TRANSITION PATH SAMPLING USING KOOPMAN OPERATORS AND EXIT-TIME OPTIMAL CONTROL

Anonymous authors
Paper under double-blind review

ABSTRACT

Sampling transitions between metastable states is a central problem in dynamical systems theory and molecular dynamics in particular. A key challenge is the existence of high free-energy barriers that separate the states, making transitions extremely rare. Recent machine learning-based methods cast transition path sampling (TPS) as an optimal stochastic control (OSC) problem over a fixed time horizon, and parameterize the drift bias via a neural network trained by simulation-in-the-loop, requiring repeated biased rollouts. To address computational and performance guarantee issues of these models, we propose a new approach for the problem based on Koopman operators. Because Koopman operators are linear, their leading eigenfunctions reveal the metastable sets and provide an estimate of the committor function with no transition path information required. Furthermore, we formulate TPS as an OSC problem up to an *exit time*. Our time horizon is the first hitting time of the target set, and our running cost penalizes time spent in nonreactive regions by encoding the estimated committor function. We derive the optimal controller in closed form and approximate it in a reproducing kernel Hilbert space (RKHS). This reduces the problem of constructing the optimal controller to solving a single equality-constrained quadratic program, whose solution can be characterized by a linear Karush-Kuhn-Tucker (KKT) system. On the two-channel double well and alanine dipeptide, our controller increases the fraction of trajectories reaching the target from 0% to 99.8% within $1{,}000$ steps, and from 0% to 93% within $1\mathrm{ps}$, respectively.

1 INTRODUCTION

Sampling transitions between metastable sets is a fundamental problem in dynamical systems theory and molecular dynamics, underlying the study of conformational changes of proteins, folding, and drug discovery (Grubmüller, 1995; Voter, 1997). A typical challenge is that the states of interest are separated by high free-energy barriers: the overwhelming majority of simulated trajectories remain within a single metastable set, while trajectories that cross from one metastable set to another without returning (reactive paths) are rare under native dynamics (Vanden-Eijnden et al., 2006).

Classical transition path sampling (TPS) performs Markov chain Monte Carlo in the path space, generating new trajectories from existing ones (Dellago et al., 1998; Bolhuis et al., 2002; Bolhuis & Swenson, 2021). However, they require an initial reactive path connecting the two metastable sets, which can be difficult to generate. In physical chemistry and molecular dynamics, rare event sampling is typically accelerated by adding bias potentials to the energy landscape (Voter, 1997; Grubmüller, 1995; Laio & Parrinello, 2002). These are usually defined over a small number of hand-crafted collective variables. Modern machine learning-based samplers such as PIPS (Holdijk et al., 2023) and TPS-DPS (Seong et al., 2025) cast transition path sampling as optimal stochastic control or path-measure matching problems over a fixed time horizon, and parameterize the drift bias with a neural network trained by simulation-in-the-loop, which calls for repeated biased rollouts. Learning a bias force through simulation, however, can be challenging. For example, early in the training, few rollouts reach the target, and hence the training signal is highly sparse. Furthermore, a fixed-horizon formulation requires a properly chosen horizon, and the transition time is a random variable dependent on the path itself. Partial mitigation of these drawbacks is possible using replay buffers and annealing, and truncating the paths at a time that maximizes specific kernel values.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To resolve the above issues, we propose an *operator-based* framework for efficient TPS that uses *Koopman spectral analysis* and optimal stochastic control (OSC) up to an *exit time*, in which the stopping rule is the first hitting time of the target set. Specifically, we cast TPS as sampling from a Gibbs measure on the path space problem, following the information-theoretic frameworks of Mitter & Newton (2003) and Raginsky (2026). By the Gibbs variational principle, the Gibbs measure associated with a reference path measure and an energy on a path is, under certain conditions, the unique minimizer of the free energy. Mitter & Newton (2003) used this insight to pose nonlinear filtering as an optimal control problem. To adapt the described sampling method to TPS, we take the reference to be the unbiased path measure of the original dynamics, and the energy to penalize the time spent in nonreactive regions via the committor function. Since we assume the time horizon to be the first time a path reaches the target set, we operate with exit rather than fixed times. We prove that the resulting Gibbs measure is the law of a controlled diffusion process and prove that the optimal controller minimizes the free energy over *all* path measures on the stopped σ-algebra rather than only those generated by admissible controls. Furthermore, the optimal controller admits a closed-form solution and can be computed without simulations in the loop.

It is worth pointing out that the proposed approach differs from the recent work of Du et al. (2026), which investigates a related hitting-time problem for committor estimation, treating the committor as a value function and training it with neural networks that require repeated simulation of controlled dynamics during training. As previously emphasized, we instead follow an *operator-theoretic* approach in which the spectrum of the Koopman generator of the model reveals the metastable sets, and where the committor function is constructed from Koopman eigenfunctions. These can be estimated from a fixed dataset without requiring any trajectory to transition between the metastable sets. The Koopman generator plays a central role within this framework: its leading eigenfunctions reveal the metastable sets and are incorporated into the committor function, which in turn defines the running cost of the OSC. The optimal controller of the OSC can be solved via a boundary value problem wrt the operator. This pipeline comes with several practical advantages. First, it requires no transition paths, as Koopman eigenfunctions can be computed without crossings between metastable sets via a generalized eigenproblem. Second, the controller admits a closed-form solution, and thus is obtained without network training but through solutions of convex quadratic programs. The approach is also interpretable, as the controller is built from physically meaningful quantities that characterize the mechanism of the transition, namely the metastable sets and the committor function. Furthermore, the sampled paths can be reweighted to recover the statistics of the original dynamics.

Our contributions are multifold: (1) We formulate TPS as an OSC up to an exit time. We show in Theorem 1 that the optimal controller admits a closed-form solution, $u^* = \nabla \log \rho$, and that its path measure minimizes the free energy over *all* path measures, not only those generated by admissible controls; (2) We construct this controller by approximating ρ in a Reproducing Kernel Hilbert Space (RKHS). The problem then reduces to solving a convex quadratic program with linear equality constraints, whose solution can be characterized by a single linear Karush–Kuhn–Tucker (KKT) system rather than repeated rollouts; (3) In TPS, the source and target metastable sets are often predefined through collective variables. Following Koopman-based approaches that identify metastable sets (Mardt et al., 2018; Pérez-Hernández et al., 2013), we construct these sets without collective variables by using the leading eigenfunctions of the Koopman generator, which capture the slow inter-basin modes. From the same eigenfunctions, we construct a committor function, which defines the running cost of the OSC; (4) We generate transition paths by simulating the controlled diffusion process until it first enters the target sets, and show the resulting paths can be reweighted via Lemma 4 to recover the statistics of the original dynamics up to the transition; (5) We evaluate our method on both synthetic and real-world datasets.

The paper is organized as follows. In Section 2, we introduce the diffusion process of interest and the corresponding Koopman generator. In Section 3, we develop the OSC up to an exit time in general form. There, we also derive a closed form of the optimal controller. The results of this section can be specialized for efficient TPS. Section 4 constructs metastable sets and the committor function from Koopman eigenfunctions, and approximates the optimal controller in an RKHS. We generate transition paths by simulating the optimally controlled SDE, and report numerical experiments in Section 5. Due to space limitations, proofs and extra experiments are deferred to the Appendix G.4.

2 PRELIMINARIES: ITÔ DIFFUSION AND THE KOOPMAN GENERATOR

Let $\mathbb{X} \subseteq \mathbb{R}^n$ denote the state space and consider the Itô diffusion process

$$dX_t = b(X_t)\,dt + \sigma(X_t)\,dW_t, \qquad X_0 = x, \tag{1}$$

where W_t is an m-dimensional Brownian motion process, and $b : \mathbb{X} \to \mathbb{R}^n$ and $\sigma : \mathbb{X} \to \mathbb{R}^{n\times m}$ are the drift and diffusion functions, respectively Karatzas & Shreve (2014); Oksendal (2000). A model of the form (1) used in MD simulations is the *Langevin dynamics,* where b and σ are specified by a potential energy U, friction γ, and bath temperature T_{emp}. In the overdamped case, the state is the atomic configuration, and the drift and diffusion coefficients are specialized as $b(x) = -\frac{1}{\gamma m}\nabla U(x)$, $\sigma(x) = \sigma_0 I$, with $\sigma_0 = \sqrt{\frac{2k_B T_{\text{emp}}}{\gamma m}}$. In the underdamped case, the state additionally carries the momenta, and friction and noise act only on the momentum component. In either case, the process is metastable whenever U has multiple deep minima (wells) separated by barriers large relative to $k_B T_{\text{emp}}$, and a trajectory spends most of its time in a single well.

Next, let $\mathbb{T} := [0, \infty)$ and fix the space $\Omega = C(\mathbb{T}; \mathbb{X})$ of continuous paths mapping $\mathbb{T}$ to $\mathbb{X}$. Also, let $X_t(\omega) = \omega(t)$ denote the evaluation of a path $\omega \in \Omega$ at time $t \in \mathbb{T}$. Let $\mathcal{F} = \sigma(X_t : t \in \mathbb{T})$ be the canonical σ-algebra, and let $(\mathcal{F}_t)_{t\in\mathbb{T}}$ with $\mathcal{F}_t = \sigma(X_s : s \le t)$ be the filtration, both generated by the process $(X_t)_{t\in\mathbb{T}}$. For $x \in \mathbb{X}$, let P^x be the law on $(\Omega, \mathcal{F})$ of the solution of (1) started at x, provided it is well posed, for example, under Assumptions 1 and 2 in Section 3. Under P^x, the canonical process $X_t(\omega) = \omega(t)$ solves (1). Define the family of Koopman operators $\{\mathcal{K}_t\}_{t\in\mathbb{T}}$ as

$$\mathcal{K}_t\varphi(x) = \mathbb{E}[\varphi(X_t)|X_0 = x], \tag{2}$$

per (Koopman & von Neumann, 1932; Lasota & Mackey, 2013). Although (1) is nonlinear, $\mathcal{K}_t$ for $t \in \mathbb{T}$ are *linear* operators and $\{\mathcal{K}_t\}_{t\in\mathbb{T}}$ defines a semigroup $\mathcal{K}_{t+s} = \mathcal{K}_t \circ \mathcal{K}_s$. The associated infinitesimal Koopman generator L is a *linear* second-order differential operator given by

$$L\varphi(x) = b(x) \cdot \nabla\varphi(x) + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}(x)\,\partial_{ij}\varphi(x), \tag{3}$$

for $\varphi \in C^2(\mathbb{X})$, where the indices $i, j = 1, \cdots, n$ range over the coordinates of $x = (x_1, \ldots, x_n) \in \mathbb{X}$, ∂_{ij} denotes the second partial derivative with respect to x_i and x_j, a_{ij} is the (i, j) entry of the diffusion matrix $a(x) := \sigma(x)\sigma(x)^\top$, and b is the drift term in (1). The spectrum of L can be utilized to characterize the global structure of the dynamics (Mezić, 2005; Mardt et al., 2018; Kostic et al., 2022). Let (λ, ψ) be an eigenpair of L, i.e., $L\psi = \lambda\psi$. Then $\mathcal{K}_t\psi = e^{\lambda t}\psi$, i.e., $\mathbb{E}_x[\psi(X_t)] = e^{\lambda t}\psi(x)$. The constant function $\psi_0 \equiv 1$ is always an eigenfunction with $\lambda_0 = 0$.

The Koopman operator and its spectrum can be approximately computed from data via numerical methods such as dynamic mode decomposition (DMD) (Schmid, 2010; Rowley et al., 2009) and extended dynamic mode decomposition (EDMD) (Williams et al., 2015), as well as nonparametric variants based on RKHS known as kernel EDMD (Williams et al., 2014; Kostic et al., 2022; Klus et al., 2020; Hou et al., 2023). We provide a brief introduction to RKHS in Appendix A. Following Kostic et al. (2022); Hou et al. (2023), we approximate the leading eigenpairs (λ_k, ψ_k) of L from samples $\{x_i\}_{i=1}^N \subset \mathbb{X}$. Since b and a are known, L can be applied directly to the kernel via (3). This yields a generalized eigenproblem in terms of Gram matrices; see, e.g., (Hou et al., 2023, Appendix M) for details.

A trajectory of (1) that crosses from one metastable set to another without returning to its origin is a (reactive) transition path (Vanden-Eijnden et al., 2006). For two disjoint sets A_i and A_j, let τ_{A_k} be the first time the system enters A_k for $k = i, j$. The committor function $q^{(i,j)}(x)$ for the transition from $x \in A_i$ to A_j is defined as $q(x) = P\big(\tau_{A_j} < \tau_{A_i} \mid x_0 = x\big)$ (Vanden-Eijnden et al., 2006; Metzner et al., 2009; Hoffmann et al., 2021). Along a transient path, $q^{(i,j)}(x_t)$ is close to 0 near A_i and close to 1 near A_j, whereas a path that lingers in A_i has $q^{(i,j)}(x_t) \approx 0$ most of the time. When the metastable sets are clear from context, we write q as a shorthand for $q^{(i,j)}$. Next, let $D_q := \mathbb{X} \setminus (A_i \cup A_j)$. The committor function q satisfies the boundary value problem

$$Lq = 0 \ \text{ in } D_q, \qquad q|_{\partial A_i} = 0, \qquad q|_{\partial A_j} = 1, \tag{4}$$

where L is the Koopman generator, and we extended q by $q \equiv 0$ on A_i and $q \equiv 1$ on A_j. We defer the proof of this identity to the Appendix C.

The above concepts lead to our pipeline for TPS: the leading eigenfunctions of the Koopman generator L identify the metastable sets and provide the spectral term of the committor function; the committor function then defines the running cost of the OSC presented in the next section, whose optimal controller can be approximated via a boundary value problem in L.

3 Transition Path Sampling via OSC up to an Exit Time

3.1 The Sampling Problem

Consider the diffusion process governed by the Itô SDE (1). To facilitate our analysis of the outlined sampling approach, we make the assumption:

Assumption 1. *$b(x)$ and $\sigma(x)$ are bounded and Lipschitz on $\mathbb{X}$.*

Under Assumption 1, the SDE (1) has a unique strong solution. Let P^x be the path measure started at x. Let $D \subset \mathbb{X}$ be open, and $\tau_D := \inf\{t \geq 0 : X_t \notin D\}$ be the exit time. Since D is open, τ_D is the hitting time of the closed set $\mathbb{X} \setminus D$ and is therefore an $(\mathcal{F}_t)$-stopping time. We observe paths until they reach the target set (equivalently, exit the set D), and we work on the stopped σ-algebra $\mathcal{F}_{\tau_D}$ defined by $\mathcal{F}_{\tau_D} := \big\{A \in \mathcal{F} :\ A \cap \{\tau_D \leq t\} \in \mathcal{F}_t \text{ for all } t \geq 0\big\}$. Then X_{τ_D} is $\mathcal{F}_{\tau_D}$-measurable.

For $x \in \bar{D}$ and each path χ under P^x, define the Hamiltonian as

$$H(\chi) = \int_0^{\tau_D} f(X_t)\, dt + \Phi(X_{\tau_D}), \tag{5}$$

where $f : \bar{D} \to [0, \infty)$ and $\Phi : \partial D \to [0, \infty)$ are measurable, and $H(\chi) := +\infty$ on $\{\tau_D = \infty\}$. Here, $f \geq 0$ is the running cost and Φ is the terminal cost. Since paths are continuous and D is open, $X_{\tau_D} \in \partial D$ on $\{\tau_D < \infty\}$. Hence, H is a well-defined $\mathcal{F}_{\tau_D}$-measurable function.

For the purpose of efficient sampling, among all path measures on $(\Omega, \mathcal{F}_{\tau_D})$, we seek the one that minimizes the free energy

$$F(\tilde{P}) := \tilde{\mathbb{E}}\big[H(\chi)\big] + D_{\mathrm{KL}}\big(\tilde{P} \,\|\, P^x\big), \qquad \tilde{P} \in \mathcal{P}(\Omega, \mathcal{F}_{\tau_D}), \tag{6}$$

where $\tilde{\mathbb{E}}$ denotes the expectation w.r.t. $\tilde{P}$, and $\mathcal{P}(\Omega, \mathcal{F}_{\tau_D})$ denotes the set of all probability measures on $(\Omega, \mathcal{F}_{\tau_D})$. Here, the first term favors paths with low energy, while the second keeps $\tilde{P}$ close to the original dynamics. We next make use of the Gibbs variational principle on $\mathcal{F}_{\tau_D}$, which is adapted from (Mitter & Newton, 2003, Proposition 2.1).

Proposition 1. *Let H be an $\mathcal{F}_{\tau_D}$-measurable function, and suppose $Z := \mathbb{E}_x\big[e^{-H}\big] > 0$. Then the probability measure $P^{*,x}$ on $(\Omega, \mathcal{F}_{\tau_D})$ defined by*

$$\left.\frac{\mathrm{d}P^{*,x}}{\mathrm{d}P^x}\right|_{\mathcal{F}_{\tau_D}} = \frac{e^{-H(\chi)}}{Z} \tag{7}$$

is the unique minimizer of the free energy (6) *over the set of all probability measures on $(\Omega, \mathcal{F}_{\tau_D})$, and $\min_{\tilde{P}} F(\tilde{P}) = F(P^{*,x}) = -\log Z$.*

It will be useful to keep in mind the following motivating example.

Example 1 (TPS)**.** Let $A_i, A_j \subset \mathbb{X}$ be disjoint closed source and target sets, and let $D := \mathbb{X} \setminus A_j$. Then $A_i \subset D$, and thus a trajectory initialized within the source set starts in the interior of D. For such a trajectory, the exit time of D is the first hitting time of A_j, so τ_D is the transition time of interest. Let $q(x) := P^x(\tau_{A_j} < \tau_{A_i})$ be the committor function. Define the running and terminal cost

$$f_\beta(x) := \beta\big(1 - q^{(i,j)}(x)\big), \quad \Phi \equiv 0 \text{ on } \partial A_j, \tag{8}$$

with $\beta > 0$. Then, $H = \beta H_0$, where $H_0 = \int_0^{\tau_D} (1 - q^{(i,j)}(X_t))\, dt$ is the time spent in nonreactive regions before reaching the target. Here, the parameter β controls the trade-off between speed and fidelity. As β increases, we increasingly prefer faster paths at the cost of deviating more from the original dynamics. In Appendix F we show how to undo this dynamics distortion via reweighting.

The sampling problem of interest is to generate paths from $P^{*,x}$. Proposition 1 identifies the target measure $P^{*,x}$ but does not provide a way to sample from it. We therefore phrase the sampling problem as a control problem where we add a control drift to the reference process (1) so that the law of the controlled diffusion equals $P^{*,x}$ on $\mathcal{F}_{\tau_D}$. We call a control $u : \bar{D} \to \mathbb{R}^n$ *admissible* if u is Lipschitz continuous on $\bar{D}$ and use $\mathcal{U}$ to denote the set of all admissible controls. A $u \in \mathcal{U}$ enters the drift through $a(x)u(x)$, resulting in the controlled process

$$d\tilde{X}_t = \big[b(\tilde{X}_t) + a(\tilde{X}_t)\, u(\tilde{X}_t)\big]\, dt + \sigma(\tilde{X}_t)\, d\tilde{W}_t, \tag{9}$$

where $\tilde{\tau}_D := \inf\{t \geq 0 : \tilde{X}_t \notin D\}$. As we show in Theorem 1 (ii), restricting the focus to a Lipschitz continuous controller does not weaken the result, since we can show that the path measure of the optimal admissible controller is optimal over *all* probability measures on $(\Omega, \mathcal{F}_{\tau_D})$.

As the time-horizon is the first time the process exits D, the exit time τ_D is itself a random variable. We first establish that it has a finite mean. We make the following Assumptions

Assumption 2. *There exists $c > 0$ such that $z^\top a(x)\, z \geq c|z|^2$ for all $x \in \mathbb{X}, z \in \mathbb{R}^n$.*

Assumption 3. *$D \subset \mathbb{X}$ is open and bounded with $C^{2,\alpha}$ boundary ∂D for some $\alpha \in (0, 1)$.*

These assumptions are standard for exit-time problems for controlled diffusions (Fleming, 1977). It ensures a unique solution of the associated elliptic boundary value problem in L. They hold, for example, for overdamped Langevin dynamics with a smooth potential, including the double-well, four-well, and Müller–Brown systems studied in Section 5 and Appendix G.4. The proof of the following lemma is presented in Appendix B.1:

Lemma 1. *Let Assumptions 1–3 hold. Then, the SDE* (1) *starting from $x \in D$ has a pathwise unique strong solution, and its exit time τ_D satisfies $\mathbb{E}_x[\tau_D] < \infty$. Consequently, $\tau_D < \infty$ a.s.*

3.2 A Candidate Optimal Controller

We now seek a controlled diffusion of the form (9) whose path measure equals $P^{*,x}$. For this purpose, we also require a regularity condition,

Assumption 4. *The running cost $f \in C^{0,\alpha}(\bar{D})$ with $f \geq 0$, and $\Phi : \partial D \to [0, \infty)$ such that $g := e^{-\Phi} \in C^{2,\alpha}(\partial D)$.*

Define

$$\rho(x) := \mathbb{E}_x \left[\exp\Big(- \int_0^{\tau_D} f(X_t)\, dt - \Phi(X_{\tau_D}) \Big) \right]. \tag{10}$$

Here, we integrate time from 0 to an exit time τ_D. Under Assumption 3, f is bounded on $\bar{D}$, and Φ is bounded on ∂D since g is continuous and positive on the compact set ∂D. By Lemma 1, $\tau_D < \infty$ a.s. under P^x, hence (10) is well-defined, and $\rho(x) \in (0, 1]$.

We next show that ρ solves a boundary value problem via elliptic existence theory (Ladyzhenskaya & Uraltseva, 1968, Ch. III). We consider the following elliptic boundary value problem, where L is the (Koopman) generator:

$$Lw - fw = 0 \text{ in } D, \qquad w|_{\partial D} = e^{-\Phi}. \tag{11}$$

Since $\partial D \in C^{2,\alpha}$ by Assumption 2 and $g = e^{-\Phi} \in C^{2,\alpha}(\partial D)$ by Assumption 3, applying (Ladyzhenskaya & Uraltseva, 1968, Ch. III, Theorem 1.1, 1.2) gives a unique solution $w \in C^{2,\alpha}(\bar{D})$ to the problem (11). The proof of the following lemma is included in the Appendix B.3.

Lemma 2. *Let Assumptions 1–4 hold and let $w \in C^{2,\alpha}(\bar{D})$ be the unique solution of the problem* (11)*. Then, w coincides on $\bar{D}$ with ρ defined in* (10)*. Moreover, $\rho > 0$ on $\bar{D}$.*

Our sampling/free energy minimization problem can now be cast as an optimal control problem. In the fixed-horizon setting, the relative entropy of the controlled law with respect to the reference law equals the expected control effort by Girsanov's theorem, so the free energy (6) of the controlled law is an optimal control cost (Raginsky, 2026, eq. (4.5)). By contrast, here we have to consider the optimal stochastic control problem up to an exit time,

$$\min_{u \in \mathcal{U}} J(u, x) := \tilde{\mathbb{E}}_x \left[\int_0^{\tilde{\tau}_D} \Big(f(\tilde{X}_t) + \tfrac{1}{2} u(\tilde{X}_t)^\top a(\tilde{X}_t)\, u(\tilde{X}_t) \Big)\, dt + \Phi(\tilde{X}_{\tilde{\tau}_D}) \right], \tag{12}$$

subject to (9), $\tilde{X}_0 = x \in D$.

We denote its value function as $v(x) := \inf_{u\in\mathcal{U}} J(u,x)$, for $x \in \bar{D}$, with $v|_{\partial D} = \Phi$, since $\tilde{\tau}_D = 0$ when $\tilde{X}_0 \in \partial D$. Consider the candidate value function $v := -\log\rho$ (this is the logarithmic, or Cole–Hopf, transform of ρ, (Fleming, 1977)). Since $\bar{D}$ is compact by Assumption 2 and $\rho \in C(\bar{D})$ with $\rho > 0$ on $\bar{D}$ as established in Lemma 2, the extreme value theorem gives $m := \min_{\bar{D}} \rho > 0$ and $M := \max_{\bar{D}} \rho < \infty$, so $\rho(\bar{D}) \subseteq [m, M] \subset (0, \infty)$. Since $\rho \in C^{2,\alpha}(\bar{D})$, and $-\log$ is smooth on $[m, M]$, the chain rule gives $v = -\log\circ\rho \in C^{2,\alpha}(\bar{D})$, and hence ∇v is bounded on $\bar{D}$. Additionally, since $\rho\big|_{\partial D} = g = e^{-\Phi}$, the boundary condition of v is $v\big|_{\partial D} = \Phi$. We have the following lemma whose proof is presented in Appendix B.4.

Lemma 3. *The function v satisfies the Hamilton-Jacobi-Bellman (HJB) equation of the OSC problem* (12) *up to an exit time, i.e.,*

$$Lv(x) + f(x) = \tfrac{1}{2}\,\nabla v(x)^\top a(x)\,\nabla v(x) \quad \textit{in } D, \qquad v\big|_{\partial D} = \Phi. \tag{13}$$

We hence have a candidate for the optimal controller of (12) given by

$$u^*(x) = -\nabla v(x) = \nabla \log\rho(x) = \nabla\rho(x)/\rho(x). \tag{14}$$

Theorem 1 below establishes that its minimal value equals $\min F = -\log Z$ and that its optimally controlled law is $P^{*,x}$.

3.3 Verification Theorem of OSC up to an Exit Time

Let $\tilde{P}^{*,x}$ be the law of the controlled process (9) under u^*, starting at x. In the remainder of the section, we establish that u^* in (14) is the optimal controller. Moreover, using a variational characterization of the optimally controlled path measure, we show that $\tilde{P}^{*,x}$ minimizes the free energy (6) over all path measures on $(\Omega, \mathcal{F}_{\tau_D})$. We begin by presenting an explicit formula for the Radon-Nikodym density of $\tilde{P}^{*,x}$ relative to P^x on $(\Omega, \mathcal{F}_{\tau_D})$. To that end, we invoke Girsanov's formula. In the classic fixed time-horizon case, one needs a suitable regularity condition, such as Novikov's condition (Oksendal, 2000, Theorem 8.6.3). By contrast, in the exit time setting, such a condition might not hold as it requires an exponential moment bound, which is hard to verify for metastable systems. To address this challenge, the proof presented in Appendix B.6 only requires $\tau_D < \infty$ a.s. guaranteed by Lemma 1.

Lemma 4. *Let $u^* = -\nabla v$ be as in* (14)*, and let $\tilde{X}_t$ be a solution to the controlled SDE* (9) *with $u = u^*$, initiated at $\tilde{X}_0 = x$, and under a probability measure $\tilde{P}^{*,x}$ on $(\Omega, \mathcal{F}_{\tau_D})$. Then, we have*

$$\left.\frac{d\tilde{P}^{*,x}}{dP^x}\right|_{\mathcal{F}_{\tau_D}} = \frac{e^{-H(\chi)}}{\rho(x)}, \qquad \mathbb{E}_x\left[\frac{e^{-H(\chi)}}{\rho(x)}\right] = 1. \tag{15}$$

We now present the verification theorem for which the proof can be found in Appendix B.7.

Theorem 1. *Let Assumptions 1–4 hold, and let u^* be the controller in* (14)*. Then,*

(i) *[Optimality] $v(x) = J(u^*, x) = \min_u J(u, x)$ for all $x \in D$.*

(ii) *[Variational inequality] For every probability measure $\tilde{P}$ on $(\Omega, \mathcal{F}_{\tau_D})$,*

$$J(u^*, x) \le D_{\mathrm{KL}}(\tilde{P}\|P^x) + \tilde{\mathbb{E}}_x[H(\chi)], \tag{16}$$

with equality if and only if $d\tilde{P}/dP^x = e^{-H(\chi)}/\rho(x)$ on $\mathcal{F}_{\tau_D}$; equivalently, if and only if $\tilde{P} = \tilde{P}^{*,x}$, corresponding to the law of the optimally-controlled process restricted to $\mathcal{F}_{\tau_D}$.

Part (i) of the above theorem states that the minimal control cost is $\min_{u\in\mathcal{U}} J(u,x) = v(x) = -\log\rho(x)$, and by Part (ii), that the free energy (6) satisfies $F(\tilde{P}) \ge v(x)$ for *every* probability measure on $(\Omega, \mathcal{F}_{\tau_D})$, with equality precisely for $\tilde{P} = \tilde{P}^{*,x}$. Hence, the two minimal values of (6) and (12) coincide, and the optimal path measure that minimizes the free energy is exactly the one generated by the optimal controller. In the proof of Part (i) of Theorem 1, we also show that the

minimizer is unique in the following sense: $J(u,x) = v(x)$ for some $u \in \mathcal{U}$ if and only if $u(\tilde{X}_t) = u^*(\tilde{X}_t)$ for Lebesgue-a.e. $t \in [0, \tilde{\tau}_D]$, $\tilde{P}^{u,x}$-a.s.

By the definition in (10), $\rho(x)$ is large where the trajectory of the original dynamics (1) tends to exit quickly and small where it dwells in nonreactive regions. The optimal controller in (14) therefore points toward states from which a trajectory is more likely to exit. The normalization by $\rho > 0$ in (14) provides a stronger push towards where ρ is small and the original dynamics is least likely to escape; thus the controller acts more aggressively when deep inside the source basin.

Using Lemma 4, we can recover statistics of the original dynamics (1) from the accelerated path alone. We present a reweighting mechanism via Girsanov's change of measure in Appendix F.

4 Koopman-Based Construction of the Optimal Controller

Theorem 1 reduces the problem of sampling from $P^{*,x}$ to computing the optimal controller $u^* = \nabla \log \rho$, where ρ solves the elliptic problem (11). For TPS in Example 1, applying this result requires three ingredients: the source and target sets A_i, A_j, the committor function q which defines the running cost f_β, and ρ itself. Each can be characterized through the Koopman generator L. Specifically, the metastable sets can be computed through the leading eigenfunctions of L, q through $Lq = 0$ in (4), and ρ through $L\rho - f_\beta \rho = 0$ in (11).

4.1 Constructing Metastable Cores and Committor Functions

The spectrum of the linear operator L reveals global geometric structures of the underlying dynamics (Mezić, 2005; Mauroy & Mezić, 2016). In molecular dynamics, for example, the metastable sets can be estimated using leading eigenfunctions whose eigenvalues are close to zero, followed by a spectral gap (Mardt et al., 2018; Hou et al., 2023; 2026). Specifically, the leading eigenvalues $0 = \lambda_0 > \lambda_1 \geq \cdots \geq \lambda_{m-1} \gg \lambda_m \cdots$ also capture the slow-mode, inter-basin dynamics. By truncating at m modes, fast modes are removed as they encode intra-basin dynamics. The leading eigenfunctions $\{\psi_k\}_{k=1}^{m-1}$ are approximately constant within each metastable set but take different values across them; we omit $k = 0$ since $\psi_0 \equiv 1$. Let $\Psi(x) = [\psi_1(x), \ldots, \psi_{m-1}(x)] \in \mathbb{R}^{m-1}$, where $\{x_n\}_{n=1}^N \subset \mathbb{X}$ are data samples and the eigenpairs are estimated from them via kernel EDMD in Section 2. We then apply k-means clustering to $\{\Psi(x_n)\}_{n=1}^N$, which partitions the samples into m clusters with index sets $\mathcal{I}_1, \ldots, \mathcal{I}_m$. For each cluster k and the set of leading eigenfunctions ψ_l, for $l = 1, \ldots, m-1$, we define the core level of the l-th eigenfunction as

$$\theta_l^{(k)} := \frac{1}{|\mathcal{I}_k|} \sum_{n \in \mathcal{I}_k} \psi_l(x_n), \qquad \theta^{(k)} := \left[\theta_1^{(k)}, \ldots, \theta_{m-1}^{(k)}\right]^\top, \tag{17}$$

so that $\theta^{(k)}$ is the k-means centroid of cluster k. Let $\varepsilon_k > 0$ be a radius. Since each ψ_l is approximately constant on a metastable set with value approximately $\theta_l^{(k)}$ on the k-th one, we define the core of the k-th metastable set for $k = 1, \cdots, m$ as

$$A_k := \{x \in \mathbb{X} : \|\Psi(x) - \theta^{(k)}\|^2 \leq \varepsilon_k^2\}. \tag{18}$$

Leveraging the global geometric structure carried by the leading eigenfunctions of the Koopman generator, we next approximate the committor function q via an ansatz $q = \tilde{\psi} + h$, where $\tilde{\psi}$ is a spectral term constructed from leading eigenfunctions and $h \in \mathcal{H}_\kappa$ is a remainder with $\mathcal{H}_\kappa$ is an RKHS equipped with kernel κ, since in general, the solution of (4) may not lie in $\text{span}\{\psi_l\}_{l=1}^{m-1}$. The spectral term $\tilde{\psi}$ contributes to the global shape of the transition, and also captures the slow coordinate along which the transition occurs, without requiring a reaction coordinate to be specified in advance. Due to space limitations, detailed construction of $\tilde{\psi}$ that maps the core of A_i to 0 and the core of A_j to 1 is presented in Appendix D. We compute $h \in \mathcal{H}_\kappa$ via $h(x) = \sum_{l \in \mathcal{I}_q} \alpha_l \, \kappa(x, x_l)$ with $\{x_l\}_{l \in \mathcal{I}_q} \subset D_q$ and $\mathcal{I}_q = \{1, \ldots, N_q\}$. The coefficients α are obtained from an equality-constrained convex quadratic program that minimizes the residual of (4) in D_q, which is also deferred to in Appendix D. Putting together, we compute the committor function via

$$\hat{q}(x) := \tilde{\psi}(x) + \sum_{l \in \mathcal{I}_q} \alpha_l \, \kappa(x, x_l), \tag{19}$$

4.2 Kernel-Based Approximation of the Optimal Controller

Unlike the fixed-horizon case studied in (Raginsky, 2026), where u^* depends on time, the exit-time optimal controller u^* in (14) is stationary. To compute u^*, we focus on approximating ρ defined in (10) so that it solves the elliptic boundary value problem (11). By definition of ρ, and $f_\beta = \beta(1 - q^{(i,j)})$, $\Phi \equiv 0$ in Example 1, we have

$$\rho(x) = \mathbb{E}_x \left[e^{-\beta H_0} \right], \qquad H_0 = \int_0^{\tau_D} \big(1 - q^{(i,j)}(X_t)\big)\, dt, \tag{20}$$

where H_0 is the time spent in nonreactive regions before reaching the target. We need to handle two regimes, capturing the transition region and regions within the basins. Directly approximating ρ in an RKHS can be challenging, as the behavior of ρ in the two regimes may be quite different. Near ∂A_j, the target is reached almost immediately and $\rho \approx 1$. Within the source basin, however, the transition is a rare event, and ρ is small as the trajectory lingers there. Since u^* depends on ρ through $\log \rho$, we seek an ansatz of the form $\rho := \rho_0(x)w(x)$ so that $\log \rho = \log \rho_0 + \log w$. Here we choose ρ_0 in closed form to capture how ρ varies across the two regimes, and use a residual w to solve for $\mathcal{H}_\kappa$. We defer additional details to Appendix E and only mention here that we approximate ρ via a closed-form factor $\hat{\rho}$ as

$$\hat{\rho}(x) := \rho_0(x)\, w(x), \quad \rho_0(x) := \hat{q}(x) + \big(1 - \hat{q}(x)\big)r, \quad w \in \mathcal{H}_\kappa, \tag{21}$$

where $\mathcal{H}_\kappa$ is an RKHS with kernel κ, $\hat{q}$ as in (19) and $r = \frac{1}{1+\beta T_0}$, with $T_0 > 0$ which encodes the time-scale of the original dynamics; for example, we set $T_0 = \frac{1}{|\lambda_1|}$, where λ_1 is the leading nontrivial Koopman eigenvalue. Let $\{x_i\}_{i=1}^N = \{x_i\}_{i \in I_D} \cup \{x_i\}_{i \in I_{\partial D}}$, with interior points $\{x_i\}_{i \in I_D} \subset D$ and boundary points $\{x_i\}_{i \in I_{\partial D}} \subset \partial D$. We approximate $w \in \mathcal{H}_\kappa$ via $w = \sum_{i=1}^N c_i\, \kappa(x, x_i)$ (Appendix E). Applying the chain rule to $\log \hat{\rho} = \log \rho_0 + \log w$, we have

$$\hat{u}(x) = \nabla \log \hat{\rho}(x) = \frac{(1 - r)\, \nabla \hat{q}(x)}{\hat{q}(x) + \big(1 - \hat{q}(x)\big)r} + \frac{\nabla w(x)}{w(x)}. \tag{22}$$

Thus, the problem of solving for u^* reduces to solving a convex quadratic program with linear equality constraints, whose solution is characterized by a single linear KKT system as in Appendix E. The method is summarized in Algorithm 1.

Algorithm 1 Summary of the Algorithm in Four Stages

Stage I. Construct leading eigenpairs $(\lambda_k, \psi_k)_{k=1}^{m-1}$ of the Koopman generator L and identify metastable cores.
Stage II. Compute the committor function $\hat{q}$ via (19) using the leading Koopman eigenfunctions.
Stage III. Construct the optimal controller $\hat{u}$ via (22).
Stage IV. Sample starting points from the source set A_i and integrate the controlled dynamics (9) with $\hat{u}$ from (22) until the paths hit the target set.

5 Numerical Experiments

We now present empirical evaluations of the proposed method on two synthetic systems, namely, the two-channel double-well system (Figure 1) and the Müller-Brown system (Figure 2), and a real-world alanine dipeptide (Figure 3). Due to space limitations, experimental details and one additional example are deferred to Appendix G.4. In all these experiments, a path is a hit at its first entrance to the target, and is truncated at the numerical censor $T_{\max}$ otherwise. We report the target-hit percentage (THP), which is the fraction of paths that hit the target set before $T_{\max}$. We also remark that alanine dipeptide does not satisfy Assumption 2; nevertheless, our algorithm can be applied beyond the confines of the theoretical analysis, and this experiment serves as a demonstration that it performs well even in such cases.

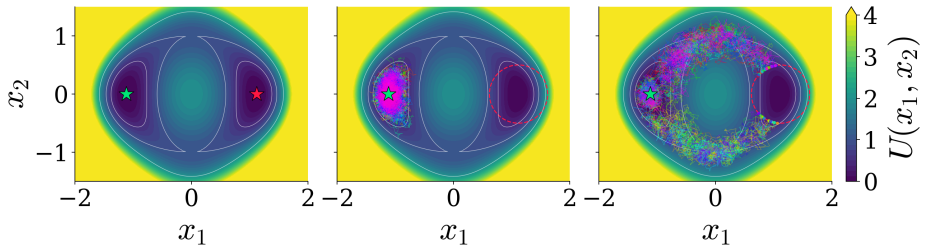
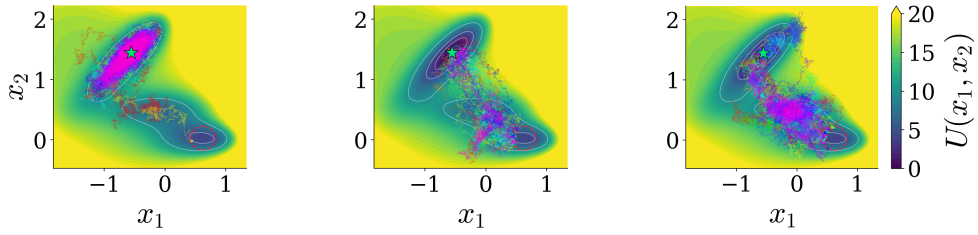
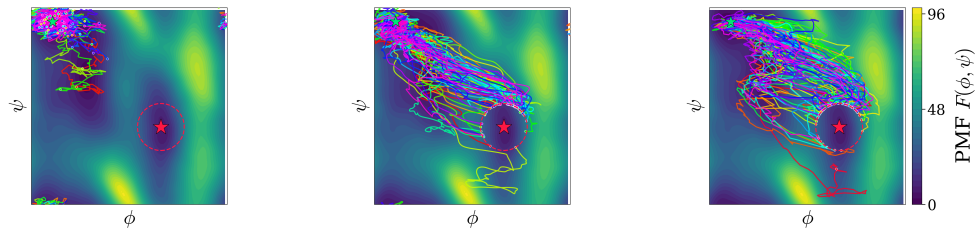


Figure 1: Two-channel double well. The potential has two minima at $(\pm\sqrt{5}/2,\ 0)$ connected by two channels around $(0, \pm 1)$. Following Algorithm 1, we compute the Koopman eigenfunctions and extract the metastable cores A_{left} and A_{right} as described in Section 4.1. Their minima are marked by the green and red stars, and the dashed red curve marks $\partial A_{\text{right}}$. We sample 1000 uncontrolled and 1000 controlled paths, all started at the minimum of A_{left}. (Left) Potential energy landscape. (Middle) Uncontrolled paths: THP($T_{\max}$) = 0%. (Right) Proposed method: THP($T_{\max}$) = 99.8%.



(a) Uncontrolled (b) TPS-DPS (c) Proposed method

Figure 2: Müller-Brown potential, with a deep well at $(-0.558, 1.442)$ (green star), a medium well at $(0.623, 0.028)$ (red circle), and a shallow well at $(-0.050, 0.467)$. We aim to sample transitions from the deep well to the medium well within $T_{\max} = 10$ps. (a) Uncontrolled paths: THP($T_{\max}$) = 1.8%. (b) TPS-DPS (Seong et al., 2025): THP($T_{\max}$) = 100%, but the paths run over the high energy barrier. (c) Proposed method: THP($T_{\max}$) = 99.0% using $\beta = 1$, with paths following the low-energy channel through the shallow well rather than crossing the barrier.



(a) Uncontrolled (b) TPS-DPS (c) Proposed method

Figure 3: Alanine dipeptide in vacuum, shown on the free-energy surface over the backbone dihedral angles (ϕ, ψ). We aim to sample the transition from $C5$ conformation (upper left) to the $C7_{\text{ax}}$ conformation (red star) within $T_{\max} = 1$ps. (a) Uncontrolled paths: $(T_{\max}) = 0\%$ (b) TPS-DPS: THP $(T_{\max}) = 100\%$, obtained by training a neural-network bias with simulation in the loop (1000 rollouts, $\sim 10^6$ gradient updates). (c) Proposed method: THP $(T_{\max}) = 93\%$, with a controller obtained in closed form. The setup and computational cost are reported in Appendix G.4.

Reproducibility statement

All detailed proofs of results and implementation details of experiments are also included in the appendix.

AI Disclosure Statement

The ideas were conceived without the use of AI, and all proofs were generated without consulting AI models. Writing was done by the authors. For a subset of coding tasks we initially used AI agents but verified and extended all codes independent from the models. We take full responsibility for the final content of this work, such as developing theoretical results, formulating mathematical claims, proving claims, and data analyses

References

Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.

Peter G Bolhuis and David W H Swenson. Transition path sampling as markov chain monte carlo of trajectories: Recent algorithms, software, applications, and future outlook. *Advanced Theory and Simulations*, 4(4):2000237, 2021.

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.

Christoph Dellago, Peter G Bolhuis, and David Chandler. Efficient transition path sampling: Application to lennard-jones cluster rearrangements. *The Journal of Chemical Physics*, 108(22): 9236–9245, 1998.

Yuanqi Du, Jiajun He, Dinghuai Zhang, Eric Vanden-Eijnden, and Carles Domingo-Enrich. Rare event analysis via stochastic optimal control. *arXiv preprint arXiv:2604.13213*, 2026.

Wendell H Fleming. Exit probabilities and optimal stochastic control. *Applied Mathematics and Optimization*, 4(1):329–346, 1977.

Helmut Grubmüller. Predicting slow structural transitions in macromolecular systems: Conformational flooding. *Phys. Rev. E*, 52:2893–2906, Sep 1995. doi: 10.1103/PhysRevE.52.2893. URL `https://link.aps.org/doi/10.1103/PhysRevE.52.2893`.

Moritz Hoffmann, Martin Konrad Scherer, Tim Hempel, Andreas Mardt, Brian de Silva, Brooke Elena Husic, Stefan Klus, Hao Wu, J Nathan Kutz, Steven Brunton, and Frank Noé. Deeptime: a python library for machine learning dynamical models from time series data. *Machine Learning: Science and Technology*, 2021.

Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochastic optimal control for collective variable free sampling of molecular transition paths. *Advances in Neural Information Processing Systems*, 36:79540–79556, 2023.

Boya Hou, Sina Sanjari, Nathan Dahlin, Subhonmesh Bose, and Umesh Vaidya. Sparse learning of dynamical systems in rkhs: An operator-theoretic approach. In *International Conference on Machine Learning*, pp. 13325–13352. PMLR, 2023.

Boya Hou, Sina Sanjari, Alec Koppel, Nathan Dahlin, and Subhonmesh Bose. Nonparametric sparse online learning of the koopman operator. *To appear at SIAM Journal on Control and Optimization*, 2026.

Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*. springer, 2014.

Stefan Klus, Feliks Nüske, and Boumediene Hamzi. Kernel-based approximation of the koopman generator and schrödinger operator. *Entropy*, 22(7):722, 2020.

Bernard O Koopman and J von Neumann. Dynamical systems of continuous spectra. *Proceedings of the National Academy of Sciences*, 18(3):255–263, 1932.

Vladimir Kostic, Pietro Novelli, Andreas Maurer, Carlo Ciliberto, Lorenzo Rosasco, and Massimiliano Pontil. Learning dynamical systems via koopman operator regression in reproducing kernel hilbert spaces. *Advances in Neural Information Processing Systems*, 35:4017–4031, 2022.

Olga A Ladyzhenskaya and Nina N Uraltseva. *Linear and Quasilinear Elliptic Equations*. Academic Press, New York, 1968.

Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99(20):12562–12566, 2002. doi: 10.1073/pnas.202427399. URL https://www.pnas.org/doi/abs/10.1073/pnas.202427399.

Andrzej Lasota and Michael C Mackey. *Chaos, fractals, and noise: stochastic aspects of dynamics*, volume 97. Springer Science & Business Media, 2013.

Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. Vampnets for deep learning of molecular kinetics. *Nature communications*, 9(1):5, 2018.

Alexandre Mauroy and Igor Mezić. Global stability analysis using the eigenfunctions of the koopman operator. *IEEE Transactions on Automatic Control*, 61(11):3356–3369, 2016.

Philipp Metzner, Christof Schütte, and Eric Vanden-Eijnden. Transition path theory for markov jump processes. *Multiscale Modeling & Simulation*, 7(3):1192–1219, 2009.

Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1):309–325, 2005.

Sanjoy K Mitter and Nigel J Newton. A variational approach to nonlinear estimation. *SIAM journal on control and optimization*, 42(5):1813–1833, 2003.

Klaus Müller and Leo D. Brown. Location of saddle points and minimum energy paths by a constrained simplex optimization procedure. *Theoretica chimica acta*, 53(1):75–93, Mar 1979. ISSN 1432-2234. doi: 10.1007/BF00547608. URL https://doi.org/10.1007/BF00547608.

Bernt Oksendal. *Stochastic differential equations: an introduction with applications (Fifth Edition)*. Springer Science & Business Media, 2000.

Guillermo Pérez-Hernández, Fabian Paul, Toni Giorgino, Gianni De Fabritiis, and Frank Noé. Identification of slow molecular order parameters for markov model construction. *The Journal of chemical physics*, 139(1), 2013.

Maxim Raginsky. A variational approach to sampling in diffusion processes. *Communications in Optimization Theory*, 2026(30):1–20, 2026. URL https://cot.mathres.org/archives/2289.

Clarence W Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S Henningson. Spectral analysis of nonlinear flows. *Journal of fluid mechanics*, 641:115–127, 2009.

Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.

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 *International Conference on Learning Representations*, volume 2025, pp. 93040–93062, 2025.

Eric Vanden-Eijnden et al. Towards a theory of transition paths. *Journal of statistical physics*, 123 (3):503–523, 2006.

Arthur F. Voter. Hyperdynamics: Accelerated molecular dynamics of infrequent events. *Phys. Rev. Lett.*, 78:3908–3911, May 1997. doi: 10.1103/PhysRevLett.78.3908. URL https://link.aps.org/doi/10.1103/PhysRevLett.78.3908.

Matthew O Williams, Clarence W Rowley, and Ioannis G Kevrekidis. A kernel-based approach to data-driven koopman spectral analysis. *Journal of Computational Dynamics*, 2014.

Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data–driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

APPENDIX

A THE REPRODUCING KERNEL HILBERT SPACE

We provide a brief introduction to the real-valued RKHS – for a more comprehensive treatment, see (Berlinet & Thomas-Agnan, 2011). Let $\mathbb{X}$ be a subset of a Euclidean space and let $\kappa : \mathbb{X}\times\mathbb{X} \to \mathbb{R}$ be a continuous, symmetric, positive semi-definite kernel. Define the RKHS associated with the kernel κ, $\mathcal{H}_\kappa$, as the completion of the span of $\{\phi(x) := \kappa(x, \cdot) : x \in \mathbb{X}\}$, equipped with the inner product $\langle\cdot,\cdot\rangle$, satisfying $\langle\phi(x), \phi(y)\rangle = \kappa(x, y)$. Here, ϕ is called the feature map for the kernel κ. The inner product satisfies the *reproducing property*, given by $\langle\phi(x), f\rangle = f(x)$, for all $x \in \mathbb{X}$ and $f \in \mathcal{H}_\kappa$. Throughout the exposition, we make the standard assumption that κ is measurable.

A.1 AUXILIARY RESULT: A LOCALIZED DYNKIN FORMULA

Several proofs in this paper follow a similar outline: we start with a real-valued function h that solves a partial differential equation on D, and we aim to evaluate the expectation of h along the trajectory at the exit time τ_D. However, Itô's formula is stated for functions that are C^2 on $\mathbb{R}^n$, whereas the functions we work with are C^2 only on the set D. Therefore, we need a mechanism to integrate up to τ_D at which time the trajectory reaches ∂D. In addition, such a τ_D is a random variable itself. We address this by localizing, following the scheme in the proof of (Oksendal, 2000, Theorem 9.1.1). Specifically, we first integrate to $\sigma_k := k \wedge \tau_{D_k}$, where D_k is an increasing sequence of open sets in D. Then on $[0, \sigma_k]$, the optional stopping theorem (Karatzas & Shreve, 2014, Theorem 1.3.22) applies. Letting $k \to \infty$ then takes the result to τ_D, given $\sigma_k \to \tau_D$. The following lemma (localized Dynkin formula) makes this construction precise.

Lemma 5. *Let Assumption 1–4 hold. Let Y be the pathwise unique strong solution of*

$$dY_t = b(Y_t)\,dt + \sigma(Y_t)\,dW_t, \qquad Y_0 = x \in D, \tag{23}$$

with generator $L_Y h := b^\top \nabla h + \frac{1}{2}\sum_{ij} a_{ij}\partial_{ij} h$, and let $\tau_D := \inf\{t > 0 : Y_t \notin D\}$. Define a sequence of open sets as

$$D_k := \{y \in D : \operatorname{dist}(y, \partial D) > 1/k\}. \tag{24}$$

Let $\sigma_k := k \wedge \tau_{D_k}$, where τ_{D_k} is the exit time of Y from D_k. Then, we have

(i) *The sequence $(\sigma_k)_{k\geq 1}$ is nondecreasing, and $\sigma_k \to \tau_D$ as $k \to \infty$.*

(ii) *Let $h \in C^2(D) \cap C(\bar{D})$, $c \in C(\bar{D})$ with $c \geq 0$, and $Z_t := \exp\big(-\int_0^t c(Y_s)\,ds\big)$. Then for every k, we have*

$$h(Y_{\sigma_k})Z_{\sigma_k} - h(x) = \int_0^{\sigma_k} Z_t\big(L_Y h - ch\big)(Y_t)\,dt + \int_0^{\sigma_k} Z_t\,\nabla h(Y_t)^\top \sigma(Y_t)\,dW_t, \tag{25}$$

and

$$\mathbb{E}\big[h(Y_{\sigma_k})Z_{\sigma_k}\big] - h(x) = \mathbb{E}\Big[\int_0^{\sigma_k} Z_t\big(L_Y h - ch\big)(Y_t)\,dt\Big]. \tag{26}$$

Proof. By construction of D_k in (24), $D_k \subset D_{k+1}$, and $\bigcup_k D_k = D$. Moreover, $\bar{D}_k$ is a compact subset of D.

Part (i) For every $k \geq 1$, since $D_k \subset D_{k+1}$, we have $\{t > 0 : Y_t \notin D_{k+1}\} \subseteq \{t > 0 : Y_t \notin D_k\}$. Hence, $\tau_{D_k} \leq \tau_{D_{k+1}}$. Since $\sigma_k = k \wedge \tau_{D_k}$ is the minimum of two nondecreasing sequences, it is itself also nondecreasing.

We next show that $\sup_k \sigma_k = \tau_D$ pathwise. Since $D_k \subset D$, we have $\sigma_k \leq \tau_D$ for every k, so $\sup_k \sigma_k \leq \tau_D$. For the reverse direction, fix ω and $t < \tau_D$. Then $\{Y_s : 0 \leq s \leq t\}$ is a compact subset of D and hence has strictly positive distance to ∂D, so it is contained in D_{k_0} for some k_0. For $k > \max(k_0, t)$ we have $D_{k_0} \subseteq D_k$, so $\tau_{D_k} > t$ and $\sigma_k = k \wedge \tau_{D_k} > t$. As $t < \tau_D$ was arbitrary, we have $\sup_k \sigma_k \geq \tau_D$.

Part (ii) Fix $k \geq 1$. We distinguish two cases, depending on whether the initial point x is in $\bar{D}_k$ or not. For the first case, consider $x \notin \bar{D}_k$. Since $\mathbb{X} \setminus \bar{D}_k$ is open and paths are continuous, there is $\epsilon > 0$ such that $Y_t \notin \bar{D}_k$, and in particular $Y_t \notin D_k$, for all $t \in (0, \epsilon)$. Hence $\tau_{D_k} = 0$ and $\sigma_k = 0$. Then $t \wedge \sigma_k = 0$ for every $t \geq 0$. We have $h(Y_0)Z_0 - h(x) = 0$ as $Z_0 = 1$ and all integrals from 0 to $\sigma_k = 0$ are zero. Hence, both sides of (25) and (26) vanish.

For the second case where $x \in \bar{D}_k$, we first show that $Y_t \in \bar{D}_k$ for all $t \in [0, \sigma_k]$. Indeed, $Y_0 = x \in \bar{D}_k$, and for $t \in (0, \sigma_k)$, we have $t < \tau_{D_k}$, hence $Y_t \in D_k$. If $\sigma_k > 0$, then $Y_{\sigma_k} = \lim_{t \to \sigma_k} Y_t \in \bar{D}_k$ by path continuity.

Since $\bar{D}_k$ is a compact subset of the open set D, there exists some $\zeta_k \in C_c^\infty(D)$ with $\zeta_k \equiv 1$ on an open neighbourhood U_k of $\bar{D}_k$, so that in particular, $U_k \subseteq \operatorname{supp} \zeta_k \subset D$. Let h_k equal $\zeta_k h$ on D and $h_k := 0$ on $\mathbb{R}^n \setminus D$. The open sets D and $\mathbb{R}^n \setminus \operatorname{supp} \zeta_k$ cover $\mathbb{R}^n$, h_k is C^2 on the first and identically zero on the second, and the two definitions agree on the overlap. Thus $h_k \in C^2(\mathbb{R}^n)$ with compact support. In addition, we have $h_k = h$, $\nabla h_k = \nabla h$ and $L_Y h_k = L_Y h$ on U_k.

Since $Z_t = \exp(-\int_0^t c(Y_s)\, ds)$ is absolutely continuous in t with $dZ_t = -Z_t\, c(Y_t)\, dt$, it has locally bounded variation and hence $dZ_t^\top\, dY_t = 0$. Applying Itô's product rule to $h_k(Y_t)Z_t$ on $[0, t \wedge \sigma_k]$, we have

$$h_k(Y_{t\wedge\sigma_k})Z_{t\wedge\sigma_k} - h_k(x) = \int_0^{t\wedge\sigma_k} Z_s\big(L_Y h_k - c\, h_k\big)(Y_s)ds + \int_0^{t\wedge\sigma_k} Z_s \nabla h_k(Y_s)^\top \sigma(Y_s) dW_s. \tag{27}$$

Since $Y_s \in \bar{D}_k \subset U_k$ for $s \leq t \wedge \sigma_k$, every term above may be written in terms of h, yielding (25).

Finally, $\sigma_k \leq k$ is a bounded stopping time, ∇h_k is continuous with compact support, hence bounded on $\mathbb{R}^n$. In addition, σ is bounded by Assumption 2, and $Z_s \in (0, 1]$ since $c \geq 0$. Hence, we have $\mathbb{E}\big[\int_0^k |Z_s\, \nabla h_k(Y_s)^\top \sigma(Y_s)|^2 ds\big] < \infty$ and the optional stopping theorem (Karatzas & Shreve, 2014, Theorem 1.3.22) applies. Taking $t = k$ in (25), so that $t \wedge \sigma_k = \sigma_k$, and taking expectations proves Eq. (26). □

B Proof of Results in Section 3

B.1 Proof of Lemma 1

Since $\bar{D}$ is bounded by Assumption 3, we have $b \in C^{0,\alpha}(\bar{D})$ for $\alpha \in (0, 1)$. We show that the SDE (1) with $X_0 = x \in D$ has a pathwise unique strong solution, and its exit time τ_D satisfies $\mathbb{E}_x[\tau_D] \leq \varphi(x) < \infty$, where $\varphi \in C^{2,\alpha}(\bar{D})$, $\varphi \geq 0$ is the unique solution of the elliptic BVP

$$L\varphi = -1 \text{ in } D, \qquad \varphi|_{\partial D} = 0. \tag{28}$$

Step 1: Consider the elliptic BVP $L\varphi = -1$ in D, $\varphi|_{\partial D} = 0$. Since $b \in C^{0,\alpha}(\bar{D})$ and $a = \sigma\sigma^\top \in C^{0,\alpha}(\bar{D})$ by Assumption 2, $\partial D \in C^{2,\alpha}$ by Assumption 3, and the boundary data $0 \in C^{2,\alpha}(\partial D)$, elliptic regularity theory (Ladyzhenskaya & Uraltseva, 1968, Ch. III, Theorem 1.1 and Theorem 1.2) guarantees existence and uniqueness of a solution $\varphi \in C^{2,\alpha}(\bar{D})$.

Step 2: To establish that $\varphi \geq 0$ on $\bar{D}$, we will show that the minimum of φ over $\bar{D}$ is attained at a boundary point and is therefore equal to 0 since $\varphi|_{\partial D} \equiv 0$. Suppose, for contradiction, that φ attains a local minimum at some interior point $x_0 \in D$. Then $\nabla\varphi(x_0) = 0$ and $\operatorname{Hess}(\varphi(x_0)) \succeq 0$ by the second-order necessary condition for a local minimum. And we have

$$L\varphi(x_0) = \underbrace{b(x_0)^\top \nabla\varphi(x_0)}_{=0} + \tfrac{1}{2} \operatorname{tr}\big(a(x_0) \operatorname{Hess}(\varphi(x_0))\big) \;\geq\; 0. \tag{29}$$

Here, by uniform ellipticity in Assumption 2, we have $a(x_0) \succ 0$, and since $\operatorname{Hess}(\varphi(x_0)) \succeq 0$, the inequality holds. This contradicts $L\varphi \equiv -1 < 0$ on D. Hence φ has no interior local minimum, so $\min_{\bar{D}} \varphi = \min_{\partial D} \varphi = 0$, i.e. $\varphi \geq 0$ on $\bar{D}$.

Step 3: Since b, σ are bounded and Lipschitz by Assumption 1, the SDE (1) has a pathwise unique strong solution for all $t \geq 0$ (Oksendal, 2000, Thm. 5.2.1).

Step 4: We next show $\mathbb{E}_x[\tau_D] \le \varphi(x)$. By the previous step, Eq. (1) has a pathwise unique strong solution, so Lemma 5 applies with $h = \varphi$ and $c \equiv 0$. Let $(\sigma_k)_{k\ge1}$ and $(D_k)_{k\ge1}$ be the stopping times and open sets defined in Lemma 5. Since $L\varphi = -1$ on D, Eq. (26) gives $\mathbb{E}_x[\varphi(X_{\sigma_k})] - \varphi(x) = -\mathbb{E}_x[\sigma_k]$. By Step 2, $\varphi \ge 0$ on $\bar{D}$. Since $D_k \subset D$, we have $\sigma_k \le \tau_{D_k} \le \tau_D$. And, since the path stays in $\bar{D}$ up to its exit time τ_D, we have $X_{\sigma_k} \in \bar{D}$. Thus we have $\mathbb{E}_x[\varphi(X_{\sigma_k})] \ge 0$, and hence

$$\mathbb{E}_x[\sigma_k] \le \varphi(x), \qquad k \in \mathbb{N}. \tag{30}$$

Since the sequence $(\sigma_k)_{k\ge1}$ is nondecreasing and $\sigma_k \to \tau_D$ from Lemma 5(i), monotone convergence applied to (30) gives $\mathbb{E}_x[\tau_D] = \lim_k \mathbb{E}_x[\sigma_k] \le \varphi(x) < \infty$. Here, since $\varphi \in C(\bar{D})$ and $\bar{D}$ is compact, $\max_{\bar{D}} \varphi < \infty$, so $\sup_{x\in D} \mathbb{E}_x[\tau_D] \le \max_{\bar{D}} \varphi < \infty$. Finally, a nonnegative random variable with finite expectation is finite almost surely, hence $\tau_D < \infty$ a.s.

B.2 Finite Mean Exit Time for Controlled Processes

By Assumption 1, b is bounded and Lipschitz on $\mathbb{X}$. More generally, for $u \in \mathcal{U}$, $b + au$ is bounded and Lipschitz on $\mathbb{X}$, as a is bounded and Lipschitz by Assumption 2, and u is Lipschitz on $\bar{D}$, which is compact, hence bounded, and is extended to $\mathbb{X}$. Hence, the following result on the finite mean exit time for every admissible control is immediate by applying Lemma 1 to $\tilde{b} = b + au$.

Corollary 1. *Consider the controlled SDE* (9)*, and denote* $\tilde{L}g := (b + au)^\top \nabla g + \frac{1}{2}\sum_{ij} a_{ij}\partial_{ij} g$. *Let Assumptions 1–3 hold and let* $u \in \mathcal{U}$*. Then, the controlled SDE* (9) *starting from* $x \in D$ *has a pathwise unique strong solution* $\tilde{X}$*, and its exit-time* $\tilde{\tau}_D$ *satisfies* $\tilde{\mathbb{E}}_x[\tilde{\tau}_D] < \infty$*. Thus we also have* $\tilde{\tau}_D < \infty$*,* $\tilde{P}^{u,x}$*-a.s.*

B.3 Proof of Lemma 2

To show $w = \rho$ on D, let $Y_t := w(X_t)$ and $Z_t := \exp(-\int_0^t f(X_s)\,ds)$. Fix $x \in D$ and we apply Lemma 5 with $h = w$, and $c = f$ that satisfies Assumption 4. Since $Lw - fw = 0$ in D, the drift term in (26) vanishes and we have

$$\mathbb{E}_x\big[w(X_{\sigma_k})Z_{\sigma_k}\big] = w(x), \qquad k \in \mathbb{N}. \tag{31}$$

We now take $k \to \infty$. By Lemma 1, $\tau_D < \infty$, P^x-a.s. Since $\sigma_k \to \tau_D$ and D is open, $X_{\sigma_k} \to X_{\tau_D} \in \partial D$. By continuity of w on $\bar{D}$ and using the boundary condition (11), we have $Y_{\sigma_k} = w(X_{\sigma_k}) \to e^{-\Phi(X_{\tau_D})}$ a.s. By continuity of $t \mapsto \int_0^t f(X_s)\,ds$, we have $Z_{\sigma_k} \to \exp(-\int_0^{\tau_D} f(X_s)\,ds)$ a.s. Moreover, since $w \in C(\bar{D})$ with $\bar{D}$ compact and $Z_t \in (0,1]$, we have $|Y_{\sigma_k}Z_{\sigma_k}| \le \max_{\bar{D}}|w| < \infty$ for every k. Hence, $Y_{\sigma_k}Z_{\sigma_k} \to e^{-\Phi(X_{\tau_D})}\exp(-\int_0^{\tau_D} f\,ds)$ a.s., and this sequence is uniformly bounded by $\max_{\bar{D}}|w|$. Thus, the bounded convergence theorem gives

$$\lim_{k\to\infty} \mathbb{E}_x\big[Y_{\sigma_k}Z_{\sigma_k}\big] = \mathbb{E}_x\left[e^{-\Phi(X_{\tau_D})}\exp\left(-\int_0^{\tau_D} f(X_s)\,ds\right)\right] = \rho(x), \tag{32}$$

where the last equality follows from the definition ρ in (10). On the other hand, $\mathbb{E}_x[Y_{\sigma_k}Z_{\sigma_k}] = w(x)$ for every k by (31), so the limit on the left is $w(x)$. Therefore, we conclude $w(x) = \rho(x)$ on D. For $x \in \partial D$, we have $\tau_D = 0$, P^x a.s. hence (10) gives $\rho(x) = e^{-\Phi(x)} = w(x)$ via the boundary condition (11). Hence, $w(x) = \rho(x)$ on $\bar{D}$.

It remains to show $\rho > 0$ on $\bar{D}$. On the boundary, $\rho|_{\partial D} = e^{-\Phi} = g > 0$ by Assumption 4. In the interior, fix $x \in D$: by Assumption 4, $f \ge 0$ is bounded on $\bar{D}$ and $\Phi = -\log g$ is bounded on ∂D, while $\tau_D < \infty$ a.s. by Lemma 1. Hence, we have

$$\exp\Big(-\int_0^{\tau_D} f(X_t)\,dt - \Phi(X_{\tau_D})\Big) > 0, \qquad P^x\text{-a.s.}, \tag{33}$$

and taking expectations gives $\rho(x) > 0$. Combining these, $\rho > 0$ on $\bar{D}$. Since $\rho \in C(\bar{D})$ and $\bar{D}$ is compact, it follows that $\min_{\bar{D}} \rho > 0$.

B.4 Proof of Lemma 3

Recall $x = (x_1, \dots, x_n) \in \mathbb{X}$, and ∂_{ij} denotes the second partial derivative with respect to x_i and x_j. The gradient of ρ is given by

$$\nabla \rho = e^{-v}(-\nabla v) = -\rho \, \nabla v. \tag{34}$$

For the second-order part $\sum_{ij} a_{ij} \partial_{ij} \rho$, we have

$$\sum_{ij} a_{ij} \partial_{ij} \rho = -\sum_{ij} a_{ij} \partial_j \rho \, \partial_i v - \rho \sum_{ij} a_{ij} \partial_{ij} v. \tag{35}$$

Substituting $\partial_j \rho = -\rho \, \partial_j v$, we have

$$\sum_{ij} a_{ij} \partial_{ij} \rho = \rho \sum_{ij} a_{ij} \partial_i v \, \partial_j v - \rho \sum_{ij} a_{ij} \partial_{ij} v = \rho \big[(\nabla v)^\top a \, \nabla v - \Delta_a v \big], \tag{36}$$

where $\Delta_a v := \sum_{ij} a_{ij} \partial_{ij} v$. Hence, $L\rho$ can be written as

$$\begin{aligned} L\rho =& b^\top (-\rho \, \nabla v) + \tfrac{1}{2} \rho \big[(\nabla v)^\top a \, \nabla v - \Delta_a v \big] \\ =& \rho \left[-b^\top \nabla v - \tfrac{1}{2} \Delta_a v + \tfrac{1}{2} (\nabla v)^\top a \, \nabla v \right] \\ =& \rho \left[-Lv + \tfrac{1}{2} (\nabla v)^\top a \, \nabla v \right], \end{aligned} \tag{37}$$

where the last line follows from $b^\top \nabla v + \frac{1}{2} \Delta_a v = Lv$. Using $L\rho - f\rho = 0$, and since $\rho > 0$, divide by ρ, we have

$$-Lv + \tfrac{1}{2} (\nabla v)^\top a \, \nabla v = f. \tag{38}$$

Rearranging the above equation, we have

$$Lv(x) + f(x) = \tfrac{1}{2} (\nabla v(x))^\top a(x) \, \nabla v(x) \quad \text{in } D, \qquad v\big|_{\partial D} = \Phi. \tag{39}$$

We next show that this coincides with the HJB equation for the OSC (12) up to an exit time. Let $\tilde{X}_t$ solve the controlled SDE (9). Since $v \in C^2(D)$, Itô's formula gives

$$dv(\tilde{X}_t) = \underbrace{\big[Lv(\tilde{X}_t) + (\nabla v(\tilde{X}_t))^\top a(\tilde{X}_t) \, u(\tilde{X}_t) \big]}_{\text{drift}} \, dt + \nabla v(\tilde{X}_t)^\top \sigma(\tilde{X}_t) \, d\tilde{W}_t. \tag{40}$$

For $u \in \mathcal{U}$, substituting (39) into the drift, and adding and subtracting $\frac{1}{2} u^\top a \, u$, we have

$$\text{drift} = -\Big(f + \tfrac{1}{2} u^\top a \, u \Big) + \tfrac{1}{2} (\nabla v + u)^\top a \, (\nabla v + u), \tag{41}$$

where the last term ≥ 0 since $a \succeq 0$. Equation (41) also shows that, for each $x \in D$, the mapping $u \mapsto u^\top a \nabla v + \frac{1}{2} u^\top a u$ is minimized uniquely at $u^* = -\nabla v$, since $a \succ 0$ by Assumption 2. Therefore, (13) is the HJB equation of the OSC (12) up to an exit time, with a candidate optimal controller u^* in (14).

B.5 Proof of Proposition 1

Since $H \geq 0$, we have $He^{-H} \geq 0$ and hence $-\mathbb{E}_x[He^{-H}] < \infty$. Together with $Z \in (0, \infty)$, hypotheses of Mitter & Newton (2003, Proposition 2.1) hold on $(\Omega, \mathcal{F}_{\tau_D})$ with reference measure $P^x|_{\mathcal{F}_{\tau_D}}$ and energy H. Hence, the claim follows.

B.6 Proof of Lemma 4

Since $\rho \in C^{2,\alpha}(\bar{D})$ as established in Lemma 2, and $-\log$ is smooth on $[m, M]$, the chain rule gives $v = -\log \circ \rho \in C^{2,\alpha}(\bar{D})$, thus we have $u^* = -\nabla v \in C^{1,\alpha}(\bar{D})$, hence $u^* \in \mathcal{U}$. By the paragraph preceding Corollary 1, $\tilde{b} := b - a \nabla v$ is bounded and Lipschitz on $\mathbb{R}^n$. Thus, (9) has a pathwise unique strong solution for all $t \geq 0$ by the standard SDE existence-uniqueness theorem for bounded, Lipschitz coefficients; see, e.g., (Oksendal, 2000, Thm. 5.2.1).

Define an exponential process as

$$Z_t := \exp\left(\int_0^t \left((\sigma^\top u^*)(X_s)\right)^\top dW_s - \tfrac{1}{2}\int_0^t |\sigma^\top u^*(X_s)|^2\, ds\right), \qquad t < \tau_D. \tag{42}$$

Since $v \in C^{2,\alpha}(\bar{D})$ and $\bar{D}$ is compact by Assumption 1, ∇v is continuous and hence bounded on $\bar{D}$. Together with σ bounded by Assumption 2, let $C := \|\sigma\|_\infty \|\nabla v\|_\infty$ and we have $\left|\sigma^\top u^*\right| \le C$ on $\bar{D}$. Since $\tau_D < \infty$ a.s. by Lemma 1, we have $\int_0^{\tau_D} |\sigma^\top u^*(X_s)|^2\, ds \le C^2 \tau_D < \infty$ a.s. Thus the stochastic integral in (42) is well defined at $t = \tau_D$, and $t \mapsto Z_t$ is continuous on $[0, \tau_D]$.

Step 1: We compute Z_t as follows. Since $a = \sigma\sigma^\top$ and $u^* = -\nabla v$, we have $(\nabla v)^\top a\, \nabla v = |\sigma^\top u^*|^2$, so the HJB equation (39) is $Lv = \frac{1}{2}|\sigma^\top u^*|^2 - f$ in D. Applying Lemma 5 with $h = v$ and $c = 0$, and taking k large enough that $\sigma_k > t$, which is possible since $\sigma_k \to \tau_D$, for $t < \tau_D$, we have,

$$dv(X_t) = Lv(X_t)\, dt + \nabla v(X_t)^\top \sigma(X_t)\, dW_t = \Big(\tfrac{1}{2}|\sigma^\top u^*(X_t)|^2 - f(X_t)\Big)dt - \left((\sigma^\top u^*)(X_t)\right)^\top dW_t, \tag{43}$$

where we used $\nabla v = -u^*$ and the HJB equation (13).

Rearranging and integrating from 0 to $t < \tau_D$, we have

$$\int_0^t \left((\sigma^\top u^*)(X_s)\right)^\top dW_s - \tfrac{1}{2}\int_0^t |\sigma^\top u^*(X_s)|^2\, ds = v(x) - v(X_t) - \int_0^t f(X_s)\, ds. \tag{44}$$

Taking exponential and using $e^{-v} = \rho$, we have

$$Z_t = \frac{\rho(X_t)}{\rho(x)} \exp\left(-\int_0^t f(X_s)\, ds\right), \qquad t < \tau_D. \tag{45}$$

Step 2: We now take $t \to \tau_D$ in (45). By Lemma 1, we have $\tau_D < \infty$, P^x-a.s., thus $X_t \to X_{\tau_D} \in \partial D$ by continuity of the path, and $\rho(X_t) \to \rho(X_{\tau_D}) = e^{-\Phi(X_{\tau_D})}$ by continuity of ρ on $\bar{D}$ together with the boundary condition $\rho|_{\partial D} = e^{-\Phi}$ in (11). Moreover, $\int_0^t f(X_s)\, ds \to \int_0^{\tau_D} f(X_s)\, ds$ by continuity of $t \mapsto \int_0^t f(X_s)\, ds$, the limit is finite since f is bounded by Assumption 4 and $\tau_D < \infty$ a.s. Hence, we have

$$Z_{\tau_D} = \frac{e^{-\Phi(X_{\tau_D})}}{\rho(x)} \exp\Big(-\int_0^{\tau_D} f(X_s)\, ds\Big) = \frac{e^{-H(\chi)}}{\rho(x)}, \qquad P^x\text{-a.s.} \tag{46}$$

Taking expectations, the Feynman–Kac identity $\rho(x) = \mathbb{E}_x[e^{-H(\chi)}]$ of (10) gives

$$\mathbb{E}_x\big[Z_{\tau_D}\big] = \frac{\rho(x)}{\rho(x)} = 1. \tag{47}$$

Step 3: Since $\bar{D}$ is compact and $\rho \in C(\bar{D})$, we have $0 < m \le \rho \le M < \infty$ on $\bar{D}$; and $\exp(-\int_0^t f\, ds) \in (0, 1]$ since $f \ge 0$. Hence (45) gives the uniform, deterministic bound

$$0 < Z_{t\wedge\tau_D} \le \frac{M}{\rho(x)} \qquad \text{for every } t \ge 0 \text{ and a.e. } \omega. \tag{48}$$

The stopped process $Z_{\cdot\wedge\tau_D}$ is a stochastic exponential, hence a local martingale; by (48) it is bounded, and a bounded local martingale is a true martingale. Therefore $\mathbb{E}_x[Z_{t\wedge\tau_D}] = 1$ for every $t \ge 0$.

Step 4: Finally, define Q on $\mathcal{F}_{\tau_D}$ via $dQ/dP^x = Z_{\tau_D}$, which is a probability measure by the above. By Girsanov's theorem, we have $W_t - \int_0^t \sigma^\top u^*(X_s)\, ds$ is a Q-Brownian motion on $[0, \tau_D]$. Thus under Q, the process solves the controlled SDE (9) with $u = u^*$ up to τ_D. By uniqueness in law for said SDE, established above, we have $Q = \tilde{P}^{*,x}$ on $\mathcal{F}_{\tau_D}$, which together with (46) proves (15).

B.7 Proof of Theorem 1

Part (i). We first consider $u \in \mathcal{U}$, and consider the optimal control u^* at the end as the special case. We proceed in four steps: (1) applying Itô's formula for $v(\tilde{X}_t)$ under $u \in \mathcal{U}$; (2) applying Lemma 5 to integrate v on $\bar{D}$; (3) passing $k \to \infty$; and (4) rearranging and specializing to $u = u^*$.

Step 1: From Itô's formula and (41), for the controlled process $\tilde{X}$,we have

$$dv(\tilde{X}_t) = -\Big(f(\tilde{X}_t) + \tfrac{1}{2}u^\top a\, u(\tilde{X}_t)\Big)dt + \tfrac{1}{2}\big(\nabla v + u\big)^\top a\,\big(\nabla v + u\big)(\tilde{X}_t)\,dt + \nabla v(\tilde{X}_t)^\top \sigma(\tilde{X}_t)\,d\tilde{W}_t. \tag{49}$$

Step 2: The above equation is derived from $v \in C^2(D)$, so it may be integrated only while the trajectory remains in D. We therefore apply Lemma 5 with drift $b + au$, $h = v$ and $c \equiv 0$ along the controlled diffusion process (9), where admissibility of u and Assumption 2 make the drift term bounded and Lipschitz. Substituting the drift decomposition in (41) into (26), we have

$$\underbrace{\tilde{\mathbb{E}}\big[v(\tilde{X}_{\sigma_k})\big]}_{:=T_1} - v(x) = \underbrace{-\tilde{\mathbb{E}}\left[\int_0^{\sigma_k}\Big(f + \tfrac{1}{2}u^\top a\, u\Big)dt\right]}_{:=T_2} + \underbrace{\tilde{\mathbb{E}}\left[\int_0^{\sigma_k}\tfrac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)\,dt\right]}_{:=T_3}. \tag{50}$$

Step 3: We next study the three terms T_1-T_3 in (50) separately. For the boundary term T_1, recall by Corollary 1, $\tilde{\tau}_D < \infty$ a.s., and we have shown before that $\sigma_k \to \tilde{\tau}_D$ as $k \to \infty$ in Lemma 5. Hence $\tilde{X}_{\sigma_k} \to \tilde{X}_{\tilde{\tau}_D}$ a.s. by continuity of the path. Since the controlled process exits D at time $\tilde{\tau}_D$, we have $\tilde{X}_{\tilde{\tau}_D} \in \partial D$ a.s., so continuity of v on $\bar{D}$ together with the boundary condition $v|_{\partial D} = \Phi$ gives $v(\tilde{X}_{\sigma_k}) \to v(\tilde{X}_{\tilde{\tau}_D}) = \Phi(\tilde{X}_{\tilde{\tau}_D})$ a.s. Moreover $v \in C(\bar{D})$ with $\bar{D}$ compact gives $|v| \le M := \max_{\bar{D}}|v| < \infty$, so $|v(\tilde{X}_{\sigma_k})| \le M$ for every k and a.e. ω. By bounded convergence theorem, we thus have

$$\lim_{k\to\infty}\tilde{\mathbb{E}}\big[v(\tilde{X}_{\sigma_k})\big] = \tilde{\mathbb{E}}\big[\Phi(\tilde{X}_{\tilde{\tau}_D})\big]. \tag{51}$$

For the running cost T_2, since $f \ge 0$ by Assumption 4 and $a \succ 0$ by Assumption 2, we have that $f + \frac{1}{2}u^\top a\, u$ is nonnegative. As σ_k is nondecreasing with $\sigma_k \to \tilde{\tau}_D$ from Lemma 5, the integrals $\int_0^{\sigma_k}(f + \frac{1}{2}u^\top a\, u)\,dt$ are nondecreasing in k and converge pointwise a.s. to $\int_0^{\tilde{\tau}_D}(f + \frac{1}{2}u^\top a\, u)\,dt$. By monotone convergence theorem, we have

$$\lim_{k\to\infty}\tilde{\mathbb{E}}\left[\int_0^{\sigma_k}\Big(f + \tfrac{1}{2}u^\top a\, u\Big)dt\right] = \tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D}\Big(f + \tfrac{1}{2}u^\top a\, u\Big)dt\right]. \tag{52}$$

For the last term T_3, proceed similarly as above, since $a \succ 0$, $\frac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)$ is nonnegative, so monotone convergence theorem gives

$$\lim_{k\to\infty}\tilde{\mathbb{E}}\left[\int_0^{\sigma_k}\tfrac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)\,dt\right] = \tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D}\tfrac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)\,dt\right]. \tag{53}$$

We next show that all three limits are finite. Indeed, by Assumption 4, $\Phi = -\log g$ is continuous on the compact set ∂D, hence bounded. Let $C := \sup_{\bar{D}}\Big(f + \frac{1}{2}u^\top a\, u\Big)$ and $C' := \sup_{\bar{D}}\frac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)$. Both terms are finite because f is bounded by Assumption 4, a is bounded by Assumption 2, $u \in \mathcal{U}$ is bounded on the compact $\bar{D}$ by definition, and ∇v is bounded on $\bar{D}$ since $v \in C^{2,\alpha}(\bar{D})$. The two integrals in (52) and (53) are dominated by $C\tilde{\tau}_D$ and $C'\tilde{\tau}_D$, respectively, whose expectations are finite since $\tilde{\mathbb{E}}\big[\tilde{\tau}_D\big] < \infty$ by Corollary 1. Combining (51)–(53) into (50), we have

$$\tilde{\mathbb{E}}\big[\Phi(\tilde{X}_{\tilde{\tau}_D})\big] - v(x) = -\tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D}\Big(f + \tfrac{1}{2}u^\top a\, u\Big)dt\right] + \tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D}\tfrac{1}{2}(\nabla v + u)^\top a\,(\nabla v + u)\,dt\right], \tag{54}$$

where every term is finite.

Step 4: Rearranging (54), and using $\nabla v + u = u - u^*$ as $u^* = -\nabla v$, we obtain

$$\begin{aligned} J(u,x) &= \tilde{\mathbb{E}}\big[\Phi(\tilde{X}_{\tilde{\tau}_D})\big] + \tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D}\Big(f + \tfrac{1}{2}u^\top a u\Big)dt\right] \\ &= v(x) + \tilde{\mathbb{E}}\left[\int_0^{\tilde{\tau}_D} \tfrac{1}{2}(u-u^*)^\top a(u-u^*)dt\right] \\ &\geq v(x), \end{aligned} \tag{55}$$

where the inequality follows from $a \succ 0$ by Assumption 2. Since $u \in \mathcal{U}$ was arbitrary, $v(x)$ is a lower bound. Here, equality holds if and only if $\tilde{\mathbb{E}}\big[\int_0^{\tilde{\tau}_D} \frac{1}{2}(u-u^*)^\top a\,(u-u^*)\,dt\big] = 0$. Since the integrand is nonnegative, this holds if and only if $\int_0^{\tilde{\tau}_D} \frac{1}{2}(u-u^*)^\top a\,(u-u^*)(\tilde{X}_t)\,dt = 0$, $\tilde{P}^{u,x}$-a.s. Again, by nonnegativity, it holds if and only if $(u-u^*)^\top a\,(u-u^*)(\tilde{X}_t) = 0$ for Lebesgue-almost every $t \in [0, \tilde{\tau}_D(\omega)]$,$\tilde{P}^{u,x}$-a.s. By uniform ellipticity in Assumption 2, we have $(u-u^*)^\top a\,(u-u^*) \geq c\,|u-u^*|^2$ with $c > 0$. So this is equivalent to $u(\tilde{X}_t) = u^*(\tilde{X}_t)$ for Lebesgue-almost every $t \in [0, \tilde{\tau}_D]$, $\tilde{P}^{u,x}$-a.s.

It remains to check that this bound is attained. Since $v \in C^{2,\alpha}(\bar{D})$, we have $u^* = -\nabla v \in C^{1,\alpha}(\bar{D})$, hence $u^* \in \mathcal{U}$ and (55) applies. The remainder term then vanishes identically, and we have $J(u^*, x) = v(x)$. Hence, the infimum is attained, and for every $x \in D$, we have

$$v(x) = J(u^*, x) = \min_{u \in \mathcal{U}} J(u, x). \tag{56}$$

This completes the proof of Part (i).

Part (ii) Since $\tau_D < \infty$, P^x-a.s. as in Lemma 1, $H(\chi)$ is well-defined and finite P^x-a.s., and by the Feynman–Kac representation of ρ, we can define

$$Z := \mathbb{E}_x\big[e^{-H(\chi)}\big] = \rho(x) \in (0, 1], \tag{57}$$

so $Z \in (0, \infty)$ as required by (Raginsky, 2026, Proposition 1). Moreover $f \geq 0$ by Assumption 4 and $\Phi = -\log g \geq 0$ since $g \leq 1$ by the same assumption, so $H(\chi) \geq 0$ P^x-a.s. Hence the integrand $-H(\chi)e^{-H(\chi)}$ is non-positive and $-\mathbb{E}_x\big[H(\chi)\,e^{-H(\chi)}\big] < \infty$. So both assumptions of Proposition 1 are satisfied, and applying it on $(\Omega, \mathcal{F}_{\tau_D})$ with $P = P^x$ and $H = H(\chi)$, for every admissible path measure $\tilde{P} \ll P^x$, we have

$$D_{\mathrm{KL}}(\tilde{P} \| P^x) + \mathbb{E}[H(\chi)] \;\geq\; -\log \rho(x) = v(x), \tag{58}$$

with equality if and only if $d\tilde{P}/dP^x = e^{-H(\chi)}/\rho(x)$. By Part (i), we have $v(x) = J(u^*, x)$, so (58) is (16). By Lemma 4, $\tilde{P}^{*,x}$ of the optimally-controlled process satisfies

$$\left.\frac{d\tilde{P}^{*,x}}{dP^x}\right|_{\mathcal{F}_{\tau_D}} = \frac{e^{-H(\chi)}}{\rho(x)} \tag{59}$$

Hence equality in (58), and so in (16), holds if and only if $\tilde{P} = \tilde{P}^{*,x}$.

C Proof of Identity (4) in Section 4.1

Let Assumption 2 hold. By construction in Section 4.1, A_i, A_j are closed, and $\partial A_i \cap \partial A_j = \emptyset$. Assume that $D_q = \mathbb{X} \setminus (A_i \cup A_j)$ satisfies Assumption 3, with $\partial D_q = \partial A_i \cup \partial A_j$. Denote $\varphi : \partial D_q \to \{0, 1\}$ with $\varphi = 0$ on ∂A_i and $\varphi = 1$ on ∂A_j. Since ∂A_i and ∂A_j are disjoint compact sets, φ is locally constant, hence $\varphi \in C^{2,\alpha}(\partial D_q)$. By (Ladyzhenskaya & Uraltseva, 1968, Ch. III, Theorem 1.1 and Theorem 1.2), the boundary value problem

$$Lw = 0 \text{ in } D_q, \qquad w|_{\partial D_q} = \varphi, \tag{60}$$

has a unique solution $w \in C^{2,\alpha}(\bar{D}_q)$. In particular, w is continuous and bounded on the compact set $\bar{D}_q$.

We next show $q = w$. For $x \in D_q$, let $\tau := \inf\{t > 0 : X_t \notin D_q\}$ be the first time the trajectory leaves D_q. Since $\mathbb{X} \setminus D_q = A_i \cup A_j$, we have $\tau = \tau_{A_i} \wedge \tau_{A_j}$. Let D_k, σ_k be as in Lemma 5 with $D = D_q$. Apply Lemma 5(ii) with $h = w$ and $c = 0$, we have $Z_t = 1$. Since $Lw = 0$ in D_q, from (26), we have

$$\mathbb{E}_x\big[w(X_{\sigma_k})\big] = w(x), \qquad k \in \mathbb{N}. \tag{61}$$

Here $w(X_{\sigma_k})$ is well defined since $X_{\sigma_k} \in \bar{D}_q$ by the case split in the proof of Lemma 5.

We next take $k \to \infty$. From Lemma 5(i), we have $\sigma_k \to \tau$. Since $\tau < \infty$ a.s., D_q is open, and by path continuity, we have $X_{\sigma_k} \to X_\tau \in \partial D_q$ a.s. Since w is continuous on $\bar{D}_q$, we have $w(X_{\sigma_k}) \to w(X_\tau) = \varphi(X_\tau)$ a.s. Furthermore, $|w(X_{\sigma_k})| \le \max_{\bar{D}_q} |w| < \infty$ for $k \in \mathbb{N}$. By the bounded convergence theorem, we then have

$$w(x) = \mathbb{E}_x\big[\varphi(X_\tau)\big] = P^x\big(X_\tau \in \partial A_j\big), \tag{62}$$

where the second equality follows since $\varphi(X_\tau) = \mathbf{1}_{\partial A_j}(X_\tau)$.

Finally, we show $P^x\big(X_\tau \in \partial A_j\big) = q(x)$. Since A_i, A_j are closed by construction, and paths are continuous, and by definition of τ, we have $X_{\tau_{A_k}} \in A_k$ whenever $\tau_{A_k} < \infty$ for $k \in \{i, j\}$. If $X_\tau \in \partial A_j \subset A_j$, then $\tau \neq \tau_{A_i}$, since otherwise $X_\tau \in A_i \cap A_j = \emptyset$, and we have $\tau = \tau_{A_j} < \tau_{A_i}$. Conversely, if $\tau_{A_j} < \tau_{A_i}$, then $\tau = \tau_{A_j}$, and we have $X_\tau \in \partial A_j$. Thus $\{X_\tau \in \partial A_j\} = \{\tau_{A_j} < \tau_{A_i}\}$, and for $x \in D_q$, we have

$$w(x) = P^x\big(\tau_{A_j} < \tau_{A_i}\big) = q(x). \tag{63}$$

By definition of q, for $x \in \partial A_i$ ($x \in \partial A_j$), we have $q(x) = 0 = w(x)$ ($q(x) = 1 = w(x)$). Hence $q = w$ on $\bar{D}_q$. Therefore, $q \in C^{2,\alpha}(\bar{D}_q)$, and q is the unique solution of (4).

D Detailed Computation of the Committor Function from Section 4.1

We first construct the spectral term $\tilde{\psi}$ that separates the two cores using the leading eigenfunctions of L. Consider A_k and the set of leading eigenfunctions ψ_l, for $l = 1, \dots, m-1$. With the core level $\theta_l^{(k)}$ defined in (17), we define $\hat{\psi}(x) := \sum_{l=1}^{m-1} c_l\, \psi_l(x)$, where $c_l := \theta_l^{(j)} - \theta_l^{(i)}$ measures how much the l-th mode varies between A_i, A_j. Define the core level of $\hat{\psi}$ as $\bar{\theta}^{(k)} := \frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \hat{\psi}(x_i)$. Substituting $\hat{\psi}$, we have $\bar{\theta}^{(k)} = \sum_l c_l \theta_l^{(k)}$. Let $c = [c_1, \cdots, c_{m-1}]$. We recenter and rescale $\hat{\psi}$ to get $\tilde{\psi}$ as

$$\tilde{\psi}(x) := \frac{\hat{\psi}(x) - \bar{\theta}^{(i)}}{\|c\|^2}, \qquad \bar{\theta}^{(i)} = \sum_l c_l \theta_l^{(i)}. \tag{64}$$

With this choice of c, we have $\bar{\theta}^{(j)} - \bar{\theta}^{(i)} = \|c\|^2$, and $\tilde{\psi}$ maps the centroid of A_i to 0 and that of A_j to 1. Since each ψ_l is an eigenfunction, we have $\nabla\hat{\psi} = \sum_{l=1}^{m-1} c_l\, \nabla\psi_l$, and $L\hat{\psi} = \sum_{l=1}^{m-1} c_l\, \lambda_l\, \psi_l$. Applying L to (64) and using that $\bar{\theta}^{(i)}$ and $\|c\|^2$ are constants, so that $L\bar{\theta}^{(i)} = 0$, we have $L\tilde{\psi} = L\hat{\psi}/\|c\|^2 = \frac{1}{\|c\|^2} \sum_{l=1}^{m-1} c_l\, \lambda_l\, \psi_l$. Using linearity of L, we obtain $Lh = Lq - L\tilde{\psi} = -L\tilde{\psi}$ in D_q. Together with the boundary conditions $q|_{\partial A_i} = 0$ and $q|_{\partial A_j} = 1$, we have that h solves

$$Lh = -\,L\tilde{\psi} \text{ in } D_q, \qquad h|_{\partial A_i} = -\tilde{\psi}|_{\partial A_i}, \qquad h|_{\partial A_j} = 1 - \tilde{\psi}|_{\partial A_j}. \tag{65}$$

We next approximate the remainder h in the ansatz $q = \tilde{\psi} + h$ in $\mathcal{H}_\kappa$. To solve (65) with $\tilde{\psi}$ in (64), let $\{x_l\}_{l \in \mathcal{I}_q} \subset D_q$ with $\mathcal{I}_q = \{1, \dots, N_q\}$, and consider solving for a $h \in \mathcal{H}_\kappa$ of the form

$$h(x) = \sum_{l \in \mathcal{I}_q} \alpha_l\, \kappa(x, x_l), \qquad \alpha = (\alpha_l)_{l \in \mathcal{I}_q}. \tag{66}$$

The problem then reduces to solving for α. To this end, we impose (65) it at interior points $\{x_n\}_{n\in\mathcal{J}_q} \subset D_q$ where the residual is minimized, and boundary points $\{x_n\}_{n\in\mathcal{J}_\partial} \subset \partial A_i \cup \partial A_j$ where the boundary condition is imposed. Specifically, define the matrix $A \in \mathbb{R}^{|\mathcal{J}_q|\times|\mathcal{I}_q|}$ and vector $g \in \mathbb{R}^{|\mathcal{J}_q|}$ with entries

$$A_{nl} := (L\kappa)(x_n, x_l), \qquad [g]_n := (L\tilde{\psi})(x_n) = \frac{1}{\|c\|^2}\sum_{l=1}^{m-1} c_l\lambda_l\psi_l(x_n), \tag{67}$$

for $n \in \mathcal{J}_q$. Similarly, define $K_\partial \in \mathbb{R}^{|\mathcal{J}_\partial|\times|\mathcal{I}_q|}$ and $b \in \mathbb{R}^{|\mathcal{J}_\partial|}$ as

$$(K_\partial)_{nl} := \kappa(x_n, x_l), \qquad (b)_n := \begin{cases} -\tilde{\psi}(x_n), & x_n \in \partial A_i, \\ 1-\tilde{\psi}(x_n), & x_n \in \partial A_j, \end{cases} \tag{68}$$

for $n \in \mathcal{J}_\partial$. Here, $L\kappa$ can be computed using the derivative of the kernel function κ. Specifically, for the Gaussian kernel $\kappa(x,y) = \exp\big(-\|x-y\|^2/2\sigma^2\big)$ with bandwidth σ, writing $z := x - y$, we have

$$\nabla_x\kappa(x,y) = -\frac{z}{\sigma^2}\,\kappa(x,y), \qquad \partial_{x_ix_j}\kappa(x,y) = \Big(\frac{z_iz_j}{\sigma^4} - \frac{\delta_{ij}}{\sigma^2}\Big)\kappa(x,y). \tag{69}$$

Hence, we have

$$(L\kappa)(x,y) = \left(-\frac{b(x)\cdot z}{\sigma^2} + \frac{z^\top a(x)\, z}{2\sigma^4} - \frac{\operatorname{tr} a(x)}{2\sigma^2}\right)\kappa(x,y). \tag{70}$$

We next solve for α in (66) by imposing the boundary condition at $\{x_n\}_{n\in\mathcal{J}_\partial}$ and minimizing the least-squares sense with Tikhonov regularization as

$$\begin{aligned} &\min_{\alpha} \ \big\|A\alpha + g\big\|^2 + \lambda_q\,\alpha^\top K\alpha, \\ &\text{s.t. } K_\partial\,\alpha = b, \end{aligned} \tag{71}$$

for $\lambda_q \geq 0$ and $K_{ll'} = \kappa(x_l, x^{l'})$. For such an equality-constrained quadratic program, define a multiplier $\mu \in \mathbb{R}^{|\mathcal{J}_\partial|}$. The Lagrangian is then $\|A\alpha+g\|^2 + \lambda_q\alpha^\top K\alpha + 2\mu^\top(K_\partial\alpha - b)$, and the KKT conditions in (α,μ) are

$$\begin{bmatrix} A^\top A + \lambda_q K & K_\partial^\top \\ K_\partial & 0 \end{bmatrix}\begin{bmatrix}\alpha\\ \mu\end{bmatrix} = \begin{bmatrix} -A^\top g \\ b\end{bmatrix}. \tag{72}$$

Solving for $\hat{\alpha}$ via the linear system (72) and plugging this into $q = \tilde{\psi} + h$, the committor function can be approximated via (19).

E Details Behind the Computation of ρ from Section 4.2

We start by constructing ρ_0. Consider $x \in D$, define $\tau_i := \inf\{t > 0 : X_t \in A_i\}$, and $\tau_j := \inf\{t > 0 : X_t \in A_j\} = \tau_D$. Let $R := \{\tau_j < \tau_i\}$ be the reactive event, i.e., the trajectory reaches the target set before returning to the source set. By definition of the committor function in Section 4.1, we have $q^{(i,j)}(x) = P^x(R)$ where P^x denotes the probability of starting from x. Conditioning $\mathbb{E}_x\left[e^{-\beta H_0}\right]$ in (20) on R and R^c, and using the law of total expectation, we have

$$\begin{aligned} \rho(x) =& \mathbb{E}_x\left[e^{-\beta H_0}\right] \\ =& P^x(R)\underbrace{\mathbb{E}_x\left[e^{-\beta H_0}\,\middle|\,R\right]}_{:=\rho_{\mathrm{R}}(x)} + P^x(R^c)\underbrace{\mathbb{E}_x\left[e^{-\beta H_0}\,\middle|\,R^c\right]}_{:=\rho_{\mathrm{NR}}(x)} \\ =& q^{(i,j)}(x)\rho_{\mathrm{R}}(x) + \big(1-q^{(i,j)}(x)\big)\,\rho_{\mathrm{NR}}(x), \end{aligned} \tag{73}$$

where $\rho_{\mathrm{R}}, \rho_{\mathrm{NR}} \in (0,1]$ since $H_0 \geq 0$, with the convention that a term is omitted when its conditioning event has probability zero. Thus $\rho(x)$ is a convex combination of $\rho_{\mathrm{R}}(x)$ and $\rho_{\mathrm{NR}}(x)$ with weight being $q(x)$. In practice, we use $\hat{q}$ from (19) as an approximation of q.

On R^c, the trajectory reaches A_i at time $\tau_i < \tau_j$. We therefore split $H_0 = H_0^i + H_0^{ij}$ with $H_0^i := \int_0^{\tau_i}(1 - q^{(i,j)})\,dt$ and $H_0^{ij} := \int_{\tau_i}^{\tau_j}(1 - q^{(i,j)})\,dt$. Conditioning on the path up to time τ_i, by the Markov property, we have $\mathbb{E}\big[e^{-\beta H_0^{ij}} \mid \mathcal{F}_{\tau_i}\big] = \rho(X_{\tau_i})$ where $X_{\tau_i} \in \partial A_i$. Since $e^{-\beta H_0} = e^{-\beta H_0^i} e^{-\beta H_0^{ij}}$, we have

$$\rho_{\mathrm{NR}}(x) = \mathbb{E}_x\left[e^{-\beta H_0^i}\,\rho(X_{\tau_i})\Big| R^c\right]. \tag{74}$$

For $\rho(X_{\tau_i})$, notice that the trajectory lingers inside A_i, and hence $\rho(X_{\tau_i}) \approx \rho|_{A_i}$. To compute the latter, let $y \in A_i$. Since $q^{(i,j)} \equiv 0$ on A_i and the trajectory spends nearly all of time $[0, \tau_j]$ inside A_i, we have $H_0^i \approx \tau_j$. Since the system is metastable, we assume the first hitting time is asymptotically exponential with mean $T_0 = 1/|\lambda_1|$, where λ_1 is the leading nontrivial eigenvalue of the Koopman generator. In other words,, τ_j has density $T_0^{-1}e^{-t/T_0}$ on $t \geq 0$, and we have

$$\rho_A \approx \mathbb{E}\left[e^{-\beta\tau_j}\right] = \int_0^\infty \frac{1}{T_0}\,e^{-t/T_0}\,e^{-\beta t}\,dt = \frac{1}{1+\beta T_0} =: r. \tag{75}$$

Hence, we have

$$\rho(x) = q^{(i,j)}(x)\rho_{\mathrm{R}}(x) + \Big[\big(1 - q^{(i,j)}(x)\big)r\Big]\;\mathbb{E}_x\left[e^{-\beta H_0^i}\,\Big|\,R^c\right]. \tag{76}$$

On R, the trajectory goes from x to A_j without visiting A_i, hence $\rho_{\mathrm{R}} = \mathbb{E}_x\left[e^{-\beta H_0}\,\middle|\,R\right]$ lies in $(0, 1]$ with H_0 the duration of reactive crossing. For term $\mathbb{E}_x\left[e^{-\beta H_0^i}\middle| R^c\right]$, on R^c, H_0^i is the time of a non-reactive path falling back into A_i. Therefore, neither term encodes the metastability of the system, and the metastability enters only through $q^{(i,j)}(x)$ and $\big(1 - q^{(i,j)}(x)\big)r$. We therefore consider the factorization of ρ as

$$\rho = \Big(q^{(i,j)}(x) + \Big(1 - q^{(i,j)}(x)\Big)\,r\Big) \times \frac{\rho}{q^{(i,j)}(x) + \big(1 - q^{(i,j)}(x)\big)r}. \tag{77}$$

Replacing $q^{(i,j)}$ by its approximation $\hat{q}$ from (19), we define $\rho_0(x) := \hat{q}(x) + \big(1 - \hat{q}(x)\big)r$, and the remainder is the H_0.

We next compute w in the ansatz $\hat{\rho}(x) := \rho_0(x)\,w(x)$. Since ρ_0 is affine in $\hat{q}$, we also have

$$\nabla\rho_0 = (1 - r)\,\nabla\hat{q}, \qquad L\rho_0 = (1 - r)\,L\hat{q}, \tag{78}$$

By the product rule, we have

$$L(\rho_0 w) = \rho_0 L w + w L\rho_0 + \nabla\rho_0^\top a \nabla w. \tag{79}$$

Thus $L\rho - f_\beta\rho = 0$ in (11) can be written in terms of w as

$$\rho_0\,Lw + w\,\big(L\rho_0 - f_\beta\,\rho_0\big) + \nabla\rho_0^\top\,a\,\nabla w = 0 \qquad \text{in } D. \tag{80}$$

To solve (80), define two matrices in $\mathbb{R}^{|I_D|\times N}$ as

$$(K_D)_{li} = \kappa(x_l, x_i), \qquad (LK)_{li} = (L\kappa)(x_l, x_i), \tag{81}$$

for $x_l \in I_D$ and $i = 1, \ldots, N$. In addition, define $\Gamma_D \in \mathbb{R}^{|I_D|\times N}$, and $P, Q \in \mathbb{R}^{|I_D|\times|I_D|}$ as

$$\begin{aligned} &(\Gamma_D)_{li} = \nabla\rho_0(x_l)^\top a(x_l)\,\nabla_x\kappa(x_l, x_i), \\ &P := \mathrm{diag}\,\big(\rho_0(x_l)\big), \quad Q := \mathrm{diag}\,\big(L\rho_0(x_l) - f_\beta(x_l)\rho_0(x_l)\big), \end{aligned} \tag{82}$$

where $\nabla_x\kappa$ is available in closed form as in (69), and $\nabla\hat{q} = \nabla\tilde{\psi} + \sum_l \alpha_l \nabla_x\kappa(\cdot, x_l)$ from (19). Using the chain rule and (21), we also have

$$L\rho_0 = \rho_0\Big(\eta\,L\hat{q} + \tfrac{\eta^2}{2}\,\nabla\hat{q}^\top a\,\nabla\hat{q}\Big). \tag{83}$$

Hence (81) together with (82), the residual of (80) at the interior points can thus be written as

$$\big[P\,LK + Q\,K_D + \Gamma_D\big]\,c =: S_D\,c, \qquad S_D \in \mathbb{R}^{|I_D|\times N}. \tag{84}$$

Likewise, with $b_l := e^{-\Phi(x_l)}/\rho_0(x_l)$ for $x_l \in I_{\partial D}$ and $(K_{\partial D})_{li} := \kappa(x_l, x_i)$, the boundary condition can be written as $K_{\partial D}\, c = b$. We then consider the following regularized optimization problem

$$\begin{aligned} \min_{c\in\mathbb{R}^N} \ & \|S_D\, c\|^2 + \lambda\, c^\top K c \\ \text{s.t. } & K_{\partial D}\, c = b, \end{aligned} \tag{85}$$

for $\lambda \geq 0$, and $K \in \mathbb{R}^{N\times N}$ is the Gram matrix, $K_{ii'} = \kappa(x_i, x^{i'})$, so that $c^\top K c = \|w\|^2_{\mathcal{H}_k}$. Here, the boundary condition is imposed exactly, while the interior residual is minimized in the least squares sense with Tikhonov regularization. Here, (85) is again an equality-constrained convex quadratic program. We define the Lagrangian $\|S_D c\|^2 + \lambda c^\top K c + 2\mu^\top (K_{\partial D} c - b)$, and obtain the KKT condition, which is a linear system in c and the Lagrange multiplier μ

$$\begin{pmatrix} S_D^\top S_D + \lambda K & K_{\partial D}^\top \\ K_{\partial D} & 0 \end{pmatrix} \begin{pmatrix} c \\ \mu \end{pmatrix} = \begin{pmatrix} 0 \\ b \end{pmatrix}. \tag{86}$$

With c determined by (86), applying the chain rule to $\log\hat{\rho} = \log\rho_0 + \log w$ with ρ_0 in (21), we obtain (22).

F Reweighting the Generated Paths

The controlled dynamics (9) generate transition paths more efficiently than the reference dynamics, but at the price of sampling from a tilted path measure rather than the reference one. To fix this, first consider starting from some initial condition $x \in D$. Recall that P^x is the path measure of (1) started at x, and $\tilde{P}^{*,x}$ is that of the optimally controlled process (9) started at x. By Girsanov's theorem, P^x and $\tilde{P}^{*,x}$ are equivalent masures on the common path space Ω. Let $\chi \in \Omega$ be a generic element of the path space. By Lemma 4, we have

$$\frac{d\tilde{P}^{*,x}}{dP^x} = \frac{e^{-H(\chi)}}{\rho(x)}, \tag{87}$$

where $\rho(x) = \mathbb{E}_x[e^{-H}]$, and $H(\chi) := \int_0^{\tilde{\tau}_D} f(\tilde{x}_t)\, dt + \Phi(\tilde{x}_{\tilde{\tau}_D})$. Hence, the correction from the controlled process (9) to the original process (1) at a fixed start is simply $\rho(x)e^{H(\chi)}$.

When the initial conditions are drawn from a probability measure ν on the source set A_i, e.g., the Boltzmann measure restricted to A_i, consider the mixture

$$M(\cdot) = \int_{A_i} P^{x_0}(\cdot)\, \nu(dx_0), \quad \tilde{M}(\cdot) = \int_{A_i} \tilde{P}^{*,x_0}(\cdot)\, \nu(dx_0). \tag{88}$$

Both M and $\tilde{M}$ share the same initial distribution ν. Since every path starts from its own initial condition $x_0 = \chi(0)$, we have

$$\frac{d\tilde{M}}{dM} = \frac{d\tilde{P}^{*,\chi(0)}}{dP^{\chi(0)}} = \frac{e^{-H(\chi)}}{\rho\big(\chi(0)\big)}. \tag{89}$$

Hence, to map $\tilde{M}$ back to M, each path is weighted by $dM/d\tilde{M}$ by Lemma 4. From (89) and replacing ρ by $\hat{\rho}$ in (21), we can approximate it as

$$\text{weight}(\chi) = \frac{dM}{d\tilde{M}} = \hat{\rho}\big(\tilde{x}_0\big)\, e^{H(\chi)}, \qquad \tilde{x}_0 := \chi(0). \tag{90}$$

Here, $\hat{\rho}(\tilde{x}_0)$ is given by (21).

We further note that as $\beta \to \infty$, the probability that a sampled path spends more than an arbitrary fixed time $\varepsilon > 0$ in nonreactive regions tends to zero. To see this, recall from Example 1 that $H = \beta H_0$. From (87), we have

$$\begin{aligned} \tilde{P}^{*,x}(H_0 > \varepsilon) = \frac{\mathbb{E}_x\left[e^{-\beta H_0}\mathbf{1}\{H_0 > \varepsilon\}\right]}{\mathbb{E}_x\left[e^{-\beta H_0}\right]} &\leq \frac{e^{-\beta\varepsilon}}{\mathbb{E}_x\left[e^{-\beta H_0}\right]} \\ &\leq \frac{e^{-\beta\varepsilon}}{\mathbb{E}_x\left[e^{-\beta H_0}\mathbf{1}\{H_0 \leq \varepsilon/2\}\right]} \\ &\leq \frac{e^{-\beta\varepsilon/2}}{P^x(H_0 \leq \varepsilon/2)}, \end{aligned} \tag{91}$$

for $\varepsilon > 0$ with $P^x(H_0 \le \varepsilon/2) > 0$. Here, the first inequality follows from $e^{-\beta H_0} \le e^{-\beta\varepsilon}$ on $\{H_0 > \varepsilon\}$. Taking $\beta \to \infty$, we have $\tilde{P}^{*,x}(H_0 > \varepsilon) \to 0$.

G Experiment Details and Additional Examples

G.1 The Double-Well System

For the two-dimensional double-well system, the drift and diffusion coefficients in (1) are specialized as $b(x) = -\nabla U(x)$, with

$$U(x,y) = \tfrac{1}{6}\Big(4(1-x^2-y^2)^2 + 2(x^2-2)^2 + \big((x+y)^2-1\big)^2 + \big((x-y)^2-1\big)^2 - 2\Big), \quad (92)$$

and $\sigma(x) = \sigma_0 I$ with $\sigma_0 = \sqrt{2k_B T_{\mathrm{emp}}} \approx 0.455$ at $T_{\mathrm{emp}} = 1200\,\mathrm{K}$. Its potential has two minima at $(\pm\sqrt{5}/2,\, 0)$ and two channels around $(0, \pm 1)$. Following Algorithm 1, we compute the Koopman eigenfunctions and extract the metastable cores A_{left} and A_{right} as described in Section 4.1. Specifically, A_{left} and A_{right} are obtained by retaining those points whose spectral embedding lies within radius $\varepsilon = 0.5$ of its cluster centroid. We sample 1000 uncontrolled and 1000 controlled paths, all initialized at the minimum of A_{left}. A path is a hit at its first entrance to A_{right}, and is truncated at the numerical censor $T_{\max}$ otherwise. We use a step size of $\Delta t = 0.01$, and $T_{\max} = 10$, equivalently, a budget of $T_{\max}/\Delta t = 1000$ integration steps, and $\beta = 1$. We adopt two metrics to evaluate performance: the target hit percentage (THP), which measures the proportion of trajectories that successfully reach the target before $T_{\max}$; and the energy of the transition state (ETS), which measures the maximum potential energy along each transition path. As shown in Figure 1 and Table 1, the controller raises the THP from 0% to 99.8%, and the controlled ensemble retains both channels rather than collapsing onto one.

Table 1: Target hit percentage (THP) and energy of the transition state (ETS) for the 2-channel double well. Uncontrolled and controlled rows are measured in our setup; TPS-DPS values are as reported by Seong et al. (2025).

Method	THP (↑) %	ETS (↓) kJmol^{-1}
Uncontrolled (1200K)	0.0	–
Uncontrolled (2400K)	2.7	1.73 ± 0.42
Uncontrolled (3600K)	18.8	2.07 ± 0.46
Uncontrolled (4800K)	42.1	2.49 ± 0.66
TPS-DPS (1200K)	99.9	1.38 ± 0.16
Proposed method (1200K)	**99.8**	1.40 ± 0.11

G.2 Muller-Brown Potential

The Müller-Brown potential (Müller & Brown, 1979) is described by

$$U(x,y) = \frac{1}{10}\sum_{k=1}^{4} A_k \exp\left(a_k(x-x_k)^2 + b_k(x-xk)(y-y_k) + c_k(y-y_k)^2\right), \quad (93)$$

with parameters specified as

$$\begin{pmatrix} A_1 & a_1 & b_1 & c_1 & x_1 & y_1 \\ A_2 & a_2 & b_2 & c_2 & x_2 & y_2 \\ A_3 & a_3 & b_3 & c_3 & x_3 & y_3 \\ A_4 & a_4 & b_4 & c_4 & x_4 & y_4 \end{pmatrix} = \begin{pmatrix} -200 & -1 & 0 & -10 & 1 & 0 \\ -100 & -1 & 0 & -10 & 0 & 0.5 \\ -170 & -6.5 & 11 & -6.5 & -0.5 & 1.5 \\ 15 & 0.7 & 0.6 & 0.7 & -1 & -1 \end{pmatrix} \quad (94)$$

We consider a particle with $m = 12\mathrm{a.m.u}$ and $\gamma = 1\mathrm{ps}^{-1}$. As shown in Figure 2, the dynamics has one deep well at $(-0.558, 1.442)$, a medium-deep well at $(0.623, 0.028)$, and a shallow well sitting

Table 2: Müller-Brown benchmark at $T = 150\,\mathrm{K}$ and $200\,\mathrm{K}$.

Method	THP (↑) %	ETS (↓) kJmol^{-1}
Uncontrolled (150K)	0.1	12.4
Uncontrolled (200K)	1.8	14.0 ± 1.3
TPS-DPS (150K)	100	15.9 ± 3.5
TPS-DPS (200K)	100	15.7 ± 1.9
Proposed method (150K, $\beta = 1$)	95.1	14.1 ± 2.1
Proposed method (150K, $\beta = 2$)	95.9	14.3 ± 2.1
Proposed method (200K, $\beta = 1$)	98.4	14.5 ± 2.1
Proposed method (200K, $\beta = 2$)	**99.0**	14.5 ± 2.1

between the two at $(-0.050, 0.467)$. The goal is to transition from the deep well to the medium well. The source and target cores are defined with radius $\varepsilon = 0.2$. We choose a stepsize of $\Delta t = 0.002\mathrm{ps}$ and $T_{\max} = 10\mathrm{ps}$. We use a Gaussian kernel with bandwidth 0.35. At temperature $T_{\mathrm{emp}} = 200K$, the controller raises the THP from 1.8% for the uncontrolled process to 99.0%.

G.3 Alanine Dipeptide

We simulate alanine dipeptide (ACE–ALA–NME) in vacuum with AMBER99SB-ILDN, hydrogen-bond constraints, and the VVVR integrator at $300\,\mathrm{K}$, friction $10^{-3}\,\mathrm{fs}^{-1}$, and timestep $1\,\mathrm{fs}$. Trajectories start from the energy-minimized $C5$ structure. Following (Seong et al., 2025), the target is defined as $\mathcal{B} = \{\mathbf{R} \mid ||\xi(\mathbf{R}) - \xi(\mathbf{R}_{\mathcal{B}})|| < 0.75\}$ (in radians), where $\xi = (\phi, \psi)$ are the backbone dihedral angles and $\mathbf{R}_{\mathcal{B}}$ is the energy-minimized $C7_{\mathrm{ax}}$ structure. A trajectory stops when it enters $\mathcal{B}$ or reaches $T_{\max} = 1\mathrm{ps}$.

The controller $\hat{u} = \nabla \log \hat{\rho}$ is evaluated on the Kabsch-aligned Cartesian configuration and applied as an external force $F_{\mathrm{bias}} = s\,\hat{u}$ in OpenMM, with $s = 30\,\mathrm{kJ\,mol^{-1}nm^{-1}}$ and $\|\hat{u}\|$ capped at $75\,\mathrm{kJ\,mol^{-1}nm^{-1}}$ per atom. We use $\beta = 1$, Gaussian bandwidth 1, and $\varepsilon = 0.01$, with 80 kernel centers and 400 interior and 48 boundary collocation points. OpenMM integrates the underdamped Langevin dynamics

$$dx_i = \frac{p_i}{m_i} dt, \quad dp_i = F_i(x)dt - \gamma_i p_i dt + \sqrt{2 m_i \gamma_i k_B T} dW_i,$$

where i indexes the atom. In this experiment, we solve the boundary value problem (11) with the isotropic generator $L = F \cdot \nabla + \frac{\sigma_0^2}{2}\Delta$ where F is the force and $\sigma_0 = 1$ as an approximation of the simulated dynamics that ignores the per-atom mobility $1/(m_i\gamma)$.

The proposed pipeline as summarized in Algorithm 1 runs once: Koopman eigenfunctions are constructed via a generalized eigenproblem in terms of Gram matrices (see, e.g., (Hou et al., 2023, Appendix M) for details), followed by solving for $\hat{q}$ and $\hat{u}$. In contrast, the baselines in (Seong et al., 2025) repeatedly simulate controlled trajectories and update the neural network, so they pay the sampling cost per rollout. Table 3 compares the cost of obtaining the bias in the drift term. All timings are measured on a single NVIDIA A100 GPU.

G.4 Additional Experiment: The Four-Well Dynamics

Consider the two-dimensional four-well potential described by

$$U(x, y) = 5((x^2 - 1)^2 + (y^2 - 1)^2), \tag{95}$$

which has four symmetric local minima at $(\pm 1, \pm 1)$. The source core is centered at pair $(-1, -1)$ and the target core is centered at $(1, 1)$, with radius $\varepsilon = 0.3$. We use mass $m = 12$ a.m.u, friction $\gamma = 1\mathrm{ps}^{-1}$, stepsize $\Delta t = 0.002\mathrm{ps}$, and $T_{\max} = 10\mathrm{ps}$. Results for various settings are reported in 4, and the sampled paths at $200\,\mathrm{K}$ are shown in Figure 6.

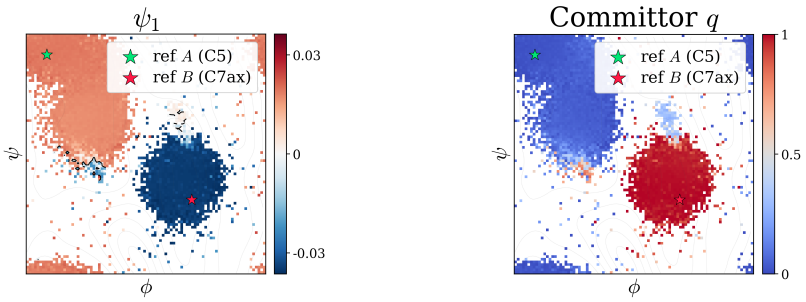
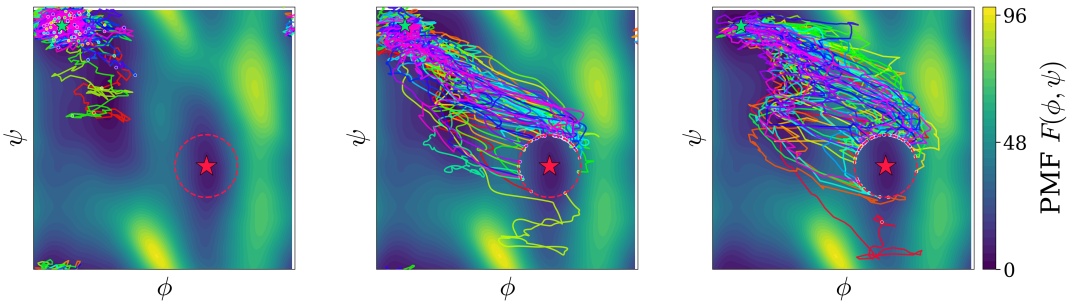


Figure 4: Alanine dipeptide on the Ramachandran plain of backbone dihedral angles (ϕ, ψ). (Left) The first nontrivial Koopman eigenfunction ψ_1 which separates the $C5$ basin from $C7_{\mathrm{ax}}$ basin. (Right) The committor q from $C5$ to $C7_{\mathrm{ax}}$.

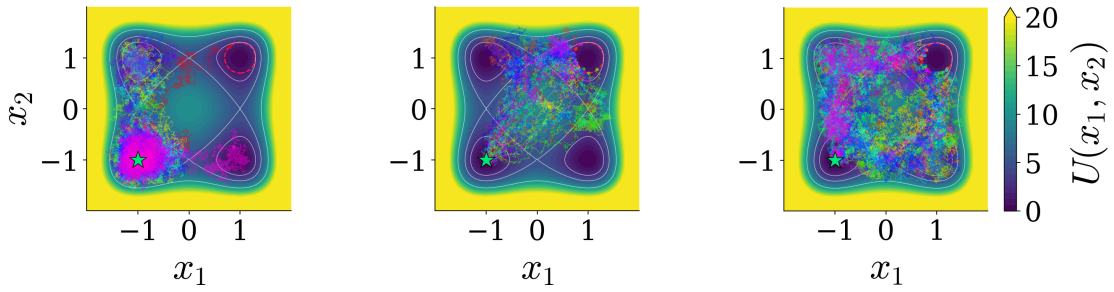


(a) Uncontrolled (b) TPS-DPS (c) Proposed method

Figure 5: Sampled transition paths for alanine dipeptide in vacuum, same as Figure 3 in Section 5. Paths start in the $C5$ basin (upper left) and end in the $C7_{\mathrm{ax}}$ basin (red star). (a) Uncontrolled paths. (b) TPS-DPS (Seong et al., 2025): THP $= 100\%$, obtained by training a neural-network bias with simulation in the loop (1000 rollouts, $\sim 10^6$ gradient updates). (c) Proposed method: THP $= 93\%$, with a closed-form controller.

Table 3: Cost of obtaining the controller on alanine dipeptide on a single NVIDIA A100 GPU, excluding the subsequent path generation, which is common to all methods.

Proposed method							
Stage			Time (s)				Total (s)
Sampling			Direct evaluation of (3)				
Koopman eigenpairs			73.8				180
Committor $\hat{q}$			56.2				
Controller $\hat{u}$			51.0				
Training-based baselines							
Method	Rollouts	$\times$	(Sampling (s)	$+$	Training (s))	$=$	Total (h)
TPS-DPS (F)	1000	$\times$	(60.6	$+$	14)	$=$	20.7
PIPS (F)	15000	$\times$	(60.6	$+$	5.4)	$=$	275



(a) Uncontrolled (b) TPS-DPS (c) Proposed method

Figure 6: Four-well dynamics at $T = 200\,\mathrm{K}$. We aim to sample transitions from the source core at $(-1, -1)$ (green star) to the target core at $(1, 1)$ (red star). (a) Uncontrolled paths: THP($T_{\max}$) = 0.9%. (b) TPS-DPS (Seong et al., 2025): THP($T_{\max}$) = 81.6%. (c) Proposed method with $\beta = 2$: THP($T_{\max}$) = 94.7%.

Table 4: Four-well benchmark at $T = 200\,\mathrm{K}$ and $300\,\mathrm{K}$.

Method	THP (↑) %	ETS (↓) kJmol^{-1}
Uncontrolled (200K)	0.9	10.6 ± 3.2
Uncontrolled (300K)	8.8	12.8 ± 3.6
TPS-DPS (200K)	81.6	$8.6 \pm .9$
TPS-DPS (300K)	87.1	$10.2pm.6$
Proposed method (200K, $\beta = 1$)	73.2	9.5 ± 1.8
Proposed method (200K, $\beta = 2$)	94.7	9.0 ± 1.8
Proposed method (300K, $\beta = 1$)	91.8	11.2 ± 3.1
Proposed method (300K, $\beta = 2$)	**100.0**	10.4 ± 2.5