Protein Conformational Trajectory Prediction with Odds Ratio Preference Optimization

Divya Nori
divnor80@mit.edu

Andrei Tyrin
atyrin@mit.edu

Abstract

Protein conformational ensembles offer valuable insights into protein function and reveal potential drug targets. In this study, we fine-tune ESMFold, a widely-used protein folding model, to predict conformational trajectories using odds ratio preference optimization (ORPO). Our approach shows initial success in generating temporally-consistent trajectories that accurately predict residue flexibility and capture the dominant conformations.

1 Introduction

Proteins, a diverse class of biological macromolecules, perform a wide range of functions critical for cellular pysiology. Understanding protein structures is essential for determining these functions. Recent advances in protein structure prediction, enabled by models like AlphaFold and ESMFold [3, 4], have marked significant progress. However, these models primarily predict static structures, whereas proteins are dynamic molecules whose movements are integral to their function. This study aims to repurpose existing protein folding models, specifically ESMFold, to generate conformational trajectories. We modify ESMFold's architecture and finetune the model using odds ratio preference optimization (ORPO) [1]. Specifically, we combine a supervised behavior cloning loss, using molecular dynamics simulations as expert data, with an odds ratio-based preference alignment loss to generate trajectories that are temporally-consistent, biophysically accurate, and conformationally diverse.

2 Methods

2.1 Problem Formulation

Consider a protein structure which is given by a 3D point cloud $s_0 \in \mathbb{R}^{N\times 37\times 3}$. N is the number of amino acid residues in the protein sequence, and each residue's geometry is given in *atom37* format where each of the 37 slots corresponds to a heavy atom of a given name. We are also given the protein's sequence which does not change over the course of our task. We aim to predict a trajectory of states $[s_1, ..., s_T]$ given s_0, which describes how this protein structure evolves over time in a room-temperature water solution with no other ligands. Thus, our **state space** is continuous, consisting of all possible point clouds with the given dimensionality. The **action space** consists of modifying any subset of the protein's atomic positions in 3D space, so this space is also continuous with dimensionality $\mathbb{R}^{N\times 37\times 3}$. Given the nature of our problem, s_{t+1} is trivially given by s_t and a_t, so we parameterize our policy such that it directly predicts a distribution over s_{t+1} given s_t.

2.2 Protein Folding Model as a Trajectory Generator

In recent years, incredible progress has been made on the task of *protein folding*. Deep learning models such as ESMFold [4] has been trained to accept protein sequences and predict all-atom 3D coordinates

6.8200 Computational Sensorimotor Learning Final Project (Spring 2024).

of single structures with high accuracy. ESMFold has gained considerable popularity because it is significantly faster than AlphaFold [3] while only slightly underperforming in terms of structure prediction accuracy. As shown in Figure 1, ESMFold begins by embedding a protein sequence with the ESM-2 language model which is pretrained on a masked language modeling objective. The resulting embedding is processed by the *folding trunk* (48 triangular self-attention blocks) which updates the internal representation of the protein. The output is processed by an equivariant structure module, from which the model predicts an all-atom structure along with confidence measures.

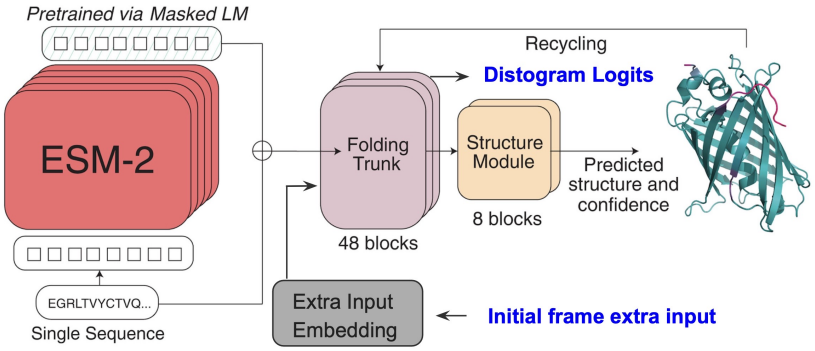

Figure 1: Architecture of ESMFold model with additional embedding module.

Folding models hold rich information about protein structures in their weights, but they are not trained to predict several conformations. Inspired by recent methods that adapt these models for conformational ensemble prediction [2], we hypothesize that pretrained protein folding models can be finetuned towards conformational trajectory prediction. Thus, we initialize our policy $\pi(s_{t+1}|s_t)$ with pretrained weights from ESMFold which has been trained on proteins deposited in the Protein Data Bank (PDB) before May 1, 2020. To adapt ESMFold for trajectory prediction, we add an extra input embedding module which conditions s_{t+1} predictions on s_t, as shown in Figure 1. The architecture of this module is similar to AlphaFold's template embedding stack, involving triangular attention and multiplicative updates.

2.3 Aligning Folding to MD Trajectories

We now focus on the problem of finetuning this model for trajectory prediction, given ground-truth molecular dynamics (MD) trajectory data. There are many potential approaches to do this. The simplest method would be to treat MD simulations as expert trajectories and apply supervised finetuning (SFT, i.e. behavior cloning). We hypothesize that this approach could work well but may sacrifice diversity if employed alone. Inspired by recent success in the language modeling domain, we hypothesize that methods rewarding the "winner" between two independently sampled completions $s_{t+1}|s_t$ could perform well. ORPO applies such as an approach and outperformed both RLHF and DPO in the language domain being significantly simpler than analogous methods [1]. Specifically, ORPO combines a supervised fine-tuning loss term with a term that maximizes the odds ratio between the likelihood of generating the favored completion $\hat{s}_{t+1}^{(w)}$ and disfavored completion $\hat{s}_{t+1}^{(l)}$. Our full objective function is given below:

$$\mathcal{L}_{TRPO} = \mathbb{E}_{(p_{t-1},\hat{p}_t^{(w)},\hat{p}_t^{(l)})}\left[\mathcal{L}_{SFT} + \lambda \mathcal{L}_{OR}\right]$$

where $\mathcal{L}_{SFT}$ is the cross-entropy loss calculated between the distogram logits and ground truth distogram distribution of the true protein conformer. The odds ratio term takes the form

$$\mathcal{L}_{OR} = -\log \sigma \left(\log \frac{\textbf{odds}_\theta(\hat{s}_t^{(w)}|s_{t-1})}{\textbf{odds}_\theta(\hat{s}_t^{(l)}|s_{t-1})} \right), \quad \textbf{odds}_\theta(s_t|s_{t-1}) = \frac{P_\theta(s_t|s_{t-1})}{1 - P_\theta(s_t|s_{t-1})}$$

where $\hat{s}_t^{(\cdot)}$ is the predicted protein conformation given s_{t-1} as the previous one. Winner is selected by the comparison of the MSE between the sampled distogram of the proposed conformer and true conformation s_t.

The central question now is how to sample the completions. We make the observation that the folding trunk predicts logits over a distogram, which is a pairwise matrix of distances between C_α atoms discretized into buckets. We can sample over this predicted distribution to arrive at two potential next-state backbone structures. The gradients propagate from the distogram logits through the folding trunk and extra input embedding module. We keep the ESM-2 language model frozen, and the structure module is not updated. We finetune the model for 400 gradient steps across 7 NVIDIA A-6000 GPUs taking significantly smaller time compared to the ESMFlow method.

3 Experiments

For all experiments, we use the ATLAS dataset [6] which contains MD trajectories for 1400 proteins. We apply a temporal split to create train, validation, and test sets (same split as used in prior work [2]). The motivation for a time-based split comes from prior work in ML for structural biology [5]. All metrics are reported on the test set. We compare our method with two baselines. First, we compare a recently proposed method named ESMFlow [2] which generates protein conformational ensembles by finetuning ESMFold on the generative flow matching objective. While this method cannot simulate trajectories with temporal consistency and instead predicts unordered states, we can compare the diversity and accuracy of conformational ensembles. Second, a simple but strong baseline is assuming that the protein simply stays static throughout the trajectory which we refer to as "Void MD." We seek to answer three experimental questions, each of which we describe below.

3.1 Temporal evolution of protein structure

Given the first frame of a true MD trajectory, we assess whether our method can generate a temporally consistent set of protein structures. Each frame-to-frame transition should resemble a 1 ns step of MD simulation, and ideally, the accuracy of each state should not diminish significantly over time.

To measure this, using our method, we sample 100-state trajectories given the true first frame of expert trajectories. We do not compare ESMFlow on this metric because it was not designed to generate temporally consistent trajectories. Void MD assumes that the true first frame is repeated 100 times through the trajectory. We compare these methods on the average RMSE aggregated across trajectory

$$RMSE_{temporal} = \frac{1}{T}\sum_{i=1}^{T} RMSE(\hat{\mathbf{X}}_i, \mathbf{X}_i)$$

where $\hat{\mathbf{X}}_i$ corresponds to the predicted all-atom coordinates at step i and $\mathbf{X}_i$ is a set of true coordinates at the same step.

We find that ESMFold-ORPO has a $RMSE_{temporal}$ of 6.859 compared to 2.411 for Void MD. This indicating that perhaps the error grows as we deviate away from the first frame, prompting further analysis. We plot $RMSE_{temporal}$ as a function of frame index and show two representative examples below.

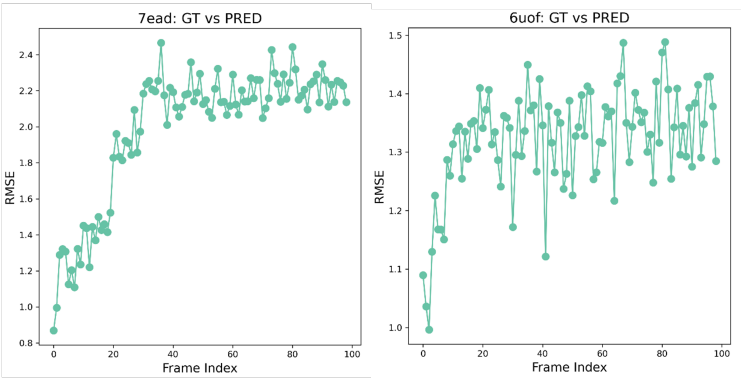

Figure 2: Temporal RMSE over time.

We observe that the temporal RMSE plateaus, showing that the error does not grow and ESMFold-ORPO can generate temporally consistent trajectories over long time horizons.

3.2 Amino acid flexibility

Next, we evaluate whether our method can accurately predict the flexibility of each residue. To measure this, we report the Pearson correlation between the Root Mean Squared Fluctuation (RMSF) of residues in true and predicted trajectories. Specifically, we report

$$RMSF_{corr} = PearsonR(RMSF_{true}, RMSF_{pred}), \quad RMSF = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(\mathbf{X}_i - \bar{\mathbf{X}})^2}$$

where $\bar{\mathbf{X}}$ is the average conformation in the trajectory, and $\mathbf{X}_i$ is a state at timestep i. The RMSF correlation values for our method and the baselines is shown below.

Table 1: Comparison of Models on Per-Target RMSF Correlation

Model	RMSF Correlation (r)
Void MD	0.167
ESMFlow-Templates	0.402
ESMFold-ORPO	0.390

We observe that ESMFold-ORPO almost matches the performance of ESMFlow-Templates, significantly outperforming Void MD.

3.3 Coverage of conformational space: case-study

Finally, we evaluate whether ESMFold-ORPO can cover the dominant conformational states present in the ground-truth MD trajectory. To identify dominant conformational states in the ground-truth trajectories, we plot the pairwise similarity between the true conformers' distance matrices on a heatmap. Highly correlated states (shown in red) are likely part of the same functional state. We define that a predicted conformer catures a particular true state if the correlation between their pairwise distance matrices is ≥ 0.8. On the heatmap, we mark the true states for which a predicted conformer captures its geometry along the diagonal with a black dot.

We focus on the cytochrome enzyme as a case-study, a common drug target for skin conditions. This protein has two dominant states, and this conformational change is critical for its function. The heatmap for this protein is shown below, with ESMFlow predictions marked on the left panel and ESMFold-ORPO marked on the right.

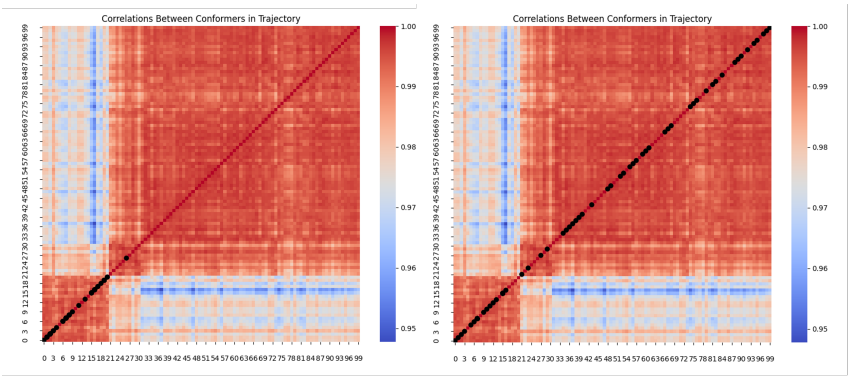

Figure 3: Coverage of conformational space for cytochrome enzyme.

Evidently, ESMFold-ORPO is able to capture both conformational states reliably, compared to ESMFlow which primarily predicts thermal fluctuations around the starting conformation. We observe this trend with many other proteins as well, indicating that perhaps ESMFold-ORPO can be used to generate ensembles of higher diversity. This analysis will be the subject of future work.

References

[1] Jiwoo Hong, Noah Lee, and James Thorne. "Reference-free monolithic preference optimization with odds ratio". In: *arXiv preprint arXiv:2403.07691* (2024).

[2] Bowen Jing, Bonnie Berger, and Tommi Jaakkola. "AlphaFold Meets Flow Matching for Generating Protein Ensembles". In: *arXiv preprint arXiv:2402.04845* (2024).

[3] John Jumper et al. "Highly accurate protein structure prediction with AlphaFold". In: *Nature* 596.7873 (2021), pp. 583–589.

[4] Zeming Lin et al. "Evolutionary-scale prediction of atomic-level protein structure with a language model". In: *Science* 379.6637 (2023), pp. 1123–1130.

[5] Hannes Stärk et al. "Equibind: Geometric deep learning for drug binding structure prediction". In: *International conference on machine learning*. PMLR. 2022, pp. 20503–20521.

[6] Yann Vander Meersche et al. "ATLAS: protein flexibility description from atomistic molecular dynamics simulations". In: *Nucleic Acids Research* 52.D1 (2024), pp. D384–D392.