Guided Generation for Developable Antibodies

Siqi Zhao [1] **Joshua Moller** [1] **Porfirio Quintero-Cadena** [1] **Lood van Niekerk** [1]

Abstract

Therapeutic antibodies require not only high-affinity target engagement, but also favorable manufacturability, stability, and safety profiles for clinical effectiveness. These properties are collectively called 'developability'. To enable a computational framework for optimizing antibody sequences for favorable developability, we introduce a guided discrete diffusion model trained on natural paired heavy- and light-chain sequences from the Observed Antibody Space (OAS) (Olsen et al., 2022) and quantitative developability measurements for 246 clinical-stage antibodies. To steer generation toward biophysically viable candidates, we integrate a Soft Value-based Decoding in Diffusion (SVDD) Module that biases sampling without compromising naturalness. In unconstrained sampling, our model reproduces global features of both the natural repertoire and approved therapeutics, and under SVDD guidance we achieve significant enrichment in predicted developability scores over unguided baselines. When combined with high-throughput developability assays, this framework enables an iterative, ML-driven pipeline for designing antibodies that satisfy binding and biophysical criteria in tandem.

1. Introduction

Therapeutic antibodies are pivotal biomolecules with applications spanning oncology (Paul et al., 2024), autoimmune disorders (Chan & Carter, 2010), infectious diseases (Sparrow et al., 2017), and metabolic conditions (Lu et al., 2020). Beyond high-affinity target binding, a developable antibody must also exhibit favorable manufacturability, formulation stability, and safety profiles to support scalable production and reliable delivery (Jain et al., 2017; Carter & Rajpal, 2022). Although high-throughput screens and *in silico* tools can identify and design candidates for binding affinity (Agarwal et al., 2024; Frey et al., 2025), comprehensive machine-driven frameworks for optimizing key developability attributes are still lacking. There is therefore an urgent need for computational methods that not only predict developability properties to triage high-quality binders for downstream validation, but also guide the redesign of existing antibody sequences toward improved developability.

Contemporary *in silico* approaches typically benchmark candidate antibodies' biophysical attributes against those of clinically approved therapeutics (Raybould et al., 2019; 2024; Park & Izadi, 2024). However, these comparative metrics often overlook the inherent variability within approved antibody repertoires and, in the absence of true negative controls, cannot establish meaningful developability thresholds. Moreover, because experimental workflows almost invariably screen solely for target binding – omitting parallel assessments of manufacturability, solubility, or stability – there is a critical need for computational methods that optimize sequences for favorable developability properties. Other model-based optimization frameworks have been proposed (Sinai et al., 2020; Stanton et al., 2022; Gruver et al., 2023; Reddy et al., 2024), and some specific to antibodies. Recently, another study (Wang et al., 2025) demonstrated the usefulness of using physical descriptors derived from predicted structures from protein language models to guide antibody design.

As a step towards addressing this challenge, we developed a machine learning-guided generative framework anchored on our newly published developability dataset. We trained our generative model on natural paired heavy- and light-chain sequences from the Observed Antibody Space (OAS) and built quantitative regressors using the comprehensive developability measurements reported in (Arsiwala et al., 2025) for 246 antibodies spanning clinical use and trial stages. We take a complementary approach to directly use experimentally generated developability measurements from diverse clinically approved antibodies to guide antibody design. We show that, without any conditional input, our model can produce novel sequences that mirror both natural repertoire diversity and features of clinically approved antibodies. Moreover, by integrating a Soft Value-based Decoding in Diffusion (SVDD) (Li et al., 2024) guidance module, we can bias generation toward candidates with predicted favorable

[1]Ginkgo Bioworks, Boston, USA. Correspondence to: Joshua Moller <jmoller@ginkgobioworks.com>, Lood van Niekerk <lvanniekerk@ginkgobioworks.com>.

Proceedings of the Workshop on Generative AI for Biology at the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

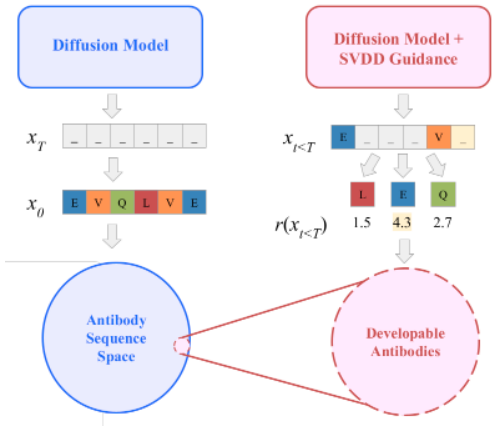

Figure 1. **Overview of our generative framework for novel antibodies.** This framework incorporates an ESM2-based diffusion model trained with paired antibodies sequences from OAS. For guided generation, soft value-based decoding in diffusion (SVDD) was used with developability predictors trained with data from (Arsiwala et al., 2025).

developability. By employing a derivative-free guidance approach, we established a flexible framework that is compatible with various types of predictors. When paired with high-throughput automated assays, this framework offers a powerful avenue to design therapeutic antibodies that meet both affinity and biophysical criteria.

1.1. Generative model

We trained an antibody-specific masked discrete diffusion model on the Observed Antibody Space (OAS) database of sequenced antibodies from over 80 studies. Specifically, we used the order-agnostic diffusion model (OADM) (Hoogeboom et al., 2021) training objective (which had previously been used in EvoDiff (Alamdari et al., 2023)). For simplicity, we re-trained an ESM-2 (8M) (Lin et al., 2023) architecture using this objective. The OAS dataset contains ∼2.4 billion unpaired and ∼1.8 million paired sequences collected from human B-cell sequencing. We used the AntiRef-90 (Briney, 2023) version of OAS to remove partial sequences and clustered using MMseqs2 (Steinegger & Söding, 2017) at a 90% threshold which yielded 1.68 million clusters.

During training we concatenated the paired heavy and light chains into a single sequence `<heavy>|<light>` by introducing a pipe token (`|`). During generation, we sampled a sequence length for the concatenated heavy and light chains from the training dataset and generated sequences according to that length. Usually order-agnostic models are decoded in random orders, but we found that decoding according to minimum entropy positions yielded better sequences. We also experimented with different sampling temperatures in the softmax function according to the formula $p(x_i) = \frac{e^{\frac{x_i}{T}}}{\sum_{j=1}^{D} e^{\frac{x_j}{T}}}$. Identically to language models, higher temperatures correspond to more diverse generations and lower temperatures approach greedy/deterministic samples for each position.

1.2. Developability dataset

Recently, we released a dataset that contains biophysical assay measurements for 9 antibody developability properties across 246 clinical antibodies (Arsiwala et al., 2025)[1], greatly expanding on the set of 137 clinical antibodies provided in (Jain et al., 2017), and running these assays at higher throughput. The 2017 dataset catalyzed the development of many developability predictors, so we sought to train new predictors on the bigger dataset and demonstrate whether we could apply SVDD guidance using these predictors as oracles.

1.3. Predictive model

As a simple oracle, we trained ridge regression models on top of ESM-2 embeddings to predict antibody developability properties using the aforementioned dataset. Heavy and light chain sequences were passed independently to ESM to generate mean-pooled embeddings, which were then concatenated and standardized. Finally, a ridge regression with alpha=0.1 was used. To evaluate performance, we performed hierarchical clustering to separate the 246 sequences into 5 roughly-equal sized folds maximally separated by pairwise sequence identity, and provide the Spearman and Pearson correlation statistics in Table 1.

For this work, we focus on two biophysical properties when training our predictors: hydrophobicity (measured by hydrophobic interaction chromatography, HIC) and self-association (measured by affinity capture self-interacting nanoparticle spectroscopy, AC-SINS, at pH 7.4). We selected these properties for two reasons. First, they directly impact the administrability of candidate antibodies. Second, they present a realistic scenario in which predictors suffer from limited performance due to data scarcity, allowing us to assess whether guidance compromises the naturalness of sequence generation.

1.4. Guided generation using the predictive models

To bias sequence generation toward favorable developability characteristics, we paired SVDD with our generative model. At each step, SVDD assesses several intermediate samples ("branches") using dedicated scoring models and selects the one with the highest composite score. We approximate the

[1]Dataset available at https://huggingface.co/datasets/ginkgo-datapoints/GDPa1

Assay	Spearman's ρ	Pearson's R
HAC RT	0.74 ± 0.22	0.80 ± 0.24
SEC %Monomer	0.54 ± 0.05	0.81 ± 0.14
AC-SINS pH 7.4	0.49 ± 0.09	0.49 ± 0.12
HIC RT	0.42 ± 0.09	0.34 ± 0.08
CHO PR Score	0.41 ± 0.09	0.41 ± 0.12
Ova PR Score	0.40 ± 0.09	0.28 ± 0.17
AC-SINS pH 6.0	0.38 ± 0.09	0.33 ± 0.17
Tm1	0.35 ± 0.16	0.32 ± 0.19
Tm2	0.20 ± 0.14	0.33 ± 0.17
SMAC RT	0.18 ± 0.17	0.00 ± 0.14
Titer	0.17 ± 0.14	0.17 ± 0.19
Purity %LC+HC	0.10 ± 0.21	0.11 ± 0.18

Table 1. Results of training an ESM-2 embedding-based regression models (oracles) on all of the properties in the dataset, as assessed by the average Spearman and Pearson correlations across 5 folds of cross-validation. Assays are ordered by Spearman rank correlation coefficient. The two chosen properties used for guidance in this work are denoted by bold.

SVDD value function via a posterior mean estimate (Figure 1): for a given position, we sample multiple branches, fully denoise each, and compute scores by equally weighting the negative normalized AC-SINS and HIC measurements. The branch whose denoised sequence achieves the top score is retained as the current state, and the process repeats. This approach obviates the need to retrain predictors on partially masked data at varying masking levels and accommodates non-differentiable scoring functions.

2. Results

2.1. Diversity/Naturalness

We first evaluated our unconditional generation model's ability to produce plausible antibodies (naturalness), and the diversity of the generated sequences. To assess naturalness, we scored generated sequences using two independent antibody-specific language models, AbLang2 (Olsen et al., 2024) and p-IgGen (Turnbull et al., 2024), under different sampling temperatures. Both models are able to take paired heavy- and light-chain antibody sequences as input. To quantify diversity, we constructed sequence-similarity networks comparing generated sequences to 246 clinical antibodies in the training set, and to natural antibodies from the Observed Antibody Space (OAS) database, employing MMseqs2 (Steinegger & Söding, 2017). We observed that higher sampling temperatures yield more diverse sequences (Figure 2a) at the expense of lower log-likelihood (naturalness) scores (Figure 2b). At $T = 1.0$, we found a favorable trade-off: generated sequences are diverse while retaining high log-likelihood scores (Figure 2a,b). This observation is also evident from the corresponding Sequence Similarity

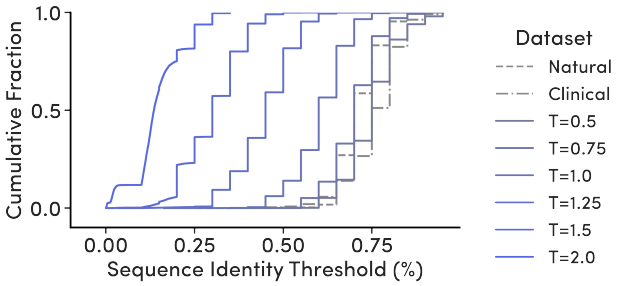

(a) Empirical Cumulative Distribution Functions (ECDFs) of pairwise sequence similarities for Natural and Clinical antibodies, and sequences generated at different sampling temperatures.

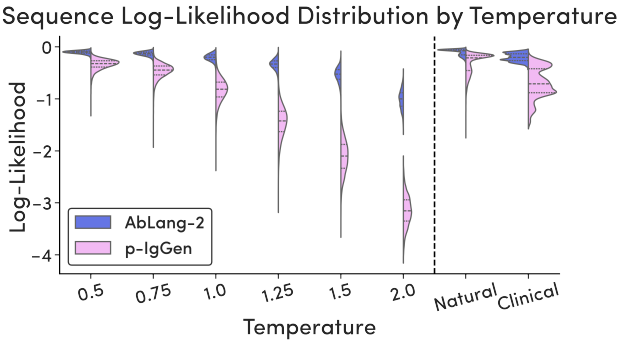

(b) Violin plot of sequence naturalness scored with AbLang-2 and p-IgGen

Figure 2. **Tunable generation of diverse antibodies.**

Network (SSN), and in contrast with sequences generated at T=2 (Figure 3). Interestingly, the natural and clinically approved antibodies are distributed throughout the sequence similarity network rather than being confined to tight clusters. This dispersion likely reflects the heterogeneous development trajectories of clinical candidates. Moreover, this sequence diversity in the training set of our predictive models, in addition to their robust cross-validation performance, supports that these models are well suited to guide antibody design.

2.2. Guided Generation

We configured SVDD to guide generation using both hydrophobicity (HIC RT) and self-association (AC SINS) predictors. In total, we generated 13,519 antibody sequences with $T = 1.0$ across two experiments, masking the 1) Complementarity Determining Region of the Heavy chain (HCDR3), a highly variable sequence that influences developability and binding; and 2) the framework regions of both chains, more conserved sequences that provide a structural scaffold for CDRs. We used the antibodies characterized in (Arsiwala et al., 2025) as starting template sequences for generation. Figure 4a shows the joint distribution of hydrophobicity and self-association scores for the 1) unconditionally generated, 2) framework guided, 3) HCDR3

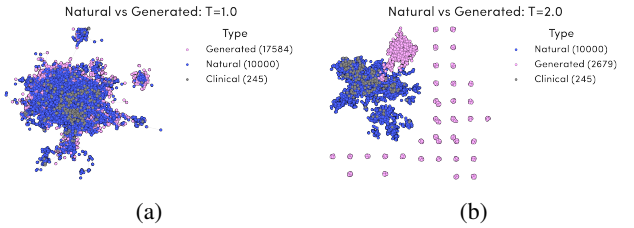

Figure 3. **Sequence Similarity Networks of Generated Antibodies with Natural and Clinical Antibodies.** a) The generated sequences cover a large region of natural sequence space, and clinical candidates are interspersed in that space. b) At high sampling temperatures, sequences are highly diverse but no longer overlap closely with natural and clinical sets. This also leads to many small disconnected clusters shown in pink.

guided, and 4) the original clinical sequences used as templates. We indicate the 10th percentile of the measured clinical sequences as dotted line references. For both scores, our guided generation successfully enriched the number of sequences in the lower left quadrant, from 3.3% in the seed sequence set, to 10.5% and 8.3% for framework and HCDR3 generation, respectively. These values were also higher than unconditional generation. In addition, generated guided sequences preserved the naturalness of unconditional and template clinical sequences (Figure 5).

3. Discussion and Future Works

In this work, we demonstrate that discrete diffusion–based models can generate plausible and diverse antibody sequences. We find that while high sampling temperatures increase diversity, they also degrade naturalness. When coupled with simple biophysical predictors, these models can effectively explore novel regions of therapeutic sequence space while preserving naturalness. A natural next step is developing better predictors for biophysical properties. Many antibody biophysical property predictors have been developed (Jain et al., 2017; Tomar et al., 2017; Khurana et al., 2018; Thumuluri et al., 2021; Zhou et al., 2022; Lai, 2022; Gentiluomo et al., 2020; Prihoda et al., 2022; Schmitt et al., 2023; Wu et al., 2025; Park & Izadi, 2024; Rollins et al., 2024) but restrictive commercial terms, difficulty installing, inconsistent benchmark comparisons or insufficient performance limit their use. Additionally, developing predictors across highly diverse sequences has been found to be much more challenging than predicting local mutational effects (Groth et al., 2023; Notin et al., 2023), but with larger amounts of high quality training data we are confident that generalized predictors across natural and clinical antibody space (and eventually, across different formats such as Fc-fusions, single domain antibodies (VHH, scFv, monobodies), and multispecifics) could achieve much higher accuracy and hence improve guided generation quality.

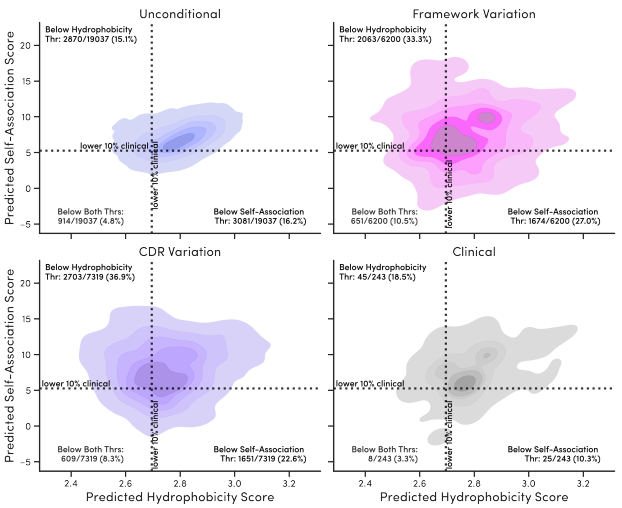

Figure 4. **Density plot of predicted self-association and hydrophobicity for various generation conditions.** Thresholds are calculated based on the lower 10% of predicted properties of clinical antibodies. We show that we can increase the proportion of antibodies in the lower left quadrant (high developability) for both properties simultaneously using guided generation.

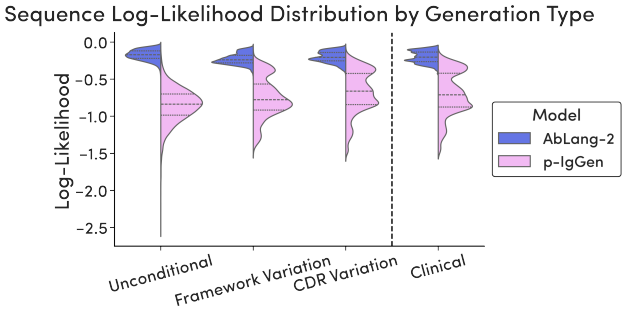

Figure 5. **Generated antibodies have high naturalness.** Violin plot of generated sequences scored with Ablang-2 and p-IgGen

To advance this framework, we plan to incorporate larger, more diverse labeled datasets covering additional biophysical assays and to investigate multi-task predictors that jointly model target binding as well as developability. Extending our guidance beyond posterior-mean SVDD may enable stronger steering with fewer denoising steps. We will also validate generated candidates experimentally to quantify real-world developability gains and refine our predictors accordingly.

Impact Statement

Antibodies comprise a major modality of therapeutics, and have been used to treat cancer, autoimmune diseases and infectious diseases. We expect that machine learning research focused on optimizing developability will have a positive

social impact in reducing clinical trial costs and enabling more therapies to come to market.

References

Agarwal, A. A., Harrang, J., Noble, D., McGowan, K. L., Lange, A. W., Engelhart, E., Lahman, M. C., Adamo, J., Yu, X., Serang, O., Minch, K. J., Wellman, K. Y., Younger, D. A., Lopez, R. M., and Emerson, R. O. AlphaBind, a domain-specific model to predict and optimize antibody-antigen binding affinity. *bioRxiv*, pp. 2024.11.11.622872, November 2024.

Alamdari, S., Thakkar, N., van den Berg, R., Lu, A. X., Fusi, N., Amini, A. P., and Yang, K. K. Protein generation with evolutionary diffusion: sequence is all you need. *bioRxiv*, pp. 2023.09.11.556673, September 2023.

Arsiwala, A., Bhatt, R., Yang, Y., Quintero Cadena, P., Anderson, K. C., Ao, X., van Niekerk, L., Rosenbaum, A., Bhatt, A., Smith, A., Grippo, L., Cao, X., Cohen, R., Patel, J., Allen, O., Faraj, A., Nandy, A., Hocking, J., Tural, B., Salvador, S., Jacobowitz, J., Schaven, K., Sherman, M., Shah, S., Tessier, P. M., and Borhani, D. A high-throughput platform for biophysical antibody developability assessment to enable AI/ML model training. *bioRxiv*, May 2025.

Briney, B. AntiRef: reference clusters of human antibody sequences. *Bioinform. Adv.*, 3(1):vbad109, August 2023.

Carter, P. J. and Rajpal, A. Designing antibodies as therapeutics. *Cell*, 185(15):2789–2805, July 2022.

Chan, A. C. and Carter, P. J. Therapeutic antibodies for autoimmunity and inflammation. *Nat. Rev. Immunol.*, 10 (5):301–316, May 2010.

Frey, N. C., Hotzel, I., Stanton, S. D., Kelly, R. L., Alberstein, R. G., Makowski, E. K., Martinkus, K., Berenberg, D., Bevers, III, J., Bryson, T., Chan, P., Czubaty, A., D'Souza, T. A., Dwyer, H., Dziewulska, A., Fairman, J. W., Goodman, A., Hofmann, J. L., Isaacson, H. H., Ismail, A. A., James, S., Joren, T., Kelow, S. P., Kiefer, J. R., Kirchmeyer, M., Kleinhenz, J., Koerber, J. T., Lafrance-Vanasse, J., Leaver-Fay, A., Lee, J. H., Lee, E., Lee, D. W., Liang, W.-C., Lin, J. Y.-Y., Lisanza, S., Loukas, A., Ludwiczak, J., Mahajan, S. P., Mahmood, O., MohammadiPeyhani, H., Nerli, S., Park, J. W., Park, J., Ra, S., Robinson, S. A., Saremi, S., Seeger, F., Sinha, I., Sokol, A. M., Spiess, C., Tagasovska, N., To, H. V., Wagstaff, E., Wang, A., Watkins, A. M., Wilson, B., Wu, S., Zadorozhny, K., Marioni, J. C., Regev, A., Wu, Y., Cho, K., Bonneau, R., and Gligorijevic, V. Lab-in-the-loop therapeutic antibody design with deep learning. *bioRxiv*, pp. 2025.02.19.639050, February 2025.

Gentiluomo, L., Roessner, D., and Frieß, W. Application of machine learning to predict monomer retention of therapeutic proteins after long term storage. *Int. J. Pharm.*, 577(119039):119039, March 2020.

Groth, P. M., Michael, R., Salomon, J., Tian, P., and Boomsma, W. FLOP: Tasks for fitness landscapes of protein wildtypes. *bioRxiv*, June 2023.

Gruver, N., Stanton, S., Frey, N. C., Rudner, T. G. J., Hotzel, I., Lafrance-Vanasse, J., Rajpal, A., Cho, K., and Wilson, A. G. Protein design with guided discrete diffusion. *arXiv [cs.LG]*, May 2023.

Hoogeboom, E., Gritsenko, A. A., Bastings, J., Poole, B., van den Berg, R., and Salimans, T. Autoregressive diffusion models. *arXiv [cs.LG]*, October 2021.

Jain, T., Sun, T., Durand, S., Hall, A., Houston, N. R., Nett, J. H., Sharkey, B., Bobrowicz, B., Caffry, I., Yu, Y., Cao, Y., Lynaugh, H., Brown, M., Baruah, H., Gray, L. T., Krauland, E. M., Xu, Y., Vásquez, M., and Wittrup, K. D. Biophysical properties of the clinical-stage antibody landscape. *Proc. Natl. Acad. Sci. U. S. A.*, 114(5):944–949, January 2017.

Khurana, S., Rawi, R., Kunji, K., Chuang, G.-Y., Bensmail, H., and Mall, R. DeepSol: a deep learning framework for sequence-based protein solubility prediction. *Bioinformatics*, 34(15):2605–2613, August 2018.

Lai, P.-K. DeepSCM: An efficient convolutional neural network surrogate model for the screening of therapeutic antibody viscosity. *Comput. Struct. Biotechnol. J.*, 20: 2143–2152, April 2022.

Li, X., Zhao, Y., Wang, C., Scalia, G., Eraslan, G., Nair, S., Biancalani, T., Regev, A., Levine, S., and Uehara, M. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. *arXiv [cs.LG]*, August 2024.

Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., Dos Santos Costa, A., Fazel-Zarandi, M., Sercu, T., Candido, S., and Rives, A. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, March 2023.

Lu, R.-M., Hwang, Y.-C., Liu, I.-J., Lee, C.-C., Tsai, H.-Z., Li, H.-J., and Wu, H.-C. Development of therapeutic antibodies for the treatment of diseases. *J. Biomed. Sci.*, 27(1):1, January 2020.

Notin, P., Kollasch, A. W., Ritter, D., Van Niekerk, L., Paul, S., Spinner, H., Rollins, N. J., Shaw, A., Weitzman, R., Frazer, J., Dias, M., Franceschi, D., Orenbuch, R., Gal, Y., and Marks, D. S. ProteinGym: Large-scale benchmarks for protein fitness prediction and design. November 2023.

Olsen, T. H., Boyles, F., and Deane, C. M. Observed antibody space: A diverse database of cleaned, annotated, and translated unpaired and paired antibody sequences. *Protein Sci.*, 31(1):141–146, January 2022.

Olsen, T. H., Moal, I. H., and Deane, C. M. Addressing the antibody germline bias and its effect on language models for improved antibody design. *Bioinformatics*, 40(11), November 2024.

Park, E. and Izadi, S. Molecular surface descriptors to predict antibody developability: sensitivity to parameters, structure models, and conformational sampling. *MAbs*, 16(1):2362788, January 2024.

Paul, S., Konig, M. F., Pardoll, D. M., Bettegowda, C., Papadopoulos, N., Wright, K. M., Gabelli, S. B., Ho, M., van Elsas, A., and Zhou, S. Cancer therapy with antibodies. *Nat. Rev. Cancer*, 24(6):399–426, June 2024.

Prihoda, D., Maamary, J., Waight, A., Juan, V., Fayadat-Dilman, L., Svozil, D., and Bitton, D. A. BioPhi: A platform for antibody design, humanization, and humanness evaluation based on natural antibody repertoires and deep learning. *MAbs*, 14(1):2020203, January 2022.

Raybould, M. I. J., Marks, C., Krawczyk, K., Taddese, B., Nowak, J., Lewis, A. P., Bujotzek, A., Shi, J., and Deane, C. M. Five computational developability guidelines for therapeutic antibody profiling. *Proc. Natl. Acad. Sci. U. S. A.*, 116(10):4025–4030, March 2019.

Raybould, M. I. J., Turnbull, O. M., Suter, A., Guloglu, B., and Deane, C. M. Contextualising the developability risk of antibodies with lambda light chains using enhanced therapeutic antibody profiling. *Commun. Biol.*, 7(1):62, January 2024.

Reddy, A. J., Geng, X., Herschl, M. H., Kolli, S., Kumar, A., Hsu, P. D., Levine, S., and Ioannidis, N. M. Designing cell-type-specific promoter sequences using conservative model-based optimization. *bioRxiv*, June 2024.

Rollins, Z. A., Widatalla, T., Waight, A., Cheng, A. C., and Metwally, E. AbLEF: antibody language ensemble fusion for thermodynamically empowered property predictions. *Bioinformatics*, 40(5), May 2024.

Schmitt, J., Razvi, A., and Grapentin, C. Predictive modeling of concentration-dependent viscosity behavior of monoclonal antibody solutions using artificial neural networks. *MAbs*, 15(1):2169440, January 2023.

Sinai, S., Wang, R., Whatley, A., Slocum, S., and others. AdaLead: A simple and robust adaptive greedy search algorithm for sequence design. *arXiv preprint arXiv*, 2020.

Sparrow, E., Friede, M., Sheikh, M., and Torvaldsen, S. Therapeutic antibodies for infectious diseases. *Bull. World Health Organ.*, 95(3):235–237, March 2017.

Stanton, S., Maddox, W., Gruver, N., Maffettone, P., Delaney, E., Greenside, P., and Wilson, A. G. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. *arXiv [cs.LG]*, March 2022.

Steinegger, M. and Söding, J. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat. Biotechnol.*, 35(11):1026–1028, November 2017.

Thumuluri, V., Martiny, H.-M., Armenteros, J. J. A., Salomon, J., Nielsen, H., and Johansen, A. R. NetSolP: predicting protein solubility in E. coli using language models. *Bioinformatics*, 38(4):941–946, November 2021.

Tomar, D. S., Li, L., Broulidakis, M. P., Luksha, N. G., Burns, C. T., Singh, S. K., and Kumar, S. In-silico prediction of concentration-dependent viscosity curves for monoclonal antibody solutions. *MAbs*, 9(3):476–489, April 2017.

Turnbull, O. M., Oglic, D., Croasdale-Wood, R., and Deane, C. M. p-IgGen: a paired antibody generative language model. *Bioinformatics*, 40(11):btae659, November 2024.

Wang, A., Sang, Z., Stanton, S. D., Hofmann, J. L., Izadi, S., Park, E., Ludwiczak, J., Kirchmeyer, M., Davidson, D., Maier, A., Pritsky, T., Frey, N. C., Watkins, A. M., and Seeger, F. A guided design framework for the optimization of therapeutic-like antibodies. April 2025.

Wu, I.-E., Kalejaye, L., and Lai, P.-K. Machine learning models for predicting monoclonal antibody biophysical properties from molecular dynamics simulations and deep learning-based surface descriptors. *Mol. Pharm.*, 22(1): 142–153, January 2025.

Zhou, Y., Xie, S., Yang, Y., Jiang, L., Liu, S., Li, W., Abagna, H. B., Ning, L., and Huang, J. SSH2.0: A better tool for predicting the hydrophobic interaction risk of monoclonal antibody. *Front. Genet.*, 13:842127, March 2022.