AbAffinity: A Large Language Model for Predicting Antibody Binding Affinity against SARS-CoV-2

Faisal Bin Ashraf¹ Animesh Ray² Stefano Lonardi¹

¹Department of Computer Science and Engineering, University of California, Riverside, 92521, CA, United States ²Riggs School of Applied Life Sciences, Keck Graduate Institute, Claremont, 91711, CA, United States

Abstract

Machine learning-based antibody design is emerging as one of the most promising approaches to combat infectious diseases, due to significant advancements in the field of artificial intelligence and an exponential surge in experimental antibody data (in particular related to COVID-19). The ability of an antibody to bind to an antigens (called *binding affinity*) is one of the the most critical properties in designing neutralizing antibodies. In this study we introduce Ab-Affinity, a new large language model that can accurately predict the binding affinity of antibodies against a target peptide, e.g., the SARS-CoV-2 spike protein.

Introduction

Antibodies are proteins composed of two identical polypeptide chains, termed "heavy chains" and two identical "light chains", respectively, connected by disulfide bonds. Each light chain consists of one variable domain and one constant domain, while each heavy chain contains one variable domain and 3-4 constant domains. Each antibody possesses an antigen binding site, called paratope, typically buried within the variable domain's complementarity-determining regions (CDRs). The paratope recognizes and binds specifically to a target region on the surface of an antigen, known as the epitope. CDRs together constitute the binding site of the antibody, and their sequence diversity allows antibodies to recognize a wide range of antigens with high specificity. The binding interaction between the antibody and antigen is highly specific. This binding may trigger various immune responses, including the neutralization of a live virus, a toxin or an antigen with catalytic activities, leading to the eventual elimination of the antigen from the organism.

Measuring the binding affinity of an antibody against an antigen is critical for understanding immune responses elicited by vaccines and infections, as well as for designing antibodies with therapeutic applications. Such measurements often involve biophysical and biochemical techniques such as Surface Plasmon Resonance (SPR), Enzyme-Linked Immunosorbent Assay (ELISA), or Bio-Layer Interferometry (BLI), which generally includes immunization of animals to obtain a host of candidate antibodies, their purifi-





Figure 1: Overview of Ab-Affinity. Ab-Affinity predicts the binding affinity against a specific target peptide; it can also provide residue-residue contact maps and an embedding of the input sequence.

cation, and testing of the candidates for choosing the optimum molecules that exhibit high-specificity and sufficient strength of epitope-paratope interaction. These approaches incur significant time and cost. Computational approaches hold promise to replace some of these time-intensive steps to narrow down the list of candidate antibodies for final testing.

Predicting the specificity and strength of antigen-antibody interaction for protein antigens is a sub-problem of predicting specific protein-protein interactions. However, the former problem is harder because the paratope and the epitope regions of many antibodies and their cognate antigens are often flexible, and contain relatively mobile intrinsically disordered protein regions (IDPRs) when not in complex with the cognate antigen (Uversky and Van Regenmortel 2021; MacRaild et al. 2016). Structures of flexible and mobile regions of proteins are relatively poorly represented in protein structure databases because they are technically harder to determine; training data for learning structural rules are consequently sparse. Predicting these interactions from sequences can be advantageous if the structural and functional roles of amino acids in the sequence are effectively identified and modeled.

The problem of predicting binding affinity has been extensively studied in the literature. For instance, ISLAND (Abbasi et al. 2020) uses a comprehensive feature extraction from protein sequence combined with regression but achieves a modest predictive accuracy. DeepDTA (Öztürk, Özgür, and Ozkirimli 2018) improves upon ISLAND by relying on Convolutional Neural Networks (CNN) for proteinligand interactions. More recent models like DG-Affinity (Yuan et al. 2023) and CSM-AB (Myung, Pires, and Ascher 2022) integrate large language models and graph-based techniques, thus enhancing the accuracy of predictions. Tag-LLM (Shen et al. 2024) and FAbCon (Barton et al. 2024) further refine these capabilities by using task-specific tags and fine-tuned large language models for antibody-antigen interactions. The use of pre-trained protein language models in DG-Affinity (Yuan et al. 2023) has shown superior performance in predicting antibody binding affinities compared to many other methods. Large language models were also used to predict the binding affinity of antibodies against the SARS-CoV-2 spike protein (Li et al. 2023). However, due to the specificity of antibodies to antigens, existing multi-target affinity prediction models are unsuitable for SARS-CoV-2. Therefore, an efficient binding affinity prediction model is essential targeting this virus.

In this study, we introduce a large language model called Ab-Affinity that was trained on single-chain fragment variable (scFv) sequences from multiple libraries of engineered antibodies. Ab-Affinity provides meaningful representation (i.e., embedding) of antibodies which align with binding properties and thermostability proprieties against SARS-CoV-2, and can be used to create the binding landscape of antibodies. Second, Ab-Affinity provides intra-residue attention maps of antibodies which can be used to explain differences of amino-acid interactions for strong and weak binding. Figure 1 illustrates the overall framework of this study.

Dataset

Methods

The training used a dataset of single-chain fragment variable (scFv) antibody sequences with the associated binding scores against a peptide in the SARS-CoV-2 HR2 region (Engelhart et al. 2022). The SARS-CoV-2 spike protein has mutated, forming variants such as Wuhan, Alpha, Delta, Omicron, and others. The selected HR2 peptide, conserved across all variants of SARS and MERS spike proteins, is crucial for evaluating antibodies against the broader coronavirus group. The dataset was obtained by introducing one, two, and three amino acid changes into the sequences of three candidate antibodies obtained from a phage display library that bound to the HR2 antigen polypeptide: Ab-14-VH and Ab-14-VL, Ab-91-VH, Ab-95-VH and Ab-95-VL. The binding affinity (K_D) of each of the 104,972 resulting variants the selected target peptide was estimated in three independent biological replicates (Engelhart et al. 2022). Each K_D value was the equilibrium dissociation constant estimated by an indirect competitive binding assay. The experimental assay was conducted in triplicate for each antibody, and the full dataset contained three K_D values for each antigen-antibody pair. We preprocessed the interaction data by taking the arithmetic mean of the two closest K_D values for each antigen-antibody interaction and by eliminating the third value to reduce the outlier effect. Antibodies that had missing values across all three replicates were disregarded from the analysis. After preprocessing, a total of 71,834 unique antibodies were utilized for model training. Ab-Affinity was trained on all antibody sequences (from each seed antibody) and their corresponding log-transformed binding affinity $log_{10}(K_d)$. Distributions of training antibodies are shown in Figure 2.

Model Architecture

Ab-Affinity's architecture is based on BERT (Devlin et al. 2018), as implemented in ESM-2 (Lin et al. 2023) (see Figure 3). We chose the BERT architecture because the dataset consists of amino acid sequences with a few point mutations. BERT effectively captures long-range dependencies, providing meaningful representations that align with binding affinity changes caused by these mutations. Ab-Affinity contains N sequential layers of encoder blocks. Each encoder block consists of multi-head attention layers (Vaswani et al. 2017) followed by feed-forward layers. The value of Ndetermines the model size. We have tested N = 6, 12, and 33, which resulted in a model with 8M, 35M, and 650Mparameters, respectively. We chose these values for N based on the ESM-2 study (Lin et al. 2023), which demonstrated the impact of model size on the performance of this BERTbased architecture. We used the last encoder layer output as the sequence representation (i.e., the embedding). Depending on the choice of N the embedding of the sequence had 320, 480 and 1280 dimensions, respectively. We added one fully connected layer to predict the binding affinity from the embedding of the sequence.

Training

We used Mean Squared Error (MSE) as the loss function, and Adam optimizer to optimize the parameters. We used 85% of the data to train the model (maintaining the distribution of affinity values). The remaining 15% was used to validate the model. We trained our model using four cores NVIDIA A100 (80GB) with a batch size of 128 for 100 epochs. For the amino-acid contact analysis, we used the method by Rao et. al. (Rao et al. 2020). We fine-tuned the pre-trained ESM-2 protein language model encoder layers, leveraging the knowledge acquired from the entire protein database to train our model. Additionally, we trained a model with randomly initiated weights to understand the impact of pretrained protein knowledge on binding affinity prediction. We saved the best-performing model based on the Pearson correlation coefficient on the validation set for each setup of training. Ab-Affinity, with 33 layers and finetuned from ESM-2, is the best-performing model among all the trained models.

t-SNE visualization of embeddings

The T-Distributed Stochastic Neighbor Embedding (t-SNE) (Van der Maaten and Hinton 2008) was used to reduce high dimensional sequence representation to 2-D space to plot the antibodies. We carried out the dimensionality reduction using Python scikit-learn package and set the perplexity at 200.



Figure 2: Dataset description. a. Distribution of Binding Affinity (log K_D); b. Distribution of antibodies from each seed; c. Distribution of antibodies with mutation.



Figure 3: Model Architecture of Ab-Affinity

Experimental Results

Ab-Affinity Predicts the Binding Affinity of Antibodies Against SARS-CoV-2

We first investigated whether Ab-Affinity creates meaningful latent space features (i.e., embeddings) for the antibodies after training. To this end, we visualized Ab-Affinity's embeddings in a 2D space using t-distributed Stochastic Neighbor Embedding (t-SNE) maps. To compare the performance of our methods with the best existing method, we also visualized the embeddings produced by ESM-2.

In the t-SNE plots in Figure 5 each point is an antibody and the color illustrates its corresponding binding affinity. Observe in Figure 5 that the embeddings produced by Ab-Affinity display a smooth gradient of binding affinity, i.e., the value of $\log K_d$ monotonically decreases. By contrast, the ESM-2 embeddings in Figure 5 do not show a clear gradient descent along the t-SNE components.

To evaluate the performance of our method, we then compared Ab-Affinity's ability to predict binding affinity with those of three other LLM-based methods, namely DG-Affinity, ESM-2, and AbLang. The same test dataset, which was held out from training and not utilized during model development, was used to predict binding affinity and compare the performance of the methods. To proceed with this comparison, We recall that DG-Affinity uses a pre-trained language model combined with a ConvNext-based architecture for prediction, which outperforms 26 other methods on an independent antibody dataset (Yuan et al. 2023). ESM-2 was pre-trained on the UniRef50 protein database and generated 640-dimensional antibody sequence embeddings, which we fed into a simple linear regression model. AbLang uses separate embeddings for the heavy and light chains of antibodies, which we concatenated into a 1534-dimensional vector and used in a linear regression model to predict the affinity score.

Figure 4 illustrates (a) Pearson and Spearman-rank correlation coefficients for the four methods on the test set, (be) scatter plots for actual vs. predicted affinity for the four methods. The highest Pearson/Spearman correlation coefficients were observed with the predictions from Ab-Affinity. Observe that DG-Affinity has the lowest Pearson correlation coefficient of 0.194, presumably because although it used an LLM pre-trained on antibodies, the architecture head itself was not trained with any SARS-CoV-2-specific data. Surprisingly, ESM-2, which was pre-trained on "generic" proteins, produced a good Pearson/Spearman correlation. AbLang, which is instead trained on antibody sequences,



Figure 4: Comparing affinity prediction models: (a) Pearson and Spearman correlation for DG-Affinity (p-values = 3.86×10^{-14} , and 3.38×10^{-15}), ESM-2 (p-values = 8.03×10^{-198} , and 8.02×10^{-198}), AbLang (p-values = 1.24×10^{-175} , and 1.02×10^{-163}) and Ab-Affinity (p-values = 4.03×10^{-261} , and 8.65×10^{-217}); scatter plot for actual vs. predicted binding affinity for (b) DG-Affinity, (c) ESM-2, (d) AbLang, (e) Ab-Affinity (includes antibodies for all three seeds)



Figure 5: t-SNE representation of the embedding produced by ESM-2 and Ab-Affinity; antibodies are colored according to their predicted binding affinity.

also achieved a good Pearson/Spearman correlation. The scatter plots of actual vs. predicted binding affinity in Figure 4(b-e) show that the best-fit lines for Ab-Affinity have the highest slope, indicating the best predictive accuracy.

To explore whether the prediction methods are sensitive to the choice of the test dataset, we compared Ab-Affinity with the other three methods using a different dataset. Table 1 compares the performance of CNN-based and LLM-based methods for predicting binding affinity to the same peptide. We used the 14H dataset (heavy chain sequences) and 14L dataset (light chain sequences) which are the mutated version of Ab-14 heavy and light chain sequences, respectively. To fit the heavy chain sequences from 14H dataset, the seed light chain sequence of Ab-14 was paired with different heavy chains from 14H to form scFv sequences. Similarly, the seed heavy chain was paired with different light chains from 14L to form additional scFv sequences. We observed that the Pearson correlation coefficient of both 14H and 14L antibodies generated Ab-Affinity were the highest among all models. While the best Spearman correlation for 14H was obtained by A2Binder, that of Ab-Affinity closely followed. Note that these correlation coefficients are specific to predicting the binding affinity for antibodies derived from the Ab-14 seed antibody. However, Ab-Affinity not only shows a superior correlation for Ab-14 but also for all three seed antibodies (Figure 4).

Ab-Affinity's Embeddings enable Downstream Classification Tasks

We used Ab-Affinity embeddings to carry out two downstream classification tasks related to antibody binding characteristics, namely (i) the problem of determining the antibody binding affinity classes (High, Medium, Low), and (ii) the problem of determining whether the binding affinity of an antibody was stronger than the binding affinity of the corresponding seed antibody (Yes, No). Figure 6 reports the Receiver Operating Characteristic (ROC) curves for the

Table 1: Comparing Ab-Affinity against other binding affinity predictors on 14H and 14L dataset

| Model | Ref | 14H | | 14L | |
|---------------|---------------------------------|---------|----------|---------|----------|
| | | Pearson | Spearman | Pearson | Spearman |
| Ens-Grad | (Liu et al. 2020) | 0.601 | 0.476 | 0.637 | 0.645 |
| ESM-F | (He et al. 2024) | 0.634 | 0.516 | 0.674 | 0.681 |
| AntiBERTa2 | (Barton, Galson, and Leem 2024) | 0.623 | 0.545 | 0.673 | 0.684 |
| AbMAP | (Singh et al. 2023) | 0.606 | 0.510 | 0.674 | 0.685 |
| A2Binder | (He et al. 2024) | 0.642 | 0.553 | 0.683 | 0.688 |
| Ensembles-14H | (Li et al. 2023) | N/A | 0.512 | N/A | N/A |
| Ensembles-14L | (Li et al. 2023) | N/A | N/A | N/A | 0.688 |
| Ab-Affinity | [this] | 0.652 | 0.526 | 0.712 | 0.713 |

two classification tasks, comparing the classifiers' AUC using the embeddings produced by Ab-Affinity against ESM-2. For both tasks the classifiers built using Ab-Affinity's embedding has much higher AUC values, suggesting that Ab-Affinity's embeddings are more informative for downstream applications.



Figure 6: ROC curves and AUC values for two classification tasks, namely (**a**,**b**) determining the binding affinity class of an antibody (High, Medium, Low) and (**c**,**d**) determining whether the binding is improved compared to the seed antibody (Y/N); (**a**,**c**) ROC curves using the ESM-2 embedding; (**b**,**d**) ROC curves using the Ab-Affinity embedding; AUC values are reported for each classifier

Ab-Affinity's Attention Maps for Strong and Weak Binding Antibodies

We extracted attention maps from the Ab-Affinity model to determine which residue-residue interactions were more important for the prediction of binding affinity. We produced residue-residue contact maps for heavy and light chains of antibodies with strong binding affinity (i.e., $\log K_d < 0.5$) and antibodies with weak binding affinity (i.e., $\log K_d > 5.5$). Figure 7 illustrates the residue-residue differences between the attention maps for strong and weak binding anti-



Figure 7: Differences between the attention residue-residue maps for strong binding (i.e. $\log K_d < 0.5$) and weak binding (i.e., $\log K_d > 5.5$) for antibodies generated from Ab-14 (a) heavy chain, and (b) light chain.

bodies for both heavy and light chains. The CDRs are highlighted in pink, green and purple in Figure 7. Observe that many of the strongest differences occur in CDR-H1, CDR-H2, CDR-L1 or the regions immediately adjacent to them.

Ab-Affinity Captures Thermostabilty Properties of Antibodies

The ability of a protein to resist degradation or changes to its physical structure at higher temperature is called thermostability. This property is essential for therapeutic antibodies as it determines their structural integrity and functionality at higher temperature. Here we investigated the correlation of Ab-Affinity's embeddings with the thermostability of the corresponding antibody. For this task, we used studies (Hie et al. 2024; Rosace et al. 2023) to create a small dataset of experimentally-determined thermostability for 26 SARS-CoV-2 antibodies. In the t-SNE plots in Figure 8 each point is an antibody and the color illustrates its corresponding thermostability. Observe that in the Ab-Affinity's embeddings, antibodies are clearly separated into two clusters, each with relatively similar thermostability values. It is not the case for the ESM-2 embeddings. One explanation of this finding is that the fine-tuned language model has learned intrinsic rules that not only govern the physics of antigenantibody interaction but also the physics of protein stability that contributes to the stable interactions over a range of temperatures.



Figure 8: t-SNE visualization of the embeddings produced by ESM-2 and Ab-Affinity; points are colored according to the experimentally-determined thermostability of those antibodies

Conclusion

We introduced Ab-Affinity, a large language model that can predict the binding affinity of antibodies against the SARS-CoV-2 spike protein. In our experiments, Ab-Affinity demonstrated stronger predictive performance than other methods in the literature. The model was able to learn the effect of mutations on binding with the SARS-CoV-2 spike protein. The t-SNE visualization of Ab-Affinity's embeddings suggested that our approach captured the relationships between antibody sequences and binding affinities. Our experiments also show that Ab-Affinity's embeddings enabled a strong performance on several classification tasks, emphasizing the utility of our model beyond affinity prediction. The residue-residue attention maps extracted from Ab-Affinity reveal that our model appears to focus on the CDR or the neighboring regions, which are the binding sites in the antibody. Ab-Affinity also demonstrated an "understanding" of thermostability, despite not being explicitly trained on this attribute.

In conclusion, Ab-Affinity represents a significant advancement in the prediction of antibody binding affinities, demonstrating both high accuracy and versatility in its applications. Our findings highlight its potential as a powerful tool for antibody design and further exploration of proteinantibody interactions.

The model is readily available for use and can be easily installed on any machine via PyPi. For convenience, users can access and install Ab-Affinity directly from the Python Package Index (PyPi) using the following link: https://pypi.org/project/AbAffinity/. This ensures seamless integration into various research environments and workflows.

Acknowledgment

The work reported here was funded by the National Institute of Allergy and Infectious Diseases (NIAID) of the National Institutes of Health (NIH) under the grant 3R01AI169543 to SL and AR.

References

Abbasi, W. A.; Yaseen, A.; Hassan, F. U.; Andleeb, S.; and Minhas, F. U. A. A. 2020. Island: in-silico proteins binding affinity prediction using sequence information. *BioData Mining* 13:1–13.

Barton, J.; Gaspariunas, A.; Yadin, D. A.; Dias, J.; Nice, F. L.; Minns, D. H.; Snudden, O.; Povall, C.; Valle Tomas, S.; Dobson, H.; et al. 2024. A generative foundation model for antibody sequence understanding. *bioRxiv* 2024–05.

Barton, J.; Galson, J. D.; and Leem, J. 2024. Enhancing antibody language models with structural information. *bioRxiv* 2023–12.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Engelhart, E.; Emerson, R.; Shing, L.; Lennartz, C.; Guion, D.; Kelley, M.; Lin, C.; Lopez, R.; Younger, D.; and Walsh, M. E. 2022. A dataset comprised of binding interactions for 104,972 antibodies against a sars-cov-2 peptide. *Scientific Data* 9(1):653.

He, H.; He, B.; Guan, L.; Zhao, Y.; Jiang, F.; Chen, G.; Zhu, Q.; Chen, C. Y.-C.; Li, T.; and Yao, J. 2024. De novo generation of sars-cov-2 antibody cdrh3 with a pre-trained generative large language model. *Nature Communications* 15(1):6867.

Hie, B. L.; Shanker, V. R.; Xu, D.; Bruun, T. U.; Weidenbacher, P. A.; Tang, S.; Wu, W.; Pak, J. E.; and Kim, P. S. 2024. Efficient evolution of human antibodies from general protein language models. *Nature Biotechnology* 42(2):275–283.

Li, L.; Gupta, E.; Spaeth, J.; Shing, L.; Jaimes, R.; Engelhart, E.; Lopez, R.; Caceres, R. S.; Bepler, T.; and Walsh, M. E. 2023. Machine learning optimization of candidate antibody yields highly diverse sub-nanomolar affinity antibody libraries. *Nature Communications* 14(1):3454.

Lin, Z.; Akin, H.; Rao, R.; Hie, B.; Zhu, Z.; Lu, W.; Smetanin, N.; Verkuil, R.; Kabeli, O.; Shmueli, Y.; et al. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* 379(6637):1123–1130. Liu, G.; Zeng, H.; Mueller, J.; Carter, B.; Wang, Z.; Schilz, J.; Horny, G.; Birnbaum, M. E.; Ewert, S.; and Gifford, D. K. 2020. Antibody complementarity determining region design using high-capacity machine learning. *Bioinformatics*

36(7):2126–2133. MacRaild, C. A.; Richards, J. S.; Anders, R. F.; and Norton, R. S. 2016. Antibody recognition of disordered antigens. *Structure* 24(1):148–157.

Myung, Y.; Pires, D. E.; and Ascher, D. B. 2022. Csm-ab: graph-based antibody-antigen binding affinity prediction and docking scoring function. *Bioinformatics* 38(4):1141–1143.

Öztürk, H.; Özgür, A.; and Ozkirimli, E. 2018. Deepdta: deep drug–target binding affinity prediction. *Bioinformatics* 34(17):i821–i829.

Rao, R. M.; Meier, J.; Sercu, T.; Ovchinnikov, S.; and Rives, A. 2020. Transformer protein language models are unsupervised structure learners. *bioRxiv*.

Rosace, A.; Bennett, A.; Oeller, M.; Mortensen, M. M.; Sakhnini, L.; Lorenzen, N.; Poulsen, C.; and Sormanni, P. 2023. Automated optimisation of solubility and conformational stability of antibodies and proteins. *Nature communications* 14(1):1937.

Shen, J.; Tenenholtz, N.; Hall, J. B.; Alvarez-Melis, D.; and Fusi, N. 2024. Tag-llm: Repurposing general-purpose llms for specialized domains. *arXiv preprint arXiv:2402.05140*.

Singh, R.; Im, C.; Qiu, Y.; Mackness, B.; Gupta, A.; Sorenson, T.; Sledzieski, S.; Erlach, L.; Wendt, M.; Nanfack, Y. F.;

et al. 2023. Learning the language of antibody hypervariability. *bioRxiv* 2023–04.

Uversky, V. N., and Van Regenmortel, M. H. 2021. Mobility and disorder in antibody and antigen binding sites do not prevent immunochemical recognition. *Critical Reviews in Biochemistry and Molecular Biology* 56(2):149–156.

Van der Maaten, L., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research* 9(11).

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems* 30.

Yuan, Y.; Chen, Q.; Mao, J.; Li, G.; and Pan, X. 2023. Dgaffinity: predicting antigen–antibody affinity with language models from sequences. *BMC bioinformatics* 24(1):430.