
Entangled Schrödinger Bridge Matching

Sophia Tang¹, Yinuo Zhang², Pranam Chatterjee^{1,3,†}

¹Department of Computer and Information Science, University of Pennsylvania

²Center of Computational Biology, Duke-NUS Medical School

³Department of Bioengineering, University of Pennsylvania

†Corresponding author: pranam@seas.upenn.edu

Abstract

Simulating trajectories of multi-particle systems on complex energy landscapes is a central task in molecular dynamics (MD) and drug discovery, but remains challenging at scale due to computationally expensive and long simulations. Previous approaches leverage techniques such as flow or Schrödinger bridge matching to implicitly learn joint trajectories through data snapshots. However, many systems, including biomolecular systems and heterogeneous cell populations, undergo *dynamic* interactions that evolve over their trajectory and cannot be captured through static snapshots. To close this gap, we introduce **Entangled Schrödinger Bridge Matching (EntangledSBM)**, a framework that learns the first- and second-order stochastic dynamics of interacting, multi-particle systems where the direction and magnitude of each particle’s path depend dynamically on the paths of the other particles. We define the Entangled Schrödinger Bridge (EntangledSB) problem as solving a coupled system of bias forces that *entangle* particle velocities. We show that our framework accurately simulates heterogeneous cell populations under perturbations and rare transitions in high-dimensional biomolecular systems.

1 Introduction

Schrödinger bridge matching (SBM) has enabled significant progress in problems ranging from molecular dynamics (MD) simulation [Holdijk et al., 2023] to cell state modeling [Tong et al., 2024] by effectively learning dynamic trajectories between initial and target distributions. Diverse SBM frameworks have been developed for learning trajectories with minimized state-cost [Liu et al., 2023] to branching trajectories that map from a single initial distribution [Tang et al., 2025a]. Most of these frameworks assume that the particle acts independently or undergoes interactions that can be implicitly captured through training on static snapshots of the system. Existing approaches for modeling interacting particle systems rely on the mean-field assumption, where all particles are exchangeable and interact only through averaged effects [Liu et al., 2022, Zhang et al., 2025a]. While this assumption may hold for homogeneous particle dynamics, it fails to describe the heterogeneous dynamics observed in more complex domains, such as conformationally dynamic proteins [Vincoff et al., 2025, Chen et al., 2023], heterogeneous cell populations [Smela et al., 2023a], or interacting token sequences [Geshkovski et al., 2023]. In these settings, the motion of each particle depends not only on its own position but also on the evolving configuration of surrounding particles.

To accurately model such dependencies, the joint evolution of an n -particle system must be described by a coupled probability distribution that transports the initial distribution $\pi_0(\mathbf{X}_0)$ to the target distribution $\pi_B(\mathbf{X}_T)$ in the phase space of both positions and velocities $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t) \in \mathcal{X}$. However, modeling these interactions requires learning to simulate second-order dynamics, where interactions between velocity fields evolve over time, which remains largely unexplored. To address this gap,

we introduce **Entangled Schrödinger Bridge Matching (EntangledSBM)**, a novel framework that learns interacting second-order dynamics of n -particle systems, capturing dependencies on both the static position and dynamic velocities of the system at each time step.

Contributions Our contributions can be summarized as follows: **(1)** We formulate the Entangled Schrödinger Bridge (EntangledSB) problem, which considers the optimal path between distributions following second-order Langevin dynamics with an entangled bias force (Sec 3). **(2)** We introduce **EntangledSBM**, a novel parameterization of the bias force that can be conditioned, *at inference-time*, on a target distribution or terminal state, enabling generalizable sampling of diverse target distributions (Sec 4). **(3)** We evaluate EntangledSBM on mapping cell cluster dynamics under drug perturbations (Sec 5.1) and transition path sampling of high-dimensional molecular systems (Sec 5.2).

Related Works We provide a comprehensive discussion on related works in App A.

2 Preliminaries

Langevin Dynamics A time-evolving n -particle molecular system can be represented as $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$, where $\mathbf{R}_t = \{\mathbf{r}_t^i \in \mathbb{R}^d\}_{i=1}^n$ denotes the set of coordinates $\mathbf{V}_t = \{\mathbf{v}_t^i \in \mathbb{R}^d\}_{i=1}^n$ denotes the set of velocities of each particle i . The evolution of the positions and velocities of the system given a potential energy function $U : \mathcal{X} \rightarrow \mathbb{R}$ can be modelled with *Langevin dynamics*, which effectively captures the motion of particles under **conservative forces** between particles and **stochastic collisions** with the surrounding environment [Bussi and Parrinello, 2007, Yang et al., 2006] using a pair of stochastic differential equations (SDEs) defined as

$$d\mathbf{r}_t^i = \mathbf{v}_t^i dt, \quad d\mathbf{v}_t^i = \frac{-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t)}{m_i} dt - \gamma \mathbf{v}_t^i dt + \sqrt{\frac{2\gamma k_B \tau}{m_i}} d\mathbf{W}_t^i \quad (1)$$

where m_i is the mass of each particle, γ is the friction coefficient, k_B is the Boltzmann constant, τ is the temperature, and $d\mathbf{W}_t$ is standard Brownian motion. In molecular dynamics (MD) simulations of biomolecules, the particles undergo *underdamped* Langevin dynamics with small γ , where *inertia* is not negligible.

Many biological systems, including cell clusters for cell-state trajectory simulation, can be modeled with *overdamped* Langevin dynamics, where inertia is negligible but the system still undergoes external forces that define its motion. This can be represented with the first-order SDE given by

$$d\mathbf{r}_t^i = \frac{-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t)}{\gamma} dt + \sqrt{\frac{2k_B \tau}{\gamma}} d\mathbf{W}_t^i \quad (2)$$

Schrödinger Bridge Matching Tasks such as simulating cell state trajectories and transition path sampling aim to simulate the Langevin dynamics from an initial distribution to a desired target state or distribution. Given an initial distribution $\pi_{\mathcal{A}}$ and a target distribution $\pi_{\mathcal{B}}$, we define the distribution of paths $\mathbf{X}_{0:T} := (\mathbf{X}_t)_{t \in [0, T]}$ satisfying the endpoint constraints $\mathbf{R}_0 \sim \pi_{\mathcal{A}}$ and $\mathbf{R}_T \sim \pi_{\mathcal{B}}$ as the *optimal bridge distribution* $\mathbb{P}^*(\mathbf{X}_{0:T})$ defined as

$$\mathbb{P}^*(\mathbf{X}_{0:T}) = \frac{1}{Z} \mathbb{P}^0(\mathbf{X}_{0:T}) \pi_{\mathcal{B}}(\mathbf{R}_T), \quad Z = \mathbb{E}_{\mathbf{X}_{0:T} \sim \mathbb{P}^0} [\pi_{\mathcal{B}}(\mathbf{R}_T)] \quad (3)$$

where $\mathbb{P}^0(\mathbf{X}_{0:T})$ is the base path distribution generated from the SDEs (1) or (2) and Z is the normalizing constant. Schrödinger Bridge Matching (SBM) aims to parameterize a **control** or **bias force** \mathbf{b}_θ that tilts the path distribution $\mathbb{P}^{\mathbf{b}_\theta}(\mathbf{X}_{0:T})$ minimizes the KL-divergence from the bridge path distribution \mathbb{P}^* given by

$$\mathbf{b}^* = \arg \min_{\mathbf{b}_\theta} D_{\text{KL}} \left(\mathbb{P}^{\mathbf{b}_\theta} \parallel \mathbb{P}^* \right) \quad \text{s.t.} \quad \begin{cases} d\mathbf{r}_t^i = \mathbf{v}_t^i dt \\ d\mathbf{v}_t^i = \frac{-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t) + \mathbf{b}_\theta(\mathbf{R}_t)}{m_i} dt - \gamma \mathbf{v}_t^i dt + \sqrt{\frac{2\gamma k_B \tau}{m_i}} d\mathbf{W}_t^i \end{cases} \quad (4)$$

In the case of transition path sampling (TPS) where there is a single initial state $\mathbf{R}_{\mathcal{A}}$ and target state $\mathbf{R}_{\mathcal{B}}$, we define the target distribution as the relaxed indicator function $\pi_{\mathcal{B}}(\mathbf{R}_T) = \mathbf{1}_{\mathcal{B}}(\mathbf{R}_T)$ centered around $\mathbf{R}_{\mathcal{B}}$.

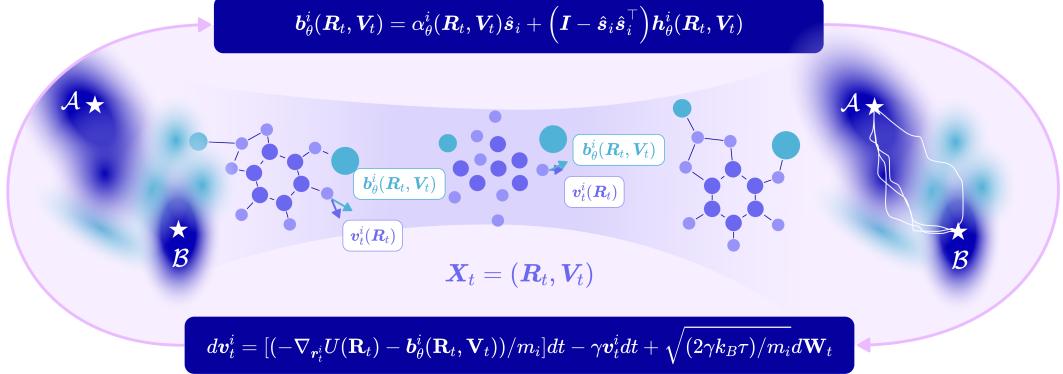


Figure 1: **Entangled Schrödinger Bridge Matching.** We consider the problem of simulating interacting multi-particle systems, where each particle’s velocity depends dynamically on the velocities of the other particles in the system, and introduce **EntangledSBM**, a framework that parameterizes an entangled bias force capturing the dynamic interactions between particles.

3 Learning Interacting Multi-Particle Dynamics

The **key challenge** with simulating the dynamics of multi-particle systems lies in the question: *how can we simulate dynamic trajectories from static snapshot data?* The emergence of flow and Schrödinger bridge matching frameworks have effectively approached this problem by defining a **parameterized velocity field** that learns feasible trajectories from data snapshots. However, current strategies remain limited in their ability to model *interacting* multi-particle systems where the velocities carry inherent dependencies that cannot be captured via static snapshots of the system. Prior attempts to capture interactions between particles rely on the *mean-field assumption*, where each particle acts as an average of its surrounding particles, which **does not hold for heterogeneous biomolecular systems**.

To address this challenge, we propose a framework that introduces an additional degree of freedom through a **bias force** that implicitly captures the dependencies between *both* the static positions and dynamic velocities of each particle in the system to control the joint trajectories. We leverage a Transformer architecture and treat each particle as an individual token with features that attend to the features of all other tokens. Crucially, our approach requires no handcrafted features, making it scalable to high-complexity systems.

In Sec 3.1, we formalize this problem as the **Entangled Schrödinger Bridge (EntangledSB)** problem, which aims to determine the optimal set of trajectories to a target state while capturing the dynamic interactions between particles. We then provide a tractable approach to solve the EntangledSB problem with stochastic optimal control (SOC) theory in Sec 3.2. We illustrate the high-level framework of learning **entangled** bias forces in Alg 1 and detail our specific implementation and parameterization in Sec 4.

3.1 Entangled Schrödinger Bridge Problem

Here, we formalize the **Entangled Schrödinger Bridge (EntangledSB) problem**, which aims to find a set of optimal bias forces for each particle in the system that depends *dynamically* on the positions and velocities across the entire multi-particle system to guide the Schrödinger bridge trajectory to a target distribution. Specifically, we consider an n -particle system $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$ with $\mathbf{R}_t = \{\mathbf{r}_t^i \in \mathbb{R}^d\}_{i=1}^n$ and $\mathbf{V}_t = \{\mathbf{v}_t^i \in \mathbb{R}^d\}_{i=1}^n$, where each particle evolves over the time horizon $t \in [0, T]$ via the **bias-controlled** SDE given by

$$d\mathbf{r}_t^i = \frac{1}{\gamma} \left(-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t) + \mathbf{b}^i(\mathbf{R}_t, \mathbf{V}_t) \right) dt + \sqrt{\frac{2k_B\tau}{\gamma}} d\mathbf{W}_t^i \quad (5)$$

The velocities follow the potential energy landscape defined over the joint coordinates of the system $U(\mathbf{R}_t)$ with an *entangled* bias force $\mathbf{b}^i(\mathbf{R}_t, \mathbf{V}_t)$ that depends dynamically on the position and velocity of *all* particles in the system. We can further extend this framework to **second-order** Langevin

Algorithm 1 Framework for Entangled Schrödinger Bridge Matching

```

1: Input: Parameterized velocity network  $\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t) : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ , initial and target
   joint distributions of the system  $\pi_0, \pi_T$ , energy function  $U(\mathbf{R}_t) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ , objective function
   in path space  $\mathcal{L}(\mathbb{P}^*, \mathbb{P}^{b_\theta})$  where  $\mathbb{P}^* = \arg \min_{\mathbb{P}^{b_\theta}} \mathcal{L}(\mathbb{P}^*, \mathbb{P}^{b_\theta})$ 
2: while not converged do
3:   Sample initial positions  $\mathbf{R}_0 \sim \pi_0$ 
4:   for  $t = 0$  to  $T$  do
5:     Evaluate unconditional velocities  $\mathbf{V}_t(\mathbf{R}_t)$ 
6:     Compute entangled bias force  $\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)$ 
7:     Simulate step over a discrete time step  $\Delta t$  to get updated coordinates  $\mathbf{R}_t$ 
8:   end for
9:   Compute reward on terminal states  $r(\mathbf{X}_T)$ 
10:  Evaluate path objective  $\mathcal{L}(\mathbb{P}^*, \mathbb{P}^{b_\theta})$  on simulated paths
11:  Update  $\theta$  using the gradient  $\nabla_\theta \mathcal{L}(\mathbb{P}^*, \mathbb{P}^{b_\theta})$ 
12: end while

```

dynamics via the pair of SDEs

$$d\mathbf{r}_t^i = \mathbf{v}_t^i dt, \quad d\mathbf{v}_t^i = \frac{-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t) + \mathbf{b}^i(\mathbf{R}_t, \mathbf{V}_t)}{m_i} dt - \gamma \mathbf{v}_t^i dt + \sqrt{\frac{2\gamma k_B \tau}{m_i}} d\mathbf{W}_t^i \quad (6)$$

Formally, we seek the optimal $\mathbf{b}^*(\mathbf{R}_t, \mathbf{V}_t) := \{\mathbf{b}^i(\mathbf{R}_t, \mathbf{V}_t)\}_{i=1}^n$ that solves the EntangledSB problem defined below.

Definition 3.1 (Entangled Schrödinger Bridge Problem). *Given an initial distribution π_A , a target distribution π_B , the EntangledSB problem aims to determine the set of optimal bias forces $\mathbf{b}^* := \{\mathbf{b}^i(\mathbf{R}_t, \mathbf{V}_t)\}_{i=1}^n$ for each particle i in the system that solves the optimization problem*

$$\mathbf{b}^*(\mathbf{R}_t, \mathbf{V}_t) = \arg \min_{\mathbf{b}_\theta} D_{KL} \left(\mathbb{P}^{b_\theta} \parallel \mathbb{P}^* \right) \quad \text{s.t.} \quad \begin{cases} \mathbb{P}^* = \frac{1}{Z} \mathbb{P}^0 \pi_B(\mathbf{X}_T) \\ \mathbb{P}_0^* = \pi_A(\mathbf{X}_0) \end{cases} \quad (7)$$

where \mathbb{P}^0 is the base path measure with joint dynamics that follow the SDEs in (6).

3.2 Solving EntangledSB with Stochastic Optimal Control

To solve the EntangledSB problem, we leverage stochastic optimal control (SOC) theory where we aim to find the set of optimal control drifts $\mathbf{b}^*(\mathbf{R}_t, \mathbf{V}_t)$ given a target distribution π_B .

Proposition 3.1 (Solving EntangledSB with Stochastic Optimal Control). *We can solve the EntangledSB problem with the stochastic optimal control (SOC) objective given by*

$$\mathbf{b}^* = \arg \min_{\mathbf{b}_\theta} \mathbb{E}_{\mathbf{X}_{0:T} \sim \mathbb{P}^{b_\theta}} \left[\int_0^T \frac{1}{2} \|\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)\|^2 dt - r(\mathbf{X}_T) \right] \quad \text{s.t.} \quad (6) \quad (8)$$

where $r(\mathbf{X}_T) := \log \pi_B(\mathbf{R}_T)$ is the terminal reward that measures the log-probability under the target distribution.

The proof is provided in App C.1. We highlight that the SOC problem is optimized over *unconstrained* trajectories that need not map to explicit samples from π_B but are iteratively refined to generate trajectories with a terminal distribution that matches π_B . In Sec 4.2, we introduce an importance-weighted cross-entropy objective that efficiently solves the SOC problem with theoretical guarantees.

4 Entangled Schrödinger Bridge Matching

In this section, we introduce **Entangled Schrödinger Bridge Matching** (EntangledSBM), a novel framework for learning to simulate trajectories of n -particle systems that with a **bias force that**

dynamically depends on the positions and velocities of the other particles in the system and can generalize to *unseen* target distributions without further training. Our unique parameterization ensures a non-increasing distance toward the target distribution without sacrificing expressivity with hard constraints (Sec 4.1). We introduce a **weighted cross-entropy** objective that enables efficient *off-policy learning* with a replay buffer of simulated trajectories (Sec 4.2).

4.1 Parameterizing the Entangled Bias Force

Bias Force Parameterization To ensure that the bias force does not increase the distance from the target distribution π_B , we ensure that the projection of the predicted bias force onto the direction of the gradient of the target distribution is positive for all particles i in the system

$$\langle b_\theta^i(\mathbf{R}_t, \mathbf{V}_t), \nabla_{\mathbf{r}_t^i} \log \pi_B \rangle \geq 0 \quad (9)$$

Let $\mathbf{s}_i = \nabla_{\mathbf{r}_t^i} \log \pi_B \in \mathbb{R}^d$ denote the direction towards the target state. To ensure that the bias force is within the cone surrounding \mathbf{s}_i , we use the following parameterization

$$b_\theta^i(\mathbf{R}_t, \mathbf{V}_t) = \underbrace{\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \hat{\mathbf{s}}_i}_{\text{parallel component}} + \underbrace{\left(\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top \right) \mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)}_{\text{orthogonal component}}, \quad \hat{\mathbf{s}}_i = \frac{\mathbf{s}_i}{\|\mathbf{s}_i\|} \quad (10)$$

where $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) := \text{softplus}(\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t)) \geq 0$ is a scaling factor for the unit vector $\hat{\mathbf{s}}_i$ and $\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \in \mathbb{R}^d$ is the per-atom correction vector that is projected onto the plane orthogonal to $\hat{\mathbf{s}}_i$, both parameterized with neural networks θ . Since the orthogonal component does not affect the dot product with $\hat{\mathbf{s}}_i$, the non-negativity constraint $\langle b_\theta^i(\mathbf{R}_t, \mathbf{V}_t), \nabla_{\mathbf{r}_t^i} \log \pi_B \rangle \geq 0$ in (9) is guaranteed by the first term. Intuitively, the orthogonal component enables the bias force to have greater flexibility in moving sideways (e.g., to avoid infeasible regions or add rotational/collective effects) while ensuring that the distance from some target state remains non-increasing.

Proposition 4.1 (Non-Increasing Distance from Target Distribution). *For small enough Δt , the distance from some target state $\mathbf{R}_B \in \pi_B$ state after an update with the bias force $\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ defined in (10) is non-increasing, such that*

$$\exists \mathbf{R}_B \in \pi_B \text{ s.t. } \|\mathbf{R}_{t+\Delta t} - \mathbf{R}_B\| \leq \|\mathbf{R}_t - \mathbf{R}_B\| \quad (11)$$

where $\mathbf{R}_{t+\Delta t} = \mathbf{R}_t + (\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)/m_i)\Delta t$.

The proof is provided in App C.2. In contrast to previous works that constrain the bias force to point strictly in the direction of a fixed target position, our approach allows greater flexibility in the direction of the orthogonal component.

Model Architecture To integrate dependencies on the positions and velocities across n particles, we leverage a Transformer-based architecture where each particle has input features $\text{Cat}(\mathbf{r}_t^i, \mathbf{v}_t^i) \in \mathbb{R}^{2d}$ and the full system features are $\text{Cat}(\mathbf{R}_t, \mathbf{V}_t) \in \mathbb{R}^{n \times 2d}$. We further input the direction $\mathbf{R}_B - \mathbf{R}_t$ and distance $\|\mathbf{R}_B - \mathbf{R}_t\|$ from the target distribution, which enables generalization to unseen target distributions at inference by learning the dependence of the bias force on the target direction. The Transformer encoder enables efficient and expressive propagation of feature information across all pairs of particles in the system to generate context-aware embeddings for bias force parameterization. For MD systems, we ensure invariance of the coordinate frame using the Kabsch algorithm [Kabsch, 1976], which aligns the position \mathbf{R}_t with the target position \mathbf{R}_B before input into the model. Further details are provided in App D and E.

4.2 Off-Policy Learning with Weighted Cross-Entropy

Log-Variance Objective To train the bias force to match the optimal b^* , we can adapt the *log-variance* (LV) divergence [Seong et al., 2025, Nüsken and Richter, 2021] defined as

$$\mathcal{L}_{\text{LV}}(\mathbb{P}^*, \mathbb{P}^{b_\theta}) = \text{Var}_{\mathbb{P}^*} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] = \mathbb{E}_{\mathbb{P}^*} \left[\left(\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} - \mathbb{E}_{\mathbb{P}^*} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] \right)^2 \right] \quad (12)$$

where \mathbf{v} is an arbitrary control that enables *off-policy learning* from trajectories that need not be generated by the current bias force \mathbf{b}_θ . While the optimal solution to the LV objective is exactly the optimal bias force \mathbf{b}^* , the non-convex nature of the LV objective with respect to the path measure \mathbb{P}^{b_θ} inhibits convergence. We provide more details in App B.

Cross-Entropy Objective To achieve a more theoretically-favorable optimization problem, we propose a cross-entropy objective that is *convex* with respect to the biased path measure \mathbb{P}^{b_θ} given by

$$\mathcal{L}_{CE}(\mathbb{P}^*, \mathbb{P}^{b_\theta}) = \mathbb{E}_{\mathbb{P}^*} \left[-\log \mathbb{P}^{b_\theta} \right] := D_{KL}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) = \mathbb{E}_{\mathbb{P}^*} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] = \mathbb{E}_{\mathbb{P}^v} \left[\frac{d\mathbb{P}^*}{d\mathbb{P}^v} \log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] \quad (13)$$

where \mathbb{P}^* is the target bridge measure and \mathbb{P}^v is the path measure generated with an arbitrary control \mathbf{v} . To avoid taking the expectation with respect to the unknown path measure \mathbb{P}^* , we define an importance weight $w^*(\mathbf{X}_{0:T}) := \frac{d\mathbb{P}^*}{d\mathbb{P}^v}(\mathbf{X}_{0:T})$ independent of \mathbf{b}_θ that enables optimization using trajectories generated from an arbitrary control \mathbf{v} . We note that similar CE objectives have been adopted in earlier work for different applications [Kappen and Ruiz, 2016, Zhu et al., 2025, Tang et al., 2025b], which we discuss in App A.

Proposition 4.2 (Convexity and Uniqueness of Cross-Entropy Objective). *The cross-entropy objective \mathcal{L}_{CE} is convex in \mathbb{P}^{b_θ} and there exists a unique minimizer \mathbf{b}^* that is the solution to the EntangledSOC problem in Proposition 3.1.*

The proof is given in App C.3. To amortize the cost of simulation and reinforce high-reward trajectories, we define the arbitrary control as the biased path measure from the previous iteration $\mathbf{v} := \bar{\mathbf{b}} = \text{stopgrad}(\mathbf{b}_\theta)$, which allows us to reuse the trajectories over multiple training steps by maintaining a replay buffer \mathcal{R} that contains the trajectories $\mathbf{X}_{0:T}$ and their importance weights.

Proposition 4.3 (Equivalence of Variational and Path Integral Objectives). *The cross-entropy objective can be expressed in path-integral form as*

$$\mathcal{L}_{CE}(\theta) = \mathbb{E}_{\mathbb{P}^v} [w^*(\mathbf{X}_{0:T}) \mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:T})] \quad (14)$$

$$\begin{cases} w^*(\mathbf{X}_{0:T}) = \frac{d\mathbb{P}^*}{d\mathbb{P}^v}(\mathbf{X}_{0:T}) = \frac{e^{r(\mathbf{X}_T)}}{Z} \frac{d\mathbb{P}^0}{d\mathbb{P}^v}(\mathbf{X}_{0:T}) \\ \mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:T}) = \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)\|^2 - \int_0^T (\mathbf{b}_\theta^\top \mathbf{v})(\mathbf{R}_t, \mathbf{V}_t) dt - \int_0^T \mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)^\top d\mathbf{W}_t \end{cases} \quad (15)$$

where we define the reference measure $\mathbf{v} = \bar{\mathbf{b}} := \text{stopgrad}(\mathbf{b}_\theta)$ is the off-policy control drift from the previous iteration.

The proof is given in App C.4. Since the normalizing constant Z is intractable in practice, we compute the importance weight as $w^*(\mathbf{X}_{0:T}) := \text{softmax}_{\mathcal{B}}(r(\mathbf{X}_T) + \log p^0(\mathbf{X}_{0:T}) - \log p^{\bar{\mathbf{b}}}(\mathbf{X}_{0:T}))$, which is a batch estimate of $\frac{d\mathbb{P}^*}{d\mathbb{P}^{\bar{\mathbf{b}}}}$. When a batch contains only a single trajectory (i.e., when the system is large), we additionally store the value of $\log r(\mathbf{X}_T)$ sample from the importance weighted distribution over the replay buffer as $\mathbf{X}_{0:T} \sim \text{Cat}(\text{softmax}_{\mathcal{R}}(w^*(\mathbf{X}_{0:T})))$.

Discrete Time Objective Since we want to aim to train on off-policy trajectories from previous iterations to dynamically update the learned bias force given the velocities of the remaining particles, we require storing the simulated trajectories in a discretized form $\mathbf{X}_{0:K} := (\mathbf{X}_k)_{k \in \{0, \dots, K\}}$ with step size Δt . To compute the term $\mathcal{F}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}$ in the CE loss, we discretize it as

$$\hat{\mathcal{F}}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}(\mathbf{X}_{0:K}) = \frac{1}{2} \sum_{k=0}^{K-1} \|\mathbf{b}_\theta(\mathbf{X}_k)\|^2 \Delta t - \sum_{k=0}^{K-1} (\mathbf{b}_\theta \cdot \bar{\mathbf{b}})(\mathbf{X}_k) \Delta t - \sum_{k=0}^{K-1} \mathbf{b}_\theta(\mathbf{X}_k) \cdot \Delta \mathbf{W}_k \quad (16)$$

where $\Delta \mathbf{W}_k = \Sigma^{-1}[\mathbf{R}_{k+1} - \mathbf{R}_k - (f(\mathbf{X}_k) + \Sigma \bar{\mathbf{b}}(\mathbf{R}_k, \mathbf{V}_k)) \Delta t]$. Now, we can establish a simplified version of the cross-entropy objective in Prop 4.3.

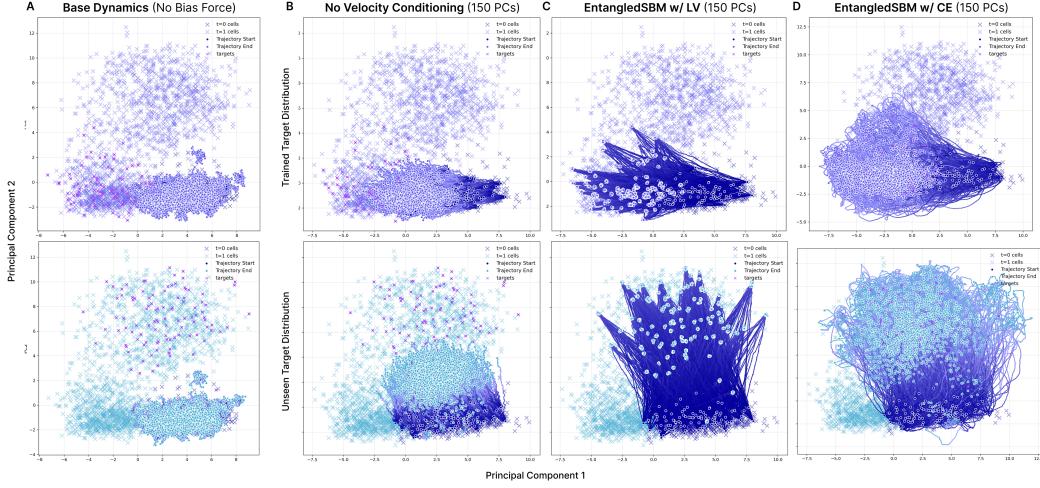


Figure 2: **Simulated cell cluster dynamics with EntangledSBM under Clonidine perturbation with 150 PCs.** Nearest neighbour cell clusters with $n = 16$ cells are simulated over 100 time steps to perturbed cells seen during training (**Top**) and an unseen perturbed population (**Bottom**). The gradient indicates the evolution of timesteps from the initial time $t = 0$ (navy) to the final time $t = T$ (purple or turquoise). (**A**) Trajectories under base dynamics with no bias force. (**B**) Trajectories simulated with base and bias forces with (**B**) no velocity conditioning, (**C**) log-variance (LV) objective, and (**D**) cross-entropy (CE) objective.

Proposition 4.4 (Discretized Cross-Entropy). *Given the discretized $\hat{\mathcal{F}}_{b_\theta, \bar{b}}(\mathbf{X}_{0:K})$ in (16), we can derive a simplified loss function as*

$$\hat{\mathcal{L}}_{CE}(\theta) = \mathbb{E}_{\mathbf{X}_{0:K} \sim \mathbb{P}^b} \left[\underbrace{\frac{d\mathbb{P}^*}{d\mathbb{P}^b}(\mathbf{X}_{0:K})}_{w^*(\mathbf{X}_{0:K})} \log \left(\frac{p^0(\mathbf{X}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{X}_{0:K})} \right) \right] \quad (17)$$

where $w^*(\mathbf{X}_{0:K})$ is the importance weight of the discrete time trajectory $\mathbf{X}_{0:K}$.

The proof is provided in App C.5. This allows us to train by sampling trajectories, tracking their log-probabilities under the biased and base path measure, computing the terminal reward $r(\mathbf{X}_K)$, storing the values in the replay buffer \mathcal{R} , and reusing the trajectories over N_{epochs} training iterations.

5 Experiments

We evaluate **EntangledSBM** on several trajectory simulation tasks involving n -particle systems with interacting dynamics, including simulating cell cluster dynamics under drug perturbation (Sec 5.1) and transition path sampling (TPS) for molecular dynamics (MD) simulation of proteins (Sec 5.2).

5.1 Simulating Interacting Cell Dynamics Under Perturbation

Populations of cells undergo dynamic signaling interactions that cause shifts in cell state. Under a drug perturbation, these interactions determine the perturbed state of the cell population, which is of significant importance in drug discovery and screening. In this experiment, we use **EntangledSBM** to parameterize an *entangled* bias force that learns **dynamic interactions between cells**, guiding the trajectories of a batch of cells to the perturbed state. We demonstrate that EntangledSBM can not only accurately reconstruct the perturbed cell states within the training distribution following trajectories on the data manifold, but can also generate paths to *unseen* target distributions that diverge from the training distribution, demonstrating its potential to **scale to diverse perturbations and cell types with sparse data**.

Setup and Baselines To model cell state trajectories with high resolution, we evaluate two perturbations (Clonidine and Trametinib at $5\mu\text{L}$) from the Tahoe-100M dataset [Zhang et al., 2025b],

each containing data for a total of 60K genes. We select the top 2000 highly variable genes (HVGs) and perform principal component analysis (PCA), to maximally capture the variance in the data via the top principal components (38% in the top-50 PCs). To evaluate the ability of EntangledSBM to simulate trajectories to *unseen* target cell clusters that diverge in its distribution from the training target distribution, we further cluster the perturbed cell data to construct multiple disjoint perturbed cell populations. For Clonidine, we generated two clusters, and for Trametinib, we generated three clusters (Figure 5). For each experiment, we trained on a single cluster, and the remaining clusters were left for evaluation.

We trained the bias force given batches of cell clusters with $n = 16$ cells, where the initial cluster is sampled from the unperturbed cell data π_A , and the target cluster is sampled from one of the perturbed cell populations π_B . Each batch of cells $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$ has positions $\mathbf{R}_t \in \mathbb{R}^{n \times d}$ and velocities $\mathbf{V}_t \in \mathbb{R}^{n \times d}$, where d is the dimension of the principal components (PCs) that we simulate. We define the energy landscape such that regions of high data density have low energy and the base dynamics \mathbb{P}^0 follow the gradient toward high data density as detailed in App D.1. We evaluate the reconstruction accuracy of the perturbed distribution at $t = T$ using the maximum mean discrepancy (**RBF-MMD**) for all d PCs and Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) of the top two PCs between ground truth and reconstructed clusters after 100 simulation steps with details in D.2. To demonstrate the significance of entangled velocity conditioning on performance, we trained our model with only the positions as input $b_\theta(\mathbf{R}_t)$. We additionally compare our cross-entropy objective \mathcal{L}_{CE} from Sec 4.2 with the log-variance divergence \mathcal{L}_{LV} in App F. Additional experiment details are provided in App D.

Clonidine Perturbation Results We demonstrate that EntangledSBM accurately guides the base dynamics, which largely remain in the initial data-dense state, to the target perturbed state across increasing PC dimensions $d = \{50, 100, 150\}$ and reconstructs the target distribution (Fig. 2; Table 1). Notably, we show that our parameterization of the bias force enables generalization to **unseen perturbed populations** that *diverge* from the training distribution by learning dependencies on the target state (Fig. 2). Comparing the reconstruction metrics of the target perturbed cell states with and without velocity conditioning, we confirm our hypothesis that introducing dependence on the dynamic velocities across a cell cluster enables more accurate reconstruction of the target distribution compared to when the bias term is trained with only positional dependence (Table 1). Furthermore, we observe that the LV objective \mathcal{L}_{LV} generates nearly straight abrupt trajectories to the target state while the CE objective \mathcal{L}_{CE} generates smooth paths along the data manifold (Table 6; Fig. 7), demonstrating its superiority as an **unconstrained objective** for learning controls for optimizing path distributions as further discussed in App F.

Trametinib Perturbation Results We further evaluate our method to predict trajectories of cell clusters under perturbation with Trametinib, which induces three distinct perturbed cell distributions (Figure 5B). Despite training on only one of the perturbed distributions, we demonstrate that EntangledSBM trained with \mathcal{L}_{CE} is capable of accurately reconstructing the remaining cell distributions *without* additional training, demonstrating significantly improved performance compared to the bias force trained without velocity conditioning (Table 1; Fig. 3). Furthermore, we observe the same phenomena observed for Clonidine when training with the LV objective, which results in trajectories that fail to capture the intermediate cell dynamics (Table 7; Fig. 8).

5.2 Transition Path Sampling of Protein Folding Dynamics

Simulating rare transitions that occur over long timescales and between metastable states on molecular dynamics (MD) landscapes remains a significant challenge due to the high feature dimensionality of biomolecules. In this experiment, we aim to simulate feasible transition paths across high-energy barriers at an *all-atom* resolution, a task that is challenging for both traditional MD and ML-based methods. We demonstrate that EntangledSBM generates feasible transition paths with higher target accuracy against a range of baselines.

Setup and Baselines We follow the setup in Seong et al. [2025] and simulate the position and velocity using OpenMM [Eastman et al., 2017] with an annealed temperature schedule. To evaluate the performance of EntangledSBM, we consider the RMSD of the Kabsch-aligned coordinates averaged across 64 paths (**RMSD**; Å) [Kabsch, 1976], percentage of simulated trajectories that hit the

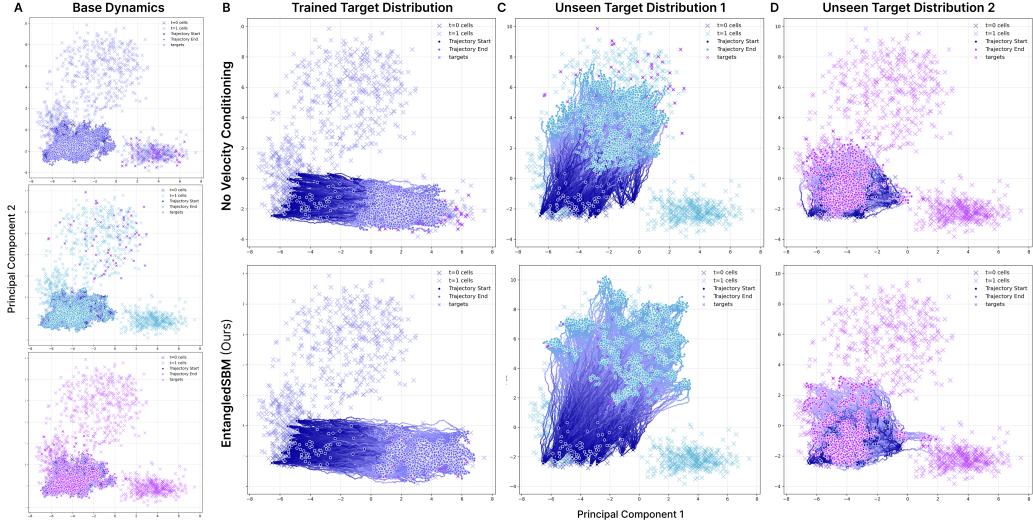


Figure 3: **Comparison of EntangledSBM with and without velocity conditioning for cell cluster simulation under Trametinib perturbation.** 50 PCs are simulated with the learned bias force trained with the CE objective without velocity conditioning (**Top**) and with velocity conditioning (**Bottom**) to (**B**) the perturbed population used for training and (**C, D**) the two unseen target populations.

Table 1: **Results for simulating cell cluster dynamics under Clonidine and Trametinib perturbation with EntangledSBM.** Best values for each PC dimension are **bolded**. We report RBF-MMD for all PCs and \mathcal{W}_1 and \mathcal{W}_2 distances of the top 2 PCs between ground truth and reconstructed clusters after 100 simulation steps and cluster size set to $n = 16$. Mean and standard deviation of metrics from 5 independent simulations are reported. We simulate PCs $d = \{50, 100, 150\}$ for Clonidine and $d = 50$ for Trametinib to cells sampled from the training target distribution and unseen target distributions, and compare against the learned bias force with no velocity conditioning.

Clonidine Perturbation						
Model	Trained Target Distribution			Unseen Target Distribution		
	RBF-MMD (↓)	\mathcal{W}_1 (↓)	\mathcal{W}_2 (↓)	RBF-MMD (↓)	\mathcal{W}_1 (↓)	\mathcal{W}_2 (↓)
Base Dynamics (50 PCs)	0.677 ± 0.001	5.947 ± 0.005	6.015 ± 0.005	0.784 ± 0.001	8.217 ± 0.005	8.384 ± 0.005
EntangledSBM w/o Velocity Conditioning						
50 PCs	0.440 ± 0.000	1.741 ± 0.003	1.857 ± 0.004	0.478 ± 0.000	2.907 ± 0.006	3.022 ± 0.006
100 PCs	0.494 ± 0.000	2.315 ± 0.004	2.423 ± 0.004	0.539 ± 0.000	4.110 ± 0.004	4.249 ± 0.003
150 PCs	0.510 ± 0.000	2.497 ± 0.006	2.620 ± 0.006	0.560 ± 0.000	4.573 ± 0.006	4.716 ± 0.007
EntangledSBM w/ CE						
50 PCs	0.401 ± 0.000	0.342 ± 0.002	0.400 ± 0.001	0.419 ± 0.000	0.538 ± 0.013	0.705 ± 0.030
100 PCs	0.455 ± 0.000	0.953 ± 0.025	1.015 ± 0.025	0.500 ± 0.001	0.899 ± 0.006	1.055 ± 0.008
150 PCs	0.478 ± 0.000	0.753 ± 0.008	0.826 ± 0.007	0.506 ± 0.000	0.700 ± 0.009	0.811 ± 0.011

Trametinib Perturbation								
Method	Trained Target Distribution			Unseen Target Distribution 1			Unseen Target Distribution 2	
	RBF-MMD (↓)	\mathcal{W}_1 (↓)	\mathcal{W}_2 (↓)	RBF-MMD (↓)	\mathcal{W}_1 (↓)	\mathcal{W}_2 (↓)	RBF-MMD (↓)	\mathcal{W}_1 (↓)
Base Dynamics	0.938 ± 0.001	7.637 ± 0.005	7.653 ± 0.006	0.900 ± 0.000	7.766 ± 0.009	7.877 ± 0.009	0.754 ± 0.001	1.201 ± 0.009
EntangledSBM w/o Velocity Conditioning	0.449 ± 0.000	1.506 ± 0.005	1.544 ± 0.005	0.476 ± 0.000	2.116 ± 0.005	2.197 ± 0.005	0.480 ± 0.000	0.505 ± 0.004
EntangledSBM w/ CE	0.428 ± 0.000	0.392 ± 0.005	0.434 ± 0.006	0.409 ± 0.000	0.453 ± 0.008	0.561 ± 0.009	0.451 ± 0.000	0.394 ± 0.003

target state (**THP**; %), and the highest energy transition state along the biased trajectories averaged across the trajectories that hit the target (**ETS**; kJ/mol). For baselines, we compare against unbiased MD (**UMD**) with temperatures of 3600K for alanine dipeptide and 1200K for the fast-folding proteins, steered MD (**SMD**; Schlitter et al. [1994], Izrailev et al. [1999]) with temperatures of 10K and 20K, path integral path sampling (**PIPS**; Holdijk et al. [2023]), transition path sampling with diffusion path samplers (**TPS-DPS**; Seong et al. [2025]) with the highest performing *scaled* parameterization. Additional experiment details are provided in App E.

Alanine Dipeptide First, we consider Alanine Dipeptide with two alanine residues and 22 atoms. We aim to simulate trajectories to the target state $\pi_B = \{R \mid \|\xi(R) - \xi(R_B)\| < 0.75\}$ defined

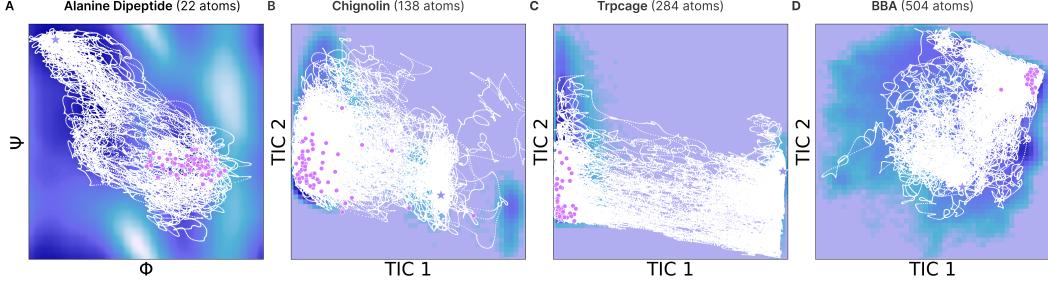


Figure 4: **Transition paths generated with EntangledSBM.** The energy landscape is colored from dark navy (low potential energy) to light purple (high potential energy) and plotted with the backbone dihedral angles (ϕ, ψ) for Alanine Dipeptide and the top two TICA components for the fast-folding proteins. The starting state \mathbf{R}_0 and target states \mathbf{R}_B are indicated with purple stars, and the final states of the simulated trajectories \mathbf{R}_T are indicated with pink circles.

Table 2: **Transition path sampling benchmarks with EntangledSBM.** Best values are **bolded**. All metrics are averaged over 64 paths. Unless specified in brackets, paths are generated at 300K for Alanine Dipeptide and Chignolin and 400K for the others. Hyperparameters and evaluation metrics are detailed in App E. \dagger denotes values taken from [Seong et al. \[2025\]](#).

Protein		Alanine Dipeptide			Chignolin			
Method		RMSD (\downarrow)	THP (\uparrow)	ETS (\downarrow)	Method	RMSD (\downarrow)	THP (\uparrow)	ETS (\downarrow)
UMD \dagger		1.19 ± 0.32	6.25	—	UMD \dagger	7.23 ± 0.93	1.56	388.17
SMD (10K) \dagger		0.86 ± 0.21	7.81	812.47 ± 148.80	SMD (10K) \dagger	1.26 ± 0.31	6.25	-527.95 ± 93.58
SMD (20K) \dagger		0.56 ± 0.27	54.69	78.40 ± 12.76	SMD (20K) \dagger	0.85 ± 0.24	34.38	179.52 ± 138.87
PIPS (Force) \dagger		0.66 ± 0.15	43.75	28.17 ± 10.86	PIPS (Force) \dagger	4.66 ± 0.17	0.00	—
TPS-DPS (Scalar) \dagger		0.25 ± 0.20	76.00	22.79 ± 13.57	TPS-DPS (Scalar) \dagger	1.17 ± 0.66	59.38	-780.18 ± 216.93
EntangledSBM (Ours)		0.18 ± 0.07	92.19	47.91 ± 22.76	EntangledSBM (Ours)	0.92 ± 0.13	64.06	2825.61 ± 318.94

Protein		Trp-cage			BBA			
Method		RMSD (\downarrow)	THP (\uparrow)	ETS (\downarrow)	Method	RMSD (\downarrow)	THP (\uparrow)	ETS (\downarrow)
UMD \dagger		8.27 ± 1.13	0.00	—	UMD \dagger	10.81 ± 1.05	0.00	-
SMD (10K) \dagger		1.68 ± 0.23	3.12	-312.54 ± 20.67	SMD (10K) \dagger	2.89 ± 0.32	0.00	-
SMD (20K) \dagger		1.20 ± 0.20	42.19	-226.40 ± 85.59	SMD (20K) \dagger	1.66 ± 0.30	26.56	-3104.95 ± 97.57
PIPS (Force) \dagger		7.47 ± 0.19	0.00	-	PIPS (Force) \dagger	9.84 ± 0.18	0.00	-
TPS-DPS (Scalar) \dagger		0.76 ± 0.12	81.25	-317.61 ± 140.89	TPS-DPS (Scalar) \dagger	1.21 ± 0.09	84.38	-3801.68 ± 139.38
EntangledSBM (Ours)		1.04 ± 0.22	82.81	765.74 ± 155.28	EntangledSBM (Ours)	0.84 ± 0.08	96.88	1453.80 ± 367.84

by the backbone dihedral angles $\xi(\mathbf{R}) = (\phi, \psi)$. We show that EntangledSBM generates feasible trajectories through both saddle points, representing the two reaction channels (Figure 4A), achieving superior target hit potential (THP) and lower root mean squared error (RMSD) from the target state than all baselines (Table 2).

Fast-Folding Proteins We evaluate EntangledSBM on the more challenging task of modeling the all-atom transition paths of four fast-folding proteins, including **Chignolin** (10 amino acids; 138 atoms), **Trp-cage** (284 atoms; 20 amino acids), and **BBA** (504 atoms; 28 amino acids). We define a hit as a trajectory where the final state reaches the metastable distribution $\pi_B = \{\mathbf{R} \mid \|\xi(\mathbf{R}) - \xi(\mathbf{R}_B)\| < 0.75\}$, where $\xi(\mathbf{R})$ are the top two time-lagged independent component analysis (TICA; [Pérez-Hernández et al. \[2013\]](#)) components. As shown in Fig. 4, EntangledSBM generates *diverse* transition paths across high-energy barriers that successfully reach the target state. We achieve a superior target hit percentage than all baselines across all proteins and a lower or comparable RMSD (Table 2). While we observe a higher average energy of transition state (ETS) compared to baselines, this can be attributed to the larger proportion and greater diversity of successful target-hitting paths. Given that the base dynamics \mathbb{P}^0 rarely move beyond the initial metastable state, we show that EntangledSBM effectively learns the target bridge dynamics \mathbb{P}^* despite high energy barriers.

6 Conclusion

In this work, we present **Entangled Schrödinger Bridge Matching (EntangledSBM)**, a principled framework for learning the second-order dynamics of interacting multi-particle systems through entangled bias forces and an unconstrained cross-entropy objective. EntangledSBM captures dependencies between particle positions and velocities, enabling the modeling of complex dynamics across biological scales. For perturbation modeling, EntangledSBM reconstructs perturbed cell states while generalizing to divergent target states not seen during training, and for molecular dynamics (MD), it generates physically plausible transition paths for fast-folding proteins at an all-atom resolution.

Limitations and Future Directions Our experiments primarily demonstrate that integrating entangled velocities and a cross-entropy objective for path distribution matching enhances performance in modeling multi-particle systems. While we demonstrate that our unique parameterization enables performance gains across diverse systems across biological scales, it remains limited to a selected set of perturbations and small, fast-folding proteins. We envision our method generalizing across a broader range of drug-induced and genetic perturbations across diverse cell types and representations, larger proteins and biomolecules, as well as other multi-particle physical systems.

Declarations

Acknowledgments We thank Mark III Systems for providing database and hardware support that has contributed to the research reported within this manuscript.

Author Contributions S.T. conceived the study. S.T. devised and developed model architectures and theoretical formulations, and trained and benchmarked models. Y.Z. advised on model design and theoretical framework. S.T. drafted the manuscript and designed the figures. P.C. supervised and directed the study, and reviewed and finalized the manuscript.

Data and Materials Availability The codebase is freely accessible to the academic community at <https://github.com/sophtang/EntangledSBM> and <https://huggingface.co/ChatterjeeLab/EntangledSBM>.

Funding Statement This research was supported by NIH grant R35GM155282 to the lab of P.C.

Competing Interests P.C. is a co-founder of Gameto, Inc., UbiquiTx, Inc., Atom Bioworks, Inc., and advises companies involved in biologics development and cell engineering. P.C.’s interests are reviewed and managed by the University of Pennsylvania in accordance with their conflict-of-interest policies. S.T. and Y.Z. have no conflicts of interest to declare.

References

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

Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, 238:1279–1287, 2024.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A. Theodorou, and Ricky T. Q. Chen. Generalized schrödinger bridge matching. *International Conference on Learning Representations*, 2023.

Sophia Tang, Yinuo Zhang, Alexander Tong, and Pranam Chatterjee. Branched schrödinger bridge matching. *arXiv preprint arXiv:2506.09007*, 2025a.

Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos A. Theodorou. Deep generalized schrödinger bridge. *Advances in Neural Information Processing Systems*, 2022.

Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and interactions with unbalanced mean field schrödinger bridge. *arXiv preprint arXiv:2505.11197*, 2025a.

Sophia Vincoff, Shrey Goel, Ksenia Kholina, Rishab Pulugurta, Pranay Vure, and Pranam Chatterjee. Fuson-plm: a fusion oncoprotein-specific language model via adjusted rate masking. *Nature Communications*, 16(1):1436, 2025.

Tianlai Chen, Lauren Hong, Vivian Yudistyra, Sophia Vincoff, and Pranam Chatterjee. Generative design of therapeutics that bind and modulate protein states. *Current Opinion in Biomedical Engineering*, 28:100496, 2023.

Merrick D Pierson Smela, Christian C Kramme, Patrick RJ Fortuna, Jessica L Adams, Rui Su, Edward Dong, Mutsumi Kobayashi, Garyk Brixi, Venkata Srikanth Kavirayuni, Emma Tysinger, et al. Directed differentiation of human ipscs to functional ovarian granulosa-like cells via transcription factor overexpression. *Elife*, 12:e83291, 2023a.

Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. The emergence of clusters in self-attention dynamics. *Advances in Neural Information Processing Systems*, 36:57026–57037, 2023.

Giovanni Bussi and Michele Parrinello. Accurate sampling using langevin dynamics. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 75(5):056707, 2007.

Sichun Yang, José N Onuchic, and Herbert Levine. Effective stochastic dynamics on a protein folding energy landscape. *The Journal of chemical physics*, 125(5), 2006.

Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of Crystallography*, 32(5):922–923, 1976.

Kiyoung Seong, Seonghyun Park, Seonghwan Kim, Woo Youn Kim, and Sungsoo Ahn. Transition path sampling with improved off-policy training of diffusion path samplers. In *The Thirteenth International Conference on Learning Representations*, 2025.

Nikolas Nüsken and Lorenz Richter. Solving high-dimensional hamilton–jacobi–bellman pdes using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2(4):48, 2021.

Hilbert Johan Kappen and Hans Christian Ruiz. Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 162(5):1244–1266, 2016.

Yuchen Zhu, Wei Guo, Jaemoo Choi, Guan-Horng Liu, Yongxin Chen, and Molei Tao. Mdns: Masked diffusion neural sampler via stochastic optimal control. *arXiv preprint arXiv:2508.10684*, 2025.

Sophia Tang, Yuchen Zhu, Molei Tao, and Pranam Chatterjee. Tr2-d2: Tree search guided trajectory-aware fine-tuning for discrete diffusion. *arXiv preprint arXiv:2509.25171*, 2025b.

Jesse Zhang, Airol A Ubas, Richard de Borja, Valentine Svensson, Nicole Thomas, Neha Thakar, Ian Lai, Aidan Winters, Umair Khan, and Matthew G. et al. Jones. Tahoe-100m: A giga-scale single-cell perturbation atlas for context-dependent gene function and cellular modeling. *bioRxiv*, 2025b.

Peter Eastman, Jason Swails, John D Chodera, Robert T McGibbon, Yutong Zhao, Kyle A Beauchamp, Lee-Ping Wang, Andrew C Simmonett, Matthew P Harrigan, Chaya D Stern, et al. Openmm 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS computational biology*, 13(7):e1005659, 2017.

Jürgen Schlitter, Michael Engels, and Peter Krüger. Targeted molecular dynamics: a new approach for searching pathways of conformational transitions. *Journal of molecular graphics*, 12(2):84–89, 1994.

Sergei Izrailev, Sergey Stepaniants, Barry Isralewitz, Dorina Kosztin, Hui Lu, Ferenc Molnar, Willy Wriggers, and Klaus Schulten. Steered molecular dynamics. In *Computational Molecular Dynamics: Challenges, Methods, Ideas: Proceedings of the 2nd International Symposium on Algorithms for Macromolecular Modelling, Berlin, May 21–24, 1997*, pages 39–65. Springer, 1999.

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

Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annual review of physical chemistry*, 53(1):291–318, 2002.

Christoph Dellago, Peter G Bolhuis, Félix S Csajka, and David Chandler. Transition path sampling and the calculation of rate constants. *The Journal of chemical physics*, 108(5):1964–1977, 1998.

Eric Vanden-Eijnden et al. Transition-path theory and path-finding algorithms for the study of rare events. *Annual review of physical chemistry*, 61:391–420, 2010.

Ferry Hooft, Alberto Perez de Alba Ortiz, and Bernd Ensing. Discovering collective variables of molecular transitions via genetic algorithms and neural networks. *Journal of chemical theory and computation*, 17(4):2294–2306, 2021.

Glenn M Torrie and John P Valleau. Nonphysical sampling distributions in monte carlo free-energy estimation: Umbrella sampling. *Journal of computational physics*, 23(2):187–199, 1977.

Johannes Kästner. Umbrella sampling. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 1(6):932–942, 2011.

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

Bernd Ensing, Marco De Vivo, Zhiwei Liu, Preston Moore, and Michael L Klein. Metadynamics as a tool for exploring free energy landscapes of chemical reactions. *Accounts of chemical research*, 39(2):73–81, 2006.

Davide Branduardi, Giovanni Bussi, and Michele Parrinello. Metadynamics with adaptive gaussians. *Journal of chemical theory and computation*, 8(7):2247–2254, 2012.

Giovanni Bussi and Davide Branduardi. Free-energy calculations with metadynamics: Theory and practice. *Reviews in Computational Chemistry Volume 28*, pages 1–49, 2015.

Jeffrey Comer, James C Gumbart, Jérôme Hénin, Tony Lelièvre, Andrew Pohorille, and Christophe Chipot. The adaptive biasing force method: Everything you always wanted to know but were afraid to ask. *The Journal of Physical Chemistry B*, 119(3):1129–1151, 2015.

Michele Invernizzi and Michele Parrinello. Rethinking metadynamics: from bias potentials to probability distributions. *The journal of physical chemistry letters*, 11(7):2731–2736, 2020.

Mohammad M Sultan and Vijay S Pande. Automated design of collective variables using supervised machine learning. *The Journal of chemical physics*, 149(9), 2018.

Jutta Rogal, Elia Schneider, and Mark E Tuckerman. Neural-network-based path collective variables for enhanced sampling of phase transformations. *Physical Review Letters*, 123(24):245701, 2019.

Wei Chen and Andrew L Ferguson. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *Journal of computational chemistry*, 39(25):2079–2102, 2018.

Lixin Sun, Jonathan Vandermause, Simon Batzner, Yu Xie, David Clark, Wei Chen, and Boris Kozinsky. Multitask machine learning of collective variables for enhanced sampling of rare events. *Journal of Chemical Theory and Computation*, 18(4):2341–2353, 2022.

Martin Sipka, Andreas Erlebach, and Lukas Grajciar. Constructing collective variables using invariant learned representations. *Journal of Chemical Theory and Computation*, 19(3):887–901, 2023a.

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

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

Martin Sipka, Johannes CB Dietschreit, Lukáš Grajciar, and Rafael Gómez-Bombarelli. Differentiable simulations for enhanced sampling of rare events. In *International Conference on Machine Learning*, pages 31990–32007. PMLR, 2023b.

Avishhek Das, Dominic C Rose, Juan P Garrahan, and David T Limmer. Reinforcement learning of rare diffusive dynamics. *The Journal of Chemical Physics*, 155(13), 2021.

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

Sanjeev Raja, Martin Šípka, Michael Psenka, Tobias Kreiman, Michal Pavelka, and Aditi S Krishnapriyan. Action-minimization meets generative modeling: Efficient transition path sampling with the onsager-machlup functional. *arXiv preprint arXiv:2504.18506*, 2025.

Denis Blessing, Julius Berner, Lorenz Richter, Carles Domingo-Enrich, Yuanqi Du, Arash Vahdat, and Gerhard Neumann. Trust region constrained measure transport in path space for stochastic optimal control and inference. *arXiv preprint arXiv:2508.12511*, 2025.

Ophir Shalem, Neville E Sanjana, Ella Hartenian, Xi Shi, David A Scott, Tarjei S Mikkelsen, Dirk Heckl, Benjamin L Ebert, David E Root, John G Doench, et al. Genome-scale crispr-cas9 knockout screening in human cells. *Science*, 343(6166):84–87, 2014.

Christian Kramme, Alexandru M. Plesa, Helen H. Wang, Bennett Wolf, Merrick Pierson Smela, Xiaoge Guo, Richie E. Kohman, Pranam Chatterjee, and George M. Church. An integrated pipeline for mammalian genetic screening. *Cell Reports Methods*, 1(6):100082, October 2021.

Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D Marjanovic, Danielle Dionne, Tyler Burks, Rakirta Raychowdhury, et al. Perturb-seq: dissecting molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *cell*, 167(7):1853–1866, 2016.

George I Gavrilidis, Vasileios Vasileiou, Aspasia Orfanou, Naveed Ishaque, and Fotis Psomopoulos. A mini-review on perturbation modelling across single-cell omic modalities. *Computational and Structural Biotechnology Journal*, 23:1886, 2024.

Mutsumi Kobayashi, Misato Kobayashi, Junko Odajima, Keiko Shioda, Young Sun Hwang, Kotaro Sasaki, Pranam Chatterjee, Christian Kramme, Richie E Kohman, George M Church, et al. Expanding homogeneous culture of human primordial germ cell-like cells maintaining germline features without serum or feeder layers. *Stem Cell Reports*, 17(3):507–521, 2022.

Merrick D Pierson Smela, Christian C Kramme, Patrick RJ Fortuna, Jessica L Adams, Rui Su, Edward Dong, Mutsumi Kobayashi, Garyk Brixi, Venkata Srikar Kavirayuni, Emma Tysinger, et al. Directed differentiation of human ipscs to functional ovarian granulosa-like cells via transcription factor overexpression. *Elife*, 12:e83291, 2023b.

Merrick Pierson Smela, Christian C Kramme, Patrick RJ Fortuna, Bennett Wolf, Shrey Goel, Jessica Adams, Carl Ma, Sergiy Velychko, Ursula Widocki, Venkata Srikar Kavirayuni, et al. Rapid human oogonia-like cell specification via transcription factor-directed differentiation. *EMBO Reports*, pages 1–30, 2025.

Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):3222, 2021.

Yuhui Zhang, Yuchang Su, Chenyu Wang, Tianhong Li, Zoe Wefers, Jeffrey Nirschl, James Burgess, Daisy Ding, Alejandro Lozano, Emma Lundberg, et al. Cellflow: Simulating cellular morphology changes via flow matching. *arXiv preprint arXiv:2502.09775*, 2025c.

Martin Rohbeck, Charlotte Bunne, Edward De Brouwer, Jan-Christian Huetter, Anne Biton, Kelvin Y. Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow matching across time and conditions. *The Thirteenth International Conference on Learning Representations*, 2025.

Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. *arXiv preprint arXiv:2408.14608*, 2024.

Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025.

Sayali Anil Alatkar and Daifeng Wang. Artemis integrates autoencoders and schrodinger bridges to predict continuous dynamics of gene expression, cell population and perturbation from time-series single-cell data. *bioRxiv*, pages 2025–01, 2025.

Kacper Kapuśnian, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. *Advances in Neural Information Processing Systems*, 2024.

Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. *International Conference on Learning Representations*, 2025d.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

David A Sivak, John D Chodera, and Gavin E Crooks. Time step rescaling recovers continuous-time dynamical properties for discrete-time langevin integration of nonequilibrium systems. *The Journal of Physical Chemistry B*, 118(24):6466–6474, 2014.

Kresten Lindorff-Larsen, Stefano Piana, Kim Palmo, Paul Maragakis, John L Klepeis, Ron O Dror, and David E Shaw. Improved side-chain torsion potentials for the amber ff99sb protein force field. *Proteins: Structure, Function, and Bioinformatics*, 78(8):1950–1958, 2010.

James A Maier, Carmenza Martinez, Koushik Kasavajhala, Lauren Wickstrom, Kevin E Hauser, and Carlos Simmerling. ff14sb: improving the accuracy of protein side chain and backbone parameters from ff99sb. *Journal of chemical theory and computation*, 11(8):3696–3713, 2015.

Hai Nguyen, Daniel R Roe, and Carlos Simmerling. Improved generalized born solvent model parameters for protein simulations. *Journal of chemical theory and computation*, 9(4):2020–2034, 2013.

Outline of Appendix

In App A, we provide an overview and discussion of related works in off-policy learning, Schrödinger bridge matching, and transition path sampling. App B provides the theoretical basis of stochastic optimal control (SOC) theory used in our work and a theoretical motivation for our entangled bias force. App C provides the theoretical formulations and proofs for EntangledSBM. We provide details on the experimental setup and training hyperparameters for the cell perturbation experiment in App D and the transition path sampling (TPS) experiment in App E. We further compare the log-variance (LV) divergence and cross-entropy (CE) objectives in App F. Finally, we provide the pseudocode for the training and inference of EntangledSBM in App G.

Notation In this work, we consider an n -particle system $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$ where $\mathbf{R}_t \in \mathbb{R}^{n \times d}$ and $\mathbf{V}_t \in \mathbb{R}^{n \times d}$ denote the positions and velocities of the full system that evolves over the time horizon $t \in [0, T]$. The notation \mathbf{X} indicates a random variable and \mathbf{x} denotes a deterministic realization. Each particle i is defined by its position $\mathbf{r}_t^i \in \mathbb{R}^d$ and velocity $\mathbf{v}_t^i \in \mathbb{R}^d$ which lie in d -dimensional Euclidean space. The base dynamics evolve according to a potential energy function denoted $U(\mathbf{R}_t) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ and the biased dynamics evolve with an additional parameterized bias force $\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t) : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ that captures the implicit dependencies within the system and controls the base dynamics towards a target distribution $\pi_{\mathcal{B}}(\mathbf{R}_T)$ from a state in the initial distribution $\mathbf{R}_A \in \pi_{\mathcal{A}}(\mathbf{R}_0)$. The corresponding path measures over trajectories $\mathbf{X}_{0:T} := (\mathbf{X}_t)_{t \in [0, T]}$ are denoted \mathbb{P}^0 for the base process and \mathbb{P}^{b_θ} for the biased process. The optimal path measure that solves the EntangledSB problem is denoted $\mathbb{P}^* \equiv \mathbb{P}^{b^*}$ with bias force $\mathbf{b}^*(\mathbf{R}_t, \mathbf{V}_t)$.

A Related Works

Off-Policy Learning Importance-weighted cross-entropy (CE) objectives have previously been used to sample the target distribution $g^* \propto h f_u$ by aligning a proposal distribution f_v to the intractable target with $\text{KL}(g^* \| f_v)$ [Kappen and Ruiz, 2016]. The main difference is that our approach aims to match a **controlled path measure** \mathbb{P}^{b_θ} with an energy-minimizing path measure \mathbb{P}^* rather than a static target distribution. Furthermore, we adapt our path CE objective to apply for discrete-time paths, similar to the approach taken by Seong et al. [2025], which discretizes the log-variance (LV) divergence [Nüsken and Richter, 2021]. While the LV objective has theoretical optimality guarantees, it is non-convex in \mathbb{P}^{b_θ} . Since the CE objective is convex in \mathbb{P}^{b_θ} , it yields an ideal optimization landscape for θ and demonstrates improved performance as shown in Sec 5.1. We also highlight that our optimization scheme is a form of *off-policy* learning, as we minimize the objective with respect to the path distribution generated from an auxiliary control \mathbf{v} which we set to the non-gradient-tracking bias force $\mathbf{v} := \bar{\mathbf{b}} = \text{stopgrad}(\mathbf{b}_\theta)$. This objective is inspired by recent work that leverages off-policy reinforcement learning in the discrete state space for discrete diffusion sampling [Zhu et al., 2025] and fine-tuning [Tang et al., 2025b].

Transition Path Sampling To overcome the challenge of generating feasible transition paths across high energy barriers for transition path sampling, several non-ML and ML-based approaches have been introduced [Bolhuis et al., 2002, Dellago et al., 1998, Vanden-Eijnden et al., 2010]. Non-ML approaches often rely on collective variables (CVs), which reduce the dimensionality of the molecular conformation to a function of coordinates that are known to be involved in transition states [Hooft et al., 2021]. These include steered MD [Schlitter et al., 1994, Izrailev et al., 1999], umbrella sampling [Torrie and Valleau, 1977, Kästner, 2011], meta-dynamics [Laio and Parrinello, 2002, Ensing et al., 2006, Branduardi et al., 2012, Bussi and Branduardi, 2015], adaptive biasing force [Comer et al., 2015], and on-the-fly probability-enhanced sampling [Invernizzi and Parrinello, 2020]. To overcome the lack of well-understood CVs for large biomolecular systems, ML-based methods have been explored for determining or constructing CVs [Sultan and Pande, 2018, Rogal et al., 2019, Chen and Ferguson, 2018, Sun et al., 2022, Sipka et al., 2023a], but modeling transition paths at all-atom resolution is still desirable due to systems of increased complexity, and CVs are uncertain.

Several ML-based approaches for simulating transition paths have been developed [Singh and Limmer, 2023, Yan et al., 2022, Sipka et al., 2023b, Das et al., 2021]. Notably, Path Integral Path Sampling (PIPS; Holdijk et al. [2023]) learns a bias force with stochastic optimal control theory, Doob's Lagrangian Du et al. [2024] defines an optimal transition path with a Lagrangian

objective that solves the Doob’s h -transform, Transition Path Sampling with Diffusion Path Samplers (TPS-DPS; [Seong et al. \[2025\]](#)) trains a bias force with an off-policy log-variance objective, and Action Minimization with Onsager-Machlup (OM) Functional [\[Raja et al., 2025\]](#) generates a discrete interpolation between endpoints by minimizing the OM functional. Furthermore, [Blessing et al. \[2025\]](#) introduces a trust-region-constrained SOC optimization algorithm, which is applied to TPS. While these approaches demonstrate the promise of ML-based TPS, their applicability remains limited to small-scale biomolecular systems (Alanine Dipeptide, tetrapeptides, etc.) for all-atom simulation [\[Holdijk et al., 2023, Du et al., 2024\]](#) or require coarse-grained representations when scaling up to larger fast-folding proteins [\[Raja et al., 2025\]](#). We directly compare our method against [Seong et al. \[2025\]](#) and [Blessing et al. \[2025\]](#), and achieve state-of-the-art performance on all-atom simulations of larger, fast-folding proteins, and [Blessing et al. \[2025\]](#). Furthermore, we highlight that our framework extends beyond molecular dynamics (MD) simulations, with applications to single-cell simulation and generalization to unseen target distributions.

Modeling Cell Dynamics Under Perturbation Predicting the dynamics of heterogeneous cell populations under perturbations such as treatment with a drug candidate, gene editing, and knockouts, or protein expression, has critical applications in drug discovery and predictive medicine [\[Shalem et al., 2014, Kramme et al., 2021, Dixit et al., 2016, Gavrilidis et al., 2024, Zhang et al., 2025b, Kobayashi et al., 2022, Smela et al., 2023b, Pierson Smela et al., 2025, Yeo et al., 2021\]](#). Generative modeling frameworks including flow matching [\[Zhang et al., 2025c, Rohbeck et al., 2025, Atanackovic et al., 2024, Tong et al., 2024, Wang et al., 2025\]](#), Schrödinger Bridge Matching [\[Alatkar and Wang, 2025, Tong et al., 2024, Kapušniak et al., 2024, Tang et al., 2025a\]](#), and stochastic optimal control [\[Zhang et al., 2025d,a\]](#) have enabled significant advances in parameterizing velocities that evolve dynamically over time on non-linear energy manifolds by learning on static snapshots across the trajectory. However, these models often require *explicit training on a specific target distribution* and fail to generalize to unseen target distributions at inference, especially if they diverge from those seen in the training data, which limits their scalability to cell types and perturbations that are not seen during training. Furthermore, many of these approaches simulate cells as populations of particles that evolve independently from each other, which fails to capture the diverse intercellular signaling and interactions that occur naturally and under perturbations. Although there have been some early developments in learning these interactions from data [\[Zhang et al., 2025a\]](#), they operate under the mean-field assumption, where each cell takes the average effect of all surrounding cells rather than simulating individual pairwise interactions. This assumption does not account for the complex interactions between cells and limits its applicability to heterogeneous cell populations.

B Extended Theoretical Preliminaries

In this section, we provide relevant theoretical background in stochastic optimal control (SOC) theory for path measures. We use the terms *bias force* and *optimal control* interchangeably, denoted with \mathbf{b} , which is less common in general SOC literature where \mathbf{u} is often used. For a more comprehensive analysis of SOC theory, we refer the reader to [Nüsken and Richter \[2021\]](#).

B.1 Stochastic Optimal Control

Controlled Path Measures First, we consider the controlled path measure \mathbb{P}^b defined by the SDE:

$$d\mathbf{X}_t = [\mathbf{f}(\mathbf{X}_t, t) + \Sigma \mathbf{b}(\mathbf{X}_t, t)]dt + \Sigma d\mathbf{W}_t \quad (18)$$

where $\mathbf{b} : [0, T] \times \mathcal{X} \rightarrow \mathcal{X} \in \mathcal{U}$ is known as the **control**¹ that tilts the path measure from the unconditional dynamics \mathbb{P}^0 . Now, we will define **Radon-Nikodym derivative (RND)** of the path measures corresponding to the controlled and unconditional SDEs, which is crucial for defining our objective.

¹ \mathcal{U} is the set of admissible controls continuously differentiable in (\mathbf{x}, t) with bounded linear growth in \mathbf{x}

Lemma B.1 (Radon-Nikodym Derivative). *Given a pair of SDEs with and without a control drift $\mathbf{b}(\mathbf{X}_t, t)$ on $t \in [0, T]$ with the same diffusion Σ defined as*

$$\mathbb{P}^0 : d\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t, t)dt + \Sigma d\mathbf{W}_t, \quad \mathbf{X}_0 = \mathbf{x}_0 \quad (19)$$

$$\mathbb{P}^b : d\mathbf{X}_t^b = [\mathbf{f}(\mathbf{X}_t, t) + \Sigma \mathbf{b}(\mathbf{X}_t, t)] dt + \Sigma d\mathbf{W}_t, \quad \mathbf{X}_0^b = \mathbf{x}_0 \quad (20)$$

Then, the Radon-Nikodym derivative satisfies

$$\log \frac{d\mathbb{P}^b}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T (\mathbf{b}^\top \Sigma^{-1})(\mathbf{X}_t, t) d\mathbf{X}_t - \int_0^T (\Sigma^{-1} \mathbf{f} \cdot \mathbf{b})(\mathbf{X}_t, t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}(\mathbf{X}_t, t)\|^2 dt$$

or equivalently:

$$\log \frac{d\mathbb{P}^b}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T \mathbf{b}(\mathbf{X}_t, t)^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}(\mathbf{X}_t, t)\|^2 dt \quad (21)$$

Proof. First, we use Girsanov's theorem to get the Radon-Nikodym derivative with respect to the zero-drift reference path measure \mathbb{Q} defined by the SDE $d\mathbf{X}_t = \Sigma(\mathbf{X}_t, t)d\mathbf{W}_t$, we have

$$\log \frac{\mathbb{P}^0}{\mathbb{Q}}(\mathbf{X}_{0:T}) = \int_0^T (\mathbf{f}(\mathbf{X}_t, t) \cdot \Sigma^{-2}(\mathbf{X}_t, t)) d\mathbf{X}_t - \frac{1}{2} \int_0^T (\mathbf{f} \cdot \Sigma^{-2} \mathbf{f})(\mathbf{X}_t, t) dt \quad (22)$$

and for the controlled path measure \mathbb{P}^b , we have

$$\begin{aligned} \log \frac{\mathbb{P}^b}{\mathbb{Q}}(\mathbf{X}_{0:T}) &= \int_0^T (\mathbf{f} + \Sigma \mathbf{b})(\mathbf{X}_t, t) \cdot \Sigma^{-2}(\mathbf{X}_t, t) d\mathbf{X}_t \\ &\quad - \frac{1}{2} \int_0^T ((\mathbf{f} + \Sigma \mathbf{b}) \cdot \Sigma^{-2}(\mathbf{f} + \Sigma \mathbf{b}))(\mathbf{X}_t, t) dt \end{aligned} \quad (23)$$

Then, using the identity

$$\log \frac{d\mathbb{P}^b}{d\mathbb{P}^0} = \log \frac{d\mathbb{P}^b}{d\mathbb{Q}} \frac{d\mathbb{Q}}{d\mathbb{P}^0} = \log \frac{d\mathbb{P}^b}{d\mathbb{Q}} + \log \frac{d\mathbb{Q}}{d\mathbb{P}^0} = \log \frac{d\mathbb{P}^b}{d\mathbb{Q}} - \log \frac{d\mathbb{P}^0}{d\mathbb{Q}} \quad (24)$$

Substituting in (22) and (23) and canceling terms, we get

$$\log \frac{d\mathbb{P}^b}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T (\mathbf{b}^\top \Sigma^{-1})(\mathbf{X}_t, t) d\mathbf{X}_t - \int_0^T (\Sigma^{-1} \mathbf{f} \cdot \mathbf{b})(\mathbf{X}_t, t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}(\mathbf{X}_t, t)\|^2 dt \quad (25)$$

Equivalently, since we can write $d\mathbf{W}_t = \Sigma^{-1}(\mathbf{X}_t, t)(d\mathbf{X}_t + \mathbf{b}(\mathbf{X}_t, t)dt)$, we can write

$$\log \frac{d\mathbb{P}^b}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T \mathbf{b}(\mathbf{X}_t, t)^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}(\mathbf{X}_t, t)\|^2 dt \quad (26)$$

which concludes the proof. \square

Remark B.1. *In the underdamped Langevin model, the diffusion acts only on the velocity coordinates, so the full diffusion matrix on $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$ is degenerate, i.e., $\Sigma = \begin{bmatrix} 0 & 0 \\ 0 & \Sigma_V \end{bmatrix}$.*

We therefore apply Girsanov's theorem on the velocity SDE only, where the diffusion Σ_V is invertible, and treat positions as deterministic integrals of velocity (i.e., $\mathbf{R}_t = \mathbf{R}_0 + \int_0^t \mathbf{V}_s ds$). Therefore, Girsanov's theorem is first defined on velocity paths as

$$\log \frac{d\mathbb{P}_V^b}{d\mathbb{P}_V^0}(\mathbf{X}_{0:T}) = \int_0^T \mathbf{b}(\mathbf{X}_t, t)^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}(\mathbf{X}_t, t)\|^2 dt \quad (27)$$

and pushed forward to the full path measure via $(\mathbf{R}_t, \mathbf{V}_t) = (\mathbf{R}_0 + \int_0^t \mathbf{V}_s ds, \mathbf{V}_t)$. For simplicity in notation, we define $\Sigma^{-1} \equiv \Sigma_V^{-1}$ in the theoretical proofs.

Stochastic Optimal Control The reward-guided **stochastic optimal control (SOC)** problem aims to determine the **optimal control** $\mathbf{b}^* = \min_{\mathbf{b} \in \mathcal{U}} J(\mathbf{b}; \mathbf{x}, t)$ that minimizes the **cost functional** $J(\mathbf{b}_\theta; \mathbf{x}, t)$ with reward $r: \mathcal{X} \rightarrow \mathbb{R}$ given by

$$J(\mathbf{b}; \mathbf{x}, t) = \mathbb{E}_{p_t} \left[\int_t^T \left(\mathbf{f}(\mathbf{X}_s, s) + \frac{1}{2} \|\mathbf{b}(\mathbf{X}_s, s)\|^2 \right) ds - r(\mathbf{X}_T) \middle| \mathbf{X}_t = \mathbf{x} \right] \quad (28)$$

which measures the *cost-to-go* from the current state $\mathbf{X}_t = \mathbf{x}$ at time t under the controlled dynamics defined in (18) to the terminal state \mathbf{X}_T . The infimum of the cost functional is defined as

$$J(\mathbf{b}^*; \mathbf{x}, t) = \inf_{\mathbf{b} \in \mathcal{U}} J(\mathbf{b}; \mathbf{x}, t) =: V_t(\mathbf{x}) \quad (29)$$

which is also referred to as the *value function* $V_t(\mathbf{x})$. To derive the path integral representation of the optimal control \mathbf{b}^* , we introduce the *work functional* which computes the running cost subtracted by the terminal reward of a state path $\mathbf{X}_{0:T}$,

$$\mathcal{W}(\mathbf{X}_{t:T}, t) := \int_t^T \mathbf{f}(\mathbf{X}_s, s) ds - r(\mathbf{X}_T) \quad (30)$$

We can now define a series of statements that connect the optimal drift \mathbf{b}^* and the value function V_t .

Theorem B.1. *Let V_t be the value function and \mathbf{b}^* be the optimal control. Then, the following are true:*

- (a) *The optimal control satisfies $\mathbf{b}^*(\mathbf{x}) = -\Sigma^\top \nabla_{\mathbf{x}} V_t(\mathbf{x})$*
- (b) *The Radon-Nikodym derivative of the optimal path measure \mathbb{P}^* can be defined by taking state paths from the unconditional path measure \mathbb{P}^0 and reweighting them with the output of the work functional $\mathcal{W}(\mathbf{X}_{0:T}, 0)$.*

$$\frac{d\mathbb{P}^*}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \frac{1}{Z} e^{-\mathcal{W}(\mathbf{X}_{0:T}, 0)}, \quad Z = \mathbb{E}_{\mathbf{X}_{0:T} \sim \mathbb{P}^0} \left[e^{-\mathcal{W}(\mathbf{X}_{0:T}, 0)} \right] \quad (31)$$

- (c) *For any $(\mathbf{x}, t) \in \mathcal{X} \times [0, T]$, the value function can be written in path integral form by taking the expectation over the uncontrolled path measure \mathbb{P}^0 as*

$$V_t(\mathbf{x}, t) = -\log \mathbb{E} \left[e^{-\mathcal{W}(\mathbf{X}_{t:T}, t)} \middle| \mathbf{X}_t = \mathbf{x} \right] \quad (32)$$

- (d) *Combining (a) and (c), we can write the optimal control as*

$$\mathbf{b}^*(\mathbf{x}) = \Sigma^\top \nabla_{\mathbf{x}} \log \mathbb{E}_{\mathbf{X}_{t:T} \sim \mathbb{P}^0} \left[e^{-\mathcal{W}(\mathbf{X}_{t:T}, t)} \middle| \mathbf{X}_t = \mathbf{x} \right] \quad (33)$$

While Theorem B.1 provides a closed-form solution to the optimal control \mathbf{b}^* , it remains impractical to compute as \mathbb{P}^0 contains an infinite number of paths from any point \mathbf{x} . This problem motivates parameterizing \mathbf{b}^* with a neural network \mathbf{b}_θ . The natural objective for obtaining an accurate approximation of \mathbf{b}^* that induces the path measure \mathbb{P}^* is the KL-divergence $\mathcal{L}(\mathbb{P}^*, \mathbb{P}^{\mathbf{b}_\theta}) := D_{\text{KL}}(\mathbb{P}^{\mathbf{b}_\theta} \parallel \mathbb{P}^*)$ as its minimizer is exactly \mathbb{P}^* .

However, taking the gradient $\nabla_\theta D_{\text{KL}}(\mathbb{P}^{\mathbf{b}_\theta} \parallel \mathbb{P}^*)$ requires differentiating through the full stochastic trajectories $\mathbf{X}_{0:T} \sim \mathbb{P}^{\mathbf{b}_\theta}$ generated by the Euler-Maruyama SDE solver due to the expectation over $\mathbb{P}^{\mathbf{b}_\theta}$, resulting in a significant computational bottleneck. To overcome this bottleneck, previous work [Nüsken and Richter, 2021, Seong et al., 2025] leverages alternative path-measure objectives that do not involve expectations over $\mathbb{P}^{\mathbf{b}_\theta}$ while preserving optimality guarantees. We focus on two objectives: the **log-variance divergence** used in Seong et al. [2025] and the **cross-entropy objective**, which we formalize in the present work for solving the EntangledSB problem.

Log-Variance Objective Here, we describe the **log-variance (LV) divergence** which is introduced in Nüsken and Richter [2021] and applied to transition path sampling (TPS) in TPS-DPS [Seong et al., 2025]. Given the path measure $\mathbb{P}^{\mathbf{b}_\theta}$ generated with the parameterized bias force \mathbf{b}_θ and the target path measure \mathbb{P}^* , the log-variance divergence \mathcal{L}_{LV} is defined as

$$\mathcal{L}_{\text{LV}}(\mathbb{P}^*, \mathbb{P}^{\mathbf{b}_\theta}) := \text{Var}_{\mathbb{P}^{\mathbf{b}_\theta}} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{\mathbf{b}_\theta}} \right] = \mathbb{E}_{\mathbb{P}^{\mathbf{b}_\theta}} \left[\left(\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{\mathbf{b}_\theta}} - \mathbb{E}_{\mathbb{P}^{\mathbf{b}_\theta}} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{\mathbf{b}_\theta}} \right] \right)^2 \right] \quad (34)$$

where \mathbb{P}^v can be defined as any arbitrary path measure generated from a control \mathbf{v} . To allow simulated trajectories to be used across multiple training iterations, it is common to define \mathbf{v} as the bias force with frozen parameters $\mathbf{v} = \bar{\mathbf{b}} := \text{stopgrad}(\mathbf{b}_\theta)$. Using Lemma B.1, we define $\mathcal{L}_{\text{LV}}(\mathbb{P}^*, \mathbb{P}^{b_\theta}) = \mathbb{E}_{\mathbb{P}^v}[(\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}} - \mathbb{E}_{\mathbb{P}^v}[\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}])^2]$, where $\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}$ is defined as

$$\begin{aligned}\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}} &:= \log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} = \log \frac{d\mathbb{P}^*}{d\mathbb{P}^0} \frac{d\mathbb{P}^0}{d\mathbb{P}^{b_\theta}} \\ &= r(\mathbf{X}_T) + \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt - \int_0^T (\mathbf{b}_\theta^\top \mathbf{v})(\mathbf{X}_t) dt - \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t\end{aligned}\quad (35)$$

using similar justification as in App C.4. Since the expectation $\mathbb{E}_{\mathbb{P}^v}[\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}]$ is computationally intractable, Seong et al. [2025] introduces a scalar parameter w that is jointly optimized with θ such that $\arg \min_w \mathcal{L}_{\text{LV}}(\theta, w) = \mathbb{E}_{\mathbb{P}^v}[\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}]$. Substituting in w for $\mathbb{E}_{\mathbb{P}^v}[\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}]$ and taking the discretization over K steps, the LV objective becomes

$$\mathcal{L}_{\text{LV}}(\theta, w) = \mathbb{E}_{\mathbf{X}_{0:K} \sim \mathbb{P}^v} [(\mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:K}) - w)^2] = \mathbb{E}_{\mathbf{x}_{0:K} \sim \mathbb{P}^v} \left[\left(\log \frac{p^0(\mathbf{x}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{x}_{0:K})} - w \right)^2 \right]$$

where $\mathbf{v} = \bar{\mathbf{b}} := \text{stopgrad}(\mathbf{b}_\theta)$.

B.2 Justification for Entangled Bias Forces

Proposition B.1 (Monotone Optimality of Entangled Bias Forces). *Let \mathcal{F}_{ind} be the hypothesis class of bias forces that depend only on particle positions $\mathbf{b}_\theta^i := \mathbf{b}_\theta^i(\mathbf{R}_t)$ and \mathcal{F}_{ent} be the hypothesis class of entangled bias forces $\mathbf{b}_\theta^i := \mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$. Then, under the cross-entropy objective,*

$$\inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ent}}} \mathcal{L}_{\text{CE}}(\mathbf{b}_\theta) \leq \inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ind}}} \mathcal{L}_{\text{CE}}(\mathbf{b}_\theta) \quad (36)$$

with strict improvement when the optimal control is non-factorizable.

Proof. Let \mathbb{P}^* denote the optimal path measure that solves the EntangledSB problem from (3.1) and \mathbb{P}^{b_θ} be the path measure induced by the bias force \mathbf{b}_θ . Let \mathcal{F}_{ind} be the hypothesis class of bias forces independent of the full system velocities $\mathbf{b}_\theta^i := \mathbf{b}_\theta^i(\mathbf{R}_t)$ and \mathcal{F}_{ent} be the hypothesis class of entangled bias forces $\mathbf{b}_\theta^i := \mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$. Clearly, we have $\mathcal{F}_{\text{ind}} \subseteq \mathcal{F}_{\text{ent}}$. It follows that

$$\{\mathbb{P}^{b_\theta} : \mathbf{b}_\theta \in \mathcal{F}_{\text{ind}}\} \subseteq \{\mathbb{P}^{b_\theta} : \mathbf{b}_\theta \in \mathcal{F}_{\text{ent}}\} \quad (37)$$

Taking the infima over the KL-divergence functional $D_{\text{KL}}(\mathbb{P}^* \parallel \cdot) \geq 0$ over a larger set of functions cannot increase the value.

$$\inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ent}}} D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) \leq \inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ind}}} D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) \quad (38)$$

Furthermore, if $\mathbb{P}^* \notin \overline{\{\mathbb{P}^{b_\theta} : \mathbf{b}_\theta \in \mathcal{F}_{\text{ind}}\}}$ which denotes the closure of the set of path measures induced by \mathcal{F}_{ind} and there exists $\mathbf{b}^* \in \overline{\mathcal{F}_{\text{ent}}}$ where $\mathbb{P}^{b^*} = \mathbb{P}^*$, then we have

$$\inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ent}}} D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) = 0 \quad \text{while} \quad \inf_{\mathbf{b}_\theta \in \mathcal{F}_{\text{ind}}} D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) > 0 \quad (39)$$

yielding strict improvement. By Proposition 4.2, we have that $\mathbb{P}^{b^*} = \mathbb{P}^*$ is the unique minimizer of $D_{\text{KL}}(\mathbb{P}^* \parallel \cdot)$. \square

C Theoretical Proofs

For notational simplicity, we will drop the explicit dependence on $(\mathbf{R}_t, \mathbf{V}_t)$ and simply denote the stochastic path of the system as $(\mathbf{X}_t)_{t \in [0, T]}$.

C.1 Proof of Proposition 3.1

Lemma C.1. *The bias force \mathbf{b}^* that minimizes the SOC objective $J(\mathbf{b}_\theta; \mathbf{x}, t)$ generates the path measure \mathbb{P}^b that minimizes the KL-divergence with the optimal tilted path measure \mathbb{P}^* that satisfies*

$$\mathbb{P}^* = \frac{1}{Z} \mathbb{P}^0 e^{r(\mathbf{X}_T)}, \quad Z = \mathbb{E}_{\mathbb{P}_T^0} [e^{r(\mathbf{X}_T)}] \quad (40)$$

where \mathbb{P}^0 is the base path measure.

Proof. Let the dynamics of $\mathbf{X}_t = (\mathbf{R}_t, \mathbf{V}_t)$ under the base path measure \mathbb{P}^0 be defined as

$$d\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t)dt + \Sigma d\mathbf{W}_t \quad (41)$$

and the dynamics under the biased path measure \mathbb{P}^{b_θ} be defined as

$$d\mathbf{X}_t = [\mathbf{f}(\mathbf{X}_t) + \Sigma \mathbf{b}_\theta(\mathbf{X}_t)] dt + \Sigma d\mathbf{W}_t \quad (42)$$

By Lemma B.1, we have that the logarithm of the Radon-Nikodym derivative of the biased measure \mathbb{P}^{b_θ} with respect to the base measure \mathbb{P}^0 is given by

$$\log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T \mathbf{b}_\theta(\mathbf{X}_t)^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \quad (43)$$

Now, taking the expectation with respect to \mathbb{P}^{b_θ} , we derive the expression for the KL-divergence

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}^{b_\theta} \|\mathbb{P}^0) &= \mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0} \right] \\ &= \mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\int_0^T \mathbf{b}_\theta(\mathbf{X}_t)^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \right] \\ &= \mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\int_0^T \mathbf{b}_\theta(\mathbf{X}_t)^\top (d\mathbf{W}_t^b + \mathbf{b}_\theta(\mathbf{X}_t)dt) - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \right] \\ &= \underbrace{\mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\int_0^T \mathbf{b}_\theta(\mathbf{X}_t)^\top d\mathbf{W}_t^b \right]}_{=0} + \mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbb{P}^{b_\theta}} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \end{aligned} \quad (44)$$

Substituting this into the SOC objective in (3.1), we get

$$\begin{aligned} J(\mathbf{b}_\theta) &= \mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\int_0^T \frac{1}{2} \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt - r(\mathbf{X}_T) \right] \\ &= D_{\text{KL}}(\mathbb{P}^{b_\theta} \|\mathbb{P}^0) - \mathbb{E}_{\mathbb{P}^{b_\theta}} [r(\mathbf{X}_T)] \end{aligned} \quad (45)$$

By adding $\log Z$ on both sides, we do not change the minimizer of $J(\mathbf{b}_\theta)$ and the SOC objective becomes the KL-divergence between the controlled path measure \mathbb{P}^{b_θ} and the optimal path measure \mathbb{P}^* :

$$\begin{aligned} J(\mathbf{b}_\theta) + \log Z &= D_{\text{KL}}(\mathbb{P}^{b_\theta} \|\mathbb{P}^0) - \mathbb{E}_{\mathbb{P}^{b_\theta}} [r(\mathbf{X}_T)] + \log Z \\ &= \underbrace{\mathbb{E}_{\mathbb{P}^{b_\theta}} \left[\log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0} - r(\mathbf{X}_T) + \log Z \right]}_{D_{\text{KL}}(\mathbb{P}^{b_\theta} \|\mathbb{P}^*)} \end{aligned} \quad (46)$$

This shows that the minimizer \mathbf{b}^* to the SOC objective also generates the path measure \mathbb{P}^{b^*} that is closest in KL-divergence to the optimal path measure \mathbb{P}^* . Therefore, the optimal control solution satisfies

$$\log \frac{d\mathbb{P}^*}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = r(\mathbf{X}_T) - \log Z \iff \mathbb{P}^* = \frac{1}{Z} \mathbb{P}^0 e^{r(\mathbf{X}_T)} \quad (47)$$

which concludes our proof. \square

Proposition 3.1 (Solving EntangledSB with Stochastic Optimal Control). *We can solve the EntangledSB problem with the stochastic optimal control (SOC) objective given by*

$$\mathbf{b}^* = \arg \min_{\mathbf{b}_\theta} \mathbb{E}_{\mathbf{X}_{0:T} \sim \mathbb{P}^{b_\theta}} \left[\int_0^T \frac{1}{2} \|\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)\|^2 dt - r(\mathbf{X}_T) \right] \quad \text{s.t.} \quad (6) \quad (8)$$

where $r(\mathbf{X}_T) := \log \pi_{\mathcal{B}}(\mathbf{R}_T)$ is the terminal reward that measures the log-probability under the target distribution.

We aim to recover the optimal path measure \mathbb{P}^* that minimizes the KL divergence from the base dynamics \mathbb{P}^0 while matching the terminal distribution $\pi_{\mathcal{B}}$ defined as

$$\mathbb{P}^*(\mathbf{X}_{0:T}) = \frac{1}{Z} \mathbb{P}^0 \pi_{\mathcal{B}}(\mathbf{X}_T) \quad (48)$$

From Lemma C.1, we have that the bias force \mathbf{b}_θ that minimizes the SOC objective defined as

$$\mathbf{b}^* = \arg \min_{\mathbf{b}_\theta} \mathbb{E}_{\mathbf{X}_{0:T} \sim \mathbb{P}^{b_\theta}} \left[\int_0^T \frac{1}{2} \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt - r(\mathbf{X}_T) \right] \quad (49)$$

also generates the path measure \mathbb{P}^{b^*} that minimizes the KL-divergence $D_{\text{KL}}(\mathbb{P}^{b_\theta} \parallel \mathbb{P}^*)$ with the optimal path measure \mathbb{P}^* defined as

$$\mathbb{P}^* = \frac{1}{Z} \mathbb{P}^0 e^{r(\mathbf{X}_T)} \quad (50)$$

Therefore, setting $r(\mathbf{X}_T) := \log \pi_{\mathcal{B}}(\mathbf{X}_T)$ recovers the solution to the EntangledSB problem. \square

C.2 Proof of Proposition 4.1

Proposition 4.1 (Non-Increasing Distance from Target Distribution). *For small enough Δt , the distance from some target state $\mathbf{R}_{\mathcal{B}} \in \pi_{\mathcal{B}}$ state after an update with the bias force $\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ defined in (10) is non-increasing, such that*

$$\exists \mathbf{R}_{\mathcal{B}} \in \pi_{\mathcal{B}} \quad \text{s.t.} \quad \|\mathbf{R}_{t+\Delta t} - \mathbf{R}_{\mathcal{B}}\| \leq \|\mathbf{R}_t - \mathbf{R}_{\mathcal{B}}\| \quad (11)$$

where $\mathbf{R}_{t+\Delta t} = \mathbf{R}_t + (\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)/m_i) \Delta t$.

Proof. We aim to show that our parameterization of the bias force in (10) ensures that the distance between the current position \mathbf{R}_t and some target state $\mathbf{R}_{\mathcal{B}} \sim \pi_{\mathcal{B}}$ in the target distribution is *non-increasing*. For each atom $i \in \{1, \dots, n\}$, we define the bias force as

$$\mathbf{b}_\theta^i := \alpha_\theta^i \hat{\mathbf{s}}_i + \left(\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top \right) \mathbf{h}_\theta^i, \quad \alpha_\theta^i \geq 0, \quad \hat{\mathbf{s}}_i = \frac{\nabla_{\mathbf{r}_t^i} \log \pi_{\mathcal{B}}}{\|\nabla_{\mathbf{r}_t^i} \log \pi_{\mathcal{B}}\|} \quad (51)$$

After the update in the direction of the potential energy $-\nabla_{\mathbf{r}_t^i} U(\mathbf{R}_t)$, the Euler update to the position with time step Δt using the bias force is

$$\mathbf{r}_{t+\Delta t}^i = \mathbf{r}_t^i + \frac{\Delta t}{m_i} \left(\alpha_\theta^i \hat{\mathbf{s}}_i + \left(\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top \right) \mathbf{h}_\theta^i \right) = \mathbf{r}_t^i + \frac{\Delta t}{m_i} \mathbf{b}_\theta^i \quad (52)$$

Let $\mathbf{r}_{\mathcal{B}}^i \in \text{supp}(\pi_{\mathcal{B}})$ be a target point chosen on the ascent ray of $\hat{\mathbf{s}}_i$ given by

$$\mathbf{r}_{\mathcal{B}}^i \in \left\{ \mathbf{r}_t^i + \lambda \hat{\mathbf{s}}_i : \lambda \geq 0 \right\} \cap \text{supp}(\pi_{\mathcal{B}}) \quad (53)$$

We denote displacement to this target as $\mathbf{d}_i := \mathbf{r}_{\mathcal{B}}^i - \mathbf{r}_t^i = \rho_i \hat{\mathbf{d}}_i$ with $\hat{\mathbf{d}}_i := \frac{\mathbf{d}_i}{\|\mathbf{d}_i\|}$ and $\rho_i := \|\mathbf{d}_i\| \geq 0$. By construction of the ray, $\hat{\mathbf{d}}_i = \hat{\mathbf{s}}_i$ (i.e., \mathbf{d}_i is parallel with $\hat{\mathbf{s}}_i$). After a single step, the squared

distance is given by

$$\begin{aligned}
\|\mathbf{r}_B^i - \mathbf{r}_{t+\Delta t}^i\|^2 &= \left\| \mathbf{r}_B^i - \left(\mathbf{r}_t^i + \frac{\Delta t}{m_i} \mathbf{b}_\theta^i \right) \right\|^2 \\
&= \left\| \mathbf{d}_i - \frac{\Delta t}{m_i} \mathbf{b}_\theta^i \right\|^2 \\
&= \|\mathbf{d}_i\|^2 - \frac{2\Delta t}{m_i} \langle \mathbf{d}_i, \mathbf{b}_\theta^i \rangle + \frac{\Delta t^2}{m_i^2} \|\mathbf{b}_\theta^i\|^2
\end{aligned} \tag{54}$$

Using the decomposition of \mathbf{b}_θ^i and the fact that $\mathbf{d}_i \parallel \hat{\mathbf{s}}_i$ and $(\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top) \mathbf{h}_\theta^i \perp \hat{\mathbf{s}}_i$, we obtain

$$\langle \mathbf{d}_i, \mathbf{b}_\theta^i \rangle = \alpha_\theta^i \langle \mathbf{d}_i, \hat{\mathbf{s}}_i \rangle + \langle \mathbf{d}_i, \underbrace{(\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top) \mathbf{h}_\theta^i}_{\perp \hat{\mathbf{s}}_i} \rangle = \alpha_\theta^i \langle \rho_i \hat{\mathbf{s}}_i, \hat{\mathbf{s}}_i \rangle = \alpha_\theta^i \rho_i \geq 0 \tag{55}$$

Thus, the first-order term in (54) always decreases (or preserves) the distance, and the orthogonal correction cannot increase it at first order because it is orthogonal to \mathbf{d}_i^i . When $\Delta t \rightarrow 0$, the second-order term becomes negligible and the first-order term dominates. We determine the exact threshold for Δt that guarantees a non-increasing distance to be

$$\frac{2\Delta t}{m_i} \alpha_\theta^i \rho_i \geq \frac{\Delta t^2}{m_i^2} \|\mathbf{b}_\theta^i\|^2 \implies \Delta t \leq \frac{2m_i \alpha_\theta^i \rho_i}{\|\mathbf{b}_\theta^i\|^2} \tag{56}$$

under which the decrease due to the first-order term dominates the quadratic increase, yielding

$$\|\mathbf{r}_{t+\Delta t}^i - \mathbf{r}_B^i\| \leq \|\mathbf{r}_t^i - \mathbf{r}_B^i\| \tag{57}$$

In the continuous-time limit as $\Delta t \rightarrow 0$, the non-increasing guarantee follows immediately from (55):

$$\frac{\partial}{\partial t} \|\mathbf{r}_B^i - \mathbf{r}_t^i\|^2 = \frac{\partial}{\partial t} \|\mathbf{d}_i\|^2 = -\frac{2}{m_i} \langle \mathbf{d}_t^i, \mathbf{b}_\theta^i \rangle = -\frac{2}{m_i} \alpha_\theta^i \rho_i \leq 0 \tag{58}$$

Summing over $i = 1, \dots, n$ (or equivalently using the concatenated coordinates), we conclude that there exists a target configuration $\mathbf{R}_B \in \pi_B$ such that the distance to \mathbf{R}_B is *non-increasing* under the bias force update, which proves the claim. \square

C.3 Proof of Proposition 4.2

Proposition 4.2 (Convexity and Uniqueness of Cross-Entropy Objective). *The cross-entropy objective \mathcal{L}_{CE} is convex in \mathbb{P}^{b_θ} and there exists a unique minimizer \mathbf{b}^* that is the solution to the EntangledSOC problem in Proposition 3.1.*

Proof. First, we show that the cross-entropy objective is convex in the path measure \mathbb{P}^{b_θ} . The functional $\mathbb{P}^{b_\theta} \mapsto D_{KL}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta})$ is convex with respect to its second argument, since $x \mapsto -\log x$ is convex. Since $\mathcal{L}_{CE}(\mathbb{P}^*, \mathbb{P}^{b_\theta}) := D_{KL}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta})$ which is strictly convex in \mathbb{P}^{b_θ} , minimizing \mathcal{L}_{CE} yields a unique minimizer $\mathbb{P}^{b^*} \equiv \mathbb{P}^*$. From Proposition 3.1, we have that the unique minimizer of \mathcal{L}_{CE} is exactly the solution to the EntangledSB problem. While this does not necessarily imply convexity in the neural network parameters θ , it yields a more favorable optimization objective. \square

C.4 Proof of Proposition 4.3

Proposition 4.3 (Equivalence of Variational and Path Integral Objectives). *The cross-entropy objective can be expressed in path-integral form as*

$$\mathcal{L}_{CE}(\theta) = \mathbb{E}_{\mathbb{P}^v} [w^*(\mathbf{X}_{0:T}) \mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:T})] \tag{14}$$

$$\begin{cases} w^*(\mathbf{X}_{0:T}) = \frac{d\mathbb{P}^*}{d\mathbb{P}^v}(\mathbf{X}_{0:T}) = \frac{e^{r(\mathbf{X}_T)}}{Z} \frac{d\mathbb{P}^0}{d\mathbb{P}^v}(\mathbf{X}_{0:T}) \\ \mathcal{F}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:T}) = \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)\|^2 - \int_0^T (\mathbf{b}_\theta^\top \mathbf{v})(\mathbf{R}_t, \mathbf{V}_t) dt - \int_0^T \mathbf{b}_\theta(\mathbf{R}_t, \mathbf{V}_t)^\top d\mathbf{W}_t \end{cases} \tag{15}$$

where we define the reference measure $\mathbf{v} = \bar{\mathbf{b}} := \text{stopgrad}(\mathbf{b}_\theta)$ is the off-policy control drift from the previous iteration.

From the cross-entropy objective, we have

$$\mathcal{L}_{\text{CE}}(\mathbb{P}^*, \mathbb{P}^{b_\theta}) := D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta}) = \mathbb{E}_{\mathbb{P}^*} \left[\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] = \mathbb{E}_{\mathbb{P}^b} \left[\frac{d\mathbb{P}^*}{d\mathbb{P}^b} \log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} \right] \quad (59)$$

Setting $w^* := \frac{d\mathbb{P}^*}{d\mathbb{P}^b}$, we can derive

$$w^*(\mathbf{X}_{0:T}) = \frac{d\mathbb{P}^*}{d\mathbb{P}^b}(\mathbf{X}_{0:T}) = \frac{d\mathbb{P}^*}{d\mathbb{P}^0} \frac{d\mathbb{P}^0}{d\mathbb{P}^b}(\mathbf{X}_{0:T}) \quad (60)$$

From Lemma C.1, we have

$$w^*(\mathbf{X}_{0:T}) = \frac{e^{r(\mathbf{X}_T)}}{Z} \frac{d\mathbb{P}^0}{d\mathbb{P}^b}(\mathbf{X}_{0:T}) \quad (61)$$

which is our definition for w^* . Now, expanding $\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}}$, we have

$$\begin{aligned} \log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}} &= \log \frac{d\mathbb{P}^*}{d\mathbb{P}^0} + \log \frac{d\mathbb{P}^0}{d\mathbb{P}^{b_\theta}} \\ &= \log \frac{e^{r(\mathbf{X}_T)}}{Z} - \log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0} \\ &= r(\mathbf{X}_T) - \log Z - \log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0} \end{aligned} \quad (62)$$

Applying Girsanov's theorem from Lemma B.1, we expand the second term $\log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0}$ as

$$\log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) = \int_0^T (\mathbf{b}_\theta(\mathbf{X}_t)^\top \Sigma^{-1}) d\mathbf{X}_t - \int_0^T (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt$$

Given that $d\mathbf{X}_t$ evolves via the SDE

$$d\mathbf{X}_t = (\mathbf{f}(\mathbf{X}_t) + \Sigma \bar{\mathbf{b}}(\mathbf{X}_t)) dt + \Sigma d\mathbf{W}_t \quad (63)$$

Applying $(\mathbf{b}_\theta^\top \Sigma^{-1})$ to both sides of the equation, we have

$$\begin{aligned} (\mathbf{b}_\theta^\top \Sigma^{-1}) d\mathbf{X}_t &= (\mathbf{b}_\theta^\top \Sigma^{-1})(\mathbf{f}(\mathbf{X}_t) + \Sigma \bar{\mathbf{b}}(\mathbf{X}_t)) dt + (\mathbf{b}_\theta^\top \Sigma^{-1}) \Sigma d\mathbf{W}_t \\ &= \mathbf{b}_\theta^\top (\Sigma^{-1} \mathbf{f})(\mathbf{X}_t) dt + (\mathbf{b}_\theta^\top \Sigma^{-1} \Sigma \bar{\mathbf{b}})(\mathbf{X}_t) dt + (\mathbf{b}_\theta^\top \Sigma^{-1} \Sigma) d\mathbf{W}_t \\ &= (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt + (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \mathbf{b}_\theta^\top d\mathbf{W}_t \end{aligned} \quad (64)$$

Substituting this into (63), we get

$$\begin{aligned} \log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) &= \int_0^T (\mathbf{b}_\theta(\mathbf{X}_t)^\top \Sigma^{-1}) d\mathbf{X}_t - \int_0^T (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \\ &= \int_0^T \left((\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt + (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \mathbf{b}_\theta^\top d\mathbf{W}_t \right) \\ &\quad - \int_0^T (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \\ &= \int_0^T (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt + \int_0^T (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t \\ &\quad - \int_0^T (\Sigma^{-1} \mathbf{f}(\mathbf{X}_t))^\top \mathbf{b}_\theta(\mathbf{X}_t) dt - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \\ &= \int_0^T (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \end{aligned} \quad (65)$$

Putting everything together into the expression for the cross-entropy objective, we get

$$\begin{aligned} \mathcal{L}_{\text{CE}}(\mathbf{b}_\theta) &= \mathbb{E}_{\mathbb{P}^v} \left[w^*(\mathbf{X}_{0:T}) \left(r(\mathbf{X}_T) - \log Z - \log \frac{d\mathbb{P}^{b_\theta}}{d\mathbb{P}^0}(\mathbf{X}_{0:T}) \right) \right] \\ &= \mathbb{E}_{\mathbb{P}^v} \left[w^*(\mathbf{X}_{0:T}) \left(r(\mathbf{X}_T) - \log Z - \left[\int_0^T (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \right] \right) \right] \\ &= \mathbb{E}_{\mathbb{P}^v} \left[w^*(\mathbf{X}_{0:T}) (r(\mathbf{X}_T) - \log Z) - w^*(\mathbf{X}_{0:T}) \left(\int_0^T (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt + \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t - \frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt \right) \right] \\ &= \mathbb{E}_{\mathbb{P}^v} \left[\underbrace{w^*(\mathbf{X}_{0:T}) (r(\mathbf{X}_T) - \log Z)}_{\text{independent of } \theta} + \underbrace{w^*(\mathbf{X}_{0:T}) \left(\frac{1}{2} \int_0^T \|\mathbf{b}_\theta(\mathbf{X}_t)\|^2 dt - \int_0^T (\mathbf{b}_\theta^\top \bar{\mathbf{b}})(\mathbf{X}_t) dt - \int_0^T \mathbf{b}_\theta^\top d\mathbf{W}_t \right)}_{:= \mathcal{F}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}(\mathbf{X}_{0:T})} \right] \end{aligned}$$

Since $w^*(\mathbf{X}_{0:T}) (r(\mathbf{X}_T) - \log Z)$ is independent of θ , we can drop it and write the cross-entropy objective as

$$\mathcal{L}_{CE}(\theta) = \mathbb{E}_{\mathbb{P}^{\bar{b}}} [w^*(\mathbf{X}_{0:T}) \mathcal{F}_{\mathbf{b}_\theta, \bar{b}}(\mathbf{X}_{0:T})] \quad (66)$$

where we define $\bar{b} := \text{stopgrad}(\mathbf{b}_\theta)$. \square

C.5 Proof of Proposition 4.4

Proposition 4.4 (Discretized Cross-Entropy). *Given the discretized $\hat{\mathcal{F}}_{\mathbf{b}_\theta, \bar{b}}(\mathbf{X}_{0:K})$ in (16), we can derive a simplified loss function as*

$$\hat{\mathcal{L}}_{CE}(\theta) = \mathbb{E}_{\mathbf{X}_{0:K} \sim \mathbb{P}^{\bar{b}}} \left[\underbrace{\frac{d\mathbb{P}^*}{d\mathbb{P}^{\bar{b}}}(\mathbf{X}_{0:K}) \log \left(\frac{p^0(\mathbf{X}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{X}_{0:K})} \right)}_{w^*(\mathbf{X}_{0:K})} \right] \quad (17)$$

where $w^*(\mathbf{X}_{0:K})$ is the importance weight of the discrete time trajectory $\mathbf{X}_{0:K}$.

We will show that the discretized expression for $\hat{\mathcal{F}}_{\mathbf{b}_\theta, \hat{b}}(\mathbf{X}_{0:K})$ can be written as

$$\hat{\mathcal{F}}_{\mathbf{b}_\theta, \mathbf{v}}(\mathbf{X}_{0:K}) = \log \left(\frac{p^*(\mathbf{X}_{0:K})}{p^{b_\theta}(\mathbf{X}_{0:K})} \right) = \log \left(\frac{p^0(\mathbf{X}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{X}_{0:K})} \right) \quad (67)$$

First, we expand the expression as

$$\begin{aligned} \log \left(\frac{p^0(\mathbf{X}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{X}_{0:K})} \right) &= \log p^0(\mathbf{X}_{0:K}) + r(\mathbf{X}_K) - \log p^{b_\theta}(\mathbf{X}_{0:K}) \\ &= \sum_{k=0}^{K-1} \left[\log p^*(\mathbf{X}_{k+1} | \mathbf{X}_k) - \log p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k) \right] + r(\mathbf{X}_K) \end{aligned} \quad (68)$$

The distribution of the position \mathbf{R}_{k+1} under the uncontrolled dynamics is a Gaussian with mean \mathbf{R}_k and co-variance $\Sigma^\top \Sigma \Delta t$. The log-density becomes

$$\begin{aligned} \log p^0(\mathbf{X}_{k+1} | \mathbf{X}_k) &= \log \mathcal{N}(\mathbf{X}_{k+1} | \mathbf{X}_k, \Sigma^\top \Sigma \Delta t) \\ &= -\frac{1}{2} (\mathbf{X}_{k+1} - \mathbf{X}_k)^\top (\Sigma^\top \Sigma \Delta t)^{-1} (\mathbf{X}_{k+1} - \mathbf{X}_k) \end{aligned} \quad (69)$$

Similarly, the log-density under the bias force \mathbf{b}_θ is given by

$$\begin{aligned} \log p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k) &= \log \mathcal{N}(\mathbf{X}_{k+1} | \mathbf{X}_k + \Sigma \mathbf{b}_\theta(\mathbf{X}_k) \Delta t, \Sigma^\top \Sigma \Delta t) \\ &= -\frac{1}{2} (\mathbf{X}_{k+1} - \mathbf{X}_k - \Sigma \mathbf{b}_\theta(\mathbf{X}_k) \Delta t)^\top (\Sigma^\top \Sigma \Delta t)^{-1} (\mathbf{X}_{k+1} - \mathbf{X}_k - \Sigma \mathbf{b}_\theta(\mathbf{X}_k) \Delta t) \end{aligned} \quad (70)$$

Now, we write the increment from time k to $k+1$ as

$$\begin{aligned} \mathbf{X}_{k+1} &= \mathbf{X}_k + \Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k \\ \mathbf{X}_{k+1} - \mathbf{X}_k &= \Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k \end{aligned} \quad (71)$$

and substitute into (69) to get

$$\begin{aligned} &\log p^0(\mathbf{X}_{k+1} | \mathbf{X}_k) - \log p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k) \\ &= -\frac{1}{2} (\Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k)^\top (\Sigma^\top \Sigma \Delta t)^{-1} (\Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k) \\ &+ \frac{1}{2} (\Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k - \Sigma \mathbf{b}_\theta(\mathbf{X}_k) \Delta t)^\top (\Sigma^\top \Sigma \Delta t)^{-1} (\Sigma \bar{b}(\mathbf{X}_k) \Delta t + \Sigma \Delta \mathbf{W}_k - \Sigma \mathbf{b}_\theta(\mathbf{X}_k) \Delta t) \end{aligned}$$

Given that $(\Sigma^\top \Sigma \Delta t)^{-1} = \frac{1}{\Delta t} (\Sigma^\top \Sigma)^{-1} = \frac{1}{\Delta t} (\Sigma^{-1} \Sigma^{-\top})$, and rewriting $\mathbf{a} = \bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k - \mathbf{b}_\theta \Delta t$ and $\bar{\mathbf{a}} = \bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k$, we have

$$\begin{aligned}
& \log p^0(\mathbf{X}_{k+1} | \mathbf{X}_k) - \log p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k) \\
&= \frac{1}{2\Delta t} \left[-(\Sigma \bar{\mathbf{a}})^\top (\Sigma^{-1} \Sigma^{-\top}) (\Sigma \bar{\mathbf{a}}) + (\Sigma \mathbf{a})^\top (\Sigma^{-1} \Sigma^{-\top}) (\Sigma \mathbf{a}) \right] \\
&= \frac{1}{2\Delta t} \left[-\bar{\mathbf{a}}^\top (\Sigma^\top \Sigma^{-1} \Sigma^{-\top} \Sigma) \bar{\mathbf{a}} + \mathbf{a}^\top (\Sigma^\top \Sigma^{-1} \Sigma^{-\top} \Sigma) \mathbf{a} \right] \\
&= \frac{1}{2\Delta t} [-\|\bar{\mathbf{a}}\|^2 + \|\mathbf{a}\|^2] \\
&= \frac{1}{2\Delta t} [-\|\bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k\|^2 + \|\bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k - \mathbf{b}_\theta \Delta t\|^2] \\
&= \frac{1}{2\Delta t} [-\|\bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k\|^2 + \|\bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k\|^2 - 2(\mathbf{b}_\theta \Delta t) (\bar{\mathbf{b}} \Delta t + \Delta \mathbf{W}_k) + \|\mathbf{b}_\theta \Delta t\|^2] \\
&= \frac{1}{2\Delta t} [-2(\mathbf{b}_\theta \cdot \bar{\mathbf{b}})(\Delta t)^2 - 2(\mathbf{b}_\theta \cdot \Delta \mathbf{W}_k)\Delta t + \|\mathbf{b}_\theta\|^2(\Delta t)^2] \\
&= -(\mathbf{b}_\theta \cdot \bar{\mathbf{b}})\Delta t - (\mathbf{b}_\theta \cdot \Delta \mathbf{W}_k) + \frac{1}{2}\|\mathbf{b}_\theta\|^2\Delta t
\end{aligned} \tag{72}$$

Summing over all K steps and adding $r(\mathbf{X}_K)$, we have

$$\begin{aligned}
& \sum_{k=0}^{K-1} \left[\log p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k) - \log p^0(\mathbf{X}_{k+1} | \mathbf{X}_k) \right] + r(\mathbf{X}_K) \\
&= \underbrace{\frac{1}{2} \sum_{k=0}^{K-1} \|\mathbf{b}_\theta(\mathbf{X}_k)\|^2 \Delta t - \sum_{k=0}^{K-1} (\mathbf{b}_\theta \cdot \bar{\mathbf{b}})(\mathbf{X}_k) \Delta t - \sum_{k=0}^{K-1} \mathbf{b}_\theta(\mathbf{X}_k) \cdot \Delta \mathbf{W}_k + r(\mathbf{X}_K)}_{:= \hat{\mathcal{F}}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}(\mathbf{X}_{0:K})} \tag{73}
\end{aligned}$$

Therefore, we can write $\hat{\mathcal{F}}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}(\mathbf{X}_{0:K})$ equal to

$$\hat{\mathcal{F}}_{\mathbf{b}_\theta, \bar{\mathbf{b}}}(\mathbf{X}_{0:K}) = \sum_{k=0}^{K-1} \log \frac{p^0(\mathbf{X}_{k+1} | \mathbf{X}_k)}{p^{b_\theta}(\mathbf{X}_{k+1} | \mathbf{X}_k)} + r(\mathbf{X}_K) = \log \left(\frac{p^0(\mathbf{X}_{0:K}) \exp(r(\mathbf{X}_K))}{p^{b_\theta}(\mathbf{X}_{0:K})} \right) \tag{74}$$

where $r(\mathbf{X}_K) := \log \pi_{\mathcal{B}}(\mathbf{X}_K)$. \square

D Cell Perturbation Experiment Details

D.1 Experiment Setup

Data Processing For this experiment, we extract the cell perturbation data from the Tahoe-100M dataset consists of 50 cell lines and over 1000 different drug-dose conditions [Zhang et al., 2025b]. Specifically, we use data on a single cell line (A-549) under two drug perturbation conditions: Clonidine at $5\mu M$ and Trametinib at $5\mu M$. The initial distribution is defined as the DMSO-treated control cells. Using centroid-based sampling, we obtain balanced training sets of 1033 cells per cluster (Figure 5). We split the cells into training and validation sets with a 0.9/0.1 ratio. All final visualizations are plotted with the first two principal components, but the cell state trajectories are simulated for $d \in \{50, 100, 150\}$ for Clonidine and $d = 50$ for Trametinib.

Table 3: Training cluster cell counts for perturbation experiments.

	Clonidine		Trametinib		
	Cluster 1	Cluster 2	Cluster 1	Cluster 2	Cluster 3
Original Cell Count	1675	1033	1622	686	381

Model Architecture To learn the dependencies across the n particles, we use Transformer blocks to learn dependencies between the positions $\mathbf{R}_t \in \mathbb{R}^{n \times d}$ and velocities $\mathbf{V}_t \in \mathbb{R}^{n \times d}$ in the system.

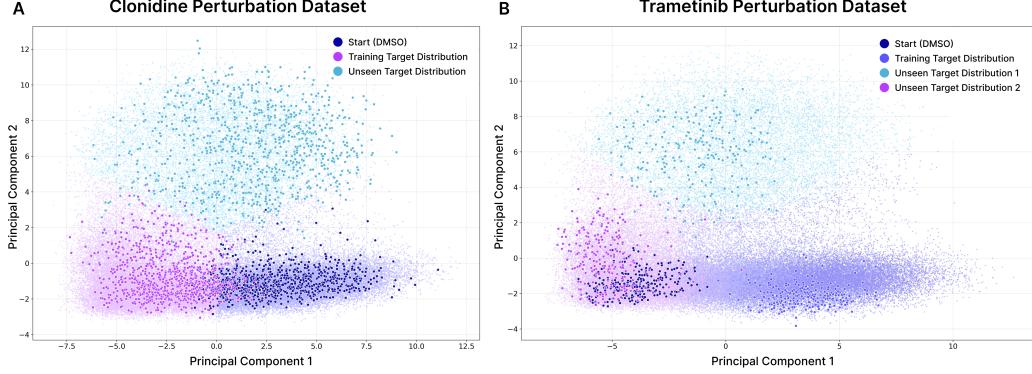


Figure 5: **Training and held-out cell clusters for the cell perturbation experiment.** All cells are plotted with the top 2 PCs. (A) Clonidine ($5\mu L$) perturbation data containing the initial DMSO-treated control cells (dark blue), perturbed cluster for training (magenta), and held-out validation cluster (turquoise). (B) Trametinib ($5\mu L$) perturbation data containing the initial DMSO-treated control cells (navy blue), the perturbed cluster for training (purple), and two held-out perturbed clusters (magenta and turquoise).

The input to the model is the concatenated position and velocity vectors $\text{Cat}(\mathbf{R}_t, \mathbf{V}_t) \in \mathbb{R}^{n \times 6}$. The architecture consists of an input projection layer that projects the input to $d_{\text{hidden}} = 256$, a 4-layer Transformer Encoder (gelu activation, feedforward dimension $d_{\text{ff}} = 512$, 8 attention heads, and dropout 0.1), and two MLP decoder heads that predict the scaling factor $\alpha_t(\mathbf{R}_t, \mathbf{V}_t)$ and correction vector $\mathbf{h}_t^i(\mathbf{R}_t, \mathbf{V}_t)$ per atom. To ensure positive scaling, the softplus activation is applied to the output of the scalar MLP head. The correction vector is projected onto the orthogonal plane of the vector pointing in the direction of the target $\hat{\mathbf{s}}_i$, and the output bias force is the sum of the scaled directional unit vector $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t)\hat{\mathbf{s}}_i$ and the projected correction vector $(\mathbf{I} - \hat{\mathbf{s}}_i\hat{\mathbf{s}}_i^\top)\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$.

Terminal Reward To solve the EntangledSB problem, we defined the terminal reward to be the log-probability under the target density $\pi_{\mathcal{B}}$ given by $r(\mathbf{X}_T) = \log \pi_{\mathcal{B}}(\mathbf{R}_T)$. We set the target density as the Gaussian centered around $\mathbf{R}_{\mathcal{B}}$ with radius σ .

$$\pi_{\mathcal{B}}(\mathbf{R}_T) = \frac{1}{(2\pi\sigma^2)^{\frac{d}{2}}} \exp\left(\frac{-\|\mathbf{R}_T - \mathbf{R}_{\mathcal{B}}\|_2^2}{2\sigma^2}\right) \quad (75)$$

Therefore, to train the bias force to reconstruct the input target state of a cell cluster $\mathbf{R}_{\mathcal{B}} \in \mathbb{R}^{n \times d}$, we define $r : \mathcal{X} \rightarrow \mathbb{R}$ as:

$$r(\mathbf{X}_T) := \frac{-\|\mathbf{R}_T - \mathbf{R}_{\mathcal{B}}\|_2^2}{2\sigma^2} \quad (76)$$

The specific values for σ are provided in Table 4.

Simulating Trajectories on the Data Manifold To define an energy landscape on the data manifold, we use the RBF metric introduced in Kapuśnian et al. [2024], which is lower in magnitude when \mathbf{x} is within the support of the data manifold and larger in magnitude when \mathbf{x} moves away from the support of the dataset into sparse regions. First, we define the elements of a function h_j^{RBF} that satisfies $h_j^{\text{RBF}}(\mathbf{x}) \approx 1$ on the data manifold as

$$h_j^{\text{RBF}}(\mathbf{x}) = \sum_{m=1}^{N_c} \omega_{m,j}(\mathbf{x}) \exp\left(-\frac{\lambda_m}{2}\|\mathbf{x} - \hat{\mathbf{x}}_m\|^2\right) \quad (77)$$

$$\lambda_m = \frac{1}{2} \left(\frac{\kappa}{|C_m|} \sum_{\mathbf{x} \in C_m} \|\mathbf{x} - \hat{\mathbf{x}}_m\|^2 \right)^{-2} \quad (78)$$

where $\{\omega_{m,j}\}_{m=1}^{N_c}$ for each cluster and each coordinate are learned given a dataset \mathcal{D} to enforce $h_j^{\text{RBF}}(\mathbf{x}) \approx 1$ for all $\mathbf{x} \in \mathcal{D}$ with the following loss function

$$\mathcal{L}_{\text{RBF}}(\{\omega_{m,j}\}) = \sum_{\mathbf{x}_i \in \mathcal{D}} (1 - h_j^{\text{RBF}}(\mathbf{x}_i))^2 \quad (79)$$

Then, we define the potential energy function of the base dynamics as

$$U(\mathbf{x}) = - \sum_{j=1}^d \log(M_j(\mathbf{x}) + \varepsilon), \quad \text{where } M_j(\mathbf{x}) = \frac{1}{(h_j^{\text{RBF}}(\mathbf{x}) + \varepsilon)^\alpha} \quad (80)$$

where $M_j(\mathbf{x})$ is small in regions of high data density and large in regions of low data density. The resulting force $-\nabla_{\mathbf{x}}U(\mathbf{x})$ is used by the simulator as the natural-gradient step.

Hyperparameters We present the hyperparameters used for the cell perturbation modeling experiment in Table 4. For each perturbation, we conducted ablations while keeping all other parameters constant on: (1) using the log-variance divergence with the learnable control variate \mathcal{L}_{LV} described in B.1 instead of the cross-entropy objective \mathcal{L}_{CE} and (2) removing the dependency on velocities as a feature input to the bias force model $\mathbf{b}_\theta(\mathbf{R}_t)$. The hyperparameters of the Transformer architecture are given in App D.1 and are kept constant across all experiments. All models are trained with the Adam optimizer [Kingma and Ba, 2014] with learning rate $\eta = 0.0001$. Due to the small batch sizes used for training, we leverage the importance weights $w^*(\mathbf{X}_{0:T})$ for categorical sampling from the replay buffer \mathcal{R} as $\mathbf{X}_{0:T} \sim \text{Cat}(\text{softmax}_{\mathcal{R}}(w^*(\mathbf{X}_{0:T})))$ to mimic the effect of the reweighting in the cross-entropy objective \mathcal{L}_{CE} over a larger batch size.

Table 4: **Hyperparameter settings for cell perturbation experiment.** The Clonidine perturbation experiment is split into three columns for each of the three dimensions of principal components (PCs) used $d \in \{50, 100, 150\}$.

Parameter	Clonidine		Trametinib
	50 PCs	100 PCs	150 PCs
number of rollouts N_{rollouts}	100	100	100
trains per rollout N_{epochs}	1000	1000	1000
step size Δt	0.01	0.01	0.01
total time steps T	100	100	100
number of samples M	64	64	64
number of particles n	16	16	16
batch size N_{batch}	64	64	64
buffer size $ \mathcal{R} $	1000	1000	1000
radius σ	0.1	0.1	0.1
friction γ	2.0	2.0	2.0
learning rate	0.0001	0.0001	0.0001
RBF N_c	150	300	300
RBF κ	1.5	2.0	3.0

D.2 Evaluation Metrics

Maximum Mean Discrepancy (RBF-MMD) We evaluate reconstruction accuracy of the target distribution and the distribution simulated with EntangledSBM using MMD with the RBF kernel (RBF-MMD) on all d principal components used during training. For Clonidine, we evaluate $d \in \{50, 100, 150\}$ and for Trametinib, we evaluate $d = 50$. Given the simulated endpoints of M paths for M cell clusters $\{\mathbf{R}_T \in \mathbb{R}^{n \times d}\}_{j=1}^M$ and the target states $\{\mathbf{R}_{\mathcal{B}} \in \mathbb{R}^{n \times d}\}_{\ell=1}^M$, the RBF-MMD is calculated as

$$\text{RBF-MMD} = \frac{1}{M^2} \sum_{j=1}^M \sum_{\ell=1}^M k_{\text{mix}}(\mathbf{R}_T^j, \mathbf{R}_{\mathcal{B}}^\ell) + \frac{1}{M^2} \sum_{j=1}^M \sum_{\ell=1}^M k_{\text{mix}}(\mathbf{R}_{\mathcal{B}}^j, \mathbf{R}_{\mathcal{B}}^\ell) - \frac{2}{M^2} \sum_{j=1}^M \sum_{\ell=1}^M k_{\text{mix}}(\mathbf{R}_T^j, \mathbf{R}_{\mathcal{B}}^\ell) \quad (81)$$

where we define the mixture of RBF kernel functions $k_{\text{mix}}(\cdot, \cdot)$ as

$$k_{\text{mix}}(\mathbf{R}, \mathbf{R}') = \frac{1}{|\Sigma|} \sum_{\sigma \in \Sigma} \exp\left(-\frac{\|\mathbf{R} - \mathbf{R}'\|^2}{2\sigma^2}\right) \quad (82)$$

for $\Sigma = \{0.01, 0.1, 1, 10, 100\}$.

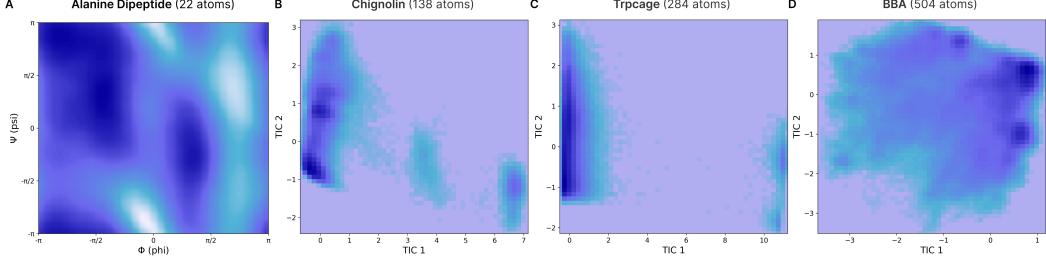


Figure 6: **Energy landscapes for transition path sampling experiments.** Potential energy plotted on dihedral angles (ϕ, ψ) for Alanine Dipeptide and the top two TICA components for fast-folding proteins.

1-Wasserstein (\mathcal{W}_1) and 2-Wasserstein (\mathcal{W}_2) Distances We further compute the \mathcal{W}_1 and \mathcal{W}_2 distances for the top two PCs of the simulated endpoints of M paths for M cell clusters $\{\mathbf{R}_T \in \mathbb{R}^{n \times d}\}_{j=1}^M$ which form the predicted distribution π_T and the *full distribution* that the target states are sampled from π_B since the \mathcal{W}_1 and \mathcal{W}_2 distances can be calculated for a pair of distributions with different sizes. Concretely, the \mathcal{W}_1 and \mathcal{W}_2 distances are calculated as

$$\mathcal{W}_1 = \left(\min_{\pi \in \Pi(\pi_T, \pi_B)} \int \|\mathbf{R}_T - \mathbf{R}_B\|_2 d\pi(\mathbf{R}_T, \mathbf{R}_B) \right) \quad (83)$$

$$\mathcal{W}_2 = \left(\min_{\pi \in \Pi(\pi_T, \pi_B)} \int \|\mathbf{R}_T - \mathbf{R}_B\|_2^2 d\pi(\mathbf{R}_T, \mathbf{R}_B) \right)^{1/2} \quad (84)$$

which quantify the minimal effort required to transform the simulated endpoint distribution into the true target distribution, demonstrating the ability of the model to capture the true perturbation dynamics.

E Transition Path Sampling Experiment Details

E.1 Experiment Setup

Model Architecture We use the same model architecture as the cell perturbation experiment described in Sec D, with the addition of Kabsch alignment [Kabsch, 1976]. We define the **aligned frame** as the frame of the target coordinates and align the input positions and velocities of the *heavy atoms* (non-hydrogen atoms) to the aligned frame. The model then predicts the optimal bias force in the aligned frame, which we transform back to the original frame of the input positions.

Terminal Reward Following the cell perturbation experiment, given the coordinates of a single target state \mathbf{R}_B , we define π_B as the log-probability under the target Gaussian centered around the coordinates of \mathbf{R}_B with radius σ :

$$r(\mathbf{X}_T) := \exp \left(\frac{-\|\mathbf{R}_T - \mathbf{R}_B\|_2^2}{2\sigma^2} \right) \quad (85)$$

The specific values for σ are provided in Table 5.

Molecular Dynamics Setup We closely follow the setup in [Seong et al., 2025]. To simulate the position and velocity via the Langevin SDEs, we use the Velocity-Verlet with Velocity Randomization (VVVR) integrator [Sivak et al., 2014] in OpenMM [Eastman et al., 2017] that updates the state with a velocity-Verlet step for the deterministic forces and a stochastic Ornstein-Uhlenbeck step that randomizes velocities. We set the friction parameter to $\gamma = 1\text{ps}^{-1}$ and the time step size to $\Delta t = 1\text{fs}$ for a total time horizon of $T = 1000\text{fs}$ for Alanine Dipeptide and $T = 5000\text{fs}$ for the fast-folding proteins. Following [Seong et al., 2025], we leverage temperature annealing with a starting temperature of $\tau_{\text{start}} = 600\text{K}$ and a final temperature of $\tau_{\text{end}} = 300\text{K}$ for Alanine Dipeptide and Chignolin and $\tau_{\text{end}} = 400\text{K}$ for the remaining fast-folding proteins.

To simulate the unconditional dynamics following the potential energy landscape $U(\mathbf{R}_t)$ visualized in Fig. 6 and the corresponding force field $F = -\nabla_{\mathbf{R}_t} U(\mathbf{R}_t)$, we use the AMBER99SB force field

with the ILDN side-chain torsion corrections (`amber99sbildn`) force field [Lindorff-Larsen et al., 2010] for Alanine Dipeptide in a vacuum and the `ff14SBonlysc` forcefield [Maier et al., 2015] paired with the `gbn2` solvation model [Nguyen et al., 2013] for the fast-folding proteins.

Hyperparameters We present the hyperparameters used for the TPS experiment in Table 5. The hyperparameters of the Transformer architecture are given in App D.1 and are kept constant across all experiments. We run the benchmark against TPS-DPS [Seong et al., 2025] with the hyperparameters below to ensure fair comparison. All models are trained with the Adam optimizer [Kingma and Ba, 2014] with learning rate $\eta = 1 \times 10^{-4}$.

Table 5: **Hyperparameter settings for transition path sampling experiment.** The setup follows that of Seong et al. [2025] to ensure fair comparison.

Parameter	Task			
	Alanine Dipeptide	Chignolin	Trp-cage	BBA
number of rollouts N_{rollouts}	100	100	100	100
trains per rollout N_{epochs}	1000	1000	1000	1000
step size Δt	1fs	1fs	1fs	1fs
total time steps T	1000	5000	5000	5000
number of samples M	64	64	64	64
number of particles n	22	138	284	504
batch size N_{batch}	16	1	1	1
buffer size $ \mathcal{R} $	1000	100	100	100
starting temperature (Kelvin)	600	600	600	600
ending temperature (Kelvin)	300	300	400	400
radius σ	0.1	0.5	0.5	0.5
friction γ	0.001	0.001	0.001	0.001
learning rate	0.0001	0.0001	0.0001	0.0001

E.2 Evaluation Metrics

We follow the evaluations of [Seong et al., 2025, Holdijk et al., 2023] and report three metrics: Root Mean Square Distance (RMSD), Target Hit Percentage (THP), and Energy of Transition State (ETS).

Root Mean Square Distance (RMSD) RMSD (\downarrow) measures the distance between the *heavy atoms* (non-hydrogen atoms) of the final position of the system at time \mathbf{R}_T and the target position $\mathbf{R}_{\mathcal{B}}$. To align the coordinate frames, we use the Kabsch algorithm [Kabsch, 1976], which determines the optimal rotation and translation such that two pairs of heavy atoms are aligned. The mean and standard deviation across the hits for M paths are reported.

Target Hit Percentage (THP) THP (\uparrow) measures the percentage of the total trajectories where the final position \mathbf{R}_T reaches the vicinity of the target state $\mathbf{R}_{\mathcal{B}}$. Specifically, we consider a hit when the two backbone dihedral angles (ϕ, ψ) for Alanine Dipeptide or the first two TICA components for the fast-folding proteins (Chignolin, Trp-cage, BBA), denoted $\xi(\mathbf{R})$, are within the 0.75-radius sphere around the target defined as $\pi_{\mathcal{B}} = \{\mathbf{R} \mid \|\xi(\mathbf{R}) - \xi(\mathbf{R}_{\mathcal{B}})\| < 0.75\}$. For M total paths, the THP is calculated as

$$\text{THP} = \frac{\sum_{i=1}^M \mathbf{1}[\mathbf{R}_T^i \in \pi_{\mathcal{B}}]}{M} \quad (86)$$

Energy of Transition State (ETS) ETS (\downarrow) measures the maximum potential energy returned by $U : \mathcal{X} \rightarrow \mathbb{R}$ along the discrete transition path $\mathbf{R}_{0:K}$ in kJmol^{-1} . It is calculated for the trajectories that reach the vicinity of the target state $\mathbf{R}_T \in \pi_{\mathcal{B}}$ and is classified as a hit by the THP metric.

$$\text{ETS}(\mathbf{R}_{0:K}) = \max_{k \in \{1, \dots, K\}} U(\mathbf{R}_k) \quad (87)$$

The mean and standard deviation across the hits for M paths are reported.

F Comparison to Log-Variance Divergence

In this section, we discuss the intuition behind the differences in cell state trajectories generated with the cross-entropy (CE) objective and log-variance (LV) objective for Clonidine, as shown in Fig. 7 and Table 6, and Trametinib as shown in Fig. 8 and Table 7.

We observe that the log-variance (LV) objective fails to accurately simulate the intermediate dynamics of cells following perturbation and generates nearly straight, abrupt trajectories connecting the initial and terminal populations. The trajectories generated from the LV-trained bias forces ignore the gradual, curved progression of cell states through intermediate data manifolds, in contrast to the trajectories generated from the cross-entropy (CE)-trained bias forces (Fig. 7, 8). While the LV-objective accurately reconstructs the target cell distribution, it fails to reconstruct the intermediate dynamics that arise from nonlinear coupling between position and velocity fields after perturbation (Tables 6, 7).

While the LV objective aims to make the log-likelihood ratio $\log \frac{d\mathbb{P}^*}{d\mathbb{P}^{b_\theta}}(\mathbf{X}_{0:T})$ constant *in expectation* to minimize the variance, it is easiest to only maximize the terminal reward $r(\mathbf{X}_T)$ rather than intermediate path alignment with \mathbb{P}^0 . In contrast, the CE objective is defined as the KL-divergence $D_{\text{KL}}(\mathbb{P}^* \parallel \mathbb{P}^{b_\theta})$ which expands into a path-integral action term $\frac{1}{2} \int_0^T \|b_\theta(\mathbf{R}_t, \mathbf{V}_t)\|^2 dt$ using Girsanov's theorem. This explicitly regularizes the whole trajectory, rewarding gradual, smooth transport through intermediate states.

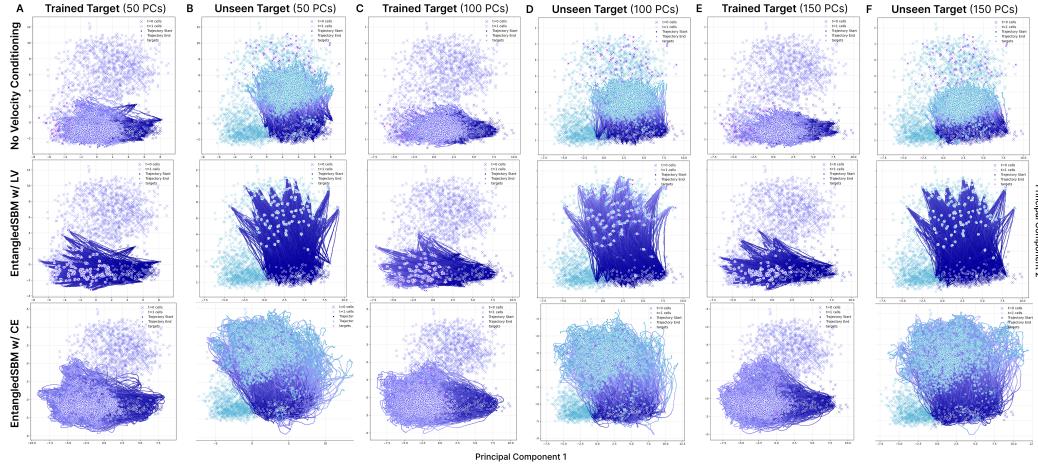


Figure 7: Full visualization of simulated cell cluster dynamics with EntangledSBM under Clonidine perturbation. Nearest neighbour cell clusters with $n = 16$ cells are simulated over 100 time steps. The gradient indicates the evolution of timesteps from the initial time $t = 0$ (navy) to the final time $t = T$ (purple or turquoise). 50 PCs simulated to (A) trained target distribution and (B) unseen target. 100 PCs simulated to (C) trained target distribution and (D) unseen target. 150 PCs simulated to (E) trained target distribution and (F) unseen target.

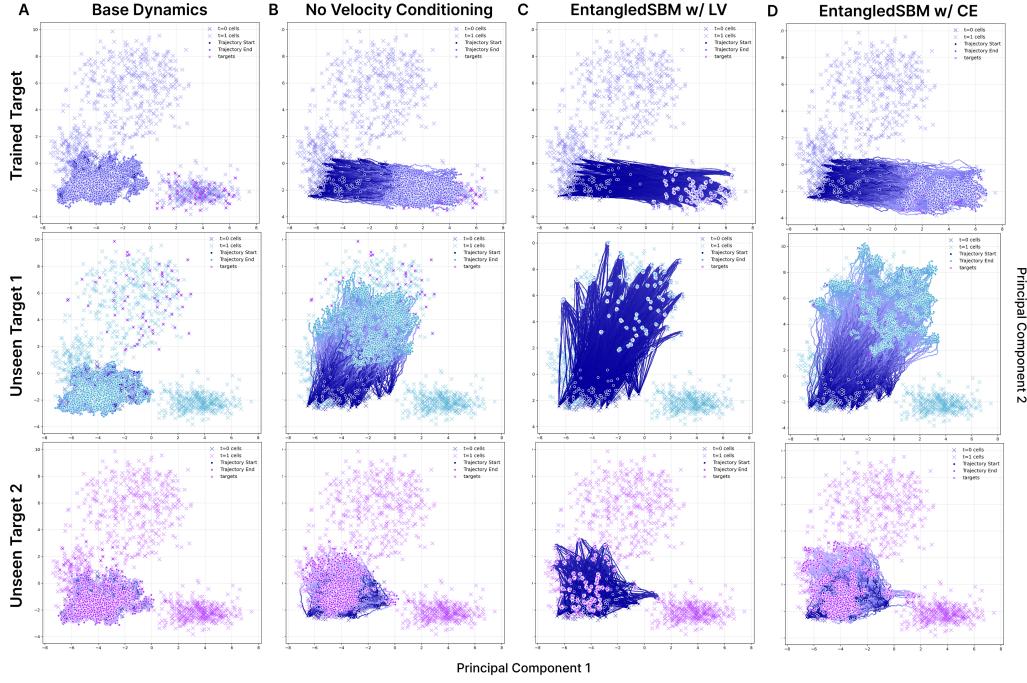


Figure 8: **Full visualization of simulated cell cluster dynamics with EntangledSBM under Trametinib perturbation.** All trajectories generated for $d = 50$ PCs. Nearest neighbour cell clusters with $n = 16$ cells are simulated over 100 time steps. The gradient indicates the evolution of timesteps from the initial time $t = 0$ (navy) to the final time $t = T$ (purple, turquoise, or magenta). **(A)** Base dynamics with no bias force. Dynamics with base and learned bias force trained with **(B)** no velocity conditioning, **(C)** log-variance (LV) objective, and **(D)** cross-entropy objective to the trained target distribution and two unseen target distributions.

Table 6: **Full comparisons for simulating cell cluster dynamics under Clonidine perturbation with EntangledSBM.** We report RBF-MMD for all d PCs with the target cluster and \mathcal{W}_1 and \mathcal{W}_2 distances of top 2 PCs against the full distribution of the perturbed cells for the seen and unseen populations after 100 simulation steps and cluster size set to $n = 16$. Mean and standard deviation of metrics from 5 independent simulations are reported. Comparisons include the base dynamics, without velocity conditioning, and EntangledSBM with the log-variance (LV) objective instead of cross-entropy (CE) for increasing principal component dimensions $d = \{50, 100, 150\}$.

Model	Seen Target Distribution			Unseen Target Distribution		
	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)
Base Dynamics (50 PCs)	0.677 ± 0.001	5.947 ± 0.005	6.015 ± 0.005	0.784 ± 0.001	8.217 ± 0.005	8.384 ± 0.005
EntangledSBM w/o Velocity Conditioning						
50 PCs	0.440 ± 0.000	1.741 ± 0.003	1.857 ± 0.004	0.478 ± 0.000	2.907 ± 0.006	3.022 ± 0.006
100 PCs	0.494 ± 0.000	2.315 ± 0.004	2.423 ± 0.004	0.539 ± 0.000	4.110 ± 0.004	4.249 ± 0.003
150 PCs	0.510 ± 0.000	2.497 ± 0.006	2.620 ± 0.006	0.560 ± 0.000	4.573 ± 0.006	4.716 ± 0.007
EntangledSBM w/ LV						
50 PCs	0.294 ± 0.000	0.205 ± 0.001	0.285 ± 0.001	0.290 ± 0.000	0.230 ± 0.001	0.382 ± 0.001
100 PCs	0.330 ± 0.000	0.277 ± 0.001	0.332 ± 0.001	0.346 ± 0.000	0.358 ± 0.001	0.492 ± 0.001
150 PCs	0.323 ± 0.000	0.128 ± 0.001	0.278 ± 0.000	0.327 ± 0.000	0.193 ± 0.001	0.425 ± 0.001
EntangledSBM w/ CE						
50 PCs	0.401 ± 0.000	0.342 ± 0.002	0.400 ± 0.001	0.419 ± 0.000	0.538 ± 0.013	0.705 ± 0.030
100 PCs	0.455 ± 0.000	0.953 ± 0.025	1.015 ± 0.025	0.500 ± 0.001	0.899 ± 0.006	1.055 ± 0.008
150 PCs	0.478 ± 0.000	0.753 ± 0.008	0.826 ± 0.007	0.506 ± 0.000	0.700 ± 0.009	0.811 ± 0.011

Table 7: **Full comparisons for simulating cell cluster dynamics under Trametinib perturbation with EntangledSBM.** We report RBF-MMD for all $d = 50$ PCs with the target cluster and \mathcal{W}_1 and \mathcal{W}_2 distances of top 2 PCs against the full distribution of the perturbed cells for the seen and unseen populations after 100 simulation steps and cluster size set to $n = 16$. Mean and standard deviation of metrics from 5 independent simulations are reported. Comparisons include the base dynamics, without velocity conditioning, and EntangledSBM with the log-variance (LV) objective instead of cross-entropy (CE).

Method	Seen Target Distribution			Unseen Target Distribution 1			Unseen Target Distribution 2		
	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)
Base Dynamics	0.938 \pm 0.001	7.637 \pm 0.005	7.653 \pm 0.006	0.900 \pm 0.000	7.766 \pm 0.009	7.877 \pm 0.009	0.754 \pm 0.001	1.201 \pm 0.009	1.455 \pm 0.012
EntangledSBM w/o Velocity Conditioning	0.449 \pm 0.000	1.506 \pm 0.005	1.544 \pm 0.005	0.476 \pm 0.000	2.116 \pm 0.005	2.197 \pm 0.005	0.480 \pm 0.000	0.505 \pm 0.004	0.627 \pm 0.005
EntangledSBM w/ LV	0.308 \pm 0.000	0.175 \pm 0.002	0.256 \pm 0.001	0.302 \pm 0.000	0.340 \pm 0.001	0.565 \pm 0.000	0.312 \pm 0.000	0.198 \pm 0.001	0.321 \pm 0.001
EntangledSBM w/ CE	0.428 \pm 0.000	0.392 \pm 0.005	0.434 \pm 0.006	0.409 \pm 0.000	0.453 \pm 0.008	0.561 \pm 0.009	0.451 \pm 0.000	0.394 \pm 0.003	0.469 \pm 0.004

G Algorithms

Here, we provide the pseudocode for training (Alg 2) and inference (Alg 3) with EntangledSBM.

Algorithm 2 EntangledSBM Training

```

1: Input: Parameterized networks  $\alpha_\theta(\mathbf{R}_t, \mathbf{V}_t) : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^n$  and  $\mathbf{h}_\theta(\mathbf{R}_t, \mathbf{V}_t) : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ , potential energy function  $U(\mathbf{R}_t) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ , distribution of target states  $\pi_B$ , buffer size  $|\mathcal{R}|$ , batch size  $N_{\text{batch}}$ , number of samples  $M$ , number of rollouts  $N_{\text{rollouts}}$ , training steps per rollout  $N_{\text{steps}}$ , number of timesteps  $T$ 
2:  $\Delta t \leftarrow \frac{1}{T}$ 
3:  $\mathcal{R} \leftarrow \{\}$   $\triangleright$  Initialize empty replay buffer
4: for rollout in  $1, \dots, N_{\text{rollouts}}$  do
5:   for  $t$  in  $1, \dots, T$  do
6:     Predict  $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \in \mathbb{R}$  and  $\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \in \mathbb{R}^d$  with parameterized neural network
       where  $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \geq 0$  is enforced with softplus activation
7:      $\mathbf{s}_i \leftarrow \nabla_{\mathbf{r}_t^i} \log \pi_B$ ,  $\hat{\mathbf{s}}_i \leftarrow \mathbf{s}_i / \|\mathbf{s}_i\|$ 
8:     Compute bias force
        $\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \leftarrow \alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \hat{\mathbf{s}}_i + (\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top) \mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ 
9:     Generate  $M$  discrete trajectories  $\{\mathbf{X}_{0:T}\}_{j=1}^M$  with current bias  $\mathbf{b}_\theta^i$  following (6) with
       Euler-Maruyama integration for each particle  $i$ 
        $\mathbf{r}_{t+1}^i = \mathbf{r}_t^i + \mathbf{v}_t^i(\mathbf{R}_t) \Delta t + \Sigma \mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \Delta t + \Sigma \epsilon_t$ 
10:    end for
11:     $\mathcal{B} \leftarrow \mathcal{B} \cup \{\mathbf{X}_{0:T}\}_{j=1}^M$   $\triangleright$  update buffer with discrete trajectories
12:    for step in  $1, \dots, N_{\text{steps}}$  do
13:      Sample batch  $\{\mathbf{X}_{0:T}\}_{j=1}^{N_{\text{batch}}}$  from buffer  $\mathcal{R}$ 
14:      Compute cross-entropy objective  $\mathcal{L}_{\text{CE}}$  with (13)
15:      Update  $\theta$  with  $\nabla_\theta \mathcal{L}_{\text{CE}}$ 
16:    end for
17:  end for
18: return parameterized  $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$  and  $\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ 

```

Algorithm 3 EntangledSBM Inference

1: **Input:** Trained networks $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ and $\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$, potential energy function $U(\mathbf{R}_t)$, initial state $\mathbf{X}_0 = (\mathbf{R}_0, \mathbf{V}_0)$, target state $\mathbf{X}_T = (\mathbf{R}_T, \mathbf{V}_T)$, time steps T , friction γ
2: $\Delta t \leftarrow \frac{1}{T}$, $\mathbf{R}_t \leftarrow \mathbf{R}_0$, $\mathbf{V}_t \leftarrow \mathbf{V}_0$
3: $\mathcal{P} \leftarrow \{\}$ \triangleright initialize path
4: **for** t in $0, \dots, T$ **do**
5: Predict $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ and $\mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$ with parameterized neural network where
 $\alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \geq 0$ is enforced with softplus activation
6: $\mathbf{s}_i \leftarrow \nabla_{\mathbf{r}_t^i} \log \pi_B$, $\hat{\mathbf{s}}_i \leftarrow \mathbf{s}_i / \|\mathbf{s}_i\|$
7: Compute bias force

$$\mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \leftarrow \alpha_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \hat{\mathbf{s}}_i + (\mathbf{I} - \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^\top) \mathbf{h}_\theta^i(\mathbf{R}_t, \mathbf{V}_t)$$

8: Generate M discrete trajectories $\{\mathbf{X}_{0:T}\}_{j=1}^M$ with current bias \mathbf{b}_θ^i following (6) with Euler-Maruyama integration for each particle i

$$\mathbf{r}_{t+1}^i = \mathbf{r}_t^i + \mathbf{v}_t^i(\mathbf{R}_t) \Delta t + \Sigma \mathbf{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t) \Delta t + \Sigma \epsilon_t$$

9: Append to path $\mathcal{P} \leftarrow \mathcal{P} \cup \{\mathbf{R}_t\}$
10: **end for**
11: **return** path \mathcal{P}
