
Path Integral Stochastic Optimal Control for Sampling Transition Paths

Lars Holdijk^{*1,2} Yuanqi Du^{*2} Ferry Hooft³ Priyank Jaini⁴ Bernd Ensing^{3,5} Max Welling²

Abstract

We consider the problem of *Sampling Transition Paths*. Given two metastable conformational states of a molecular system, e.g. a folded and unfolded protein, we aim to sample the most likely transition path between the two states. Sampling such a transition path is computationally expensive due to the existence of high free energy barriers between the two states. To circumvent this, previous work has focused on simplifying the trajectories to occur along specific molecular descriptors called Collective Variables (CVs). However, finding CVs is not trivial and requires chemical intuition. For larger molecules, where intuition is not sufficient, using these CV-based methods biases the transition along possibly irrelevant dimensions. Instead, this work proposes a method for sampling transition paths that consider the entire geometry of the molecules. To achieve this, we first relate the problem to recent work on the Schrödinger bridge problem and stochastic optimal control. Using this relation, we construct a method that takes into account important characteristics of molecular systems such as second-order dynamics and invariance to rotations and translations. We demonstrate our method on the commonly studied Alanine Dipeptide, but also consider larger proteins such as Polyproline and Chignolin.

1. Introduction

Modeling non-equilibrium systems in natural sciences involves analyzing dynamical behaviour that occur with very low probability known as *rare events*, i.e. particular instances of the dynamical system that are atypical. The kinetics of many important molecular processes, such

^{*}Equal contribution ¹OATML, University of Oxford ²AMLab, University of Amsterdam ³Computational Chemistry Group, University of Amsterdam ⁴Google Brain ⁵AI4Science, University of Amsterdam. Correspondence to: Lars Holdijk <larsholdijk@gmail.com>.

as phase transitions, protein folding, conformational changes, and chemical reactions, are all dominated by these rare events. One way to sample these rare events is to follow the time evolution of the underlying dynamical system using Molecular Dynamic (MD) simulations until a reasonable number of events has been observed. However, this is computationally highly inefficient due to the large time-scales involved, which are typically related to the presence of high energy and/or entropy barriers between the metastable states. Thus, the main problem is: *How can we efficiently sample trajectories that give rise to these rare but interesting transition events?*

Numerous enhanced sampling methods such as steered MD (Jarzynski, 1997), umbrella sampling (Torrie & Valleau, 1977), constrained MD (Carter et al., 1989), transition path sampling (Dellago & Bolhuis, 2009), and many more, have been developed to deal with the problem of *rare events* in molecular simulation. Most of these methods bias the dynamical system with well-chosen geometric descriptors of the transition, so-called *collective variables* (CVs), that allow the system to overcome high-energy transition barriers and sample these rare events. The performance of these enhanced sampling techniques is critically dependent on the choice of these CVs. However, choosing appropriate CVs for all but the simplest molecular systems is fraught with difficulty, as it relies on human intuition, insights about the molecular system, and trial and error.

A key alternative to sampling these rare transition paths is to model an alternate dynamical system that allows sampling these rare trajectories in an optimal manner (Ahmed et al., 2006; Jack, 2020; Todorov, 2009). (Rose et al., 2021) proposes an RL solution to modelling such a system.

In this paper, we consider the problem of sampling *rare* transition paths by developing an alternative dynamical system using *path integral stochastic optimal control* (Kappen, 2005; 2007; Kappen & Ruiz, 2016; Theodorou et al., 2010). Our method models this alternative dynamics of the system by applying an external control policy to each of the atoms in the molecule. In particular, the method learns the external control policy such that it minimizes the amount of external work needed to overcome the lowest barrier and transition the molecular system from an initial meta-stable state to a final one. The method does not require any knowledge

of CVs to sample these rare trajectories. Furthermore, we draw connections between sampling rare transition paths and the Schrödinger bridge problem (Schrödinger, 1931; 1932). Subsequently, we demonstrate that stochastic optimal control is well suited to solving these problems by extending the work of (Kappen & Ruiz, 2016) for molecular systems by incorporating Hamiltonian dynamics and equivariance constraints.

Our main contributions in this paper are:

- We demonstrate the equivalence between sampling transition paths, the Schrödinger bridge problem, and path integral stochastic optimal control (§2).
- We extend the PICE algorithm (Kappen & Ruiz, 2016) for molecular dynamics by incorporating inductive biases in the form of equivariances and Hamiltonian dynamics (§3).
- Our proposed method does not require any prior knowledge of CVs, which is important for modeling large and complex molecular transitions for which CVs are unknown (§2-3).
- We demonstrate the efficacy of our method on conformational transitions in three molecular systems of varying complexity, namely Alanine Dipeptide, Polyproline and, Chignolin (§4).

2. Preliminaries and Problem Setup

In the following sections we will first define notations and provide key definitions, before discussing a mathematical framework to solve the main problem of sampling rare events.

Consider a system evolving over time where $\pi(\mathbf{x})$ is the distribution of states \mathbf{x} and $\pi_i(\mathbf{x}_i|\mathbf{x}_{i-1})$ a Markovian transition kernel. We can define a distribution over trajectories generated by said evolving system as

$$\pi(\mathbf{x}(\tau)) := \pi(\mathbf{x}_0) \cdot \prod_{i=1}^{\tau} \pi_i(\mathbf{x}_i|\mathbf{x}_{i-1}). \quad (1)$$

Here, $\mathbf{x}(\tau)$ defines a trajectory of states of length τ . While the evolution of the system is continuous, for convenience we discretize time and represent the trajectory as an ordered sequence of states $\mathbf{x}(\tau) = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\tau\}$.

In our application of interest, we are interested in sampling trajectories from this distribution with the boundary condition that the initial \mathbf{x}_0 and terminal state \mathbf{x}_τ are drawn from a pre-specified marginal distributions π_0 and π_τ , respectively. In the chemistry context, these marginal distributions describe stable states of the molecular system located at local minima of the free energy surface. For example, these

stable states can be reactants and products of chemical reactions, or native and unfolded states of protein. Unfortunately, these stable states are often separated by high free energy barriers. Trajectories $\mathbf{x}(\tau)$ sampled starting in \mathbf{x}_0 are thus unlikely to terminate in the target state \mathbf{x}_τ .

In this paper, we therefore construct an alternative sampling approach that generates trajectories that are still likely under the distribution $\pi(\mathbf{x}(\tau))$ while also adhering to the boundary conditions by crossing the high free energy barrier. Formally, we aim to find an alternative Markovian dynamical system $\hat{\pi}(\mathbf{x}(\tau))$ with pre-defined marginals π_0 and π_τ that is as close to $\pi(\mathbf{x}(\tau))$ as possible, i.e.

$$\hat{\pi}^*(\mathbf{x}(\tau)) := \arg \min_{\hat{\pi}(\mathbf{x}(\tau)) \in \mathcal{D}(\pi_0, \pi_\tau)} \mathbb{D}_{\text{KL}}(\hat{\pi}(\mathbf{x}(\tau)) \| \pi(\mathbf{x}(\tau))) \quad (2)$$

Here, $\mathcal{D}(\pi_0, \pi_\tau)$ is the space of path measures with marginals π_0 and π_τ .

This problem, of finding the alternative dynamical system, is also known as the Schrödinger Bridge Problem (SBP) (Schrödinger, 1931; 1932). To solve the problem of sampling transition paths that cross the high free energy barrier we can take inspiration from computational methods for solving the Schrödinger Bridge Problem such as discussed in (Vargas et al., 2021; De Bortoli et al., 2021).

However, in this work we specifically explore an alternative approach to solving the SBP that lends itself well to modelling the chemical nature of our problem by learning a biased force to be added to the system. In the next section, to set the stage for this novel approach, we relate the problem of sampling transition paths as a path integral stochastic optimal control problem and establish an equivalence between learning alternative dynamical systems for sampling rare events, Schrödinger bridge problem, and stochastic optimal control.

2.1. Sampling Transition Paths through Stochastic Optimal Control

The original dynamics of the system, as given in Equation (1), can be reformulated as a stochastic process:

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t) dt + \mathbf{G}(\mathbf{x}_t, t) \cdot d\boldsymbol{\varepsilon}_t, \quad t \in [0, \tau] \quad (3)$$

where $\mathbf{f} : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}^d$ and $\mathbf{G} : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}^{d \times d}$ are deterministic functions representing the drift and volatility of the system. The stochastic process $\boldsymbol{\varepsilon}_t$ is a Brownian motion with variance ν .

We approximate the system dynamics described in Equation (3) with another dynamical system by applying an external bias potential (or control) $\mathbf{u}(\mathbf{x}_t, t) \in \mathbb{R}^d$ to the system that pushes it over the transition state barriers. We write

these dynamics as follows:

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t) dt + \mathbf{G}(\mathbf{x}_t, t) \cdot \left(\mathbf{u}(\mathbf{x}_t, t) dt + d\boldsymbol{\varepsilon}_t \right), \quad t \in [0, \tau] \quad (4)$$

The control $\mathbf{u}(\mathbf{x}_t, t)$ does not directly operate on the system dynamics here but is rather controlled through the same control matrix \mathbf{G} as the Brownian motion. We will see later on that this formulation is important when considering the second order dynamics of the system.

Given a trajectory $\mathbf{x}(\tau) = (\mathbf{x}_0, \dots, \mathbf{x}_\tau) \in \mathbb{R}^{\tau \times d}$ generated through the SDE in Equation (4), we define the cost of this trajectory under control \mathbf{u} following (Kappen, 2007; Theodorou et al., 2010) as:

$$C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t) = \frac{1}{\lambda} \left(\varphi(\mathbf{x}_\tau) + \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \mathbf{u}(\mathbf{x}_t, t) + \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \boldsymbol{\varepsilon}_t \right) \quad (5)$$

where φ denotes the terminal cost, λ is a constant and \mathbf{R} is the cost of taking action \mathbf{u} in the current state and is generally given as a weight matrix for a quadratic control cost. The goal then is to find the optimal control \mathbf{u}^* that minimizes the expected cost:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \mathbb{E}_{\tau, \boldsymbol{\varepsilon}_t} [C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t)] \quad (6)$$

where the expectation is taken over trajectories τ sampled using the SDE under control \mathbf{u} .

Note that the last cost term in Equation (5) relating the Brownian motion and the control is unusual and does not have a clear intuition. However, as we will see in the next section, this term plays an important role when relating the cost to a KL-divergence. Additionally, as discussed in Thijssen & Kappen (2015), the additional cost vanishes under expectation ($\mathbb{E}_{\tau, \boldsymbol{\varepsilon}_t} [\mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \boldsymbol{\varepsilon}_t] = 0$) and does thus not influence the optimal control \mathbf{u}^* given by Equation (6).

Relation to sampling transition paths: Interestingly, the objective in Equation (6) is exactly related to the problem of sampling transition paths as given in Equation (2). As (Kappen & Ruiz, 2016) establish, Equation (4) defines a probability distribution $\pi_{\mathbf{u}}(\mathbf{x}(\tau))$ over trajectories $\mathbf{x}(\tau)$ through:

$$\pi_{\mathbf{u}}(\mathbf{x}(\tau)) = \prod_{t=0}^{\tau} \mathcal{N}(\mathbf{x}_{t+1} | \boldsymbol{\mu}_t, \Sigma_t) \quad (7)$$

with $\boldsymbol{\mu}_s = \mathbf{x}_s + \mathbf{f}(\mathbf{x}_s, s) dt + \mathbf{G}(\mathbf{x}_s, s)(\mathbf{u}(\mathbf{x}_s, s) dt)$ and $\Sigma_s = \mathbf{G}(\mathbf{x}_s, s)^T \nu \mathbf{G}(\mathbf{x}_s, s)$.

For different \mathbf{u} , these distributions are related through the Girsanov Theorem (Cameron & Martin, 1944). As shown

in Appendix A, if we make the common assumption that control cost \mathbf{R} and the variance of the Brownian motion are inversely correlated as $\lambda \mathbf{R}^{-1} = \nu$, we can obtain:

$$\log \frac{\pi_{\mathbf{u}}(\mathbf{x}(\tau))}{\pi_0(\mathbf{x}(\tau))} = \frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \mathbf{u}(\mathbf{x}_t, t) + \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \boldsymbol{\varepsilon}_t \quad (8)$$

Here $\pi_0(\mathbf{x}(\tau))$ denotes the distribution over trajectories with no control (i.e. $\mathbf{u} = 0$) as $\pi_0(\mathbf{x}(\tau))$. For reference, the assumption relating the control cost and the variance of the Brownian motion is a common trait of the set of control problems referred to as *Path Integral Stochastic Optimal Control* (Kappen, 2005).

We observe that the right-hand side of Equation (8) can also be found in the definition of the control cost in Equation (5), including the additional cost term related to the Brownian noise. As Kappen & Ruiz (2016) show, we can thus use Equation (8) to rewrite the objective in Equation (6) as:

$$\pi_{\mathbf{u}^*} = \arg \min_{\pi_{\mathbf{u}}} \mathbb{E}_{\mathbf{x}(\tau) \sim \pi_{\mathbf{u}}} \left[\frac{1}{\lambda} \varphi(\mathbf{x}_\tau) + \mathbb{D}_{\text{KL}}(\pi_{\mathbf{u}}(\mathbf{x}(\tau)) \| \pi_0(\mathbf{x}(\tau))) \right] \quad (9)$$

This objective is an approximation of the Schrödinger Bridge formulation in Equation (2) where the constraints on the marginal distributions are replaced by a regularization term in the form of the terminal cost. When the expected terminal cost dominates the KL-divergence the found distribution should be very similar.

Indeed, this connection between the Schrödinger Bridge Problem and stochastic optimal control has previously been established (Chen et al., 2016; Pavon et al., 2021). Here, through the formulations in Equations (2) and (9), we also establish the equivalence between sampling transition paths, Schrödinger bridge problem, and stochastic optimal control. This allows us to utilize solutions for finding the optimal control in Equation (9) for the aforementioned problems.

Optimal Control Policy: (Kappen & Ruiz, 2016) introduced the Path Integral Cross Entropy (PICE) method for solving Equation (9). The PICE method derives an explicit expression for the optimal policy and distribution $\pi_{\mathbf{u}^*}$ when $\lambda = \nu \mathbf{R}$ given by

$$\pi_{\mathbf{u}^*} = \frac{1}{\eta(\mathbf{x}, t)} \pi_{\mathbf{u}}(\mathbf{x}(\tau)) \exp(-C(\mathbf{x}(\tau), \mathbf{u})) \quad (10)$$

where $\eta(\mathbf{x}, t) = \mathbb{E}_{\tau \sim \pi_{\mathbf{u}}} [\exp(-C(\mathbf{x}(\tau), \mathbf{u}))]$ is the normalization constant. This establishes the optimal distribution $\pi_{\mathbf{u}^*}$ as a reweighing of any distribution induced by an arbitrary control \mathbf{u} . Similar to importance sampling, depending

on the choice of the proposal distribution $\pi_{\mathbf{u}}$ the estimator variance can greatly differ. Thus, the objective is to find the \mathbf{u} that best approximates \mathbf{u}^* .

PICE, subsequently, achieves this by minimizing the KL-divergence between the optimal controlled distribution $\pi_{\mathbf{u}^*}$ and a parameterized distribution $\pi_{\mathbf{u}_\theta}$ using gradient descent as follows:

$$\frac{\partial \mathbb{D}_{\text{KL}}(\pi_{\mathbf{u}^*} | \pi_{\mathbf{u}_\theta})}{\partial \theta} = -\frac{1}{\eta} \mathbb{E}_{\mathbf{x}(\tau) \sim \pi_{\mathbf{u}_\theta}} \left[\exp(-C(\mathbf{x}(\tau), \mathbf{u}_\theta)) \sum_{t=0}^{\tau} (\mathbf{R}\varepsilon_t \cdot \frac{\partial \mathbf{u}_\theta}{\partial \theta}) \right] \quad (11)$$

Similar to the optimal control in Equation (10), the gradient used to minimize the KL-divergence is found by reweighing for each sampled trajectory the gradient of the control policy \mathbf{u}_θ by the cost of said trajectory. A full algorithm for determining this gradient and training the policy \mathbf{u}_θ is provided in Algorithm 1.

Thus, PICE provides an iterative gradient descent method to learn a parameterized policy (or control) \mathbf{u}_θ and subsequently a distribution over paths $\mathbf{x}(\tau)$. We can then use this learned control, \mathbf{u}_θ , to approximate the solution for sampling transition paths as well as the Schrödinger bridge problem.

In this section, we have set up our main problem of sampling rare transition paths and established its relationship to both the Schrödinger bridge problem and stochastic optimal control. Subsequently, we discussed an iterative gradient descent based method for solving the optimal control problem. In the next section, we will extend this iterative algorithm to consider the entire geometry of the molecular system by incorporating Hamiltonian dynamics using an augmented state space \mathbf{x}_t , and symmetries by learning a policy \mathbf{u}_θ .

3. Path Integral Optimal Control for Sampling Transition Paths

We consider a molecule consisting of n atoms with an initial and final configuration $\mathbf{r}_0 \in \mathbb{R}^{3 \times n}$ and $\mathbf{r}_\tau \in \mathbb{R}^{3 \times n}$ i.e. we are given a vector defining the 3D positions of each atom in the molecule. Thus, a direct method to sample transition paths $\mathbf{r}(\tau)$ for this problem is to learn a control \mathbf{u}_θ acting directly on the positions \mathbf{r} of the molecule using the iterative gradient descent method discussed in Section 2. However, the collective behaviour of the atoms and molecules are governed by classical molecular dynamics i.e. Newtonian equations of motion:

$$d\mathbf{r} = \mathbf{v}(t) dt, \quad \text{and}, \quad d\mathbf{v} = \mathbf{a}(t) dt$$

Algorithm 1: Training Policy \mathbf{u}_θ

Input: $\mathbf{r}_0, \mathbf{r}_\tau$: Initial and target molecular positions,
 $U(\cdot)$: Potential Energy function,
 $\varphi(\cdot)$: Terminal cost,
 $\mathbf{u}_\theta(\cdot, \cdot)$: Initial parameterized policy,
 N : Number of trajectories sampled per update,
 τ : Time horizon,
 ν : Variance of Brownian noise,
 \mathbf{R} : Control cost matrix,
 μ : Learning rate,
 dt : Time discretization step

while not converged do
 ▷ Generate trajectories with current policy \mathbf{u}_θ
 $\lambda \leftarrow \mathbf{R}\nu$;
 $n \leftarrow 0$;
while $n < N$ **do**
 ▷ Initialize initial trajectory state
 $(\mathbf{r}_{n,0}, \mathbf{v}_{n,0}, t) \leftarrow (\mathbf{r}_0, \mathbf{0}, 0)$;
while $t < (\tau/dt)$ **do**
 ▷ Sample Brownian noise and action
 $\varepsilon_{n,t} \sim \mathcal{N}(0, \nu)$;
 $\mathbf{u}_{n,t} \leftarrow \mathbf{u}_\theta(\mathbf{r}_{n,t}, t)$;
 ▷ Update positions and velocity
 $\mathbf{r}_{n,t+1} \leftarrow \mathbf{r}_{n,t} + \mathbf{v}_{n,t} \cdot dt$;
 $\mathbf{v}_{n,t+1} \leftarrow \mathbf{v}_{n,t} - (\nabla_{\mathbf{r}_{n,t}} U(\mathbf{r}_{n,t}) + \mathbf{u}_{n,t} + \varepsilon_{n,t}) \cdot dt$;
 $t \leftarrow t + 1$;
end
 ▷ Determine trajectory cost and gradient
 $C_n \leftarrow \frac{1}{\lambda} (\varphi(\mathbf{r}_{n,\tau}) + \sum_{i=0}^{\tau} \mathbf{u}_{n,i}^T \mathbf{R} \mathbf{u}_{n,i} + \mathbf{u}_{n,i}^T \mathbf{R} \varepsilon_{n,i})$;
 $\Delta \theta_n \leftarrow \exp(-C_n) + \sum_{i=0}^{\tau} \partial \mathbf{u}_{n,i} \partial \theta$;
 $n \leftarrow n + 1$;
end
 ▷ Determine gradient normalization and perform policy update
 $\eta \leftarrow \sum_{i=0}^N \exp(-C_i)$;
 $\theta \leftarrow \theta + \frac{\mu}{\eta} \sum_{i=0}^N \Delta \theta_i$;
end

where $\mathbf{v}(t) \in \mathbb{R}^{3 \times n}$ is the velocity and $\mathbf{a}(t) \in \mathbb{R}^{3 \times n}$ is acceleration given by $\mathbf{a}(t) = \frac{\nabla_{\mathbf{r}} U(\mathbf{r})}{m}$ where $U(\mathbf{r})$ is the potential energy of the system m is the mass. The potential energy of a system is defined by a heavily parameterized sum of pairwise empirical potential functions, such as harmonic bonds, angle potentials, inter-molecular electrostatic and Van der Waals potentials. Here, we compute this using the OpenMM framework (Eastman et al., 2017). As such, we need to adapt the dynamical system defined in Equation (4) to incorporate these molecular dynamics. Intuitively, we achieve this by limiting the control $\mathbf{u} \in \mathbb{R}^{3 \times n}$ to act as an additive force on the velocity of the system. This change in velocity via the control \mathbf{u} ultimately changes the position of the system.

Incorporating second order dynamics: Formally, we incorporate the second order dynamics of the system defined above by considering an augmented state space: Let

$\mathbf{x}_0 := (\mathbf{r}_0, \mathbf{v}_0) \in \mathbb{R}^{3 \times n} \times \mathbb{R}^{3 \times n}$ be the initial configuration of the system defining the initial positions and velocities of each atom and $\mathbf{x}_\tau := (\mathbf{r}_\tau, \mathbf{v}_\tau)$ be the final configuration. We, thus, model the dynamical system in Equation (4) as:

$$\begin{aligned} \underbrace{\begin{pmatrix} d\mathbf{r}_t \\ d\mathbf{v}_t \end{pmatrix}}_{d\mathbf{x}_t} &= \underbrace{\begin{pmatrix} \mathbf{v}_t \\ -\nabla_{\mathbf{r}_t} U(\mathbf{r}_t) \end{pmatrix}}_{\mathbf{f}(\mathbf{x}_t, t)} dt \\ &+ \underbrace{\begin{pmatrix} \mathbf{0}_{3n} \\ \mathbb{I}_{3n} \end{pmatrix}}_{\mathbf{G}(\mathbf{x}_t, t)} \cdot (\mathbf{u}(\mathbf{x}_t, t) dt + d\boldsymbol{\varepsilon}_t), \end{aligned} \quad (12)$$

$t \in [0, \tau]$. We note that due to the choice of $\mathbf{G}(\mathbf{x}_t, t)$ in Equation (12) above, the additional bias force, $\mathbf{u}(\mathbf{x}_t, t)$, applied to the system only influences the acceleration and velocity of the atoms. $d\mathbf{r}_t$ is solely influenced by the velocity \mathbf{v}_t , thus conforming to the classical molecular dynamics of the system.

Invariance to rotations and translations: Secondly, the molecules in consideration are invariant w.r.t. translations and 3D rotations i.e. the molecular orientations achieved along a transition path need to incorporate this equivariance w.r.t. SE(3) group. For this purpose, we need to make the terminal cost function, $\varphi(\mathbf{x}_\tau)$, in Equation (5) to be equivariant. We enforce this by defining the terminal cost as the exponentiated pairwise distance between atoms which is commonly used distance metric (Shi et al., 2021) that is invariant to rotations and translations i.e. $\varphi(\mathbf{r}_t) = \exp \sum_{i,j}^n (d_{ij}(\mathbf{r}_t) - d_{ij}(\mathbf{r}_\tau))^2$ where $d_{ij}(\mathbf{r}_t) = \|(\mathbf{r}_t)_i - (\mathbf{r}_t)_j\|_2^2$.

Physics inspired policy network (\mathbf{u}_θ): The main learnable component of our path integral stochastic optimal control method for sampling transition paths is the policy network \mathbf{u}_θ . Following the discussion above and formalized in Equation (12), we can interpret the control \mathbf{u}_θ as an additive bias force applied to the system. In this work, we consider two different design approaches to modelling \mathbf{u}_θ . In our first approach, we model \mathbf{u}_θ as a neural network that predicts the bias force on the system in which case the velocity evolves as $d\mathbf{v} = (\nabla_{\mathbf{r}_t} U(\mathbf{r}_t) + \mathbf{u}_{\theta,t}) dt$. Alternatively, in our second approach, we model \mathbf{u}_θ as a network predicting the bias potential energy. In this case, the corresponding force, $\mathbf{F}(\mathbf{r}_t)$, applied to the system is calculated by backpropagating through the network, $\mathbf{F}(\mathbf{r}_t) := \nabla_{\mathbf{r}_t} u_{\theta,t}$. The change in velocity is then given by $d\mathbf{v} = (\nabla_{\mathbf{r}_t} U(\mathbf{r}_t) + \mathbf{F}(\mathbf{r}_t)) dt$. Additionally, \mathbf{u}_θ or u_θ could be implemented using recent advances in physics inspired equivariant neural networks (Cohen & Welling, 2016; Satorras et al., 2021) that take into account the SE(3) symmetry of the system.

4. Experiments

We evaluate our path integral stochastic optimal control method for sampling transition paths with three different molecular systems, namely (i) **Alanine Dipeptide**, a small amino acid with well-studied transition paths, (ii) **Polyproline**, a small protein with two distinct conformations with different helix orientations, and (iii) **Chignolin**, an artificial mini-protein studied to understand the folding process of proteins. We begin by detailing the experimental setup.

Molecular Dynamics Simulation: We simulate the molecular dynamics following Equation (12). For the drift function $\mathbf{f}(\mathbf{x}_t, t)$, we evaluate the potential function (and subsequently the gradient) using the OpenMM engine (Eastman et al., 2017). We define the forcefield used by OpenMM for each considered molecular system in their respective following sections. Furthermore, we use the Velocity Verlet with Velocity Randomization (VVVR) integrator (Sivak et al., 2013) within OpenMM at a temperature of 300 K with a collision rate of 1.0 ps^{-1} wherein the bias potential, $\mathbf{u}(\mathbf{x}_t, t)$, given by the policy network is integrated as a custom external force.

Policy Network, $\mathbf{u}(\mathbf{x}_t, t)$: We implement the policy network as a 6 layer MLP with ReLU activation for all our experiments below. The width of the layers of the policy network is dependent on the number of atoms in the molecule under consideration. We implement all code in Pytorch. We ran the experiments on a single GPU (either an NVIDIA RTX3080 or RTX2080) and our code is available on [GitHub](#).

4.1. Alanine Dipeptide

Alanine Dipeptide is an extensively studied molecule (Tobias & Brooks III, 1992; Rossky & Karplus, 1979; Head-Gordon et al., 1991; Swenson et al., 2018) for developing and testing enhanced sampling methods due to ready availability of its two CVs (ϕ, ψ). The conformation transition for Alanine Dipeptide can thus be understood in terms of these two dihedral angles ϕ and ψ as displayed in Figure 1 A. Prior work has, thus, focused on transforming from the initial configuration (see Figure 1.A) to the final configuration (Figure 1.E) by rotating these CVs. As we discussed previously, a major advantage of our method is that we do not require the knowledge of CVs to sample a transition path. However, in our experiment for Alanine Dipeptide, we will use these CVs to compare the quality of the trajectory sampled by our method.

OpenMM and policy network setup: We use the amber 99sb-ildn force field (Lindorff-Larsen et al., 2010) without any solvent, a time-step of 1.0 fs for the VVVR integrator and a cutoff of 1 nm for the Particle Mesh Ewald (PME) method (Essmann et al., 1995). For our experiment, we

	τ fs	Temp. K	EPD (\downarrow) $\text{nm} \times 10^{-3}$	THP (\uparrow) %	ETP (\downarrow) kJ mol^{-1}
Force Prediction	500	300	2.07	41.1 %	0.68
Energy Prediction	500	300	1.25	89.2 %	-5.21
MD w. fixed timescale	500	300	7.92	0%	-
	500	1500	7.47	0%	-
	500	4500	6.33	0%	-
	500	9000	6.82	1.7 %	1019.83
MD w/ fixed timescale	34810	1500	1.88	100%	551.51
	48683	4500	2.01	100%	1647.35

Table 1. Benchmark scores for the proposed method and extended MD baselines. From-left-to-right: Time-horizon τ representing the trajectory length (note that we take one policy step every 1 fs), simulation temperature, Expected Pairwise distance (EPD), Target Hit Percentage (THP), and Energy Transition Point (ETP). ETP can only be calculate when a trajectory reaches the target. All metrics are averaged over 1000 trajectories except for MD w/ fixed timescale which is ran only for 10 trajectories.

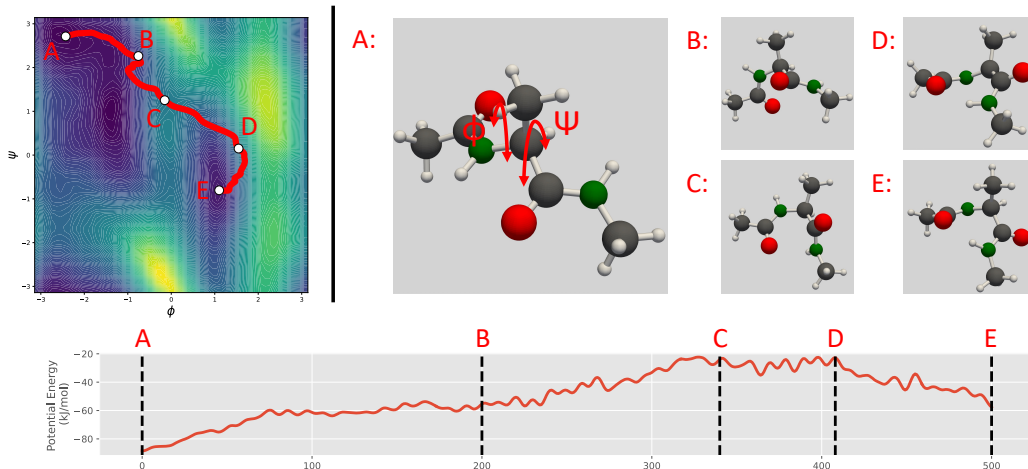


Figure 1. Visualization of a trajectory sampled with the proposed method. *Left*: The sampled trajectory projected on the free energy landscape of Alanine Dipeptide as a function of two CVs *Right*: Conformations along the sampled trajectory: A) starting conformation showing the CV dihedral angles, B-D) intermediate conformations with C being the highest energy point on the trajectory, and E) final conformation, which closely aligns with the target conformation. *Bottom*: Potential energy during transition. Letters represent the same configurations in the transition.

consider both the design choices for the policy network, $u_\theta(x_t, t)$, discussed in Section 3 i.e. directly predicting the force and predicting the energy. We trained the policy networks for 15,000 roll-outs with a time horizon of 500 fs each consisting of 16 samples. A gradient update was made to the policy network after each roll-out with a learning rate of 10^{-5} . The Brownian motion has 0.1 standard deviation.

Baseline and evaluation metrics: We compare our method to MD simulations with extended time-horizon and increased system temperatures to sample transition paths. To our knowledge, there are no fixed quantitative metrics in the literature to compare different methods that sample transition paths. Thus, we introduce here three metrics to evaluate the quality of transition paths: (i) *Expected Pairwise Distance* (EPD) measures the euclidean distance between the final conformation in the trajectory and the target conforma-

tion, reflecting the goal of the transition to end in the target state, (ii) *Target Hit Percentage* (THP) assures that the final configuration is also close in terms of CVs by measuring the percentage of trajectories correctly transforming these CVs, and (iii) *Energy Transition Point* (ETP) which evaluates the capacity of each method to find transition paths that cross the high-energy barrier at a low point by taking the maximum potential energy of the molecule along the trajectory. A good trajectory will be one that passes through the minimal high-energy barrier and ETP aims to measure this. We provide more details in Appendix B.1.1.

Results: We first visualize the trajectory generated by the *energy prediction* policy in Figure 1. The trajectory in Figure 1 demonstrates that the control policy transforms the molecule from the initial position (A) to the final position (E) by transitioning over the barrier with the least energy at (C).

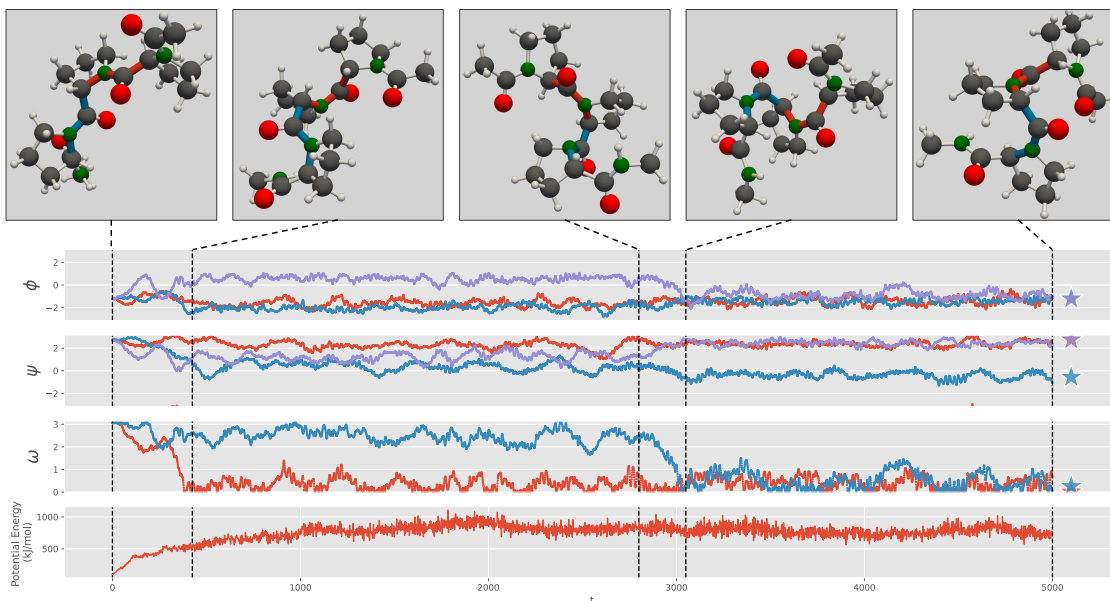


Figure 2. Visualization of the Polyproline transformation from PP-II to PP-I. From-top-to-bottom 5 stages of the transition, ψ CVs, ϕ CVs, ω CVs, and Potential Energy. For the CVs multiple instances of the same dihedral angles can be found in a single molecule. Stars indicate target CV states. Colored bonds represent the bonds involved in the ω CV.

Interestingly, the trajectory follows the expected transitions in the CVs without them being explicitly specified e.g. the transition path visualized on the left in Figure 1 shows that the molecule first rotates the dihedral angle associated with CV ϕ in (A \rightarrow B), then gradually rotates along both ψ and ϕ in (B \rightarrow C \rightarrow D), and finally rotates ψ in (D \rightarrow E) to reach the final configuration. As expected, we observe that the potential energy goes up during the transition until it reaches the top of the energy barrier (C). After this point, the molecule settles down in its new low-energy state.

Next, we compare the performance of the trajectories sampled using the *force* and *energy* predicting policy networks with MD simulations on the metrics introduced before in Section 3. We find that the trajectories generated by both the policy networks outperform the MD baselines, but the more physics-aligned energy predicting policy performs best under our metrics. This policy network consistently reaches the target conformation both in terms of full geometry and the CVs orientation. Furthermore, our policy network generates these trajectories in a significantly shorter time than temperature enhanced MD simulations without a fixed timescale. When we do limit MD to run for the same timescale as the proposed method, we found that, in contrast to the proposed method, temperature enhanced MD simulations are unable to generate successful trajectories.

4.2. Polyproline Helix

Polyproline is a helix-shaped protein structure that consists of repeating proline residues. Polyproline helix can form

two different conformations namely *Polyproline-I* (PP-I) and *Polyproline-II helix* (PP-II) (Moradi et al., 2009; 2010). These conformations can be distinguished by their respective helix rotation. PP-I forms a compact right-handed helix due to its peptide bonds having cis-isomers while PP-II has trans-isomer peptide bonds and forms a left-handed helix. Furthermore, the backbone of the polyproline helix also contains two different dihedral angles. We will refer to these peptide bonds and dihedral-angles as the ω , ϕ and ψ CVs respectively. Polyproline can have varying lengths due to its repeated structure. In our experiment, we consider the polyproline trimer with 3 proline residues transitioning from PP-II to PP-I.

OpenMM and policy network setup: We initialize OpenMM with the amber protein.ff14SBonlysc forcefield and gbn2 as the implicit solvent forcefield. The VVVR integrator had a timestep of 2.0 fs and a cutoff of 5 nm for PME. The proposed method was ran for a total of 10,000 fs (resulting in 5,000 policy steps). The policy network was trained over 500 rollouts with 25 samples each using a learning rate of 3×10^{-5} and a standard deviation of 0.1 for the Brownian motion.

Results: We visualize the transformation of the three collective variables (ω , ϕ , ψ) as well as the corresponding potential energy of the conformation in Figure 2 for a sampled transition path from our trained policy network. The ω CV admits the biggest change for the transition from PP-I going from 180° to 0° . We observe that the transition path sampled by our method aligns with the expected changes in

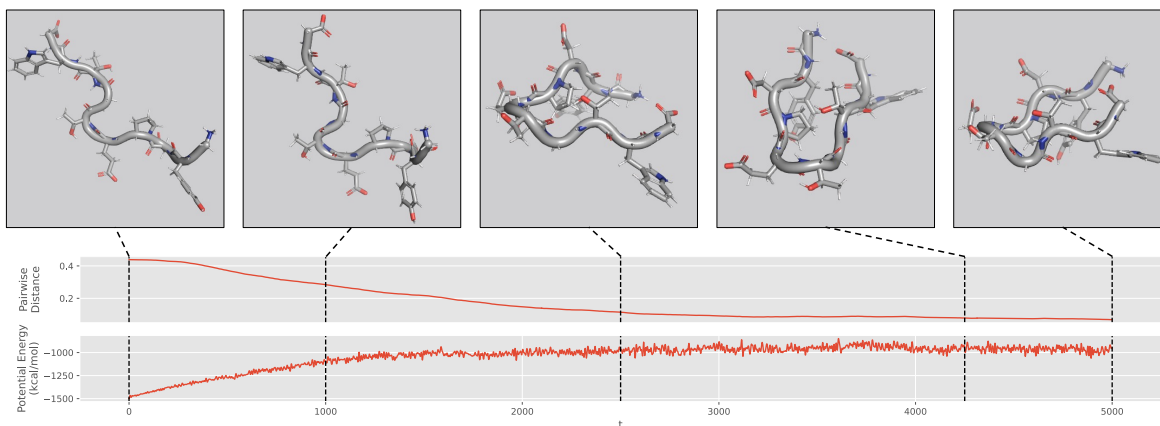


Figure 3. Visualization of the Chignolin folding process. *Top*: 5 stages of the folding process, *Middle*: Pairwise distance wrt to the target conformation of the molecule, *Bottom*: Potential Energy.

CVs in spite of our method not containing any knowledge about these CVs. Figure 2 shows that the peptide bonds transition from a trans-isomer to a cis-isomer state at steps 450 and 3,000. We notice the biggest changes in CVs at these steps in Figure 2. We also note that in addition to the change in the peptide bonds, the final conformation differs from the initial in one of the ψ -dihedral angles. Technically, PP-I has ψ -dihedral angles similar to PP-II, but as a result of the inherent noise of MD our target conformation was sampled with a slight rotation here as well. Interestingly, our method successfully learned to sample transition paths terminating in a similar perturbed state. This indicates that our proposed method is resilient to target states not having a minimal-energy configuration.

4.3. Chignolin

Chignolin is a small β -hairpin protein constructed artificially to study protein folding mechanisms (Honda et al., 2004; Seibert et al., 2005). Due to its small size, its folding process is easier to study than larger scale proteins while being similar enough to shed light on this complex process. In contrast to Alanine Dipeptide and Polyproline, there is no agreement on the transition mechanism describing the folding of Chignolin. Both the CVs involved (Satoh et al., 2006; Paissoni & Camilloni, 2021), as well as the sequence of steps (Harada & Kitao, 2011; Satoh et al., 2006; Suenaga et al., 2007; Enemark et al., 2012) describing the folding process have multiple different interpretations. Thus, methods that do not require prior knowledge of CVs are particularly useful to study this protein.

OpenMM and policy network setup: To sample transition paths between the folded and unfolded state of the Chignolin protein, we initialize OpenMM using the same forcefield and VVVR integrator as for Polyproline with the exception that we sample a new force from our policy network every

1.0 fs. We do this 5000 times for each rollout for a total time horizon of 5000 fs. Note that the typical folding time of Chignolin is recorded to be $0.6 \mu\text{s}$ (Lindorff-Larsen et al., 2011). The policy network is trained over 500 rollouts of 16 samples with a learning rate of 1×10^{-4} and standard deviation of 0.05 for the Brownian motion.

Results: In Figure 3, we visualize the transition of Chignolin at 5 different timesteps during the transition path. We observe that to transition the protein from its low energy unfolded state to the folded conformation, the proposed method guides the protein into a region of higher energy. This increase is initially more steep (0→1500) than in the later stages. Additionally, most of the finer-grained folding (2500→4000) occurs with a high potential energy before settling into the lower-energy folded state. We notice that for the restricted folding time we use in our experiments (5000 fs vs $0.6 \mu\text{s}$), the molecule does not end at the final configuration but reaches close to it as shown by the plot on pairwise distance. Furthermore, the learned policy network is able to transition through the high energy transition barrier in this restricted time. We do not encounter this for molecules with a shorter natural transition time (as illustrated by the potential energy of Alanine Dipeptide in Figure 1).

5. Discussion, Limitations, and Future Work

In this work, we proposed a path integral stochastic optimal control method for the problem of sampling rare transition paths for molecular systems that incorporates the Hamiltonian dynamics and equivariance of the system. In passing, we showed an equivalence between the problem of sampling transition paths, stochastic optimal control, and the Schrödinger bridge problem. We empirically tested our method on three different molecular systems of varying sizes and demonstrated that it was able to sample transition

paths on the full geometry of the system without biasing along CVs.

One observed limitation of the proposed method is that for molecules with long natural transition times, we observe the transitions to not converge to the configuration of minimal energy after crossing the high-energy transition barrier. We hypothesize that this is due to the method operating on a reduced time horizon (e.g. 5000 fs instead of 0.6 μ s in the case of Chignolin), or due to the terminal control cost function not requiring the velocity to be zero at the end of the transition. Nevertheless, we note that the method is successful in transitioning the molecules over the high energy barriers as exemplified by the known CVs changing appropriately.

There are many exciting directions for future work. Our experiments showed that the molecules transitioned along with the CVs correctly in spite of not having any information about the CVs. It will be interesting to see if we can infer these CVs from the learned policy and dynamics of the systems. Secondly, our method can have implications for training diffusion models within a fixed time-scale by additionally learning the control policy to transform one distribution into another.

References

- Ahamed, T. I., Borkar, V. S., and Juneja, S. [Adaptive importance sampling technique for Markov chains using stochastic approximation](#). *Operations Research*, 54(3): 489–504, 2006.
- Cameron, R. H. and Martin, W. T. [Transformations of weiner integrals under translations](#). *Annals of Mathematics*, pp. 386–396, 1944.
- Carter, E., Ciccotti, G., Hynes, J. T., and Kapral, R. [Constrained reaction coordinate dynamics for the simulation of rare events](#). *Chemical Physics Letters*, 156(5):472–477, 1989.
- Chen, Y., Georgiou, T. T., and Pavon, M. [On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint](#). *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016.
- Cohen, T. and Welling, M. [Group equivariant convolutional networks](#). In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.
- De Bortoli, V., Thornton, J., Heng, J., and Doucet, A. [Diffusion Schrödinger bridge with applications to score-based generative modeling](#). *Advances in Neural Information Processing Systems*, 34, 2021.
- Dellago, C. and Bolhuis, P. G. [Transition path sampling and other advanced simulation techniques for rare events](#). *Advanced computer simulation approaches for soft matter sciences III*, pp. 167–233, 2009.
- 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., et al. [OpenMM 7: Rapid development of high performance algorithms for molecular dynamics](#). *PLoS computational biology*, 13(7): e1005659, 2017.
- Enemark, S., Kurniawan, N. A., and Rajagopalan, R. [\$\beta\$ -Hairpin forms by rolling up from C-terminal: Topological guidance of early folding dynamics](#). *Scientific Reports*, 2(1):1–6, 2012.
- Essmann, U., Perera, L., Berkowitz, M. L., Darden, T., Lee, H., and Pedersen, L. G. [A smooth particle mesh Ewald method](#). *The Journal of chemical physics*, 103(19):8577–8593, 1995.
- Harada, R. and Kitao, A. [Exploring the folding free energy landscape of a \$\beta\$ -hairpin miniprotein, chignolin, using multiscale free energy landscape calculation method](#). *The Journal of Physical Chemistry B*, 115(27):8806–8812, 2011.
- Head-Gordon, T., Head-Gordon, M., Frisch, M. J., Brooks III, C. L., and Pople, J. A. [Theoretical study of blocked glycine and alanine peptide analogs](#). *Journal of the American chemical society*, 113(16):5989–5997, 1991.
- Honda, S., Yamasaki, K., Sawada, Y., and Morii, H. [10 residue folded peptide designed by segment statistics](#). *Structure*, 12(8):1507–1518, 2004.
- Hooft, F., Pérez de Alba Ortíz, A., and Ensing, B. [Discovering collective variables of molecular transitions via genetic algorithms and neural networks](#). *Journal of chemical theory and computation*, 17(4):2294–2306, 2021.
- Jack, R. L. [Ergodicity and large deviations in physical systems with stochastic dynamics](#). *The European Physical Journal B*, 93(4):1–22, 2020.
- Jarzynski, C. [Nonequilibrium equality for free energy differences](#). *Physical Review Letters*, 78(14):2690, 1997.
- Kappen, H. J. [Path integrals and symmetry breaking for optimal control theory](#). *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.
- Kappen, H. J. [An introduction to stochastic control theory, path integrals and reinforcement learning](#). In *AIP conference proceedings*, volume 887, pp. 149–181. American Institute of Physics, 2007.

- Kappen, H. J. and Ruiz, H. C. Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 162(5):1244–1266, 2016.
- Lindorff-Larsen, K., Piana, S., Palmo, K., Maragakis, P., Klepeis, J. L., Dror, R. O., and Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Structure, Function, and Bioinformatics*, 78(8):1950–1958, 2010.
- Lindorff-Larsen, K., Piana, S., Dror, R. O., and Shaw, D. E. How fast-folding proteins fold. *Science*, 334(6055):517–520, 2011.
- Moradi, M., Babin, V., Roland, C., Darden, T. A., and Sagui, C. Conformations and free energy landscapes of polyproline peptides. *Proceedings of the National Academy of Sciences*, 106(49):20746–20751, 2009.
- Moradi, M., Babin, V., Roland, C., and Sagui, C. A classical molecular dynamics investigation of the free energy and structure of short polyproline conformers. *The Journal of chemical physics*, 133(12):09B614, 2010.
- Paissoni, C. and Camilloni, C. How to determine accurate conformational ensembles by metadynamics metainference: a chignolin study case. *Frontiers in molecular biosciences*, 8:491, 2021.
- Pavon, M., Trigila, G., and Tabak, E. G. The Data-Driven Schrödinger Bridge. *Communications on Pure and Applied Mathematics*, 74(7):1545–1573, 2021.
- Rose, D. C., Mair, J. F., and Garrahan, J. P. A reinforcement learning approach to rare trajectory sampling. *New Journal of Physics*, 23(1):013013, 2021.
- Rosky, P. J. and Karplus, M. Solvation. A molecular dynamics study of a dipeptide in water. *Journal of the American Chemical Society*, 101(8):1913–1937, 1979.
- Satoh, D., Shimizu, K., Nakamura, S., and Terada, T. Folding free-energy landscape of a 10-residue mini-protein, chignolin. *FEBS letters*, 580(14):3422–3426, 2006.
- Satorras, V. G., Hoogeboom, E., and Welling, M. E(n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021.
- Schrödinger, E. *Über die umkehrung der naturgesetze*. Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyter u . . . , 1931.
- Schrödinger, E. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pp. 269–310, 1932.
- Seibert, M. M., Patriksson, A., Hess, B., and Van Der Spoel, D. Reproducible polypeptide folding and structure prediction using molecular dynamics simulations. *Journal of molecular biology*, 354(1):173–183, 2005.
- Shi, C., Luo, S., Xu, M., and Tang, J. Learning gradient fields for molecular conformation generation. In *International Conference on Machine Learning*, pp. 9558–9568. PMLR, 2021.
- Sivak, D. A., Chodera, J. D., and Crooks, G. E. Using nonequilibrium fluctuation theorems to understand and correct errors in equilibrium and nonequilibrium simulations of discrete Langevin dynamics. *Physical Review X*, 3(1):011007, 2013.
- Suenaga, A., Narumi, T., Futatsugi, N., Yanai, R., Ohno, Y., Okimoto, N., and Taiji, M. Folding dynamics of 10-residue β -hairpin peptide chignolin. *Chemistry—An Asian Journal*, 2(5):591–598, 2007.
- Swenson, D. W., Prinz, J.-H., Noe, F., Chodera, J. D., and Bolhuis, P. G. OpenPathSampling: A Python framework for path sampling simulations. 1. Basics. *Journal of chemical theory and computation*, 15(2):813–836, 2018.
- Theodorou, E., Buchli, J., and Schaal, S. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.
- Thijssen, S. and Kappen, H. J. Path integral control and state-dependent feedback. *Phys. Rev. E*, 91:032104, Mar 2015. doi: 10.1103/PhysRevE.91.032104. URL <https://link.aps.org/doi/10.1103/PhysRevE.91.032104>.
- Tobias, D. J. and Brooks III, C. L. Conformational equilibrium in the alanine dipeptide in the gas phase and aqueous solution: A comparison of theoretical results. *The Journal of Physical Chemistry*, 96(9):3864–3870, 1992.
- Todorov, E. Efficient computation of optimal actions. *Proceedings of the national academy of sciences*, 106(28):11478–11483, 2009.
- Torrie, G. M. and Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics*, 23(2):187–199, 1977.
- Vargas, F., Thodoroff, P., Lamacraft, A., and Lawrence, N. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

A. Stochastic Optimal Control

In this section we expand on Section 2.1. Specifically, we expand on the derivation of the Stochastic Optimal Control (SOC) objective in terms of a KL-divergence (Appendix A.1) and the derivation of the iterative gradient descent procedure (Appendix A.2). Note that the derivations presented here are primarily a rephrasing of those given in (Kappen & Ruiz, 2016) using notation similar to the remainder of the paper.

Let us start by restating the objective of Path Integral Stochastic Optimal Control. Given a control \mathbf{u} and the Brownian motion $\boldsymbol{\varepsilon}_t$ with variance ν , Equation (4) defines a trajectory $\mathbf{x}(\tau) = (\mathbf{x}_0, \dots, \mathbf{x}_\tau) \in \mathbb{R}^{T \times d}$. We can define the cost for said trajectory as

$$C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t) = \frac{1}{\lambda} \left(\varphi(\mathbf{x}_\tau) + \sum_{t=0}^{\tau-1} \frac{1}{2} \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \mathbf{u}(\mathbf{x}_t, t) + \mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \boldsymbol{\varepsilon}_t \right) \quad (13)$$

where φ denotes the terminal cost, λ is a constant and \mathbf{R} defines a weighted control cost.

This is a restatement of Equation (5), included here to make future reference easier. We note a number of important observations.

1. Following Equation (4), we observe that the control \mathbf{u} acts *linearly* on the dynamics of the system.
2. The cost of the control itself is *quadratic*, weighted by the matrix \mathbf{R} .
3. Under expectation the final term vanished; $\mathbb{E}_{\boldsymbol{\varepsilon}}[\mathbf{u}(\mathbf{x}_t, t)^T \mathbf{R} \boldsymbol{\varepsilon}_t] = 0$

The first two observations are what classify the current control problem in the family of *Path Integral Stochastic Optimal Control* (Kappen, 2005) and are a requirement to be able to derive the explicit expression for the optimal control policy given in Equation (10). The third observation, while unusual in the context of SOC, is needed to rewrite the SOC objective in terms of the KL-divergence as we will see next. Additionally, if we restate the SOC-objective

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \mathbb{E}_{\tau, \boldsymbol{\varepsilon}_t} [C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t)] \quad (14)$$

we observe that observation 3 shows that the additional cost does not change the optimal control \mathbf{u}^* .

Lastly, we note that the family of Path Integral Stochastic Optimal Control problems assumes that $\lambda = \mathbf{R}\nu$. This assumption is needed both for rewriting the SOC objective as a KL-divergence and to find an explicit expression for the solution.

A.1. SOC objective as a KL-divergence

As noted in the main body of the paper, Equation (4) defines a probability distribution $\pi_{\mathbf{u}}(\mathbf{x}(\tau))$ over trajectories $\mathbf{x}(\tau)$ as

$$\pi_{\mathbf{u}}(\mathbf{x}(\tau)) = \prod_{t=0}^{\tau-1} \mathcal{N}(\mathbf{x}_{t+1} | \boldsymbol{\mu}_t, \Sigma_t) \quad (15)$$

with $\boldsymbol{\mu}_s = \mathbf{x}_s + f(\mathbf{x}_s, s) dt + \mathbf{G}(\mathbf{x}_s, s)(\mathbf{u}(\mathbf{x}_s, s) dt)$ and $\Sigma_s = \mathbf{G}(\mathbf{x}_s, s)^T \nu \mathbf{G}(\mathbf{x}_s, s)$. Similarly, Equation (3) defines a probability distribution $\pi_0(\mathbf{x}(\tau))$, where now $\mathbf{u} = 0$:

$$\pi_0(\mathbf{x}(\tau)) = \prod_{t=0}^{\tau-1} \mathcal{N}(\mathbf{x}_{t+1} | \hat{\boldsymbol{\mu}}_t, \hat{\Sigma}_t) \quad (16)$$

with $\hat{\boldsymbol{\mu}}_s = \boldsymbol{x}_s + f(\boldsymbol{x}_s, s) dt$ and $\hat{\Sigma}_s = \Sigma_s$. These distributions are related through the Girsanov Theorem (Cameron & Martin, 1944), giving:

$$\begin{aligned}
 \pi_{\boldsymbol{u}}(\boldsymbol{x}(\tau)) &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} -\frac{1}{2} \frac{\boldsymbol{u}_t^T \boldsymbol{G}_t^T \boldsymbol{G}_t \boldsymbol{u}_t}{\Sigma_t} + \frac{\boldsymbol{G}_t \boldsymbol{u}_t (f_t + \boldsymbol{x}_t - \boldsymbol{x}_{t-1})}{\Sigma_t}\right) \\
 &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} -\frac{1}{2} \frac{\boldsymbol{u}_t^T \boldsymbol{G}_t^T \boldsymbol{G}_t \boldsymbol{u}_t}{\Sigma_t} + \frac{\boldsymbol{G}_t \boldsymbol{u}_t (\boldsymbol{G}_t (\boldsymbol{u}_t + \boldsymbol{\varepsilon}_t))}{\Sigma_t}\right) \\
 &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} \frac{1}{2} \frac{\boldsymbol{u}_t^T \boldsymbol{G}_t^T \boldsymbol{G}_t \boldsymbol{u}_t}{\Sigma_t} + \frac{\boldsymbol{u}_t^T \boldsymbol{G}_t^T \boldsymbol{G}_t \boldsymbol{\varepsilon}_t}{\Sigma_t}\right) \\
 &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} \frac{1}{2} \boldsymbol{u}_t^T \boldsymbol{G}_t^T \Sigma_t^{-1} \boldsymbol{G}_t \boldsymbol{u}_t + \boldsymbol{u}_t^T \boldsymbol{G}_t^T \Sigma_t^{-1} \boldsymbol{G}_t \boldsymbol{\varepsilon}_t\right) \\
 &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} \frac{1}{2} \boldsymbol{u}_t^T \boldsymbol{\nu}^{-1} \boldsymbol{u}_t + \boldsymbol{u}_t^T \boldsymbol{\nu}^{-1} \boldsymbol{\varepsilon}_t\right) \\
 &= \pi_0(\boldsymbol{x}(\tau)) \exp\left(\frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \boldsymbol{u}_t^T \boldsymbol{R} \boldsymbol{u}_t + \boldsymbol{u}_t^T \boldsymbol{R} \boldsymbol{\varepsilon}_t\right), \tag{17}
 \end{aligned}$$

where we use the assumption $\lambda = \boldsymbol{R}\boldsymbol{\nu}$ in the last step. We use shorthand notation to simplify $\boldsymbol{u}_t = \boldsymbol{u}(\boldsymbol{x}_t, t)$, $\boldsymbol{G}_t = \boldsymbol{G}(\boldsymbol{x}_t, t)$, and $f_t = f(\boldsymbol{x}_t, t)$. From here we can obtain the relation in Equation (8).

Now, as again show in (Kappen & Ruiz, 2016), we can use this relation to rewrite the cost in Equation (13) as

$$C(\boldsymbol{x}(\tau), \boldsymbol{u}, \boldsymbol{\varepsilon}_t) = \frac{1}{\lambda} \varphi(\boldsymbol{x}_\tau) + \log \frac{\pi_{\boldsymbol{u}}(\boldsymbol{x}(\tau))}{\pi_0(\boldsymbol{x}(\tau))}, \tag{18}$$

and thus, the distribution over trajectories under optimal control \boldsymbol{u}^* can now be defined as

$$\begin{aligned}
 \pi_{\boldsymbol{u}^*} &= \arg \min_{\pi_{\boldsymbol{u}}} \mathbb{E}_{\boldsymbol{x}(\tau) \sim \pi_{\boldsymbol{u}}} [C(\boldsymbol{x}(\tau), \boldsymbol{u}, \boldsymbol{\varepsilon}_t)] \\
 &= \arg \min_{\pi_{\boldsymbol{u}}} \mathbb{E}_{\boldsymbol{x}(\tau) \sim \pi_{\boldsymbol{u}}} \left[\frac{1}{\lambda} \varphi(\boldsymbol{x}_\tau) + \log \frac{\pi_{\boldsymbol{u}}(\boldsymbol{x}(\tau))}{\pi_0(\boldsymbol{x}(\tau))} \right] \\
 &= \arg \min_{\pi_{\boldsymbol{u}}} \mathbb{E}_{\boldsymbol{x}(\tau) \sim \pi_{\boldsymbol{u}}} \left[\frac{1}{\lambda} \varphi(\boldsymbol{x}_\tau) \right] + \mathbb{E}_{\boldsymbol{x}(\tau) \sim \pi_{\boldsymbol{u}}} \left[\log \frac{\pi_{\boldsymbol{u}}(\boldsymbol{x}(\tau))}{\pi_0(\boldsymbol{x}(\tau))} \right] \\
 &= \arg \min_{\pi_{\boldsymbol{u}}} \mathbb{E}_{\boldsymbol{x}(\tau) \sim \pi_{\boldsymbol{u}}} \left[\frac{1}{\lambda} \varphi(\boldsymbol{x}_\tau) \right] + \mathbb{D}_{\text{KL}} \left(\pi_{\boldsymbol{u}}(\boldsymbol{x}(\tau)) \parallel \pi_0(\boldsymbol{x}(\tau)) \right) \tag{19}
 \end{aligned}$$

This objective is an approximation of the Schrodinger Bridge formulation in Equation (2) where the constraints on the marginal distributions are replaced by a regularization term in the form of the terminal cost. When the expected terminal cost dominates the KL-divergence the found distribution should be similar.

A.2. Iterative Gradient Descent

As mentioned earlier, the specific control problem we are considering here (linear acting control and weighted quadratic control cost) is known as *Path Integral Stochastic Optimal Control*. Work on this control problem has established that under the additional assumption that $\lambda = \boldsymbol{R}\boldsymbol{\nu}$ there exists an explicit solution describing the optimal control \boldsymbol{u}^* . While there are a number of different papers establishing this result (Kappen, 2005; Theodorou et al., 2010; Kappen, 2007), we note that (Kappen & Ruiz, 2016) is most in line with our work. As such, we refer the interested reader to this work to find the proof for the following statement that defines the distribution over optimal trajectories as a reweighing of the distributions over trajectories under no control:

$$\pi_{\boldsymbol{u}^*} = \frac{1}{\eta} \pi_0(\boldsymbol{x}(\tau)) \exp\left(-\frac{1}{\lambda} \varphi(\boldsymbol{x}_\tau)\right), \tag{20}$$

where $\eta = \mathbb{E}_{\mathbf{x}(\tau) \sim \pi_0} [\exp(-\frac{1}{\lambda} \varphi(\mathbf{x}_\tau))]$ is the normalization constant. Given the previously established relation (Equation (8)) between π_0 and $\pi_{\mathbf{u}}$, we can equivalently express the optimal control \mathbf{u}^* in terms of any control \mathbf{u} using importance sampling

$$\begin{aligned}
 \pi_{\mathbf{u}^*} &= \frac{1}{\eta} \frac{\pi_0(\mathbf{x}(\tau))}{\pi_{\mathbf{u}}(\mathbf{x}(\tau))} \pi_{\mathbf{u}}(\mathbf{x}(\tau)) \exp(-\frac{1}{\lambda} \varphi(\mathbf{x}_\tau)) \\
 &= \frac{1}{\eta} \frac{1}{\exp\left(\frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}_t^T \mathbf{R} \mathbf{u}_t + \mathbf{u}_t^T \mathbf{R} \boldsymbol{\varepsilon}_t\right)} \pi_{\mathbf{u}}(\mathbf{x}(\tau)) \exp(-\frac{1}{\lambda} \varphi(\mathbf{x}_\tau)) \\
 &= \frac{1}{\eta} \pi_{\mathbf{u}}(\mathbf{x}(\tau)) \exp\left(-\frac{1}{\lambda} \varphi(\mathbf{x}_\tau) - \frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}_t^T \mathbf{R} \mathbf{u}_t - \mathbf{u}_t^T \mathbf{R} \boldsymbol{\varepsilon}_t\right) \\
 &= \frac{1}{\eta} \pi_{\mathbf{u}}(\mathbf{x}(\tau)) \exp(-C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t))
 \end{aligned} \tag{21}$$

With an explicit expression for the optimal control policy given, the PICE method aims to find a parameterized distribution $\pi_{\mathbf{u}_\theta}$ that is close to the optimal control in terms of KL-divergence

$$\pi_{\mathbf{u}_\theta^*} = \arg \min_{\pi_{\mathbf{u}_\theta}} \mathbb{D}_{\text{KL}}\left(\pi_{\mathbf{u}^*}(\mathbf{x}(\tau)) \parallel \pi_{\mathbf{u}_\theta}(\mathbf{x}(\tau))\right). \tag{22}$$

Using the explicit expression for the optimal control, the KL-divergence is given as follows:

$$\begin{aligned}
 &\mathbb{D}_{\text{KL}}\left(\pi_{\mathbf{u}^*}(\mathbf{x}(\tau)) \parallel \pi_{\mathbf{u}_\theta}(\mathbf{x}(\tau))\right) \\
 &\propto -\mathbb{E}_{\pi_{\mathbf{u}^*}} [\log \pi_{\mathbf{u}_\theta}] \\
 &= -\mathbb{E}_{\pi_{\mathbf{u}^*}} \left[\log \pi_0(\mathbf{x}(\tau)) \exp\left(\sum_{t=0}^{\tau} -\frac{1}{2} \frac{\mathbf{u}_\theta(t)^T \mathbf{G}_t^T \mathbf{G}_t \mathbf{u}_\theta(t)}{\Sigma_t} + \frac{\mathbf{G}_t \mathbf{u}_\theta(t) (f_t + \mathbf{x}_t - \mathbf{x}_{t-1})}{\Sigma_t}\right) \right] \\
 &\propto -\mathbb{E}_{\pi_{\mathbf{u}^*}} \left[\sum_{t=0}^{\tau} -\frac{1}{2} \frac{\mathbf{u}_\theta(t)^T \mathbf{G}_t^T \mathbf{G}_t \mathbf{u}_\theta(t)}{\Sigma_t} + \frac{\mathbf{G}_t \mathbf{u}_\theta(t) (\mathbf{G}_t (\mathbf{u}^*(t) + \boldsymbol{\varepsilon}_t))}{\Sigma_t} \right] \\
 &= \mathbb{E}_{\pi_{\mathbf{u}^*}} \left[\frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}_\theta(t)^T \mathbf{R} \mathbf{u}_\theta(t) - \mathbf{u}_\theta(t)^T \mathbf{R} \mathbf{u}^*(t) - \mathbf{u}_\theta(t)^T \mathbf{R} \boldsymbol{\varepsilon}_t \right] \\
 &= \frac{1}{\eta} \mathbb{E}_{\pi_{\mathbf{u}^*}} \left[e^{-C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t)} \frac{1}{\lambda} \sum_{t=0}^{\tau} \frac{1}{2} \mathbf{u}_\theta(t)^T \mathbf{R} \mathbf{u}_\theta(t) - \mathbf{u}_\theta(t)^T \mathbf{R} \mathbf{u}(t) - \mathbf{u}_\theta(t)^T \mathbf{R} \boldsymbol{\varepsilon}_t \right]
 \end{aligned} \tag{23}$$

We use shorthand notation to simplify $\mathbf{u}(t) = \mathbf{u}(\mathbf{x}_t, t)$, $\mathbf{G}_t = \mathbf{G}(\mathbf{x}_t, t)$, and $f_t = f(\mathbf{x}_t, t)$. Line 1 we discard the constant term $\mathbb{E}_{\pi_{\mathbf{u}^*}} [\log \pi_{\mathbf{u}^*}]$. Line 2 we make use of the established relation between $\pi_{\mathbf{u}}$ and π_0 for any control \mathbf{u} . Line 3 we discard the constant term $\mathbb{E}_{\pi_{\mathbf{u}^*}} [\log \pi_0]$ and note that the expectation is over trajectories sampled from $\pi_{\mathbf{u}^*}$. Line 4 we rewrite using the assumption that $\lambda = \mathbf{R}^{-1}$. Line 5 we use Equation (21) to rewrite the distribution over an arbitrary control \mathbf{u} using.

We can minimize this explicit expression using Gradient Descent, to do this, we derive the gradient of the KL-divergence

$$\begin{aligned}
 &\frac{\partial \mathbb{D}_{\text{KL}}\left(\pi_{\mathbf{u}^*}(\mathbf{x}(\tau)) \parallel \pi_{\mathbf{u}_\theta}(\mathbf{x}(\tau))\right)}{\partial \theta} \\
 &= \frac{1}{\eta} \mathbb{E}_{\pi_{\mathbf{u}^*}} \left[e^{-C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t)} \frac{1}{\lambda} \sum_{t=0}^{\tau} (\mathbf{R} \mathbf{u}_\theta(t) - \mathbf{R} \mathbf{u}(t) - \mathbf{R} \boldsymbol{\varepsilon}_t) \frac{\partial \mathbf{u}_\theta(t)}{\partial \theta} \right].
 \end{aligned} \tag{24}$$

Finally, we note that the expectation is over trajectories of any distribution $\pi_{\mathbf{u}}$, and as such, this distribution can also be

chosen to be equal to the parameterized distribution $\pi_{\mathbf{u}} = \pi_{\mathbf{u}_\theta}$. This gives us the final gradient

$$\begin{aligned} & \frac{\partial \mathbb{D}_{\text{KL}}\left(\pi_{\mathbf{u}^*}(\mathbf{x}(\tau)) \parallel \pi_{\mathbf{u}_\theta}(\mathbf{x}(\tau))\right)}{\partial \theta} \\ &= \frac{1}{\eta} \mathbb{E}_{\pi_{\mathbf{u}_\theta}} \left[e^{-C(\mathbf{x}(\tau), \mathbf{u}, \boldsymbol{\varepsilon}_t)} \frac{1}{\lambda} \sum_{t=0}^{\tau} (\mathbf{R}\mathbf{u}_\theta(t) - \mathbf{R}\mathbf{u}_\theta(t) - \mathbf{R}\boldsymbol{\varepsilon}_t) \frac{\partial \mathbf{u}_\theta(t)}{\partial \theta} \right] \end{aligned} \quad (25)$$

$$= -\frac{1}{\eta} \mathbb{E}_{\pi_{\mathbf{u}_\theta}} \left[e^{-C(\mathbf{x}(\tau), \mathbf{u}_\theta, \boldsymbol{\varepsilon}_t)} \frac{1}{\lambda} \sum_{t=0}^{\tau} \mathbf{R}\boldsymbol{\varepsilon}_t \frac{\partial \mathbf{u}_\theta(t)}{\partial \theta} \right]. \quad (26)$$

B. Extension Experimental section

B.1. Alanine Dipeptide

B.1.1. DISCUSSION BASELINES AND EVALUATION METRICS

Metrics Three different metrics are used for the comparison covering multiple desiderata for the sampled transition trajectories. For each metric we report the score over 1000 trajectories with the exception of the *Molecular Dynamics without fixed timescale* baseline which is only ran until 10 trajectories are successfully generated.

Expected Pairwise Distance (EPD) The EPD measures the similarity between the final conformation in the trajectory and the target conformation taking into account the full 3D geometry of the molecule. Note that the expected pairwise distance for uncontrolled MD with the target as the starting conformation has a EPD of 2.25×10^{-3} . All trajectories with an EPD of less than this can thus be considered to transition the molecule within one standard deviation of the target distribution.

Target Hit Percentage (THP): The second metric under which we evaluate the proposed Transition Path Sampler measures the similarity of the final and target conformation in terms of the collective variables. The THP measures the percentage of generated trajectories/paths that reach the target state. As such, higher hit percentages are preferred. We determine a trajectory to have hit the target in CV space when ϕ and ψ are both within 0.75 of the target.

Energy Transition Point (ETP): The final metric looks at the potential energy of the transition point—the conformation in the trajectory with the highest potential energy. This directly evaluates the capability of the method to find the transition path that crosses the boundary at the lowest saddle point.

Baselines We compare the proposed Transition Path Sampling method with extended Molecular Dynamics simulation using different time-scales and temperature points. As discussed earlier, there are currently no other methods available for Transition Path Sampling using the full 3D geometry of the molecules.

Molecular Dynamics with fixed timescale: This set of baselines is limited to the same timescale as the proposed Transition Path Sampler, 500 femtoseconds, but uses varying temperatures. With higher temperatures we should have a higher probability of crossing the barrier and hitting the target configuration.

Molecule Dynamics without fixed timescales: In contrast to the other set of baselines, the MD simulation for this set is not limited to 500 femtoseconds, but is instead ran until the target conformation is reached. We consider a trajectory to have reached its target if the following two conditions have been met: 1) the current conformation classifies as having hit the target under the conditions of the metric described above and 2) the current conformation is within one standard deviation of the target distributions mean.

By running the MD simulations until the target is reached we aim to gain intuition into the speed-up that it achieved by the fixed timescale of the proposed Transition Path Sampler.

B.1.2. ADDITIONAL RESULTS: VISUALIZATION FORCE PREDICTION

We observe that the force predicting policy has learned a different trajectory than the energy predicting model presented in the main body of the paper. While different, both of the trajectories pass the high energy barrier in a locally low point. Previous work on finding transition path has also observed that multiple viable paths can be found for Alanine Dipeptide (Hooft et al., 2021).

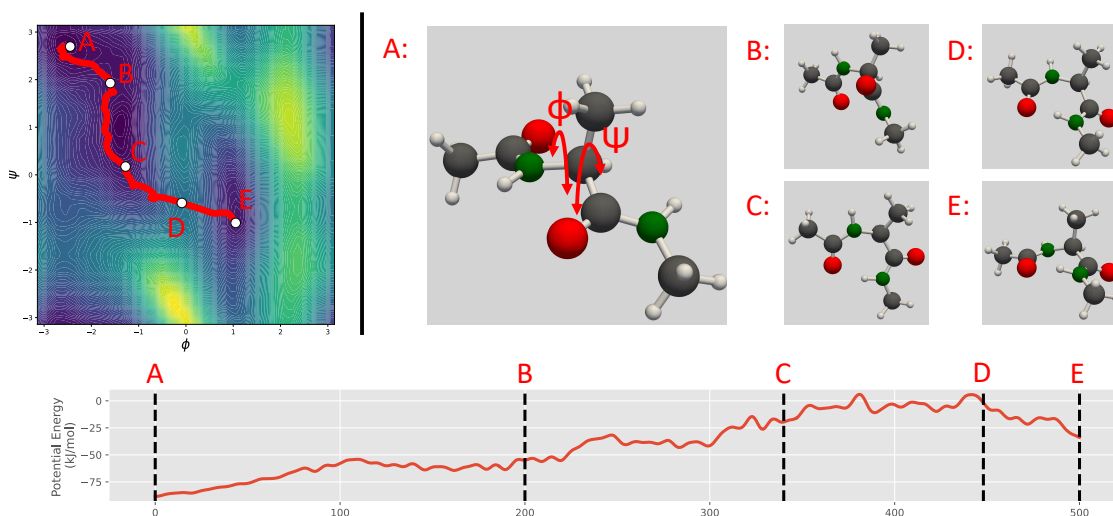


Figure 4. Visualization of a trajectory sampled with the proposed force prediction method. *Left:* The sampled trajectory projected on the free energy landscape of Alanine Dipeptide as a function of two CVs *Right:* Conformations along the sampled trajectory: A) starting conformation showing the CV dihedral angles, B-D) intermediate conformations with D being the highest energy point on the trajectory, and E) final conformation, which closely aligns with the target conformation. *Bottom:* Potential energy during transition. Letters represent the same configurations in the transition.