
Active Learning for Energy-Based Antibody Optimization and Enhanced Screening

Kairi Furui Masahito Ohue*

Institute of Science Tokyo

furui@i.c.titech.ac.jp, ohue@c.titech.ac.jp

* Corresponding author

Abstract

Accurate prediction and optimization of protein-protein binding affinity is crucial for therapeutic antibody development. Although machine learning-based prediction methods $\Delta\Delta G$ are suitable for large-scale mutant screening, they struggle to predict the effects of multiple mutations for targets without existing binders. Energy function-based methods, though more accurate, are time consuming and not ideal for large-scale screening. To address this, we propose an active learning workflow that efficiently trains a deep learning model to learn energy functions for specific targets, combining the advantages of both approaches. Our method integrates the RDE-Network deep learning model with Rosetta’s energy function-based Flex ddG to efficiently explore mutants that bind to Flex ddG. In a case study targeting HER2-binding Trastuzumab mutants, our approach significantly improved the screening performance over random selection and demonstrated the ability to identify mutants with better binding properties without experimental $\Delta\Delta G$ data. This workflow advances computational antibody design by combining machine learning, physics-based computations, and active learning to achieve more efficient antibody development.

1 Introduction

Optimizing the binding affinity between proteins is an important aspect of drug development, including antibody engineering [1, 2]. In recent years, deep learning-based *de novo* molecular design methods have been proposed [3, 4, 5, 6, 7, 8]. However, approaches that utilize existing antibody information to design improved variants are also important, as *de novo* methods may not fully leverage the knowledge of existing antibodies [9, 10, 11, 12]. This is particularly crucial when dealing with known antibodies that require improvements in desirable properties such as binding affinity, specificity, and developability. In recent years, many mutation-induced binding free energy change ($\Delta\Delta G$) prediction methods based on machine learning and deep learning have been proposed [9, 13, 14, 15, 16, 17, 18], and their prediction accuracy has improved. However, in cases where there are insufficient data on existing antibodies, the prediction results may not be reliable. In such situations, even if the predicted $\Delta\Delta G$ for mutants, especially multiple mutations, is optimized, overfitting may occur, making it difficult to propose improved sequences in practice. In addition, traditional methods for calculating $\Delta\Delta G$ of mutants based on energy functions [19, 20, 21] are applicable even when sufficient $\Delta\Delta G$ information about known mutants is not available, and it is known that energy functions with structural sampling, such as Rosetta’s energy function-based $\Delta\Delta G$ prediction method Flex ddG, have high prediction accuracy. However, such methods that involve structural sampling require enormous computational costs compared to machine learning-based methods, making them unsuitable for large-scale screening of mutants.

In recent years, the pharmaceutical industry has shown an increasing interest in using a machine learning approach called active learning [22, 23, 24] to streamline time-consuming calculations such as compound docking and molecular simulations [25, 26, 27]. Active learning is a method that iteratively selects the next data sample to be experimented based on a machine learning model called a surrogate model [22, 23, 24], with the aim of efficiently conducting costly experiments such as experimental verification and molecular simulations. Here, it is important to balance the exploration phase, which reduces uncertainty and selects data with high information content, with the exploitation phase, which discovers data with high activity or desirable properties, to select cost-effective data samples. These methods are useful for exploring promising candidate samples even in the absence of extensive existing experimental data. By iteratively adding data, active learning can guide the optimization process more efficiently than conventional high-throughput screening methods. For example, in the field of small molecule drug discovery, active learning combined with molecular simulations and docking simulations has already succeeded in improving screening performance without relying on experimental information [25, 26, 27]. Similarly, several active learning approaches have been proposed in the field of antibody engineering [28, 29, 30]. Seo *et al.* [28] take an approach that combines a deep learning model based on the CTRL method with experimental validation by ELISA. Lu *et al.* [29] propose a method called K-Nearest Robust Active Learning, KRAL, and apply it to epitope prediction in B cell data. Gessner *et al.* [30] discover antibody sequences with improved binding affinity based on relative binding free energies calculated using Bayesian optimization. Such active learning approaches for experimental information or relative binding free energy calculations provide accurate predictions of binding affinity but require enormous costs, time, and difficult setups. Also, to achieve sufficient performance with these approaches, a sufficient number of samples may be required, so a more convenient approach for faster pre-screening would be useful. Furthermore, among these existing approaches, methods based on relative free energy perturbation calculations can be executed faster than experimental validation approaches if computational resources are available, but it is difficult to handle multiple mutations [31].

Therefore, we propose a new workflow that efficiently learns the scoring of the Flex ddG method [21] through active learning, in contrast to conventional methods, to construct a fast pre-screening model. Flex ddG achieves performance comparable to state-of-the-art models [32] but requires a relatively long computation time for structural sampling. Our workflow guides the selection of promising antibody variants by combining the predictive performance of machine learning with the physics-based knowledge of Flex ddG. In this study, we achieve this by incorporating both experimental $\Delta\Delta G$ values and computational binding information predictions through multitask learning based on the architecture of the RDE-Network developed by Luo *et al.* This allows us to identify antibody sequences that may improve binding affinity while minimizing the need for large-scale experiments and enhancing the predictive performance of Flex ddG’s computational $\Delta\Delta G$. To demonstrate the effectiveness of our method, we optimize Trastuzumab (Herceptin[®]), an antibody targeting the human epidermal growth factor receptor (HER2), and the results showed a significant improvement in screening performance compared to random selection. Furthermore, an interesting result was obtained that the binding classification performance could be improved even without experimental information on HER2 by using additional data based on Flex ddG. This study contributes to the growing field of efficient and effective computational antibody design by presenting a novel workflow combining machine learning, physics-based computation, and active learning.

2 Materials and Methods

Figure 1 shows an overview of the active learning workflow proposed in this study.

Trastuzumab, an antibody targeting the HER2 protein, was used as the antibody to be optimized, and the complementarity-determining regions (CDR-H) of the heavy chain were optimized. In this process, mutant data was sampled by randomly mutating the amino acid sequence of Trastuzumab’s CDR-H (8 residues for CDR-H1, 8 residues for CDR-H2, and 13 residues for CDR-H3) to 19 amino acid residues excluding cysteine with a mutation probability of 0.2, generating 100,000 mutants (of which 98,567 were unique sequences).

Luo *et al.*’s RDE-Network [16] was used as the Flex ddG surrogate model. RDE-Network enables $\Delta\Delta G$ prediction performance by learning information on mutants included in the SKEMPI2 dataset [33] based on the predicted probability distribution, using a flow-based rotamer probability distribution model pretrained on crystal structures. RDE-Network is one of the state-of-the-art models

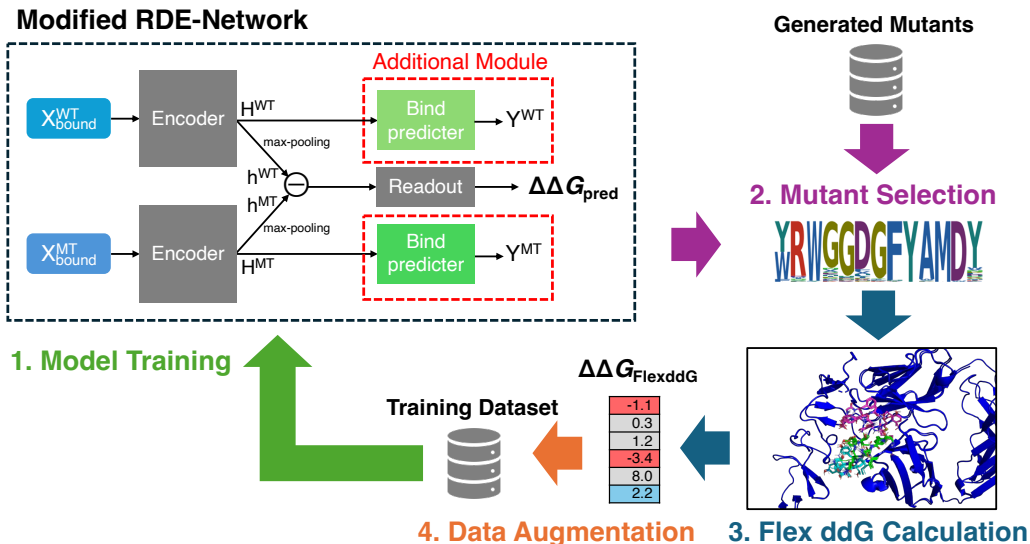


Figure 1: Overview of the proposed active learning workflow.

in terms of $\Delta\Delta G$ prediction performance for mutants. Here, in addition to the $\Delta\Delta G$ prediction module existing in the original RDE-Network prediction model, a Binding Module that predicts whether to bind directly from the output layer of the wildtype and mutation RDE modules was added. In other words, multitask learning was performed using experimental $\Delta\Delta G$ and computational binding information. For the Rotamer Density Estimator, pre-trained parameters were used and fixed during surrogate model training. The learning parameters other than the newly added parts of the model were the same as those of Luo *et al.*, and the cross-entropy loss for binding was added with a weight of 0.5. Unlike the data split of Luo *et al.*, the data split of the learning dataset was based on the Holdout information included in the SKEMPI2 dataset to ensure that PPBs similar to the PPBs of the validation set were not included in the training set. Based on the holdout information, a relationship network of complexes included in the SKEMPI2 dataset was constructed, and the connected components of this network were treated as a single group. Based on these groups, a 3-fold group split was performed, using the fold containing HER2-Trastuzumab data as validation data and the remaining folds as training data. This excluded all $\Delta\Delta G$ information related to HER2 and antibodies from the training set prior to active learning.

In each cycle, 200 trastuzumab mutants were selected and a total of 1200 were selected over 6 cycles of active learning. In each cycle, only mutants that differed by 2 or more mutations from those selected so far, including previous cycles, were selected in descending order of the surrogate model's predictions, to ensure an appropriate balance between exploration and exploitation of the selected sequences. The parameters of Flex ddG were based on those of [21] with the number of structural ensembles set to $N = 10$. The $\Delta\Delta G$ calculated by Flex ddG was added to the learning data in the next active learning cycle as binding data if it was less than 0 and as non-binding data if it was 2 or greater. Data with $0 < \Delta\Delta G < 2$ were excluded from the learning data. Also, in the first cycle, the model was trained by the surrogate model without any HER2 mutation data.

As a test set for model evaluation, surface plasmon resonance (SPR) assay data of Trastuzumab mutants designed by Shanehsazzadeh *et al.* [6] was used. At this time, $\Delta\Delta G$ was obtained from the K_D values of the SPR experiment, and only mutants with the same length of CDR-H3 as Trastuzumab were selected, which will be referred to as the Trastuzumab mutant dataset hereafter. The Trastuzumab mutant dataset includes 419 binding mutants and 846 non-binding mutants.

3 Results and Discussion

Figure 2(a) shows the transition of the distribution of the top 200 Flex ddG values of the data selected in previous cycles. For comparison, the results were compared with the case where 1200 data




Figure 2: (a) Transition of the calculated top 200 Flex ddG values of the selected mutants at each active learning cycle. (b) Number of selected mutants that bound and unbound based on Flex ddG at each active learning cycle.

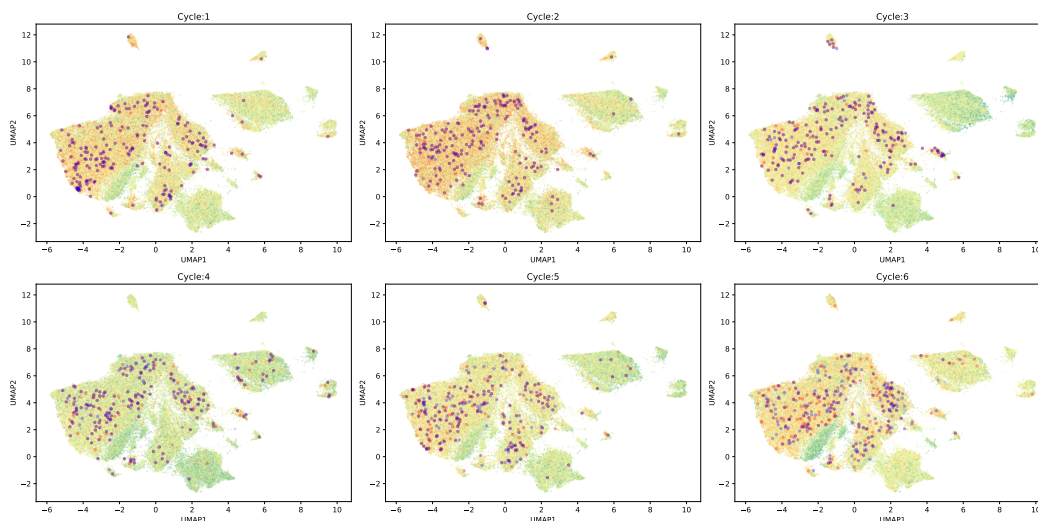


Figure 3: The H-chain sequences of the HER2 mutant dataset were embedded using AbLang [34] and compressed to 2D using UMAP [35] for visualization. The red dots represent the selected mutants at each active learning cycle, and the blue dots represent the top 200 selections for that cycle. In the early cycles of active learning, the blue and red dots overlap, but in the later cycles, the red dots spread out more than the blue dots.

points were randomly selected from the data pool and evaluated by Flex ddG. The figure shows that the proposed method can discover Trastuzumab mutants with significantly lower Flex ddG values compared to random selection in the final cycle. In the early stages of learning, the top predictions of RDE-Network tended to select worse mutants than Random, but by using them for learning in subsequent cycles, the enrichment was improved. Furthermore, in Figure 2(b), the number of selected mutants that bound and unbound based on Flex ddG at each active learning cycle, it can be seen that in the initial learning, many unbound data were sampled due to inaccurate model predictions, but in the latter half of learning, almost no non-binding data were selected. Figure 2(c) shows the transition of the rank of the predicted $\Delta\Delta G$ of the selected mutants. In the early cycles, the selection is close to greedy selection, but in the latter half of learning, mutants with worse ranks are selected, which means that more diverse mutants were selected, i.e., more exploratory sampling was performed, due to the constraint of selecting only those that differed by 2 or more mutations from the previously selected ones. This can also be seen from the difference between greedy selection and actual selection in the embedded space in Figure 3. Thus, balancing exploitation and exploration is important for more efficient learning of active learning models.

Furthermore, Figures 4(a) to (d) show the Spearman correlation with experimental $\Delta\Delta G$ or Flex ddG's $\Delta\Delta G$ and the ROC-AUC for experimental binding and Flex ddG-based binding at each active learning cycle, respectively. The Spearman correlation was calculated only for the binding mutants in

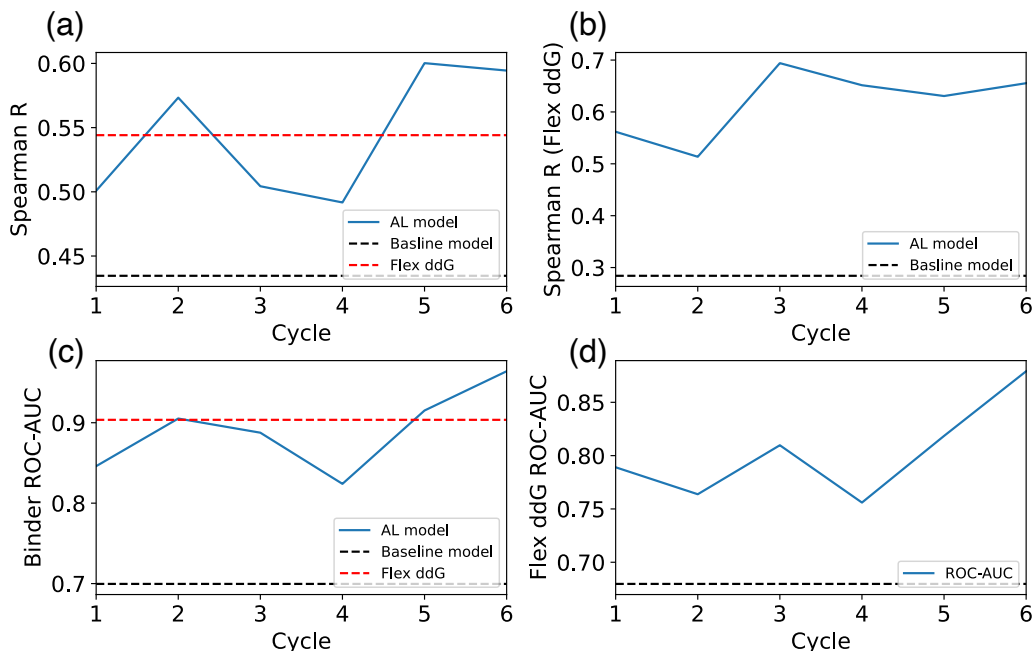


Figure 4: (a) Transition of Spearman correlation of the surrogate model for the Trastuzumab mutant dataset at each active learning cycle. (b) Transition of Spearman correlation on Flex ddG of the surrogate model for the Trastuzumab mutant dataset at each active learning cycle. Each correlation was calculated only for those labeled as binders. (c) Transition of ROC-AUC scores for binding classification for the Trastuzumab mutant dataset at each active learning cycle. (d) Transition of ROC-AUC scores for binding classification on Flex ddG for the Trastuzumab mutant dataset at each active learning cycle. The target variable was set as “binder” when $\Delta\Delta G_{\text{FlexddG}} < 0$.

the Trastuzumab mutant dataset, and for the binding ROC-AUC, 50 data points each were selected from the experimental binding data ($\Delta\Delta G_{\text{FlexddG}} < 0$) and randomly selected non-binding data for a total of 100 data points for comparison with Flex ddG calculation. From Figures 4(a) and 4(b), it can be seen that the predictive performance of the active learning model slightly improved from the baseline and reached a level comparable to Flex ddG. However, the correlation with experimental $\Delta\Delta G$ did not change even as the number of active learning cycles increased. Also, Figures 4(c) and 4(d) show a significant improvement in performance compared to the baseline model. Furthermore, in the latter half of active learning, the discriminative performance surpassed that of Flex ddG. This is an interesting result showing that the actual screening performance can be improved without using experimental information, just by adding the computational values of Flex ddG to the learning. From the above, the proposed method seems to be effective in improving the binding classification performance in particular.

4 Conclusions

In this study, we developed an active learning workflow for optimizing antibody sequences. The proposed method successfully screened Trastuzumab variants against HER2 efficiently. By utilizing active learning models as proposed in this study, we succeeded in efficiently screening by evaluating only a subset of the vast number of Trastuzumab variants against human HER2 using Flex ddG, without calculating the energy of all sequences. Furthermore, by combining the SKEMPI2 dataset and the computational energy information of variants’ Flex ddG, RDE-Network was able to more accurately discriminate the binding of human HER2 without learning the experimental $\Delta\Delta G$ information. This suggests that it may be useful for exploring variants with higher binding in the early stages of antibody design, where only a few weak binders are obtained, or for optimizing variants with better pharmacological properties. In this study, RDE-Network was used as a surrogate model for active learning, but this framework is applicable to any model and can be useful for easily constructing

models tailored to specific targets. However, the active learning method proposed in this study is primarily for optimization and requires the preparation of complex structures of known binding ligands. Furthermore, since the sequences evaluated in this study have not been experimentally validated, the effect may be limited if Flex ddG calculations do not correlate well with the experimental $\Delta\Delta G$, and it should be noted that in actual drug development phases, the developability of the selected variants should also be considered. This study investigated offline-based optimization, but combining it with genetic algorithms, as in other cases, may cover a more diverse mutant space. In the future, by combining active learning with more accurate affinity prediction methods that require even more extensive calculations, such as relative binding free energy perturbation calculations, we aim to construct models with higher binding affinity prediction accuracy for specific targets and improve screening accuracy.

References

- [1] Jisun Kim, Matthew McFee, Qiao Fang, Osama Abdin, and Philip M Kim. Computational and artificial intelligence-based methods for antibody development. *Trends Pharmacol. Sci.*, 44(3):175–189, 2023.
- [2] Sara Joubbi, Alessio Micheli, Paolo Milazzo, Giuseppe Maccari, Giorgio Ciano, Dario Cardamone, and Duccio Medini. Antibody design using deep learning: from sequence and structure design to affinity maturation. *Brief. Bioinform.*, 25(4):bbae307, 07 2024.
- [3] Richard W Shuai, Jeffrey A Ruffolo, and Jeffrey J Gray. IgLM: Infilling language modeling for antibody sequence design. *Cell Syst.*, 14(11):979–989, 2023.
- [4] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Antibody-antigen docking and design via hierarchical equivariant refinement. *arXiv preprint arXiv:2207.06616*, 2022.
- [5] Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, et al. AbDiffuser: full-atom generation of in-vitro functioning antibodies. *Adv. Neural Inf. Process. Syst.*, 36, 2024.
- [6] Amir Shanehsazzadeh, Sharrol Bachas, Matt McPartlon, George Kasun, John M Sutton, Andrea K Steiger, Richard Shuai, Christa Kohnert, Goran Rakocevic, Jahir M Gutierrez, et al. Unlocking de novo antibody design with generative artificial intelligence. *bioRxiv*, 2023.
- [7] Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific antibody design and optimization with diffusion-based generative models for protein structures. *Adv. Neural Inf. Process. Syst.*, 35:9754–9767, 2022.
- [8] Takafumi Ueki and Masahito Ohue. Antibody complementarity-determining region design using AlphaFold2 and DDG predictor. *J. Supercomput.*, 80:11989–12002, 2024.
- [9] Sisi Shan, Shitong Luo, Ziqing Yang, Junxian Hong, Yufeng Su, Fan Ding, Lili Fu, Chenyu Li, Peng Chen, Jianzhu Ma, et al. Deep learning guided optimization of human antibody against SARS-CoV-2 variants with broad neutralization. *Proc. Natl. Acad. Sci.*, 119(11):e2122954119, 2022.
- [10] Yimeng Zeng, Hunter Elliott, Phillip Maffettone, Peyton Greenside, Osbert Bastani, and Jacob R Gardner. Antibody design with constrained bayesian optimization. In *ICLR 2024 Workshop on GEM*.
- [11] Magnus Høie, Alissa Hummer, Tobias Olsen, Morten Nielsen, and Charlotte Deane. AntiFold: Improved antibody structure design using inverse folding. In *NeurIPS 2023 GenBio Workshop*, 2023.
- [12] Ariel Tennenhouse, Lev Khmelnsky, Razi Khalaila, Noa Yeshaya, Ashish Noronha, Moshit Lindzen, Emily K Makowski, Ira Zaretsky, Yael Fridmann Sirkis, Yael Galon-Wolfenson, et al. Computational optimization of antibody humanness and stability by systematic energy-based ranking. *Nat. Biomed. Eng.*, 8(1):30–44, 2024.

- [13] Minghui Li, Franco L Simonetti, Alexander Goncarenco, and Anna R Panchenko. MutaBind estimates and interprets the effects of sequence variants on protein–protein interactions. *Nucleic Acids Res.*, 44(W1):W494–W501, 2016.
- [14] Ning Zhang, Yuting Chen, Haoyu Lu, Feiyang Zhao, Roberto Vera Alvarez, Alexander Goncarenco, Anna R. Panchenko, and Minghui Li. MutaBind2: Predicting the impacts of single and multiple mutations on protein-protein interactions. *iScience*, 23(3):100939, 2020.
- [15] Carlos HM Rodrigues, Yoochan Myung, Douglas EV Pires, and David B Ascher. mCSM-PPI2: predicting the effects of mutations on protein–protein interactions. *Nucleic Acids Res.*, 47(W1):W338–W344, 2019.
- [16] Shitong Luo, Yufeng Su, Zuofan Wu, Chenpeng Su, Jian Peng, and Jianzhu Ma. Rotamer density estimator is an unsupervised learner of the effect of mutations on protein-protein interaction. *bioRxiv*, 2023.
- [17] Guanglei Yu, Qichang Zhao, Xuehua Bi, and Jianxin Wang. DDAffinity: predicting the changes in binding affinity of multiple point mutations using protein 3d structure. *Bioinformatics*, 40(Supplement_1):i418–i427, 06 2024.
- [18] Shiwei Liu, Tian Zhu, Milong Ren, Chungong Yu, Dongbo Bu, and Haicang Zhang. Predicting mutational effects on protein-protein binding via a side-chain diffusion probabilistic model. *Adv. Neural Inf. Process. Syst.*, 36, 2024.
- [19] P Benjamin Stranges and Brian Kuhlman. A comparison of successful and failed protein interface designs highlights the challenges of designing buried hydrogen bonds. *Prot. Sci.*, 22(1):74–82, 2013.
- [20] Oliver Buß, Jens Rudat, and Katrin Ochsenreither. FoldX as protein engineering tool: better than random based approaches? *Comput. Struct. Biotechnol. J.*, 16:25–33, 2018.
- [21] Kyle A. Barlow, Shane Ó Conchúir, Samuel Thompson, Pooja Suresh, James E. Lucas, Markus Heinonen, and Tanja Kortemme. Flex ddG: Rosetta ensemble-based estimation of changes in protein–protein binding affinity upon mutation. *J. Phys. Chem. B*, 122(21):5389–5399, 2018. PMID: 29401388.
- [22] Burr Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin-Madison, 2009.
- [23] Daniel Reker and Gisbert Schneider. Active-learning strategies in computer-assisted drug discovery. *Drug Discov. Today*, 20(4):458–465, 2015.
- [24] Daniel Reker, Petra Schneider, Gisbert Schneider, and J B Brown. Active learning for computational chemogenomics. *Future Med. Chem.*, 9(4):381–402, 2017.
- [25] Lei Wang, Shao-Hua Shi, Hui Li, Xiang-Xiang Zeng, Su-You Liu, Zhao-Qian Liu, Ya-Feng Deng, Ai-Ping Lu, Ting-Jun Hou, and Dong-Sheng Cao. Reducing false positive rate of docking-based virtual screening by active learning. *Brief. Bioinform.*, 24(1):bbac626, 2023.
- [26] Kyle D Konze, Pieter H Bos, Markus K Dahlgren, Karl Leswing, Ivan Tubert-Brohman, Andrea Bortolato, Braxton Robbason, Robert Abel, and Sathesh Bhat. Reaction-based enumeration, active learning, and free energy calculations to rapidly explore synthetically tractable chemical space and optimize potency of cyclin-dependent kinase 2 inhibitors. *J. Chem. Inf. Model.*, 59(9):3782–3793, 2019.
- [27] Hannes Loeffler, Shunzhou Wan, Marco Klähn, Agastya Bhati, and Peter Coveney. Optimal molecular design: Generative active learning combining REINVENT with absolute binding free energy simulations. *ChemRxiv*, 2024.
- [28] Seung-woo Seo, Min Woo Kwak, Eunji Kang, Cha Eun Kim, Eunyoung Park, Tae Hyun Kang, and Jinhan Kim. Accelerating antibody design with active learning. *BioRxiv*, 2022.
- [29] Tianchi Lu. K-nearest robust active learning on big data and application in epitope prediction. *Wirel. Commun. Mob. Comput.*, 2021(1):8752022, 2021.

- [30] Alexandra Gessner, Sebastian W Ober, Owen Vickery, Dino Ogljić, and Talip Uçar. Active learning for affinity prediction of antibodies. *arXiv preprint arXiv:2406.07263*, 2024.
- [31] Fangqiang Zhu, Feliza A Bourguet, William FD Bennett, Edmond Y Lau, Kathryn T Arrildt, Brent W Segelke, Adam T Zemla, Thomas A Desautels, and Daniel M Faissol. Large-scale application of free energy perturbation calculations for antibody design. *Sci. Rep.*, 12(1):12489, 2022.
- [32] Wei Lu, Jixian Zhang, Jihua Rao, Zhongyue Zhang, and Shuangjia Zheng. AlphaFold3, a secret sauce for predicting mutational effects on protein-protein interactions. *bioRxiv*, 2024.
- [33] Justina Jankauskaitė, Brian Jiménez-García, Justas Dapkūnas, Juan Fernández-Recio, and Iain H Moal. SKEMPI 2.0: an updated benchmark of changes in protein–protein binding energy, kinetics and thermodynamics upon mutation. *Bioinformatics*, 35(3):462–469, 2019.
- [34] Tobias H Olsen, Iain H Moal, and Charlotte M Deane. AbLang: an antibody language model for completing antibody sequences. *Bioinform. Adv.*, 2(1):vbac046, 2022.
- [35] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.