DE NOVO DESIGN OF ANTIGEN-SPECIFIC ANTIBODIES USING STRUCTURAL CONSTRAINT-BASED GENERATIVE LANGUAGE MODEL

Yuran Jia[1,2]*, **Bing He**[2]*, **Tianxu Lv**[2], **Yang Xiao**[2], **Tianyi Zhao**[1]†, **Jianhua Yao**[2]†

[1]Faculty of Computing, Harbin Institute of Technology, Harbin, China
[2]Tencent AI for Life Sciences Lab, Shenzhen, China

ABSTRACT

Despite significant advances in computational antibody design, the limited availability of high-quality binding data continues to constrain the exploration of diverse antibody syntax and uncharted evolutionary landscapes. To overcome these challenges, we developed PALM-PA (Pre-trained Antibody Generative Large Language Model–Preference Alignment), which integrates antibody linguistic patterns with structural constraints to explore novel sequence spaces. Experimental validation on influenza A hemagglutinin and programmed death-ligand 1 (PD-L1) demonstrated nanomolar binding affinities (30.2 nM and 1.29 nM, respectively), underscoring the feasibility of using structure-guided language models for the *de novo* design of antibodies.

1 INTRODUCTION

Traditional monoclonal antibody development is limited by the low throughput and high cost of experimental methods, which have driven the rapid development of computational paradigms (structure-guided and sequence-guided) in antibody engineering (Wang et al. (2024); Wu et al. (2024); Ruffolo et al. (2021)). Structure-guided methods design the target antibody scaffold to ensure folding stability and binding specificity, followed by reverse optimization of the amino acid sequence to match preset conformational constraints, demonstrating advantages in single-domain antibodies design (Bennett et al. (2024)). Sequence-guided methods rely on language models trained on large-scale natural antibody repertoires to efficiently explore the vast sequence space by learning the intrinsic grammatical rules and evolutionary preferences of antibodies, and have shown potential in the design of CDRH3 targeting SARS-CoV-2 (He et al. (2024)).

Due to the multilevel nonlinear coupling in the mapping of antibody sequence–structure–function, the two paradigms face complementary challenges. Structure-based methods, which rely on preset scaffolds, can ensure structural compatibility but often compress the sequence search space into local optima, thus limiting the exploration of antibody variants that lie within the underexplored "dark sequence space." In contrast, sequence-based methods, although capable of generating highly heterogeneous candidates, may fail to guarantee structural compatibility and functional feasibility due to their inability to decode implicit structural semantics. To bridge this gap, it is necessary to explore how to internalize the core structural constraint information flow for functional antibody folding and binding while inheriting the global evolutionary exploration patterns of sequence models, so as to systematically mine the noncanonical binding motifs hidden in the deep sequence space and provide a high - confidence pre - enriched library for experimental screening. This "learning–constraint synergy" paradigm will shift antibody design from local optimization toward a more global and emergent exploration.

Motivated by these insights, we propose PALM-PA, which integrates antibody language model with iterative structural optimization algorithms to achieve *de novo* design of antigen-specific CDR loops. Our key contributions include:

*Equal contribution
†Corresponding authors

- PALM-PA: We introduce a novel large language model for antibodies that leverages a multiobjective joint learning strategy guided by antibody syntax.
- Iterative Structural Anchoring Alignment: The Kahneman–Tversky Optimization (KTO) process is integrated into PALM-PA framework, establishing a dual-track feedback mechanism between sequence and structure.
- Disease-related Targets Validation: PALM-PA was validated on disease-relevant targets, including influenza A hemagglutinin and programmed death-ligand 1 (PD-L1), for which high-affinity binding antibodies were designed, achieving highest affinities of 30.2 nM and 1.29 nM, respectively.

2 METHODS

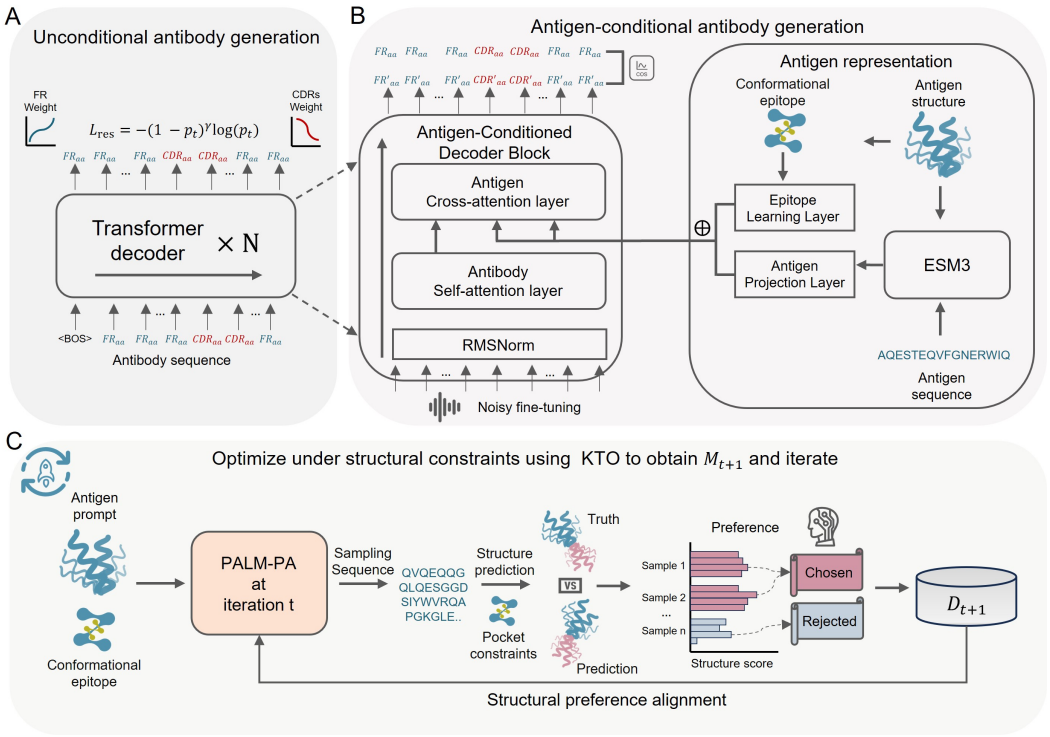

Figure 1: The workflow of PALM-PA. (A) The unconditional antibody generation model is trained on large-scale antibody sequences data, with dynamic adjustment of the weights of FR and CDRs, emphasizing hard-to-learn residues. (B) Antigen information is encoded using ESM3, which is fused with epitope embeddings via a projection layer and integrated through cross-attention for antigen-specific antibody generation. Noise perturbation is applied as a data augmentation strategy, and residue semantic substitution cost is used to guide the learning of design rules. (C) Iterative structural constraint alignment.

2.1 ARCHITECTURE OF BASELINE PALM

As shown in Figure 1, we first constructed an unconditional antibody generation model using the LLaMA Transformer decoder. To address differences in sequence patterns between the antibody framework regions (FR) and complementarity-determining regions (CDRs), we introduced a region-adaptive focal function. This function leverages homoscedastic uncertainty to dynamically balance the contributions of each region to the training loss while emphasizing more challenging residues. To capture antigen–antibody binding patterns, we encoded antigen sequences and structures using the pre-trained ESM3 (Hays et al., 2024), fused these representations with learnable conformational

epitope information via a projection layer, and fed the combined representation into the decoder's cross-attention layer. The cross-attention layer was randomly initialized, whereas the decoder's self-attention layers inherited pre-trained weights from the unconditional antibody model.

Given the limited antigen–antibody binding data, conventional cross-entropy loss functions fail to capture the substitutability of functional residues, thereby hindering exploration of the "dark" sequence space. To overcome this limitation, we redesigned the region-adaptive focal function based on optimal transport principles by incorporating a residue semantic cost. This approach leverages evolutionary information from the pre-trained antibody model to more accurately characterize amino acid dependencies and learn residue similarity and substitutability. In addition, a noise injection fine-tuning strategy was applied at each time step to further enhance the model's adaptability to natural antibody diversity and overall robustness.

2.2 Iterative Alignment with Structural Constraints using Evolutionary-Aware KTO

Direct Preference Optimization (DPO) (Rafailov et al. (2024)) has achieved remarkable success in aligning large language models using explicit paired preference data. However, its direct application to antigen–antibody structural optimization faces inherent challenges. DPO is sensitive to imperfect preference annotations. The subtle and interdependent conformational features of antigen–antibody complexes—such as the conformational flexibility at the antigen binding site and the dynamic coordination between epitope and paratope interfaces—render the generation of reliable preference annotations exceedingly challenging. These characteristics induce non-transitive preferences, which often lead to optimization bias, overfitting, and a reduction in the sequence diversity that is critical for robust antibody function.

To address these issues, we propose an improved iterative KTO (Ethayarajh et al. (2024)) that incorporates an evolutionarily-aware semantic cost term. Building upon the traditional KTO loss, the semantic cost term quantifies the evolutionary feasibility of residue substitutions, thus preserving biologically plausible sequence variation while optimizing for structural compatibility. First, we establish a structural validity criterion $I_{\text{des}}(x, y)$ that evaluates antigen–antibody complexes using predefined geometric compatibility metrics within the binding pocket. Samples that exceed a composite quality threshold τ_{des} are labeled as desirable ($I_{\text{des}}(x, y) = 1$); otherwise, they are labeled as undesirable ($I_{\text{des}}(x, y) = 0$).

During the $(t + 1)$th iteration, our overall optimization objective integrates two key components:

$$\begin{aligned}\mathcal{L}_{t+1}(\pi_{t+1}, \pi_t) =& \mathbb{E}_{(x,y)\sim\mathcal{D}_{t+1}}\left[w(y)\left(1 - \sigma\left(\beta\,\Delta(x, y)\right)\right)\right] \\ &+ \alpha \inf_{\gamma\in\Pi(Q_\theta(\cdot|x),\, P(\cdot|x))} \mathbb{E}_{(v_i,v_j)\sim\gamma}\left[\chi_{(x,y)\in\mathcal{S}} C(v_i, v_j)\right]\end{aligned}$$

The dataset $\mathcal{D}_{t+1}$ is constructed by sampling from the policy π_t at iteration t and incorporating historical data. The first term of the overall loss represents the KTO loss, which drives the model update by maximizing the reward difference between desirable and undesirable outputs. The second term is the semantic cost term that quantifies the evolutionary feasibility of residue substitutions, with its weight controlled by the hyperparameter α. This term is computed by finding the optimal joint distribution between the model distribution $Q_\theta(\cdot \mid x)$ and the data distribution $P(\cdot \mid x)$, where $C(v_i, v_j)$ denotes the cost of transporting probability mass between tokens v_i and v_j in the semantic embedding space.

3 Results

3.1 Structural preference optimization can improve the quality of sampled antibodies

First, we performed structure-guided iterative optimization on the baseline PALM over five rounds. In each round, 2,000 optimization steps were executed, generating 500 candidate sequences per antigen–antibody complex. The three-dimensional conformations of these candidates were predicted using both Chai-1 (Discovery et al. (2024)) and tFold (Wu et al. (2024)). To provide feedback signals

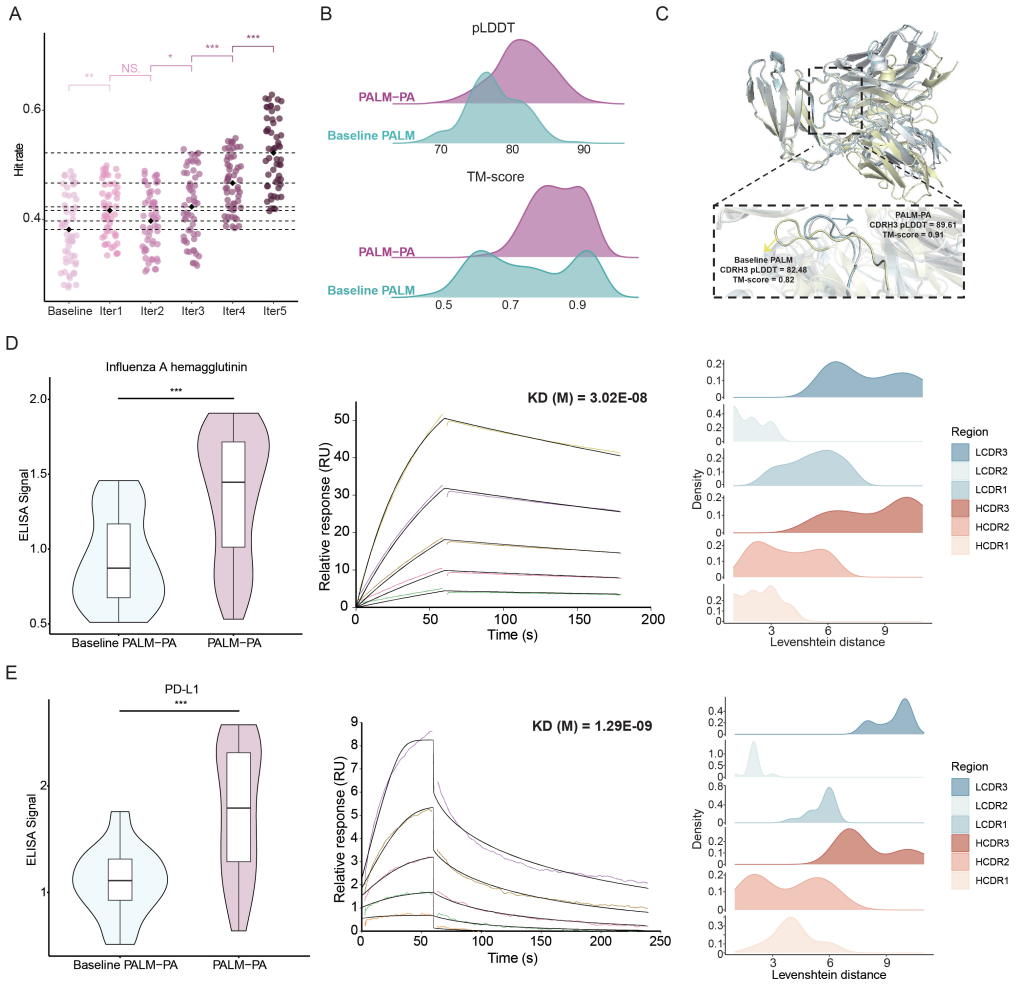

Figure 2: In silico simulation and experimental validation of PALM–PA. (A) Iterative optimization trajectory on the RAbD benchmark dataset. The proportion of "desirable" samples in antigen–antibody complexes progressively increases. (B) Comparison of pLDDT and TM-score distributions before and after structural optimization for baseline PALM and PALM–PA in the RAbD dataset. (C) Comparison of the designed antibody from baseline PALM (yellow) and PALM–PA (blue) against the ground truth (grey) complex (PDB: 5ggs). The structures shown are the highest pLDDT conformations among the predicted complex structures of sampled sequences from baseline PALM and PALM–PA. (D–E) For influenza A hemagglutinin and PD-L1, from left to right: the ELISA binding signals for both baseline PALM and PALM–PA; the SPR binding curve of the binding antibody exhibiting the strongest affinity; and the minimum edit distance between the antibody's CDRs and those of the positive antibodies in the training set. Note: "baseline PALM" indicates the model without structural constraints, and "PALM–PA" refers to the model after iterative structural constraint optimization; ELISA absorbance was measured at 450 nm, with specific binding defined as an absorbance greater than 0.5 and at least three times the background absorbance; Statistical significance was determined using t-test. $*$ indicates $p < 0.05$, $**$ indicates $p < 0.01$, and $***$ indicates $p < 0.001$.

for model refinement, the generated candidates were evaluated using two metrics: TM-score exceeding 0.75 and the average predicted local distance difference test (pLDDT) score for the CDR loops above 80. A candidate was considered desirable only if the conformations predicted by both Chai-1 and tFold met the thresholds. Candidates meeting these criteria were labeled as "desirable samples," while those failing to meet both thresholds were designated as "rejected samples." This classifica-

tion feedback was then incorporated into subsequent optimization rounds. To evaluate the structural optimization performance of PALM-PA, we conducted experiments on the antigen–antibody complex dataset RAbD. For each complex, we sampled 100 sequences using both baseline PALM and PALM-PA at each iteration, enabling a direct comparison of the generated candidate sequences.

As demonstrated in Figure 2A, PALM-PA enhanced the average hit rate of candidate antibodies by 36% (from 38.7% to 52.6%) over five iterative refinement cycles. Comparative structural analysis revealed substantial advantages of PALM-PA over baseline PALM (Figure 2B). The constrained optimization protocol significantly improved complementarity-determining region (CDR) quality metrics, elevating median pLDDT score from 76.78 to 81.34 ($p < 0.001$) and enhancing structural fidelity as evidenced by TM-score improvement from 0.74 to 0.85 ($p < 0.001$). These quantitative enhancements indicate superior structural stability and native-like conformational sampling in PALM-PA generated antibodies. Figure 2C illustrates a representative example (PDB: 5ggs): before optimization, the highest-pLDDT sequence from baseline PALM showed significant deviation in the CDRH3 loop, whereas after five rounds, the corresponding sequence from PALM-PA achieved a TM-score of 0.93, demonstrating markedly improved alignment. These results suggest that the iterative refinement process enables PALM-PA to internalize critical antigen-antibody interface constraints through structural-aware feedback learning, effectively bridging sequence-structure relationships.

3.2 De Novo Design of Antigen-Specific Antibodies

To evaluate the practicality of PALM-PA, we employed a widely used standard antibody framework (Ewert et al. (2003)) as the starting point for de novo design and validated our approach by designing CDRs for influenza A hemagglutinin and PD-L1. Specifically, we established screening criteria by integrating TM-score > 0.8, CDRs conformational stability (pLDDT > 80), and atomic-level accuracy (PAE < 6). In each iteration, samples that met these thresholds and whose CDRH3 exhibited a Levenshtein distance within the top 50% relative to a positive control were selected as "desirable samples" and fed back into the optimization loop, thereby guiding PALM-PA to sample within a high-quality subspace of the antibody deep sequence space.

For influenza A hemagglutinin (Figures 2D) and PD-L1 (Figures 2E), we conducted sequence sampling using both baseline PALM and PALM-PA. Subsequent evaluation using Enzyme-Linked ImmunoSorbent Assay (ELISA) revealed that the structural optimization implemented in PALM-PA led to stronger binding signals. Subsequent surface plasmon resonance (SPR) experiments on the designed antibodies revealed that the highest affinity antibodies had K_d values of 30.2 nM for influenza A hemagglutinin and 1.29 nM for PD-L1. These results demonstrate that, guided by structural optimization, PALM-PA effectively explores the deep sequence space to design antibodies that achieve specific antigen interactions.

4 Conclusion

In this work, we introduced PALM-PA, a framework that integrates pre-trained large language models with iterative structural preference optimization for antigen-specific antibody *de novo* design. Structural information, while valuable, often fails to provide sufficient insights into antigen-antibody binding interfaces, particularly in flexible and conformationally variable regions (Towse & Daggett (2012); Zheng et al. (2023)). In this context, language models exhibit an inherent capability to generate and evaluate sequences beyond the boundaries of natural training datasets, effectively overcoming these limitations. By introducing a paradigm driven by iterative structural constraints, PALM-PA systematically guides antibody sequence design. This approach significantly enhances the capacity of language models to explore novel syntactic combinations within defined structural constraints, facilitating the efficient discovery of diverse and high-quality antibody sequences tailored to specific antigenic targets. Validation across disease-related targets demonstrated the robustness and efficiency of PALM-PA, highlighting its capacity to guide antibody design under specific structural constraints and expand the possibilities of antibody engineering.

Despite these advancements, PALM-PA has limitations that require further exploration. Its reliance on structural feedback highlights the need for more versatile and robust strategies to guide sequence optimization, particularly in scenarios where structural information is sparse or unavailable. This

limitation poses challenges for designing antibodies targeting antigens without existing antibody binders or structural references. Future efforts should aim to reduce dependence on structural inputs and develop generalized optimization approaches to expand the applicability of PALM-PA to uncharacterized or novel antigens.

REFERENCES

Brennan Abanades, Tobias H Olsen, Matthew I J Raybould, Broncio Aguilar-Sanjuan, Wing Ki Wong, Guy Georges, Alexander Bujotzek, and Charlotte M Deane. The patent and literature antibody database (plabdab): An evolving reference set of functionally diverse, literature-annotated antibody sequences and structures. *Nucleic Acids Research*, 52(D1):D545–D551, 2024. doi: 10.1093/nar/gkad1056.

Jared Adolf-Bryfogle, Oleks Kalyuzhniy, Michael Kubitz, Brian D. Weitzner, Xiaozhen Hu, Yumiko Adachi, William R. Schief, and Roland L. Dunbrack Jr. RosettaAntibodyDesign (RAbD): A general framework for computational antibody design. *PLOS Computational Biology*, 14(4): e1006112, 2018. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1006112.

Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. Efficient training of language models to fill in the middle. *arXiv*, 2022. doi: 10.48550/arXiv.2207.14255.

Nathaniel R. Bennett, Joseph L. Watson, Robert J. Ragotte, Andrew J. Borst, Déjenaé L. See, Connor Weidle, Riti Biswas, Ellen L. Shrock, Philip J. Y. Leung, Buwei Huang, Inna Goreshnik, Russell Ault, Kenneth D. Carr, Benedikt Singer, Cameron Criswell, Dionne Vafeados, Mariana Garcia Sanchez, Ho Min Kim, Susana Vázquez Torres, Sidney Chan, and David Baker. Atomically accurate de novo design of single-domain antibodies. *bioRxiv*, 2024. doi: 10.1101/2024.03.14.585103.

Piotr Deszyński, Jakub Młokosiewicz, Adam Volanakis, Igor Jaszczyszyn, Natalie Castellana, Stefano Bonissone, Rajkumar Ganesan, and Konrad Krawczyk. Indi—integrated nanobody database for immunoinformatics. *Nucleic Acids Research*, 50(D1):D1273–D1281, 2022. doi: 10.1093/nar/gkab1021.

Chai Discovery, Jacques Boitreaud, Jack Dent, Matthew McPartlon, Joshua Meier, Vinicius Reis, Alex Rogozhnikov, and Kevin Wu. Chai-1: Decoding the molecular interactions of life. *bioRxiv*, 2024. doi: 10.1101/2024.10.10.615955.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization. *arXiv*, 2024. doi: 10.48550/arXiv.2402.01306.

Stefan Ewert, Thomas Huber, Annemarie Honegger, and Andreas Plückthun. Biophysical Properties of Human Antibody Variable Domains. *Journal of Molecular Biology*, 325(3):531–553, 2003. ISSN 0022-2836. doi: 10.1016/S0022-2836(02)01237-8.

Haohuai He, Bing He, Lei Guan, Yu Zhao, Feng Jiang, Guanxing Chen, Qingge Zhu, Calvin Yu-Chian Chen, Ting Li, and Jianhua Yao. De novo generation of sars-cov-2 antibody cdrh3 with a pre-trained generative large language model. *Nature Communications*, 15(1):6867, 2024. doi: 10.1038/s41467-024-50903-y.

Tobias H. Olsen, Fergus Boyles, and Charlotte M. Deane. Observed antibody space: A diverse database of cleaned, annotated, and translated unpaired and paired antibody sequences. *Protein Science*, 31(1):141–146, 2022. doi: 10.1002/pro.4205.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv*, 2024. doi: 10.48550/arXiv.2305.18290.

Jeffrey A. Ruffolo, Jeffrey J. Gray, and Jeremias Sulam. Deciphering antibody affinity maturation with language models and weakly supervised learning. *arXiv*, 2021. doi: 10.48550/arXiv.2112.07782.

Constantin Schneider, Matthew I J Raybould, and Charlotte M Deane. Sabdab in the age of biotherapeutics: Updates including sabdab-nano, the nanobody structure tracker. *Nucleic Acids Research*, 50(D1):D1368–D1372, 2022. doi: 10.1093/nar/gkab1050.

Clare-Louise Towse and Valerie Daggett. When a domain isn't a domain, and why it's important to properly filter proteins in databases. *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology*, 34(12):1060–1069, 2012. doi: 10.1002/bies.201200116.

Rubo Wang, Fandi Wu, Xingyu Gao, Jiaxiang Wu, Peilin Zhao, and Jianhua Yao. Iggm: A generative model for high-resolution antibody structural prediction with imputation. *arXiv*, 2024. doi: 10.48550/arXiv.2402.03672.

Xiang Wu, Benjamin Wilson, Inna Goreshnik, M. Srivastava, and Jonathan Pappas. Fast and accurate modeling of de novo designed antibodies using deep learning. *bioRxiv*, 2024. doi: 10.1101/2024.02.05.420264.

Zaixiang Zheng, Yifan Deng, Dongyu Xue, Yi Zhou, Fei YE, and Quanquan Gu. Structure-informed Language Models Are Protein Designers. *bioRxiv*, 2023. doi: 10.48550/arXiv.2302.01649.

A Appendix

A.1 Dataset

To pretrain the unconditional antibody generation model, we curated a comprehensive dataset from Observed Antibody Space (OAS) (Olsen et al. (2022)), INDI (Deszyński et al. (2022)) and PLAbDab (Abanades et al. (2024)), enabling the model to learn universal antibody syntax. Rigorous data-cleaning steps were applied, including removing sequences with invalid amino acid characters, incomplete CDRs, and redundant CDRH3 sequences based on 90% sequence identity and 80% sequence coverage thresholds. This resulted in a dataset comprising 87,138,696 single-chain antibodies and 166,917 paired antibodies.

To train the baseline PALM, antigen-antibody complexes were gathered from the SAbDab (Schneider et al. (2022), entries updated as of December 31, 2023) and INDI. To reserve the RAbD dataset (Adolf-Bryfogle et al. (2018)) for testing, complexes from it and entries with a CDRH3 sequence similarity greater than 0.4 to those in it were excluded. Further filtering criteria included the exclusion of entries lacking heavy chains or antigens, those with incomplete structural information (CDRH3 shorter than 7 residues or longer than 26 residues, and antigen residue coverage less than 90%). Additionally, a masked language modeling paradigm was introduced (Bavarian et al. (2022)), where masked CDRs were appended to sequence ends, enabling the model to focus on specific regions for antibody design and ensuring the effective generation of targeted antibody sequences.