HOW TO MAKE THE MOST OF YOUR MASKED LANGUAGE MODEL FOR PROTEIN ENGINEERING

Calvin McCarter
BigHat Biosciences
cmccarter@bighatbio.com

Nick Bhattacharya
BigHat Biosciences

Sebastian W. Ober
BigHat Biosciences

Hunter Elliott
BigHat Biosciences

arXiv:2603.10302v1 [cs.LG] 11 Mar 2026

ABSTRACT

A plethora of protein language models have been released in recent years. Yet comparatively little work has addressed how to best sample from them to optimize desired biological properties. We fill this gap by proposing a flexible, effective sampling method for masked language models (MLMs), and by systematically evaluating models and methods both in silico and in vitro on actual antibody therapeutics campaigns. Firstly, we propose sampling with stochastic beam search, exploiting the fact that MLMs are remarkably efficient at evaluating the pseudo-perplexity of the entire 1-edit neighborhood of a sequence. Reframing generation in terms of entire-sequence evaluation enables flexible guidance with multiple optimization objectives. Secondly, we report results from our extensive in vitro head-to-head evaluation for the antibody engineering setting. This reveals that choice of sampling method is at least as impactful as the model used, motivating future research into this under-explored area.

1 INTRODUCTION

Antibodies are a class of natural immune proteins that have served as the starting point for the most successful class of biologic drugs (Carter & Lazar, 2018), and are thus a focus of efforts to utilize machine learning to accelerate drug discovery. Developing an antibody therapeutic typically involves first identifying a candidate via animal immunization (Köhler & Milstein, 1975), *in vitro* screening libraries (McCafferty et al., 1990), or *in silico* "de novo" design methods (Bennett et al., 2024). Second, because candidates usually require substantial improvement, *in vitro* iterative optimization is used to refine the candidate into a valid therapeutic (Lu et al., 2020).

The space of possible antibody sequence mutations is combinatorially vast, and only a small number of mutations are typically advantageous; meanwhile, optimization experiments are limited to hundreds of sequences and are costly and slow. Machine learning-guided iterative design improves this process by prioritizing mutated sequences which are likely to yield improvements along desired criteria. In each round, a language model receives top-performing candidates (seeds), and optionally receives additional sequence scoring functions (such as supervised models trained on sequences previously measured in the lab). The generative model, combined with a sampling algorithm, outputs mutated sequences, some of which are selected for the next round of lab measurements. This process is depicted in Figure 1.

Many protein language models (Madani et al., 2020; Lin et al., 2022; Alamdari et al., 2023; Hayes et al., 2025; Bhatnagar et al., 2025) and antibody-specific language models (Shuai et al., 2021; Olsen et al., 2024; Fournier et al., 2024; Turnbull et al., 2024; Burbach & Briney, 2025) have been released in recent years. Because protein sequences lack causal directionality, masked language models (MLMs) rather than causal language models (CLMs) are the typical choice. Depending on the goals of the optimization campaign, the language model chosen may be one trained on a large human antibody dataset (The Observed Antibody Space, OAS; Kovaltsuk et al., 2018), on a database of antibodies with known structures (enriched for therapeutic antibodies, Dunbar et al., 2014), or on a more generic protein sequence database (Consortium, 2015). Regardless of the choice of model,

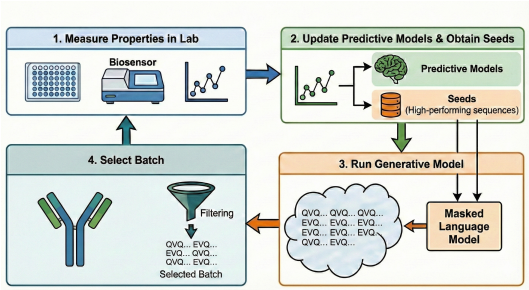

Figure 1: Machine learning-guided iterative design of therapeutic antibodies. This work focuses on Step 3, wherein a MLM is combined with seed sequences and (optionally) predictive models, and proposes mutated sequences.

the premise is that model likelihoods will positively correlate with the probability of the sequence yielding a valid, functional antibody.

The goal of a sampling algorithm is to elicit from the MLM a large, diverse set of sequences with relatively high model likelihood and some notion of proximity to the seed sequences. Unfortunately, systematic evaluation of sampling MLM algorithms is essentially non-existent. Furthermore, previously-adopted MLM sampling algorithms employ mutation-centric iterative sampling procedures, essentially any-order analogues of greedy sampling methods for causal language model generation (Hoogeboom et al., 2021; Sahoo et al., 2024). These methods are computationally costly and (as we will show later) empirically tend to elicit unlikely, dysfunctional sequences.

As mentioned previously, it is also frequently desirable to bias the language model's distribution based on additional scoring functions. Sometimes these are differentiable functions, such as neural networks trained on binding affinity measurements. In other cases these scores, such as the OASis percentile (a risk score for immunogenic adverse events; Prihoda et al., 2022) or isoelectric point (predictive of stability and pharmacokinetics quality), are both non-differentiable and require as input a clean (rather than partially-masked) sequence. Mutation-centric sampling methods poorly handle such scenarios without additional approximations (Raghu et al., 2025) or costly computations (Tang et al., 2025).

In this work, we propose to instead take a sequence-centric approach, where we evaluate full sequences according to the MLM, and search sequence space starting from the seed. We show that this is surprisingly computationally efficient and effective at producing sequences preferred by the MLM and optional scoring functions. Next, we evaluate a variety of models and algorithms, first *in silico* and then *in vitro* on actual antibody therapeutics programs. Surprisingly, we find that choice of sampling algorithm matters at least as much as choice of model. We also report other interesting findings, such as the fact that ESM-2 (Lin et al., 2022), though trained on generic protein sequences, is still very effective for antibody optimization.

1.1 Related Work

CLM sampling algorithms Sampling algorithms for causal language models currently tend to follow the greedy decoding paradigm, making decisions solely based on the distribution for the next token. Argmax decoding often produces repetitive text, so methods like softmax temperature annealing (Ackley et al., 1985) and nucleus (top-p) filtering (Holtzman et al., 2019) are widely used to spread out probability mass among the most likely tokens, while eliminating the least likely tokens. Methods based on (partial) sequence likelihoods, such as beam search (Gimpel et al., 2013; Meister et al., 2020) and stochastic beam search (SBS) (Kool et al., 2019) have been proposed, but are less frequently used. These approaches have complexity $O(L^2)$, where L is the sequence length.

MLM sampling algorithms Sampling from masked language models (MLMs) for protein engineering has a very different design space from that of text CLMs. The chief differences are the lack of a pre-defined left-to-right order, the notion of a seed sequence, and the desired proposal of a diverse batch for subsequent *in silico* filtering and finally *in vitro* measurement. Also, sometimes

specific regions are constrained to be unchanged, such as the complementarity-determining region (CDR) if binding is satisfactory, or the framework region (FR) if only binding requires further modification.

Sampling methods for protein optimization tend to be mutation-centric. Denoising sampling (Hoogeboom et al., 2021) is the closest analogue to CLM greedy decoding, and was adopted for protein optimization by Hayes et al. (2025). To edit up to $E \leq L$ residues, conditioned on the remaining $L - E$ fixed positions, it first masks all E positions then runs unmasking for E iterations. Alternatively, Gibbs sampling (Ackley et al., 1985) has been adopted for protein (Johnson et al., 2021) engineering, and tends to propose sequences which are closer to the seed sequence. Gibbs-like sampling masks-and-unmasks for E iterations; at each iteration 1 position is masked, while the remaining $L - 1$ positions are unmasked. For both autoregressive denoising sampling and Gibbs-like sampling, the sampling step can be replaced with an argmax; this was previously employed for antibody humanization (Gordon et al., 2024). Mutation-centric samplers also require a strategy to choose the next position to decode; for example, the ESM-3 work (Hayes et al., 2025) optimized proteins via selection by either lowest entropy or max probability.

Mutation-centric MLM samplers have complexity $O(EL^3)$ per generated sequence. To obtain N unique sequences, one must run the sampling method at least N times; independent runs are not guaranteed to produce unique sequences.

MLM guidance with scoring functions Prior work on incorporation of additional scoring functions is also limited. The ESM-3 work (Hayes et al., 2025) proposes using derivative-free guidance (Li et al., 2024) but requires scoring functions to accept unclean (partially-masked) sequences. Alternatively, Raghu et al. (2025) does not force scoring functions to evaluate unclean sequences, but to remain compatible with mutation-centric sampling, assumes that the scoring function is a product of experts over positions in the sequence (Schiff et al., 2024). One recent work (Tang et al., 2025) uses Monte Carlo tree search (MCTS) to avoid these constraints by performing a complex search and sampling process that scores many clean sequences to estimate the value of one partially-masked sequence; this is problematic for scores like OASis percentile which are very expensive.

Benchmarking Benchmarking on sampling for protein optimization is scarce. The most extensive sampler benchmarking results for protein optimization have been reported in (Darmawan et al., 2025). However, it analyzed sampling methods for a causal protein language model (Notin et al., 2022), and it involved only *in silico* evaluation. The argmax versions of denoising and Gibbs-like sampling were evaluated *in vitro* in (Gordon et al., 2024); argmax Gibbs was marginally better.

2 Methods

2.1 Background on protein MLMs

Denote the set of 20 amino acids by $\mathcal{A} = \{a_1, a_2, ..., a_{20}\}$. Let $\boldsymbol{x} = [x_1, ..., x_L] \in \mathcal{A}^L$ denote a protein sequence, representing a sequence of L amino acids with each $x_i \in \mathcal{A}$. A masked language model (MLM) takes an input sequence $\boldsymbol{x}$ and outputs an $L \times 20$ matrix $\boldsymbol{Y} = [\boldsymbol{y}_1, \ldots, \boldsymbol{y}_L]^\top$. Each $\boldsymbol{y}_i$ can naively be interpreted as approximate log-probabilities over the set of amino acids, with the "probability" for each residue at location i given by

$$\hat{p}(x_i = a_r | \boldsymbol{x}_{j \neq i}) \approx \frac{\exp(\boldsymbol{Y}_{ir}/\tau)}{\sum_{k=1}^{20} \exp(\boldsymbol{Y}_{ik}/\tau)} = \texttt{softmax}_\tau(\boldsymbol{Y}_{i,:}), \tag{1}$$

where $\tau > 0$ is the temperature of the softmax function. Note that $\boldsymbol{Y}$ was computed from the full sequence $\boldsymbol{x}$ including the value of x_i, so this only approximately holds. This approximation allows a single $O(L^2)$ forward-pass[1] and was used with the Sapiens antibody MLM (Prihoda et al., 2022). This approximation replaces $\hat{p}(x_i = a_r | \boldsymbol{x}_{j \neq i})$ with $\hat{p}(x_i = a_r | \boldsymbol{x})$.

To obtain the true conditional probability distribution $\hat{p}(x_i | \boldsymbol{x}_{j \neq i})$, one instead masks x_i such that $\tilde{\boldsymbol{x}} = [x_1, ..., x_{i-1}, \texttt{MASK}, x_{i+1}, ..., x_L]$, and computes MLM output $\tilde{\boldsymbol{Y}}^{(i)}$, thus obtaining

$$\hat{p}(x_i = a_r | \boldsymbol{x}_{j \neq i}) = \texttt{softmax}_\tau(\tilde{\boldsymbol{Y}}^{(i)}_{i,:}). \tag{2}$$

[1] We assume that MLMs use quadratic complexity scaled dot-product attention.

A separate forward-pass is needed for each i, and each time $L-1$ rows of $\boldsymbol{Y}$ are discarded. Thus each iteration of mutation-centric sampling has cost $O(L^3)$.

Unlike CLMs, MLMs do not offer an efficient way to compute a sequence log-likelihood or perplexity. Instead, a pseudo-log-likelihood (PLL) is computed by averaging over positions $\frac{1}{L}\sum_i \log \hat{p}(x_i = a_r|\boldsymbol{x}_{j\neq i})$, for a total cost of $O(L^3)$; in contrast, the CLM true log-likelihood (LL) costs only $O(L^2)$.

2.2 Temperature-annealed stochastic beam search

Instead of asking the MLM to generate mutations, here we propose to ask the MLM to evaluate sequences via its PLL, which allows us to reduce our task to a search problem. This would appear to be infeasible given the cost of computing the PLL, but it turns out that having computed the PLL for a single sequence, the (approximate) PLLs for all sequences that are one substitution away come for free. Let $\boldsymbol{x}'$ be a sequence that is one substitution away from template $\boldsymbol{x}$ at position k. We then use an approximation to the PLL given by

$$PLL(\boldsymbol{x}';\tau) \approx \sum_{i=1}^{L} \log\left(\texttt{softmax}_\tau(\tilde{\boldsymbol{Y}}^{(i)}_{i,\texttt{index}(x'_i)})\right) \tag{3}$$

where we use the exact conditional probabilities at position k, while using the template's conditional probabilities $\hat{p}(x_i|\mathbf{x}_{j\neq i})$ elsewhere. This latter approximation is known as the wild-type marginal approximation, widely used for zero-shot mutation effect prediction (Meier et al., 2021). More discussion of our proposed approximation is in Appendix A.1.1.

The ability to quickly compute approximate PLLs motivates running beam search for E steps, starting from the initial seed sequence to naturally enforce proximity. Each time we expand a sequence from the beam, it becomes the new template sequence $\boldsymbol{x}$, and we evaluate its PLL, thus limiting the approximation error. With a beam of size B, we incur a cost of $O(BEL^3)$, obtaining $20BEL$ unique sequences. For a realistic scenario where $E = 5, L = 100$ and one desires a large number of sequences pre-filtering, the $20EL\times$ speedup compared to mutation-centric sampling is enormous. And as we will see later, this approach tends to elicit higher-quality sequences from MLMs than mutation-centric sampling.

In practice, we would also like to obtain a diverse set of sequences. We adopt stochastic beam search (Kool et al., 2019), which adds Gumbel noise before ranking (Gumbel, 1954; Maddison et al., 2014) to balance sequence likelihood and diversity. The aforementioned softmax temperature, because it scales the likelihood term but not the Gumbel term, adjusts the balance.

Our approach can be applied to CLM LLs instead of MLM PLLs, if one wants to use a protein or antibody CLM to generate sequences near to a seed sequence, rather than to generate sequences from scratch. However, this requires $20L$ separate LL calculations, for a net $20\times$ cost increase.

2.3 Multi-objective optimization (MOO) with gradient-free guidance

Our proposed search framework treats both the MLM and optional scoring functions as black boxes, neither providing partially-masked sequences nor requesting gradients. Any multi-objective scalarization technique is thus applicable for our framework. Recent works addressing protein multi-objective optimization have employed Pareto non-dominated sorting (NDS) (Chen et al., 2025; Tang et al., 2025; Zhang & Chatterjee, 2025). We consider that, as well as smooth Tchebycheff scalarization (STS) (Lin et al., 2024), which attempts to improve the performance of all objectives, rather than any objective as does NDS. NDS does not admit weights over objectives; in some of our experiments, we employed a weighted extension of STS. See Appendix A.1.2 for more details).

3 Experiments

3.1 *In silico* experiments

We first evaluated a wide variety of combinations of models and sampling algorithms *in silico*. We evaluated nine MLMs: ESM-2 (35M, 150M, and 650M params) (Lin et al., 2022), Sapiens

(500K) (Prihoda et al., 2022), AbLang-2 (45M) (Olsen et al., 2024), AMPLIFY (120M and 350M) (Fournier et al., 2024), DiffAbOpt (Raghu et al., 2025) (a CDR-only model that also conditions on a seed's predicted structure (Passaro et al., 2025)), and an internally-trained 13M parameter BERT-style model (Devlin et al., 2019) trained on (Dunbar et al., 2014); all but ESM-2 were trained on (OAS; Kovaltsuk et al., 2018) antibody sequences. We also evaluated three CLMs: pIgGen and pIgGen-developability (17M) (Turnbull et al., 2024), and CloneLM (377M) (Amin et al., 2024), all of which were trained on antibodies. These methods were evaluated on a real scFv antibody therapeutics program for which we have supervised oracles for predicting synthesizability, binding affinity, and thermostability (melting temperature); we also evaluated OASis percentile and isoelectric point. Each method was given the same set of 10 seeds; we evaluated generated variant sequences with exactly 3 substitutions.

An initial evaluation of sampling algorithms on AbLang2 found that Gibbs sampling outperformed denoising in predicted synthesizability, and produced more diversity than Gibbs-argmax (Appendix Figure 5). We therefore focused our evaluations of MLMs with these samplers, and of CLMs with our beam search. Overall, the AbLang2 and ESM2-650M models distinguished themselves with excellent metrics across a variety of criteria, and our beam search outperformed Gibbs sampling. See Appendix A.2 for further experimental details and results.

3.2 *In vitro* EXPERIMENTS

When choosing a smaller set of models for *in vitro* evaluation, we brought forward AbLang2 and ESM2-650M for their *in silico* performance. We also evaluated Sapiens as a lightweight baseline model, pIgGen-dev as a representative of CLMs, SAbDabMLM due to its non-OAS training data, and DiffAb+ to represent structure-based methods. We applied these models to a fresh FAb antibody therapeutic campaign, evaluating chiefly on their ability to pass synthesizability and binding QC criteria. We used an initial candidate hit as seed; for each method 2/3 of the sequences sent to the wet lab had exactly 4 CDR-only substitutions, and 1/3 had exactly 3 substitutions in any location. (Since DiffAb+ is a CDR-only method, all its variants had exactly 4 CDR-only substitutions.)

Later, after obtaining 729 samples and training a supervised classification model for synthesizability and binding success, we evaluated the effect of combining this with AbLang2 when proposing batches. We use the model for post-MLM filtering ($p_{success} > 0.6$), post-MLM ranking, and finally NDS and STS guidance with post-MLM ranking.

Results, comprising 289 samples across 13 methods, each with at least 21 samples, are shown in Figure 2. The overall success rate of each method is shown in Fig 2A. Consistent with *in silico* results, ESM2-650M and AbLang2 had the best unsupervised success rates. Also, our proposed beam search outperformed Gibbs on all three models where both were tried. On Sapiens, Gibbs-argmax outperformed Gibbs, but it struggled to propose many unique new sequences. We found that using supervision to rank the unguided AbLang2 outputs improved success rate considerably. Using NDS and STS MOO guidance during generation improved results even further, with STS MOO guidance leading to a perfect 100% success rate.

In Fig 2B, we see that methods with higher success rate usually had tighter binding on their successful antibodies. Notable exceptions were SAbDabMLM and pIgGen-dev, perhaps due to their idiosyncratic training data. Both NDS and STS guidance not only improved success rate, but eliminated generation of very weak binders. In Fig 2C, for yield relative to CFPS input volume (a quantitative measure of synthesizability), we do not see any substantial trends among unguided methods. We do see, though, that our guidance methods produced antibodies with less variance in yield, particularly on the higher end of the distribution.

We show additional *in silico* evaluations of the *in vitro* experiment sequences in Fig 3 and Fig 4. In Fig 3A, showing OASis humanness, we once again witness a pattern of superiority for our proposed beam search methods over their Gibbs counterparts. On the other hand, we surprisingly see that AbLang2 tended to produce less human antibodies, despite being trained on human sequences; meanwhile, ESM2 650M had fairly high humanness outputs, despite not even being trained on antibody-specific data. SAbDabMLM sequences have low humanness, unsurprisingly since it was not trained on OAS. We do see that the guidance methods produced less human antibodies, a risk of using a single guidance model which can in principle be ameliorated by guidance with a humanness metric.

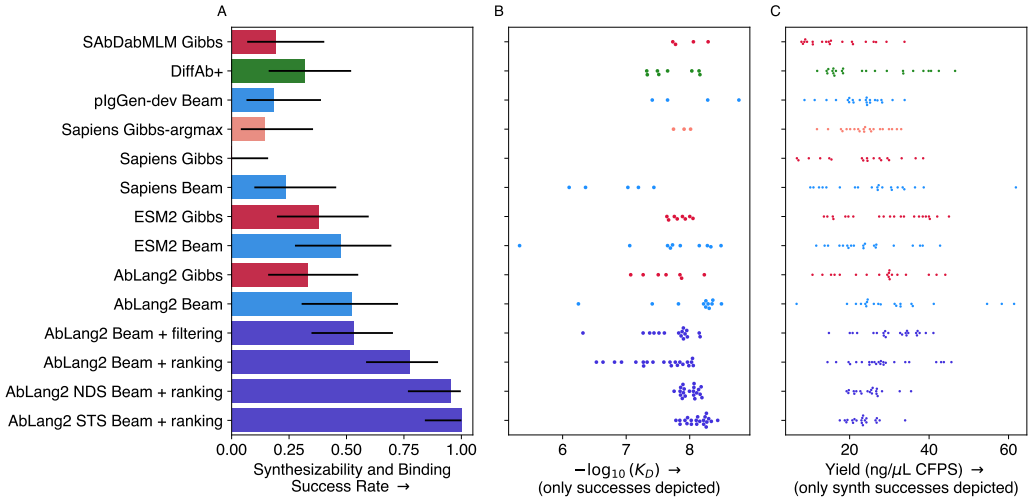

Figure 2: *In vitro* results. For success rate, we show 95% CIs via binomial test inversion.

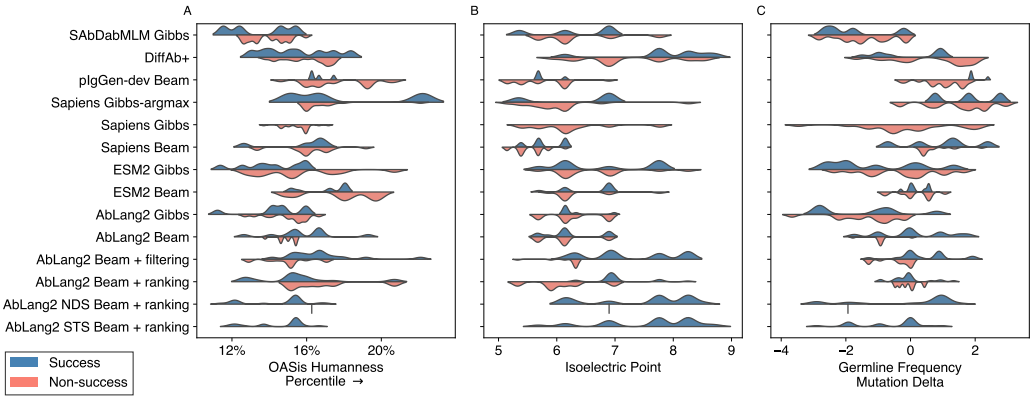

Figure 3: Additional developability-related *in silico* evaluations of the sequences in the *in vitro* experiment. For each method, observations are split by successful synthesizability and binding.

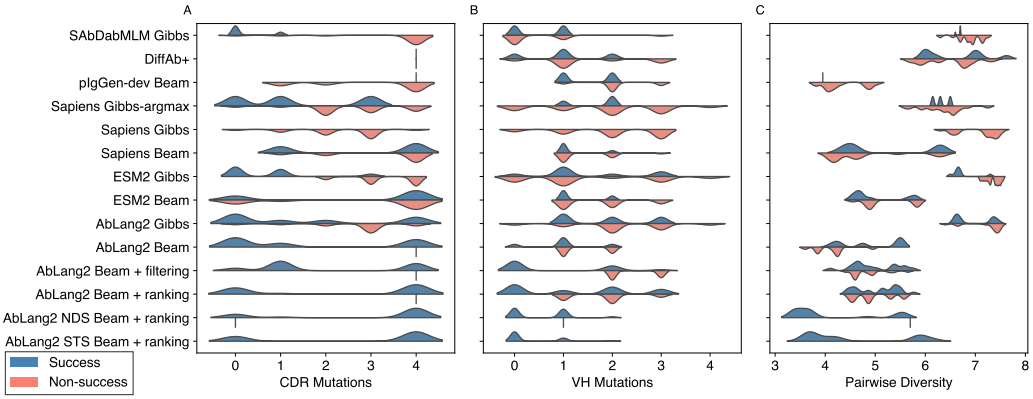

Figure 4: Additional analysis of mutational preferences of different methods for the *in vitro* experiment. or each method, observations are split by successful synthesizability and binding.

In Fig 3C, we show germline bias. Positive values indicate mutations toward more germline-typical residues. Recall that while increased humanness is usually good for antibody therapeutics optimization (i.e. reduced immunogenicity), positive germline bias can be problematic for binding specificity (Olsen et al., 2024). We see that Sapiens Gibbs-argmax and Sapiens Beam, but not Sapiens Gibbs, have high germline bias. This is consistent with the observation in Olsen et al. (2024) that Sapiens has high germline bias, and with our observation that Gibbs sampling often fails to elicit sequences that are consistent with that model's preferences.

In Fig 4A, we reassuringly see that the successful supervised methods still heavily mutated the CDR regions, like the unsupervised methods. On the other hand, in Fig 4B, we see that they were systematically different from unsupervised methods in that they more frequently mutated the VL rather than the VH chain. In Fig 4C, showing pairwise diversity, we see that our beam search method and especially our guidance-based methods had less pairwise diversity. However, it's important to note that this experiment reflects diversity when using a single seed. As shown in our *in silico* evaluation, beam search had less intra-seed diversity, while maintaining inter-seed diversity. Low intra-seed diversity means that children of the same seed tend to share the same mutated positions and especially the same new residues, whereas low inter-seed diversity means that variants from different seeds are converging to the same part of sequence space; the latter is more problematic.

4 Conclusion

We proposed a new MLM sampling method for iterative protein optimization. We benchmarked sampling methods and several language models on real antibody therapeutics optimization campaigns. Our results yield new insights into the effects of models, sampling algorithms, and guidance.

We conclude with several practical recommendations. First, we advocate for the use of supervision, for both ranking and guidance, whenever labeled data may be acquired. Second, we are able to recommend ESM2-650M and especially AbLang2 for antibody engineering. Third, we advise using our proposed stochastic beam search over Gibbs-type sampling. Fourth, we highlight the need for care when using supervised guidance, as it can lead to undesirable side effects in multi-objective optimization settings. Fifth, we suggest examining smooth Tchebycheff scalarization as an alternative to non-dominated Pareto-based sorting, for settings where one wants to simultaneously maximally-satisfy multiple objectives, rather than make progress along each objective.

REFERENCES

David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.

Sarah Alamdari, Nitya Thakkar, Rianne Van Den Berg, Neil Tenenholtz, Robert Strome, Alan M Moses, Alex X Lu, Nicolò Fusi, Ava P Amini, and Kevin K Yang. Protein generation with evolutionary diffusion: sequence is all you need. *BioRxiv*, pp. 2023–09, 2023.

Alan Nawzad Amin, Nate Gruver, Yilun Kuang, Lily Li, Hunter Elliott, Calvin McCarter, Aniruddh Raghu, Peyton Greenside, and Andrew Gordon Wilson. Bayesian optimization of antibodies informed by a generative model of evolving sequences. *arXiv preprint arXiv:2412.07763*, 2024.

Nathaniel R Bennett, Joseph L Watson, Robert J Ragotte, Andrew J Borst, Déjenaé L See, Connor Weidle, Riti Biswas, Ellen L Shrock, Philip JY Leung, Buwei Huang, et al. Atomically accurate de novo design of single-domain antibodies. *bioRxiv: The Preprint Server for Biology*, 2024.

Aadyot Bhatnagar, Sarthak Jain, Joel Beazer, Samuel C Curran, Alexander M Hoffnagle, Kyle S Ching, Michael Martyn, Stephen Nayfach, Jeffrey A Ruffolo, and Ali Madani. Scaling unlocks broader generation and deeper functional understanding of proteins. *bioRxiv*, pp. 2025–04, 2025.

Sarah M Burbach and Bryan Briney. A curriculum learning approach to training antibody language models. *PLOS Computational Biology*, 21(9):e1013473, 2025.

Paul J Carter and Greg A Lazar. Next generation antibody drugs: pursuit of the'high-hanging fruit'. *Nature Reviews Drug Discovery*, 17(3):197–223, 2018.

Tong Chen, Yinuo Zhang, Sophia Tang, and Pranam Chatterjee. Multi-objective-guided discrete flow matching for controllable biological sequence design. *arXiv preprint arXiv:2505.07086*, 2025.

Cyrus Chothia and Arthur M Lesk. Canonical structures for the hypervariable regions of immunoglobulins. *Journal of molecular biology*, 196(4):901–917, 1987.

Peter JA Cock, Tiago Antao, Jeffrey T Chang, Brad A Chapman, Cymon J Cox, Andrew Dalke, Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, et al. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25(11):1422, 2009.

UniProt Consortium. Uniprot: a hub for protein information. *Nucleic acids research*, 43(D1): D204–D212, 2015.

Jeremie Theddy Darmawan, Yarin Gal, and Pascal Notin. Sampling protein language models for functional protein design. In *Learning Meaningful Representations of Life (LMRL) Workshop at ICLR 2025*, 2025.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

James Dunbar and Charlotte M Deane. Anarci: antigen receptor numbering and receptor classification. *Bioinformatics*, 32(2):298–300, 2016.

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.

Quentin Fournier, Robert M Vernon, Almer van der Sloot, Benjamin Schulz, Sarath Chandar, and Christopher James Langmead. Protein language models: Is scaling necessary? *bioRxiv*, pp. 2024–09, 2024.

Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. A systematic exploration of diversity in machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1100–1111, 2013.

Cade Gordon, Aniruddh Raghu, Peyton Greenside, and Hunter Elliott. Generative humanization for therapeutic antibodies. *arXiv preprint arXiv:2412.04737*, 2024.

Emil Julius Gumbel. *Statistical theory of extreme values and some practical applications: a series of lectures*, volume 33. US Government Printing Office, 1954.

Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J. Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q. Tran, Jonathan Deaton, Marius Wiggert, Rohil Badkundri, Irhum Shafkat, Jun Gong, Alexander Derry, Raul S. Molina, Neil Thomas, Yousuf A. Khan, Chetan Mishra, Carolyn Kim, Liam J. Bartie, Matthew Nemeth, Patrick D. Hsu, Tom Sercu, Salvatore Candido, and Alexander Rives. Simulating 500 million years of evolution with a language model. *Science*, 2025. doi: 10.1126/science.ads0018. URL http://dx.doi.org/10.1126/science.ads0018.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.

Emiel Hoogeboom, Alexey A Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and Tim Salimans. Autoregressive diffusion models. *arXiv preprint arXiv:2110.02037*, 2021.

Sean R Johnson, Sarah Monaco, Kenneth Massie, and Zaid Syed. Generating novel protein sequences using gibbs sampling of masked language models. *bioRxiv*, pp. 2021–01, 2021.

Elvin A Kabat and Tai Te Wu. Attempts to locate complementarity-determining residues in the variable positions of light and heavy chains. *Annals of the New York Academy of Sciences*, 190 (1):382–393, 1971.

Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. Neural machine translation in linear time. *arXiv preprint arXiv:1610.10099*, 2016.

Georges Köhler and Cesar Milstein. Continuous cultures of fused cells secreting antibody of predefined specificity. *nature*, 256(5517):495–497, 1975.

Wouter Kool, Herke Van Hoof, and Max Welling. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International conference on machine learning*, pp. 3499–3508. PMLR, 2019.

Aleksandr Kovaltsuk, Jinwoo Leem, Sebastian Kelm, James Snowden, Charlotte M Deane, and Konrad Krawczyk. Observed antibody space: a resource for data mining next-generation sequencing of antibody repertoires. *The Journal of Immunology*, 201(8):2502–2509, 2018.

Marie-Paule Lefranc, Christelle Pommié, Manuel Ruiz, Véronique Giudicelli, Elodie Foulquier, Lisa Truong, Valérie Thouvenin-Contet, and Gérard Lefranc. Imgt unique numbering for immunoglobulin and t cell receptor variable domains and ig superfamily v-like domains. *Developmental & Comparative Immunology*, 27(1):55–77, 2003.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, et al. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*, 2024.

Xi Lin, Xiaoyuan Zhang, Zhiyuan Yang, Fei Liu, Zhenkun Wang, and Qingfu Zhang. Smooth tchebycheff scalarization for multi-objective optimization. *arXiv preprint arXiv:2402.19078*, 2024.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Sal Candido, et al. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022.

Ruei-Min Lu, Yu-Chyi Hwang, I-Ju Liu, Chi-Chiu Lee, Han-Zen Tsai, Hsin-Jung Li, and Han-Chung Wu. Development of therapeutic antibodies for the treatment of diseases. *Journal of biomedical science*, 27(1):1, 2020.

Ali Madani, Bryan McCann, Nikhil Naik, Nitish Shirish Keskar, Namrata Anand, Raphael R Eguchi, Po-Ssu Huang, and Richard Socher. Progen: Language modeling for protein generation. *arXiv preprint arXiv:2004.03497*, 2020.

Chris J Maddison, Daniel Tarlow, and Tom Minka. A* sampling. *Advances in neural information processing systems*, 27, 2014.

John McCafferty, Andrew D Griffiths, Greg Winter, and David J Chiswell. Phage antibodies: filamentous phage displaying antibody variable domains. *nature*, 348(6301):552–554, 1990.

Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in neural information processing systems*, 34:29287–29303, 2021.

Clara Meister, Ryan Cotterell, and Tim Vieira. If beam search is the answer, what was the question? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2173–2185, 2020.

Pascal Notin, Mafalda Dias, Jonathan Frazer, Javier Marchena-Hurtado, Aidan N Gomez, Debora Marks, and Yarin Gal. Tranception: protein fitness prediction with autoregressive transformers and inference-time retrieval. In *International Conference on Machine Learning*, pp. 16990–17017. PMLR, 2022.

Tobias H Olsen, Iain H Moal, and Charlotte M Deane. Addressing the antibody germline bias and its effect on language models for improved antibody design. *Bioinformatics*, 40(11):btae618, 2024.

Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, et al. Boltz-2: Towards accurate and efficient binding affinity prediction. *BioRxiv*, 2025.

David Prihoda, Jad Maamary, Andrew Waight, Veronica Juan, Laurence Fayadat-Dilman, Daniel Svozil, and Danny A Bitton. Biophi: A platform for antibody design, humanization, and humanness evaluation based on natural antibody repertoires and deep learning. In *MAbs*, volume 14, pp. 2020203. Taylor & Francis, 2022.

Aniruddh Raghu, Sebastian Ober, Maxwell Kazman, and Hunter Elliott. Guided sequence-structure generative modeling for iterative antibody optimization. *arXiv preprint arXiv:2509.16357*, 2025.

Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.

Yair Schiff, Subham Sekhar Sahoo, Hao Phung, Guanghan Wang, Sam Boshar, Hugo Dalla-torre, Bernardo P de Almeida, Alexander Rush, Thomas Pierrot, and Volodymyr Kuleshov. Simple guidance mechanisms for discrete diffusion models. *arXiv preprint arXiv:2412.10193*, 2024.

Richard W Shuai, Jeffrey A Ruffolo, and Jeffrey J Gray. Generative language modeling for antibody design. *BioRxiv*, pp. 2021–12, 2021.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *ArXiv*, pp. arXiv–2412, 2025.

Oliver M Turnbull, Dino Oglic, Rebecca Croasdale-Wood, and Charlotte M Deane. p-iggen: a paired antibody generative language model. *Bioinformatics*, 40(11):btae659, 2024.

Kevin K Yang, Nicolo Fusi, and Alex X Lu. Convolutions are competitive with transformers for protein sequence pretraining. *Cell systems*, 15(3):286–294, 2024.

Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.

Ruoxi Zhang and Pranam Chatterjee. Multi-objective nanobody design via masked discrete diffusion with simplex refinement. In *NeurIPS 2025 Workshop on Structured Probabilistic Inference {\&} Generative Modeling*, 2025.

A APPENDIX

A.1 ADDITIONAL DETAILS ON OUR PROPOSED METHOD

A.1.1 APPROXIMATIONS TO THE PSEUDO-LOG-LIKELIHOOD IN BEAM SEARCH

Recall that our beam search scores each candidate sequence $\boldsymbol{x}'$ (which differs from the template $\boldsymbol{x}$ by a substitution at position k) using an approximate PLL (Eq. 3). Here we discuss the space of approximations and their computational tradeoffs.

The exact PLL of $\boldsymbol{x}'$ is

$$PLL(\boldsymbol{x}') = \sum_{i=1}^{L} \log \hat{p}(x'_i \mid \boldsymbol{x}'_{j\neq i}), \tag{4}$$

where each term requires masking position i in the *child* sequence $\boldsymbol{x}'$ and running a forward pass. Since $\boldsymbol{x}'$ differs from $\boldsymbol{x}$ at position k, the conditional distributions $\hat{p}(\cdot \mid \boldsymbol{x}'_{j\neq i})$ differ from those of the template $\hat{p}(\cdot \mid \boldsymbol{x}_{j\neq i})$ at every position $i \neq k$ (because the context includes x'_k instead of x_k). Computing the exact PLL therefore requires L forward passes *per child*. With $O(L)$ children per beam search step (one per position per amino acid), the total cost is $O(L^4)$.

The key insight behind our beam search is that having computed the template's PLL – which requires L forward passes yielding the matrix $\boldsymbol{A}$ where $\boldsymbol{A}_{i,:} = \texttt{softmax}_\tau(\tilde{\boldsymbol{Y}}^{(i)}_{i,:})$ – one can cheaply approximate the PLL of all single-substitution neighbors. The question is what approximation to use at positions $i \neq k$.

Table 1 summarizes the five approaches. They are distinguished by what the model observes at positions i and k when computing the conditional probability at position $i \neq k$, as well as their computational cost. We note that position k always uses the exact single-mask conditional $\boldsymbol{A}_{k,\texttt{index}(x'_k)}$ in all approaches except the no-masking variants.

Table 1: Comparison of PLL approximations for scoring all single-substitution neighbors during one step of beam search. Position k is the substitution site; position $i \neq k$ is any other position. Each forward pass has cost $O(L^2)$.

Approximation	**Pos. i sees at i**	**Pos. i sees at k**	**Total cost**
Exact PLL	MASK	x'_k (child)	$O(L^4)$
Double-mask	MASK	MASK	$O(L^4)$
Wild-type marginal	MASK	x_k (template)	$O(L^3)$
No-masking (child)	x'_i (visible)	x'_k (child)	$O(L^3)$
No-masking (template)	x_i (visible)	x_k (template)	$O(L^2)$

Wild-type marginal (our approach). At position $i \neq k$, we use the template's own conditional distribution, i.e., the distribution obtained by masking position i in $\boldsymbol{x}$:

$$PLL_{\text{wt}}(\boldsymbol{x}'; \tau) = \sum_{i=1}^{L} \log \boldsymbol{A}_{i,\texttt{index}(x'_i)}. \tag{5}$$

Since $x'_i = x_i$ for $i \neq k$, only the term at $i = k$ differs from the template's PLL, so all neighbor PLLs can be read off from $\boldsymbol{A}$ with no additional forward passes. The approximation error is that position i is scored conditioned on seeing the *template* token x_k at position k, rather than the child's substituted token x'_k. The total cost is $O(L^3)$ (for the L stepwise-masked forward passes to compute $\boldsymbol{A}$), amortized over all $O(L)$ neighbors.

Double-mask. To avoid conditioning on the wrong token at position k, one can instead mask *both* positions i and k simultaneously. Let $\tilde{\boldsymbol{Y}}^{(i,k)}$ denote the MLM output when both positions i and k are masked. The approximate PLL becomes

$$PLL_{\text{dm}}(\boldsymbol{x}'; \tau) = \log \boldsymbol{A}_{k,\texttt{index}(x'k)} + \sum_{i\neq k} \log \texttt{softmax}_\tau(\tilde{\boldsymbol{Y}}^{(i,k)}_{i,\texttt{index}(x_i)}). \tag{6}$$

At position k we use the same single-mask conditional as the wild-type marginal; at positions $i \neq k$, we are agnostic about the substitution site. Computing all pairwise double-mask outputs requires $\binom{L}{2}$ forward passes. Together with L single-mask forward passes, the total cost is $O(L^4)$. This is as expensive as the exact PLL, while introducing approximation error.

No-masking. At the other extreme, one can skip masking at position i entirely and use $\hat{p}(x_i \mid \boldsymbol{x})$ as an approximation to $\hat{p}(x_i \mid \boldsymbol{x}_{j \neq i})$ (Prihoda et al., 2022). This can be applied to either the template sequence (one forward pass shared by all neighbors, $O(L^2)$ total) or each child sequence individually ($O(L^3)$ total, since each child requires its own forward pass). In either case, the model observes x_i itself when predicting position i, yielding probabilities near 1 that contribute little discriminative signal. At position k, the child no-masking variant at least conditions on the correct token x'_k (just without masking it), while the template no-masking variant conditions on the wrong token x_k without masking.

Discussion. The approaches are ordered in Table 1 from most to least principled. The exact PLL and the double-mask approach have the same asymptotic cost but differ in that the exact PLL conditions each position on the child's actual token at k, while the double-mask marginalizes over k via the mask token. In practice, the wild-type marginal offers an appealing balance: it is the cheapest approach that still masks position i (yielding meaningful conditional probabilities), while the error from conditioning on the template token at k is empirically small in practice (Meier et al., 2021).

A.1.2 Multi-objective optimization guidance

For multi-objective optimization, we begin by converting the pseudo-log-likelihood scores to pseudo-perplexity scores. (Each sequence's pseudo-log-likelihood was already perturbed with independent Gumbel noise $g \sim \text{Gumbel}(0, 1)$ prior to computing the pseudo-perplexity.) We multiply other objectives by -1 as needed to make the criteria also minimization. Then, for each objective, we standardize scores to z-scores using the mean and standard deviation computed across all candidate neighbors of the current template sequence. To produce a single ranking, we aggregate the z-scored objectives using a smooth Tchebycheff scalarization (Lin et al., 2024), computing $\log \sum_i c_i \exp(\tilde{f}_i)$ for each sequence, where $\tilde{f}_i$ are the z-scored objectives and $c_i > 0$ are per-objective weights, with larger c_i placing greater emphasis on objective i. Sequences are then ranked by this score, with lower values preferred.

A.2 Full details for in silico experiments

Experimental setup We obtained predicted probabilities of synthesizability (i.e. protein quantification, PQ) from an oracle trained on 6038 samples, split 60%-20%-20% into train-valid-test sets. We trained a 5-model ensemble of 1d dilated convolutional neural networks (Yu & Koltun, 2015) with mean test AUROC 0.9057. We obtained melting temperature (T_m) predictions by finetuning ESM2-150M (Lin et al., 2022) on internal phage display and CFPS data. The model was trained to predict first melting temperature (in Celsius) measured by the Uncle (differential scanning fluorimetry) instrument's peak-calling algorithm on the Barycentric Mean (BCM) of intrinsic fluorescence. For samples from the lineage from which our seeds were drawn, the train-valid-test split was 384-188-219 samples, with test Spearman correlation of 0.7693. On a larger test-set of 1111 samples from different lineages for the same therapeutics program, the test Spearman was 0.8486.

To compute intra-seed (inter-seed) pairwise diversity, for each child sequence, we align it to all other child sequences of the same (different) seed, and compute the number of mutations. For each child antibody, we compute the mean number of mutations over all its siblings; we do not aggregate over all children. Thus, for each method, pairwise diversity is represented by a distribution (each child generated by the method being a sample) rather than a single number.

For computing OASis humanness percentile scores, (Prihoda et al., 2022) we use the medium ($\geq$ 50% subjects) threshold. For each child antibody, we computed the germline frequency mutation delta as the sum over all mutated positions (relative to the seed) of the difference between the child residue's germline family frequency and the seed residue's germline family frequency at that position, using IMGT numbering-aware germline gene family frequency tables. For computing isoelectric point, we used Biopython v1.85 (Cock et al., 2009) with the concatenated VH and VL

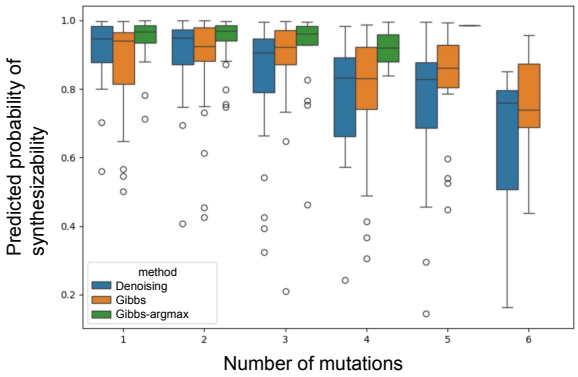

Figure 5: *In silico* predicted probability of synthesizability for AbLang2 generated sequences, as a function of number of edits.

sequences. For the number of CDR mutations metric, for each child sequence we align (Dunbar & Deane, 2016) to the seed, and classify mutations using the Chothia (Chothia & Lesk, 1987) scheme.

Experimental results We first show results for predicted synthesizability (Fig. 6A) and predicted thermostability (Fig 6B). Methods using our proposed beam search method consistently performed better on predicted synthesizability, compared to other sampling methods; in contrast, there was no discernable improvement for thermostability. Pairwise diversity results are shown in Fig. 6C for inter-seed diversity and Fig 6D for intra-seed diversity. We see that beam search has much less intra-seed diversity than Gibbs, but this pattern is much less strong for inter-seed diversity. More generally, we see that intra-seed and inter-seed diversity are not necessarily correlated; CloneLM Beam and Amplify Beam have decent intra-seed diversity but low inter-seed diversity.

Results for humanness and germline bias are shown in Fig. 7A and Fig. 7B, respectively. The most noteworthy observations are that SAbDabMLM Gibbs produces sequences with low humanness. Meanwhile, CloneLM Beam, Sapiens Gibbs-argmax, and Amplify Beam have high germline bias. For isoelectric point in Fig. 7C, we observe that ESM2-650M Beam is the only method that simultaneously performed well in general, while also generating variants with low pI; thus it may be useful in scenarios where pI reduction is needed. For CDR mutation count in Fig. 7D, we see that SAbDabMLM Gibbs is similar to DiffAb+ in being biased to CDR-only mutations, despite not being intentionally designed to do so. This suggests that this behavior is due to biases in their shared training data. Conversely, we notice that pIgGen tends to produce mutations in framework regions.

A.3 Additional details on the in vitro experiment

A.3.1 Experimental setup

When generating variants with CDR-only mutations, we defined CDRs by the union of masks based on the Kabat (Kabat & Wu, 1971), Chothia (Chothia & Lesk, 1987), and IMGT (Lefranc et al., 2003) definitions.

Beam search was run with a beam size of 5, with the exception of the guidance methods which used a beam size of 20. All methods used a stochastic beam search temperature of 1.5. For unsupervised methods, the Gumbel-perturbed pseudo-log-likelihoods were used to rank sequences.

For supervised filtering, guidance, and ranking, we split the labeled 729 samples 60%-20%-20% into train-valid-test sets, and trained an ensemble of 5 models. We aligned sequences to a reference antibody sequence, and provided one-hot encodings of the aligned sequences as input features to the models. We trained CARP/Bytenet (Kalchbrenner et al., 2016; Yang et al., 2024) classification models to predict simultaneous synthesizability and binding success. The ensemble's models had test mean AUROC 0.9342 and standard deviation 0.0101. For filtering, we filtered generated se-

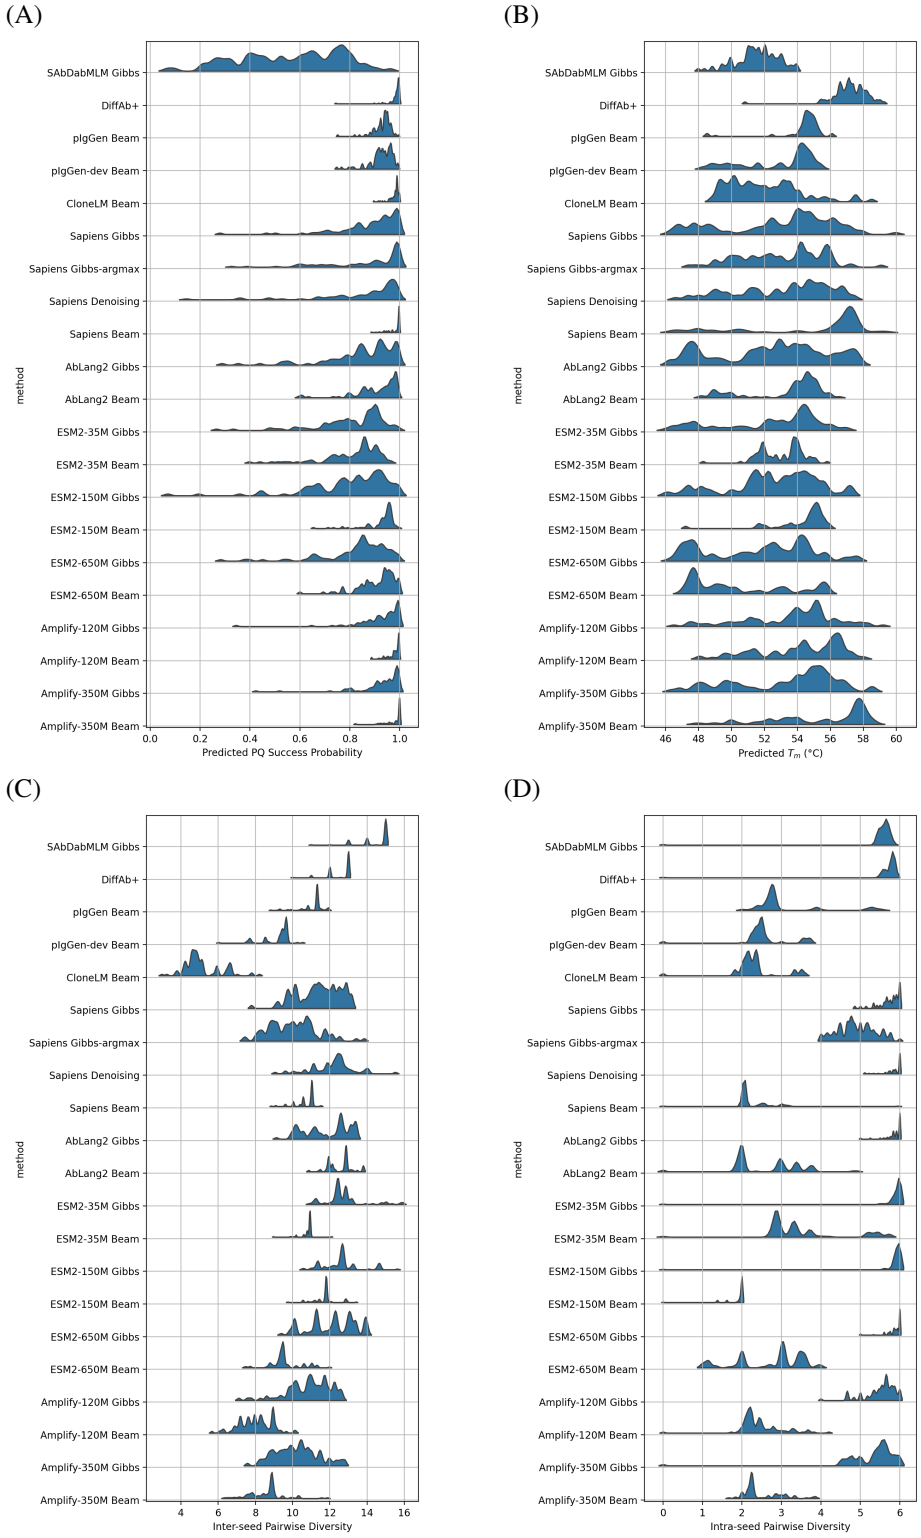

Figure 6: *In silico* metrics of quality and diversity for generated sequences: (A) predicted probability of synthesizability, (B) predicted melting temperature, (C) inter-seed pairwise diversity, and (D) intra-seed pairwise diversity.

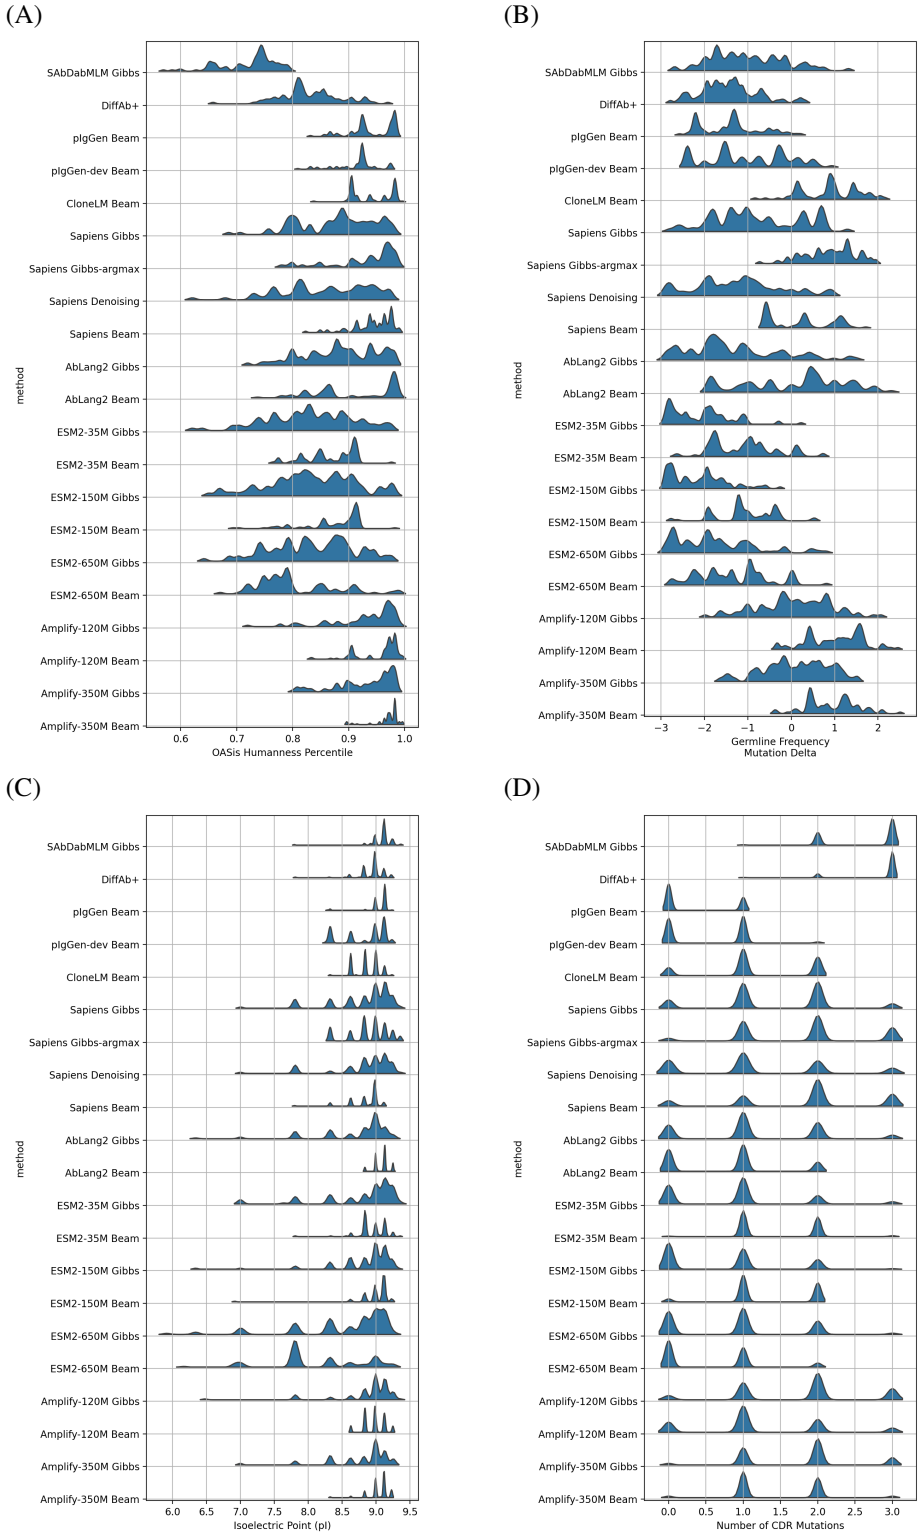

Figure 7: Additional *in silico* evaluation results for generated sequences: (A) OASis humanness percentile, (B) germline mutation bias, (C) isoelectric point, and (D) number of CDR mutations.

quences by ($p_{success} > 0.6$); for ranking, we sorted generated sequences by $p_{success}$. For NDS guidance and ranking, during beam search we combined AbLang2 pseudo-perplexity and supervised $p_{success}$, breaking NDS ties with the pseudo-perplexity; at ranking time, sequences were sorted by $p_{success}$. For STS guidance and ranking, during beam search we combined AbLang2 pseudo-perplexity (with $c_{\text{AbLang2}} = 1$) and supervised $p_{success}$ (with $c_{\text{supervised}} = 2$); at ranking time, sequences were sorted by $p_{success}$. Note that at ranking time, the ranker only had access to the top-1000 generated sequences, as sorted by the beam search criteria: for the unguided methods, these are the Gumbel-perturbed pseudo-log-likelihoods; for the guided methods, these are the aforementioned multiobjective scalarization criteria.

Before sending sequences to the wet-lab, designed sequences were filtered to have isoelectric point ≤ 9. Sequences were also filtered to exclude those introducing any of the following standard sequence liabilities relative to the seed: Asp-Pro cleavage sites, asparagine deamidation motifs (NG, NH, NN, NS, NT), aspartate isomerization motifs (DG, DD, DH, DS, DT), N-linked glycosylation sequons (N-X-S/T, where X $\neq$ P), high-risk oxidation sites (exposed Met or Trp), and unpaired cysteine residues.

A.3.2 In vitro binder classification

Binding kinetics were measured by bio-layer interferometry (BLI) and fit to a 1:1 binding model. Each sample was subjected to automated quality control to determine whether it produced a reliable binding measurement. Before kinetics-specific QC was applied, samples were checked for upstream protein production failures based on protein quantification metrics. A sample was failed if its protein concentration was below 50 ng/µL, or if its A260/A280 purity ratio fell outside the range of 0.5 to 0.85. Samples were classified as non-binders if the maximum binding response signal at the highest analyte concentration was less than 0.05 nm. We manually reviewed select BLI curves, verifying that these automated checks were indeed failing apparent non-binders, while passing apparent binders.