

Molecular Free Energies, Rates, and Mechanisms from Data-Efficient Path Sampling Simulations

Published as part of *Journal of Chemical Theory and Computation* virtual special issue “Machine Learning and Statistical Mechanics: Shared Synergies for Next Generation of Chemical Theory and Computation”.

Gianmarco Lazzeri, Hendrik Jung, Peter G. Bolhuis, and Roberto Covino*



Cite This: *J. Chem. Theory Comput.* 2023, 19, 9060–9076



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: Molecular dynamics is a powerful tool for studying the thermodynamics and kinetics of complex molecular events. However, these simulations can rarely sample the required time scales in practice. Transition path sampling overcomes this limitation by collecting unbiased trajectories and capturing the relevant events. Moreover, the integration of machine learning can boost the sampling while simultaneously learning a quantitative representation of the mechanism. Still, the resulting trajectories are by construction non-Boltzmann-distributed, preventing the calculation of free energies and rates. We developed an algorithm to approximate the equilibrium path ensemble from machine-learning-guided path sampling data. At the same time, our algorithm provides efficient sampling, mechanism, free energy, and rates of rare molecular events at a very moderate computational cost. We tested the method on the folding of the mini-protein chignolin. Our algorithm is straightforward and data-efficient, opening the door to applications in many challenging molecular systems.



I. INTRODUCTION

Molecules are everywhere. They constitute biological structures, chemical reactions, and materials. Molecules are also inherently dynamical. Molecular dynamics (MD) simulations are accurate physics-based models that give access to the time evolution of molecular systems with atomic resolution, including complex biomolecules, materials, and chemical reactions.^{1–3} MD can provide a thermodynamic, kinetic, and mechanistic characterization of a wide range of phenomena, such as conformational changes,^{4,5} folding,^{6,7} ligand binding,⁸ oligomerization,^{9,10} protein–membrane interactions,¹¹ nucleation,¹² and ion permeation.¹³

Ideally, long MD simulations would produce equilibrium trajectories extensively exploring the configuration space of a molecular system.^{14,15} In these simulations, the trajectories would enter metastable states and spend the most time there. Rarely, they would go on brief excursions in the transition region outside metastable states. Even more rarely, those excursions would result in an actual transition, crossing an energy barrier to reach an alternative state.¹⁶

Only by repeatedly transitioning between metastable states would these long trajectories sample the stationary Boltzmann distribution that describes the system’s thermodynamics.¹⁷ One could then count how often the trajectories undergo a transition to obtain the reaction rate constants that describe

the kinetics. Moreover, one could isolate the trajectory segments that connect two metastable states—the transition paths (TPs)—and collect them in the transition path ensemble (TPE), which describes the mechanism of the transition.^{18,19}


However, many interesting phenomena are rare events, stochastic transitions occurring on exponentially longer time scales than the MD integration time step.²⁰ This makes them practically inaccessible by typical MD simulations.²¹ Many enhanced sampling strategies have been developed to overcome these challenges without resorting to “brute-force” simulations.²² One popular solution is to apply an unphysical biasing force that steers the dynamics and enhances the exploration of the configuration space.²³ These methods require accurate prior knowledge of the system—a low-dimensional feature representation or, ideally, a reaction coordinate.^{22,24} Additionally, the bias distorts the system’s dynamics, complicating their reconstruction.²⁵

Received: July 28, 2023

Revised: October 24, 2023

Accepted: October 24, 2023

Published: November 21, 2023






Figure 1. Illustration of the algorithm. We simulate transition paths connecting two metastable states (top center) with AIMMD, which also generates excursions (top left and right) in the process. Meanwhile, we perform equilibrium simulations around the states (bottom left and right). Configurations in the trajectories generated by path sampling do not follow the Boltzmann distribution (dashed profile) by construction. After reweighting them based on a machine-learned committor $p_B(x)$, we can merge all trajectories and recover the equilibrium kinetics and thermodynamics of the entire configuration space (bottom center).

Alternative techniques do not bias the dynamics but instead, exploit a judicious initialization of the sampling. In other words, the initial configurations are not sampled from the Boltzmann distribution but follow alternative criteria to facilitate sampling in specific regions of the configuration space. Adaptive sampling,²⁶ swarm of trajectories,²⁷ the weighted ensemble approach,²⁸ and transition path sampling²⁹ (TPS) all belong to this category. An advantage of these methods is that they provide not only ensembles of configurations but entire trajectories that are the realization of the system's dynamics. Therefore, these are perfectly reasonable trajectories that can be reweighted to take into account the biased selection of the starting configuration.^{30,31} In particular, TPS is a rigorous way to collect the TPs of a hypothetical extensive equilibrium trajectory.^{18,32} By avoiding sampling in the metastable states, TPS can be computationally very efficient²⁹ and has enabled the characterization of several molecular processes.^{33–38} The resulting TPE contains mechanistic information.³⁹ In practice, generating TPs with high efficiency can be challenging, limiting the effectiveness of TPS schemes.⁴⁰

Recently, we integrated deep learning with TPS to automatize and significantly improve the sampling of TPs in complex molecular systems.⁴¹ We called this approach “AI for molecular mechanism discovery” (AIMMD). In AIMMD, a neural network controls TPS and boosts the production of TPs connecting two states; at the same time, it autonomously learns the transition mechanism by learning the committor.⁴¹ The committor is the ideal reaction coordinate describing a general stochastic transition between two states, i.e., an optimal one-dimensional projection that quantitatively monitors the progress along a transition.⁴² However, TPS and AIMMD do not directly provide free energy profiles and rates.⁴⁰ Configurations in the TPE are by construction, not Boltzmann-distributed.

Let us focus on a transition between two states, A and B, even though the following considerations are general. Thinking again at an ideal extensive trajectory, we can collect all trajectory segments that leave a state and end as soon as they reach a state.⁴³ We can classify them according to their origin and destination: A-to-A, A-to-B, B-to-A, and B-to-B. The TPE contains only those segments that connect the two states (A-to-B and B-to-A); it does not contain any excursions (A-to-A and B-to-B), those trajectory segments that temporarily leave a state and come back to it before reaching another one. However, these excursions significantly contribute to the Boltzmann distribution in the transition region between A and B. Regarding kinetics, the TPE gives access to the transition path time (the average duration of a TP) but not to the waiting times in the metastable states, which dominate the transition rate.¹⁷

An effective strategy for obtaining the free energy is to use short unbiased simulations. One solution is offered by transition interfaces sampling (TIS),^{44,45} a popular extension of TPS. Rogal et al. introduced the reweighted path ensemble (RPE), which reweights the individual TIS trajectories to approximate the free energy in the transition region.⁴³ However, TIS is computationally relatively expensive and depends effectively on the knowledge of a reasonably good reaction coordinate. Recently, Brotzakis and Bolhuis proposed an algorithm to approximate the RPE by “waste-recycling” TPS simulations.^{46,47}

In this study, inspired by concepts introduced in refs 43 and 46, we propose a new computational scheme that enables us to simultaneously access mechanisms, thermodynamics, and kinetics of stochastic rare event transitions. We show that by building on the committor estimated by AIMMD, we can estimate free energy profiles and rates from just a few short simulations. We also developed a procedure for extending the estimate of the free energy in the transition region to the metastable states—hence to the whole accessible configuration

space—with minimal additional computational cost (Figure 1). We illustrated our method on two benchmark 2-dimensional systems with high energy barriers and multiple reactive channels and with the folding of the mini-protein chignolin.⁴⁸ In all cases, we successfully determined the complete free energy profiles and rates in a small fraction of the computational resources required for a typical MD simulation.

The paper is organized as follows. In section II, we provide a detailed explanation of the algorithm. In section III, we introduce the studied systems and present computational methods. In section IV, we illustrate our results with a particular emphasis on the computational performance. We end with concluding remarks and a future outlook.

II. THEORY

II.A. AIMMD Sampling Scheme. For the sake of completeness, we briefly summarize the theory behind the AIMMD sampling scheme.⁴¹ Let us consider a system with two metastable states A and B separated by an energy barrier.^{17,18} We assume that the system's dynamics (in the full configuration space) are Markovian. The system is described by its configuration x . A trajectory, or path, is a sequence $\mathbf{x} = \{x_0, \dots, x_t, \dots, x_L\}$ sampled at regular time intervals of length $L[\mathbf{x}]$.

TPS is a Markov chain Monte Carlo technique that generates a series of paths $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(i)}$ connecting A and B.^{29,49} In this work, $\mathbf{x}^{(i)}$ is the trajectory simulated at step i , and $\mathbf{y}^{(i)}$ is the last accepted path after that step. While the $\mathbf{x}^{(i)}$ values are always different, the $\mathbf{y}^{(i)}$ values can repeat in the case of rejection. Therefore, $\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(i)}\}$ is a subset of $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(i)}\}$. With an increasing number of steps, the chain of paths converges to the TPE (the equilibrium ensemble of all the system's TPs):

$$\mathcal{P}_{\text{TP}} \approx \{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(n)}\} \quad (1)$$

Element $\mathbf{y}^{(i)}$ in the chain is generated from $\mathbf{y}^{(i-1)}$. First, we select the shooting point $x_{\text{sp}}^{(i)}$ from $\mathbf{y}^{(i)}$. Then, we produce a trial path $\mathbf{x}^{(i)}$ by a two-way shooting move:⁵⁰ we evolve two subtrajectories from $x_{\text{sp}}^{(i)}$ backward and forward in time until they hit either A or B, time-reverse the former, and join the two subtrajectories together. To satisfy the fundamental requirement of detailed balance, the acceptance probability of $\mathbf{x}^{(i)}$ as the next element in the chain is

$$p_{\text{acc}}[\mathbf{y}^{(i-1)} \rightarrow \mathbf{x}^{(i)}] = \tilde{h}_{\text{AB}}[\mathbf{x}^{(i)}] \min \left[1, \frac{p_{\text{sel}}(\mathbf{x}_{\text{sp}}^{(i)}; \mathbf{x}^{(i)})}{p_{\text{sel}}(\mathbf{x}_{\text{sp}}^{(i)}; \mathbf{y}^{(i-1)})} \right] \quad (2)$$

The indicator functional $\tilde{h}_{\text{AB}}[\mathbf{x}]$ equals unity if the trajectory connects A and B, and zero otherwise. $p_{\text{sel}}(\mathbf{x}_{\text{sp}}^{(i)}; \mathbf{x})$ is the probability of selecting $x_{\text{sp}}^{(i)}$ among the configurations of trajectory \mathbf{x} ; it can be any selection criterion function and can even change at different steps.⁴⁹ The $\mathbf{y}^{i-1} \rightarrow \mathbf{x}^i \equiv \mathbf{y}^i$ move is accepted or rejected according to p_{acc} ; in the latter case, we repeat $\mathbf{y}^{(i)} \equiv \mathbf{y}^{(i-1)}$. Note that the trial TPs may have a p_{acc} lower than one and, therefore, get rejected, modifying the weights of the accepted trajectories. On average, the acceptance of a new TP varies with the system and the selection bias. A good TPS algorithm increases the acceptance probability of the trial paths while preserving their heterogeneity.^{40,51}

In AIMMD, a neural network adaptively controls TPS in a data-driven way.⁴¹ The network models the committor

$p_{\text{B}}(x)$ —the probability that a trajectory initiated with random velocities at x reaches B before A.^{17,19,52} The committor quantifies the progress along the transition and is considered the optimal reaction coordinate.^{52,53} This enables us to quantify the transition mechanism and it also allows us to control the sampling. In fact, in the limit of Markovian dynamics, the probability of sampling a TP by a two-way shooting from x is⁵⁴

$$P(\text{TP}|x) = 2p_{\text{B}}(x)(1 - p_{\text{B}}(x)) \quad (3)$$

Since we do not apply bias forces to accelerate the transition in any direction, the dynamics remain time-reversible, such that we can always exchange A and B, p_{B} and $p_{\text{A}} = 1 - p_{\text{B}}$.

In AIMMD, we control sampling by modeling the shooting point selection probability $p_{\text{sel}}(x_{\text{sp}}^{(i)}; \mathbf{x})$ as a function of the committor. This choice enables us to control the exploitation–exploration dilemma. A selection probability peak around the transition state, $p_{\text{B}} = 0.5$, would result in a high TP-generating efficiency (exploitation). On the other hand, discovering new reaction channels—new transition mechanisms—often requires selecting points close to the state boundaries (exploration). Here, we strike a balance between the two requirements by selecting shooting points following a uniform distribution as a function of p_{B} . Consequently, at convergence, the optimal sampling rate of new TPs will be $\langle P(\text{TP} | x_{\text{sp}}) \rangle_{p_{\text{B}}} = 1/3$.

A selection probability that is a function of the committor of the shooting configuration is justified by the fact that the TPE and the equilibrium Boltzmann distributions are proportional along an isocommittor surface:⁵⁴

$$P(x|\text{TP}) \propto \rho(x)P(\text{TP}|x) \propto \rho(x)p_{\text{B}}(x)(1 - p_{\text{B}}(x)) \quad (4)$$

which follows from eq 3. Consequently, the shooting points are also Boltzmann-distributed along the isocommittor surfaces at convergence, which ensures an appropriate exploration of the transition region.

We learn the committor by training a neural network. The training set contains only the shooting points $x_{\text{sp}}^{(i)}$, hence it has the same size as the number of steps and therefore generated trajectories. To each $x_{\text{sp}}^{(i)}$, we associate a result $r^{(i)} = (r_{\text{A}}^{(i)}, r_{\text{B}}^{(i)})$ that comes from two-way shooting. The trial path $\mathbf{x}^{(i)}$ is the union of two subtrajectories (backward and forward) shot from $x_{\text{sp}}^{(i)}$. Note that we can include both backward and forward subtrajectories due to the diffusive nature of the trajectories, i.e., both backward and forward subtrajectories are valid instances of a committor shot. $r_{\text{A}}^{(i)}$ and $r_{\text{B}}^{(i)}$ count how many $\mathbf{x}^{(i)}$ subtrajectories reach the boundaries of A and B, respectively, with $r_{\text{A}}^{(i)} + r_{\text{B}}^{(i)} \leq 2$. The probability of those outcomes depends on the committor of $x_{\text{sp}}^{(i)}$. Maximizing the likelihood of the results is equivalent to minimizing the log-binomial loss^{41,55}

$$\mathcal{L}^{(n_{\text{TPS}})} = - \sum_{i=1}^{n_{\text{TPS}}} w_{\text{train}}^{(i)} [r_{\text{A}}^{(i)} \log(1 - p_{\text{B}}(x_{\text{sp}}^{(i)})) + r_{\text{B}}^{(i)} \log p_{\text{B}}(x_{\text{sp}}^{(i)})] \quad (5)$$

by learning the committor model $p_{\text{B}}(x)$. n_{TPS} is the number of trial shots, and $w_{\text{train}}^{(i)}$ is the importance in the training of the i -th shooting point (for the choice of $w_{\text{train}}^{(i)}$, see section III.A). Thus, the network learns the committor with no prior information and simultaneously enhances TPS.

II.B. Approximating the Equilibrium Path Ensemble. AIMMD directly provides a valid estimate of the TPE.

However, the TPs are only a small subset of the more general equilibrium path ensemble (PE) $\mathcal{P}[\mathbf{x}]$, which consists of all unbiased trajectories that start and end as soon as they cross any state boundary. In addition to transitions and excursions, as defined in the Introduction, the PE also has trajectories entirely contained in either A or B. We can split $\mathcal{P}[\mathbf{x}]$ into $\mathcal{P}_A[\mathbf{x}]$ and $\mathcal{P}_B[\mathbf{x}]$ —the path ensembles “gravitating” around the basins of attraction of states A and B, respectively.⁵⁶ \mathcal{P}_A (or \mathcal{P}_B) contains all trajectories that start either entering or leaving A (or B) and end upon crossing any state boundary (A or B). Trajectories in $\mathcal{P}_A[\mathbf{x}]$ do not have configurations in B and vice versa.

Under the ergodic hypothesis, sampling $\mathcal{P}[\mathbf{x}] = \mathcal{P}_A[\mathbf{x}] \cup \mathcal{P}_B[\mathbf{x}]$ is equivalent to simulating and then splitting an infinitely long unbiased trajectory (Figure 2a). Thus, the PE contains both thermodynamic and kinetic information about the studied transition.

Our goal is to approximate the PE with a set of n dynamically unbiased short trajectory segments, which are not necessarily distributed a priori according to the equilibrium PE,




Figure 2. (a) Ensemble definitions. Given an infinitely long unbiased simulation, the path ensemble (PE) collects all the trajectory segments starting and ending upon crossing the boundary of a state (cuts at the white dots). The trajectories in \mathcal{P}_A (blue) start at the boundary of A; their configurations populate $\mathcal{P}_{A,\lambda_A}^+$ (light blue) and $\mathcal{P}_{A,\lambda_A}^-$ (dark blue), based on whether they have committor values bigger or smaller than λ_A . \mathcal{P}_B and $\mathcal{P}_{B,\lambda_B}^\pm$ are defined analogously. (b) Extension to the metastable states. The reweighted AIMMD trajectories approximate the PE in the cyan ($\mathcal{P}_{A,\lambda_A}^+$) and yellow ($\mathcal{P}_{B,\lambda_B}^-$) regions. In the blue ($\mathcal{P}_{A,\lambda_A}^-$) and red ($\mathcal{P}_{B,\lambda_B}^+$) regions, we use (potentially) short simulations initiated around the states (horizontal lines); their occasional excursions in the AIMMD trajectories’ regions are crucial to the extension, as they determine the trajectories’ relative weight. (c) AIMMD trial trajectories to and from A before reweighting. Each path is represented by a rectangle between its minimum and maximum committor values (TPs span the whole transition region); the height of the rectangle corresponds to its weight; a white dot indicates the committor value of its shooting point. The trajectories’ “crossing statistics” do not match the expected crossing probability for $\mathcal{P}_{A,\lambda_A}^+$ (red line). (d) The same trajectories as in (c) after applying the reweighting scheme of eq 12a. The crossing statistics now match $P_A(\lambda|\lambda_A) = \lambda_A/\lambda$.

and need to be reweighted. For this purpose, we use the following notation:

$$\mathcal{P} \approx \{(\mathbf{w}^{(1)}, \mathbf{x}^{(1)}), \dots, (\mathbf{w}^{(i)}, \mathbf{x}^{(i)}), \dots, (\mathbf{w}^{(n)}, \mathbf{x}^{(n)})\} \quad (6)$$

where $\mathbf{w}^{(i)}$ is the vector of weights associated with each configuration in trajectory $\mathbf{x}^{(i)}$. The aim of this approximation is that the distribution of the configurations in the reweighted trajectories, $\rho(x)$, must follow the Boltzmann distribution of the configurations in the PE. More generally, the ensemble average of any thermodynamic observable $O(x)$ must be

$$\langle O \rangle \approx \sum_{i=1}^n \sum_{t=0}^{L[\mathbf{x}^{(i)}]} w_t^{(i)} O(x_t^{(i)}) \quad (7)$$

Given a (small) committor threshold $\lambda_A > 0$, we further split $\mathcal{P}_A[\mathbf{x}]$ into $\mathcal{P}_{A,\lambda_A}^-$ and $\mathcal{P}_{A,\lambda_A}^+$, according to the following definition: the configurations in $\mathcal{P}_{A,\lambda_A}^-$ have committor values smaller than λ_A , while those in $\mathcal{P}_{A,\lambda_A}^+$ have $p_B(x) \geq \lambda_A$ (Figure 2a). Given a (large) committor threshold $\lambda_B < 1$, we obtain $\mathcal{P}_{B,\lambda_B}^-$ and $\mathcal{P}_{B,\lambda_B}^+$ analogously. These new ensembles are strictly speaking not proper path ensembles but are proper configurational ones.

$\mathcal{P}_{A,\lambda_A}^-$, $\mathcal{P}_{A,\lambda_A}^+$, $\mathcal{P}_{B,\lambda_B}^-$, and $\mathcal{P}_{B,\lambda_B}^+$ form a partition of configurations in the PE. We will approximate each one separately and join them together once their relative weights are determined.

The configurations in $\mathcal{P}_{A,\lambda_A}^+$, $\mathcal{P}_{B,\lambda_B}^-$ are more difficult to sample, since they contain the rare event of interest. Our approach is to approximate the $\mathcal{P}_{A,\lambda_A}^+$, $\mathcal{P}_{B,\lambda_B}^-$ ensembles with the n_{TPS} trial paths $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n_{\text{TPS}})}$ produced by AIMMD to sample the transition between A and B. Crucially, these include the paths that were not reactive (A-to-A and B-to-B). The justification for this lies in the path-recycling method introduced in ref 46, which established that trial trajectories created by two-way shooting, whether transitions or not, are proper paths that take part in the equilibrium PE. By creating the trial paths along the entire order parameter range, one ensures proper coverage of the PE. The resulting trial paths are naturally not distributed according to the equilibrium ensemble because they were created from a biased selection and thus must be properly reweighted.

For $\mathcal{P}_{A,\lambda_A}^-$ and $\mathcal{P}_{B,\lambda_B}^+$, which entirely include the metastable states, we will use $n_A + n_B$ short unbiased trajectories $\mathbf{z}_A^{(1)}, \dots, \mathbf{z}_A^{(n_A)}$, and $\mathbf{z}_B^{(1)}, \dots, \mathbf{z}_B^{(n_B)}$ initialized around A and B, respectively (Figure 2b). In this way, we complement the AIMMD trajectories with short equilibrium simulations and extend \mathcal{P} to the metastable states. To avoid confusion, we call $\mathbf{v}_A^{(i)}$ and $\mathbf{v}_B^{(j)}$ the weights associated with the trajectories $\mathbf{z}_A^{(i)}$ and $\mathbf{z}_B^{(j)}$, respectively.

Our estimate of the properly weighted configurations in the PE thus becomes

$$\begin{aligned} \mathcal{P} \approx & \{(\mathbf{v}_A^{(1)}, \mathbf{z}_A^{(1)}), \dots, (\mathbf{v}_A^{(n_A)}, \mathbf{z}_A^{(n_A)})\} \\ & \cup \{(\mathbf{w}_A^{(1)}, \mathbf{x}^{(1)}), \dots, (\mathbf{w}_A^{(n_{\text{TPS}})}, \mathbf{x}^{(n_{\text{TPS}})})\} \\ & \cup \{(\mathbf{w}_B^{(1)}, \mathbf{x}^{(1)}), \dots, (\mathbf{w}_B^{(n_{\text{TPS}})}, \mathbf{x}^{(n_{\text{TPS}})})\} \\ & \cup \{(\mathbf{v}_B^{(1)}, \mathbf{z}_B^{(1)}), \dots, (\mathbf{v}_B^{(n_B)}, \mathbf{z}_B^{(n_B)})\} \end{aligned} \quad (8)$$

with $n = n_{\text{TPS}} + n_A + n_B$. The first two sets jointly approximate \mathcal{P}_A , and the latter two approximate \mathcal{P}_B , with the corresponding distributions of configurations ρ_A and ρ_B .

II.C. Reweighting the Trial Trajectories. In this section, we derive an explicit solution for reweighting the trajectories sampled by AIMMD. AIMMD generates all of the trajectory types that can be extracted from long equilibrium simulations (excursions and transitions). Individually, AIMMD trajectories are indistinguishable from those extracted from long equilibrium runs, because they are dynamically unbiased. However, they are biased in their starting configurations, which means that large excursions and transitions are significantly more frequent than equilibrium (Figure 2c). In other words, the path ensemble generated by AIMMD has the same elements as the equilibrium ensemble but with different weights. Our strategy is to use importance sampling to correct this imbalance: we associate a crossing probability distribution along a progress coordinate and reweight the distribution of the ensemble of paths generated by AIMMD to match that of the equilibrium PE (Figure 2d). We will show that by using the committor as a progress coordinate, we can write an explicit solution of the latter, which enables a closed-form solution for the weights.

Our algorithm determines the optimal weights $w_A^{(i)}$ and $w_B^{(i)}$ for each trajectory $\mathbf{x}^{(i)}$. However, there are portions of the transition region ($\lambda(x) < \lambda_A$ for $\mathcal{P}_{A,\lambda_A}^+$ and $\lambda(x) > \lambda_B$ for $\mathcal{P}_{B,\lambda_B}^-$) where we use the equilibrium segments for the density estimates. There, the $\mathbf{x}^{(i)}$ trajectories have zero weight. For this reason, we keep the vectorial notation:

$$w_{A,t}^{(i)} = w_A^{(i)} \theta(p_B(\mathbf{x}^{(i)}(t)) \geq \lambda_B) \quad (9a)$$

$$w_{B,t}^{(i)} = w_B^{(i)} \theta(p_B(\mathbf{x}^{(i)}(t)) \leq \lambda_B), \quad \forall i, t \quad (9b)$$

where $\theta(x)$ is the Heaviside function, which is 1 if $x > 0$ and 0 otherwise. Therefore, each $w_{A,t}^{(i)}$ value is always either 0 or $w_A^{(i)}$. Furthermore, $w_A^{(i)} = 0$ if $\mathbf{x}^{(i)}$ does not originate and/or terminate in A, and $w_B^{(i)} = 0$ if it does not start and/or end in B.

Each TP has both $w_A^{(i)} > 0$ and $w_B^{(i)} > 0$. It contributes to both $\mathcal{P}_{A,\lambda_A}^+$ and $\mathcal{P}_{B,\lambda_B}^-$ because of microscopic time reversibility: to our purposes, an unbiased trajectory that goes from A to B is equally valid to the time-reversed counterpart from B to A. By allowing for time-reversed trajectory segments, we improve the accuracy of the $\mathcal{P}_{A,\lambda_A}^+$ and $\mathcal{P}_{B,\lambda_B}^-$ estimates. Moreover, the weights are halved to avoid double counting.

We introduce the (equilibrium) crossing probability $P_A(\lambda | \lambda_A)$:^{44,56,57} the probability that a trajectory segment from equilibrium sampling, starting in A, and crossing ($\lambda_A \in (0, \lambda]$) reaches λ before returning to A or ending in B. By assigning the proper importance to each trial path starting and/or ending in A, we ensure that the weighted fraction of paths that reach λ follows $P_A(\lambda | \lambda_A)$. While λ is usually interpreted as an order parameter that is a reasonable proxy for the reaction coordinate, here we take λ to be the best possible reaction coordinate: the committor itself. In trajectory-based approaches such as TIS and ours, the crossing probability is the key to reweighting. By using the committor we obtain a simple closed solution for the crossing probability (see Appendix A for proof):

$$P_A(\lambda | \lambda_A) = \frac{\lambda_A}{\lambda}, \quad \forall \lambda_A \in (0, \lambda], \forall \lambda \in [\lambda_A, 1] \quad (10)$$

The trajectory segments in $\mathcal{P}_{A,\lambda_A}^+$ start and/or end in A and cross $\lambda_A \ll 1$ by definition. Therefore, the fraction of them reaching λ must correspond to eq 10. In particular, most trajectories would only make small excursions in the transition region ($\lambda \ll 1$) before returning to A. In our case, we approximate $\mathcal{P}_{A,\lambda_A}^+$ with a finite set of trial paths. Again, the fraction of them reaching λ (the “crossing statistics”) should also match $P_A(\lambda | \lambda_A)$.

In AIMMD, however, we initialize trajectories at higher committor values by controlling the selection probability of the shooting points. For example, let $\lambda^{(i)} \equiv p_B(\mathbf{x}_{\text{sp}}^{(i)})$ be the value of the i -th shooting point. If $\mathbf{x}_{\text{sp}}^{(i)}$ is at the transition state ($\lambda^{(i)} = 0.5$), then $\mathbf{x}^{(i)}$ would start by construction at $\lambda = 0.5$. This selection biases the crossing statistics (Figure 2c). Only by appropriately weighting the trajectories can we match the observed statistics with the expected one and reconstruct $\mathcal{P}_{A,\lambda_A}^+$.

The weights $w_A^{(i)}$ should increase the contribution of small excursions and decrease the contribution of large excursions and TPs (Figure 2c). The RPE theory⁴³ demonstrates that the weight of each path $\mathbf{x}^{(i)}$ depends on the furthestmost value of the committor $\lambda_{\text{max}}^{(i)}$ along that path (the magnitude of the associated excursion).

The reweighting is an important sampling procedure: the weight of $\mathbf{x}^{(i)}$ should be of the form $E(\lambda_{\text{max}}^{(i)})/S(\lambda_{\text{max}}^{(i)})$, where $E(\lambda_{\text{max}}^{(i)})$ is the fraction of expected paths that should touch at least $\lambda_{\text{max}}^{(i)}$, and $S(\lambda_{\text{max}}^{(i)})$ is the fraction of simulated paths that touched at least $\lambda_{\text{max}}^{(i)}$. From eq 10, we know that $E(\lambda_{\text{max}}^{(i)})$ is proportional to $1/\lambda_{\text{max}}^{(i)}$. In this way, smaller excursions in the transition region become increasingly promoted as $\lambda_{\text{max}}^{(i)}$ approaches λ_A , while the TPs ($\lambda_{\text{max}} = 1$) have the lowest weights. Conversely, $S(\lambda_{\text{max}}^{(i)})$ is proportional to $m_A(\lambda_{\text{max}}^{(i)})$: the number of paths from A that touched $\lambda_{\text{max}}^{(i)}$. However, for assessing S we can only consider the trajectories shot from committor values $p_B(\mathbf{x}_{\text{sp}}^{(j)}) = \lambda^{(j)}$ lower than $\lambda_{\text{max}}^{(i)}$ —the only ones following the crossing statistics at $\lambda_{\text{max}}^{(i)}$, since paths with a shooting point beyond $\lambda_{\text{max}}^{(i)}$ are pushed closer to state B by construction. For these paths, we compute

$$m_A(\lambda_{\text{max}}^{(i)}) = \sum_{j=1}^{n_{\text{TPS}}} \tilde{h}_A[\mathbf{x}^{(j)}] \theta(\lambda_{\text{max}}^{(i)} - \lambda^{(j)}) \theta(\lambda_{\text{max}}^{(j)} - \lambda_{\text{max}}^{(i)}) \quad (11)$$

which counts the trajectories generated from shooting points with committor values smaller than $\lambda_{\text{max}}^{(i)}$ and reaching (at least) $\lambda_{\text{max}}^{(i)}$. The indicator functional \tilde{h}_A ensures that $\mathbf{x}^{(j)}$ starts/ends in A and crosses λ_A .

By matching the expected and observed crossing statistics, we finally obtain an explicit form for the weights

$$w_A^{(i)} = \tilde{h}_A[\mathbf{x}^{(i)}] \frac{c_A}{\lambda_{\text{max}}^{(i)} m_A(\lambda_{\text{max}}^{(i)})} \quad (12a)$$

where c_A is a normalizing constant. A consequence of eq 12a is that TPs are all reweighted by the same amount. In Appendix B, we show that eq 12a can be obtained rigorously as a limit case of the RPE theory.⁴³ The reweighting algorithm is insensitive in the direction orthogonal to the committor: we can modify the relative free energies only among different isocommittor surfaces and not along the same surface. Thus, the trajectory segments must already be Boltzmann distributed on the isocommittor surfaces before reweighting for obtaining a good PE estimate. In AIMMD, this condition is fulfilled by

construction once the algorithm converged to an accurate committor model as granted by eq 4.

Switching states A and B, we must exchange $p_B(x) = \lambda$ with $p_A(x) = 1 - \lambda$, but the derivation remains the same. Thus

$$w_B^{(i)} = \tilde{h}_B[x^{(i)}] \frac{c_B}{(1 - \lambda_{\min}^{(i)})m_B(\lambda_{\min}^{(i)})} \quad (12b)$$

and analogously, $m_B(\lambda)$ counts the trajectories shot at committor values bigger than λ , starting and/or ending in B, and reaching λ , while $\lambda_{\min}^{(i)}$ is the furthestmost committor value reached by $x^{(i)}$ from B ($\lambda_{\min} = 0$ if the path is reactive). In Figure 2c,d, we show how this procedure recovers the expected crossing probability from synthetic data. A uniform shooting point distribution in committor space homogenizes m_A and m_B in eq 12 and improves the accuracy of the estimate.

The relative importance of \mathcal{P}_A and \mathcal{P}_B is set by the normalizing constants c_A and c_B . We impose the fixed ratio:

$$\frac{c_A}{c_B} = \frac{\rho'_B(\lambda = 0.5)}{\rho'_A(\lambda = 0.5)} \quad (13)$$

where $\rho'_A(\lambda = 0.5)$ and $\rho'_B(0.5)$ are the unnormalized \mathcal{P}_A and \mathcal{P}_B densities at the transition state. (For a justification of eq 13, see Appendix C). In practice, we can compute the weighted population of the ensembles between $\lambda = 0.45$ and $\lambda = 0.55$. The scaling in the densities depends on the choice of the committor as a coordinate but factors out by taking the ratio. Using these constants thus ensures that the transition state interface has equal amounts of trajectories going to A and B in the PE. An alternative solution to determine the c_A and c_B constants would make use of the crossing probabilities, which are unitless and require that they match at the destination state. For the systems discussed here, using densities or probabilities did not lead to any significant difference.

II.D. Extension to the Metastable States. So far, we determined the weights for the paths in the transition region. Now we will derive a solution for the $v_A^{(i)}$ and $v_B^{(i)}$ vectors containing the weights for the short, unbiased simulations in the wells. As the $z_A^{(i)}$ trajectories are the outcome of equilibrium sampling around state A, their configurations must have equal weight throughout $\mathcal{P}_{A,\lambda_A}^-$; the same argument holds for the $z_B^{(i)}$. Thus:

$$v_{A,t}^{(i)} = \gamma_A \theta(p_B(z_A^{(i)}(t)) < \lambda_A) \quad (14a)$$

$$v_{B,t}^{(i)} = \gamma_B \theta(p_B(z_B^{(i)}(t)) > \lambda_B), \quad \forall i, t \quad (14b)$$

where γ_A and γ_B are positive constants. The occasional excursions of the $z_A^{(i)}$ and $z_B^{(i)}$ above and below the λ_A and λ_B thresholds are instrumental for determining γ_A and γ_B and therefore extending our $\mathcal{P}[x]$ evaluation to the metastable states. The number of $z_A^{(i)}$ configurations that go beyond λ_A , when multiplied by γ_A , must match the total population of $\mathcal{P}_{A,\lambda_A}^+$. Conversely, the number of $z_B^{(i)}$ configurations that cross λ_B , when multiplied by γ_B , must correspond to $\mathcal{P}_{B,\lambda_B}^+$. We enforce the above statements by setting

$$\gamma_A(\lambda_A) = \frac{\sum_{i=1}^{n_{\text{TFS}}} w_A^{(i)} \sum_{t=0}^{L[x^{(i)}]} \theta(p_B(x^{(i)}(t)) \geq \lambda_A)}{\sum_{i=1}^{n_A} \sum_{t=0}^{L[z_A^{(i)}]} \theta(p_B(z_A^{(i)}(t)) \geq \lambda_A)} \quad (15a)$$

$$\gamma_B(\lambda_B) = \frac{\sum_{i=1}^{n_{\text{TFS}}} w_B^{(i)} \sum_{t=0}^{L[x^{(i)}]} \theta(p_B(x^{(i)}(t)) \leq \lambda_B)}{\sum_{i=1}^{n_B} \sum_{t=0}^{L[z_B^{(i)}]} \theta(p_B(z_B^{(i)}(t)) \leq \lambda_B)} \quad (15b)$$

γ_A and γ_B should be constant for all choices of λ_A , λ_B . In practice, they become inaccurate when λ_A and λ_B are too close to A and B (due to relatively large error in the committor estimate and the noninfinitesimal interval between trajectory frames) or too close to the transition state (due to the inadequate equilibrium sampling). In particular, from eq 10 a trajectory leaving A reaches λ_A a factor of λ_A^{-1} times more frequently than undergoing a transition.

It is convenient to determine the value of λ_A and λ_B by fixing the number of equilibrium configurations that go beyond those thresholds:

$$M_A = \sum_{i=1}^{n_A} \sum_{t=0}^{L[z_A^{(i)}]} \theta(p_B(z_A^{(i)}(t)) \geq \lambda_A) \quad (16a)$$

$$M_B = \sum_{i=1}^{n_B} \sum_{t=0}^{L[z_B^{(i)}]} \theta(p_B(z_B^{(i)}(t)) \leq \lambda_B) \quad (16b)$$

In this way, $\lambda_A = \lambda_A(M_A)$ and $\lambda_B = \lambda_B(M_B)$ follow from inversion. By setting M_A and M_B , we ensure that enough equilibrium sampling contributes to the calculation of $\gamma_A(\lambda_A(M_A))$ and $\gamma_B(\lambda_B(M_B))$. We can then optimize M_A and M_B as the parameters returning the most robust γ_A and γ_B to small boundary changes (Figure S3,S4b). The weights of eq 12a are independent of the choice of λ_A and λ_B and thus are computed only once.

In general, it is always possible to match the distributions around the wells and in the transitions region by using WHAM or analogous approaches.^{58–60}

Finally, we enforce global normalization by rescaling all weights such that they sum to one.

II.E. Free Energy Profiles along Arbitrary Variables.

We can project the PE and get the equilibrium distribution ρ as a function of any set of collective variables $q = \{q_1, \dots, q_k\}$.^{43,61} Starting from eq 7, making the weights explicit, and applying the density operator $\delta(q(x) - q')$:

$$\begin{aligned} \rho(q') \propto & \gamma_A \sum_{i=1}^{n_A} \sum_{t=0}^{L[z_A^{(i)}]} \theta(p_B(z_A^{(i)}(t)) < \lambda_A) \delta(q(z_A^{(i)}(t)) - q') \\ & + \sum_{i=1}^{n_{\text{TFS}}} w_A^{(i)} \sum_{t=0}^{L[x^{(i)}]} \theta(p_B(x^{(i)}(t)) \geq \lambda_A) \delta(q(x^{(i)}(t)) - q') \\ & + \sum_{i=1}^{n_{\text{TFS}}} w_B^{(i)} \sum_{t=0}^{L[x^{(i)}]} \theta(p_B(x^{(i)}(t)) \leq \lambda_B) \delta(q(x^{(i)}(t)) - q') \\ & + \gamma_B \sum_{i=1}^{n_B} \sum_{t=0}^{L[z_B^{(i)}]} \theta(p_B(z_B^{(i)}(t)) > \lambda_B) \delta(q(z_B^{(i)}(t)) - q') \end{aligned} \quad (17)$$

In practice, we bin along q and count the weighted population of each contribution to the PE approximation in each bin.

The corresponding free energy profile is $F(q) = -k_B T \log \rho(q)$ up to an additive constant. The free energy difference between A and B is

$$\Delta F_{AB} = F_B - F_A = k_B T \log \frac{\rho_B}{\rho_A} \quad (18)$$

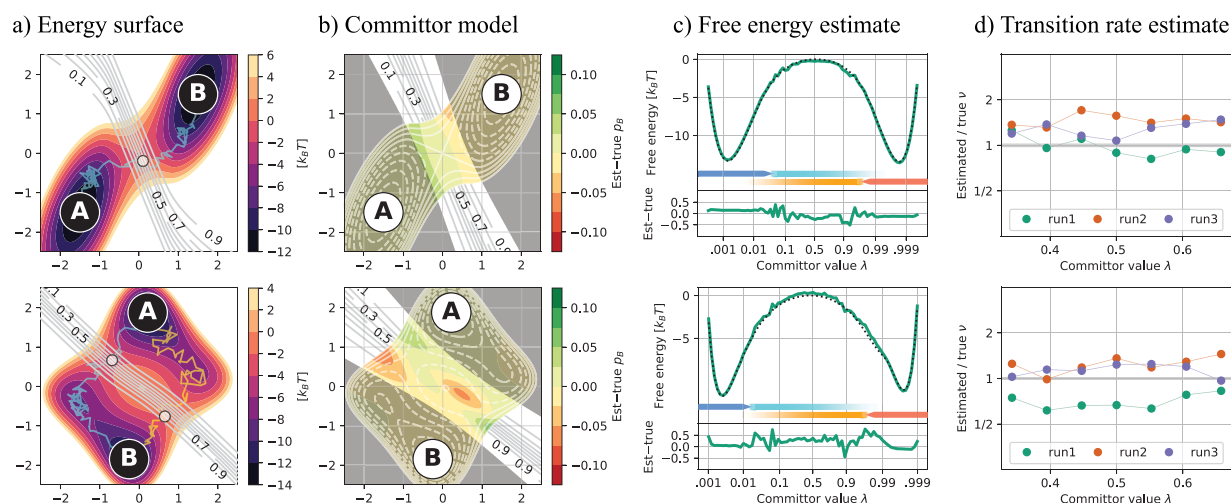


Figure 3. Validation on the two-dimensional benchmark systems (top, double well; bottom, Wolfe-Quapp), results after 500 AIMMD steps. (a) Potential energy surface (filled contour), example of TPs in different channels (blue and orange lines, shooting points in white), and true committor (contour lines). (b) AIMMD run1, committor model (contour lines), error of the model (filled contour), and region between the λ_A and λ_B thresholds (light area). (c) Run1, free energy as a function of the estimated committor (solid line, top axis), free energy from numerical computation (dotted line), and error of the estimate (bottom axis). The arrows indicate the contributions of the simulations around A (blue), the AIMMD trial trajectories from A (cyan), the simulations around B (red), and the AIMMD trial trajectories from B (orange). (d) Bayesian rate estimate of ν at different committor values, with each color denoting a different run. The gray area is the 95% confidence interval of ν from the equilibrium simulations.

where ρ_A and ρ_B are the sum of $\rho(q)$ over all configurations in A and B, respectively.

II.F. Rate Constants. For the reaction rate constants, we employ the Bayesian framework developed by Hummer:^{54,62}

$$\nu = \frac{2}{k_{AB}^{-1} + k_{BA}^{-1}} = \frac{\rho(\lambda)}{\rho_{TP}(\lambda)} \frac{2\lambda(1-\lambda)}{\langle t_{AB} \rangle_{TP}} \quad (19)$$

where k_{AB} and k_{BA} are the A \rightarrow B and B \rightarrow A rates, respectively, and we project all densities on the committor. The transition path density ρ_{TP} comes from the previous AIMMD run. $\langle t_{AB} \rangle_{TP}$ is the average duration of the TPs, and $2\lambda(1-\lambda)$ is the probability $P(TP | \lambda)$ of an equilibrium trajectory crossing λ to be reactive, as in eq 3.

Equation 19 holds for any value of λ , as follows from eq 4. However, it produces more accurate results around the transition state ($\lambda = 0.5$) due to a better estimate of the committor. The outcome is the quantity ν combining both k_{AB} and k_{BA} ; it is the inverse of the average mean first passage time for A \rightarrow B and B \rightarrow A, and the reciprocal of half the average first return time.⁶³ The individual rate constants

$$k_{AB} = \frac{1 + e^{-\Delta F_{AB}/k_B T}}{2} \nu \quad (20a)$$

$$k_{BA} = \frac{1 + e^{+\Delta F_{AB}/k_B T}}{2} \nu \quad (20b)$$

follow from eq 19 and the fact that $k_{AB} = e^{-\Delta F_{AB}/k_B T} k_{BA}$.

We could have also estimated the rates multiplying the reactive fluxes through the interfaces defined by λ_A and λ_B with the expected crossing probabilities.⁵⁶ However, we found that eq 19 provides a numerically more robust estimate.

II.G. The Complete Algorithm. We summarize the entire algorithm as follows:

1. Perform AIMMD simulations and learn the committor $p_B(x)$ from the $\{x_{sp}^{(i)}, r^{(i)}\}$ training set. Collect the $x^{(1)}$,

$x^{(2)}, \dots, x^{(n)}$ trial paths (note that this includes the rejected paths).

2. At the same time, run MD from multiple equilibrium configurations in states A and B. Collect the sampled $z_A^{(i)}$ and $z_B^{(i)}$ trajectories.
3. Evaluate p_B on all the simulated configurations; for each $x^{(i)}$, save $\lambda^{(i)} = p_B(x_{sp}^{(i)})$, $\lambda_{\min}^{(i)}$, and $\lambda_{\max}^{(i)}$.
4. Weight the $x^{(i)}$ for approximating $\mathcal{P}_{A,\lambda_A}^+$ and $\mathcal{P}_{B,\lambda_B}^-$ according to eqs 12, with preliminary $c_A = c_B = 1$. Obtain the unnormalized $w_A^{(i)}$ and $w_B^{(i)}$ (Figure 2d).
5. Evaluate γ_A and γ_B from eqs 15 and determine the optimal λ_A and λ_B parameters. Assign the weights $v_A^{(i)}$ and the $v_B^{(i)}$ trajectories according to $\gamma_A(\lambda_A)$ and $\gamma_B(\lambda_B)$.
6. Project \mathcal{P}_A and \mathcal{P}_B on the transition state: $\rho_A(\lambda = 0.5)$ and $\rho_B(\lambda = 0.5)$. Impose the condition of eq 13 by rescaling $c_A = 1/\rho_A(0.5)$ and $c_B = 1/\rho_B(0.5)$, and thus $w_A^{(i)}$, $w_B^{(i)}$, γ_A , and γ_B .
7. Merge all the simulated trajectories as in eq 8. Normalize the weights over all configurations.
8. Obtain a free energy profile as a function of the estimated committor, or as a function of arbitrary variables, and estimate the rate constants.

III. METHODS

III.A. AIMMD and PE Calculations. We used and extended the AIMMD Python package developed by Jung⁴¹ to run the path sampling simulations. For each AIMMD run, we initialized a deep neural network in PyTorch⁶⁴ with 4 hidden linear layers of size 8192, 2048, 512, and 128 with ELU activation functions, 4 residual units⁶⁵ with 4 layers, and 128 neurons per layer, and a final linear layer (Figure S1). The architecture is expressive enough to approximate the committor in many-dimensional systems. The encoder structure encourages the pruning of unimportant features.

The network's output is $q(x)$, which is a one-to-one function of the committor:⁶⁶

$$p_B(x) = \sigma(q(x)) = \frac{1}{1 + e^{-q(x)}} \quad (21)$$

In AIMMD, we define a step as a shooting-attempt, i.e., starting two simulations with opposite initial velocities from a selected configuration. A step can lead to the production of an excursion or a transition path. Training does not have to be done after every step. We performed 100 training epochs by minimizing the binomial loss of eq 5 with the ADAM optimizer.⁶⁷ We found that the learning rates $l_r = 10^{-5}$ (2D systems) and $l_r = 3.5 \times 10^{-6}$ (chignolin), dependent on the network architecture, were good hyperparameter choices to prevent the model from overfitting.⁶⁸ We observed that the optimal l_r value decreased with increased input feature dimensionality with no system-specific dependency. To regularize the training set, we set the shooting points' importance $w_{\text{loss}}^{(i)}$ such that the A-to-A, B-to-B, and A-to-B and B-to-A results would have each equal cumulative weight. We saved the neural network parameters at regular intervals.

To achieve the target uniform selection probability, we determined p_{sel} based on the committor values of the origin trajectory (the last accepted trajectory in the Markov chain built by TPS). We organized the candidate shooting points into 10 equally p_B -spaced bins and scaled the probability by the population of the bins' population. The probability of an empty bin was distributed to the adjacent ones. For each trajectory, the selection probability is a function of the committor only. To ensure detailed balance in the Markov chain, we kept the rule consistent within a step when calculating the acceptance probability of eq 2.

We wrote the PathEnsemble Python class to automate the PE estimation summarized in section II.G. A PathEnsemble instance collects the features vectors, committor values, and complementary information on a set of trajectories. It can extract the TPE from the trials of an AIMMD run, weight the trajectories for estimating \mathcal{P}_A and \mathcal{P}_B , and combine ensembles together. It can also project the free energy onto an arbitrary set of collective variables.

For assessing the accuracy of the rate estimates, we considered the maximum relative error $E(k, \hat{k}) = |k - \hat{k}| / \min\{k, \hat{k}\}$, where k is the reference value and \hat{k} is the estimated value of the rate.

III.B. 2D Systems. The two-dimensional (2D) systems are defined by their energy surface on the (x, y) plane. The double well energy surface has equation:

$$U(x, y) = \begin{cases} -2\Delta G(x/\delta)^2 + k_0(x - y)^2/2 & \text{if } x/\delta < 0.5, \\ \Delta G[2(|x/\delta| - 1)^2 - 1] + k_0(x - y)^2/2 & \text{if } x/\delta \geq 0.5 \end{cases} \quad (22)$$

with the barrier height $\Delta G = 12 k_B T$, $k_0 = 10.4 k_B T$, and $\delta = 1.5$ (Figure 3a, top). The Wolfe-Quapp energy surface⁶⁹ is defined by

$$U(x, y) = \frac{\Delta G}{5}(x^4 + y^2 - 2x^2 - 3y^2 + xy + 0.3x + 0.1y) \quad (23)$$

we set $\Delta G = 10 k_B T$, and by rotating the x and y coordinates by 45 deg (Figure 3a, bottom).

In each system, we evolved a particle with overdamped Langevin dynamics¹⁷ (diffusion coefficient $D = 10^{-5}$ with

unitary distance, energy, integration time step, and mass), and saved the trajectories every 500 (double well) and 1,000 (Wolfe-Quapp) integration steps. In this way, TPs contain approximately 100 frames. As metastable states, we picked circles of radius $r = 0.5$ around the local minima. We computed the reference committor by numerically solving the stationary Fokker-Plank equation⁷⁰ and derived the reference k_{AB} and k_{BA} rate constants by fitting the exponential decay¹⁷ of 40,000 replicas initiated in A and B.

For each system, we performed 3 AIMMD runs of 5,000 steps each, directly feeding the x, y coordinates to three different neural networks. To assess the speedup given by AIMMD, we performed a standard TPS (run0) as a benchmark. As the initial trajectory (not representative of the dynamics, therefore excluded from the TPE), we drew a straight line connecting the minima; the first trial TP is always accepted.

From each AIMMD run, we computed and extended the PE with 20 equilibrium trajectories initiated from the energy minima (10 each). The trajectories are at most 500,000 (double well) and 25,000 (Wolfe-Quapp) frames long and were trimmed in case they reached the other state. We determined the λ_A, λ_B thresholds according to eqs 16 such that 100 configurations from A and from B went beyond those committor values (Figure S3). For the ideal scenario of optimal sampling around the basins, we numerically computed the $\rho_A(x, y) = \rho(x, y) (1 - p_B(x, y))$ and $\rho_B(x, y) = \rho(x, y) p_B(x, y)$ distributions, and scaled their weights such that λ_A, λ_B were consistent with the previous case.

III.C. Chignolin. We obtained the folded structure of CLN025 (amino acid sequence YYDPETGTWY) from the 2RVD entry of the Protein Data Bank^{71,72} (Figure 5b). We solvated the peptide with TIP3 water in a 4 nm cubic periodic box and generated a topology file with Charmm-GUI;⁷³ the final system has 6,468 atoms, 166 belonging to the peptide. We reproduced the settings of Lindorff-Larsen et al.⁷⁴ and chose the CHARMM22* force-field.⁷⁵ We ran the simulations at $T = 340$ K with GROMACS 2022.4⁷⁶ and the velocity Verlet integrator; we fixed the volume after 1 ns of equilibration and maintained the temperature with the velocity rescale thermostat.⁷⁷ We set a 0.95 nm threshold for the short-range interactions and left the remaining GROMACS parameters unchanged. We integrated the positions and momenta every $dt = 2$ fs and saved the former every $\Delta t = 100$ ps in the XTC trajectory files.

We calculated the reference free energy profiles and rate constants from 4 equilibrium MD simulations, totaling 120 μ s. We visualized the trajectories with VMD⁷⁸ and analyzed them with MDTraj.⁷⁹

We defined the folded (A) and unfolded (B) states based on the fraction of native contacts:⁸⁰

$$A = \{x|Q(x) \geq 0.99\} \quad (24a)$$

$$B = \{x|Q(x) \leq 0.01\} \quad (24b)$$

where the reference configuration ($t = 55.1$ ns of the first equilibrium MD simulation) is the centroid of the C_α -RMSD folded state cluster.⁷⁴ We also considered the following additional collective variables (see also Figure 6):

1. the distance between Asp3O and Gly7N (d_1), forming a hydrogen bond in the native state;⁸¹
2. the distance between Asp3N and Gly7O (d_2);
3. the distance between Asp3N and Thr8O (d_3);

- the fraction of native contacts between Tyr2 and Trp9 (Q_{29});⁴⁸
- the fraction of native contacts between Pro4 and Gly7 (Q_{47});
- the $C\alpha$ -RMSD with respect to the reference structure;⁷⁴
- the radius of gyration of the protein's heavy atoms (r_g).

We performed 3 AIMMD runs of 250 steps each. As in the 2D systems, we ran further standard TPS (run0, run0b, and run0c) with random selection probability as a benchmark to compare performances. As the input for the neural network, we featurized the system calculating 2064 distances between heavy atoms at least 4 residues apart and min-max normalized the distances according to the values sampled in the short equilibrium simulations in A and B.

To obtain the initial TP trajectory, we set the temperature to $T = 600$ K and quickly unfolded the native configuration in 0.9 ns. All successive simulations were instead performed at $T = 340$ K. We always accepted the first TP sampled after initialization and discarded the initial high-temperature trajectory from the TPE.

The path length limit for two-way-shooting is 100 ns, which resulted in only 2 trajectories being rejected across the AIMMD runs. Those trajectories were excluded from the PE estimates but not from the training. In general, we strive for a limit that is much higher than the average length of TPs and still much lower than the average first return time.

In extending the PE, we integrated each AIMMD run with 20 short simulations initiated from two original 20 ns equilibrium trajectories around A and B (Figure 2b). The simulations were terminated as soon as they hit 50 (500 frames, for A) or 5 ns (for B). We picked the λ_A and λ_B thresholds such that 10 configurations around A and 50 around B crossed those committor values (Figure S4b).

IV. RESULTS

IV.A. Double Well Potential. We illustrated our method on the double well benchmark system. The potential has a $12 k_B T$ energy barrier, resulting in an average first return time of about 10^5 times larger than the average TP time. The committor varies significantly in a small portion of the transition region, with most equilibrium configurations highly committed to either A or B (Figure 3a, top). Many configurations of TPE are also far from the barrier. Hence, biasing the shooting point selection probability toward the transition state is essential for good sampling performance.⁸²

AIMMD generated 4,698 TPs and 2,456 accepted ones in 15,000 steps across 3 independent runs. The resulting TPEs match the reference (Figure S2a). To mimic a data-poor regime, we calculated the committor, free energy, and rates using only the first 500 steps, corresponding to 163 TPs. The networks quickly learned the committor (Figure 3b, top row), with an absolute error of $p_B(x)$ below 0.05 in the reactive channel. We computed the PE from the AIMMD run1 data and projected the free energy on the committor estimated by the network (Figure 3c, top). The absolute error of $F(q)$ remains below $0.2 k_B T$ once it is aligned to the target.

We estimated the kinetics of the system by calculating ν with the Bayesian approach of eq 19 at different committor values λ . Again, we stopped at 500 steps and plotted the results of all the 3 runs to show the statistics (Figure 3d, top). The estimates from the same run are stable between the λ_A and λ_B thresholds (light area in Figure 3b). They range between 0.8 and 1.5 times

the reference rate. Each run took about 0.0035 cumulative simulations in ν^{-1} units. By addition of the sampling around the states for the PE extension, the total simulated time reaches $0.742 \nu^{-1} = 0.742 k_{AB}^{-1}$. We expect no transitions at all from an equilibrium simulation of the same length. In contrast, our method successfully provided accurate free energy and rates and also learned the reaction coordinate for the transition.

IV.B. Wolfe-Quapp Potential. The Wolfe-Quapp potential (Figure 3, bottom row) contains two alternative reaction channels, posing an additional challenge for the sampling. The two channels have different energy profiles and travel times across the isocommittor surfaces. Switching channels requires crossing a separation barrier of $2k_B T$ in the transition state.

AIMMD substantially increased the switching frequency compared to standard TPS, considering both the total steps and the accepted trajectories alone (Figure 4). As a result, it




Figure 4. Reactive channel switch during AIMMD simulations on the Wolfe-Quapp energy surface. The plots show the channel containing the last accepted TP with the number of steps (a) or accepted moves (b). The red lines represent AIMMD run1, while the blue lines are from the standard (uniform selection) TPS run0. The channel is determined by computing the Hausdorff distance to a reference path in Channel 1. The asterisks mark the channel switches.

took 400 steps on average to switch between channels. The switching frequency is closely related to the path decorrelation and indicates that AIMMD generates diverse trajectories.⁴¹ Occasionally selecting shooting points close to the states helped, as that promoted exploration of different configurations with a reasonable toll on exploitation: at convergence, the expected TPs' production rate (0.33) is 66% of the theoretical maximum (0.5, when selecting shooting points only at the transition state). We stress that we only used the instantaneous p_B in determining the selection probability. One could tune the exploration–exploitation trade-off by shaping $p_{sel}(p_B)$, although we found the uniform solution optimal for improving both the committor mode and the estimated PE accuracy. One could further optimize the selection bias to enhance the exploration at the cost of harvesting fewer TPs.

We put ourselves in a data-poor regime and took the first 500 steps (and 172 TPs) of the AIMMD run1. The resulting committor is less accurate in the low-energy channel (Figure 3b, bottom), albeit on par with the double well system overall. The speed boost of AIMMD allowed us to explore both channels; however, we put ourselves in a data-poor condition, where their relative importance is hard to infer. This especially affects the evaluated TPE (Figure S2b). Notwithstanding, we obtained an excellent estimate of the free energy profile as a function of the estimated committor, and the estimated rates remain within a 2-fold error between λ_A and λ_B (Figure 3d, bottom). In each run, we simulated approximately 0.071 of cumulative time in ν^{-1} units, $0.722 \nu^{-1}$ by including the additional sampling in the metastable states.

IV.C. Chignolin. Chignolin is the smallest peptide folding into a β -hairpin structure.⁷² The CLN025 variant exhibits a two-state behavior with the folded state showing remarkable stability^{83,84} (Figure 5b). The small size and short transition

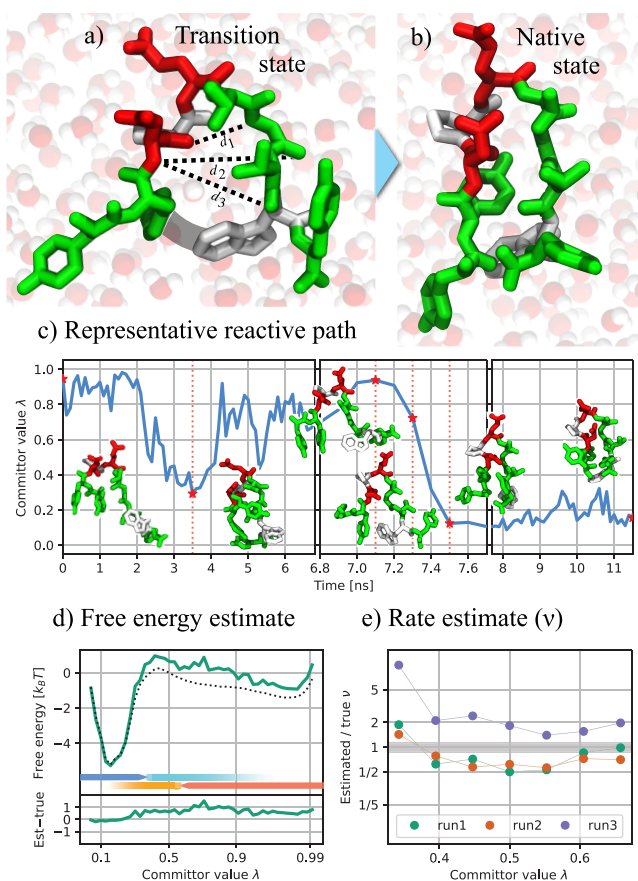


Figure 5. Characterizing the folding of chignolin. (a) Chignolin's transition state configuration, no-hydrogen licorice representation colored in white (nonpolar residues), red (acidic), and green (polar). We highlighted the formation of d_1 , d_2 , and d_3 hydrogen bonds and the association of Tyr2 and Trp9. (b) Chignolin's native structure represented as in a. (c) Representative folding trajectory (step 100 of run1), committor time series with renders of highlighted configurations. The model was trained on the first 50 steps of run1. (d) AIMMD run1, free energy as a function of the committor after 50 steps (solid line, top axis), free energy from long equilibrium simulations (dotted line, with uncertainty $<0.1k_B T$), and difference between the two (bottom axis). The arrows indicate the contributions of the simulations around A (blue), the AIMMD trial trajectories from A (cyan), the simulations around B (red), and the AIMMD trial trajectories from B (orange). (e) Bayesian rate estimate of ν at different committor values, each color denoting a different run. The gray area is the 95% confidence interval of ν from the long equilibrium simulations.

time, together with the formation of nonlocal structures, make it a good molecular system for testing our method and an entry point to studying more complex molecular events. From our equilibrium MD simulations at 340 K, totaling 120 μ s, we estimated a folding rate $k_{BA} = 2.5 \pm 0.5 \mu$ s⁻¹, an unfolding rate $k_{AB} = 0.28 \pm 0.05 \mu$ s⁻¹, a combined $\nu = 0.50 \pm 0.07 \mu$ s⁻¹, and a transition path time $\langle t_{AB} \rangle_{TP} = 11 \pm 2$ ns, in agreement with Lindorff-Larsen et al.⁷⁴

AIMMD provided an accurate transition mechanism, free energy, and rates of folding with only a handful of trajectories. We focus on AIMMD run1 (the other two runs yielded consistent results, see also Figure S5 for their free energy estimates). After only 50 steps (containing 20 TPs), the committor clearly distinguishes between the folded and unfolded state (Figure 5c) and is accurate when validated on an independent data set (Figure S4a). This limited number of trajectories produces a folding free energy profile within $1 k_B T$ of the expected value at the barrier (Figure 5d, uncertainty of the reference $<0.1k_B T$, see also Figure S4c for the individual contributions to the PE). Also, we could estimate $\nu = 0.25 \mu$ s⁻¹, which is less than a factor of 2 away from the reference value. The folding rate $k_{BA} = 2.49 \mu$ s⁻¹ is compatible with the estimate from very long equilibrium simulations; the two other runs were less accurate for this specific observable (13.4 and 10.1 μ s⁻¹), but still within an order of magnitude from the reference (Figure S4d). The run took 1.10 μ s of cumulative simulated time, corresponding to $0.548 \nu^{-1}$. We chose not to average the estimates from each run, because our primary focus was to assess the variance of independent runs. If enough resources are available, then this is always an effective strategy to assess convergence. Bootstrapping can be instead a practical way of estimating the variance for more challenging systems, where independent runs are not feasible.

When extending the estimate of the PE to the metastable states, we chose the largest M_A and M_B that ensured the stability of the reweighting factors (Figure S4b). We emphasize that the simulations used for the PE extension are short and are confined to the states. The AIMMD trial trajectories are thus essential to combine the two equilibrium path ensembles associated with each state with the proper weights. The additional data added up to 0.90 μ s, or $0.448 \nu^{-1}$.

Our method enables the characterization of the folding mechanism beyond free energy and rates. One could directly inspect the TPs or obtain explicit models of the committor with AIMMD (Figure 5c). Another way is to project the estimated PE on selected collective variables through eq 17. This has the advantage of allowing the expert to choose among domain-specific features and comes with no extra computational cost. It also produces a multifaceted representation of the process, putting the accent on different aspects and therefore rendering a more complete picture of this complex reorganization. We obtained the free energy on 6 pairs of collective variables discussed in the literature^{48,81} (Figure 6), all in good agreement with the equilibrium MD data.

Another advantage of estimating the PE is that it provides the effective (generalized) committor⁶¹ in any reduced space through eq C1. p_B encodes the progress of the reaction and complements the free energy information. For example, Figure 6d reveals alternative pathways to the “turn zipper” folding mechanism. Here, the Tyr2-Trp9 contacts, independent of the complete formation of Pro4-Gly7, are the real limiting factors in the reaction. Similar behavior was already reported in the literature and associated with hydrophobic collapse.^{83,85} The Asp3,N-Tyr8,O H-bond formation⁸⁶ is another crucial event at the barrier (Figure 6a): this, along with the fraction of native contacts and the $C\alpha$ -RMSD, stand out as the most important features in separating the folded and unfolded state.

IV.D. Performance Evaluation. We assessed the performance of our proposed algorithm under data-poor and data-rich regimes, in terms of both required computational resources and the quality of our estimates. In the data-poor scenario, we

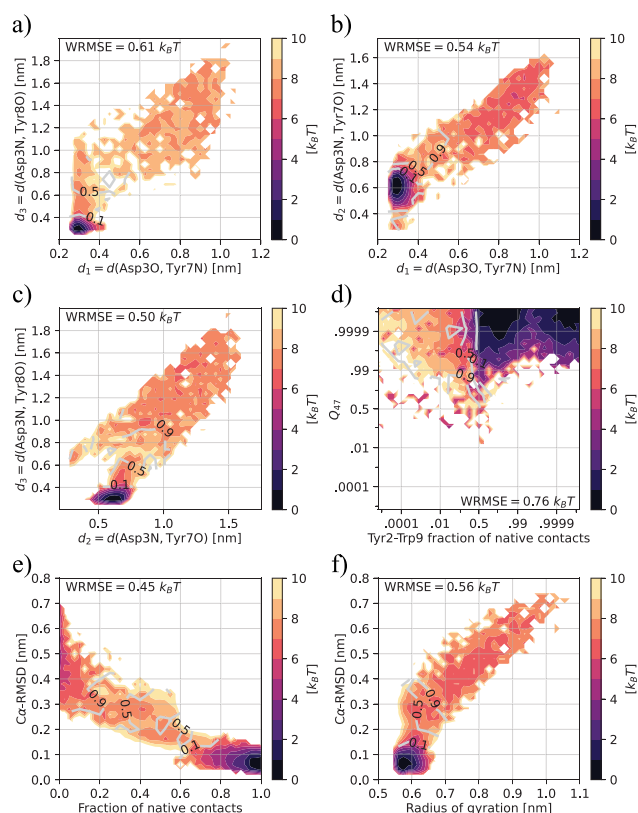


Figure 6. Chignolin's free energy profiles (colored contour) and effective committor (contour lines) extracted from the PE estimated from the first 250 steps of AIMMD run1, projected on 6 pairs of collective variables. We annotated the root-mean-square deviation $\sqrt{\langle (F - \bar{F})^2 \rangle_p}$ from the reference free energy \bar{F} obtained from long equilibrium MD simulations (WRMSE).

restricted the total simulations below ν^{-1} (half the average first return time). In this data-rich situation, we imposed no limitation on the computational resources to evaluate the highest expected accuracy achievable by our method. We stress that all of the estimates presented until now fall in the data-poor regime. The results including the complete simulations are collected in Figures S6, S7, and S8.

We verified AIMMD's capability of accelerating TPS from its early stages (Figure 7). The gain in production depends on how quickly the network converges to a reasonable committor

model and how peaked the TPE density is at the transition state. All of the systems promptly increased their production rate of TPs. Chignolin's higher chance of generating TPs from random configurations resulted in no significant difference in the number of accepted trajectories; in general, the system's complexity led to a higher variability within and among the runs. However, the most significant advantage of AIMMD is in obtaining the transition mechanism, free energies, and rates. To achieve this objective, learning the committor is crucial. Notably, training the network *a posteriori* on standard TPS results led to worse committor models than using AIMMD (Figure S9).

To evaluate the gain of the full procedure, we focused on the rate estimates as the most illustrative example since they require evaluation of the PE and TPE across the entire configuration space; plus, they are notably difficult to obtain with state-of-the-art techniques. To assess the computational resources, we also considered the total simulated time in units of ν^{-1} but excluded the equilibrium simulations around the states as they can be executed in an embarrassingly parallel way while doing path sampling. In the 2D systems, we replaced the simulations in the states with the reference ρ_A and ρ_B to isolate the error arising from the underlying approximations of our method and see how the accuracy scales with sampling. Training the networks on GPU and estimating the PE took a negligible fraction of the resources dedicated to MD.

At every stage of the AIMMD runs, the accuracy of the rate estimates consistently outperformed the predictions from equilibrium simulations of matching duration (Figure S10). Specifically, when the total duration is less than ν^{-1} (data-poor scenario), no back-and-forth transitions are expected to happen, and estimating the rates from a long equilibrium simulation cannot yield accurate results. The dashed lines in Figure S10 show the expected accuracy when the time surpasses ν^{-1} . In all of the systems, the relative error quickly dropped to a factor 2 after a few TPs and consistently reduced up to 10% with the increasing number of AIMMD trajectories. The directional rate constants (k_{fold} of chignolin, shown in Figures 8d and S10d) are less precise because of their further dependence on F_{AB} through eqs 20. At convergence, a small systematic error emerges from the discrete time interval between the trajectory points. In particular, the true λ_{max} is always bigger than the recorded one for an excursion from A, slightly altering the free energy profiles (Figure S6b).

When considering the total simulated time (Figure S10), the performance gain depends on the factor $\eta = \nu \cdot \langle t_{AB} \rangle_{\text{TP}}$ (the




Figure 7. Efficiency of generating TPs by AIMMD compared to standard TPS, for the double well (a), Wolfe-Quapp (b), and chignolin (c). The plots show the cumulative number of generated (solid lines) and accepted TPs (dotted lines), the latter based on the probabilities of eq 2. The red lines denote AIMMD run1, while the blue lines are from TPS run0. The dashed line (target) is the optimal upper-bound number of TPs given the chosen AIMMD selection probability.




Figure 8. Accuracy of estimated kinetics as a function of an increasing number of simulations for the double well (a), Wolfe-Quapp (b), and chignolin (c). The plot shows the maximum relative error of the ν estimate at $\lambda = 0.5$ (transition state) as a function of the number of steps. Each color corresponds to a different AIMMD run. For chignolin, we also plot the error of the folding rate, k_{fold} (d). The vertical lines mark the steps considered in Figures 3 and 5. The gray area is the 95% confidence interval of ν from equilibrium simulations.

“rareness” of the event). Despite longer TPs, AIMMD enabled a significant computational gain applied to the study of chignolin’s folding, also yielding reasonable folding rates (Figure 8d). Moreover, the two competing pathways did not compromise the results in the Wolfe-Quapp system, even in the case of no channel switches throughout the simulations.

Remarkably, substituting the instantaneous committor with its numerical computation did not significantly improve the rate estimate for the 2D systems aside from the early steps (Figure S11). We believe this is due to the robustness of the Bayesian approach while accounting for the ensembles’ fluctuations. In particular, deviations in the TPE density $\rho_{\text{TP}}(\lambda)$ are likely to reflect on $\rho(\lambda)$ and counterbalance in eq 19. Although it is possible to apply the method on TPS data trained *a posteriori*, the estimate is generally worse (Figure S11). Finally, the adaptive λ_A and λ_B thresholds improved the accuracy of the results when the network underfits the committor close to the states.

V. DISCUSSION

Understanding how molecules dynamically organize is key to revealing how they function and enabling technological and biomedical breakthroughs. This understanding comes in two ways: an accurate quantitative description and a qualitative explanation that allows us to obtain an intuitive insight and paves the way to formulating hypotheses and models. MD simulations can provide both—the first as free energies and rates and the second as mechanisms. However, standard simulation schemes usually cannot sample the time scales required for either goal.

Here, we present a path sampling algorithm that gives access at the same time to mechanisms, free energy, and rates for rare events in molecular systems. Our algorithm is general and straightforward and produces good free energy and rate estimates at a moderate computational cost. In essence, we run AIMMD simulations⁴¹ to sample trajectories that explore the transition region between two metastable states, which we then reweight and integrate with short, unbiased simulations in the states. Our algorithm provides a free energy profile that can be projected on any collective variable. The only requirement of our algorithm is a definition of the two states and an initial trajectory connecting them. AIMMD will adaptively learn how to optimally simulate TPs and learn the committor. In the basins, we run simple, unbiased simulations that can be conducted parallel to path sampling.

Our algorithm builds on established path sampling approaches and overcomes some of their limitations. TIS is a powerful method to obtain rates by seeding paths at interfaces

between two states. While able to yield very accurate rate estimates, TIS is computationally expensive. Here, we combined a path recycling scheme using straightforward two-way shooting simulations⁴⁶ with the RPE theory⁴³ to approximate the equilibrium path distribution between the states. From a more abstract viewpoint, next to a regular Markov chain sampling from the constrained TP distribution, our method creates a set of trajectories containing all trial paths that are all acceptable in the equilibrium path ensemble. Giving each path in this set the correct weight yields an approximation of the equilibrium PE. Our method thus builds on regular TPS, but it takes advantage of the trial paths in an unorthodox way, setting it apart from other path sampling approaches.

The high efficiency of our algorithm relies on two factors: AIMMD samples TPs with near-to-optimal efficiency; the committor model learned by AIMMD is the ideal reaction coordinate that simplifies the reweighting and makes it numerically more robust. Using the committor in combination with straightforward two-way-shooting simulations radically simplifies the algorithm in practice and enables us to recycle existing TPS simulation campaigns to extract free energy and rates *a posteriori*.

Rates are among the most challenging quantities to estimate in MD simulations. While many techniques exist to evaluate free energy profiles, rate calculations are much less established. On the other hand, free energy profiles are not observables—only free energy differences between metastable states are, while rates can often be measured in experiments, providing a natural way of comparing experiments and simulations. We anticipate that our algorithm and analogous approaches⁸⁷ will make the calculation of rates from MD simulation more accessible. By comparing calculated and measured rates, we can assess the systematic uncertainties arising from using semiempirical force fields, which generally were not parametrized on kinetic measurements.

AIMMD learns the committor in a self-consistent way. An accurate committor model guarantees by construction that the expected and actual sampling rates of TPs are equal (Figure 7). Learning an accurate committor is key to obtaining accurate estimates of free energy and rates. The convergence of these can be assessed by monitoring their variations as a function of an increasing number of steps.

Selecting good features to describe a molecular system is based on physical symmetries, experience, and trial and error. System-specific features are available in many scientific domains, and using them facilitates training and enables a more insightful mechanistic interpretation. When these are

unavailable, one can use general features, such as pairwise distances (as we did to describe the folding of chignolin). The neural network that models the committor in AIMMD can be fed with all of the possible reasonable features simultaneously. There are no practical restrictions on the dimensionality of the input featurization. For instance, we could use more than 2,000 pairwise distances to describe chignolin. Scaling up to hundreds of thousands of features should not be a problem. Learning relevant features from a Cartesian representation of molecular systems remains an outstanding problem.

Despite the many advancements, important challenges remain. Our algorithm focuses on characterizing rare molecular events between the two states. While AIMMD and the underlying committor theory generalize to transition between multiple states,^{41,88} in practice, it might be more efficient to reduce this problem to a collection of pairwise transitions. The definition of two states is not always straightforward. However, it requires only order parameters that do not have to resolve the transition. Also, state definitions can be iteratively refined by using the committor. After a first simulation campaign, configurations with committor values close to 0 and 1 can be used as new, more accurate state boundaries. The correlation along the chain of sampled TPs is still a great challenge.⁸⁹ We showed how AIMMD alleviates this problem by speeding up the switching between alternative reactive channels, but many steps are still required. Integrating generative AI approaches, as recently proposed by Dellago and co-workers, might provide a solution.^{51,82}

Our algorithm relies on many short, unbiased simulations. The clear advantage is that the dynamics are not distorted, and reweighting is necessary only to obtain the correct stationary distribution in the transition region. This also means that the longest time scale that one must be able to simulate is the duration of the TPs. This time scale ultimately determines the applicability of our method. Importantly, the typical TP duration is usually exponentially shorter than the typical lifetime in the states and can be on the scales of nanoseconds, even for large and complex systems.⁹⁰ In practice, one should be able to perform at least ~ 100 steps of our algorithm, taking into account that it can be easily parallelized.⁴¹ For some systems, it will still be challenging to sample even a few TPs in a reasonable time. In addition, the energy wells of the states could be so deep that sampling excursions that overlap with the reweighted paths in the transition region could be impractical. Using a static biasing potential can help in both cases.²³

Our algorithm is simple to use and data-efficient. It builds on highly efficient simulation packages like GROMACS⁹¹ and OpenMM.⁹² In this way, it seamlessly capitalizes on the new software and conventional force-field developments. But it will also take advantage of the latest exciting developments in generative AI for conformational sampling.⁹³ Path sampling simulations are increasingly more attractive for investigating rare molecular events.

APPENDIX A. CROSSING PROBABILITY ALONG THE COMMITTOR

We demonstrate eq 10 in the nonrestrictive assumption of Markovian dynamics.^{95,96} If $\lambda = \lambda_A$ or $\lambda = 1$, the proof is trivial—the case of $\lambda = 1$ follows from the definition of committor. Assume now that $0 < \lambda_A < \lambda < 1$; \mathbf{x} is a trajectory leaving A at $t = 0$ and crossing the committor value λ_A . Let t'

be the first time when $p_B(\mathbf{x}(t')) = \lambda_A$. \mathbf{x} can continue from t' in 3 possible ways:

1. it reaches A before B without crossing λ ;
2. it reaches A before B after crossing λ ;
3. it reaches B before A, hence crossing λ ,

with probability p_1 , p_2 , and p_3 , respectively. We find that $p_1 = 1 - P_A(\lambda | \lambda_A)$, $p_3 = p_B(\mathbf{x}(t')) = \lambda_A$, and $p_2 = P_A(\lambda | \lambda_A) (1 - \lambda)$ as combination of two independent events: \mathbf{x} crossing λ from λ_A before A and \mathbf{x} reaching A from λ before B. Since $p_1 + p_2 + p_3 = 1$:

$$\begin{aligned} 1 - P_A(\lambda | \lambda_A) + \lambda_A + P_A(\lambda | \lambda_A)(1 - \lambda) &= 1 \\ P_A(\lambda | \lambda_A)(-1 + 1 - \lambda) + \lambda_A &= 0 \end{aligned} \quad (\text{A1})$$

from which $P_A(\lambda | \lambda_A) = \lambda_A/\lambda$. A related result was obtained in eq 9 of ref 97.

APPENDIX B. CONNECTION TO THE RPE THEORY

We explain the connection between the weighting scheme described in eq 12a with TIS⁴⁴ and the RPE approach.⁴³ In previous studies, the transition region between A and B was partitioned into interfaces defined by a progress coordinate, which generally was not the committor. TIS then required sampling a large number of unbiased trajectories at each interface. From these simulations, one can estimate an ensemble for each interface and then merge them according to the global crossing probability estimated, e.g., with WHAM^{58–60}). This method proved to be computationally demanding and highly sensitive to the interface selection.⁹⁸

Recently, Brotzakis and Bolhuis developed the virtual interface exchange (VIE) algorithm, which populates the TIS interfaces with TPS trial trajectories.^{46,47,99} Having access to the committor, the optimal reaction coordinate,⁴² our method can be seen as a limit case of VIE with an infinite number of interfaces defined after the committor.

Carrying the analogy with the RPE approach, each AIMMD trial trajectory $\mathbf{x}^{(i)}$ is now the unique representative of the interface

$$I_{\lambda}^{(i)} = \{x | p_B(x) = \lambda^{(i)}\} \quad (\text{B1})$$

defined by its shooting point, with $p_B(x_{\text{sp}}^{(i)}) = \lambda^{(i)}$. Once the AIMMD sampling has converged, $x_{\text{sp}}^{(i)}$ is correctly sampled from the Boltzmann distribution restricted to $I_{\lambda}^{(i)}$. This result has two reasons: first, the shooting point is drawn from the TPE distribution biased along the committor (the selection probability does not alter ρ_{TPE} within $I_{\lambda}^{(i)}$); second, from eq 4 the Boltzmann and TPE distribution are proportional when projected onto the committor. As a consequence, the configurations in the trial paths are already Boltzmann-distributed when restricting ourselves to $I_{\lambda}^{(i)}$. As we explain in the main text, this feature is essential for a good reweighting and PE estimate. We retain the initial AIMMD steps (before convergence) in the computations to optimize resource usage.

The RPE combines all the associated interfaces by assigning weights related to the crossing probability.⁴³ Each interface $I_{\lambda}^{(i)}$ gives its own “crossing statistics” (or histogram) starting from $\lambda^{(i)}$:

$$H_A^{(i)}(\lambda) = \tilde{h}_A[\mathbf{x}^{(i)}] \theta(\lambda - \lambda^{(i)}) \theta(\lambda_{\text{max}}^{(i)} - \lambda) \quad (\text{B2})$$

in which $\mathbf{x}^{(i)}$ is the only contributor. The statistics before $\lambda^{(i)}$ do not provide meaningful information because we deliberately forced the trajectory to reach that point.

The total crossing statistics from λ_A can be obtained in two ways:

1. from the $H_A^{(i)}(\lambda)$, through a function $f(\lambda)$:

$$H_A(\lambda) = f(\lambda) \sum_{i=1}^{n_{\text{TPS}}} H_A^{(i)}(\lambda) \quad (\text{B3})$$

2. from the individually weighted trajectories (as in Figure 2d):

$$K_A(\lambda) = \sum_{i=1}^{n_{\text{TPS}}} w_A^{(i)} \theta(\lambda_{\text{max}}^{(i)} - \lambda) \quad (\text{B4})$$

The two approaches are apparently very different: $H_A^{(i)}$ does not contribute to the total statistics before its shooting interface, as $H_A^{(i)}(\lambda < \lambda^{(i)}) = 0$, whereas $\mathbf{x}^{(i)}$ does.

Rogal et al.^{43,46} showed that $K_A(\lambda)$ converges to $H_A(\lambda)$ if we set

$$w_A^{(i)} = \tilde{h}_A[\mathbf{x}^{(i)}] f(\lambda_{\text{max}}^{(i)}) \quad (\text{B5})$$

in the limit of infinite trajectories. Therefore, eq B5 gives the optimal weights for reconstructing $\mathcal{P}_{A,\lambda_A}^+$. Here, we impose $H_A(\lambda)$ to match $P_A(\lambda | \lambda_A)$. By comparing eqs 10 and B3, we obtain

$$f(\lambda) \propto \frac{1}{\lambda m_A(\lambda)} \quad (\text{B6})$$

where m_A is defined as in eq 11. By injecting eq B5, we finally get eq 12a.

APPENDIX C. NORMALIZATION CONSTANTS OF THE RPE

We derive the constraint of eq 13. Bolhuis and Lechner⁶¹ proved that the effective committor function of the coordinates q is

$$p_B(q) = \frac{\rho_B(q)}{\rho_A(q) + \rho_B(q)} \quad (\text{C1})$$

If the $q = \lambda$ are committor values themselves, then $p_B(\lambda) = \lambda$. At the transition state ($\lambda = 0.5$) we obtain the identity:

$$\frac{1}{2} = \frac{\rho_B(\lambda = 0.5)}{\rho_A(0.5) + \rho_B(0.5)} \quad (\text{C2})$$

satisfied by $\rho_A(0.5) = \rho_B(0.5)$ and $c_A \tilde{\rho}_A(0.5) = c_B \tilde{\rho}_B(0.5)$.

ASSOCIATED CONTENT

Data Availability Statement

For the purpose of Open Access, the author has applied a CC-BY license to any Author Accepted Manuscript version arising from this submission. All data needed to evaluate the conclusions in the paper are openly available in the paper and the Supporting Information, and in the “Source code and data for AIMMD and PE estimate” repository at <http://doi.org/10.5281/zenodo.8048453>. **Code Availability Statement:** We performed path sampling simulations adapting the AIMMD Python package developed by Jung,⁴¹ which builds upon OpenPathSampling (OPS), a Python library for TPS simulations.⁹⁴ We performed the reweighting and projections described in section II.G with the custom-written PathEnsemble Python package. The PathEnsemble code, the data featured in this paper, and the scripts for running the simulations and

analyzing the results are available at the repository [10.5281/zenodo.8048453](https://zenodo.org/record/8048453).

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.3c00821>.

Supplementary Figures S1–S11 (PDF)

AUTHOR INFORMATION

Corresponding Author

Roberto Covino – Frankfurt Institute for Advanced Studies, Frankfurt am Main 60438, Germany; Goethe University Frankfurt, Frankfurt am Main 60438, Germany; orcid.org/0000-0003-0884-0402; Email: covino@fias.uni-frankfurt.de

Authors

Gianmarco Lazzeri – Frankfurt Institute for Advanced Studies, Frankfurt am Main 60438, Germany; Goethe University Frankfurt, Frankfurt am Main 60438, Germany; orcid.org/0000-0002-0593-5572

Hendrik Jung – Goethe University Frankfurt, Frankfurt am Main 60438, Germany; Department of Theoretical Biophysics, Max Planck Institute of Biophysics, Frankfurt am Main 60438, Germany; orcid.org/0000-0002-2159-0391

Peter G. Bolhuis – Van't Hoff Institute for Molecular Sciences, University of Amsterdam, Amsterdam 1090GD, The Netherlands; orcid.org/0000-0002-3698-9258

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.3c00821>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

We thank Drs. Jutta Rogal, Gerhard Hummer, and Attila Szabo for stimulating discussions and helpful comments. G.L. and R.C. acknowledge the support of the Frankfurt Institute of Advanced Studies, the LOEWE Center for Multiscale Modelling in Life Sciences of the state of Hesse, the CRC 1507: Membrane-associated Protein Assemblies, Machineries, and Supercomplexes, and computational resources and support by the SURFsara National Supercomputing and e-Science Support Center in The Netherlands, the Center for Scientific Computing of the Goethe University, and the Jülich Supercomputing Centre. G.L. was supported by a grant from the HPC-Europa3 program and acknowledges support of the iQbio graduate school of the Goethe University. R.C. acknowledges the support of the International Max Planck Research School on Cellular Biophysics. H.J. acknowledges support by the Max Planck Society.

REFERENCES

- (1) Dror, R. O.; Dirks, R. M.; Grossman, J.; Xu, H.; Shaw, D. E. Biomolecular Simulation: A Computational Microscope for Molecular Biology. *Annual Review of Biophysics* **2012**, *41*, 429–452.
- (2) Massobrio, C.; Du, J.; Bernasconi, M.; Salmon, P. S., Eds. *Molecular Dynamics Simulations of Disordered Materials*, 1st ed.; Springer International Publishing, 2015.
- (3) Frenkel, D.; Smit, B. *Understanding molecular simulation: from algorithms to applications*, 3rd ed.; Elsevier, 2023.
- (4) Adcock, S. A.; McCammon, J. A. Molecular Dynamics: Survey of Methods for Simulating the Activity of Proteins. *Chem. Rev.* **2006**, *106*, 1589–1615.

- (5) Shaw, D. E.; Maragakis, P.; Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Eastwood, M. P.; Bank, J. A.; Jumper, J. M.; Salmon, J. K.; Shan, Y.; Wriggers, W. Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science* **2010**, *330*, 341–346.
- (6) Swope, W. C.; Pitera, J. W.; Suits, F. Describing Protein Folding Kinetics by Molecular Dynamics Simulations. 1. Theory. *J. Phys. Chem. B* **2004**, *108*, 6571–6581.
- (7) Best, R. B. Atomistic molecular simulations of protein folding. *Curr. Opin. Struct. Biol.* **2012**, *22*, 52–61.
- (8) Deng, Y.; Roux, B. Computations of Standard Binding Free Energies with Molecular Dynamics Simulations. *J. Phys. Chem. B* **2009**, *113*, 2234–2246.
- (9) Psachoulia, E.; Marshall, D. P.; Sansom, M. S. P. Molecular Dynamics Simulations of the Dimerization of Transmembrane α -Helices. *Acc. Chem. Res.* **2010**, *43*, 388–396.
- (10) Urbanc, B.; Betnel, M.; Cruz, L.; Bitan, G.; Teplow, D. B. Elucidation of Amyloid β -Protein Oligomerization Mechanisms: Discrete Molecular Dynamics Study. *J. Am. Chem. Soc.* **2010**, *132*, 4266–4280.
- (11) Nishimura, T.; Lazzeri, G.; Mizushima, N.; Covino, R.; Tooze, S. A. Unique amphipathic α -helix drives membrane insertion and enzymatic activity of ATG3. *Science Advances* **2023**, *9*, eadh1281.
- (12) Matsumoto, M.; Saito, S.; Ohmine, I. Molecular dynamics simulation of the ice nucleation and growth process leading to water freezing. *Nature* **2002**, *416*, 409–413.
- (13) Allen, T. W.; Andersen, O. S.; Roux, B. Molecular dynamics — potential of mean force calculations as a tool for understanding ion permeation and selectivity in narrow channels. *Biophys. Chem.* **2006**, *124*, 251–267.
- (14) Wales, D. J. The energy landscape as a unifying theme in molecular science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2005**, *363*, 357–377.
- (15) Onuchic, J. N.; Wolynes, P. G. Theory of protein folding. *Curr. Opin. Struct. Biol.* **2004**, *14*, 70–75.
- (16) Freddolino, P. L.; Harrison, C. B.; Liu, Y.; Schulten, K. Challenges in protein-folding simulations. *Nat. Phys.* **2010**, *6*, 751–758.
- (17) Peters, B. *Reaction rate theory and rare events*, 1st ed.; Elsevier, 2017.
- (18) E, W.; Vanden-Eijnden, E. Transition-Path Theory and Path-Finding Algorithms for the Study of Rare Events. *Annu. Rev. Phys. Chem.* **2010**, *61*, 391–420.
- (19) Roux, B. Transition rate theory, spectral analysis, and reactive paths. *J. Chem. Phys.* **2022**, *156*, 134111.
- (20) Hartmann, C.; Banisch, R.; Sarich, M.; Badowski, T.; Schütte, C. Characterization of Rare Events in Molecular Dynamics. *Entropy* **2014**, *16*, 350–376.
- (21) Lane, T. J.; Shukla, D.; Beauchamp, K. A.; Pande, V. S. To milliseconds and beyond: challenges in the simulation of protein folding. *Curr. Opin. Struct. Biol.* **2013**, *23*, 58–65.
- (22) Bernardi, R. C.; Melo, M. C.; Schulten, K. Enhanced sampling techniques in molecular dynamics simulations of biological systems. *Biochimica et Biophysica Acta (BBA) - General Subjects* **2015**, *1850*, 872–877.
- (23) Hénin, J.; Lelièvre, T.; Shirts, M. R.; Valssson, O.; Delemotte, L. Enhanced sampling methods for molecular dynamics simulations. *Living Journal of Computational Molecular Science* **2022**, *4*, 1583.
- (24) Fiorin, G.; Klein, M. L.; Hénin, J. Using collective variables to drive molecular dynamics simulations. *Mol. Phys.* **2013**, *111*, 3345–3362.
- (25) Gershenson, A.; Gosavi, S.; Faccioli, P.; Wintrode, P. L. Successes and challenges in simulating the folding of large proteins. *J. Biol. Chem.* **2020**, *295*, 15–33.
- (26) Kleiman, D. E.; Nadeem, H.; Shukla, D. Adaptive Sampling Methods for Molecular Dynamics in the Era of Machine Learning. *arXiv* **2023**, No. 2307.09664v1, DOI: 10.48550/arXiv.2307.09664.
- (27) Chen, H.; Ogdén, D.; Pant, S.; Cai, W.; Tajkhorshid, E.; Moradi, M.; Roux, B.; Chipot, C. A companion guide to the string method with swarms of trajectories: Characterization, performance, and pitfalls. *J. Chem. Theory Comput.* **2022**, *18*, 1406–1422.
- (28) Zuckerman, D. M.; Chong, L. T. Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annual Review of Biophysics* **2017**, *46*, 43–57.
- (29) Bolhuis, P. G.; Chandler, D.; Dellago, C.; Geissler, P. L. Transition Path Sampling: Throwing Ropes Over Rough Mountain Passes, in the Dark. *Annu. Rev. Phys. Chem.* **2002**, *53*, 291–318.
- (30) Krivov, S. V. Nonparametric Analysis of Nonequilibrium Simulations. *J. Chem. Theory Comput.* **2021**, *17*, 5466–5481.
- (31) Russo, J. D.; Copperman, J.; Zuckerman, D. M. Iterative trajectory reweighting for estimation of equilibrium and non-equilibrium observables. *arXiv* **2020**, No. 2006.09451v1, DOI: 10.48550/arXiv.2006.09451.
- (32) Metzner, P.; Schütte, C.; Vanden-Eijnden, E. Transition Path Theory for Markov Jump Processes. *Multiscale Modeling & Simulation* **2009**, *7*, 1192–1219.
- (33) Bolhuis, P. G. Transition-path sampling of β -hairpin folding. *Proc. Natl. Acad. Sci. U. S. A.* **2003**, *100*, 12129–12134.
- (34) Knott, B. C.; Haddad Momeni, M.; Crowley, M. F.; Mackenzie, L. F.; Gotz, A. W.; Sandgren, M.; Withers, S. G.; Stahlberg, J.; Beckham, G. T. The Mechanism of Cellulose Hydrolysis by a Two-Step, Retaining Cellobiohydrolase Elucidated by Structural and Transition Path Sampling Studies. *J. Am. Chem. Soc.* **2014**, *136*, 321–329.
- (35) Bolhuis, P. G.; Dellago, C.; Chandler, D. Reaction coordinates of biomolecular isomerization. *Proc. Natl. Acad. Sci. U. S. A.* **2000**, *97*, 5877–5882.
- (36) Vlugt, T. J. H.; Dellago, C.; Smit, B. Diffusion of isobutane in silicalite studied by transition path sampling. *J. Chem. Phys.* **2000**, *113*, 8791–8799.
- (37) Bouleffelf, S. E.; Ravikovitch, P. I.; Sholl, D. S. Modeling Diffusion of Linear Hydrocarbons in Silica Zeolite LTA Using Transition Path Sampling. *J. Phys. Chem. C* **2015**, *119*, 15643–15653.
- (38) Escobedo, F. A.; Borrero, E. E.; Araque, J. C. Transition path sampling and forward flux sampling. Applications to biological systems. *J. Phys.: Condens. Matter* **2009**, *21*, 333101.
- (39) E, W.; Ren, W.; Vanden-Eijnden, E. Transition pathways in complex systems: Reaction coordinates, isocommittor surfaces, and transition tubes. *Chem. Phys. Lett.* **2005**, *413*, 242–247.
- (40) Bolhuis, P. G.; Dellago, C. Practical and conceptual path sampling issues. *European Physical Journal Special Topics* **2015**, *224*, 2409–2427.
- (41) Jung, H.; Covino, R.; Arjun, A.; Leitold, C.; Dellago, C.; Bolhuis, P. G.; Hummer, G. Machine-guided path sampling to discover mechanisms of molecular self-organization. *Nature Computational Science* **2023**, *3*, 334–345.
- (42) Berezhkovskii, A. M.; Szabo, A. Relations among Unidirectional Fluxes at Equilibrium, Committors, and First Passage and Transition Path Times. *J. Phys. Chem. B* **2022**, *126*, 6624–6628.
- (43) Rogal, J.; Lechner, W.; Juraszek, J.; Ensing, B.; Bolhuis, P. G. The reweighted path ensemble. *J. Chem. Phys.* **2010**, *133*, 174109.
- (44) van Erp, T. S.; Bolhuis, P. G. Elaborating transition interface sampling methods. *J. Comput. Phys.* **2005**, *205*, 157–181.
- (45) Hall, S. W.; Díaz Leines, G.; Sarupria, S.; Rogal, J. Practical guide to replica exchange transition interface sampling and forward flux sampling. *J. Chem. Phys.* **2022**, *156*, 200901.
- (46) Brotzakis, Z. F.; Bolhuis, P. G. Approximating free energy and committor landscapes in standard transition path sampling using virtual interface exchange. *J. Chem. Phys.* **2019**, *151*, 174111.
- (47) Frenkel, D. *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Vol. 1*; Springer Berlin Heidelberg, 2006; pp 127–137.
- (48) Harada, R.; Kitao, A. Exploring the Folding Free Energy Landscape of a β -Hairpin Mini-protein, Chignolin, Using Multiscale Free Energy Landscape Calculation Method. *J. Phys. Chem. B* **2011**, *115*, 8806–8812.
- (49) Bolhuis, P. G.; Swenson, D. W. H. Transition Path Sampling as Markov Chain Monte Carlo of Trajectories: Recent Algorithms,

Software, Applications, and Future Outlook. *Advanced Theory and Simulations* **2021**, *4*, 2000237.

(50) Mullen, R. G.; Shea, J.-E.; Peters, B. Easy Transition Path Sampling Methods: Flexible-Length Aimless Shooting and Permutation Shooting. *J. Chem. Theory Comput.* **2015**, *11*, 2421–2428.

(51) Falkner, S.; Coretti, A.; Dellago, C. Enhanced Sampling of Configuration and Path Space in a Generalized Ensemble by Shooting Point Exchange. *arXiv* **2023**, No. 2302.08757v2, DOI: 10.48550/arXiv.2302.08757.

(52) Berezhkovskii, A. M.; Szabo, A. Committors, first-passage times, fluxes, Markov states, milestones, and all that. *J. Chem. Phys.* **2019**, *150*, 054106.

(53) Chen, H.; Roux, B.; Chipot, C. Discovering Reaction Pathways, Slow Variables, and Commitor Probabilities with Machine Learning. *J. Chem. Theory Comput.* **2023**, *19*, 4414–4426.

(54) Hummer, G. From transition paths to transition states and rate coefficients. *J. Chem. Phys.* **2004**, *120*, 516–523.

(55) Peters, B. Recent advances in transition path sampling: accurate reaction coordinates, likelihood maximisation and diffusive barrier-crossing dynamics. *Mol. Simul.* **2010**, *36*, 1265–1281.

(56) van Erp, T. S.; Moroni, D.; Bolhuis, P. G. A novel path sampling method for the calculation of rate constants. *J. Chem. Phys.* **2003**, *118*, 7762–7774.

(57) Cabriolu, R.; Skjelbred Refsnes, K. M.; Bolhuis, P. G.; van Erp, T. S. Foundations and latest advances in replica exchange transition interface sampling. *J. Chem. Phys.* **2017**, *147*, 152722 DOI: 10.1063/1.4989844.

(58) Kumar, S.; Rosenberg, J. M.; Bouzida, D.; Swendsen, R. H.; Kollman, P. A. weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *J. Comput. Chem.* **1992**, *13*, 1011–1021.

(59) Stelzl, L. S.; Kells, A.; Rosta, E.; Hummer, G. Dynamic Histogram Analysis To Determine Free Energies and Rates from Biased Simulations. *J. Chem. Theory Comput.* **2017**, *13*, 6328–6342.

(60) Ferguson, A. L. BayesWHAM: A Bayesian approach for free energy estimation, reweighting, and uncertainty quantification in the weighted histogram analysis method. *J. Comput. Chem.* **2017**, *38*, 1583–1605.

(61) Bolhuis, P. G.; Lechner, W. On the Relation Between Projections of the Reweighted Path Ensemble. *J. Stat. Phys.* **2011**, *145*, 841–859.

(62) Best, R. B.; Hummer, G. Reaction coordinates and rates from transition paths. *Proc. Natl. Acad. Sci. U. S. A.* **2005**, *102*, 6732–6737.

(63) Berezhkovskii, A.; Hummer, G. Single-file transport of water molecules through a carbon nanotube. *Physical review letters* **2002**, *89*, 064503.

(64) Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **2019**, *32*, 8026–8037.

(65) He, K.; Zhang, X.; Ren, S.; Sun, J. *Computer Vision – ECCV 2016*; Springer International Publishing, 2016; pp 630–645.

(66) Daqi, G.; Yan, J. Classification methodologies of multilayer perceptrons with sigmoid activation functions. *Pattern Recognition* **2005**, *38*, 1469–1482.

(67) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2014**, No. 1412.6980v1, DOI: 10.48550/arXiv.1412.6980.

(68) Ying, X. An Overview of Overfitting and its Solutions. *Journal of Physics* **2019**, *1168*, 022022.

(69) Quapp, W. A growing string method for the reaction pathway defined by a Newton trajectory. *J. Chem. Phys.* **2005**, *122*, 174106.

(70) Covino, R.; Woodside, M. T.; Hummer, G.; Szabo, A.; Cossio, P. Molecular free energy profiles from force spectroscopy experiments by inversion of observed committors. *J. Chem. Phys.* **2019**, *151*, 154115.

(71) Honda, S.; Akiba, T.; Kato, Y.; Sawada, Y.; Sekijima, M.; Ishimura, M.; Ooishi, A.; Watanabe, H.; Odahara, T.; Harata, K. Crystal structure of a ten-amino acid protein. *J. Am. Chem. Soc.* **2008**, *130*, 15327–31.

(72) Yasuda, S.; Hayashi, T.; Kinoshita, M. Physical origins of the high structural stability of CLN025 with only ten residues. *J. Chem. Phys.* **2014**, *141*, 105103.

(73) Jo, S.; Kim, T.; Iyer, V. G.; Im, W. CHARMM-GUI: A web-based graphical user interface for CHARMM. *J. Comput. Chem.* **2008**, *29*, 1859–1865.

(74) Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. How Fast-Folding Proteins Fold. *Science* **2011**, *334*, 517–520.

(75) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. How Robust Are Protein Folding Simulations with Respect to Force Field Parameterization? *Biophys. J.* **2011**, *100*, L47–L49.

(76) Bauer, P.; Hess, B.; Lindahl, E. *GROMACS 2022 Manual*; Zenodo, 2022; DOI: 10.5281/zenodo.6103568.

(77) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **2007**, *126*, 014101.

(78) Humphrey, W.; Dalke, A.; Schulten, K. VMD: Visual molecular dynamics. *J. Mol. Graphics* **1996**, *14*, 33–38.

(79) McGibbon, R. T.; Beauchamp, K. A.; Harrigan, M. P.; Klein, C.; Swails, J. M.; Hernández, C. X.; Schwantes, C. R.; Wang, L.-P.; Lane, T. J.; Pande, V. S. MDTraj: A Modern Open Library for the Analysis of Molecular Dynamics Trajectories. *Biophys. J.* **2015**, *109*, 1528–1532.

(80) Best, R. B.; Hummer, G.; Eaton, W. A. Native contacts determine protein folding mechanisms in atomistic simulations. *Proc. Natl. Acad. Sci. U. S. A.* **2013**, *110*, 17874–17879.

(81) Satoh, D.; Shimizu, K.; Nakamura, S.; Terada, T. Folding free-energy landscape of a 10-residue mini-protein, chignolin. *FEBS Lett.* **2006**, *580*, 3422–3426.

(82) Falkner, S.; Coretti, A.; Romano, S.; Geissler, P. L.; Dellago, C. Conditioning Boltzmann generators for rare event sampling. *Machine Learning: Science and Technology* **2023**, *4*, 035050.

(83) Davis, C. M.; Xiao, S.; Raleigh, D. P.; Dyer, R. B. Raising the Speed Limit for β -Hairpin Formation. *J. Am. Chem. Soc.* **2012**, *134*, 14476–14482.

(84) Sumi, T.; Koga, K. Theoretical analysis on thermodynamic stability of chignolin. *Sci. Rep.* **2019**, *9*, 5186.

(85) Dinner, A. R.; Lazaridis, T.; Karplus, M. Understanding β -hairpin formation. *Proc. Natl. Acad. Sci. U. S. A.* **1999**, *96*, 9068–9073.

(86) McKiernan, K. A.; Husic, B. E.; Pande, V. S. Modeling the mechanism of CLN025 beta-hairpin formation. *J. Chem. Phys.* **2017**, *147*, 104107.

(87) Palacio-Rodríguez, K.; Pietrucci, F. Free Energy Landscapes, Diffusion Coefficients, and Kinetic Rates from Transition Paths. *J. Chem. Theory Comput.* **2022**, *18*, 4639–4648.

(88) Rogal, J.; Bolhuis, P. G. Multiple state transition path sampling. *J. Chem. Phys.* **2008**, *129*, 224107.

(89) Ghamari, D.; Hauke, P.; Covino, R.; Faccioli, P. Sampling rare conformational transitions with a quantum computer. *Sci. Rep.* **2022**, *12*, 16336.

(90) Okazaki, K.; Wöhlert, D.; Warnau, J.; Jung, H.; Özkan Yildiz; Kühlbrandt, W.; Hummer, G. Mechanism of the electroneutral sodium/proton antiporter PaNhaP from transition-path shooting. *Nat. Commun.* **2019**, *10*, 1742.

(91) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, *1–2*, 19–25.

(92) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology* **2017**, *13*, e1005659.

(93) Noé, F.; Olsson, S.; Köhler, J.; Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science* **2019**, *365*, eaaw1147.

(94) Swenson, D. W. H.; Prinz, J.-H.; Noe, F.; Chodera, J. D.; Bolhuis, P. G. OpenPathSampling: A Python Framework for Path Sampling Simulations. 1. Basics. *J. Chem. Theory Comput.* **2019**, *15*, 813–836.

(95) Best, R. B.; Hummer, G. Diffusion models of protein folding. *Phys. Chem. Chem. Phys.* **2011**, *13*, 16902.

(96) Berezhkovskii, A.; Szabo, A. Time scale separation leads to position-dependent diffusion along a slow coordinate. *J. Chem. Phys.* **2011**, *135*, 074108.

(97) Vanden-Eijnden, E.; Venturoli, M.; Ciccotti, G.; Elber, R. On the assumptions underlying milestoning. *J. Chem. Phys.* **2008**, *129*, 174102.

(98) Kratzer, K.; Arnold, A.; Allen, R. J. Automatic, optimized interface placement in forward flux sampling simulations. *J. Chem. Phys.* **2013**, *138*, 164112.

(99) Coluzza, I.; Frenkel, D. Virtual-Move Parallel Tempering. *ChemPhysChem* **2005**, *6*, 1779–1783.