ANTIGEN-SPECIFIC ANTIBODY MULTI-MODAL FOUNDATION MODEL FOR FUNCTIONAL ANTIBODY DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Antibodies play a key role in immune recognition by binding specific antigens. Although recent protein language models have enabled progress in single-chain protein modeling and generation, they often fall short in antigen-specific antibody design, where effective modeling requires explicit pairing between antibody and antigen, particularly at the epitope level. To address these limitations, we introduce AAMFM, an **A**ntigen-specific **A**ntibody **M**ultimodal **F**oundation **M**odel that learns unified representations of antibody sequences and structures conditioned on antigen context. AAMFM incorporates rich antigen information including geometric interfaces and epitope annotations via a cross-modal adapter, enabling joint modeling of antibody-antigen interactions in a shared latent space. To further guide the model toward functional relevance, we fine-tune AAMFM using Calibrated Direct Preference Optimization (Cal-DPO), leveraging preference signals extracted from a strong structural prior (AlphaFold3) to align learning with binding-specific objectives. Extensive experiments demonstrate that AAMFM achieves state-of-the-art performance in functional antibody design, revealing its potential for antigen-specific antibody engineering. Our code is available at https://anonymous.4open.science/r/AAMFM.

1 INTRODUCTION

Antibodies are key components of the immune system responsible for recognizing and neutralizing antigens. Their precise targeting capability predominantly arises from the complementarity-determining regions (CDRs), which are the primary determinants of binding affinity (Jones et al., 1986; Ewert et al., 2004; Xu & Davis, 2000; Akbar et al., 2021). Given the central role of therapeutic antibodies in modern medicine, developing methods to design CDRs that are both functional and biologically plausible is critical for creating novel and effective treatments. However, the extreme sequence diversity of CDRs and their tight coupling between sequence, structure, and function make rational design particularly challenging.

While computational approaches have shown promise in accelerating protein engineering, antibody design remains difficult. General protein language models (PLMs) (Lin et al., 2023; Nijkamp et al., 2023; Ruffolo et al., 2024; Hayes et al., 2025) largely focus on sequence modeling and fail to capture protein complex interactions, whereas specialized antibody co-design models (Jin et al., 2021; Luo et al., 2022; Kong et al., 2022; 2023; Zhu et al., 2024; Wang et al., 2024) are limited by scarce data, leading to overfitting and poor generalization. Moreover, unreliable energy-based scoring functions such as Rosetta (Alford et al., 2017) show weak correlation with experimental affinities (Luo et al., 2023), hindering effective guidance. Although AlphaFold3 (AF3) (Abramson et al., 2024) provides structure confidence metrics (e.g., ipTM) that correlate with interface quality (Lu et al., 2024; Hitawala & Gray, 2024; Wee & Wei, 2024), how to leverage such structural priors to guide functional, antigen-specific antibody generation remains an open challenge.

To address these challenges, we introduce AAMFM, an Antigen-specific Antibody MultiModal Foundation Model built upon the multimodal protein language model ESM3 (Hayes et al., 2025). AAMFM is first adapted to the antibody domain using large-scale sequence–structure antibody pairs and subsequently fine-tuned on high-resolution antibody–antigen structure data, where antigen geometry and epitope information are incorporated via a cross-modal adapter. Crucially, AAMFM further enhances functional relevance through preference optimization, in which candidate anti-

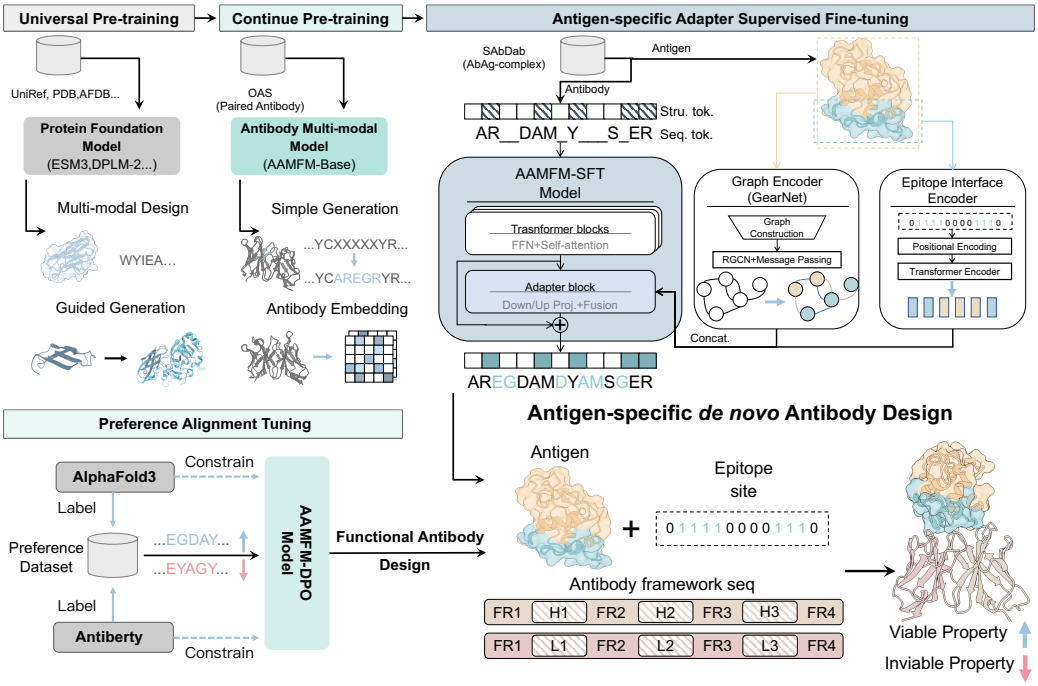

Figure 1: Hierarchical training and inference pipeline of the proposed model, including pre-training, adapter-based SFT and preference optimization.

bodies are evaluated by predicting their complex structures with the target antigen. The resulting structural quality metrics, specifically the AF3 score, capture both intrinsic foldability and predicted antigen-binding capability and are used to define preference signals. Finally, AAMFM is aligned using calibrated direct preference optimization on this preference data, enabling the model to generate antibodies that satisfy both structural plausibility and antigen-specific binding requirements.

2 METHODS

2.1 MULTI-LEVEL PRETRAINING AND FINE-TUNE STRATEGIES

To jointly design the sequences of all CDRs and the full antibody structure given antigen information, epitope position, and framework, we adopt a multi-stage learning framework initialized with ESM3, as shown in Figure 1. While ESM3 captures broad protein representations, antibodies exhibit unique structural constraints. To bridge this domain gap and enhance functional generation, the following hierarchical strategies are implemented:

Antibody Domain Adaptation Pre-training. To address the gap between the general model and the antibody domain, we perform adaptation using 1.4 million paired antibody sequence-structure data from the Observed Antibody Space (OAS) database (Olsen et al., 2022). The structure of the corresponding sequences is predicted by AbodyBuilder2 (Abanades et al., 2023) and IgFold (Ruffolo et al., 2023), which has been proved useful in previous works (Dreyer et al., 2023; Høie et al., 2024). This stage aligns the model with the underlying sequence–structure distribution of antibodies.

Antigen-specific Supervised Fine-tuning and Preference Optimization. Following domain adaptation, we introduce higher-quality data for supervised fine-tuning(SFT) utilizing a large number of high-resolution antibody-antigen complexes from Structural Antibody Database (SAbDab) (Dunbar et al., 2014) to learn more precise atomic-level interactions. In parallel, antigen information is incorporated through an adapter (Section 2.2). Building on this, an additional preference optimization stage is applied using a DPO-style method (Section 2.3) to further improve generation quality and better align the model with functional and plausible design objectives.

2.2 ANTIGEN GEOMETRIC-EPITOPE-AWARE ADAPTER

To integrate antigen geometric and antigen-antibody interface information, we propose the Antigen Geometric-Epitope-aware Adapter. This module is designed as a lightweight adapter (Houlsby et al., 2019), integrated via a residual connection (He et al., 2016) after the final attention block's output h. It computes an adjustment term Δh to produce a refined output $h' = h + \Delta h$. The designed adapter primarily consists of the following modules:

Antigen Feature Encoder The antigen graph G_{ag} is input into GearNet (Zhang et al., 2023) to obtain geometric node feature f_{geo} which are followingly encoded by a convolutional network into the geometric antigen representationh_{ag} to address the variability in antigen lengths. If the complex does not have antigen (f_{geo} is not provided), a learnable default feature is used instead.

Epitope Interface Encoder This module processes binary information I representing the epitope site on the antigen. Similar to previous work (Deng et al., 2024), the interface information is first embedded and combined with learnable positional encodings. Subsequently, a Transformer encoder layer captures contextual information and outputs position-wise interface features h_{ep}.

Feature Fusion The backbone features h are down-projected and concatenated with the antigen geometry and interface features h_{ag}, h_{ep}. These combined features are fused via an MLP and subsequently up-projected to the original dimension, yielding the final adjustment term Δh for the residual connection.

This module effectively incorporates the antigen's geometric information and epitope interface positions into the model's representations, enhancing its antigen specificity in antibody design tasks.

2.3 ALPHAFOLD3-BASED PREFERENCE OPTIMIZATION

Training Objective While Reinforcement Learning effectively align models with specific objectives (Ouyang et al., 2022), Direct Preference Optimization (DPO) (Rafailov et al., 2023) has emerged as a more stable, contrastive, and simple alternative. However, standard DPO focuses mainly on the relative values of implicit rewards associated with two responses while ignoring their actual values, resulting in suboptimal alignment with human preferences. We employ Calibrated DPO (Cal-DPO) (Xiao et al., 2024) to address this limitation by incorporating an additional calibration loss, encouraging the model's output confidence to better reflect actual sequence quality.

Let π_θ and π_{ref} denote the policy (fine-tuned) and reference (SFT) models, respectively. The dataset $\mathcal{D}$ consists of preference tuples (c, y_w, y_l), where c represents the context (antibody framework s_{fr}, antigen G_{ag}, and interface I), while y_w and y_l denote the preferred and rejected CDR sequences. The standard DPO loss is formulated as:

$$\mathcal{L}_{\text{DPO}}(\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(c,y_w,y_l)\sim\mathcal{D}} \left[\log \sigma \left(\beta \left(\hat{r}_\theta(c, y_w) - \hat{r}_\theta(c, y_l)\right)\right)\right] \tag{1}$$

where σ is the sigmoid function and $\hat{r}_\theta(c, y) = \log \frac{\pi_\theta(y|c)}{\pi_{\text{ref}}(y|c)}$ represents the log-likelihood ratio implicit reward. Crucially, we compute log-likelihoods exclusively over tokens within the CDR regions, normalized by CDR length to mitigate length bias.

To ensure the model's confidence reflects the quality of the generated samples, Cal-DPO adds a calibration term. Defining a target margin $M = 1/(2\beta)$, this loss forces the likelihood ratios to align with the margin:

$$\mathcal{L}_{\text{Cal}}(\theta; \pi_{\text{ref}}) = \mathbb{E}_{(c,y_w,y_l)\sim\mathcal{D}} \left[\left(\hat{r}_\theta(c, y_w) - M\right)^2 + \left(\hat{r}_\theta(c, y_l) + M\right)^2\right] \tag{2}$$

The final objective is a weighted sum (in our implementation, with a weighting factor $\lambda = 1$):

$$\mathcal{L}_{\text{Cal-DPO}} = \mathcal{L}_{\text{DPO}} + \lambda \cdot \mathcal{L}_{\text{Cal}} \tag{3}$$

Preference Definition To construct the preference dataset, we sample candidate sequences using the SFT model. After filtering duplicates, we score the remaining sequences using a reproduced AF3 (Protenix (ByteDance et al., 2025)) and AntiBERTy (Ruffolo et al., 2021). The AF3 score serves as a proxy for functionality, assessing both antibody foldability and antigen-binding confidence. The pseudo-log-likelihood (PLL) from AntiBERTy indicates plausibility by reflecting the sequence's similarity to natural antibodies. To align the model with both objectives, we define a strict pairwise preference: for two candidates targeting the same antigen, one is labeled as preferred only if it exceeds the other by at least 0.2 in AF3 score and 0.1 in PLL; otherwise, the pair is discarded.

3 EXPERIMENTS

To evaluate both the generative capabilities of our model, we conduct experiments on two primary tasks: *de novo* antibody co-design and epitope-binding antibody CDR-H3 co-design. In addition, we perform comprehensive ablation studies to assess the contributions of each component. Experiment details and ablation are proved in Appendix B.2.

3.1 *De novo* ANTIBODY SEQUENCE-STRUCTURE CO-DESIGN

Setup We adopt the open-source version of ESM3(1.4B) as the base model for training. The first-stage training is performed on a subset of OAS dataset for 2 epochs. In the second-stage SFT, a curated version of the SAbDab dataset is utilized to train the model for 45 epochs. Antibodies are clustered based on 50% sequence identity in the CDR-H3 region and RAbD dataset (Adolf-Bryfogle et al., 2018) is used as the test set. The resulting fine-tuned model is referred to as **AAMFM-SFT**. We then perform preference alignment via Cal-DPO. Specifically, we construct preference pairs by sampling 13,000 candidate sequences from AAMFM-SFT and evaluating them with Protenix and Antiberty. Following 4 epochs of fine-tuning on these pairs, the model is termed **AAMFM-CalDPO**.

Baselines For antibody sequence-structure co-design, we compare AAMFM with a series of methods, including: **ESM3** (Hayes et al., 2025), a multimodal protein foundation model; **Diffab** (Luo et al., 2022); **AbX** (Zhu et al., 2024), generative diffusion models; **dyMEAN** (Kong et al., 2023), a graph discriminative model.

Metrics Designs are evaluated at both the individual CDR and complete antibody levels. For full antibodies, the following metrics are reported: (1) AntiBERTy pseudo-log-likelihood (Pll) to gauge sequence plausibility; (2) AF3Score (ranking score), pTM, and ipTM to reflect foldability and antigen-binding functionality; and (3) the proportion of hydrophobic residues (PHR) to assess specificity. For CDRs, Amino Acid Recovery (AAR) measures sequence identity, and Root Mean Square Deviation (RMSD) of C_α atoms quantifies structural deviation from the native state.

Results We fix the antibody framework sequence and design the sequences of all six CDR regions along with the complete antibody structure, generating 10 antibody samples per antigen.

Table 1: Results of full antibody design on SAbDab dataset

Method	Pll ↑	AF3Score ↑	PHR ↓	pTM ↑	ipTM ↑
ESM3-open	-1.40	0.846	**0.400**	0.876	0.839
Diffab	-1.36	0.850	0.445	0.881	0.842
dyMEAN	-1.23	0.842	0.465	0.874	0.840
AbX	-0.96	0.862	0.480	<u>0.892</u>	0.855
AAMFM-SFT	<u>-0.91</u>	<u>0.870</u>	0.465	0.890	<u>0.865</u>
AAMFM-CalDPO	**-0.87**	**0.883**	<u>0.443</u>	**0.909**	**0.877**

Table 1 shows the evaluation results for full antibody design. Both variants of AAMFM show strong performance compared to existing SoTA models. Notably, AAMFM-CalDPO achieves the best scores across all four key metrics—Pll, AF3Score, pTM, and ipTM—highlighting its strength in generating functionally sound, plausible, and antigen-specific sequences.

Table 2 presents a fine-grained evaluation of each CDR region. AAMFM achieves performance comparable to state-of-the-art models in terms of AAR and RMSD, with AAMFM-SFT version outperforming them in 6 out of the 12 evaluated subregions. This indicates its superior capability in capturing local structural and sequence features, producing designs that closely match the natural distribution of antibodies.

3.2 EPITOPE-BINDING CDR-H3 CO-DESIGN

Setup, Baselines and Metrics The overall experimental setup, baselines, and evaluation metrics largely follow Section 3.1, with three key differences. First, only the CDR-H3 sequence is designed, while the full antibody structure is still generated. Second, we introduce an additional baseline, GeoAB (Lin et al., 2024), a graph model specifically designed for CDR-H3 co-design. Third, to emphasize local CDR-H3 properties, the evaluation protocol is augmented by (i) computing the PLL score only within the CDR-H3 region (Ye et al., 2024) and (ii) reporting A-RMSD, which measures the RMSD of the aligned CDR-H3 region using the Kabsch algorithm.

Table 2: Results of CDRs design on SAbDab dataset

Method	CDR-H1		CDR-H2		CDR-H3	
	AAR(%) ↑	RMSD ↓	AAR(%) ↑	RMSD ↓	AAR(%) ↑	RMSD ↓
ESM3-open	67.06	0.94	49.54	0.89	32.19	3.31
Diffab	62.32	0.92	44.82	0.96	32.33	3.15
dyMEAN	75.71	1.09	66.67	0.98	38.46	3.51
AbX	79.92	**0.85**	69.85	**0.76**	**43.24**	2.80
AAMFM-SFT	**81.94**	0.90	**70.86**	0.90	40.32	**2.54**
AAMFM-CalDPO	77.00	0.92	66.18	0.93	36.23	2.62
Method	CDR-L1		CDR-L2		CDR-L3	
	AAR(%) ↑	RMSD ↓	AAR(%) ↑	RMSD ↓	AAR(%) ↑	RMSD ↓
ESM3-open	49.83	0.82	51.72	0.70	44.43	1.27
Diffab	48.07	0.86	52.22	0.56	43.31	1.39
dyMEAN	69.86	0.71	80.00	**0.41**	57.78	1.79
AbX	**79.37**	**0.70**	82.96	0.42	65.96	1.18
AAMFM-SFT	75.57	0.95	**83.87**	0.73	**72.70**	**1.11**
AAMFM-CalDPO	70.48	1.02	79.01	0.74	68.24	1.16

Table 3: Results of CDR-H3 antibody design on SAbDab dataset

Method	Pll↑	AF3Score↑	A-RMSD↓	RMSD↓	AAR↑	PHR↓	ipTM ↑
ESM3-open	-1.88	0.854	2.02	3.32	33.04	**0.392**	0.848
Diffab	-1.69	0.863	2.12	2.99	32.90	0.443	0.856
dyMEAN	-1.57	0.863	1.69	3.32	39.49	0.433	0.857
GeoAB	-1.60	0.856	**1.41**	**1.67**	40.40	0.413	0.850
AbX	-1.46	0.877	2.06	3.06	**43.14**	0.509	0.871
AAMFM-SFT	-1.20	0.879	1.55	2.43	40.55	0.432	0.878
AAMFM-CalDPO	**-1.18**	**0.887**	1.62	2.57	37.16	0.440	**0.882**

Results In this setting, we employ two variants of our model to design only the CDR-H3 sequence and the full antibody structure. Table 3 compares our models with baseline methods across several evaluation metrics. Our models consistently outperform state-of-the-art approaches in terms of sequence plausibility (Pll) and functional quality (AF3Score and ipTM). They also exhibit competitive performance on other metrics, such as AAR and RMSD, further demonstrating their effectiveness in viable, realistic antibody design.

4 Conclusion

In this work, we present an antigen-specific antibody multi-modal foundation model AAMFM for antibody functional design. This model leverages a powerful pre-trained protein language base model (ESM3) and uniquely integrates antibody sequence, structure, and specific antigen information, including geometric features and epitope sites, into a unified representation space using a lightweight adapter. Furthermore, AAMFM incorporates a novel finetune stage using Cal-DPO, leveraging preference signals extracted from a strong structural prior(AF3), thereby aligning the generation process with key objectives of functional and plausible antibody design. Experimental results demonstrate that AAMFM achieves state-of-the-art performance across multiple antibody design tasks.

REFERENCES

Brennan Abanades, Wing Ki Wong, Fergus Boyles, Guy Georges, Alexander Bujotzek, and Charlotte M Deane. Immunebuilder: Deep-learning models for predicting the structures of immune proteins. *Communications Biology*, 6(1):575, 2023.

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

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.

Rahmad Akbar, Philippe A Robert, Milena Pavlović, Jeliazko R Jeliazkov, Igor Snapkov, Andrei Slabodkin, Cédric R Weber, Lonneke Scheffer, Enkelejda Miho, Ingrid Hobæk Haff, et al. A compact vocabulary of paratope-epitope interactions enables predictability of antibody-antigen binding. *Cell Reports*, 34(11), 2021.

Rebecca F Alford, Andrew Leaver-Fay, Jeliazko R Jeliazkov, Matthew J O'Meara, Frank P DiMaio, Hahnbeom Park, Maxim V Shapovalov, P Douglas Renfrew, Vikram K Mulligan, Kalli Kappel, et al. The rosetta all-atom energy function for macromolecular modeling and design. *Journal of chemical theory and computation*, 13(6):3031–3048, 2017.

AML AI4Science Team ByteDance, Xinshi Chen, Yuxuan Zhang, Chan Lu, Wenzhi Ma, Jiaqi Guan, Chengyue Gong, Jincai Yang, Hanyu Zhang, Ke Zhang, et al. Protenix-advancing structure prediction through a comprehensive alphafold3 reproduction. *bioRxiv*, pp. 2025–01, 2025.

Juntao Deng, Miao Gu, Pengyan Zhang, Mingyu Dong, Tao Liu, Yabin Zhang, and Min Liu. Nanobody–antigen interaction prediction with ensemble deep learning and prompt-based protein language models. *Nature Machine Intelligence*, 6(12):1594–1604, 2024.

Frédéric A Dreyer, Daniel Cutting, Constantin Schneider, Henry Kenlay, and Charlotte M Deane. Inverse folding for antibody sequence design using deep learning. *arXiv preprint arXiv:2310.19513*, 2023.

James Dunbar, Konrad Krawczyk, Jinwoo Leem, Terry Baker, Angelika Fuchs, Guy Georges, Jiye Shi, and Charlotte M Deane. Sabdab: the structural antibody database. *Nucleic acids research*, 42(D1):D1140–D1146, 2014.

Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. Prottrans: Toward understanding the language of life through self-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):7112–7127, 2021.

Stefan Ewert, Annemarie Honegger, and Andreas Plückthun. Stability improvement of antibodies for extracellular and intracellular applications: Cdr grafting to stable frameworks and structure-based framework engineering. *Methods*, 34(2):184–199, 2004.

Noelia Ferruz, Steffen Schmidt, and Birte Höcker. Protgpt2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):4348, 2022.

Kaiyuan Gao, Lijun Wu, Jinhua Zhu, Tianbo Peng, Yingce Xia, Liang He, Shufang Xie, Tao Qin, Haiguang Liu, Kun He, et al. Pre-training antibody language models for antigen-specific computational antibody design. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 506–517, 2023.

Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wiggert, et al. Simulating 500 million years of evolution with a language model. *Science*, pp. eads0018, 2025.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Fatima N Hitawala and Jeffrey J Gray. What has alphafold3 learned about antibody and nanobody docking, and what remains unsolved? *bioRxiv*, 2024.

Magnus Haraldson Høie, Alissa Hummer, Tobias H Olsen, Broncio Aguilar-Sanjuan, Morten Nielsen, and Charlotte M Deane. Antifold: Improved antibody structure-based design using inverse folding. *arXiv preprint arXiv:2405.03370*, 2024.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.

Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi S Jaakkola. Iterative refinement graph neural network for antibody sequence-structure co-design. In *International Conference on Learning Representations*, 2021.

Peter T Jones, Paul H Dear, Jefferson Foote, Michael S Neuberger, and Greg Winter. Replacing the complementarity-determining regions in a human antibody with those from a mouse. *Nature*, 321 (6069):522–525, 1986.

Xiangzhe Kong, Wenbing Huang, and Yang Liu. Conditional antibody design as 3d equivariant graph translation. In *The Eleventh International Conference on Learning Representations*, 2022.

Xiangzhe Kong, Wenbing Huang, and Yang Liu. End-to-end full-atom antibody design. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 17409–17429, 2023.

Gideon D Lapidoth, Dror Baran, Gabriele M Pszolla, Christoffer Norn, Assaf Alon, Michael D Tyka, and Sarel J Fleishman. Abdesign: A n algorithm for combinatorial backbone design guided by natural conformations and sequences. *Proteins: Structure, Function, and Bioinformatics*, 83 (8):1385–1406, 2015.

Haitao Lin, Lirong Wu, Huang Yufei, Yunfan Liu, Odin Zhang, Yuanqing Zhou, Rui Sun, and Stan Z Li. Geoab: Towards realistic antibody design and reliable affinity maturation. *bioRxiv*, pp. 2024–05, 2024.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.

Wei Lu, Jixian Zhang, Jiahua Rao, Zhongyue Zhang, and Shuangjia Zheng. Alphafold3, a secret sauce for predicting mutational effects on protein-protein interactions. *bioRxiv*, pp. 2024–05, 2024.

Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific antibody design and optimization with diffusion-based generative models for protein structures. *Advances in Neural Information Processing Systems*, 35:9754–9767, 2022.

Shitong Luo, Yufeng Su, Zuofan Wu, Chenpeng Su, Jian Peng, and Jianzhu Ma. Rotamer density estimator is an unsupervised learner of the effect of mutations on protein-protein interaction. *bioRxiv*, pp. 2023–02, 2023.

Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, et al. Abdiffuser: full-atom generation of in-vitro functioning antibodies. *Advances in Neural Information Processing Systems*, 36, 2024.

Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in neural information processing systems*, 34:29287–29303, 2021.

Erik Nijkamp, Jeffrey A Ruffolo, Eli N Weinstein, Nikhil Naik, and Ali Madani. Progen2: exploring the boundaries of protein language models. *Cell systems*, 14(11):968–978, 2023.

Tobias H Olsen, Iain H Moal, and Charlotte M Deane. Ablang: an antibody language model for completing antibody sequences. *Bioinformatics Advances*, 2(1):vbac046, 2022.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744, 2022.

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, 2023.

Roshan M Rao, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. Msa transformer. In *International Conference on Machine Learning*, pp. 8844–8856. PMLR, 2021.

Milong Ren, ZaiKai He, and Haicang Zhang. Multi-objective antibody design with constrained preference optimization. In *The Thirteenth International Conference on Learning Representations*, 2025.

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

Jeffrey A Ruffolo, Lee-Shin Chu, Sai Pooja Mahajan, and Jeffrey J Gray. Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies. *Nature communications*, 14(1):2389, 2023.

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 modeling the universe of crispr-cas sequences. *BioRxiv*, pp. 2024–04, 2024.

Danqing Wang, YE Fei, and Hao Zhou. On pre-training language model for antibody. In *The eleventh international conference on learning representations*, 2023.

Rubo Wang, Fandi Wu, Xingyu Gao, Jiaxiang Wu, Peilin Zhao, and Jianhua Yao. Iggm: A generative model for functional antibody and nanobody design. *bioRxiv*, pp. 2024–09, 2024.

Xinyou Wang, Zaixiang Zheng, YE Fei, Dongyu Xue, Shujian Huang, and Quanquan Gu. Dplm-2: A multimodal diffusion protein language model. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Zichen Wang, Yaokun Ji, Jianing Tian, and Shuangjia Zheng. Retrieval augmented diffusion model for structure-informed antibody design and optimization. *The Thirteenth International Conference on Learning Representations*, 2025b.

JunJie Wee and Guo-Wei Wei. Benchmarking alphafold3's protein-protein complex accuracy and machine learning prediction reliability for binding free energy changes upon mutation. *ArXiv*, pp. arXiv–2406, 2024.

Teng Xiao, Yige Yuan, Huaisheng Zhu, Mingxiao Li, and Vasant Honavar. Cal-dpo: Calibrated direct preference optimization for language model alignment. *Advances in Neural Information Processing Systems*, 37:114289–114320, 2024.

John L Xu and Mark M Davis. Diversity in the cdr3 region of vh is sufficient for most antibody specificities. *Immunity*, 13(1):37–45, 2000.

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. *arXiv preprint arXiv:2409.06744*, 2024.

Mingze Yin, Hanjing Zhou, Jialu Wu, Yiheng Zhu, Yuxuan Zhan, Zitai Kong, Hongxia Xu, Chang-Yu Hsieh, Jintai Chen, Tingjun Hou, et al. S^2alm: Sequence-structure pre-trained large language model for comprehensive antibody representation learning. *arXiv preprint arXiv:2411.15215*, 2024.

Zuobai Zhang, Minghao Xu, Arian Rokkum Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein representation learning by geometric structure pretraining. In *The Eleventh International Conference on Learning Representations*, 2023.

Zaixiang Zheng, Yifan Deng, Dongyu Xue, Yi Zhou, Fei Ye, and Quanquan Gu. Structure-informed language models are protein designers. In *International conference on machine learning*, pp. 42317–42338. PMLR, 2023.

Xiangxin Zhou, Dongyu Xue, Ruizhe Chen, Zaixiang Zheng, Liang Wang, and Quanquan Gu. Antigen-specific antibody design via direct energy-based preference optimization. *arXiv preprint arXiv:2403.16576*, 2024a.

Xiangxin Zhou, Dongyu Xue, Ruizhe Chen, Zaixiang Zheng, Liang Wang, and Quanquan Gu. Antigen-specific antibody design via direct energy-based preference optimization. In *ICML 2024 AI for Science Workshop*, 2024b.

Tian Zhu, Milong Ren, and Haicang Zhang. Antibody design using a score-based diffusion model guided by evolutionary, physical and geometric constraints. In *Forty-first International Conference on Machine Learning*, 2024.

A Experiment Details

A.1 Baseline Details

DiffAb (Luo et al., 2022) We used DiffAb from the official GitHub repository (https://github.com/luost26/diffab). To ensure a fair comparison, we retrained DiffAb on the same dataset as our model, using the configuration file `codesign_multicdrs.yml` for 6 CDRs design and `codesign_single.yml` for CDR-H3 design.

dyMEAN (Kong et al., 2023) We used dyMEAN from the official GitHub repository (https://github.com/THUNLP-MT/dyMEAN). For consistency, dyMEAN was retrained on our dataset using the configuration file `multi_cdr_design.json` for 6 CDRs design and `single_cdr_design.yml` for CDR-H3 design.

AbX (Zhou et al., 2024a) We used AbX from the official GitHub repository (https://github.com/CarbonMatrixLab/AbX). The provided checkpoints (https://zenodo.org/records/14577013) were used for evaluation.

GeoAb (Lin et al., 2024) We used GeoAb from the official GitHub repository (https://github.com/Edapinenut/GeoAB). To ensure a fair comparison, GeoAb was retrained on the same dataset as our model and only can perform CDR-H3 codesign.

A.2 Metric Details

Pll Pll is calculated by AntiBERTy (Ruffolo et al., 2021) to gauge sequence plausibility. For CDR-H3, we modify the calculation of the Pll to apply only within the CDR-H3 region same as previous work(Ye et al., 2024) for fine-grained comparison. AntiBERTy can be accessed at https://github.com/jeffreyruffolo/AntiBERTy.

AF3-related scores We adopted Protenix (ByteDance et al., 2025) for evaluation due to its fully open-source implementation and its ability to extract evolutionary features using ESM2-3B

embeddings, thereby avoiding the time-consuming MSA search. Specifically, we used version 0.4.0 from the `constraint_esm` branch of the GitHub repository (https://github.com/bytedance/Protenix/tree/constraint_esm) to evaluate the sequences. The evaluation metrics include **pTM**, **ipTM**, and the **AF3Score (ranking score)**.

Proportion of Hydrophobic Residues (PHR) Following privious works (Zhou et al., 2024b; Ye et al., 2024), we use PHR to assess the antibodies specificity designed by the models. This metric is calculted only at CDR-H3 region. PHR is defined as the proportion of hydrophobic amino acids—specifically Alanine (A), Valine (V), Isoleucine (I), Leucine (L), Methionine (M), Phenylalanine (F), Tryptophan (W), and Tyrosine (Y). The lower PHR means better specifically.

A.3 Dataset Preparation

For the synthetic dataset, we used a subset of the OAS database, which includes approximately 1.4 million paired antibody sequences predicted by IgFold (https://data.graylab.jhu.edu/OAS_paired.tar.gz, https://data.graylab.jhu.edu/Jaffe2022.tar.gz) and ABodyBuilder (https://zenodo.org/records/7258553). We added Gaussian noise with a standard deviation of 0.1 Å to the atomic coordinates of all predicted protein structures. For the DPO dataset, we used AAMFM-SFT to generate sequences based on all PDB structures in the RAbD dataset. We sampled sequences across different temperatures ranging from 0.1 to 1.5, generating 30 CDR-region sequences per complex at each temperature. Each generated sequence was then evaluated using both AntiBERTy and Protenix to obtain scoring metrics for downstream finetune.

A.4 Implementation Details

Our model was developed and implemented using the PyTorch framework. All experiments are run on a single A800 GPU, with a memory storage of 80GB. We use ESM-3 (version 3.1.6) as the base model, specifically the open-source 1.4B parameter version available at https://github.com/evolutionaryscale/esm.

In the first training stage (on the OAS subset), we used the AdamW optimizer with a learning rate of 5e-5 and a batch size of 8. Mixed-precision training with bfloat16 (bf16) was employed to accelerate training and reduce memory consumption. We trained the model for 2 epochs.

In the second stage (Supervised Fine-Tuning, SFT), both the antigen and antibody sequences were fed into the model. During this phase, we trained additional adapter layers. To prevent overfitting, we froze the base model and trained only the adapters for the first 40 epochs using the AdamW optimizer with a learning rate of 1e-4 and a warm-up over 10% of the training steps. After 40 epochs, we unfroze the base model and continued training for an additional 5 epochs with a learning rate of 1e-5 using AdamW (the adapter's learning rate is still 1e-4). Training was also conducted with bf16 mixed precision. Due to the increased input sequence length in this stage, we employed gradient accumulation with an effective batch size of 8 (batch size of 1 with accumulation steps set to 8).

In the third stage (Cal-DPO), we used the AdamW optimizer with a learning rate of 5e-5 and trained the model for 4 epochs, using bf16 mixed precision. Due to the increased input sequence length in this stage, we employed gradient accumulation with an effective batch size of 8 (batch size of 1 with accumulation steps set to 8).

B Detailed Reults

B.1 Analysis

We compared the training process of DPO and Cal-DPO algorithms, as illustrated in Figure 2. While both methods exhibit a increasing margin (the difference between the log probabilities of preferred and dispreferred samples), we observed a key difference in their behavior. During standard DPO training, the log probabilities of both preferred and dispreferred sequences decrease simultaneously, indicating that the model may not fully capture the preference signal. In contrast, Cal-DPO maintains a relatively stable log probability for preferred sequences while decreasing the log probability

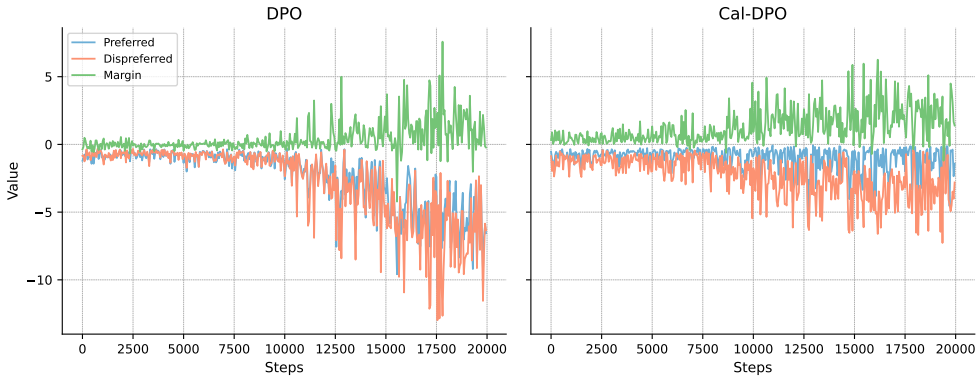

Figure 2: Comparative visualization of the training dynamics of DPO and Cal-DPO

Table 4: Ablation on CDRs design

Method	Pll ↑	AF3Score ↑	ipTM ↑	RMSD ↓	AAR ↑	PHR ↓
ESM3-open	-1.40	0.846	0.839	3.31	32.19	**0.400**
w/o SFT	-0.89	0.876	0.869	**2.49**	39.59	0.445
w/o adapter	-0.94	0.868	0.862	2.57	39.22	0.443
w/o DPO (SFT)	-0.91	0.870	0.865	2.54	**40.32**	0.465
AAMFM-CalDPO	**-0.87**	**0.881**	**0.877**	2.62	36.23	0.443

of dispreferred ones. This behavior suggests that Cal-DPO more effectively enhances the model's ability to favor preferred samples during generation.

B.2 Ablation

Setup We identified three key algorithmic components in our training pipeline: adapter training for incorporating antigen and epitope information, supervised fine-tuning (SFT), and Cal-DPO. To evaluate the effectiveness of each component, we conducted an ablation study starting from the ESM-3B base model. Specifically, we examined the following variants: (1) pretraining on the OAS dataset without subsequent SFT (**w/o SFT**), (2) performing SFT on SAbDab data without using the antigen adapter (**w/o adapter**), and (3) omitting the Cal-DPO stage (**w/o DPO**). Each variant was independently evaluated to assess the contribution of the corresponding component to overall model performance.

Results As shown in Table 4, we conduct an ablation study to demonstrate the effectiveness of each training stage and module by respectively removing the SFT stage, the antigen adapter, and the DPO training. Following the experimental setup described in Section 4.1, we design all six CDR regions along with the full antibody structure. We report the Pll of the full antibody sequence, AF3Score, ipTM, as well as the RMSD and AAR of the CDR-H3 region.

Experimental results show that models trained with SFT generate antibody sequences that more closely resemble the natural distribution. Furthermore, models aligned through preference aligning (Cal-DPO) generate antibody sequences that are more plausible and functional, indicating greater potential for real-world applications.

C Related Work

Protein Multimodal Language Models Protein Language Models (PLMs) have made significant strides in recent years by learning from large-scale protein sequence data to effectively capture protein interactions, structures, and functions (Meier et al., 2021; Rao et al., 2021; Elnaggar et al., 2021; Ferruz et al., 2022; Lin et al., 2023; Nijkamp et al., 2023; Ruffolo et al., 2024). Recently, the

research trend has shifted from unimodal to multimodal approaches, enhancing model capabilities by incorporating additional structural or functional information. LMDesign (Zheng et al., 2023) integrates geometric information into language models using an adapter, demonstrating superior performance compared to sequence-only models. ESM3 (Hayes et al., 2025), a representative large-scale multimodal model, is jointly pre-trained on vast amounts of protein sequence, structure, and function data, enabling the generation of protein data across different modalities. DPLM-2 (Wang et al., 2025a) extends discrete diffusion protein language model to the multimodal domain, focusing on learning the joint distribution of sequences and structures. In specific applications, S^2ALM(Yin et al., 2024) is an antibody-specific language model combining sequence and structure; however, it can't design antibody binding to specific antigen. In contrast to these methods, our approach utilizes the powerful base model and fine-tunes it using multi-level finetune strategy. This strategy aims not only to achieve accurate antibody representation but also to effectively perform antigen-specific functional antibody design.

Computational Antibody Design Early antibody design methods were often limited to energy-based approaches (Lapidoth et al., 2015; Adolf-Bryfogle et al., 2018). Recently, deep learning approaches mainly follows two directions: antibody sequence design and sequence-structure co-design. The methods used for antibody sequence design include language models (Ruffolo et al., 2021; Olsen et al., 2022; Wang et al., 2023; Gao et al., 2023) and inverse folding models (Dreyer et al., 2023; Høie et al., 2024; Wang et al., 2025b). Antibody sequence-structure co-design methods mainly taking antibody-antigen complex as a graph and using graph networks, or using diffusion-based model co-design structure and sequence of antibody CDR (Jin et al., 2021; Kong et al., 2022; 2023; Lin et al., 2024; Luo et al., 2022; Zhu et al., 2024; Martinkus et al., 2024; Zhou et al., 2024a; Wang et al., 2024; Ren et al., 2025). Albeit powerful, existing co-design models are often constrained by the limited availability of training data, which makes it challenging for them to fully learn and capture the deep, complex features in protein. On the other hand, while language models generally possess strong representational capabilities, they typically exhibit weaker performance in generating functional antibodies specifically targeted specific to a given antigen. To overcome these limitations, we propose a novel multimodal antibody language model.

D Method Details

D.1 Antigen-aware Model Architecture

We use ESM3-1.4B as our base model (Hayes et al., 2025). ESM3 is a bidirectional transformer architecture that integrates multimodal information from protein sequence, structure, and function. These modalities are first embedded and fused at the input layer, then processed through a stack of transformer blocks. At the output, shallow multi-layer perceptron (MLP) heads project the final hidden states into token probabilities corresponding to each modality. Instead of using modality-specific architectural components, ESM3 adopts a unified tokenization approach to represent the complexity of proteins in a shared multimodal feature space. This design enables efficient and scalable training. In particular, protein structures are tokenized using a discrete autoencoder that compresses 3D spatial information into discrete structural tokens.

To incorporate antigen geometry and antigen-antibody interface features for antigen-specific antibody generation, we introduce the Antigen Geometric-Epitope-aware Adapter. Antigen 3D coordinates are first passed through a GNN (GearNet (Zhang et al., 2023)) to obtain node embeddings, which are used as geometric features. These embeddings are projected into a unified feature space through a combination of MLPs, convolutional layer, and pooling layer. Meanwhile, epitope information is encoded as a binary matrix aligned with the antigen sequence (1 for epitope residues, 0 for other region), which is then passed through a learnable embedding layer that maps the binary values to an 8-dimensional space. These embeddings are combined with positional encodings and input into a Transformer encoder to obtain contextualized epitope representations. The resulting geometric and epitope features are concatenated with the last-layer output of ESM3's base transformer (after projecting it from 1536 to 64 dimensions). The concatenated features are then mapped back to the original 1536-dimensional space and added to the base model's final hidden states.

D.2 Direct Preference Optimization

Given a preference pair (y_w, y_l) under input context c (e.g., masked antibody sequence s_{fr}, antigen features G_{ag}, and interface information I), the Bradley–Terry (BT) model assumes a latent reward function $r(y|c)$ and models the preference as:

$$P(y_w \succ y_l \mid c) = \frac{\exp(r(y_w \mid c))}{\exp(r(y_w \mid c)) + \exp(r(y_l \mid c))}. \tag{4}$$

In the standard reinforcement learning from human feedback (RLHF) framework (Ouyang et al., 2022), a reward model r_ϕ is trained by maximizing the expected reward under the optimized policy π_θ, while regularizing it towards the reference policy π_{ref} with a KL-divergence:

$$\max_{\pi_\theta} \mathbb{E}_{y\sim\pi_\theta}\left[r_\phi(y \mid c)\right] - \beta D_{\text{KL}}(\pi_\theta(y \mid c) \| \pi_{\text{ref}}(y \mid c)). \tag{5}$$

The optimal solution for the reward model has a closed form:

$$\pi^*_\theta(y \mid c) = \frac{\pi_{\text{ref}}(y \mid c) \cdot \exp(r^*(y \mid c)/\beta)}{Z} \tag{6}$$

where Z is a normalization constant. This leads to the optimal reward model:

$$r^*(y \mid c) = \beta \log \frac{\pi_\theta(y \mid c)}{\pi_{\text{ref}}(y \mid c)} + \beta \log Z. \tag{7}$$

Substituting Equation 7 into the Bradley-Terry model in Equation 4 and optimizing the log-likelihood, the Direct Preference Optimization (DPO) objective (Rafailov et al., 2023) becomes:

$$\mathcal{L}_{\text{DPO}}(\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(c,y_w,y_l)\sim\mathcal{D}}\left[\log\sigma\left(\beta \log \frac{\pi_\theta(y_w \mid c)}{\pi_{\text{ref}}(y_w \mid c)}\right) - \log\sigma\left(\beta \log \frac{\pi_\theta(y_l \mid c)}{\pi_{\text{ref}}(y_l \mid c)}\right)\right], \tag{8}$$

where $\sigma(\cdot)$ is the sigmoid function and β is a temperature parameter controlling preference sharpness.

Cal-DPO introduces an additional calibration loss term. Defining the target margin $M = 1/(2\beta)$, the calibration loss encourages the log-likelihood ratios to align with this margin:

$$\mathcal{L}_{\text{Cal}}(\theta; \pi_{\text{ref}}) = \mathbb{E}_{(c,y_w,y_l)\sim\mathcal{D}}\left[\left(\log\sigma(\beta \log \frac{\pi_\theta(y_w \mid c)}{\pi_{\text{ref}}(y_w \mid c)} - M\right)^2 + \left(\log\sigma\left(\beta \log \frac{\pi_\theta(y_l \mid c)}{\pi_{\text{ref}}(y_l \mid c)} + M\right)^2\right] \tag{9}$$

The final training loss is a weighted sum of the DPO and calibration terms:

$$\mathcal{L}_{\text{Cal-DPO}} = \mathcal{L}_{\text{DPO}} + \lambda \cdot \mathcal{L}_{\text{Cal}} \tag{10}$$

E Broader Impacts

Our work has broad potential applications in biotechnology and therapeutic development. For instance, the ability to design functional antibodies that bind to specific antigens could contribute to the development of vaccines, or the engineering of agonists and antagonists that modulate specific immune pathways or signaling proteins. Such applications have the potential to advance precision medicine, immunotherapy, and the treatment of various diseases.

We explicitly prohibit the use of our model for potentially harmful purposes, such as the development of biological weapons, targeting of human proteins for malicious intent, or any application that may pose significant ethical, safety, or biosecurity risks.

F LIMITATIONS

The main limitation of this work is that the designed antibodies have not yet been fully validated through in vitro experiments while we are actively collaborating with experimental research groups to pursue this. Another key future direction is to generalize the AAMFM framework beyond antibodies, extending it to the design of a broader class of protein complexes.

G CODES

Our code is avilable at https://anonymous.4open.science/r/AAMFM.