

IMPROVING PROTEIN SEQUENCE DESIGN THROUGH DESIGNABILITY PREFERENCE OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Protein sequence design methods have demonstrated strong performance in sequence generation for *de novo* protein design. However, as the training objective was sequence recovery, it does not guarantee designability—the likelihood that a designed sequence folds into the desired structure. To bridge this gap, we redefine the training objective by steering sequence generation toward high designability. To do this, we integrate Direct Preference Optimization (DPO), using AlphaFold pLDDT scores as the preference signal, which significantly improves the *in silico* design success rate. To further refine sequence generation at a finer, residue-level granularity, we introduce Residue-level Designability Preference Optimization (ResiDPO), which applies residue-level structural rewards and decouples optimization across residues. This enables direct improvement in designability while preserving regions that already perform well. Using a curated dataset with residue-level annotations, we fine-tune LigandMPNN with ResiDPO to obtain EnhancedMPNN, which achieves a nearly 3-fold increase in *in silico* design success rate (from 6.56% to 17.57%) on a challenging enzyme design benchmark.

1 INTRODUCTION

The computational design of proteins with specific functions has considerable potential for solving outstanding challenges in medicine and biotechnology Huang et al. (2016); Watson et al. (2023); Vázquez Torres et al. (2025); Kim et al. (2024). Current computational protein design pipelines often decouple the problem into two stages—first generating a backbone structure, then “inverse folding” to assign a sequence predicted to adopt that backbone Watson et al. (2023); Dauparas et al. (2022). While significant progress has been made in protein sequence design (PSD) using deep learning Ingraham et al. (2019); Jing et al. (2020); Liu et al. (2022); Wang (2022); Gao et al. (2022); Anand et al. (2022); Jain et al. (2022); Zhou et al. (2023); Zheng et al. (2023); Gao et al. (2023; 2024), surpassing traditional physics-based methods like Rosetta Zaidman et al. (2020), there are remaining challenges. Existing methods primarily optimize for sequence recovery—the ability to reproduce native sequences given their backbones. However, for protein structure and function design success, what is most important is how closely the designed sequence folds to the designed target structure Ruffolo et al. (2025); Zhu et al. (2025); Johansen et al. (2025); Hou et al. (2025). As expected, there is a strong correlation between designability and experimental success for both protein binder design Bennett et al. (2023) and enzyme design Kim et al. (2024). A key challenge is that state-of-the-art PSD models, while optimized for sequence recovery, often exhibit poor designability Ye et al. (2024). These models can introduce subtle yet critical errors, such as placing structure-breaking residues like Proline and Glycine within a desired β -sheet, causing it to misfold into a flexible loop (Fig. 1). This low reliability necessitates generating thousands of candidates to find a few viable ones, leading to success rates as low as 3% for enzyme design with standard RFDiffusion and ProteinMPNN pipelines Watson et al. (2023). This brute-force screening is computationally inefficient, slows experimental validation, and highlights an urgent need for methods that directly optimize for designability.

We set out to directly optimize a ProteinMPNN-like model inspired by the success of alignment techniques like Reinforcement Learning from Human Feedback (RLHF) Christiano et al. (2017) used to bridge objective gaps in large language models (LLMs) Lla; DeepSeek-AI (2025). We aimed to explicitly align PSD models towards generating sequences with high designability. Proteins offer advantages over natural language for this adaptation. First, instead of relying on subjective human preferences, we can leverage physics-based simulation Wang et al. (2024) or high-accuracy structure

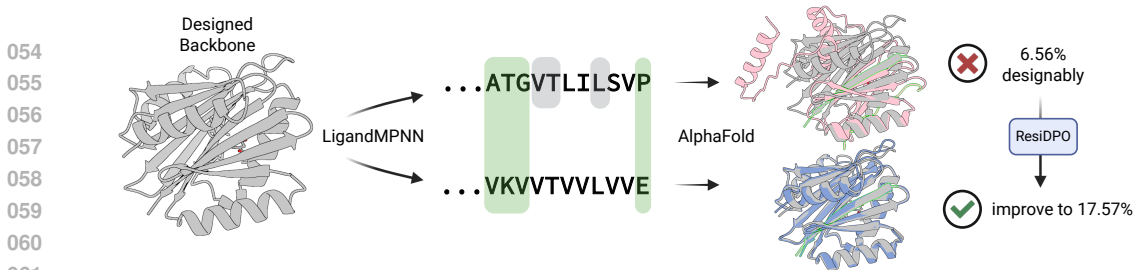


Figure 1: ResiDPO directly optimizes protein designability by decoupling optimization at the residue level. While standard models often generate sequences with local structural flaws (low pLDDT), ResiDPO selectively applies a preference reward to improve these problematic regions (such as Proline and Glycine in the figure which break the helix). To avoid catastrophic forgetting, it simultaneously preserves high-confidence regions using KL regularization. This decoupled strategy boosts the *in silico* design success rate from 6.56% to 17.57%.

predictors Senior et al. (2020); Baek et al. (2021). We chose the predicted Local Distance Difference Test (pLDDT) score from AlphaFold2 Jumper et al. (2021), which correlates well with structural accuracy (Fig. 6) and provides a quantitative, objective reward signal for designability. Second, the fixed length of sequences designed for a specific backbone allows for fine-grained, residue-level reward assignment, unlike the sequence-level rewards typically used in LLMs.

Building on these insights, here we introduce Residue-level Designability Preference Optimization (ResiDPO), an adaptation of the Direct Preference Optimization (DPO) Rafailov et al. (2023) framework tailored for protein design. Standard DPO, while effective for aligning language models using sequence-level preferences, faces challenges when applied naively to PSD. It optimizes a single loss function balancing preference learning against a KL-divergence term that regularizes towards the original model distribution. This can create conflicting gradients, especially when trying to significantly increase the probability of high-designability sequences. ResiDPO overcomes this by leveraging residue-level pLDDT scores as rewards. It *decouples* the DPO loss: for residues predicted to improve designability (e.g., low initial pLDDT), it prioritizes maximizing the preference reward signal; for residues already contributing positively to the structure (e.g., high pLDDT and high confidence from the base model), it prioritizes KL regularization to maintain learned structural features. This decoupling provides a clearer, more stable optimization target, directly enhancing designability without catastrophic forgetting (Fig. 1). We use ResiDPO to develop an improved model, EnhancedMPNN, which increases the design success rate nearly 3-fold on a range of challenging enzyme and binder design challenges.

Our contributions are summarized as follows:

- We frame the challenge of low designability in PSD as a critical objective misalignment problem, amenable to targeted alignment strategies inspired by reinforcement learning.
- We propose ResiDPO, a novel alignment algorithm that adapts and extends DPO for protein design by leveraging objective, residue-level structural feedback (pLDDT) to decouple the optimization objective, enabling more effective learning of designability.
- We curate and will release a large-scale dataset featuring residue-level pLDDT labels, providing a valuable resource for training and evaluating designability-focused PSD models.
- We demonstrate empirically that our ResiDPO-finetuned model, EnhancedMPNN, achieves substantial improvements in *in silico* design success rates (nearly 3x for enzymes, 2x for binders), significantly reducing computational cost and accelerating the design-build-test cycle for functional proteins.

2 RELATED WORK

Protein Sequence Design (Inverse Folding). Protein sequence design methods aim to generate amino acid sequences that fold into a desired 3D structure. These methods have seen significant progress, driven by advancements in graph neural networks and attention mechanisms. Early work like GraphTrans Ingraham et al. (2019) extended the Transformer architecture Vaswani et al. (2017)

to capture long-range interactions in protein sequences by modeling sparse relationships between residues distant in sequence but proximal in 3D space. Geometric Vector Perceptrons (GVPs) Jing et al. (2020) further enhanced it by incorporating geometric information, representing residues and their relationships as vectors and scalars. This approach was adopted by ESM-IF1 Hsu et al. (2022), which combined a GVP-based structure encoder with a GVP-Transformer to learn inter-residue interactions. GCA Tan et al. (2023) focused on incorporating global context. ProteinMPNN Dauparas et al. (2022) extends GraphTrans by introducing more geometry features and random decoding. PiFold Gao et al. (2022) introduced a novel residue featurization, a PiGNN module, and further improved efficiency by dispensing with autoregressive decoding. More recent efforts have explored knowledge integration (KW-Design Gao et al. (2023)) and adaptation of pre-trained language models (LM-Design Zheng et al. (2023), ZymCTRL Munsamy et al. (2024), CarbonDesign Ren et al. (2024)) for sequence design. These methods have collectively improved natural sequence recovery rates from approximately 40% to over 60%.

However, maximizing natural sequence recovery does not directly address the crucial aspect of designability—the ability of a designed sequence to fold into the target structure. As highlighted by Huang et al. (2016), the space of natural sequences represents only a small fraction of possible sequences, and multiple sequences can fold into the same structure. Therefore, ensuring high designability is paramount. Unlike existing methods that primarily focus on recovering natural sequences, this work explicitly aims to improve the designability of protein sequence design methods.

Many proteins, particularly enzymes, do not function in isolation. They often require cofactors, such as ligands and metals, for activity. LigandMPNN Dauparas et al. (2025) extends ProteinMPNN to incorporate ligand awareness. Given the challenges inherent in *de novo* enzyme design, we adopt LigandMPNN as our base model.

Preference learning. Self-supervised language models have demonstrated promising capability in completing zero-shot tasks. Before applying them to downstream tasks, aligning with human preferences has proven crucial for improving performance on downstream tasks. Early approaches employed reinforcement learning from human feedback (RLHF) Christiano et al. (2017) to fine-tune models based on relative judgments of response quality. However, RLHF can be complex and challenging to implement. Direct Preference Optimization (DPO) Rafailov et al. (2023) significantly simplified this process by eliminating the need for an explicit reward model and instead directly optimizing a pairwise logistic loss based on preferences. Subsequent work has extended DPO by removing the Bradley-Terry assumption (IPO Azar et al. (2024)), enabling listwise optimization (LiPO Liu et al. (2024)), removing the requirement for a reference model (SimPO Meng et al. (2024)), and so on. Recently, DPO has been applied to protein sequence design. Park et al. (2024) used DPO for peptide design, demonstrating improvements of 8% in structural similarity and 20% in sequence diversity. In this work, we focus on improving DPO to improve the designability of the designed sequences.

3 METHODS

This section details Residue-level Designability Preference Optimization (ResiDPO), a novel training framework specifically designed to align protein sequence design (PSD) models with the objective of *designability*. We first contextualize PSD within the preference optimization paradigm, adapting concepts from Direct Preference Optimization (DPO) used in large language models (LLMs). We then describe our strategy for generating preference data leveraging quantitative protein quality metrics. Finally, we introduce the ResiDPO objective function, which overcomes limitations of standard DPO by decoupling preference learning and model regularization at the residue level.

3.1 PROTEIN DESIGN AS PREFERENCE OPTIMIZATION

The goal of protein sequence design is to generate an amino acid sequence $y = (y_1, \dots, y_L)$ that folds into a target backbone structure x . While conventional PSD models are trained via sequence recovery (maximizing $P(y_{\text{native}}|x)$), this objective is often misaligned with the desired outcome: generating novel sequences y that fold precisely to the structure x (designability). Here, we frame PSD as a preference optimization problem, aiming to increase the probability of sequences with high designability.

Drawing inspiration from alignment techniques in LLMs Christiano et al. (2017), we first adapt the Direct Preference Optimization (DPO) framework Rafailov et al. (2023). DPO optimizes a policy model π_θ to better satisfy preferences compared to a reference model π_{ref} (typically the initial pre-trained model), using a dataset \mathcal{D} of preference pairs (y_w, y_l) for a given prompt x , where y_w is preferred over y_l . The DPO objective maximizes the likelihood of preferred responses while regularizing against large deviations from the reference model via a KL divergence penalty:

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

where σ is the sigmoid function, and β is a hyperparameter controlling the strength of the preference relative to the regularization. In our context, x is the target protein backbone structure, y_w and y_l are candidate amino acid sequences, π_{ref} is the pre-trained PSD model (e.g., LigandMPNN), and π_θ is the model being fine-tuned.

For a quantitative measure of designability, we utilize the predicted Local Distance Difference Test (pLDDT) score derived from AlphaFold2 (AF2) Jumper et al. (2021). The pLDDT score serves as a proxy for folding accuracy and stability, correlating well with experimental success and metrics like C α RMSD (Fig. 6). Thus, a sequence y_w with a higher pLDDT score is considered preferable to a sequence y_l with a lower score for the same target structure x .

3.2 PREFERENCE PAIR GENERATION

Effective DPO training relies on high-quality preference pairs (y_w, y_l) . While LLM alignment often requires a separate reward model or human labeling to score outputs, PSD benefits from objective quality metrics like pLDDT. To generate preference pairs for a given structure x , we first sample multiple candidate sequences using the reference model π_{ref} . We then predict the structure and pLDDT score for each sampled sequence using AF2. Based on pLDDT scores, we explored different strategies for selecting (y_w, y_l) pairs:

Rejection Sampling: This approach is analogous to methods used in the post-training alignment of LLMs Grattafiori et al. (2024); DeepSeek-AI (2025), where outputs are scored by a reward model. In our context, for a given structure x , we identify the sequence with the highest pLDDT score in the sampled batch as the preferred sequence y_w . Other sequences from the same batch, with lower pLDDT scores, are then randomly selected to serve as less preferred sequences y_l , forming multiple pairs (y_w, y_l) . We adopt this method as a baseline for comparison.

Application Sampling: Inspired by practical protein design workflows where sequences below a certain quality threshold are discarded Watson et al. (2023); Kim et al. (2024), this method selects y_w as a sequence with pLDDT ≥ 80 and y_l as a sequence with pLDDT ≥ 75 . However, the scarcity of high-pLDDT generations (pLDDT ≥ 80) for many structures using current PSD models severely limits the number of preference pairs obtainable with this approach (Fig. 7), reducing dataset diversity.

Relative Sampling: To generate more preference pairs and better utilize the available data, this method selects any pair of sequences (y_i, y_j) generated for the same structure x such that their pLDDT scores differ by more than a predefined threshold δ : $pLDDT(y_i) - pLDDT(y_j) > \delta$. The sequence with the higher pLDDT is designated y_w , and the one with the lower score is y_l .

Ablation studies indicated that Relative Sampling with a pLDDT difference threshold $\delta = 10$ yielded the best performance in downstream tasks. Therefore, we adopted this strategy for generating the preference dataset \mathcal{D} used in subsequent experiments.

3.3 RESIDUE-LEVEL DESIGNABILITY PREFERENCE OPTIMIZATION (RESIDPO)

The standard DPO objective (Eq. 1) optimizes a single loss, implicitly balancing preference maximization and regularization (KL divergence from π_{ref}). In PSD, optimizing the entire sequence probability to favor y_w can conflict with the KL constraint, particularly when substantial sequence changes are needed for improved designability. Modifying probabilities at certain residues to improve predicted local structure (higher pLDDT) inevitably increases the KL divergence, potentially hindering optimization or leading to instability.

Unlike variable-length text generation, protein design for a fixed backbone structure yields fixed-length sequences (L). This enables a residue-level analysis of designability via per-residue pLDDT scores, $pLDDT(y, i)$ for residue i in sequence y . ResiDPO leverages this granularity to decouple the optimization objectives. Instead of a single sequence-level objective, it applies targeted updates at the residue level, separating preference learning from constraint enforcement.

ResiDPO decomposes the loss into two components: Residue-level Preference Learning (RPL) and Residue-level Constraint Learning (RCL).

3.3.1 RESIDUE-LEVEL PREFERENCE LEARNING (RPL)

RPL focuses the preference optimization on specific residue positions where the preferred sequence y_w shows significant improvement in local predicted structure compared to the less preferred sequence y_l . The RPL loss is defined as:

$$\mathcal{L}_{\text{RPL}} = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \sum_{i \in \mathcal{I}} \left(\frac{\log \pi_{\theta}(y_w^i | x) - \log \pi_{\theta}(y_l^i | x)}{|\mathcal{I}|} \right) \right], \quad (2)$$

where \mathcal{I} is the set of residue indices i where the pLDDT difference exceeds a margin α :

$$\mathcal{I} = \{i \mid pLDDT(y_w, i) - pLDDT(y_l, i) > \alpha\}. \quad (3)$$

Here, y_k, i denotes the amino acid at position i of sequence y_k . This loss term encourages the model π_{θ} to increase the relative log-probability of preferred residues, specifically at positions deemed important for improving designability.

If no residues satisfy the condition ($\mathcal{I} = \phi$), indicating similar local pLDDT profiles or global differences, we revert to a sequence-level preference signal by setting $\mathcal{I} = \{1, \dots, L\}$ for that pair, effectively applying the standard DPO logic over the entire sequence.

3.3.2 RESIDUE-LEVEL CONSTRAINT LEARNING (RCL)

RCL aims to preserve the knowledge of the reference model π_{ref} at positions that are already well-structured and confidently predicted, preventing catastrophic forgetting and maintaining structural integrity. It applies a KL-divergence penalty at specific residue positions:

$$\mathcal{L}_{\text{RCL}} = \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\sum_{j \in \mathcal{J}} \frac{\pi_{ref}(y_w^j | x) \cdot \log \frac{\pi_{ref}(y_w^j | x)}{\pi_{\theta}(y_w^j | x)}}{|\mathcal{J}|} \right], \quad (4)$$

where

$$\mathcal{J} = \{j \mid pLDDT(y_w, j) > \beta \cap \pi_{ref}(y_w, j | x) > \gamma\} \quad (5)$$

is the set of residue positions j where the pLDDT is above a threshold β and the reference model’s probability exceeds a confidence threshold γ . The RCL loss encourages π_{θ} to stay close to π_{ref} at these “high-quality, high-confidence” positions \mathcal{J} .

3.3.3 COMBINED RESIDPO LOSS

The final ResiDPO loss function combines the residue-level preference learning and constraint learning terms:

$$\mathcal{L}_{\text{ResiDPO}} = \mathcal{L}_{\text{RPL}} + \lambda \mathcal{L}_{\text{RCL}}, \quad (6)$$

where λ is a hyperparameter balancing the strength of the preference signal against the constraint encouraging preservation of the reference model’s reliable predictions. This decoupled, residue-level optimization allows ResiDPO to more effectively navigate the complex trade-off between improving designability (via RPL) and maintaining model stability and learned structural features (via RCL), leading to improved performance in generating designable protein sequences.

4 EXPERIMENTS

In this section, we present the results of our proposed ResiDPO for protein sequence design. We begin by describing the construction of our training and evaluation dataset, followed by an ablation study of key hyperparameters. Then we demonstrate the practical impact of ResiDPO on a challenging enzyme design benchmark. Finally, we present a detailed analysis of ResiDPO’s performance compared to standard DPO.

4.1 PDB-D DATASET AND EVALUATION METRICS

To train and evaluate our models, we curated PDB-D, a diverse and high-quality dataset of protein structures with corresponding per-residue pLDDT labels using AlphaFold2 (AF2) Jumper et al. (2021). To ensure the data quality, we curated monomeric structures from the Protein Data Bank (PDB) Berman et al. (2000) determined by X-ray crystallography with a resolution better than 3.5 Å. We restricted our dataset to proteins shorter than 1,000 residues for three principal reasons: (1) AF2 structure prediction for larger proteins is computationally demanding; (2) current backbone generation models often encounter challenges in designing very large structures; and (3) shorter protein sequences are frequently preferred in practical applications. To prevent data leakage, we adopted a rigorous data-splitting strategy. Structures released after September 30, 2021, were designated as the validation set, consistent with prior work Abramson et al. (2024). We performed structure-based clustering and treated clusters containing any validation structure as validation clusters. All structures in the remaining clusters formed the training set. This process resulted in a training set of 19,203 structures and a validation set of 1,690 representative structures. For each structure, we used LigandMPNN Dauparas et al. (2025) to generate eight sequences at a temperature of 1.0 to encourage diversity. These sequences were then predicted by AF2 to obtain per-residue pLDDT scores (Fig. 1).

As described above, to evaluate the designability of sequences on design benchmarks, we employ the gold standard: running full AF2 predictions. We evaluated design success based on the quality of the predicted AF2-structures and pLDDT scores. However, this approach incurs significant computational cost. To enable efficient iteration and evaluation during our extensive ablation studies, we propose the new metric **pLDDT Accuracy**. Unlike the full AF2 evaluation used for the final benchmark, pLDDT Accuracy is an efficient proxy that measures the correlation between the model’s output likelihood for a validation sequence and the sequence’s actual pLDDT score. This allows us to effectively evaluate the model’s success in learning to generate sequences likely to be designable (i.e., achieve high pLDDT) without the computational burden of running AF2 for every ablation configuration. We also track sequence recovery on the LigandMPNN validation set to understand how far the model moves from the original output distribution.

4.2 EXPERIMENTAL DETAILS

Due to the importance of ligand context in enzyme design, we employed LigandMPNN Dauparas et al. (2025), a ligand-aware variant of ProteinMPNN, as our base model. We first conducted a comprehensive grid search to optimize hyperparameters for standard DPO. Subsequently, we utilized these optimized hyperparameters in our proposed ResiDPO method. Unless otherwise stated, all training runs utilized the Adam Kingma & Ba (2015) optimizer with a learning rate of 5e-7 for 100,000 iterations. We implemented a learning rate schedule consisting of a 3% warmup period followed by cosine annealing decay. The training was distributed across two L40 GPUs, utilizing a total batch size of 8 with a gradient accumulation factor of 16.

For ResiDPO, we employed the following parameter settings: $\alpha = 10$, $\beta = 80$, $\gamma = 0.5$, and $\lambda = 0.01$. We set β to 80 based on the observation that, in most practical protein design applications, a pLDDT score greater than 80 generally indicates sufficient designability. Since β and γ are highly correlated, this choice helps reduce the search space, allowing us to only look for the best value for γ .

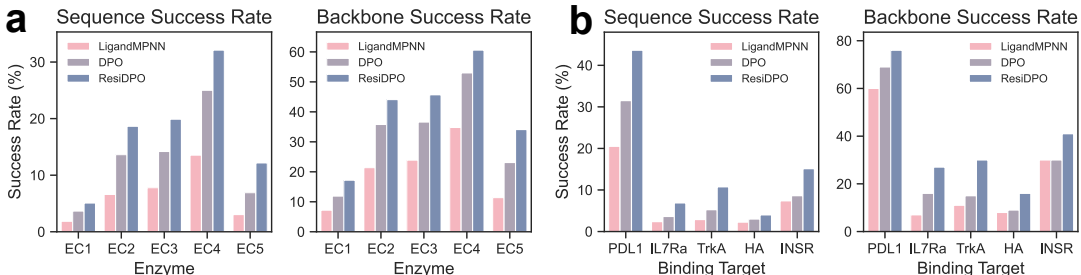


Figure 2: Design success rates for LigandMPNN, DPO-finetuned LigandMPNN, and ResiDPO-finetuned LigandMPNN (EnhancedMPNN) on the enzyme design (a) and binder design (b) benchmarks. EnhancedMPNN, trained with our ResiDPO method to directly optimize for designability, demonstrates significantly higher design success rates. This enhanced designability allows for the successful sequence design of traditionally “undesignable” backbones. The consistent improvement observed in the binder design benchmark highlights the generality of the ResiDPO approach for improving designability across diverse protein design problems, including protein-protein interaction.

4.3 DESIGN BENCHMARK RESULTS

4.3.1 ENZYME DESIGN BENCHMARK

To evaluate the performance of ResiDPO, we used the enzyme active site scaffolding benchmark from RFDiffusion2 Ahern et al. (2025). This benchmark provides five enzymes from five EC classes with defined catalytic pockets (We expanded the benchmark into 41 enzymes in Appendix Sec. A.7). For each enzyme, we generated 1,000 backbones based on the catalytic pocket (Tab. 3) using RFDiffusion2 and designed eight sequences per backbone using LigandMPNN, DPO-finetuned LigandMPNN, and ResiDPO-finetuned LigandMPNN (EnhancedMPNN) at a temperature of 0.1. The catalytic residues were fixed (undesigned) throughout both backbone and sequence generation processes. We assessed design success based on the criteria of pLDDT 80 and C α RMSD 1.5 Å.

As illustrated in Fig. 2a, ResiDPO significantly improved the sequence design success rate compared to both LigandMPNN and DPO. EnhancedMPNN achieved an average sequence design success rate of 17.57%, a nearly threefold improvement over ligandMPNN (6.56%). This improvement was consistent across all five enzymes. Furthermore, the increased designability also led to a significant increase in backbone success rate (i.e., fraction of backbones with at least one successful sequence), from 19.74% for LigandMPNN to 40.34% for EnhancedMPNN. This indicates that ResiDPO *not only improves sequence design but also expands the set of designable backbones*, potentially recalls more designable *de novo* proteins, and saves significant computational resources.

4.3.2 BINDER DESIGN BENCHMARK

To assess the generalization capabilities of our approach beyond the monomeric proteins prevalent in the PDB-D training dataset, we evaluated ResiDPO on a challenging protein binder design task. We utilized the binder benchmark set introduced by RFDiffusion Watson et al. (2023), for which models must design sequences that promote specific inter-chain interactions. For each target protein in the benchmark, 100 backbone structures were generated using RFDiffusion. Subsequently, for each backbone, the baseline LigandMPNN, the DPO-finetuned LigandMPNN, and the ResiDPO-finetuned LigandMPNN were tasked with designing eight sequences. Following RFDiffusion, the *in silico* success of a designed sequence was defined by three criteria: inter-chain PAE ≤ 10 , C α RMSD ≤ 1 Å, and pLDDT ≥ 80 . As shown in Fig. 2b, the baseline LigandMPNN achieved a sequence design success rate of 7.07%. Fine-tuning with standard DPO modestly increased this rate to 10.40%. EnhancedMPNN, leveraging ResiDPO, achieved a sequence design success rate of 16.07%. This represents a substantial improvement, more than doubling (approximately 2.27-fold) the success rate of the baseline LigandMPNN. These results demonstrate that ResiDPO, despite not being explicitly trained on protein complexes or interaction data, confers a strong generalization for improving designability in complex multichain systems.

Table 1: Comparison of different sampling methods and loss components of ResiDPO on the validation set. The best results of each section are labeled in bold.

Sampling	Method	Seq. Recovery	pLDDT Acc.
-	LigandMPNN	57.63	57.71
Rejection Sampling	DPO	56.86	61.23
Application Sampling	DPO	56.29	61.67
Relative Sampling	DPO	57.03	62.11
Relative Sampling	RPL	54.23	63.44
Relative Sampling	ResiDPO	55.56	66.08

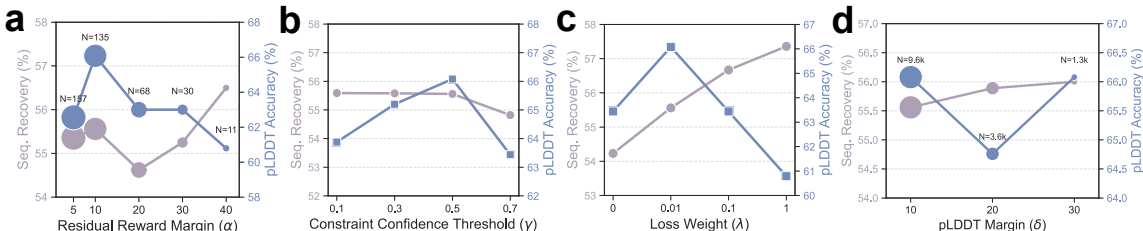


Figure 3: Ablations on the hyperparameters of ResiDPO. The marker size N in panels a and d indicates the number of selected residue pairs ($|Z|$) and remaining backbones, separately.

4.4 ABLATION STUDY

Our ablation studies validate the design of ResiDPO by systematically evaluating its core components and hyperparameter sensitivities. We first established that Relative Sampling is the optimal strategy for generating preference pairs for the DPO baseline, yielding a pLDDT accuracy of 62.11% (Tab. 1). Isolating our Residue-level Preference Learning (RPL) component improved pLDDT to 63.44% but at the cost of sequence recovery (54.23%). The full ResiDPO algorithm, which integrates RPL with Residue-level Constraint Learning (RCL), achieved the highest pLDDT accuracy of 66.08% while maintaining strong sequence recovery (55.56%). This demonstrates that RCL is crucial for preventing knowledge degradation from the pretrained model, synergizing with RPL to effectively learn preferences without catastrophic forgetting.

Key hyperparameters were meticulously tuned for ResiDPO to balance performance and robustness (Fig. 3). The PRL margin ($\alpha = 10$), RCL confidence threshold ($\gamma = 0.5$), and loss weight ($\lambda = 0.01$) were identified to maximize pLDDT accuracy with decent sequence recovery. For preference pair construction, a pLDDT margin of $\delta = 10$ was chosen over $\delta = 30$; while both yielded similar top pLDDT scores, $\delta = 10$ offered a significantly larger and more diverse training set (9,557 vs. 1,283 pairs), promoting better generalization.

4.5 WHAT MAKES IT DIFFERENT: HOW RESIDPO ENHANCES PROTEIN DESIGNABILITY?

To understand the mechanisms by which ResiDPO enhances designability, we analyzed the changes in amino acid composition induced by EnhancedMPNN on both natural (validation set from PDB-D) and designed (enzyme benchmark) backbones.

We generated confusion matrices (Fig. 10) to visualize the residue substitutions introduced by EnhancedMPNN. A consistent trend emerged across both natural and designed backbones: EnhancedMPNN tended to introduce more charged residues, particularly Glutamic Acid (E), Lysine (K), and Arginine (R). Almost 50% of the original amino acids that were originally Alanine (A), Glutamine (Q), Serine (S), and Threonine (T) are mutated to others by EnhancedMPNN. Proline (P) and Glycine (G) remained largely unchanged.

We also compare the overall amino acid distributions generated by different methods for both types of backbones (Fig. 4). The observed trends aligned with the residue-level analysis. Notably, Alanine (A) abundance is significantly reduced in designed backbones, particularly with ResiDPO. While DPO and ResiDPO generally exhibit similar trends, ResiDPO often induces more pronounced changes. For instance, both methods tend to decrease the frequency of Alanine (A) and increase the frequency of Glutamic acid (E) in both natural and designed backbones, but ResiDPO’s effect is considerably stronger.

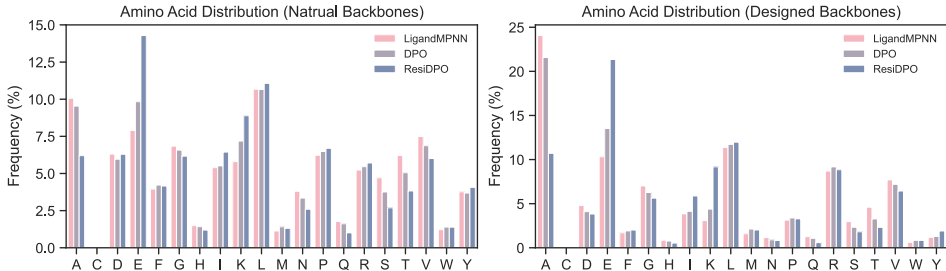


Figure 4: Comparison of amino acid distributions generated by different methods on natural backbones (left) and designed backbones (right).

The overall sequence distribution changes can be understood by considering how amino acid sequences determine protein 3D structures. To a first approximation, polar and charged residues are on the protein surface, where they can interact with water, while hydrophobic residues are buried away from water in the protein core. Structure prediction methods such as AF2, learn these (and many other) patterns. Alanine residues can be buried or exposed, and hence are somewhat ambiguous, as are the small polar residues serine and threonine. In contrast, glycine and proline have unique conformational preferences, making them relatively unambiguous (glycine can adopt conformations inaccessible to other residues, while proline is limited to a subset of local conformations). Hence the EnhancedMPNN overall can be viewed as reducing the ambiguity in the sequence-structure mapping, which increases the AF2 predicted confidence.

4.6 HOW MANY BACKBONES ARE NEEDED FOR EFFECTIVE OPTIMIZATION?

Beyond enhancing sequence designability, ResiDPO offers promise for optimizing other functional properties, particularly those constrained by the high cost of wet-lab experiments and data acquisition. We assessed its data efficiency by comparing ResiDPO against DPO using varying numbers of randomly selected training backbones (Fig. 5).

ResiDPO demonstrated a clear advantage, particularly in low-data regimes. With only 100 samples, DPO’s pLDDT accuracy declined, while ResiDPO, leveraging its targeted residual-level supervision, learned effectively. Notably, ResiDPO with just 1k samples achieved comparable pLDDT accuracy (63.44%) to DPO with 19k samples—a significant data efficiency gain. Given that only 502 of the randomly sampled backbones met our selection criterion ($\delta = 10$), our findings suggest that ResiDPO can achieve substantial performance improvements with a modest set of ~500 carefully chosen backbones, highlighting its practical applicability for broader, resource-constrained optimization tasks.

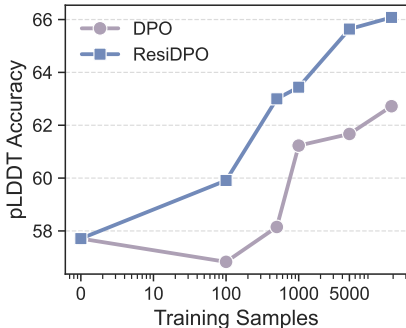


Figure 5: ResiDPO achieves good performance with only a thousand training samples.

5 CONCLUSION

In this work, we introduced Residue-level Designability Preference Optimization (ResiDPO), a novel training strategy designed to bridge the gap between training sequence recovery and inference designability. By performing per-residue comparison, we effectively decouple the traditionally conflicting objectives of Direct Preference Optimization (DPO) into two distinct components: reward learning from improved positions and constraint learning for already well-predicted positions. This refined approach provides a more targeted and effective optimization process, leading to a nearly three-fold increase in design success for enzymes and a two-fold improvement for protein binders. What’s more, by making previously undesignable backbones designable, ResiDPO meaningfully expands the designable space for both backbones and sequences, potentially prompting advancements in backbone design methodologies. While we have demonstrated the efficacy of ResiDPO using LigandMPNN, the method is architecture-agnostic and readily generalizable to other sequence design models. ResiDPO can also be readily generalized to a wide range of design tasks to improve properties such as stability and expressibility. Experimental characterization of EnhancedMPNN-generated sequences will reveal whether the enhanced in silico foldability translates to greater control over structure and higher design success rates.

REFERENCES

- Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/>.
- Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O’Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Židek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, pp. 1–3, May 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w.
- Woody Ahern, Jason Yim, Doug Tischer, Saman Salike, Seth M. Woodbury, Donghyo Kim, Indrek Kalvet, Yakov Kipnis, Brian Coventry, Han Raut Altae-Tran, Magnus Bauer, Regina Barzilay, Tommi S. Jaakkola, Rohith Krishna, and David Baker. Atom level enzyme active site scaffolding using RFdiffusion2, April 2025.
- Namrata Anand, Raphael Eguchi, Irimpan I. Mathews, Carla P. Perez, Alexander Derry, Russ B. Altman, and Po-Ssu Huang. Protein sequence design with a learned potential. *Nature Communications*, 13(1):746, February 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-28313-9.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A General Theoretical Paradigm to Understand Learning from Human Preferences. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pp. 4447–4455. PMLR, April 2024.
- Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N. Kinch, R. Dustin Schaeffer, Claudia Millán, Hahnbeom Park, Carson Adams, Caleb R. Glassman, Andy DeGiovanni, Jose H. Pereira, Andria V. Rodrigues, Alberdina A. van Dijk, Ana C. Ebrecht, Diederik J. Opperman, Theo Sagmeister, Christoph Buhlheller, Tea Pavkov-Keller, Manoj K. Rathinaswamy, Udit Dalwadi, Calvin K. Yip, John E. Burke, K. Christopher Garcia, Nick V. Grishin, Paul D. Adams, Randy J. Read, and David Baker. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, August 2021. doi: 10.1126/science.abj8754.
- Nathaniel R. Bennett, Brian Coventry, Inna Goreschnik, Buwei Huang, Aza Allen, Dionne Vafeados, Ying Po Peng, Justas Dauparas, Minkyung Baek, Lance Stewart, Frank DiMaio, Steven De Munck, Savvas N. Savvides, and David Baker. Improving de novo protein binder design with deep learning. *Nature Communications*, 14(1):2625, May 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-38328-5.
- Helen M. Berman, John Westbrook, Zukang Feng, Gary Gilliland, T. N. Bhat, Helge Weissig, Ilya N. Shindyalov, and Philip E. Bourne. The Protein Data Bank. *Nucleic Acids Research*, 28(1):235–242, January 2000. ISSN 0305-1048. doi: 10.1093/nar/28.1.235. URL <http://www.rcsb.org/>.
- Huayu Chen, Guande He, Lifan Yuan, Ganqu Cui, Hang Su, and Jun Zhu. Noise Contrastive Alignment of Language Models with Explicit Rewards. *NeurIPS*, 37:117784–117812, December 2024.
- Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. Provably Robust DPO: Aligning Language Models with Noisy Feedback. In *ICML*, pp. 42258–42274. PMLR, July 2024.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *NeurIPS, NIPS’17*, pp. 4302–4310, Red Hook, NY, USA, December 2017. Curran Associates Inc. ISBN 978-1-5108-6096-4.

- J. Dauparas, I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. M. Wicky, A. Courbet, R. J. de Haas, N. Bethel, P. J. Y. Leung, T. F. Huddy, S. Pellock, D. Tischler, F. Chan, B. Koepnick, H. Nguyen, A. Kang, B. Sankaran, A. K. Bera, N. P. King, and D. Baker. Robust deep learning-based protein sequence design using ProteinMPNN. *Science*, 378(6615):49–56, October 2022. doi: 10.1126/science.add2187.
- Justas Dauparas, Gyu Rie Lee, Robert Pecoraro, Linna An, Ivan Anishchenko, Cameron Glasscock, and David Baker. Atomic context-conditioned protein sequence design using LigandMPNN. *Nature Methods*, 22(4):717–723, April 2025. ISSN 1548-7105. doi: 10.1038/s41592-025-02626-1.
- DeepSeek-AI. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning, January 2025.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Model Alignment as Prospect Theoretic Optimization. In *ICML*, pp. 12634–12651. PMLR, July 2024.
- Zhangyang Gao, Cheng Tan, and Stan Z. Li. PiFold: Toward effective and efficient protein inverse folding. In *ICLR*, September 2022.
- Zhangyang Gao, Cheng Tan, Xingran Chen, Yijie Zhang, Jun Xia, Siyuan Li, and Stan Z. Li. KW-Design: Pushing the Limit of Protein Design via Knowledge Refinement. In *The Twelfth International Conference on Learning Representations*, October 2023.
- Zhangyang Gao, Jue Wang, Cheng Tan, Lirong Wu, Yufei Huang, Siyuan Li, Zhirui Ye, and Stan Z. Li. UniIF: Unified Molecule Inverse Folding. In *NeurIPS*. arXiv, May 2024. doi: 10.48550/arXiv.2405.18968.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, and et al. The Llama 3 Herd of Models, November 2024.
- Kaipeng Hou, Wei Huang, Miao Qi, Thomas H Tugwell, Turki M Alturaifi, Yuda Chen, Xingjie Zhang, Lei Lu, Samuel I Mann, Peng Liu, et al. De novo design of porphyrin-containing proteins as efficient and stereoselective catalysts. *Science*, 388(6747):665–670, 2025.
- Chloe Hsu, Robert Verkuil, Jason Liu, Zeming Lin, Brian Hie, Tom Sercu, Adam Lerer, and Alexander Rives. Learning inverse folding from millions of predicted structures. In *ICML*, pp. 8946–8970. PMLR, June 2022.
- Po-Ssu Huang, Scott E. Boyken, and David Baker. The coming of age of de novo protein design. *Nature*, 537(7620):320–327, September 2016. ISSN 1476-4687. doi: 10.1038/nature19946.
- John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative Models for Graph-Based Protein Design. In *NeurIPS*, volume 32. Curran Associates, Inc., 2019.
- Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure F. P. Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghui Zhang, Lena Simine, Payel Das, and Yoshua Bengio. Biological Sequence Design with GFlowNets. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 9786–9801. PMLR, June 2022.
- Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror. Learning from Protein Structure with Geometric Vector Perceptrons. In *ICLR*, October 2020.
- Kristoffer Haurum Johansen, Darian Stephan Wolff, Beatrice Scapolo, Monica L Fernández-Quintero, Charlotte Risager Christensen, Johannes R Loeffler, Esperanza Rivera-de Torre, Max D Overath, Kamilla Kjærgaard Munk, Oliver Morell, et al. De novo-designed pmhc binders facilitate t cell-mediated cytotoxicity toward cancer cells. *Science*, 389(6758):380–385, 2025.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman,

- Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstern, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, August 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2.
- Donghyo Kim, Seth M. Woodbury, Woody Ahern, Indrek Kalvet, Nikita Hanikel, Saman Salike, Samuel J. Pellock, Anna Lauko, Donald Hilvert, and David Baker. Computational Design of Metallohydrolases, November 2024.
- D. P. Kingma and L. J. Ba. Adam: A Method for Stochastic Optimization. In *ICLR*. Ithaca, NYArXiv, 2015.
- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J. Liu, and Jialu Liu. Statistical Rejection Sampling Improves Preference Optimization. In *ICLR*, October 2023.
- Tianqi Liu, Zhen Qin, Junru Wu, Jiaming Shen, Misha Khalman, Rishabh Joshi, Yao Zhao, Mohammad Saleh, Simon Baumgartner, Jialu Liu, Peter J. Liu, and Xuanhui Wang. LiPO: Listwise Preference Optimization through Learning-to-Rank, May 2024.
- Yufeng Liu, Lu Zhang, Weilun Wang, Min Zhu, Chenchen Wang, Fudong Li, Jiahai Zhang, Houqiang Li, Quan Chen, and Haiyan Liu. Rotamer-free protein sequence design based on deep learning and self-consistency. *Nature Computational Science*, 2(7):451–462, July 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00273-6.
- Yu Meng, Mengzhou Xia, and Danqi Chen. SimPO: Simple Preference Optimization with a Reference-Free Reward. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024.
- Geraldene Munsamy, Ramiro Illanes-Vicioso, Silvia Funcillo, Ioanna T. Nakou, Sebastian Lindner, Gavin Ayres, Lesley S. Sheehan, Steven Moss, Ulrich Eckhard, Philipp Lorenz, and Noelia Ferruz. Conditional language models enable the efficient design of proficient enzymes, May 2024.
- Ryan Park, Darren J. Hsu, C. Brian Roland, Maria Korshunova, Chen Tessler, Shie Mannor, Olivia Viessmann, and Bruno Trentini. Improving Inverse Folding for Peptide Design with Diversity-regularized Direct Preference Optimization, October 2024.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. *Advances in Neural Information Processing Systems*, 36:53728–53741, December 2023.
- Milong Ren, Chungong Yu, Dongbo Bu, and Haicang Zhang. Accurate and robust protein sequence design with CarbonDesign. *Nature Machine Intelligence*, 6(5):536–547, May 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00838-2.
- Jeffrey A Ruffolo, Stephen Nayfach, Joseph Gallagher, Aadyot Bhatnagar, Joel Beazer, Riffat Hussain, Jordan Russ, Jennifer Yip, Emily Hill, Martin Pacesa, et al. Design of highly functional genome editors by modelling crispr-cas sequences. *Nature*, pp. 1–8, 2025.
- Andrew W. Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander W. R. Nelson, Alex Bridgland, Hugo Penedones, Stig Petersen, Karen Simonyan, Steve Crossan, Pushmeet Kohli, David T. Jones, David Silver, Koray Kavukcuoglu, and Demis Hassabis. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, January 2020. ISSN 1476-4687. doi: 10.1038/s41586-019-1923-7.
- Cheng Tan, Zhangyang Gao, Jun Xia, Bozhen Hu, and Stan Z. Li. Global-Context Aware Generative Protein Design. In *ICASSP*, pp. 1–5, June 2023. doi: 10.1109/ICASSP49357.2023.10095229.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon and U. V. Luxburg and S. Bengio and H. Wallach and R. Fergus and S. Vishwanathan and R. Garnett (ed.), *NeurIPS*, pp. 5998–6008. Curran Associates, Inc., 2017.

- Susana Vázquez Torres, Melisa Benard Valle, Stephen P. Mackessy, Stefanie K. Menzies, Nicholas R. Casewell, Shirin Ahmadi, Nick J. Burlet, Edin Muratspahić, Isaac Sappington, Max D. Overath, Esperanza Rivera-de-Torre, Jann Ledergerber, Andreas H. Laustsen, Kim Boddum, Asim K. Bera, Alex Kang, Evans Brackenbrough, Iara A. Cardoso, Edouard P. Crittenden, Rebecca J. Edge, Justin Decarreau, Robert J. Ragotte, Arvind S. Pillai, Mohamad Abedi, Hannah L. Han, Stacey R. Gerben, Analisa Murray, Rebecca Skotheim, Lynda Stuart, Lance Stewart, Thomas J. A. Fryer, Timothy P. Jenkins, and David Baker. De novo designed proteins neutralize lethal snake venom toxins. *Nature*, 639(8053):225–231, March 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08393-x.
- Jue Wang. Protein sequence design by deep learning. *Nature Computational Science*, 2(7):416–417, July 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00274-5.
- Tong Wang, Xinheng He, Mingyu Li, Yatao Li, Ran Bi, Yusong Wang, Chaoran Cheng, Xiangzhen Shen, Jiawei Meng, He Zhang, Haiguang Liu, Zun Wang, Shaoning Li, Bin Shao, and Tie-Yan Liu. Ab initio characterization of protein molecular dynamics with AI2BMD. *Nature*, 635(8040):1019–1027, November 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-08127-z.
- Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with RFdiffusion. *Nature*, pp. 1–3, July 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06415-8.
- Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. Self-Play Preference Optimization for Language Model Alignment. In *ICLR*, January 2025.
- Fei Ye, Zaixiang Zheng, Dongyu Xue, Yuning Shen, Lihao Wang, Yiming Ma, Yan Wang, Xinyou Wang, Xiangxin Zhou, and Quanquan Gu. ProteinBench: A Holistic Evaluation of Protein Foundation Models. <https://arxiv.org/abs/2409.06744v1>, September 2024.
- Daniel Zaidman, Jaime Prilusky, and Nir London. PROsettaC: Rosetta Based Modeling of PROTAC Mediated Ternary Complexes. *Journal of Chemical Information and Modeling*, 60(10):4894–4903, October 2020. ISSN 1549-9596, 1549-960X. doi: 10.1021/acs.jcim.0c00589.
- Zaixiang Zheng, Yifan Deng, Dongyu Xue, Yi Zhou, Fei Ye, and Quanquan Gu. Structure-informed Language Models Are Protein Designers. In *ICML*, June 2023.
- Xinyi Zhou, Guangyong Chen, Junjie Ye, Ercheng Wang, Jun Zhang, Cong Mao, Zhanwei Li, Jianye Hao, Xingxu Huang, Jin Tang, and Pheng Ann Heng. ProRefiner: An entropy-based refining strategy for inverse protein folding with global graph attention. *Nature Communications*, 14(1):7434, November 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-43166-6.
- Jingyi Zhu, Mingfu Liang, Ke Sun, Yu Wei, Ruiying Guo, Lijing Zhang, Junhui Shi, Dan Ma, Qi Hu, Gaoxingyu Huang, et al. De novo design of transmembrane fluorescence-activating proteins. *Nature*, pp. 1–9, 2025.