

Zero-Shot Generation of Protein Conformational Ensembles Through AlphaFold Latent Flooding

Jing Huang

huangjing@westlake.edu.cn

Westlake University <https://orcid.org/0000-0001-9639-2907>

Runtong Qian

Westlake University <https://orcid.org/0009-0005-8743-4833>

Rui Zhan

Westlake University <https://orcid.org/0009-0005-5055-7610>

Zilin Song

Westlake University <https://orcid.org/0000-0002-3127-3971>

Article

Keywords:

Posted Date: June 5th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9485266/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Zero-Shot Generation of Protein Conformational Ensembles Through AlphaFold Latent Flooding

 Runtong Qian^{1,2,3}

 Rui Zhan^{2,3}

 Zilin Song^{2,3,*}

 Jing Huang^{2,3,*,‡}

¹*School of Life Sciences, Fudan University, Shanghai 200433, China;* ²*State Key Laboratory of Gene Expression, School of Life Sciences, Westlake University, Hangzhou, Zhejiang 310030, China;* ³*Westlake AI Therapeutics Lab, Westlake Laboratory of Life Sciences and Biomedicine, Hangzhou, Zhejiang 310024, China.*

*Correspondence: songzilin@westlake.edu.cn (Z.S.) and huangjing@westlake.edu.cn (J.H.).

‡Lead contact.

Abstract

Despite transformative advances in protein structure prediction, generating conformational ensembles directly from sequence in an efficient and accessible manner remains a central challenge. We introduce a heuristic importance sampling framework for the zero-shot prediction of functionally relevant, conformationally diverse structures, bypassing the need for extensive physics-based modeling or prior domain knowledge. Guided by an analysis of massive activations in AlphaFold (AF), we develop a latent flooding algorithm that enables adversarial exploration of the latent space while preserving local geometric integrity. We show that the AF latent flooding (AFLF) method robustly recovers experimental structural fluctuations, captures function-relevant conformational states, and reveals cryptic binding sites across diverse protein systems. These results suggest that AF latent features implicitly encode the biophysical principles of protein thermodynamics, which AFLF exploits. Finally, we show that AFLF is computationally accessible and interoperable, offering a test-time generalization of AF for investigating protein dynamics and accelerating structure-based ligand discovery.

1 Introduction

Deep learning has catalyzed seminal advances in computational structural biology by facilitating accurate protein structure predictions and capturing essential biophysical interactions. Exemplified by the AlphaFold (AF) series of models and related methods,^{1–6} these breakthroughs have delivered the predictions of static protein structures and their complexes with unprecedented accuracy.⁷ However, proteins are inherently dynamic entities that often adopt multiple conformational states in response to environmental perturbations or molecular interactions, a crucial aspect of their biological function.⁸

While methods based on neural networks have proven highly effective in predicting static, folded protein structures, they often struggle to enumerate the complete spectrum of functionally relevant states.⁹ This limitation significantly hinders the applicability of structural predictors like AF for crucial downstream tasks such as mechanistic investigation and allosteric ligand discovery.¹⁰ Recent advances have sought to bridge this gap by incorporating large-scale molecular dynamics (MD) simulations or leveraging deep generative methods to explore hidden protein conformations.^{11–13} Yet, these emerging approaches in general require substantial computational resources to generate labeled data and train models. The nontrivial challenge remains the development of methods that can reliably generate diverse, biologically relevant structural ensembles solely from the primary sequence, while preserving the predictive accuracy and computational accessibility that have made biological foundation models like AF transformative.¹⁴

Feature perturbation strategies, including subsampling or clustering the multi-sequence alignments (MSAs), template manipulations, and test-time interventions, have nudged the AF models into producing alternative, biologically plausible conformations without retraining.^{15–20} Although approaches as such are demonstrably capable of expanding the conformational landscape accessible to a single sequence input, the stacked neural transformations render the underlying processes opaque to human perception. Consequently, the absence of mechanistic explanations for why and how such perturbation steers the network to predict new functional states precludes the conversion of AF latent space into an enumerable conformational repository, and its learned structural landscape is left beyond systematic exploration.

Here, we introduce AF Latent Flooding (AFLF), a heuristic multiscale framework that repurposes the AF2 model as a zero-shot engine for generating protein conformational ensembles. We dissect how the perturbations on the AF2 learned latent representations propagate into diverse structural outputs. We assemble key AF2 components to enable direct latent inference and integrate Low-Rank Adaptation (LoRA) for efficient protein structural enumeration. We develop a self-repelling, self-adaptive sampler to efficiently explore the broad yet relevant conformational space embedded in the AF2 latent landscape. Through applications including quantification of structural fluctuations, prediction of large-scale conformational changes, and detection of cryptic pockets, we demonstrate that AFLF enables highly accessible ensemble generation, extending the AF foundation models into readily deployable assets for structure-based ligand discovery and related downstream applications.

2 Results

2.1 Massive activations in AF latent representations

To probe the latent space in the AF2 model, we modified its inference pipeline to retrieve the latent features: the MSA/pair tensors fed to the Evoformer stack and the single/pair tensors fed to the structural module (StructMod). Interestingly, all four latent tensors exhibited the same “massive activation” pattern reported in transformer-based large language models (LLMs),²¹ where a minute fraction of tensor elements deviates by orders of magnitude from the median (**Figure 1a**). Empirically, the massively activated elements in general encode input-agnostic attention biases that dominate layer-wise attention probabilities, yet their magnitudes are strictly required for model performance. Recent works have shown that protein language models (PLMs) and LLMs adopt similar patterns of geometry in their latent representations.^{22,23} We therefore expect that the massive activations in AF2 serve an analogous function, as the Evoformer stack, which is trained with a masked language (MSA) modeling objective, effectively renders the model also a PLM.

To delineate the contribution of the massively activated elements, we systematically ablated the four latent tensors. We flagged the tensor elements whose values fell within the extreme tails (top and bottom fractions) of the overall distribution as massively activated and subjected them to three perturbation schemes. Briefly, scheme1 resampled the flagged entries from a Gaussian distribution fitted to the full-tensor statistics and scheme2 zero-masked the same flagged entries. The reference scheme0 masked an equal number of randomly chosen entries to provide a baseline for the information loss. Across all four latent tensors and repeated random seeds, the targeted perturbations of the massive activations (scheme1 and scheme2) produced statistically significant structural deviations from the reference (scheme0, **Figure 1b**). A recent study on a community reproduction of AF2²⁴ showed that global zero-masking of the pair representations severely depresses the template-modeling score on the predicted structures.²⁵ As the geometric learning objectives of AF2 were imposed almost exclusively on the Structural Module, it is intuitive that its pair representations primarily encode the latent manifold of structural topologies. Here, we further show that corrupting only the massive pair activations, either in the Evoformer or in the StructMod, was sufficient to cause mode-collapse in the predicted coordinates, whereas distributing the same number of zeros at random preserved topologically plausible predictions (**Figures S1, S2, S3**). In contrast, perturbing the massive Evoformer MSA activations warped the global fold and softened the local secondary structures, without precipitating structural collapse, and the reference scheme0 that randomly masked identical numbers of tensors elements maintained the overall fold architecture. Moreover, perturbations on the StructMod single representation, regardless of schemes, produced no discernible difference in the predicted structure. A similar mode of action is also observed for larger proteins (**Figures S4, S5, S6, S7**).




Figure 1. Latent features distribution and ubiquitin structural deviations from perturbed latent features. From left to right: Evoformer MSA, Evoformer pair, structural module (StructMod) single, and StructMod pair representations. **(a)** Permillage value distributions of the AF2 latent tensor elements used to predict ubiquitin conformations. The 1st and 99th percentiles are indicated and n is the total number of elements in each tensor. **(b)** Backbone root mean squared deviations ($\log_{10}$ RMSDs) of ubiquitin conformations predicted from perturbed latent tensors relative to the native unperturbed prediction. Statistical tests: scheme0 ($n = 5$) versus scheme1 ($n = 5$) by independent samples t -test; scheme0 ($n = 5$) versus scheme2 ($n = 1$) by one-sample t -test. ***: $p < 0.001$; **: $p < 0.01$; *: $p < 0.05$; ns: not significant.

Collectively, the latent tensors in AF2, like those of LLMs, contain a small subset of massively activated elements. Across the four latent representations, the massively activated entries serve distinct functions in the structural inference pipeline: corrupting them in the two pair tensors collapses the predicted coordinates; in the Evoformer MSA tensor it reshapes the global fold while preserving local topology; and in the StructMod single tensor it produces no measurable difference. Our perturbation experiments suggest that AF2 embeds the mechanism of determining specific folding patterns within the Evoformer MSA massive activations. In parallel, the null effect of perturbing the StructMod single representations indicates that the Evoformer stack progressively distills coevolutionary information from the MSA representation into the pair representation, such that, once the distillation is complete, the residual first-row MSA tensor (the single representation) becomes too uninformative to impact the structural output.

2.2 Latent Flooding for Protein Conformation Generation

To access the embedded protein structural state in the AF2 latent space, based on the above analysis, numerical interventions should preferably target the Evoformer MSA representations whose extreme valued elements constitute the decisive mechanism of folding. Our preliminary experiments indicated that the latent tensors cannot be universally partitioned into the massively activated and the residual subsets as the layered nonlinear mixing couples their contributions; therefore, the entire Evoformer MSA tensor was adapted to ensure system-agnostic behavior. The AFLF modeling proceeds in two consecutive stages: (1) AFLF checkpoint modeling, in which the AF2 embedding module fuses the MSA and the template input features into the latent MSA and pair representations that are cached via a modified AF2 inference pipeline; and (2) AFLF inference modeling, where we initialize the Low-Rank Adaptation (LoRA)²⁶ tensors on the cached checkpoints and propagate the augmented representations through the Evoformer and the StructMod stacks to generate new structures (**Figure 2a**). We note that the identical latent inference mechanism can also be applied to any intermediate layer as the checkpoint.

To explore the protein conformational landscape latent in AF, the AFLF inference modeling implements a self-repelling, self-adaptive enhanced sampling protocol (see **Methods**). The AFLF method characterizes the protein conformation by the inter-centroid distances between predefined atom groups, thereby providing a conformational space of reduced dimensions while retaining essential geometric relations. The AFLF sampler adopts a memory-guided strategy that maintains a running stack of previously visited repellent states which derives a series of Gaussian repelling potentials to penalize revisiting previous conformations




Figure 2. AlphaFold Latent Flooding (AFLF) model. **(a)** The AFLF model. **(b)** The repelling and the anchoring losses. The repelling loss drives the latent state away from previously visited conformations; The anchoring loss penalizes excessive deviation of repelled coordinates from their reference distances. **(c)** The local geometric loss enforces rigid-body restraints on conserved motifs via optimally aligned RMSDs. **(d)** The global geometric loss regulates protein-scale plasticity by minimizing the cross-entropy between reference and predicted inter-residue distance distributions in kernel space.

(Figure 2b). To overcome the non-ergodicity inherent to naïve self-repelling walks in the AF latent space, AFLF implements an adaptive importance sampling scheme based on the coefficient of variations tracked along the trajectory. The adaptive coefficients scale inversely with the variation statistics: distant pairs exhibiting low variability are deemed insufficiently sampled and are assigned higher repulsive weights, whereas high variability distances with adequate sampling are downscaled. This self-adaptive mechanism dynamically redirects online the sampling effort toward underexplored regions of the latent conformational space.

To constrain the latent exploration on the manifold of native-like folds, AFLF incorporates a multiscale geometric regularization strategy comprising three complementary loss terms. The anchoring loss applies boundary penalties to prevent excessive deviations of repelled coordinates from their reference (Figure 2b). The local geometric loss enforces the structural integrity on essential motifs (such as disulfide bonds and proline rings) via optimally aligned root mean squared deviations (RMSDs, Figure 2c). The global geometric loss controls the protein-scale structural plasticity through the cross-entropy between the reference (defaulting to the AF naïve prediction) and predicted inter-residue distance distributions (Figure 2d).

Conceptually, the online modulation of AFLF dynamics by the sampling statistics parallels the conformational flooding²⁷ approach, in which Gaussian biases are planted along slow collective coordinate variations to accelerate classical MD sampling. The AFLF continuously rescales the repulsive potentials according to the instantaneous sampling heterogeneity of the inter-centroid distances, thereby implementing a memory-guided adaptive enumeration of the AF latent manifold. The multiscale geometric regularization schemes (anchoring, local motif rigidity, and global distance distribution entropy) are fused with the adaptive repelling potential into a single differentiable objective. The gradient flow is realized by low-rank adaptation of the Evoformer MSA checkpoint, enabling direct exploration of the latent space without retraining the full network. Equipped with the AFLF protocol, we next apply it to reproduce experimental structural fluctuations, to explore alternative functional states, and to identify cryptic binding sites.

2.3 Generating Protein Structural Fluctuations

To validate AFLF for reproducing experimental structural fluctuations, we selected ubiquitin, a model protein whose high-resolution crystal structure exhibits a well-characterized flexibility gradient from a rigid core to a highly mobile C-terminus. The AFLF sampling was initialized from the latent state produced by the native AF prediction and was advanced for 20,000 steps with Gaussian flooding applied on all inter-centroid distances between residues separated by at least four sequential positions, yielding a continuous conformational trajectory (Movie S1). For both the MD baseline and the AFLF trajectory, we computed per-residue *B*-factors based on root mean square fluctuations (RMSFs).²⁸

Remarkably, the structural flexibility of ubiquitin captured by AFLF reproduced the crystallographic *B*-factors at the same amplitude and preserved the ranking correlation as faithfully as the conventional MD reference: Kendall $\tau = 0.46$ (~73% concordant pairs) for AFLF versus $\tau = 0.49$ (~74.5% concordant




Figure 3. Ubiquitin structural fluctuations generated by AFLF. **(a)** The B -factors from the crystal structure (PDB id: 1UBQ), the baseline MD simulation, and the AFLF trajectory. The **(b)** AFLF-derived and **(c)** experimental ubiquitin structural fluctuations colored by low (pink) and high (purple) B -factors.

pairs) for MD (**Figure 3a**, **Table S1**). Visual inspection of the residue-level flexibility maps indicated that the AFLF ensemble reconstructed the experimentally observed pattern of rigid (purple) and mobile (pink) regions (**Figures 3b**, **3c**). Importantly, the agreement emerges without reweighting or recalibration, showing that AFLF, as a zero-shot AF approach, can deliver experimentally consistent structural dynamics.

2.4 Generating Protein Functional States

To evaluate AFLF for exploring large-scale conformational transitions, we focused on adenylate kinase (AdK), a canonical two-state model enzyme whose catalytic cycle involves the opening of the ATP-binding domain (LID, residues 121–159) and the nucleotide-monophosphate-binding domain (NMP, residues 30–59) with respect to the structurally rigid CORE domain (remaining residues). Starting from the AF-predicted closed conformation, we initiated a 10,000-step AFLF simulation after decomposing the protein into segments defined by the secondary structure elements: each α -helix, β -strand, and loop longer than six residues was treated as an independent segment. The AFLF flooding potentials were applied exclusively on the segment pairs whose minimal inter-residue distance fell below 4.5 Å and the non-contacting distances were left unbiased. The AFLF anchoring thresholds were set to -1 to $+25$ Å for contacting segment pairs, except for intra-domain pairs and those involving the CORE domain, which were restrained within -1 to $+1$ Å. Referring to the crystallographic closed and open structures, we projected the AFLF trajectory onto the two-dimensional LID-CORE/NMP-CORE plane for quantifying the AdK closed $\leftrightarrow$ open transitions.²⁹




Figure 4. Adenylate kinase (AdK) conformational transitions generated by AFLF. **(a)** The AFLF trajectory projected on the two dimensional plane of the LID-CORE angle and the NMP-CORE angle. The experimental closed (PDB id: 1AKE) and open (PDB id: 4AKE) structures are shown. The representative AFLF generated conformations (pink) for **(b)** the closed and **(c)** the open states are superposed with the experimental crystal structures (gray). **(d)** The AFLF-generated trajectory projected on the LID-NMP centroid distance.

Notably, the AFLF ensemble populated a contiguous density of conformations between the closed and open crystallographic structures (**Figure 4a**, **Movie S2**). The representative snapshots recovered both crystal structures with the C α RMSDs being 1.18 Å (closed, **Figure 4b**) and 3.06 Å (open, **Figure 4c**). Along the simulated trajectory, we recorded nine closed $\leftrightarrow$ open transition events whose LID-NMP centroid distance traversed the full experimental interval (**Figure 4d**). The resulting structural distribution confirmed that the AFLF ensemble autonomously captured both native endpoints and interpolated the intermediate transient space that connects them.

2.5 Generating Protein Cryptic Cavities

To assess AFLF for detecting cryptic cavities, we targeted five proteins whose cryptic binding sites have been experimentally documented (**Table S2**):³⁰ the class A β -lactamase TEM-1, the anti-apoptotic protein Bcl-x_L, the bovine β -lactoglobulin (β LG), the human nucleoside diphosphate kinase A (NDPK-A), and the human glycolipid transfer protein (GLTP). In all cases, the apo crystal structure captures the cryptic cavity in the occluded state and the exposed conformations emerge only under mutational or ligand-induced perturbations. A recent benchmark has shown that the native AF predictions capture both the occluded and the exposed cryptic binding site conformations only when the training data carries balanced populations of each,³¹ suggesting that the native AF inference pipeline is constrained to recapitulate the structural distributions already embedded in its learned repertoire. Here, we demonstrate that the AFLF approach effectively broadens the predicted conformational landscape and robustly predicts both the occluded and the exposed states without physics-based refinement. Similar to the AdK case, we began from the AF native predictions, partitioned the proteins into secondary structure segments, activated the AFLF flooding potentials for contacting segment pairs, and retained the unbiased treatment of non-contacting pairs. The anchoring bounds were uniformly set to -1 Å and $+10$ Å. Using the crystallographically resolved conformations as structural references, we monitored the volume of the cryptic cavities throughout the 10,000-step AFLF trajectories for all five test proteins.




Figure 5. Cryptic cavity profiles generated by AFLF. **(a)** The cavity volume distributions of the AFLF generated conformations; **(b)** The backbone C α RMSD profiles to the reference apo and holo conformations; The representative AFLF generated structures of the class A β -lactamase TEM-1 for **(c)** the apo state and the holo state that exposes H11 cryptic binding site. The minimum backbone C α RMSDs between the AFLF generated conformations (pink) and the experimental crystal structures (proteins in gray and ligands in purple) are listed. **(d)** The AFLF generated TEM-1 conformation (pink) that exposes the Ω loop cryptic binding site (purple arrow), the apo state TEM-1 crystal structure (gray), and their cavity volumes. See also **Figure S8**.

Conspicuously, the AFLF ensembles spanned the full volumetric space between the apo-occluded and the holo-exposed conformations in all five test proteins (**Figure 5a**), indicating that the model did not merely recognize the two endpoint states but also captured the intervening transient states. In parallel, the C α RMSDs relative to both the apo and the holo states consistently remained below ~ 3 Å (**Figure 5b**), indicating that the sampled conformations preserved native folds. At the same time, the trajectories effectively accessed conformational states in which the cryptic pockets are fully exposed. Detailed inspection of the AFLF snapshots that best superimposed on the crystallographic conformations confirmed that the model recovered both states without being preconditioned on the ligand moiety (**Figure 5c, 5d, S8**). The minimal structural deviations (C α RMSDs < 2 Å) and the close resemblance of the optimally aligned AFLF conformations to the holo states demonstrated that the method produced apo-exposed conformations that are immediately exploitable for structure-based ligand discovery.

Among the five systems, TEM-1 is an important therapeutic target for combating antimicrobial resistance³², and its cryptic binding sites remain challenging to identify using either large-scale MD simulations or AF-based workflows^{31,33,34}. TEM-1 has two reported cryptic binding sites. The H11 site, located between helices H11 and H12, has been structurally characterized in its exposed form (**Figure 5c**). In contrast, the site proximal to the Ω loop has not yet been experimentally observed in an exposed conformation, despite functional evidence supporting its role in allosteric regulation³³. AFLF successfully identified both cryptic sites. For the H11 site, the majority of the sampled conformations did not contain this cavity, while full opening was captured as a rare event along the AFLF trajectory. For the Ω loop site, AFLF sampled an exposed state (**Figure 5d**), with the cavity volume increasing from 287.07 Å³ to 846.98 Å³. Together, these results underscored AFLF’s ability to extrapolate plausible functional conformations beyond the training data based on experimentally resolved structures.

3 Discussions

The unprecedented progress made in deep-learning protein structure prediction has culminated in biological foundation models that deliver coordinates at atomic resolution directly from the primary sequence. However, the algorithms of single-state predictors such as AF were optimized to reproduce the most probable conformations in the protein structural database and not the geometrical heterogeneity underlying catalysis, allostery, or ligandability. Here, we provide a democratized and interoperable ensemble generation protocol AFLF that autonomously explores the structural landscapes encrypted within the AF latent feature space.

To dissect how the latent dynamics driven by AFLF evolves the massive activations, we traced every latent state visited along the AdK trajectory. For each AFLF step, we extracted the (top and bottom) 1% most extreme entries in the Evoformer MSA representation, computed the Jaccard index (JI) between each pair of latent states, and embedded the latent trajectory using the precomputed JI distance matrix with the Uniform Manifold Approximation and Projections (UMAP) method (**Figure 6**). Although the AdK backbone repeatedly visited the open and the closed conformations, the UMAP embedding revealed a unidirectional path on which no cyclic pattern was observed. The same monotonic trend was recovered when the optimal-transport distances (d_{OT}) between each latent state and the initial state were plotted against the AFLF time, with d_{OT} increasing almost linearly along the AFLF trajectory (**Figure S9**). Evidently, the identity and the magnitude of the massive activations continuously evolve throughout the AFLF dynamics whereas the structural output can revert to an earlier macroscopic state. The many-to-one mapping of the latent states to the structural outputs justifies our choice of operating LoRA on the entire MSA tensor rather than selectively on the initial massive activation set. In addition, the dynamical pattern also explains why earlier heuristic approaches – such as MSA subsampling or clustering – succeed only sporadically: these methods rely on manipulating the discrete MSA inputs and lack the continuously tunable latent coordinates to systematically traverse the AF latent manifold of topologically accessible ensemble of conformations. AFLF, by contrast, bypasses MSA curation altogether and re-tunes its perturbations in response to intermediate sampling feedback, providing a stable, on-the-fly enumeration through the latent conformational ensemble.

Technically, three lines of evidence collectively establish the mechanistic validity and the empirical effectiveness of the AFLF protocol. First, the systematic ablation of the four AF latent tensors reveals that the massively activated elements in the Evoformer MSA representation act as the decisive fold maker, whereas corruption of the pair tensors precipitates coordinate collapse and perturbation of the StructMod single tensor is effectively silent. Second, the AFLF flooding potentials are modulated through the statistical variation




Figure 6. AFLF latent dynamics of adenylate kinase (AdK) projected on the two-dimensional Uniform Manifold Approximation and Projection (UMAP) embeddings using the Jaccard index as the distance metric. All latent states are colored by its optimal transport distance (d_{OT}) to the initial state.

rescaling to circumvent the non-ergodicity that stalls the naïve self-repelling walks and distances that have not yet varied are automatically penalized, ensuring that the sampler schedule allocates computational efforts to under-explored regions rather than oscillating within a local basin. Third, the interoperability of AFLF is substantiated by its consistent performance across spatial resolutions: residue-level inter-centroid distances recover crystallographic B-factors, whereas segment-level distances capture functional state transitions and expose cryptic binding sites. Moreover, the handful of transient cryptic cavities documented to date represent only a minute fraction of the proteome and remain numerically overwhelmed by the orthosteric sites that dominate contemporary structure-based ligand design. Here, AFLF produces conformational ensembles without recourse to seeded ligand poses, mutational profiles, or coevolutionary restraints; the protocol recovers the structural ensemble solely by proposing potential configurations that are intrinsically sustainable by the apo architecture. As the AFLF inference is decoupled from retrospective knowledge, every cavity it reports is a topological attribute distilled from the structural fluctuations embedded in the AF model rather than the artifact of inductive bias.

In summary, we present AFLF to bridge the gap between the single state predictions of AF and the conformational dynamics that underlies protein function. By grafting a lightweight, self-repelling, self-adaptive importance sampler onto the LoRA perturbations to the Evoformer latent representation, we transform the otherwise opaque AF latent embeddings into a navigable manifold and generate topologically accessible ensembles at a moderate computational cost. The AFLF protocol offers a fully modular in-place augmentation to AF, thereby delivering an immediately deployable enhancement to the foundational model without *post hoc* fine-tuning. Of course, combining physics-based modeling with AFLF may still confer measurable advantages in selected scenarios, such as determining the optimal transition paths.³⁵ Finally, beyond its immediate practical utility, the AFLF protocol also furnishes a conceptual template for transforming a discriminative foundation model into a generative engine that continuously produces diverse conformational hypotheses. The same “latent flooding” philosophy can, in principle, be transposed to AF-Multimer to sample protein-protein complexes or to co-folding frameworks for exploring biomolecular interactions.

4 Methods

4.1 AFLF Checkpoint Modeling

To retrieve the initial latent states from AF, we execute a lightly modified AF inference pipeline and collect the latent input representations to the Evoformer, $\mathbf{h}_{\text{msa}} \in \mathbb{R}^{N_{\text{seq}} \times N_{\text{res}} \times 256}$ and $\mathbf{h}_{\text{pair}} \in \mathbb{R}^{N_{\text{res}} \times N_{\text{res}} \times 128}$, as well as to the StructMod, $\mathbf{g}_{\text{single}} \in \mathbb{R}^{N_{\text{res}} \times 384}$ and $\mathbf{g}_{\text{pair}} \in \mathbb{R}^{N_{\text{res}} \times N_{\text{res}} \times 128}$. For each of the representations, we instantiate the LoRA perturbative tensors (rank = 8) from the uniform distribution $\mathcal{U}(-0.1, 0.1)$. The

initially adapted latents are,

$$\mathbf{h}_{\text{msa}}^{(0)} = \mathbf{h}_{\text{msa}} - (\mathbf{C}_{\text{h,msa}}^{(0)} \mathbf{B}_{\text{h,msa}}^{(0)} \mathbf{A}_{\text{h,msa}}^{(0)}), \quad (1)$$

$$\mathbf{h}_{\text{pair}}^{(0)} = \mathbf{h}_{\text{pair}} - (\mathbf{C}_{\text{h,pair}}^{(0)} \mathbf{B}_{\text{h,pair}}^{(0)} \mathbf{A}_{\text{h,pair}}^{(0)}), \quad (2)$$

$$\mathbf{g}_{\text{single}}^{(0)} = \mathbf{g}_{\text{single}} - (\mathbf{B}_{\text{g,single}}^{(0)} \mathbf{A}_{\text{g,single}}^{(0)}), \quad (3)$$

$$\mathbf{g}_{\text{pair}}^{(0)} = \mathbf{g}_{\text{pair}} - (\mathbf{C}_{\text{g,pair}}^{(0)} \mathbf{B}_{\text{g,pair}}^{(0)} \mathbf{A}_{\text{g,pair}}^{(0)}). \quad (4)$$

Using the LoRA tensors as the input, the AFLF inference pipeline can use the input checkpoint to the Evoformer or the StructMod as the forward entry. To infer from the Evoformer checkpoint, we use,

$$\mathcal{F}_{\text{evoformer}}(\mathbf{C}_{\text{h,msa}}, \mathbf{B}_{\text{h,msa}}, \mathbf{A}_{\text{h,msa}}, \mathbf{C}_{\text{h,pair}}, \mathbf{B}_{\text{h,pair}}, \mathbf{A}_{\text{h,pair}}) \quad (5)$$

$$= \mathcal{F}_{\text{evoformer}}^{\text{AF}} \left(\mathbf{h}_{\text{msa}}^{(0)} + (\mathbf{C}_{\text{h,msa}} \mathbf{B}_{\text{h,msa}} \mathbf{A}_{\text{h,msa}}), \mathbf{h}_{\text{pair}}^{(0)} + (\mathbf{C}_{\text{h,pair}} \mathbf{B}_{\text{h,pair}} \mathbf{A}_{\text{h,pair}}) \right), \quad (6)$$

where $\mathcal{F}_{\text{evoformer}}^{\text{AF}}$ is the AF native inference function from the Evoformer inputs. To infer from the StructMod checkpoint, we use,

$$\mathcal{F}_{\text{structmod}}(\mathbf{B}_{\text{g,single}}, \mathbf{A}_{\text{g,single}}, \mathbf{C}_{\text{g,pair}}, \mathbf{B}_{\text{g,pair}}, \mathbf{A}_{\text{g,pair}}) \quad (7)$$

$$= \mathcal{F}_{\text{structmod}}^{\text{AF}} \left(\mathbf{g}_{\text{single}}^{(0)} + (\mathbf{B}_{\text{g,single}} \mathbf{A}_{\text{g,single}}), \mathbf{g}_{\text{pair}}^{(0)} + (\mathbf{C}_{\text{g,pair}} \mathbf{B}_{\text{g,pair}} \mathbf{A}_{\text{g,pair}}) \right), \quad (8)$$

where $\mathcal{F}_{\text{structmod}}^{\text{AF}}$ is the corresponding AF native inference function from the StructMod inputs.

4.2 AFLF Inference Modeling

In this study, without loss of generality, the AFLF model applies LoRA selectively to the Evoformer MSA representations,

$$\mathcal{F}^{\text{AFLF}}(\mathbf{C}_{\text{h,msa}}, \mathbf{B}_{\text{h,msa}}, \mathbf{A}_{\text{h,msa}}) = \mathcal{F}_{\text{evoformer}}^{\text{AF}} \left(\mathbf{h}_{\text{msa}}^{(0)} + (\mathbf{C}_{\text{h,msa}} \mathbf{B}_{\text{h,msa}} \mathbf{A}_{\text{h,msa}}), \mathbf{h}_{\text{pair}} \right), \quad (9)$$

and outputs the coordinates $\mathbf{r} \in \mathbb{R}^{N \times 3}$ as well as the confidence scores $\mathbf{y}^{\text{PLDDT}} \in [0, 1]^{N_{\text{res}}}$. The output heavy atom coordinates from the main AF trunk are end-to-end differentiable with regard to all inputs, including the LoRA tensors introduced by AFLF. Accordingly, we operate most of the AFLF objectives on the pairwise distances between the centroids of the M atom groups formed by partitioning the total number of N predicted atomic coordinates. Practically, the user-defined lookup table $\mathcal{G} := \{G_i \mid i \in \{0, \dots, M-1\}\}$ maps the heavy atom $n \in G_i$ with coordinates $\mathbf{r}_n$ to the i -th atom group, for which the centroid coordinates $\mathbf{c} \in \mathbb{R}^{M \times 3}$ are computed,

$$\mathbf{c} = \left\{ \mathbf{c}_i \mid \mathbf{c}_i = \frac{1}{\#G_i} \sum_{n \in G_i} \mathbf{r}_n \text{ and } i \in \{0, \dots, M-1\} \right\}. \quad (10)$$

The AFLF structural state can then be characterized by the ordered set of inter-centroid distances,

$$\mathcal{D}(\mathbf{c}) := \{d_{ij} \mid d_{ij} = \|\mathbf{c}_i - \mathbf{c}_j\|_F \text{ and } i, j \in \{0, \dots, M-1\}\}. \quad (11)$$

The notation is also used explicitly on the predicted atomic coordinates, $\mathcal{D}(\mathbf{c}) = \mathcal{D}(\mathbf{r}; \mathcal{G}) =: \mathcal{D}_{\mathcal{G}}(\mathbf{r})$.

To facilitate the AFLF structural exploration, we propose a self-repelling, self-adaptive walk guided by the Gaussian flooding potentials in the continuum of the AFLF inferred conformations. Let τ be the instant (artificial) time and Δt the interval of taking repelling states, we keep a running stack of K memorized repellents $\{\mathcal{D}_{\mathcal{G}}(\mathbf{r}(t_k)) \mid k \in \{0, \dots, K-1\}\}$ taken at time stamps $t_k = \lfloor \max(0, \tau/\Delta t - k) \rfloor \Delta t$, where $\lfloor \cdot \rfloor$ denotes the floor rounding function. The flooding potential contributed by each of the repellent distances $d_{ij}(t_k) \in \mathcal{D}_{\mathcal{G}}(\mathbf{r}(t_k))$ on the instantaneous AFLF prediction $d_{ij}(\tau) \in \mathcal{D}_{\mathcal{G}}(\mathbf{r}(\tau))$ reads,

$$\mathcal{L}_{\text{repelling}, ij}(d_{ij}(\tau)) := \sum_{k=0}^{K-1} w_k \exp \left(-\frac{1}{2\varsigma^2} (d_{ij}(\tau) - d_{ij}(t_k) - \varepsilon_k)^2 \right), \quad (12)$$

where w_k denotes the weight per repellent, ς the Gaussian width, and $\varepsilon_k \sim \mathcal{U}(-\epsilon, \epsilon)$ the perturbative noise capped at ϵ to prevent instantaneous ($\tau = t_k$) gradient vanishing. The reduced repelling loss is developed,

$$\mathcal{L}_{\text{repelling}}(\mathbf{r}(\tau)) := \frac{1}{M} \sum_{i=0}^{M-1} \sum_{j=0}^{M-1} \psi(\alpha_{ij} \beta_{ij}(\tau) \mathcal{L}_{\text{repelling},ij}), \quad (13)$$

where $\alpha_{ij} \in \{0, 1\}$ is the predefined interaction mask, $\beta_{ij}(\tau) \in \mathbb{R}_{>0}$ the self-adaptive coefficient, and ψ the random element-wise dropout function.

The self-adaptive coefficient $\beta_{ij}(\tau)$ is of particular interest and deserves further discussion. Prior works have postulated that AF models could distill the core biophysical dynamics encrypted in the protein structural data on which they were trained.^{36,37} However, the structural ensemble generated by mutating AF latent states captures neither realistic nor artificial times, reflecting only the static time-independent thermodynamics of proteins. Consequently, the dynamics on the latent landscape of the AF foundational model lacks the ergodic freedom needed for basin escaping. Prolonged naïve self-repelling samplers therefore in general fail to reproducibly explore diverse protein conformations. While the AF models *per se* are not tied to physical temporal scales, we propose that the structural ensembles learned by AF can nevertheless be decoded by enhanced sampling across multiple spatial scales. The AFLF sampling is enhanced by a self-adaptive scheme facilitated through online importance rescaling, as detailed below. Taking the arithmetic average and the variance of each d_{ij} along the AFLF trajectory, $\mu_{ij}(\tau) = (1/\tau) \int_0^\tau dt d_{ij}(t)$ and $\sigma_{ij}^2(\tau) = (1/\tau) \int_0^\tau dt (d_{ij}(t) - \mu_{ij}(\tau))^2$, we estimate the instantaneous sampling progress using the α_{ij} -masked coefficient of variation (CV),

$$\text{CV}_{ij}(\tau) = \begin{cases} \sigma_{ij}(\tau)/\mu_{ij}(\tau) & \text{if } \alpha_{ij} = 1, \\ +\infty & \text{if } \alpha_{ij} = 0, \end{cases} \quad (14)$$

and define,

$$\beta_{ij}(\tau) := \frac{\exp(-b \text{CV}_{ij}(\tau))}{\sum_{i',j'} \exp(-b \text{CV}_{i'j'}(\tau))} \sum_{i'',j''} \alpha_{i''j''}, \quad (15)$$

where $b \in \mathbb{R}_{\geq 0}$ is the reciprocal softmax tempering parameter and the constant scaling $\sum_{i'',j''} \alpha_{i''j''}$ equates to 1 the average of all $\beta_{ij}(\tau)$ that are unmasked ($\alpha_{ij} = 1$). Briefly, distances with low $\text{CV}_{ij}(\tau)$ are deemed under-sampled and are driven to explore further by a stronger Gaussian penalty with large $\beta_{ij}(\tau)$, and *vice versa*. The tempering parameter b controls how peaked the softmax distribution is. For example, taking $b \rightarrow +\infty$ would focus the sampling only on the lowest fluctuating d_{ij} whereas $b = 0$ leads to uniformly weighted sampling on all d_{ij} . It should be noted that the $\beta_{ij}(\tau)$ -modulated importance sampling implicitly assumes homogeneous thermal distribution among all d_{ij} , ignoring that, for example, the distance between the centroids of two α -helices is intrinsically less variable than one spanning two disordered loops. In cases where the relevant $\{d_{ij}\}$ exhibit nontrivial heterogeneous scales in terms of their physical distribution, in addition to calibrating $\beta_{ij}(\tau)$ for selected entries *a priori*, one could instead apply the centroid-wise rescaling,

$$\beta_i(\tau) := \frac{\exp(-b \text{CV}_i(\tau))}{\sum_{i'} \exp(-b \text{CV}_{i'}(\tau))} \frac{\sum_{i'',j''} \alpha_{i''j''}}{\sum_{j''} \alpha_{ij''}}, \quad (16)$$

where $\text{CV}_i(\tau) = \sum_{\{j|\alpha_{ij}=1\}} \text{CV}_{ij}(\tau)$ and $\forall j : \beta_{ij}(\tau) = \beta_i(\tau)$. We note that adopting $\beta_i(\tau)$ entails that the centroids concerning fewer distances ($\sum_{j''} \alpha_{ij''}$ small) are permanently assigned proportionally larger rescaling weights. When $\{\alpha_{ij}\}$ is derived from local interaction networks such as contact maps, the $\beta_i(\tau)$ rescaling mirrors the empirical knowledge that higher contact density stabilizes local conformational dynamics. Of course, an unbiased measure of “sampling adequacy” — were it available without further unjustified statistical or thermodynamical hypothesis — would be the ideal metric for computing $\beta_{ij}(\tau)$ online. In practice, the CV-derived estimators add negligible $\mathcal{O}(1)$ computational overhead and are likely sufficient for most applications.

To contain the AFLF LoRA-ed latent states on the manifold of native-like folds, we introduce a multiscale strategy for regulating the structural outputs. We note that all of the following loss terms targeting distinct spatial scales can be selectively disabled according to the intended scope of application.

On the spatial scale of the atoms, the local structural features should be preserved by rigid body restraints for all AFLF generated conformations. Denote $\mathbf{r}_p(\tau)$ as the subset of atomic coordinates that sufficiently characterizes the p -th local spatial feature and $\hat{\mathbf{r}}_p$ the corresponding reference from, e.g., the initial prediction $\mathbf{r}(0)$. The local topological loss is the optimally aligned deviation of the P subsets from the reference,

$$\mathcal{L}_{\text{local}}(\mathbf{r}(\tau)) := \frac{1}{P} \sum_{p=0}^{P-1} \|\psi_{\text{atom}}(\mathbf{r}_p(\tau) - \hat{\mathbf{r}}_p \mathbf{R}_p - \mathbf{T}_p)\|_F^2, \quad (17)$$

where ψ_{atom} denotes the atom-wise dropout and $(\mathbf{R}_p, \mathbf{T}_p) = \text{argmin}_{\mathbf{R}_p \in \text{SO}(3), \mathbf{T}_p \in \mathbb{R}^3} \|\mathbf{r}_p(\tau) - \hat{\mathbf{r}}_p \mathbf{R}_p - \mathbf{T}_p\|_F$. We compute each $(\mathbf{R}_p, \mathbf{T}_p)$ pair with the closed-form Kabsch solver,³⁸ which does not need to be differentiable since the optimal alignment acts only on the constant reference. Typically, we apply $\mathcal{L}_{\text{local}}$ to maintain the structural constraints of molecular topologies that are not explicitly encoded by the internal topology used in AF inference, such as the disulfide bonds between cysteines, the pyrrolidine ring in prolines, and essential hydrogen bonding interactions.

On the spatial scale of the residues, the global structural plasticity could be controlled according to the AF predicted native fold. We represent each residue by the coordinates of its centroid and collect a total number of Q inter-centroid distances from the instantaneous prediction $d_q(\tau) \in \mathcal{D}_{\mathcal{G}_{\text{res}}}(\mathbf{r}(\tau))$ and the reference structure $\hat{d}_q = d_q(0) \in \mathcal{D}_{\mathcal{G}_{\text{res}}}(\mathbf{r}(0))$, where $\mathcal{G}_{\text{res}}$ is the residue-wise lookup table. The global plasticity loss is the cross-entropy between the density estimate on the reference and the prediction,

$$\mathcal{L}_{\text{global}}(\mathbf{r}(\tau)) := \frac{1}{\sum_{q=0}^{Q-1} \alpha_q} \sum_{q=0}^{Q-1} \sum_{l=0}^L \psi \left(-\alpha_q \rho_l(\hat{d}_q) \log(\rho_l(d_q(\tau))) \right), \quad (18)$$

where $\alpha_q \in \{0, 1\}$ is the predefined binary mask that discards selected entries. Here, ρ_l is the kernel density estimator obtained from the histogram of $L + 1$ bins. By partitioning the support of the distance ensemble into bins of uniform width $\Delta d \in \mathbb{R}_{>0}$, we denote the l -th bin as $[l\Delta d, (l+1)\Delta d)$, except for the last bin $[L\Delta d, +\infty)$. The distance density is estimated through the Gaussian kernel,

$$\rho_l(d_q(\tau)) = \Phi \left(\frac{(l+1)\Delta d - d_q(\tau)}{\varsigma'} \right) - \Phi \left(\frac{l\Delta d - d_q(\tau)}{\varsigma'} \right), \quad (19)$$

where $\Phi(x) = \frac{1}{2} (1 + \text{erf}(x/\sqrt{2}))$ is the differentiable Gaussian cumulative density function and ς' the smearing width. Typically, given arbitrary upper bound Δd , we adapt ς' to control the tolerated residue-wise fluctuations, and thus effectively the global structural plasticity. We note that $\mathcal{L}_{\text{global}}$ is computed from kernel density estimations on calculated geometrical distances, as opposed to the softmax normalized latent representation in the ‘‘distogram’’ learning of AF.

On the spatial scale where the repelling potential is applied, it is possible that some of the repelling distances are not controlled by $\mathcal{L}_{\text{local}}$ and $\mathcal{L}_{\text{global}}$. We thus introduce the boundary penalty for each repelling distance to prevent large deviations $\Delta d_{ij}(\tau) = d_{ij}(\tau) - \hat{d}_{ij}$ from the reference $\hat{d}_{ij} = d_{ij}(0) \in \mathcal{D}_{\mathcal{G}}(\mathbf{r}(0))$,

$$\mathcal{L}_{\text{anchoring}} := \frac{1}{M} \sum_{i=0}^{M-1} \sum_{j=0}^{M-1} \psi \left(\alpha_{ij} \phi \left(\max(0, \Delta d_{\text{thresh}}^{\text{lower}} - \Delta d_{ij}(\tau)) + \max(0, \Delta d_{ij}(\tau) - \Delta d_{\text{thresh}}^{\text{upper}}) \right) \right), \quad (20)$$

where ϕ denotes the Huber loss function³⁹ and $\Delta d_{\text{thresh}}^{\text{lower}} < \Delta d_{\text{thresh}}^{\text{upper}}$ the lower and the upper bounds between which the boundary loss is inactive. Typically, $\mathcal{L}_{\text{anchoring}}$ sets the permitted space of conformations for AFLF to explore.

Finally, the AF model confidence loss $\mathcal{L}_{\text{pLDDT}} := (1/N_{\text{res}}) \sum \mathbf{1} - \mathbf{y}_{\text{pLDDT}}$ can be combined with the geometry losses into the time-dependent AFLF loss,

$$\mathcal{L}_{\text{AFLF}} := \mathcal{L}_{\text{repelling}} + w_{\text{local}} \mathcal{L}_{\text{local}} + w_{\text{global}} \mathcal{L}_{\text{global}} + w_{\text{anchoring}} \mathcal{L}_{\text{anchoring}} + w_{\text{pLDDT}} \mathcal{L}_{\text{pLDDT}}, \quad (21)$$

where w_{local} , w_{global} , $w_{\text{anchoring}}$, and w_{pLDDT} are the corresponding scalar weights. Dynamical propagation on $\mathcal{L}_{\text{AFLF}}$ can then be most conveniently realized with momentum-accelerated, gradient-following samplers such as Adam.⁴⁰ The detailed AFLF configuration is listed in **Table S3**. The AFLF dynamics are

primarily driven by the repelling loss term, as shown in the loss profiles along the trajectory for each system (Figures S10, S11, S12, S13, S14, S15, S16).

4.3 Computational Details

The AFLF program was implemented as an extension to the AF 2.3.2 source (<https://github.com/google-deepmind/alphafold>) based on the JAX 0.3.25 ecosystem (<http://github.com/google-deepmind/jax>) and the Optax 0.1.7 library (<https://github.com/google-deepmind/optax>).

The AFLF checkpoints were retrieved from the default AF inference pipeline with 3 recycling iterations that embeds both the features and the positions from the previous prediction. The MSAs were resampled and clustered into 508 clusters and the structural templates were disabled. The AFLF inference modeling uses only the Evoformer and the StructMod blocks, and the original AF test-time configuration is thus obsolete. As the reverse-mode gradient computation in the Evoformer stack is memory-intensive, we executed the AFLF inference in the “jax.numpy.bfloat16” precision and rematerialized intermediate activations on demand during the backward sweep, such that the entire workload fit on a single 40 GB NVIDIA A40 GPU. The “model_1_ptm” set of model parameters was adopted for all systems. The computational wall time of our AFLF implementation scales almost linearly with system size (Table S4).

The CHARMM⁴¹ and the OpenMM⁴² programs were used for the reference MD simulations of ubiquitin using the CHARMM36m⁴³ classical potential. The MDAnalysis⁴⁴, the MDpocket⁴⁵, and the umap-learn⁴⁶ libraries were applied for the structural analysis. All molecular graphics were prepared with ChimeraX.⁴⁷

Acknowledgement

The authors thank Profs. Yiqin Gao, Wei Yang, and Yaoqi Zhou for insightful discussions. This work is supported by the Shenzhen Medical Research Fund (B2404003), the National Natural Science Foundation of China (T2596084, 22503068), the Hangzhou Joint Fund of Zhejiang Provincial Natural Science Foundation of China (LHZQN26B030001), the Pioneer and Leading Goose R&D Program of Zhejiang (2023C03109), and the State Key Laboratory of Gene Expression. We thank the Westlake University Supercomputer Center for computational resources and related assistance.

Author Contributions

Runtong Qian: Software, Validation, Formal analysis, Investigation, Data curation, Writing – review & editing, Visualization; **Rui Zhan:** Validation, Formal analysis, Investigation, Data curation, Writing – review & editing, Visualization; **Zilin Song:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition; **Jing Huang:** Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Data Availability

All data supporting the findings of this study will be deposited in a public repository upon completion of the peer-review process and incorporation of required revisions.

Code Availability

The computer code developed in the current study is publicly accessible at the following repository: <https://github.com/JingHuangLab/AFLF.git>.

References

- 1 Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).

- 2 Evans, R. et al. Protein complex prediction with AlphaFold-Multimer. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.10.04.463034> (2021).
- 3 Abramson, J. et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).
- 4 Baek, M. et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science* **373**, 871–876 (2021).
- 5 Mirdita, M. et al. ColabFold: Making protein folding accessible to all. *Nat. Methods* **19**, 679–682 (2022).
- 6 Lin, Z. et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).
- 7 Lane, T. J. Protein structure prediction has reached the single-structure frontier. *Nat. Methods* **20**, 170–173 (2023).
- 8 Wei, G., Xi, W., Nussinov, R. & Ma, B. Protein ensembles: How does nature harness thermodynamic fluctuations for life? The diverse functional roles of conformational ensembles in the cell. *Chem. Rev.* **116**, 6516–6551 (2016).
- 9 Xie, T. & Huang, J. Can protein structure prediction methods capture alternative conformations of membrane transporters? *J. Chem. Inf. Model.* **64**, 3524–3536 (2024).
- 10 Nussinov, R., Regev, C. & Jang, H. Leveraging conformational ensembles in allosteric drug discovery. *Trends Pharmacol. Sci.* **47**, 276–289 (2026).
- 11 Lu, J., Zhong, B., Zhang, Z. & Tang, J. Str2str: A score-based framework for zero-shot protein conformation sampling. In *Proceedings of the 12th International Conference on Learning Representations* (OpenReview.net, 2024).
- 12 Jing, B., Berger, B. & Jaakkola, T. AlphaFold meets flow matching for generating protein ensembles. In *Proceedings of the 41st International Conference on Machine Learning* (PMLR, 2024).
- 13 Lewis, S. et al. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science* **389**, eadv9817 (2025).
- 14 Sala, D., Engelberger, F., Mchaourab, H. S. & Meiler, J. Modeling conformational states of proteins with AlphaFold. *Curr. Opin. Struct. Biol.* **81**, 102645 (2023).
- 15 Heo, L. & Feig, M. Multi-state modeling of G-protein coupled receptors at experimental accuracy. *Proteins*. **90**, 1873–1885 (2022).
- 16 Del Alamo, D., Sala, D., Mchaourab, H. S. & Meiler, J. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *eLife* **11**, e75751 (2022).
- 17 Wallner, B. AFsample: Improving multimer prediction with AlphaFold using massive sampling. *Bioinformatics* **39**, btad573 (2023).
- 18 Wayment-Steele, H. K. et al. Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature* **625**, 832–839 (2024).
- 19 Xie, T., Song, Z. & Huang, J. Conditioned protein structure prediction. *PRX Life* **2**, 043001 (2024).
- 20 Lee, M. et al. Large-scale predictions of alternative protein conformations by AlphaFold2-based sequence association. *Nat. Commun.* **16**, 5622 (2025).
- 21 Sun, M., Chen, X., Kolter, J. Z. & Liu, Z. Massive activations in large language models. In *Proceedings of the 1st Conference on Language Modeling* (COLM Organization, 2024).
- 22 Valeriani, L. et al. The geometry of hidden representations of large transformer models. In *Proceedings of the 37th Conference on Neural Information Processing Systems* 51234–51252 (Curran Associates Inc., 2023).
- 23 Lu, A. X. et al. Tokenized and continuous embedding compressions of protein sequence and structure. *Patterns* **6**, 101289 (2025).
- 24 Ahdritz, G. et al. OpenFold: Retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization. *Nat. Methods* **21**, 1514–1524 (2024).
- 25 Hayes, T. L. & Krishnan, G. P. Quantifying the role of OpenFold components in protein structure prediction. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2511.14781> (2025).
- 26 Hu, E. J. et al. LoRA: Low-rank adaptation of large language models. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2106.09685> (2021).

- 27 Grubmüller, H. Predicting slow structural transitions in macromolecular systems: Conformational flooding. *Phys. Rev. E* **52**, 2893–2906 (1995).
- 28 Wlodek, S. T., Clark, T. W., Scott, L. R. & McCammon, J. A. Molecular dynamics of acetylcholinesterase dimer complexed with tacrine. *J. Am. Chem. Soc.* **119**, 9513–9522 (1997).
- 29 Beckstein, O., Denning, E. J., Perilla, J. R. & Woolf, T. B. Zipping and unzipping of adenylate kinase: Atomistic insights into the ensemble of open $\leftrightarrow$ closed transitions. *J. Mol. Biol.* **394**, 160–176 (2009).
- 30 Beglov, D. et al. Exploring the structural origins of cryptic sites on proteins. *Proc. Natl. Acad. Sci. U. S. A* **115**, E3416–E3425 (2018).
- 31 Lazou, M. et al. Predicting multiple conformations of ligand binding sites in proteins suggests that AlphaFold2 may remember too much. *Proc. Natl. Acad. Sci. U. S. A* **121**, e2412719121 (2024).
- 32 Song, Z., Zhou, H., Tian, H., Wang, X. & Tao, P. Unraveling the energetic significance of chemical events in enzyme catalysis via machine-learning based regression approach. *Commun. Chem.* **3**, 134 (2020).
- 33 Meller, A., Bhakat, S., Solieva, S. & Bowman, G. R. Accelerating cryptic pocket discovery using AlphaFold. *J. Chem. Theory Comput.* **19**, 4355–4363 (2023).
- 34 Qu, T., Austin, S. L., Zheng, L., Zhang, J. & Yang, W. Identification of protein cryptic sites via conformational dynamics capturing and water-based pocket characterization in molecular dynamics simulations. *J. Chem. Theory Comput.* **21**, 8571–8582 (2025).
- 35 Song, Z., Xu, Y., Zhang, H., Ding, Y. & Huang, J. Identifying the most probable transition path with constant advance replicas. *J. Chem. Theory Comput.* **20**, 9857–9870 (2024).
- 36 Roney, J. P. & Ovchinnikov, S. State-of-the-art estimation of protein model accuracy using AlphaFold. *Phys. Rev. Lett.* **129**, 238101 (2022).
- 37 Zhang, Z. et al. Unifying constraints linking protein folding and native dynamics decoded from AlphaFold. *Phys. Rev. Lett.* **136**, 068401 (2026).
- 38 Kabsch, W. A solution for the best rotation to relate two sets of vectors. *Acta Cryst. A* **32**, 922–923 (1976).
- 39 Huber, P. J. Robust estimation of a location parameter. *Ann. Math. Stat.* **35**, 73–101 (1964).
- 40 Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations* (OpenReview.net, 2015).
- 41 Hwang, W. et al. CHARMM at 45: Enhancements in accessibility, functionality, and speed. *J. Phys. Chem. B* **128**, 9976–10042 (2024).
- 42 Eastman, P. et al. OpenMM 8: Molecular dynamics simulation with machine learning potentials. *J. Phys. Chem. B* **128**, 109–116 (2023).
- 43 Huang, J. et al. CHARMM36m: An improved force field for folded and intrinsically disordered proteins. *Nature methods* **14**, 71–73 (2017).
- 44 Michaud-Agrawal, N., Denning, E. J., Woolf, T. B. & Beckstein, O. MDAnalysis: A toolkit for the analysis of molecular dynamics simulations. *J. Comput. Chem.* **32**, 2319–2327 (2011).
- 45 Schmidtke, P., Bidon-Chanal, A., Luque, F. J. & Barril, X. MDpocket: open-source cavity detection and characterization on molecular dynamics trajectories. *Bioinformatics* **27**, 3276–3285 (2011).
- 46 McInnes, L., Healy, J. & Melville, J. UMAP: Uniform manifold approximation and projection for dimension reduction. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.1802.03426> (2020).
- 47 Meng, E. C. et al. UCSF ChimeraX: Tools for structure building and analysis. *Protein Sci.* **32**, e4792 (2023).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [MovieS1.mp4](#)
- [si.pdf](#)
- [MovieS2.mp4](#)
- [rs.pdf](#)