

---

# Antigen-Specific Antibody Design via Direct Energy-based Preference Optimization

---

Xiangxin Zhou<sup>1,2,4,\*</sup>

Dongyu Xue<sup>4,\*</sup>

Ruizhe Chen<sup>3,4,\*</sup>

Zaixiang Zheng<sup>4</sup>

Liang Wang<sup>1,2</sup>

Quanquan Gu<sup>4,†</sup>

<sup>1</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>2</sup>New Laboratory of Pattern Recognition (NLPR),

State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS),  
Institute of Automation, Chinese Academy of Sciences (CASIA)

<sup>3</sup>College of Computer Science and Electronic Engineering, Hunan University

<sup>4</sup>ByteDance Research

## Abstract

Antibody design, a crucial task with significant implications across various disciplines such as therapeutics and biology, presents considerable challenges due to its intricate nature. In this paper, we tackle antigen-specific antibody sequence-structure co-design as an optimization problem towards specific preferences, considering both rationality and functionality. Leveraging a pre-trained conditional diffusion model that jointly models sequences and structures of antibodies with equivariant neural networks, we propose *direct energy-based preference optimization* to guide the generation of antibodies with both rational structures and considerable binding affinities to given antigens. Our method involves fine-tuning the pre-trained diffusion model using a residue-level decomposed energy preference. Additionally, we employ gradient surgery to address conflicts between various types of energy, such as attraction and repulsion. Experiments on RAbD benchmark show that our approach effectively optimizes the energy of generated antibodies and achieves state-of-the-art performance in designing high-quality antibodies with low total energy and high binding affinity simultaneously, demonstrating the superiority of our approach.

## 1 Introduction

Antibodies, vital proteins with an inherent Y-shaped structure in the immune system, are produced in response to an immunological challenge. Their primary function is to discern and neutralize specific pathogens, typically referred to as antigens, with a significant degree of specificity [39]. The specificity mainly comes from the Complementarity Determining Regions (CDRs), which accounts for most binding affinity to specific antigens [24, 15, 49, 2]. Hence, the design of CDRs is a crucial step in developing potent therapeutic antibodies, which plays an important role in drug discovery.

Traditional *in silico* antibody design methods rely on sampling or searching protein sequences over a large search space to optimize the physical and chemical energy, which is inefficient and easily trapped in bad local minima [1, 31, 47]. Recently, deep generative models have been employed to model protein sequences in nature for antibody design [5, 17]. Following the fundamental biological principle

---

\*Equal contribution (this work was done during Xiangxin and Ruizhe’s internship at ByteDance Research).

†Correspondence to: Quanquan Gu <quanquan.gu@bytedance.com>.

that structure determines function numerous efforts have been focused on antibody sequence-structure co-design [22, 21, 36, 29, 30, 37], which demonstrate superiority over sequence design-based methods. However, the main evaluation metrics in the aforementioned works are amino acid recovery (AAR) and root mean square deviation (RMSD) between the generated antibody and the real one. This is controversial because AAR is susceptible to manipulation and does not precisely gauge the quality of the generated antibody sequence. Meanwhile, RMSD does not involve side chains, which are vital for antigen-antibody interaction. Besides, it is biologically plausible that a specific antigen can potentially bind with multiple efficacious antibodies [45, 12]. This motivates us to examine the generated structures and sequences of antibodies through the lens of energy, which reflects the rationality of the designed antibodies and their binding affinity to the target antigens. We have noted that nearly all antibody sequence-structure co-design methods struggle to produce antibodies with low energy. This suggests the presence of irrational structures and inadequate binding affinity in antibodies designed by these methods (see Fig. 1). We attribute this incapability to the insufficient model training caused by a scarcity of high-quality data.




Figure 1: The third CDR in the heavy chain, CDR-H3 (colored in yellow), of real antibody (left) and synthetic antibody (right) designed by MEAN [29] for a given antigen (PDB ID: 4cmh). The rest parts of antibodies except CDR-H3 are colored in blue. The antigens are colored in gray. We use red (resp. black) dotted lines to represent clashes between a CDR-H3 atom and a framework/antigen atom (resp. another CDR-H3 atom). We consider a clash occurs when the overlap of the van der Waals radii of two atoms exceeds  $0.6\text{\AA}$ .

To tackle the above challenges and bridge the gap between *in silico* antibody sequence-structure co-design methods and the intrinsic need for drug discovery, we formulate the antibody design task as an antibody optimization problem with a focus on better rationality and functionality. Inspired by direct preference optimization [DPO, 41] and self-play fine-tuning techniques [10] that achieve huge success in the alignment of large language models (LLMs), we proposed a direct energy-based preference optimization method named ABDPO for antibody optimization. More specifically, we first pre-train a conditional diffusion model on real antigen-antibody datasets, which simultaneously captures sequences and structures of complementarity-determining regions (CDR) in antibodies with equivariant neural networks. We then progressively fine-tune this model using synthetic antibodies generated by the model itself given an antigen with energy-based preference. This preference is defined at a fine-grained residue level, which promotes the effectiveness and efficiency of the optimization process. To fulfill the requirement of various optimization objectives, we decompose the energy into multiple types so that we can incorporate prior knowledge and mitigate the interference between conflicting objectives (e.g., repulsion and attraction energy) to guide the optimization process. Fine-tuning with self-synthesized energy-based antibody preference data represents a revolutionary solution to address the limitation of scarce high-quality real-world data, a significant challenge in this domain. We highlight our main contributions as follows:

- We tackle the antibody sequence-structure co-design problem through the lens of energy from the perspectives of both rationality and functionality.
- We propose direct residue-level energy-based preference optimization to fine-tune diffusion models for designing antibodies with rational structures and high binding affinity to specific antigens.
- We introduce energy decomposition and conflict mitigation techniques to enhance the effectiveness and efficiency of the optimization process.
- Experiments show ABDPO’s effectiveness in generating antibodies with energies resembling natural antibodies and generality in optimizing multiple preferences.

## 2 Related Work

**Antibody Design.** The application of deep learning to antibody design can be traced back to at least [35, 43, 3]. In recent years, sequence-structure co-design of antibodies has attracted increasing attention. Jin et al. [22] proposed to simultaneously design sequences and structures of CDRs in an autoregressive way and iteratively refine the designed structures. Jin et al. [21] further utilized the epitope and focused on designing CDR-H3 with a hierarchical message passing equivariant

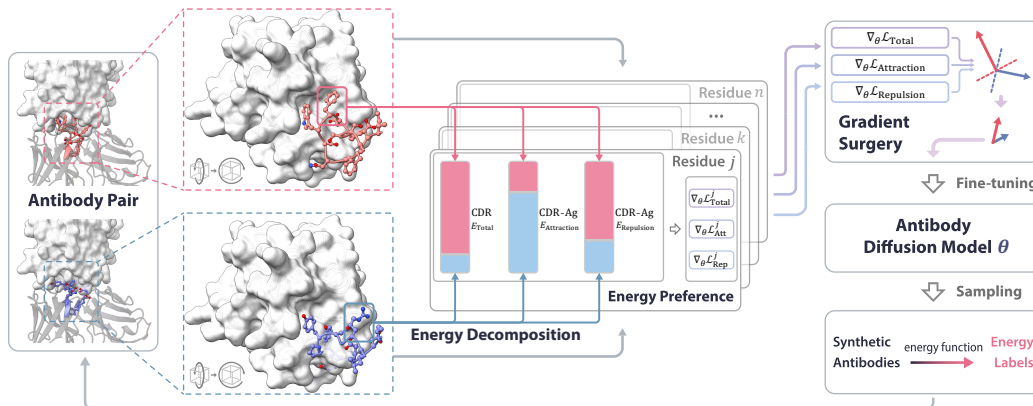


Figure 2: Overview of ABDPO. This process can be summarized as: (a) Generate antibodies with the pre-trained diffusion model; (b) Evaluate the multiple types of residue-level energy and construct preference data; (c) Compute the losses for energy-based preference optimization and mitigate the conflicts between losses of multiple types of energy; (d) Update the diffusion model.

network. Kong et al. [29] incorporated antigens and the light chains of antibodies as conditions and designed CDRs with E(3)-equivariant graph networks via a progressive full-shot scheme. Luo et al. [36] proposed a diffusion model that takes residue types, atom coordinates and side-chain orientations into consideration to generate antigen-specific CDRs. Kong et al. [30] focused on epitope-binding CDR-H3 design and modelled full-atom geometry. Recently, Martinkus et al. [37] proposed AbDiffuser, a novel diffusion model for antibody design, that incorporates more domain knowledge and physics-based constraints and also enables side-chain generation. Besides, Wu and Li [48], Gao et al. [19] and Zheng et al. [52] introduced pre-trained protein language model to antibody design. Distinct from the above works, our method places a stronger emphasis on designing and optimizing antibodies with low energy and high binding affinity.

**Alignment of Generative Models.** Solely maximizing the likelihood of training data does not always lead to a model that satisfies users’ preferences. Recently, many efforts have been made on the alignment of the generative models to human preferences. Reinforcement learning has been introduced to learning from human/AI feedback to large language models, such as RLHF [40] and RLAI [33]. Typically, RLHF consists of three phases: supervised fine-tuning, reward modeling, and RL fine-tuning. Similar ideas have also been introduced to text-to-image generation, such as DDPO [7], DPOK [16] and DiffAC [53]. They view the generative processes of diffusion models as a multi-step Markov Decision Process (MDP) and apply policy gradient for fine-tuning. Rafailov et al. [41] proposed direct preference optimization (DPO) to directly fine-tune language models on preference data, which matches RLHF in performance. Recently, DPO has been introduced to text-to-image generation [46, 6]. Notably, in the aforementioned works, models pre-trained with large-scale datasets have already shown strong performance, in which case alignment further increases users’ satisfaction. In contrast, in our work, the model pre-trained with limited real-world antibody data is insufficient in performance. Therefore, preference optimization in our case is primarily used to help the model understand the essence of nature and meet the requirement of antibody design.

### 3 Method

In this section, we present ABDPO, a direct energy-based preference optimization method for designing antibodies with reasonable rationality and functionality (Fig. 2). We first define the antibody generation task and introduce the diffusion model for this task in Sec. 3.1. Then we introduce residue-level preference optimization for fine-tuning the diffusion model and analyze its advantages in effectiveness and efficiency in Sec. 3.2. Finally, in Sec. 3.3, we introduce the energy decomposition and describe how to mitigate the conflicts when optimizing multiple types of energy.

#### 3.1 Preliminaries

We focus on designing CDR-H3 of the antibody given antigen structure as CDR-H3 contributes the most to the diversity and specificity of antibodies [49, 2] and the rest part of the antibody including the frameworks and other CDRs. Following Luo et al. [36], each amino acid is represented by its type  $s_i \in \{\text{ACDEFGHIKLMNPQRSTVWY}\}$ ,  $C_\alpha$  coordinate  $\mathbf{x}_i \in \mathbb{R}^3$ , and frame orientation

$\mathbf{O}_i \in \text{SO}(3)$  [28], where  $i = 1, \dots, N$  and  $N$  is the number of the amino acids in the protein complex. We assume the CDR-H3 to be generated has  $m$  amino acids, which can be denoted by  $\mathcal{R} = \{(s_j, \mathbf{x}_j, \mathbf{O}_j) | j = n+1, \dots, n+m\}$ , where  $n+1$  is the index of the first residue in CDR-H3 sequence. The rest part of the antigen-antibody complex can be denoted by  $\mathcal{P} = \{(s_i, \mathbf{x}_i, \mathbf{O}_i) | i \in \{1, \dots, N\} \setminus \{n+1, \dots, n+m\}\}$ . The antibody generation task can be then formulated as modeling the conditional distribution  $P(\mathcal{R}|\mathcal{P})$ .

Denosing Diffusion Probabilistic Model [DDPM, 20] have been introduced to antibody generation by Luo et al. [36]. This approach consists of a forward *diffusion* process and a reverse *generative* process. The diffusion process gradually injects noises into data as follows:

$$\begin{aligned} q(s_j^t | s_j^0) &= \mathcal{C}(\mathbb{1}(s_j^t) | \bar{\alpha}^t \mathbb{1}(s_j^0) + \bar{\beta}^t \mathbb{1}/K), \\ q(\mathbf{x}_j^t | \mathbf{x}_j^0) &= \mathcal{N}(\mathbf{x}_j^t | \sqrt{\bar{\alpha}^t} \mathbf{x}_j^0, \bar{\beta}^t \mathbf{I}), \\ q(\mathbf{O}_j^t | \mathbf{O}_j^0) &= \mathcal{IG}_{\text{SO}(3)}(\mathbf{O}_j^t | \text{ScaleRot}(\sqrt{\bar{\alpha}^t} \mathbf{O}_j^0), \bar{\beta}^t), \end{aligned}$$

where  $(s_j^0, \mathbf{x}_j^0, \mathbf{O}_j^0)$  are the noisy-free amino acid at time step 0 with index  $j$ , and  $(s_j^t, \mathbf{x}_j^t, \mathbf{O}_j^t)$  are the noisy amino acid at time step  $t$ .  $\mathbb{1}(\cdot)$  is the one-hot operation.  $\{\beta^t\}_{t=1}^T$  is the noise schedule for the diffusion process [20], and we define  $\bar{\alpha}^t = \prod_{\tau=1}^t (1 - \beta^\tau)$  and  $\bar{\beta}^t = 1 - \bar{\alpha}^t$ .  $K$  is the number of amino acid types. Here,  $\mathcal{C}(\cdot)$ ,  $\mathcal{N}(\cdot)$ , and  $\mathcal{IG}_{\text{SO}(3)}(\cdot)$  are categorical distribution, Gaussian distribution on  $\mathbb{R}^3$ , and isotropic Gaussian distribution on  $\text{SO}(3)$  [32] respectively. `ScaleRot` scales the rotation angle with fixed rotation axis to modify the rotation matrix [18].

Correspondingly, the reverse generative process learns to recover data by iterative denosing. The denosing process  $p(\mathcal{R}^{t-1} | \mathcal{R}^t, \mathcal{P})$  from time step  $t$  to time step  $t-1$  is defined as follows:

$$p(s_j^{t-1} | \mathcal{R}^t, \mathcal{P}) = \mathcal{C}(s_j^{t-1} | \mathbf{f}_{\theta_1}(\mathcal{R}^t, \mathcal{P})[j]), \quad (1)$$

$$p(\mathbf{x}_j^{t-1} | \mathcal{R}^t, \mathcal{P}) = \mathcal{N}(\mathbf{x}_j^{t-1} | \mathbf{f}_{\theta_2}(\mathcal{R}^t, \mathcal{P})[j], \beta^t \mathbf{I}), \quad (2)$$

$$p(\mathbf{O}_j^{t-1} | \mathcal{R}^t, \mathcal{P}) = \mathcal{IG}_{\text{SO}(3)}(\mathbf{f}_{\theta_3}(\mathcal{R}^t, \mathcal{P})[j], \beta^t), \quad (3)$$

where  $\mathcal{R}^t = \{s_j, \mathbf{x}_j, \mathbf{O}_j\}_{j=n+1}^{n+m}$  is the noisy sequence and structure of CDR-H3 at time step  $t$ ,  $\mathbf{f}_{\theta_1}, \mathbf{f}_{\theta_2}, \mathbf{f}_{\theta_3}$  are parameterized by denosing neural networks. We utilize SE(3)-equivariant networks [23, 25] as denosing networks because proteins are structures in three-dimensional space, and their properties and characteristics should remain invariant regardless of their observation views.  $\mathbf{f}(\cdot)[j]$  denotes the output that corresponds to the  $j$ -th amino acid. The training objective of the reverse generative process is to minimize the Kullback–Leibler (KL) divergence between the variational distribution  $p$  and the posterior distribution  $q$  as follows:

$$L = \mathbb{E}_{\mathcal{R}^t \sim q} \left[ \frac{1}{m} \sum_{j=n+1}^{n+m} \mathbb{D}_{\text{KL}} \left( q(\mathcal{R}^{t-1}[j] | \mathcal{R}^t, \mathcal{R}^0, \mathcal{P}) \| p_{\theta}(\mathcal{R}^{t-1}[j] | \mathcal{R}^t, \mathcal{P}) \right) \right]. \quad (4)$$

With some algebra, we can simplify the above objective and derive the reconstruction loss at time step  $t$  as follows:

$$L_s^t = \mathbb{E}_{\mathcal{R}^t} \left[ \frac{1}{m} \sum_{j=n+1}^{n+m} \mathbb{D}_{\text{KL}} \left( q(s_j^{t-1} | s_j^t, s_j^0) \| p(s_j^{t-1} | \mathcal{R}^t, \mathcal{P}) \right) \right], \quad (5)$$

$$L_x^t = \mathbb{E}_{\mathcal{R}^t} \left[ \frac{1}{m} \sum_{j=n+1}^{n+m} \|\mathbf{x}_j^0 - \mathbf{f}_{\theta_2}(\mathcal{R}^t, \mathcal{P})\|^2 \right], \quad (6)$$

$$L_O^t = \mathbb{E}_{\mathcal{R}^t} \left[ \frac{1}{m} \sum_{j=n+1}^{n+m} \|(\mathbf{O}_j^0)^\top \mathbf{f}_{\theta_3}(\mathcal{R}^t, \mathcal{P})[j] - \mathbf{I}\|_F^2 \right], \quad (7)$$

where  $\mathcal{R}^t \sim q(\mathcal{R}^t | \mathcal{R}^0)$  and  $\mathcal{R}^0 \sim P(\mathcal{R} | \mathcal{P})$ , and  $\|\cdot\|_F$  is the matrix Frobenius norm. Note that as Luo et al. [36] mentioned, Eqs. (1) and (3) are an empirical perturbation-denosing process instead of a rigorous one. Thus the terminology *KL-divergence* may not be proper for orientation  $\mathbf{O}$ . Nevertheless, we can still approximately derive an empirical reconstruction loss for orientation  $\mathbf{O}$  as above that works in practice. The overall loss is  $L \approx \mathbb{E}_{t \sim \mathcal{U}[1, T]} [L_s^t + L_x^t + L_O^t]$ . After optimizing this loss, we can start with the noises from the prior distribution and then apply the reverse process to generate antibodies.

### 3.2 Direct Energy-based Preference Optimization

Only the antibodies with considerable sequence-structure rationality and binding affinity can be used as effective therapeutic candidates. Fortunately, these two properties can be estimated by biophysical energy. Thus, we introduce direct energy-based preference optimization to fine-tune the pre-trained diffusion models for antibody design.

Inspired by RLHF [40], we can fine-tune the pre-trained model to maximize the reward as:

$$\max_{\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_{\theta}} [r(\mathcal{R}^0)] - \beta \mathbb{D}_{\text{KL}}(p_{\theta}(\mathcal{R}^0) \| p_{\text{ref}}(\mathcal{R}^0)),$$

where  $p_{\theta}$  (resp.  $p_{\text{ref}}$ ) is the distribution induced by the model being fine-tuned (resp. the fixed pre-trained model),  $\beta$  is a hyperparameter that controls the KL divergence regularization, and  $r(\cdot)$  is the reward function. The optimal solution to the above objective takes the form:

$$p_{\theta^*}(\mathcal{R}^0) = \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right).$$

Following Rafailov et al. [41], we turn to the DPO objective as follows:

$$L_{\text{DPO}} = -\mathbb{E}_{\mathcal{R}_1^0, \mathcal{R}_2^0} \left[ \log \sigma \left( \beta \text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) \left[ \log \frac{p_{\theta}(\mathcal{R}_1^0)}{p_{\text{ref}}(\mathcal{R}_1^0)} - \log \frac{p_{\theta}(\mathcal{R}_2^0)}{p_{\text{ref}}(\mathcal{R}_2^0)} \right] \right) \right],$$

where  $\sigma(\cdot)$  is sigmoid and  $\text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0)$  indicate the preference over  $\mathcal{R}_1^0$  and  $\mathcal{R}_2^0$ . We use “ $\succ$ ” to denote the preference. Specifically,  $\text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) = 1$  (resp.  $-1$ ) if  $\mathcal{R}_1^0 \succ \mathcal{R}_2^0$  (resp.  $\mathcal{R}_2^0 \prec \mathcal{R}_1^0$ ) in which case we call  $\mathcal{R}_1^0$  (resp.  $\mathcal{R}_2^0$ ) the “winning” sample and  $\mathcal{R}_2^0$  (resp.  $\mathcal{R}_1^0$ ) the “losing” sample, and  $\text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) = 0$  if they tie.  $\mathcal{R}_1^0$  and  $\mathcal{R}_2^0$  are a pair of data sampled from the Bradley-Terry [BT, 8] model with reward  $r(\cdot)$ , i.e.,  $p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = \sigma(r(\mathcal{R}_1^0) - r(\mathcal{R}_2^0))$ . Please refer to Appendix C for more detailed derivations.

Due to the intractable  $p_{\theta}(\mathcal{R}^0)$ , following Wallace et al. [46], we introduce latent variables  $\mathcal{R}^{1:T}$  and utilize the evidence lower bound optimization (ELBO). In particular,  $L_{\text{DPO}}$  can be modified as follows:

$$L_{\text{DPO-Diffusion}} = -\mathbb{E}_{\mathcal{R}_1^0, \mathcal{R}_2^0} \left[ \log \sigma \left( \beta \mathbb{E}_{\mathcal{R}_1^{1:T}, \mathcal{R}_2^{1:T}} \left[ \text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) \left( \log \frac{p_{\theta}(\mathcal{R}_1^{0:T})}{p_{\text{ref}}(\mathcal{R}_1^{0:T})} - \log \frac{p_{\theta}(\mathcal{R}_2^{0:T})}{p_{\text{ref}}(\mathcal{R}_2^{0:T})} \right) \right] \right) \right],$$

where  $\mathcal{R}_1^{1:T} \sim p_{\theta}(\mathcal{R}_1^{1:T} | \mathcal{R}_1^0)$  and  $\mathcal{R}_2^{1:T} \sim p_{\theta}(\mathcal{R}_2^{1:T} | \mathcal{R}_2^0)$ .

Following Wallace et al. [46], we can utilize Jensen’s inequality and convexity of function  $-\log \sigma$  to derive the following upper bound of  $L_{\text{DPO-Diffusion}}$ :

$$\tilde{L}_{\text{DPO-Diffusion}} = -\mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \log \sigma \left( \beta T \text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) \left[ \log \frac{p_{\theta}(\mathcal{R}_1^{t-1} | \mathcal{R}_1^t)}{p_{\text{ref}}(\mathcal{R}_1^{t-1} | \mathcal{R}_1^t)} - \log \frac{p_{\theta}(\mathcal{R}_2^{t-1} | \mathcal{R}_2^t)}{p_{\text{ref}}(\mathcal{R}_2^{t-1} | \mathcal{R}_2^t)} \right] \right) \right],$$

where  $t \sim U(0, T)$ ,  $(\mathcal{R}_1^{t-1}, \mathcal{R}_1^t)$  and  $(\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)$  are sampled from reverse generative process of  $\mathcal{R}_1^0$  and  $\mathcal{R}_2^0$ , respectively, i.e.,  $(\mathcal{R}_1^{t-1}, \mathcal{R}_1^t) \sim p_{\theta}(\mathcal{R}_1^{t-1}, \mathcal{R}_1^t | \mathcal{R}_1^0)$  and  $(\mathcal{R}_2^{t-1}, \mathcal{R}_2^t) \sim p_{\theta}(\mathcal{R}_2^{t-1}, \mathcal{R}_2^t | \mathcal{R}_2^0)$ .

In our case, the antibodies with low energy are desired. Thus, we define the reward  $r(\cdot)$  as  $-\mathcal{E}(\cdot)/T$ , where  $\mathcal{E}(\cdot)$  is the energy function and  $T$  is the temperature. Different from the text-to-image generation where the (latent) reward is assigned to a complete image instead of a pixel [46], we know more fine-grained credit assignment. Specifically, it is known that  $\mathcal{E}(\mathcal{R}^0) = \sum_{j=n+1}^{n+m} \mathcal{E}(\mathcal{R}^0[j])$ , i.e., the energy of an antibody is the summation of the energy of its amino acids [4]. Thus the preference can be measured **at the residue level** instead of the entire CDR level. Besides, we have  $\log p_{\theta}(\mathcal{R}^{t-1} | \mathcal{R}^t) = \sum_{j=n+1}^{n+m} \log p_{\theta}(\mathcal{R}^{t-1}[j] | \mathcal{R}^t)$ , which is a common assumption of diffusion models. Thus we can derive a residue-level DPO-Diffusion loss:

$$L_{\text{residue-DPO-Diffusion}} = -\mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \log \sigma \left( \beta T \sum_{j=n+1}^{n+m} \text{sgn}(\mathcal{R}_1^0[j], \mathcal{R}_2^0[j]) \left[ \log \frac{p_{\theta}(\mathcal{R}_1^{t-1}[j] | \mathcal{R}_1^t)}{p_{\text{ref}}(\mathcal{R}_1^{t-1}[j] | \mathcal{R}_1^t)} - \log \frac{p_{\theta}(\mathcal{R}_2^{t-1}[j] | \mathcal{R}_2^t)}{p_{\text{ref}}(\mathcal{R}_2^{t-1}[j] | \mathcal{R}_2^t)} \right] \right) \right].$$

Thus, by Jensen’s inequality and the convexity of  $-\log \sigma$ , we can further derive  $\tilde{L}_{\text{residue-DPO-Diffusion}}$ , which is an upper bound of  $L_{\text{residue-DPO-Diffusion}}$ :

$$\begin{aligned} \tilde{L}_{\text{residue-DPO-Diffusion}} &= -\mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \right. \\ &\quad \left. \sum_{j=n+1}^{n+m} \log \sigma \left( \beta T \text{sgn}(\mathcal{R}_1^0[j], \mathcal{R}_2^0[j]) \left[ \log \frac{p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t)}{p_{\text{ref}}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t)} - \log \frac{p_{\theta}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t)}{p_{\text{ref}}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t)} \right] \right) \right]. \end{aligned}$$

The gradients of  $\tilde{L}_{\text{DPO-Diffusion}}$  and  $\tilde{L}_{\text{residue-DPO-Diffusion}}$  w.r.t the parameters  $\theta$  can be written as:

$$\begin{aligned} \nabla_{\theta} \tilde{L}_{\text{DPO-Diffusion}} &= -\beta T \mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \sum_{j=n+1}^{n+m} \text{sgn}(\mathcal{R}_1^0, \mathcal{R}_2^0) \right. \\ &\quad \left. \cdot \sigma(\hat{r}(\mathcal{R}_2^0) - \hat{r}(\mathcal{R}_1^0)) \left( \nabla_{\theta} \log p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t) - \nabla_{\theta} \log p_{\theta}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t) \right) \right], \end{aligned}$$

and

$$\begin{aligned} \nabla_{\theta} \tilde{L}_{\text{residue-DPO-Diffusion}} &= -\beta T \mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \sum_{j=n+1}^{n+m} \text{sgn}(\mathcal{R}_1^0[j], \mathcal{R}_2^0[j]) \right. \\ &\quad \left. \cdot \sigma(\hat{r}(\mathcal{R}_2^0[j]) - \hat{r}(\mathcal{R}_1^0[j])) \left( \nabla_{\theta} \log p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t) - \nabla_{\theta} \log p_{\theta}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t) \right) \right], \end{aligned}$$

where  $\hat{r}(\cdot) := \log(p_{\theta}(\cdot)/p_{\text{ref}}(\cdot))$ , which can be viewed as the estimated reward by current policy  $p_{\theta}$ .

We can see that  $\nabla_{\theta} \tilde{L}_{\text{DPO-Diffusion}}$  actually reweight  $\nabla_{\theta} \log p_{\theta}(\mathcal{R}^{t-1}[j]|\mathcal{R}^t)$  with the estimated reward of the complete antibody while  $\nabla_{\theta} \tilde{L}_{\text{residue-DPO-Diffusion}}$  does this with the estimated reward of the amino acid itself. In this case,  $\nabla_{\theta} \tilde{L}_{\text{DPO-Diffusion}}$  will increase (resp. decrease) the likelihood of all amino acids of the “winning” sample (resp. “losing”) at the same rate, which may mislead the optimization direction. In contrast,  $\nabla_{\theta} \tilde{L}_{\text{residue-DPO-Diffusion}}$  does not have this issue and can fully utilize the residue-level signals from estimated reward to effectively optimize antibodies.

We further approximate the objective  $\tilde{L}_{\text{residue-DPO-Diffusion}}$  by sampling from the forward diffusion process  $q$  instead of the reverse generative process  $p_{\theta}$  to achieve diffusion-like efficient training. With further replacing  $\log \frac{p_{\theta}}{p_{\text{ref}}}$  with  $-\log \frac{q}{p_{\theta}} + \log \frac{p_{\text{ref}}}{q}$  which is exactly  $-\mathbb{D}_{\text{KL}}(q||p_{\theta}) + \mathbb{D}_{\text{KL}}(q||p_{\text{ref}})$  when taking expectation with respect to  $q$ , we can derive the final loss for fine-tuning the diffusion model as follows:

$$\begin{aligned} L_{\text{ABDPO}} &= -\mathbb{E}_{t, \mathcal{R}_1^0, \mathcal{R}_2^0, (\mathcal{R}_1^{t-1}, \mathcal{R}_1^t), (\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)} \left[ \sum_{j=n+1}^{n+m} \log \sigma \left( -\beta T \text{sgn}(\mathcal{R}_1^0[j], \mathcal{R}_2^0[j]) \right. \right. \\ &\quad \left. \left. \cdot \left\{ \mathbb{D}_{\text{KL},1}^t(q||p_{\theta})[j] - \mathbb{D}_{\text{KL},1}^t(q||p_{\text{ref}})[j] - \mathbb{D}_{\text{KL},2}^t(q||p_{\theta})[j] + \mathbb{D}_{\text{KL},2}^t(q||p_{\text{ref}})[j] \right\} \right) \right], \quad (8) \end{aligned}$$

where  $\mathcal{R}_1^0, \mathcal{R}_2^0 \sim p_{\theta}(\mathcal{R})$ ,  $(\mathcal{R}_1^{t-1}, \mathcal{R}_1^t)$  and  $(\mathcal{R}_2^{t-1}, \mathcal{R}_2^t)$  are sampled from forward diffusion process of  $\mathcal{R}_1^0$  and  $\mathcal{R}_2^0$ , respectively, which can be much more efficient than the reverse generative process that involves hundreds of model forward estimation. Here we use  $\mathbb{D}_{\text{KL},1}^t(q||p_{\theta})[j]$  to denote  $\mathbb{D}_{\text{KL}}(q(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^{t-1}, \mathcal{R}_1^0)||p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^0))$ . Similar for  $\mathbb{D}_{\text{KL},1}^t(q||p_{\text{ref}})[j]$ ,  $\mathbb{D}_{\text{KL},2}^t(q||p_{\theta})[j]$ , and  $\mathbb{D}_{\text{KL},2}^t(q||p_{\text{ref}})[j]$ . These KL divergence can be estimated as in Eqs. (5) to (7).

### 3.3 Energy Decomposition and Conflict Mitigation

The energy usually consists of different types, such as attraction and repulsion. Empirically, direct optimization on single energy will lead to some undesired “shortcuts”. Specifically, in some cases, repulsion dominates the energy of the antibody so the model will push antibodies as far from the antigen as possible to decrease the repulsion during optimization, and finally fall into a bad local minima. This effectively reduces the repulsion, but also completely eliminates the attraction between antibodies and antigens, which seriously impairs the functionality of the antibody. This motivates us to explicitly express the energy with several distinct terms and then control the optimization process towards our preference.

Inspired by Yu et al. [51], we utilize “gradient surgery” to alleviate interference between different types of energy during energy preference optimization. More specifically, we have  $\mathcal{E}(\cdot) = \sum_{v=1}^V w_v \mathcal{E}_v(\cdot)$ , where  $V$  is the number of types of energy, and  $w_v$  is a constant weight for the  $v$ -th kind of energy. For each type of energy  $\mathcal{E}_v(\cdot)$ , we compute its corresponding energy preference gradient  $\nabla_{\theta} L_v$  as

Table 1: Summary of AAR, RMSD, CDR  $E_{\text{total}}$ , CDR-Ag  $\Delta G$  (kcal/mol), pLL, PHR, and  $N_{\text{success}}$  of antibodies designed by our model and baselines. ( $\downarrow$ ) / ( $\uparrow$ ) denotes a smaller / larger number is better.

Methods	AAR ( $\uparrow$ )	RMSD ( $\downarrow$ )	CDR $E_{\text{total}}$ ( $\downarrow$ )	CDR-Ag $\Delta G$ ( $\downarrow$ )	pLL ( $\uparrow$ )	PHR ( $\downarrow$ )	$N_{\text{success}}$ ( $\uparrow$ )
HERN	32.38%	9.18	10887.77	2095.88	-2.02	40.46%	0
MEAN	36.20%	<b>1.69</b>	7162.65	1041.43	<b>-1.79</b>	<b>30.62%</b>	0
dyMEAN	<b>40.04%</b>	1.82	3782.67	1730.06	-1.82	43.72%	0
DiffAb	34.92%	1.92	1729.51	1297.25	-2.10	41.27%	0
ABDPO	31.25%	1.98	<b>629.44</b>	<b>307.56</b>	-2.18	69.67%	<b>9</b>
ABDPO+	36.27%	2.01	1106.48	637.62	-2.00	44.21%	5

Eq. (8), and then alter the gradient by projecting it onto the normal plane of the other gradients (in a random order) if they have conflicts. This process works as follows:

$$\nabla_{\theta} L_v \leftarrow \nabla_{\theta} L_v - \frac{\min(\nabla_{\theta} L_v^{\top} \nabla_{\theta} L_u, 0)}{\|\nabla_{\theta} L_u\|^2} \nabla_{\theta} L_u, \quad (9)$$

where  $v \in \{1, \dots, V\}$  and  $u = \text{Shuffle}(1, \dots, V)$ .

## 4 Experiments

### 4.1 Experimental Setup

**Dataset Curation** To pre-train the diffusion model for antibody generation, we use the Structural Antibody Database [SAbDab, 13] under IMGT [34] scheme as the dataset. We collected antigen-antibody complexes with both heavy and light chains and protein antigens and discarded the duplicate data with the same CDR-L3 and CDR-H3 sequence. The remaining complexes are used to cluster via MMseqs2 [44] with 40% sequence similarity as the threshold based on the CDR-H3 sequence of each complex. We then select the clusters that do not contain complexes in RAbD benchmark [1] and split the complexes into training and validation sets with a ratio of 9:1 (1786 and 193 complexes respectively). Specifically, the validation set is composed of clusters that only contain one complex. The test set consists of 55 eligible complexes from the RAbD benchmark (details in Appendix D.2).

For the synthetic data used in ABDPO fine-tuning, 10,112 samples are randomly sampled for each antigen-antibody complex in the test set using the aforementioned pre-trained diffusion model. Then, we use pyRosetta [9] to apply the side-chain packing for these samples.

**Preference Definition** To apply ABDPO, we need to build the preference dataset and construct the “winning” and “losing” pair. The accurate relationship between preferences based on *in silico* with wet-lab experimental results is a scientific issue that remains unresolved, with a wide range of opinions. ABDPO’s solution to this open question is to provide a generic framework that allows for arbitrary definitions and combinations of preferences to satisfy various requirements in antibody design.

To demonstrate the effectiveness of ABDPO, we define the preferences as lower total energy and lower binding energy. The two energies are defined on residue level, specifically, (1)  $\text{Res}_{\text{CDR}} E_{\text{total}}$  is the total energy of each residue within the designed CDR, and is used to represent the overall rationality of the corresponding residue; (2)  $\text{Res}_{\text{CDR-Ag}} \Delta G$  is the interaction energy between each designed CDR residue and the target antigen, representing the functionality of the corresponding residue.  $\text{Res}_{\text{CDR-Ag}} \Delta G$  is further decomposed into (2.1)  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$ , the sum of the interaction energies except repulsion between the designed CDR residue and the antigen, and (2.2)  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$ , the repulsion energy between the design CDR residue and the antigen.

As a generic framework, ABDPO also supports non-energy-based preferences. To verify this, we demonstrate an advanced version named ABDPO+. ABDPO+ incorporates two additional preferences: pseudo log-likelihood (pLL) from AntiBERTy [42] and the percent of hydrophobicity residues (PHR). Different from the previously mentioned energy-based preferences, pLL and PHR are defined on the whole CDR level. For pLL, a higher value is considered better and is designated as “winning”, conversely; for PHR, a lower value is preferable.

**Baselines** We compare our model with various representative antibody sequence-structure co-design baselines. **HERN** [21] designs sequences of antibodies autoregressively with the iterative refinement




Figure 3: Visualization of reference antibodies in RABD and antibodies designed by ABDPO given specific antigens (PDB ID: **1iqd** (left), **1ic7** (middle), and **2dd8** (right)). The unit of energy annotated is kcal/mol and omitted here for brevity.

of structures; **MEAN** [29] generates sequences and structures of antibodies via a progress full-shot scheme; **dyMEAN** [30] designs antibodies sequences and structures with full-atom modeling; **DiffAb** [36] models antibody distributions with a diffusion model that considers the amino acid type,  $C_{\alpha}$  positions and side-chain orientations, which is a more rigorous *generative* model than the above baselines. The side-chain atoms are packed by pyRosetta. For **dyMEAN**, we (1) provide the ground-truth framework structure as input like other methods, (2) only use its generated backbones and pack the side-chain atoms by pyRosetta for a more fair comparison.

**Evaluation** Following the previous studies, we preliminarily evaluate the generated sequence and structure with AAR and  $C_{\alpha}$  RMSD. Besides, we carry out a series of more reasonable metrics. We utilize the preferences aforementioned to evaluate the designed antibodies from multiple perspectives, but at the whole CDR level. Specifically, (1) CDR  $E_{total}$ , the total energy of the designed CDR, is utilized to evaluate the rationality by aggregating all  $\text{Res}_{CDR} E_{total}$  of residues within the CDR; (2) CDR-Ag  $\Delta G$  denotes the difference in total energy between the bound state and the unbound state of the CDR and antigen, which is calculated to evaluate the functionality. PHR and pLL remain the same definition as above. All methods are able to generate multiple antibodies for a specific antigen (a randomized version of MEAN, rand-MEAN, is used here). We employ each method to design 192 antibodies for each complex, and we report the mean metrics across all 55 complexes. We further report the number of successfully designed antibody-antigen complexes,  $N_{success}$ , to evaluate their rationality and functionality comprehensively. The design for an antibody-antigen complex is considered as “successful” when at least one generated sample holds energies close to or lower than the natural one, i.e., for both of the two energy types,  $E_{generated} < E_{natural} + \text{std}(E_{natural}^{all-complexes})$ .

## 4.2 Main Results

We report the evaluation metrics in Tab. 1. As the results show, ABDPO performs significantly superior to other antibody sequence-structure co-design methods in the two energy-based metrics, CDR  $E_{total}$  and CDR-Ag  $\Delta G$ , while maintaining the AAR and RMSD. With the two additional preferences, ABDPO+ avoids the expense of the increased PHR while achieving better performance than DiffAb in remaining metrics (even surpassing DiffAb in AAR). This demonstrates the effectiveness and compatibility of ABDPO in terms of optimizing multi-objectives simultaneously. We have also provided the detailed evaluation results for each complex in Appendix E.2.

We do not consider AAR and RMSD as the main reference evaluation metrics as their inadequacy (refer to Appendix A for more details). With the new evaluation methods, issues that used to be hidden by AAR and RMSD are exposed. It is observed that structural clashes can not be avoided completely in any method, resulting in the high energy values of generated antibodies, even for ABDPO and ABDPO+. The structural clashes between CDR and the antigen finally lead to the unreasonable high CDR-Ag  $\Delta G$ . However, the primary goal in antibody design is to generate at least one effective antibody. Given the complexity of protein interactions, it is not plausible that every generated antibody will yield effectiveness. Therefore,  $N_{success}$  is a more valuable metric. ABDPO and ABDPO+ are the only two to achieve successful cases, with 9 and 5 successful cases out of 55 complexes, respectively. Following this concept, we also rank the designed antibodies for each complex by a uniform strategy (see Appendix D.3), calculate the metrics of the highest-ranked design for each complex, and report the mean metrics across the 55 complexes (see Appendix E.1). Notably, ABDPO is the only method that achieves CDR-Ag  $\Delta G$  lower than 0.






Figure 4: Changes of median CDR  $E_{total}$ ,  $E_{nonRep}$ ,  $E_{Rep}$ , and CDR-Ag  $\Delta G$  (kcal/mol) over-optimization steps, shaded to indicate interquartile range (from 25-th percentile to 75-th percentile).

We also visualize three cases (PDB ID: 1iqd, 1ic7, and 2dd8) in Fig. 3. It is shown that ABDPO can design CDRs with both fewer clashes and proper relative spatial positions towards the antigens, and even better energy performance than that of natural antibodies.

We conduct another two experiments to demonstrate further the generality of ABDPO: (1) directly incorporate auxiliary training losses for those properties of which gradients are computable; (2) introduce energy minimization before energy calculation, which is more in line with the real workflow. ABDPO shows consistent performance and demonstrates its generality. Please refer to Appendix F for related details.

### 4.3 Ablation Studies

Our approach comprises three main novel designs, including residue-level direct energy-based preference optimization, energy decomposition, and conflict mitigation by gradient surgery. Thus we perform comprehensive ablation studies to verify our hypothesis on the effects of each respective design component. Here we take the experiment on one complex (PDB ID: 1a14) as the example. Here, we apply more fine-tuning steps and additionally introduce  $E_{nonRep}$  (aggregation of Res<sub>CDR-Ag</sub>  $E_{nonRep}$  within the designed CDR),  $E_{Rep}$  (aggregation of Res<sub>CDR-Ag</sub>  $E_{Rep}$ ) for a more obvious and detailed comparison. More cases of ablation studies can be found in Appendix G.

**Effects of Residue-level Energy Preference Optimization** We hypothesize that residue-level DPO leads to more explicit and intuitive gradients that can promote effectiveness and efficiency compared with the vanilla DPO [46] as the analysis in Sec. 3.2. To validate this, we compare ABDPO with its counterpart with the CDR-level preference instead of residue-level. As Fig. 4 shows, regarding the counterpart (blue dotted line), the changes in all metrics are not obvious, while almost all metrics rapidly converge to an ideal state in ABDPO (red line). This demonstrated the effects of residue-level energy preference in improving the optimization efficiency.

**Effects of Energy Decomposition** In generated antibodies, the huge repulsion caused by clashes accounts for the majority of the two types of energy. This prevents us from using the  $\Delta G$  as an optimization objective directly as the model is allowed to minimize repulsion by keeping antibodies away from antigens, quickly reducing the energies. To verify this, we compared ABDPO with a version that directly optimize  $\Delta G$ . As shown in Fig. 4, without energy decomposition (green dashed line), both  $E_{Rep}$  and  $E_{nonRep}$  quickly diminish to 0, indicating that there is no interaction between the generated antibodies and antigens. Conversely, ABDPO (red line) can minimize  $E_{Rep}$  to 0 while maintaining  $E_{nonRep}$ , which means the interactions are preserved.

**Effects of Gradient Surgery** To show the effectiveness of gradient surgery in mitigating conflicts when optimizing multiple objectives, we compare ABDPO and its counterpart without gradient surgery. As Fig. 4 shows, the counterpart (purple dashed line) can only slightly optimize CDR-Ag  $E_{nonRep}$  but incurs strong repulsion (i.e.,  $E_{Rep}$ ), learning to irrational structures. ABDPO (red line) can converge to a state where CDR  $E_{total}$  and  $E_{Rep}$  achieve a conspicuously low point, suggesting the generated sequences and structures are stable, and  $E_{nonRep}$  is still significantly less than zero, showing that considerable binding affinity is kept.

**Comparison with Supervised Fine-tuning** Supervised Fine-tuning (SFT) can be an alternative way of generating antibodies with lower energy. For SFT, we first select the top 10% high-quality samples from ABDPO training data on a complex (PDB ID: 1a14). We fine-tune the diffusion model under the same settings as ABDPO. Results in Tab. 2 show that SFT only marginally surpasses the pre-trained diffusion model, and ABDPO performs significantly superior to SFT. We attribute the performance of ABDPO to the preference optimization scheme and the fine-grained residue-level energy rather than the entire CDR.

Table 2: Comparison of ABDPO and supervised fine-tuning (SFT) on 1a14.

Methods	CDR $E_{\text{total}}$ ( $\downarrow$ )		CDR-Ag $\Delta G$ ( $\downarrow$ )	
	Avg.	Med.	Avg.	Med.
DiffAb	1314.20	1133.36	534.21	248.28
DiffAb <sub>SFT</sub>	1053.82	869.37	374.27	144.25
ABDPO	<b>336.02</b>	<b>226.25</b>	<b>88.64</b>	<b>0.10</b>

## 5 Conclusions

In this work, we rethink antibody sequence-structure co-design through the lens of energy and propose ABDPO for designing antibodies meeting multi-objectives like rationality and functionality. The introduction of direct energy-based preference optimization along with energy decomposition and conflict mitigation by gradient surgery shows promising results in generating antibodies with low energy and high binding affinity. With ABDPO, existing computing software and domain knowledge can be easily combined with deep learning techniques, jointly facilitating the development of antibody design. Limitations and future work are discussed in Appendix H.

## References

- [1] Jared Adolf-Bryfogle, Oleks Kalyuzhnyi, Michael Kubitz, Brian D Weitzner, Xiaozhen Hu, Yumiko Adachi, William R Schief, and Roland L Dunbrack Jr. 2018. RosettaAntibodyDesign (RABD): A general framework for computational antibody design. *PLoS computational biology*, 14(4):e1006112.
- [2] Rahmad Akbar, Philippe A. Robert, Milena Pavlović, Jeliasko R. Jeliaskov, Igor Snapkov, Andrei Slabodkin, Cédric R. Weber, Lonneke Scheffer, Enkelejda Miho, Ingrid Hobæk Haff, Dag Trygve Tryslew Haug, Fridtjof Lund-Johansen, Yana Safonova, Geir K. Sandve, and Victor Greiff. 2021. A compact vocabulary of paratope-epitope interactions enables predictability of antibody-antigen binding. *Cell Reports*, 34(11):108856.
- [3] Rahmad Akbar, Philippe A Robert, Cédric R Weber, Michael Widrich, Robert Frank, Milena Pavlović, Lonneke Scheffer, Maria Chernigovskaya, Igor Snapkov, Andrei Slabodkin, et al. 2022. In silico proof of principle of machine learning-based antibody design at unconstrained scale. In *MABs*, volume 14, page 2031482. Taylor & Francis.
- [4] Rebecca F. Alford, Andrew Leaver-Fay, Jeliasko R. Jeliaskov, Matthew J. O’Meara, Frank P. DiMaio, Hahnbeom Park, Maxim V. Shapovalov, P. Douglas Renfrew, Vikram K. Mulligan, Kalli Kappel, Jason W. Labonte, Michael S. Pacella, Richard Bonneau, Philip Bradley, Roland L. Jr. Dunbrack, Rhiju Das, David Baker, Brian Kuhlman, Tanja Kortemme, and Jeffrey J. Gray. 2017. The Rosetta All-Atom Energy Function for Macromolecular Modeling and Design. *Journal of Chemical Theory and Computation*, 13(6):3031–3048. PMID: 28430426.
- [5] Ethan C Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M Church. 2019. Unified rational protein engineering with sequence-based deep representation learning. *Nature methods*, 16(12):1315–1322.
- [6] Anonymous. 2023. Proximal Preference Optimization for Diffusion Models.
- [7] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. 2023. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*.
- [8] Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- [9] Sidhartha Chaudhury, Sergey Lyskov, and Jeffrey J. Gray. 2010. PyRosetta: a script-based interface for implementing molecular modeling algorithms using Rosetta. *Bioinformatics*, 26(5):689–691.
- [10] Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*.
- [11] Gavin E Crooks, Gary Hon, John-Marc Chandonia, and Steven E Brenner. 2004. WebLogo: a sequence logo generator. *Genome research*, 14(6):1188–1190.

- [12] Jinhui Dong, Seth J Zost, Allison J Greaney, Tyler N Starr, Adam S Dingens, Elaine C Chen, Rita E Chen, James Brett Case, Rachel E Sutton, Pavlo Gilchuk, et al. 2021. Genetic and structural basis for SARS-CoV-2 variant neutralization by a two-antibody cocktail. *Nature microbiology*, 6(10):1233–1244.
- [13] James Dunbar, Konrad Krawczyk, Jinwoo Leem, Terry Baker, Angelika Fuchs, Guy Georges, Jiye Shi, and Charlotte M Deane. 2014. SAbDab: the structural antibody database. *Nucleic acids research*, 42(D1):D1140–D1146.
- [14] Peter Eastman, Jason Swails, John D Chodera, Robert T McGibbon, Yutong Zhao, Kyle A Beauchamp, Lee-Ping Wang, Andrew C Simmonett, Matthew P Harrigan, Chaya D Stern, et al. 2017. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS computational biology*, 13(7):e1005659.
- [15] Stefan Ewert, Annemarie Honegger, and Andreas Plückthun. 2004. Stability improvement of antibodies for extracellular and intracellular applications: CDR grafting to stable frameworks and structure-based framework engineering. *Methods*, 34(2):184–199. Intrabodies.
- [16] Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. 2023. Reinforcement Learning for Fine-tuning Text-to-Image Diffusion Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [17] Noelia Ferruz, Steffen Schmidt, and Birte Höcker. 2022. ProtGPT2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):4348.
- [18] Jean Gallier and Dianna Xu. 2003. Computing exponentials of skew-symmetric matrices and logarithms of orthogonal matrices. *International Journal of Robotics and Automation*, 18(1):10–20.
- [19] Kaiyuan Gao, Lijun Wu, Jinhua Zhu, Tianbo Peng, Yingce Xia, Liang He, Shufang Xie, Tao Qin, Haiguang Liu, Kun He, et al. 2023. 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*, pages 506–517.
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc.
- [21] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2022. Antibody-antigen docking and design via hierarchical structure refinement. In *International Conference on Machine Learning*, pages 10217–10227. PMLR.
- [22] Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi S. Jaakkola. 2022. Iterative Refinement Graph Neural Network for Antibody Sequence-Structure Co-design. In *International Conference on Learning Representations*.
- [23] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror. 2021. Learning from Protein Structure with Geometric Vector Perceptrons. In *International Conference on Learning Representations*.
- [24] Peter T Jones, Paul H Dear, Jefferson Foote, Michael S Neuberger, and Greg Winter. 1986. Replacing the complementarity-determining regions in a human antibody with those from a mouse. *Nature*, 321(6069):522–525.
- [25] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589.
- [26] Kazutaka Katoh and Daron M Standley. 2013. MAFFT multiple sequence alignment software version 7: improvements in performance and usability. *Molecular biology and evolution*, 30(4):772–780.

- [27] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [28] Miltiadis Kofinas, Naveen Shankar Nagaraja, and Efstratios Gavves. 2021. Roto-translated Local Coordinate Frames For Interacting Dynamical Systems. In *Advances in Neural Information Processing Systems*.
- [29] Xiangzhe Kong, Wenbing Huang, and Yang Liu. 2023. Conditional Antibody Design as 3D Equivariant Graph Translation. In *The Eleventh International Conference on Learning Representations*.
- [30] Xiangzhe Kong, Wenbing Huang, and Yang Liu. 2023. End-to-End Full-Atom Antibody Design. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17409–17429. PMLR.
- [31] Gideon D Lapidoth, Dror Baran, Gabriele M Pszolla, Christoffer Norn, Assaf Alon, Michael D Tyka, and Sarel J Fleishman. 2015. Abdesign: A n algorithm for combinatorial backbone design guided by natural conformations and sequences. *Proteins: Structure, Function, and Bioinformatics*, 83(8):1385–1406.
- [32] Adam Leach, Sebastian M Schmon, Matteo T Degiacomi, and Chris G Willcocks. 2022. Denoising diffusion probabilistic models on so (3) for rotational alignment. In *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*.
- [33] Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*.
- [34] Marie-Paule Lefranc, Christelle Pommié, Manuel Ruiz, Véronique Giudicelli, Elodie Foulquier, Lisa Truong, Valérie Thouvenin-Contet, and Gérard Lefranc. 2003. IMGT unique numbering for immunoglobulin and T cell receptor variable domains and Ig superfamily V-like domains. *Developmental & Comparative Immunology*, 27(1):55–77.
- [35] Ge Liu, Haoyang Zeng, Jonas Mueller, Brandon Carter, Ziheng Wang, Jonas Schilz, Geraldine Horny, Michael E Birnbaum, Stefan Ewert, and David K Gifford. 2020. Antibody complementarity determining region design using high-capacity machine learning. *Bioinformatics*, 36(7):2126–2133.
- [36] Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. 2022. Antigen-Specific Antibody Design and Optimization with Diffusion-Based Generative Models for Protein Structures. In *Advances in Neural Information Processing Systems*.
- [37] Karolis Martinkus, Jan Ludwiczak, WEI-CHING LIANG, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, and Andreas Loukas. 2023. AbDiffuser: full-atom generation of in-vitro functioning antibodies. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [38] Sanzo Miyazawa and Robert L Jernigan. 1985. Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation. *Macromolecules*, 18(3):534–552.
- [39] Kenneth Murphy and Casey Weaver. 2016. *Janeway’s immunobiology*. Garland science.
- [40] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*.
- [41] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In *Thirty-seventh Conference on Neural Information Processing Systems*.

- [42] Jeffrey A Ruffolo, Jeffrey J Gray, and Jeremias Sulam. 2021. Deciphering antibody affinity maturation with language models and weakly supervised learning. *arXiv preprint arXiv:2112.07782*.
- [43] Koichiro Saka, Taro Kakuzaki, Shoichi Metsugi, Daiki Kashiwagi, Kenji Yoshida, Manabu Wada, Hiroyuki Tsunoda, and Reiji Teramoto. 2021. Antibody design using LSTM based deep generative model from phage display library for affinity maturation. *Scientific reports*, 11(1):5852.
- [44] Martin Steinegger and Johannes Söding. 2017. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature biotechnology*, 35(11):1026–1028.
- [45] Gabriel D Victora and Michel C Nussenzweig. 2012. Germinal centers. *Annual review of immunology*, 30:429–457.
- [46] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. 2023. Diffusion Model Alignment Using Direct Preference Optimization. *arXiv preprint arXiv:2311.12908*.
- [47] Shira Warszawski, Aliza Borenstein Katz, Rosalie Lipsh, Lev Khmelnsky, Gili Ben Nissan, Gabriel Javitt, Orly Dym, Tamar Unger, Orli Knop, Shira Albeck, et al. 2019. Optimizing antibody affinity and stability by the automated design of the variable light-heavy chain interfaces. *PLoS computational biology*, 15(8):e1007207.
- [48] Fang Wu and Stan Z. Li. 2023. A Hierarchical Training Paradigm for Antibody Structure-sequence Co-design. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [49] John L Xu and Mark M Davis. 2000. Diversity in the CDR3 Region of VH Is Sufficient for Most Antibody Specificities. *Immunity*, 13(1):37–45.
- [50] Jason Yim, Brian L. Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. 2023. SE(3) diffusion model with application to protein backbone generation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 40001–40039. PMLR.
- [51] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836.
- [52] Zaixiang Zheng, Yifan Deng, Dongyu Xue, Yi Zhou, Fei YE, and Quanquan Gu. 2023. Structure-informed Language Models Are Protein Designers. In *International Conference on Machine Learning*.
- [53] Xiangxin Zhou, Liang Wang, and Yichi Zhou. 2024. Stabilizing Policy Gradients for Stochastic Differential Equations via Consistency with Perturbation Process. *arXiv preprint arXiv:2403.04154*.

## A Motivation for Choosing Energy as Evaluation

There are many inadequacies in using AAR and RMSD as the main evaluation metrics in AI-based antibody design. Antibody design is a typical function-oriented protein design task, necessitating a more fine-grained measure of discrepancy compared to general protein design tasks. Especially when the part of the antibody to be designed and evaluated, CDR-H3, is usually shorter, more precise evaluation becomes particularly important.




Figure 5: **A:** Tyr (Y) and Phe (F) differ by only one oxygen atom. In contrast, there is a substantial difference between Gly (G) and Trp (W). Gly lacks a side chain, whereas Trp possesses the largest side chain of all amino acids. **B:** the visualization of the frequency of occurrence of each amino acid at various positions in RAbD CDR-H3 sequences. The sequences are initially aligned using MAFFT [26] and subsequently visualized with WebLogo [11]. The width of each column corresponds to the frequency of occurrence at that position.

For AAR, there are two main limitations in measuring the similarity between the generated sequence and the reference sequence. The first limitation is located in measuring the difference in different incorrect recoveries. Among the 20 common amino acids, some have high similarity between them, such as Tyr and Phe, while others have significant differences, such as Gly and Trp (Fig. 5A). When an amino acid in CDR is erroneously recovered to different amino acids, their impact will also vary. However, AAR does not differentiate between these different types of errors, only identifying them as “incorrect”.

A further, more serious issue is that AAR is easily hacked. Although the CDR region is often considered hypervariable, a mild conservatism in sequence still exists (Fig. 5B), which allows the model to obtain satisfactory AAR using a simple but incorrect way - directly generating the amino acids with the highest probability of occurrence at each position, while ignoring the condition of the given antigen which is extremely harmful to the specificity of antibodies. We made a simple attempt by simply counting the amino acids with the highest frequency of occurrence at various positions in all samples in SAbDab, and then composing them into a CDR-H3 sequence, which looks roughly like “ARD + rand(Y, G)\* + FDY”, achieving an AAR of **38.77%** on the RAbD dataset.




Figure 6: The distribution of CDR-Ag  $E_{nonRep}$  (left) and CDR-Ag  $E_{Rep}$  (right) formed by the whole CDR atoms (colored in red) and solely by CDR side-chain atoms (colored in blue) among SAbDab dataset.

While RMSD fails to measure the discrepancies on side-chain atoms, in general, the calculation of RMSD focuses on the alpha carbon atom or the four backbone atoms due to their stable existence in any type of amino acid and thus ignores the side-chain atoms. However, side-chain atoms in the CDR region are extremely important as they contribute to most of the interactions between the CDR and the antigen. Our analyses on the SAbDab dataset also prove the importance of the side chain in CDR-Antigen interaction in terms of energy. As shown in Fig. 6, the distribution of energies formed

by the whole residues in CDR is colored in red while the distribution of energies formed only by side-chain atoms of CDR is colored in blue. The interaction energy formed by side-chain atoms accounts for the vast majority of the total interaction energy in both types of energy.

The above reasons have led us to abandon AAR and RMSD as learning objectives and evaluation metrics, and instead use energy as our goal. Energy can simultaneously consider the relationship between structure and sequence, distinguish different generation results in more detail, and importantly, reflect the rationality and functionality of antibodies in a more fundamental way. Despite the various shortcomings of AAR and RMSD, we have demonstrated that the antibodies generated by ABDPO achieve lower AAR and comparable RMSD compared to those generated by other methods. However, in practice, ABDPO-generated antibodies exhibit distinct binding patterns to antigens, differing from reference antibodies, and demonstrate significantly better energy performance than those produced by other methods. This further highlights the inadequacies of using AAR and RMSD as evaluation metrics in antibody design tasks, exposing their vulnerability to being “hacked”.

## B Energy Calculation

In ABDPO, we conduct the calculation on  $\text{Res}_{\text{CDR}} E_{\text{total}}$  at residue level, and a more fine-grained calculation on the two functionality-associated energies at the sub-residue level. We use Rosetta to calculate all types of energies in this paper.

We denote the residue with the index  $i$  in the antibody-antigen complex as  $A_i$ , then  $A_i^{sc}$  and  $A_i^{bb}$  represent the **side chain** and **backbone** of the residue respectively.

For the energies in the proposed preference, we describe the function for energies of a **Single** residue as ES, and  $\text{ES}_{\text{total}}$  is the sum of all types of energy with the default weight in REF15 [4]. The function for interaction energies between **Paired** residues is described as EP, which consists of six different energy types:  $\text{EP}_{\text{hbond}}$ ,  $\text{EP}_{\text{att}}$ ,  $\text{EP}_{\text{rep}}$ ,  $\text{EP}_{\text{sol}}$ ,  $\text{EP}_{\text{elec}}$ , and  $\text{EP}_{\text{lk}}$ .

Following the settings previously mentioned in Sec. 3.1, the indices of residues within the CDR-H3 range from  $n + 1$  to  $n + m$ , and the indices of residues within the antigen range from  $g + 1$  to  $g + k$ . Then, for the CDR residue with the index  $j$ , the three types of energy are defined as:

$$\text{Res}_{\text{CDR}} E_{\text{total}}^j = \text{ES}_{\text{total}}(A_j), \quad (10)$$

$$\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}^j = \sum_{i=g+1}^{g+k} \sum_{e \in \{\text{hbond}, \text{att}, \text{sol}, \text{elec}, \text{lk}\}} \left( \text{EP}_e(A_j^{sc}, A_i^{sc}) + \text{EP}_e(A_j^{sc}, A_i^{bb}) \right), \quad (11)$$

$$\begin{aligned} \text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}^j = & \sum_{i=g+1}^{g+k} \left( \text{EP}_{\text{rep}}(A_j^{sc}, A_i^{sc}) + \text{EP}_{\text{rep}}(A_j^{sc}, A_i^{bb}) \right. \\ & \left. + 2 \times \text{EP}_{\text{rep}}(A_j^{bb}, A_i^{sc}) + 2 \times \text{EP}_{\text{rep}}(A_j^{bb}, A_i^{bb}) \right). \end{aligned} \quad (12)$$

It can be observed from Eqs. (11) and (12) that the two functionality-associated energies, namely  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$  and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$ , which collectively describe the interaction energy between CDR and the antigen, are computed at the level of side-chain and backbone.  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$  is only calculated on the interactions caused by the side-chain atoms in the CDR-H3 region, while  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$  assigns a greater cost to the repulsions caused by the backbone atoms in the CDR-H3 region. This modification is carried out according to the fact that the side-chain atoms contribute the vast majority of energy to the interaction between CDR-H3 and antigens (Fig. 6), and  $E_{\text{nonRep}}$  exhibits a benefit in interactions, while  $E_{\text{Rep}}$  could be regarded as a cost.

The fine-grained calculation of  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$  and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$  is indispensable. Without the fine-grained calculation, the model tends to generate poly-G CDR-H3 sequences, such as “GGGGGGGGGG” for any given antigen and the rest of the antibody. The most likely reason for this is that G, Glycine, can maximize the reduction of clashes and gain satisfactory CDR  $E_{\text{total}}$  and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$  as it doesn’t contain side chain and simultaneously form a weak attraction to the antigen solely relying on its backbone atoms.

We emphasize that the two functionality-associated energies,  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$  and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$  are calculated exclusively at the sub-residue level when serving as the determination of preference in

guiding the direct energy-based preference optimization process. However, when these energies are used as evaluation metrics, they are calculated at the residue level, in which the greater cost to the repulsions attributed to the backbone atoms is negated.

## C Theoretical Justification

In this section, we show the detailed mathematical derivations of formulas in Sec. 3.2. Although many of them are similar to Rafailov et al. [41], we still present them in detail for the sake of completeness. Besides, we will also present the details of preference data generation.

First, we will show the derivation of the optimal solution of the KL-constrained reward-maximization objective, i.e.,  $\max_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} [r(\mathcal{R}^0)] - \beta \mathbb{D}_{\text{KL}}(p_\theta(\mathcal{R}^0) \| p_{\text{ref}}(\mathcal{R}^0))$  as follows:

$$\begin{aligned} & \max_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} [r(\mathcal{R}^0)] - \beta \mathbb{D}_{\text{KL}}(p_\theta(\mathcal{R}^0) \| p_{\text{ref}}(\mathcal{R}^0)) \\ &= \max_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} \left[ r(\mathcal{R}^0) - \beta \log \frac{p_\theta(\mathcal{R}^0)}{p_{\text{ref}}(\mathcal{R}^0)} \right] \\ &= \min_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} \left[ \log \frac{p_\theta(\mathcal{R}^0)}{p_{\text{ref}}(\mathcal{R}^0)} - \frac{1}{\beta} r(\mathcal{R}^0) \right] \\ &= \min_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} \left[ \log \frac{p_\theta(\mathcal{R}^0)}{\frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right)} - \log Z \right] \end{aligned}$$

where  $Z$  is the partition function that does not involve the model being trained, i.e.,  $p_\theta$ . And we can define

$$p^*(\mathcal{R}^0) := \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right).$$

With this, we can now arrive at

$$\begin{aligned} & \min_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} \left[ \log \frac{p_\theta(\mathcal{R}^0)}{p^*(\mathcal{R}^0)} \right] - \log Z \\ &= \min_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} [\mathbb{D}_{\text{KL}}(p_\theta \| p^*)] + Z \end{aligned}$$

Since  $Z$  does not depend on  $p_\theta$ , we can directly drop it. According to Gibb's inequality that KL-divergence is minimized at 0 if and only if the two distributions are identical. Hence we arrive at the optimum as follows:

$$p_{\theta^*}(\mathcal{R}^0) = p^*(\mathcal{R}^0) = \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right). \quad (13)$$

Then we will show that the objective that maximizes likelihood on preference data sampled from  $p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = \sigma(r(\mathcal{R}_1^0) - r(\mathcal{R}_2^0))$ , which is exactly  $L_{\text{DPO}}$ , leads to the same optimal solution. For this, we need to express the pre-defined reward  $r(\cdot)$  with the optimal policy  $p^*$ :

$$r(\mathcal{R}^0) = \beta \log \frac{p^*(\mathcal{R}^0)}{p_{\text{ref}}(\mathcal{R}^0)} + Z$$

The we plugin the expression of  $r(\cdot)$  into  $p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = \sigma(r(\mathcal{R}_1^0) - r(\mathcal{R}_2^0))$  as follows:

$$\begin{aligned} p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) &= \sigma(r(\mathcal{R}_1^0) - r(\mathcal{R}_2^0)) \\ &= \sigma\left(\beta \log \frac{p^*(\mathcal{R}_1^0)}{p_{\text{ref}}(\mathcal{R}_1^0)} - \beta \log \frac{p^*(\mathcal{R}_2^0)}{p_{\text{ref}}(\mathcal{R}_2^0)}\right), \end{aligned}$$

where  $Z$  is canceled out. For brevity, we use the following notation for brevity:

$$p_\theta(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = \sigma\left(\beta \log \frac{p_\theta(\mathcal{R}_1^0)}{p_{\text{ref}}(\mathcal{R}_1^0)} - \beta \log \frac{p_\theta(\mathcal{R}_2^0)}{p_{\text{ref}}(\mathcal{R}_2^0)}\right).$$

With this, we have

$$\begin{aligned} \min_{p_\theta} L_{\text{DPO}} &= \min_{p_\theta} -\mathbb{E}_{\mathcal{R}_1^0, \mathcal{R}_2^0 \sim p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0)} p_\theta(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) \\ &= \max_{p_\theta} \mathbb{E}_{\mathcal{R}_1^0, \mathcal{R}_2^0 \sim p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0)} p_\theta(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) \\ &= \min_{p_\theta} \mathbb{D}_{\text{KL}}\left(p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) \parallel p_\theta(\mathcal{R}_1^0 \succ \mathcal{R}_2^0)\right) \end{aligned}$$



Again with Gibb’s inequality, we can easily identify that  $p_{\theta}(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = p(\mathcal{R}_1^0 \succ \mathcal{R}_2^0)$  achieves the minimum. Thus  $p^*(\mathcal{R}^0) = \frac{1}{2}p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta}r(\mathcal{R}^0)\right)$  is also the optimal solution of  $L_{\text{DPO}}$ .

## D Implementation Details

### D.1 Model Details

The architecture of the diffusion model used in our method is the same as Luo et al. [36]. The input of the model is the perturbed CDR-H3 and its surrounding context, i.e., 128 nearest residues of the antigen or the antibody framework around the residues of CDR-H3. The input is composed of single residue embeddings and pairwise embeddings. The single residue embedding encodes the information of its amino acid types, torsional angles, and 3D coordinates of all heavy atoms. The pairwise embedding encodes the Euclidean distances and dihedral angles between the two residues. The sizes of the single residue feature and the residue-pair features are 1285 and 64, respectively. Then the features are processed by Multiple Layer Perceptrons (MLPs). The number of layers is 6. The size of the hidden state in the layers is 128. The output of the model is the predicted categorical distribution of amino acid types,  $C_{\alpha}$  coordinates, and a  $so(3)$  vector for the rotation matrix.

The number of diffusion steps is 100. We use the cosine  $\beta$  schedule with  $s = 0.01$  suggested in Ho et al. [20] for amino acid types,  $C_{\alpha}$  coordinates, and orientations.

### D.2 Training Details

**Pre-training** Following Luo et al. [36], the diffusion model is first trained via the gradient descent method Adam [27] with `init_learning_rate=1e-4`, `betas=(0.9,0.999)`, `batch_size=16`, and `clip_gradient_norm=100`. During the training phase, the weight of rotation loss, position loss, and sequence loss are each set to 1.0. We also schedule to decay the learning rate multiplied by a factor of 0.8 and a minimum learning rate of  $5e - 6$ . The learning rate is decayed if there is no improvement for the validation loss in 10 evaluations. The evaluation is performed for every 1000 training steps. We trained the model on one NVIDIA A100 80G GPU and it could converge within 30 hours and 200k steps.

**Test set** The original RABD dataset contains 60 antibody-antigen complexes. In this study, we hope all the complex consists of an antibody heavy chain and a light chain, and at least one protein antigen chain. In practice, **2ghw** and **3uzq** lack light chains, while **3h3b** lacks heavy chains. **5d96** was excluded because of the incorrect chain ID information in `rabd_summary.jsonl`<sup>3</sup>, where heavy chain  $J$  and light chain  $I$  do not bind to antigen chain  $A$ . As for **4etq**, we actually conducted the training (CDR  $E_{\text{total}}=70.55$ , CDR-Ag  $\Delta G=-4.57$ ), but HERN reported an error when running for this complex, so we did not report it.

**Pair data construction** In terms of the construction of “winning” and “losing” data pair, we did not pre-define “preferred” and “non-preferred” datasets but rather constructed a unified data pool. During each training step, the paired data used for DPO training is randomly sampled from the data pool. Although their energies and properties have been pre-calculated, the “winning” and “losing” labels are determined in real time. In practice, we used several labels, involving three different preferences related to energy and two preferences related to non-energy-based properties. The “winning” and “losing” labels among these preferences are not necessarily consistent. Therefore, the loss for each type of energy/preference is calculated separately and then aggregated with different weights to update the entire model. Moreover, as the training progresses, we continuously sample new data, calculate their energy, add them to the data pool, and discard some of the older post-added data simultaneously to ensure that the data stays in sync with the policy.

**Fine-tuning** For ABDPO fine-tuning, the pre-trained diffusion model is further fine-tuned via the gradient descent method Adam with `init_learning_rate=1e-5`, `betas=(0.9,0.999)`, and `clip_gradient_norm=100`. The batch size is 48. More specifically, in a batch, there are 48 pairs of preference data. We do not use a decay learning rate and do not use weight decay in the fine-tuning process. And we use  $\beta = 0.01$  and  $0.005$  in Eq. (8). We use the hyperparameter search space as follows. As for the three energies introduced in Sec. 4.1, we use 8:8:2 to reweight them (i.e.,  $\text{Res}_{\text{CDR}} E_{\text{total}}$ ,  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$ , and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$ ), and reweight pLL and PHR in ABDPO+

<sup>3</sup>[https://github.com/THUNLP-MT/MEAN/blob/main/summaries/rabd\\_summary.jsonl](https://github.com/THUNLP-MT/MEAN/blob/main/summaries/rabd_summary.jsonl)

to 1. In practice, different antibody-antigen complexes prefer different hyperparameters. For a fair comparison with baselines, we do not carefully pick the optimal hyperparameter for each complex but use a uniform hyperparameter. We fine-tune the pre-trained diffusion model on four NVIDIA A800 40G GPUs for 1,800 steps for each antigen, separately.

### D.3 Ranking Strategy

To rank the numerous generated antibodies with multiple energy labels, we applied a simple ranking strategy based on single energy metrics. The CDR  $E_{\text{total}}$  and the CDR-Ag  $\Delta G$  of each antibody are ranked independently. Then, a composite ranking score for each antibody is defined as the sum of its CDR  $E_{\text{total}}$  rank and CDR-Ag  $\Delta G$  rank (for ABDPO+, PHR and pLL are also involved). Finally, the antibodies are ranked according to these composite scores. We acknowledge that this ranking strategy has several limitations. For instance:

1. Equal weights are assigned to all energy types and properties, despite them having differing importance in reality.
2. The distribution patterns of different energy types and properties can vary, with these distributions usually being non-uniform. This could result in scenarios where minor numerical differences in the top-ranking CDR-Ag  $\Delta G$  values coincide with larger differences in CDR  $E_{\text{total}}$ , potentially leading to the selection of samples with suboptimal CDR  $E_{\text{total}}$ .

However, addressing these issues would require extensive and in-depth exploration of antibody binding mechanisms and energy calculation methodologies. We chose this straightforward, yet impartial, ranking strategy for two key reasons:

1. The primary goal of this work is to reformulate the antibody design task as an energy-focused optimization problem and propose a feasible implementation, rather than to delve into the mechanisms of antibody-antigen binding;
2. Our approach is designed to avoid introducing statistical biases or preferences based on potentially erroneous prior knowledge or favoritism towards particular antibody design methods.

## E More Evaluation Results

### E.1 Evaluation Results for Ranked Top-1 Design

In Tab. 1, we have reported the average results of all antibodies designed by our method and other baselines. Here we provide the evaluation results for the ranked top-1 design in Tab. 3 (refer to the ranking strategy in Appendix D.3).

Table 3: Average performance of top-1 designs of 55 complexes designed by baselines and our model.

Methods	CDR $E_{\text{total}}$ ( $\downarrow$ )	CDR-Ag $\Delta G$ ( $\downarrow$ )	PHR ( $\downarrow$ )	pLL ( $\uparrow$ )	AAR ( $\uparrow$ )	RMSD ( $\downarrow$ )
RAbD	5.25	-13.04	45.78%	-2.20	100.00%	0.00
HERN	8495.56	1296.22	48.18%	-2.01	33.29%	9.21
MEAN	3867.47	207.99	36.91%	-1.72	35.18%	1.70
dyMEAN	2987.93	1283.97	46.27%	-1.79	<b>40.74%</b>	1.81
DiffAb	381.82	58.84	49.19%	-2.03	37.99%	1.62
ABDPO	<b>68.51</b>	<b>-4.96</b>	69.97%	-2.15	32.92%	<b>1.58</b>
ABDPO+	332.10	29.27	<b>32.81%</b>	<b>-1.54</b>	39.55%	1.67

### E.2 Detailed Evaluation Results for each Complex

In Tab. 4 and Tab. 5, we list the CDR  $E_{\text{total}}$ , CDR-Ag  $\Delta G$ , PHR and pLL of the reference antibody in RAbD and the average/ranked top-1 antibodies designed by HERN, MEAN, dyMEAN, DiffAb, ABDPO, and ABDPO+ for each complex in the test set separately. In Tab. 5, we highlight the energy values of the designed complexes that surpass the natural one in terms of two energies simultaneously with **bold text**.








Figure 7: **Left**: the distribution of peptide bond length within CDR-H3 in the SAbDab dataset; **Right**: the kernel density estimation (KDE) function fit on the natural peptide bond length distribution.

## F Arbitrary Preferences

### F.1 Incorporating Auxiliary Loss

A predominant advantage of the ABDPO is its unique capacity to seamlessly integrate traditional bioinformatics, computational biology, and computational chemistry tools — those incapable of directly computing gradients — into the training regimen of AI models. This integration significantly broadens the ABDPO’s applicability and versatility in antibody design. However, it is pertinent to acknowledge the existence of antibody energies/properties for which gradient calculations are feasible. Indeed, fundamental geometric characteristics, such as bond lengths, angles, and torsion angles, alongside more intricate properties predicted by deep-learning models, are gradient-computable. These gradient-computable features offer an explicit direction for optimization, potentially enhancing the effectiveness and efficiency of the model optimization process.

In light of this, we initiated another experiment aimed at exploring ABDPO’s compatibility with traditional gradient-based losses, extending beyond the DPO loss. Specifically, we propose a special version based on ABDPO+, ABDPO++, which incorporates an auxiliary loss about peptide bond length. As a covalent bond, the variation range of peptide bond lengths is very limited, and thus we can consider the length of peptide bonds to be a fixed value and then utilize an MSE loss to directly penalize the unreasonable peptide bond length in generated antibodies.

In practice, we consider the ground truth peptide bond length to be 1.3310 (the average length of peptide bonds within CDR-H3 in SAbDab, the distribution could be seen in Fig. 7 left) and apply the auxiliary loss only when the sampled  $t$  is near 0 ( $t < 15$  in this experiment while  $T$  is 100), and the weight is set to 0.25. The peptide bond length is calculated based on the predicted  $(s_j^0, \mathbf{x}_j^0, \mathbf{O}_j^0)$  which is denoised with one step from  $(s_j^t, \mathbf{x}_j^t, \mathbf{O}_j^t)$ , then an MSE loss of peptide bond length can be calculated. Finally, this auxiliary loss, together with various DPO losses, updates the model through the conflict mitigation mentioned in Sec. 3.3.

Table 6: Summary of CDR  $E_{\text{total}}$ , CDR-Ag  $\Delta G$  (kcal/mol), pLL, PHR, C-N<sub>score</sub>, AAR, and RMSD of reference antibodies and antibodies designed by ABDPOW/O and baselines in the experiment involves auxiliary loss. (↓) / (↑) denotes a smaller / larger number is better.

Methods	CDR $E_{\text{total}}$ (↓)	CDR-Ag $\Delta G$ ↓	pLL (↑)	PHR (↓)	C-N <sub>score</sub> (↑)	AAR (↑)	RMSD (↓)
HERN	10887.77	2095.88	-2.02	40.46%	0.12	32.38%	9.18
MEAN	7162.65	1041.43	<b>-1.79</b>	<b>36.20%</b>	1.68	36.30%	<b>1.69</b>
dyMEAN	3782.67	1730.06	-1.82	43.72%	2.08	<b>40.04%</b>	1.82
DiffAb	1729.51	1297.25	-2.10	41.27%	3.85	34.92%	1.92
ABDPO	<b>629.44</b>	<b>307.56</b>	-2.18	69.67%	2.55	31.25%	1.98
ABDPO+	1106.48	637.62	-2.00	44.21%	2.95	36.27%	2.01
ABDPO++	1349.39	747.89	-1.99	44.46%	<b>4.51</b>	36.30%	1.95

To evaluate the consistency of generated antibodies’ peptide bond length to the natural antibodies, we fit a Kernel Density Estimation function using the length of peptide bonds found within the CDR-H3 region of natural antibodies (shown in Fig. 7 right), then the density of the generated peptide bond length,  $C-N_{\text{score}}$ , is used to represent the consistency. We report the average experiment result in Tab. 6. It can be observed that ABDPO++ significantly optimized the length of the peptide bond, achieving the best  $C-N_{\text{score}}$  of 4.51, while maintaining the optimization to the other 4 preferences. The experimental result demonstrates the compatibility of ABDPO with traditional gradient-based losses, indicating that ABDPO has a wider scope in actual application.

## F.2 Incorporating Energy Minimization

Energy minimization is indispensable in the standard protein design protocol and is typically applied to the raw co-crystal structure and the generated structure. Most existing AI-based antibody design methods have not undergone similar operations, but to verify the performance of ABDPO in a more realistic workflow environment, we have also proposed another version based on ABDPO+ that integrates energy minimization, ABDPOW/O.

For the minimization of the raw co-crystal structure, we compared the performance of baseline methods trained with and without minimized co-crystal structure but observed no significant difference. A possible reason for this is that most of the methods do not generate the side chain and thus are not sensitive to energy minimization, which mainly optimizes the side-chain conformation. Thus we follow the previous studies, and directly use raw co-crystal structure to train the baseline models and the pre-trained model in ABDPO.

We carry out minimization during the evaluation phase and apply the minimization to the generated antibodies before energy calculation. Therefore, the preference dataset used in ABDPOW/O is built upon the minimized energy. The energy minimization process consists of two parts, **peptide bond length rectification** and **loop refinement**. We first set the length of the peptide bond to 1.3310, the average length of the peptide bonds within CDR-H3 in the SAbDab dataset. Then we use *LoopMover\_Refine\_CCD* from pyRosetta to refine the structure of the designed CDR loop. To reduce time consumption in loop refinement, we set the *outer\_cycles* to 1 and *max\_inner\_cycles* to 10 (a bigger number of cycles will lead to better energy performance undoubtedly, but also makes the time consumption uncontrollable).

Another modification of ABDPOW/O compared to ABDPO+ is that the decomposition of  $\text{Res}_{\text{CDR-Ag}} \Delta G$  into  $\text{Res}_{\text{CDR-Ag}} E_{\text{nonRep}}$  and  $\text{Res}_{\text{CDR-Ag}} E_{\text{Rep}}$  is canceled. Energy decomposition is indispensable in the main experiment because of the huge repulsion, and is not necessary in this experiment as the repulsion would be diminished by the post-minimization process.

Table 7: Summary of CDR  $E_{\text{total}}$ , CDR-Ag  $\Delta G$  (kcal/mol), PHR, and pLL of reference antibodies and antibodies designed by ABDPOW/O and baselines in the experiment involves energy minimization. (↓) / (↑) denotes a smaller / larger number is better.

Methods	CDR $E_{\text{total}}$ (↓)	CDR-Ag $\Delta G$ (↓)	PHR (↓)	pLL (↑)
RAbD	-0.6699	-10.2772	0.4578	-2.2046
HERN	2765.5834	0.8332	41.41%	-2.0409
MEAN	1162.0961	0.0508	<b>30.63%</b>	<b>-1.7936</b>
dyMEAN	611.1203	-2.051	43.73%	-1.8187
DiffAb	82.6216	-0.2734	38.58%	-2.0963
ABDPOW/O	<b>69.8181</b>	<b>-3.0007</b>	36.71%	-2.0251

In Tab. 7, we report the average values of the evaluation metrics for all the generated antibodies in this experiment. Given that the peptide bond length has been rectified, measuring the C-N score is deemed unnecessary in this context. It can be observed that the post-minimization eliminates most of the clashes between the designed antibodies and the corresponding antigens, making CDR-Ag  $\Delta G$  fall within a reasonable range of value. ABDPOW/O still achieves the best performance in the two energy-based metrics, CDR  $E_{\text{total}}$  and CDR-Ag  $\Delta G$ , and surpasses DiffAb in all metrics. This experiment proves (1) the effectiveness of ABDPO in a more realistic setting, and (2) the ability of ABDPO to optimize the energies/properties not directly calculated from the generated antibodies. The values of the two sequence-related metrics, PHR and pLL, for the baseline methods slightly

differ from those in Tab. 1. This discrepancy arises because we imposed a maximum processing time during the loop refinement phase, leading to the exclusion of samples whose refinement was incomplete within the allocated time.

## G Extended Ablation Studies

Due to the massive training cost in the RABD benchmark, we investigate the effectiveness and necessity of each proposed component on five representative antigens, whose PDB IDs are 1a14, 2dd8, 3cx5, 4ki5, and 5mes. From the results in Fig. 8, it is clear that ABDPO can significantly boost the overall performance of ablation cases. Note that moving averages are applied to smooth out the curves to help in identifying trends, including Fig. 4. We present observations and constructive insights of the three proposed components as follows:

1. The residue-level DPO is vital for training stability specifically for CDR  $E_{\text{total}}$ . As aforementioned in Section 3.2, the residue-level DPO implicitly provides fine-grained and rational gradients. In contrast, vanilla DPO (without residue-level DPO) may impose unexpected gradients on stable residues, which incurs the adverse direction of optimization. According to each energy curve in Figure 8, we observe that residue-level DPO surpasses vanilla DPO by at least one energy term.
2. Without Energy Decomposition, all five cases appear undesired “shortcuts” aforementioned in Section 3.3. We observe that the energy of CDR  $E_{\text{total}}$  exhibits a slight performance improvement over the ABDPO after the values of attraction and repulsion reach zero. We suppose that is the result of the combined effects of low attraction and repulsion. Because the generated CDR-H3 is far away from the antigen in this case, the model can concentrate on refining CDR  $E_{\text{total}}$  without the interference of attraction and repulsion.
3. The Gradient Surgery can keep a balance between attraction and repulsion. We can see the curves of  $E_{\text{nonRep}}$  are consistently showing a decline, while the curves of  $E_{\text{Rep}}$  are showing an increase. This observation verifies that ABDPO without Gradient Surgery is unable to optimize  $E_{\text{nonRep}}$  and  $E_{\text{Rep}}$  simultaneously. Additionally, the increase in attraction significantly impacts the repulsion, causing the repulsion to fluctuate markedly.

## H Limitations and Future Work

**Diffusion Process of Orientations** As Luo et al. [36] stated and we have mentioned in Sec. 3.1, Eq. (1) is not a rigorous diffusion process. Thus the loss in Eq. (7) cannot be rigorously derived from the KL-divergence in Eq. (4), though they share the idea of reconstructing the ground truth data by prediction. However, due to the easy implementation and fair comparison with the generative baseline, i.e., DiffAb [36], we adopt Eq. (7) in the ABDPO loss in Eq. (8). In practice, we empirically find that it works well. FrameDiff [50], a protein backbone generation model, adopts a noising process and a rotation loss that are well compatible with the theory of score-based generative models (also known as diffusion models). In the future, we modify the diffusion process of orientations as Yim et al. [50] for potential further improvement.

**Energy Estimation** In this work, we utilize Rosetta/pyRosetta to calculate energy, although it is already one of the most authoritative energy simulation software programs and widely used in protein design and structure prediction, the final energy value is still difficult to perfectly match the actual experimental results. In fact, any computational energy simulation software, whether it is based on force field methods such as OpenMM [14] or statistical methods like the Miyazawa-Jernigan potential [38], will exhibit certain biases and cannot fully simulate reality. Sometimes there is a significant difference between the energy calculated by the software and the results observed experimentally. One possible reason is that theoretical calculations often rely on the designed sequence and structure of antibodies; meanwhile, in actual experiments, the actual folding of the CDR region into the designed structure can be difficult, which leads to significant discrepancies in theoretical calculations. An in vitro experiment is the only way to verify the effectiveness of the designed antibodies. However, due to the significant amount of time consumed by in vitro experiments and considering that the main goal of our work is to propose a novel view of antibody design, we did not perform the in vitro experiment.

**Future Work on Preference Definition** The preferences used in ABDPO determine the tendency of antibody generation, and we will strive to continue exploring the definition of preference to more

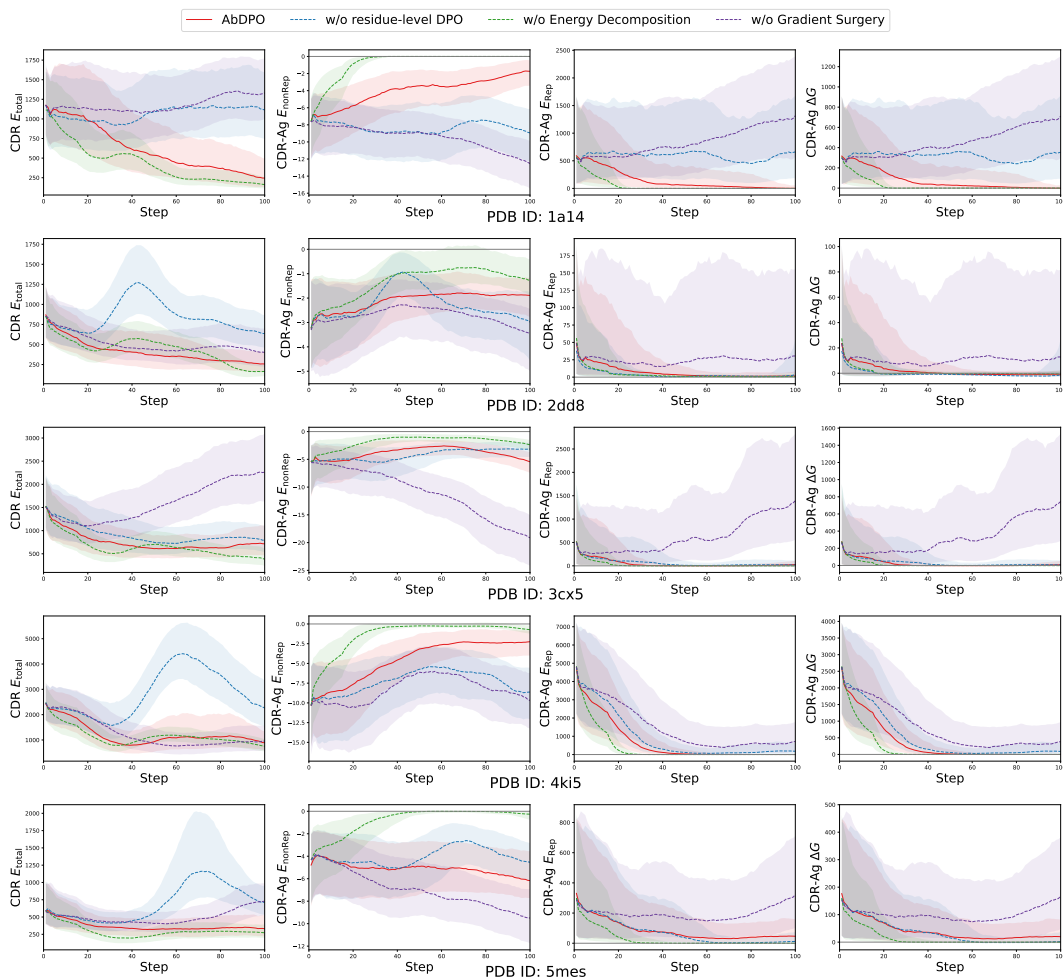


Figure 8: Changes of median CDR  $E_{total}$ , CDR-Ag  $E_{nonRep}$ , CDR-Ag  $E_{Rep}$ , and CDR-Ag  $\Delta G$  (kcal/mol) over-optimization steps, shaded to indicate interquartile range (from 25-th percentile to 75-th percentile). The rows represent PDB 1a14, 2dd8, 3cx5, 4ki5, and 5mes respectively, in a top-down order.

closely align the antibody design process with the real-world environment of antibody activity. Further, we aim to synchronize the preference with the outcomes of in vitro experiments and expect that our method will ultimately generate effective antibodies in real-world applications. The exploration of preference can be divided into two aspects: enhancing existing preferences and integrating new components or energies.

1. The improvement to current preference: (1) performing more fine-grained calculations on the current three types of energy, such as decomposing CDR  $E_{total}$  into interactions between the CDR and the rest of the antibody, interactions within the CDR, and energy at the single amino acid level; (2) exploring the varying importance of preferences for antibodies and determining the relative weights of each preference during the optimization and ranking of generated antibodies.
2. The incorporation of new components or energies is intended to address additional challenges in antibody engineering, focusing on aspects such as antibody stability, solubility, immunogenicity, and expression level. Additionally, we consider integrating components that target antibody specificity.



## **I Potential Societal Impacts**

Our work on antibody design can be used in developing potent therapeutic antibodies and accelerate the research process of drug discovery. The generality of our method extends beyond its current application; it is adaptable for various computer-aided design scenarios including, but not limited to, small molecule, material, and chip design. It is also needed to ensure the responsible use of our method and refrain from using it for harmful purposes.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper discusses the limitations of the work in Appendix H.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper provides the full set of assumptions and a complete (and correct) proof for each theoretical result.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper fully discloses all the information needed to reproduce the main experimental results of the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The release of code requires approval.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specifies all the training and test details in Appendix D.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We lack sufficient computing power to complete such statistical tests.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: This paper provides sufficient information on the computer resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper discusses both positive and negative impacts of the work performed in Appendix I.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The creators or original owners of assets (e.g., code, data, models), used in the paper, are properly credited.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.