

Generative Co-Design of Antibody Sequences and Structures via Black-Box Guidance in a Shared Latent Space

Yinghua Yao^{1,2}, Yuangang Pan^{1,2*}, Xixian Chen³

¹Center for Frontier AI Research, Agency for Science, Technology and Research, Singapore

²Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore

³Singapore Institute of Food and Biotechnology Innovation, Agency for Science, Technology and Research, Singapore

eva.yh.yao, yuangang.pan@gmail.com, xixian_chen@sifbi.a-star.edu.sg

Abstract

Advancements in deep generative models have enabled the joint modeling of antibody sequence and structure, given the antigen-antibody complex as context. However, existing approaches for optimizing complementarity-determining regions (CDRs) to improve developability properties operate in the raw data space, leading to excessively costly evaluations due to the inefficient search process. To address this, we propose LatEnt blAck-box Design (LEAD), a sequence-structure co-design framework that optimizes both sequence and structure within their shared latent space. Optimizing shared latent codes can not only break through the limitations of existing methods, but also ensure synchronization of different modality designs. Particularly, we design a black-box guidance strategy to accommodate real-world scenarios where many property evaluators are non-differentiable. Experimental results demonstrate that our LEAD achieves superior optimization performance for both single and multi-property objectives. Notably, LEAD reduces query consumption by a half while surpassing baseline methods in property optimization. The code is available at <https://github.com/EvaFlower/LatEnt-blAck-box-Design>.

1 Introduction

Antibodies are glycoproteins produced by the immune system that specifically bind to antigens—foreign molecules that trigger immune responses. Antigen recognition and binding arise primarily from the complementarity-determining regions (CDRs) of antibodies [Janeway *et al.*, 2001; Presta, 1992]. As such, the design of CDRs is a pivotal step in developing effective therapeutic antibodies and plays a key role in advancing drug discovery.

The space of possible antibodies is vast and discrete, and experimental validation is slow and costly [Larman *et al.*, 2012]. This drives a significant development of *in silico* antibody design approaches to help build a small, enriched library

of candidates to identify viable options with minimal experimental measurements. Traditional computational methods sample antibody sequences/structures based on hand-crafted and statistical energy functions [Lapidoth *et al.*, 2015; Adolf-Bryfogle *et al.*, 2018; Warszawski *et al.*, 2019], which are inefficient and easily trapped in local optima. Recent advancements in diffusion-based deep generative models enable the simultaneous distribution modeling of both the sequence and structure of CDRs [Jin *et al.*, 2022; Luo *et al.*, 2022; Kong *et al.*, 2023], thereby improving design efficiency. These models have the advantage of incorporating both antigen epitopes and antibody framework structures during the design process, which is essential to enhance binding affinity.

The clinical effectiveness of antibody-based therapeutics relies not only on their biological performance—such as target specificity and potency—but also on their biophysical and chemical characteristics, collectively referred to as “developability”. Critical aspects of developability, including solubility, thermal stability, aggregation tendency, and immunogenicity, must be carefully optimized to guarantee both manufacturability and safety in clinical applications [Villegas-Morcillo *et al.*, 2024; Sumida *et al.*, 2024]. While most efforts primarily focus on enhancing antigen-binding properties, some recent works [Villegas-Morcillo *et al.*, 2024; Kulytè *et al.*, 2024; Zhou *et al.*, 2024] incorporate developability into antibody design by guiding either the sequence or structure in the raw data space, or fine-tuning the model, which usually requires extensive property evaluations.

In this work, we employ deep generative models that learn a joint antibody sequence and structure distribution to design antibody CDRs *de novo*, optimizing not only for antigen affinity but also for enhanced developability. Following the principle that sequence determines structure, and structure determines function [Branden and Tooze, 2012], we propose a novel LatEnt blAck-box Design (LEAD) framework (Fig. 1), which guides property optimization through the shared latent space of sequences and structures. Specifically, to address real-world scenarios where many property evaluators are not differentiable, we design our guidance strategy based on black-box techniques, allowing flexible and general-purpose optimization across diverse antibody design tasks. The contributions are summarized as follows:

- We propose LatEnt blAck-box Design (LEAD), a

*Corresponding author

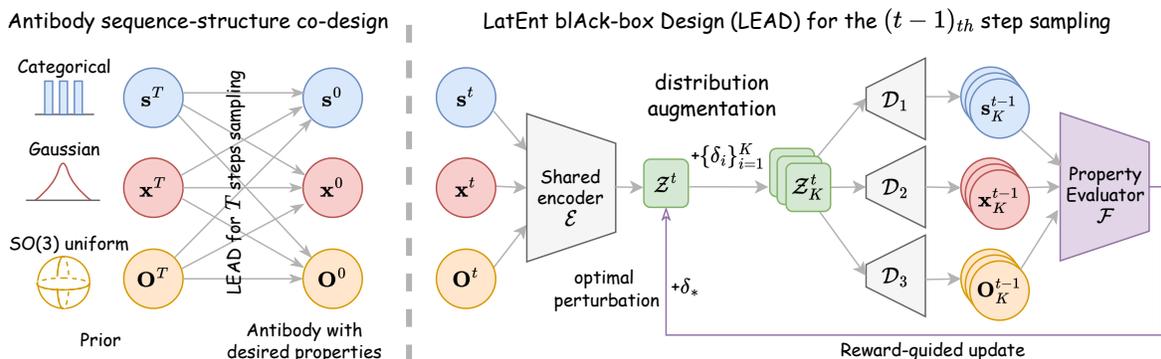


Figure 1: Overall framework. **Left panel** describes the whole process of our guided sampling. **Right panel** details the black-box guidance incorporated in the shared latent space of sequence and structure.

method that jointly optimizes sequence and structure within the shared latent space of a diffusion model for antigen-specific antibody design. LEAD establishes a general framework, and different variants of LEAD are proposed for single and multiple property optimization.

- LEAD enables property guidance by incorporating black-box optimization into the denoising process, without any additional training of diffusion models.
- Experimental results demonstrate that LEAD achieves superior optimization efficiency compared to existing approaches that perform guidance in the raw antibody space. Specifically, it requires fewer property evaluations while maintaining or exceeding performance in both single and multiple-property optimization.

2 Related Work

Antibody Design Deep learning methods have achieved great success in antibody design. Large Language Model (LLM)-based approaches are efficient in generating protein sequences [Saka *et al.*, 2021; Shuai *et al.*, 2023], but they struggle to design antibodies for specific antigens since protein function is usually shaped by a complex interplay between sequence and structural features. The first CDR sequence-structure co-design model is introduced for designing antibodies against SARS-CoV-2 [Jin *et al.*, 2022], but its reliance on an antigen-specific predictor limits its generalizability to other antigens. Recently, diffusion models have shown significant potential in both continuous and categorical data domains and have been successfully applied to generating high-quality protein sequences or protein backbones [Watson *et al.*, 2023; Ingraham *et al.*, 2023; Lisanza *et al.*, 2024; Gruver *et al.*, 2024]. DiffAb [Luo *et al.*, 2022] proposed the first diffusion model to perform joint design of sequence and structure of the antibody CDR regions while conditioning on the antigen-antibody complex. AbDifuser [Martinkus *et al.*, 2024] improved this further by incorporating strong priors and being more memory efficient with side chain generation. DiffAbXL [Ucar *et al.*, 2024b] is an advanced version of DiffAb, with an improved architecture.

Property Optimization Sequence-structure co-design methods enhance the antigen-targeting performance of

antibodies. Recent studies further incorporate property optimization to improve their developability. DIFFFORCE [Kulytė *et al.*, 2024] introduces a differentiable force field w.r.t. the 3D atom coordinates to guide the sampling process. ABDPO [Zhou *et al.*, 2024] fine-tunes DiffAb via direct preference optimization to optimize energies of generated antibodies. [Villegas-Morcillo *et al.*, 2024] performs guidance in the raw space of antibodies according to developability properties during diffusion sampling. However, these methods require extensive property evaluations, limiting their scalability and efficiency.

3 Black-box Antibody Design in the Shared Latent Space of Diffusion Model

In this section, we first introduce the background of diffusion model for antibody sequence and structure co-design. Then we develop our black-box guidance strategy in the shared latent space of this co-design model.

3.1 Diffusion Model for Antibody Co-design

Our work builds upon the diffusion model for antibody sequence and structure co-design. It enables the joint generation of Complementarity-Determining Region (CDR) sequences and structures conditioning on the antibody framework and its corresponding bound antigen. Specifically, each amino acid is represented by its type $s_i \in \{\text{ACDEFGHIKLMNPQRSTVWY}\}$, C_α atom coordinate $x_i \in \mathbb{R}^3$, and amino acid orientation $O_i \in \text{SO}(3)$, where $i = 1, 2, \dots, N$ and N is the total number of the amino acids in the antibody-antigen complex.

The target CDR loop is denoted as $\mathcal{A} = \{(s_i, x_i, O_i) \mid i = 1, 2, \dots, m\}$, given the rest of the antibody-antigen complex $\mathcal{R} = \{(s_j, x_j, O_j) \mid j \neq i, j \in \{1, 2, \dots, N\}\}$. For the sake of notation, let $\mathbf{s} = \{s_i\}_{i=1}^m$, $\mathbf{x} = \{x_i\}_{i=1}^m$, $\mathbf{O} = \{O_i\}_{i=1}^m$ denote the whole feature of each modality of the target CDR loop, respectively. Then, we have $\mathcal{A} = (\mathbf{s}, \mathbf{x}, \mathbf{O})$. The antibody generation task can be then formulated as modeling the conditional distribution $P(\mathcal{A}|\mathcal{R})$. DiffAb [Luo *et al.*, 2022] represents the first application of Denoising Diffusion Probabilistic Model (DDPM) [Ho *et al.*, 2020] to antibody generation, which consists of a forward diffusion process and a reverse generative process.

Forward diffusion process ($t = 1, \dots, T$) The antibody sequence-structure co-design framework initially assumes the three modalities are conditionally independent. Namely,

$$q(\mathcal{A}^t | \mathcal{A}^0) = q(\mathbf{s}^t | \mathbf{s}^0) q(\mathbf{x}^t | \mathbf{x}^0) q(\mathbf{O}^t | \mathbf{O}^0), \quad (1)$$

where $\mathcal{A}^0 = (\mathbf{s}^0 = \{s_i^0\}_{i=1}^m, \mathbf{x}^0 = \{x_i^0\}_{i=1}^m, \mathbf{O}^0 = \{O_i^0\}_{i=1}^m)$ is the noisy-free CDR loop at time step 0. Similarly, $\mathcal{A}^t = (\mathbf{s}^t, \mathbf{x}^t, \mathbf{O}^t)$ is the noisy CDR loop at time step t .

Then, the forward diffusion process gradually introduces noise into each modality through different distributions towards the prior distributions. Namely,

$$q(s_i^t | s_i^0) = \mathcal{C}(\mathbb{I}(s_i^t) | \bar{\alpha}^t \mathbb{I}(s_i^0) + \bar{\beta}^t \mathbf{1}/20), \quad (2a)$$

$$q(x_i^t | x_i^0) = \mathcal{N}(x_i^t | \sqrt{\alpha^t} x_i^0, \bar{\beta}^t \mathbf{I}), \quad (2b)$$

$$q(O_i^t | O_i^0) = \mathcal{IG}_{\text{SO}(3)}(O_i^t | \text{ScaleRot}(\sqrt{\alpha^t}, O_i^0), \bar{\beta}^t), \quad (2c)$$

where $i = 1, 2, \dots, m$ is the index for the amino acid. $\mathbb{I}(\cdot)$ is the one-hot operation that converts each amino acid type to a 20-dimensional one-hot vector and $\mathbf{1}$ is an all-one vector. $\{\beta^t\}_{t=1}^T$ is the noise schedule for the diffusion process, and we define $\bar{\alpha}^t = \prod_{\tau=1}^t (1 - \beta^\tau)$ and $\bar{\beta}^t = 1 - \bar{\alpha}^t$. Here, $\mathcal{C}(\cdot), \mathcal{N}(\cdot)$, and $\mathcal{IG}_{\text{SO}(3)}(\cdot)$ are the 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 a fixed rotation axis to modify the rotation matrix [Gallier and Xu, 2003].

Reverse generative process ($t = T, T - 1, \dots, 1$) The reverse generative diffusion process learns to recover data by iteratively denoising from the prior distributions. The parametric models $p_\theta(\mathcal{A}^{t-1} | \mathcal{A}^t, \mathcal{R})$ are employed to approximate the posterior distributions at each generation time step. To model the interplay between sequences and structures, a three-output neural network for the three modalities with a shared encoder and separate decoders [Luo *et al.*, 2022] is adopted, which takes the whole sequence-structure context \mathcal{R} and the noisy antibody at the previous denoising step $\mathcal{A}^t = (\mathbf{s}^t, \mathbf{x}^t, \mathbf{O}^t)$ as input and predicts $\mathcal{A}^{t-1} = (\mathbf{s}^{t-1}, \mathbf{x}^{t-1}, \mathbf{O}^{t-1})$. Namely,

$$p_\theta(\mathcal{A}^{t-1} | \mathcal{A}^t, \mathcal{R}) = p_\theta(\mathbf{s}^{t-1} | \mathcal{A}^t, \mathcal{R}) p_\theta(\mathbf{x}^{t-1} | \mathcal{A}^t, \mathcal{R}) \times p_\theta(\mathbf{O}^{t-1} | \mathcal{A}^t, \mathcal{R}). \quad (3)$$

In particular, the denoising process of three modalities from time step t to time step $t - 1$ can be defined as follows:

$$p_\theta(s_i^{t-1} | \mathcal{A}^t, \mathcal{R}) = \mathcal{C}(\mathbb{I}(s_i^{t-1}) | f_{1,i}(\mathcal{A}^t, \mathcal{R})), \quad (4a)$$

$$p_\theta(x_i^{t-1} | \mathcal{A}^t, \mathcal{R}) = \mathcal{N}(x_i^{t-1} | f_{2,i}(\mathcal{A}^t, \mathcal{R}), \beta^t), \quad (4b)$$

$$p_\theta(O_i^{t-1} | \mathcal{A}^t, \mathcal{R}) = \mathcal{IG}_{\text{SO}(3)}(O_i^{t-1} | f_{3,i}(\mathcal{A}^t, \mathcal{R}), \beta^t), \quad (4c)$$

where $i = 1, 2, \dots, m$ is the index for the amino acid. f_1, f_2, f_3 are parameterized by SE(3)-equivariant neural networks, whose i -th entry denotes the output for the i -th amino acid, respectively.

3.2 Function-guided Latent Optimization

With the pre-trained joint diffusion model defined in Section 3.1, we explore how to incorporate the functional guidance $\mathcal{F}(\mathcal{A}^t)$ (without loss of generality, we assume higher

function reward is better), defined with respect to sequence and structure, into the denoising sampling process to generate antibodies with targeted properties.

To enable more efficient functional guidance, we no longer implement guidance in the raw data space. Instead, we consider the shared latent space of sequences and structures.

First of all, we define the shared latent space of multiple modalities. According to the definition of the three-output denoising neural network, we reformulate it as follows:

$$f_1(\mathcal{A}^t, \mathcal{R}) = \mathcal{D}_1(\mathcal{E}(\mathcal{A}^t, \mathcal{R})) = \mathcal{D}_1(\mathcal{Z}^t), \quad (5a)$$

$$f_2(\mathcal{A}^t, \mathcal{R}) = \mathcal{D}_2(\mathcal{E}(\mathcal{A}^t, \mathcal{R})) = \mathcal{D}_2(\mathcal{Z}^t), \quad (5b)$$

$$f_3(\mathcal{A}^t, \mathcal{R}) = \mathcal{D}_3(\mathcal{E}(\mathcal{A}^t, \mathcal{R})) = \mathcal{D}_3(\mathcal{Z}^t), \quad (5c)$$

where $\mathcal{E}(\cdot)$ denotes the shared encoder while $\mathcal{D}_*(\cdot)$ denotes the modality-specific decoder. Denoting $\mathcal{E}(\mathcal{A}^t, \mathcal{R})$ as \mathcal{Z}^t , we consider the space where \mathcal{Z}^t resides as the shared latent space of multiple modalities.

To avoid extra noise distraction, we adopt the deterministic version of the reverse generative process, a.k.a. Denoising Diffusion Implicit Models (DDIM) [Song *et al.*, 2020]. For simplicity of notation, let DDIM also denote a one-step reverse generative process consisting of a three-output denoising network with a shared encoder ($\mathcal{E}, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$).

Definition 1. [DDIM for deterministic sampling] Given the noisy antibody \mathcal{A}^t and the rest antigen-antibody context \mathcal{R} , the subsequent denoised antibody \mathcal{A}^{t-1} can be obtained deterministically via DDIM, i.e., $\mathcal{A}^{t-1} \sim \text{DDIM}(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{Z}^t, \beta_t)$ where $\mathcal{Z}^t = \mathcal{E}(\mathcal{A}^t, \mathcal{R})$ and β_t is the t -step noise schedule as in Eq. (4). For simplicity, we denote $\mathcal{A}^{t-1} = \mathcal{G}(\mathcal{Z}^t), t = 1, 2, \dots, T$.

In order to obtain high reward design, we propose optimizing the shared latent code \mathcal{Z}^t to obtain the high-reward denoising CDR loop $\mathcal{A}^{t-1} = \mathcal{G}(\mathcal{Z}^t)$ at each sampling step. Namely,

$$\max_{\mathcal{Z}^t} \mathcal{F}(\mathcal{G}(\mathcal{Z}^t)), \forall t = T, T - 1, \dots, 1. \quad (6)$$

Since the shared embedding \mathcal{Z}^t encodes both sequence and structure, any updates to \mathcal{Z}^t affect all three modalities simultaneously, ensuring their synchronization throughout the optimization process.

3.3 Reward-guided Updates

Most property evaluators and denoising sampling for the CDRs are not differentiable, so we cannot perform white-box functional guidance. Motivated by the black-box optimization technique [Wierstra *et al.*, 2014; Nesterov and Spokoiny, 2017], we propose learning the distribution of latent codes \mathcal{Z}^t to enable black-box gradient updates, thereby facilitating an efficient optimization of \mathcal{Z}^t . Let $\pi(\mathcal{Z}^t; \mathcal{W}^t)$ denote the distribution of latent code, where \mathcal{W}^t is the distribution parameter. We can reformulate Eq. (6) as

$$\max_{\pi(\mathcal{Z}^t; \mathcal{W}^t)} J(\mathcal{W}^t) = \mathbb{E}_{\mathcal{Z}^t \sim \pi(\mathcal{Z}^t; \mathcal{W}^t)} [\mathcal{F}(\mathcal{G}(\mathcal{Z}^t))], \quad (7)$$

$$\forall t = T, T - 1, \dots, 1.$$

Algorithm 1 LatEnt bLack-box Design (LEAD)

1: **Input:** a pre-trained antibody diffusion model featuring three-output denoising network with shared encoder ($\mathcal{E}, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$), the antibody-antigen backbone \mathcal{R} , the dynamic step size $\{\beta_t\}_{t=1}^T$, the property evaluator $\mathcal{F}(\uparrow)$, number of evaluations per time step K , initial guidance step $T_{\text{init}} \leq T$, the step size δ .

2: **Initialize:** $\mathbf{A}^T = (\mathbf{s}^T \sim \mathcal{C}(\mathbb{I}(\mathbf{s}^T)|1/M), \mathbf{x}^T \sim \mathcal{N}(\mathbf{x}^T|\mathbf{0}, \mathbf{I}), \mathbf{O}^T \sim \mathcal{IG}_{\text{SO}(3)}(\mathbf{I}, 1))$.

3: **for** $t = T, T-1, \dots, 1$ **do**

4: **if** $t > T_{\text{init}}$ **then**

5: $\mathbf{A}^{t-1} \sim \text{DDPM}(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{E}(\mathbf{A}^t, \mathcal{R}), \beta_t)$;

6: **else**

7: $\mathcal{Z}^t = \mathcal{E}(\mathbf{A}^t, \mathcal{R})$;

8: **for** $k = 1, 2, \dots, K$ **do**

9: $\mathcal{Z}_k^t = \mathcal{Z}^t + \sigma \delta_k^t$ where $\delta_k^t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$;

10: $\mathbf{A}_k^{t-1} \sim \text{DDIM}(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{Z}_k^t, \beta_t)$;

11: **end for**

12: Obtain the optimal noise δ_\star^t based on Section 3.3;

13: $\mathbf{A}_\star^{t-1} \sim \text{DDIM}(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{Z}^t + \sigma \delta_\star^t, \beta_t)$;

14: **end if**

15: **end for**

16: **Output:** generated CDR loop \mathbf{A}_\star^0 .

Using the log-likelihood trick, the Monte Carlo gradient estimation w.r.t. \mathcal{W}^t is estimated as:

$$\nabla_{\mathcal{W}^t} J(\mathcal{W}^t) \approx \frac{1}{K} \sum_{k=1}^K \mathcal{F}(\mathcal{G}(\mathcal{Z}_k^t)) \nabla_{\mathcal{W}^t} \log \pi(\mathcal{Z}_k^t; \mathcal{W}^t), \quad (8)$$

where \mathcal{Z}_k^t is sampled from $\pi(\mathcal{Z}^t; \mathcal{W}^t)$. K is the number of function queries used to refine \mathcal{W}^t .

To further simplify, we instantiate the population distribution π as an isotropic multivariate Gaussian with mean \mathcal{W}^t and fixed covariance $\sigma^2 \mathbf{I}$, allowing us to write $J(\mathcal{W}^t)$ in terms of \mathcal{W}^t directly. Denoting $\mathcal{W}^t = \mathcal{Z}^t = \mathcal{E}(\mathbf{A}^t, \mathcal{R})$, Eq. (7) can be reformulated as:

$$\max_{\mathcal{Z}^t} J(\mathcal{Z}^t) = \mathbb{E}_{\delta^t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta^t))], \quad (9)$$

$$\forall t = T, T-1, \dots, 1.$$

Based on Eq. (8) and Eq. (9), we obtain the gradient w.r.t. \mathcal{Z}^t :

$$\nabla_{\mathcal{Z}^t} \mathbb{E} [\mathcal{F}(\mathcal{G}(\mathcal{Z}^t))] \approx \frac{1}{\sigma K} \sum_{k=1}^K \delta_k^t \mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_k^t)). \quad (10)$$

Thus, the latent code \mathcal{Z}^t can be updated as follows:

$$\mathcal{Z}^t \leftarrow \mathcal{Z}^t + \frac{1}{\sigma K} \sum_{k=1}^K \delta_k^t \mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_k^t)). \quad (11)$$

However, we find that Eq. (11) fails to achieve an effective optimization in the case of sequence derived properties (see empirical study). We hypothesize that it may result from the discrete nature of sequence space. Inspired from the fact that Eq. (11) actually updates \mathcal{Z}^t with a reward-weighted noise, we propose sampling an optimal noise based on probability

derived from function rewards to be compatible with the discrete nature of data. Specifically, the K reward values can be converted into probabilities via the softmax function:

$$P_k^t = \frac{\exp(\mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_k^t)))}{\sum_{j=1}^K \exp(\mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_j^t)))}. \quad (12)$$

The optimized noise index is sampled from this distribution:

$$\gamma^t \sim \text{Categorical}(P_1^t, P_2^t, \dots, P_K^t), \quad (13)$$

which is a one-hot vector. The optimized noise can be derived by $\delta_\star^t = \sum_{j=1}^K \gamma_j^t \delta_j^t$. The latent code is updated by:

$$\mathcal{Z}^t \leftarrow \mathcal{Z}^t + \sigma \delta_\star^t. \quad (14)$$

It is intuitive to derive a hard selection version of the above strategy by defining:

$$\delta_\star^t = \arg \max_{\delta_k^t} \{\mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_k^t))\}_{k=1}^K. \quad (15)$$

We refer to this strategy as the hard selection version of LEAD, termed LEAD-H, and to Eq. (13) as the soft selection version of LEAD, termed LEAD-S.

Additionally, we find that combining LEAD-H/LEAD-S with the weighted version of LEAD (termed LEAD-W) yields better results, respectively. We hypothesize that LEAD-W can explore a larger search space, thereby enhancing the advantages of LEAD-H/LEAD-S. The combining strategy is as the following. With a slight abuse of annotation, we re-denote the optimal noise obtained by LEAD-H or LEAD-S as ζ_i^t . Defining $\zeta_2^t = \frac{1}{\sigma K} \sum_{k=1}^K \delta_k^t \mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \delta_k^t))$, we further define the optimal noise as:

$$\delta_\star^t = \arg \max_{\zeta_i^t} \{\mathcal{F}(\mathcal{G}(\mathcal{Z}^t + \sigma \zeta_i^t))\}_{i=1}^2. \quad (16)$$

We perform LEAD guidance from the later stage of reverse diffusion process, namely, from $T_{\text{init}} (\leq T)$ steps. Before the guidance ($t > T_{\text{init}}$), the denoising process replicates the original one, namely

$$\mathbf{A}^{t-1} \sim \text{DDPM}(\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \mathcal{E}(\mathbf{A}^t, \mathcal{R}), \beta_t). \quad (17)$$

The whole algorithm is summarized in Algorithm 1.

4 Experiment

We follow [Villegas-Morcillo *et al.*, 2024] to consider two developability properties, i.e., hydrophathy score and folding energy for antibody design.

- **Hydrophathy score** (abbr. Hydro). It averages the hydrophathy values over the generated CDR sequences and serves as a proxy for solubility and aggregation [Kyte and Doolittle, 1982]. Negative scores indicate hydrophilicity while positive scores indicate hydrophobicity, and small score is preferred.
- **Folding energy difference** (abbr. $\Delta \Delta G$). It denotes the structural stability of generated antibodies relative to the reference. Following [Villegas-Morcillo *et al.*, 2024], we predict $\Delta \Delta G$ (abbr. Pred $\Delta \Delta G$) using an existing predictor released by [Shan *et al.*, 2022], which takes as input the amino acid sequence and C_α atom coordinates. A lower value suggests improved structural stability.

Guidance	Model	CDR H1				CDR L1			
		AAR (%)	RMSD (Å)	Hydro	Pred $\Delta\Delta G$	AAR (%)	RMSD (Å)	Hydro	Pred $\Delta\Delta G$
Unconditional	DiffAb	65.8	0.88	-0.23	-0.04	56.4	1.22	-0.54	-0.57
	GuideRaw	65.1	1.14	-0.31	-0.38	56.5	1.34	-0.65	-1.28
Pred $\Delta\Delta G$	LEAD	38.7	0.89	-0.88*	-0.67	31.1	1.25	-0.72	-3.11*
	GuideRaw	24.0	0.91	-2.93	-0.07	37.4	1.45	-2.78	-1.14
Hydro	LEAD	16.1	0.92	-3.22*	-0.05	24.4	1.28	-3.04*	-1.45
	GuideRaw	58.4	1.05	-0.78	-0.56	55.2	1.24	-0.99	-1.32
($\Delta\Delta G$ + Hydro)	LEAD	27.0	0.86	-2.31*	-0.66	30.7	1.26	-1.79*	-2.83*
		CDR H2				CDR L2			
Unconditional	DiffAb	49.0	0.77	-1.18	0.45	58.2	1.06	-0.46	-0.02
	GuideRaw	47.6	1.59	-1.31	0.02	55.6	1.76	-0.65	-0.36
Pred $\Delta\Delta G$	LEAD	21.5	0.70	-1.33	-1.34*	26.1	1.34	-0.66	-1.52*
	GuideRaw	6.10	0.76	-4.07	-0.24	22.9	1.05	-3.77	-0.46
Hydro	LEAD	3.68	0.74	-4.17*	-0.39	11.2	1.34	-4.01*	-0.57
	GuideRaw	37.6	1.24	-2.17	-0.05	52.2	1.30	-1.83	-0.41
($\Delta\Delta G$ + Hydro)	LEAD	16.4	0.73	-3.02*	-0.95*	18.9	1.18	-2.68*	-1.45*
		CDR H3				CDR L3			
Unconditional	DiffAb	26.3	3.35	-0.73	0.14	46.5	1.07	-1.31	-2.19
	GuideRaw	25.8	3.41	-1.03	-1.99	44.5	1.16	-1.39*	-2.94
Pred $\Delta\Delta G$	LEAD	21.1	3.38	-1.15	-2.16	20.1	1.18	-1.12	-3.52*
	GuideRaw	11.7	3.39	-3.87	-1.06	32.2	1.10	-3.63	-2.04
Hydro	LEAD	8.88	3.37	-3.96*	-0.94	21.4	1.20	-3.87*	-2.22
	GuideRaw	26.0	3.63	-1.62	-2.20	45.3	1.16	-1.93	-2.84
($\Delta\Delta G$ + Hydro)	LEAD	20.8	3.39	-2.26*	-2.35	26.5	1.20	-2.28*	-3.25*

Table 1: Average AAR, RMSD, Hydro, and Pred $\Delta\Delta G$ on 19 test complexes, evaluated for each CDR. ($\Delta\Delta G$ + Hydro) refers to a two-objective optimization with equal weight. The best results w.r.t different types of guidance are highlighted in bold. “*” denotes the difference is considered to be statistically significant with a significance level $\alpha = 0.05$. Results of GuideRaw are copied from the original paper, and its property guidance is incorporated from the first denoising step, i.e., $T_{\text{init}} = 100$. In contrast, our LEAD adopts $T_{\text{init}} = 50$, using fewer evaluation queries.

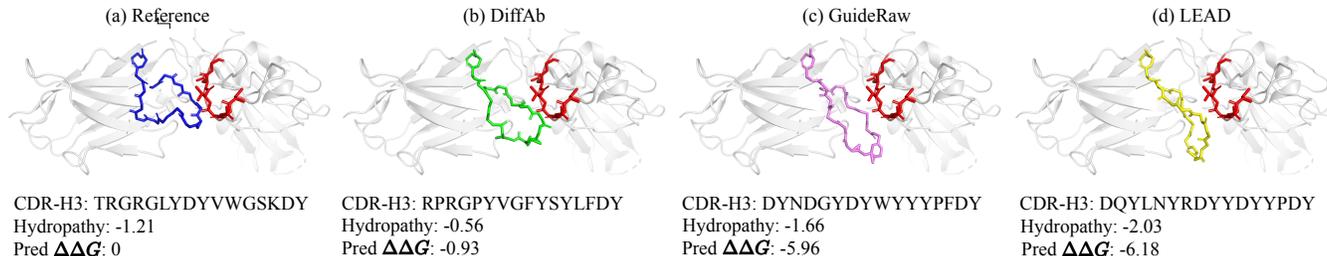


Figure 2: Visualization of sequence-structure designs generated by various methods for optimizing Pred $\Delta\Delta G$. The presented test complex is 7chf_A_B_R. The antigen epitope is displayed in red color. The other colored one is the designed CDR-H3.

Dataset We utilize the test set described in [Luo *et al.*, 2022], which consists of 19 antibody-antigen complexes derived from the SabDab database [Dunbar *et al.*, 2014]. The test set includes protein antigens from various pathogens, such as influenza and SARS-CoV-2.

Hyperparameters Unless otherwise specified, the latent noise scale σ is adaptively set to β_t (as in Eq. (3)) and the number of evaluation per time step K is set to 20, the initial guidance step T_{init} is set to 50. It is more effective to add guidance at a later stage of the denoising process as samples from the earlier steps can be too noisy to evaluate accurately.

Baselines We employ the widely adopted antibody sequence-structure co-design diffusion model framework,

DiffAb [Luo *et al.*, 2022] as the pretrained model. We consider the relevant work, property optimization under the framework of sequence-structure co-design as baselines¹. [Villegas-Morcillo *et al.*, 2024] performs guidance similar to our hard/soft selection strategy, but on the raw data space. We name it as GuideRaw-H and GuideRaw-S for hard and soft selection version, respectively. Except where mentioned, the experimental setup follows the setting of the original

¹Other baselines DIFFFORCE [Kulytė *et al.*, 2024] and ABDPO [Zhou *et al.*, 2024] do not release their codes. Nevertheless, DIFFFORCE requires a differential force field model, which is inapplicable to the black-box setting of this work. ABDPO needs to fine-tunes DIFFAB, requiring a large amount of property evaluations.

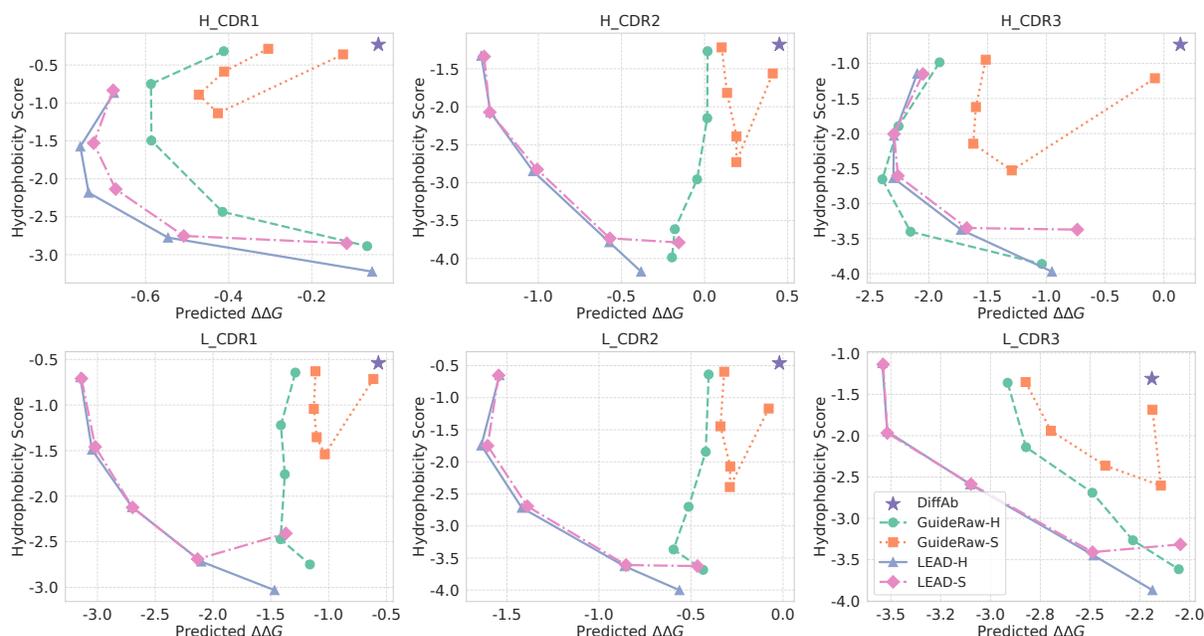


Figure 3: Trade-off in multi-property optimization. Each point represents the average Hydro and Pred $\Delta\Delta G$ on 19 test complexes, evaluated for each CDR. Different weights are used to combine the two objectives, i.e., $w * \Delta\Delta G + (1 - w) * \text{Hydro}$, with $w = 0, 0.25, 0.5, 0.75, 1$.

work. That is, the initial guidance step T_{init} is set to 100 and the number of evaluation per time step K is set to 20.

Evaluation In addition to the properties mentioned above, we evaluate these designs using the following metrics: (i) **Amino Acid Recovery (AAR)**, which measures the sequence identity between the reference and generated CDR sequences; (ii) **Root Mean Square Deviation (RMSD)**, which computes the C_{α} atom distance between the reference and generated CDR structures. AAR and RMSD are reported for reference purposes and are not bolded in Table results. For a detailed explanation, please refer to point (4) in Section 4.1.

4.1 Single Property Optimization

For all methods, we generate 100 designs per CDR (six CDRs in total) for each of the 19 test antigen-antibody complexes. The guidance method is to optimize the hydrophobicity score and the predicted $\Delta\Delta G$, respectively. Here, both GuideRaw and our LEAD adopt the hard selection version for a fair comparison. Other guidance strategies are studied in Section 4.4.

We report the evaluation results in Table 1. (1) The results show that our LEAD outperforms GuideRaw in terms of both target properties, which is attributed to our optimization over the shared latent code ensuring the synchronization of different modalities. (2) LEAD conducts guidance in only 50 steps, compared to the 100 steps required by GuideRaw. This demonstrates that our method achieves better performance while requiring fewer black-box function evaluations. (3) DiffAb performs unconditional generation and exhibits strong AAR and RMSD scores. This indicates that the antibody generated by this sequence-structure co-design model can recover CDR structures similar to the given reference. However, the target properties of the antibodies produced by DiffAb fall below expectations. (4) Property optimization often leads to a degradation of AAR and RMSD, as it requires

adjusting the generated CDRs to meet property requirements. Specifically, the hydrophobicity score is determined by the sequences only. As a result, we observe that antibodies optimized by all property guidance methods exhibit significant AAR degradation but improved hydrophobicity scores. Notably, the RMSD scores remain relatively low across all generations, indicating that the generated CDR structures are still similar to the references.

Fig. 2 visualizes the generated antibodies. It shows that all methods can generate valid sequences and structures. GuideRaw and our LEAD obtain better hydrophobicity score and predicted $\Delta\Delta G$ than natural antibodies.

4.2 Multi-property Optimization

Optimizing multiple properties presents challenges because each property can have different value scales and units. To address this, we normalize property values to ensure they contribute comparably during optimization. We adopt a weighted sum with adjustable weights to balance trade-offs according to design priorities. Here, we optimize both the hydrophobicity score and $\Delta\Delta G$ with $w * \Delta\Delta G + (1 - w) * \text{Hydro}$. The weight w is set as 0, 0.25, 0.5, 0.75, 1, respectively. For a fair comparison, we set up the same number of function queries for our LEAD and GuideRaw in this setting. That is, both methods set $T_{\text{init}} = 50$ and $K = 20$. Still, we generate 100 generated designs per CDR for each of the 19 test antigen-antibody complexes.

The results are summarized in Fig. 3. The hydrophobicity and energy properties exhibit a trade-off, where improving one often comes at the expense of the other. Our LEAD effectively navigates this balance, achieving superior solutions compared to the baseline approaches.

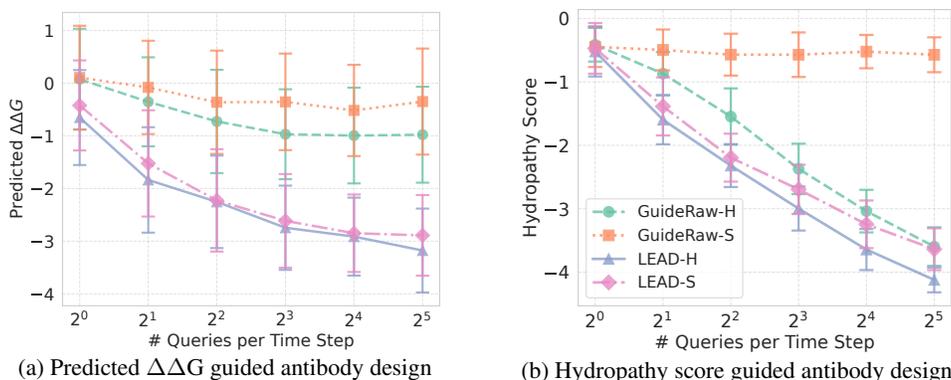


Figure 4: Query-efficient comparison on the test complex 5xku.C.B.A (design H-CDR1) between GuideRaw and our LEAD w.r.t. two selection strategies. For a fair comparison, property guidance is incorporated for all methods after 50 time steps respectively.

Guidance Metric	Pred $\Delta\Delta G$				Hydro			
	AAR (%)	RMSD (\AA)	Hydro	Pred $\Delta\Delta G$	AAR (%)	RMSD (\AA)	Hydro	Pred $\Delta\Delta G$
DiffAb	58.4 \pm 13.9	1.28 \pm 0.20	-0.43 \pm 0.39	-0.75 \pm 0.87	58.4 \pm 13.9	1.28 \pm 0.20	-0.43 \pm 0.39	-0.75 \pm 0.87
GuideRaw-H	66.1 \pm 11.2	1.62 \pm 0.40	-0.57 \pm 0.34	-1.01 \pm 0.93	15.2 \pm 12.3	1.37 \pm 0.29	-3.25 \pm 0.28	-0.99 \pm 0.79
GuideRaw-S	66.8 \pm 12.2	1.53 \pm 0.33	-0.50 \pm 0.35	-0.81 \pm 0.97	68.4 \pm 10.8	1.39 \pm 0.29	-0.59 \pm 0.31	-0.11 \pm 0.91
LEAD-H	13.3 \pm 5.12	1.35 \pm 0.24	-1.10 \pm 0.56	-2.66 \pm 0.69	4.78 \pm 5.84	1.33 \pm 0.24	-3.75 \pm 0.25	-1.73 \pm 0.76
LEAD-S	12.2 \pm 2.29	1.35 \pm 0.23	-0.83 \pm 0.56	-2.65 \pm 0.71	10.4 \pm 9.05	1.33 \pm 0.27	-3.22 \pm 0.33	-1.65 \pm 0.84
LEAD-W	37.0 \pm 17.4	1.67 \pm 0.21	-0.28 \pm 0.62	-2.64 \pm 0.93	35.1 \pm 17.2	1.47 \pm 0.26	-0.77 \pm 0.44	-1.21 \pm 0.93
LEAD-(W+H)	14.1 \pm 7.70	1.43 \pm 0.23	<u>-0.97 \pm 0.55</u>	<u>-2.89 \pm 0.81</u>	3.80 \pm 3.72	1.30 \pm 0.26	-3.78 \pm 0.23	<u>-1.78 \pm 0.70</u>
LEAD-(W+S)	14.4 \pm 8.48	1.45 \pm 0.24	-0.83 \pm 0.65	-2.90 \pm 0.80	7.34 \pm 6.48	1.29 \pm 0.21	-3.38 \pm 0.34	-1.88 \pm 0.68

Table 2: Average AAR, RMSD, Hydro, and Pred $\Delta\Delta G$ on the test complex 5xku.C.B.A (design H-CDR1). For a fair comparison, property guidance is incorporated for all methods after 50 time steps respectively.

4.3 Query Efficiency

By guiding optimization in the shared latent space, our LEAD captures property semantics effectively and preserves the correspondence between sequence and structure, leading to more efficient guidance. We conduct experiments with different numbers of queries, $2^0, 2^1, 2^2, 2^3, 2^4, 2^5$ per guided denoising step. In this experiment, all methods generate 100 designs for H-CDR1 from one of the test complexes and perform property-guided denoising after 50 denoising steps.

The results in Fig. 4 demonstrate the efficiency of our method. To be specific, with an increased number of black-box evaluations, the optimization process explores a wider range of possibilities, leading to a more thorough refinement of the target property. It is interesting to see that our LEAD achieves a much larger performance gain in optimizing the predicted $\Delta\Delta G$ than in optimizing the hydropathy score. The predicted $\Delta\Delta G$ is relevant to both sequences and structures, while the hydropathy score is only determined by sequences. This particularly demonstrates our method’s superiority on those sequence-structure derived properties.

4.4 Guidance Strategies

Here, we explore the efficacy of different guidance strategies demonstrated in Section 3.3, namely, LEAD-H, LEAD-S, LEAD-W, LEAD-(W+H), LEAD-(W+S).

(1) Our LEAD with different guidance strategies achieves better results in terms of optimizing predicted $\Delta\Delta G$. This property is relevant to both antibody sequence and structure.

The guidance over the shared latent space indeed enhances its optimization. (2) In terms of the hydropathy score, LEAD-W does not achieve efficient optimization, resulting in a score similar to that of unconditional generation. This is because the hydropathy score is a sequence derived property, and its gradient space is not continuous. As a result, the weighted strategy fails to perform an effective search. Instead, the discrete search strategies, LEAD-H and LEAD-S, are effective in this case. (3) The weighted strategy (Eq. (16)) can slightly boost the discrete search strategies, as demonstrated by the results of LEAD-(W+H) and LEAD-(W+S).

5 Conclusion

In this work, we introduce LEAD, a black-box guidance framework that operates in a shared latent space of sequence and structure, enabling efficient property optimization for antibody CDRs without model retraining. LEAD is generally effective at maintaining sequence-structure synchronization; however, its performance is limited by the capacity of DiffAb, which may fail to preserve meaningful alignment between sequence and structure. More advanced co-design frameworks [Yang *et al.*, 2025; Uçar *et al.*, 2024a] may offer improved capabilities in this regard. Future work will also incorporate wet-lab experiments to validate computational designs and provide feedback to further guide optimization. Beyond antibody design, LEAD is broadly applicable to other AI-driven co-design tasks involving coupled sequence-structure modalities, e.g., protein, RNA, and small-molecule design.

Acknowledgments

This paper was supported by the Manufacturing Trade and Connectivity (MTC) Individual Research Grant (IRG) (Grant No. M24N7c0091). It is also supported by the National Research Foundation, Singapore under its National Large Language Models Funding Initiative (AISG Award No: AISG-NMLP-2024-004). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- [Adolf-Bryfogle *et al.*, 2018] Jared Adolf-Bryfogle, Oleks Kalyuzhnyi, Michael Kubitz, Brian D Weitzner, Xiaozhen Hu, Yumiko Adachi, William R Schief, and Roland L Dunbrack Jr. Rosettaantibodydesign (rabd): A general framework for computational antibody design. *PLoS computational biology*, 14(4):e1006112, 2018.
- [Branden and Tooze, 2012] Carl Ivar Branden and John Tooze. *Introduction to protein structure*. Garland Science, 2012.
- [Dunbar *et al.*, 2014] James Dunbar, Konrad Krawczyk, Jinwoo Leem, Terry Baker, Angelika Fuchs, Guy Georges, Jiye Shi, and Charlotte M Deane. Sabdab: the structural antibody database. *Nucleic acids research*, 42(D1):D1140–D1146, 2014.
- [Gallier and Xu, 2003] Jean Gallier and Dianna Xu. Computing exponentials of skew-symmetric matrices and logarithms of orthogonal matrices. *International Journal of Robotics and Automation*, 18(1):10–20, 2003.
- [Gruver *et al.*, 2024] Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete diffusion. *Advances in neural information processing systems*, 36, 2024.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [Ingraham *et al.*, 2023] John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.
- [Janeway *et al.*, 2001] Charles Janeway, Paul Travers, Mark Walport, Mark Shlomchik, et al. *Immunobiology: the immune system in health and disease*, volume 2. Garland Pub. New York, 2001.
- [Jin *et al.*, 2022] Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi S. Jaakkola. Iterative refinement graph neural network for antibody sequence-structure co-design. In *International Conference on Learning Representations*, 2022.
- [Kong *et al.*, 2023] Xiangzhe Kong, Wenbing Huang, and Yang Liu. End-to-end full-atom antibody design. In *Proceedings of the 40th International Conference on Machine Learning*, pages 17409–17429, 2023.
- [Kulytė *et al.*, 2024] Paulina Kulytė, Francisco Vargas, Simon Valentin Mathis, Yu Guang Wang, José Miguel Hernández-Lobato, and Pietro Liò. Improving antibody design with force-guided sampling in diffusion models. *arXiv preprint arXiv:2406.05832*, 2024.
- [Kyte and Doolittle, 1982] Jack Kyte and Russell F Doolittle. A simple method for displaying the hydropathic character of a protein. *Journal of molecular biology*, 157(1):105–132, 1982.
- [Lapidoth *et al.*, 2015] Gideon D Lapidoth, Dror Baran, Gabriele M Pszolla, Christoffer Norn, Assaf Alon, Michael D Tyka, and Sarel J Fleishman. Abdesign: A n algorithm for combinatorial backbone design guided by natural conformations and sequences. *Proteins: Structure, Function, and Bioinformatics*, 83(8):1385–1406, 2015.
- [Larman *et al.*, 2012] H Benjamin Larman, George Jing Xu, Natalya N Pavlova, and Stephen J Elledge. Construction of a rationally designed antibody platform for sequencing-assisted selection. *Proceedings of the National Academy of Sciences*, 109(45):18523–18528, 2012.
- [Leach *et al.*, 2022] Adam Leach, Sebastian M Schmon, Matteo T. Degiacomi, and Chris G. Willcocks. Denoising diffusion probabilistic models on SO(3) for rotational alignment. In *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022.
- [Lisanza *et al.*, 2024] Sidney Lyayuga Lisanza, Jacob Merle Gershon, Samuel WK Tipps, Jeremiah Nelson Sims, Lucas Arnoldt, Samuel J Hendel, Miriam K Simma, Ge Liu, Muna Yase, Hongwei Wu, et al. Multistate and functional protein design using rosettafold sequence space diffusion. *Nature biotechnology*, pages 1–11, 2024.
- [Luo *et al.*, 2022] Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific antibody design and optimization with diffusion-based generative models for protein structures. *Advances in Neural Information Processing Systems*, 35:9754–9767, 2022.
- [Martinkus *et al.*, 2024] Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, et al. Abdiffuser: full-atom generation of in-vitro functioning antibodies. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Nesterov and Spokoiny, 2017] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.
- [Presta, 1992] Leonard G Presta. Antibody engineering. *Current Opinion in Structural Biology*, 2(4):593–596, 1992.
- [Saka *et al.*, 2021] Koichiro Saka, Taro Kakuzaki, Shoichi Metsugi, Daiki Kashiwagi, Kenji Yoshida, Manabu Wada,

- Hiroyuki Tsunoda, and Reiji Teramoto. Antibody design using lstm based deep generative model from phage display library for affinity maturation. *Scientific reports*, 11(1):5852, 2021.
- [Shan *et al.*, 2022] Sisi Shan, Shitong Luo, Ziqing Yang, Junxian Hong, Yufeng Su, Fan Ding, Lili Fu, Chenyu Li, Peng Chen, Jianzhu Ma, et al. Deep learning guided optimization of human antibody against sars-cov-2 variants with broad neutralization. *Proceedings of the National Academy of Sciences*, 119(11):e2122954119, 2022.
- [Shuai *et al.*, 2023] Richard W Shuai, Jeffrey A Ruffolo, and Jeffrey J Gray. Iglm: Infilling language modeling for antibody sequence design. *Cell Systems*, 14(11):979–989, 2023.
- [Song *et al.*, 2020] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [Sumida *et al.*, 2024] Kiera H Sumida, Reyes Núñez-Franco, Indrek Kalvet, Samuel J Pellock, Basile IM Wicky, Lukas F Milles, Justas Dauparas, Jue Wang, Yakov Kipnis, Noel Jameson, et al. Improving protein expression, stability, and function with proteinmpnn. *Journal of the American Chemical Society*, 146(3):2054–2061, 2024.
- [Uçar *et al.*, 2024a] Talip Uçar, Cedric Malherbe, and Ferran Gonzalez. Benchmarking generative models for antibody design & exploring log-likelihood for sequence ranking. *bioRxiv*, pages 2024–10, 2024.
- [Ucar *et al.*, 2024b] Talip Ucar, Cedric Malherbe, and Ferran Gonzalez. Exploring log-likelihood scores for ranking antibody sequence designs. In *NeurIPS 2024 Workshop on AI for New Drug Modalities*, 2024.
- [Villegas-Morcillo *et al.*, 2024] Amelia Villegas-Morcillo, Jana M. Weber, and Marcel J. T. Reinders. Guiding diffusion models for antibody sequence and structure co-design with developability properties. *PRX Life*, 2:033012, Sep 2024.
- [Warszawski *et al.*, 2019] Shira Warszawski, Aliza Borenstein Katz, Rosalie Lipsh, Lev Khmelnsky, Gili Ben Nissan, Gabriel Javitt, Orly Dym, Tamar Unger, Orli Knop, Shira Albeck, et al. Optimizing antibody affinity and stability by the automated design of the variable light-heavy chain interfaces. *PLoS computational biology*, 15(8):e1007207, 2019.
- [Watson *et al.*, 2023] Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- [Wierstra *et al.*, 2014] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural evolution strategies. *The Journal of Machine Learning Research*, 15(1):949–980, 2014.
- [Yang *et al.*, 2025] Nianzu Yang, Jian Ma, Songlin Jiang, Huaijin Wu, Shuangjia Zheng, Wengong Jin, and Junchi Yan. Repurposing alphafold3-like protein folding models for antibody sequence and structure co-design. In *ICLR 2025 Workshop on Generative and Experimental Perspectives for Biomolecular Design*, 2025.
- [Zhou *et al.*, 2024] Xiangxin Zhou, Dongyu Xue, Ruizhe Chen, Zaixiang Zheng, Liang Wang, and Quanquan Gu. Antigen-specific antibody design via direct energy-based preference optimization. *arXiv preprint arXiv:2403.16576*, 2024.