Initializing Doob's Lagrangian with the Committor: A Freidlin–Wentzell Perspective on Rare Event Sampling

Anonymous authors
Paper under double-blind review

Abstract

Sampling rare transitions between metastable states is a fundamental challenge in computational chemistry because physically relevant transition paths occur on timescales far exceeding those accessible to direct simulation. Variational path optimization techniques based on minimizing a stochastic path action have been developed for decades, but most of these approaches require propagating the dynamics of multiple replicas of the system of interest. Doob's Lagrangian (DL) casts the problem as the minimization of an action over controlled path measures, with a neural control field and a training objective that does not require simulating the dynamics, but, in practice, it is sensitive to initialization and unstable under a misspecified transit time T. In this work, we establish a mathematical equivalence between the DL framework and Freidlin–Wentzell (FW) large-deviation theory, showing that minimizing the Doob action over both the control field and T is asymptotically equivalent to computing the quasipotential governing Arrhenius–Kramers transition rates. This equivalence motivates a principled initialization strategy based on the stationary committor function, for which we prove, when a single transition channel dominates, that the initial control lies exponentially close in T to the optimal Doob control field and that the joint (v, T) optimization is stable. On alanine dipeptide and chignolin at finite temperature, the committor initialization combined with the optimization of T removes the collapse of the transit time to unphysically small values observed under linear initialization, and recovers physically meaningful transition paths where the original DL method fails.

1 Introduction

Molecular dynamics (MD) can elucidate physical phenomena with high spatial and temporal precision, well beyond what is typically accessible experimentally (Hollingsworth & Dror, 2018). As such, MD has become a fundamental tool in a variety of contexts, from understanding reaction mechanisms in catalysis (Piccini et al., 2022) to characterizing metastable states of important drug targets (Shukla et al., 2014). Nevertheless, capturing *dynamical transitions* between metastable states is often impractical because these transitions are exponentially rare on simulation timescales, making direct equilibrium sampling statistically intractable (Freidlin & Wentzell, 1984; Hänggi et al., 1990). Enhanced sampling methods attempt to overcome this challenge by either biasing the dynamics directly, as umbrella sampling and metadynamics do along collective variables supplied *a priori* (Laio & Parrinello, 2002; Torrie & Valleau, 1977), or resampling trajectories, as in sequential Monte Carlo (Doucet et al., 2001; Del Moral, 2004). Learned generative approaches arose as an alternative, in which a control or a distribution over transition paths is optimized against energetic or dynamical objectives (Holdijk et al., 2023; Vargas et al., 2023). Recent work has proposed reformulating statistical conditioning in the sense of Doob using a Lagrangian formulation (Du et al., 2024; Lee et al., 2025), which casts transition path sampling as a variational problem over controlled path measures. The method steers a reference diffusion toward a target state B within a fixed transit time T, using a flow-matching parametrization (Lipman et al., 2023) to produce a time-dependent ensemble with Gaussian marginals and optimizes the control field by minimizing an action functional related to Doob's h-transform. In practice, despite this

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.

theoretical appeal, the method exhibits two interrelated failure modes: (i) *poor initialization*, in which linear interpolation between states converges to a high-energy local minimum rather than the physical transition path; and (ii) *poor transit-time selection*, in which a misspecified fixed T drives the system non-physically, producing paths that do not represent equilibrium barrier crossings.

We resolve both failure modes by establishing a theoretical equivalence between the Doob's Lagrangian (DL) and Freidlin–Wentzell (FW) large-deviation frameworks (Freidlin & Wentzell, 1984; Bouchet & Reygner, 2016). This equivalence shows that minimizing the Doob action over both the control field *and* the transit time T is asymptotically equivalent to computing the FW quasipotential. It simultaneously motivates a committor-based initialization (Rotskoff et al., 2022; Mitchell & Rotskoff, 2024) that places the optimizer in a provably locally convex basin, eliminating the collapse of the transit time to unphysically small values that we observed under linear initialization. Stochastic optimal control and large-deviation formulations of rare events are classical (Fleming, 1977; Freidlin & Wentzell, 1984), and neural networks enter this framework as a representation of the control field. The contribution of this work is an optimization strategy for a physically meaningful objective on path space, designed together with the theory that identifies its minimizer and the initialization from which it is reached.

Contributions. We summarize our contributions as follows.

- **FW–Doob equivalence.** We prove that the Doob action evaluated on characteristic paths is proportional to the Freidlin-Wentzell action (Lemma 3.1), and that their respective minimizers coincide in the vanishing noise limit (Theorem 3.1). Minimizing the Doob action over (v, T) therefore recovers the quasipotential barrier that governs Arrhenius–Kramers rates (Corollary 3.1).
- **Committor-based initialization.** We initialize the Doob control field from the stationary committor function, and prove that, when a single transition channel dominates, this places the optimizer exponentially close (in T) to the optimal Doob control field v^*, ensuring stable joint convergence over (v, T) (Theorem 3.2).
- **Diagnosis of failure modes.** We identify and demonstrate, both theoretically and empirically, the two failure modes of the original DL method and show through the learned transit times that the proposed modifications eliminate them while retaining the simulation-free training objective, on benchmark molecular systems at finite temperature.

2 Related Work

The DL framework belongs to the family of methods that learn a control field, parametrized by a neural network, through the minimization of a variational objective on path measures.

Rare-event sampling. Umbrella sampling (Torrie & Valleau, 1977), metadynamics (Laio & Parrinello, 2002), and forward-flux sampling (Allen et al., 2005) require collective variables or an order parameter supplied *a priori*, and their statistical efficiency relies on these variables resolving the slow degrees of freedom of the transition. Transition path sampling (Bolhuis et al., 2002) performs Monte Carlo in trajectory space, given metastable states, at the cost of integrating the dynamics for each proposal path. Minimum-action methods require only the drift, the noise model, and the endpoints of the transition, and rest on the large-deviation principle of Freidlin & Wentzell (1984). In the vanishing-noise limit, reactive trajectories concentrate along the minimizer of the FW action; in particular, the infimum over all transit times and paths gives access to the quasipotential, a quantity related to the Arrhenius–Kramers exit time estimate (Hänggi et al., 1990; Bouchet & Reygner, 2016). The minimum action method (E et al., 2004), and its geometric formulation (Heymann & Vanden-Eijnden, 2008) compute the minimizer on a discretized path, with a number of unknowns equal to the number of discretization points times the dimension, and relax it with gradient updates. DL (Du et al., 2024) minimizes an action on path space at finite noise over a continuous representation of the path. The time marginals of the controlled process are Gaussians, or mixtures of Gaussians, whose means and covariances are parametrized using neural nets, and the time integral in the action is estimated by sampling times uniformly and configurations from these marginals, so that training requires neither integration of the dynamics nor a time grid. We show that, when the transit time is jointly optimized with the control, the minimizers of this action coincide with those of the FW action, in the vanishing noise limit.

Doob h-transform and controlled diffusions. The connection between rare-event sampling and stochastic optimal control is classical (Fleming, 1977; Day, 1987). Du et al. (2024) cast Doob's h transform as a variational objective over path measures, parameterized via a flow matching method (Lipman et al., 2023; Albergo & Vanden-Eijnden, 2023), and optimized simulation-free; Lee et al. (2025) extended this to broader boundary conditions and larger peptide systems. A complementary line of work directly integrates on the path-action to generate transition trajectories. Adjanor & Athenes (2005) gave early path-sampling estimates of free energies via direct integration of the path action, and Raja et al. (2025) recently showed that pre-trained generative models can be repurposed zero-shot for transition path sampling by minimizing the Onsager–Machlup action functional. The variational path sampling framework of Das et al. (2022); Singh & Limmer (2023); Singh et al. (2025) takes yet another route, parameterizing a time-dependent control force and optimizing the KL divergence between the controlled and reactive ensembles; Singh & Limmer (2024) further showed that the committor gradient provides the optimal such controller, closely connecting this line of work to ours. Yan et al. (2022) demonstrated that neural-network control forces can characterize dynamical phase transitions at scale, working directly with the large-deviation cost in trajectory space. Related controlled-diffusion samplers include diffusion Schrödinger bridges (de Bortoli et al., 2021), score-based models (Song et al., 2021), and SMC approaches (Doucet et al., 2001; Del Moral, 2004). We identify the FW equivalence in the DL formulation and use it to resolve the initialization and transit-time failure modes of (Du et al., 2024).

Committor functions. The committor $\tilde{q}(x)$, i.e., the probability of reaching state B before A from x, is the central quantity of transition path theory (E & Vanden-Eijnden, 2006). Rotskoff et al. (2022) proposed a variational neural committor amenable to stochastic gradient optimization, and Mitchell & Rotskoff (2024) used learned committors to guide the estimation of molecular transition rates. To our knowledge, Theorem 3.2 provides the first rigorous guarantee for initializing the DL optimizer with a learned committor, bounding the distance between the committor-based control and the optimal control v^* by a term exponentially small in T.

3 Theory

The DL framework (Du et al., 2024) and the more established FW large-deviation theory (Freidlin & Wentzell, 1984; Bouchet & Reygner, 2016; Grafke et al., 2017) have been developed independently for transition path sampling. Our central theoretical contribution is to show that these two perspectives are equivalent in the vanishing-noise limit, i.e., minimizing the Doob action over the control field and the transit time T recovers the FW quasipotential, and therefore the Arrhenius–Kramers rate exponent. This equivalence explains why DL recovers physically meaningful instantons, exposes T as a free optimization variable on equal footing with v, and motivates a committor-based initialization which in our experiments removes the collapse of the transit time to unphysically small values observed under linear initialization. Proofs of all results in this section are given in Appendix A.

3.1 Background and notation

Stochastic system. We consider an Itô diffusion on $\mathbb{R}^d$,

$$dX_t^\epsilon = b(X_t^\epsilon)\,dt + \sqrt{2\epsilon}\,\sigma(X_t^\epsilon)\,dW_t\,, \tag{1}$$

with smooth drift $b\colon \mathbb{R}^d \to \mathbb{R}^d$, smooth diffusion coefficient $\sigma\colon \mathbb{R}^d \to \mathbb{R}^{d\times m}$, diffusion tensor $a(x) := \sigma(x)\sigma(x)^\top$, and noise level $\epsilon > 0$. We assume b admits two stable equilibria $A, B \subset \mathbb{R}^d$ separated by a saddle $x_\star$, so that $A \to B$ transitions are exponentially rare for small ϵ. As $\epsilon \downarrow 0$ the behavior of the process is described by FW large-deviation theory (Freidlin & Wentzell, 1984).

Freidlin–Wentzell action. The probability of rare excursions can be estimated in the vanishing-noise limit using the FW functional (Freidlin & Wentzell, 1984). For a path ϕ in the Sobolev space $H^1([T_1, T_2])$ of absolutely continuous paths with square-integrable velocity such that $\phi_{T_1} = x$, the FW action is defined as

$$S_{[T_1,T_2],x}(\phi) = \frac{1}{4}\int_{T_1}^{T_2} \big\langle \dot{\phi}_t - b(\phi_t),\ a(\phi_t)^{-1}\big(\dot{\phi}_t - b(\phi_t)\big)\big\rangle\,dt\,, \tag{2}$$

where $S_{[T_1,T_2],x}(\phi) = +\infty$ if $\phi \notin H^1([T_1,T_2])$. Paths that follow the drift have zero action; deviations cost $\|\dot{\phi}_t - b(\phi_t)\|^2_{a^{-1}}$. The *quasipotential*

$$V(A,x) = \inf\{S_{[T_1,T_2],A}(\phi)\colon \phi_{T_1} = A,\ \phi_{T_2} = x,\ T_1 < T_2\} \tag{3}$$

is the central object of FW theory, and it sets the exponential rate of $A \to B$ barrier crossings.

Doob's h-transform and the controlled SDE. The DL framework (Du et al., 2024) steers (1) toward a target event by solving a variational problem over path measures. The reference and controlled SDEs read

$$dX_t = b(X_t)\,dt + \Xi_t\,dW_t\,, \qquad \text{(reference } \mathbb{P}\text{)} \tag{4}$$

$$d\tilde{X}_t = \big[b(\tilde{X}_t) + 2G_t v_t(\tilde{X}_t)\big]dt + \Xi_t\,dW_t\,, \qquad \text{(controlled } \mathbb{Q}^v\text{)} \tag{5}$$

with noise amplitude Ξ_t, diffusion matrix $G_t = \frac{1}{2}\Xi_t\Xi_t^\top$, and control field $v_t(x)$. The FW setting is recovered by the choice $\Xi_t = \sqrt{2\epsilon}\,\sigma(X_t)$, giving $G_t = \epsilon\,a(X_t)$, in which case (4) coincides with (1). From now on we drop the superscript ϵ.

Action functionals. Two action functionals appear throughout the paper. The *stochastic Doob action* is a functional of a path of probability densities $q = (q_t)_{t\in[0,T]}$, a control field v, and $T > 0$,

$$S^{\mathrm{DL}}_{[0,T]}[q,v] = \int_0^T \int_{\mathbb{R}^d} q_t(x)\,\langle v_t(x),\, G_t\,v_t(x)\rangle\,dx\,dt\,. \tag{6}$$

We write $q^v = (q^v_t)_{t\in[0,T]}$ for the time marginals of the controlled path measure $\mathbb{Q}^v$. The transit time T enters the Doob action as a free parameter, which we optimize jointly with v. In the vanishing-noise limit q^v_t concentrates on the characteristic path φ^v_t defined in (10), and the Doob action reduces to its value on this path,

$$S^{\mathrm{DL}}_{[0,T],A}(\varphi^v) := S^{\mathrm{DL}}_{[0,T]}[\delta_{\varphi^v}, v] = \int_0^T \big\langle v(\varphi^v_t),\ G_t\,v(\varphi^v_t)\big\rangle\,dt\,. \tag{7}$$

3.2 KL divergence and the Doob action

Applying Girsanov's theorem (Girsanov, 1960), the Radon-Nikodym derivative of $\mathbb{Q}^v$ with respect to the reference path measure $\mathbb{P}$ is explicit, and the resulting KL divergence corresponds to the cost of the control, as stated in Eq. 11 of Du et al. (2024).

Proposition 3.1 (KL–Doob action equivalence)**.** *Let $\mathbb{P}$ and $\mathbb{Q}^v$ denote the path measures associated with* (4) *and* (5) *respectively, and let q^v_t be the marginal density of $\mathbb{Q}^v$ at time t. Then*

$$D_{\mathrm{KL}}(\mathbb{Q}^v\|\mathbb{P}) = S^{\mathrm{DL}}_{[0,T]}[q^v, v]\,. \tag{8}$$

The endpoint-constrained problem $\min_v D_{\mathrm{KL}}(\mathbb{Q}^v\|\mathbb{P})$ subject to $\mathbb{Q}^v_0 = \delta_A$, $\mathbb{Q}^v_T = \delta_B$ is therefore equivalent to $\min_{q,v} S^{\mathrm{DL}}_{[0,T]}[q,v]$, with unique minimizer $v^*_t(x) = \nabla_x \log h(x,t)$, where h solves the backward Kolmogorov equation

$$\partial_t h + \langle \nabla h,\, b\rangle + \mathrm{tr}(G_t\nabla^2 h) = 0\,, \qquad h(x,T) = \mathbf{1}_B(x)\,. \tag{9}$$

3.3 From Doob's Lagrangian to the Freidlin–Wentzell action

We prove the equivalence of the two action functionals in three stages. Along characteristic paths, the two functionals differ only by the factor ϵ (Lemma 3.1). Each problem then admits a value function solving a PDE, a first-order Hamilton-Jacobi-Bellman equation for FW and, through the Cole-Hopf transform of h, a viscous Hamilton-Jacobi equation for DL. A vanishing-viscosity argument shows that the two value functions coincide as $\epsilon \to 0$ (Lemma 3.2 and Lemma 3.3, Theorem 3.1). As a final step, targeting the metastable state B is equivalent to targeting the saddle point $x_\star$ (Proposition 3.2), so the minimal action yields the Arrhenius-Kramers rate exponent (Corollary 3.1).

Characteristic paths. For any v, the *characteristic path* $\varphi^v = (\varphi_t^v)_{t\in[0,T]}$ is the vanishing-noise limit of (5),

$$\dot{\varphi}_t^v = b(\varphi_t^v) + 2G_t v(\varphi_t^v)\,, \qquad \varphi_0^v = A\,. \tag{10}$$

In this limit, q_t^v concentrates on $\delta_{\varphi_t^v}$, and $S^{\mathrm{DL}}_{[0,T]}[q^v, v]$ reduces to (7).

Lemma 3.1 (Action scaling). *For any admissible control field v and its associated characteristic path φ^v, the FW action and the deterministic Doob action on the characteristic path are related by*

$$S_{[0,T],A}(\varphi^v) = \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^v)\,. \tag{11}$$

Since (11) holds pointwise, minimizing S^{DL} and minimizing S over the same path family yield identical first-order optimality conditions; and identifying their minima requires the PDE argument of Section 3.4.

3.4 The quasipotential as the Doob energy

Following a standard vanishing-viscosity argument (Fleming, 1977; Day, 1987), which derives the value-function PDE of each problem, we show that the FW value function I and the vanishing-noise limit W^0 of the Doob energy coincide.

Lemma 3.2 (FW value function and HJB equation). *Define the FW value function*

$$I(t,y) = \inf_{\substack{\phi\in H^1([t,T])\\ \phi_t=y,\ \phi_T=x_{\mathrm{target}}}} S_{[t,T],y}(\phi)\,. \tag{12}$$

Assuming $I \in C^1([0,T)\times\mathbb{R}^d)$, it is the unique classical solution of the Hamilton–Jacobi–Bellman equation

$$\frac{\partial I}{\partial t} + \langle b(y), \nabla_y I\rangle - \langle \nabla_y I,\ a(y)\,\nabla_y I\rangle = 0\,, \tag{13}$$

with terminal condition $I(T,\, x_{\mathrm{target}}) = 0$.

Lemma 3.3 (Doob energy and vanishing-viscosity limit). *Define the* Doob energy *via the Cole–Hopf transformation*

$$W^\epsilon(t,y) = -\epsilon \log h(y,t)\,, \qquad h(y,t) = e^{-W^\epsilon(t,y)/\epsilon}\,. \tag{14}$$

Then W^ϵ satisfies the viscous Hamilton–Jacobi equation

$$\frac{\partial W^\epsilon}{\partial t} + \langle b(y), \nabla W^\epsilon\rangle - \langle \nabla W^\epsilon,\ a(y)\,\nabla W^\epsilon\rangle + \epsilon\,\mathrm{tr}\big(a(y)\,\nabla^2 W^\epsilon\big) = 0\,. \tag{15}$$

In the limit $\epsilon\to 0$, assuming pointwise convergence $W^\epsilon \to W^0$, the viscosity term vanishes and W^0 satisfies

$$\frac{\partial W^0}{\partial t} + \langle b(y), \nabla W^0\rangle - \langle \nabla W^0,\ a(y)\,\nabla W^0\rangle = 0\,. \tag{16}$$

Theorem 3.1 Optimality equivalence

Since $W^\epsilon(T, x_{\mathrm{target}}) = 0$, the optimal Doob action satisfies $\epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = W^\epsilon(0,A)$, and in the vanishing-noise limit it converges to the FW value function,

$$\lim_{\epsilon\to 0}\ \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = I(0,A)\,. \tag{17}$$

Therefore, minimizing the Doob action over (v,T) is asymptotically equivalent to computing the quasipotential,

$$\lim_{\epsilon\to 0}\ \inf_{T>0}\ \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = V(A, x_{\mathrm{target}})\,. \tag{18}$$

3.5 Saddle-point property and the rate exponent

For metastable systems the physical bottleneck is the saddle $x_\star$, the following proposition shows that targeting B or $x_\star$ gives exactly the same quasipotential, since the trajectory from $x_\star$ to B can follow the drift, along which the FW action vanishes.

Proposition 3.2 (Saddle-point property). *Let A and B be two stable equilibria of b, and let $x_\star$ be the saddle point on the basin boundary. Then*

$$V(A,B) = V(A, x_\star)\,. \tag{19}$$

Corollary 3.1 Doob action recovers the rate exponent

Under the assumptions of Theorem 3.1 and Proposition 3.2, minimizing the Doob action over (v, T) recovers the quasipotential barrier $V(A, x_\star)$ that sets the Arrhenius–Kramers rate,

$$\lim_{\epsilon \to 0} \inf_{T>0} \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = V(A, B) = V(A, x_\star)\,. \tag{20}$$

Proof follows immediately from Theorem 3.1 with $x_{\text{target}} = B$, and Proposition 3.2.

Corollary 3.1 places DL in the same family as classical instanton methods, e.g., gMAM (Heymann & Vanden-Eijnden, 2008), MAM (E et al., 2004), which compute the same minimizer over an explicit discretization of the path, whereas DL represents it through a learned control field v trained by automatic differentiation.

3.6 Committor-based initialization

So far we have established that, in the vanishing-noise limit, the minimizer of the Doob action coincides with the FW minimizer (Theorem 3.1). We now address a question left open by Du et al. (2024), namely whether there is an initialization for which convergence is theoretically guaranteed. Naively initialized to a linear interpolation $v^{\mathrm{lin}}_t(x) \propto (B-A)/T$, the joint optimization drives T to unphysically small values, forcing the path across high-energy barriers, which is highly unphysical. We therefore initialize from the stationary committor (Rotskoff et al., 2022; Mitchell & Rotskoff, 2024).

Notation. Let $\Omega \subseteq \mathbb{R}^d$ be a physical domain, with $A, B \subset \Omega$, let $\tilde{q} : \Omega \to [0, 1]$ denote the stationary committor, and let $\tilde{v}(x) := \nabla \log \tilde{q}(x)$ the associated control. Write $v^*_t(x) := \nabla \log h(x, t)$ for the optimal time-dependent Doob control, $q^*_t := q^{v^*}_t$ for the time marginal densities of $\mathbb{Q}^{v^*}$, $\mathcal{S} := \bigcup_{t\in(0,T]} \mathrm{supp}(q^*_t)$ for their support, and $\tau^v(A) := \mathbb{E}^v_A[\tau_{\partial B}]$ for the mean first-passage time to ∂B under $\mathbb{Q}^v$ started at A, where $\tau_{\partial B}$ denotes the first hitting time of ∂B. The weighted norm of a time dependent field $f = (f_t)_{t\in[0,T]}$ is

$$\|f\|^2_{L^2(q^*)} := \int_0^T \int_\Omega q^*_t(x)\, |f(x)|^2\, dx\, dt\,, \tag{21}$$

and L is the Lipschitz constant of $\nabla^2[(\tilde{q} - h(\cdot, 0))/\tilde{q}]$ on $\mathcal{S}$.

Assumption 3.1 (No path metastability)**.** *The transition from A to B is mediated by a unique dominant instanton, i.e. the FW action has a unique minimizer up to time reparametrization. Equivalently, the generator $\mathcal{L}$ of the reference SDE, with absorbing boundary conditions on $\partial A \cup \partial B$, has smallest eigenvalue modulus $\delta := |\mu_1|$ satisfying $\delta = O(1)$ as $\epsilon \to 0$.*

Assumption 3.1 holds for landscapes with a single dominant transition channel; it fails when two or more channels compete with comparable action, in which case $\delta = O(e^{-c/\epsilon})$ and the bounds below become vacuous. The assumption enters these bounds, and the algorithm runs unchanged on landscapes with competing channels, which the Gaussian-mixture parameterization of Du et al. (2024) can represent with one component per channel.

Theorem 3.2 Committor initialization

Let $\tilde{q}$ be the stationary committor, i.e., the solution of $\mathcal{L}\tilde{q} = 0$ on $\Omega \setminus (\partial A \cup \partial B)$ with $\tilde{q}|_{\partial A} = 0$, $\tilde{q}|_{\partial B} = 1$. Suppose Assumption 3.1 holds with spectral gap $\delta > 0$, and that the transit time satisfies $T \geq C_* \delta^{-1} \log L$ for a sufficiently large constant C_*. Then the following hold.

(i) For every $t \in [0, T]$ the h-transform satisfies

$$\sup_{x\in\Omega} \big|h(x, t) - \tilde{q}(x)\big| \;\leq\; C_0\, e^{-\delta(T-t)}\,. \tag{22}$$

(ii) The committor-based control satisfies

$$\|\tilde{v} - v^*\|_{L^2(q^*)} \;\leq\; C_{\delta v}\, e^{-\delta T/2}\,. \tag{23}$$

(iii) The mean FPT satisfies

$$\left|\tau^{\tilde{v}}(A) - \tau^{v^*}(A)\right| \;\le\; C_\tau\, e^{-\delta T}\,. \tag{24}$$

(iv) The control $\tilde{v}$ lies in a basin of v^* in which $v \mapsto S^{\mathrm{DL}}_{[0,T]}[q^v, v]$ is locally convex with respect to the $L^2(q^*)$ norm. Gradient descent on (v, T) initialized at $(\tilde{v}, \tau^{\tilde{v}}(A))$ is stable and converges to (v^*, T^*).

The constants C_0, $C_{\delta v}$, $C_\tau > 0$ depend on δ, $\inf_{x\in\mathcal{S}} \tilde{q}(x)$, and Ω, not on T or ϵ.

The rate $\delta/2$ in (23) relative to (22) arises from converting a pointwise L^∞ estimate into a weighted $L^2(q^*)$ norm via a time integral; under Assumption 3.1, $\delta = O(1)$ and the condition on T reduces to $T \gtrsim \log L$. A learned neural committor (Rotskoff et al., 2022; Mitchell & Rotskoff, 2024) provides a realization of $\tilde{v}$ in high dimensions.

4 Improvements to the Doob's Lagrangian algorithm

Having established the equivalence between the Doob action and Freidlin–Wentzell large-deviation theory (Freidlin & Wentzell, 1984), we now identify two practical drawbacks of the original algorithm (Du et al., 2024) and propose targeted fixes. First, fixing the transit time T prevents true minimization of the FW action (7): when T is much shorter than the mean first-passage time, the path is forced across unphysically high energy barriers (Hänggi et al., 1990). Second, the assumption of a unique reaction channel breaks down in molecular systems with complex energy landscapes (Bolhuis et al., 2002); a naive linear initialization then relaxes into the nearest metastable configuration rather than a physically meaningful pathway.

We address both issues with two modifications to Du et al. (2024). We jointly optimize the Doob action (6) over (v, T) with a dual optimization scheme, alternating between updates to the path parameters and to T. For initialization, we use the committor-based method of Mitchell & Rotskoff (2024) to generate a physically reasonable reactive path, pre-train the generative model to match it, and set T to a value slightly larger than the length of that path before minimizing (6).

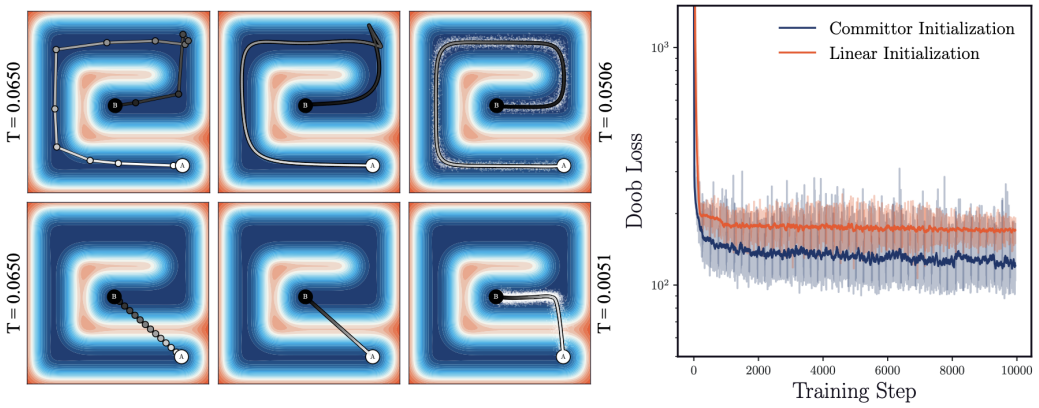


Figure 1: Comparison of a physically reasonable path ansatz with a naive linear initialization. With flexible T, the model exhibits unphysical barrier hopping in the linear case. Conversely, the model learns a more physically reasonable transition path and transit time when initialized with a reactive path generated by a committor-based method. This initialization also finds a lower minimum in the Doob action compared to the linear initialization.

4.1 Two-dimensional maze potential

Figure 1 shows the effect of these modifications on a two-dimensional maze potential. With committor-based initialization, the model learns a diffusive transition through the maze, recovering a transit time close to that of the initial path ansatz. With linear initialization, the path collapses into the nearest metastable state; to minimize (6), the optimizer shortens T and hops over an unphysically high energy barrier.

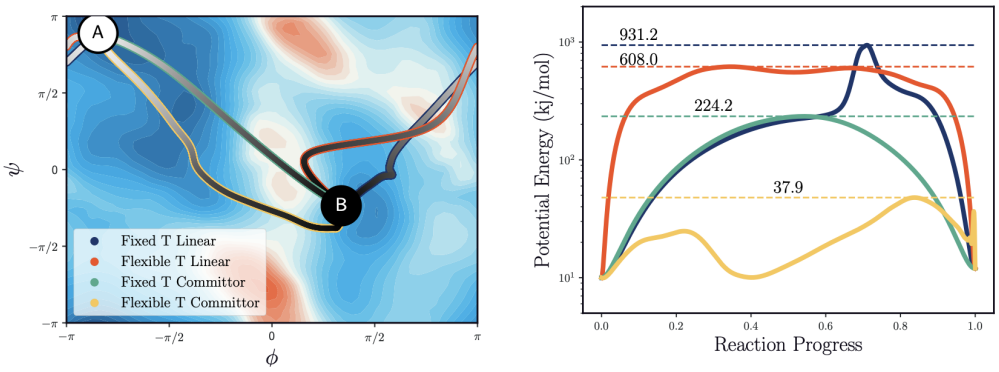


Figure 2: Performance of the Doob's Lagrangian algorithm in a first-order alanine dipeptide system. (Left) Mean paths generated by the algorithm in various regimes are shown in dihedral space. (Right) The potential energy profiles along each mean path are plotted on a common timescale, shifted to begin at an energy of 10 kJ/mol. The barrier height relative to the initial energy is shown for each method. Poor choices of T and poor initialization generate high-energy structures, while smart initialization and a flexible T generate paths with physically reasonable barrier height.

We observe that, while the learned transit time in the committor-initialized case is close to the time of the initial path, both times represent very fast transitions compared to unbiased simulations. We attribute this to the method's reliance on Gaussian time marginals, which prevent the path from significant backtracking during the transition, which is typical in diffusive systems like these. However, we note that because the path ansatz is a path sampled directly from simulation, the resulting transit time learned by the model is still physically reasonable, albeit rare.

4.2 Alanine dipeptide with first-order dynamics

We now demonstrate the utility of our improvements on molecular systems, beginning with overdamped dynamics for the transition between the $C7_{ax}$ and $C7_{eq}$ conformations of alanine dipeptide. We use the same initial and final states as Du et al. (2024) and operate in Cartesian coordinates. In the molecular systems we study, we focus only on the optimization of the mean paths, and thus apply the Doob's Lagrangian algorithm with zero noise, as we observed that in the case of nonzero noise and Cartesian coordinates, single-Gaussian paths quickly collapse to minimal variance.

The results for overdamped alanine dipeptide are shown in Figure 2. When initialized from a linear interpolation, a fixed-T approach ($T = 0.01$ ps) relaxes to the most proximal local minimizer of the Lagrangian, which in the rugged Cartesian space represents a very high-energy transition (roughly 1000 kJ/mol). This problem persists even when T is allowed to fluctuate (beginning at $T = 1$ ps), though the energy barrier estimate is slightly lower. Initializing the algorithm with committor-based paths generated from the method introduced in Mitchell & Rotskoff (2024) further improves performance; however, setting a fixed small T (also 0.01 ps) pulls the path out of its initial basin and into a high-energy path that continues to overestimate the energy barrier. Only when we initialize the algorithm with a committor-based path *and* allow T to relax from a relatively high initial time ($T = 1$ ps) do we recover a path that crosses a saddle point in free energy space and produces a reasonable estimate of the barrier height for the transition. The estimated 37.9 kJ/mol agrees well with the barrier height estimated by other free energy methods Šućur & Spiwok (2016).

The differential performance between the methods presented in Figure 2 is incredibly significant. Because the rate of transitions is exponential in the height of the energy barrier between the two states considered, an order-of-magnitude overestimation of barrier height results in a *many*-orders-of-magnitude underestimation of the molecular transition rate, depending on the temperature of the system. For example, at 300K, an overestimate of 10 kJ/mol in the barrier height corresponds to a rate underestimate by a factor of 50; a

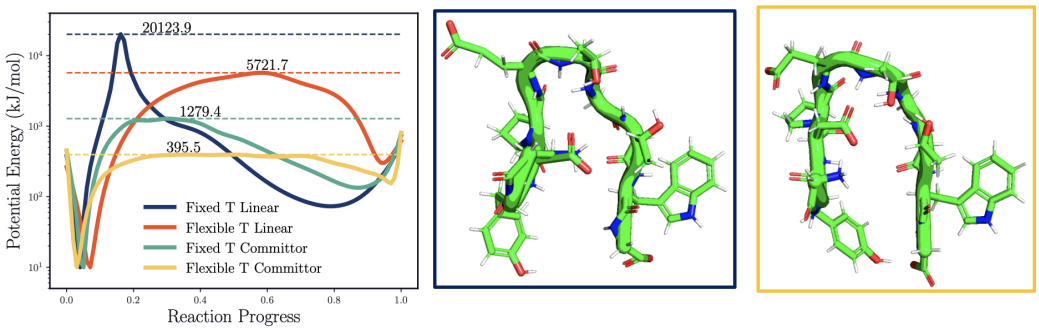


Figure 3: Performance of the Doob's Lagrangian algorithm in a second-order chignolin system. (Left) Energy profiles of the mean path learned by minimizing the objective in different regimes. Note that all methods are effective at minimizing high-energy initial structures. (Center) Fixing $T = 1$ ps with a linear interpolation generates a transition path where bonds overlap, leading to an extreme overestimate of the energy barrier. (Right) Initializing with a committor-based path and allowing T to fluctuate generates a much more reasonable transition state structure.

100 kJ/mol overestimate of barrier height corresponds to an error of roughly 18 orders of magnitude in the transition rate.

4.3 Chignolin with second-order dynamics

We also demonstrate dramatically improved performance in a more complicated molecular system: the unfolding of the miniprotein chignolin simulated under second-order dynamics. In order to facilitate the construction of committor-based paths, we ran our committor algorithm at a temperature of 500K and selected starting and ending configurations consistent with this temperature. Figure 3 demonstrates that when implemented with zero noise, the Lagrangian optimization quickly relaxes to energy-minimized states near both configurations regardless of implementation (as observed in Du et al. (2024)); we thus calculate the height of the energy barriers obtained from each method relative to the energy minimum along each path. Like alanine dipeptide, the fixed methods operated with $T = 1$ ps, while the flexible-T methods relaxed from an initial $T = 1$ ns.

Similar to alanine dipeptide, the combination of an optimization over T and a committor-based transition path dramatically lowers the energy of the transitions found by the Lagrangian algorithm. We further note that the barrier height found by the committor-based transition is on the same energetic scale as the initial and final conformations sampled from the equilibrium ensemble at 500K. A structural illustration of the superiority of the committor-based flexible-T approach is also illustrated in Figure 3. The transition state found by the linear interpolant involves the overlapping of chemical bonds near the N-terminus; an artifact of the naive interpolation in Cartesian coordinates. The committor-based transition state has no steric clashes and represents the concerted breaking of intra-peptide hydrogen bonds, a physically realistic mechanism.

5 Conclusion

We established a formal equivalence between Doob's Lagrangian and Freidlin-Wentzell large deviation theory, which motivates joint optimization over the path distribution and transit time T, and identifies a mechanism for extracting rate constants in single-barrier systems. We showed that committor-guided sampling methods simultaneously provide physically reasonable path initialization and reliable estimates of T. On benchmark peptide systems, these modifications dramatically improve the quality of sampled transition paths at negligible additional cost. The FW-DL correspondence opens several directions: refining the committor from DL-generated paths in a closed loop. This could be extended to multiple transition channels via Gaussian mixture formulations, and to estimate reaction rates in low-temperature systems directly from the quasipotential.

AI use statement

In this work, generative AI tools assisted with revising the manuscript, implementing portions of the codebase, searching and summarizing the literature, and formatting equations. Generative AI tools were not used for research ideation or experimental design. All theoretical results and their proofs were conceived and derived by the authors. Generative AI tools were used to double-check calculations for errors. The authors reviewed and edited all AI-assisted text, verified all cited references, and reviewed and tested all AI-assisted code. The authors take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Reproducibility statement

Complete proofs of all theoretical results not included in the main text, together with the assumptions they rely on, are given in Appendix A. Appendix B describes the network architectures and gives an overview of the algorithm. Committor estimation, simulation, and optimization parameters for the molecular systems are described in Appendix C. Source code that reproduces all experiments is included in the supplementary materials.

References

Gilles Adjanor and Manuel Athenes. Gibbs free-energy estimates from direct path-sampling computations. *The Journal of Chemical Physics*, 123(23), 2005.

Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *International Conference on Learning Representations*, 2023.

Rosalind J. Allen, Patrick B. Warren, and Pieter Rein ten Wolde. Sampling rare switching events in biochemical networks. *Physical Review Letters*, 94(1):018104, 2005.

Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, C. K. Luk, Bert Maher, Yunjie Pan, Christian Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Shunting Zhang, Michael Suo, Phil Tillet, Xu Zhao, Eikan Wang, Keren Zhou, Richard Zou, Xiaodong Wang, Ajit Mathews, William Wen, Gregory Chanan, Peng Wu, and Soumith Chintala. Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ASPLOS '24, pp. 929–947, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703850. doi: 10.1145/3620665.3640366. URL `https://doi.org/10.1145/3620665.3640366`.

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*, 53:291–318, 2002.

Freddy Bouchet and Julien Reygner. Generalisation of the eyring–kramers transition rate formula to irreversible diffusion processes. In *Annales Henri Poincaré*, volume 17, pp. 3499–3532. Springer, 2016.

David A. Case, Hasan Metin Aktulga, Kellon Belfon, David S. Cerutti, G. Andrés Cisneros, Vinícius Wilian D. Cruzeiro, Negin Forouzesh, Timothy J. Giese, Andreas W. Götz, Holger Gohlke, Saeed Izadi, Koushik Kasavajhala, Mehmet C. Kaymak, Edward King, Tom Kurtzman, Tai-Sung Lee, Pengfei Li, Jian Liu, Tyler Luchko, Ray Luo, Madushanka Manathunga, Matias R. Machado, Hai Minh Nguyen, Kurt A. O'Hearn, Alexey V. Onufriev, Feng Pan, Sergio Pantano, Ruxi Qi, Ali Rahnamoun, Ali Risheh, Stephan Schott-Verdugo, Akhil Shajan, Jason Swails, Junmei Wang, Haixin Wei, Xiongwu Wu, Yongxian Wu,

Shi Zhang, Shiji Zhao, Qiang Zhu, Thomas E. Cheatham III, Daniel R. Roe, Adrian Roitberg, Carlos Simmerling, Darrin M. York, Maria C. Nagan, and Kenneth M. Merz Jr. Ambertools. *Journal of Chemical Information and Modeling*, 63(20):6183–6191, 2023. doi: 10.1021/acs.jcim.3c01153. PMID: 37805934.

Michael G Crandall, Hitoshi Ishii, and Pierre-Louis Lions. User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1):1–67, 1992.

Avishek Das, Benjamin Kuznets-Speck, and David T Limmer. Direct evaluation of rare events in active matter from variational path sampling. *Physical Review Letters*, 128(2): 028005, 2022.

Martin V. Day. Recent progress on the small parameter exit problem. *Stochastics*, 20(2): 121–150, 1987.

Valentin de Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Pierre Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, New York, 2004.

Arnaud Doucet, Nando de Freitas, and Neil Gordon. An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo Methods in Practice*, pp. 3–14. Springer, New York, 2001.

Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Aspuru-Guzik, and Kirill Neklyudov. Doob's lagrangian: A sample-efficient variational approach to transition path sampling. *Advances in Neural Information Processing Systems*, 37:65791–65822, 2024.

Weinan E and Eric Vanden-Eijnden. Towards a theory of transition paths. *Journal of Statistical Physics*, 123(3):503–523, 2006.

Weinan E, Weiqing Ren, and Eric Vanden-Eijnden. Minimum action method for the study of rare events. *Communications on pure and applied mathematics*, 57(5):637–656, 2004.

Weinan E, Weiqing Ren, and Eric Vanden-Eijnden. Finite temperature string method for the study of rare events. *The Journal of Physical Chemistry B*, 109(14):6688–6693, 2005.

Peter Eastman, Raimondas Galvelis, Raúl P. Peláez, Charlles R. A. Abreu, Stephen E. Farr, Emilio Gallicchio, Anton Gorenko, Michael M. Henry, Frank Hu, Jing Huang, Andreas Krämer, Julien Michel, Joshua A. Mitchell, Vijay S. Pande, João PGLM Rodrigues, Jaime Rodriguez-Guerra, Andrew C. Simmonett, Sukrit Singh, Jason Swails, Philip Turner, Yuanqing Wang, Ivy Zhang, John D. Chodera, Gianni De Fabritiis, and Thomas E. Markland. Openmm 8: Molecular dynamics simulation with machine learning potentials. *The Journal of Physical Chemistry B*, 128(1):109–116, 2024. doi: 10.1021/acs.jpcb.3c06662. PMID: 38154096.

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

Wendell H Fleming and H Mete Soner. *Controlled Markov processes and viscosity solutions*. Springer, 2006.

Mark I. Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, New York, 1984.

David. Gilbarg and Neil S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics, 224. Springer Berlin Heidelberg, Berlin, Heidelberg, 2nd ed. 2001. edition, 2001. ISBN 3-642-61798-0.

Igor V. Girsanov. On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory of Probability & Its Applications*, 5(3):285–301, 1960.

Tobias Grafke, Tobias Schäfer, and Eric Vanden-Eijnden. Long term effects of small random perturbations on dynamical systems: Theoretical and computational tools. In *Recent progress and modern challenges in applied mathematics, modeling and computational science*, pp. 17–55. Springer, 2017.

Peter Hänggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after Kramers. *Reviews of Modern Physics*, 62(2):251–341, 1990.

Matthias Heymann and Eric Vanden-Eijnden. The geometric minimum action method: A least action principle on the space of curves. *Communications on Pure and Applied Mathematics*, 61(8):1052–1117, 2008.

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. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 79540–79556. Curran Associates, Inc., 2023. doi: 10.52202/075280-3481. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/fb7f55f36c53247a704792a721272706-Paper-Conference.pdf.

Scott A Hollingsworth and Ron O Dror. Molecular dynamics simulation for all. *Neuron*, 99 (6):1129–1143, 2018.

Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99(20):12562–12566, 2002.

Jungyoon Lee, Michael Plainer, Yuanqi Du, Lars Holdijk, Rob Brekelmans, Dominique Beaini, and Kirill Neklyudov. Scaling deep learning solutions for transition path sampling. In *ICLR 2025 Workshop on Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Robert T. McGibbon, Kyle A. Beauchamp, Matthew P. Harrigan, Christoph Klein, Jason M. Swails, Carlos X. Hernández, Christian R. Schwantes, Lee-Ping Wang, Thomas J. Lane, and Vijay S. Pande. Mdtraj: A modern open library for the analysis of molecular dynamics trajectories. *Biophysical Journal*, 109(8):1528–1532, 2015. ISSN 0006-3495. doi: https://doi.org/10.1016/j.bpj.2015.08.015. URL https://www.sciencedirect.com/science/article/pii/S0006349515008267.

Andrew R Mitchell and Grant M Rotskoff. Committor guided estimates of molecular transition rates. *Journal of Chemical Theory and Computation*, 20(21):9378–9393, 2024.

Klaus Müller. Reaction paths on multidimensional energy hypersurfaces. *Angewandte Chemie International Edition in English*, 19(1):1–13, 1980.

B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Universitext. Springer Berlin Heidelberg, 2010. ISBN 9783642143946. URL https://books.google.it/books?id=EQZEAAAAQBAJ.

GiovanniMaria Piccini, Mal-Soon Lee, Simuck F Yuk, Difan Zhang, Greg Collinge, Loukas Kollias, Manh-Thuong Nguyen, Vassiliki-Alexandra Glezakou, and Roger Rousseau. Ab initio molecular dynamics with enhanced sampling in heterogeneous catalysis. *Catalysis Science & Technology*, 12(1):12–37, 2022.

Ross G Pinsky. *Positive harmonic functions and diffusion*, volume 45. Cambridge University Press, 1995.

Sanjeev Raja, Martin Sipka, Michael Psenka, Tobias Kreiman, Michal Pavelka, and Aditi S Krishnapriyan. Action-minimization meets generative modeling: Efficient transition path sampling with the onsager-machlup functional. In *International Conference on Machine Learning*, pp. 50972–51008. PMLR, 2025.

Grant M Rotskoff, Andrew R Mitchell, and Eric Vanden-Eijnden. Active importance sampling for variational objectives dominated by rare events: Consequences for optimization and generalization. In *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, pp. 757–780. PMLR, 2022.

Schrödinger, LLC. The PyMOL molecular graphics system, version 1.8. November 2015.

Diwakar Shukla, Yilin Meng, Benoît Roux, and Vijay S Pande. Activation pathway of src kinase reveals intermediate states as targets for drug design. *Nature communications*, 5 (1):3397, 2014.

Aditya N Singh and David T Limmer. Variational deep learning of equilibrium transition path ensembles. *The Journal of Chemical Physics*, 159(2), 2023.

Aditya N Singh and David T Limmer. Splitting probabilities as optimal controllers of rare reactive events. *The Journal of Chemical Physics*, 161(5), 2024.

Aditya N Singh, Avishek Das, and David T Limmer. Variational path sampling of rare dynamical events. *Annual Review of Physical Chemistry*, 76, 2025.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Glenn M. Torrie and John P. Valleau. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics*, 23(2): 187–199, 1977.

Fredi Tröltzsch. *Optimal control of partial differential equations: theory, methods, and applications*, volume 112. American Mathematical Soc., 2010.

Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. In *The Eleventh International Conference on Learning Representations*, 2023. URL `https://openreview.net/forum?id=8pvnfTAbu1f`.

Jiawei Yan, Hugo Touchette, and Grant M Rotskoff. Learning nonequilibrium control forces to characterize dynamical phase transitions. *Physical Review E*, 105(2):024115, 2022.

Zoran Šućur and Vojtěch Spiwok. Sampling enhancement and free energy prediction by the flying gaussian method. *Journal of Chemical Theory and Computation*, 12(9):4644–4650, 2016. doi: 10.1021/acs.jctc.6b00551.

A Theoretical Supplements

A.1 Proof of Proposition 3.1 (KL–Doob equivalence)

Proposition A.1 (KL–Doob action equivalence)**.** *Let $\mathbb{P}$ and $\mathbb{Q}^v$ denote the path measures associated with* (4) *and* (5) *respectively, and let q_t^v be the marginal density of $\mathbb{Q}^v$ at time t. Then*

$$D_{\mathrm{KL}}(\mathbb{Q}^v \| \mathbb{P}) = S^{\mathrm{DL}}_{[0,T]}[q^v, v]\,, \tag{25}$$

where $S^{\mathrm{DL}}_{[0,T]}[\,\cdot\,,\,\cdot\,]$ is the stochastic Doob action defined in (6)*.*

Proof. Since $\mathbb{P}$ and $\mathbb{Q}^v$ share the same diffusion Ξ_t and their drifts differ by $2G_t v_t = \Xi_t \Xi_t^\top v_t$, Girsanov's theorem (Girsanov, 1960; Øksendal, 2010) gives the Radon–Nikodým derivative

$$\frac{d\mathbb{Q}^v}{d\mathbb{P}} = \exp\!\left(-\int_0^T \langle G_t v_t(X_t),\, v_t(X_t)\rangle\, dt + \int_0^T \langle \Xi_t^\top v_t(X_t),\, dW_t\rangle \right). \tag{26}$$

The KL divergence is the expectation of the log-Radon–Nikodým derivative under $\mathbb{Q}^v$:

$$D_{\mathrm{KL}}(\mathbb{Q}^v \| \mathbb{P}) = \mathbb{E}_{\mathbb{Q}^v}\left[-\int_0^T \langle G_t v_t,\, v_t\rangle\, dt + \int_0^T \langle \Xi_t^\top v_t,\, dW_t\rangle \right]. \tag{27}$$

The stochastic integral above is not a martingale under $\mathbb{Q}^v$. We perform a change of measure: under $\mathbb{Q}^v$ the Brownian motion satisfies $dW_t = dW_t^{\mathbb{Q}} + \Xi_t^\top v_t\, dt$, so

$$\mathbb{E}_{\mathbb{Q}^v}\left[\int_0^T \langle \Xi_t^\top v_t,\, dW_t\rangle\right] = \underbrace{\mathbb{E}_{\mathbb{Q}^v}\left[\int_0^T \langle \Xi_t^\top v_t,\, dW_t^{\mathbb{Q}}\rangle\right]}_{=0} + 2\,\mathbb{E}_{\mathbb{Q}^v}\left[\int_0^T \langle G_t v_t,\, v_t\rangle\, dt\right], \tag{28}$$

where the first term vanishes since $W_t^{\mathbb{Q}}$ is a $\mathbb{Q}^v$-martingale, and $|\Xi_t^\top v_t|^2 = 2\langle G_t v_t, v_t\rangle$. Combining both terms yields

$$D_{\mathrm{KL}}(\mathbb{Q}^v \| \mathbb{P}) = \mathbb{E}_{\mathbb{Q}^v}\left[\int_0^T \langle G_t v_t(X_t),\, v_t(X_t)\rangle\, dt\right] = S^{\mathrm{DL}}_{[0,T]}[q^v, v]\,, \tag{29}$$

which is the claimed identity. □

A.2 Proof of Lemma 3.1 (Action scaling)

Lemma A.1 (Action scaling)**.** *For any admissible control field v and its associated characteristic path φ^v, the FW action and the Doob action on the characteristic path are related by*

$$S_{[0,T],A}(\varphi^v) = \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^v)\,. \tag{30}$$

Proof. Along the characteristic path (10), the velocity deviation from the drift is $\dot\varphi_t^v - b(\varphi_t^v) = 2G_t v(\varphi_t^v)$. Substituting into the FW action (2) and using the identity $a(\varphi_t^v)^{-1} = (G_t/\epsilon)^{-1} = \epsilon G_t^{-1}$, which follows from $G_t = \epsilon a(\varphi_t^v)$:

$$\begin{aligned} S_{[0,T],A}(\varphi^v) &= \frac{1}{4}\int_0^T \left\langle 2G_t v(\varphi_t^v),\ \epsilon G_t^{-1}\, 2G_t v(\varphi_t^v)\right\rangle dt \\ &= \epsilon \int_0^T \left\langle v(\varphi_t^v),\ G_t\, v(\varphi_t^v)\right\rangle dt = \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^v)\,. \end{aligned}$$

□

A.3 Proof of Lemma 3.2 (FW value function and HJB equation)

Lemma A.2 (FW value function and HJB equation)**.** *Define the FW value function*

$$I(t,y) = \inf_{\substack{\phi\in H^1([t,T])\\ \phi_t=y,\ \phi_T=x_{\text{target}}}} S_{[t,T],y}(\phi)\,. \tag{31}$$

Assuming $I \in C^1([0,T)\times\mathbb{R}^d)$, it is the unique classical solution of the Hamilton–Jacobi–Bellman equation

$$\frac{\partial I}{\partial t} + \langle b(y), \nabla_y I\rangle - \langle \nabla_y I,\ a(y)\,\nabla_y I\rangle = 0\,, \tag{32}$$

with terminal condition $I(T,\, x_{\text{target}}) = 0$.

Proof. The derivation follows from Bellman's dynamic programming principle (Fleming & Soner, 2006). For any $\Delta t > 0$, optimality implies

$$I(t,y) = \inf_u \{\ell(y,u)\,\Delta t + I(t+\Delta t,\ y+u\Delta t)\} + o(\Delta t)\,, \tag{33}$$

where $\ell(y,u) = \frac{1}{4}\langle u-b,\ a^{-1}(u-b)\rangle$ is the FW Lagrangian and $u = \dot\phi$ is the control velocity. Taylor expanding $I(t+\Delta t, y+u\Delta t)$, cancelling $I(t,y)$, dividing by Δt, and taking $\Delta t \to 0$ gives

$$-\partial_t I = \inf_u \big\{\ell(y,u) + \langle \nabla_y I,\ u\rangle\big\}\,. \tag{34}$$

The right-hand side is the negative Legendre transform of ℓ evaluated at $-\nabla_y I$. The Hamiltonian $H(y,p) = \sup_u\{\langle p,u\rangle - \ell(y,u)\}$ is computed by setting $\partial_u[\langle p,u\rangle - \ell] = 0$, which gives the optimal velocity $u^* = b(y) + 2a(y)p$ and the explicit form $H(y,p) = \langle p, b(y)\rangle + \langle p, a(y)p\rangle$. Substituting $p = -\nabla_y I$ yields (32). □

A.4 Proof of Lemma 3.3 (Doob energy and vanishing-viscosity limit)

Lemma A.3 (Doob energy and vanishing-viscosity limit)**.** *Define the* Doob energy *via the Cole–Hopf transformation*

$$W^\epsilon(t,y) = -\epsilon \log h(y,t)\,, \qquad h(y,t) = e^{-W^\epsilon(t,y)/\epsilon}\,. \tag{35}$$

Then W^ϵ satisfies the viscous Hamilton–Jacobi equation

$$\frac{\partial W^\epsilon}{\partial t} + \langle b(y), \nabla W^\epsilon\rangle - \langle \nabla W^\epsilon,\ a(y)\,\nabla W^\epsilon\rangle + \epsilon\,\mathrm{tr}\big(a(y)\,\nabla^2 W^\epsilon\big) = 0\,. \tag{36}$$

In the limit $\epsilon \to 0$, assuming pointwise convergence $W^\epsilon \to W^0$, the viscosity term vanishes and W^0 satisfies

$$\frac{\partial W^0}{\partial t} + \langle b(y), \nabla W^0\rangle - \langle \nabla W^0,\ a(y)\,\nabla W^0\rangle = 0\,. \tag{37}$$

Proof. Substitute (35) into the backward Kolmogorov equation (9). Computing $\partial_t h = -\epsilon^{-1} h\,\partial_t W^\epsilon$, $\nabla h = -\epsilon^{-1} h\,\nabla W^\epsilon$, and $\nabla^2 h = \epsilon^{-1} h\big[\epsilon^{-1}(\nabla W^\epsilon)(\nabla W^\epsilon)^\top - \nabla^2 W^\epsilon\big]$, then substituting into (9) and dividing through by $-h/\epsilon$ gives (36). Using $G_t = \epsilon a(y)$ simplifies the coefficient of the second-order term to $\epsilon\,\mathrm{tr}(a\nabla^2 W^\epsilon)$, which vanishes as $\epsilon \to 0$, leaving (37). □

A.5 Proof of Theorem 3.1 (Optimality equivalence)

Theorem A.1 Optimality equivalence

Since $W^\epsilon(T, x_{\text{target}}) = 0$, the optimal Doob action satisfies

$$\epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = W^\epsilon(0,A)\,, \tag{38}$$

and in the vanishing-noise limit it converges to the FW value function,

$$\lim_{\epsilon\to 0} \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = I(0,A)\,. \tag{39}$$

Therefore, minimizing the Doob action over (v,T) is asymptotically equivalent to computing the quasipotential,

$$\lim_{\epsilon\to 0} \inf_{T>0} \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = V(A, x_{\mathrm{target}})\,. \tag{40}$$

Proof. By Lemmas 3.2 and 3.3, the functions $I(t,y)$ and $W^0(t,y)$ satisfy the identical first-order Hamilton–Jacobi PDE (32)=(37), and both satisfy the terminal condition $f(T, x_{\mathrm{target}}) = 0$ (with $+\infty$ elsewhere on the boundary). By the uniqueness of viscosity solutions to first-order Hamilton–Jacobi equations (Crandall et al., 1992; Fleming & Soner, 2006), $W^0 = I$ everywhere.

It remains to verify the integral representation (38). Setting $v_t^*(x) = \nabla_x \log h = -\epsilon^{-1}\nabla_x W^\epsilon$ and evaluating the total time derivative of W^ϵ along φ^{v^*}:

$$\begin{aligned}\frac{d}{dt}W^\epsilon(t,\varphi_t^{v^*}) &= \partial_t W^\epsilon + \langle \nabla W^\epsilon,\, \dot\varphi^{v^*}\rangle \\ &= \big(\langle \nabla W^\epsilon,\, a\nabla W^\epsilon\rangle - \langle b,\, \nabla W^\epsilon\rangle\big) + \langle \nabla W^\epsilon,\, b - 2a\nabla W^\epsilon\rangle \\ &= -\langle \nabla W^\epsilon,\, a(y)\,\nabla W^\epsilon\rangle\,,\end{aligned} \tag{41}$$

where we used the HJB relation $\partial_t W^\epsilon = \langle \nabla W^\epsilon, a\nabla W^\epsilon\rangle - \langle b, \nabla W^\epsilon\rangle$ (from (15) at leading order) and the optimal dynamics (10). Since the DL action integrand equals $\langle \nabla W^\epsilon, a\nabla W^\epsilon\rangle = -\frac{d}{dt}W^\epsilon(t,\varphi_t^{v^*})$, integrating from 0 to T and applying the terminal condition $W^\epsilon(T,x) = 0$ gives the identity (38). Taking $\epsilon \to 0$ and optimizing over T recovers the quasipotential by definition. □

A.6 Proof of Proposition 3.2 (Saddle-point property)

Proposition A.2 (Saddle-point property)**.** *Let A and B be two stable equilibria of the drift b, and let $x_\star$ be the saddle point on the basin boundary separating their respective basins of attraction. Then*

$$V(A,B) = V(A,x_\star)\,. \tag{42}$$

Proof. Any continuous path connecting A to B must cross the separatrix; the instanton ϕ^* does so at some finite time T_{cross}. Decompose ϕ^* into an uphill fluctuation segment ϕ^*_{up} on $[0, T_{\mathrm{cross}}]$ and a downhill relaxation segment ϕ^*_{down} on $[T_{\mathrm{cross}}, +\infty)$:

$$V(A,B) = S_{[0,T_{\mathrm{cross}}],A}(\phi^*_{\mathrm{up}}) + S_{[T_{\mathrm{cross}},\infty),x_\star}(\phi^*_{\mathrm{down}}) = V(A,x_\star) + S_{[0,\infty),x_\star}(\phi^*_{\mathrm{down}})\,. \tag{43}$$

The downhill segment $\psi := \phi^*_{\mathrm{down}}$ satisfies $\dot\psi = b(\psi)$ exactly, so $\dot\psi - b(\psi) \equiv 0$ and the FW action integrand vanishes identically. Hence $S_{[0,\infty),x_\star}(\phi^*_{\mathrm{down}}) = 0$. Since $S \geq 0$ by definition, this value is also the infimum over all paths from $x_\star$ to B, giving $\inf_\phi S_{[0,\infty),x_\star}(\phi) = 0$. Substituting back yields $V(A,B) = V(A,x_\star)$. □

A.7 Proof of Theorem 3.2 (Committor initialization)

Assumption A.1 (No path metastability)**.** *The transition from A to B is mediated by a unique dominant instanton, i.e. the FW action has a unique minimizer up to time reparametrization. Equivalently, the generator $\mathcal{L}$ of the reference SDE, with absorbing boundary conditions on $\partial A \cup \partial B$, has smallest eigenvalue modulus $\delta := |\mu_1|$ satisfying $\delta = O(1)$ as $\epsilon \to 0$.*

We work throughout under Assumption 3.1 (no path metastability, spectral gap $\delta = \mu_1 = O(1)$) and the condition $T \geq C_*\delta^{-1}\log L$. We use the notation of Section 3.6: $\tilde q$ is the

stationary committor, $\tilde{v}(x) = \nabla_x \log \tilde{q}(x)$, $v_t^*(x) = \nabla_x \log h(x,t)$, $q_t^* = q_t^{v^*}$ is the time marginal of $\mathbb{Q}^{v^*}$, $\mathcal{S} := \bigcup_{t\in(0,T]} \mathrm{supp}(q_t^*)$, and $\tilde{q}_{\min} := \inf_{x\in\mathcal{S}} \tilde{q}(x)$.

Theorem A.2 Committor initialization

Let $\tilde{q}$ be the stationary committor, i.e., the solution of $\mathcal{L}\tilde{q} = 0$ on $\Omega \setminus (\partial A \cup \partial B)$ with $\tilde{q}|_{\partial A} = 0$, $\tilde{q}|_{\partial B} = 1$. Suppose Assumption 3.1 holds with spectral gap $\delta > 0$, and that the transit time satisfies $T \geq C_* \delta^{-1} \log L$ for a sufficiently large constant C_*. Then the following hold.

(i) For every $t \in [0,T]$ the h-transform satisfies

$$\sup_{x\in\Omega} \big|h(x,t) - \tilde{q}(x)\big| \;\leq\; C_0\, e^{-\delta(T-t)}\,. \tag{44}$$

(ii) The committor-based control satisfies

$$\|\tilde{v} - v^*\|_{L^2(q^*)} \;\leq\; C_{\delta v}\, e^{-\delta T/2}\,. \tag{45}$$

(iii) The mean first-passage time satisfies

$$\big|\tau^{\tilde{v}}(A) - \tau^{v^*}(A)\big| \;\leq\; C_\tau\, e^{-\delta T}\,. \tag{46}$$

(iv) The control $\tilde{v}$ lies in a basin of v^* in which $v \mapsto S^{\mathrm{DL}}_{[0,T]}[q^v, v]$ is locally convex with respect to the $L^2(q^*)$ norm. Gradient descent on (v,T) initialized at $(\tilde{v}, \tau^{\tilde{v}}(A))$ is stable and converges to (v^*, T^*), where $T^* := \tau^{v^*}(A)$.

The constants C_0, $C_{\delta v}$, $C_\tau > 0$ depend on δ, $\inf_{x\in\mathcal{S}} \tilde{q}(x)$, and Ω, but not on T or ϵ.

1. Spectral bound on $h(x,t) - \tilde{q}(x)$

We aim to show that, for every $t \in [0,T]$, $\sup_x |h(x,t) - \tilde{q}(x)| \leq C_0 e^{-\delta(T-t)}$ with rate $\delta = O(1)$.

Change of variables to time-to-go. The function h solves the backward Kolmogorov equation (BKE)

$$\partial_t h + \mathcal{L}h = 0 \quad \text{on } \Omega \times [0,T), \qquad h(x,T) = \mathbf{1}_B(x), \tag{47}$$

with absorbing boundary conditions $h|_{\partial A} = 0$, $h|_{\partial B} = 1$. Introduce the time-to-go $s := T - t \in [0,T]$ and set $H(x,s) := h(x, T-s)$. Differentiating through the definition and substituting the BKE, which states $\partial_t h = -\mathcal{L}h$:

$$\partial_s H(x,s) = -\partial_t h\big|_{t=T-s} = \mathcal{L}H(x,s).$$

The initial condition in s is $H(x,0) = \mathbf{1}_B(x)$, giving the forward Cauchy problem

$$\partial_s H = \mathcal{L}H, \qquad H(x,0) = \mathbf{1}_B(x), \qquad H\big|_{\partial A\cup\partial B} = \text{(Dirichlet)}. \tag{48}$$

Decomposition around $\tilde{q}$. The stationary committor satisfies $\mathcal{L}\tilde{q} = 0$ with $\tilde{q}|_{\partial A} = 0$, $\tilde{q}|_{\partial B} = 1$. Write $H(x,s) = \tilde{q}(x) + \Phi(x,s)$. Substituting into (48) and using $\mathcal{L}\tilde{q} = 0$ gives $\partial_s \Phi = \mathcal{L}\Phi$. The boundary values of H and $\tilde{q}$ agree on $\partial A \cup \partial B$, so Φ satisfies homogeneous Dirichlet conditions. The initial condition is $\Phi(x,0) = \mathbf{1}_B(x) - \tilde{q}(x)$. Thus:

$$\partial_s \Phi = \mathcal{L}\Phi, \qquad \Phi\big|_{\partial A\cup\partial B} = 0, \qquad \Phi(x,0) = \mathbf{1}_B(x) - \tilde{q}(x). \tag{49}$$

Eigenfunction expansion. Under uniform ellipticity of a on the compact domain Ω, the operator $\mathcal{L}$ with homogeneous Dirichlet conditions has discrete spectrum with strictly negative eigenvalues (there is no zero eigenvalue because the killed process loses mass through absorption). Write

$$\mathcal{L}\psi_k = -\mu_k \psi_k, \quad k = 1,2,\ldots, \quad 0 < \mu_1 \leq \mu_2 \leq \cdots \nearrow \infty.$$

Expanding the initial condition in this basis, $\Phi(x,0) = \sum_{k=1}^{\infty} c_k \psi_k(x)$ with $c_k = \langle \mathbf{1}_B - \tilde{q}, \psi_k \rangle_\pi$, the unique solution to (49) is

$$\Phi(x,s) = \sum_{k=1}^{\infty} c_k \, \psi_k(x) \, e^{-\mu_k s}. \tag{50}$$

One verifies directly that $\partial_s \Phi = \sum_k (-\mu_k) c_k \psi_k e^{-\mu_k s} = \mathcal{L}\Phi$ and that the boundary conditions and initial condition are satisfied. Reverting to (x,t) via $s = T - t$:

$$h(x,t) = \tilde{q}(x) + \sum_{k=1}^{\infty} c_k \, \psi_k(x) \, e^{-\mu_k (T-t)}. \tag{51}$$

Pointwise bound. Taking absolute values in (51) and factoring out the slowest-decaying mode, using $\mu_k \geq \mu_1 = \delta$ for all k:

$$|h(x,t) - \tilde{q}(x)| \leq e^{-\delta(T-t)} \sum_{k=1}^{\infty} |c_k| \, \|\psi_k\|_{L^\infty} =: C_0 \, e^{-\delta(T-t)}.$$

The constant $C_0 := \sum_k |c_k| \, \|\psi_k\|_{L^\infty}$ is finite: each ψ_k is in $C^1(\Omega)$ by Schauder elliptic regularity (Gilbarg & Trudinger, 2001) applied to $\mathcal{L}\psi_k + \mu_k \psi_k = 0$, so $\|\psi_k\|_{L^\infty} < \infty$, and the coefficients c_k decay fast enough (controlled by the L^2 regularity of the initial data and the Sobolev embedding $W^{2,p} \hookrightarrow L^\infty$ for $p > d/2$) for the sum to converge. By the same Schauder estimates applied to the gradient:

$$|\nabla h(x,t) - \nabla \tilde{q}(x)| \leq C_1 \, e^{-\delta(T-t)}, \tag{52}$$

for a constant C_1 with the same dependencies as C_0. This establishes conclusion (i).

Assumption 3.1 ensures $\delta = \mu_1 = O(1)$. Under path metastability, two or more transition channels produce nearly degenerate leading eigenvalues and $\mu_1 = O(e^{-c/\epsilon})$, making $C_0 e^{-\delta(T-t)} \approx C_0$ for $T = O(1)$ and the bound vacuous.

2. Bounding $\|\delta v\|_{L^2(q^*)}$

In this stage we aim to translate the spectral bound of Stage 1 into the control perturbation bound.

Explicit expression for δv. Let $R(x) := h(x,0) - \tilde{q}(x)$. By (51) at $t = 0$ and the bounds of Stage 1:

$$\|R\|_{L^\infty} \leq C_0 \, e^{-\delta T}, \qquad \|\nabla R\|_{L^\infty} \leq C_1 \, e^{-\delta T}. \tag{53}$$

The control perturbation at $t = 0$ is

$$\delta v(x) := \tilde{v}(x) - v_0^*(x) = \nabla \log \tilde{q}(x) - \nabla \log(\tilde{q}(x) + R(x)) = -\nabla \log\left(1 + \frac{R(x)}{\tilde{q}(x)}\right). \tag{54}$$

Because the bound $|h(x,t) - \tilde{q}(x)| \leq C_0 e^{-\delta(T-t)}$ holds for all $t \in [0,T]$, using $\tilde{v}$ as a time-independent control incurs only an additional $O(e^{-\delta T/2})$ error uniformly in t, which is of the same order as the bound we derive below. Thus δv is an admissible additive perturbation of v^* on the full interval $[0,T]$.

Positivity of $\tilde{q}$ on $\mathcal{S}$. We need to bound the relative error $\rho := R/\tilde{q}$, which requires $\tilde{q}(x) > 0$ on the support $\mathcal{S}$ of q^*. We use the probabilistic interpretation of the committor directly. For any $x \in \mathcal{S}$, there exists $t > 0$ such that $q_t^*(x) > 0$, meaning the optimally controlled process has positive probability of passing through x at time t. By definition, the controlled process started at such x reaches B in finite time (it is on the instanton, the unique dominant channel under Assumption 3.1). The committor $\tilde{q}(x)$ is the probability of reaching B before A from x under the *reference* process; since the reference dynamics also connect x to B with positive probability (the controlled process shows a path exists), we have $\tilde{q}(x) > 0$. The zero set of $\tilde{q}$ is exactly ∂A, which is disjoint from $\mathcal{S}$ since the instanton leaves A immediately for $t > 0$. Therefore $\tilde{q}_{\min} := \inf_{x \in \mathcal{S}} \tilde{q}(x) > 0$.

L^∞ bound on δv. Writing $\delta v = -\nabla\rho/(1+\rho)$ and computing $\nabla\rho$ via the quotient rule:

$$\nabla\rho = \frac{\tilde{q}\,\nabla R - R\,\nabla\tilde{q}}{\tilde{q}^2}.$$

From (53) and $\tilde{q} \geq \tilde{q}_{\min}$ on $\mathcal{S}$:

$$\|\nabla\rho\|_{L^\infty(\mathcal{S})} \leq \left(\frac{C_1}{\tilde{q}_{\min}} + \frac{C_0\|\nabla\tilde{q}\|_{L^\infty}}{\tilde{q}_{\min}^2}\right) e^{-\delta T} =: C_\rho\, e^{-\delta T}. \tag{55}$$

Define $\eta := C_0 e^{-\delta T}/\tilde{q}_{\min}$. For $T > \delta^{-1}\log(C_0/\tilde{q}_{\min})$ we have $\eta < 1$, so $1+\rho(x) \geq 1-\eta > 0$ uniformly on $\mathcal{S}$. Therefore:

$$\|\delta v\|_{L^\infty(\mathcal{S})} = \left\|\frac{\nabla\rho}{1+\rho}\right\|_{L^\infty(\mathcal{S})} \leq \frac{C_\rho}{1-\eta}\, e^{-\delta T}. \tag{56}$$

Since $\int_\Omega q_t^*\,dx = 1$ for all t and q_t^* is supported on $\mathcal{S}$:

$$\|\delta v\|^2_{L^2(q^*)} \leq \int_0^T \|\delta v\|^2_{L^\infty(\mathcal{S})}\,dt \leq T\left(\frac{C_\rho}{1-\eta}\right)^2 e^{-2\delta T}.$$

Taking square roots gives $\|\delta v\|_{L^2(q^*)} \leq \frac{C_\rho}{1-\eta}\sqrt{T}\,e^{-\delta T}$. The function $T \mapsto \sqrt{T}\,e^{-\delta T/2}$ attains its maximum at $T = 1/\delta$, so that $\sup_{T>0}\sqrt{T}\,e^{-\delta T/2} = C(\delta) := \delta^{-1/2}e^{-1/2}$. For any $T > 0$ we can write $\sqrt{T}\,e^{-\delta T} = (\sqrt{T}\,e^{-\delta T/2})\cdot e^{-\delta T/2} \leq C(\delta)\,e^{-\delta T/2}$, giving a T-independent prefactor at the cost of halving the rate:

$$\|\delta v\|_{L^2(q^*)} \leq \underbrace{\frac{C_\rho\,C(\delta)}{1-\eta}}_{=:C_{\delta v}}\, e^{-\delta T/2}. \tag{57}$$

The constant $C_{\delta v}$ depends on δ, $\tilde{q}_{\min}$, and Ω, but not on T or ϵ. This establishes conclusion (ii).

3. Sensitivity of the mean first-passage time

In this stage we prove conclusion (iii) by bounding $|\tau^{\tilde{v}}(A) - \tau^{v^*}(A)|$ in terms of $\|\delta v\|_{L^\infty}$.

Mean first-passage time as a PDE solution. For any admissible v, the mean first-passage time $\tau^v(x) := \mathbb{E}_x^v[\tau_{\partial B}]$ solves the Dirichlet problem (Øksendal, 2010)

$$\mathcal{L}^v\tau^v(x) = -1 \quad \text{on } \Omega\setminus\partial B, \qquad \tau^v\big|_{\partial B} = 0, \tag{58}$$

where $\mathcal{L}^v = \langle b + 2Gv, \nabla\rangle + \epsilon\,\mathrm{tr}(a\nabla^2)$ is the controlled generator. Under uniform ellipticity of $\mathcal{L}^v$, a unique classical solution $\tau^v \in C^2(\Omega)\cap C(\bar{\Omega})$ exists by the Fredholm alternative (Gilbarg & Trudinger, 2001). The control $v_t^* = \nabla\log h(\cdot,t)$ depends on time. For a time-dependent control, the expected remaining time $u(x,t)$ solves $\partial_t u + \mathcal{L}_t^v u = -1$ with $u|_{\partial B} = 0$, and the argument below applies with the Green's function replaced by the transition density of the process killed at ∂B, integrated over y and $s \geq t$. We present the time-homogeneous case to simplify notation.

Derivative of $v \mapsto \tau^v(A)$. Write $v_\alpha := v^* + \alpha\,\delta v$ and $\dot{\tau} := \partial_\alpha \tau^{v_\alpha}|_{\alpha=0}$. Differentiating (58) with respect to α at $\alpha = 0$, noting that $\partial_\alpha\mathcal{L}^{v_\alpha}|_{\alpha=0} = 2G\,\delta v\cdot\nabla$, yields:

$$\mathcal{L}^{v^*}\dot{\tau}(x) = -2\langle G\,\delta v(x), \nabla\tau^{v^*}(x)\rangle \quad \text{on } \Omega\setminus\partial B, \qquad \dot{\tau}\big|_{\partial B} = 0. \tag{59}$$

Green's function representation. Let $\mathcal{G}_{v^*}(x,y) \geq 0$ be the Green's function of $-\mathcal{L}^{v^*}$ (Pinsky, 1995; Øksendal, 2010) on Ω with Dirichlet conditions at ∂B (non-negative by the maximum principle). The solution to (59) is

$$\dot{\tau}(x) = 2\int_\Omega \mathcal{G}_{v^*}(x,y)\,\langle G\,\delta v(y), \nabla\tau^{v^*}(y)\rangle\,dy,$$

so at $x = A$:

$$\dot{\tau}(A) = 2\int_\Omega \mathcal{G}_{v^*}(A,y)\,\langle G\,\delta v(y), \nabla\tau^{v^*}(y)\rangle\,dy. \tag{60}$$

Bound on $|\dot\tau(A)|$. *Triangle inequality.* Since $\mathcal{G}_{v^*}(A,y) \geq 0$: $|\dot\tau(A)| \leq 2\int_\Omega \mathcal{G}_{v^*}(A,y)\,|\langle G\,\delta v(y), \nabla\tau^{v^*}(y)\rangle|\,dy$.

For any vectors $u, w \in \mathbb{R}^d$: $|\langle Gu, w\rangle| \leq \|G\|_{\mathrm{op}}\,|u|\,|w|$. Bounding both factors by their L^∞ suprema over Ω and pulling these constants outside the integral:

$$|\dot\tau(A)| \leq 2\,\|G\|_\infty\,\|\delta v\|_{L^\infty}\,\|\nabla\tau^{v^*}\|_{L^\infty}\int_\Omega \mathcal{G}_{v^*}(A,y)\,dy.$$

Since $-\mathcal{L}^{v^*}\tau^{v^*} = 1$ with $\tau^{v^*}|_{\partial B} = 0$, the Green's function representation gives

$$\int_\Omega \mathcal{G}_{v^*}(A,y)\,dy = \tau^{v^*}(A) =: \tau^*. \tag{61}$$

Define $\kappa^{-1} := 2\|G\|_\infty\,\tau^*\,\|\nabla\tau^{v^*}\|_{L^\infty}$, which is $O(1)$ under Assumption 3.1 since the controlled process follows the instanton with $\tau^* = O(1)$. Then:

$$|\dot\tau(A)| \leq \kappa^{-1}\|\delta v\|_{L^\infty}. \tag{62}$$

First-order Taylor expansion. The map $\alpha \mapsto \tau^{v_\alpha}(A)$ is differentiable by the smooth dependence of the solution to (58) on the coefficients of $\mathcal{L}^{v_\alpha}$ (implicit function theorem applied to the operator equation). Hence:

$$\left|\tau^{\tilde v}(A) - \tau^{v^*}(A)\right| \leq \kappa^{-1}\|\delta v\|_{L^\infty} + O(\|\delta v\|_{L^\infty}^2) \leq \underbrace{\frac{C_\rho}{\kappa(1-\eta)}}_{=:C_\tau} e^{-\delta T}, \tag{63}$$

using (56). This establishes conclusion (iii).

4. Self-consistency of the spectral bound

We verify that the bound of Stage 1, derived for a fixed T, remains valid when T is replaced by $\tilde T := \tau^{\tilde v}(A)$, the natural transit time under the perturbed control.

Write $h^T(x,t)$ to make the dependence on the horizon explicit. Stage 3 gives $|\tilde T - T^*| = |\delta\tau| \leq C_\tau e^{-\delta T^*}$. Evaluate (44) at $T = \tilde T$, $t = 0$:

$$|h^{\tilde T}(x,0) - \tilde q(x)| \leq C_0\,e^{-\delta\tilde T} = C_0\,e^{-\delta T^*}\,e^{-\delta\,\delta\tau}.$$

Case $\delta\tau \geq 0$: $e^{-\delta\,\delta\tau} \leq 1$, so the bound is no worse than $C_0 e^{-\delta T^*}$.

Case $\delta\tau < 0$: $e^{-\delta\,\delta\tau} = e^{\delta|\delta\tau|} = 1 + \delta|\delta\tau| + O(|\delta\tau|^2)$. Using (63): $\delta|\delta\tau| \leq \delta C_\tau e^{-\delta T^*} =: Me^{-\delta T^*}$ with $M = O(1)$. Therefore

$$|h^{\tilde T}(x,0) - \tilde q(x)| \leq C_0\,e^{-\delta T^*}\left(1 + O(e^{-\delta T^*})\right). \tag{64}$$

In both cases the leading-order bound is unchanged; the correction at the shifted time is $O(e^{-2\delta T^*})$, quadratically smaller. The argument is self-consistent, and the bounds of Stages 1–3 hold at the shifted horizon.

5. Perturbative analysis, stationarity, and local convexity

We prove conclusion (iv), that $\tilde v$ lies in a basin of v^* in which $v \mapsto S^{\mathrm{DL}}_{[0,T]}[q^v, v]$ is locally convex with respect to the $L^2(q^*)$ norm, and gradient descent on (v,T) converges to (v^*, T^*).

By Stage 2, $\tilde v = v^* + \delta v$ with $\|\delta v\|_{L^2(q^*)} \leq C_{\delta v}e^{-\delta T/2}$. Write $v_\alpha := v^* + \alpha\,\delta v$ and expand the KL divergence $D_{\mathrm{KL}}(\mathbb{Q}^{v_\alpha}\|\mathbb{P}) = S^{\mathrm{DL}}_{[0,T]}[q^{v_\alpha}, v_\alpha]$ to second order in α:

$$\begin{aligned} D_{\mathrm{KL}}(\mathbb{Q}^{v^*+\alpha\delta v}\|\mathbb{P}) = S^{\mathrm{DL}}_{[0,T]}[q^*, v^*] + 2\alpha\int_0^T\!\!\int_\Omega q_t^*(x)\langle\delta v(x), G_t v_t^*(x)\rangle\,dx\,dt \\ + \alpha^2 S^{\mathrm{DL}}_{[0,T]}[q^*, \delta v] + O(\alpha^3). \end{aligned} \tag{65}$$

Stationarity. Minimizing $S^{\mathrm{DL}}_{[0,T]}[q^v, v]$ subject to the Fokker–Planck constraint

$$\partial_t q_t^v + \nabla \cdot \big((b + 2G_t v_t)\, q_t^v\big) - \mathrm{tr}(G_t\, \nabla^2 q_t^v) = 0, \quad q_0^v = \delta_A, \tag{66}$$

via the Lagrangian with adjoint variable $\lambda_t(x)$ (Tröltzsch, 2010), the first variation in $v_t(x)$ yields the optimality condition $v_t^*(x) = \nabla_x \lambda_t(x)$ wherever $q_t^*(x) > 0$, and the first variation in $q_t(x)$ gives the adjoint PDE for λ_t with terminal condition $\lambda_T = 0$ on ∂B. One verifies directly that $\lambda_t(x) = \log h(x,t)$ satisfies both conditions via the BKE (47). Substituting into the linear term of (65):

$$\frac{\delta}{\delta v} D_{\mathrm{KL}}(\mathbb{Q}^v \| \mathbb{P})\Big|_{v=v^*} = 2 \int_0^T \int_\Omega q_t^*(x) \langle \delta v(x), G_t \underbrace{(v_t^*(x) - \nabla_x \lambda_t(x))}_{=\,0} \rangle \, dx\, dt = 0 \quad \forall\, \delta v. \tag{67}$$

The linear term vanishes identically, hence v^* is a critical point of $S^{\mathrm{DL}}_{[0,T]}[q^v, v]$.

Local convexity. The second-order term in (65) is:

$$\frac{1}{2} \frac{d^2}{d\alpha^2} D_{\mathrm{KL}}(\mathbb{Q}^{v^* + \alpha \delta v} \| \mathbb{P})\Big|_{\alpha=0} = S^{\mathrm{DL}}_{[0,T]}[q^*, \delta v] = \int_0^T \int_\Omega q_t^*(x) \langle \delta v(x), G_t\, \delta v(x) \rangle \, dx\, dt. \tag{68}$$

Since $G_t = \epsilon\, a(x)$ is symmetric positive definite and $q_t^*(x) \geq 0$ with $\int q_t^*\, dx = 1$, the integrand is non-negative for any $\delta v \not\equiv 0$:

$$\frac{d^2}{d\alpha^2} D_{\mathrm{KL}}(\mathbb{Q}^{v^* + \alpha \delta v} \| \mathbb{P})\Big|_{\alpha=0} = 2 S^{\mathrm{DL}}_{[0,T]}[q^*, \delta v] \geq 0. \tag{69}$$

Stationarity (67) and positive semi-definiteness (69) confirm that v^* is a critical point at the centre of a locally convex basin of $S^{\mathrm{DL}}_{[0,T]}[q^v, v]$.

The committor initialization falls inside the convex basin. The second-order approximation (65) is controlled by the cubic remainder. If $\nabla^2[(\tilde{q} - h(\cdot, 0))/\tilde{q}]$ is L-Lipschitz, the remainder satisfies $|\text{remainder}| \leq C\, L\, \|\delta v\|^3_{L^2(q^*)}$ for a constant C depending only on d. The quadratic term dominates whenever $\|\delta v\|_{L^2(q^*)} \leq c/\sqrt{L}$ for c sufficiently small. From (57), this holds when $C_{\delta v}\, e^{-\delta T/2} \leq c/\sqrt{L}$, i.e. for $T \geq 2\delta^{-1} \log(C_{\delta v}\sqrt{L}/c) = C_* \delta^{-1} \log L$ as assumed. The initialization $(\tilde{v}, \tau^{\tilde{v}}(A))$ therefore lies strictly inside the convex basin, and gradient descent on (v, T) is stable and converges to (v^*, T^*). This establishes conclusion (iv). □

The accuracy requirement on the committor translates to $\|\nabla \log(1 + \rho)\|_{L^2(q^*)} \leq c/\sqrt{L}$, which can be verified a posteriori from the quality of the approximate committor used to generate the initialization.

B Analytical verification on benchmark potentials

Cor. 3.1 places DL in the same family as classical instanton methods, e.g., gMAM (Heymann & Vanden-Eijnden, 2008), MAM (E et al., 2004), and the string method (E et al., 2005), all of which compute minimum-action paths; the key differentiator is that DL parameterizes a learned *control field* v rather than an explicit path discretization, making it amenable to automatic differentiation and high-dimensional systems. We verify this equivalence explicitly on two canonical test systems, assuming a pure gradient drift $b(x) = -\nabla U(x)$ with isotropic diffusion $a(x) = I$, for which the FW action simplifies to $S_{[0,T],A}(\phi) = \frac{1}{4} \int_0^T \|\dot{\phi}_t + \nabla U(\phi_t)\|^2\, dt$.

Potentials. The *symmetric double-well* potential is

$$U_{\mathrm{DW}}(x, y) = \beta\big((x^2 - 1)^2 + 2y^2\big), \qquad \beta = 1, \tag{70}$$

with stable equilibria $A = (-1, 0)$, $B = (1, 0)$, saddle $x_\star = (0, 0)$, and barrier height $\Delta U = \beta = 1$. The *Müller–Brown* potential (Müller, 1980) is

$$U_{\mathrm{MB}}(x, y) = \sum_{k=1}^{4} A_k \exp\big[a_k (x - x_k^0)^2 + b_k (x - x_k^0)(y - y_k^0) + c_k (y - y_k^0)^2\big]. \tag{71}$$

This yields two stable minima A and B, an intermediate metastable minimum C, and two saddles S_1 and S_2 separating $A \to C$ and $C \to B$ respectively.

Symmetric double-well. Expanding the squared norm in the FW action and applying the AM–GM and Cauchy–Schwarz inequalities yields

$$S_{[0,T],A}(\phi) \;\geq\; \int_A^{x_\star} \nabla U(\phi) \cdot d\phi = U(x_\star) - U(A) =: \Delta U\,, \tag{72}$$

attained precisely when $\dot{\phi}_t = \nabla U(\phi_t)$ (the time-reversed gradient flow), which defines the Minimum Energy Path (MEP). Hence $V(A, x_\star) = \Delta U$ and, by Proposition 3.2, $V(A, B) = \Delta U$. The DL optimization must therefore converge to

$$\inf_{T>0} \epsilon\, S^{\mathrm{DL}}_{[0,T],A}(\varphi^{v^*}) = \Delta U\,. \tag{73}$$

Müller–Brown potential. The optimal path threads $A \to S_1 \to C \to S_2 \to B$. Applying the MEP argument to each uphill segment and using the zero-cost of downhill relaxations (Proposition 3.2), the quasipotential decomposes additively:

$$V(A, B) = \big[U(S_1) - U(A)\big] + \big[U(S_2) - U(C)\big]\,. \tag{74}$$

Baseline methods. We compare against four baselines, all of which target the same quasipotential $V(A, B)$ and are therefore directly comparable to the DL action S_{DL}. The *finite-temperature string* (FTS) (E et al., 2005) evolves a discretized path under the mean force at finite temperature, converging to the MEP in the zero-noise limit; it serves as a natural reference since it operates at the same noise levels ϵ as our method. The *geometric minimum action method* (gMAM) (Heymann & Vanden-Eijnden, 2008) minimizes the FW action directly over curves in configuration space by decoupling path geometry from transit time via arc-length reparametrization, and provides the closest methodological analogue to DL among the baselines. *Direct path optimization* (E et al., 2004) minimizes the action over an explicit path discretization without the geometric reparametrization of gMAM, and is included to assess the benefit of the arc-length reparametrization. The *semiclassical approximation* (S_{SC}) (Hänggi et al., 1990) evaluates the quadratic expansion of the FW action around the instanton; it is asymptotically exact as $\epsilon \to 0$ but degrades rapidly at larger noise, providing a useful upper bound on the regime where single-instanton approximations are valid.

Model Parameterization. Here we provide a description of our method. Our initial paths are extracted as reactive trajectories from the committor estimation algorithm described in Mitchell & Rotskoff (2024). We use the Gaussian path parameterization of Du et al. (2024, Eq. 15), in which a neural network takes as input a time $0 < t < T$ and outputs the mean $\mu^{(\theta)}_{t|0,T}$ and per-dimension standard deviation $\sigma^{(\theta)}_{t|0,T}$ of the marginal $q_{t|0,T} = \mathcal{N}\big(\mu^{(\theta)}_{t|0,T}, \Sigma^{(\theta)}_{t|0,T}\big)$, where $\Sigma^{(\theta)}_{t|0,T} = \mathrm{diag}\,\big((\sigma^{(\theta)}_{t|0,T})^2\big)$. Points are sampled as

$$x_{t|0,T} = \mu^{(\theta)}_{t|0,T} + \sigma^{(\theta)}_{t|0,T} \odot \xi, \qquad \xi \sim \mathcal{N}(0, I_d), \tag{75}$$

where, with $s = t/T$,

$$\mu^{(\theta)}_{t|0,T} = (1-s)A + sB + s(1-s)\,\mathrm{NNET}_\theta(s)_{[:d]}, \tag{76}$$

$$\sigma^{(\theta)}_{t|0,T} = \sigma_{\min} + s(1-s)\exp\big(\mathrm{NNET}_\theta(s)_{[d:]}\big). \tag{77}$$

The drift u_t and control v_t are computed from this parameterization as in Du et al. (2024, Eqs. 13–14). Stage 2 of Algorithm 1 follows Du et al. (2024, Alg. 1), except that T is trainable, the action is scaled by T, and θ and T are updated in alternation.

Reported quantities. All methods are evaluated through the quasipotential $V(A, B)$ (Eq. (3)); the specific quantities reported in Table 1 are: S_{DL}, the action recovered by our neural control field; $V(A, B)$, the analytically computed quasipotential ($\Delta U = 1$ for the double-well; Eq. (74) for Müller–Brown); S_{FTS}, S_{gMAM}, S_{Direct}, and S_{SC} the actions along the paths produced by each baseline. The ratio $S_{\mathrm{DL}}/V(A, B)$ quantifies degradation of the neural method relative to the exact answer as ϵ increases and the single-instanton approximation weakens.

Algorithm 1 Training with committor path initialization and flexible T

Require: Reference drift b, diffusion matrix G_t; reactive path $\gamma = (y_0, \dots, y_K)$ with $y_0 = A$, $y_K = B$; initial time T_0

Stage 1: fit the mean path to the committor path

for $n = 1, \dots, N_{\text{fit}}$ **do**

  $\mathcal{J}_{\text{fit}}(\theta) = \sum_{k=0}^{K} \left\| \mu^{(\theta)}_{t_k|0,T} - y_k \right\|^2, \quad t_k = kT/K$

  $\theta \leftarrow \text{optimizer}_\theta\big(\theta, \nabla_\theta \mathcal{J}_{\text{fit}}\big)$

end for

Stage 2: variational Doob objective with trainable T, adapted from Du et al. (2024, Alg. 1)

Initialize $T \leftarrow T_0$

for step $= 0, 1, 2, \dots$ until converged **do**

  Sample $t \sim \mathcal{U}(0, T)$, $\xi \sim \mathcal{N}(0, I_d)$

  $x_t = \mu^{(\theta)}_{t|0,T} + \sigma^{(\theta)}_{t|0,T} \odot \xi$

  $u_t = \partial_t \mu^{(\theta)}_{t|0,T} + \Big(\frac{1}{2}\, \partial_t \Sigma^{(\theta)}_{t|0,T} - G_t \Big) \big(\Sigma^{(\theta)}_{t|0,T} \big)^{-1} \big(x_t - \mu^{(\theta)}_{t|0,T} \big)$

  $v_t = \frac{1}{2}\, G_t^{-1} \big(u_t - b(x_t) \big)$

  $S^{\text{DL}}(\theta, T) = T \, \langle v_t,\, G_t\, v_t \rangle$

  if step mod $200 < 100$ **then**

    $T \leftarrow \text{optimizer}_T\big(T, \partial_T S^{\text{DL}}\big)$ ▷ θ frozen

  else

    $\theta \leftarrow \text{optimizer}_\theta\big(\theta, \nabla_\theta S^{\text{DL}}\big)$ ▷ T frozen

  end if

end for

return θ, T

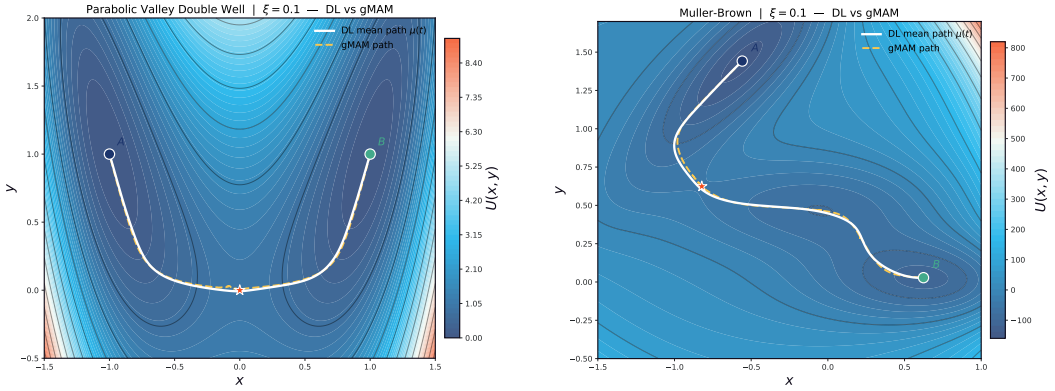
Table 1: Quasipotential estimates across noise scales ϵ on the double-well (left) and Müller–Brown (right) potentials. See Appendix B for definitions of all quantities.

	Double-well potential					Müller–Brown potential				
ϵ	S_{DL}	S_{SC}	S_{FTS}	$V(A,B)$	$S_{\text{DL}}/V(A,B)$	S_{DL}	S_{FTS}	$V(A,B)$	S_{Direct}	S_{gMAM}
0.01	19956.1	20052.7	20000.1	20000.0	1.00	2.79×10^6	3.07×10^6	2.29×10^6	2.30×10^6	2.31×10^6
0.05	801.2	801.8	800.1	800.0	1.00	1.10×10^5	1.19×10^5	0.92×10^5	0.92×10^5	0.93×10^5
0.10	200.9	200.6	200.1	200.0	1.00	2.77×10^4	3.10×10^4	2.29×10^4	2.30×10^4	2.31×10^4
0.50	11.8	8.4	8.1	8.0	1.48	1094.8	1270.8	916.4	919.1	925.7
1.00	7.3	2.2	2.1	2.0	3.64	262.9	364.5	229.1	229.8	231.4
3.00	8.1	0.7	0.2	0.2	36.3	34.1	42.9	25.5	25.5	25.7
5.00	8.3	0.7	0.2	0.1	103.5	16.7	17.2	9.2	9.2	9.3

C Molecular Simulation Details

All molecular simulations were performed in OpenMM 8.0 Eastman et al. (2024) using the AMBER 14 suite of force fields Case et al. (2023). MDTraj McGibbon et al. (2015) was used to manipulate and analyze molecular trajectories, and PyMol Schrödinger, LLC (2015) was used to visualize structures and trajectories. Experiments were run on single NVIDIA 4080 GPUs; wall-clock execution time of each optimization experiment was on the order of 2-3 hours maximum.

We based our implementation of the Doob's Lagrangian algorithm on the repository associated with Du et al. (2024), with minimal changes. Importantly, to backpropagate through the MD force calculations, we built our implementation in PyTorch Ansel et al. (2024) and calculated forces using a custom PyTorch function that mirrors the force calculation from OpenMM. For the construction of committor-based reaction paths, we used the implementation from Mitchell & Rotskoff (2024) out-of-the-box for both systems. We report the action values and final learned T for all flexible-T molecular systems in table 2. System-specific details are provided below.



(a) Double-well potential.

(b) Müller–Brown potential.

Figure 4: Mean trajectory $\mu(t)$ learned by the DL algorithm on the two benchmark potentials, overlaid on the energy landscape. In both cases the learned path connects A to B through the saddle (star) and recovers the MEP predicted analytically above.

Table 2: Learned transit times and final Doob action values for flexible-T optimization under linear and committor-based initialization.

System	Initialization	T_{init} (ps)	T_{learned} (ps)	Action value
Alanine dipeptide	Linear	1.0	2.49×10^{-3}	1.27×10^{5}
	Committor	1.0	1.44×10^{-2}	3.53×10^{4}
Chignolin	Linear	1000	0.172	1.33×10^{8}
	Committor	1000	2.15	6.36×10^{7}

C.1 Alanine Dipeptide

Initial structures for alanine dipeptide were obtained from the repository associated with Du et al. (2024). We ran the committor-based sampling algorithm using a circular boundary with radius 10 degrees in $\phi - \psi$ dihedral space around each configuration as definitions for states A and B. We used a swarm size of 100 and a swarm trajectory length of 1 fs simulated at 300K using OpenMM's Brownain integrator with friction coefficient 1/ps and timestep 0.01 fs. We ran the committor algorithm for 1000 iterations, collecting several dozen $A \to B$ transition paths; from these, we selected the path with the lowest maximum energy for use as the initial interpolant in the Lagrangian optimization.

We parameterized the mean and variance of the single-Gaussian transition paths with a 5-layer feed-forward neural network with 256 hidden units and a SiLU activation function. Because we analyzed molecular systems in the absence of noise, the parameterization of the mean was used exclusively. For each set of optimizations, we ran 10^5 training steps with an Adam optimizer of learning rate 10^{-5}. At each iteration, we sampled configurations uniformly from the interval $[0, T]$ with a batch size of 100 points.

Fixed-T simulations were run with $T = 0.01$ ps. For the flexible-T simulations, we initialized $T = 1$ ps and trained it with its own Adam optimizer of learning rate 10^{-4}. During training, we alternated between optimizing T with fixed neural network parameters and optimizing the parameters with fixed T every 100 training steps. To initialize the algorithm with path obtained by the committor, we trained the network to linearly interpolate between the sampled points along the path by using a mean-squared error on the cartesian coordinates aligned to the initial structure.

C.2 Chignolin

Initial structures for chignolin were also obtained from the repository associated with Du et al. (2024). We ran the committor-based sampling algorithm using the Euclidean distance between the flattened pairwise distance arrays of backbone C_α atoms as basin definitions (a radius of 0.25nm for the folded state and 1.0nm for the unfolded state). We used a swarm size of 100 and a swarm trajectory length of 100 fs simulated at 500K using OpenMM's BAOAB integrator with friction coefficient 1/ps and timestep 1 fs. We ran the committor algorithm until a single $A \to B$ transition was obtained, which took 729 iterations of the algorithm.

Fixed-T simulations were run with $T = 1$ ps, as in Du et al. (2024). For the flexible-T simulations, we initialized $T = 1$ ns and trained it with its own Adam optimizer of learning rate 10^{-1}. The optimization and initialization strategies for chignolin were identical to alanine dipeptide.