On Interpretable Deep Learning and Protein Language Models for Rational Antibody Design

Aurelien Pelissier[1], **Maria Rodriguez Martinez**[2,§]

[1] *Institute of Computational Life Sciences, ZHAW, 8400 Winterthur, Switzerland*
[3] *Yale School of Medicine, 06510 New Haven, United States*
† *Corresponding author.* maria.rodriguezmartinez@yale.edu

Abstract

The ability of antibodies to bind to a wide variety of targets in a highly specific and selective manner has led to an increasing interest in their use as therapeutics for a broad range of diseases. However, computational methods have struggled to accurately predict the impact of mutations in antibody-antigen complexes on binding affinity, which has limited their effectiveness in antibody engineering and development. Nevertheless, recent breakthroughs in protein language models and deep learning, along with increased availability of training data, have enabled precise induction of several epitope-focused immunogens, sparkling hope in the field of rational antibody design for the development of vaccines and therapeutic antibodies. Here, we leverage protein language models, interpretable deep learning models, and the growing repository of antibody structures and binding data to identify and characterize structural and sequence motifs relevant to various antibody properties. This approach can guide the efficient design of antibodies tailored to specific needs. We demonstrate our method in the context of SARS-CoV-2 binding, broadly neutralizing antibodies against HIV strains, and anti-citrullinated protein antibodies (ACPAs) in rheumatoid arthritis.

1 Introduction

Antibodies (Ab) are central components of our immune system that provide protection against foreign pathogens by recognizing their target antigen (Ag) with exquisite specificity and remarkable affinity [1]. Through V(D)J recombination and somatic hypermutation in the bone marrow and lymphatic tissues [2], they have the capacity for binding an extraordinary variety of epitopes as a result of their sequence diversity, which is estimated at 10^{13} unique molecules in the human antibody repertoire [3]. This ability to bind to a wide variety of targets in a highly specific and selective manner has led to an increasing interest in their use as therapeutics for a broad range of diseases including several types of cancer [4] and rheumatoid arthritis [5]. Thanks to the significant improvement in antibody engineering over the past decade, targeted neutralizing antibodies now account for over half of the therapeutic market [6].

Considering the huge diversity of antibodies, in which each small variation potentially affects multiple biological properties, the computational design of immunogens that elicit precise and focused antibody responses remains a major challenge. Nevertheless, recent approaches, leveraging the Rosetta energy function [7, 8, 9], geometric deep learning [10] or protein language models [11], has enabled the precise induction of several epitope-focused immunogens, sparkling hope in the field of rational antibody design for the development of vaccines and therapeutic antibodies.

Protein language models. In recent years, self-supervised protein language models (PLMs) have emerged as a powerful paradigm for a large number of protein-related tasks, including biological and molecular property prediction [12]. At their core, these models treat amino acid sequences as a biological language. This language can be decoded using deep learning models trained on vast numbers of protein sequences (approximately 250 million), enabling the translation of specific sequences into meaningful vector representations of proteins (embeddings) in a high-dimensional latent space. Some of the best-known PLMs include protBERT (Bidirectional Encoder Representations from Transformers) [13], ESM (Evolutionary Scale Modeling) [14], aminoBERT [15] and ProGen [16]. Although neither of them was specifically trained for molecular property or structure prediction, their immense scale (15 billion parameters for the largest) allows them to distill fundamental qualities of the biological language. This is demonstrated by their ability to predict protein 3D-structures [17, 18], binding events [19, 20], and identifying functional sites [21]. Importantly, the latent space representations learned by these models can serve as inputs for subsequent predictive models, which helps reduce training time and model complexity while enhancing performance on downstream tasks [22]. Additionally, this learned latent space is significant for generative AI, where *de-novo* protein sequences with desired binding properties can be directly generated without iterative optimization processes [23, 24], a method that has also been specifically applied to antibodies [25, 26].

We note that, given the success of these PLMs in many bioinformatics-related tasks, similar models have been trained directly on antibodies (AbLang [27], Antibert [28], AbMAP [29]). The rationale behind these models is that they might improve representation capabilities for immune-related applications [30]. Still, both general and antibody specific PLMs encode different type of information and they may both be relevant for any particular downstream task [22].

Interpretability. Many machine learning models, especially PLMs, are inherently non-interpretable due to the encoding of amino acid information in highly abstract and complex latent spaces. However, valuable insights can still be obtained by examining the model's layers and utilizing attention mechanisms [31]. For antibodies, attention mechanisms and interpretability can assist in assessing the impact of mutations, even when the antigen is unknown. As mutating an amino acid predicted to be crucial for the function or structure of an antibody is likely to have a more significant effect on antigen binding than mutating less critical residues, interpretability can directly aid in rational antibody design. This has been shown through protein language models utilizing evolutionary scores for residues [11] and attention scores that identify important amino acids for the antibody's structure [32].

Additionally, interpretability methods applied to deep learning models trained explicitly on antigen-antibody complexes [33, 34] can yield more antigen-specific insights. These techniques can uncover the specific motifs and structural elements that contribute to the binding, providing valuable information for understanding and improving antibody design targeted to specific antigens.

Proposed work. PLMs have already proven to be highly useful for immunology-related tasks such as antigen binding [30, 22, 35, 36]. However, the rationale behind their decisions in the context of antigen binding remains mostly unexplored. Here, we leverage attention mechanisms [37] to interpret the amino acid pairs responsible for the model's predictions in various tasks. Additionally, we incorporate structural information from predicted and, where available, X-ray crystallography structures to further characterize the structural and sequence motifs relevant to the predictions. These learned motifs will enhance our understanding of antibody development and the roles of different antibody subtypes in the body.

More specifically, we will leverage the motifs learned from antibodies with specific properties, such as antigen specificity to SARS-CoV-2, broadly neutralizing antibodies (bnAbs) to HIV [38], or anti-citrullinated protein antibodies (ACPAs) [39], for efficient repertoire mining, i.e. to identify new potential antibodies with these properties in existing repertoires. Then, we will demonstrate that the attention scores from the learned predictive model are highly informative for rational antibody design. Specifically, we show that mutations in amino acid motifs with high attention scores are significantly more likely to result in notable changes in binding affinity. To evaluate the advantages of these attention-based features, we train a model to predict the effects of mutations on antibody binding affinity [40], and compare its performance to other benchmarks that do not utilize these scores.

2 Proposal

2.1 Available training Data

In this work, we will focus on the binding of antibody to a given target antigen. Either as a binary classification task (binder *vs* non-binder), or as a regression task (binding affinity). We leverage different data types:

(i) **General:** A recent study [20] combined several existing database to obtain a large collection of curated antibody-antigen binding affinity values, including both protein sequences and structures. The curated datasets express sufficient generalizability since they contain numerous antigens such as SARS-CoV-2, HIV, MERS, and flu as well as many others. The study provide both sequence only and structure based information (P2PXML-Seq and P2PXML-PDB), which contain 111,845 and 8,475 datapoints, respectively.

(ii) **Binders to Sars-Cov2:** The spike protein of SARS-CoV-2 and its variants are among the most extensively studied antigens in terms of antibody binding. Numerous databases and datasets have been developed since the pandemic to catalog this information, including CovEpiAb [41], CovAbDab [42], Ab-CoV [43], Alpha-seq [44], and S3AI-Cov [34]. These resources provide valuable data on the interactions between antibodies and the spike protein, with affinity typically quantified by metrics such as the dissociation constant (Kd) or the half-maximal inhibitory concentration (IC50).

(iii) **Binders to HIV-1:** Antibody binding to the HIV-1 protein, which is known for its high mutation rate, has also been extensively studied. The CATNAP database contains published IC50/IC80 data for anti-HIV neutralizing antibodies [45]. This data has been used to evaluate the out-of-distribution (OOD) performance in antigen binding predictions [46, 34](`https://github.com/enai4bio/DeepAAI/tree/main`).

(iv) **HIV bnAbs:** The International AIDS Vaccine Initiative (IAVI) has documented the discovery of over 200 bNAbs for HIV, targeting various conserved regions of the HIV-1 envelope glycoprotein [38, 47]. These bnAbs are recognized for their ability to neutralize a wide range of HIV strains by targeting specific epitopes that remain relatively constant despite the virus's high mutation rate. HIResist [48] is a database of HIV-1 Resistance to broadly neutralizing antibodies (bnAbs).

(v) **RA ACPAs:** Anti-citrullinated protein antibodies (ACPAs) are a type of autoantibody associated with Rheumatoid Arthritis (RA). Citrullination is a post-translational modification of proteins, where the amino acid arginine is converted into citrulline. This process can alter protein structure and

function, potentially leading to the immune system recognizing these modified proteins as foreign and generating an autoimmune response. Approximately 100 ACPA clones have now been identified in the context of RA [39]. These ACPAs are notable for their ability to bind a diverse array of citrullinated peptides [49]. This binding poly-specificity plays a crucial role in the pathogenesis of RA, as it drives the autoimmune attack on joint tissues.

(vi) **Nanobodies:** Nanobodies, also known as single-domain antibodies (sdAbs), are a distinctive class of antibodies derived from camelids such as camels, llamas, and alpacas [50]. They consist solely of the variable domain of the heavy chain of heavy-chain antibodies (VHH). They are relevant because they offer several unique advantages over conventional antibodies, including smaller size, high stability, and the ability to bind to epitopes that are often inaccessible to larger antibodies. These properties make nanobodies particularly valuable in therapeutic, diagnostic, and research applications. Extensive binding data exists for these nanobodies, including AVIDa-SARS-CoV-2 [51] for SARS-CoV-2 and AVIDa-hIL6 [52] for IL-6 binding. Additionally, there are specialized resources dedicated to nanobodies, such as the NanoLAS database [53] and the Integrated Database of Nanobodies for Immunoinformatics (INDI [54]). These resources provide comprehensive data and tools for nanobody research and applications.

2.2 Models inputs

In this study, we utilize two types of inputs for our binding prediction task (Figure 1). First, we employ state-of-the-art protein language models, including both general (ESM2-650M [17]) and antibody-specific models (Ablang [27]). We encode our sequences at the residue level to maintain the context-specific role of each amino acid. We note that each PLM has a different dimensionality, and concatenating different PLMs together is also a viable option.

Then, we leverage the structural features of the antibody. Specifically, we consider dihedral and planar angles as structural features, which have been proposed in previous studies [55, 56, 32]. For two residues i and j, the relative orientation is defined by six parameters (d , ω, θ_{ij}, θ_{ji}, ϕ_{ij} and ϕ_{ji}. We also incorporate the normal mode correlation map [57], which considers the structure as an elastic network rather than relying solely on rigid geometric potentials. This method captures both structural features and the energetic patterns of local and global residue fluctuations, an approach that has been shown to enhance predictive performance compared to traditional contact maps [33].

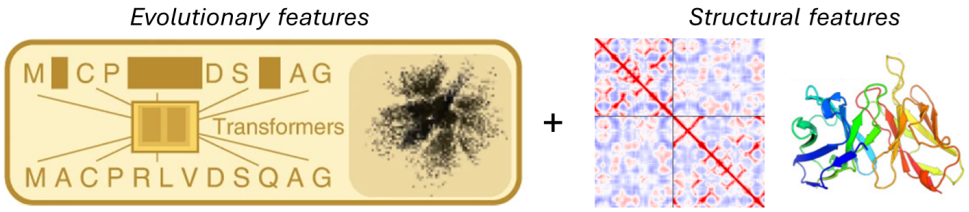

Figure 1: Protein language models (i.e. evolutionary features) and structural features as input for antibody binding predictions.

Antibody structures. Approximately 1000 antibody structures have been characterized via X-ray crystallography [58], including ACPAs, bnAbs to HIV, and Sars-CoV-2 antibodies [42]. As this may not be enough to train our model, we will infer the structures for other antibodies using Igfold [59], tFold-Ab [60], or DeepAB [32]. These methods provide a confidence score for the predicted structure, so we will only retain structures with high-confidence predictions to train our model.

2.3 Models

We consider two models: one utilizing only sequence data with protein language models (i.e., evolutionary features), and another that incorporates both structural and evolutionary features. We will train

and compare their performance to evaluate the advantage of including structural features, even if they are inferred from the sequence.

First, we train a model based solely on sequences and PLM embeddings (Figure 2A). The prediction model uses concatenated heavy and light chain sequences as input. These sequences are encoded at the residue level using a protein language model. The heavy and light chains are numbered according to the Chothia alignment, with missing residues (gaps) encoded as zeros in the PLM embeddings. This process results in an embedding with dimensions (L x N), where L = 292 represents the length of the antibody sequence and N represents the length of the embeddings (2048 for ESM2-650M and 1048 for Ablang). Subsequently, the embeddings are transformed into a tensor (L x L x 2N) by concatenating the embeddings of each pair of residues.

The 2D embedding is then passed through a 2D ResNet. The 2D ResNet starts with a 2D convolution that reduces the dimensionality to L × L × 64, followed by 25 2D ResNet blocks. Each block consists of two 2D convolutions with a kernel size of 5 × 5, maintaining the same dimensionality. These blocks cycle through convolution dilation values of 1, 2, 4, 8, and 16, repeating this cycle five times. Following the 2D ResNet, the output branch includes a 2D convolution that reduces the dimensionality to L × L × 37, followed by a recurrent criss-cross attention (RCCA) module [37]. The RCCA module uses two criss-cross attention operations with shared weights, enabling each residue pair to gather information across the entire spatial dimension. Attention queries and keys are projected to a dimension of L × L × 1 (one attention head) for the subsequent task.

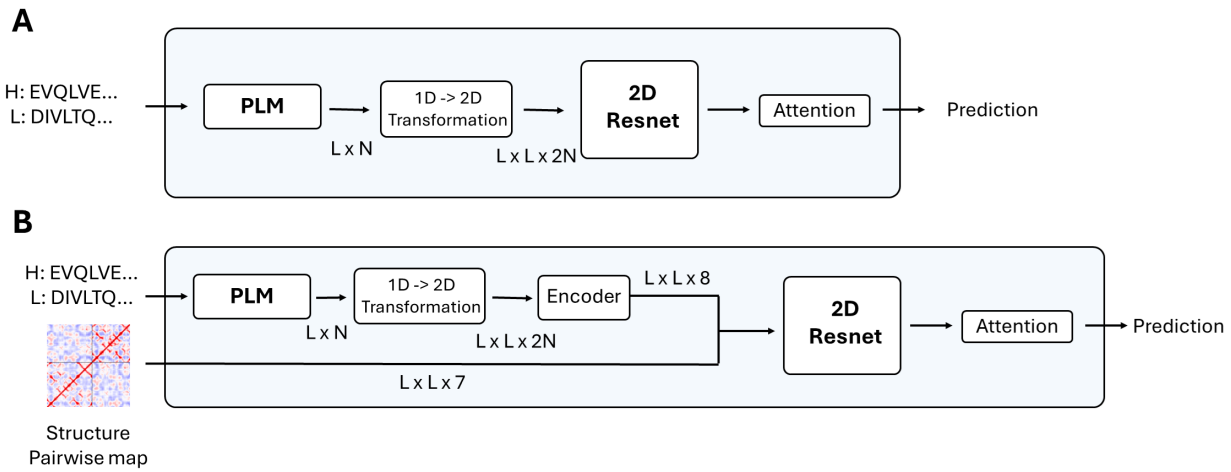

Figure 2: Interpretable Deep learning model architecture for antigen binding prediction. (A) Model using protein language only as input (i.e. evolutionary features). (B) Model incorporating together both structural and evolutionary features.

We then then consider another model that also take structural information as input (Figure 2B). To combine both PLM embeddings and structural features in the predictions, we concatenate the L × L × 7 structural feature matrix with the 2D PLM embeddings. However, since the PLM embeddings have a dimensionality much larger than the structural features, we first reduce the PLM size to L × L × 8 with an encoder before concatenation. This approach ensures a better balance between sequence and structural features.

Model training. We follow a similar approach for the binary classification task (binder vs non-binder) and the affinity regression task (Kd or IC50). We trained five models on random 90/10% training/validation splits and averaged the model logits to make predictions. Models are trained using focal loss and the Adam optimizer with a learning rate of 0.01, reducing the learning rate upon plateauing of the validation loss.

2.4 Application 1: Repertoire mining

In the past few years, paired BCRs from patient repertoires under various immune conditions have been sequenced. However, not all BCRs in these datasets are antigen-specific, meaning they are not necessarily relevant to the immune condition, and many different epitopes may be involved. By efficiently identifying the antibodies relevant to different immune conditions, we aim to improve our understanding of their formation and population evolution during various stages of disease progression or vaccination.

Publicly available repertoires include SARS-CoV-2 mRNA vaccination repertoires [61], 22 HIV-infected patients [62], 76 Rheumatoid Arthritis (RA) patients, 10 osteoarthritis patients, and 78 healthy volunteers [63] as well as 13 ACPA-positive RA patients from another study with 6 healthy controls [64]. We will leverage these repertoires to identify new broadly neutralizing antibodies (bnAbs) for HIV and anti-citrullinated protein antibodies (ACPAs) for Rheumatoid arthritis, which could not be identified through conventional methods [65, 66, 67]. These sequences will be validated experimentally, leading to a better understanding of how these antibodies develop and the roles of different antibody subtypes in the body.

2.5 Application 2: Generating *de-novo* antibodies with desired properties

Instead of mining repertoires to identify new binders, we can aim to generate an antibody directly with the desired properties through generative AI [25, 26]. To achieve this, we can utilize the latent space learned by our model and employ it to design antibody sequences that meet our criteria. This approach leverages the model's understanding of antibody structures and functions to predict and generate sequences that are likely to exhibit the desired characteristics. To generate a sequence directly from a chosen point in the latent space, we first train a decoder from our learned latent space to efficiently reconstruct the sequence from the latent space. Then, we identify a region in the latent space that corresponds to high specificity characteristics. This identification is based on training data where the model has learned to map specific properties to certain regions in the latent space. Once the region is identified, we can sample points within this area. Each sampled point represents a potential antibody sequence in the encoded form. We then use the decoder part of our model to translate these points back into antibody sequences. The decoder reconstructs the sequence by interpreting the latent representation, ensuring that the generated sequence retains the high specificity property encoded in the latent space. This process will allow us to systematically generate and evaluate numerous antibody candidates with the desired traits, streamlining the discovery and optimization phases.

2.6 Application 3: Mutation impact prediction

Reliably predicting the impact of mutations on binding affinity remains a significant challenge due to limited data and overfitting issues [68]. We aim to evaluate if the attention scores inferred from our binding model models can be relevant in mutation impact prediction, and if it can outperform current benchmarks [40, 69]. If this proves too difficult, we might consider predicting whether a mutation increases or decreases affinity, a simpler binary task [70].

Database for mutation impact on affinity. The AB-bind [71] and SKEMPI2.0 [72] databases are a collection of mutations over different antibody–antigen structures that include experimentally-determined changes in binding free energy associated with each mutant. After data cleaning and removal of redundant measurements between the two databases, a total of 905 single mutations over 60 AbAg complexes are obtained [40]. Then, there are several studies that provide additional measurements not included in these databases [73, 74, 75, 11].

Additionally, some studies focus on mutating the antigen rather than the antibody. For instance, extensive research on Covid-19 examines mutations in the Receptor Binding Domain (RBD) of the spike protein instead of antibodies [76, 77, 78]. Combined together, these studies includes affinity data on more than 5 thousands RBD mutations.

3 Follow-up work

3.1 Leveraging the Antigen-Antibody complexes for OOD predictions

Upon successfully identifying binding motifs related to antibody-antigen (Ab-Ag) interactions, we will evaluate our model's generalization capability for predicting interactions with other antigens, known as out-of-distribution (OOD) predictions [34]. To accomplish this, antigen information must be integrated into our predictive model (Figure 3). While directly concatenating PLM and structure embeddings from the antigen and antibody is relatively straightforward, it carries the risk of the model treating the antigen merely as a class label, leading to poor generalization on out-of-distribution (OOD) antigens, as it was observed in the context of TCR binding predictions [79]. Therefore, a more explicit modeling of Ab-Ag interactions is preferable. Although previous studies have demonstrated the relevance of Ab-Ag complexes structures in reliably predicting binding affinity [33], the availability of X-ray crystallography data for these complexes is limited and dispersed across various antigens [58, 80, 81]. Tools do exist to infer Ab-Ag complex structures directly from sequences (tFold-AbAg [60], Alphafold-multimer [82, 83]), but accurately identifying the correct epitope on the antigen remains challenging, and the predicted Ab-Ag complex structures may not always be of sufficient quality to benefit our model.

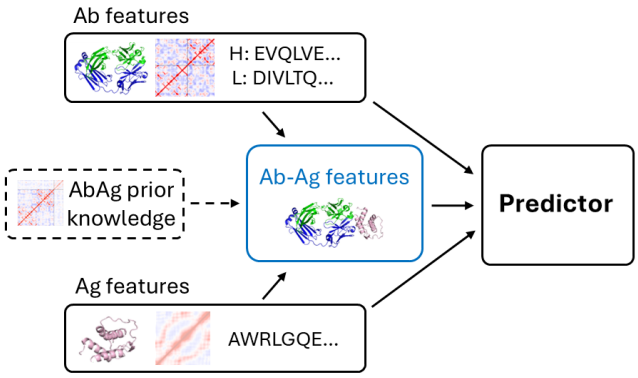

Figure 3: General Ab-Ag affinity model predictor. In addition to the information of the antigen and the antibody, the model considers the Ab-Ag interaction module, encoding how the Ab and Ag interact. This can be encoded in different ways, as shown in the studies [20, 34, 35]

In this work, inspired by previous studies [20, 34, 35], we will explore various approaches to model Ab-Ag interactions within our framework. Specifically, we will evaluate whether explicitly modeling Ab-Ag complexes using existing tools enhances performance compared to directly encoding Ab-Ag interactions within our model.

3.2 Characterizing TCR binding

T cell receptors (TCRs) share many similarities with antibodies, and the methods discussed in this proposal can be easily extended to TCRs. T cells typically bind to peptides, which are small pieces of antigen presented by the MHC complex, involving linear epitopes [84]. This is simpler than the conformal non-linear epitopes typically encountered in B cell binding. On the other hand, when considering TCR binding, we must account for the trio of MHC, peptide, and TCR (pMHC-complex). The TCR must recognize both the MHC molecule and the specific peptide it presents, requiring a precise fit between all three components. This adds an additional layer of complexity to the binding interaction compared to the antibody-antigen complex, as the peptide must be correctly presented by the MHC to be recognized by the TCR.

Tools and datasets for T cells: TCR structures can be inferred with good precision using TCR-model2 [85], and the direct prediction of TCR-pMHC complexes is also feasible [86]. Several TCR structure databases are available, including sTCRDab [87] and TCR3D [88]. For TCR specificity data,

various databases are available, such as VDJdb [89], which contains over 70,000 TCR sequences and approximately 1,100 different epitopes, and McPas-TCR [90], which offers a manually curated set of about 40,000 pairs. Additionally, newer datasets are rapidly becoming available, such as the MIRA dataset published during the COVID-19 pandemic, which includes over 135,000 TCRs binding to various COVID-19 epitopes [91]. Finally, affinity data linked with structure can be found in the TCRatlas database [92].

Tasks for TCR binding models: Similar to autoantibodies, we can characterize the sequence and structural motifs associated with self-reactive TCRs [93, 94, 95] involved in autoimmune diseases. Self-reactive TCRs can be found in the IEDB [96] database, with many TCRs associated with Type 1 Diabetes (2000 TCRs), Celiac Disease (800 TCRs), and Rheumatoid Arthritis (150 TCRs).

Additionally, we can characterize the binding motif associated with polyreactive TCRs. While most TCRs exhibit specificity for a particular peptide-MHC complex, some TCRs have the ability to recognize and bind multiple distinct antigens or peptide-MHC complexes, hence the term polyreactive [97]. These typically recognize multiple viral epitopes [98, 99], or recognize self-antigens and contribute to the breakdown of self-tolerance [100]. They are characterized by highly flexible and diverse CDR3 regions [101], which allow them to adapt to various peptide-MHC conformations, a specificity that could be captured by our trained model. Interestingly, TCRs found across multiple individuals and often show a degree of polyreactivity, recognizing common microbial peptides and sometimes self-antigens [99].

Bibliography

[1] Ruei-Min Lu et al. “Development of therapeutic antibodies for the treatment of diseases”. In: *Journal of biomedical science* 27.1 (2020), pp. 1–30.

[2] Aleksandr Kovaltsuk et al. “How B-cell receptor repertoire sequencing can be enriched with structural antibody data”. In: *Frontiers in immunology* 8 (2017), p. 1753.

[3] Victor Greiff et al. “Bioinformatic and statistical analysis of adaptive immune repertoires”. In: *Trends in immunology* 36.11 (2015), pp. 738–749.

[4] Zehra Elgundi et al. “The state-of-play and future of antibody therapeutics”. In: *Advanced drug delivery reviews* 122 (2017), pp. 2–19.

[5] Toshio Tanaka, Yoshihiro Hishitani, and Atsushi Ogata. “Monoclonal antibodies in rheumatoid arthritis: comparative effectiveness of tocilizumab with tumor necrosis factor inhibitors”. In: *Biologics: targets & therapy* 8 (2014), p. 141.

[6] Lisa Urquhart. *Market watch: top drugs and companies by sales in 2017*. 2018.

[7] Jared Adolf-Bryfogle et al. “RosettaAntibodyDesign (RAbD): A general framework for computational antibody design”. In: *PLoS computational biology* 14.4 (2018), e1006112.

[8] Fabian Sesterhenn et al. “De novo protein design enables the precise induction of RSV-neutralizing antibodies”. In: *Science* 368.6492 (2020).

[9] Nathaniel R Bennett et al. “Atomically accurate de novo design of single-domain antibodies”. In: *bioRxiv* (2024).

[10] Sisi Shan et al. “Deep learning guided optimization of human antibody against SARS-CoV-2 variants with broad neutralization”. In: *Proceedings of the National Academy of Sciences* 119.11 (2022), e2122954119.

[11] Brian L Hie et al. “Efficient evolution of human antibodies from general protein language models”. In: *Nature Biotechnology* 42.2 (2024), pp. 275–283.

[12] Tristan Bepler and Bonnie Berger. “Learning the protein language: Evolution, structure, and function”. In: *Cell systems* 12.6 (2021), pp. 654–669.

[13] Nadav Brandes et al. “ProteinBERT: a universal deep-learning model of protein sequence and function”. In: *Bioinformatics* 38.8 (2022), pp. 2102–2110.

[14] Alexander Rives et al. “Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences”. In: *Proceedings of the National Academy of Sciences* 118.15 (2021), e2016239118.

[15] Ratul Chowdhury et al. “Single-sequence protein structure prediction using a language model and deep learning”. In: *Nature Biotechnology* 40.11 (2022), pp. 1617–1623.

[16] Ali Madani et al. “Large language models generate functional protein sequences across diverse families”. In: *Nature Biotechnology* (2023), pp. 1–8.

[17] Zeming Lin et al. “Evolutionary-scale prediction of atomic-level protein structure with a language model”. In: *Science* 379.6637 (2023), pp. 1123–1130.

[18] Tristan Bepler and Bonnie Berger. “Learning protein sequence embeddings using information from structure”. In: *International Conference on Learning Representations*. 2019.

[19] Zilong Hou et al. “Learning the protein language of proteome-wide protein-protein binding sites via explainable ensemble deep learning”. In: *Communications Biology* 6.1 (2023), p. 73.

[20] Nuwan Sriyantha Bandara et al. “Deep Geometric Framework to Predict Antibody-Antigen Binding Affinity”. In: *bioRxiv* (2024), pp. 2024–06.

[21] Wayland Yeung et al. “Alignment-free estimation of sequence conservation for identifying functional sites using protein sequence embeddings”. In: *Briefings in Bioinformatics* 24.1 (2023), bbac599.

[22] Nicolas Deutschmann et al. “Do domain-specific protein language models outperform general models on immunology-related tasks?” In: *ImmunoInformatics* 14 (2024), p. 100036.

[23] Matt Sternke and Joel Karpiak. “ProteinRL: Reinforcement learning with generative protein language models for property-directed sequence design”. In: *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*. 2023.

[24] Viacheslav Meshchaninov et al. “Diffusion on language model embeddings for protein sequence generation”. In: *arXiv preprint arXiv:2403.03726* (2024).

[25] Amir Shanehsazzadeh et al. “Unlocking de novo antibody design with generative artificial intelligence”. In: *bioRxiv* (2023), pp. 2023–01.

[26] Richard W Shuai, Jeffrey A Ruffolo, and Jeffrey J Gray. “Generative language modeling for antibody design”. In: *BioRxiv* (2021), pp. 2021–12.

[27] Tobias H Olsen, Iain H Moal, and Charlotte M Deane. “AbLang: an antibody language model for completing antibody sequences”. In: *Bioinformatics Advances* 2.1 (2022), vbac046.

[28] Jinwoo Leem et al. “Deciphering the language of antibodies using self-supervised learning”. In: *Patterns* 3.7 (2022).

[29] Rohit Singh et al. “Learning the Language of Antibody Hypervariability”. In: *bioRxiv* (2023), pp. 2023–04.

[30] Danqing Wang, YE Fei, and Hao Zhou. “On pre-training language model for antibody”. In: *The Eleventh International Conference on Learning Representations*. 2023.

[31] Jesse Vig et al. “Bertology meets biology: Interpreting attention in protein language models”. In: *arXiv preprint arXiv:2006.15222* (2020).

[32] Jeffrey A Ruffolo, Jeremias Sulam, and Jeffrey J Gray. “Antibody structure prediction using interpretable deep learning”. In: *Patterns* 3.2 (2022).

[33] Kevin Michalewicz, Mauricio Barahona, and Barbara Bravi. “ANTIPASTI: interpretable prediction of antibody binding affinity exploiting Normal Modes and Deep Learning”. In: *bioRxiv* (2023), pp. 2023–12.

[34] Yutian Liu et al. “Interpretable antibody-antigen interaction prediction by introducing route and priors guidance”. In: *bioRxiv* (2024), pp. 2024–03.

[35] Ye Yuan et al. “DG-Affinity: predicting antigen–antibody affinity with language models from sequences”. In: *BMC bioinformatics* 24.1 (2023), p. 430.

[36] Meng Wang et al. “Language model-based B cell receptor sequence embeddings can effectively encode receptor specificity”. In: *Nucleic acids research* 52.2 (2024), pp. 548–557.

[37] Zilong Huang et al. “Ccnet: Criss-cross attention for semantic segmentation”. In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2019, pp. 603–612.

[38] Elise Landais and Penny L Moore. “Development of broadly neutralizing antibodies in HIV-1 infected elite neutralizers”. In: *Retrovirology* 15 (2018), pp. 1–14.

[39] Bruno Raposo et al. “The peculiar features, diversity and impact of citrulline-reactive autoantibodies”. In: *Nature Reviews Rheumatology* (2024), pp. 1–18.

[40] Yoochan Myung et al. “mCSM-AB2: guiding rational antibody design using graph-based signatures”. In: *Bioinformatics* 36.5 (2020), pp. 1453–1459.

[41] Xue Zhang et al. “CovEpiAb: a comprehensive database and analysis resource for immune epitopes and antibodies of human coronaviruses”. In: *Briefings in Bioinformatics* 25.3 (2024), bbae183.

[42] Matthew IJ Raybould et al. “CoV-AbDab: the coronavirus antibody database”. In: *Bioinformatics* 37.5 (2021), pp. 734–735.

[43] Puneet Rawat et al. “Ab-CoV: a curated database for binding affinity and neutralization profiles of coronavirus-related antibodies”. In: *Bioinformatics* 38.16 (2022), pp. 4051–4052.

[44] Emily Engelhart et al. “A dataset comprised of binding interactions for 104,972 antibodies against a SARS-CoV-2 peptide”. In: *Scientific Data* 9.1 (2022), p. 653.

[45] Hyejin Yoon et al. “CATNAP: a tool to compile, analyze and tally neutralizing antibody panels”. In: *Nucleic acids research* 43.W1 (2015), W213–W219.

[46] Jie Zhang et al. “Predicting unseen antibodies’ neutralizability via adaptive graph neural networks”. In: *Nature Machine Intelligence* 4.11 (2022), pp. 964–976.

[47] Devin Sok and Dennis R Burton. “Recent progress in broadly neutralizing antibodies to HIV”. In: *Nature immunology* 19.11 (2018), pp. 1179–1188.

[48] Milind Misra et al. “HIResist: a database of HIV-1 resistance to broadly neutralizing antibodies”. In: *Bioinformatics* 40.3 (2024), btae103.

[49] Peter Sahlström et al. “Different hierarchies of anti–modified protein autoantibody reactivities in rheumatoid arthritis”. In: *Arthritis & Rheumatology* 72.10 (2020), pp. 1643–1657.

[50] Gemma L Gordon et al. “A comparison of the binding sites of antibodies and single-domain antibodies”. In: *Frontiers in Immunology* 14 (2023), p. 1231623.

[51] Hirofumi Tsuruta et al. “A SARS-CoV-2 Interaction Dataset and VHH Sequence Corpus for Antibody Language Models”. In: *arXiv preprint arXiv:2405.18749* (2024).

[52] Hirofumi Tsuruta et al. “AVIDa-hIL6: a large-scale VHH dataset produced from an immunized alpaca for predicting antigen-antibody interactions”. In: *Advances in Neural Information Processing Systems* 36 (2024).

[53] Shuchang Xiong et al. “NanoLAS: a comprehensive nanobody database with data integration, consolidation and application”. In: *Database* 2024 (2024), baae003.

[54] Piotr Deszyński et al. “INDI—integrated nanobody database for immunoinformatics”. In: *Nucleic Acids Research* 50.D1 (2022), pp. D1273–D1281.

[55] Jianyi Yang et al. “Improved protein structure prediction using predicted interresidue orientations”. In: *Proceedings of the National Academy of Sciences* 117.3 (2020), pp. 1496–1503.

[56] Jeffrey A Ruffolo et al. “Geometric potentials from deep learning improve prediction of CDR H3 loop structures”. In: *Bioinformatics* 36.Supplement_1 (2020), pp. i268–i275.

[57] Eric C Dykeman and Otto F Sankey. “Normal mode analysis and applications in biological physics”. In: *Journal of Physics: Condensed Matter* 22.42 (2010), p. 423202.

[58] James Dunbar et al. “SAbDab: the structural antibody database”. In: *Nucleic acids research* 42.D1 (2014), pp. D1140–D1146.

[59] Jeffrey A Ruffolo et al. “Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies”. In: *Nature communications* 14.1 (2023), p. 2389.

[60] Fandi Wu et al. “Fast and accurate modeling and design of antibody-antigen complex using tFold”. In: *bioRxiv* (2024), pp. 2024–02.

[61] Wooseob Kim et al. “Germinal centre-driven maturation of B cell response to mRNA vaccination”. In: *Nature* 604.7904 (2022), pp. 141–145.

[62] Deniz Cizmeci et al. “Distinct clonal evolution of B-cells in HIV controllers with neutralizing antibody breadth”. In: *Elife* 10 (2021), e62648.

[63] Fanlei Hu et al. “Impaired CD27+ IgD+ B cells with altered gene signature in rheumatoid arthritis”. In: *Frontiers in immunology* 9 (2018), p. 626.

[64] Yan Wang et al. “Rheumatoid arthritis patients display B-cell dysregulation already in the naïve repertoire consistent with defects in B-cell tolerance”. In: *Scientific reports* 9.1 (2019), p. 19995.

[65] Aurelien Pelissier et al. “Convergent evolution and B-cell recirculation in germinal centers in a human lymph node”. In: *Life Science Alliance* 6.11 (2023).

[66] Wing Ki Wong et al. “Ab-Ligity: identifying sequence-dissimilar antibodies that bind to the same epitope”. In: *MAbs*. Vol. 13. 1. Taylor & Francis. 2021, p. 1873478.

[67] Eve Richardson et al. “A computational method for immune repertoire mining that identifies novel binders from different clonotypes, demonstrated by identifying anti-pertussis toxoid antibodies”. In: *MAbs*. Vol. 13. 1. Taylor & Francis. 2021, p. 1869406.

[68] Alissa M Hummer et al. “Investigating the volume and diversity of data needed for generalizable antibody-antigen ddG prediction”. In: *bioRxiv* (2023), pp. 2023–05.

[69] Ruofan Jin et al. “AttABseq: an attention-based deep learning prediction method for antigen–antibody binding affinity changes based on protein sequences”. In: *Briefings in Bioinformatics* 25.4 (2024).

[70] Thomas Clark et al. “Enhancing antibody affinity through experimental sampling of non-deleterious CDR mutations predicted by machine learning”. In: *Communications Chemistry* 6.1 (2023), p. 244.

[71] Sarah Sirin et al. “AB-Bind: antibody binding mutational database for computational affinity predictions”. In: *Protein Science* 25.2 (2016), pp. 393–409.

[72] Justina Jankauskaitė et al. “SKEMPI 2.0: an updated benchmark of changes in protein–protein binding energy, kinetics and thermodynamics upon mutation”. In: *Bioinformatics* 35.3 (2019), pp. 462–469.

[73] Kyle A Barlow et al. “Flex ddG: Rosetta ensemble-based estimation of changes in protein–protein binding affinity upon mutation”. In: *The Journal of Physical Chemistry B* 122.21 (2018), pp. 5389–5399.

[74] Alexander Benedix et al. “Predicting free energy changes using structural ensembles”. In: *Nature methods* 6.1 (2009), pp. 3–4.

[75] Shira Warszawski et al. “Optimizing antibody affinity and stability by the automated design of the variable light-heavy chain interfaces”. In: *PLoS computational biology* 15.8 (2019), e1007207.

[76] Allison J Greaney et al. “Complete mapping of mutations to the SARS-CoV-2 spike receptor-binding domain that escape antibody recognition”. In: *Cell host & microbe* 29.1 (2021), pp. 44–57.

[77] Tyler N Starr et al. “Deep mutational scanning of SARS-CoV-2 receptor binding domain reveals constraints on folding and ACE2 binding”. In: *cell* 182.5 (2020), pp. 1295–1310.

[78] Alief Moulana et al. “The landscape of antibody binding affinity in SARS-CoV-2 Omicron BA. 1 evolution”. In: *Elife* 12 (2023), e83442.

[79] Anna Weber, Jannis Born, and Maria Rodriguez Martinez. “TITAN: T-cell receptor specificity prediction with bimodal attention networks”. In: *Bioinformatics* 37.Supplement_1 (2021), pp. i237–i244.

[80] Nan Zhao et al. “ABAG-docking benchmark: a non-redundant structure benchmark dataset for antibody–antigen computational docking”. In: *Briefings in Bioinformatics* 25.2 (2024), bbae048.

[81] Andreas V Madsen et al. “Structural trends in antibody-antigen binding interfaces: a computational analysis of 1833 experimentally determined 3D structures”. In: *Computational and Structural Biotechnology Journal* 23 (2024), pp. 199–211.

[82] Richard Evans et al. “Protein complex prediction with AlphaFold-Multimer”. In: *biorxiv* (2021), pp. 2021–10.

[83] Rui Yin and Brian G Pierce. “Evaluation of AlphaFold antibody–antigen modeling with implications for improving predictive accuracy”. In: *Protein Science* 33.1 (2024), e4865.

[84] Jose L Sanchez-Trincado, Marta Gomez-Perosanz, and Pedro A Reche. “Fundamentals and methods for T-and B-cell epitope prediction”. In: *Journal of immunology research* 2017.1 (2017), p. 2680160.

[85] Rui Yin et al. “TCRmodel2: high-resolution modeling of T cell receptor recognition using deep learning”. In: *Nucleic Acids Research* 51.W1 (2023), W569–W576.

[86] Philip Bradley. “Structure-based prediction of T cell receptor: peptide-MHC interactions”. In: *Elife* 12 (2023), e82813.

[87] Jinwoo Leem et al. “STCRDab: the structural T-cell receptor database”. In: *Nucleic acids research* 46.D1 (2018), pp. D406–D412.

[88] Ragul Gowthaman and Brian G Pierce. “TCR3d: The T cell receptor structural repertoire database”. In: *Bioinformatics* 35.24 (2019), pp. 5323–5325.

[89] Mikhail Goncharov et al. “VDJdb in the pandemic era: a compendium of T cell receptors specific for SARS-CoV-2”. In: *Nature methods* 19.9 (2022), pp. 1017–1019.

[90] Nili Tickotsky et al. “McPAS-TCR: a manually curated catalogue of pathology-associated T cell receptor sequences”. In: *Bioinformatics* 33.18 (2017), pp. 2924–2929.

[91] Sean Nolan et al. “A large-scale database of T-cell receptor beta (TCRβ) sequences and binding associations from natural and synthetic exposure to SARS-CoV-2”. In: *Research square* (2020).

[92] Tyler Borrman et al. “ATLAS: A database linking binding affinities with structures for wild-type and mutant TCR-pMHC complexes”. In: *Proteins: Structure, Function, and Bioinformatics* 85.5 (2017), pp. 908–916.

[93] Johannes Textor et al. “Machine learning analysis of the T cell receptor repertoire identifies sequence features of self-reactivity”. In: *Cell Systems* 14.12 (2023), pp. 1059–1073.

[94] Darina Paprckova and Ondrej Stepanek. “Narcissistic T cells: reactivity to self makes a difference”. In: *The FEBS Journal* 288.6 (2021), pp. 1778–1788.

[95] Nelli Heikkilä et al. “Generation of self-reactive, shared T-cell receptor α chains in the human thymus”. In: *Journal of Autoimmunity* 119 (2021), p. 102616.

[96] Randi Vita et al. “The immune epitope database (IEDB): 2018 update”. In: *Nucleic acids research* 47.D1 (2019), pp. D339–D343.

[97] Don Mason. “A very high level of crossreactivity is an essential feature of the T-cell receptor”. In: *Immunology today* 19.9 (1998), pp. 395–404.

[98] Linda Wooldridge et al. “A single autoimmune T cell receptor recognizes more than a million different peptides”. In: *Journal of Biological Chemistry* 287.2 (2012), pp. 1168–1177.

[99] Pradyot Dash et al. “Paired analysis of TCRα and TCRβ chains at the single-cell level in mice”. In: *The Journal of clinical investigation* 121.1 (2011), pp. 288–295.

[100] Jarrett J Adams et al. “T cell receptor signaling is limited by docking geometry to peptide-major histocompatibility complex”. In: *Immunity* 35.5 (2011), pp. 681–693.

[101] Michael E Birnbaum et al. “Deconstructing the peptide-MHC specificity of T cell recognition”. In: *Cell* 157.5 (2014), pp. 1073–1087.

[102] Cunliang Geng et al. “Finding the $\Delta\Delta$G spot: Are predictors of binding affinity changes upon mutations in protein–protein interactions ready for it?” In: *Wiley Interdisciplinary Reviews: Computational Molecular Science* 9.5 (2019), e1410.

[103] Priyashree Chaudhary, Athi N Naganathan, and M Michael Gromiha. “Folding RaCe: a robust method for predicting changes in protein folding rates upon point mutations”. In: *Bioinformatics* 31.13 (2015), pp. 2091–2097.

[104] Shuichi Kawashima et al. “AAindex: amino acid index database, progress report 2008”. In: *Nucleic acids research* 36.suppl_1 (2007), pp. D202–D205.

[105] Srivamshi Pittala and Chris Bailey-Kellogg. “Mixture of Experts for Predicting Antibody-Antigen Binding Affinity from Antigen Sequence”. In: *bioRxiv* (2019), p. 511360.

[106] Sherlyn Jemimah, Masakazu Sekijima, and M Michael Gromiha. “ProAffiMuSeq: sequence-based method to predict the binding free energy change of protein–protein complexes upon mutation using functional classification”. In: *Bioinformatics* 36.6 (2020), pp. 1725–1730.

[107] Derek M Mason et al. “Deep learning enables therapeutic antibody optimization in mammalian cells by deciphering high-dimensional protein sequence space”. In: *bioRxiv* (2019), p. 617860.

[108] Mostafa Karimi et al. “DeepAffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks”. In: *Bioinformatics* 35.18 (2019), pp. 3329–3338.

[109] Rebecca F Alford et al. “The Rosetta all-atom energy function for macromolecular modeling and design”. In: *Journal of chemical theory and computation* 13.6 (2017), pp. 3031–3048.

[110] Sidhartha Chaudhury, Sergey Lyskov, and Jeffrey J Gray. “PyRosetta: a script-based interface for implementing molecular modeling algorithms using Rosetta”. In: *Bioinformatics* 26.5 (2010), pp. 689–691.

[111] Michel E Goldberg and Lisa Djavadi-Ohaniance. “Methods for measurement of antibody/antigen affinity based on ELISA and RIA”. In: *Current opinion in immunology* 5.2 (1993), pp. 278–281.

[112] Hung-Pin Peng et al. “Origins of specificity and affinity in antibody–protein interactions”. In: *Proceedings of the National Academy of Sciences* 111.26 (2014), E2656–E2665.

[113] Kouhei Tsumoto and Jose MM Caaveiro. “Antigen–Antibody Binding”. In: *eLS* (2001), pp. 1–8.