GAP: Guided Diffusion for A Priori Transition State Sampling

Hyukjun Lim[1*†] **Soojung Yang**[2*] **Lucas Pinède**[3*†] **Miguel Steiner**[2] **Yuanqi Du**[4]
Rafael Gómez-Bombarelli[2]
[1]Seoul National University [2]MIT [3]Chimie ParisTech, PSL University [4]Cornell University
rafagb@mit.edu, soojungy@mit.edu

Abstract

Transition states, the first-order saddle points on high-dimensional potential energy surfaces, govern the kinetics and thus the mechanisms of chemical reactions and conformational changes. Thus, identifying them is a central challenge in molecular simulations for physical and life sciences. Existing methods for locating them often require knowledge about the transition, such as a good initial guess of the transition pathway or reaction coordinates. We introduce ASTRA, **A** Priori **S**ampling of **TRA**nsition States with Guided Diffusion, a workflow that reframes this search problem as a direct generative task. ASTRA utilizes a score-based diffusion model trained exclusively on configurations from known metastable states, requiring no prior data from the transition region or prior information about the reaction coordinate. During inference, ASTRA guides the generative process to draw samples from the isodensity surface that separates the basis of stable states with a principled composition of conditional scores. This process is coupled with a Score-Aligned Ascent mechanism that maximizes the energy along the score-based reaction coordinate approximation, effectively collapsing the sampling onto the transition state ensemble. We validate our approach on a series of benchmarks, from 2D potentials to the high-dimensional conformational changes of alanine dipeptide, the folding of the chignolin protein, and a chemical reaction. Our results demonstrate that ASTRA not only locates transition states with high precision but also discovers competing reaction pathways in mechanistic studies.

1 Introduction

Understanding and predicting the dynamics of molecular systems, from chemical reactions to protein folding, is fundamentally dependent on characterizing their transition states (TS) Eyring [1935], Wigner [1938]. These states, defined as first-order saddle points on the potential energy surface (PES), represent the highest energy configurations along a minimum energy path and act as the kinetic bottlenecks that determine reaction rates and mechanisms. Despite their central importance, locating transition states is a notoriously difficult problem. Their short lives and scarcity make them experimentally elusive and computationally expensive to locate, requiring specialized algorithms to navigate complex, high-dimensional landscapes.

Over the past decade, there have been major advances in the development of systematic TS optimization campaigns across diverse chemical systems Dewyer and Zimmerman [2017], Simm et al. [2018], Unsleber and Reiher [2020], Baiardi et al. [2022], Ismail et al. [2022], Wen et al. [2023], Margraf et al. [2023], Steiner and Reiher [2022]. However, these methods remain computationally intensive

*Equal contribution.
†Work done at MIT.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: AI for Accelerated Materials Discovery.

and the underlying algorithms often require chemically informed initial guesses of reaction pathways or TS structures that are difficult to generalize across different chemical systems.

Data-driven transferable models can bypass this need for prior knowledge of the system, but their applicability is limited to domains where large curated datasets of known transition states Pattanaik et al. [2020], Duan et al. [2023], Kim et al. [2024] or connected minima Nam et al. [2025] are available for training. Recent alternatives that sample reaction paths with diffusion models Raja et al. [2025] rely on accurately learning the PES near the TS, limiting their utility for unexplored systems.

To address these respective limitations of classical and data-driven methods, we introduce **A** Priori **S**ampling of **TRA**nsition States with Guided Diffusion (ASTRA), a workflow that integrates generative sampling with optimization steps based on physical forces to achieve robust generalizability. Our approach combines recent advances in generative machine learning with principles from computational chemistry. The central idea is that a useful initial guess of the TS ensemble can be inferred from the learned probability distributions of the surrounding metastable states, i.e., reactants and products. ASTRA employs a single, conditionally trained score-based diffusion model that learns the data manifold of metastable states Song et al. [2020], Ho and Salimans [2022]. At inference, we guide the reverse diffusion process to sample configurations on the dividing surface where the probability of belonging to either state is equal. This is achieved through a principled composition of the conditional scores with the method introduced in Ref. 18, which we term **Score-Based Interpolation** (SBI). These samples are then refined by force-based updates on the PES, ascending along the reaction coordinate and descending along orthogonal directions, leveraging our second key insight that the reaction coordinate can be effectively approximated from the conditional diffusion scores. This process, which we call **Score-Aligned Ascent** (SAA), alleviates the need for the diffusion model to perfectly learn the PES. Crucially, this entire process is carried out *a priori*, without any knowledge or training data from the transition region itself.

Our contributions are threefold: (1) We propose a workflow for direct TS sampling that leverages guided diffusion, eliminating the need for path-finding algorithms or prior TS data. (2) We introduce the combination of Score-Based Interpolation and Score-Aligned Ascent as a principled mechanism for guiding a diffusion process to sample first-order saddle points. (3) We show that ASTRA successfully identifies known transition states and competing pathways in systems ranging from 2D potentials to high-dimensional chemical systems.

2 Method

2.1 Score-Based Interpolation of Diffusion Models

Our objective is to infer the TS, given MD data of each metastable state. First, we annotate each data point with its metastable state c and train a denoising score matching model [Song et al., 2020], with Classifier-Free Guidance (CFG) [Ho and Salimans, 2022]. The model learns to approximate both the conditional, $\nabla \log p_t(\mathbf{x}|c)$, and unconditional score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ that can be used to sample from the reverse process stochastic differential equation (SDE):

$$d\mathbf{x} = [f(\mathbf{x}, \tau) - g(\tau)^2 \nabla \log p_t(\mathbf{x}|c) d\tau + g(\tau) d\bar{\mathbf{w}}, \tag{1}$$

from a randomly initialized $\mathbf{x}_\tau \sim \mathcal{N}(0, I)$. The scores are parametrized by the CFG-based score $s^{\mathrm{c}}_\theta(\mathbf{x}, t) \approx \nabla \log p_t(\mathbf{x}|c)$, given by extrapolation:

$$s^{\mathrm{c}}_\theta(\mathbf{x}, t) = s_\theta(\mathbf{x}, t, \emptyset) + \gamma(s_\theta(\mathbf{x}, t, c) - s_\theta(\mathbf{x}, t, \emptyset)), \tag{2}$$

where c is the condition, replaced by a null token $\emptyset$ for the unconditional score, and γ is the guidance scale that controls the strength of conditioning.

Our key intuition is that the TS region can be effectively sampled by interpolating between the learned conditional scores of the two metastable states, A and B. We formulate this using a combined score function for the reverse SDE:

$$s^{\mathrm{comb}}_\theta(\mathbf{x}, t) = s^{\mathrm{B}}_\theta(\mathbf{x}, t) + \kappa(s^{\mathrm{A}}_\theta(\mathbf{x}, t) - s^{\mathrm{B}}_\theta(\mathbf{x}, t)), \tag{3}$$

where the interpolation weight κ controls the contribution of each state's score. We propose two strategies for determining κ: **Isodensity Interpolation (II)** and **Simple Averaging (SA)**. The II approach is inspired by the Superposition of Diffusion Models (SuperDiff) AND operator, specifically

designed to sample the equal density region between states [Skreta et al., 2024]. The SA approach provides a simpler baseline by directly averaging the two conditional scores, which corresponds to setting $\kappa = 0.5$ [Liu et al., 2022].

2.2 Score-Aligned Ascent for Saddle Point Search

To further refine the TS samples and push them towards the true saddle point, we propose the Score-Aligned Ascent (SAA) algorithm. Transition states are first-order saddle points on a multi-dimensional potential energy surface (PES). Locating a transition state corresponds to an optimization problem, where the energy needs to be maximized along the given transformation often referred to as the reaction coordinate $\mathbf{r}$, and minimized in all orthogonal directions. This is achieved by decomposing the negative energy gradient ($-\nabla U(\mathbf{x})$) into components parallel and perpendicular to the reaction coordinate. We then invert the parallel component to create an ascent force that pushes the system uphill along $\mathbf{r}$, while preserving the original components that minimize the energy in all orthogonal directions. This operation yields the idealized SAA force: $\mathbf{F} = -\nabla U(\mathbf{x}) + 2(\nabla U(\mathbf{x}) \cdot \mathbf{r})\mathbf{r}/\|\mathbf{r}\|_2^2$. This force projection-based optimization to find the saddle point has been introduced in Henkelman and Jónsson [1999]. One of our key findings is that we can dynamically approximate $\mathbf{r}$ at each step by taking the difference of scores conditioned on each states: $\mathbf{r}_{\text{SAA}} = s_\theta^{\text{B}}(\mathbf{x}, t) - s_\theta^{\text{A}}(\mathbf{x}, t)$. This vector naturally points from one metastable state to the other (see Appendix D).

Summarizing the workflow, we first solve the reverse-time SDE with interpolated score up to a chosen fixed timestep close to the SDE endpoint to ensure reliable force calculation, pause the reverse SDE process, and use $\mathbf{F}_{\text{SAA}} = -\nabla U(\mathbf{x}) + 2(\nabla U(\mathbf{x}) \cdot \mathbf{r}_{\text{SAA}})\mathbf{r}_{\text{SAA}}/\|\mathbf{r}_{\text{SAA}}\|_2^2$ to optimize the configurations. We then resume the reverse SDE sampling, which corrects nonphysical configurations (see Algorithm 1).

3 Experiments and Results

3.1 Validation on 2D Model Potentials

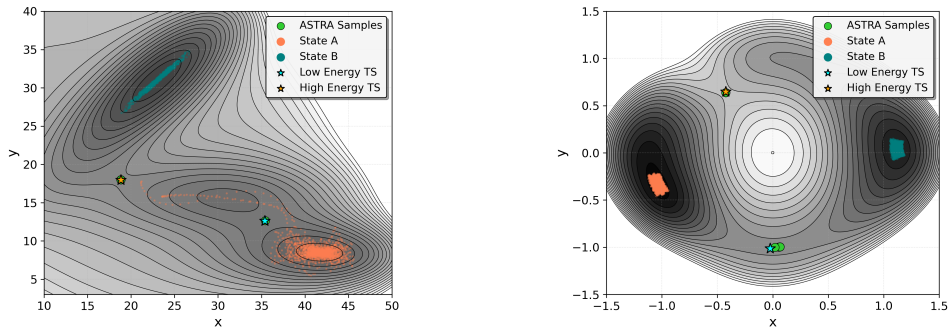

(a) Müller-Brown potential. (b) Double-path potential.

Figure 1: Our method (ASTRA) simultaneously discovers multiple transition states on two benchmark potentials, (a) Müller-Brown and (b) double-path potential.

Müller-Brown Potential. The Müller-Brown potential is a common two-dimensional benchmark for evaluating transition state (TS) search algorithms. Its energy surface features three local minima connected by two first-order saddle points, offering a controlled environment to assess a method's ability to identify distinct transition pathways. For our analysis, we assigned the two lowest-energy minima as the initial and final states (State A and State B in Figure 1a). Despite this setup, the ASTRA algorithm successfully locates configurations clustering around both accessible transition states connecting A and B. Ablation studies shown in Figure 6 indicate that this ability to discover multiple, topologically distinct saddle points is primarily driven by the Score-Aligned Ascent (SAA) component of our method.

Double-Path Potential. We next evaluated our method's robustness on a synthetic 2D potential engineered to possess two topologically distinct transition pathways connecting the same reactant and

product states. As shown in Figure 1b, the ASTRA algorithm successfully identifies and populates both TS regions corresponding to the two competing pathways. This simultaneous discovery of distinct reaction channels, achieved without any explicit path-based guidance, demonstrates a key advantage of our method. Whereas traditional algorithms are designed to converge to a single pathway, our approach can explore the entire TS manifold, thereby providing mechanistic insights into complex energy landscapes that are often inaccessible to conventional techniques or avoids being trapped in local maxima that are irrelevant for the studied chemical process.

3.2 Application to Chemical Systems

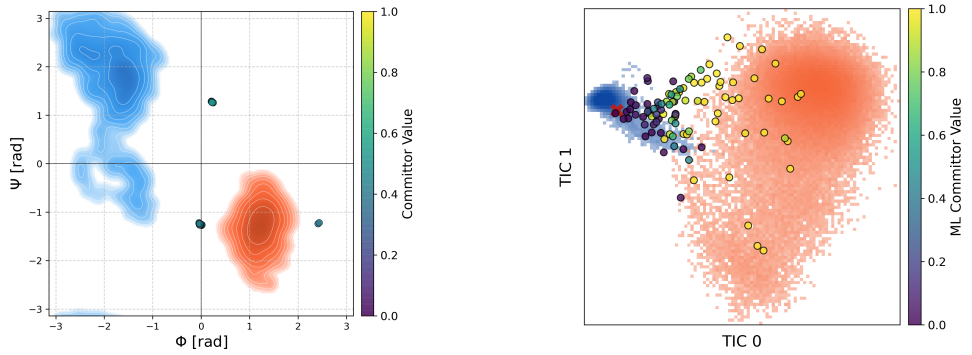

(a) Alanine dipeptide conformational change.

(b) Chignolin folding.

Figure 2: Application to chemical systems of (a) Alanine dipeptide conformational change, and (b) Chignolin folding.

Alanine Dipeptide. This system involves a 22-atom molecule implying a high dimensional configuration space. The conformational landscape of alanine dipeptide is conventionally analyzed through its Ramachandran plot, defined by the dihedral angles ϕ and ψ. We focus on the isomerization between the two conformational states, the C5 and C7ax conformations, which are separated by multiple free energy barriers [Chekmarev et al., 2004]. The configurations generated by our algorithm lie in three narrow regions on the Ramachandran plot. The agreement with reference data is validated by a narrow distribution of committor values around 0.5, calculated by running short Langevin Dynamics simulations from the samples, as shown in Figure 2a and 10. Both Figure 2a and 10 demonstrate a peaked distribution of committor values around 0.5, validating our method.

Chignolin. We conclude our experiments by studying coarse-grained chignolin, a 10-residue fast-folding protein with multiple competing folding pathways, and transition timescales spanning microseconds, using the DESRES dataset [Lindorff-Larsen et al., 2011]. To visualize the ASTRA-generated configurations we project them onto the two-dimensional space of time-lagged independent components (TICA) [Molgedey and Schuster, 1994, Pérez-Hernández et al., 2013]. We compute the committor values with an ML model trained by Kang et al. [2024]. Figure 2b and 11 demonstrates that ASTRA samples concentrate along the $q \approx 0.5$ isosurface, providing quantitative and qualitative confirmation of our method. Crucially, the samples resolve into two distinct clusters, consistent with the two known competing reaction pathways termed TS^{down} and TS^{up}, shown in Figure 14.

4 Conclusion

We introduce ASTRA, a workflow that reconceptualizes TS sampling as a direct generative modeling task by leveraging a score-based model conditioned on known metastable states. Our method composes learned distributions with Score-Based Interpolation to target directly the dividing surface of equal probability density, while Score-Aligned Ascent provides a mechanism based on physical forces to guide sampling. This approach successfully locates transition state regions with high precision and discovers competing reaction pathways in biomolecular systems, as demonstrated on alanine dipeptide and chignolin. Future work will focus on testing this workflow to larger systems with production-level force fields and navigating more complex reaction networks with multiple intermediate states.

Acknowledgments and Disclosure of Funding

This paper was supported by Korea Institute for Advancement of Technology(KIAT) grant funded by the Korea Government(Ministry of Education)(P0025681-G02P22450002201-10054408, Semiconductor-Specialized University)

References

Henry Eyring. The activated complex in chemical reactions. *The Journal of chemical physics*, 3(2): 107–115, 1935.

Eugene Wigner. The transition state method. *Transactions of the Faraday Society*, 34:29–41, 1938.

Amanda L Dewyer and Paul M Zimmerman. Finding Reaction Mechanisms, Intuitive or Otherwise. *Org. Biomol. Chem.*, 15(3):501–504, 2017. doi: 10.1039/C6OB02183B.

Gregor N Simm, Alain C Vaucher, and Markus Reiher. Exploration of Reaction Pathways and Chemical Transformation Networks. *J. Phys. Chem. A*, 2018.

Jan P. Unsleber and Markus Reiher. The Exploration of Chemical Reaction Networks. *Annu. Rev. Phys. Chem.*, 71(1):121–142, 2020. doi: 10.1146/annurev-physchem-071119-040123.

Alberto Baiardi, Stephanie A. Grimmel, Miguel Steiner, Paul L. Türtscher, Jan P. Unsleber, Thomas Weymuth, and Markus Reiher. Expansive Quantum Mechanical Exploration of Chemical Reaction Paths. *Acc. Chem. Res.*, 55(1):35–43, 2022. doi: 10.1021/acs.accounts.1c00472.

Idil Ismail, Raphael Chantreau Majerus, and Scott Habershon. Graph-Driven Reaction Discovery: Progress, Challenges, and Future Opportunities. *J. Phys. Chem. A*, 126(40):7051–7069, 2022. doi: 10.1021/acs.jpca.2c06408.

Mingjian Wen, Evan Walter Clark Spotte-Smith, Samuel M. Blau, Matthew J. McDermott, Aditi S. Krishnapriyan, and Kristin A. Persson. Chemical reaction networks and opportunities for machine learning. *Nat. Comput. Sci.*, 3(1):12–24, 2023. doi: 10.1038/s43588-022-00369-z.

Johannes T. Margraf, Hyunwook Jung, Christoph Scheurer, and Karsten Reuter. Exploring Catalytic Reaction Networks with Machine Learning. *Nat. Catal.*, 6(2):112–121, 2023. doi: 10.1038/s41929-022-00896-y.

Miguel Steiner and Markus Reiher. Autonomous reaction network exploration in homogeneous and heterogeneous catalysis. *Top. Catal.*, 65(1):6–39, 2022.

Lagnajit Pattanaik, John B Ingraham, Colin A Grambow, and William H Green. Generating transition states of isomerization reactions with deep learning. *Physical Chemistry Chemical Physics*, 22 (41):23618–23626, 2020.

Chenru Duan, Yuanqi Du, Haojun Jia, and Heather J Kulik. Accurate transition state generation with an object-aware equivariant elementary reaction diffusion model. *Nature computational science*, 3 (12):1045–1055, 2023.

Seonghwan Kim, Jeheon Woo, and Woo Youn Kim. Diffusion-based generative ai for exploring transition states from 2d molecular graphs. *Nature Communications*, 15(1):341, 2024.

Juno Nam, Miguel Steiner, Max Misterka, Soojung Yang, Avni Singhal, and Rafael Gómez-Bombarelli. Transferable learning of reaction pathways from geometric priors. *arXiv preprint arXiv:2504.15370*, 2025.

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.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

Marta Skreta, Lazar Atanackovic, Joey Bose, Alexander Tong, and Kirill Neklyudov. The superposition of diffusion models using the itô density estimator. In *The Thirteenth International Conference on Learning Representations*, 2024.

Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *European conference on computer vision*, pages 423–439. Springer, 2022.

Graeme Henkelman and Hannes Jónsson. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *J. Chem. Phys.*, 111(15):7010–7022, 1999.

Dmitriy S Chekmarev, Tateki Ishida, and Ronald M Levy. Long-time conformational transitions of alanine dipeptide in aqueous solution: Continuous and discrete-state kinetic models. *The Journal of Physical Chemistry B*, 108(50):19487–19495, 2004.

Kresten Lindorff-Larsen, Stefano Piana, Ron O Dror, and David E Shaw. How fast-folding proteins fold. *Science*, 334(6055):517–520, 2011.

Lutz Molgedey and Heinz Georg Schuster. Separation of a mixture of independent signals using time delayed correlations. *Physical review letters*, 72(23):3634, 1994.

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.

Peilin Kang, Enrico Trizio, and Michele Parrinello. Computing the committor with the committor to study the transition state ensemble. *Nature Computational Science*, 4(6):451–460, 2024.

Chenru Duan, Guan-Horng Liu, Yuanqi Du, Tianrong Chen, Qiyuan Zhao, Haojun Jia, Carla P Gomes, Evangelos A Theodorou, and Heather J Kulik. Optimal transport for generating transition states in chemical reactions. *Nature Machine Intelligence*, 7(4):615–626, 2025.

Haibo Wang, Yuxuan Qiu, Yanze Wang, Rob Brekelmans, and Yuanqi Du. Generalized flow matching for transition dynamics modeling. *arXiv preprint arXiv:2410.15128*, 2024.

Yilun Du, Conor Durkan, Robin Strudel, Joshua B Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, and Will Sussman Grathwohl. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc. In *International conference on machine learning*, pages 8489–8510. PMLR, 2023.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correctors in diffusion: Annealing, guidance, and product of experts. In *Forty-second International Conference on Machine Learning*, 2025.

Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.

Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *International Conference on Machine Learning*, pages 22825–22855. PMLR, 2023.

Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John P Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. *Advances in Neural Information Processing Systems*, 36:31372–31403, 2023.

Jiajun He, José Miguel Hernández-Lobato, Yuanqi Du, and Francisco Vargas. Rne: a plug-and-play framework for diffusion density estimation and inference-time control. *arXiv preprint arXiv:2506.05668*, 2025.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. pmlr, 2015.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

RA Olsen, GJ Kroes, G Henkelman, A Arnaldsson, and H Jónsson. Comparison of methods for finding saddle points without knowledge of the final states. *J. Chem. Phys.*, 121(20):9776–9792, 2004.

Andreas Heyden, Alexis T Bell, and Frerich J Keil. Efficient methods for finding transition states in chemical reactions: Comparison of improved dimer method and partitioned rational function optimization method. *The Journal of chemical physics*, 123(22), 2005.

Johannes Kästner and Paul Sherwood. Superlinearly converging dimer method for transition state search. *The Journal of chemical physics*, 128(1), 2008.

Cheng Shang and Zhi-Pan Liu. Constrained broyden minimization combined with the dimer method for locating transition state of complex reactions. *Journal of Chemical Theory and Computation*, 6 (4):1136–1144, 2010.

Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023.

Marloes Arts, Victor Garcia Satorras, Chin-Wei Huang, Daniel Zugner, Marco Federici, Cecilia Clementi, Frank Noé, Robert Pinsler, and Rianne van den Berg. Two for one: Diffusion models and force fields for coarse-grained molecular dynamics. *Journal of Chemical Theory and Computation*, 19(18):6151–6159, 2023.

Jim Lawrence, Javier Bernal, and Christoph Witzgall. A purely algebraic justification of the kabsch-umeyama algorithm. *Journal of research of the National Institute of Standards and Technology*, 124:1, 2019.

Xiaolei Zhu, Keiran C Thompson, and Todd J Martínez. Geodesic interpolation for reaction pathways. *The Journal of Chemical Physics*, 150(16), 2019.

A Related Works

A.1 Machine Learning for Transition State Search

Numerous machine learning approaches have been developed to study rare events and in particular transition search. Early works leveraged a dataset of accumulated transition states to train generative models [Pattanaik et al., 2020] with diffusion models [Duan et al., 2023, Kim et al., 2024] and flow matching [Duan et al., 2025]. However, to alleviate the requirement of creating a transition state dataset, later works aimed to leverage geometric optimization and energy evaluations during training to find transition paths and extract transition states from it [Nam et al., 2025]. Wang et al. [2024] also leveraged the metric induced by local molecular dynamics simulation to learn a generalized flow matching model, while Raja et al. [2025] relied on the learned forcefield from pre-trained diffusion models with the Onsager–Machlup functional to approximate the transition path.

A.2 Inference-time Control for Diffusion Models

Diffusion models excel at generating high-quality data, spanning images, texts to molecular structures. Beyond sampling from the learned distribution, they can be guided to generate from modified target distributions that encode design objectives, constraints, or rewards. The early and well-known examples are classifier guidance and classifier-free guidance to sample from conditional distributions [Ho and Salimans, 2022]. Later on, inference-time control has been extended to compositional, reward-tilted, annealed, equal density distributions and more [Du et al., 2023, Skreta et al., 2024, 2025]. One popular branch of methods develop upon the heuristics and approximate guidance [Chung et al., 2022]. Another branch of methods study exact guidance where expensive Monte Carlo (MC) estimations are required [Lu et al., 2023]. Recently, sequential Monte Carlo methods have become relevant to reweight on path space to reduce the variance of inference-time control [Wu et al., 2023, Skreta et al., 2025, He et al., 2025].

B Background

B.1 Score-Based Diffusion Models

Given a forward process that transforms a data distribution $p_0(\mathbf{x}_0)$ over a continuous time variable $t \in [0, T]$ into a known prior distribution $p_T(\mathbf{x}_T)$, typically the standard Normal distribution $\mathcal{N}(0, I)$, a diffusion model $s_\theta(\mathbf{x}, t)$ can be trained such that a reverse process transforms samples from the prior distribution to the data distribution [Sohl-Dickstein et al., 2015]. The model can learn the added noise in each time step t [Ho et al., 2020] or the gradient of the log-probability of the data distribution $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$, the so-called score function [Song and Ermon, 2019]. Both training strategies can be related as discretizing the same continuous reverse process described by the stochastic differential equation (SDE)

$$d\mathbf{x} = [f(\mathbf{x}, \tau) - g(\tau)^2 \nabla_{\mathbf{x}} \log p_\tau(\mathbf{x})]d\tau + g(\tau)d\bar{\mathbf{w}}, \tag{4}$$

where $f(\mathbf{x}, t)$ is the drift coefficient, $g(t)$ is the diffusion coefficient, $\bar{\mathbf{w}}$ is a reverse Wiener process, and the reverse time $\tau \in [T, 0]$ [Song et al., 2020]. Once the model $s_\theta(\mathbf{x}, t)$ is trained, it can be applied as a generative model by initializing the reverse SDE with noise $\mathbf{x}_\tau \sim \mathcal{N}(0, I)$ and solving it numerically.

B.2 Classifier-Free Guidance

Classifier-Free Guidance (CFG) [Ho and Salimans, 2022] is a technique for conditional generation that trains a single diffusion model to handle both conditional and unconditional generation scenarios. The model is trained to predict the score of the data distribution, conditioned on the label $c \in \{C_1, C_2, ..., \emptyset\}$. During training, the conditioning information is randomly dropped with a given probability p_{drop}. This technique enables the single model to learn both the conditional score functions, $\nabla \log p_t(\mathbf{x}|c)$, for each label, and the unconditional score function, $\nabla \log p_t(\mathbf{x})$, over the entire data manifold. During inference, the CFG-based score $s^{\text{c}}_\theta(\mathbf{x}, t)$ is given by interpolation

$$s^{\text{c}}_\theta(\mathbf{x}, t) = s_\theta(\mathbf{x}, t, \emptyset) + \gamma(s_\theta(\mathbf{x}, t, c) - s_\theta(\mathbf{x}, t, \emptyset)), \tag{5}$$

where $s_\theta(\mathbf{x}, t, c)$ is the conditional score, $s_\theta(\mathbf{x}, t, \emptyset)$ is the unconditional score, where the condition c is replaced by a null token $\emptyset$, and γ is the guidance scale that controls the strength of conditioning.

This approach provides fine-grained control over conditioning strength without requiring separately trained conditional models.

B.3 Superposition of Diffusion Models

The Superposition of Diffusion Models (SuperDiff) [Skreta et al., 2024] provides a principled way to combine the score functions from multiple models, each trained on a different data distribution. Given two score models, s_A and s_B, trained on distributions p_A and p_B respectively, a composite score function can be constructed by a weighted average:

$$s_\theta^{\text{comb}}(\mathbf{x}, t) = s_\theta^{\text{B}}(\mathbf{x}, t) + \kappa(s_\theta^{\text{A}}(\mathbf{x}, t) - s_\theta^{\text{B}}(\mathbf{x}, t)). \tag{6}$$

The interpolation factor $\kappa(t)$ determines the nature of the composition, enabling the generation of samples from distributions that represent logical combinations of the base distributions. For example, setting $\kappa(t)$ based on the relative likelihoods of a sample under each model can produce samples from the union (OR) of the distributions, while setting κ based on equal likelihoods can produce samples from the intersection (AND). In the context of transition state sampling, we are interested in the latter and a $\kappa(t)$ that satisfies the equal density:

$$d \log p_t^{\text{A}}(\mathbf{x}) = d \log p_t^{\text{B}}(\mathbf{x}) \tag{7}$$

can be found by solving a set of three linear equations [Skreta et al., 2024]. As a comparison, we also include in our ablations Simple Averaging of the scores to interpolate and steer the diffusion process [Liu et al., 2022].

B.4 Local mode maximization

Transition states are first-order saddle points on a multi-dimensional potential energy surface (PES). Localizing a transition state corresponds to an optimization problem, where the energy needs to be maximized along the given transformation and minimized in all orthogonal directions. If the mode of transformation $\mathbf{r}$, commonly referred to as reaction coordinate, is known, a gradient-based optimization to the transition state can be achieved with a sufficiently close start condition. The optimization is driven by the step $\mathbf{\Delta x}_i$ that is determined by the gradient of the potential energy with respect to nuclear coordinates $\mathbf{g}(\mathbf{x})$

$$\mathbf{\Delta x}_i = -\mathbf{g}(\mathbf{x}) + 2\left(\mathbf{g}(\mathbf{x})^T \mathbf{r}\right)\mathbf{r}. \tag{8}$$

However, the definition of $\mathbf{r}$ is non-trivial for a given chemical system. Existing approaches, such as the Dimer algorithm Henkelman and Jónsson [1999], Olsen et al. [2004], Heyden et al. [2005], Kästner and Sherwood [2008], Shang and Liu [2010], find local approximations to $\mathbf{r}$ by locating the direction of least curvature with gradient calculations for multiple close-lying configurations on the PES.

C Additional Details on Experiments

For all systems, we train a score-based diffusion model on a dataset containing metastable state configurations, and apply our ASTRA algorithm to identify the TS ensemble. For chemical systems, we report the committor function value (q) which gives the probability that a trajectory starting from a given state will reach one metastable state before reaching another metastable state.

1D PES Example. The 1D PES example is a simple potential with two distinct states. The analytical formula for this potential energy surface is given by:

$$V(x) = 0.1x^4 - x^2 + 0.1x + 1.0 \tag{9}$$

Müller-Brown Potential. For the Müller-Brown potential, our methodology adheres to the training procedure detailed in Raja et al. [2025]. A score-based diffusion model was trained on the 'tiny' subset of the provided dataset, which consists of 4,000 samples. The two metastable states, A and B, are defined by the coordinate criteria y < 20 and y > 20, respectively.

The analytical surface of this system is defined as:

$$V(x,y) = \sum_{i=1}^{4} A_i \exp\left[\alpha_i(x-a_i)^2 + \beta_i(x-a_i)(y-b_i) + \gamma_i(y-b_i)^2\right] \tag{10}$$

with the parameters taken directly from the simulation code:

$$\begin{aligned} A &= 10 \times \text{barrier} \times [-1.73, -0.87, -1.47, 0.13] \\ \alpha &= 10^{-2} \times [-0.39, -0.39, -2.54, 0.273] \\ a &= [48, 32, 24, 16] \\ \beta &= 10^{-2} \times [0, 0, 4.30, 0.23] \\ b &= [8, 16, 32, 24] \\ \gamma &= 10^{-2} \times [-3.91, -3.91, -2.54, 0.273] \end{aligned}$$

where 'barrier' is a scaling factor for the height of the potential barrier, for which we used 1.0.

Double-Path Example. For the double-path potential, the dataset was generated via Langevin Dynamics simulations, following the procedure outlined in Raja et al. [2025]. A score-based diffusion model was subsequently trained on a curated subset of these simulation data, with 4,000 samples. State A and B are defined by the regions x < 0 and x > 0, respectively.

The analytical formula for the potential energy surface is given by:

$$\begin{aligned} V(x,y) = 10 \Bigg(& 2 + \frac{4}{3}x^4 - 2y^2 + y^4 + \frac{10}{3}x^2(y^2-1) \Bigg) \\ & + 7 \exp\left(-\frac{(x+0.7)^2 + (y-0.8)^2}{0.4^2}\right) \\ & + \exp\left(-\frac{(x-1.0)^2 + (y+0.3)^2}{0.4^2}\right) \\ & - 6 \exp\left(-\frac{(x+1.0)^2 + (y+0.6)^2}{0.4^2}\right) \end{aligned} \tag{11}$$

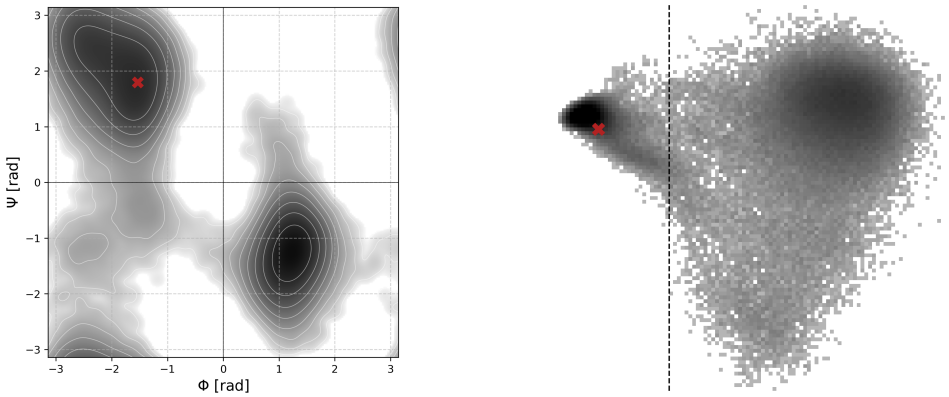

(a) Alanine dipeptide. (b) Chignolin.

Figure 3: Potential Energy Surface (PES) of alanine dipeptide and chignolin.

Alanine Dipeptide. We generate the dataset by running Langevin Dynamics from each state (C5 and C7ax) for 500 ps with a 1 fs time step saving every 10 steps. We then slice parts of the trajectory that covers the widest region in Ramachandran space (ϕ, ψ) and subsample randomly 2.5k configurations for each of the two states. In Figure 3a, we show the PES of alanine dipeptide. We propagate the dynamics using a custom force field taken from Raja et al. [2025] at T=300K. We then train a diffusion model based on the EquiformerV2 architecture [Liao et al., 2023] with up

to L=2 representations, four attention heads and 64 channels. The radius graph is computed with $r_{cutoff} = 5.0$ Angstrom. We use the AdamW optimizer with a constant learning rate 6.10^{-4}, 0.001 weight decay. We use an 0.999 Exponential Moving Average (EMA) decay with an effective batch size of 128, and train for 500 epochs. For diffusion, we leverage the Denoising Diffusion Probabilistic Model framework (DDPM) with 1000 time steps and a cosine schedule [Ho et al., 2020]. For SAA, we use 1000 optimization steps by Adam optimizer with a learning rate of 0.01, $\beta_1 = 0.0$, and $\beta_2 = 0.999$.

The committor probabilities are computed from 100 configurations sampled from the method. We first run a 2 ps long simulation and filter out all the configurations leading to instability that are labelled as non physical. For all reported results, more than 98% of samples were kept and used for the committor calculations. We compute the committor probabilities by running 100 replicas stochastic Langevin Dynamics with randomly initialized velocities using the same force field that generated the dataset to stay consistent with the PES distribution. The simulations are run for 1 ps with 1 fs at T=300K. We define each state limits and record the region first reach for each replica to compute the probability of reaching one state or the other.

Chignolin. For our analysis of Chignolin, which is one of the fast-folding proteins, we employed the DRSRES coarse-grained (CG) dataset [Lindorff-Larsen et al., 2011]. In Figure 3b, we show the PES of chignolin training dataset. We removed all the conformations with ML committor value between 0.0001 and 0.9999 in the training dataset. Our conditional diffusion model is an adaptation of the pre-trained architecture from Arts et al. [2023]. We finetuned this base model to enable conditional generation by incorporating a group embedding that specifies the state, i.e., folded or unfolded, at the beginning of the forward process. Sampling was then performed using SAA hyperparameters with 100 optimization steps, a learning rate of 0.0005, and a pause ratio of 0.05. Finally, the generated configurations were validated by computing their committor probabilities with the machine-learned committor predictor model from Kang et al. [2024].

D Validation of Score Difference Approximation to Reaction Coordinate

To validate the claim that the score difference approximates the reaction coordinate, i.e., $s_{\mathrm{A}}(\mathbf{x}, t) - s_{\mathrm{B}}(\mathbf{x}, t) \approx \mathbf{r}$, we compare the two vectors in Figure 4. The visualization shows strong alignment between our approximation (blue arrows) and the true reaction coordinate (red arrows). This is quantified by the high average cosine similarity: 0.9275 on the Müller-Brown potential and 0.9584 on the Double-Path example.

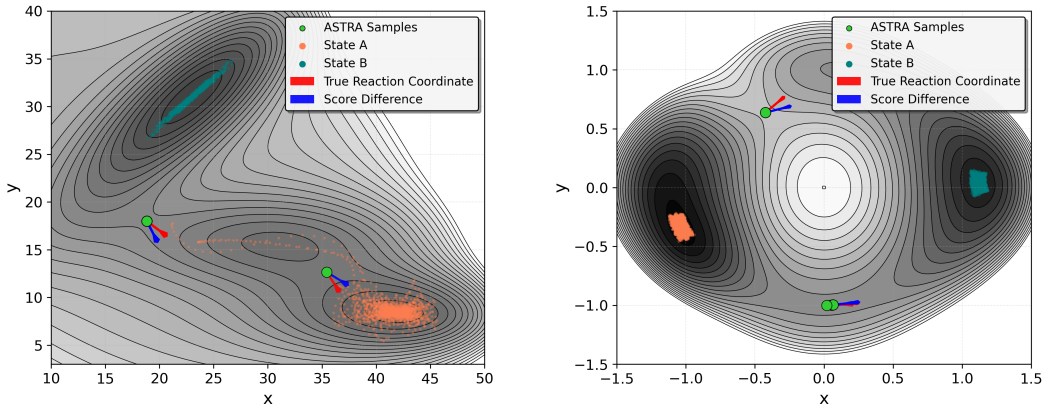

(a) Müller-Brown Potential.

(b) Double-Path Example.

Figure 4: Validation of Score Difference Approximation. The score difference vector (red) closely aligns with the true reaction coordinate (blue) on (a) the Müller-Brown potential and (b) a Double-Path example.

E ASTRA Sampling Algorithm

Algorithm 1 Sampling with Score-Based Interpolation and Score-Aligned Ascent

Require: Denoising model $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, c)$ with classifier-free guidance; Energy-based force field $\mathbf{F}(\mathbf{x})$; Total number of time steps T; Guidance pause time step T_{pause}; Number of optimization steps N_{opt}; Optimizer hyperparameters Θ_{opt}.
Output: Sampled molecular conformation $\mathbf{x}_0$.

function SBI_COMBINE($\mathbf{x}_t, t$)
  Compute conditional noise estimates: $\boldsymbol{\epsilon}_t^A \leftarrow \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, A)$, $\boldsymbol{\epsilon}_t^B \leftarrow \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, B)$.
  Compute the interpolation factor: $\kappa_t \leftarrow \text{SBI}(\boldsymbol{\epsilon}_t^A, \boldsymbol{\epsilon}_t^B, \mathbf{x}_t, t)$.
  Combine noise estimates: $\boldsymbol{\epsilon}_{\text{comb}} \leftarrow \boldsymbol{\epsilon}_t^B + \kappa_t(\boldsymbol{\epsilon}_t^A - \boldsymbol{\epsilon}_t^B)$.
  return $\boldsymbol{\epsilon}_{\text{comb}}, \boldsymbol{\epsilon}_t^A, \boldsymbol{\epsilon}_t^B$
end function

Initialize positions $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$.
// Phase 1: Denoising with Score-Based Interpolation
for $t \leftarrow T, T_{\text{pause}} + 1$ **do**
  $\boldsymbol{\epsilon}_{\text{comb}}, \_, \_ \leftarrow \text{SBI\_Combine}(\mathbf{x}_t, t)$
  Perform one reverse diffusion step to obtain $\mathbf{x}_{t-1}$ from $\mathbf{x}_t$ and $\boldsymbol{\epsilon}_{\text{comb}}$.
end for

// Phase 2: Score-Aligned Ascent Optimization at latent time T_{pause}
Let $\mathbf{z} \leftarrow \mathbf{x}_{T_{\text{pause}}}$. Initialize an optimizer $\mathcal{O}$ with Θ_{opt}.
for $k \leftarrow 1, N_{\text{opt}}$ **do**
  $\boldsymbol{\epsilon}_{\text{comb}}, \boldsymbol{\epsilon}_{T_{\text{pause}}}^A, \boldsymbol{\epsilon}_{T_{\text{pause}}}^B \leftarrow \text{SBI\_Combine}(\mathbf{x}_{T_{\text{pause}}}, T_{\text{pause}})$
  Approximate the reaction coordinate: $\boldsymbol{r}_{\text{SAA}} \leftarrow \boldsymbol{\epsilon}_{T_{\text{pause}}}^B - \boldsymbol{\epsilon}_{T_{\text{pause}}}^A$.
  Predict the clean sample: $\hat{\mathbf{x}}_0 \leftarrow \text{predict\_x}_0(\mathbf{z}, T_{\text{pause}}, \boldsymbol{\epsilon}_{\text{comb}})$.
  Evaluate the force from the external field: $\mathbf{f} \leftarrow \mathbf{F}\mathbf{F}(\hat{\mathbf{x}}_0)$.
  Compute the Score-Aligned Ascent force: $\mathbf{F}_{\text{SAA}} \leftarrow \mathbf{f} - 2\frac{\mathbf{f} \cdot \boldsymbol{r}_{\text{SAA}}}{\|\boldsymbol{r}_{\text{SAA}}\|_2^2} \boldsymbol{r}_{\text{SAA}}$.
  Update positions with gradient descent: $\mathbf{z} \leftarrow \mathcal{O}(\mathbf{z}, \nabla_{\mathbf{z}}\mathcal{L})$, where $\nabla_{\mathbf{z}}\mathcal{L} = -\mathbf{f}_{\text{SAA}}$.
end for
$\mathbf{x}_{T_{\text{pause}}} \leftarrow \mathbf{z}$.

// Phase 3: Resumed Denoising
for $t \leftarrow T_{\text{pause}}, 1$ **do**
  $\boldsymbol{\epsilon}_{\text{comb}}, \_, \_ \leftarrow \text{SBI\_Combine}(\mathbf{x}_t, t)$
  Perform one reverse diffusion step to obtain $\mathbf{x}_{t-1}$ from $\mathbf{x}_t$ and $\boldsymbol{\epsilon}_{\text{comb}}$.
end for
return $\mathbf{x}_0$.

F Additional Figures for Ablation Study

F.1 Effect of Score-Based Interpolation and Score-Aligned Ascent

F.1.1 1D PES Example

We first investigate the effects of Score-Based Interpolation and Score-Aligned Ascent on a one-dimensional potential energy surface, with the results presented in Figure 5. When employing simple averaging for interpolation, it fails to capture the true transition state, instead averaging points to a higher energy region. In contrast, isodensity interpolation generates samples that more accurately close to the true transition state. The introduction of score-aligned ascent further refines these results. For both interpolation methods, score-aligned ascent pushes the diffusion process towards TS regions, correcting deviations and ensuring the final states closely align with the true TS. Notably, the combination of isodensity interpolation and score-aligned ascent yields the most accurate and physically meaningful transition path.

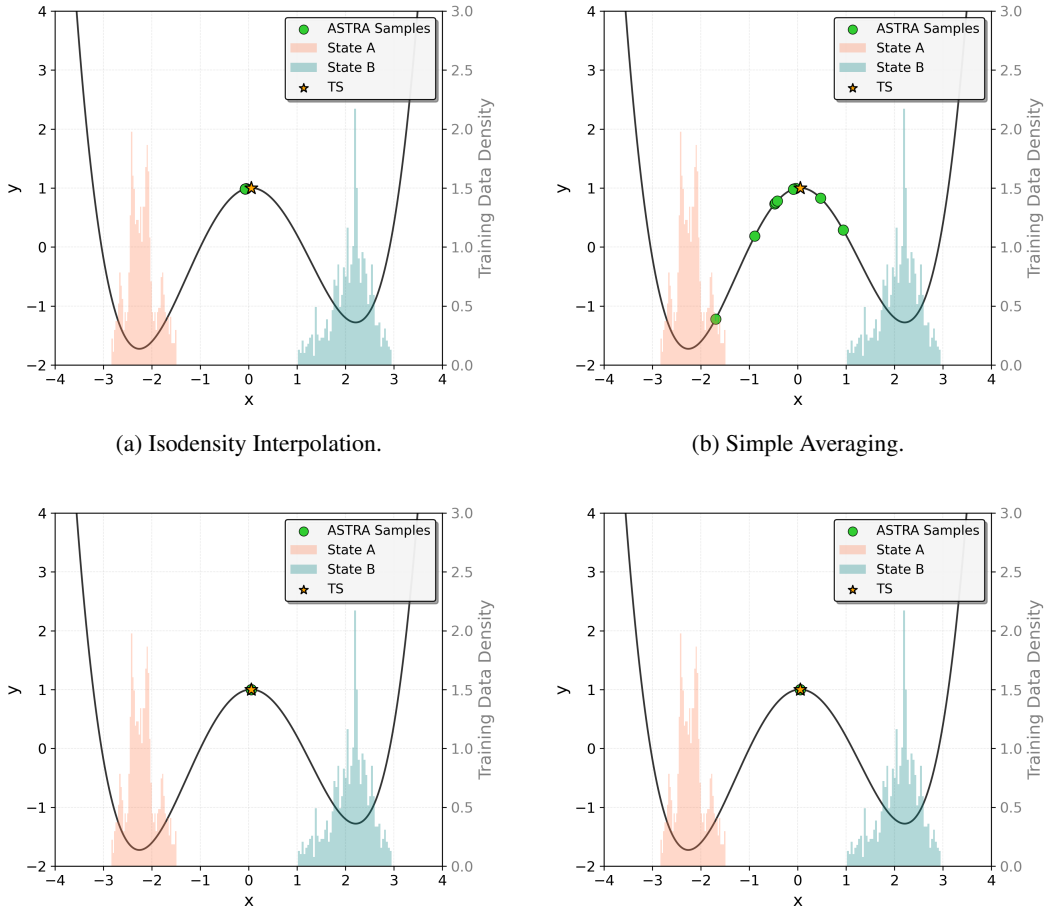

(a) Isodensity Interpolation.

(b) Simple Averaging.

(c) Isodensity Interpolation + Score-Aligned Ascent.

(d) Simple Averaging + Score-Aligned Ascent.

Figure 5: Impact of Score-Based Interpolation and Score-Aligned Ascent on 1D PES Example.

F.1.2 2D PES Example

We perform the ablation study to two-dimensional potential energy surfaces of Müller-Brown potential and Double-path potential to evaluate the distinct contributions of the Score-Based Interpolation (SBI) and Score-Aligned Ascent (SAA) components. On both systems, as shown in Figure 6 and 7, using Isodensity interpolation draws a line between two states, while Simple Averaging (SA) alone being less effective, producing a more scattered distribution of samples, with several located in nonphysical high-energy regions. Neither interpolation method on its own is sufficient for precise TS localization. The introduction of the Score-Aligned Ascent proves crucial, which refines the samples, collapsing the broad distributions onto the precise locations of the low- and high-energy transition states.

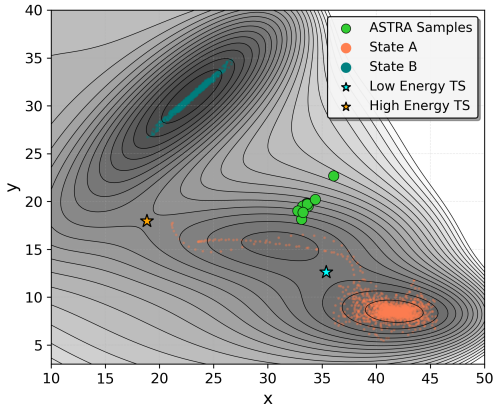

(a) Isodensity Interpolation.

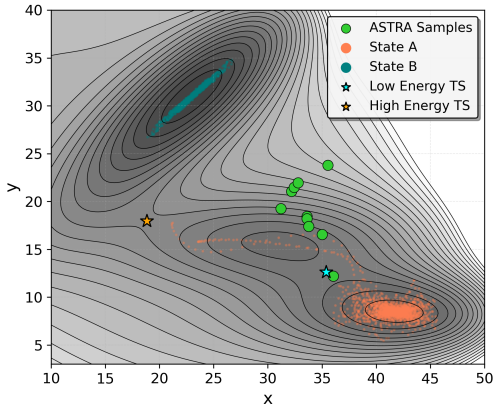

(b) Simple Averaging.

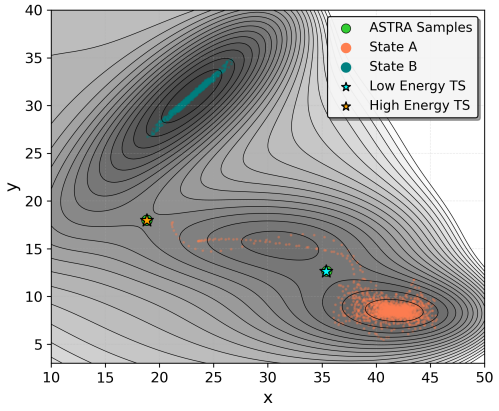

(c) Isodensity Interpolation + Score-Aligned Ascent.

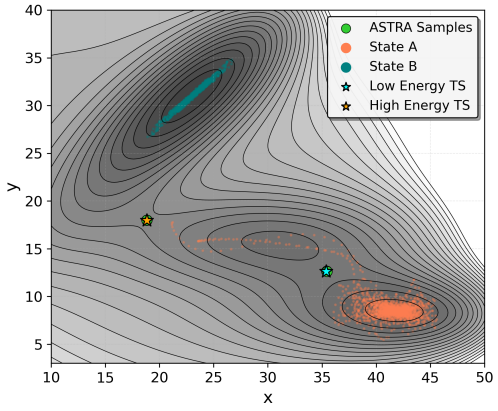

(d) Simple Averaging + Score-Aligned Ascent.

Figure 6: Impact of Score-Based Interpolation and Score-Aligned Ascent on Müller-Brown potential.

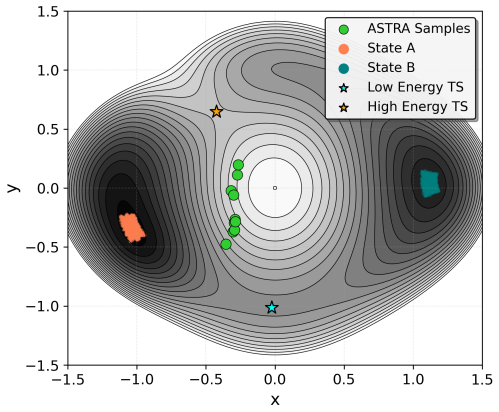

(a) Isodensity Interpolation.

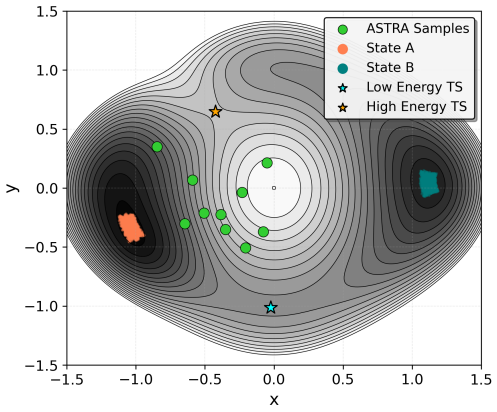

(b) Simple Averaging.

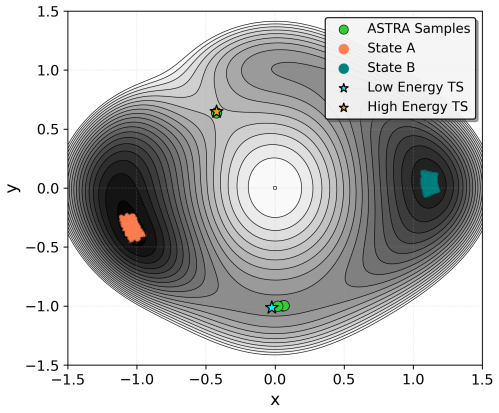

(c) Isodensity Interpolation + Score-Aligned Ascent.

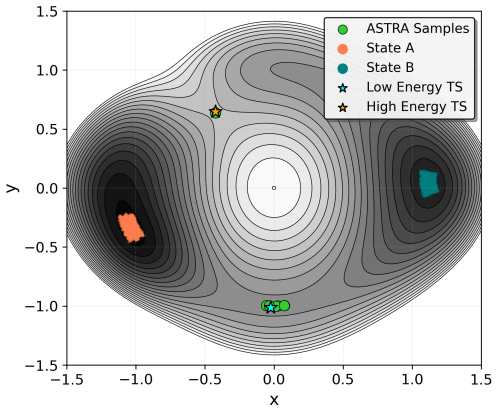

(d) Simple Averaging + Score-Aligned Ascent.

Figure 7: Impact of Score-Based Interpolation and Score-Aligned Ascent on Double-path potential.

F.1.3 Chemical Systems

We further extend our ablation study to chemical systems, alanine dipeptide and chignolin, to assess the performance of our method in high-dimensional cases.

For alanine dipeptide, as shown in Figure 8, the results highlight the critical role of score-aligned ascent. Isodensity Interpolation and Simple Averaging, when used alone, incorrectly identify a region of very high potential energy as the transition state. Upon introducing score-aligned ascent, both methods are improved, which drives generated samples are successfully guided away from high-energy regions and converge precisely onto the known, physically meaningful transition states based on the PES.

A similar trend is observed for the folding of chignolin, depicted in Figure 9. While both interpolation methods can generate a coarse path between the folded and unfolded states, the resulting samples are diffuse and do not clearly define the transition pathway. The addition of score-aligned ascent is essential for refining this pathway, driving the scattered points near-TS region. This demonstrates that the combination of isodensity interpolation and score-aligned ascent is robust and effective for identifying transition states in complex biomolecular systems.

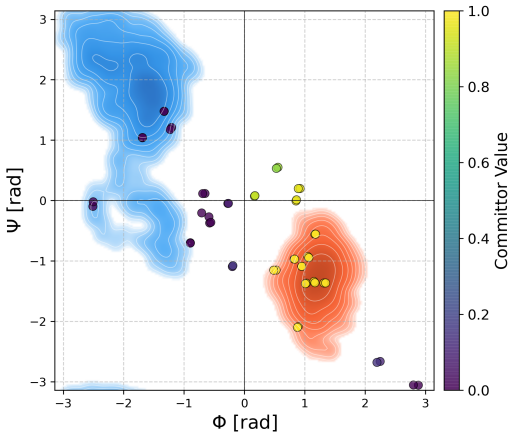

(a) Isodensity Interpolation.

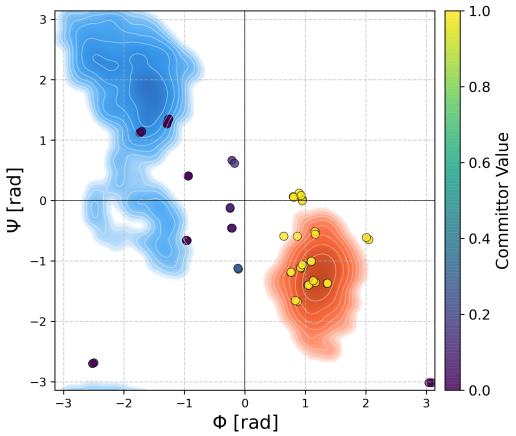

(b) Simple Averaging.

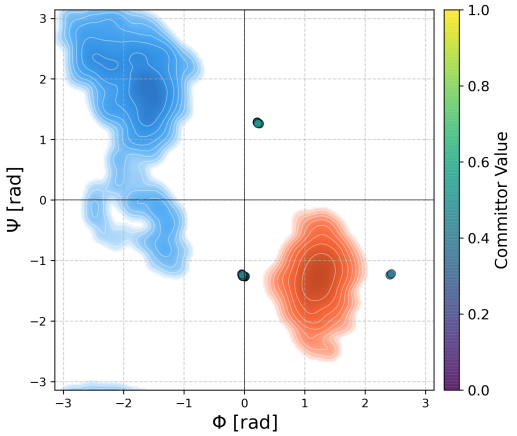

(c) Isodensity Interpolation + Score-Aligned Ascent.

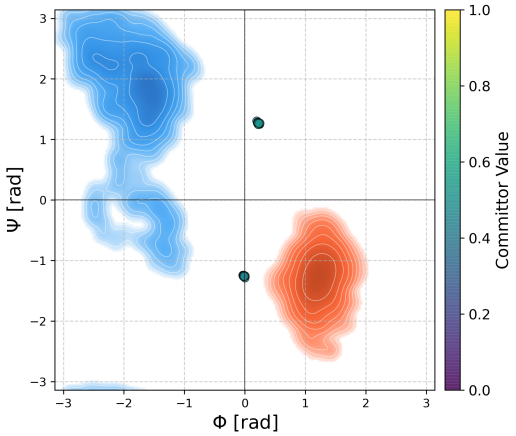

(d) Simple Averaging + Score-Aligned Ascent.

Figure 8: Impact of Score-Based Interpolation and Score-Aligned Ascent on alanine dipeptide.

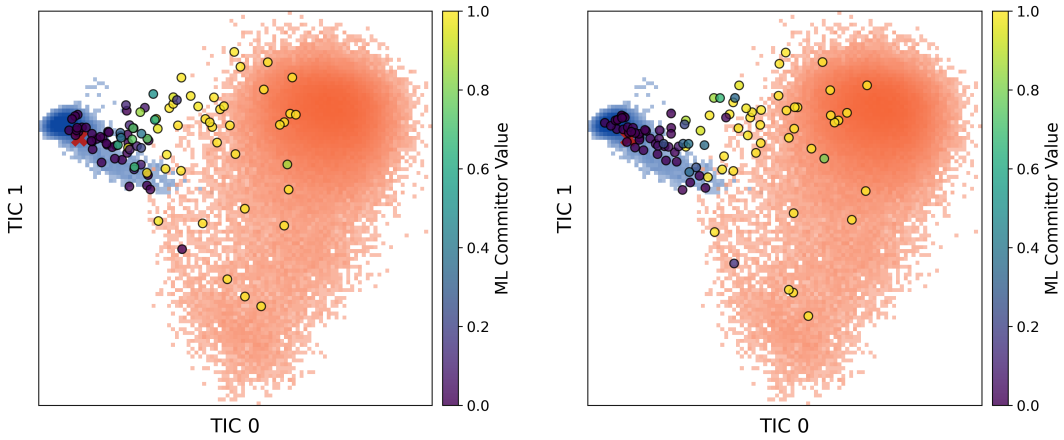

(a) Isodensity Interpolation.

(b) Simple Averaging.

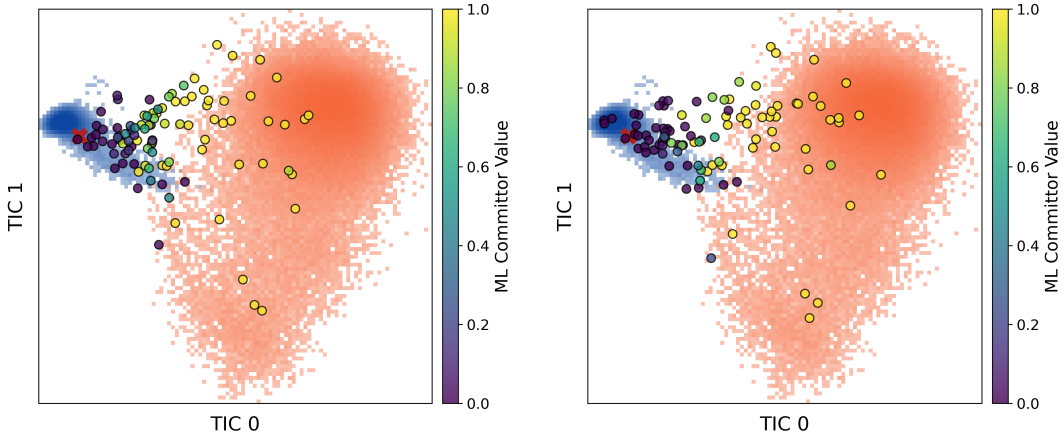

(c) Isodensity Interpolation + Score-Aligned Ascent.

(d) Simple Averaging + Score-Aligned Ascent.

Figure 9: Impact of Score-Based Interpolation and Score-Aligned Ascent on Chignolin.

F.2 Committor Analysis of Chemical Systems

F.2.1 Alanine Dipeptide

In Figure 10, we report histograms of the distributions of committor values calculated based on the samples in Figure 8. We observe that given the partial coverage of the metastable states by our MD simulations at 300K, simply using Score-Based Interpolation allows to sample an intermediate region but not the TSs themselves. The SAA algorithm successfully directs the samples towards the true TSs as demonstrated by a peaked distribution of the committor values around 0.5. SuperDiff AND surpasses SA in our experiments.

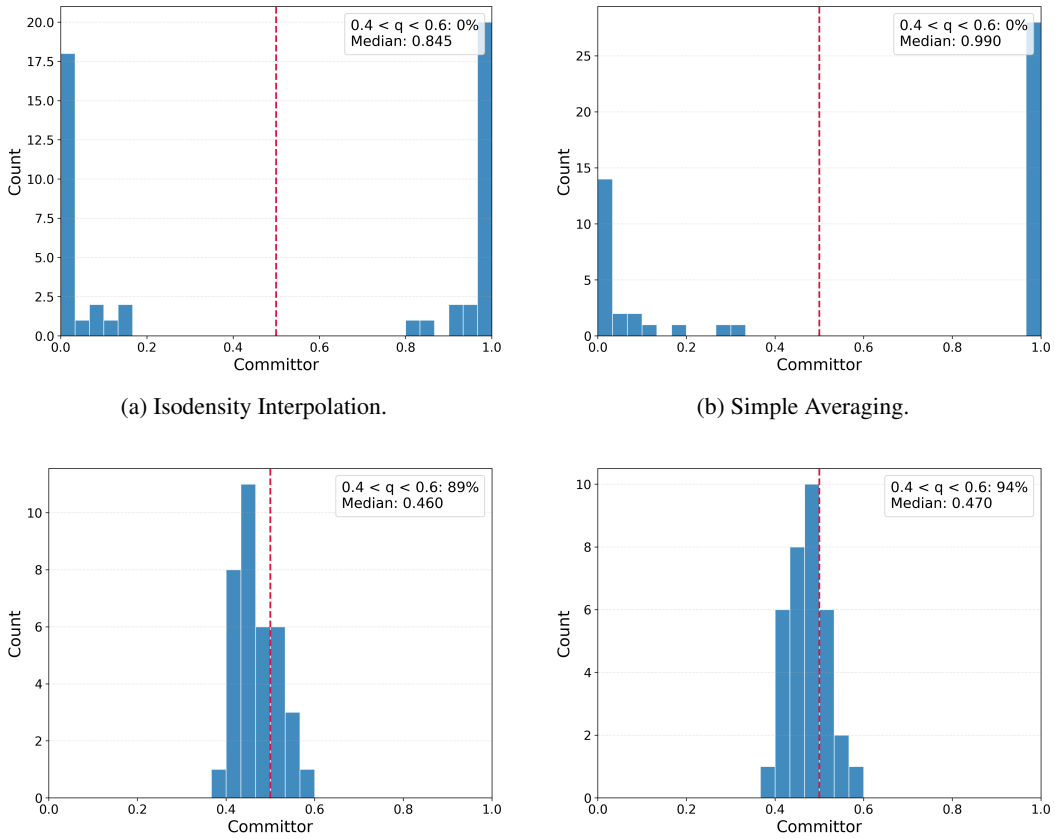

(a) Isodensity Interpolation.

(b) Simple Averaging.

(c) Isodensity Interpolation + Score-Aligned Ascent.

(d) Simple Averaging + Score-Aligned Ascent.

Figure 10: Histogram of calculated committors on alanine dipeptide.

F.2.2 Chignolin

We report ML committor values computed with the trained models from [Kang et al., 2024] to compare the effects of SBI and SAA algorithm. We observe that SAA consistently increases the number of samples close to the $q = 0.5$ region. We note that the ML committor model draws much sharper isocommittor surfaces compared to our MD evaluation for alanine dipeptide. We refer to F.3 for all TS sampling capabilities of our method applied to this system.

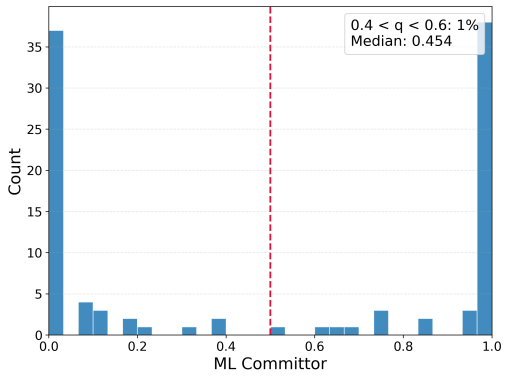

(a) Isodensity Interpolation.

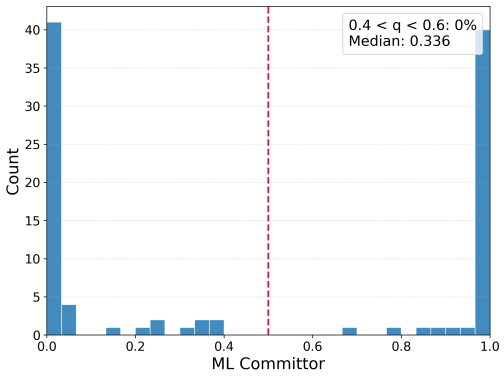

(b) Simple Averaging.

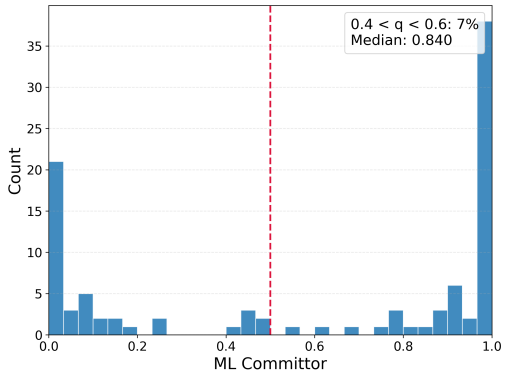

(c) Isodensity Interpolation + Score-Aligned Ascent.

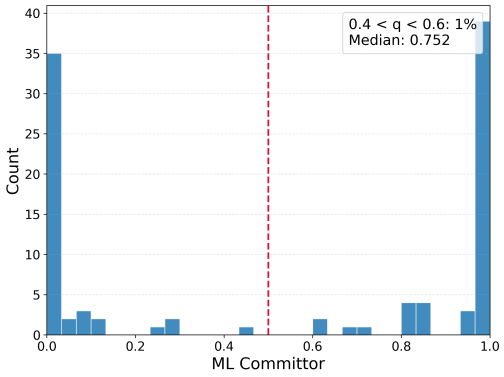

(d) Simple Averaging + Score-Aligned Ascent.

Figure 11: Histogram of ML predicted committors on Chignolin.

F.2.3 Comparison with Other Methods

We evaluate two baseline methods for generating conformational pathways: linear interpolation, performed after applying the Kabsch algorithm [Lawrence et al., 2019], and a geodesic interpolation approach based on the radius of gyration, as proposed by Zhu et al. [2019], in Figure 12 and 13 Specifically for the alanine dipeptide, both methods fail to produce physically realistic pathways, as they traverse significant energy barriers on the Potential Energy Surface (PES). Linear interpolation, in particular, is prone to generating intermediate conformations with severe steric clashes, leading to geometrically invalid structures. While the geodesic method avoids such direct structural inconsistencies, the pathway it defines remains energetically prohibitive and is therefore an unviable representation of the transition.

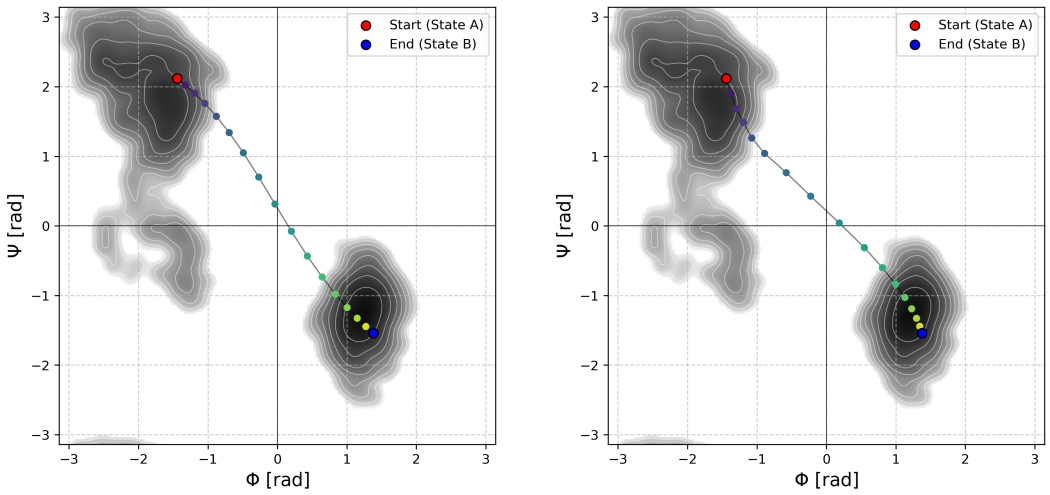

(a) Linear interpolation.

(b) Geodesic interpolation.

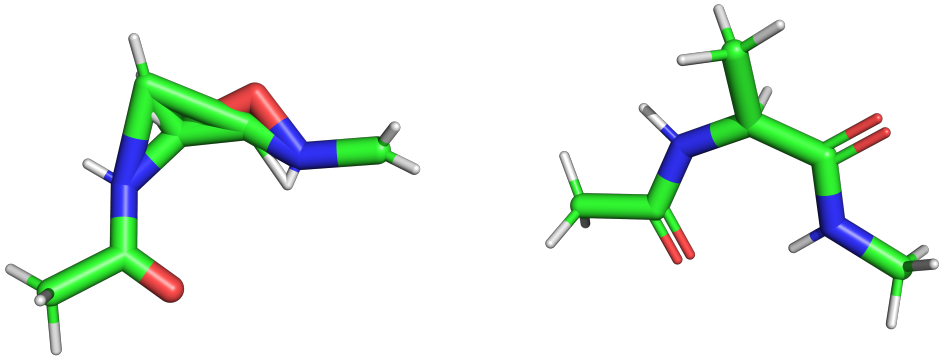

(c) Structure of linear interpolation.

(d) Structure of geodesic interpolation.

Figure 12: Comparison of linear and geodesic interpolation for alanine dipeptide. The top row visualizes the interpolated pathways on the potential energy surface (PES), while the bottom row shows representative intermediate structures.

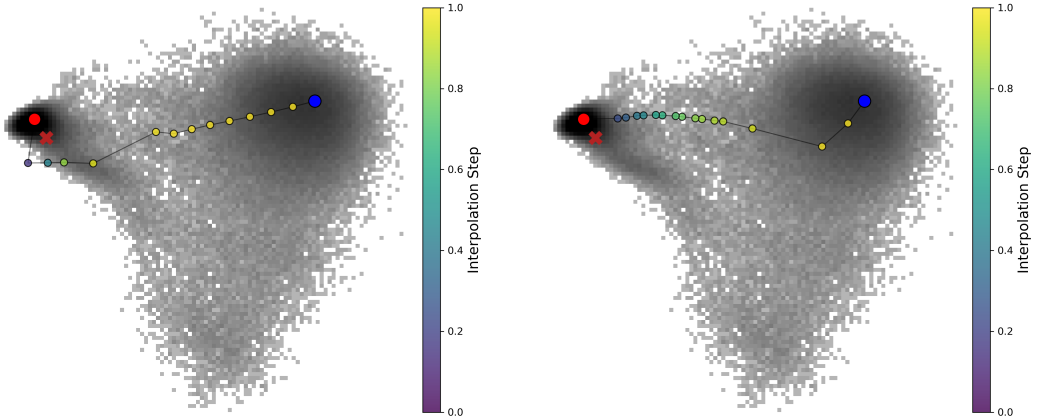

(a) Linear interpolation.

(b) Geodesic interpolation.

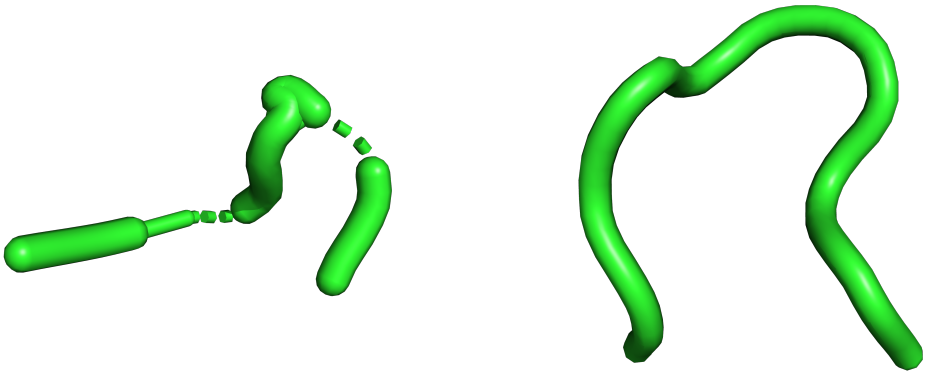

(c) Structure of linear interpolation.

(d) Structure of geodesic interpolation.

Figure 13: Comparison of linear and geodesic interpolation for the folding of chignolin. The top row displays the pathways on the PES, and the bottom row shows intermediate molecular structures.

F.3 Qualitative Analysis of Chignolin Transition Mechanism

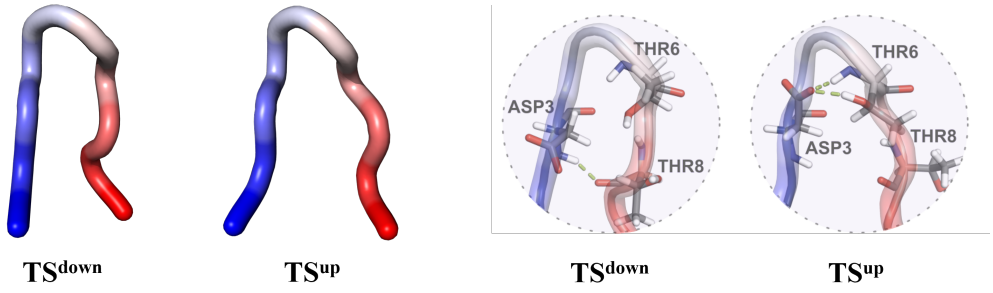

(a) ASTRA generated samples. (b) References.

Figure 14: Visualization of chignolin transition state mechanisms. Representative structures from ASTRA-sampled TS^{down} and TS^{up} (**left**) ensembles are overlaid with transparent tubes on the reference conformations from Kang et al. [2024] (**right**). The structural agreement validates our method's ability to resolve distinct folding pathways at atomic resolution.

To further validate the mechanistic accuracy of our approach, we conduct detailed structural analysis of the identified transition state sub-ensembles. As shown in Figure 14, representative structures from our ASTRA-generated TS^{down} and TS^{up} ensembles exhibit remarkable concordance with reference configurations from Kang et al. [2024]. This structural validation confirms that our method not only localizes the correct transition region but also faithfully reproduces the atomic-level details that distinguish competing folding mechanisms, establishing ASTRA as a powerful tool for mechanistic discovery in complex biomolecular systems.