Learning from B Cell Evolution: Adaptive Multi-Expert Diffusion for Antibody Design via Online Optimization

Hanqi Feng[1]*, Peng Qiu[1]*, Mengchun Zhang[2]*, Yiran Tao[1], You Fan[3], Jingtao Xu[4], Barnabas Poczos[1]*

[1]School of Computer Science, Carnegie Mellon University
[2]Department of Biostatistics, University of Pittsburgh
[3]Department of Mathematics, King's College London
[4]Department of Statistics, London School of Economics and Political Science
{hanqif, pengq, bpoczos}@cs.cmu.edu, andy0814xu@gmail.com

Abstract

Recent diffusion models show strong potential for antibody design, but most use uniform strategies that ignore antigen-specific needs. Inspired by B-cell affinity maturation—balancing affinity, stability, and self-avoidance—we present a biologically motivated framework that brings physics-based priors into an online meta-learning loop. Specialized experts (van der Waals, molecular recognition, energy balance, interface geometry) adapt their weights during generation via iterative feedback, mimicking natural refinement. Instead of fixed protocols, the system discovers target-specific guidance. In experiments, it: (1) learns SE(3)-equivariant guidance for different antigen classes without pretraining, preserving molecular symmetries; (2) improves hotspot coverage and interface quality through target-specific adaptation; (3) enables iterative refinement where each complex learns its own optimization profile via online evaluation; and (4) generalizes from small epitopes to large interfaces, supporting precision campaigns for individual targets.

1 Introduction

Computational antibody design remains a fundamental challenge in therapeutic development, requiring simultaneous optimization of hotspot coverage, structural stability, and binding interface quality [10, 23]. While recent diffusion-based methods like RFdiffusion [33] and RFAntibody [4] show promise for generating protein structures, they lack mechanisms to incorporate real-time constraints during generation, often producing designs that fail to meet multiple competing objectives [8].

Current approaches face three fundamental limitations: (1) they generate structures through diffusion without target-specific guidance [32, 33], (2) they cannot balance multiple objectives during generation, requiring inefficient post-hoc filtering [17, 31], and (3) they either ignore physical constraints or require extensive labeled data to train property predictors [7, 15]. These limitations significantly reduce their effectiveness for therapeutic antibody design, where each target presents unique challenges [2, 22].

We present an adaptive guidance framework that addresses these limitations by introducing physics-aware constraints into the SE(3)-equivariant diffusion process [6, 14]. Our contributions include:

*Corresponding author

The NeurIPS 2025 Workshop on AI Virtual Cells and Instruments: A New Era in Drug Discovery and Development (AI4D3 2025), San Diego, California, USA, 2025.

- **Multi-expert guidance system**: We develop guidance modules for van der Waals interactions, molecular recognition [11], energy balance [16], and interface geometry [28], each providing targeted gradients during diffusion while maintaining equivariance [27].
- **Novel expert routing**: A dynamic routing system activates appropriate experts based on real-time structural metrics and diffusion timestep, ensuring optimal constraint application throughout generation [24, 35].
- **Online parameter adaptation**: Using Bayesian optimization with Gaussian processes, our framework learns optimal guidance strategies for each antigen-antibody pair through iterative batch evaluation, automatically discovering target-specific temporal profiles without requiring pre-training [25, 29].

Experiments across diverse targets demonstrate substantial improvements: 7% reduction in CDR-H3 RMSD, 9% increase in hotspot coverage, 12% better interface pAE, and 5% higher shape complementarity, with enhanced metrics across all evaluation criteria. This balanced optimization addresses the critical "weakest link" problem [25] in antibody design. By combining physics-based guidance with online learning, our work opens new directions for adaptive approaches in biomolecular design.

2 Related Work

Computational Antibody Design The field has evolved from physics-based approaches using Rosetta [19] and MODELLER [34] to deep learning methods inspired by AlphaFold2 [18]. DeepAb [26] pioneered deep learning for antibody structure prediction, while ABlooper [1] focused on CDR-H3 modeling. Recent generative approaches shifted to de novo design. DiffAb [21] introduced diffusion models for joint sequence-structure generation but operates in internal coordinates, limiting inter-chain modeling. RFAntibody [4] addressed this using SE(3)-equivariant backbone generation in Cartesian space. However, current methods lack both physicochemical constraint enforcement and target-specific adaptation during generation. This results in high-throughput campaigns producing substantial fractions of structurally unreasonable or experimentally nonviable designs [30].

3 Preliminaries

3.1 SE(3)-Equivariant Diffusion Models

Protein structure modeling requires respecting 3D spatial symmetries. RFdiffusion [33] combines SE(3)-equivariant networks with diffusion models, performing diffusion directly on backbone coordinates while maintaining rotational and translational equivariance.

Frame Representations and Forward Diffusion Process Each residue i is represented by a rigid body frame $\mathbf{T}_i = (\mathbf{R}_i, \mathbf{t}_i) \in \mathrm{SE}(3)$, where $\mathbf{R}_i \in \mathrm{SO}(3)$ is a rotation matrix and $\mathbf{t}_i \in \mathbb{R}^3$ is a translation vector. For N residues, the complete structure is $\mathbf{T} = \{\mathbf{T}_1, \ldots, \mathbf{T}_N\}$.

The forward diffusion process adds noise over time steps $t \in 0, 1, \ldots, T$. Rotations follow the IGSO3 (heat-kernel) distribution, $q(\mathbf{R}_t|\mathbf{R}_0) = \mathrm{IGSO3}(\mathbf{R}_t; \mathbf{R}_0, \sigma_t)$ with density:

$$f(\mathbf{R}_t|\mathbf{R}_0, \sigma_t) = \frac{1}{Z(\sigma_t)} \exp\left(-\frac{\|\log(\mathbf{R}_0^\top \mathbf{R}_t)\|_F^2}{2\sigma_t^2}\right)$$

where σ_t is the noise level at timestep t, controlling the variance of the distribution [36]. Translations follow Gaussian diffusion with center-of-mass constraint:

$$q(\mathbf{t}_t \mid \mathbf{t}_0) = \mathcal{N}(\mathbf{t}_t; \mathbf{t}_0, \sigma_t^2 \mathbf{I}_{3N}), \quad \text{s.t.} \sum_{i=1}^{N} \mathbf{t}_{t,i} = \sum_{i=1}^{N} \mathbf{t}_{0,i}.$$

Reverse Diffusion and Network Architecture The reverse process learns to denoise the forward diffusion process by estimating clean structure $\mathbf{T}_0$ from noisy $\mathbf{T}_t$ [13]:

$$p_\theta(\mathbf{T}_{t-1}|\mathbf{T}_t) = p(\mathbf{T}_{t-1}|\mathbf{T}_t, \hat{\mathbf{T}}_0),$$
$$\hat{\mathbf{T}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{T}_t - \sqrt{1-\bar{\alpha}_t}\epsilon_\theta(\mathbf{T}_t, t)\right)$$

where $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$ is the cumulative product of noise schedule parameters, and $\epsilon_\theta(\mathbf{T}_t, t)$ is the neural network with parameters θ that predicts the noise added at timestep t.

4 Problem Formulation and Motivation

While state-of-the-art generative models excels at generating diverse protein structures, it faces a critical challenge in antibody design: maintaining balanced performance across multiple essential metrics. In antibody engineering, success requires simultaneous optimization of structural accuracy (CDR-H3 RMSD), prediction confidence (pAE, ipAE), interface quality (shape complementarity, buried surface area), and biophysical properties (VDW interactions, stability) [5, 23, 25].

The fundamental issue is that RFdiffusion's purely data-driven approach lacks mechanisms to enforce multi-objective balance, producing the "weakest link" problem—one failing metric makes the entire design experimentally nonviable [5, 25]; in high-throughput campaigns, even if 90% of metrics are excellent, failure in the remaining 10% removes the design from consideration. We address this with an adaptive guidance system that steers diffusion toward regions where all objectives are met simultaneously; unlike post-hoc filtering that discards failures, our method proactively maintains metric balance during generation, increasing the yield of experimentally viable candidates per computational cycle.

5 Method

We present an adaptive physics-guided diffusion framework for antibody design that integrates multi-expert guidance with online learning. Our approach enhances the standard SE(3) equivariant diffusion process [14, 33] by incorporating domain-specific physical constraints through a dynamically weighted ensemble of expert modules. Figure 1 presents our adaptive framework. The system combines physics-based expert guidance (left) with the antibody design pipeline, where guided diffusion generates structures, ProteinMPNN designs sequences, and RoseTTAFold2 [3] validates results. A Bayesian optimization feedback loop (bottom) continuously learns optimal guidance parameters from evaluation metrics.

5.1 Theoretical Foundation

Biophysical Principles of Antibody–Antigen Recognition Our multi-expert system is grounded in the physics of molecular recognition. First, the energetic landscape of binding can be expressed as the free-energy decomposition [12]:

$$\Delta G_{\text{bind}} = \Delta H_{\text{vdW}} + \Delta H_{\text{elec}} + \Delta H_{\text{hbond}} - T\Delta S_{\text{conf}} + \Delta G_{\text{solv}},$$

where the terms denote van der Waals, electrostatic, hydrogen-bonding, conformational-entropy, and solvation contributions, respectively. Our expert modules directly target these physical components. Second, optimal binding also reflects shape complementarity, consistent with the lock-and-key model extended by induced fit [9, 20]. We quantify complementarity as

$$SC = \frac{1}{2}\left(\frac{S_{\text{buried,Ab}}}{S_{\text{total,Ab}}} + \frac{S_{\text{buried,Ag}}}{S_{\text{total,Ag}}}\right) \cdot \exp\left(-\frac{d_{\text{gap}}}{\lambda}\right),$$

where Ab and Ag denote antibody and antigen, S_{buried} is the buried surface area upon complex formation, S_{total} the total surface area of each partner, and d_{gap} penalizes interface gaps with decay length λ.

5.2 Physics-Guided SE(3) Diffusion

Our framework modifies the standard reverse diffusion process by incorporating physics-based guidance gradients. The reverse step is formulated as:

$$\mathbf{T}_{t-1} = \boldsymbol{\mu}_\theta(\mathbf{T}_t, t) + \boldsymbol{\Sigma}_\theta^{1/2}(\mathbf{T}_t, t) \cdot \mathbf{z} - \lambda(t) \cdot \mathbf{g}(\mathbf{T}_t, t)$$

where $\boldsymbol{\mu}_\theta$ and $\boldsymbol{\Sigma}_\theta$ are the predicted mean and covariance from the denoising network, $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, and $\mathbf{g}(\mathbf{T}_t, t)$ represents the combined guidance gradient:

$$\mathbf{g}(\mathbf{T}_t, t) = \sum_{i=1}^{4} w_i(t) \cdot \lambda_i(t) \cdot \nabla \mathcal{L}_i(\mathbf{T}_t)$$

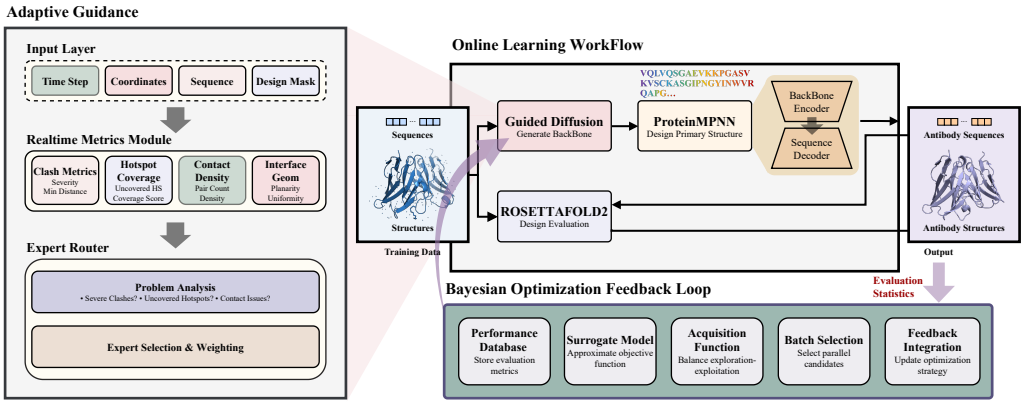

Figure 1: Adaptive multi-expert diffusion framework with online learning for antibody design.

where $w_i(t)$ are adaptive weights for each expert and $\nabla\mathcal{L}_i$ are gradients of expert-specific physics-based loss functions.

The guidance strength for each expert i at timestep t is determined as:

$$\lambda_i(t) = \lambda_{\text{base},i} \cdot f_{\text{temporal}}(t, \alpha, \beta)$$

where $\lambda_{\text{base},i}$ is the base strength for expert i, and $f_{\text{temporal}}(t, \alpha, \beta)$ controls when each expert is most active during denoising. The shape parameters α and β are learned via Bayesian optimization to discover optimal temporal profiles—early-peaking ($\alpha < \beta$) for global structure experts, late-peaking ($\alpha > \beta$) for atomic refinement:

$$f_{\text{temporal}}(t, \alpha, \beta) = \frac{B(t_{\text{norm}}; \alpha, \beta)}{B(t_{\text{mode}}; \alpha, \beta)} \cdot \lambda_{\text{peak}}$$

where $t_{\text{norm}} = (t-1)/(T-1)$ normalizes the timestep to $[0, 1]$, and $B(\cdot; \alpha, \beta)$ is the Beta density function. The shape parameters α and β are learned via Bayesian optimization to discover optimal activation timing for each target.

5.3 Multi-Expert Guidance System

We employ four specialized expert modules, each addressing critical aspects of antibody-antigen interactions:

VDW Balance Expert This expert prevents steric clashes while maintaining favorable van der Waals interactions. The loss function penalizes atomic overlaps:

$$\mathcal{L}_{\text{vdw}} = \sum_{\{i,j\} \subseteq \text{CDR} \cup \text{target}} \max(0, r_{\text{clash}} - d_{ij})^2$$

where r_{clash} is the clash threshold and $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$ is the distance between backbone atoms. The gradient is:

$$\nabla\mathcal{L}_{\text{vdw}} = 2 \sum_{d_{ij} < r_{\text{clash}}} \frac{r_{\text{clash}} - d_{ij}}{d_{ij}} \cdot (\mathbf{x}_i - \mathbf{x}_j)$$

Molecular Recognition Expert This expert ensures proper CDR coverage of epitope hotspots. For each hotspot $h \in \mathcal{H}$, we encourage proximity to CDR residues:

$$\mathcal{L}_{\text{hotspot}} = \sum_{h \in \mathcal{H}} \min_{c \in \text{CDR}} \|\mathbf{x}_h - \mathbf{x}_c\|^2$$

Energy Balance Expert This expert maintains optimal contact density at the binding interface. The loss function is defined as:

$$\mathcal{L}_{\text{contact}} = \begin{cases} \phi^-(n_c) & \text{if } n_c < \tau^- \\ \phi^+(n_c) & \text{if } n_c > \tau^+ \\ 0 & \text{if } \tau^- \leq n_c \leq \tau^+ \end{cases}$$

where n_c denotes the number of inter-molecular contacts within threshold d_c, $[\tau^-, \tau^+]$ defines the target range, and $\phi^-(\cdot)$, $\phi^+(\cdot)$ are penalty functions for sparse and dense packing, respectively. The gradient modulates CDR positioning to optimize interface contacts.

Interface Quality Expert This expert optimizes the geometric properties of the binding interface through multiple physically motivated objectives:

$$\mathcal{L}_{\text{geom}} = w_u \cdot \mathcal{L}_{\text{uniformity}} + w_c \cdot \mathcal{L}_{\text{cavity}}$$

where w_u and w_c are weighting coefficients that balance the contribution of each geometric criterion. The uniformity of residue spacing is evaluated to ensure well-distributed interface contacts:

$$\mathcal{L}_{\text{uniformity}} = \frac{\sigma(d_{ij})}{\mu(d_{ij})}, \quad \forall (i,j) \in \mathcal{I},\ d_{ij} < d_{\text{cutoff}}$$

where d_{ij} represents pairwise distances between interface residues $\mathcal{I}$, and $\sigma(\cdot)$, $\mu(\cdot)$ denote standard deviation and mean, respectively. This metric prevents local clustering or sparse regions at the interface. The cavity penalty identifies isolated residues lacking sufficient neighbors:

$$\mathcal{L}_{\text{cavity}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \mathbf{1}[\mathcal{N}_i < n_{\text{threshold}}]$$

where $\mathcal{N}_i$ counts neighbors within radius r_{neighbor}, and $\mathbf{1}[\cdot]$ is the indicator function that equals 1 if the condition is true and 0 otherwise. This penalty prevents the formation of hydrophobic cavities and ensures a well-packed interface.

5.4 Adaptive Expert Routing

Our routing mechanism dynamically activates experts based on real-time structural analysis, ensuring that computational resources focus on the most critical problems while maintaining SE(3)-equivariance. The system computes problem severity scores:

$$s_i = f_i(\text{metrics}_t) \in [0, 1]$$

where f_i evaluates current structural metrics to determine the severity of problems relevant to expert i.

The activation is problem-driven rather than time-dependent. Experts are weighted proportionally to problem severity:

$$w_i(t) = \frac{s_i}{\sum_{j: s_j > \theta_{\min}} s_j} \cdot \mathbf{1}[s_i > \theta_{\min}]$$

This adaptive weighting ensures that only experts addressing detected problems are activated, preventing unnecessary or conflicting guidance.

Critically, SE(3)-equivariance is preserved because all expert gradients are computed from invariant geometric features. Each expert's gradient satisfies:

$$\nabla \mathcal{L}_i(\mathbf{R}\mathbf{T} + \mathbf{t}) = \mathbf{R} \nabla \mathcal{L}_i(\mathbf{T})$$

for any rotation $\mathbf{R} \in \text{SO}(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$. This is achieved by basing all computations on pairwise distances, relative orientations, and other invariant descriptors rather than absolute coordinates. The weighted combination of equivariant gradients maintains equivariance: $\mathbf{g}(\mathbf{T}_t, t) = \sum_i w_i(t) \cdot \nabla \mathcal{L}_i(\mathbf{T}_t)$ remains equivariant since weights $w_i(t)$ are scalar functions of invariant metrics.

5.5 Online Parameter Learning

We employ Bayesian optimization with Gaussian processes (GP) to adaptively learn the optimal Beta distribution parameters for each antibody-antigen system post-generation. Our approach models a two-dimensional parameter space $\boldsymbol{\theta} = (\alpha, \beta) \in \Theta$, where α and β are the shape parameters controlling the temporal modulation profile. The optimization process maintains a probabilistic surrogate model of the objective function using a Gaussian process:

$$f(\boldsymbol{\theta}) \sim \mathcal{GP}(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}'))$$

where $\mu(\boldsymbol{\theta})$ represents the predicted mean performance and $k(\boldsymbol{\theta}, \boldsymbol{\theta}')$ is the covariance function. We employ a Matérn kernel with automatic noise estimation:

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = k_{\text{Matérn}}(\boldsymbol{\theta}, \boldsymbol{\theta}') + \sigma_n^2(\boldsymbol{\theta})\delta_{ij}$$

After each design evaluation, the GP posterior is updated using the observed loss:

$$L(\boldsymbol{\theta}) = \sum_{m=1}^{M} \omega_m \cdot \frac{\text{metric}_m}{\nu_m}$$

where the metrics include CDR-H3 backbone RMSD, predicted aligned error (pAE), and interaction pAE (ipAE), each normalized by appropriate constants ν_m.

The next parameter configuration is selected by maximizing the Expected Improvement (EI) acquisition function:

$$\boldsymbol{\theta}_{n+1} = \arg\max_{\boldsymbol{\theta} \in \Theta} \text{EI}(\boldsymbol{\theta})$$

To handle stochastic evaluations, we aggregate results from similar parameters to reduce noise and periodically re-evaluate promising configurations. This approach automatically discovers optimal Beta parameters for each target without manual tuning.

6 Experiments

Datasets We evaluated our framework on six antibody-antigen pairs: IL17A (PDB: 6PPG), ACVR2B (PDB: 5NGV), FXI (PDB: 6HHC), TSLP (PDB: 5J13), IL36R (PDB: 6U6U), and TNFRSF9 (PDB: 6A3W). Prior work including RFAntibody [4] typically evaluated four targets; we expanded this to six therapeutically relevant targets spanning inflammation, thrombosis, and cancer applications (17-40 kDa). This extended benchmark from IgDesign [29] includes 1,243 SPR-validated designs across diverse binding modes. We annotated 5 hotspot residues per target and generated 1,000 designs each, following the field's established protocol for rigorous benchmarking of antibody design methods.

Baselines Most recent methods are unavailable due to lack of publicly accessible code. We selected: **DiffAb** [21]: torsion-space CDR generation; **RFAntibody** [4]: the current state-of-the-art method without physics guidance, widely recognized for its exceptional performance across diverse antibody design tasks. For fair comparison, all methods used identical inputs (target structure, hotspot annotations, and Trastuzumab framework).

Evaluation Metrics We employ a comprehensive set of metrics to evaluate antibody design quality across structural accuracy, binding interface characteristics, and biophysical properties. Structural accuracy is measured by **CDR-H3 RMSD** (Å; lower is better). Confidence uses AlphaFold2-Multimer predicted alignment error over the whole complex (**pAE**) and across antibody–antigen residue pairs (**ipAE**; both lower is better), plus **pLDDT** averaged over CDRs (higher is better). Interface quality is assessed by shape complementarity (**SC**; higher), buried surface area (**BSA**; larger), **hotspot coverage** (fraction of epitope hotspots in contact; higher), **CDR participation** (fraction of CDR residues contacting antigen; higher), and a **van der Waals** packing term (more favorable indicates fewer clashes). We define success criteria as meeting three thresholds at once: CDR-H3 RMSD $<$ 3.0 Å, mean pAE $<$ 10.0, and mean interaction pAE (ipAE) $<$ 10.0. This multi-metric evaluation ensures structural accuracy and binding confidence.

Table 1: Implementation details of our adaptive physics-guided framework

Parameter	Value	Parameter	Value
Diffusion & Guidance		*Beta Distribution*	
Timesteps T	50	α range	[0.5, 10.0]
Expert modules	4	β range	[0.5, 10.0]
VDW threshold	2.8 Å	Peak factor λ_{peak}	5.0
Gradient clip	2.0	Initial (α, β)	(2.0, 2.0)
Distance cutoff	8.0 Å		
Expert Activation		*Online Learning*	
Min threshold $\theta_{\min}$	0.1	Kernel	Matérn-5/2
VDW strength	0.5	Length scale	2.0
Recognition strength	2.0	GP noise	0.3
Activation	Adaptive	Acquisition ξ	0.01
Evaluation Metrics		*Implementation*	
RMSD norm ν_1	1.5 Å	Framework	PyTorch
pAE norm ν_2	7.0	GPU memory	48 GB
ipAE norm ν_3	10.0	CDR design	All 6

Table 2: Overall performance comparison across six targets highlighting balanced optimization

Method	Success(%)	H3-RMSD(Å)	pAE	iPAE
DiffAb	11.1±17.0	9.1±10.1	9.4±1.2	10.1±2.2
RFAntibody	36.7±14.0	2.1±1.8	7.1±0.5	8.5±1.7
Ours	**41.1±10.5**	**1.9±1.4**	**6.9±0.4**	**8.1±1.3**
Method	SC	BSA	VDW(kcal/mol)	pLDDT
DiffAb	0.55±0.14	**5277±2341**	15.2±9.7	0.886±0.01
RFAntibody	0.55±0.03	3583±888	−1.2±1.3	**0.896±0.01**
Ours	**0.58±0.02**	3870±780	**−1.3±1.0**	0.893±0.01
Method	Hotspot(%)	CDR Int.(%)	Seq. Div.	Elec. Energy
DiffAb	16.2±5.5	0.0±0.0	0.46±0.08	**−6.3±11.1**
RFAntibody	48.5±18.1	54.7±13.6	**0.53±0.00**	−4.0±7.1
Ours	**57.5±12.6**	**61.3±11.2**	0.53±0.00	−5.6±6.1

Implementation Details All experiments used 4 NVIDIA RTX A6000 GPUs with the following settings in Table 1.

6.1 Main Results

Table 2 presents comprehensive evaluation across all targets. Our method achieved improvements in most metrics while maintaining a more balanced performance profile compared to both baselines, effectively addressing the critical "weakest link" problem in antibody design where poor performance in any single metric can compromise experimental viability. As shown in Figure 2, our adaptive guidance consistently improves performance across all key metrics (CDR-H3 RMSD, pAE, ipAE) for multiple targets, with notably reduced variance in the distributions.

6.1.1 Online Learning Impact

To evaluate the contribution of our online learning framework, we conducted ablation studies comparing three configurations: (1) the full system with both multi-expert guidance and online

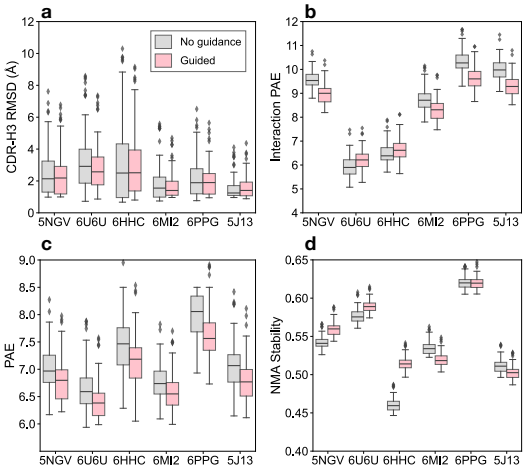

Figure 2: Adaptive guidance improves antibody design quality across multiple metrics. Box plots compare baseline RFAntibody (gray) with our guided approach (pink) on six targets for (a) CDR-H3 RMSD, (b) interaction pAE, (c) mean pAE, and (d) NMA stability.

Table 3: Ablation study demonstrating the importance of online learning

Config	Success	H3-RMSD	pAE	iPAE
RFAntibody	36.7±14.0	2.11±1.81	7.13±0.46	8.49±1.71
+ Expert	38.9±12.5	2.01±1.58	6.95±0.42	8.28±1.65
+ Full	**41.1±10.5**	**1.94±1.36**	**6.85±0.39**	**8.12±1.31**
Config	SC	BSA	VDW	pLDDT
RFAntibody	0.549±.026	3583±888	−1.18±1.27	0.896±.011
+ Expert	0.562±.026	3712±885	−1.25±1.15	0.885±.014
+ Full	**0.576±.024**	**3870±780**	**−1.33±0.95**	**0.893±.008**
Config	Hotspot	CDR	Seq.D	Elec.E
RFAntibody	48.5±18.1	54.7±13.6	0.534±.004	−3.96±7.1
+ Expert	55.6±15.2	60.8±13.5	0.532±.005	−5.31±8.1
+ Full	**57.5±12.6**	**61.3±11.2**	**0.532±.004**	**−5.64±6.1**

learning, (2) multi-expert guidance with fixed parameters (no online learning), and (3) the baseline RFAntibody without any guidance.

The results reveal the complementary benefits of multi-expert guidance and online learning:

Multi-Expert Guidance Alone: Adding physics-based guidance to RFAntibody improves the success rate from 36.7% to 38.9% (Table 3). While all metrics show improvement, the gains are modest and the standard deviations remain relatively high. This suggests that fixed guidance parameters, though helpful, cannot adapt to the diverse characteristics of different antibody-antigen systems.

Online Learning Enhancement: The addition of online parameter optimization yields substantial improvements across all metrics. Most notably, the standard deviations decrease significantly, indicating more consistent and reliable design quality. This variance reduction is crucial for practical applications where predictability is as important as average performance.

6.1.2 Parameter Adaptation Analysis

Our online learning framework employs Gaussian Process-based Bayesian optimization to discover optimal Beta distribution parameters for each antibody-antigen system (Table 4).

Table 4: Target-specific Beta distribution parameters discovered through Bayesian optimization

Target Type	Beta Params		Key Expert Emphasis
	α	β	
Small epitopes (IL17A, IL36R)	3.2±0.4	1.8±0.3	VDW + Recog.
Large interfaces (ACVR2B, TNFRSF9)	2.1±0.3	3.5±0.5	Interface + Energy
Buried epitopes (FXI)	2.8±0.4	2.2±0.3	Recog. + Energy
Flexible targets (TSLP)	2.0±0.5	2.0±0.5	All experts balanced

7 Conclusion

We presented an adaptive physics-guided framework that mimics B cell affinity maturation to optimize antibody design through online learning. Our key contributions include: (1) formulating antibody design as an adaptive process inspired by B cell affinity maturation, (2) pioneering online optimization for diffusion model guidance without retraining, and (3) developing a multi-expert system that achieves balanced optimization across competing objectives, mirroring natural selection pressures.

The biological significance extends beyond technical metrics. Like B cells iteratively refining antibodies through mutation and selection, our framework continuously learns from each design attempt. This parallel to immune evolution enables rapid adaptation to new antigens—critical for emerging pathogen response and personalized therapeutics. By integrating biological principles with machine learning, we bridge immunology and AI to accelerate therapeutic development.

References

[1] Brennan Abanades, Guy Georges, Alexander Bujotzek, and Charlotte M. Deane. ABlooper: Fast accurate antibody CDR loop structure prediction with accuracy estimation. *Bioinformatics*, 38(7):1877–1880, 2022.

[2] Rahmad Akbar, Philippe A Robert, Chloe R Weber, Michael Widrich, Robert Frank, Milena Pavlović, Lonneke Scheffer, Maria Chernigovskaya, Igor Snapkov, Andrei Slabodkin, et al. Progress and challenges for the machine learning-based design of fit-for-purpose monoclonal antibodies. *mAbs*, 14(1):2008790, 2022.

[3] Minkyung Baek, Ivan Anishchenko, Ian R Humphreys, Qian Cong, David Baker, and Frank DiMaio. Efficient and accurate prediction of protein structure using rosettafold2. *BioRxiv*, pages 2023–05, 2023.

[4] Nathaniel R Bennett, Joseph L Watson, Robert J Ragotte, Andrew J Borst, DéJenaé L See, Connor Weidle, Riti Biswas, Yutong Yu, Ellen L Shrock, Russell Ault, et al. Atomically accurate de novo design of antibodies with rfdiffusion. *bioRxiv*, pages 2024–03, 2025.

[5] Yoonjoo Choi, Christian Ndong, Karl E Griswold, Chris Bailey-Kellogg, and Charlotte M Deane. Predicting antibody developability profiles through early stage discovery screening. *mAbs*, 10(8):1225–1239, 2018.

[6] Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. In *International Conference on Learning Representations*, 2023.

[7] Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J Ragotte, Lukas F Milles, Basile IM Wicky, Alexis Courbet, Rob J de Haas, Neville Bethel, et al. Robust deep learning–based protein sequence design using proteinmpnn. *Science*, 378(6615):49–56, 2022.

[8] Raphael R Eguchi, Namrata Anand, Christian A Choe, and Po-Ssu Huang. Ig-vae: Generative modeling of protein structure by direct 3d coordinate generation. *PLoS Computational Biology*, 18(6):e1010271, 2022.

[9] Emil Fischer. Einfluss der configuration auf die wirkung der enzyme. *Berichte der deutschen chemischen Gesellschaft*, 27(3):2985–2993, 1894.

[10] Sarel Fischman and Yanay Ofran. Computational design of antibodies. *Current Opinion in Structural Biology*, 51:156–162, 2018.

[11] Pablo Gainza, Freyr Sverrisson, Federico Monti, Emanuele Rodolà, Davide Boscaini, Michael M Bronstein, and Bruno E Correia. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192, 2020.

[12] Michael K Gilson and Huan-Xiang Zhou. Calculation of protein-ligand binding affinities. *Annual review of biophysics and biomolecular structure*, 36:21–42, 2007.

[13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[14] Emiel Hoogeboom, Vıctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International Conference on Machine Learning*, pages 8867–8887. PMLR, 2022.

[15] Chloe Hsu, Robert Verkuil, Jason Liu, Zeming Lin, Brian Hie, Tom Sercu, Adam Lerer, and Alexander Rives. Learning inverse folding from millions of predicted structures. In *International Conference on Machine Learning*, pages 8946–8970. PMLR, 2022.

[16] John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-based protein design. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[17] Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi Jaakkola. Iterative refinement graph neural network for antibody sequence-structure co-design. In *International Conference on Learning Representations*, 2022.

[18] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

[19] Kristian W. Kaufmann, Gordon H. Lemmon, Samuel L. DeLuca, Jonathan H. Sheehan, and Jens Meiler. Practically useful: what the Rosetta protein modeling suite can do for you. *Biochemistry*, 49(14):2987–2998, 2010.

[20] Daniel E Koshland. Application of a theory of enzyme specificity to protein synthesis. *Proceedings of the National Academy of Sciences*, 44(2):98–104, 1958.

[21] 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.

[22] Derek M Mason, Simon Friedensohn, Cédric R Weber, Christian Jordi, Bastian Wagner, Simon M Meng, Roy A Ehling, Lucia Bonati, Jan Dahinden, Pablo Gainza, et al. Optimization of therapeutic antibodies by predicting antigen specificity from antibody sequence via deep learning. *Nature Biomedical Engineering*, 5(6):600–612, 2021.

[23] Richard A Norman, Francesco Ambrosetti, Alexandre MJJ Bonvin, Lucy J Colwell, Sebastian Kelm, Sandeep Kumar, and Konrad Krawczyk. Computational approaches to therapeutic antibody design: established methods and emerging trends. *Immunological Reviews*, 284(1): 65–96, 2020.

[24] Myron Olympiou, Arjun Sarkar, Kristyna V Lexa, Elizabeth Twomey, Dima Kozakov, and Sandor Vajda. Antibody interface prediction with 3d zernike descriptors and svm. *Bioinformatics*, 38(7):1881–1887, 2022.

[25] Matthew IJ Raybould, Claire Marks, Konrad Krawczyk, Bruck Taddese, Jaroslaw Nowak, Alan P Lewis, Alexander Bujotzek, Jiye Shi, and Charlotte M Deane. Five computational developability guidelines for therapeutic antibody profiling. *Proceedings of the National Academy of Sciences*, 116(10):4025–4030, 2019.

[26] Jeffrey A. Ruffolo, Jeremias Sulam, and Jeffrey J. Gray. DeepAb: Antibody structure prediction using interpretable deep learning. *Patterns*, 3(2):100406, 2022.

[27] Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In *International Conference on Machine Learning*, pages 9323–9332. PMLR, 2021.

[28] Constantin Schneider, Andrew Buchanan, Bruck Taddese, and Charlotte M Deane. Dlab: deep learning methods for structure-based virtual screening of antibodies. *Bioinformatics*, 38(2): 377–383, 2022.

[29] Amir Shanehsazzadeh, Matt McPartlon, George Kasun, Andrea K Steiger, John M Sutton, Edriss Yassine, Rylee Shuai, Cailen Kohnert, Matthew Osborn, Shaheed Mustafa, et al. Unlocking de novo antibody design with generative artificial intelligence. *bioRxiv*, pages 2023–01, 2023.

[30] Yuchen Shen, Chenhao Zhang, Sijie Fu, Chenghui Zhou, Newell Washburn, and Barnabás Póczos. Chemistry-inspired diffusion with non-differentiable guidance. *arXiv preprint arXiv:2410.06502*, 2024.

[31] Richard W Shuai, Jeffrey A Ruffolo, and Jeffrey J Gray. Generative language modeling for antibody design. *bioRxiv*, page 2021.12.13.472419, 2021.

[32] Brian L Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and Tommi Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. In *International Conference on Learning Representations*, 2023.

[33] 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.

[34] Benjamin Webb and Andrej Sali. Comparative protein structure modeling using MODELLER. *Current Protocols in Bioinformatics*, 54:5–6, 2016.

[35] Li C Xue, Drena Dobbs, Alexandre MJJ Bonvin, and Vasant Honavar. Computational prediction of protein interfaces: A review of data driven methods. *FEBS Letters*, 589(23):3516–3526, 2015.

[36] Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023.