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

Anonymous Authors¹

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 (Murphy & Weaver, 2016). The specificity mainly comes from the Complementarity Determining Regions (CDRs), which accounts for most binding affinity to specific antigens (Jones

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.




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 (Kong et al., 2023a) 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 0.6Å.

et al., 1986; Ewert et al., 2004; Xu & Davis, 2000; Akbar et al., 2021). 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 (Adolf-Bryfogle et al., 2018; Lapidoth et al., 2015; Warszawski et al., 2019). Recently, deep generative models have been employed to model protein sequences in nature for antibody design (Alley et al., 2019; Ferruz et al., 2022). Following the fundamental biological principle that structure determines function numerous efforts have been focused on antibody sequence-structure co-design (Jin et al., 2022b;a; Luo et al., 2022; Kong et al., 2023a;b; Martinkus et al., 2023), 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 (Victora & Nussenzweig,

2012; Dong et al., 2021). 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.

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, Rafailov et al., 2023) and self-play fine-tuning techniques (Chen et al., 2024) 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 (Liu et al., 2020; Saka et al., 2021; Akbar et al., 2022). In recent years, sequence-structure co-design of antibodies has attracted increasing attention. Jin et al. (2022b) proposed to simultaneously design sequences and structures of CDRs in an autoregressive way and iteratively refine the designed structures. Jin et al. (2022a) further utilized the epitope and focused on designing CDR-H3 with a hierarchical message passing equivariant network. Kong et al. (2023a) 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. (2022) proposed a diffusion model that takes residue types, atom coordinates and side-chain orientations into consideration to generate antigen-specific CDRs. Kong et al. (2023b) focused on epitope-binding CDR-H3 design and modelled full-atom geometry. Recently, Martinkus et al. (2023) proposed Ab-Diffuser, a novel diffusion model for antibody design, that incorporates more domain knowledge and physics-based constraints and also enables side-chain generation. Besides, Wu & Li (2023); Gao et al. (2023) and Zheng et al. (2023) 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 (Ouyang et al., 2022) and RLAIF (Lee et al., 2023). 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 (Black et al., 2023), DPOK (Fan et al., 2023) and DiffAC (Zhou et al., 2024). 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. (2023) 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 (Wallace et al., 2023; Anonymous, 2023). 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 (Xu & Davis, 2000; Akbar et al., 2021) and the rest part of the antibody including the frameworks and other CDRs. Following Luo et al. (2022), each amino acid is represented by its type $s_i \in \{\text{ACDEFGHIKLMNPQRSTVWY}\}$, C_α coordinate $\mathbf{x}_i \in \mathbb{R}^3$, and frame orientation $\mathbf{O}_i \in \text{SO}(3)$ (Kofinas et al., 2021), 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 $P(\mathcal{R}|\mathcal{P})$.

Denoising Diffusion Probabilistic Model (DDPM, Ho et al., 2020) have been introduced to antibody generation by Luo et al. (2022). This approach consists of a forward *diffusion* process and a reverse *generative* process. The diffusion process gradually injects noises into data as follows:

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

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 . $\mathbf{1}(\cdot)$ is the one-hot operation. $\{\beta^t\}_{t=1}^T$ is the noise schedule for the diffusion process (Ho et al., 2020), 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 $\text{SO}(3)$ (Leach et al., 2022) respectively. ScaleRot scales the rotation angle with fixed rotation axis to modify the rotation matrix (Gallier & Xu, 2003).

Correspondingly, the reverse generative process learns to recover data by iterative denoising. The denoising 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}_{θ_1} , \mathbf{f}_{θ_2} , \mathbf{f}_{θ_3} are parameterized by SE(3)-equivariant neural networks (Jing et al., 2021; Jumper et al., 2021). $\mathbf{f}(\cdot)[j]$ 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:

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

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

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

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

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

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. (2022) 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 (Ouyang et al., 2022), we can fine-tune the pre-trained model to maximize the reward as:

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




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.

where p_θ (resp. p_{ref}) is the distribution induced by the model being fine-tuned (resp. the fixed pre-trained model), β 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. (2023), 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, Bradley & Terry, 1952) 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. (2023), we introduce latent variables $\mathcal{R}^{1:T}$ and utilize the evidence lower bound optimization (ELBO). In particular, L_{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. (2023), 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 $t \sim \mathcal{U}(0, T)$, $(\mathcal{R}_1^{t-1}, \mathcal{R}_1^t)$ 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 (Wallace et al., 2023), 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 (Alford et al., 2017). 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:

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

$$\log \frac{p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t)}{p_{\text{ref}}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}_1^t)} - \log \frac{p_{\theta}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t)}{p_{\text{ref}}(\mathcal{R}_2^{t-1}[j]|\mathcal{R}_2^t)} \Bigg) \Bigg].$$

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 $L_{\text{residue-DPO-Diffusion}}$:

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

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

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

and

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

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

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_{θ} 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 $-\mathbb{D}_{KL}(q||p_{\theta}) + \mathbb{D}_{KL}(q||p_{\text{ref}})$ when taking expectation with respect to q , we can derive the final loss for fine-tuning the diffusion model as follows:

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

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 $\mathbb{D}_{\text{KL}}(q(\mathcal{R}_1^{t-1}[j]|\mathcal{R}^{t-1}, \mathcal{R}^0)||p_{\theta}(\mathcal{R}_1^{t-1}[j]|\mathcal{R}^0))$. 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. (2020), 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 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, Dunbar et al., 2014) under IMGT (Lefranc et al., 2003) 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 MM-seqs2 (Steinegger & Söding, 2017) 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 (Adolf-Bryfogle

et al., 2018) 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 (Chaudhury et al., 2010) 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 (Ruffolo et al., 2021) 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** (Jin et al., 2022a) designs sequences of antibodies autoregressively with the iterative refinement of structures; **MEAN** (Kong et al., 2023a) generates sequences and structures of antibodies via a progress full-shot scheme; **dyMEAN** (Kong et al., 2023b) designs antibodies sequences and structures with full-atom modeling; **DiffAb** (Luo et al., 2022) models antibody distributions with a diffusion model

that considers the amino acid type, C_{α} 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_{α} 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 $\text{Res}_{\text{CDR}} E_{\text{total}}$ of residues within the CDR; (2) CDR-Ag ΔG 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, $E_{\text{generated}} < E_{\text{natural}} + \text{std}(E_{\text{natural}}^{\text{all-complexes}})$.

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 ΔG , 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 ΔG . However, the

Table 1. Summary of AAR, RMSD, CDR E_{total} , CDR-Ag ΔG (kcal/mol), pLL, PHR, and $N_{success}$ of antibodies designed by our model and baselines. (↓) / (↑) denotes a smaller / larger number is better.

Methods	AAR (↑)	RMSD (↓)	CDR E_{total} (↓)	CDR-Ag ΔG (↓)	pLL (↑)	PHR (↓)	$N_{success}$ (↑)
HERN	32.38%	9.18	10887.77	2095.88	-2.02	40.46%	0
MEAN	36.20%	1.69	7162.65	1041.43	-1.79	30.62%	0
dyMEAN	40.04%	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	629.44	307.56	-2.18	69.67%	9
ABDPO+	36.27%	2.01	1106.48	637.62	-2.00	44.21%	5




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.

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 ΔG lower than 0.

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 $Res_{CDR-Ag} E_{nonRep}$ within the designed CDR), E_{Rep} (aggregation of $Res_{CDR-Ag} E_{Rep}$) 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 (Wallace et al., 2023) 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 dashed 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 major




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

ity of the two types of energy. This prevents us from using the ΔG 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 ΔG . 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_{total} (\downarrow)		CDR-Ag ΔG (\downarrow)	
	Avg.	Med.	Avg.	Med.
DiffAb	1314.20	1133.36	534.21	248.28
DiffAb _{SFT}	1053.82	869.37	374.27	144.25
ABDPO	336.02	226.25	88.64	0.10

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.

Impact Statements

Our work on antibody design can be used in developing potent therapeutic antibodies and accelerate the research process of drug discovery. Our method may be adapted to other scenarios of computer-aided design, such as small molecule design, material design, and chip design. It is also needed to ensure the responsible use of our method and refrain from using it for harmful purposes.

References

- Adolf-Bryfogle, J., Kalyuzhnyi, O., Kubitz, M., Weitzner, B. D., Hu, X., Adachi, Y., Schief, W. R., and Dunbrack Jr, R. L. Rosettaantibodydesign (rabd): A general framework for computational antibody design. *PLoS computational biology*, 14(4):e1006112, 2018.
- Akbar, R., Robert, P. A., Pavlović, M., Jeliazkov, J. R., Snapkov, I., Slabodkin, A., Weber, C. R., Scheffer, L., Miho, E., Haff, I. H., Haug, D. T. T., Lund-Johansen, F., Safonova, Y., Sandve, G. K., and Greiff, V. A compact vocabulary of paratope-epitope interactions enables predictability of antibody-antigen binding. *Cell Reports*, 34(11):108856, 2021. ISSN 2211-1247. doi: <https://doi.org/10.1016/j.celrep.2021.108856>. URL <https://www.sciencedirect.com/science/article/pii/S2211124721001704>.
- Akbar, R., Robert, P. A., Weber, C. R., Widrich, M., Frank, R., Pavlović, M., Scheffer, L., Chernigovskaya, M., Snapkov, I., Slabodkin, A., et al. In silico proof of principle of

- 440 machine learning-based antibody design at unconstrained
441 scale. In *MAbs*, volume 14, pp. 2031482. Taylor & Fran-
442 cis, 2022.
- 443
444 Alford, R. F., Leaver-Fay, A., Jeliaskov, J. R., O’Meara,
445 M. J., DiMaio, F. P., Park, H., Shapovalov, M. V., Renfrew,
446 P. D., Mulligan, V. K., Kappel, K., Labonte, J. W., Pacella,
447 M. S., Bonneau, R., Bradley, P., Dunbrack, R. L. J., Das,
448 R., Baker, D., Kuhlman, B., Kortemme, T., and Gray, J. J.
449 The rosetta all-atom energy function for macromolecular
450 modeling and design. *Journal of Chemical Theory and*
451 *Computation*, 13(6):3031–3048, 2017. doi: 10.1021/acs.
452 jctc.7b00125. URL [https://doi.org/10.1021/](https://doi.org/10.1021/acs.jctc.7b00125)
453 [acs.jctc.7b00125](https://doi.org/10.1021/acs.jctc.7b00125). PMID: 28430426.
- 454
455 Alley, E. C., Khimulya, G., Biswas, S., AlQuraishi, M.,
456 and Church, G. M. Unified rational protein engineering
457 with sequence-based deep representation learning. *Nature*
458 *methods*, 16(12):1315–1322, 2019.
- 459
460 Anonymous. Proximal preference optimization for diffusion
461 models, 2023. URL [https://openreview.net/](https://openreview.net/forum?id=u8fg8acFsT)
462 [forum?id=u8fg8acFsT](https://openreview.net/forum?id=u8fg8acFsT).
- 463
464 Black, K., Janner, M., Du, Y., Kostrikov, I., and Levine, S.
465 Training diffusion models with reinforcement learning.
466 *arXiv preprint arXiv:2305.13301*, 2023.
- 467
468 Bradley, R. A. and Terry, M. E. Rank analysis of incom-
469 plete block designs: I. the method of paired comparisons.
470 *Biometrika*, 39(3/4):324–345, 1952.
- 471
472 Chaudhury, S., Lyskov, S., and Gray, J. J. PyRosetta: a
473 script-based interface for implementing molecular mod-
474 eling algorithms using Rosetta. *Bioinformatics*, 26(5):
475 689–691, 01 2010. ISSN 1367-4803. doi: 10.1093/
476 bioinformatics/btq007. URL [https://doi.org/10.](https://doi.org/10.1093/bioinformatics/btq007)
477 [1093/bioinformatics/btq007](https://doi.org/10.1093/bioinformatics/btq007).
- 478
479 Chen, Z., Deng, Y., Yuan, H., Ji, K., and Gu, Q. Self-play
480 fine-tuning converts weak language models to strong lan-
481 guage models. *arXiv preprint arXiv:2401.01335*, 2024.
- 482
483 Crooks, G. E., Hon, G., Chandonia, J.-M., and Brenner, S. E.
484 Weblogo: a sequence logo generator. *Genome research*,
485 14(6):1188–1190, 2004.
- 486
487 Dong, J., Zost, S. J., Greaney, A. J., Starr, T. N., Dingens,
488 A. S., Chen, E. C., Chen, R. E., Case, J. B., Sutton, R. E.,
489 Gilchuk, P., et al. Genetic and structural basis for sars-
490 cov-2 variant neutralization by a two-antibody cocktail.
491 *Nature microbiology*, 6(10):1233–1244, 2021.
- 492
493 Dunbar, J., Krawczyk, K., Leem, J., Baker, T., Fuchs, A.,
494 Georges, G., Shi, J., and Deane, C. M. Sabdab: the
495 structural antibody database. *Nucleic acids research*, 42
496 (D1):D1140–D1146, 2014.
- 497
498 Eastman, P., Swails, J., Chodera, J. D., McGibbon, R. T.,
499 Zhao, Y., Beauchamp, K. A., Wang, L.-P., Simmonett,
500 A. C., Harrigan, M. P., Stern, C. D., et al. Openmm 7:
501 Rapid development of high performance algorithms for
502 molecular dynamics. *PLoS computational biology*, 13(7):
503 e1005659, 2017.
- 504
505 Ewert, S., Honegger, A., and Plückthun, A. Stabil-
506 ity improvement of antibodies for extracellular and
507 intracellular applications: Cdr grafting to stable frame-
508 works and structure-based framework engineering.
509 *Methods*, 34(2):184–199, 2004. ISSN 1046-2023.
510 doi: <https://doi.org/10.1016/j.ymeth.2004.04.007>.
511 URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1046202304000738)
512 [science/article/pii/S1046202304000738](https://www.sciencedirect.com/science/article/pii/S1046202304000738).
513 Intrabodies.
- 514
515 Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier,
516 C., Abbeel, P., Ghavamzadeh, M., Lee, K., and Lee, K.
517 Reinforcement learning for fine-tuning text-to-image dif-
518 fusion models. In *Thirty-seventh Conference on Neural*
519 *Information Processing Systems*, 2023. URL [https://](https://openreview.net/forum?id=80TPepXzeh)
520 openreview.net/forum?id=80TPepXzeh.
- 521
522 Ferruz, N., Schmidt, S., and Höcker, B. Protgpt2 is a deep
523 unsupervised language model for protein design. *Nature*
524 *communications*, 13(1):4348, 2022.
- 525
526 Gallier, J. and Xu, D. Computing exponentials of skew-
527 symmetric matrices and logarithms of orthogonal matri-
528 ces. *International Journal of Robotics and Automation*,
529 18(1):10–20, 2003.
- 530
531 Gao, K., Wu, L., Zhu, J., Peng, T., Xia, Y., He, L., Xie,
532 S., Qin, T., Liu, H., He, K., et al. Pre-training antibody
533 language models for antigen-specific computational anti-
534 body design. In *Proceedings of the 29th ACM SIGKDD*
535 *Conference on Knowledge Discovery and Data Mining*,
536 pp. 506–517, 2023.
- 537
538 Ho, J., Jain, A., and Abbeel, P. Denoising diffusion
539 probabilistic models. In Larochelle, H., Ranzato,
540 M., Hadsell, R., Balcan, M., and Lin, H. (eds.),
541 *Advances in Neural Information Processing Systems*,
542 volume 33, pp. 6840–6851. Curran Associates, Inc.,
543 2020. URL [https://proceedings.neurips.](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf)
544 [cc/paper_files/paper/2020/file/](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf)
545 [4c5bcfec8584af0d967f1ab10179ca4b-Paper.](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf)
546 [pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf).
- 547
548 Jin, W., Barzilay, R., and Jaakkola, T. Antibody-antigen
549 docking and design via hierarchical structure refinement.
550 In *International Conference on Machine Learning*, pp.
551 10217–10227. PMLR, 2022a.
- 552
553 Jin, W., Wohlwend, J., Barzilay, R., and Jaakkola, T. S.
554 Iterative refinement graph neural network for antibody

- sequence-structure co-design. In *International Conference on Learning Representations*, 2022b. URL https://openreview.net/forum?id=LI2bhrE_2A.
- Jing, B., Eismann, S., Suriana, P., Townshend, R. J. L., and Dror, R. Learning from protein structure with geometric vector perceptrons. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=1YLJDvSx6J4>.
- Jones, P. T., Dear, P. H., Foote, J., Neuberger, M. S., and Winter, G. Replacing the complementarity-determining regions in a human antibody with those from a mouse. *Nature*, 321(6069):522–525, 1986.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnoy, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- Katoh, K. and Standley, D. M. Mafft multiple sequence alignment software version 7: improvements in performance and usability. *Molecular biology and evolution*, 30(4):772–780, 2013.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kofinas, M., Nagaraja, N. S., and Gavves, E. Roto-translated local coordinate frames for interacting dynamical systems. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=c3RKZas9am>.
- Kong, X., Huang, W., and Liu, Y. Conditional antibody design as 3d equivariant graph translation. In *The Eleventh International Conference on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=LFHFQbjxIiP>.
- Kong, X., Huang, W., and Liu, Y. End-to-end full-atom antibody design. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 17409–17429. PMLR, 23–29 Jul 2023b.
- Lapidoth, G. D., Baran, D., Pszolla, G. M., Norn, C., Alon, A., Tyka, M. D., and Fleishman, S. J. Abdesign: A n algorithm for combinatorial backbone design guided by natural conformations and sequences. *Proteins: Structure, Function, and Bioinformatics*, 83(8):1385–1406, 2015.
- Leach, A., Schmon, S. M., Degiacomi, M. T., and Willcocks, C. G. Denoising diffusion probabilistic models on so (3) for rotational alignment. In *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022.
- Lee, H., Phatale, S., Mansoor, H., Lu, K., Mesnard, T., Bishop, C., Carbune, V., and Rastogi, A. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023.
- Lefranc, M.-P., Pommié, C., Ruiz, M., Giudicelli, V., Foulquier, E., Truong, L., Thouvenin-Contet, V., and Lefranc, G. Imgt unique numbering for immunoglobulin and t cell receptor variable domains and ig superfamily v-like domains. *Developmental & Comparative Immunology*, 27(1):55–77, 2003.
- Liu, G., Zeng, H., Mueller, J., Carter, B., Wang, Z., Schilz, J., Horny, G., Birnbaum, M. E., Ewert, S., and Gifford, D. K. Antibody complementarity determining region design using high-capacity machine learning. *Bioinformatics*, 36(7):2126–2133, 2020.
- Luo, S., Su, Y., Peng, X., Wang, S., Peng, J., and Ma, J. Antigen-specific antibody design and optimization with diffusion-based generative models for protein structures. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=jSorGn2Tjg>.
- Martinkus, K., Ludwiczak, J., LIANG, W.-C., Lafrance-Vanasse, J., Hotzel, I., Rajpal, A., Wu, Y., Cho, K., Bonneau, R., Gligorijevic, V., and Loukas, A. Abdifuser: full-atom generation of in-vitro functioning antibodies. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=7GyYpomkEa>.
- Miyazawa, S. and Jernigan, R. L. Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation. *Macromolecules*, 18(3):534–552, 1985.
- Murphy, K. and Weaver, C. *Janeway’s immunobiology*. Garland science, 2016.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Gray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., and Lowe, R. Training language models to follow instructions with human feedback. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=TG8KACxEON>.

- 550 Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D.,
551 Ermon, S., and Finn, C. Direct preference optimization:
552 Your language model is secretly a reward model. In *Thirty-*
553 *seventh Conference on Neural Information Processing*
554 *Systems*, 2023. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2305.18290)
555 [2305.18290](https://arxiv.org/abs/2305.18290).
- 556 Ruffolo, J. A., Gray, J. J., and Sulam, J. Deciphering
557 antibody affinity maturation with language models
558 and weakly supervised learning. *arXiv preprint*
559 *arXiv:2112.07782*, 2021.
- 560 Saka, K., Kakuzaki, T., Metsugi, S., Kashiwagi, D., Yoshida,
561 K., Wada, M., Tsunoda, H., and Teramoto, R. Anti-
562 body design using lstm based deep generative model from
563 phage display library for affinity maturation. *Scientific*
564 *reports*, 11(1):5852, 2021.
- 565 Steinegger, M. and Söding, J. Mmseqs2 enables sensi-
566 tive protein sequence searching for the analysis of mas-
567 sive data sets. *Nature biotechnology*, 35(11):1026–1028,
568 2017.
- 569 Victora, G. D. and Nussenzweig, M. C. Germinal centers.
570 *Annual review of immunology*, 30:429–457, 2012.
- 571 Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Pu-
572 rushwalkam, S., Ermon, S., Xiong, C., Joty, S., and Naik,
573 N. Diffusion model alignment using direct preference
574 optimization. *arXiv preprint arXiv:2311.12908*, 2023.
- 575 Warszawski, S., Borenstein Katz, A., Lipsh, R., Khmelnit-
576 sky, L., Ben Nissan, G., Javitt, G., Dym, O., Unger, T.,
577 Knop, O., Albeck, S., et al. Optimizing antibody affinity
578 and stability by the automated design of the variable light-
579 heavy chain interfaces. *PLoS computational biology*, 15
580 (8):e1007207, 2019.
- 581 Wu, F. and Li, S. Z. A hierarchical training paradigm for
582 antibody structure-sequence co-design. In *Thirty-seventh*
583 *Conference on Neural Information Processing Systems*,
584 2023. URL [https://openreview.net/forum?](https://openreview.net/forum?id=hV52oj0Sik)
585 [id=hV52oj0Sik](https://openreview.net/forum?id=hV52oj0Sik).
- 586 Xu, J. L. and Davis, M. M. Diversity in the cdr3 region
587 of vh is sufficient for most antibody specificities.
588 *Immunity*, 13(1):37–45, 2000. ISSN 1074-7613.
589 doi: [https://doi.org/10.1016/S1074-7613\(00\)00006-6](https://doi.org/10.1016/S1074-7613(00)00006-6).
590 URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1074761300000066)
591 [science/article/pii/S1074761300000066](https://www.sciencedirect.com/science/article/pii/S1074761300000066).
- 592 Yim, J., Trippe, B. L., De Bortoli, V., Mathieu, E., Doucet,
593 A., Barzilay, R., and Jaakkola, T. SE(3) diffusion model
594 with application to protein backbone generation. In
595 Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato,
596 S., and Scarlett, J. (eds.), *Proceedings of the 40th Interna-*
597 *tional Conference on Machine Learning*, volume 202 of
598 *Proceedings of Machine Learning Research*, pp. 40001–
599 40039. PMLR, 23–29 Jul 2023.
- 600 Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K.,
601 and Finn, C. Gradient surgery for multi-task learning.
602 *Advances in Neural Information Processing Systems*, 33:
603 5824–5836, 2020.
- 604 Zheng, Z., Deng, Y., Xue, D., Zhou, Y., YE, F., and Gu, Q.
Structure-informed language models are protein design-
ers. In *International Conference on Machine Learning*,
2023.
- Zhou, X., Wang, L., and Zhou, Y. Stabilizing policy
gradients for stochastic differential equations via
consistency with perturbation process. *arXiv preprint*
arXiv:2403.04154, 2024.

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 (Kato & Standley, 2013) and subsequently visualized with WebLogo (Crooks et al., 2004). 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 ES_{total} is the sum of all types of energy with the default weight in REF15 (Alford et al., 2017). The function for interaction energies between **Paired** residues is described as EP, which consists of six different energy types: EP_{hbond} , EP_{att} , EP_{rep} , EP_{sol} , EP_{elec} , and EP_{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_{nonRep} exhibits a benefit in interactions, while E_{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_{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. (2023), 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.,

715 $\max_{p_\theta} \mathbb{E}_{\mathcal{R}^0 \sim p_\theta} [r(\mathcal{R}^0)] - \beta \mathbb{D}_{\text{KL}}(p_\theta(\mathcal{R}^0) \| p_{\text{ref}}(\mathcal{R}^0))$ as follows:

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

720 where Z is the partition function that does not involve the model being trained, i.e., p_θ . And we can define

$$721 \quad p^*(\mathcal{R}^0) := \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right).$$

722 With this, we can now arrive at

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

725 Since Z does not depend on p_θ , 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:

$$726 \quad p_{\theta^*}(\mathcal{R}^0) = p^*(\mathcal{R}^0) = \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right). \quad (13)$$

727 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_{DPO} , leads to the same optimal solution. For this, we need to express the pre-defined reward $r(\cdot)$ with the optimal policy p^* :

$$728 \quad r(\mathcal{R}^0) = \beta \log \frac{p^*(\mathcal{R}^0)}{p_{\text{ref}}(\mathcal{R}^0)} + Z$$

729 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:

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

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

$$733 \quad p_{\theta}(\mathcal{R}_1^0 \succ \mathcal{R}_2^0) = \sigma\left(\beta \log \frac{p_\theta(\mathcal{R}_1^0)}{p_{\text{ref}}(\mathcal{R}_1^0)} - \beta \log \frac{p_\theta(\mathcal{R}_2^0)}{p_{\text{ref}}(\mathcal{R}_2^0)}\right).$$

734 With this, we have

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

738 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 $p^*(\mathcal{R}^0) = \frac{1}{Z} p_{\text{ref}}(\mathcal{R}^0) \exp\left(\frac{1}{\beta} r(\mathcal{R}^0)\right)$ is also the optimal solution of L_{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. (2022). 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_α coordinates, and a $so(3)$ vector for the rotation matrix.

The number of diffusion steps is 100. We use the cosine β schedule with $s = 0.01$ suggested in Ho et al. (2020) for amino acid types, C_α coordinates, and orientations.

D.2 Training Details

Pre-training. Following Luo et al. (2022), the diffusion model is first trained via the gradient descent method Adam (Kingma & Ba, 2014) 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`¹, 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+ 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_{total} and the CDR-Ag ΔG of each antibody are ranked independently. Then, a composite ranking score for each antibody is defined as the sum of its CDR E_{total} rank and CDR-Ag ΔG 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.

¹https://github.com/THUNLP-MT/MEAN/blob/main/summaries/rabd_summary.jsonl

- 825 2. The distribution patterns of different energy types and properties can vary, with these distributions usually being
826 non-uniform. This could result in scenarios where minor numerical differences in the top-ranking CDR-Ag ΔG values
827 coincide with larger differences in CDR E_{total} , potentially leading to the selection of samples with suboptimal CDR
828 E_{total} .
829

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

- 832 1. The primary goal of this work is to reformulate the antibody design task as an energy-focused optimization problem
833 and propose a feasible implementation, rather than to delve into the mechanisms of antibody-antigen binding;
834
835 2. Our approach is designed to avoid introducing statistical biases or preferences based on potentially erroneous prior
836 knowledge or favoritism towards particular antibody design methods.
837

838 E More Evaluation Results

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

840 In Tab. 1, we have reported the average results of all antibodies designed by our method and other baselines. Here we
841 provide the evaluation results for the ranked top-1 design in Tab. 3 (refer to the ranking strategy in Appendix D.3).
842

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

844 Methods	CDR E_{total} (\downarrow)	CDR-Ag ΔG (\downarrow)	PHR (\downarrow)	pLL (\uparrow)	AAR (\uparrow)	RMSD (\downarrow)
845 RAbD	5.25	-13.04	45.78%	-2.20	100.00%	0.00
846 HERN	8495.56	1296.22	48.18%	-2.01	33.29%	9.21
847 MEAN	3867.47	207.99	36.91%	-1.72	35.18%	1.70
848 dyMEAN	2987.93	1283.97	46.27%	-1.79	40.74%	1.81
849 DiffAb	381.82	58.84	49.19%	-2.03	37.99%	1.62
850 ABDPO	68.51	-4.96	69.97%	-2.15	32.92%	1.58
851 ABDPO+	332.10	29.27	32.81%	-1.54	39.55%	1.67

852 E.2 Detailed Evaluation Results for each Complex

853 In Tab. 4 and Tab. 5, we list the CDR E_{total} , CDR-Ag ΔG , PHR and pLL of the reference antibody in RAbD and the
854 average/ranked top-1 antibodies designed by HERN, MEAN, dyMEAN, DiffAb, ABDPO, and ABDPO+ for each complex
855 in the test set separately. In Tab. 5, we highlight the energy values of the designed complexes that surpass the natural one in
856 terms of two energies simultaneously with **bold text**.
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879




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 ($t \leq 15$ 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 $(s_j^0, x_j^0, \mathbf{O}_j^0)$ which is denoised with one step from $(s_j^t, x_j^t, \mathbf{O}_j^t)$, 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_{total} , CDR-Ag ΔG (kcal/mol), pLL, PHR, C- N_{score} , AAR, and RMSD of reference antibodies and antibodies designed by ABDPOW/O and baselines in the experiment involves auxiliary loss. (\downarrow) / (\uparrow) denotes a smaller / larger number is better.

Methods	CDR E_{total} (\downarrow)	CDR-Ag ΔG (\downarrow)	pLL (\uparrow)	PHR (\downarrow)	C- N_{score} (\uparrow)	AAR (\uparrow)	RMSD (\downarrow)
HERN	10887.77	2095.88	-2.02	40.46%	0.12	32.38%	9.18
MEAN	7162.65	1041.43	-1.79	36.20%	1.68	36.30%	1.69
dyMEAN	3782.67	1730.06	-1.82	43.72%	2.08	40.04%	1.82
DiffAb	1729.51	1297.25	-2.10	41.27%	3.85	34.92%	1.92
ABDPO	629.44	307.56	-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%	4.51	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_{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_{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 Res_{CDR-Ag} ΔG into Res_{CDR-Ag} E_{nonRep} and Res_{CDR-Ag} E_{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_{total} , CDR-Ag ΔG (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_{total} (↓)	CDR-Ag ΔG (↓)	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	30.63%	-1.7936
dyMEAN	611.1203	-2.051	43.73%	-1.8187
DiffAb	82.6216	-0.2734	38.58%	-2.0963
ABDPOW/O	69.8181	-3.0007	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 ΔG fall within a reasonable range of value. ABDPOW/O still achieves the best performance in the two energy-based metrics, CDR E_{total} and CDR-Ag ΔG , 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_{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_{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_{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_{nonRep} are consistently showing a decline, while the curves of E_{Rep} are showing an increase. This observation verifies that ABDPO without Gradient Surgery is unable to optimize E_{nonRep} and E_{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. (2022) 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 (Luo et al., 2022), we adopt Eq. (7) in the ABDPO loss in Eq. (8). In practice, we empirically find that it works well. FrameDiff (Yim et al., 2023), 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. (2023) 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 (Eastman et al., 2017) or statistical methods like the Miyazawa-Jernigan potential (Miyazawa & Jernigan, 1985), 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 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.

1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181

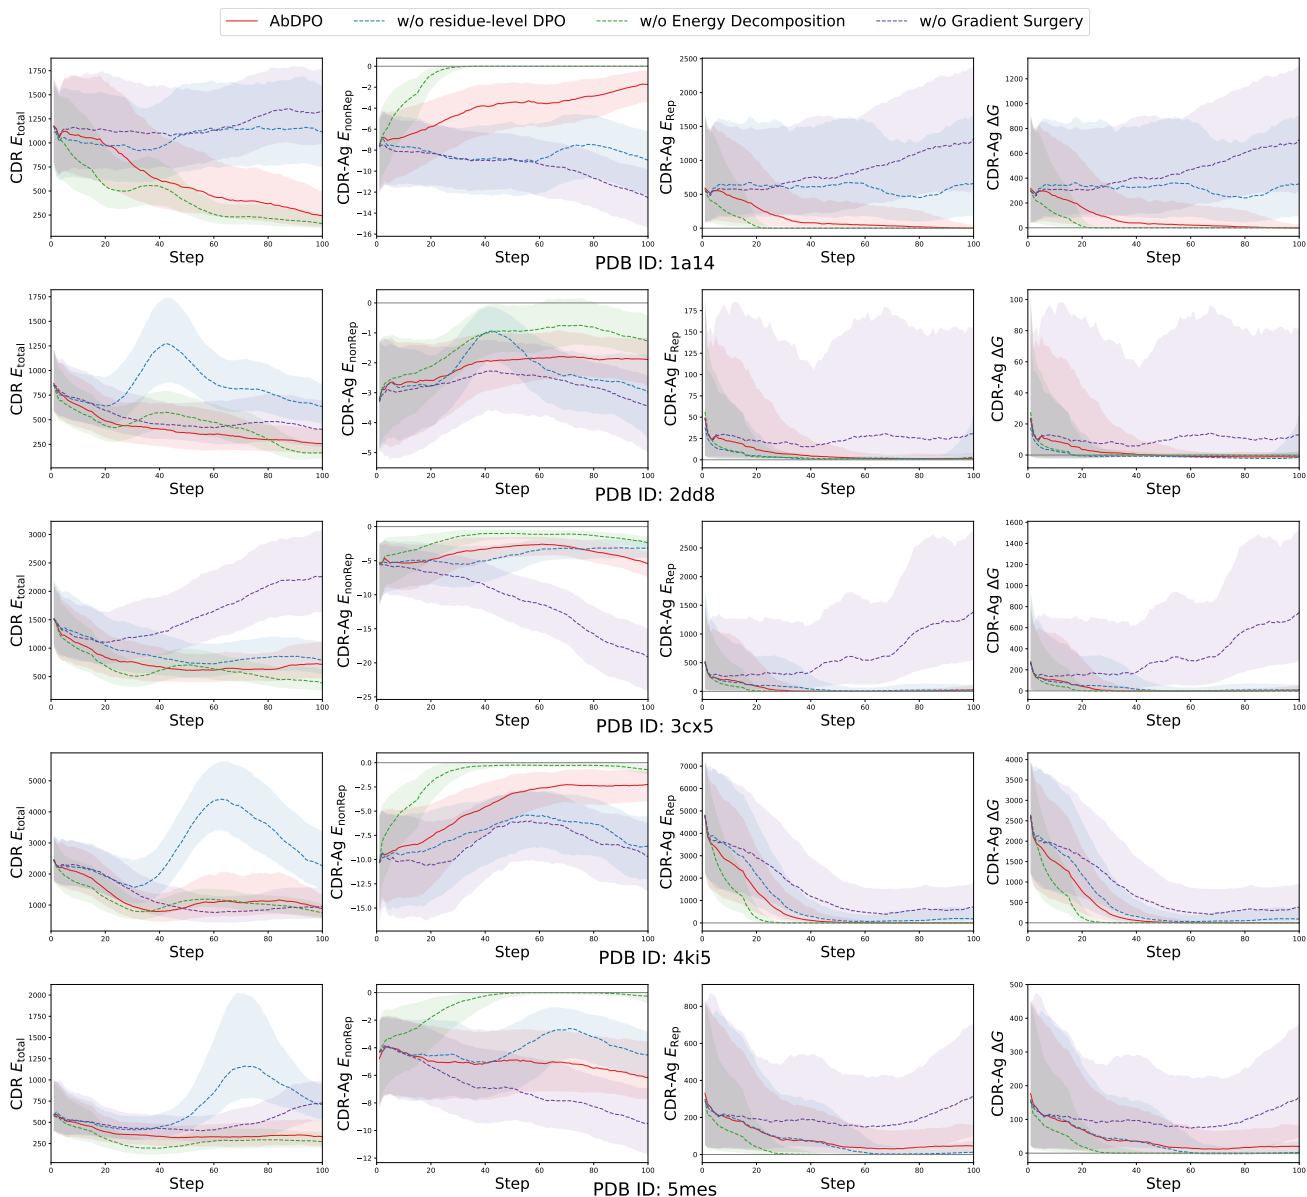


Figure 8. Changes of median CDR E_{total} , CDR-Ag E_{nonRep} , CDR-Ag E_{Rep} , and CDR-Ag ΔG (kcal/mol) over-optimization steps, shaded to indicate interquartile range (from 25-th percentile to 75-th percentile). The rows represent PDB 1a14, 2dd8, 3cx5, 4ki5, and 5mes respectively, in a top-down order.