DLAB - Deep learning methods for structure-based virtual screening of antibodies

Constantin Schneider[1*], Andrew Buchanan[2], Bruck Taddese[3,i], Charlotte M. Deane[1],

1 Department of Statistics, University of Oxford, Oxford, UK
2 Antibody Discovery & Protein Engineering, R&D, AstraZeneca, Cambridge, UK
3 Discovery Sciences, R&D, AstraZeneca Cambridge, UK

* constantin.schneider@stats.ox.ac.uk

Abstract

Antibodies are one of the most important classes of pharmaceuticals, with over 80 approved molecules currently in use against a wide variety of diseases. The drug discovery process for antibody therapeutic candidates however is time- and cost-intensive and heavily reliant on *in-vivo* and *in-vitro* high throughput screens. Here, we introduce a framework for structure-based deep learning for antibodies (DLAB) which can virtually screen putative binding antibodies against antigen targets of interest. DLAB is built to be able to predict antibody-antigen binding for antigens with no known antibody binders.

We demonstrate that DLAB can be used both to improve antibody-antigen docking and structure-based virtual screening of antibody drug candidates. DLAB enables improved pose ranking for antibody docking experiments as well as selection of antibody-antigen pairings for which accurate poses are generated and correctly ranked. DLAB also outperforms baseline methods at identifying binding antibodies against specific antigens in a series of case studies. Our results demonstrate the promise of deep learning methods for structure-based virtual screening of antibodies.

Introduction

Antibodies are the most successful class of biotherapeutics, with 85 monoclonal antibody drugs on the US market at the time of writing. Global sales in monoclonal antibody therapeutics reached an estimated $98 billion in 2018 [12], making them one of the fastest growing and largest segments of the pharmaceutical industry. The potential and use of antibodies as therapeutics for a wide range of diseases is due to the high specificity and affinity of their binding, facilitated through the variability in their complementarity determining regions (CDRs) [21].

[i] Bruck Taddese has left AstraZeneca since the completion of the work detailed in this manuscript

The development process for novel antibody therapeutics already benefits from computational tools predicting properties of the antibody and the antibody-antigen complex [27]. Here, we demonstrate that machine learning approaches using structural antibody data can enable large scale computational screening in the antibody development pipeline.

In order to develop a successful therapeutic antibody, several features have to be optimised. Antibodies need to bind with high efficacy, specificity and affinity to the target of interest, while at the same time avoiding an immune reaction in the patient [31] as well as avoiding properties that lead to poor developability, such as self-association, viscosity or immunogenicity [28]. To achieve these goals, large-scale experimental screens are usually employed in the pre-clinical stages of antibody-drug development [3, 31].

Typically, initial leads for therapeutic human antibodies are generated using either in vitro display platforms or in vivo transgenic animals. Further improvement of these initial hits is then achieved through affinity maturation, either via generation and screening of further, hit-based mutagenesis libraries or via rational engineering. High-affinity antibodies generated in this way can be further engineered to achieve desirable properties for antibody therapeutics, for example changes to the constant region to modulate the effector functions as well as *in-vivo* half-life and the improvement of developability [7].

These experimental methods are often effective at generating high-affinity antibodies for downstream development, but are cost- and time-intensive. Furthermore, they do not as standard generate insights into the binding mode of the generated antibodies, or if they are binding to the target epitope on the antigen.

The development pipeline detailed above can be supplemented using in-silico methods, particularly once initial binding candidates have been identified. Computational tools have been used to rationally engineer the binding site in several different ways. As antibody binding affinity is defined by the three-dimensional structure of the antigen-binding region on the antibody (the paratope, which is mainly comprised of the CDRs), antibody modelling tools can aid the design process by rapidly generating models of antibodies (e.g. ABodyBuilder [17], Rosetta [33] and Kotai Antibody Builder [36]). Some methods combine antibody model generation with docking against target antigens and engineering of the antibody (e.g. the Rosetta antibody design tool [1]).

The antibody models generated by those methods, though commonly high quality, are often not accurate across the CDRs, particularly the CDR H3 [4, 17]. This presents a challenge for rigid-body docking methods, which do not allow flexibility of the binding partners, to be able to recapitulate the true antibody-antigen interface. More accurate interface models can be generated using computationally intensive docking approaches [33], however those are too slow to be used on high-throughout screens, where thousands of interface models need to be generated [27].

The caveats described above (CDR model inaccuracy and the speed/accuracy trade-off for interface prediction) mean that currently, no effective structure driven

computational tools are available for the early high-throughput screening stages of the antibody development pipeline.

Machine learning approaches in the field of antibody therapeutics discovery have so far mainly focused on sequence rather than on structure as input [6, 20, 22, 23]. These machine learning approaches have been shown to be highly efficient for prediction tasks which depend only on the antibody (rather than on the antibody-antigen interface) or are specific to one antigen: random forests have been used both for paratope prediction [23] and for the prediction of VH-VL interface angles in antibody modelling [6]. A current state-of-the art paratope prediction tool, Parapred, is based upon a recurrent neural network approach using the variable region amino acid sequence of the antibody as input data [20]. Another recent study has highlighted the potential of large scale, high throughput machine learning approaches for early stage antibody therapeutics development [22]. A large scale (50k samples) mutagenesis study on one antibody target was used to train a recurrent neural network, which was then used to retrieve new antibody sequences from a sub-sample of antibody sequence space, all of which showed binding affinity *in vitro*. For a review of deep learning approaches in antibody research see [11].

Sequence-based machine learning approaches have yet to provide generalisable predictions across different antigens in one model. An approach in this direction has recently been explored in a study by Akbar *et al.* [2], in which machine-learning accessible descriptions of the binding interface were generated using structure-derived interaction motifs. Further, graph convolutional neural networks have recently been used to encode structural information for the prediction of antibody and antigen interface residues [25], demonstrating the ability of machine learning approaches to successfully utilise structural information derived from both the antibody and the antigen.

Here, we describe a structure-based deep learning approach for early stage virtual screening of antibody therapeutics. Our approach is able to make generalisable predictions across different antigens. Adopting a three-dimensional gridding method which has been used successfully alongside convolutional neural network methods to predict small molecule/protein binding [13, 26], we implement a similar convolutional neural network trained on rapidly generated rigid-body docking poses of modelled antibody structures in complex with antigen epitopes. We use this **D**eep **L**earning approach for **A**nti**B**ody screening (**DLAB**) to both improve the ranking of docks from the ZDock docking algorithm [24] and, in combination with docking scores generated by ZDock, for the prediction of antibody-antigen binding.

The DLAB source code is available from the corresponding author on request.

Methods

Crystal structure data set

The structural antibody database (SAbDab) [9] contains an up-to-date collection of all antibody structures deposited in the PDB [5]. We selected a data set of structures of VH-VL paired antibodies in complex with protein or peptide antigens with a resolution of less than 3 Å from a snapshot of the SAbDab, downloaded on 19/12/2018. The data set consisted of 1216 pdb files of antibody-antigen complexes, of which 759 were non-redundant. Here, we considered an antibody to be non-redundant if its CDR sequence (the concatenated sequence across all CDR regions according to IMGT definition) was only present once in the data set. In the following, this data set of 759 complexes is called the crystal structure data set. The PDB accession codes for the crystal structure data set can be found in Supplementary File 1.

Model data set

All antibody structures in the crystal structure data set were modelled using ABodyBuilder [17]. Only non-identical template structures were used for modelling. All sidechains were modelled using PEARS [18]. This set of modelled antibodies is referred to as the model data set. Model quality was assessed by calculating the Cα root mean square deviation (RMSD) of each modelled antibody to the crystal structure of the antibody, either across the entire Fv region or across the antibody CDRs by aligning the modelled and the crystal structure antibody by their framework regions and calculating RMSD across the CDR Cα atoms.

Binding and non-binding antibody/antigen examples

We considered every antibody-antigen pair in the crystal structure data set as a binding pair. Since antibodies bind with high specificity, we generated non-binding antibody-antigen pairs by randomly sampling 50 non-cognate antibodies per antigen from the crystal structure/model antibody data set. Our definition of non-cognate required the sampled antibodies to share less than 90% of their CDR sequence with the binding antibody.

Docking pose generation

For both the crystal structure and model data set, docked antibody-antigen pairs were generated using ZDock [24].

The antigen epitope was defined by separating the antibody and antigen in the crystal structure and calculating surface exposed residues for each binding partner separately using the PSA algorithm [16]. All atoms belonging to surface residues on the antigen less than 4 Å from a surface exposed residue on the antibody were considered part of the epitope. Further, all atoms belonging to a surface exposed

residue on the antigen within 4 Å of the defined epitope were included in the allowed docking interface to model a realistic level of access to the epitope.

The paratope of crystal structure antibodies was defined in the same way, using the interacting residues from the crystal structure.

For modelled antibody structures, the paratope was predicted using Parapred [20]. All residues with a Parapred score over 1% were considered paratope residues. Different score cutoffs were tested (including the 67% value suggested in the Parapred paper). Using a Parapred score of 1% without adding neighbouring residues to the predicted paratope yielded comparable docking results to using the 67% cutoff and adding neighbouring residues to the allowed docking contact area (see Supplementary Figure 4) and was used throughout the study.

Clustering for train-test splits

To avoid similarity between binding modes in the train and test sets, CD-Hit [19] was used to cluster antibodies by CDR sequence identity as defined above, using a clustering cutoff of 90% sequence identity. For all learning tasks set out below, train-test splits were performed using clustered cross-validation, assigning all members of a cluster to either the train or the test set.

Data input for convolutional neural networks

Following the method of Ragoza *et al.* [26], the docking poses were prepared for input into CNNs by discretising the atom information into four-dimensional grids, where three dimensions describe the spatial arrangement of the interaction site and the fourth dimension is used to indicate atom types (see Supplementary Figure 1).

The centre of the interaction site of docking poses was calculated using the PSA algorithm by averaging the coordinates of all surface-exposed atoms within 4 Å of the interaction partner on both the antibody and the antigen and taking the mean of the two center points. Poses which after docking had no interactions under 4 Å were discarded.

The grid contained only atoms that were within 24 Å of the interaction centre. These interaction site specifications were found to cover on average 96% of all interacting atoms on both antibody and antigen (see Supplementary Figure 3). The grid resolution was set to 0.5 Å, leading to a total grid size of 96^3 voxels.

Reranking ZDock docking poses with DLAB-Re

We trained a CNN, termed DLAB-Re(scoring) (architecture shown in Supplementary Figure 2), to predict the *fnat* score (fraction of contacts between the interaction partners in the crystal structure that are recapitulated in the docked pose). The network generates a probability distribution over 11 *fnat* intervals (10 steps of size 0.1 over [0, 1] and one bin for *fnat* 0.0 poses), which are used to generate a *fnat* prediction via weighted averaging. For model training and testing, the top 500 poses as ranked by ZDock for each pairing in the binder set were annotated with their

respective *fnat* score interval. For any given antibody-antigen pairing, there are considerably more poses with low *fnat* scores in the top 500 ZDock poses than high *fnat* poses. To avoid biasing the network towards predicting all poses into low *fnat* intervals, we used a stratified sampling scheme, sampling poses from each interval at the same rate during training (but not during testing).

During training, the input data was augmented by random rotation around the interaction centre, followed by random translations along the x, y and z axis between -2 Å and 2 Å. Models were trained for 200,000 parameter update steps using categorical cross-entropy and the rectified Adam optimiser.

For downstream training and evaluation of the DLAB-VS models, the DLAB-Re models generated during cross-validated training on the cognate antibody-antigen pairings were used to rerank the top 500 poses for all antibody-antigen pairings. For non-cognate pairings, the poses were reranked using the DLAB-Re model not trained on the target antigen.

To identify antibody-antigen pairings with low quality docking poses, we determined the highest DLAB-Re score given to any of the top 500 poses generated by ZDock for each antibody-antigen pairing. This score (DLAB-Re-max) was used to discard particular pairings by ranking all pairings by their DLAB-Re-max score and discarding the bottom 40%, 60% or 80%. To contrast the performance of ZDock on the same task, this score thresholding was also applied to the ZDock output score of the top pose as ranked by ZDock.

Virtual screening with DLAB-VS and ZDock

The goal of virtual antibody screening is to discern binding antibodies against a given epitope from a pool of candidate antibodies. To generate a classification model able to accomplish this task, the two CNN architectures depicted in Supplementary Figure 2 were used to generate a binder/non-binder classifier for individual docking poses of antibody-antigen complexes, termed DLAB-VS (virtual screening).

The input poses for training were selected as follows. For non-binding pairs, the highest ranked pose after DLAB-Re rescoring was selected as a non-binding pose. For binding pairs, we selected up to 50 poses with $fnat > 0.7$ where those were available. Further, following the approach taken by [29], 5 poses of the same binding pair with $fnat < 0.1$ were selected as non-binding poses in order to force the networks to learn from the interaction between antibody and antigen by providing for the same antibody-antigen pairing both good and bad binding poses.

Data augmentation was performed in the same manner as for DLAB-Re. Models were trained for 50,000 parameter update steps using the rectified Adam optimiser (as the smaller input data set resulted in earlier convergence). A validation set comprising 10% of the total data set was created using the same CD-Hit clustering as for the training set creation. The validation set was used to select a snapshot of the model during training by choosing the model snapshot with the highest average precision on the validation set. To counteract class imbalance, binder and non-binder poses were sampled so that each batch contained equal numbers of both classes.

For each train/test split and network architecture, two different validation sets were used to train an ensemble of 4 models per fold.

For binding classification at test time, the DLAB-VS scores of the 10 highest ranked poses after DLAB-Re rescoring for each antibody-antigen pairing were averaged to form the DLAB-VS score for that pairing. If the model ensemble was used, the output scores by the ensemble members were averaged to arrive at the DLAB-VS score for the pairing. For each antigen target, the antibodies docked against that target (correct and decoys) were ranked by their respective DLAB-VS score.

For the ZDock based classifier, the ZDock score of the top-ranked binding pose was used to rank the antibodies docked against a particular target.

For the DLAB-VS+ZDock model, the DLAB-VS output scores and the ZDock output score were normalised per antigen target via minmax scaling and averaged to arrive at the final scores for each target antigen.

The DLAB-Re-max score, output from the reranking method, of each antibody-antigen pairing was used to discard antigens for which the binding antibody was not predicted to have produced any satisfactory docking poses.

DOVE rescoring

We compared the DLAB-Re results to the DOVE method for CNN-based docking pose ranking [32]. Input file preparation and score generation were performed according to the tutorials on the author's github page. As detailed on the author's github page, only the GOAP and ATOM20/ATOM40 scores were used, as the IT-scores were unavailable.

Additional test sets

To create an unseen test set, which was not used at any point during model choice and hyperparameter optimisation, we used SAbDab entries deposited after the snapshot used for the training data set creation to create an unseen test set. This data set, referred to in the following as the post-snapshot model data set, contained 222 antibody/protein antigen complexes, which formed 173 CDR clusters after clustering the CDR-sequences using CD-Hit at 90% identity. On this test set, we performed modelling, docking, rescoring and binder classification as described above, using the models trained on the model data set, with the exception of the ensemble DLAB-VS score calculation, for which we combined all 40 previously trained models into one ensemble from which the DLAB-VS score for each pairing was averaged. The PDB accession codes for this data set can be found in Supplementary File 2.

An additional test set was created using publicly available COVID-19 patient antibody sequencing data [10] alongside confirmed binder sequences against SARS-Cov1 [30] to create a potential realistic test case, referred to in the following as the SARS-Cov model set. Models of the anti-Cov1 antibody 80R as well as 5 confirmed light chain variants of 80R and 13 clonotype matches against 80R (same VJ gene assignment via ANARCI [8] and $>$ 80% CDR H3 sequence identity) derived from

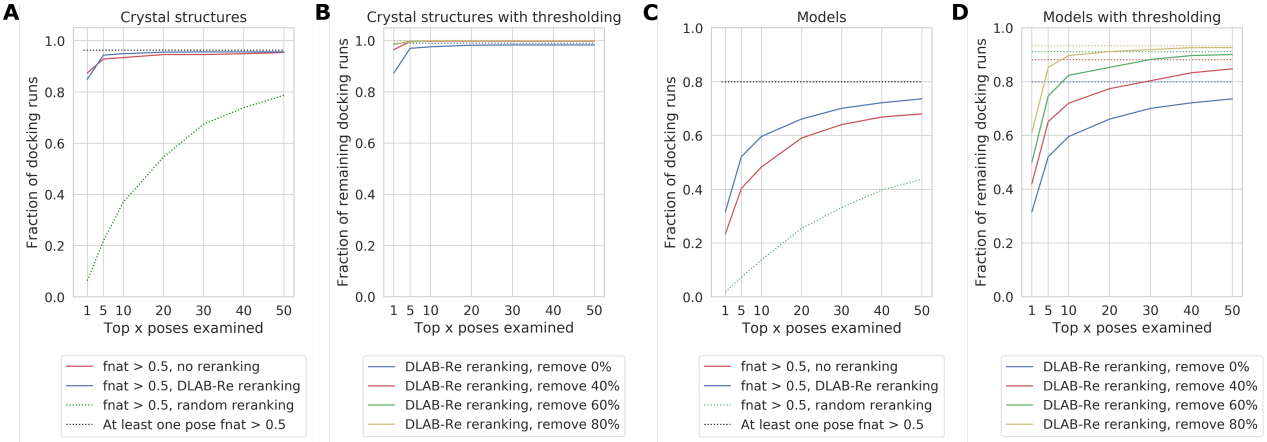

Figure 1. DLAB-Re improves docking performance on the crystal structure data set (**A, B**) and the model data set (**C, D**). On crystal structure data, ZDock ranking and DLAB-Re ranking perform similarly and well. On models, the ZDock baseline performance is considerably worse and DLAB-Re significantly improves ranking performance. **(A, C)**: DLAB-Re ranks the top 500 poses generated by ZDock better than ZDock, enriching the ratio of pairings with *fnat* > 0.5 poses ranked highly. (**B, D**) Using the DLAB-Re-max score to remove 40%, 60% or 80% of the antibody-antigen pairings respectively can remove antibody-antigen pairings which did not yield high-*fnat* poses. This selects for pairings for which *fnat* > 0.5 poses exist in the top 500 poses generated by ZDock (dashed line) and for which DLAB-Re ranks the top 500 poses well (solid line).

COVID-19 patient antibody repertoires were docked against the 80R epitope on the SARS-Cov1 RBD as well as the corresponding area on the SARS-Cov2 RBD. To generate negative samples, 300 samples from a antibody model library (AML) derived from the observable antibody space (OAS) database [15] were also docked against the two epitopes, both with and without the additional constraint that the AML antibodies had to match the 80R H3 length +/- 1 amino acid to avoid the trivial criterion of H3 length clustering. We performed docking, rescoring and binder classification as described above, using the large ensemble of 40 models since neither 80R nor any antibody with a CDR sequence with $> 80\%$ sequence identity to 80R was included in the original training set. The heavy and light chain sequences for the SARS-Cov model set can be found in Supplementary File 3.

Results

Crystal structure docking yields high-quality poses

In order to establish a baseline for ZDock performance on antibody and antigen crystal structures, we re-docked the complexes in the crystal structure data set and quantified the docking performance through the *fnat* score of the docking poses (Supplementary Figure 5). ZDock yielded high-quality docking poses, ranking at least one pose with *fnat* > 0.5 in the top ten poses for 93% of the pairings.

Model docking yields low-quality poses

Docking the model data set antibodies against their cognate antigens on the other hand yields lower quality docking poses. Here, ZDock created a pose with *fnat* > 0.5 in the top ten poses for only 48% of antibody-antigen pairings but 80% of pairings had a pose with *fnat* > 0.5 in their 500 highest ranked poses (Supplementary Figure 5).

DLAB-Re can improve ZDock docking pose ranking

Given these results, the first stage was to create a method that is able to identify good docking poses and rank these correctly. DLAB-Re is a CNN trained to predict the *fnat* of docking poses of antibody-antigen pairings. In order to determine the ability of our method, DLAB-Re, to improve docking pose ranking, we ranked the top 500 docking poses generated by ZDock for each binder pair by the predicted *fnat* value, using the clustered, cross-validated train-test procedure set out in the Methods section.

This rescoring procedure improves upon the performance of ZDock ranking. On the crystal structure data set, DLAB-Re recapitulates the ranking performance of native ZDock. On the model data set, DLAB-Re increased the number of antibody-antigen pairings for which a pose with *fnat* > 0.5 is ranked in the top ten poses by 25% (see Fig. 1 C).

DLAB-Re enables identification of successfully docked antibody-antigen pairings

Using the maximum predicted *fnat* score generated by DLAB-Re, we can discard poorly docked antibody-antigen pairings (see methods and Fig 1B, 1D). Choosing thresholds so that 40%, 60% or 80% of pairings respectively are discarded, the remaining pairings are increasingly enriched both in pairings for which the docking poses are ranked well by DLAB-Re as well as in pairings which have at least one pose with a high *fnat* score in the top 500 poses. Discarding 80% of the pairings in the model data set raises the proportion of pairings for which a pose with at least 0.5 *fnat* was ranked by DLAB-Re in the top ten poses from 56% to 90%, meaning that using this thresholding approach eliminated 95.5 % of the antibody-antigen pairings for which DLAB-Re did not manage to rank a pose with at least 0.5 *fnat* in the top ten poses while retaining 32 % of the pairings for which it did.

Using this approach with the ZDock output scores does not yield the same improvement, only raising the the proportion of pairings for which a pose with at least 0.5 *fnat* was ranked by ZDock in the top ten poses from 48% to 59% (see Supplementary Figure 6).

A CNN docking rescoring tool trained on crystal structure data does not replicate the DLAB-Re performance

We compared the performance of DLAB-Re with the DOVE tool developed by Wang *et al.* [32]. DOVE is a CNN-based docking pose evaluation tool, which is designed to

predict docking pose quality according to CAPRI criteria on crystal structure based general protein-protein docking poses. We used the publicly available ATOM40+GOAP model to test whether this training generalises to the antigen - model antibody docking case. Using the DOVE ATOM40+GOAP score to rerank antibody-antigen docking poses, DOVE performs considerably worse than both ZDock and DLAB-Re. This holds for both the crystal structure and the model data set (see Supplementary Figure 6A and 6B).

DOVE is trained to classify docking poses into one of two classes: CAPRI acceptable and not CAPRI acceptable [34, 35], where CAPRI acceptable poses have $fnat > 0.1$ and interface RMSD < 4A or ligand RMSD < 10A. Using this classification to evaluate both DOVE and DLAB-Re results, DOVE still performed worse than ZDock and DLAB-Re (see Supplementary Figure 6C). These results highlight the added value from training DLAB-Re both on a domain specific (antibody-antigen) as well as task specific (modelled antibodies) data set.

ZDock easily retrieves binders from the crystal structure data set but not from the model data set

Virtual screening for antibody discovery aims to find binding antibodies for a given epitope from a large set of potential binders. We attempted to retrieve correct binders from the crystal structure data set by docking both the cognate antibody crystal structure as well as 50 non-cognate antibody crystal structures against each antigen crystal structure in the data set. Using ZDock, the cognate binder is found in the top 2% (i.e. top ranked) of antibodies for 49.7% of antigen targets, and in the top 10% in 65.6% of antigen targets, outperforming the random baseline. On the model data set, using the same procedure, ZDock ranks the binder in the top 2% of antibodies for only 5.0% of antigen targets, and in the top 10% for only 19.9% of antigen targets (see Supplementary Figure 8).

DLAB-VS and ZDock can be combined to improve performance on the model data set

In order to improve our ability to perform virtual screening on modelled antibodies, we trained a new model, termed DLAB-VS (virtual screening) to classify antibody-antigen pairings as binders or non-binders, as detailed in the methods section. For a large proportion of model antibody structures, the docking pipeline does not yield high quality complex structures (see Fig. 1). Therefore, at train time only docking poses with $fnat > 0.7$ were shown to the network as positive binders and poses with $fnat < 0.1$ were shown to the network as non-binders. Furthermore, at test time we averaged the network output over the ten highest ranked poses for each antibody-antigen pairing. Lastly, rather than training a single model for each train/test fold, we trained four models using two different architectures for each of ten clustered cross-validation folds as detailed in the methods section and used the averaged output as the DLAB-VS score. Using this training approach, the DLAB-VS

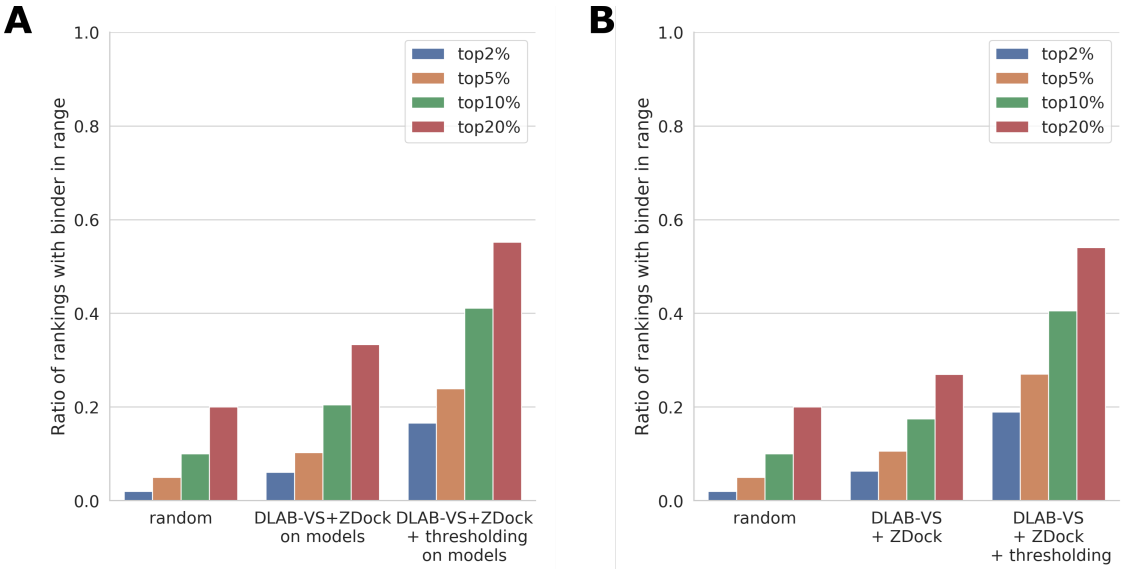

Figure 2. DLAB-VS and ZDock binder classification performance. For each approach, the ratio of pairings for which the binding antibody was ranked in the top 2%, top 5%, top 10% and top 20% respectively is shown. **(A)**: Comparison of the performance of ZDock and DLAB-VS binder classification on the model data set to the random expectation ("random") of finding the binder in the top N%. Using the combination of DLAB-VS and ZDock scores ("DLAB-VS+ZDock") detailed in the methods section and supplementing it with the DLAB-Re-max thresholding ("DLAB-VS+ZDock + thresholding"), the classification performance on the model data set can be improved significantly. **(B)**: Performance on the post-snapshot model data set after removing any CDR sequences with overlap to the model data set (at 90% CDR sequence identity as defined above) both without ("DLAB-VS+ZDock") and with ("DLAB-VS+ZDock + thresholding") DLAB-Re-max score thresholding. The DLAB-VS performance on this test set is better than the performance on the model data set, as discussed in the Results section. Performance for the CDR sequences clustering with sequences in the model data set is shown in Supplementary Figure 10.

model achieved classification performance comparable to ZDock: on the model data set, the cognate antibody was ranked in the top 2% of antibodies for 4.8% of antigens and in the top 10% for 17.2% of antigens (see Supplementary Figure 8).

However, the two approaches, ZDock and DLAB-VS, did not perform equally across antigen targets. It was possible to improve classification performance by using the mean of the two scores to rank putative binders (see methods). Using this approach, termed DLAB-VS+ZDock, the binder was ranked in the top 2% of antibodies for 6.1% of antigen targets, and in the top 10% in 20.4% of antigen targets (see Fig. 2). In the following, we use this DLAB-VS+ZDock approach.

Using DLAB-Re to discard antigen targets enables selection of well-performing models

The performance of the DLAB-VS+ZDock model is highly dependent on the quality of the docking poses from which the score is derived as well as the quality of the antibody model (see Supplementary Figure 9). As described above, DLAB-Re enables the selection of well-docked antibody-antigen pairings, therefore, to further improve the performance of the DLAB-VS+ZDock classifier, we used the output from DLAB-Re to identify antibody models with high likelihood of being well docked from the model data set. To test if this would improve results, for each of the training cross-validation folds, we only considered antigen targets for which the DLAB-Re-max score of the cognate antibody was within the top 20% of DLAB-Re-max scores of antibody-antigen pairings within that fold. On these targets, where the cognate antibody was predicted to be well docked by DLAB-Re, the cognate antibody was ranked in the top 2% for 16.6% of antigens and in the top 10% for 41.1% of antigens (see Fig. 2A). This improvement was reliant on using the combined DLAB-VS+ZDock model, using the same approach while ranking by the ZDock output scores alone only increased the classification performance marginally (see Supplementary Figure 8).

DLAB-VS+ZDock performs well on the post-snapshot model data set

In order to test the performance of the DLAB pipeline on a completely unseen test set, we ran the pipeline on the post-snapshot model data set (see methods). On this set, an ensemble of all 40 DLAB-Vs models trained during the clustered cross-validation training on the model data set and the ZDock output scores were used to rank the cognate antibody model as well as 50 non-cognate antibody models. The ensemble achieved higher performance on this test set than on the previous test cases (both with and without the DLAB-Re selection criterion).

In order to calculate how different the post-snapshot model data set was from the training set and the potential for this to influence model performance, we clustered the antibody CDR sequences in both the model data set and the post-snaphot model data set at 90% identity. Of the new CDR sequences, 17.3% clustered with at least one CDR sequence in the snapshot. On the cognate antigen targets for those

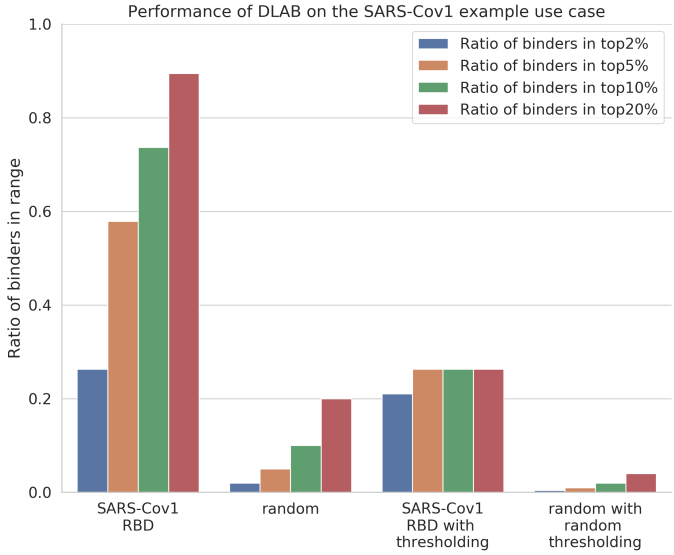

Figure 3. Performance of the DLAB pipeline on the SARS-Cov1 model data set. For each configuration, we show the ratio of the binders which are ranked in the top 2%, top 5%, top 10% and top 20% of the data set. Shown here is the DLAB-VS+ZDock performance against the epitope on the SARS-Cov1 receptor binding domain without and with DLAB-Re-max score thresholding to remove 80% of the antibody-antigen pairings, compared to the random expectation for both cases. The DLAB-VS+ZDock classifier strongly enriches putative binders at the top of the ranking, both before and after thresholding.

antibodies, the large ensemble performs exceptionally well both before and after DLAB-Re thresholding (binder in top 2% for 20% and 50% of antigen targets respectively, see Supplementary Figure 10). On the subset of new additions to SAbDab without overlap to the snapshot, the performance is still slightly better using the 40-model ensemble DLAB-VS+ZDock than on the snapshot using 10 folds of 4-model ensembles (see Fig. 2B).

DLAB-VS can retrieve binders from a random sample of the OAS

In a realistic use case, antibody virtual screening would be performed in order to find binders to a single target antigen with a large set of candidate antibody models, which would be likely to contain several putative binders. In order to test the DLAB pipeline in this setting, we generated a data set for which several binders against the same epitope are known.

The test case was created using antibodies binding to the receptor binding domain on the SARS-Cov1 or SARS-Cov2 spike protein. Using this test set, which contains antibodies with short H3 loops, the DLAB pipeline achieves good performance ranking the binders against both a randomly chosen subset (see Supplementary Figure 11) and a H3 length-matched subset of an antibody model library (AML) (see Fig. 3). On the length-matched subset, for the SARS-Cov1 epitope, 8 of the top 10 ranked antibodies before thresholding and the 4 of the 5 top ranked antibodies after

thresholding with DLAB-Re (discarding the 80% of antibodies in the set with the lowest DLAB-Re-max score) are experimentally confirmed SARS-Cov1 binders or clonotype matches in SARS-Cov2 infected patients for the confirmed binding antibody 80R. Similarly, for the SARS-Cov2 epitope, 7 of the top 10 ranked antibodies before thresholding and 3 of the top 5 ranked antibodies after thresholding are 80R-derived confirmed binders or clonotype matches to 80R.

Discussion

One of the major shortcomings of current computational antibody drug discovery is the lack of structure-based, early-pipeline screening tools to identify promising candidate antibodies.

Here, we have shown how DLAB, our structure-based deep learning approach, can be used to improve pose selection in antibody-antigen docking experiments and can enable the identification of antibody-antigen pairings for which accurate poses have been generated and selected. DLAB-Re is able to identify pairings for which a pose with $fnat > 0.5$ is in the top ten poses for 90% of the pairings, which can be used to improve binder classification performance downstream.

We have furthermore demonstrated that our DLAB tool can identify putative binders to a given epitope in several different settings. The complete DLAB pipeline of docking followed by DLAB-Re and DLAB-VS enriched binders both against the background of non-binding SAbDab-deposited sequences as well as in a more realistic usage scenario against H3 length-matched antibody sequences drawn from antibody repertoire data.

On the crystal data set with highly accurate antibody structures and docking poses, both DLAB-VS and ZDock are able to strongly enrich binders. In the case of model antibodies docked to antigens, where both model and docking quality have to be considered, ZDock and DLAB-VS approaches fail to achieve strong discrimination between binders and non-binders. However combining the two scores improved performance. These results are in line with previously published work on the ability to classify cognate antibodies through cross-docking analysis [14].

On a realistic use case, using the SARS-CoV receptor binding domain epitope as the target epitope, we have demonstrated the utility of the DLAB pipeline, ranking 8 of 19 binding antibodies in the top ten against a set of 300 H3 length-matched decoys.

We have demonstrated the applicability of structure-based deep learning approaches both to antibody research in general and to the virtual screening task specifically. Methods such as DLAB will improve with increasing availability of structural antibody data as well as improved antibody modelling and improved fast docking methods. DLAB demonstrates the potential of structure-based deep learning approaches to supplement traditional experimental screening approaches and sets a course for structure-based virtual screening methods for antibody drug discovery.

Acknowledgements

This work was supported by funding from the Engineering and Physical Sciences Research Council (EPSRC) and the Medical Research Council (MRC) [grant number EP/L016044/1] and AstraZeneca.

References

[1] J. Adolf-Bryfogle, O. Kalyuzhniy, M. Kubitz, B. D. Weitzner, X. Hu, Y. Adachi, W. R. Schief, and R. L. Dunbrack. RosettaAntibodyDesign (RAbD): A general framework for computational antibody design. *PLoS Computational Biology*, 14(4):e1006112, Apr 2018.

[2] R. Akbar, P. A. Robert, M. Pavlović, J. R. Jeliazkov, I. Snapkov, A. Slabodkin, C. R. Weber, L. Scheffer, E. Miho, I. H. Haff, D. T. T. Haug, F. Lund-Johansen, Y. Safonova, G. K. Sandve, and V. Greiff. A compact vocabulary of paratope-epitope interactions enables predictability of antibody-antigen binding. *bioRxiv*, bioRxiv:759498, 2020.

[3] J. C. Almagro, T. R. Daniels-Wells, S. M. Perez-Tapia, and M. L. Penichet. Progress and challenges in the design and clinical development of antibodies for cancer therapy. *Frontiers in Immunology*, 8:1751, Jan 2018.

[4] J. C. Almagro, A. Teplyakov, J. Luo, R. W. Sweet, S. Kodangattil, F. Hernandez-Guzman, and G. L. Gilliland. Second Antibody Modeling Assessment (AMA-II). *Proteins: Structure, Function and Bioinformatics*, 82(8):1553–1562, 2014.

[5] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne. The Protein Data Bank. *Nucleic Acids Research*, 28(1):235–242, 01 2000.

[6] A. Bujotzek, J. Dunbar, F. Lipsmeier, W. Schäfer, I. Antes, C. M. Deane, and G. Georges. Prediction of VH-VL domain orientation for antibody variable domain modeling. *Proteins: Structure, Function and Bioinformatics*, 83(4):681–695, Apr 2015.

[7] M. L. Chiu and G. L. Gilliland. Engineering antibody therapeutics. *Current Opinion in Structural Biology*, 38:163–173, Jun 2016.

[8] J. Dunbar and C. M. Deane. ANARCI: Antigen receptor numbering and receptor classification. *Bioinformatics*, 32(2):298–300, Sep 2015.

[9] J. Dunbar, K. Krawczyk, J. Leem, T. Baker, A. Fuchs, G. Georges, J. Shi, and C. M. Deane. SAbDab: The structural antibody database. *Nucleic Acids Research*, 42(D1):D1140–6, Jan 2014.

[10] J. D. Galson, S. Schaetzle, R. J. Bashford-Rogers, M. I. Raybould, A. Kovaltsuk, G. J. Kilpatrick, R. Minter, D. K. Finch, J. Dias, L. K. James, G. Thomas, W. Y. J. Lee, J. Betley, O. Cavlan, A. Leech, C. M. Deane, J. Seoane, C. Caldas, D. J. Pennington, P. Pfeffer, and J. Osbourn. Deep Sequencing of B Cell

Receptor Repertoires From COVID-19 Patients Reveals Strong Convergent Immune Signatures. *Frontiers in Immunology*, 11:3283, Dec 2020.

[11] J. Graves, J. Byerly, E. Priego, N. Makkapati, S. V. Parish, B. Medellin, and M. Berrondo. A Review of Deep Learning Methods for Antibodies. *Antibodies*, 9(2):12, Apr 2020.

[12] A. L. Grilo and A. Mantalaris. The Increasingly Human and Profitable Monoclonal Antibody Market. *Trends in biotechnology*, 37(1):9–16, Jan 2019.

[13] F. Imrie, A. Bradley, M. van der Schaar, and C. M. Deane. Protein Family-specific Models using Deep Neural Networks and Transfer Learning Improve Virtual Screening and Highlight the Need for More Data. *Journal of Chemical Information and Modeling*, acs.jcim.8b00350, Oct 2018.

[14] K. P. Kilambi and J. J. Gray. Structure-based cross-docking analysis of antibody-antigen interactions. *Scientific Reports*, 7(1):8145, Dec 2017.

[15] A. Kovaltsuk, J. Leem, S. Kelm, J. Snowden, C. M. Deane, and K. Krawczyk. Observed antibody space: A resource for data mining next-generation sequencing of antibody repertoires. *The Journal of Immunology*, 201(8):2502–2509, 2018.

[16] B. Lee and F. M. Richards. The interpretation of protein structures: Estimation of static accessibility. *Journal of Molecular Biology*, 55(3):379–IN4, Feb 1971.

[17] J. Leem, J. Dunbar, G. Georges, J. Shi, and C. M. Deane. ABodyBuilder: Automated antibody structure prediction with data–driven accuracy estimation. *mAbs*, 8(7):1259–1268, Oct 2016.

[18] J. Leem, G. Georges, J. Shi, and C. M. Deane. Antibody side chain conformations are position-dependent. *Proteins: Structure, Function and Bioinformatics*, 86(4):383–392, Apr 2018.

[19] W. Li and A. Godzik. Cd-hit: A fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics*, 22(13):1658–1659, Jul 2006.

[20] E. Liberis, P. Veličković, P. Sormanni, M. Vendruscolo, and P. Liò. Parapred: antibody paratope prediction using convolutional and recurrent neural networks. *Bioinformatics*, 34(17):2944–2950, Apr 2018.

[21] J. K. H. Liu. The history of monoclonal antibody development - Progress, remaining challenges and future innovations. *Annals of Medicine and Surgery*, 3(4):113–116, Dec 2014.

[22] D. M. Mason, S. Friedensohn, C. R. Weber, C. Jordi, B. Wagner, S. Meng, and S. T. Reddy. Deep learning enables therapeutic antibody optimization in mammalian cells. *bioRxiv*, bioRxiv:617860, 2019.

[23] P. P. Olimpieri, A. Chailyan, A. Tramontano, and P. Marcatili. Prediction of site-specific interactions in antibody-antigen complexes: The proABC method and server. *Bioinformatics*, 29(18):2285–2291, Sep 2013.

[24] B. G. Pierce, Y. Hourai, and Z. Weng. Accelerating protein docking in ZDOCK using an advanced 3D convolution library. *PLoS ONE*, 6(9):e24657, Sep 2011.

[25] S. Pittala and C. Bailey-Kellogg. Learning context-aware structural representations to predict antigen and antibody binding interfaces. *Bioinformatics*, 36(13):3996–4003, Jul 2020.

[26] M. Ragoza, J. Hochuli, E. Idrobo, J. Sunseri, and D. R. Koes. Protein-Ligand Scoring with Convolutional Neural Networks. *Journal of Chemical Information and Modeling*, 57(4):942–957, Apr 2017.

[27] M. I. Raybould, W. K. Wong, and C. M. Deane. Antibody-antigen complex modelling in the era of immunoglobulin repertoire sequencing. *Molecular Systems Design and Engineering*, 4(4):679–688, 2019.

[28] M. I. J. Raybould, C. Marks, K. Krawczyk, B. Taddese, J. Nowak, A. P. Lewis, A. Bujotzek, J. Shi, and C. M. Deane. Five computational developability guidelines for therapeutic antibody profiling. *Proceedings of the National Academy of Sciences*, 116(10):4025–4030, 2019.

[29] J. Scantlebury, N. Brown, F. Von Delft, and C. M. Deane. Data Set Augmentation Allows Deep Learning-Based Virtual Screening to Better Generalize to Unseen Target Classes and Highlight Important Binding Interactions. *Journal of Chemical Information and Modeling*, 60(8):3722–3730, Aug 2020.

[30] J. Sui, D. R. Aird, A. Tamin, A. Murakami, M. Yan, A. Yammanuru, H. Jing, B. Kan, X. Liu, Q. Zhu, Q.-a. Yuan, G. P. Adams, W. J. Bellini, J. Xu, L. J. Anderson, and W. A. Marasco. Broadening of neutralization activity to directly block a dominant antibody-driven sars-coronavirus evolution pathway. *PLOS Pathogens*, 4(11):1–14, Nov 2008.

[31] T. J. Suscovich and G. Alter. In situ production of therapeutic monoclonal antibodies. *Expert Review of Vaccines*, 14(2):205–219, Feb 2014.

[32] X. Wang, G. Terashi, C. W. Christoffer, M. Zhu, and D. Kihara. Protein docking model evaluation by 3D deep convolutional neural networks. *Bioinformatics*, 36(7):2113–2118, Nov 2019.

[33] B. D. Weitzner, J. R. Jeliazkov, S. Lyskov, N. Marze, D. Kuroda, R. Frick, J. Adolf-Bryfogle, N. Biswas, R. L. Dunbrack, and J. J. Gray. Modeling and docking of antibody structures with Rosetta. *Nature Protocols*, 12(2):401–416, 2017.

[34] S. J. Wodak and R. Méndez. Prediction of protein-protein interactions: The CAPRI experiment, its evaluation and implications. *Current Opinion in Structural Biology*, 14(2):242–249, Apr 2004.

[35] W. K. Wong, J. Leem, and C. M. Deane. Comparative analysis of the cdr loops of antigen receptors. *Frontiers in Immunology*, 10:2454, 2019.

[36] K. Yamashita, K. Ikeda, K. Amada, S. Liang, Y. Tsuchiya, H. Nakamura, H. Shirai, and D. M. Standley. Kotai Antibody Builder: Automated high-resolution structural modeling of antibodies. *Bioinformatics*, 30(22):3279–3280, 2014.

Supplementary Information

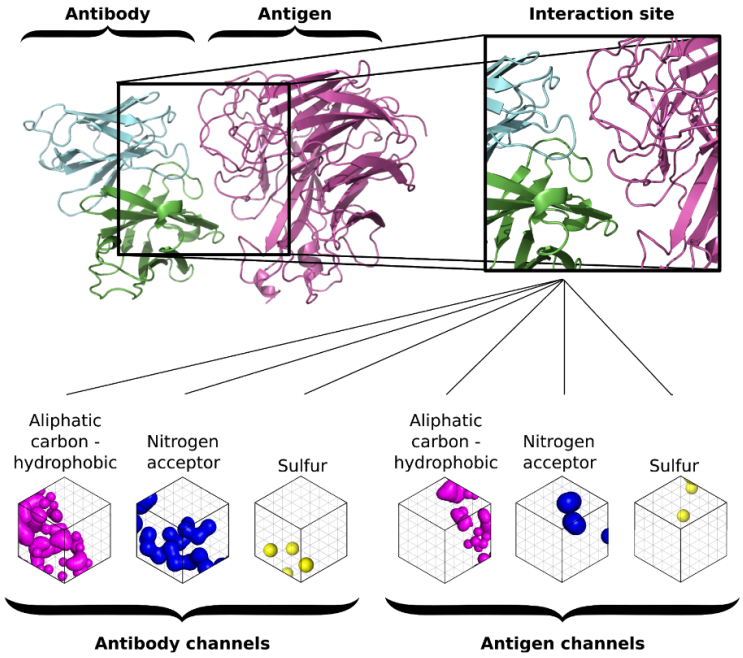

Supplementary Figure 1. The binding interface discretisation algorithm. (**Top**) The interface is defined as a 48 Å cube centered on the interaction center between antibody and antigen. (**Bottom**) The atoms occupying the interface in the docking pose are discretised using the libmolgrid python API into 3D arrays of atom type densities (16 types for both antibody and antigen, of which 3 are depicted here as an example).

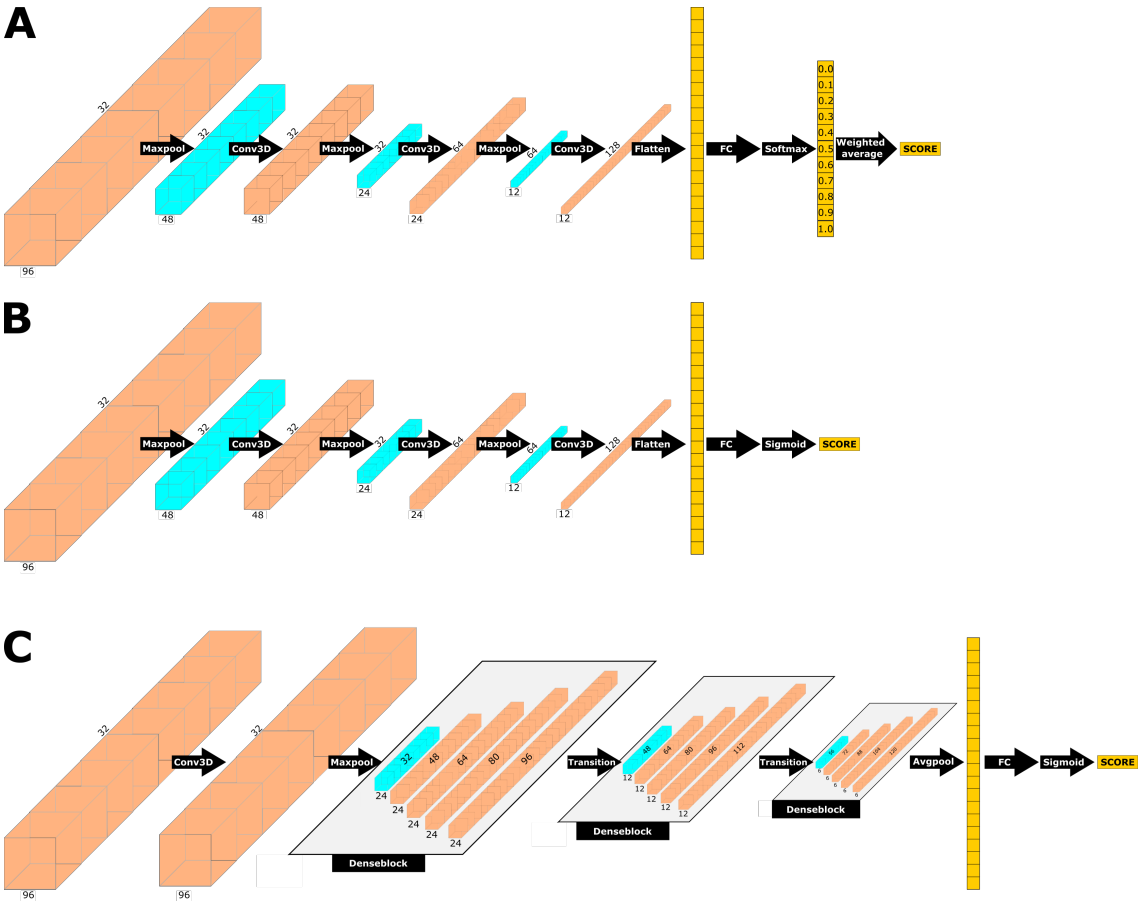

Supplementary Figure 2. CNN architectures used in the study. (**A**) DLAB-Re architecture. A 3-layer CNN followed by a fully connected layer is used to predict membership of one of 11 *fnat* intervals. To make *fnat* predictions, the output of the softmax layer is transformed into a single score via weighted averaging. (**B, C**) The DLAB-VS architectures forming the ensemble of DLAB-VS models. Both architectures output a binary binder/non-binder classification for an antibody/antigen pairing. (**B**) The same CNN architecture as in (A) with a different output layer. (**C**) A CNN architecture consisting of a single convolutional layer followed by 3 Denseblocks, using the same output layer as (B).

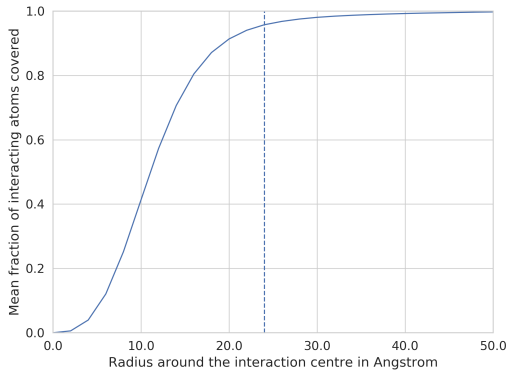

Supplementary Figure 3. The fraction of interacting atoms included in the interaction box depending on the radius around the interaction center. The dashed vertical line indicates 24 Å, the radius used in this study.

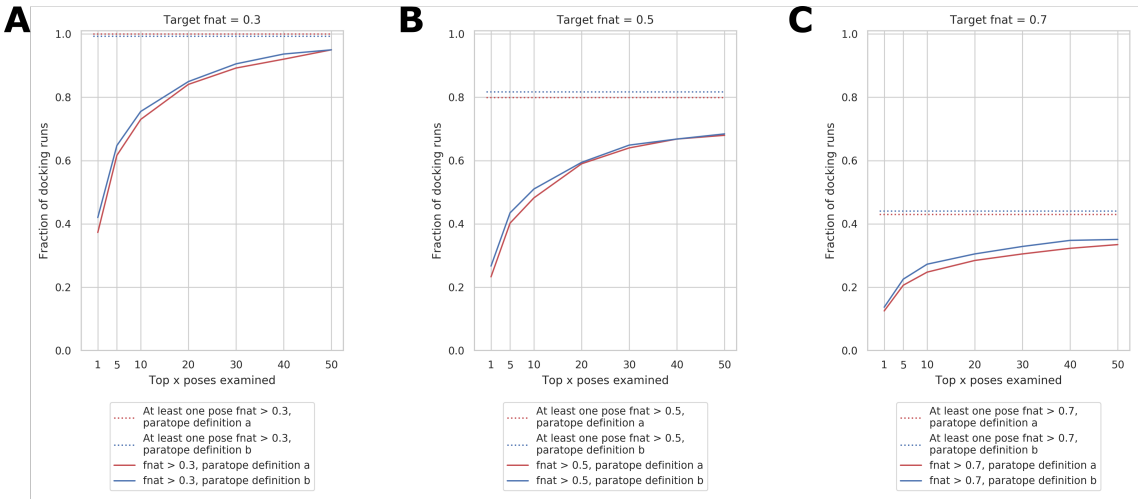

Supplementary Figure 4. Comparison of two paratope definitions for docking using ZDock. Definition a uses the Parapred output at 1% score threshold to determine the paratope for the purpose of defining the allowed interaction area during docking. Definition b uses a score threshold of 67%, but also includes residues within 4 Å of the Parapred defined paratope within the allowed interaction area. The two approaches are compared by the fraction of antibody-antigen pairings for which the top x poses as ranked by ZDock contain at least one pose with *fnat* over 0.3 (A), 0.5 (B) or 0.7 (C) respectively (solid lines) as well as the fraction of pairings for which a pose meeting the target *fnat* is present in the top 500 poses (dashed line).

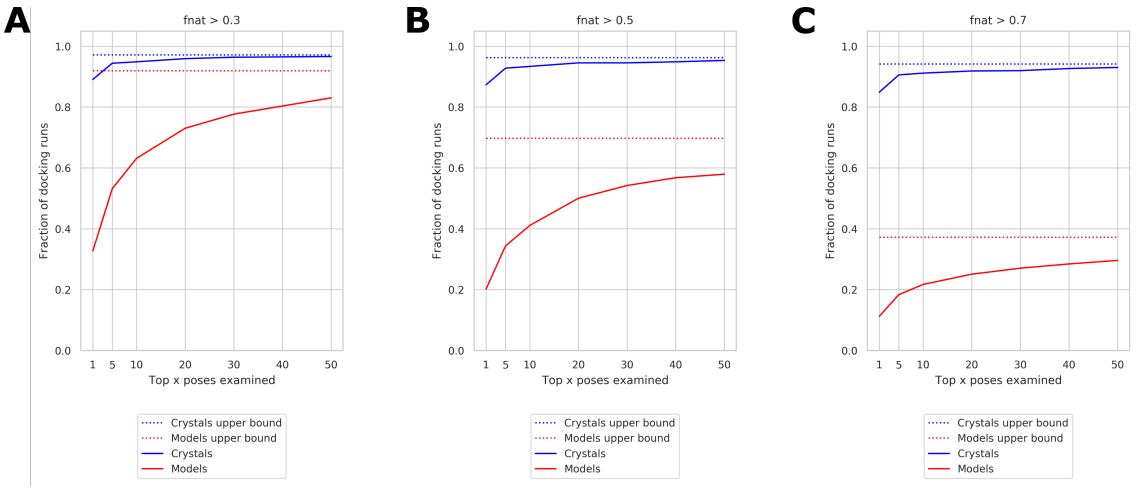

Supplementary Figure 5. Comparison of ZDock runs on crystal structures and models. Each of the three plots shows the fraction of antibody-antigen pairings which have at least one pose meeting the threshold *fnat* score 0.3 (A), 0.5 (B) or 0.7 (C) in their x highest ranked docks. The dashed line indicates the fraction of docks which have such a pose in their top500 poses (the upper bound of the pose ranking algorithm.

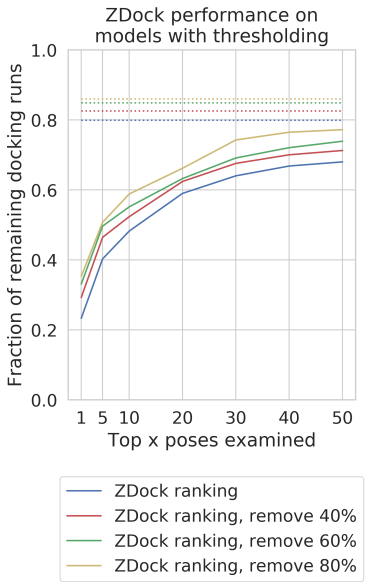

Supplementary Figure 6. Performance of ZDock-based max score thresholding on ranking performance. While higher overall ZDock scores correlate with improved ranking performance, the effect is much less pronounced than using DLAB-Re-max scores. The solid lines indicate the fraction of antibody-antigen pairings which have at least one pose with *fnat* over 0.5 in their x highest ranked docks, the dotted line indicates the fraction of pairings for which a pose with *fnat* over 0.5 exists in top 500 poses.

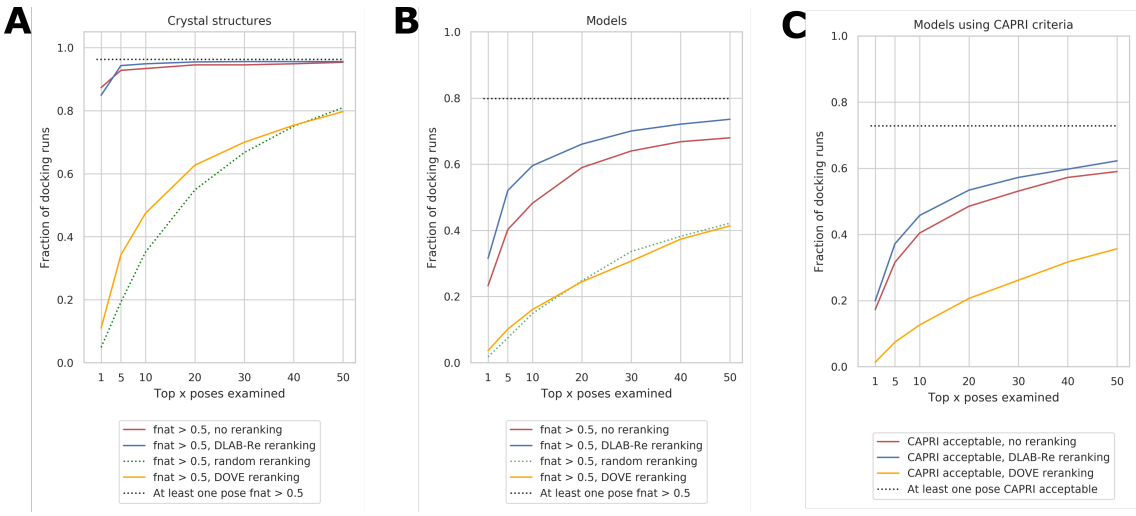

Supplementary Figure 7. Comparison of DLAB-Re performance to DOVE performance. (A, B) Comparison using the pose *fnat* scores on the crystal structure data set (A) and the model data set (B). The fraction of antibody-antigen pairings which have at least one pose with *fnat* over 0.5 in their x highest ranked docks for ZDock ranking ("no reranking") as well as DOVE and DLAB-Re reranking is shown in solid lines, the baseline of randomly shuffling the top 500 poses is shown with a dotted green line and the ratio of pairings with a pose with *fnat* over 0.5 is indicated with a dotted black line. (C) Comparison using the CAPRI acceptability criterion for each pose. Legend explanation see (A, B). In all three cases (A, B, C), DLAB-Re achieves an improvement over native ZDock, while DOVE does not replicate the ranking performance achieved by ZDock.

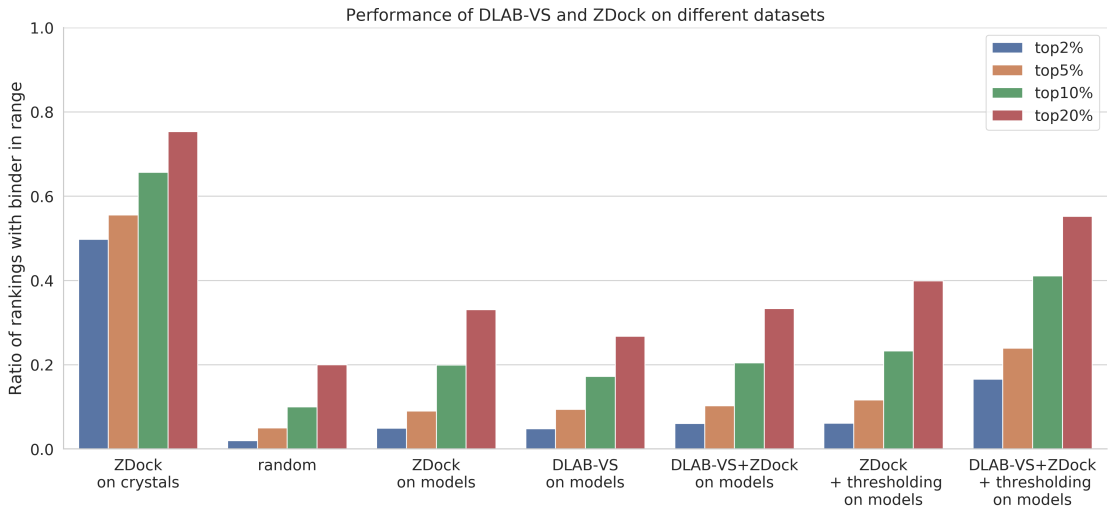

Supplementary Figure 8. DLAB-VS and ZDock binder classification performance on different data sets. For each approach, the ratio of pairings for which the binding antibody was ranked in the top 2%, top 5%, top 10% and top 20% respectively is shown. Comparison of the performance of ZDock binder classification on the crystal data set to the random expectation ("random") of finding the binder in the top N% and the performance of ZDock, DLAB-VS and the combination of ZDock and DLAB-VS generated as detailed in the Methods section ("DLAB-VS+Zdock") on the model data set, both with and without using the DLAB-Re-max thresholding approach ("+ thresholding").

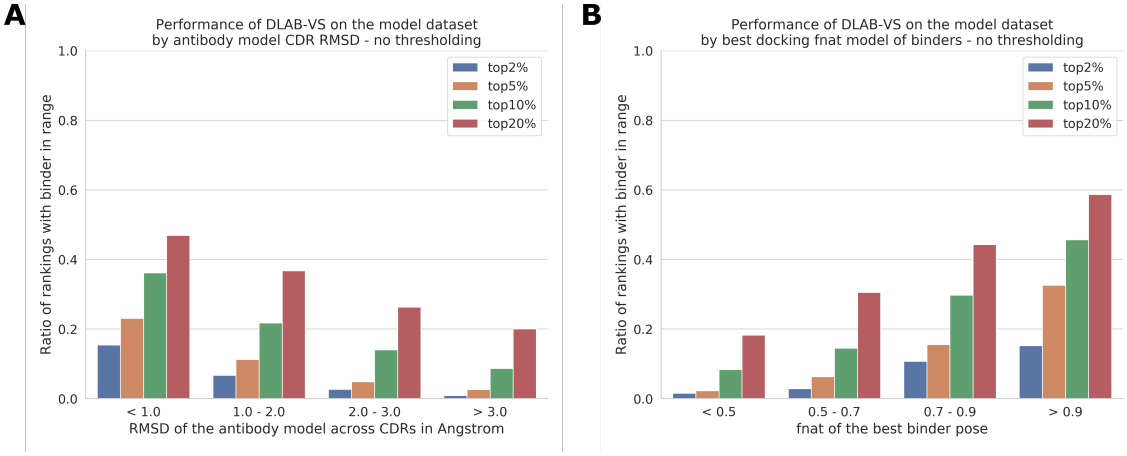

Supplementary Figure 9. Dependence of the performance of DLAB-VS+ZDock on antibody model quality (A), measured via RMSD of the CDR regions to the corresponding crystal structure, and docking quality (B), measured via the highest fnat achieved in the top500 docking poses generated by ZDock. Both good antibody models and high quality docking poses correlate with high classification performance.

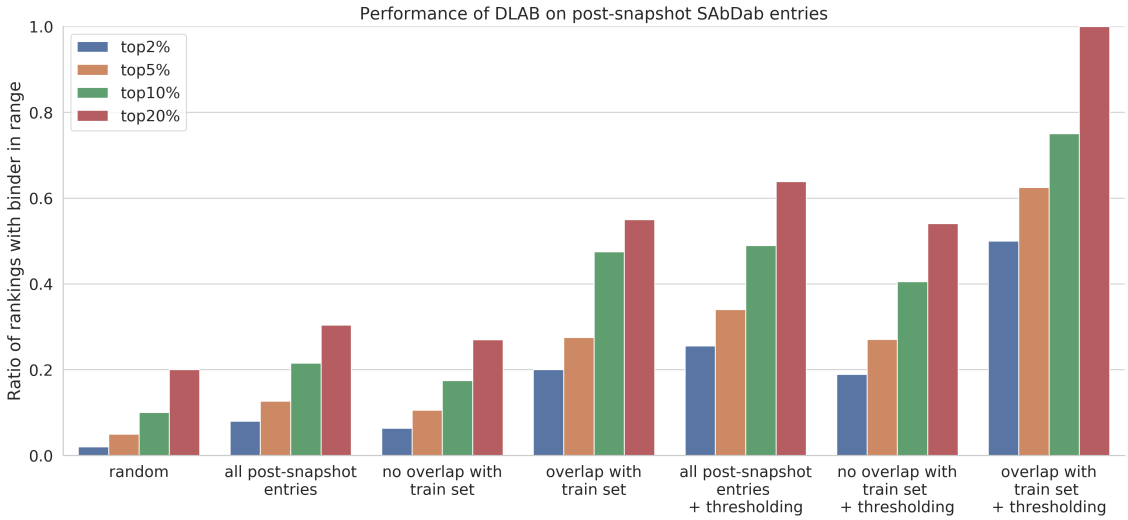

Supplementary Figure 10. Full overview of the performance of DLAB-VS+ZDock on the post-snapshot model data set. As in the main text figure, the bars show the fraction of antigens within the data set for which the correct binder is ranked by DLAB-VS within the top 2%, top 5%, top 10% or top 20% respectively. We compare the random expectation value ("random") to the whole post-snapshot model set ("all post-snapshot entries"), the antigen targets within the set for which the binding antibody does ("overlap with train set") or does not ("no overlap with train set") cluster with at least one antibody from the model data set at 90% CDR sequence identity. For each of these three options, the improved performance upon using the DLAB-Re-max score to discard 80% of antigen targets as described above is shown as well ("+ thresholding").

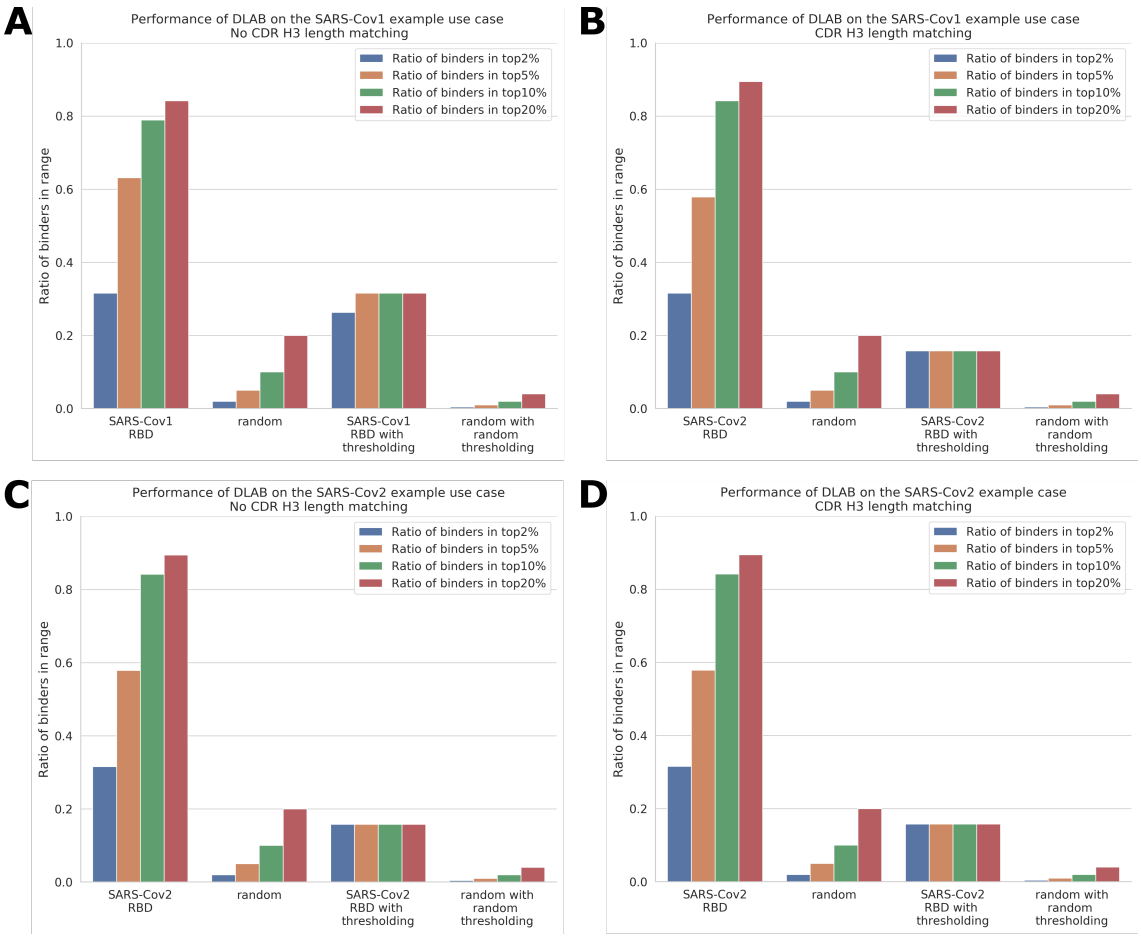

Supplementary Figure 11. Performance of the DLAB pipeline on the SARS-CoV test sets, docking against the SARS-CoV1 (A, B) or the SARS-CoV2 (C, D) RBD while either choosing a completely random decoy set (A, C) or matching the length of the CDR H3 in the decoys to the binders in the test set (B, D)