ANTIBODY DESIGN WITH CONSTRAINED BAYESIAN OPTIMIZATION

Yimeng Zeng[1]*, **Hunter Elliott**[2], **Phillip Maffettone**[2],
Peyton Greenside[2], **Osbert Bastani**[1], **Jacob R. Gardner**[1]

[1] University of Pennsylvania, [2] BigHat Biosciences,

ABSTRACT

In therapeutic antibody design, achieving a balance between optimizing binding affinity subject to multiple constraints, and sequence diversity within a batch for experimental validation presents an important challenge. Contemporary methods often fall short in simultaneously optimizing these attributes, leading to inefficiencies in experimental exploration and validation. In this work, we tackle this problem using the latest developments in constrained latent space Bayesian optimization. Our methodology leverages a deep generative model to navigate the discrete space of potential antibody sequences, facilitating the selection of diverse, high-potential candidates for synthesis. We also propose a novel way of training VAEs that leads to a lower dimensional latent space and achieves excellent performance under the data-constrained setting. We validate our approach ***in vitro*** by synthesizing optimized antibodies, demonstrating consistently high binding affinities and preserved thermal stability.

1 INTRODUCTION

Generative modeling has emerged as a transformative approach in computational biology, particularly in the discovery and optimization of biological sequences such as molecules and protein sequences. With recent developments in latent space Bayesian optimization (BO) (Jin et al., 2018; Maus et al., 2022), it is now feasible to search for sequences with desired properties from this exponentially large combinatorial space.

Recent progress in this topic has focused on both improving the deep generative model that is used for optimization, and the optimization algorithm itself. Earlier works of Jin et al. (2018) and Maus et al. (2022) investigated specific VAE architectures and string representations that reliably produced valid molecules. The work of Stanton et al. (2022) and Maus et al. (2023b) used masked language models and VAEs in the domain of optimizing protein sequences. Notably, Gruver et al. (2023) proposed `LaMBO-2`, a method that uses saliency maps to guide discrete diffusion models to generate edits on antibody sequences. While it showed promising results for generating an enriched library without excessive *in vitro* screening, it faces similar problems as genetic methods (Ren et al., 2022) in incorporating the kind of constraints often required for real-world therapeutic antibody optimization, with an increasing number of sampling steps leading to violation of the constraints.

Addressing these challenges, our work proposes a comprehensive framework for latent space BO within the context of antibody sequence optimization. We make the following contributions:

- We introduce a novel method for the antibody optimization setting. Our method consistently produces a batch of sequences that are diverse from each other and always satisfies given constraints.
- We propose a novel VAE model training approach that efficiently generates a mutation-based dataset from a seed sequence, minimizing training data requirements and latent space dimensions while preserving optimization efficacy.

*yimengz@seas.upenn.edu

- **We validate this approach *in vitro*, producing antibodies with improved binding affinities and preserved thermostability.**

2 Background and Related Work

Black-box and Bayesian optimization. In black-box optimization, we aim to optimize an *oracle* objective function $f(\mathbf{x})$ over a space of candidates $\mathbf{x}^* = \text{argmax}_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x})$. Examples of such problems include molecule activity maximization for drug discovery (Trabucco et al., 2022; Maus et al., 2022), and binding affinity of DNA sequences or proteins (Barrera et al., 2016; Gruver et al., 2023). Commonly, $f(x)$ is assumed to be expensive to evaluate or even completely unknown.

Bayesian optimization is a sample-efficient framework to solve these costly to evaluate model-based optimization problems (Osborne et al., 2009; Mockus, 1982; Snoek et al., 2012). At iteration t of BO, one has access to observations $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^t$, where y_i denotes the objective value of the input x_i. Typically, a Gaussian process (Rasmussen, 2003) is employed as the surrogate model to approximate the objective function using these inputs and values. This surrogate model aids the optimization by employing an acquisition function, which strategically proposes the next candidates for evaluation. After querying these candidates through the true oracle, the surrogate model is updated with the new observations. This process gradually builds a more comprehensive dataset and refines the surrogate model, thereby improving the quality of the proposed samples in future iterations.

Bayesian optimization over latent spaces. Due to the discrete and structural nature of biological sequences, we utilize recent developments in latent space BO that adapt BO from continuous black-box optimization problems to the discrete domain (Tripp et al., 2020; Gómez-Bombarelli et al., 2018). Latent space BO leverages the capabilities of deep generative models, most commonly variational autoencoders (VAEs) (Kingma & Welling, 2013) to aid optimization. Concretely, a VAE is composed of two networks: an encoder $\mathcal{E}(\mathbf{z} \mid \mathbf{x}) : \mathcal{X} \to \mathcal{P}(\mathcal{Z})$ mapping from amino acid sequences to latent space $\mathcal{Z}$, and a decoder $\mathcal{D}(\mathbf{x} \mid \mathbf{z}) : \mathcal{Z} \to \mathcal{P}(\mathcal{X})$ that probabilistically decodes latent space vectors back into amino acid sequences. The search space is now over the continuous latent space $\mathcal{Z}$ of the VAE instead of the discrete space of amino acid sequences $\mathcal{X}$, we can now formulate our optimization problem as:

$$\mathbf{x}^\star \approx \mathcal{D}(\mathbf{z}^\star) \quad \text{where} \quad \mathbf{z}^\star = \arg\max_{\mathbf{z}\in\mathcal{Z}} f'\left(\mathcal{D}(\mathbf{z})\right) \tag{1}$$

The objective function we are optimizing now takes in a latent vector $\mathbf{z}$ and decodes it into an amino acid sequence, which is then evaluated with the provided surrogate model f'. We then select a batch of sequences that are then sent to the lab for synthesis and experimental validation. Our optimization algorithm is based off `LOL-BO` (Maus et al., 2022), with adaptations for the constrained setting.

Bayesian optimization for biological discovery. Due to the combinatorial space of possible biological sequences and structures, Bayesian optimization (BO) is a powerful tool for biological discovery. Earlier works utilized BO without deep generative models over a fixed list of molecules for accelerating drug screening (Graff et al., 2021; Hernández-Lobato et al., 2017). However, the space of all possible molecules is so large that any pre-defined list of molecules is negligible (Kirkpatrick & Ellis, 2004). Latent space BO methods resolve this by utilizing the capabilities of deep generative models such as VAEs and Diffusion models so that it is possible to generate any possible molecules during optimization (Maus et al., 2022; Gruver et al., 2023). These methods employ either straightforward string representations (Stanton et al., 2022) or more complex graphical or grammatical structures (Kusner et al., 2017; Jin et al., 2018). The works of Khan et al. (2023) and Romero et al. (2012) explored designing novel antibodies by exploration of the protein fitness landscape, while Maus et al. (2023b) optimized solutions to the inverse protein folding problem with BO. These developments highlight the capability and success of BO in advancing biological discovery.

Antibody design and engineering. Antibodies are the fastest growing class of therapeutics (Carter & Lazar, 2018), and thus represent an important application domain for sequence optimization. To be functional, antibodies must bind their target with strong affinity, but to be viable therapeutics they must also meet a range of criteria collectively termed "developability" (Jarasch et al., 2015). Because many of these developability properties, such as thermostability, need not be maximized but only

exceed a threshold, this lends to framing antibody sequence design as a constrained optimization problem *e.g.* maximizing binding affinity subject to the constraint of acceptable thermostability.

3 METHODS

Our antibody design task seeks to generate sequences that exhibit optimal binding affinity towards a target, while maintaining a lower bound on thermostability based on program requirements. We start with a batch of seed amino acid sequences and their binding affinities to the target as initial observations. We are also given surrogate models that predict the binding affinity and melting temperature of the antibody sequences. The objective is to generate a collection of antibody sequences characterized by elevated affinity and preserved thermostability, which are subsequently synthesized and subjected to empirical validation in a laboratory setting.

Optimizing with constraints. In the domain of antibody engineering, it is standard practice to impose constraints on the generated sequences to ensure their viability. For example, these constraints can be used to preserve developability by requiring thermostability exceed a threshold value, or to restrict the introduced mutations to specific regions such as the CDRs. Additionally, there is a preference for sequences that do not deviate excessively from given seed sequences (Fowler et al., 2014; Storici & Resnick, 2006), as experience has shown that sequences too divergent from known seeds make synthesis and validation significantly harder. The extension of Equation 1 to include these constraints can be represented as follows:

$$\mathbf{x}^\star \approx \mathcal{D}(\mathbf{z}^\star) \quad \text{where} \quad \mathbf{z}^\star = \arg\max_{\mathbf{z}\in\mathcal{Z}} f'\left(\mathcal{D}(\mathbf{z})\right) \text{ s.t. } \forall i\colon c_i\left(\mathcal{D}(\mathbf{z})\right) \leq 0,$$

In this scenario, $c_i(\cdot)$ represents black-box constraints that are applied to the decoded antibody sequences $\mathcal{D}(\mathbf{z})$. We integrated a constrained BO algorithm `SCBO` (Eriksson & Poloczek, 2021) with `LOL-BO` in a straightforward manner, as they are all based on the `TurBO` algorithm (Eriksson et al., 2019). It is known that training the VAE and GP surrogate model end-to-end substantially improves optimization performance (Maus et al., 2022). Following a similar approach, we regularize the latent space of the VAE by training the constraint surrogate models jointly with both the VAE and the objective surrogate models. With a number of m constraints, we now have the modified VAE ELBO with $m+1$ GPs written as:

$$\mathcal{L}_{\text{joint}}(\theta_{\mathcal{E}}, \theta_{\mathcal{D}}, \theta_{\text{GP}_{0:m}}) = \mathbb{E}_{\mathcal{E}(\mathbf{z}|\mathbf{x})}\left[\sum_{i=0}^{m} \mathcal{L}_{\text{GP}_i}\left(\theta_{\text{GP}_i}, \theta_{\mathcal{E}}; \mathbf{z}, \mathbf{y}\right)\right] + \mathcal{L}_{\text{VAE}}(\theta_{\mathcal{E}}, \theta_{\mathcal{D}}; \mathbf{x}),$$

Where $\mathbf{y}$ denotes the objective values obtained so far during optimization. Because the encoder parameters $\theta_{\mathcal{E}}$ are updated jointly with the GP parameters θ_{GP_i}, the encoder $\mathcal{E}$ not only acts as an encoder for the VAE but is also acting as a deep kernel for the surrogate and constraint GP models (Wilson et al., 2016).

Producing diverse solutions. Given the constraint of approximately 70 sequences for laboratory validation, it is desirable to identify a set of solutions that not only exhibit high binding affinity but also encapsulate sufficient diversity to fully explore the antibody design space. To this end, the optimization algorithm has been augmented with `ROBOT` (Maus et al., 2023a), which focuses on discovering a batch of diverse solutions subject to a specific diversity criterion. We use the Levenshtein edit distance $\delta(x, x')$ as the diversity constraint in our optimization process, and solve the following optimization problems:

$$\mathbf{z}_1^\star = \arg\max_{\mathbf{z}\in\mathcal{Z}} f'\left(\mathcal{D}(\mathbf{z})\right) \text{ s.t. } \forall i\colon c_i\left(\mathcal{D}(\mathbf{z})\right) \leq 0,$$

$$\mathbf{z}_n^\star = \arg\max_{\mathbf{z}\in\mathcal{Z}} f'\left(\mathcal{D}(\mathbf{z})\right) \text{ s.t. } \forall i\colon c_i\left(\mathcal{D}(\mathbf{z})\right) \leq 0, \;\&\; \delta\left(\mathcal{D}(z), z_j^\star\right) \geq \tau \text{ for } j = 1, ..., n-1$$

Solving the above problems produces a set of n high-affinity antibody sequences $\{x_i^\star\}_{i=1}^n = \{\mathcal{D}(z_i)^\star\}_{i=1}^n$ so that each antibody is at least τ edits away from any other antibody.

Single Sequence VAE. To address the challenges of latent space BO for antibody optimization, there are two desirable properties when designing a VAE:

- **String reconstruction accuracy**: A high string reconstruction accuracy can ensure that the latent vectors identified by the BO algorithm can be accurately decoded into viable antibody sequences.

- **Latent space dimensionality**: It is known that the performance of BO deteriorates in higher dimensional spaces (Shahriari et al., 2016), a VAE with a low latent space dimension is more desirable for optimization.

In this work, we introduce a novel approach to VAE training, specifically targeted at the antibody optimization problem. Our strategy is grounded in the fact that we aim to produce antibodies with maximal binding affinity, preserved thermostability, and which are only a few edits away from a given "seed" sequence, with mutations predominantly within the CDRs. Following a similar approach to Nikankin et al. (2023), we trained the VAE on a dataset comprising 1.8 million variants of a seed sequence. Each variant is generated through 1-10 random mutations strictly within the allowed editable regions, and then passed through the oracles to obtain a label during optimization. We also removed the non-editable regions of the antibody during training. Therefore, the VAE is only responsible for modeling the mutations within the editable regions, relieving the burden of reconstructing the entire sequence.

4 EXPERIMENTS

4.1 CONSTRAINED LATENT SPACE OPTIMIZATION SETUP

Model details. In this work, we pretrain an autoregressive VAE using the Transformer architecture (Vaswani et al., 2017) with 5 encoder and 5 decoder layers. The model was trained using the VAE ELBO with the KL divergence term multiplied by a constant factor of 0.001 to emphasize reconstruction accuracy (Higgins et al., 2017), achieving a reconstruction accuracy of 99.8% for the targeted antibody sequence, with a 48-dimensional latent space.

Constraints. In all of our experiments, we enforce the following constraints:

- The predicted melting temperature of the antibodies should be above T°C for $T \in \{60, 65\}$ for improved thermal stability
- All of the edits to the seed sequences should be within the complementarity-determining regions (CDRs)
- For any sequence x and batch of n seed sequences $\{x_i\}_{i=1}^n$, we have $\min_i (\delta(x, x_i)) \leq k$ for $k \in \{1, 2, 3, 4, 5\}$

For the diversity criteria, we set $\tau = \lfloor \frac{k}{2} \rfloor$ for each edit distance constraint k to ensure maximum diversity while not degrading optimization performance. For the experiments in this work, we did not set an oracle call budget, instead we ran the optimization algorithm until convergence for each optimization run, and aggregated the best sequences from each temperature and edit distance threshold for experimental validation.

4.2 AFFINITY AND THERMOSTABILITY ORACLE MODELS

The black box oracle models we optimized against consisted of an ensemble of 10 regressors for each task (K_D and T_m). Each ensemble was comprised of CNNs trained on distinct 0.8/0.1/0.1 fractional train/validation/test splits for cross validation. The K_d models were pre-trained on ∼650,000 data points derived from phage display affinity maturation, then fine tuned on 1425 affinity measurements derived from Bio-Layer Interferometry (BLI) measurement of antibodies synthesized via Cell Free Protein Synthesis (CFPS). These same CFPS produced antibodies were characterized for thermostability via NanoDSF, and this data was used to train an ensemble of dilated CNNs for T_m modeling. The K_d model ensemble had a measured vs. predicted spearman correlation in crossvalidation 0.73, while the the T_m model ensemble had an spearman correlation of 0.62. Both models used dilated convolutional architectures, with the affinity model using a CARP/ByteNet architecture (Yang et al., 2022) with 32 residual ByteNet blocks and a model dimension of 64. The ligher-weight thermostability model has 8 dilated convolutional layers with stnadard 2^l progression in dilation rate, two stride-2 convolutional layers and a fully connected layer. All models were trained on sequences aligned to a reference sequence.

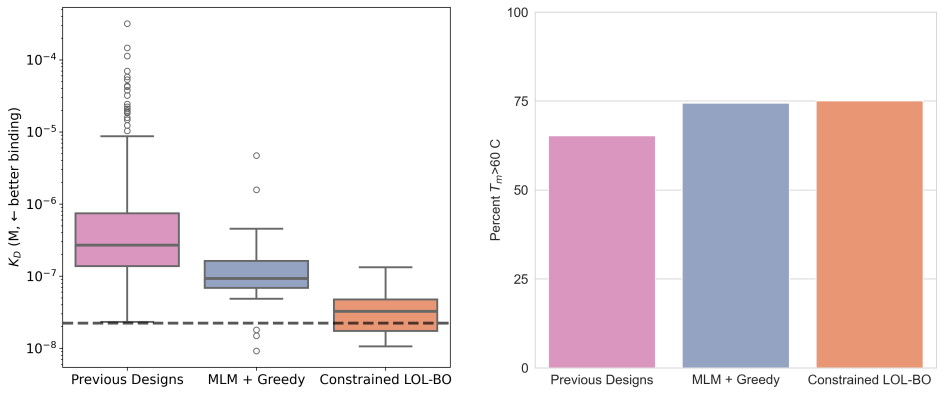

(a) Optimized Target Binding Affinity

(b) Thermostability Constraint Satisfaction

Figure 1: **Lab validation confirms optimized antibodies have strong target binding and largely satisfy thermostability constraints.** Resulting variants have the best median affinities (a) and represent 7 of the top 10 designs overall, many exceeding the affinity of the best optimization seed (a, dashed line). Measured T_{m}s (b) are preserved at or above the 60 C program requirement at a rate equal to or better than prior methods, including greedy optimization with the same oracles.

4.3 Lab validated results

As part of an ongoing therapeutic optimization campaign a range of sequence optimization methods were deployed, ranging from simple greedy optimization to several single- and multi-objective BayesOpt methods (Stanton et al., 2022; Maus et al., 2022; Gruver et al., 2021). However due to the iterative nature of this optimization campaign, these earlier methods were deployed with smaller, differently distributed training data and so to avoid incongruous comparison we group these together as "Previous Designs". For the most recent iteration, we synthesized 70 antibodies from constrained `LOL-BO` + single sequence VAE optimized sequences. We also synthesized antibodies from a baseline greedy optimization method, where CDR variants were proposed using a masked language model (Prihoda et al., 2022), and then ranked via predicted affinities and thermostabilities using the same oracles and a non-dominated sort, with the top 128 selected for synthesis. Variants which were synthesizable in sufficient quantities were then characterized for target binding affinity K_D and thermostability T_m. The constrained `LOL-BO` variants had the best median affinity of any method deployed to date on this program, also exceeding that of the baseline method 1a. The thermostability of the designed variants was preserved at or above the 60 °C T_m program requirement (similar to the stability of the optimization seed) at a rate as good or better than other designs.

5 Discussion

Generative modeling is playing an increasingly pivotal role in therapeutic discovery, with Bayesian optimization emerging as a promising approach in this domain. We have shown how a specifically designed deep generative model can be integrated with latent space BO algorithms to obtain better optimization results, achieving good sequence diversity while supporting an arbitrary number of constraints relevant to real-world therapeutics development. Moreover, we successfully optimized antibody binding while constraining thermal stability to be maintained, demonstrating superiority ***in vitro*** to previous design rounds consisting of various optimization methods.

One limitation of our approach is that our method of reducing latent space dimensions relies on knowing which parts of the biological sequence we want to optimize over, for example, CDRs of a given antibody. While we have focused our experimental evaluation on antibody optimization tasks alone, further work is needed to adapt our method to any arbitrary biological sequence optimization problem. Future work may also examine how non-autoregressive models, such as diffusion models and GANs, can be integrated with our algorithm.

6 ACKNOWLEDGMENTS

JG and YZ are supported by NSF award IIS-2145644. OB is supported by NSF Award CCF-1917852. We would like to thank the BigHat Production team, especially Elisabeth Visser, Caleb Ling, Leon Giang, and Zach Ormsby and the Data Science team, especially Lindsey Osimiri. We'd also like to thank program leads Jamie Haddon and Emily Delaney as well as Aaron Solomon and the rest of the BigHat ML team for productive discussions and insightful suggestions.

REFERENCES

Luis A. Barrera, Anastasia Vedenko, Jesse V. Kurland, Julia M. Rogers, Stephen S. Gisselbrecht, Elizabeth J. Rossin, Jaie Woodard, Luca Mariani, Kian Hong Kock, Sachi Inukai, Trevor Siggers, Leila Shokri, Raluca Gordân, Nidhi Sahni, Chris Cotsapas, Tong Hao, Song Yi, Manolis Kellis, Mark J. Daly, Marc Vidal, David E. Hill, and Martha L. Bulyk. Survey of variation in human transcription factors reveals prevalent dna binding changes. *Science*, 351(6280):1450–1454, 2016. doi: 10.1126/science.aad2257.

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.

David Eriksson and Matthias Poloczek. Scalable constrained bayesian optimization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *PMLR*, pp. 730–738. PMLR, 13–15 Apr 2021. doi: 10.48550/arXiv.2002.08526.

David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. doi: 10.48550/arXiv.1910.01739.

Douglas M Fowler, Jason J Stephany, and Stanley Fields. Measuring the activity of protein variants on a large scale using deep mutational scanning. *Nat. Protoc.*, 9(9):2267–2284, September 2014. doi: 10.1038/nprot.2014.153.

Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, January 2018. doi: 10.1021/acscentsci.7b00572.

David E. Graff, Eugene I. Shakhnovich, and Connor W. Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.*, 12:7866–7881, 2021. doi: 10.1039/D0SC06805E.

Nate Gruver, Samuel Stanton, Polina Kirichenko, Marc Finzi, Phillip Maffettone, Vivek Myers, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Effective surrogate models for protein design with bayesian optimization. In *2021 ICML Workshop on Computational Biology*, 2021.

Nate Gruver, Samuel Stanton, Nathan C. Frey, Tim G. J. Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew Gordon Wilson. Protein design with guided discrete diffusion. *arXiv Preprint*, 2023. doi: 10.48550/arXiv.2305.20009.

José Miguel Hernández-Lobato, James Requeima, Edward O. Pyzer-Knapp, and Alán Aspuru-Guzik. Parallel and distributed thompson sampling for large-scale accelerated exploration of chemical space. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *PMLR*, pp. 1470–1479. PMLR, 06–11 Aug 2017.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.

Alexander Jarasch, Hans Koll, Joerg T Regula, Martin Bader, Apollon Papadimitriou, and Hubert Kettenberger. Developability assessment during the selection of novel therapeutic antibodies. *Journal of pharmaceutical sciences*, 104(6):1885–1898, 2015.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *PMLR*, pp. 2323–2332. PMLR, 10–15 Jul 2018.

Asif Khan, Alexander I. Cowen-Rivers, Antoine Grosnit, Derrick-Goh-Xin Deik, Philippe A. Robert, Victor Greiff, Eva Smorodina, Puneet Rawat, Rahmad Akbar, Kamil Dreczkowski, Rasul Tutunov, Dany Bou-Ammar, Jun Wang, Amos Storkey, and Haitham Bou-Ammar. Toward real-world automated antibody design with combinatorial bayesian optimization. *Cell Reports Methods*, 3(1):100374, January 2023. ISSN 2667-2375. doi: 10.1016/j.crmeth.2022.100374.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint*, 2013. doi: 10.48550/arXiv.1312.6114.

Peter Kirkpatrick and Clare Ellis. Chemical space. *Nature*, 432(7019):823–823, December 2004. ISSN 1476-4687. doi: 10.1038/432823a.

Matt J. Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *PMLR*, pp. 1945–1954. PMLR, 06–11 Aug 2017.

Natalie Maus, Haydn Jones, Juston Moore, Matt J Kusner, John Bradshaw, and Jacob Gardner. Local latent space bayesian optimization over structured inputs. In *Advances in Neural Information Processing Systems*, volume 35, pp. 34505–34518. Curran Associates, Inc., 2022. doi: 10.48550/arXiv.2201.11872.

Natalie Maus, Kaiwen Wu, David Eriksson, and Jacob Gardner. Discovering many diverse solutions with bayesian optimization. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *PMLR*, pp. 1779–1798. PMLR, 25–27 Apr 2023a. doi: 10.48550/arXiv.2210.10953.

Natalie Maus, Yimeng Zeng, Daniel Allen Anderson, Phillip Maffettone, Aaron Solomon, Peyton Greenside, Osbert Bastani, and Jacob R. Gardner. Inverse protein folding using deep bayesian optimization. *arXiv Preprint*, 2023b. doi: 10.48550/arXiv.2305.18089.

Jonas Mockus. The Bayesian approach to global optimization. In *System Modeling and Optimization*, pp. 473–481. Springer, 1982.

Yaniv Nikankin, Niv Haim, and Michal Irani. SinFusion: Training diffusion models on a single image or video. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *PMLR*, pp. 26199–26214. PMLR, 23–29 Jul 2023.

M. A. Osborne, R. Garnett, and S. J. Roberts. Gaussian processes for global optimization. In *3rd International Conference on Learning and Intelligent Optimization (LION3)*, pp. 1–15, 2009.

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.

Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pp. 63–71. Springer, 2003.

Zhizhou Ren, Jiahan Li, Fan Ding, Yuan Zhou, Jianzhu Ma, and Jian Peng. Proximal exploration for model-guided protein sequence design. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *PMLR*, pp. 18520–18536. PMLR, 17–23 Jul 2022.

Philip A. Romero, Andreas Krause, and Frances H. Arnold. Navigating the protein fitness landscape with gaussian processes. *Proceedings of the National Academy of Sciences*, 110(3), December 2012. ISSN 1091-6490. doi: 10.1073/pnas.1215251110.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2016. doi: 10.1109/JPROC.2015.2494218.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Proc NeurIPS*, volume 25, pp. 2951–9, 2012. doi: 10.48550/arXiv.1206.2944.

Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating Bayesian optimization for biological sequence design with denoising autoencoders. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 20459–20478. PMLR, 17–23 Jul 2022.

Francesca Storici and Michael A Resnick. The delitto perfetto approach to in vivo site-directed mutagenesis and chromosome rearrangements with synthetic oligonucleotides in yeast. pp. 329–345, 2006. doi: 10.1016/S0076-6879(05)09019-1.

Brandon Trabucco, Xinyang Geng, Aviral Kumar, and Sergey Levine. Design-bench: Benchmarks for data-driven offline model-based optimization. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *PMLR*, pp. 21658–21676. PMLR, 17–23 Jul 2022. doi: 10.48550/arXiv.2202.08450.

Austin Tripp, Erik Daxberger, and José Miguel Hernández-Lobato. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. In *Advances in Neural Information Processing Systems*, volume 33, pp. 11259–11272. Curran Associates, Inc., 2020. doi: 10.48550/arXiv.2006.09191.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. doi: 10.48550/arXiv.1706.03762.

Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *PMLR*, pp. 370–378, Cadiz, Spain, 09–11 May 2016. PMLR.

Kevin K Yang, Alex Xijie Lu, and Nicolo Fusi. Convolutions are competitive with transformers for protein sequence pretraining. In *ICLR2022 Machine Learning for Drug Discovery*, 2022.

A COMPUTATIONAL COST

We trained the single sequence VAE using the pytorch-lightning library with DDP to distribute the training process across 6×48GB GPUs (NVIDIA RTX A6000) for 20 hours. For our constrained `LOL-BO` runs, we use 1×24GB or 1×48GB GPU (NVIDIA RTX A5000/A6000) and optimize until convergence, each run taking approximately 1 to 3 days.