^t D_s d\mathbf{W}_s^Q \quad (62)$$

Then, for all $t \geq 0$, almost surely, we get:

$$\langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}} = \langle \mathbf{X}_0, \mathbf{Y}_0 \rangle_{\mathcal{H}} + \int_0^t \langle \mathbf{Y}_s, A_s \rangle_{\mathcal{H}} ds + \int_0^t \langle \mathbf{X}_s, C_s \rangle_{\mathcal{H}} ds \quad (63)$$

$$+ \int_0^t \langle \mathbf{Y}_s, B_s d\mathbf{W}_s^Q \rangle_{\mathcal{H}} + \int_0^t \langle \mathbf{X}_s, D_s d\mathbf{W}_s^Q \rangle_{\mathcal{H}} + \int_0^t \text{Tr} [B_s Q D_s^\dagger] ds. \quad (64)$$

Equivalently, in differential form,

$$d\langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}} = \langle \mathbf{Y}_t, A_t \rangle_{\mathcal{H}} dt + \langle \mathbf{X}_t, C_t \rangle_{\mathcal{H}} dt + \langle \mathbf{Y}_t, B_t d\mathbf{W}_t^Q \rangle_{\mathcal{H}} + \langle \mathbf{X}_t, D_t d\mathbf{W}_t^Q \rangle_{\mathcal{H}} + \text{Tr} [B_t Q D_t^\dagger] dt.$$

Proof. Let us define $\mathbb{R}$ -valued process $\mathbf{Z}_t := \langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}}$ and C^2 mapping $f : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ such that $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{H}}$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{H}$. Since

$$D_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) = \mathbf{y}, D_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) = \mathbf{x}, D_{\mathbf{x}\mathbf{x}} f(\mathbf{x}, \mathbf{y}) = 0, D_{\mathbf{y}\mathbf{y}} f(\mathbf{x}, \mathbf{y}) = 0, \quad (65)$$

and $D_{\mathbf{x}\mathbf{y}} f(\mathbf{x}, \mathbf{y})$ is the continuous bilinear form $(h, k) \rightarrow \langle h, k \rangle_{\mathcal{H}}$. Now, write the semi-martingale decomposition on $\mathcal{H} \times \mathcal{H}$:

$$(\mathbf{X}_t, \mathbf{Y}_t) = (\mathbf{X}_0, \mathbf{Y}_0) + \int_0^t (A_s, B_s) ds + \int_0^t \tilde{B}_s d\mathbf{W}_s^Q, \quad (66)$$

where $\tilde{B}_s : \mathcal{U} \rightarrow \mathcal{H} \times \mathcal{H}$ such that $\tilde{B}_s u = (B_s u, D_s u)$. By applying the infinite-dimensional Itô formula (Da Prato and Zabczyk, 2014, Chapter 4.4) for f on $\mathcal{H} \times \mathcal{H}$ gives:

$$d\mathbf{Z}_t = \langle D_{\mathbf{x}} f(\mathbf{X}_t, \mathbf{Y}_t), d\mathbf{X}_t \rangle_{\mathcal{H}} + \langle D_{\mathbf{y}} f(\mathbf{X}_t, \mathbf{Y}_t), d\mathbf{Y}_t \rangle_{\mathcal{H}} + \frac{1}{2} \text{Tr} \left[D_{\mathbf{x}\mathbf{y}} f(\mathbf{X}_t, \mathbf{Y}_t) (\tilde{B}_t Q \tilde{B}_t^\dagger) \right] dt, \quad (67)$$

where we compute the trace by expanding the function in any orthonormal basis $\{\phi^{(k)}\}_{k \geq 1}$ of $\mathcal{U}$:

$$\frac{1}{2} \text{Tr} \left[D_{\mathbf{x}\mathbf{y}} f(\mathbf{X}_t, \mathbf{Y}_t) (\tilde{B}_t Q \tilde{B}_t^\dagger) \right] = \sum_{k \geq 1} \lambda^{(k)} \left(\langle B_t, \phi^{(k)} \rangle_{\mathcal{H}} + \langle D_t, \phi^{(k)} \rangle_{\mathcal{H}} \right) = \text{Tr} [B_t Q D_t^\dagger], \quad (68)$$

where $\lambda^{(k)}$ is eigen-value corresponding eigen-basis $\phi^{(k)}$ such that $Q\phi^{(k)} = \lambda^{(k)}\phi^{(k)}$. Therefore,

$$d\langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}} = \langle \mathbf{Y}_t, d\mathbf{X}_t \rangle_{\mathcal{H}} + \langle \mathbf{X}_t, d\mathbf{Y}_t \rangle_{\mathcal{H}} + \text{Tr} [B_t Q D_t^\dagger] dt. \quad (69)$$

By plugging $d\mathbf{X}_t = A_t dt + B_t d\mathbf{W}_t^Q$ and $d\mathbf{Y}_t = C_t dt + D_t d\mathbf{W}_t^Q$, we get the desired results. $\square$

Lemma B.5 (Verification Theorem). *Let $\mathcal{V}$ be a solution of Hamilton-Jacobi-Bellman (HJB) equation:*

$$\partial_t \mathcal{V}_t + \mathcal{A} \mathcal{V}_t + \inf_{\alpha \in \mathbb{A}} \left[\langle \alpha, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle + \frac{1}{2} \|\alpha\|_{\mathcal{H}}^2 \right] = 0, \quad \mathcal{V}(T, \mathbf{x}) = g(\mathbf{x}) \quad (70)$$

which satisfying the Theorem B.3. Then, we have $\mathcal{V}(0, \mathbf{x}_0) \leq \mathcal{J}(\alpha, \mathbb{P}^\alpha)$ for every $\alpha \in \mathbb{A}$. Now, let $(\alpha^*, \mathbf{X}^*)$ be an admissible pair such that

$$\alpha_t^*(t, \mathbf{X}_t^*) = \arg \inf_{\alpha \in \mathbb{A}} \left[\langle \alpha_t, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 \right] = -\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^*) \quad (71)$$

for almost every $t \in [0, T]$ and $\mathbb{P}$ -almost surely. Then $(\alpha^*, \mathbf{X}^*)$ satisfying $\mathcal{V}(0, \mathbf{x}_0) = \mathcal{J}(\alpha^*, \mathbb{P}^*)$.

Proof. To start the proof, we formally minimize $F(D_{\mathbf{x}} \mathcal{V})$. Using results from (Da Prato and Zabczyk, 1997), we obtain the candidate minimizer and proceed accordingly:

$$F(\mathbf{x}) = \inf_{\alpha \in \mathbb{A}} \left[\langle \mathbf{x}, \alpha \rangle + \frac{1}{2} \|\alpha\|_{\mathcal{H}}^2 \right] = -\frac{1}{2} \|\mathbf{x}\|_{\mathcal{H}}^2. \quad (72)$$

Thus, $F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V})$ is minimized at $\alpha^* = -\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}$. Next, applying Itô's formula (Da Prato and Zabczyk, 2014, Chapter 4.4) to $\mathcal{V}$ and take the expectation on both sides to obtain

$$\mathbb{E}_{\mathbb{P}^\alpha} [\mathcal{V}(T, \mathbf{X}_T^\alpha)] = \mathbb{E}_{\mathbb{P}^\alpha} [g(\mathbf{X}_T^\alpha)] = \quad (73)$$

$$= \mathcal{V}(0, \mathbf{x}_0) + \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \left(\partial_t \mathcal{V}(t, \mathbf{X}_t^\alpha) + \mathcal{A}\mathcal{V}(t, \mathbf{X}_t^\alpha) + \langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} \right) dt \right], \quad (74)$$

Using that $\mathcal{V}$ satisfies (71), add $\mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 dt \right]$ to both sides. Then the LHS of (73) becomes:

$$\mathbb{E}_{\mathbb{P}^\alpha} \left[g(\mathbf{X}_T^\alpha) + \int_0^T \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 dt \right] = \mathcal{J}(\alpha, \mathbb{P}^\alpha). \quad (75)$$

Now for the RHS of the equation (73):

$$\mathcal{V}(0, \mathbf{x}_0) + \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \left(\frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 + \partial_t \mathcal{V}(t, \mathbf{X}_t^\alpha) + \mathcal{A}\mathcal{V}(t, \mathbf{X}_t^\alpha) + \langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} \right) dt \right] \quad (76)$$

$$\stackrel{(i)}{=} \mathcal{V}(0, \mathbf{x}_0) + \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \left(\langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 - F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha)) \right) dt \right]. \quad (77)$$

Here (i) holds because we add and subtract $\mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T F(D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha)) dt \right]$ and use again that $\mathcal{V}$ satisfies (71). Therefore, we obtain the following equation:

$$\mathcal{J}(\alpha, \mathbb{P}^\alpha) = \mathcal{V}(0, \mathbf{x}_0) \quad (78)$$

$$+ \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \left(\langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 - F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha)) \right) dt \right]. \quad (79)$$

By definition, $\left[\langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 \right] - F(D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha)) \geq 0$. Thus, taking the infimum over $\alpha \in \mathbb{A}$ on the RHS of (78) yields $\mathcal{J}(\alpha, \mathbb{P}^\alpha) \geq \mathcal{V}(0, \mathbf{x}_0)$. We already verified that $F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V})$ attains the infimum at $\alpha_t^* = -\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}$, so choose $u = \alpha^*$. Then,

$$\left[\langle \sigma Q^{1/2} D_{\mathbf{x}} \mathcal{V}(s, \mathbf{X}_s^u), u_s \rangle_{\mathcal{H}} + \frac{1}{2} \|u_s\|_{\mathcal{H}}^2 \right] - F(\sigma Q^{1/2} D_{\mathbf{x}} \mathcal{V}(s, \mathbf{X}_s^u)) = 0. \quad (80)$$

Thus, we get $\mathcal{J}(u, \mathbb{P}^u) = \mathcal{V}(0, \mathbf{x}_0)$. Combining with (78), we conclude that $(\alpha^*, \mathbf{X}^{\alpha^*})$ is optimal for $(t, \mathbf{x}) \in [0, T] \times \mathcal{H}$. This completes the proof. $\square$

B.2. Derivation of Divergence Between Path Measures

Here we present an infinite-dimensional generalization of Girsanov theorem (Da Prato and Zabczyk, 2014, Theorem 10.14), which plays a crucial role in estimating the divergence between two paths.

Theorem B.6 (Generalized Girsanov). *Let $\{\zeta_t\}_{0 \leq t \leq T}$ be a $\mathcal{U}$ -valued predictable process such that*

$$\mathbb{P} \left(\int_0^T \|\zeta_t\|_{\mathcal{U}}^2 dt < \infty \right) = 1, \quad \mathbb{E} \left[\exp \left(\int_0^T \langle \zeta_t, d\mathbf{W}_t^Q \rangle_{\mathcal{U}} - \frac{1}{2} \int_0^T \|\zeta_t\|_{\mathcal{U}}^2 dt \right) \right] = 1. \quad (81)$$

Then the process $\tilde{\mathbf{W}}_t^Q = \mathbf{W}_t^Q - \int_0^t \zeta_s ds$ is a Q -Wiener process with respect to the filtration on the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$ where

$$d\mathbb{Q} = \exp \left(\int_0^T \langle \zeta_t, d\mathbf{W}_t^Q \rangle_{\mathcal{U}} - \frac{1}{2} \int_0^T \|\zeta_t\|_{\mathcal{U}}^2 dt \right) d\mathbb{P}. \quad (82)$$

Or alternatively, by substituting $\mathbf{W}_t^Q = \tilde{\mathbf{W}}_t^Q + \int_0^t \zeta_s ds$ to (87), we obtain

$$d\mathbb{Q} = \exp \left(\int_0^T \langle \zeta_t, d\tilde{\mathbf{W}}_t^Q \rangle_{\mathcal{U}} + \frac{1}{2} \int_0^T \|\zeta_t\|_{\mathcal{U}}^2 dt \right) d\mathbb{P}. \quad (83)$$

Proof. Proof can be found in (Da Prato and Zabczyk, 2014, Theorem 10.14) $\square$

Based on Theorem B.6, we derive the KL-objective in (12). Let $\eta_t = Q^{1/2}\alpha_t \in \mathcal{H}_0 \subseteq \mathcal{H}$ and $\tilde{\mathbf{W}}_t \in \mathcal{U}$ be a $\mathbb{P}^\alpha$ -cylindrical Wiener process and set $\mathbf{W}_t := \tilde{\mathbf{W}}_t - \int_0^t \alpha_s ds$. Then, substitute $\mathbf{W}_t^Q = \tilde{\mathbf{W}}_t^Q - \int_0^t \eta_t ds$ to a controlled SDEs in $\mathbb{P}^\alpha$:

$$d\mathbf{X}_t^\alpha = \mathcal{A}\mathbf{X}_t^\alpha dt + \sigma_t \eta_t dt + \sigma_t d\mathbf{W}_t^Q \quad (84)$$

$$= \mathcal{A}\mathbf{X}_t dt + \sigma_t \eta_t dt + \sigma_t [d\tilde{\mathbf{W}}_t^Q - \eta_t dt] \quad (85)$$

$$= \mathcal{A}\mathbf{X}_t dt + \sigma_t d\tilde{\mathbf{W}}_t^Q \quad (86)$$

It implies that $\tilde{\mathbf{W}}_t^Q$ is Q-Wiener process with respect to $\mathbb{P}^\alpha$. Moreover, by Girsanov theorem,

$$d\mathbb{P} = \exp \left(\int_0^T \langle \alpha_t, d\tilde{\mathbf{W}}_t \rangle_{\mathcal{U}} - \frac{1}{2} \int_0^T \|\alpha_t\|_{\mathcal{U}}^2 dt \right) d\mathbb{P}^\alpha \quad (87)$$

Now, by change of measure in (10), we get the following relation:

$$d\mathbb{P}^* = \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{X}_T)} d\mathbb{P} \quad (88)$$

$$= \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T^\alpha) - \log q_T(\mathbf{X}_T^\alpha)} \frac{d\mathbb{P}}{d\mathbb{P}^\alpha} d\mathbb{P}^\alpha, \quad (89)$$

where $\tilde{\mathcal{Z}} = \mathbb{E} \left[\frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{X}_T)} \right]$ is normalization constant. Combining (87) and (89) and take logarithm on both sides, we get

$$\log \frac{d\mathbb{P}^*}{d\mathbb{P}^\alpha} = \log \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T^\alpha) - \log q_T(\mathbf{X}_T^\alpha)} + \log \frac{d\mathbb{P}}{d\mathbb{P}^\alpha} \quad (90)$$

$$= -U(\mathbf{X}_T^\alpha) - \log q_T(\mathbf{X}_T^\alpha) + \int_0^T \langle \alpha_t, d\tilde{\mathbf{W}}_t \rangle_{\mathcal{U}} - \frac{1}{2} \int_0^T \|\alpha_t\|_{\mathcal{U}}^2 dt - \log \tilde{\mathcal{Z}}. \quad (91)$$

Since $\tilde{\mathbf{W}}_t$ is $\mathbb{P}^\alpha$ -Wiener process, taking expectation with respect to $\mathbb{P}^\alpha$ on (91) we get

$$\mathbb{E}_{\mathbb{P}^\alpha} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^\alpha} \right] = \mathbb{E}_{\mathbb{P}^\alpha} \left[\frac{1}{2} \int_0^T \|\alpha_t\|_{\mathcal{U}}^2 dt + (U(\mathbf{X}_T^\alpha) + \log q_T(\mathbf{X}_T^\alpha)) \right] + \log \tilde{\mathcal{Z}}. \quad (92)$$

Then, since $\tilde{\mathcal{Z}}$ is constant, we get the desired KL-divergence minimization objective:

$$\min_{\alpha} D_{\text{KL}}(\mathbb{P}^\alpha | \mathbb{P}^*) = \min_{\alpha} \mathbb{E}_{\mathbb{P}^\alpha} \left[\frac{1}{2} \int_0^T \|\alpha_t\|_{\mathcal{U}}^2 dt + (U(\mathbf{X}_T^\alpha) + \log q_T(\mathbf{X}_T^\alpha)) \right]. \quad (93)$$

B.3. Proof of Lemma 1

Lemma 1 (Finite Normalization Constant). *Let us assume $-U(\mathbf{x}) \leq \beta \|\mathbf{x}\|_{\mathcal{H}}^2 - C$ with constants $\beta, C > 0$. Then, for any $\text{Tr}(Q) < \infty$, choosing $\nu = \mathcal{N}(0, Q)$ in (6) results in a finite normalization constant $\mathcal{Z} < \infty$.*

Proof. The Fernique's integrability theorem (Da Prato and Zabczyk, 2014, Theorem 2.7) states that for any centred Gaussian measure $\nu = \mathcal{N}(0, Q)$ with any $\text{Tr}(Q) < \infty$ on a separable Banach space $\mathcal{B}$ (in particular on $\mathcal{H}$) there exists $\beta > 0$ such that $\int_{\mathcal{H}} e^{\beta \|\mathbf{x}\|_{\mathcal{H}}^2} d\nu(\mathbf{x}) < \infty$. Because the energy satisfies quadratic lower bound $U(\mathbf{x}) \geq -\beta \|\mathbf{x}\|_{\mathcal{H}}^2 + C$, we obtain $e^{-U(\mathbf{x})} \leq \tilde{C} e^{\beta \|\mathbf{x}\|_{\mathcal{H}}^2}$, where $\tilde{C} = e^{-C}$. Hence, we get:

$$\mathcal{Z} = \int_{\mathcal{H}} e^{-U(\mathbf{x})} d\nu(\mathbf{x}) \leq \tilde{C} \int_{\mathcal{H}} e^{\beta \|\mathbf{x}\|_{\mathcal{H}}^2} d\nu(\mathbf{x}) < \infty, \quad (94)$$

which confirms that the partition function $\mathcal{Z}$ is finite. $\square$

B.4. Proof of Theorem 2

Theorem 2 (Explicit RND). *Suppose Theorem B.1 holds. Then, for any $0 < t < \infty$ and initial condition $\mathbf{x}_0 \in \mathcal{H}$, $\nu_t^{\mathbf{x}_0} = \mathcal{N}(e^{tA}\mathbf{x}_0, Q_t)$ and $\nu_\infty = \mathcal{N}(0, Q_\infty)$ are mutually absolutely continuous. Moreover, for any $\mathbf{x} \in \mathcal{H}$, we obtain the explicit RND $q_t(\mathbf{x}_0, \mathbf{x}) := \frac{d\nu_t^{\mathbf{x}_0}}{d\nu_\infty}(\mathbf{x})$ defined as follows:*

$$q_t(\mathbf{x}_0, \mathbf{x}) = \det(\Theta_t)^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0} \rangle_\infty + \langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x} \rangle_\infty - \frac{1}{2} \langle e^{2tA} \Theta_t^{-1} \mathbf{x}, \mathbf{x} \rangle_\infty \right], \quad (9)$$

where we define $\mathbf{m}_t^{\mathbf{x}_0} = e^{tA}\mathbf{x}_0$, $\Theta_t = 1 - e^{2tA}$ and Q_∞ norm $\langle u, v \rangle_\infty := \langle Q_\infty^{-\frac{1}{2}}u, Q_\infty^{-\frac{1}{2}}v \rangle_{\mathcal{H}}$.

Proof. From (41), we get the covariance identity:

$$Q_t = \int_0^t \sigma_{t-r}^2 e^{rA} Q e^{rA^\dagger} dr \quad (95)$$

$$= \int_0^t \sigma_\infty^2 e^{rA} Q e^{rA^\dagger} dr + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{rA} Q e^{rA^\dagger} dr \quad (96)$$

$$= Q_\infty - \int_t^\infty e^{rA} Q e^{rA^\dagger} dr + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{rA} Q e^{rA^\dagger} dr \quad (97)$$

$$= Q_\infty - e^{tA} Q_\infty e^{tA^\dagger} + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{rA} Q e^{rA^\dagger} dr. \quad (98)$$

Since, $\int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{rA} Q e^{rA^\dagger} dr \rightarrow o(1)$ as we choose $\sigma_t \rightarrow \sigma_\infty$, we get the relation:

$$Q_t \approx Q_\infty - e^{tA} Q_\infty e^{tA^\dagger} = Q_\infty^{1/2} \Theta_t Q_\infty^{1/2}. \quad (99)$$

Hence, by applying (Da Prato and Zabczyk, 2002, Proposition 1.3.11), for two centered Gaussians $\mathcal{N}(0, Q_t)$ and $\mathcal{N}(0, Q_\infty)$,

$$q_t(0, \mathbf{x}) = \frac{d\mathcal{N}(0, Q_t)}{d\mathcal{N}(0, Q_\infty)}(\mathbf{x}) = \det(\Theta)^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \langle e^{2tA} \Theta^{-1} \mathbf{x}, \mathbf{x} \rangle_\infty \right]. \quad (100)$$

Now, for the shift from centered case to the general mean function $e^{tA}\mathbf{x}_0$, we utilize the Cameron-Martin theorem (Da Prato and Zabczyk, 2002, Theorem. 1.3.6). It implies that:

$$\frac{d\mathcal{N}(e^{tA}\mathbf{x}_0, Q_t)}{d\mathcal{N}(0, Q_t)}(\mathbf{x}) = \exp \left[\langle Q_t^{-\frac{1}{2}} e^{tA}\mathbf{x}_0, Q_t^{-\frac{1}{2}} \mathbf{x} \rangle_{\mathcal{H}} - \frac{1}{2} \langle Q_t^{-\frac{1}{2}} e^{tA}\mathbf{x}_0, Q_t^{-\frac{1}{2}} e^{tA}\mathbf{x}_0 \rangle_{\mathcal{H}} \right] \quad (101)$$

$$= \exp \left[\langle Q_\infty^{\frac{1}{2}} Q_t^{-1} e^{tA}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} \mathbf{x} \rangle_{\mathcal{H}} - \frac{1}{2} \langle Q_\infty^{\frac{1}{2}} Q_t^{-1} e^{tA}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} e^{tA}\mathbf{x}_0 \rangle_{\mathcal{H}} \right] \quad (102)$$

$$\stackrel{(i)}{=} \exp \left[\langle \Theta^{-1} Q_\infty^{-\frac{1}{2}} e^{tA}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} \mathbf{x} \rangle_{\mathcal{H}} - \frac{1}{2} \langle \Theta^{-1} Q_\infty^{-\frac{1}{2}} e^{tA}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} e^{tA}\mathbf{x}_0 \rangle_{\mathcal{H}} \right] \quad (103)$$

$$= \exp \left[\langle \Theta^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x} \rangle_\infty - \frac{1}{2} \langle \Theta^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0} \rangle_\infty \right], \quad (104)$$

where (i) follows from (99) such that $\Theta^{-1} Q_\infty^{-1/2} = (Q_t^{-1/2} Q_\infty^{1/2})^\dagger Q_t^{-1/2} = Q_\infty^{1/2} Q_t^{-1}$. Hence, by chain-rule, we get the desired result:

$$q_t(\mathbf{x}_0, \mathbf{x}) := \frac{d\mathcal{N}(e^{tA}\mathbf{x}_0, Q_t)}{d\mathcal{N}(0, Q_t)} \frac{d\mathcal{N}(0, Q_t)}{d\mathcal{N}(0, Q_\infty)}(\mathbf{x}) \quad (105)$$

$$= \det(\Theta)^{-\frac{1}{2}} \exp \left[-\frac{1}{2} \langle \Theta^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0} \rangle_\infty + \langle \Theta^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x} \rangle_\infty - \frac{1}{2} \langle e^{2tA} \Theta^{-1} \mathbf{x}, \mathbf{x} \rangle_\infty \right] \quad (106)$$

It concludes the proof. $\square$

B.5. Proof of Proposition 4

Proposition 4 (Adjoint Matching in $\mathcal{H}$). *Consider the following infinite-dimensional matching objective:*

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\frac{1}{2} \|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}\|_{\mathcal{H}}^2 \right] dt \quad (18)$$

$$d\mathbf{Y}_t^{\bar{\alpha}} = -\mathcal{A}^\dagger \mathbf{Y}_t^{\bar{\alpha}} dt + \mathbf{Z}_t^{\bar{\alpha}} d\mathbf{W}_t^Q, \quad \mathbf{Y}_T^{\bar{\alpha}} = D_{\mathbf{x}} g(\mathbf{X}_T^{\bar{\alpha}}). \quad (19)$$

where $\bar{\alpha} := \text{stopgrad}(\alpha^\theta)$. Then, the critical point of $\mathcal{L}$ is optimal control α^* for the SOC problem in (12).

Proof. Let $\mathbf{X}^\theta := \mathbf{X}^{\alpha^\theta}$ be a controlled process with parametrized control policy α_θ and $\xi_t := \partial_\theta \mathbf{X}_t^\theta$ be the variation process of $\mathbf{X}^\theta$ with respect to θ . Since diffusion coefficient σ_t is independent to θ , we get the differential equation for the variation process by differentiate (11) with respect to θ as follows:

$$d\xi_t = \mathcal{A}\xi_t dt + \sigma_t Q^{1/2} (D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t)) dt, \quad \xi_t = 0. \quad (107)$$

Now, let us define the cost functional with $\mathbb{P}^\theta := \mathbb{P}^{\alpha^\theta}$:

$$\mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) = \mathbb{E}_{\mathbb{P}^\theta} \left[\int_0^T \frac{1}{2} \|\alpha^\theta(t, \mathbf{X}_t^\theta)\|_{\mathcal{H}}^2 dt + g(\mathbf{X}_T^\theta) \right]. \quad (108)$$

Then, differentiating cost functional (108) with respect to θ is given by,

$$\frac{d}{d\theta} \mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) = \mathbb{E}_{\mathbb{P}^\theta} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^\theta, t), D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t) \rangle_{\mathcal{H}} dt + \langle D_{\mathbf{x}} g(\mathbf{X}_T^\theta), \xi_T \rangle_{\mathcal{H}} \right]. \quad (109)$$

Next, by applying the Theorem B.4 to $\langle \mathbf{Y}_t^\theta, \xi_t \rangle_{\mathcal{H}}$, we obtain:

$$d\langle \mathbf{Y}_t^\theta, \xi_t \rangle_{\mathcal{H}} = \langle d\mathbf{Y}_t^\theta, \xi_t \rangle_{\mathcal{H}} + \langle \mathbf{Y}_t^\theta, d\xi_t \rangle_{\mathcal{H}} \quad (110)$$

$$\stackrel{(i)}{=} \langle -\mathcal{A}^\dagger \mathbf{Y}_t^\theta, \xi_t \rangle_{\mathcal{H}} dt + \langle \mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \xi_t \rangle_{\mathcal{H}} \quad (111)$$

$$+ \langle \mathbf{Y}_t^\theta, \mathcal{A}\xi_t + \sigma_t Q^{1/2} (D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t)) \rangle_{\mathcal{H}} \quad (112)$$

$$\stackrel{(ii)}{=} \langle \mathbf{Y}_t^\theta, \sigma_t Q^{1/2} (D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t)) \rangle_{\mathcal{H}} dt + \langle \mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \xi_t \rangle_{\mathcal{H}} \quad (113)$$

where (i) follows since the quadratic covariation $[\xi, \mathbf{Y}^\theta]_t$ canceled out because ξ has no martingale term, and (ii) follows from the adjointness of $\mathcal{A}$ and $Q^{1/2}$ such that i.e., $\langle \mathcal{A}^\dagger u, v \rangle_{\mathcal{H}} = \langle u, \mathcal{A}v \rangle_{\mathcal{H}}$. Since $\xi_0 = 0$ and $\mathbf{Y}_T^\theta = D_{\mathbf{x}} g(\mathbf{X}_T^\theta)$, by integrating (110) and takes expectation, we obtain:

$$\mathbb{E}_{\mathbb{P}^\theta} [\langle D_{\mathbf{x}} g(\mathbf{X}_T^\theta), \xi_T \rangle] = \mathbb{E}_{\mathbb{P}^\theta} \left[\int_0^T d\langle \mathbf{Y}_t^\theta, \xi_t \rangle_{\mathcal{H}} \right] \quad (114)$$

$$= \mathbb{E}_{\mathbb{P}^\theta} \left[\int_0^T \langle \mathbf{Y}_t^\theta, \sigma_t Q^{1/2} (D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t)) \rangle_{\mathcal{H}} dt \right]. \quad (115)$$

Finally, substituting (114) into (109), we get:

$$\frac{d}{d\theta} \mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) \stackrel{(i)}{=} \mathbb{E}_{\mathbb{P}^\theta} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^\theta, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^\theta, D_{\mathbf{x}} \alpha^\theta(\mathbf{X}_t^\theta, t) \xi_t + \partial_\theta \alpha^\theta(\mathbf{X}_t^\theta, t) \rangle_{\mathcal{H}} dt \right], \quad (116)$$

where (i) follows from the adjointness of $Q^{1/2}$ and scalar σ_t . Since $\bar{\alpha} = \text{stopgrad}(\alpha^\theta)$ only blocks the gradient, $\bar{\alpha}(t, \mathbf{x}) =$

$\alpha^\theta(t, \mathbf{x})$ in pointwise manner. Hence, the forward law satisfies $\mathbb{P}^\theta = \mathbb{P}^{\bar{\alpha}}$. Therefore, from (116), we can define

$$\frac{d}{d\theta} \mathcal{J}(\alpha^\theta, \mathbb{P}^{\bar{\alpha}}) \stackrel{(i)}{=} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}} dt \right] \quad (117)$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}} + \langle \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}} dt \right] \quad (118)$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\int_0^T \frac{1}{2} \partial_\theta \left\| \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}} \right\|_{\mathcal{H}}^2 dt \right] \quad (119)$$

$$\stackrel{(ii)}{=} \partial_\theta \int_0^T \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\frac{1}{2} \left\| \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}} \right\|_{\mathcal{H}}^2 \right] dt \quad (120)$$

$$= \frac{d}{d\theta} \mathcal{L}(\theta), \quad (121)$$

where (i) follows by substituting $\mathbb{P}^{\bar{\alpha}}$ into (116) while $D_{\mathbf{x}} \alpha^\theta \xi_t$ disappears due to stopgrad operation blocks the gradient through $\mathbf{X}_t^{\bar{\alpha}}$ in (107) (i.e., $\xi_t = 0$) and (ii) follows from dominate convergence. Therefore, from the SMP stationarity:

$$\alpha^\theta(\mathbf{X}_t^\theta, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^\theta = \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}} = 0, \quad t \in [0, T], \quad (122)$$

we obtain $\frac{d}{d\theta} \mathcal{J}(\alpha^\theta, \mathbb{P}^{\bar{\alpha}}) = \frac{d}{d\theta} \mathcal{L}(\theta) = 0$. This implies that the SOC objective in (12) and matching objective in (18) share the same critical points.

Now, let us verify that the critical point $\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) = -\sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}$ in (122) is optimal control α^* for the SOC problem in (12). To do so, let us assume that there exists $\mathcal{V}$ be a solution of HJB equation (70) satisfying Theorem B.3. Then, by applying Itô formula, we have:

$$d\mathcal{V}(t, \mathbf{X}_t^\theta) = [\partial_t \mathcal{V}_t + \langle \mathcal{A} \mathbf{X}_t^\theta, D_{\mathbf{x}} \mathcal{V}_t \rangle_{\mathcal{H}}] \quad (123)$$

$$+ \frac{1}{2} \text{Tr}(\sigma_t^2 Q D_{\mathbf{xx}} \mathcal{V}_t) + \langle \sigma_t Q^{1/2} \alpha_t^\theta, D_{\mathbf{x}} \mathcal{V}_t \rangle_{\mathcal{H}} dt + \langle D_{\mathbf{x}} \mathcal{V}, \sigma_t d\mathbf{W}_t^Q \rangle_{\mathcal{H}} \quad (124)$$

$$\stackrel{(i)}{=} [\langle \alpha_t^\theta, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle_{\mathcal{H}} - F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t)] dt + \langle D_{\mathbf{x}} \mathcal{V}, \sigma_t d\mathbf{W}_t^Q \rangle_{\mathcal{H}}, \quad (125)$$

where (i) follows from the definition of HJB equation (70) and we define $F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t) := \inf_{\alpha \in \mathbb{A}} [\langle \alpha_t, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle_{\mathcal{H}} + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2]$. Now, assume $\mathbf{Y}_t^\theta := D_{\mathbf{x}} \mathcal{V}_t$ and $\mathbf{Z}_t^\theta := D_{\mathbf{xx}} \mathcal{V}_t$ and apply the Itô formula, then we get

$$d\mathbf{Y}_t^\theta = [\partial_t D_{\mathbf{x}} \mathcal{V}_t + \langle \mathcal{A} \mathbf{X}_t^\theta, D_{\mathbf{xx}} \mathcal{V}_t \rangle_{\mathcal{H}}] \quad (126)$$

$$+ \frac{1}{2} \text{Tr}(\sigma_t Q D_{\mathbf{xx}} (D_{\mathbf{x}} \mathcal{V}_t)) + \langle \sigma_t Q^{1/2} \alpha_t^\theta, D_{\mathbf{xx}} \mathcal{V}_t \rangle_{\mathcal{H}}] + \sigma_t \mathbf{Z}_t^\theta d\mathbf{W}_t^Q. \quad (127)$$

Here, by differentiate the HJB in (70) with respect to $\mathbf{x}$, we get:

$$\partial_t D_{\mathbf{x}} \mathcal{V}_t + \mathcal{A}^\dagger D_{\mathbf{x}} \mathcal{V} + \langle \mathcal{A} \mathbf{X}_t^\theta, D_{\mathbf{xx}} \mathcal{V}_t \rangle_{\mathcal{H}} \quad (128)$$

$$+ \frac{1}{2} \text{Tr}(\sigma_t^2 Q D_{\mathbf{xx}} (D_{\mathbf{x}} \mathcal{V}_t)) + D_{\mathbf{xx}} \mathcal{V}_t \sigma_t Q^{1/2} D_{\mathbf{x}} F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}) = 0. \quad (129)$$

Now, by substituting (128) into (126), we get the following:

$$d\mathbf{Y}_t^\theta = [-\mathcal{A}^\dagger \mathbf{Y}_t + \langle \sigma_t Q^{1/2} \alpha_t^\theta - \sigma_t Q^{1/2} D_{\mathbf{x}} F(\sigma_t Q^{1/2} \mathbf{Y}_t^\theta), D_{\mathbf{xx}} \mathcal{V} \rangle] dt + \sigma_t \mathbf{Z}_t^\theta d\mathbf{W}_t^Q \quad (130)$$

$$\stackrel{(i)}{=} [-\mathcal{A}^\dagger \mathbf{Y}_t^\theta + \langle \sigma_t Q^{1/2} \alpha_t^\theta + \sigma_t^2 Q \mathbf{Y}_t^\theta, D_{\mathbf{xx}} \mathcal{V} \rangle] dt + \sigma_t \mathbf{Z}_t^\theta d\mathbf{W}_t^Q \quad (131)$$

$$\stackrel{(ii)}{=} -\mathcal{A}^\dagger \mathbf{Y}_t^\theta dt + \sigma_t \mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \quad (132)$$

where (i) follows from $F(\sigma_t Q^{1/2} \mathbf{Y}_t^\theta) = -\frac{1}{2} \|\sigma_t Q^{1/2} \mathbf{Y}_t^\theta\|_{\mathcal{H}}^2$ by definition of F and (ii) follows from $\alpha_t^\theta = -\sigma_t Q^{1/2} \mathbf{Y}_t^\theta$. Since the stopgrad operation does not affect any step in the derivation, therefore, if $\mathbf{Y}_t^\theta = D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\theta)$, then $\mathbf{Y}_t^\theta$ is the solution of (19) and the critical point in (122) satisfies

$$\alpha_t^\theta = -\sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}} = -\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t. \quad (133)$$

Hence, by Lemma B.5, this control satisfies the HJB and is optimal, hence $\alpha^\theta = \alpha^*$. Conversely, if we assume that the critical point $\alpha_t^\theta = -\sigma_t Q^{1/2} \mathbf{Y}_t^\alpha$ is optimal where $\mathbf{Y}_t^\alpha$ is the solution of (19), then the same verification immediately yields $\mathbf{Y}_t^\alpha = D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha)$. This shows that the critical point in (122) is the optimal control α^* for the SOC problem in (12). This completes the proof. $\square$

Sufficient condition Indeed, in the special case *i.e.*, when g is convex, the converse becomes a sufficient condition. One may ask whether solving the adjoint equation and minimizing the Hamiltonian ensures optimality. In general, SOC sampling needs non convex g , so this condition often fails. We therefore do not apply it in our study. Instead, we formulate and prove a sufficient optimality condition tailored to specific SOC problem in (12).

Proposition B.7 (Stochastic Maximum Principle). *Assume the terminal cost g is convex. Let us fix an admissible control $\bar{\alpha} \in \mathbb{A}$ and consider the corresponding infinite dimensional coupled FBSDEs:*

$$d\mathbf{X}_t^\alpha = [\mathcal{A}\mathbf{X}_t^\alpha + \sigma_t Q^{1/2} \bar{\alpha}_t] dt + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{X}_0^\alpha = \mathbf{x}_0, \quad (134)$$

$$d\mathbf{Y}_t^\alpha = -[\mathcal{A}^\dagger \mathbf{Y}_t^\alpha + D_{\mathbf{x}} H(t, \mathbf{X}_t^\alpha, \bar{\alpha}_t, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha)] dt + \mathbf{Z}_t^\alpha d\mathbf{W}_t^Q, \quad \mathbf{Y}_T^\alpha = D_{\mathbf{x}} g(\mathbf{X}_T^\alpha). \quad (135)$$

Then, under the state variables is fixed $(\mathbf{X}_t^\alpha, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha)$ for a choosen control $\bar{\alpha}$, if following holds:

$$H(t, \mathbf{X}_t^\alpha, \alpha_t^*, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha) = \min_{\alpha \in \mathbb{A}} H(t, \mathbf{X}_t^\alpha, \alpha, \mathbf{Y}_t^\alpha, \mathbf{Z}_t^\alpha), \quad \forall t \in [0, T], \quad (136)$$

the resulting sequence of controls $\{\alpha_t^*\}_{t \in [0, T]}$ is the optimal control α^* for a given SOC problem (12).

Proof. We adapt the proof of the sufficient condition in (Carmona, 2016, Theorem 4.14) into our infinite-dimensional control problem. Fix $\beta \in \mathcal{U}$ be a admissible control, let us write $\mathbf{X}_t^\beta := \mathbf{X}_t^\beta$ and define $\Delta \mathbf{X}_t := \mathbf{X}_t^\beta - \mathbf{X}_t^*$. By convexity of g , we get

$$g(\mathbf{X}_T^\beta) - g(\mathbf{X}_T^*) \geq \langle D_{\mathbf{x}} g(\mathbf{X}_T^\beta), \mathbf{X}_T^\beta - \mathbf{X}_T^* \rangle_{\mathcal{H}} = \langle \mathbf{Y}_T^*, \Delta \mathbf{X}_T \rangle_{\mathcal{H}}. \quad (137)$$

Now, apply the Hilbert space product rule B.4 to $\langle \mathbf{Y}_T^*, \Delta \mathbf{X}_T \rangle_{\mathcal{H}}$. Then, by using the dynamics of $\mathbf{X}^\beta$ and $\mathbf{X}^*$ and the adjoint equation for $\mathbf{Y}^*$, we get:

$$\mathbb{E} [\langle \mathbf{Y}_T^*, \Delta \mathbf{X}_T \rangle_{\mathcal{H}}] = \mathbb{E} \int_0^T \langle -\mathcal{A}^\dagger \mathbf{Y}_t^* - D_{\mathbf{x}} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \Delta \mathbf{X}_t) \rangle_{\mathcal{H}} dt \quad (138)$$

$$+ \mathbb{E} \int_0^T \langle \mathbf{Y}_t^*, \mathcal{A} \Delta \mathbf{X}_t + \sigma_t Q^{1/2} \beta_t - \sigma_t Q^{1/2} \alpha_t^* \rangle_{\mathcal{H}} dt \quad (139)$$

$$\stackrel{(i)}{=} -\mathbb{E} \int_0^T \langle D_{\mathbf{x}} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \Delta \mathbf{X}_t) \rangle_{\mathcal{H}} dt \quad (140)$$

$$+ \mathbb{E} \int_0^T \langle \mathbf{Y}_t^*, \sigma_t Q^{1/2} \beta_t - \sigma_t Q^{1/2} \alpha_t^* \rangle_{\mathcal{H}} dt, \quad (141)$$

where (i) follows from the adjointness of $\mathcal{A}$. On the other hand, using the definition of $\mathcal{J}$ in (108):

$$\mathcal{J}(\beta) - \mathcal{J}(\alpha^*) = \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \|\beta_t\|_{\mathcal{H}}^2 - \frac{1}{2} \|\alpha_t^*\|_{\mathcal{H}}^2 \right) dt + g(\mathbf{X}_T^\beta) - g(\mathbf{X}_T^*) \right] \quad (142)$$

$$\geq \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \|\beta_t\|_{\mathcal{H}}^2 - \frac{1}{2} \|\alpha_t^*\|_{\mathcal{H}}^2 \right) + \langle \mathbf{Y}_t^*, \sigma_t Q^{1/2} \beta_t - \sigma_t Q^{1/2} \alpha_t^* \rangle_{\mathcal{H}} dt \right] \quad (143)$$

$$- \mathbb{E} \int_0^T \langle D_{\mathbf{x}} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \Delta \mathbf{X}_t) \rangle_{\mathcal{H}} dt \quad (144)$$

$$\stackrel{(i)}{=} \mathbb{E} \int_0^T \left[H(t, \mathbf{X}_t^\beta, \beta_t, \mathbf{Y}_t^*, \mathbf{Z}_t^*) - H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*) - \langle D_{\mathbf{x}} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \Delta \mathbf{X}_t) \rangle_{\mathcal{H}} \right] dt, \quad (145)$$

where (i) follows by the definition of Hamiltonian (16). By convexity of $\mathcal{H}$ with respect to $(\mathbf{x}, \alpha)$,

$$H(t, \mathbf{X}_t^\beta, \beta_t, \mathbf{Y}_t^*, \mathbf{Z}_t^*) \geq H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*) \quad (146)$$

$$+ \langle D_{\mathbf{x}} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \Delta \mathbf{X}_t) \rangle_{\mathcal{H}} + \langle D_{\alpha} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \beta_t - \alpha_t^*) \rangle_{\mathcal{H}}. \quad (147)$$

Then, the point-wise minimality of α^* implies that $\langle D_{\alpha} H(t, \mathbf{X}_t^*, \alpha^*, \mathbf{Y}_t^*, \mathbf{Z}_t^*, \beta_t - \alpha_t^*) \rangle_{\mathcal{H}} \geq 0$ for any admissible control β . Therefore, it implies that $\mathcal{J}(\beta) \geq \mathcal{J}(\alpha^*)$ because of (136). Finally, since we choose β arbitrary, resulting α^* is optimal control for all $t \in [0, T]$. This concludes the proof. $\square$

B.6. Proof of Proposition 5

Proposition 5 (Unbiased Estimator). *Define the sample-wise adjoint matching objective with conditional expectation $\tilde{\mathbf{Y}}_t^{\bar{\alpha}} := \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\mathbf{Y}_t^{\bar{\alpha}} | \mathbf{X}_t^{\bar{\alpha}}]$ for any sample trajectory $\mathbf{X}_t^{\bar{\alpha}} \sim \mathbb{P}^{\bar{\alpha}}$:*

$$\tilde{\mathcal{L}}(\theta) := \int_0^T \frac{1}{2} \|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}\|_{\mathcal{H}}^2 dt \quad (20)$$

Then, we get unbiased gradient estimator $\frac{d}{d\theta} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\tilde{\mathcal{L}}(\theta)] = \frac{d}{d\theta} \mathcal{L}(\theta)$ with same critical point in (18).

Proof. Under the stopgrad operation, the forward law $\mathbb{P}^{\bar{\alpha}}$ does not depend on θ . By the dominated convergence theorem we pass the derivative through the expectation and the integral:

$$\frac{d}{d\theta} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\tilde{\mathcal{L}}(\theta)] = \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\frac{d}{d\theta} \int_0^T \frac{1}{2} \|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}\|_{\mathcal{H}}^2 dt \right] \quad (148)$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}} dt \right] \quad (149)$$

$$= \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] \right] dt \quad (150)$$

$$= \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \sigma_t Q^{1/2} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\mathbf{Y}_t^{\bar{\alpha}} | \mathbf{X}_t^{\bar{\alpha}}], \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] \right] dt \quad (151)$$

$$\stackrel{(i)}{=} \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} [\langle \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}}] \right] dt \quad (152)$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}} \left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) \rangle_{\mathcal{H}} dt \right] \quad (153)$$

$$\stackrel{(ii)}{=} \frac{d}{d\theta} \mathcal{L}(\theta), \quad (154)$$

$$(155)$$

where (i) follows from the tower property and (ii) follows from (117). Hence, we get $\frac{d}{d\theta} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\tilde{\mathcal{L}}(\theta)] = \frac{d}{d\theta} \mathcal{L}(\theta) = 0$ at $\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) = -\sigma_t Q^{1/2} \mathbf{Y}_t^{\bar{\alpha}}$. It shows that the critical point remains unchanged to $\frac{d}{d\theta} \mathcal{L}(\theta)$. It concludes the proof. $\square$

B.7. Proof of Proposition 6

Proposition 6 (Laplacian Operator). *For the negative Laplacian $\mathcal{A} = -\Delta$ on a bounded domain with Dirichlet boundary condition, take $Q = \mathcal{A}^{-s}$ with $s > \frac{d}{2}$. Then, $\mathcal{A}$ and Q satisfy the conditions in Theorem B.1.*

Proof. (A) **Trace-class** Let $\{\phi^{(k)}\}_{k \geq 1}$ be an eigen-basis with $\mathcal{A}\phi^{(k)} = -\lambda^{(k)}\phi^{(k)}$. Then, we have

$$Q\mathbf{x} = \sum_{k \geq 1} (\lambda^{(k)})^{-s} \mathbf{x}^{(k)} \phi^{(k)}. \quad (156)$$

Hence the eigenvalues of Q are $(\lambda^{(k)})^{-s}$ and it implies that $\text{Tr}(Q) = \sum_{k \geq 1} (\lambda^{(k)})^{-s}$. By the Weyl law $\lambda^{(k)} \asymp k^{\frac{2}{d}}$ on a bounded C^∞ domain in $\mathbb{R}^d$, therefore $\text{Tr}(Q) < \infty$ if $s > \frac{d}{2}$.

(B) Self-adjoint Since Laplacian Δ is self-adjoint, we get adjointness of $\mathcal{A}$ directly, and for Q :

$$Q^\dagger = ((-\mathcal{A})^{-s})^\dagger = ((-\mathcal{A})^{-s}) = Q \quad (157)$$

(C) Dissipative Since $0 < \lambda_1 < \lambda_2, \dots < \lambda_k < \dots$ for $k \in \mathbb{N}$, we obtain:

$$\langle \mathcal{A}\mathbf{x}, \mathbf{x} \rangle_{\mathcal{H}} = - \sum_{k \geq 1} \lambda^{(k)} \left\| \langle \mathbf{x}, \phi^{(k)} \rangle_{\mathcal{H}} \right\|^2 \leq -\lambda^{(1)} \sum_{k \geq 1} \left\| \langle \mathbf{x}, \phi^{(k)} \rangle_{\mathcal{H}} \right\|^2 = -\lambda^{(1)} \|\mathbf{x}\|_{\mathcal{H}}^2. \quad (158)$$

(D) Commute For any $\mathbf{x} \in \mathcal{H}$, we get:

$$\mathcal{A}Q\mathbf{x} = \mathcal{A} \left(\sum_{k \geq 1} (\lambda^{(k)})^{-s} \mathbf{x}^{(k)} \phi^{(k)} \right) = \sum_{k \geq 1} (\lambda^{(k)})^{-s} \mathbf{x}^{(k)} (\mathcal{A}\phi^{(k)}) \quad (159)$$

$$= \sum_{k \geq 1} (\lambda^{(k)})^{-s} \mathbf{x}^{(k)} (-\lambda^{(k)} \phi^{(k)}) = - \sum_{k \geq 1} (\lambda^{(k)})^{1-s} \mathbf{x}^{(k)} \phi^{(k)} \quad (160)$$

and

$$Q\mathcal{A}\mathbf{x} = -Q \left(\sum_{k \geq 1} \lambda^{(k)} \mathbf{x}^{(k)} \phi^{(k)} \right) = - \sum_{k \geq 1} \lambda^{(k)} \mathbf{x}^{(k)} (Q\phi^{(k)}) \quad (161)$$

$$= \sum_{k \geq 1} \lambda^{(k)} \mathbf{x}^{(k)} ((\lambda^{(k)})^{-s} \phi^{(k)}) = - \sum_{k \geq 1} (\lambda^{(k)})^{1-s} \mathbf{x}^{(k)} \phi^{(k)}. \quad (162)$$

Therefore, we get $\mathcal{A}Q^{1/2}\mathbf{x} = Q^{1/2}\mathcal{A}\mathbf{x}$. It completes the proof. $\square$

B.8. Endpoint Disintegration

From (10), let us denote $\mathbf{X} := \{\mathbf{X}_t\}_{t \in [0, T]} \in \Omega$ and define the IS as follows:

$$w(\mathbf{X}) = e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{x}_0, \mathbf{X}_T)}. \quad (163)$$

Then, the RND for $\mathbb{P}^*$ with respect to reference measure $\mathbb{P}$ is given by:

$$\frac{d\mathbb{P}^*}{d\mathbb{P}}(\mathbf{X}) = \frac{w(\mathbf{X})}{\mathbb{E}_{\mathbb{P}}[w(\mathbf{X})]}. \quad (164)$$

Then, for any bounded functional $F : \Omega \rightarrow \mathbb{R}$, we have the expectation:

$$\mathbb{E}_{\mathbb{P}^*}[F(\mathbf{X}^*)] = \frac{\mathbb{E}_{\mathbb{P}}[F(\mathbf{X}) w(\mathbf{X})]}{\mathbb{E}_{\mathbb{P}}[w(\mathbf{X})]}. \quad (165)$$

Since terminal law of $\mathbb{P}$ admits the density q_T with respect to ν_∞ based on Theorem 2, we can disintegrate $\mathbb{P}$ with respect to $\mathbf{X}_T$ to obtain:

$$\mathbb{E}_{\mathbb{P}}[F(\mathbf{X})w(\mathbf{x})] = \int_{\mathcal{H}} \left(\int_{\mathcal{H}} F(\mathbf{X}) \mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y}) \right) e^{-U(\mathbf{y}) - \log q_T(\mathbf{x}_0, \mathbf{x}_T)} q_T(\mathbf{x}_0, \mathbf{y}) d\nu_\infty(\mathbf{y}) \quad (166)$$

$$= \int_{\mathcal{H}} \left(\int_{\mathcal{H}} F(\mathbf{X}) \mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y}) \right) e^{-U(\mathbf{y})} d\nu_\infty(\mathbf{y}) \quad (167)$$

$$\stackrel{(i)}{=} \int_{\mathcal{H}} \left(\int_{\mathcal{H}} F(\mathbf{X}) \mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y}) \right) \mathbb{E}_{\mathbb{P}}[w(\mathbf{X})] d\pi(\mathbf{y}), \quad (168)$$

where (i) follows from the definition of target distribution π in (6). Hence, we obtain that

$$\mathbb{E}_{\mathbb{P}^*}[F(\mathbf{X}^*)] = \int \left(\int F(\mathbf{X}) \mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y}) \right) \pi(d\mathbf{y}) \quad (169)$$

$$= \mathbb{E}_{\mathbb{P}_{\cdot|T} \mathbb{P}_T^*}[F(\mathbf{X})]. \quad (170)$$

C. Experimental Details

C.1. Bridge sampling

Because the solution of the reference process in (3) defines an Ornstein-Uhlenbeck semigroup, the conditional law at time t given time T is:

$$\mathbb{P}_{t|T}(\mathbf{X}_t|\mathbf{x}_T) = \mathcal{N}(\mathbf{m}_{t|T}, Q_{t|T}), \quad (171)$$

where the conditional mean function and covariance operator is given by:

$$\mathbf{m}_{t|T} = \mathbf{x}_0 + Q_{t|T}Q_T^{-1}(\mathbf{x}_T - \mathbf{x}_0), \quad Q_{t|T} = Q_T - Q_{t|T}Q_T^{-1}Q_{t|T} \quad (172)$$

$$Q_t = \int_0^t e^{(t-s)\mathcal{A}} \sigma_s^2 Q e^{(t-s)\mathcal{A}^\dagger} ds, \quad (173)$$

$$Q_{t|T} = \int_0^t e^{(t-s)\mathcal{A}} \sigma_s^2 Q e^{(T-s)\mathcal{A}^\dagger} ds = Q_t e^{(T-t)\mathcal{A}^\dagger}, \quad Q_{T|t} = (Q_{t|T})^\dagger. \quad (174)$$

Moreover, the coordinate process $\mathbf{X}_t^{(k)} = \langle \mathbf{X}_t, \phi^{(k)} \rangle$ of the Ornstein-Uhlenbeck semigroup in (3) is also Gaussian and the coordinates are independent. Therefore we can compute coordinate-wise:

$$\mathbb{P}_{t|T}(\mathbf{X}_t^{(k)}|\mathbf{x}_T^{(k)}) = \mathcal{N}(\mathbf{m}_{t|T}^{(k)}, q_{t|T}^{(k)}) \quad (175)$$

where we denote $q_t^{(k)} = \langle Q_t, \phi^{(k)} \rangle$ for any $t \in [0, T]$ and

$$\mathbf{m}_{t|T}^{(k)} = \mathbf{x}_0^{(k)} + \frac{q_t^{(k)}}{q_\infty^{(k)}} e^{-(T-t)\lambda^{(k)}} (\mathbf{x}_T^{(k)} - \mathbf{x}_0^{(k)}), \quad (176)$$

$$q_{t|T}^{(k)} = q_t^{(k)} - \frac{(q_t^{(k)})^2}{q_\infty^{(k)}} e^{-2(T-t)\lambda^{(k)}} \quad (177)$$

$$q_t^{(k)} = \int_0^t \sigma_s^2 e^{-2(t-s)\lambda^{(k)}} (\lambda^{(k)})^{-s} ds, \quad \forall t \in [0, T]. \quad (178)$$

C.2. Enforcing boundary condition

To properly define a transition path between two metastable states, say $\mathbf{A}$ and $\mathbf{B}$, we must fix these states as the boundary conditions for a desired path. Our method achieves this by decomposing the path $\mathbf{X}_t$ into two parts: a fixed reference path $\mathbf{x}_0$ and a fluctuating residual path $\mathbf{R}_t$. Then, the total path is given by $\mathbf{X}_t = \mathbf{x}_0 + \mathbf{R}_t$, where $\mathbf{x}_0$ lifting the residual path (boundary values are zero) to the path that have desired boundary values. The reference path $\mathbf{x}_0$ provides the direct connection from $\mathbf{A}$ to $\mathbf{B}$, while the residual path $\mathbf{R}_t$ is forced to be zero at the boundaries. This construction guarantees that the full path $\mathbf{X}_t$ always begins at $\mathbf{A}$ and ends at $\mathbf{B}$.

For a residual path $\mathbf{R}_t$, we enforce the sample-path endpoints to be a absorbed states *i.e.*, $\mathbf{R}_t[0] = 0, \mathbf{R}_t[L] = 0$ for a function $\mathbf{R}_t : \mathbb{R}_{>0} \rightarrow \mathbb{R}^d$ by choosing Dirichlet eigen-basis:

$$\phi^{(k)}[u] = \left(\frac{2}{L}\right)^{1/2} \sin\left(\frac{\pi k u}{L}\right), \quad \forall k \in \mathbb{N}, u \in \mathbf{L} := [0, L] \quad (179)$$

$$\mathbf{R}_t = \sum_{k \geq 1} \mathbf{R}_t^{(k)} \phi^{(k)}, \quad \text{where } \mathbf{R}_t^{(k)} := \langle \mathbf{R}_t, \phi^{(k)} \rangle_{\mathcal{H}} = \int_{\mathbf{L}} \mathbf{R}_t[u] \phi^{(k)}[u] du. \quad (180)$$

Therefore, we get $\mathbf{R}_t[0] = \mathbf{R}_t[L] = 0$ for all $t \in [0, T]$ since $\phi^{(k)}[0] = \phi^{(k)}[L] = 0$ from (180). For a reference path $\mathbf{x}_0$, we define the as a continuous function $\mathbf{x}_0 : \mathbf{L} \rightarrow \mathbb{R}^d$:

$$\mathbf{x}_0[u] = c_0[u]\mathbf{A} + c_1[u]\mathbf{B}, \quad \text{with } \begin{cases} c_0[0] = 1, c_0[L] = 0 \\ c_1[0] = 0, c_1[L] = 1 \end{cases} \quad (181)$$

Here, the residual process $\mathbf{R}_t$ on the Dirichlet space is given by:

$$d\mathbf{R}_t = \left[\mathcal{A}\mathbf{R}_t + \sigma_t Q^{1/2} \alpha(t, \mathbf{x}_0 + \mathbf{R}_t) \right] dt + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{R}_0 = 0. \quad (182)$$

It results that since $\mathbf{X}_t = \mathbf{x}_0 + \mathbf{R}_t$, we can deduce the following $\mathcal{H}$ -valued SDE:

$$d\mathbf{X}_t = \left[\mathcal{A}(\mathbf{X}_t - \mathbf{x}_0) + \sigma_t Q^{1/2} \alpha(t, \mathbf{X}_t) \right] dt + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{X}_0 = \mathbf{x}_0. \quad (183)$$

C.3. DST-based Galerkin SDE simulation

Given Dirichlet eigen-basis in (179), we can express the infinite-dimensional SDEs into infinite system of real-valued SDEs for each coordinate k :

$$d\mathbf{R}_t^{(k)} = \left[-\lambda^{(k)} \mathbf{R}_t^{(k)} + \sigma_t \langle \alpha(t, \mathbf{x}_0 + \mathbf{R}_t), (Q^{1/2})^\dagger \phi^{(k)} \rangle \right] dt + \sigma_t (\lambda^{(k)})^{-\frac{s}{2}} d\beta_t^{(k)}, \quad \mathbf{R}_t^{(k)} = 0, \quad (184)$$

$$= \left[-\lambda^{(k)} \mathbf{R}_t^{(k)} + \sigma_t (\lambda^{(k)})^{-\frac{s}{2}} \langle \alpha(t, \mathbf{X}_t), \phi^{(k)} \rangle \right] dt + \sigma_t (\lambda^{(k)})^{-\frac{s}{2}} d\beta_t^{(k)} \quad (185)$$

where $\{\beta_t^{(k)}\}_{k \geq 1}$ are standard Wiener processes. Note that, to compute $(Q^{1/2})^\dagger \phi^{(k)}$, we choose $\mathcal{U} = \mathcal{H}$, means $\mathbf{W}_t = \sum_{k \geq 1} \phi_k \beta_t^{(k)}$. Even though $\mathbf{W}_t$ is not an $\mathcal{H}$ -valued random variable, it is a cylindrical process indexed by $\mathcal{H}$, applying Hilbert-Schmidt $Q^{1/2}$ before stochastic integral yields the resulting process remain square integrable.

If we treat the finite path evaluations as samples on a uniform grid $\mathbf{L} = [0, L]$, we project onto the sine basis with the discrete sine transform (DST). Since a finite-resolution path has a natural cutoff frequency K as a number of evaluation points, projection to and from the sine basis is exact. Concretely, approximate the Laplacian by $\Delta \triangleq \mathbf{E} \mathbf{D} \mathbf{E}^\top$, where $\mathbf{E}^\top$ is the DST projection consists of eigen functions $\{\phi^{(k)}\}_{k=1}^K$ so that $\tilde{\mathbf{R}}_t = \mathbf{E}^\top \mathbf{R}_t$ and $\mathbf{D}$ is a diagonal matrix containing the eigen values $\{\lambda^{(k)}\}_{k=1}^K$. Hence, stochastic evolution in (182) is described by the finite-dimensional model (configuration to spectral space):

$$d\tilde{\mathbf{R}}_t = \left[-\mathbf{D} \tilde{\mathbf{R}}_t + \sigma_t \mathbf{D}^{-\frac{s}{2}} \mathbf{E}^\top \alpha(t, \mathbf{X}_t) \right] dt + \sigma_t \mathbf{D}^{-\frac{s}{2}} d\tilde{\mathbf{W}}_t. \quad (186)$$

The simulation algorithm for FAS is provided in detail in Algorithm 3. Additionally, we introduce a spectral precision scale κ on the Dirichlet Laplacian, replacing $\mathbf{D}$ by $\kappa^2 \mathbf{D}$ to control stiffness and smoothness of the reference dynamics.

C.4. overNumerical Computation

Correction RND From Proposition 6, we can calculate the RND in Theorem 2 as follows:

$$\log q_t(\mathbf{x}_0, \mathbf{x}) = -\frac{1}{2} \sum_{k \geq 1} \frac{1}{q_\infty^{(k)} (e^{2t\lambda_k} - 1)} \left[\|\mathbf{x}_0^{(k)}\|^2 - 2e^{-t\lambda_k} \langle \mathbf{x}_0^{(k)}, \mathbf{x}^{(k)} \rangle + \|\mathbf{x}^{(k)}\|^2 \right], \quad (187)$$

where $q_\infty^{(k)}$ is diagonal component of Q_∞ such that $Q_\infty \phi^{(k)} = q_\infty^{(k)} \phi^{(k)}$. From (45), we get $Q_\infty = -\frac{1}{2} \sigma_\infty^2 \mathcal{A}^{-1} Q$, therefore we get $q_\infty^{(k)} = \frac{1}{2} \sigma_\infty^2 (\lambda^{(k)})^{(1+s)}$.

Noise schedule σ_t The closed-form simulation may be infeasible because computing $q_t^{(k)}$ in (178) requires numerical integration for arbitrary noise schedule σ_t . In this case, the two noise schedules studied for adjoint sampling in (Havens et al., 2025; Liu et al., 2026) can be adapted to our setting.

- The constant noise schedule sets $\sigma(t) := \sigma$. Then, we can compute the integral in (178):

$$q_t^{(k)} = (\lambda^{(k)})^{-s} \frac{(1 - e^{-2t\lambda_k})}{2\lambda_k}. \quad (188)$$

- The geometric noise schedule (Song et al., 2021; Karras et al., 2022) decreases monotonically by predefined $\beta_{\min}$ and $\beta_{\max}$:

$$\sigma_t = \begin{cases} \beta_{\min} \left(\frac{\beta_{\max}}{\beta_{\min}} \right)^{T-t} \sqrt{2 \log \frac{\beta_{\max}}{\beta_{\min}}}, & t < T \\ \beta_{\min} \sqrt{2 \log \frac{\beta_{\max}}{\beta_{\min}}}, & t \geq T. \end{cases} \quad (189)$$

Because choosing the noise schedule $\sigma_t = \beta e^{\alpha t}$ gives a closed-form for the integral in (178), we rearrange it as follows:

$$\sigma_t = \beta_{\min} \left(\frac{\beta_{\max}}{\beta_{\min}} \right)^{T-t} \sqrt{2 \log \frac{\beta_{\max}}{\beta_{\min}}} = \beta_{\min} \left(\frac{\beta_{\max}}{\beta_{\min}} \right)^T \sqrt{2 \log \frac{\beta_{\max}}{\beta_{\min}}} e^{-\log \left(\frac{\beta_{\max}}{\beta_{\min}} \right) t}. \quad (190)$$

Table 4. Training Hyper-parameters

	Synthetic potential	Alanine dipeptide	Chignolin
$\beta_{\min}$	0.1	0.01	0.01
$\beta_{\max}$	10	15	5
κ	10^{-2}	10^{-5}	10^{-5}
M	10^3	10^4	10^4
L	100	100	100
N	512	16	16
$ \mathcal{B} $	10^4	10^3	10^3
$\alpha_{\max}$	100	10^4	10^4
$ \mathbf{U} $	100	100	100
Learning rate	10^{-4}	10^{-5}	10^{-4}
Base channel	32	64	128
Base mode	8	8	8
Embed dim	128	256	512

Therefore, we get $\beta = \beta_{\min} \left(\frac{\beta_{\max}}{\beta_{\min}} \right)^T \sqrt{2 \log \frac{\beta_{\max}}{\beta_{\min}}}$, $\alpha = -\log \left(\frac{\beta_{\max}}{\beta_{\min}} \right)$. This choice yields that:

$$q_t^{(k)} = \beta^2 (\lambda^{(k)})^{-s} \frac{e^{2\alpha t} - e^{-2\lambda_k t}}{2(\alpha + \lambda_k)} \quad (191)$$

Replay buffer $\mathcal{B}$ Following previous diffusion samplers (Havens et al., 2025; Liu et al., 2026; Choi et al., 2026), we maintain a replay buffer $\mathcal{B}$ for the computation of the objective in (22). In this case, the expectation under $\mathbb{P}^\alpha$ is approximated by an average over $\mathcal{B}$, which holds the most recent $|\mathcal{B}|$ samples. We refresh the buffer $\mathcal{B}$ by inserting N new samples after every L gradient steps.

Parametrization fo control network We parameterize the control α^θ with a neural operator to preserve its functional form. We parametrize the control as $\alpha^\theta(t, \mathbf{x}) := \sigma_t Q^{1/2} e^{-(T-t)\mathcal{A}^\dagger} u^\theta(t, \mathbf{x})$, which absorbs the noise schedule σ_t and Laplacian contraction $Q^{1/2} e^{-(T-t)\mathcal{A}^\dagger}$ into the control and removes its explicit appearance from the matching objective in (22), yielding more stable training. For all experiment, we take u^θ to be a modified U-shaped Neural Operator (UNO) (Rahman et al., 2023) tailored to diffusion-style conditioning. Specifically, the control takes a 1D sequence $\mathbf{x} \in \mathbb{R}^{|\mathbf{L}| \times \mathbb{R}^d}$ and a time input t . We embed t with a sinusoidal MLP and build simple Fourier grid features, then compress these into a single conditioning vector. Each encoder and decoder block applies a spectral convolution on low Fourier modes with a point-wise path, then uses adaptive groupnorm modulated by conditioning vector. Conditioning enters only through normalization, which keeps the operator discretization agnostic and suitable for diffusion training. We set a base channel dimension c and set channel multiplier $[1, 2, 4]$. Each stage uses two spectral residual blocks. Conversely, we set a base Fourier modes and set channel multiplier $[1, 0.5, 0.25]$. The timestep embedding and the conditioning vector use dimension equal to the highest channel.

Clipping $\alpha_{\max}$ We clip the energy gradient so its maximum norm does not exceed $\alpha_{\max}$

Algorithms The overall algorithm of the implementation of FAS is summarized in the Algorithm 4.

C.5. Possible parameterization strategie for other domain

Rectangular domain Let $\mathbf{X} : \mathbb{R}^2 \rightarrow \mathbb{R}^d$ be a field on a rectangular subset of $\mathbb{R}^2$ with zero Neumann BCs, cosine modes from the natural basis. Because real data are sampled on a finite grid, the spectrum already contains a built-in cut-off frequency. We can therefore write the operator with a finite eigen decomposition via the discrete cosine transform (DCT) (Rissanen et al., 2023; Lim et al., 2023b; Park et al., 2024). Concretely, we can approximate the Laplacian operator as $\mathcal{A} \triangleq \mathbf{E} \mathbf{D} \mathbf{E}^\top$ with cosine projection matric $\mathbf{E}^\top$ and diagonal eigen values $\{\lambda^{(k)}\}_{k=1}^K$.

Graph domain For a function $\mathbf{X} : \mathcal{G} \rightarrow \mathbb{R}^3$ of a discrete graph $\mathcal{G} = \langle \mathbf{V}, \mathcal{E} \rangle$, where $\mathbf{V} = \{v_i\}_{i=1}^n$ is the set of vertices and $\mathcal{E} = \{e_{ij} \mid (i, j) \subseteq |\mathbf{V}| \times |\mathbf{V}|\}$ is the set of edges, we can choose the eigenvectors of the symmetric graph Laplacian $\mathcal{A}_{\text{sym}}$

Algorithm 3 FAS Sampling with path lifting

Require: Initial condition $\mathbf{x}_0$, linear operator $\mathcal{A}$, trace-class operator Q , control α , discrete and inverse sine transforms DST , iDST , set of evaluation points $\mathbf{U} = \{u_i\}_{i=1}^K$.

- 1: Set initial condition $\mathbf{X}_0^\alpha = \mathbf{x}_0$ and $\mathbf{R}_0^\alpha = \mathbf{0}$
- 2: **for** time t **in range** $[0, T]$ **do**
- 3: Estimate control $\alpha_t = \alpha(\mathbf{X}_t^\alpha, t)$
- 4: Sample Gaussian noise $\varepsilon \sim \mathcal{N}(0, \mathbf{I}_K)$
- 5: Project state, control $\{\tilde{\mathbf{R}}_t^{(k)}\}_{k \in \mathbf{U}} = \text{DST}(\mathbf{R}_t^\alpha), \{\tilde{\alpha}_t^{(k)}\}_{k \in \mathbf{U}} = \text{DST}(\alpha_t)$
- 6: **for** $k = 1, \dots, |\mathbf{U}|$ **do in parallel**
- 7: $\tilde{\mathbf{R}}_{t+\delta_t}^{(k)} = \left[-\lambda^{(k)} \tilde{\mathbf{R}}_t^{(k)} + \sigma_t(\lambda^{(k)})^{-\frac{s}{2}} \tilde{\alpha}_t^{(k)} \right] \delta_t + \sigma_t(\lambda^{(k)})^{-\frac{s}{2}} \varepsilon \sqrt{\delta_t}$
- 8: **end for**
- 9: Get residual $\mathbf{R}_{t+\delta_t}^\alpha = \text{iDST}(\{\tilde{\mathbf{R}}_{t+\delta_t}^{(k)}\}_{k \in \mathbf{U}})$
- 10: Update $\mathbf{R}_t^\alpha = \mathbf{R}_{t+\delta_t}^\alpha$
- 11: Get path $\mathbf{X}_t^\alpha = \mathbf{x}_0 + \mathbf{R}_t^\alpha$
- 12: **end for**
- 13: **return** $\mathbf{X}_T^\alpha \sim \mathbb{P}_T^\alpha$

as the spectral basis. Each graph signal is already finite-dimensional, with $|\mathbf{V}|$ modes acting as an inherent spectral cut-off. The operator admits a graph Fourier decomposition $\mathcal{A}_{\text{sym}} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$, where $\mathbf{U}^\top$ holds the Laplacian eigenvectors and $\mathbf{\Lambda}$ is diagonal with eigen values $\{\lambda^{(k)}\}_{k=1}^{|\mathbf{V}|}$. Specifically, in molecule conformal generation, $\mathbf{X} : \mathcal{G} \rightarrow \mathbb{R}^3$ represents a conformal field (Wang et al., 2024). Note that unlike the rectangular or time-series domain, where every function shares a common basis, each molecule $\mathcal{G}_i$ has its own graph, so we should computed the corresponding eigen-system for each training sample.

C.6. Potential Functions and Initial conditions

For consistency and fair comparison, we follow the respective experimental setups established by baseline methods. For the Müller Brown potential, we follow the experimental setup and reported results from (Du et al., 2024). for the molecule potentials, we follow (Seong et al., 2025) and reported results from (Seong et al., 2025; Blessing et al., 2026).

- The synthetic Müller–Brown potential $V : \mathbb{R}^2 \rightarrow \mathbb{R}$ used in our experiments is given by

$$V(x, y) = -200 \cdot \exp(-(x_1 - 1)^2 - 10y^2) - 100 \cdot \exp(-x^2 - 10 \cdot (y - 0.5)^2) \quad (192)$$

$$- 170 \cdot \exp(-6.5 \cdot (x + 0.5)^2 + 11 \cdot (x + 0.5) \cdot (y - 1.5) - 6.5 \cdot (y - 1.5)^2) \quad (193)$$

$$+ 15 \cdot \exp(0.7 \cdot (x + 1)^2 + 0.6 \cdot (x + 1) \cdot (y - 1) + 0.7 \cdot (y - 1)^2). \quad (194)$$

For synthetic potential, we set initial condition $\mathbf{x}_0$ as linear interpolation between two local minimums $\mathbf{A} = (-0.558, 1.442)$ and $\mathbf{B} = (0.624, 0.028)$.

- For real-world molecule system, we implement the potential with OpenMM library (Eastman et al., 2023). Since OpenMM library dose not provide automatic differentiation, we wrap it in PyTorch (Paszke et al., 2019) and expose batched energy and force calls efficiently.

We use `amber99sbildn` (Lindorff-Larsen et al., 2010) for alanine dipeptide in vacuum and set $\mathbf{A} = \text{C5}$ and $\mathbf{B} = \text{C7ax}$. For Chignolin we use `ff14SBonlysc` (Maier et al., 2015) with the `gbn2` solvent (Nguyen et al., 2013) and we set $\mathbf{A}$ as unfolded protein and $\mathbf{B}$ as folded protein by following the setting in (Seong et al., 2025).

C.7. Transition Path Sampling on Molecule Potential

Path Initialization Due to the rugged and stiff nature of molecular potentials, linear interpolation alone can produce clashing and high energy images along the path, which might produces sub-optimal solutions. To avoid this problem, we introduce additional refinement strategy for molecule potentials.

Given an initial condition $\mathbf{x}_0$, we linearly interpolate Cartesian coordinates between $\mathbf{A}$ and $\mathbf{B}$, then refine the path using the *image dependent pair potential* (IDPP) (Smidstrup et al., 2014) before training. Let $\mathbf{C} = (\mathbf{r}_1, \dots, \mathbf{r}_N)$ denote a configuration

Algorithm 4 Function Space Adjoint Sampler (FAS)

Require: Initial condition $\mathbf{x}_0$, terminal cost $g(x)$, control network $\alpha^\theta(\mathbf{x}, t)$, replay buffer $\mathcal{B}$, total epoch M , resample sizes N , gradient steps L , maximum energy norm $\alpha_{\max}$.

- 1: **for** epoch m in $1, 2, \dots, M$ **do**
- 2: Sample N paths $\mathbf{X}_T^{\bar{\alpha}^m}$ from Algorithm 3 with $\bar{\alpha}^m = \text{stopgrad}(\alpha_{\theta}^{\theta_m})$
- 3: Compute adjoint $\mathbf{Y}_T^{\bar{\alpha}^m} = \text{clip}(\nabla g(\mathbf{X}_T^{\bar{\alpha}^m}), \alpha_{\max})$
- 4: Update buffer $\mathcal{B} \leftarrow \mathcal{B} \cup (\mathbf{X}_T^{\bar{\alpha}^m}, \mathbf{Y}_T^{\bar{\alpha}^m})$
- 5: **for** step in $1, 2, \dots, L$ **do**
- 6: Sample $t \sim \mathcal{U}[0, T]$, $(\mathbf{X}_T, \mathbf{Y}_T) \sim \mathcal{B}$ and $\mathbf{X}_t \sim \mathbb{P}_{t|T}(\cdot | \mathbf{X}_T)$
- 7: Compute the training objective:

$$\mathcal{L}_{\text{FAS}}(\theta) = \mathbb{E}_{t, (\mathbf{x}_T, \mathbf{y}_T), \mathbf{x}_t} \left[\left\| \alpha^{\theta_s}(\mathbf{X}_t, t) + \sigma_t Q^{1/2} e^{-(T-t)\mathbf{A}^\dagger} \mathbf{Y}_T \right\|_{\mathcal{H}}^2 \right]$$

- 8: Update θ_m^{l+1} with $\nabla_{\theta} \mathcal{L}_{\text{FAS}}(\theta_m^l)$.
- 9: **end for**
- 10: **end for**
- 11: **return** Approximated optimal control $\alpha^* \approx \alpha^{\theta_M}$

Algorithm 5 IDPP initialization

Require: Transition states $\mathbf{A}$ and $\mathbf{B}$, discretization $\mathbf{U} = \{u_i\}_{i=1}^K$ with $0 = u_1 < \dots < u_K = 1$.

- 1: Set initial condition $\mathbf{x}_0 = (1 - t)\mathbf{A} + t\mathbf{B}$
- 2: **for** implicit time u in range $[u_1, u_K]$ **do**
- 3: Compute target distance $\tilde{d}_{ij}[u]$ in (195) for all $(i, j) \in N \times N$.
- 4: Take L gradient steps $\nabla_{\mathbf{x}_0[u]} \mathbf{E}_{\text{IDPP}}(\mathbf{x}_0[u])$ w.r.t on IDPP energy (196).
- 5: **end for**
- 6: **return** Refined initial condition $\tilde{\mathbf{x}}_0$

with atom positions $\mathbf{r}_i \in \mathbb{R}^3$. Given endpoints $\mathbf{A}$ and $\mathbf{B}$, and a path with images $\mathbf{x}[u] = \mathbf{0}_{N \times 3}$ and the target distance for pair (i, j) at time u as linear interpolation between pair-wise distance for a given endpoints:

$$\tilde{d}_{ij}[u] = (1 - u)d_{ij}(\mathbf{A}) + ud_{ij}(\mathbf{B}), \quad \text{where} \quad d_{ij}(\mathbf{C}) = |\mathbf{r}_i - \mathbf{r}_j|. \quad (195)$$

Then, the IDPP minimizes the mismatch between current and target distances via

$$\mathbf{E}_{\text{IDPP}}(\mathbf{x}[u]) = \sum_{i < j} w_{ij}^{(u)} (d_{ij}(\mathbf{x}[u]) - \tilde{d}_{ij}[u])^2, \quad (196)$$

where $w_{ij}^{(u)} = 1/(\tilde{d}_{ij}[u])^4$ is the weight over all images $u \in \mathbf{U}$. This initialization yields a path near a minimum-energy path with negligible time cost e.g., less than a minute. Algorithm 5 summarizes the IDPP initialization. For chignolin, we take additional gradient flow step on energy (28) with small step size $\epsilon = 10^{-6}$ to make the initial path more stable after IDPP initialization:

$$\mathbf{x}_{t+\epsilon} = \mathbf{x}_t - \nabla U(\mathbf{X}_t)\epsilon. \quad (197)$$

Regularization The IDPP improves the starting path by producing realistic images for both alanine dipeptide and chignolin. However, for alanine dipeptide, diffusion in Cartesian coordinates may move the path into energy troughs that stretch or compress bonds beyond physical ranges. We therefore add a regularization term to the terminal cost to limit bond-length deviations:

$$\tilde{U}(\mathbf{X}) = U(\mathbf{x}) + \lambda_{\text{reg}} \sum_{u \in \mathbf{L}} \mathbf{E}_{\text{IDPP}}(\mathbf{x}[u]), \quad (198)$$

where $\lambda_{\text{reg}} > 0$ is constant. This regularization cost penalizes deviations of pairwise distances from the IDPP reference at each image and suppresses short-range clashes and over-stretched bonds. We set $\lambda_{\text{reg}} = 1$ for alanine dipeptide.




Figure 4. Sampled transition path on alanine dipeptide over various discretization steps.




Figure 6. Ablation studies on alanine dipeptide. (Left) Ablation study on $\beta_{\max}$. (Right) Ablation study on model Scaling.

C.8. Additional Experiments

Discretization invariance sampling Beyond effective boundary enforcement, the functional formulation offers a key benefit: discretization invariance. Once the model is trained on a fixed and usually coarse grid, it can sample paths at arbitrary resolution. This capability has increased interest in functional formulations across scientific domains.

In this section, we evaluate the discretization invariance of FAS on TPS. For TPS, this property is valuable when more fine configuration image near saddle points is required. Classical Langevin dynamics based samplers can refine a path by reducing the step size. However, such a refinement requires additional force and energy evaluations that grow with the number of steps, which increases computational cost and inference time. In contrast, FAS is trained once on a small discretization. At inference it generates target paths on any chosen refined grid without extra energy evaluations because the energy cost is amortized in the learned control.

Figure 5 shows zero shot TPS on a synthetic potential at extremely fine resolutions. Specifically, We train on 100 number of uniform discretization over implicit time interval $[0, 1]$, then evaluate at $\times 10$ and $\times 100$ finer grids with 10^3 and 10^4 discretization steps. Despite zero shot inference, FAS generates reliable paths consistently. Note that sampling on high resolution requires a larger diffusion scale than originally chosen $(\beta_{\min}, \beta_{\max})$ for desired performance, so we scale $(\beta_{\min}, \beta_{\max})$ with the refinement factor. Concretely, if the grid is refined by factor r , we use $\hat{\beta}_{\min} = (1 + 2\log_{10}(r))\beta_{\min}$ and $\hat{\beta}_{\max} = (1 + 2\log_{10}(r))\beta_{\max}$ for high-resolution sampling.




Figure 5. Sampled transition path on synthetic potential over various discretization steps.

Figure 4 shows the same behavior on alanine dipeptide. Due to the complex potential landscape of molecule system, we observed that a few evaluation points can show short spikes. We remove these artifacts by applying 10 number of post processing gradient flow steps as in (197) with $\epsilon = 10^{-6}$.

Ablation studies For FAS training, the matching objective in (22) depend on sufficient exploration of the path-energy landscape g . Moreover, broad exploration in early stage is key to locating the critical point of $\mathcal{L}_{\text{FAS}}$ because higher early

diffusion schedule σ_t enlarges the support of sampled paths in g effectively. We therefore ablate the effect of diffusion magnitude σ_t helps the performance by increasing $\beta_{\max} \in [10, 15]$ with fixed $\beta_{\min} = 0.01$. As shown in Figure 6, we observed that resulting **ETS** consistently decreases as $\beta_{\max}$ increases.

Additionally, we ablate how FAS scales with model size. Compared to the original path-wise SOC objective in (12), the point-wise computation nature of matching-type objective in (22) can scales better without heavy computation. We vary network size by setting the base dimension to 32 (UNO-S), 48 (UNO-M), and 64 (UNO-L). Note that although the number of parameters increases, memory usage remains nearly unchanged. Figure 6 shows that larger models yield better **ETS**, which demonstrates the FAS scales with model size.