

EPO: Diverse and Realistic Protein Ensemble Generation via Energy Preference Optimization

Yuancheng Sun^{1, 2, 3}, Yuxuan Ren³, Zhaoming Chen³, Xu Han^{1, 2, 3}, Kang Liu^{1, 2}, Qiwei Ye^{3*}

¹Institute of Automation, Chinese Academy of Sciences

²University of Chinese Academy of Sciences

³Beijing Academy of Artificial Intelligence

sunyuancheng2021@ia.ac.cn, chivee.ye@gmail.com

Abstract

Accurate exploration of protein conformational ensembles is essential for uncovering function but remains hard because molecular-dynamics (MD) simulations suffer from high computational costs and energy-barrier trapping. This paper presents Energy Preference Optimization (EPO), an on-line refinement algorithm that turns a pretrained protein ensemble generator into an energy-aware sampler without extra MD trajectories. Specifically, EPO leverages stochastic differential equation sampling to explore the conformational landscape and incorporates a novel energy-ranking mechanism based on list-wise preference optimization. Crucially, EPO introduces a practical upper bound to efficiently approximate the intractable probability of long sampling trajectories in continuous-time generative models, making it easily adaptable to existing pretrained generators. On Tetrapeptides, ATLAS, and Fast-Folding benchmarks, EPO successfully generates diverse and physically realistic ensembles, establishing a new state-of-the-art in nine evaluation metrics. These results demonstrate that energy-only preference signals can efficiently steer generative models toward thermodynamically consistent conformational ensembles, providing an alternative to long MD simulations and widening the applicability of learned potentials in structural biology and drug discovery.

Introduction

Proteins operate through continual interconversion among multiple conformational states. This dynamic landscape underlies fundamental biological processes, including allostery, molecular recognition, and catalysis (Nussinov 2016; Raisinghani et al. 2024; Henzler-Wildman and Kern 2007). Consequently, mechanistic insight (Nussinov et al. 2023; Kalakoti and Wallner 2025; Schafer et al. 2025)—and, by extension, structure-based drug discovery—requires models that recover the entire Boltzmann ensemble rather than a single static conformation (Teixeira et al. 2022; Nussinov 2016; Raisinghani et al. 2024).

Molecular dynamics (MD) remains the standard tool for exploring these ensembles (Stevens et al. 2023; Al-Rawashdeh and Barakat 2023; Ghahremanian et al. 2022; Badar et al. 2022). Yet MD trajectories are constrained by

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

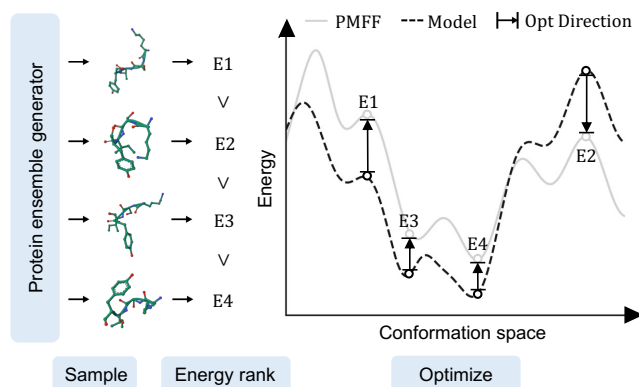


Figure 1: EPO aligns the model’s energy landscape with the ground truth by leveraging energy rankings from a physics-based molecular force field (PMFF), guiding the model to assign lower energy to more favorable conformations.

rugged energy landscapes whose high barriers render functional transitions rare on simulation timescales (Héning et al. 2022; Ray et al. 2022; Wang et al. 2021; Souza et al. 2021). Conventional runs therefore become trapped in local minima and undersample transient or high-energy states that are often functionally critical.

Recent advancements in deep generative models offer a computationally efficient approach to achieving conformational diversity. Early efforts perturb single-state predictors such as AlphaFold2 (Jumper et al. 2021) with dropout, template shuffling, or MSA subsampling (Del Alamo et al. 2022; Stein and Mchaourab 2022; Wayment-Steele et al. 2024), but yield only modest variations near the dominant conformation. A more prevalent strategy involves fine-tuning generative models on extensive MD trajectories (Lewis et al. 2025; Team and Ye 2025; Klein and Noe 2024). However, supervised fine-tuning may suffer from misalignment out of inadequate sampling. This issue is two-fold. First, producing equilibrated MD trajectories necessitates prohibitively long simulations that may still undersample the true ensemble in large biomolecular systems (Héning et al. 2022; Ray et al. 2022). Second, even when fine-tuned on well-sampled trajectories, generative models tend to revisit the same local regions of conformational space, leading

to generated conformations that fail to accurately follow the MD-distribution (Wang et al. 2024; Jing et al. 2024).

Recent advances in large language models have revealed the effectiveness of post-training refinement techniques in addressing the misalignment issues in supervised fine-tuning (Wallace et al. 2024; Guo et al. 2025). Direct Preference Optimization (DPO) (Rafailov et al. 2023) stands out as a representative example of such methods. Drawing a parallel, we posit that pretrained protein ensemble generators can be viewed as competent yet misaligned samplers. However, directly applying DPO to these generators is non-trivial. Ensemble generation aims to reproduce the full Boltzmann distribution of relevant states, whereas the pairwise comparisons used in naïve DPO tend to steer the model toward a single, energetically favorable basin, undermining distributional fidelity (Wang et al. 2023; Han et al. 2025; Lanchantin et al. 2025).

To address these challenges, this paper introduces Energy Preference Optimization (EPO), a novel physics-guided online framework for refining protein ensemble generators. Specifically, EPO integrates stochastic differential equation (SDE) sampling into an online, iterative refinement process. This allows the model to dynamically explore the conformational landscape, thereby mitigating the issue of inadequate sampling. Furthermore, to overcome the limitations of pairwise comparisons, EPO introduces a physics-based energy ranking mechanism. This mechanism employs listwise preference optimization to guide the generator towards a diverse and physically realistic ensemble rather than a single low-energy state. Crucially, EPO derives a practical upper bound for the listwise preference objective, effectively approximating the intractable transition probabilities of long sampling trajectories inherent in continuous-time generative models. This formulation integrates seamlessly with modern denoising score-matching and flow-matching architectures and can be adopted directly by existing pretrained generators.

Extensive experiments demonstrate that preference signals derived solely from physical energy are sufficient to correct the misalignment of pretrained generators. As a result, this relieves the need for expensive post-hoc MD trajectories. Empirically, EPO establishes a new state-of-the-art across nine distinct distributional metrics on the Tetrapeptides (Jing et al. 2024), ATLAS (Vander Meersche et al. 2024), and Fast-Folding (Lindorff-Larsen et al. 2011) benchmarks. Furthermore, visualizations and ablation studies validate the effectiveness of EPO’s design components.

Our contributions are summarized as follows:

- This paper proposes EPO, an online framework that directly aligns pretrained generators with the target Boltzmann distribution using listwise energy preferences.
- A practical upper bound for the intractable listwise preference objective is derived, yielding tractable gradients compatible with modern continuous-time generative models.
- Extensive experiments show that EPO attains state-of-the-art results on several metrics and benchmarks, and delivers competitive performance elsewhere, while producing diverse and physically realistic ensembles with-

out requiring additional MD simulations.

Preliminaries

Direct Preference Optimization. Let a prompt $x \in \mathcal{X}$ be given and let a policy $\pi_\theta(y|x)$ assign probabilities to responses $y \in \mathcal{Y}$ (e.g. molecular conformations, textual answers). The learning goal is to adjust θ so that the induced distribution matches externally provided preferences.

Direct Preference Optimization (DPO) (Rafailov et al. 2023) learns from a dataset $\mathcal{D} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_{i=1}^N$, where each triple contains a winning sample y_w and a losing sample y_l . Under the Bradley–Terry comparison model, the DPO loss is

$$\mathcal{L}_{\text{DPO-BT}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(y_w|x)}{\pi_{\text{ref}}(y_w|x)} - \beta \log \frac{\pi_\theta(y_l|x)}{\pi_{\text{ref}}(y_l|x)} \right) \right] \quad (1)$$

where σ is the logistic function and $\beta > 0$ controls the implicit KL regularization with respect to a fixed reference policy π_{ref} , i.e., the initial pretrained checkpoint.

In practice, human preferences often extend beyond pairwise comparisons to ranked lists, introducing a key challenge that the magnitude of preference between items is typically non-uniformly weighted. Consequently, when a prompt is associated with an ordered list $\tau = (y_{(1)}, \dots, y_{(m)})$ sorted by preference, we can generalize Eq. 1 to a listwise variant, which can be derived from the Plackett–Luce model (Rafailov et al. 2023; Liu et al. 2024) or from ListMLE (Xia et al. 2008):

$$\mathcal{L}_{\text{DPO-PL}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(\tau, x, y_1, \dots, y_k) \sim \mathcal{D}} \left[-\sum_{i=1}^K \log \frac{\exp(s_\theta(x^{\tau(i)}))}{\sum_{j=i}^K \exp(s_\theta(x^{\tau(j)}))} \right], \quad (2)$$

where $s_\theta(x^i) \triangleq \beta \log \frac{\pi_\theta(y|x^i)}{\pi_{\text{ref}}(y|x^i)}$.

Flow Matching and SDE Sampling. Flow Matching (FM) (Lipman et al. 2023; Albergo and Vanden-Eijnden 2023) trains a time-dependent velocity network $v_\theta(x, t)$ that transports samples from a simple prior distribution $p_0(x)$ (e.g., $\mathcal{N}(0, I)$) to a target data distribution $p_1(x)$. This process is defined by constructing a coupling path between pairs of samples $(x_0, x_1) \sim p_0 \times p_1$. For any such pair, the path x_t for $t \in [0, 1]$ is defined as: $x_t = \alpha_t x_1 + \sigma_t x_0$. Here, α_t and σ_t are scheduling functions designed to ensure the path smoothly interpolates from the prior sample to the data sample. Let $\dot{x}_t = \dot{\alpha}_t x_1 + \dot{\sigma}_t x_0$ denote the ground-truth velocity along this path. FM fits v_θ by minimizing the expectation of a mean squared error (MSE) loss:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{x_0, x_1, t} [\|v_\theta(x_t, t) - \dot{x}_t\|_2^2]. \quad (3)$$

After training, samples can be generated by solving a probability flow ordinary differential equation (ODE): $x_1 = \text{odeint}(v_\theta(t), x_0, t : 0 \rightarrow 1)$. Following Ma et al., we

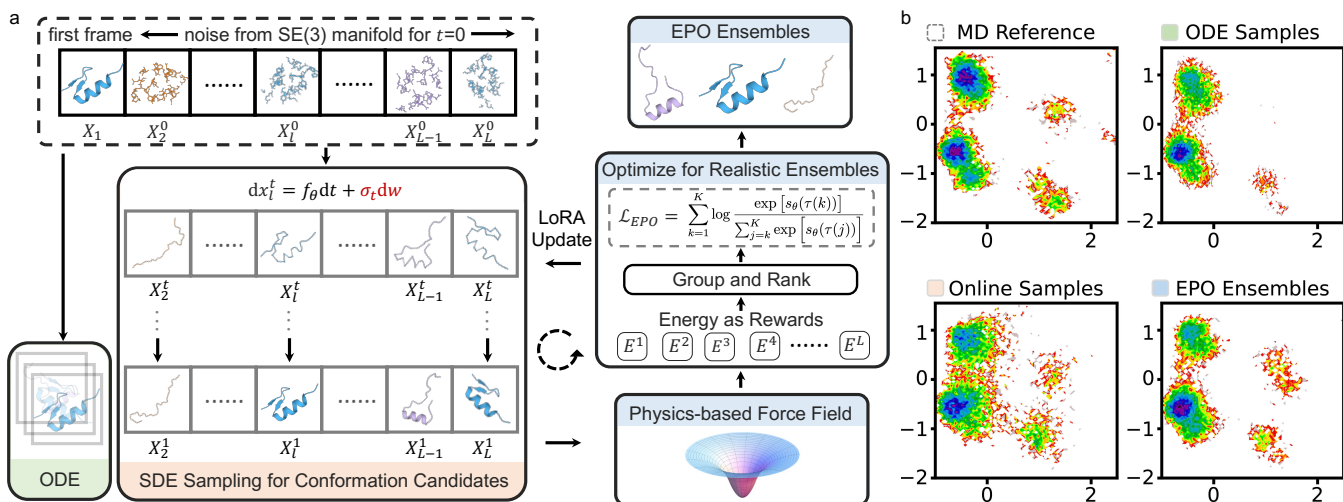


Figure 2: Overview of EPO. (a) Given the first frame (static conformation) of a protein, we leverage a pretrained ensemble generator and introduce an ODE-to-SDE strategy to enable stochastic sampling for online optimization. Energies of the online samples are used as rewards to update the model with LoRA. (b) SDE-based sampling effectively overcomes the energy barriers inherent in the original model (sequence: ASRE). Consequently, the optimized EPO model generates diverse and physically realistic protein ensembles, bypassing the need for any post-hoc MD simulations.

construct a reverse-time SDE formulation that preserves the same marginal distribution as the ODE:

$$dx_t = v(x_t, t)dt + \frac{1}{2}w_t s(x_t, t)dt + \sqrt{w_t}d\bar{W}_t, \quad (4)$$

$$s(x_t, t) = \sigma_t^{-1} \frac{\alpha_t v_\theta(x_t, t) - \dot{\alpha}_t x_t}{\dot{\alpha}_t \sigma_t - \alpha_t \dot{\sigma}_t}, \quad (5)$$

where $d\bar{W}_t$ denotes the Wiener process and $w_t > 0$ is the score norm controlling the level of stochasticity.

This equivalence allows us to employ existing pretrained velocity networks within an SDE sampling framework to enhance sampling diversity and aid optimization, as has been demonstrated in recent literature (Lu et al. 2023; Xue et al. 2025; Liu et al. 2025).

FlowDPO. For autoregressive language or diffusion models, $\log \pi_\theta(y|x)$ is directly available as token logits or trajectory log-probabilities, making DPO an efficient, black-box-free alternative to RLHF. Unlike language models, closed-form $\log \pi_\theta(y|x)$ of FM is computational; therefore, the pairwise DPO loss (Eq. 1) appears intractable because it requires trajectory-wide log-likelihoods. FlowDPO (Jiao et al. 2024) resolves this by amortizing the pairwise DPO objective over time: under a Gaussian assumption on marginal states, the global preference loss decomposes into a weighted difference of per-time $\text{MSE}_t(x_0, x_1; \theta)$, each of which is computable from v_θ in Eq. 3 alone:

$$\begin{aligned} \mathcal{L}_{\text{DPO-FM}} = & -\mathbb{E}_{x_{0,1}^w, x_{0,1}^l, t} \log \sigma \left(\right. \\ & \beta \left[\text{MSE}_t(x_0^w, x_1^w; \theta_{\text{ref}}) - \text{MSE}_t(x_0^w, x_1^w; \theta_{\text{opt}}) \right. \\ & \left. \left. - \text{MSE}_t(x_0^l, x_1^l; \theta_{\text{ref}}) + \text{MSE}_t(x_0^l, x_1^l; \theta_{\text{opt}}) \right] \right). \quad (6) \end{aligned}$$

A detailed proof can be found in Appendix.

Method

Prevailing pairwise preference optimization techniques fall short in modeling complex systems, as they cannot capture the free energy hierarchy of the conformational ensemble. Listwise objectives offer a more powerful alternative but are typically computationally intractable. To bridge this gap, we introduce EPO, a new framework that leverages a practical upper bound on the listwise objective. This formulation re-frames the problem as a direct energy minimization task, unlocking the potential of the listwise approach by ensuring both its effectiveness and computational feasibility.

Suboptimality of Pairwise DPO for Energy Alignment

Generally, a protein free-energy landscape is characterized by several metastable basins whose local minima may lie at markedly different absolute energies. When a generator is trained with pairwise Direct Preference Optimization (Eq. 1), every comparison favors the conformation of lower energy, gradually steering the model toward the single, globally best basin. This pairwise pressure undermines the very goal of ensemble modelling—namely, to reproduce the full distribution of relevant states—because higher-energy but kinetically important basins are progressively ignored. A listwise loss, which updates the parameters with respect to the ordering of an entire batch, treats these basins more evenly and therefore better preserves diversity.

Beyond diversity, listwise objectives possess a functional form—typically a log-sum-exp over batch scores (Eq. 2)—that can be interpreted as soft importance reweighting. Low-energy samples receive exponentially larger gradients, yet the contribution of higher-energy samples does not vanish entirely. Recent work on energy-based fine-tuning

has shown that such adaptive weighting is beneficial for matching the target Boltzmann distribution (Uehara et al. 2025; Li et al. 2024; Wang et al. 2024).

Finally, the computational route to parity between the two settings is impractical. Although evaluating all $\binom{K}{2}$ comparisons in each batch allows a pairwise objective to be equivalent to a listwise objective, this is prohibitive and impractical for realistic values of K in flow-matching models. Restricting the loss to a subset of pairs—such as adjacent elements after energy sorting—reduces cost but at the price of higher gradient variance and information loss. This simplified training setting is denoted as EPO-Pair and its underperformance compared to the full listwise version, EPO-List, is demonstrated in the experimental section.

For these reasons, EPO is designed with a listwise structure to guide the generative model towards the underlying physical energy landscape in a manner that is both robust and diversity-promoting.

Extension on Flow Matching

Adapting the Listwise Preference Optimization (LiPO) framework, originally designed for discrete outputs from language models, to continuous-time generative models like flow matching presents a significant challenge. The core of this adaptation lies in redefining the reward mechanism. While standard LiPO evaluates a model’s final, static output, the continuous dynamics of flow matching necessitate a reward defined over the entire generation trajectory. Therefore, drawing inspiration from recent work on preference optimization for generative models (Wallace et al. 2024; Jiao et al. 2024), we propose a trajectory-based reward function:

$$r(x, y_0) = \beta \mathbb{E}_{p_\theta(y_{1:T} | y_0, x)} \left[\log \frac{p_\theta^*(y_{0:T} | x)}{p_{\text{ref}}(y_{0:T} | x)} \right] + \beta \log Z(x). \quad (7)$$

Substituting this into the Listwise Preference Optimization (LiPO) framework, we obtain the Listwise preference losses:

$$L_{\text{LiPO-Flow}}(\theta) = -\mathbb{E}_{x, y_1, \dots, y_K \sim \mathcal{D}} \sum_{k=1}^K \log \frac{\exp \left(\mathbb{E}_{y_{1:T}^{\tau(k)}} \left[\beta \log \frac{\pi^*(y_{0:T}^{\tau(k)} | x)}{\pi_{\text{ref}}(y_{0:T}^{\tau(k)} | x)} \right] \right)}{\sum_{j=k}^K \exp \left(\mathbb{E}_{y_{1:T}^{\tau(j)}} \left[\beta \log \frac{\pi^*(y_{0:T}^{\tau(j)} | x)}{\pi_{\text{ref}}(y_{0:T}^{\tau(j)} | x)} \right] \right)}, \quad (8)$$

where $y_{1:T}^{\tau(k)} \sim p_\theta(y_{1:T} | y_0^{\tau(k)})$.

Practical Upper Bound

The listwise loss derived above is intractable to optimize directly for two primary reasons. First, the functional form of the Plackett-Luce choice probability is difficult to handle. EPO addresses this by leveraging the function’s convexity; applying Jensen’s inequality yields a tractable upper bound. Second, sampling from the reverse-time transition distribution, $p_\theta(y_{t-1,t} | y_0)$, is infeasible during training. Following established practice for continuous-time generative models (Jiao et al. 2024), we approximate this intractable

reverse transition with the corresponding forward process probability.

By combining these two solutions—the upper bound on the choice probability and the forward process approximation—and further simplifying a key resulting term with a Mean Squared Error (MSE) proxy between the predicted and reference trajectories, the final, practical objective is arrived:

$$\begin{aligned} L_{\text{LiPO-Flow Matching}}(\theta) &\leq -\mathbb{E}_{(x, \mathcal{Y}), t, y_{t^*}^k} \sum_{k=1}^K \log \frac{\exp \left(\beta \log \frac{\pi^*(y_{T-1}^{\tau(k)} | y_T^{\tau(k)})}{\pi_{\text{ref}}(y_{T-1}^{\tau(k)} | y_T^{\tau(k)})} \right)}{\sum_{j=k}^K \exp \left(\beta \log \frac{\pi^*(y_{T-1}^{\tau(j)} | y_T^{\tau(j)})}{\pi_{\text{ref}}(y_{T-1}^{\tau(j)} | y_T^{\tau(j)})} \right)} \\ &= -\mathbb{E}_{(x, \mathbf{Y}), t, y_{t^*}^k} \sum_{k=1}^K \log \frac{\exp [s_\theta(\tau(k))]}{\sum_{j=k}^K \exp [s_\theta(\tau(j))]}, \quad (9) \end{aligned}$$

$$s_\theta(\tau(i)) \triangleq \beta \left(\text{MSE}_t(y_0^{\tau(i)}, y_1^{\tau(i)}; \theta_{\text{ref}}) - \text{MSE}_t(y_0^{\tau(i)}, y_1^{\tau(i)}; \theta_{\text{opt}}) \right), \quad (10)$$

where $(x, \mathbf{Y}) \triangleq (x, y_1, \dots, y_K) \sim \mathcal{D}$, $t \sim \mathcal{U}(0, T)$, $y_{t^*}^k \triangleq y_{t-1,t}^k \sim p_\theta(y_{t-1,t} | y_0^k)$, $\forall k \in \{1, \dots, K\}$.

Ensemble Sampling as Forward Simulation

EPO uses MDGen (Jing et al. 2024) as the backbone model to perform forward simulation of molecular trajectories $\chi = [\mathbf{X}_1, \dots, \mathbf{X}_T]$ for a given L-residue amino acid sequence A . Each $\mathbf{X}_t \in \mathbb{R}^{3N}$ represents all-atom coordinates. MDGen processes proteins by defining local reference frames for each residue (Yim et al. 2023), and parameterizes molecular structures using per-residue SE(3) roto-translations (R, \mathbf{t}) and seven internal torsion angles $(\psi, \phi, \omega, \chi_1, \dots, \chi_4)$, leading to $\chi_t^l = ((R, \mathbf{t})_t^l, (\psi, \phi, \omega, \chi_1, \dots, \chi_4)_t^l) \in ([\text{SE}(3) \times \mathbb{T}^7]^L)^T$. MDGen addresses the computationally prohibitive updates in each frame by transforming the absolute poses of the trajectories into a key-frame-based relative representation. Given K key frames ($K = 1$ indicates the first trajectory frame in forward simulation, which is the setting in this paper), the state of residue j at frame t (defined by its absolute roto-translation g_t^j and torsions τ_t^j) is featurized into a token ξ_t^j . This token comprises K relative roto-translations and the residue’s torsions, embedded into a fixed-dimension vector:

$$\begin{aligned} \xi_t^j &= \left([g_t^j]^{-1} g_t^j, \tau_t^j \right) \in \text{SE}(3) \times \mathbb{T}^7 \\ &\equiv (\hat{\mathbb{Q}}^+ \oplus \mathbb{R}^3)^1 \times (S^2)^7 \subset \mathbb{R}^{7+14}. \quad (11) \end{aligned}$$

A transformer-based model architecture guided by SiT (Ma et al. 2024) is used in MDGen, parameterizing a final velocity network $v_\theta: v_\theta(\cdot, t | \{g_{t_k}\}_{k=1}^K, A) : \mathbb{R}^{T \times L \times (7K+14)} \times [0, 1] \rightarrow \mathbb{R}^{T \times L \times (7K+14)}$.

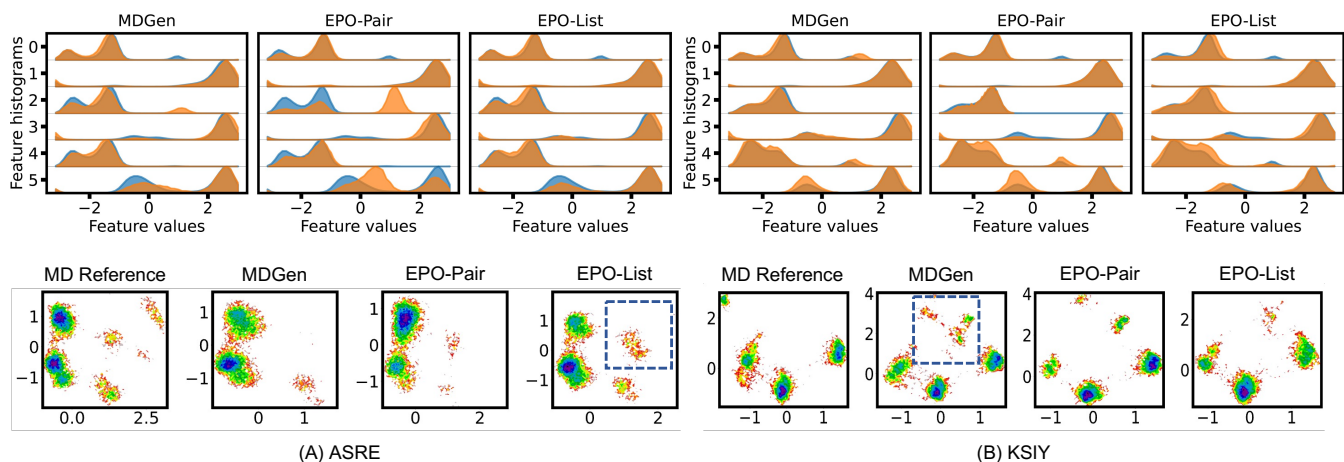


Figure 3: Comparison of torsion angle distributions and free energy surfaces (FES) for two tetrapeptide sequences. EPO successfully captures crucial metastable states absent in pretrained model outputs and effectively corrects high-energy biases observed in pretrained distributions (highlighted in dotted boxes, respectively), demonstrating strong alignment with the ground-truth energy landscapes.

	Pairwise RMSD r (\uparrow)	Global RMSF r (\uparrow)	Per-target RMSF r (\uparrow)	Root \mathcal{W}_2 (\downarrow)	MD PCA \mathcal{W}_2 (\downarrow)	%PC-sim>0.5 (\uparrow)	Weak contacts J (\uparrow)	Exposed residue J (\uparrow)
MDGen	0.48/0.42	0.50/0.49	0.71/0.70	2.69/3.47	1.89/2.43	10	0.51	0.29
EPO-Pair	0.49/0.49	0.49/0.49	0.74/0.69	2.88/3.53	1.88/2.42	9	0.42	0.43
EPO-List	0.51/0.51	0.50/0.50	0.75/0.70	2.79/3.43	1.77/2.35	13	0.45	0.44

Table 1: Comparison of generated ensembles against ground-truth MD simulations on the ATLAS benchmark. Values reported are the median, with mean values shown after a backslash where applicable.

Experiments

In this section, we conduct comprehensive experiments to rigorously evaluate the effectiveness of the proposed EPO framework. Our main findings can be summarized as follows: (1) EPO establishes new state-of-the-art performance across nine evaluation metrics on three widely adopted protein conformation benchmarks. (2) By relying solely on energy-based labels, EPO promotes more effective exploration of the conformational space and improves alignment of the generated ensembles with the Boltzmann energy distribution. (3) Despite theoretical equivalence, the listwise approach outperforms its pairwise counterpart in practical experiments.

Implementation Details

Datasets. The adopted benchmarks span a range of molecular systems—from small peptides to large, structurally varied proteins and fast-folding systems—allowing us to assess our model’s capabilities across different scales and complexities. Specifically, we first evaluate the Tetrapeptides (Jing et al. 2024) dataset, which contains all-atom molecular dynamics (MD) trajectories for 3,000 training, 100 validation, and 100 test tetrapeptides, each simulated for 100 ns. We then conduct finetuning on ATLAS (Vander Meersche et al. 2024) dataset, which is constructed to maximize domain-

level structural coverage across ECOD X-classes. It consists of explicit-solvent, all-atom MD simulations, providing three independent 100 ns trajectories for each of its 1,390 structurally diverse protein entries. We also evaluate the Fast-Folding benchmark (Lindorff-Larsen et al. 2011), which includes MD simulation data for 12 small proteins exhibiting rapid folding-unfolding dynamics.

We follow the standard dataset splits as in MDGen (Jing et al. 2024) and AlphaFlow (Jing, Berger, and Jaakkola 2024) for the three datasets.

Baselines. EigenFold (Jing et al. 2023) is the first to leverage diffusion models for this task, generating conformational ensembles by sampling from a learned distribution over protein eigenmodes. MDGen (Jing et al. 2024) directly models MD trajectories, utilizing a Scalable Interpolant Transformer and key-frame conditioning for improved computational efficiency. Str2Str (Lu et al. 2023) uses online stochastic perturbations during generation to enhance diversity, which is similar to EPO, but it does not incorporate explicit physical guidance from energy or force signals. ConfDiff (Wang et al. 2024) are fine-tuned on static, offline datasets annotated with pre-computed energy and force labels, which limits their ability to dynamically adapt to distributional shifts during training. A more detailed discussion of these related works is provided in the Appendix.

Models	JSD (\downarrow)				RMSE _{contact} (\downarrow)	RMSF
	PwD	Rg	TIC	TIC-2D	(\AA)	(\AA)
EigenFold	0.53/0.56	0.52/0.55	0.50/0.50	0.64/0.66	6.18/6.22	1.6/1.1
Str2Str-SDE	0.34/0.32	0.30/0.24	0.39/0.38	0.56/0.58	3.68/4.01	7.8/8.0
Str2Str-ODE	0.37/0.38	0.33/0.30	0.40/0.39	0.57/0.59	4.14/4.36	6.4/6.3
ConfDiff-Energy	0.34/0.34	0.31/0.29	0.39/0.40	0.54/0.56	3.65/3.80	7.1/6.1
ConfDiff-Force	<u>0.29/0.27</u>	0.26/0.24	<u>0.38/0.38</u>	<u>0.54/0.54</u>	3.25/3.38	6.2/5.7
EPO-Pair	0.30/0.30	0.29/0.29	0.39/0.40	0.57/0.61	<u>2.68/2.54</u>	6.6/6.6
EPO-List	0.28/0.26	<u>0.28/0.26</u>	0.32/0.30	0.49/0.47	2.36/2.28	6.3/6.6

Table 2: Results on Fast-Folding proteins. All values are shown as mean/median.

	χ_{bb}	χ_{sc}	χ_{all}	TICA-0	TICA-0,1
Reference	0.103	0.055	0.076	0.201	0.268
MDGen	0.130	0.093	0.109	0.230	0.316
EPO-Pair	0.127	0.098	0.110	0.237	0.318
EPO-List	0.125	0.093	0.107	0.226	0.311

Table 3: Jensen-Shannon divergence between ground-truth and generated tetrapeptide distributions. Lower is better.

Experimental Settings The reference policy model, π_{ref} , is based on MDGen checkpoints separately pretrained on the Tetrapeptides and ATLAS datasets. For the Tetrapeptides experiments, β is set to 1 and the score norm of SDE sampling is 0.01; while for the ATLAS experiments, β is set to 250 and the score norm is 0.0001. The three repeats for each ATLAS protein are merged with equal weight and randomly sampled during training. The denoising step is both set to 50, and the learning rate is both set to 1e-5. We use grid search to select the above hyperparameters, and the sensitivity analysis can be found in ablation studies. We adopt Madrax (Orlando et al. 2024), a differentiable empirical force field implemented in PyTorch, as the physics-based energy reward. Low-Rank Adaptation (LoRA) (Hu et al. 2022) is employed for parameter-efficient fine-tuning of pretrained models.

Training the EPO model on the Tetrapeptides dataset requires 30 hours, while the ATLAS dataset requires 72 hours. This process utilizes four A100 40GB GPUs on a Linux 5.4.0 platform with PyTorch 1.12 and CUDA 11.3. All reported metrics are from a single experimental run, a protocol consistent with the baseline models we compare against.

Tetrapeptides

Table 3 assesses the Jensen-Shannon divergence (JSD) between the generated and ground-truth trajectories across two sets of collective variables: (1) the individual backbone and sidechain torsion angles (χ_{bb} , χ_{sc} , χ_{all}) for each tetrapeptide; (2) the leading independent components extracted via time-lagged independent component analysis (TICA), capturing the slowest dynamical modes of peptides. EPO consistently achieves superior distributional alignment with the reference MD trajectories, closely matching the accuracy obtained by replicate 100-ns simulations.

We further visualize the torsion angle distributions and




Figure 4: Sample distributions projected onto the first two time-lagged independent components for four proteins from the Fast-Folding dataset. EPO-List demonstrates improved diversity by exploring a broader conformational landscape.

the free energy surface (FES) along the two primary TICA components in Figure 3. The torsion angle distribution plots clearly demonstrate that EPO-generated trajectories (blue) closely approximate the ground-truth MD trajectories (orange). Notably, the FES analysis reveals that EPO not only induces novel metastable states absent in the pretrained baseline (in ASRE) but also effectively corrects high-energy biases present in the pretrained model (in KSIY). These results validate our central hypothesis: EPO’s online sampling framework, combined with energy-based preference optimization, provides an effective and reliable mechanism for crossing critical energy barriers, thus enabling the accurate recovery of physically realistic conformational states.

ATLAS

Next, we evaluate EPO on the 82 test proteins in ATLAS. We report the mean and median values of the distributional metrics to quantify discrepancies between the generated and reference MD ensembles. For detailed definitions and descriptions of the evaluation metrics, please refer to Appendix. As shown in Table 1, EPO exhibits advantages over baseline methods in several key distributional metrics, including higher correlations in pairwise RMSD and per-target RMSF, as well as lower PCA \mathcal{W}_2 distances.

We attribute the modest performance gains on the ATLAS benchmark primarily to the small score norm used in SDE sampling for ATLAS proteins. Larger score norms caused




Figure 5: Analysis of key components of EPO. (A) Impact of the temperature hyperparameter β . χ represents torsion angle. (B) Effect of the number of denoising steps in SDE sampling. An insufficient number of steps causes training instability, whereas an excessive number provides no significant improvement. (C) Torsion angle distributions for SDE vs. ODE strategies on the SPFH sequence. The SDE-based model successfully identifies distant peaks.

	χ_{bb}	χ_{sc}	χ_{all}	TICA-0	TICA-0,1
EPO-ODE	0.167	0.107	0.132	0.247	0.337
EPO-SDE	0.125	0.093	0.107	0.226	0.311

Table 4: Comparison between SDE and ODE sampling strategy on tetrapeptides experiments. Lower is better.

sampling instability and structural collapse, making this adjustment necessary when applying SDE to longer protein chains. However, the reduced stochastic exploration likely limited the model’s ability to discover a more diverse and accurate conformational ensemble, underscoring an important direction for future research.

Fast-Folding

We further evaluate the transferability and generalization of EPO by performing zero-shot testing on the Fast-Folding protein benchmark. Specifically, we directly apply the checkpoint optimized on the ATLAS dataset to generate ensembles without any fine-tuning or adaptation.

Table 2 summarizes the performance of EPO-generated conformational ensembles in comparison to several established baseline methods. We employ the evaluation metrics introduced in ConfDiff (Wang et al. 2024), where detailed explanations can be found in Appendix.

EPO establishes four new state-of-the-art results and achieves second-best performance on the remaining distributional metrics. Notably, comparing with closely related baselines, i.e., Str2Str and ConfDiff, highlights the core advantages of EPO. Unlike Str2Str, EPO’s online exploration is effectively steered by an explicit physical signal. And unlike ConfDiff, EPO’s physical guidance is applied dynamically in real-time, avoiding the constraints and biases of training on a static, offline dataset. This synergy of online guidance and online exploration is what allows EPO to continuously discover novel and physically favorable states.

Ablation Studies

In this section, we investigate and validate EPO’s sensitivity to key components.

First, we examine the temperature hyperparameter β , critical for scaling the preference objective. As shown in Figure 5(A), setting β to a very small value is detrimental and can be counterproductive. This highlights that β ’s optimal value is context-dependent and must be carefully selected.

Next, we analyze the effect of denoising steps on SDE sampling. Figure 5(B) shows that insufficient steps cause unstable training and mode collapse, as the reverse trajectory fails to generate valid conformations. Conversely, increasing the number of steps beyond a certain point yields diminishing returns; it significantly increases computational cost without a corresponding improvement in sample quality or model performance. This highlights the need for a balance between stability and tractability.

Finally, we compare the exploratory power of stochastic (SDE) versus deterministic (ODE) sampling. Figure 5(C) presents torsion angle distributions on the SPFH sequence for both. The SDE approach shows a significant advantage, identifying distant peaks corresponding to distinct metastable states separated by high energy barriers. In contrast, the ODE sampler converges to limited local minima, failing to discover these crucial conformations. More distributional visualizations can be found in the Appendix.

Conclusions and Limitations

In this work, we have presented Energy Preference Optimization (EPO), an online refinement framework that aligns pretrained protein-conformation generators with the Boltzmann distribution using direct physics-based feedback. By combining an online SDE sampling strategy with a listwise preference optimization, EPO successfully produces diverse and physically realistic ensembles while eliminating the need for additional molecular dynamics simulations.

Two primary limitations motivate future research. First, the computationally intensive online refinement loop necessitates more efficient protocols. Coarse-grained models may reduce this cost. Second, our listwise preference objective, while empirically strong, lacks theoretical guarantees of convergence to the Boltzmann distribution, motivating investigation into optimization frameworks with provable convergence. Looking ahead, we plan to scale EPO to more expressive generative models for large proteins, and leverage engineering technologies in recent RLHF-style frameworks.

Acknowledgements

This work is supported by the Administrative Committee of Zhongguancun Science City. This work is also supported by Beijing Natural Science Foundation (L243006).

References

- Albergo, M. S.; and Vanden-Eijnden, E. 2023. Building Normalizing Flows with Stochastic Interpolants. In *The Eleventh International Conference on Learning Representations*.
- AlRawashdeh, S.; and Barakat, K. H. 2023. Applications of molecular dynamics simulations in drug discovery. *Computational Drug Discovery and Design*, 127–141.
- Badar, M. S.; Shamsi, S.; Ahmed, J.; and Alam, M. A. 2022. *Molecular dynamics simulations: concept, methods, and applications*. Springer.
- Del Alamo, D.; Sala, D.; Mchaourab, H. S.; and Meiler, J. 2022. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *Elife*, 11: e75751.
- Ghahremanian, S.; Rashidi, M. M.; Raeisi, K.; and Toghraie, D. 2022. Molecular dynamics simulation approach for discovering potential inhibitors against SARS-CoV-2: A structural review. *Journal of molecular liquids*, 354: 118901.
- Guo, D.; et al. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948.
- Han, J.; Jiang, M.; Song, Y.; Ermon, S.; and Xu, M. 2025. *f*-PO: Generalizing Preference Optimization with *f*-divergence Minimization. In *International Conference on Artificial Intelligence and Statistics*, 1144–1152. PMLR.
- Hénin, J.; Lelièvre, T.; Shirts, M.; Valsson, O.; and Delemotte, L. 2022. Enhanced sampling methods for molecular dynamics simulations. *Living Journal of Computational Molecular Science*, 4(1).
- Henzler-Wildman, K.; and Kern, D. 2007. Dynamic personalities of proteins. *Nature*, 450(7172): 964–972.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Jiao, R.; Kong, X.; Huang, W.; and Liu, Y. 2024. 3D Structure Prediction of Atomic Systems with Flow-based Direct Preference Optimization. *Advances in Neural Information Processing Systems*, 37: 110197–110217.
- Jing, B.; Berger, B.; and Jaakkola, T. 2024. AlphaFold meets flow matching for generating protein ensembles. In *Proceedings of the 41st International Conference on Machine Learning*, 22277–22303.
- Jing, B.; Erives, E.; Pao-Huang, P.; Corso, G.; Berger, B.; and Jaakkola, T. S. 2023. EigenFold: Generative Protein Structure Prediction with Diffusion Models. In *ICLR 2023-Machine Learning for Drug Discovery workshop*.
- Jing, B.; Stark, H.; Jaakkola, T.; and Berger, B. 2024. Generative Modeling of Molecular Dynamics Trajectories. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; et al. 2021. Highly accurate protein structure prediction with AlphaFold. *nature*, 596(7873): 583–589.
- Kalakoti, Y.; and Wallner, B. 2025. AFsample2 predicts multiple conformations and ensembles with AlphaFold2. *Communications Biology*, 8(1): 373.
- Klein, L.; and Noe, F. 2024. Transferable Boltzmann Generators. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Lanchantin, J.; Chen, A.; Dhuliawala, S.; Yu, P.; Weston, J.; Sukhbaatar, S.; and Kulikov, I. 2025. Diverse Preference Optimization. arXiv:2501.18101.
- Lewis, S.; Hempel, T.; Jiménez-Luna, J.; Gastegger, M.; Xie, Y.; Foong, A. Y.; Satorras, V. G.; Abdin, O.; Veeling, B. S.; Zaporozhets, I.; et al. 2025. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science*, eadv9817.
- Li, X.; Zhao, Y.; Wang, C.; Scalia, G.; Eraslan, G.; Nair, S.; Biancalani, T.; Ji, S.; Regev, A.; Levine, S.; and Uehara, M. 2024. Derivative-Free Guidance in Continuous and Discrete Diffusion Models with Soft Value-Based Decoding. arXiv:2408.08252.
- Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; and Shaw, D. E. 2011. How fast-folding proteins fold. *Science*, 334(6055): 517–520.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Liu, J.; Liu, G.; Liang, J.; Li, Y.; Liu, J.; Wang, X.; Wan, P.; Zhang, D.; and Ouyang, W. 2025. Flow-GRPO: Training Flow Matching Models via Online RL. arXiv:2505.05470.
- Liu, T.; Qin, Z.; Wu, J.; Shen, J.; Khalman, M.; Joshi, R.; Zhao, Y.; Saleh, M.; Baumgartner, S.; Liu, J.; et al. 2024. Lipo: Listwise preference optimization through learning-to-rank. *arXiv preprint arXiv:2402.01878*.
- Lu, J.; Zhong, B.; Zhang, Z.; and Tang, J. 2023. Str2Str: A Score-based Framework for Zero-shot Protein Conformation Sampling. In *The Twelfth International Conference on Learning Representations*.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, 23–40. Springer.
- Nussinov, R. 2016. Introduction to protein ensembles and allostery.
- Nussinov, R.; Liu, Y.; Zhang, W.; and Jang, H. 2023. Protein conformational ensembles in function: roles and mechanisms. *RSC chemical biology*, 4(11): 850–864.
- Orlando, G.; Serrano, L.; Schymkowitz, J.; and Rousseau, F. 2024. Integrating physics in deep learning algorithms: a force field as a PyTorch module. *Bioinformatics*, 40(4): btae160.

- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- Raisinghani, N.; Alshahrani, M.; Gupta, G.; Tian, H.; Xiao, S.; Tao, P.; and Verkhivker, G. 2024. Probing Functional Allosteric States and Conformational Ensembles of the Allosteric Protein Kinase States and Mutants: Atomistic Modeling and Comparative Analysis of AlphaFold2, OmegaFold, and AlphaFlow Approaches and Adaptations. *The Journal of Physical Chemistry B*, 128(45): 11088–11107.
- Ray, D.; Ansari, N.; Rizzi, V.; Invernizzi, M.; and Parrinello, M. 2022. Rare event kinetics from adaptive bias enhanced sampling. *Journal of Chemical Theory and Computation*, 18(11): 6500–6509.
- Schafer, J. W.; Lee, M.; Chakravarty, D.; Thole, J. F.; Chen, E. A.; and Porter, L. L. 2025. Sequence clustering confounds AlphaFold2. *Nature*, 638(8051): E8–E12.
- Souza, P. C.; Alessandri, R.; Barnoud, J.; Thallmair, S.; Faustino, I.; Grünewald, F.; Patmanidis, I.; Abdizadeh, H.; Bruininks, B. M.; Wassenaar, T. A.; et al. 2021. Martini 3: a general purpose force field for coarse-grained molecular dynamics. *Nature methods*, 18(4): 382–388.
- Stein, R. A.; and Mchaourab, H. S. 2022. SPEACH_AF: Sampling protein ensembles and conformational heterogeneity with AlphaFold2. *PLOS Computational Biology*, 18(8): e1010483.
- Stevens, J. A.; Grünewald, F.; van Tilburg, P.; König, M.; Gilbert, B. R.; Brier, T. A.; Thornburg, Z. R.; Luthey-Schulten, Z.; and Marrink, S. J. 2023. Molecular dynamics simulation of an entire cell. *Frontiers in Chemistry*, 11: 1106495.
- Team, O.; and Ye, Q. 2025. Towards Unraveling Biomolecular Conformational Landscapes with a Generative Foundation Model. *bioRxiv*, 2025–05.
- Teixeira, J. M.; Liu, Z. H.; Namini, A.; Li, J.; Vernon, R. M.; Krzeminski, M.; Shamandy, A. A.; Zhang, O.; Haghghat-lari, M.; Yu, L.; et al. 2022. IDPCConformerGenerator: a flexible software suite for sampling the conformational space of disordered protein states. *The Journal of Physical Chemistry A*, 126(35): 5985–6003.
- Uehara, M.; Zhao, Y.; Wang, C.; Li, X.; Regev, A.; Levine, S.; and Biancalani, T. 2025. Inference-Time Alignment in Diffusion Models with Reward-Guided Generation: Tutorial and Review. *arXiv:2501.09685*.
- Vander Meersche, Y.; Cretin, G.; Gheeraert, A.; Gelly, J.-C.; and Galochkina, T. 2024. ATLAS: protein flexibility description from atomistic molecular dynamics simulations. *Nucleic acids research*, 52(D1): D384–D392.
- Wallace, B.; Dang, M.; Rafailov, R.; Zhou, L.; Lou, A.; Purushwalkam, S.; Ermon, S.; Xiong, C.; Joty, S.; and Naik, N. 2024. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8228–8238.
- Wang, C.; Jiang, Y.; Yang, C.; Liu, H.; and Chen, Y. 2023. Beyond Reverse KL: Generalizing Direct Preference Optimization with Diverse Divergence Constraints. In *The Twelfth International Conference on Learning Representations*.
- Wang, J.; Arantes, P. R.; Bhattacharai, A.; Hsu, R. V.; Pawnikar, S.; Huang, Y.-m. M.; Palermo, G.; and Miao, Y. 2021. Gaussian accelerated molecular dynamics: Principles and applications. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 11(5): e1521.
- Wang, Y.; Wang, L.; Shen, Y.; Wang, Y.; Yuan, H.; Wu, Y.; and Gu, Q. 2024. Protein Conformation Generation via Force-Guided SE (3) Diffusion Models. In *International Conference on Machine Learning*, 56835–56859. PMLR.
- Wayment-Steele, H. K.; Ojoawo, A.; Otten, R.; Apitz, J. M.; Pitsawong, W.; Hömberger, M.; Ovchinnikov, S.; Colwell, L.; and Kern, D. 2024. Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature*, 625(7996): 832–839.
- Xia, F.; Liu, T.-Y.; Wang, J.; Zhang, W.; and Li, H. 2008. Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th international conference on Machine learning*, 1192–1199.
- Xue, Z.; Wu, J.; Gao, Y.; Kong, F.; Zhu, L.; Chen, M.; Liu, Z.; Liu, W.; Guo, Q.; Huang, W.; and Luo, P. 2025. DanceGRPO: Unleashing GRPO on Visual Generation. *arXiv:2505.07818*.
- Yim, J.; Trippe, B. L.; De Bortoli, V.; Mathieu, E.; Doucet, A.; Barzilay, R.; and Jaakkola, T. 2023. SE (3) diffusion model with application to protein backbone generation. In *Proceedings of the 40th International Conference on Machine Learning*, 40001–40039.