Entangled Schrödinger Bridge Matching

Sophia Tang[1], Yinuo Zhang[2], Pranam Chatterjee[1,3 †]

[1]Department of Computer and Information Science, University of Pennsylvania
[2]Center of Computational Biology, Duke-NUS Medical School
[3]Department of Bioengineering, University of Pennsylvania

[†]Correspondence to: pranam@seas.upenn.edu

Motivation

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. Flow and Schrödinger bridge matching have been used to implicitly learn joint trajectories through data snapshots. However, **many systems undergo dynamic interactions that evolve over their trajectory and cannot be captured through static snapshots.**

Solving the Entangled Schrödinger Bridge Problem with Stochastic Optimal Control

$$\boldsymbol{b}^\star(\boldsymbol{R}_t, \boldsymbol{V}_t) = \arg\min_{\boldsymbol{b}_\theta} D_{\mathrm{KL}}\left(\mathbb{P}^{b_\theta} \| \mathbb{P}^\star\right) \quad \text{s.t.} \quad \begin{cases} \mathbb{P}^\star = \frac{1}{Z}\mathbb{P}^0 \pi_{\mathcal{B}}(\boldsymbol{X}_T) \\ \mathbb{P}_0^\star = \pi_{\mathcal{A}}(\boldsymbol{X}_0) \end{cases}$$

Entangled Schrödinger Bridge Problem

Stochastic Optimal Control with Entangled Langevin Dynamics

$$\boldsymbol{b}^\star = \arg\min_{\boldsymbol{b}_\theta} \mathbb{E}_{\boldsymbol{X}_{0:T}\sim\mathbb{P}^{b_\theta}}\left[\int_0^T \frac{1}{2}\|\boldsymbol{b}_\theta(\boldsymbol{R}_t, \boldsymbol{V}_t)\|^2 dt - r(\boldsymbol{X}_T)\right] \quad \text{s.t.} \quad d\boldsymbol{r}_t^i = \boldsymbol{v}_t^i dt, \quad d\boldsymbol{v}_t^i = \frac{-\nabla_{\boldsymbol{r}_t^i} U(\boldsymbol{R}_t) + \boldsymbol{b}^i(\boldsymbol{R}_t, \boldsymbol{V}_t)}{m_i} dt - \gamma \boldsymbol{v}_t^i dt + \sqrt{\frac{2\gamma k_B \tau}{m_i}} d\boldsymbol{W}_t^i$$

EntangledSB Problem

We formulate the **Entangled Schrödinger Bridge** (EntangledSB) **problem** that aims to parameterize a bias force that **dynamically depends on the system's positions and velocities as they evolve over time.**

$$\boldsymbol{b}^\star(\boldsymbol{R}_t, \boldsymbol{V}_t) = \arg\min_{\boldsymbol{b}_\theta} D_{\mathrm{KL}}\left(\mathbb{P}^{b_\theta} \| \mathbb{P}^\star\right)$$

Bias Force Parameterization

We introduce a novel parameterization of the bias force that can be **conditioned**, at inference time, on a target distribution or terminal state, enabling the generation of trajectories to diverse target distributions.

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

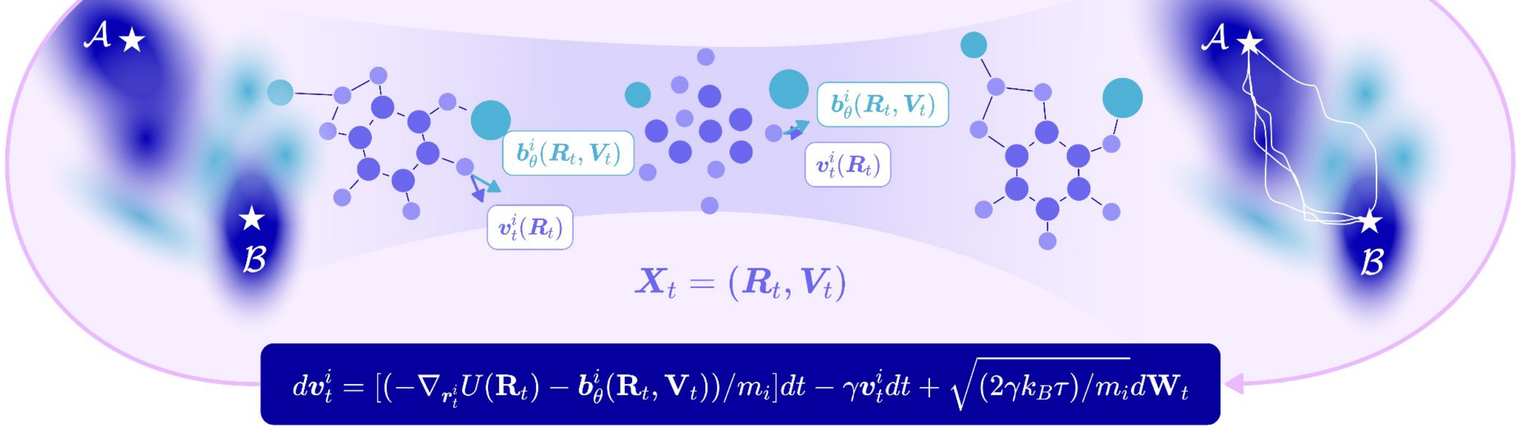

$$d\boldsymbol{v}_t^i = [(-\nabla_{\boldsymbol{r}_t^i} U(\mathbf{R}_t) - \boldsymbol{b}_\theta^i(\mathbf{R}_t, \mathbf{V}_t))/m_i]dt - \gamma \boldsymbol{v}_t^i dt + \sqrt{(2\gamma k_B \tau)/m_i} d\mathbf{W}_t$$

Transformer Architecture

We parameterize the bias force with a multi-head Transformer architecture to **learn dynamic dependencies between particles** in the system with the position and velocity as input features.

Cross-Entropy Objective

We minimize the divergence of the simulated path distribution from the optimal bridge distribution using a **weighted cross-entropy objective**.

$$\mathbb{E}_{\boldsymbol{X}_{0:K}\sim\mathbb{P}^b}\left[\frac{\mathrm{d}\mathbb{P}^\star}{\mathrm{d}\mathbb{P}^b}(\boldsymbol{X}_{0:K}) \log\left(\frac{p^0(\boldsymbol{X}_{0:K})\exp(r(\boldsymbol{X}_K))}{p^{b_\theta}(\boldsymbol{X}_{0:K})}\right)\right]$$

Framework for Learning Entangled Dynamics

Algorithm 2 Framework for **Entangled Schrödinger Bridge Matching**

Input: Parameterized velocity network $\boldsymbol{b}_\theta(\boldsymbol{R}_t, \boldsymbol{V}_t) : \mathbb{R}^{n\times d} \times \mathbb{R}^{n\times d} \to \mathbb{R}^{n\times d}$, initial and target joint distributions of the system π_0, π_T, energy function $U(\boldsymbol{R}_t) : \mathbb{R}^{n\times d} \to \mathbb{R}$, objective function in path space $\mathcal{L}(\mathbb{P}^\star, \mathbb{P}^{b_\theta})$ where $\mathbb{P}^\star = \arg\min_{\mathbb{P}^{b_\theta}} \mathcal{L}(\mathbb{P}^\star, \mathbb{P}^{b_\theta})$
while not converged **do**
  `Sample` initial positions $\boldsymbol{R}_0 \sim \pi_0$
  for $t = 0$ to T **do**
    `Evaluate` unconditional velocities $\boldsymbol{V}_t(\boldsymbol{R}_t)$
    `Compute` entangled bias force $\boldsymbol{b}_\theta(\boldsymbol{R}_t, \boldsymbol{V}_t)$
    `Simulate` step over a discrete time step Δt to get updated coordinates $\boldsymbol{R}_t$
  end for
  `Compute` reward on terminal states $r(\boldsymbol{X}_T)$
  `Evaluate` path objective $\mathcal{L}(\mathbb{P}^\star, \mathbb{P}^{b_\theta})$ on simulated paths
  `Update` θ using the gradient $\nabla_\theta \mathcal{L}(\mathbb{P}^\star, \mathbb{P}^{b_\theta})$
end while

Predicting Cell Perturbation Dynamics with Unseen Endpoints

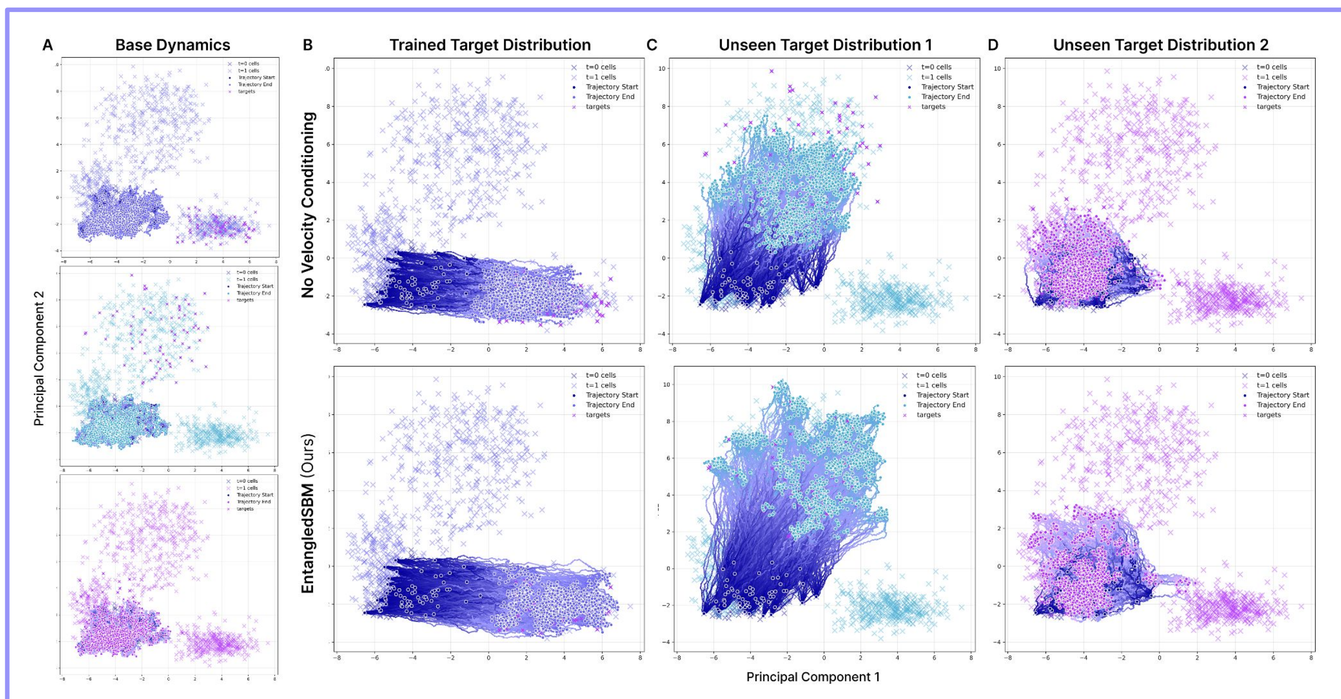

Fig 1. Comparison of EntangledSBM with and without velocity conditioning for cell cluster simulation under Trametinib perturbation. 50 PCs are simulated with the learned bias force trained with the CE objective **without velocity conditioning (Top)** and **with velocity conditioning (Bottom)** to **(B)** the perturbed population used for training and **(C, D)** the two unseen target populations.

Transition Path Sampling at All-Atom Resolution

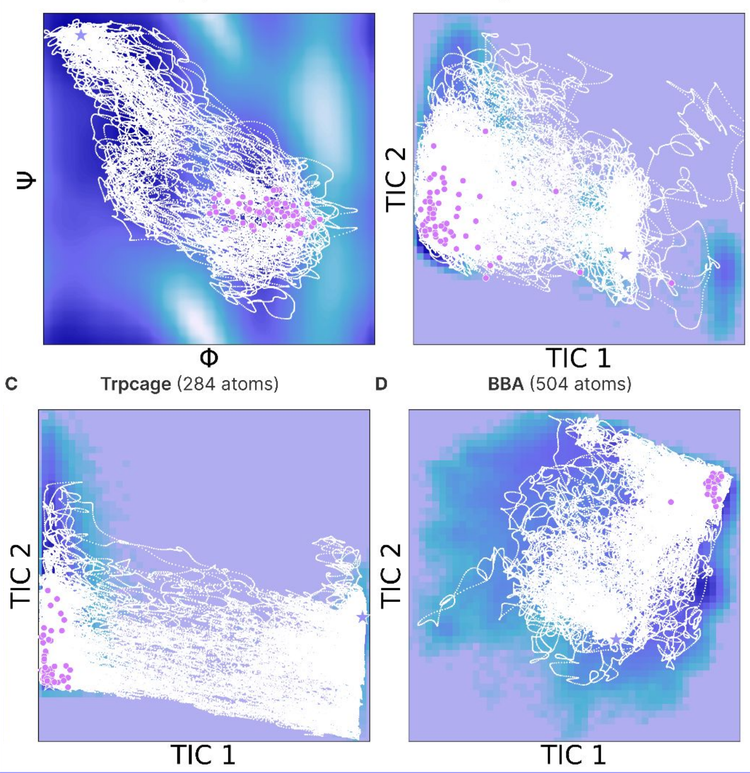

Table 1: Transition path sampling benchmarks with EntangledSBM. Best values are bolded. All metrics are averaged over 64 paths. Unless specified in brackets, paths are generated at 300K for Alanine Dipeptide and Chignolin and 400K for the others.

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

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

Conclusions

1. 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.
2. For **perturbation modeling**, EntangledSBM **reconstructs perturbed cell states while generalizing to divergent target states not seen during training**. For **molecular dynamics (MD)**, EntangledSBM generates physically plausible transition paths for fast-folding proteins at an **all-atom resolution.**

Preprint

Code