REACTIVE FLUX MATCHING: MECHANISM DISCOVERY AND ADAPTIVE SAMPLING OF RARE EVENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Path sampling methods generate ensembles of reactive trajectories connecting metastable states, but extracting mechanistic insight from these data remains nontrivial. We introduce *Reactive Flux Matching*, a framework that learns two complementary objects directly from reactive trajectory data: a *current velocity* u whose streamlines trace the dominant reaction pathways, and a scalar *potential* h, obtained from a weighted Helmholtz–Hodge decomposition of the reactive current, that serves as a data-driven reaction coordinate. Both minimize quadratic functionals over the reactive path ensemble, analogous to the flow matching loss in generative modeling, and require no knowledge of the underlying dynamics or stationary distribution. Unlike committor-based methods, u and h remain well-defined under projection onto non-Markovian collective variables, and their level sets in turn provide adaptive interfaces for improved sampling with enhanced sampling methods. Flux Matching is validated through the generation of current velocity trajectories, adaptive sampling, and rate constant calculations on molecular systems.

1 INTRODUCTION

Understanding the mechanisms of rare transitions between metastable states and extreme events is a central challenge across computational chemistry, biology, materials science, and climate science. Examples range from protein folding (Bowman et al., 2011; Dill and MacCallum, 2012) and chemical reactions to (Gispen and Dijkstra, 2023) and extreme weather phenomena such as heat waves (Meehl and Tebaldi, 2004; Stott et al., 2004), sudden stratospheric warmings (Baldwin and Dunkerton, 2001), and El Niño onset (Cai et al., 2014). These events are rare because the system must traverse regions of phase space that are seldom visited under typical dynamics: a molecular dynamics trajectory may run for microseconds while the folding event of interest occurs on millisecond timescales, and obtaining robust statistics on century-scale climate extremes requires ensembles of long simulations that are computationally prohibitive.

Path sampling methods address this timescale problem by focusing computational effort on the rare fluctuations that successfully connect a reactant region A to a product region B. Transition path sampling (TPS) (Bolhuis et al., 2002), transition interface sampling (TIS) (van Erp and Bolhuis, 2005), forward flux sampling (FFS) (Allen et al., 2009), and weighted ensemble (WE) (Huber and Kim, 1996; Aristoff, 2018) all allow to generate ensembles of *reactive trajectories*, paths that connect A to B, providing direct access to the statistics of transition events.

However, extracting mechanistic insight from reactive trajectory data remains nontrivial. The central object of interest is traditionally the *committor function* $q(x)$, the probability that a trajectory initiated at x reaches B before A (E and Vanden-Eijnden, 2006; Vanden-Eijnden, 2006; Hummer, 2004; Best and Hummer, 2005). The committor is often regarded as the ideal reaction coordinate, but it is fundamentally tied to the Markov property of the underlying dynamics. When trajectories are projected onto collective variables, as is inevitable in high-dimensional systems, the projected dynamics are generally non-Markovian, and the committor of the full system cannot be expressed as a function of the reduced variables alone. Methods that learn a committor in the reduced space must therefore make uncontrolled approximations.

In this work, we take a different approach. Rather than targeting the committor directly, we characterize the reactive path ensemble through two complementary objects, both estimated by regression from trajectory data alone.

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.

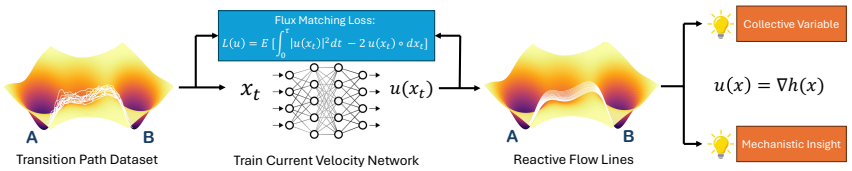

Figure 1: **Overview of Flux Matching.** From an ensemble of reactive trajectories, we learn the probability current and a scalar potential whose level sets define a data-driven reaction coordinate and milestones for refining the sampler.

The first is the *current velocity* u, the ratio of the reactive current to the reactive density. It is the average instantaneous velocity of reactive trajectories, and its streamlines describe the deterministic skeleton of the transition mechanism from A to B. The second is a scalar *potential* h, obtained from a weighted Helmholtz–Hodge decomposition of the reactive current into an irrotational gradient component and a divergence-free remainder. The potential h carries the net flux from A to B and serves as a learned reaction coordinate. For systems satisfying detailed balance, h reduces to $\log[q/(1-q)]$ where q is the forward committor, recovering the optimal committor-based coordinate without solving any boundary value problem; but the construction extends to non-Markovian settings where the committor itself is not well-defined. In particular, both u and h remain meaningful under projection onto collective variables, where they yield the exact marginal current and potential.

Both objects are unique minimizers of quadratic functionals over the reactive path ensemble (Theorem 2.4), with empirical estimators built directly from trajectory data and requiring no knowledge of the underlying drift or stationary distribution. The loss for u is structurally identical to the flow matching / stochastic interpolant loss (Lipman et al., 2022; Liu et al., 2022; Albergo et al., 2023), with the reactive path ensemble in place of a coupling between distributions; the loss for h is a Benamou–Brenier-type functional. We call this framework *Flux Matching*. Beyond direct mechanistic analysis, the level sets of h provide a principled scalar collective variable that can replace hand-chosen ones in adaptive samplers—interfaces for TIS or FFS, bin boundaries for WE, stopping surfaces for milestoning—enabling an iterative scheme in which improved sampling and improved estimation reinforce each other.

Main contributions.

- *Variational characterization of the reactive path ensemble.* We characterize the reactive path ensemble through a current velocity u and a scalar potential h obtained from a weighted Helmholtz–Hodge decomposition of the reactive current, each satisfying a quadratic variational principle (Theorem 2.4) estimable by regression from trajectory data alone.
- *Exact treatment of projected dynamics.* u and h remain well-defined under projection onto non-Markovian collective variables, in contrast to committor-based methods, which must assume the committor depends only on the observed variables.
- *Data-driven reaction coordinates and sampling loop.* The streamlines of u reveal the dominant transition pathways; the level sets of h provide a principled reaction coordinate and adaptive milestones for TIS, FFS, WE, or milestoning, enabling iterative refinement of the sampler itself.
- *Numerical validation.* We demonstrate the framework on the Müller–Brown potential under overdamped and underdamped dynamics, Alanine Dipeptide (ADP), and AIB9 molecular systems.

1.1 Related Work

Path sampling methods. Reactive Flux Matching is agnostic to how reactive trajectories are generated; suitable ensembles can be obtained from transition path sampling (TPS) (Bolhuis et al., 2002), transition interface sampling (TIS) (van Erp and Bolhuis, 2005), forward flux sampling (FFS) (Allen et al., 2009), weighted ensemble (WE) (Huber and Kim, 1996), or milestoning (Faradjian and Elber, 2004). Adaptive multilevel splitting (Cérou and Guyader, 2007; Bréhier and Lelièvre, 2019) and related rare event algorithms have been applied to climate and geophysical problems (Ragone et al., 2018; Lucente et al., 2022); see Lelièvre et al. (2010) and Hénin et al. (2022) for reviews.

Flow matching and probability currents. Flow matching (Lipman et al., 2023), stochastic interpolants (Albergo and Vanden-Eijnden, 2023; Albergo et al., 2023), and rectified flows (Liu et al., 2023) learn velocity fields that transport a simple distribution to a complex target. Our variational

principle (Theorem 2.4) has the same quadratic structure, but the underlying measure is the reactive path ensemble rather than a coupling between distributions. Related ideas have been used to learn probability currents and entropy production rates in active matter systems (Boffi and Vanden-Eijnden, 2024; 2025; Maddu et al., 2025).

Committor estimation. Transition path theory (E and Vanden-Eijnden, 2006; Vanden-Eijnden, 2006; Metzner et al., 2009) expresses the reactive density and current in terms of the committor function, and is applied to molecular simulations (Noé et al., 2009; Berezhkovskii and Makarov, 2019). Likelihood-based methods (Peters and Trout, 2006; Ma and Dinner, 2005) estimate the committor by fitting shooting outcomes, while neural network approaches (Li et al., 2019; Khoo et al., 2019; Rotskoff et al., 2022; Mitchell and Rotskoff, 2024; Lin and Ren, 2024) solve the committor PDE variationally. Recent work also leverages stochastic optimal control to estimate the committor (Du et al., 2026). These rely on the committor being well-defined as a function of the observed state—a Markov property that fails under projection onto collective variables. Reactive Flux Matching sidesteps this issue by targeting the current directly, which remains well-defined regardless.

2 Theoretical Framework

2.1 Reactive Trajectories

We consider a system evolving stochastically on $\mathbb{R}^d$ in statistical steady state, not necessarily Markovian — this generality covers coarse-grained or projected dynamics (Appendix B). Transitions occur between two disjoint closed sets with smooth boundaries: the *reactant* $A \subset \mathbb{R}^d$ and the *product* $B \subset \mathbb{R}^d$, each possibly with multiple connected components. We denote by $\Omega_R = \mathbb{R}^d \setminus (A \cup B)$ the *reactive region*.

Definition 2.1 (Reactive Trajectory and Path Ensemble)**.** A *reactive trajectory* is a continuous path $\omega = (z_t)_{t\in[0,\tau(\omega)]}$ satisfying:

$$z_0 \in \partial A, \qquad z_\tau \in \partial B, \qquad z_t \in \Omega_R \quad \text{for all } t \in (0,\tau), \tag{1}$$

where $\tau = \tau(\omega) > 0$ is the duration of the path, which varies from one trajectory to another. A *reactive path ensemble* is a probability measure $\mathbb{P}$ on the space of such paths. We denote by $\mathbb{E}$ the expectation with respect to $\mathbb{P}$.

The measure $\mathbb{P}$ encodes all statistical information about the reactive dynamics: correlations, memory effects, and the distribution of transition times τ. We assume that $\mathbb{P}$-almost every path has finite quadratic variation (so Stratonovich integrals are well-defined), and that we have access to an empirical sample $\{(z_t^{(i)})_{t\in[0,\tau_i]}\}_{i\in[N]}$.

2.2 Reactive Density, Current, and Surface Densities

The reactive path ensemble is characterized by four statistical quantities, defined through their duality pairings against test functions:

Definition 2.2 (Reactive density, current, and surface densities)**.** For any test functions $\phi : \mathbb{R}^d \to \mathbb{R}$ and $\psi : \mathbb{R}^d \to \mathbb{R}^d$,

$$\textit{reactive density } \rho_R : \Omega_R \to \mathbb{R}_{\geqslant 0}: \qquad \mathbb{E}\left[\int_0^\tau \phi(z_t)\, dt\right] = \int_{\Omega_R} \phi(z)\, \rho_R(z)\, dz, \tag{2}$$

$$\textit{reactive current } j_R : \Omega_R \to \mathbb{R}^d: \qquad \mathbb{E}\left[\int_0^\tau \psi(z_t) \circ dz_t\right] = \int_{\Omega_R} \psi(z) \cdot j_R(z)\, dz, \tag{3}$$

$$\textit{surface density on } \partial A,\ \rho_{\partial A}: \qquad \mathbb{E}[\phi(z_0)] = \int_{\partial A} \phi(z)\, \rho_{\partial A}(z)\, d\sigma_{\partial A}(z), \tag{4}$$

$$\textit{surface density on } \partial B,\ \rho_{\partial B}: \qquad \mathbb{E}[\phi(z_\tau)] = \int_{\partial B} \phi(z)\, \rho_{\partial B}(z)\, d\sigma_{\partial B}(z), \tag{5}$$

where $\circ$ denotes the Stratonovich integral and $d\sigma_{\partial A}$, $d\sigma_{\partial B}$ the surface elements on ∂A, ∂B.

Intuitively, ρ_R records the time-averaged occupation of reactive trajectories — so $\int_{\Omega_R} \rho_R \, dz = \mathbb{E}[\tau] \equiv T_R$, the mean transition time — while j_R records their average flux and $\rho_{\partial A}, \rho_{\partial B}$ record where they exit A and enter B. The Stratonovich convention in (3) is natural because it respects the chain rule: for smooth paths, it reduces to $\int_0^\tau \psi(z_t) \cdot \dot{z}_t \, dt$. All four objects are estimable from a sample by empirical averages: e.g., $\frac{1}{N} \sum_i \int_0^{\tau_i} \phi(z_t^{(i)}) \, dt \to \int_{\Omega_R} \phi \, \rho_R \, dz$ as $N \to \infty$.

Since reactive trajectories enter Ω_R through ∂A and exit through ∂B, the current is divergence-free in Ω_R with source and sink at its boundaries:

Lemma 2.3. *The current j_R satisfies*

$$\begin{aligned} \operatorname{div} j_R &= 0 && \text{in } \Omega_R, \\ \hat{n}_{\partial A} \cdot j_R &= -\rho_{\partial A} && \text{on } \partial A, \\ \hat{n}_{\partial B} \cdot j_R &= \rho_{\partial B} && \text{on } \partial B, \end{aligned} \tag{6}$$

where $\hat{n}_{\partial A}$ and $\hat{n}_{\partial B}$ denote the unit normal vectors pointing inward into A and B, respectively.

Proof. Evaluating (3) with $\psi = \nabla\phi$, the left-hand side becomes

$$\mathbb{E}\left[\int_0^\tau \nabla\phi(z_t) \circ dz_t\right] = \mathbb{E}[\phi(z_\tau)] - \mathbb{E}[\phi(z_0)] = \int_{\partial B} \phi \, \rho_{\partial B} \, d\sigma_{\partial B} - \int_{\partial A} \phi \, \rho_{\partial A} \, d\sigma_{\partial A}, \tag{7}$$

where we used (4) and (5). The right-hand side, by integration by parts, is

$$\int_{\Omega_R} \nabla\phi \cdot j_R \, dz = -\int_{\Omega_R} \phi \, \operatorname{div} j_R \, dz + \int_{\partial B} \phi \, \rho_{\partial B} \, d\sigma_{\partial B} - \int_{\partial A} \phi \, \rho_{\partial A} \, d\sigma_{\partial A}. \tag{8}$$

Equating (7) and (8) yields (6). □

2.3 Current Velocity and Potential

Given the reactive density and current, we define the *current velocity* as

$$u(z) = \frac{j_R(z)}{\rho_R(z)}. \tag{9}$$

This velocity field captures the average instantaneous velocity of reactive trajectories passing through z. It allows us to compute the streamlines of the probability current by integrating the *probability flow ODE*

$$\dot{Z}_t = u(Z_t), \tag{10}$$

either forward or backward in time. A natural choice of initial condition is to draw Z_0 from the reactive trajectory ensemble, i.e., as a snapshot z_t for some $t \in [0, \tau]$.

We also define the *potential* h via the weighted Helmholtz–Hodge decomposition of j_R relative to the density ρ_R:

$$j_R = \underbrace{\rho_R \, D\nabla h}_{\text{irrotational part}} + \underbrace{j_R^\perp}_{\text{solenoidal part}}, \tag{11}$$

where $D(z)$ is a user-specified positive-definite symmetric tensor that can be chosen to reflect the local geometry of the dynamics; in the examples below we set $D = \mathrm{Id}$. In this decomposition, the irrotational part carries the entire net flux from A to B: for any dividing surface $\Sigma \subset \Omega_R$ separating A from B, the divergence theorem applied to the regions between ∂A and Σ and ∂B and Σ implies

$$\int_\Sigma j_R \cdot \hat{n} \, d\sigma = \int_\Sigma \rho_R D\nabla h \cdot \hat{n} \, d\sigma = \int_{\partial A} \rho_{\partial A} \, d\sigma_{\partial A} = \int_{\partial B} \rho_{\partial B} \, d\sigma_{\partial B} = 1,$$

Equivalently, requiring the solenoidal part $j_R^\perp$ to have zero normal flux on $\partial A \cup \partial B$ uniquely determines the decomposition, and $\rho_R D\nabla h$ then satisfies the same boundary value problem (6) as j_R. The potential h thus admits a Benamou–Brenier characterization as the unique (up to a constant) minimizer of $\int_{\Omega_R} |\nabla h|_D^2 \, \rho_R \, dz$ (with $|v|_D^2 = v^\top D v$) subject to these prescribed boundary fluxes. The potential h thus serves as a data-driven reaction coordinate: its level sets define natural milestones for adaptive samplers such as TIS, FFS, or milestoning. For Markovian dynamics satisfying detailed balance, h reduces to $\log[q/(1-q)]$, with q the committor (Section 2.5); but the construction extends to non-Markovian settings where the committor itself is not well-defined.

2.4 Variational Characterization

The current velocity and the potential admit the following variational characterizations:

Theorem 2.4 (Flux Matching Losses)**.** *Given any symmetric positive-definite tensor $D(z)$, the current velocity u in (9) is the unique minimizer of*

$$\mathcal{L}_u(u) = \mathbb{E}\left[\int_0^\tau \left(|u(z_t)|^2_{D^{-1}(z_t)}\, dt - 2\, u(z_t)^\top D^{-1}(z_t) \circ dz_t\right)\right], \tag{12}$$

and the potential h defined in (11) is the unique (up to a constant) minimizer of

$$\mathcal{L}_h(h) = \mathbb{E}\left[\int_0^\tau |\nabla h(z_t)|^2_{D(z_t)}\, dt + 2h(z_0) - 2h(z_\tau)\right]. \tag{13}$$

Proof. By definition of the probability density and current of reactive trajectories, the objective (12) can be written as

$$\mathcal{L}_u(u) = \int_{\Omega_R} \left[|u|^2_{D^{-1}} \rho_R - 2\, u^\top D^{-1} j_R\right] dz,$$

whose unique minimizer is (9). Similarly, (13) reads

$$\mathcal{L}_h(h) = \int_{\Omega_R} |\nabla h|^2_D\, \rho_R\, dz + 2\int_{\partial A} h\, \rho_{\partial A}\, d\sigma_{\partial A} - 2\int_{\partial B} h\, \rho_{\partial B}\, d\sigma_{\partial B},$$

whose unique minimizer is the potential h from (11). □

The loss (12) is a D^{-1}-weighted version of the flow matching / stochastic interpolants objective (Albergo and Vanden-Eijnden, 2023; Albergo et al., 2023), using reactive trajectories as the interpolants; a related formulation was used in (Boffi and Vanden-Eijnden, 2025) to characterize probability currents and entropy production rates in nonequilibrium systems. The two losses in Theorem 2.4 are intimately connected: $\mathcal{L}_h$ is the restriction of $\mathcal{L}_u$ to gradient fields. Indeed, substituting $u = D\nabla h$ in (12), the weighted norm becomes $|D\nabla h|^2_{D^{-1}} = |\nabla h|^2_D$, and the Stratonovich integrand reduces to $(D\nabla h)^\top D^{-1} \circ dz_t = \nabla h^\top \circ dz_t = dh(z_t)$, whose integral telescopes to $h(z_\tau) - h(z_0)$; this yields (13) exactly.

2.5 Connection to Transition Path Theory

When the dynamics is Markovian, our framework recovers the objects from transition path theory (TPT). Consider an SDE

$$dx_t = b(x_t)\, dt + \sqrt{2}\, \sigma(x_t)\, dW_t, \tag{14}$$

with stationary density ρ and diffusion tensor $D = \sigma\sigma^\top$, and let $q(x)$ and $\tilde{q}(x)$ be the forward and backward committors — the probabilities of reaching B before A, and of having come from A rather than B, respectively. As shown in Appendix A,

$$\rho_R = \rho\, q\, \tilde{q}, \qquad j_R = q\tilde{q}\, j + \rho\big[\tilde{q}\, D\nabla q - q\, D\nabla\tilde{q}\big], \tag{15}$$

where $j = b\rho - \mathrm{div}(D\rho)$ is the stationary current of (14). The current velocity follows as

$$u = b - \mathrm{div}\, D - D\nabla \log\rho + D\nabla \log q - D\nabla \log\tilde{q}. \tag{16}$$

Under detailed balance, $j = 0$ and $\tilde{q} = 1 - q$, so (16) simplifies to $u = D\nabla \log[q/(1-q)]$; using the diffusion tensor D from (14) in (13) then gives $h = \log[q/(1-q)]$, the optimal committor-based reaction coordinate. Concrete examples of reactive path ensembles (including overdamped and underdamped Langevin dynamics, and the non-Markovian case of position-only observation) are given in Appendix B.

2.6 Non-Markovian Projection

When a system is observed through a projection $z_t = \phi(x_t)$ for some $\phi : \mathbb{R}^d \to \mathbb{R}^n$ with $n < d$, the dynamics is generally non-Markovian and the committor cannot be expressed as a function of the reduced variables alone. Nevertheless, the current velocity remains well-defined.

Algorithm 1: Reactive Flux Matching for the current velocity u and/or the potential h

Input: Reactive trajectories $\{(z_t^{(m)})_{t\in[0,\tau_m]}\}_{m=1}^M$, time lag Δt, learning rate η, neural network u_θ with initial parameters θ_0 and/or neural network h_ϕ with initial parameters ϕ_0, diffusion tensor $D(z)$ (e.g. $D = \mathrm{Id}$).

for *iteration* $= 1, 2, \ldots$ **do**

  Sample minibatch $\mathcal{B} \subset \{1, \ldots, M\}$

  $\mathcal{K} \leftarrow \{k \in \mathcal{B} : \tau_k > 2\Delta t\}$

  for $k \in \mathcal{K}$ **do**

    Sample $t_k \sim \mathcal{U}([\Delta t,\, \tau_k - \Delta t])$

    Set $z_k = z_{t_k}^{(k)},\ \Delta z_k = \frac{1}{2}\big(z_{t_k+\Delta t}^{(k)} - z_{t_k-\Delta t}^{(k)}\big),\ D_k = D(z_k)$

  $$L_u = \frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}}\Big(\big|u_\theta(z_k)\big|^2_{D_k^{-1}}\,\Delta t \;-\; 2\,u_\theta(z_k)^\top D_k^{-1}\,\Delta z_k\Big) \quad \text{and/or}$$

  $$L_h = \frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}}\Big(\big|\nabla h_\phi(z_k)\big|^2_{D_k}\,(\tau_k - 2\Delta t)\Big) + \mathcal{L}_{\mathrm{CE}}(h_\phi)$$

  $\theta \leftarrow \theta - \eta\,\nabla_\theta L_u$ and/or $\phi \leftarrow \phi - \eta\,\nabla_\phi L_h$

return u_θ and/or h_ϕ

Theorem 2.5. *Let $z_t = \phi(x_t)$ where $\phi : \mathbb{R}^d \to \mathbb{R}^n$ is C^2 with full rank Jacobian and x_t solves the SDE (14). The current velocity in the reduced space is*

$$u(z) = \mathbb{E}\big[\,(b(x) + 2D(x)\nabla \log q(x)) \cdot \nabla\phi(x) + D(x) : \nabla\nabla\phi(x)\,\big|\,\phi(x) = z\big], \tag{17}$$

where the expectation is over ρ_R conditional on $\phi(x) = z$.

The proof, given in Appendix C, follows from applying Itô's lemma to $z_t = \phi(x_t^R)$, where x_t^R evolves according to the Doob-transformed dynamics

$$dx_t^R = \big[b(x_t^R) + 2D(x_t^R)\nabla \log q(x_t^R)\big]\,dt + \sqrt{2}\,\sigma(x_t^R)\,dW_t, \tag{18}$$

which governs trajectories conditioned on being reactive. When $\phi = \mathrm{id}$, this recovers (16). For general ϕ, expression (17) does not require solving for the committors explicitly: it is the conditional expectation of $\dot{z}_t$ along reactive trajectories, which Reactive Flux Matching learns directly from data.

3 Numerical Experiments

We minimize the losses of Theorem 2.4 by stochastic gradient descent on minibatches of reactive trajectories; the procedure is summarized in Algorithm 1. Further implementation details are deffered to Sections D and F.

Reactive Flux Matching is investigated in several different settings. First, we demonstrate the effectiveness of Flux Matching on the toy 2D Müller-Brown (MB) system with both overdamped (first-order) and underdamped (second-order) dynamics. We further conduct experiments on the higher-dimensional Alanine Dipeptide (ADP) and AIB9 molecular systems.

The datasets for all systems in this study are generated using Transition Path Sampling (TPS) with one-way shooting moves. Briefly, shooting moves perform MCMC sampling in path space to obtain an ensemble of reactive trajectories under a given system's dynamics. Initial trajectories are generated at high temperature, followed by a brief equilibration period, after which the production ensemble of reactive paths is collected. System-specific details are deferred to Section E.

3.1 Reactive Flux Matching learns the reactive current and potential on the 2D Müller-Brown system

Figure 2 demonstrates the effectiveness of learning the reactive current for both overdamped and underdamped dynamics. Figure 3 further shows the level sets of the learned potential h together with the learned current velocity field.

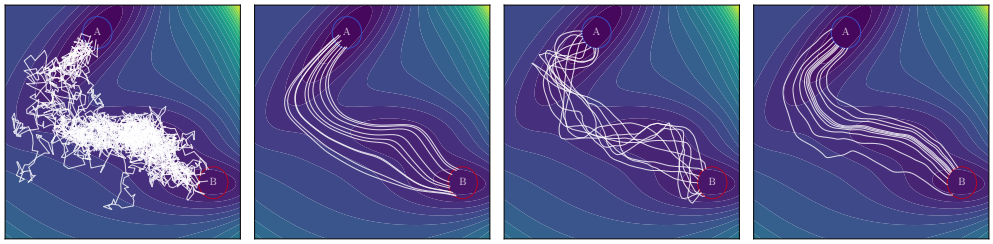

Figure 2: **Transition path and current trajectories on the Müller-Brown potential.** Reactive trajectories and learned current-velocity streamlines for overdamped (first pair) and underdamped (second pair) dynamics. The streamlines smoothly track the reactive channel in both regimes; underdamped paths show greater spread due to inertia.

Figure 3 also compares the predicted velocity with the exact solution obtained using the finite element method and observe close agreement. This indicates that Reactive Flux Matching recovers the intended target of the objective. Moreover, the level sets of the learned potential progress monotonically along the reaction direction, so h provides a natural one-dimensional reaction coordinate.

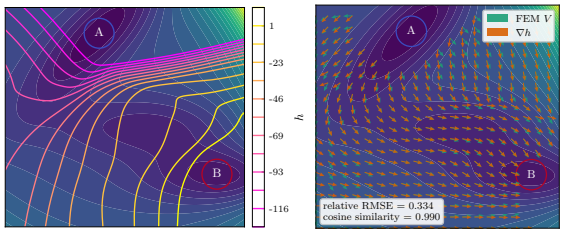

Figure 3: **Learned potential and current on the Müller–Brown potential.** *Left:* Level sets of the learned potential h. *Right:* Learned current velocity field (arrows) overlaid on the exact reactive current. The two fields agree closely.

3.2 Reactive Flux Matching estimates the current and potential on Alanine Dipeptide

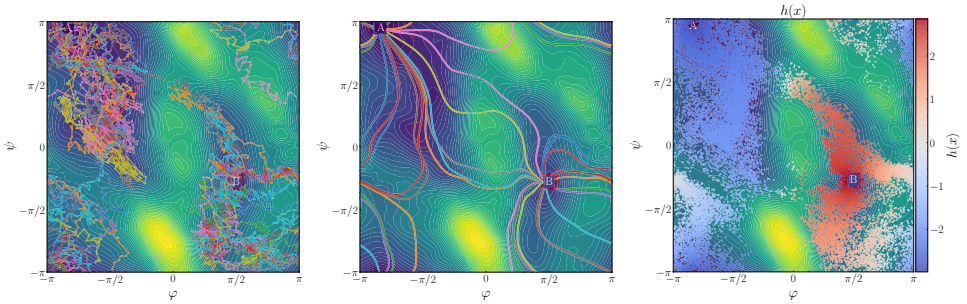

Figure 4: **Transition path, current trajectories and potential h on ADP.** Reactive trajectories (left), integrated streamlines of the learned current velocity (center), and the learned potential $h(x)$ (right), projected onto the backbone dihedral angles (φ, ψ). The streamlines reliably reach B from initial conditions in A and resolve the two main reaction channels; the potential h varies monotonically along the reaction, providing a one-dimensional reaction coordinate.

Alanine Dipeptide (N-acetyl-L-alanine-N'-methylamide) is a canonical benchmark system for testing rare event methods in molecular simulation (Bolhuis et al., 2000; Ma and Dinner, 2005; Ren et al., 2005). The free energy landscape is typically characterized using two backbone dihedral angles, φ (C-N-C_α-C) and ψ (N-C_α-C-N), which provide a natural low-dimensional projection (Ramachandran et al., 1963).

We study the transition between the main metastable states, C_7^{eq} (equatorial) and C_7^{ax} (axial) (Tobias et al., 1993; Smith, 1999). Reactive trajectories are generated using transition path sampling in vacuum at 300 K using OpenPathSampling (Swenson et al., 2019a;b).[1]. We train separate models on the full Cartesian space $\mathbb{R}^{3N}$ space, and on the full set of 7 dihedral angles of the molecule. Further details of model architecture and representations are deffered to Sections F and G.1.

[1] http://openpathsampling.org

3.2.1 Reactive Flux Matching provides insights into ADP reaction mechanism

Figure 4 shows the TPS reactive trajectories (left) and the streamlines of the learned current velocity (center), both projected onto the (φ, ψ) dihedral plane. The projected flow lines yield smooth transitions between states and are visibly more interpretable than the raw reactive path ensemble. Figure 4 additionally displays the value of h on snapshots from reactive trajectories. The potential h increases monotonically along the reaction channels connecting the two states, indicating that it captures the relevant motions of the reaction and serves as a natural collective variable for further analysis. Example transition-path structures colored by h are provided in Section G.2.

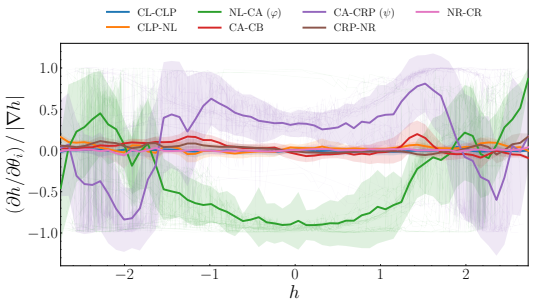

Figure 5: **Normalized gradients along the reaction.** Normalized gradients of the potential h along 100 reference TPS paths of alanine dipeptide. h picks out the backbone dihedrals φ and ψ as the essential reaction components.

As the model learns an explicit scalar reaction potential h, whose gradient field generates the transition paths, we can ask which input coordinates h actually depends on. At each frame along reactive paths we compute the unit-normalized gradient $\nabla h/|\nabla h|$ and average each component within bins of h. Normalizing removes the steep, position-dependent changes in $|\nabla h|$ and leaves only the direction in which the reaction advances.

The results shown in Figure 5 indicates that the learned reaction potential is almost entirely a function of the backbone dihedrals φ and ψ. Across all frames, φ and ψ together account for 95–98% of the squared norm of $\nabla h/|\nabla h|$. Without being told which coordinates matter, the model learns a reaction potential h that depends on φ during barrier crossing and on ψ near the product basin. It gives essentially no weight to the other five dihedrals, which yields an interpretable, low-dimensional reaction coordinate.

3.2.2 Reactive Flux Matching accelerates sampling and property calculation

The learned potential h measures progress from A to B, so it can guide its own data collection. We use it to choose shooting points in transition path sampling (TPS), which creates a **positive feedback loop**: a better h yields more informative shooting moves, which yield more diverse paths, which in turn yield a better h. On ADP, we seed the path pool with 300 transition paths and repeat the following for 10 rounds. We perform 300 one-way shooting moves, choosing each shooting point with probability inversely proportional to the density of h over the path ensemble (10 equal-width bins), so that shots are spread evenly in h between A and B. We then add the sampled paths to the pool, retrain h from scratch, and recompute the selection density. After 10 rounds the pool contains $\sim$3.3k paths. As a baseline, we select shooting points uniformly under an otherwise identical setup, and we consider two paths decorrelated once they share no configuration.

Table 1: **Shooting-point selection on ADP.** One-way TPS with uniform vs. h-flat selection; mean $\pm$ block s.e. over 10×300 moves. h-flat reduces MD cost by $2.6\times$ per independent transition path.

Metric	Uniform	h-flat	Improvement
TPS Acceptance rate	0.562 ± 0.012	0.698 ± 0.016	$+24\%$
TPS Moves / decorrelation	16.7 ± 0.9	8.74 ± 0.66	$1.91\times$
MD frames / decorrelation	2682 ± 233	1047 ± 101	$2.56\times$

As shown in Table 1, h-flat selection raises the acceptance rate by 24% and nearly halves the number of TPS moves needed per decorrelated path, reducing the MD cost per independent transition path by a factor of **2.6**. Beyond describing the reaction, h is therefore directly useful for simulation: spreading effort evenly along the learned reaction coordinate makes path sampling substantially cheaper, and the sampler and the model improve each other without additional supervision. Furthermore, Section G.3 shows that h can be combined with enhanced sampling methods such as weighted ensemble to accelerate the computation of properties such as rate constants.

3.3 Reactive Flux Matching learns the reaction on the higher dimensional AIB9 molecular system

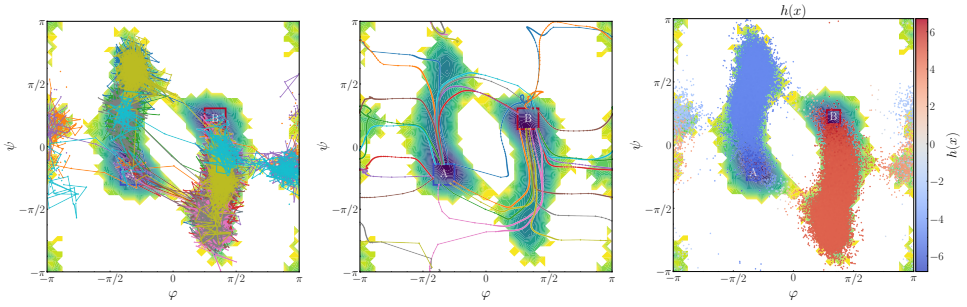

Figure 6: **Transition paths, current trajectories, and potential h for AIB9.** Reactive trajectories (left), integrated streamlines of the learned current velocity (center), and the reactive potential (right), all projected onto the backbone dihedral angles. The streamlines reliably reach B from initial conditions in A, and the potential h varies monotonically along the reaction, providing a one-dimensional reaction coordinate.

We next consider the synthetic AIB9 peptide. This system is more challenging, as it exhibits intermediate metastable states that complicate the learning of both the reactive current and the collective variable. Reactive trajectories are generated with TPS in vacuum at 500 K using OpenPathSampling. The system is also higher-dimensional, with 129 atoms; we therefore operate directly on the 29 backbone dihedral angles.

Figure 6 shows the reactive path ensemble (left), generated flow lines (center), and learned collective variable h (right), projected onto the dihedral angles of the central residue. Reactive-trajectory snapshots are saved every 100 fs, producing visibly large jumps in this projection, whereas the flow lines exhibit smooth transitions between states—further supporting our claim that the learned currents yield more interpretable reaction mechanisms. As with ADP, h varies monotonically along the reaction direction. Additional structures colored by h are provided in Section G.2.

4 Conclusion

We have presented Reactive Flux Matching, a framework for learning the probability current of reactive trajectories directly from path sampling data. The method is agnostic to how reactive trajectories are generated: it can be applied to ensembles from transition path sampling, transition interface sampling, forward flux sampling, weighted ensemble, or any other method that produces reactive paths. This flexibility allows Reactive Flux Matching to leverage well-established sampling techniques while using modern flow matching tools to estimate the current velocity field. By operating directly in the full state space, the method bypasses the need for collective variables, which are often difficult to identify in complex systems. For systems satisfying detailed balance, the learned current connects to the committor function through the gradient structure $u \propto \nabla \log[q/(1-q)]$; for general non-Markovian systems, the current velocity remains well-defined even when the committor is not. The learned current amd potential provide both quantitative information—coordinates involved in the reation, milestone placement—and qualitative mechanistic insight into reaction pathways and transition state regions.

Scaling the method to larger systems—from complex biomolecular dynamics to atmosphere-ocean flows—remains an important direction for future work. This will require neural network architectures capable of handling high-dimensional state spaces, as well as strategies to improve the quality of the reactive trajectory ensemble. A natural approach is to iterate between path sampling and current estimation: the milestones extracted from a learned current can be used to define interfaces for transition interface sampling or bins for weighted ensemble, which in turn generate trajectories that improve the current estimate. Such an iterative scheme would combine the strengths of classical rare event methods with the representational power of modern generative models.

AI USE STATEMENT

In this work, we used generative AI tools for research execution, namely implementing methods. We have not used generative AI tools for proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, formulating mathematical claims, providing ingredients for or writing proofs, developing theoretical models or conceptual frameworks, generating synthetic data sets, or interpreting results, and assisting with translation, cleaning and reformatting datasets, and supporting qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools for drafting parts of the research paper, and polishing the writing to improve readability. We have reviewed all AI-assisted work. LLM-generated code was verified and tested for correctness by the authors, and LLM-drafted and LLM-polished text was reviewed and revised by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Pytorch implementation of Reactive Flux Matching is provided as part of Supplementary Material.

ETHICS STATEMENT

This work does not involve human participants, personal or sensitive data, or other applications that raise specific ethical concerns. The authors acknowledge and adhere to the ICLR Code of Ethics.

REFERENCES

Gregory R Bowman, Vincent A Voelz, and Vijay S Pande. Taming the complexity of protein folding. *Current opinion in structural biology*, 21(1):4–11, 2011.

Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *science*, 338(6110): 1042–1046, 2012.

Willem Gispen and Marjolein Dijkstra. Brute-force nucleation rates of hard spheres compared with rare-event methods and classical nucleation theory. *The Journal of chemical physics*, 159(8), 2023.

Gerald A Meehl and Claudia Tebaldi. More intense, more frequent, and longer lasting heat waves in the 21st century. *Science*, 305(5686):994–997, 2004.

Peter A Stott, D A Stone, and M R Allen. Human contribution to the European heatwave of 2003. *Nature*, 432(7017):610–614, 2004. Landmark paper on extreme event attribution.

Mark P Baldwin and Timothy J Dunkerton. Stratospheric harbingers of anomalous weather regimes. *Science*, 294(5542):581–584, 2001.

Wenju Cai, Simon Borlace, Matthieu Lengaigne, Peter Van Rensch, Mat Collins, Gabriel Vecchi, Axel Timmermann, Agus Santoso, Michael J McPhaden, Lixin Wu, et al. Increasing frequency of extreme El Niño events due to greenhouse warming. *Nature Climate Change*, 4(2):111–116, 2014.

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.

Titus S van Erp and Peter G Bolhuis. Elaborating transition interface sampling methods. *Journal of Computational Physics*, 205(1):157–181, 2005.

Rosalind J Allen, Chantal Valeriani, and Pieter Rein ten Wolde. Forward flux sampling for rare event simulations. *Journal of Physics: Condensed Matter*, 21(46):463102, 2009.

Gary A Huber and Sangtae Kim. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophysical Journal*, 70(1):97–110, 1996.

David Aristoff. Analysis and optimization of weighted ensemble sampling. *ESAIM: Mathematical Modelling and Numerical Analysis*, 52(4):1219–1238, 2018.

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

Eric Vanden-Eijnden. Transition path theory. *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, pages 453–493, 2006.

Gerhard Hummer. From transition paths to transition states and rate coefficients. *The Journal of Chemical Physics*, 120(2):516–523, 2004.

Robert B Best and Gerhard Hummer. Reaction coordinates and rates from transition paths. *Proceedings of the National Academy of Sciences*, 102(19):6732–6737, 2005.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hoch, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL https://arxiv.org/abs/2303.08797.

Anton K Faradjian and Ron Elber. Computing time scales from reaction coordinates by milestoning. *The Journal of Chemical Physics*, 120(23):10880–10889, 2004.

Frédéric Cérou and Arnaud Guyader. Adaptive multilevel splitting for rare event analysis. *Stochastic Analysis and Applications*, 25(2):417–443, 2007.

Charles-Edouard Bréhier and Tony Lelièvre. On a new class of score functions to estimate tail probabilities of some stochastic processes with adaptive multilevel splitting. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(3):033126, 2019.

Francesco Ragone, Jeroen Wouters, and Freddy Bouchet. Computation of extreme heat waves in climate models using a large deviation algorithm. *Proceedings of the National Academy of Sciences*, 115(1):24–29, 2018.

Dario Lucente, Joran Rolland, Corentin Herbert, and Freddy Bouchet. Coupling rare event algorithms with data-based learned committor functions using the analogue Markov chain. *Journal of Statistical Mechanics: Theory and Experiment*, 2022(8):083201, 2022.

Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations: A Mathematical Perspective*. Imperial College Press, 2010.

Jérôme Hénin, Tony Lelièvre, Michael R Shirts, Omar Valsson, and Lucie Delemotte. Enhanced sampling methods for molecular dynamics simulations. *Living Journal of Computational Molecular Science*, 4(1):1583, 2022.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants, 2023. URL https://arxiv.org/abs/2209.15571.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=XVjTT1nw5z.

Nicholas M Boffi and Eric Vanden-Eijnden. Deep learning probability flows and entropy production rates in active matter. *Proceedings of the National Academy of Sciences*, 121(25):e2318106121, 2024.

Nicholas M Boffi and Eric Vanden-Eijnden. Model-free learning of probability flows: Elucidating the nonequilibrium dynamics of flocking. *Physical Review Letters*, 133, 2025. DOI: 10.1103/1zqs-nbt7.

Suryanarayana Maddu, Victor Chardès, and Michael J Shelley. Learning stochastic processes with intrinsic noise from cross-sectional biological data. *Proceedings of the National Academy of Sciences*, 122(37):e2420621122, 2025. doi: 10.1073/pnas.2420621122.

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

Frank Noé, Christof Schütte, Eric Vanden-Eijnden, Lothar Reich, and Thomas R Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences*, 106(45):19011–19016, 2009.

Alexander M Berezhkovskii and Dmitrii E Makarov. Single-molecule test for Markovianity of the dynamics along a reaction coordinate. *The Journal of Physical Chemistry Letters*, 10(10): 2512–2517, 2019.

Baron Peters and Bernhardt L Trout. Obtaining reaction coordinates by likelihood maximization. *The Journal of Chemical Physics*, 125(5):054108, 2006.

Ao Ma and Aaron R Dinner. Automatic method for identifying reaction coordinates in complex systems. *The Journal of Physical Chemistry B*, 109(14):6769–6779, 2005.

Qianxiao Li, Bo Lin, and Weiqing Ren. Computing committor functions for the study of rare events using deep learning. *The Journal of Chemical Physics*, 151(5):054112, 2019.

Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. Solving for high-dimensional committor functions using artificial neural networks. *Research in the Mathematical Sciences*, 6(1):1–13, 2019.

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*, volume 145 of *Proceedings of Machine Learning Research*, pages 757–780. PMLR, 2022.

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.

Bo Lin and Weiqing Ren. Deep learning method for computing committor functions with adaptive sampling. *arXiv preprint arXiv:2404.06206*, 2024.

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.

Peter G Bolhuis, Christoph Dellago, and David Chandler. Reaction coordinates of biomolecular isomerization. *Proceedings of the National Academy of Sciences*, 97(11):5877–5882, 2000.

Weiqing Ren, Eric Vanden-Eijnden, Paul Maragakis, and Weinan E. Transition pathways in complex systems: Application of the finite-temperature string method to the alanine dipeptide. *The Journal of Chemical Physics*, 123(13):134109, 2005.

G. N. Ramachandran, C. Ramakrishnan, and V. Sasisekharan. Stereochemistry of polypeptide chain configurations. *Journal of Molecular Biology*, 7(1):95–99, 1963.

Douglas J Tobias, Glenn J Martyna, and Michael L Klein. Molecular dynamics simulations of a protein in the canonical ensemble. *The Journal of Physical Chemistry*, 97(49):12959–12966, 1993.

Paul E Smith. The alanine dipeptide free energy surface in solution. *The Journal of chemical physics*, 111(12):5568–5579, 1999.

David W. H. Swenson, Jan-Hendrik Prinz, Frank Noé, John D. Chodera, and Peter G. Bolhuis. Openpathsampling: A flexible, open framework for path sampling simulations. 1. basics. *Journal of Chemical Theory and Computation*, 15(2):813–836, 2019a. doi: 10.1021/acs.jctc.8b00626.

David W. H. Swenson, Jan-Hendrik Prinz, Frank Noé, John D. Chodera, and Peter G. Bolhuis. Openpathsampling: A flexible, open framework for path sampling simulations. 2. building and customizing path ensembles and sample schemes. *Journal of Chemical Theory and Computation*, 15(2):837–856, 2019b. doi: 10.1021/acs.jctc.8b00627.

Peter Eastman, Raimondas Galvelis, Raúl P Peláez, Charlles RA Abreu, Stephen E Farr, Emilio Gallicchio, Anton Gorenko, Michael M Henry, Frank Hu, Jing Huang, et al. Openmm 8: molecular dynamics simulation with machine learning potentials. *The Journal of Physical Chemistry B*, 128 (1):109–116, 2023.

Anthony T Bogetti, Hannah E Piston, Jeremy MG Leung, Chino C Cabalteja, Darian T Yang, Alex J DeGrave, Karl T Debiec, David S Cerutti, David A Case, W Seth Horne, et al. A twist in the road less traveled: The amber ff15ipq-m force field for protein mimetics. *The Journal of chemical physics*, 153(6), 2020.

Yu Shi, Shuxin Zheng, Guolin Ke, Yifei Shen, Jiacheng You, Jiyan He, Shengjie Luo, Chang Liu, Di He, and Tie-Yan Liu. Benchmarking graphormer on large-scale molecular modeling datasets. *arXiv preprint arXiv:2203.04810*, 2022. URL https://arxiv.org/abs/2203.04810.

Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.

Daniel M Zuckerman and Lillian T Chong. Weighted ensemble simulation: review of methodology, applications, and software. *Annual review of biophysics*, 46:43–57, 2017.

John D Russo, She Zhang, Jeremy MG Leung, Anthony T Bogetti, Jeff P Thompson, Alex J DeGrave, Paul A Torrillo, AJ Pratt, Kim F Wong, Junchao Xia, et al. Westpa 2.0: High-performance upgrades for weighted ensemble simulations and analysis of longer-timescale applications. *Journal of Chemical Theory and Computation*, 18(2):638–649, 2022.

Paul A Torrillo, Anthony T Bogetti, and Lillian T Chong. A minimal, adaptive binning scheme for weighted ensemble simulations. *The Journal of Physical Chemistry A*, 125(7):1642–1649, 2021.

A Connection to Transition Path Theory

In this appendix, we show that when the dynamics is Markovian, our framework recovers the objects from transition path theory (TPT). Consider an SDE of the form

$$dz_t = b(z_t)\,dt + \sqrt{2}\,\sigma(z_t)\,dW_t, \tag{19}$$

with stationary density $\rho(z)$ and diffusion tensor $D(z) = \sigma(z)\sigma^\top(z)$.

Proposition A.1 (Forward and backward committors)**.** *Let $q(z)$ denote the probability of reaching B before A starting from z, and $\tilde{q}(z)$ the probability of having come from A rather than B. These satisfy the boundary value problems*

$$0 = b \cdot \nabla q + D : \nabla\nabla q, \tag{20}$$

$$0 = (-b + 2D\nabla \log \rho) \cdot \nabla \tilde{q} + D : \nabla\nabla \tilde{q}, \tag{21}$$

with boundary conditions $q|_{\partial A} = \tilde{q}|_{\partial B} = 0$ and $q|_{\partial B} = \tilde{q}|_{\partial A} = 1$.

Proof. The forward committor $q(z) = \mathbb{P}_z(\tau_B < \tau_A)$ satisfies the backward Kolmogorov equation for the process (19). Let $\mathcal{L} = b \cdot \nabla + D : \nabla\nabla$ denote the generator. For $z \in \Omega_R$, the function q is harmonic with respect to $\mathcal{L}$, i.e., $\mathcal{L}q = 0$, with the stated boundary conditions.

The backward committor $\tilde{q}(z)$ is the probability that the time-reversed process came from A. Under time reversal, the SDE (19) with stationary measure ρ becomes

$$d\tilde{z}_t = \tilde{b}(\tilde{z}_t)\,dt + \sqrt{2}\,\sigma(\tilde{z}_t)\,dW_t, \tag{22}$$

where the process $\tilde{z}_t$ has the same law as z_{-t} and the reversed drift is $\tilde{b} = -b + 2D\nabla \log \rho$. The backward committor satisfies the backward Kolmogorov equation for this reversed process, giving (21). □

Proposition A.2 (Reactive density and current)**.** *The reactive density and current are given by*

$$\rho_R = C_R^{-1}\,\rho\, q\, \tilde{q}, \tag{23}$$

$$j_R = C_R^{-1}\,\rho\,[\tilde{q}\, D\nabla q - q\, D\nabla \tilde{q}]\,, \tag{24}$$

where $C_R = \int_{\Omega_R} \rho(z)\, q(z)\, \tilde{q}(z)\, dz$.

Proof. Using the Markovian property, a point $z \in \Omega_R$ belongs to a reactive trajectory if and only if the trajectory will reach B before A (probability $q(z)$) and came from A rather than B (probability $\tilde{q}(z)$). The reactive density is therefore proportional to $\rho(z)q(z)\tilde{q}(z)$, giving (23).

For the current, we use the probability flux formula. The flux of the stationary process is $j = \rho b - D\nabla\rho$. The reactive current counts only trajectories that are reactive, weighted by the probability of being so. Using the product rule and the committor equations (20)–(21), one can verify that (24) satisfies $\operatorname{div} j_R = 0$ with the correct boundary flux, and that $j_R = \rho_R u$ with u given below. □

Proposition A.3 (Current velocity)**.** *The current velocity is*

$$u(z) = D(z)\nabla \log q(z) - D(z)\nabla \log \tilde{q}(z). \tag{25}$$

For systems satisfying detailed balance, $\tilde{q} = 1 - q$, and this simplifies to

$$u(z) = D(z)\nabla \log \frac{q(z)}{1 - q(z)f} f d. \tag{26}$$

Proof. Dividing (24) by (23):

$$u = \frac{j_R}{\rho_R} = \frac{\tilde{q} D\nabla q - q D\nabla \tilde{q}}{q\tilde{q}} \tag{27}$$

$$= D\frac{\nabla q}{q} - D\frac{\nabla \tilde{q}}{\tilde{q}} \tag{28}$$

$$= D\nabla \log q - D\nabla \log \tilde{q}. \tag{29}$$

For detailed balance, the stationary measure is $\rho \propto e^{-\beta U}$ and the drift satisfies $b = -D\nabla U + \nabla \cdot D$. In this case, the forward and backward committors are related by $\tilde{q} = 1 - q$, which gives the stated simplification. □

Proposition A.4 (Doob transform)**.** *Reactive trajectories, i.e., trajectories of* (19) *conditioned on reaching B before A, evolve according to*

$$dz_t^R = b(z_t^R)\,dt + 2D(z_t^R)\nabla \log q(z_t^R)\,dt + \sqrt{2}\,\sigma(z_t^R)\,dW_t. \tag{30}$$

Proof. This is the Doob h-transform with $h = q$. For a diffusion with generator $\mathcal{L}$, conditioning on reaching a set where $h = 1$ before a set where $h = 0$ (with $\mathcal{L}h = 0$ in between) transforms the drift by adding $2D\nabla \log h$. Since q is harmonic for $\mathcal{L}$ with $q|_{\partial B} = 1$ and $q|_{\partial A} = 0$, the conditioned process has drift $b + 2D\nabla \log q$. □

B Examples of reactive path ensembles

Here, we describe some canonical examples of reactive path ensembles that arise in applications from chemistry and physics.

Overdamped Langevin Dynamics. Consider the overdamped Langevin equation

$$dx_t = -\nabla U(x_t)\,dt + \sqrt{2\beta^{-1}}\,dW_t, \tag{31}$$

where $U : \mathbb{R}^d \to \mathbb{R}$ is a C^2 potential such that $Z = \int_{\mathbb{R}^d} e^{-\beta U(x)} dx < \infty$, and $\beta = 1/(k_B T)$ is the inverse temperature with k_B the Boltzmann constant and T the temperature. The invariant density is $\rho(x) = Z^{-1} e^{-\beta U(x)}$.

Taking $z_t = x_t$, the reactive density and current are given by

$$\rho_R(x) = C_R^{-1}\, e^{-\beta U(x)} q(x)(1 - q(x)), \tag{32}$$

$$j_R(x) = C_R^{-1} \beta^{-1} e^{-\beta U(x)} \nabla q(x), \tag{33}$$

where $C_R = \int_{\mathbb{R}^d} e^{-\beta U(x)} q(x)(1 - q(x))\,dx$ is a normalization constant, and $q(x)$ is the *committor function*: the probability that the solution to (31), started from x, reaches B before A. It satisfies the backward Kolmogorov equation

$$-\nabla U(x) \cdot \nabla q(x) + \beta^{-1} \Delta q(x) = 0, \tag{34}$$

with boundary conditions $q|_{\partial A} = 0$ and $q|_{\partial B} = 1$. Note that because detailed balance holds, the backward committor is simply $1 - q(x)$.

The current velocity is therefore

$$u(x) = \frac{\beta^{-1} \nabla q(x)}{q(x)(1 - q(x))}. \tag{35}$$

Underdamped Langevin Dynamics. Consider the underdamped Langevin equation

$$dx_t = v_t\,dt, \tag{36}$$

$$dv_t = -\nabla U(x_t)\,dt - \gamma v_t\,dt + \sqrt{2\beta^{-1} m^{-1} \gamma}\,dW_t, \tag{37}$$

where $\gamma > 0$ is the friction coefficient and m is the (diagonal) mass matrix. The invariant density is $\rho(x, v) = Z_H^{-1} e^{-\beta H(x,v)}$, where $H(x, v) = U(x) + \frac{1}{2} v^\top m\, v$ is the Hamiltonian and $Z_H = \int_{\mathbb{R}^d \times \mathbb{R}^d} e^{-\beta H(x,v)} dx dv$ is the partition function.

Taking $z_t = (x_t, v_t)$, the reactive density and current are given by

$$\rho_R(x, v) = C_R^{-1} e^{-\beta H(x,v)} q(x, v)(1 - q(x, -v)), \tag{38}$$

$$j_R(x, v) = C_R^{-1} e^{-\beta H(x,v)} \begin{pmatrix} v \\ -\nabla U(x) - g_R(x, v) \end{pmatrix} \tag{39}$$

where the normalization constant is

$$C_R = \int_{\mathbb{R}^d \times \mathbb{R}^d} e^{-\beta H(x,v)} q(x, v)(1 - q(x, -v))\,dx\,dv, \tag{40}$$

and we defined

$$g_R(x,v) = \gamma\beta^{-1}m^{-1}[(1-q(x,-v))\nabla_v q(x,v) - q(x,v)\nabla_v q(x,-v)]. \tag{41}$$

The committor $q(x,v)$ now depends on both position and velocity, and satisfies the backward Kolmogorov equation

$$v\cdot\nabla_x q - \nabla U(x)\cdot\nabla_v q - \gamma v\cdot\nabla_v q + \beta^{-1}m^{-1}\gamma\Delta_v q = 0, \tag{42}$$

with boundary conditions $q|_{\partial A} = 0$ and $q|_{\partial B} = 1$.

Remark B.1. The factor $(1-q(x,-v))$ in (38), with velocity reversed, reflects the time-reversal symmetry of reactive trajectories: under time reversal, velocities change sign, so a trajectory reactive from A to B with velocity v corresponds to a trajectory from B to A with velocity $-v$.

Non-Markovian Case: Position-Only Observation. The framework applies equally to non-Markovian processes. As an example, consider the underdamped Langevin dynamics (36)–(37), but now take $z_t = x_t$ (position only), marginalizing over the velocity. The reactive density becomes

$$\rho_R(x) = C_R^{-1}e^{-\beta U(x)}\int e^{-\frac{\beta}{2}v^\top mv}q(x,v)(1-q(x,-v))\,dv, \tag{43}$$

and the current velocity is

$$u(x) = \frac{\int v\,e^{-\frac{\beta}{2}v^\top mv}q(x,v)(1-q(x,-v))\,dv}{\int e^{-\frac{\beta}{2}v^\top mv}q(x,v)(1-q(x,-v))\,dv}. \tag{44}$$

Remark B.2. Equations (43)–(44) are exact expressions for the reactive density and current velocity of the non-Markovian process $z_t = x_t$. No approximation is involved; the velocity marginalization is built into Definition 2.2.

C Proof of Theorem 2.5

Proof. Reactive trajectories in the full space evolve according to the Doob-transformed dynamics (18):

$$dx_t^R = b(x_t^R)\,dt + 2D(x_t^R)\nabla\log q(x_t^R)\,dt + \sqrt{2}\,\sigma(x_t^R)\,dW_t. \tag{45}$$

Applying Itô's lemma to $z_t = \phi(x_t^R)$, we obtain

$$dz_t = \tilde{b}(x_t^R)\,dt + \sqrt{2}\,\nabla\phi(x_t^R)^\top\sigma(x_t^R)\,dW_t, \tag{46}$$

where

$$\tilde{b}(x) = (b(x) + 2D(x)\nabla\log q(x))\cdot\nabla\phi(x) + D(x):\nabla\nabla\phi(x). \tag{47}$$

To identify the current velocity, we use Definition 2.2 with a test function $\psi:\mathbb{R}^n\to\mathbb{R}^n$. The Stratonovich integral satisfies

$$\int_0^\tau \psi(z_t)\circ dz_t = \int_0^\tau \psi(z_t)\cdot dz_t + \frac{1}{2}\int_0^\tau d\langle\psi(z),z\rangle_t, \tag{48}$$

where the first term on the right is an Itô integral. The quadratic covariation is

$$d\langle\psi(z),z\rangle_t = 2\nabla\psi(z_t):\nabla\phi(x_t^R)^\top D(x_t^R)\nabla\phi(x_t^R)\,dt. \tag{49}$$

Substituting the SDE for dz_t into the Itô integral:

$$\int_0^\tau \psi(z_t)\cdot dz_t = \int_0^\tau \psi(z_t)\cdot\tilde{b}(x_t^R)\,dt + \sqrt{2}\int_0^\tau \psi(z_t)\cdot\nabla\phi(x_t^R)^\top\sigma(x_t^R)\,dW_t. \tag{50}$$

The stochastic integral is a martingale, so its expectation vanishes. Therefore,

$$\mathbb{E}\left[\int_0^\tau \psi(z_t)\circ dz_t\right] = \mathbb{E}\left[\int_0^\tau \psi(z_t)\cdot\tilde{b}(x_t^R)\,dt\right] + \mathbb{E}\left[\int_0^\tau \nabla\psi(z_t):\nabla\phi(x_t^R)^\top D(x_t^R)\nabla\phi(x_t^R)\,dt\right]. \tag{51}$$

By Definition 2.2, this equals

$$T_R \int_{\mathbb{R}^n} [\psi(z) \cdot u(z) + \nabla\psi(z) : D_z(z)]\, \rho_R(z)\, dz, \tag{52}$$

where $D_z(z) = \mathbb{E}\left[\nabla\phi(x)^\top D(x)\nabla\phi(x) \,|\, \phi(x) = z\right]$ is the projected diffusion tensor. Comparing with Definition 2.2, the term involving $\nabla\psi$ cancels between the Stratonovich correction and the projected diffusion, and we identify

$$u(z) = \mathbb{E}\big[\, (b(x) + 2D(x)\nabla \log q(x)) \cdot \nabla\phi(x) + D(x) : \nabla\nabla\phi(x) \,\big|\, \phi(x) = z\big], \tag{53}$$

where the expectation is over $\rho_R(x)$ conditional on $\phi(x) = z$. □

D Flux Matching Implementation

We provide here further practical consideration in implementation of Reactive Flux Matching.

D.1 The time lag Δt

The time lag Δt controls a bias–variance tradeoff in the Stratonovich approximation. Small Δt gives a more accurate increment but yields noisier displacements Δz_k, since the diffusive component dominates on short timescales; large Δt reduces this noise but introduces bias when u varies on scales below Δt. In practice, Δt should be small compared to the timescale on which u varies, yet large enough that displacements remain well-resolved relative to noise. The centered form $\Delta z_k = \frac{1}{2}(z^{(k)}_{t_k+\Delta t} - z^{(k)}_{t_k-\Delta t})$ approximates the Stratonovich integral to second order. One can also sample Δt randomly within a range or anneal it during training—starting larger for stability and decreasing for accuracy.

D.2 Boundary regularization for the potential

The bulk term $|\nabla h|^2_D$ in $\mathcal{L}_h$ is well-behaved, but the boundary terms $2h(z_0) - 2h(z_\tau)$ have constant gradients that push h toward $-\infty$ on ∂A and $+\infty$ on ∂B. This is consistent with the theoretical limit $h = \log[q/(1-q)]$ under detailed balance but is a recipe for exploding gradients in practice. We therefore replace these boundary terms by a bounded logistic surrogate, the cross-entropy loss

$$\mathcal{L}_{\mathrm{CE}}(h) = -\frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}}\Big[\log\sigma\big(-h(z_0^{(k)})\big) + \log\sigma\big(h(z_{\tau_k}^{(k)})\big)\Big], \tag{54}$$

where σ denotes the sigmoid. The minimization direction is the same as the original boundary penalty, with $h(z_0)$ driven low and $h(z_\tau)$ driven high, but the gradients saturate as $|h|$ grows, preventing runaway updates. The loss minimized in practice is the bulk term plus $\mathcal{L}_{\mathrm{CE}}$, as shown in Algorithm 1.

D.3 Neural network parameterization

We parameterize the velocity field and the potential as neural networks $u_\theta : \mathbb{R}^n \to \mathbb{R}^n$ and $h_\phi : \mathbb{R}^n \to \mathbb{R}$. For molecular systems, the inputs are either the full 3D atomic coordinates or internal dihedral angles; in the Cartesian setting, translational and rotational symmetries are handled via mean-centering and rotational data augmentation. Full architectural details are deferred to Section F.1.

E Test system

E.1 Müller-Brown

The Müller-Brown system follows the energy function given by:

$$\begin{aligned} U(x,y) = &-200\cdot\exp\left(-(x-1)^2 - 10y^2\right) \\ &-100\cdot\exp\left(-x^2 - 10\cdot(y-0.5)^2\right) \\ &-170\cdot\exp\left(-6.5\cdot(0.5+x)^2 + 11\cdot(x+0.5)\cdot(y-1.5) - 6.5\cdot(y-1.5)^2\right) \\ &+15\cdot\exp\left(0.7\cdot(1+x)^2 + 0.6\cdot(x+1)\cdot(y-1) + 0.7\cdot(y-1)^2\right). \end{aligned} \tag{55}$$

The basins A and B are circles of radius 0.1 centered at [-1.5, 0.9] and [-0.5, 1.7] respectively

Overdamped Langevin dynamics was simulated using an Euler-Maruyama integrator, while underdamped dynamics was simulated using the BAOAB integrator. Both dynamics used a time step of $1e^-4$ and $\beta^{-1} = 12.5$. The underdamped dynamics also used a friction coefficient of $\gamma = 10$.

TPS shooting moves were implemented in a straightforward manner. Shooting points were selected uniformly along each trajectory, and the shooting direction was randomly chosen. If the backward direction was chosen, velocities were reversed. If a reactive trajectory was generated from the shooting point, it was accepted or rejected based on the Metropolis criterion, which is given by the ratio of the new and old transition path lengths. A total of 1,000 transition paths were generated for both overdamped and underdamped dynamics

Ground truth for Müller–Brown. For overdamped Langevin dynamics $\mathrm{d}x_t = -\nabla U(x_t)\,\mathrm{d}t + \sqrt{2\beta^{-1}}\,\mathrm{d}W_t$ on the Müller–Brown potential ($\beta = 0.08$), we compute the exact committor q by solving the backward Kolmogorov equation $\nabla \cdot (e^{-\beta U}\nabla q) = 0$ with the finite element method. The equation is solved on the box $[-1.5, 0.9] \times [-0.5, 1.7]$ with the two states cut out as disks of radius 0.15, with $q = 0$ on ∂A, $q = 1$ on ∂B, and zero flux on the outer edges. We use the self-adjoint, Boltzmann-weighted weak form, which keeps the discrete solution within $[0, 1]$. We discretise it with P_2 elements on a mesh of 30k triangles (61k degrees of freedom) using FEniCS. The computed rate constant changes by about 10^{-5} in relative terms under mesh refinement.

From q we obtain the mean velocity of reactive trajectories predicted by transition-path theory, $V = J_{AB}/\rho_R = \beta^{-1}\nabla q/\big(q(1-q)\big)$, using an element-local gradient reconstruction to preserve relative accuracy near the states. We compare the learned ∇h with V at the 376,946 frames of the training transition paths that satisfy $10^{-3} < q < 1 - 10^{-3}$, where V is well defined. We report the relative RMSE, $\|\nabla h - V\|/\|V\|$ (root mean square over frames), and the mean cosine similarity between the two fields. ∇h is used exactly as the network outputs it, with no rescaling.

E.2 Alanine Dipeptide

Transition paths for Alanine Dipeptide were generated using the OpenPathSampling Python package with the OpenMM (Eastman et al., 2023) engine and the LangevinMiddleIntegrator. The system was simulated with the CHARMM27 force field at a time step of 1 fs, friction coefficient of 1/ps and hydrogen bond restaints. Molecular states are stored every 10 fs of simulation. Basins A and B follow the same definitions as in Mitchell and Rotskoff (2024) with squares of side 10° centered at $[\varphi, \psi] = [-150^\circ, 170^\circ]$ and $[90^\circ, 50^\circ]$, respectively. A total of 100,000 transition paths were generated for this system. We conduct a dataset ablation with dataset sizes of both 10,000 and 100,000 (Section G.4).

E.3 AIB9

Similar to ADP, the reactive path ensemble for AIB9 is also generated using the OpenPathSampling Python package with the OpenMM engine and the LangevinMiddleIntegrator. The system is simulated with the Amber15-IPQ force field (Bogetti et al., 2020) in vaccuum at 500K with time steps of 1fs, friction coefficients of 1/ps and Hbond constraints. Basins A and B are defined as squares of side 10° centered at $[\varphi, \psi] = [-60^\circ, -40^\circ]$ and $[60^\circ, 40^\circ]$, respectively. Simulation frames are saved every 100fs. We generate a dataset of 100,000 paths for AIB9 as well.

F Experimental Details

F.1 Model Architecture

Muller Brown: Current velocity model for the Muller Brown is a simple MLP model with 5 layers, hidden dimension of 128 and SiLU activation functions

ADP $\mathbb{R}^{3N}$: Alanine Dipeptide is featurized through a simple atom typing scheme where all the atom types are made distinguishable. The velocity model is parameterized as a non-equivariant transformer. Therefore the node features also contain an embedding of the (mean-centered) position.

The attention coefficients A are calculated as

$$A = \sigma\Big(\frac{QK^{\mathsf{T}}}{\sqrt{d}} + B\Big) \tag{56}$$

where B is a learnable bias computed from pairs of atoms. This pairwise attention bias takes as input the atom types and the interatomic distance. The overall architecture is heavily inspired by the Graphormer architecture (Shi et al., 2022). The transformer consists of 4 self-attention layers, each with 16 attention heads. The hidden dimension is 512, and layer normalization is applied both before and after the attention blocks. The final potential prediction is produced by an MLP applied to a sum of the hidden node features.

ADP and AIB9 Dihedrals. The dihedral angles for ADP and AIB9 are bounded in $[-\pi, \pi)$ and exhibit periodic structure that must be respected by the input representation. To handle this, we lift each dihedral x into a continuous, differentiable embedding by concatenating $[\sin(x),\, \sin(x + \pi/4)]$. This featurization removes the discontinuity at the $\pm\pi$ boundary and ensures non-vanishing gradients across the full angular domain, since the two sinusoids are phase-shifted so that their derivatives never vanish simultaneously. Both dihedral models are 4-layer MLPs with SiLU activations.

F.2 Training Details

Training data and time-step convention. Across all systems, the current-velocity models are trained on triplets $(x_{t-\Delta t},\, x_t,\, x_{t+\Delta t})$ sampled from transition-path trajectories. The training time step Δt is chosen such that the central finite-difference velocity estimate

$$\frac{x_{t+\Delta t} - x_{t-\Delta t}}{2\Delta t} \;\sim\; \mathcal{O}(1)$$

is well-scaled for optimization, avoiding the vanishing- and exploding-target regimes that arise for very small or very large Δt. All models are optimized with AdamW. At the end of each epoch, we sample 100 initial points from the reactive ensemble and integrate the learned velocity field both forward and backward in time to generate full trajectories (Section F.3); the checkpoint achieving the highest fraction of trajectories that successfully connect the two endpoint states is retained for inference. Below, we report only the system-specific hyperparameters that deviate from this shared protocol.

Müller–Brown. We use $\Delta t = 10^{-4}$ and train for 25 epochs at a constant learning rate of 10^{-3}.

ADP, Cartesian coordinates ($\mathbb{R}^{3N}$). Triplets are drawn from saved trajectories with consecutive configurations separated by 10 fs (every other frame in the saved trajectories), and we set $\Delta t = 10^{-3}$ to satisfy the scaling condition above. Training runs for 400 epochs, with a linear learning-rate warmup to 5×10^{-4} over 5,000 steps followed by cosine annealing to 10^{-5} over 100,000 steps.

ADP, dihedral coordinates. Configurations are again separated by 10 fs, and we find that $\Delta t = 10^{-1}$ satisfies the scaling condition. Training runs for 1,000 epochs at a constant learning rate of 5×10^{-4}, with a `ReduceLROnPlateau` scheduler that decays the learning rate by a factor of 0.1 whenever the validation loss has not improved for 30 consecutive epochs.

AIB9, backbone dihedrals. Configurations are separated by 100 fs; all remaining hyperparameters match the ADP dihedral setup above.

F.3 Generation of Flow Lines

To generate a flow line, we sample an initial point x_0 from the reactive ensemble, $x_0 \sim \rho_{\mathcal{R}}$, and integrate the learned velocity field independently in the forward and backward directions:

$$x_s^+ = x_0 + \int_0^s u(x_r^+)\, dr, \quad s \in [0, \tau_+], \qquad x_s^- = x_0 - \int_0^s u(x_r^-)\, dr, \quad s \in [0, \tau_-], \tag{57}$$

where the forward and backward stopping times are

$$\tau_+ := \min\{s \geqslant 0 \mid x_s^+ \in \partial B\} \wedge T_{\max}, \qquad \tau_- := \min\{s \geqslant 0 \mid x_s^- \in \partial A\} \wedge T_{\max}. \tag{58}$$

The forward integration terminates when the trajectory first enters the product boundary ∂B, and the backward integration terminates when it first enters the reactant boundary ∂A; in either case, integration also halts if the maximum time $T_{\max}$ is exceeded. The two segments are concatenated to form a complete flow line connecting ∂A to ∂B through x_0. The boundary shells ∂A and ∂B used to terminate inference-time integration are taken larger than those used to define the training data, since the learned velocity is not well-defined at the boundaries of A and B.

Müller–Brown. Integration uses an explicit Euler scheme with step size $\Delta t = 10^{-4}$.

ADP, Cartesian coordinates ($\mathbb{R}^{3N}$). Integration uses a fourth-order Runge–Kutta solver with $\Delta t = 10^{-4}$. After each step, a post-processing routine enforces the same hydrogen-bond constraints applied during MD data generation, ensuring consistency with the constrained dynamics and improving sample quality in practice.

ADP and AIB9, dihedral coordinates. Integration uses a fourth-order Runge–Kutta solver with $\Delta t = 10^{-2}$. After each step, dihedral angles are wrapped back into $[-\pi, \pi)$ to preserve the periodic structure of the coordinate space.

F.4 Metrics

We evaluate the quality of generated flow lines along two complementary axes: whether trajectories successfully traverse the transition, and how closely the resulting ensemble of configurations matches the reference distribution of reactive paths.

Completion rate. A generated flow line is considered *complete* if it connects the reactant and product states within the maximum integration time $T_{\max}$, i.e., if both stopping times τ_+ and τ_- are attained before reaching $T_{\max}$. The completion rate is the fraction of sampled initial points $x_0 \sim \rho_{\mathcal{R}}$ for which this condition is met.

Torsional Wasserstein-2 distance. To assess the distributional fidelity of completed flow lines, we compute the Wasserstein-2 ($\mathcal{W}_2$) distance between the empirical distribution of generated configurations μ and a reference distribution ν drawn from MD reactive trajectories, with both distributions represented through backbone torsion angles. Because torsions live on the periodic domain $[-\pi, \pi)$, we use the circular ground distance

$$d_{\mathrm{circ}}(\theta, \theta') := \min\big(|\theta - \theta'|,\ 2\pi - |\theta - \theta'|\big),$$

extended coordinate-wise to the product of circles $\mathbb{T}^d$, and define the torsional $\mathcal{W}_2$ as

$$\mathcal{W}_2^{\mathrm{tors}}(\mu, \nu) := \left(\inf_{\gamma \in \Gamma(\mu,\nu)} \int d_{\mathrm{circ}}(\theta, \theta')^2 \, d\gamma(\theta, \theta')\right)^{1/2},$$

where $\Gamma(\mu, \nu)$ denotes the set of couplings between μ and ν. The infimum is computed exactly by solving the discrete optimal-transport linear program with the POT library (Flamary et al., 2021). Lower $\mathcal{W}_2^{\mathrm{tors}}$ values indicate that the model reproduces the geometry of the reactive ensemble more faithfully.

G Additional Results

G.1 Flow line Generation results

Once integrated, the learned current velocity should produce paths whose statistics match those of the reactive path ensemble. We therefore evaluate generated trajectories using the following metrics, in addition to visualizing the level sets of h:

Table 2: **Quantitative results across systems.** Values are computed over 256 trajectories. In both systems, the generated trajectories closely reproduce the values observed in the reactive trajectory ensemble.

System	Method	Features	Completion rate	$\mathbb{T}$-W_2
ADP	TPS	$\mathbb{R}^{3N}$	1	0.0875 ± 0.006
ADP	FM	$\mathbb{R}^{3N}$	0.7734	0.6453 ± 0.016
ADP	FM	Dihedrals	0.9804	0.7220 ± 0.017
AIB9	TPS	$\mathbb{R}^{3N}$	1	0.1011 ± 0.034
AIB9	FM	Backbone Dihedrals	0.7578	0.9137 ± 0.007

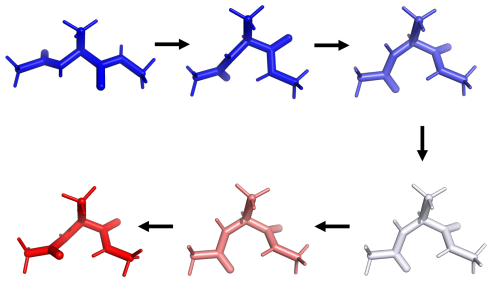

Figure 7: **ADP flow trajectory.** Snapshots of generated trajectory colored by the value of h, from dark blue at A and red at B. The transition state is correctly identified as the white snapshot.

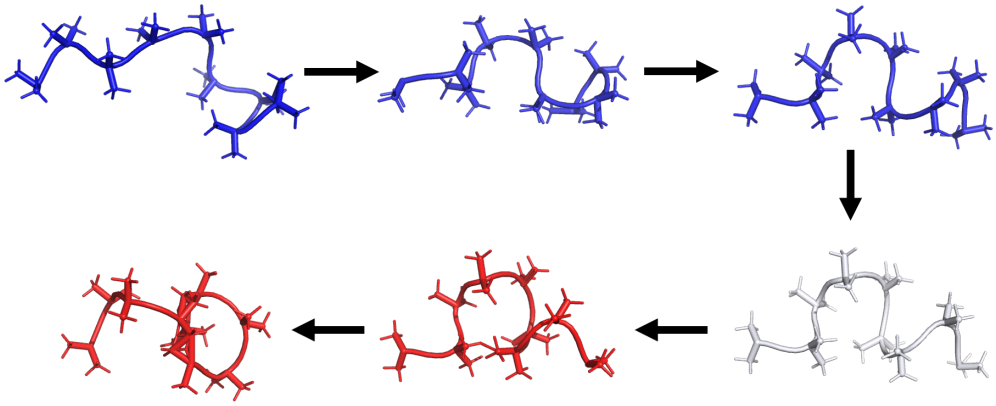

Figure 8: **AIB9 flow trajectory.** Snapshots of generated trajectory colored by the value of ϕ. The start (A) and end (B) states can be visualized as dark blue and red and the transition state, identified correctly, is white.

G.2 Example Transition path structures

G.3 Rate Constant Estimation for the ADP $C_7^{\text{eq}} \to C_7^{\text{ax}}$ Transition

Weighted ensemble setup. To assess the utility of the learned potential h as a CV, we compute the rate constant for the $C_7^{\text{eq}} \to C_7^{\text{ax}}$ transition in alanine dipeptide (ADP) using the weighted ensemble (WE) method (Huber and Kim, 1996; Zuckerman and Chong, 2017), as implemented in the

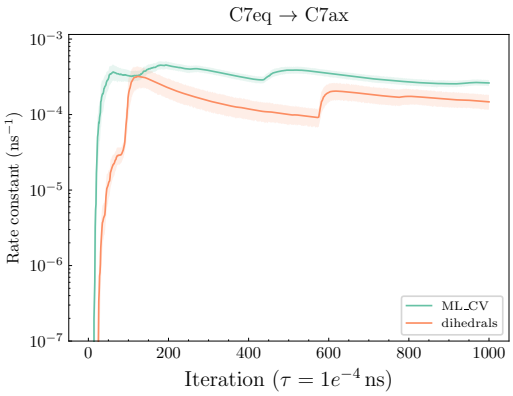

Figure 9: **Rate constant estimation for the $C_7^{\mathrm{eq}} \to C_7^{\mathrm{ax}}$ transition in ADP.** We compare WE rate estimates obtained using the learned potential h as the collective variable against those obtained using the backbone-dihedral CV ξ_{dih}. Both runs use the WESTPA 2.0 implementation with minimal adaptive binning. Using h yields earlier first-passage events, faster convergence of the running rate estimate, and tighter confidence intervals, indicating that h provides a more informative collective variable for driving the transition.

WESTPA 2.0 package (Russo et al., 2022), and compare against a CV constructed directly from the backbone dihedrals.

Weighted ensemble partitions the chosen CV space into bins and maintains a target number of trajectory replicas per bin via splitting and merging, while preserving unbiased trajectory weights. The flux of weight into the product state C_7^{ax} provides an unbiased estimator of the forward rate constant. To avoid hand-tuning bin boundaries for each CV, we use the minimal adaptive binning (MAB) scheme (Torrillo et al., 2021), which dynamically places bins along the CV based on the leading and trailing replicas at each iteration and concentrates resolution near the progressing front of the ensemble. We initialize all replicas in C_7^{eq} and run WE until the running rate estimate stabilizes.

Collective variables. We compare two choices of CV for MAB binning:

- **Backbone-dihedral CV.** Following standard practice for ADP, we define a scalar progress coordinate as the maximum angular displacement along the (ϕ, ψ) backbone dihedrals from the reactant minimum:

$$\xi_{\mathrm{dih}}(x) := \max\big(d_\phi(\phi(x), \phi_{\mathrm{ax}}),\ d_\psi(\psi(x), \psi_{\mathrm{ax}}) \big),$$

  where d_ϕ, d_ψ denote angular distances on the circle and $(\phi_{\mathrm{ax}}, \psi_{\mathrm{ax}})$ is the C_7^{ax} minimum. This CV is a natural baseline since (ϕ, ψ) are the canonical slow degrees of freedom for ADP.

- **Learned-potential CV.** We use $h_\theta(x)$ directly as a scalar CV, with MAB operating between the values of h at the two endpoint states.

For both CVs, we use 40 MAB bins with a target of 10 walkers per bin and a resampling interval of 100 fs.

Results. Figure 9 compares the running WE rate estimates obtained with the two CVs. Using h as the binning coordinate produces earlier first-passage events, faster convergence of the running rate estimate, and tighter confidence intervals than the backbone-dihedral CV. We attribute this improvement to two properties of h: it is a globally smooth, scalar-valued progress coordinate that aligns with the dominant transition direction, whereas the dihedral-based CV reduces a two-dimensional periodic space to a single scalar via a $\max$ operation that can cause replicas with different mechanistic character to fall into the same bin. Together, these results suggest that h serves as a more informative CV for driving the $C_7^{\mathrm{eq}} \to C_7^{\mathrm{ax}}$ transition in WE simulations.

G.4 DATASET SIZE ABLATION

The difference in performance on training on 100K vs 10K trajectories is reported in the table below:

Table 3: **Quantitative results across systems.** Values are computed over 256 trajectories. In both systems, the generated trajectories closely reproduce the values observed in the reactive trajectory ensemble.

System	Method	No. of trajectories	Features	Completion rate	$\mathbb{T}$-W_2
ADP	TPS	100K	$\mathbb{R}^{3N}$	1	0.0875 ± 0.006
ADP	FM	10K	$\mathbb{R}^{3N}$	0.6250	0.3683 ± 0.009
ADP	FM	100K	$\mathbb{R}^{3N}$	0.7734	0.6453 ± 0.016
ADP	FM	10K	Dihedrals	0.9531	0.8573 ± 0.005
ADP	FM	100K	Dihedrals	0.9804	0.7220 ± 0.017
AIB9	TPS	100K	$\mathbb{R}^{3N}$	1	0.1011 ± 0.034
AIB9	FM	10K	Backbone Dihedrals	0.7343	1.096 ± 0.014
AIB9	FM	100K	Backbone Dihedrals	0.7578	0.9137 ± 0.007