
ContactNet: Geometric-Based Deep Learning Model for Predicting Protein-Protein Interactions

Matan Halfon

School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
matan.halfon@mail.huji.ac.il

Tomer Cohen

School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
tomer.cohen13@mail.huji.ac.il

Raanan Fattal

School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
raanan.fattal@mail.huji.ac.il

Dina Schneidman-Duhovny

School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
dina.schneidman@mail.huji.ac.il

Abstract

Deep learning approaches achieved significant progress in predicting protein structures. These methods are often applied to protein-protein interactions (PPIs) yet require Multiple Sequence Alignment (MSA) which is unavailable for various interactions, such as antibody-antigen. Computational docking methods are capable of sampling accurate complex models, but also produce thousands of invalid configurations. The design of scoring functions for identifying accurate models is a long-standing challenge. We develop a novel attention-based Graph Neural Network (GNN), ContactNet, for classifying PPI models obtained from docking algorithms into accurate and incorrect ones. When trained on docked antigen and modeled antibody structures, ContactNet doubles the accuracy of current state-of-the-art scoring functions, achieving accurate models among its Top-10 at 43% of the test cases. When applied to unbound antibodies, its Top-10 accuracy increases to 65%. This performance is achieved without MSA and the approach is applicable to other types of interactions, such as host-pathogens or general PPIs.

1 Introduction

Experimental methods for structure determination of macromolecular complexes, such as X-ray crystallography and cryo- Electron Microscopy, can not be applied in a high-throughput manner [4]. Recent progress in deep learning models for protein folding also enabled modeling of protein-protein complexes with high accuracy [2, 9, 14]. However, these methods rely on a co-evolutionary signal from MSA which is not available for many complexes, such as antibody-antigen, leading to inaccurate predictions [31]. Computational docking methods for modeling PPIs, while applicable on a larger scale, still suffer from low accuracy [18]. Docking usually involves two parts: sampling and scoring. Sampling methods, such as Fast Fourier Transform (FFT) or geometry-based approaches [24, 11, 15], usually sample accurate models for the “rigid” docking cases, where there are no significant conformational changes between the bound and unbound structures. However, the scoring presents a major challenge in identifying a few accurate models among the thousands that are generated by the sampling methods.

Much effort has been devoted in developing scoring functions for assessing docked complexes [5, 1, 23]. Despite this effort, the accuracy remains fairly low: state-of-the-art scoring functions

rank correct models among its Top-10 models in 10% to 40% of the cases for unbound docking [12], depending on the dataset and the scoring function. The accuracy drops significantly, below 15%, when the comparative models are used for docking [26]. Recently, there has been progress in designing scoring functions using deep learning models, including 3D convolutional [29] and hierarchical rotation-equivariant networks [8]. However, they were not able to surpass traditional scoring functions, such as ZRANK [22] and SOAP-PP [7]. Graph representation and GNN models are a well-suited for protein structures and protein-protein interactions resulting in some success in scoring tasks [6, 30]. Here, we describe a new GNN-based model, ContactNet, which is specifically designed to model biological protein-protein recognition interfaces consisting of short interacting linear fragments with the appropriate invariances. The model detects and analyzes contacting patches in protein-protein interfaces to resolve the docked model validity. Specifically, the network consists of the following processing stages. First, it encodes an embedding of both the structural (geometric) and chemical properties of each residue in each input protein. The geometry is introduced by a variant of graph attention mechanism that limits the interaction to spatially neighboring amino acids. Second, the inter-protein distance matrix is used for detecting potentially interacting linear segments, which are also analyzed and encoded into low-dimensional contact descriptor vectors. Finally, a sequence of encoding-transformer layers integrates the information gathered from all the interacting linear segments to score input docked models. ContactNet was trained and tested on antibody-antigen data and achieved highly accurate classifications of antibody-antigen models generated by docking, where it nearly doubled the success rate of state-of-the-art methods. Specifically, it detected an accurate model among the Top-10 best scoring for 43% of the test set, compared to 22% by current methods.

2 Methods

2.1 ContactNet architecture

ContactNet starts by extracting low-level features at the level of neighboring amino acids in each protein and uses them to analyze larger scale contact regions and resolve their global interaction across the interface between the two proteins. This processing pipeline is performed by a sequence of three modules (Fig. 1):

Single-protein embedding module This module uses the chemical properties of each amino acid in each protein along with its neighboring amino acids and contextualizes this information in an embedded representation. We implement this module using a distance aware graph attention layer [27] where C atoms are the nodes and the edges correspond to amino acids in close proximity ($C-C$ distance $< 16\text{\AA}$). Specifically, we calculate the attention weight between two residues as $w_{ij} = \text{softmax}(\langle k_i, q_j \rangle - d_{ij})$ where $k_i = MLP(r_i)$ and $q_j = MLP(r_j)$ for residue i , and obtain a context-aware representation by $e_i = \sum_0^N w_{i,j} v_j$. allows us to introduce an inductive bias that decouples non-neighboring residues. The embedded representation is computed for each protein in the complex separately and is denoted by e_i^{receptor} and e_i^{ligand} .

Contacts embedding module The association between the proteins (binding or not) is mainly determined by the type of interaction between localized linear segments. Typically, there are around 6 peptide segments in an average interface of $1,000\text{\AA}^2$ with most of the segments up to 13 amino acids in length [21, 19]. We carefully design our network architecture to induce its inference according to these structural considerations.

The inter-protein distogram D_{ij} (distance matrix of the receptor-ligand complex bipartite-graph) is a fairly sparse, and its non-trivial elements correspond to contacts between the two proteins. Consequently, we focus our classification on the neighboring segments that are identified using a non-maximum-suppression algorithm over the inter-protein distogram matrix D_{ij} . We then extract short linear segments around these amino acids (10 amino acids in each chain direction) from their embedded single-protein representation, i.e., $[e_{i-L}^{\text{receptor}}; \dots; e_{i+L}^{\text{receptor}}]$ and $[e_{j-L}^{\text{ligand}}; \dots; e_{j+L}^{\text{ligand}}]$, where the length of these segments is $2L+1$. To allow the network to account for the interaction between every pair of amino acids in the two segments, we generate an interaction descriptor (Fig. 1B) by tiling each segment along with the other and concatenating the two, as given by $E_{i,j}(s; t) = [e_{i+s}^{\text{receptor}}; e_{j+t}^{\text{ligand}}]$ [28]. Thus, the spatial coordinates $s; t$ of the resulting 3D tensor, $E_{i,j}(s; t)$, correspond to different pairs of amino acids, and its third coordinate parametrizes their

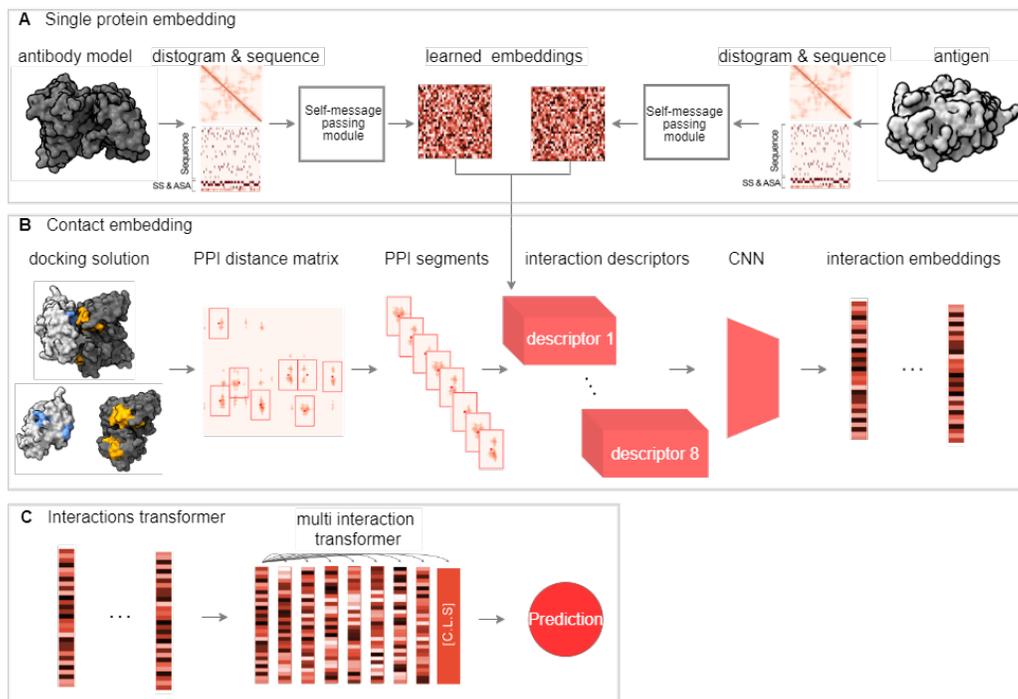


Figure 1: **ContactNet architecture.** **A** Single protein embedding module. This module uses sequence, secondary structure, solvent accessibility, and the $C \times C$ protein distance matrix. The network learns a new representation of physico-chemical features using an encoder transformer. **B** Contacts embedding module. This module extracts the most contacting and possibly interacting linear segments from the single protein embedding stage and encodes them into interaction descriptors. **C** Interaction transformer module. The transformer module is trained to classify the entire interaction interfaces based on these embedded contact descriptors.

embedded representation. By feeding each $E_{i,j}(S; t)$ through a 2D convolutional network, containing multiple pooling operators, we allow the model to analyze multiple sub-segments at different scales and lengths and produce a low-dimensional encoding vector c_{ij} for each i,j contact. This process is applied over the top 8 non-overlapping interacting segment pairs detected from the inter-protein distogram D_{ij} .

Interaction transformer module The resulting contact descriptor vectors c_{ij} are then passed, as separate tokens, to a classification transformer network (along with a special classification token). This global operator allows the model to integrate data from all the contacting segments to derive a final prediction of the validity of the complex. This module is capable of learning general non-linear relations between the contacts while being invariant to their order. The resulting classification token is then passed through an MLP to produce the final prediction.

3 Results

Prediction protocol for antigen-antibody complexes. We trained and tested ContactNet on 875 antibody-antigen structures from the AbDb dataset (see A.1). The antibodies were modeled ab initio using AlphaFold2-multimer-v1 (AFM) [9] to simulate real-life scenario where only antibody sequences are available. Overall, five models were obtained for most of the antibodies. Docking was run on the antigen and five antibody AFM models using the antibody-antigen PatchDock protocol (see A.5, Fig. 2). Docking models were re-ranked using SOAP-PP statistical potential [7] and the

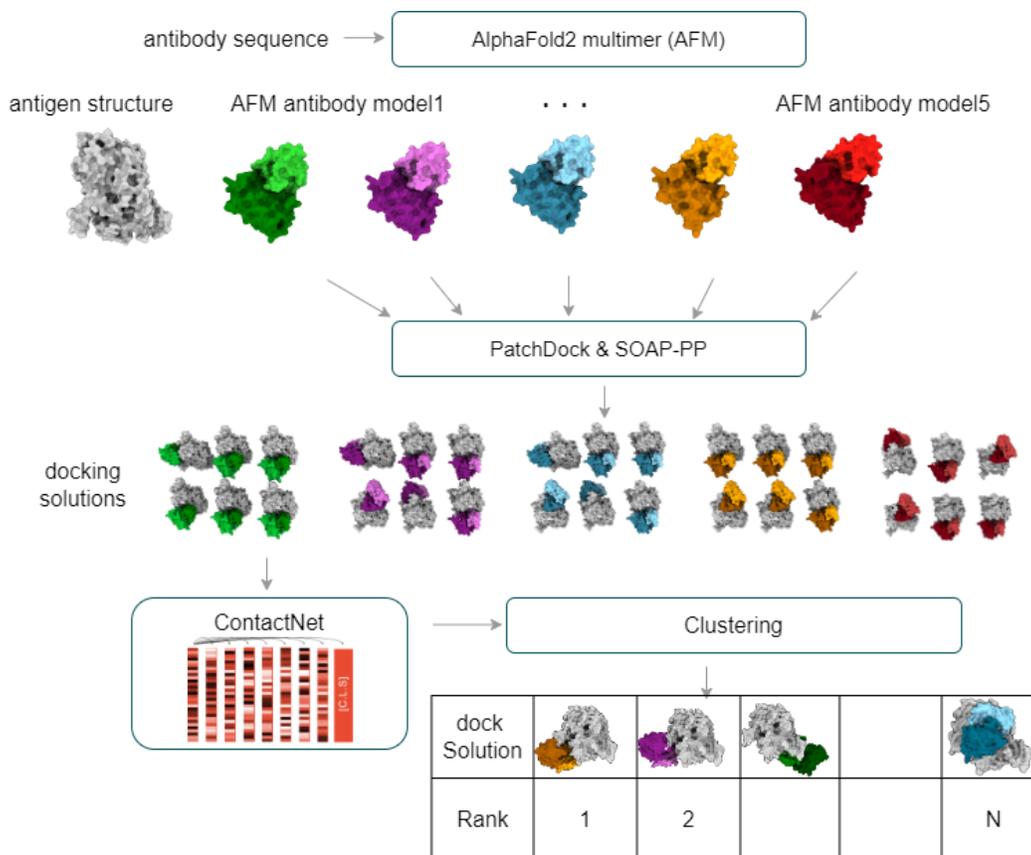


Figure 2: Prediction protocol for antigen-antibody complexes. The antibody sequence is modeled via AFM to obtain five modeled structures that are docked using PatchDock. The top docking models ranked by SOAP-PP score are reevaluated by ContactNet. Finally, the docking models from the five AFM antibody models are clustered to reduce redundancy.

top 3000 were assessed by ContactNet. Finally, we clustered the docking models among the five antibody AFM models using interface clustering [5] with an RMSD threshold of 5Å.

ContactNet performance. Overall, due to the sampling protocol (antibody modeling and docking), acceptable or higher accuracy docking models according to CAPRI criteria (see A.2) were obtained for 88% of the complexes (out of 85 in the test set), which is the upper bound for our success rate (Fig. 4A). Most of the correct models produced by docking were of acceptable and medium accuracy, with high accuracy models sampled rarely. We compared ContactNet to SOAP-PP statistical potential which defines the state-of-the-art on several benchmarks and to the five antibody-antigen complex models obtained from the AFM, as it was shown to outperform state-of-the-art docking methods on several benchmarks [9, 31].

We find that ContactNet almost tripled the Top-1 success rate (see A.2) compared to AFM and SOAP-PP (Table 1, Fig. 4A) for modeled antibodies. The Top-5 success rate was also significantly higher compared to both AFM and SOAP-PP. AFM had a better performance compared to SOAP-PP on our modeled antibodies dataset. Most of the hits selected by ContactNet in the Top-5 were of medium and acceptable accuracy (67% and 33%), while most of AFM models were of high accuracy (Fig. 4B). This can be explained by the low number of high-accuracy models sampled by the docking algorithms.

When tested on 32 unbound antibody-antigen complexes from Benchmark 5.0 [13], ContactNet achieved a higher performance with a success rate of 68% and 75% for Top-1 and Top-5 docking models, respectively (Fig. 4C). In contrast, SOAP-PP had a success rate of 20% and 35% and AFM

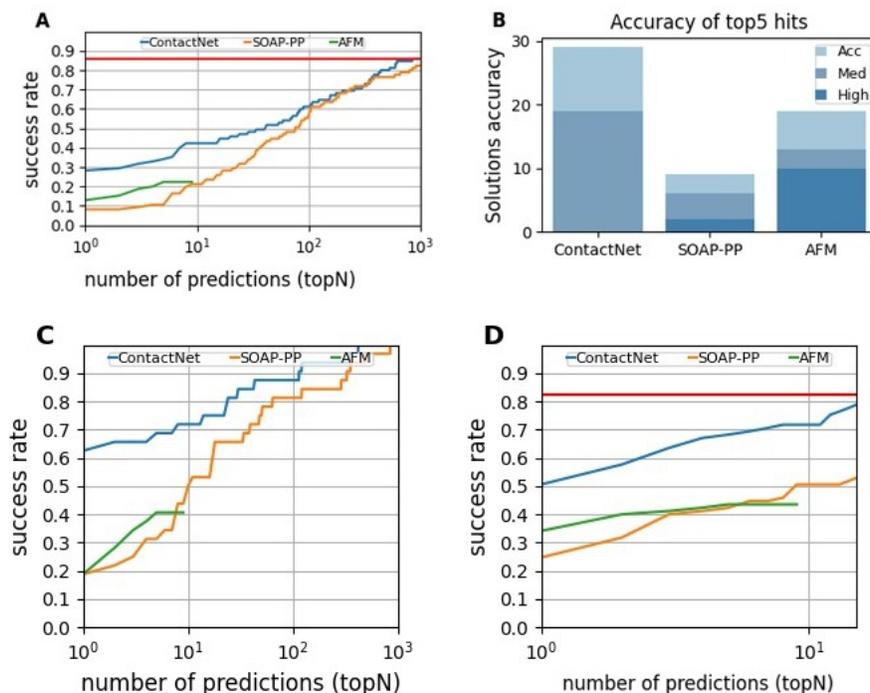


Figure 3: **Performance of ContactNet, AFM, and SOAP-PP on the modeled antibodies test set.** **A.** Success rate for TopN predictions for ContactNet (blue), AFM (green), and SOAP-PP (orange). The red line indicates the upper bound for ContactNet and SOAP-PP due to missing acceptable or higher accuracy models. **B.** Top5 models of each method divided into high, medium, and acceptable accuracy. **C.** Success rate for topN predictions on the unbound test set for ContactNet, AFM, and SOAP-PP. **D.** Success rate for topN in predicting epitopes for ContactNet, AFM, and SOAP-PP.

had a success rate of 20% and 40%, for Top-1 and Top-5 respectively. The use of unbound antibodies should not affect AFM performance because only the sequence was used as an input.

We also tested ContactNet for epitope prediction using interface assessment criteria (see A.3) [16]. ContactNet reached a success rate of 50% and 70% for Top-1 and Top-10 docking models, respectively (Fig. 4D). In contrast, SOAP-PP had a success rate of 25% (Top-1) and 50% (Top-10) and AFM had a success rate of 38% (Top-1) and 42% (Top-5).

4 Conclusions

We present a novel deep learning architecture to address the problem of scoring antibody-antigen docking models. We train ContactNet for antibody-antigen complexes which present significant challenges, including the need to model antibodies and the lack of MSA that facilitates AFM prediction. ContactNet has significantly higher accuracy compared to state-of-the-art methods.

Several key factors contribute to the high performance of ContactNet. First, the compact 2D distogram representation for the protein and for the antibody-antigen interface makes the method invariant to transformations. Second, the network architecture is designed to simulate the biological process of protein-protein association where the interface contacts are formed by small complementary patches on the surfaces of the two proteins. The use of a problem-tailored architecture enables us to derive meaningful information out of the limited amount of data and create meaningful inductive bias. Finally, the residue level protein representation not only enables efficient training of the model using less GPU memory and runtime but also makes the network less sensitive to the differences between bound and unbound structures which often have different side-chain orientations.

References

- [1] Nelly Andrusier, Ruth Nussinov, and Haim J Wolfson. Firedock: fast interaction refinement in molecular docking. *Proteins: Structure, Function, and Bioinformatics*, 69(1):139–159, 2007.
- [2] Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.
- [3] Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- [4] Yifan Cheng. Single-particle cryo-em at crystallographic resolution. *Cell*, 161(3):450–457, 2015.
- [5] Stephen R Comeau, David W Gatchell, Sandor Vajda, and Carlos J Camacho. Cluspro: a fully automated algorithm for protein–protein docking. *Nucleic acids research*, 32(suppl_2):W96–W99, 2004.
- [6] Bowen Dai and Chris Bailey-Kellogg. Protein interaction interface region prediction by geometric deep learning. *Bioinformatics*, 37(17):2580–2588, 2021.
- [7] Guang Qiang Dong, Hao Fan, Dina Schneidman-Duhovny, Ben Webb, and Andrej Sali. Optimized atomic statistical potentials: assessment of protein interfaces and loops. *Bioinformatics*, 29(24):3158–3166, 2013.
- [8] Stephan Eismann, Raphael JL Townshend, Nathaniel Thomas, Milind Jagota, Bowen Jing, and Ron O Dror. Hierarchical, rotation-equivariant neural networks to select structural models of protein complexes. *Proteins: Structure, Function, and Bioinformatics*, 89(5):493–501, 2021.
- [9] Richard Evans, Michael O’Neill, Alexander Pritzel, Natasha Antropova, Andrew W Senior, Timothy Green, Augustin Židek, Russell Bates, Sam Blackwell, Jason Yim, et al. Protein complex prediction with alphafold-multimer. *BioRxiv*, 2021.
- [10] Saba Ferdous and Andrew CR Martin. Abdb: antibody structure database—a database of pdb-derived antibody structures. *Database*, 2018, 2018.
- [11] Jeffrey J Gray, Stewart Moughon, Chu Wang, Ora Schueler-Furman, Brian Kuhlman, Carol A Rohl, and David Baker. Protein–protein docking with simultaneous optimization of rigid-body displacement and side-chain conformations. *Journal of molecular biology*, 331(1):281–299, 2003.
- [12] Johnathan D Guest, Thom Vreven, Jing Zhou, Iain Moal, Jeliasko R Jeliaskov, Jeffrey J Gray, Zhiping Weng, and Brian G Pierce. An expanded benchmark for antibody-antigen docking and affinity prediction reveals insights into antibody recognition determinants. *Structure*, 29(6):606–621, 2021.
- [13] Howook Hwang, Thom Vreven, Joël Janin, and Zhiping Weng. Protein–protein docking benchmark version 4.0. *Proteins: Structure, Function, and Bioinformatics*, 78(15):3111–3114, 2010.
- [14] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [15] Dima Kozakov, Ryan Brenke, Stephen R Comeau, and Sandor Vajda. Piper: an fft-based protein docking program with pairwise potentials. *Proteins: Structure, Function, and Bioinformatics*, 65(2):392–406, 2006.
- [16] Marc F Lensink and Shoshana J Wodak. Blind predictions of protein interfaces by docking calculations in capri. *Proteins: Structure, Function, and Bioinformatics*, 78(15):3085–3095, 2010.

- [17] Marc F Lensink, Raúl Méndez, and Shoshana J Wodak. Docking and scoring protein complexes: Capri 3rd edition. *Proteins: Structure, Function, and Bioinformatics*, 69(4):704–718, 2007.
- [18] Marc F Lensink, Sameer Velankar, Minkyung Baek, Lim Heo, Chaok Seok, and Shoshana J Wodak. The challenge of modeling protein assemblies: the casp12-capri experiment. *Proteins: Structure, Function, and Bioinformatics*, 86:257–273, 2018.
- [19] Nir London, Barak Raveh, Dana Movshovitz-Attias, and Ora Schueler-Furman. Can self-inhibitory peptides be derived from the interfaces of globular protein–protein interactions? *Proteins: Structure, Function, and Bioinformatics*, 78(15):3140–3149, 2010.
- [20] Raúl Méndez, Raphaël Leplae, Marc F Lensink, and Shoshana J Wodak. Assessment of capri predictions in rounds 3–5 shows progress in docking procedures. *Proteins: Structure, Function, and Bioinformatics*, 60(2):150–169, 2005.
- [21] Arumay Pal, Pinak Chakrabarti, Ranjit Bahadur, Francis Rodier, and Joël Janin. Peptide segments in protein-protein interfaces. *Journal of Biosciences*, 32(1):101–111, 2007.
- [22] Brian Pierce and Zhiping Weng. Zrank: Reranking protein docking predictions with an optimized energy function. *Proteins: Structure, Function, and Bioinformatics*, 67(4):1078–1086, 2007. doi: <https://doi.org/10.1002/prot.21373>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/prot.21373>.
- [23] Nicolas Renaud, Cunliang Geng, Sonja Georgievska, Francesco Ambrosetti, Lars Ridder, Dario F Marzella, Manon F Réau, Alexandre MJJ Bonvin, and Li C Xue. DeepPrank: a deep learning framework for data mining 3d protein-protein interfaces. *Nature communications*, 12(1):1–8, 2021.
- [24] David W Ritchie and Graham JL Kemp. Protein docking using spherical polar fourier correlations. *Proteins: Structure, Function, and Bioinformatics*, 39(2):178–194, 2000.
- [25] Dina Schneidman-Duhovny, Yuval Inbar, Ruth Nussinov, and Haim J Wolfson. Patchdock and symmdock: servers for rigid and symmetric docking. *Nucleic acids research*, 33(suppl_2):W363–W367, 2005.
- [26] Amar Singh, Taras Dazhenka, Petras J Kundrotas, Michael JE Sternberg, and Ilya A Vakser. Application of docking methodologies to modeled proteins. *Proteins: Structure, Function, and Bioinformatics*, 88(9):1180–1188, 2020.
- [27] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [28] Sheng Wang, Siqi Sun, Zhen Li, Renyu Zhang, and Jinbo Xu. Accurate de novo prediction of protein contact map by ultra-deep learning model. *PLoS computational biology*, 13(1):e1005324, 2017.
- [29] Xiao Wang, Genki Terashi, Charles W Christoffer, Mengmeng Zhu, and Daisuke Kihara. Protein docking model evaluation by 3d deep convolutional neural networks. *Bioinformatics*, 36(7):2113–2118, 2020.
- [30] Xiao Wang, Sean T Flannery, and Daisuke Kihara. Protein docking model evaluation by graph neural networks. *Frontiers in Molecular Biosciences*, page 402, 2021.
- [31] Rui Yin, Brandon Y Feng, Amitabh Varshney, and Brian G Pierce. Benchmarking alphafold for protein complex modeling reveals accuracy determinants. *Protein Science*, 31(8):e4379, 2022.

A Appendix

A.1 Dataset preparation

A dataset of 1,800 antibody-antigen complexes from the AbDb was used [10]. Complexes with antigens larger than 700 or smaller than 20 amino acids were discarded, resulting in 1,214 complexes. Because this is a dataset of bound complex structures, we produced noisy unbound-like models for antibodies by modeling the antibodies via AFM, thus exposing the docking and the training to modeled structures. The complexes were divided into train and test using 97% sequence identity cut-off for antibody sequences. The final train set included 790 complexes, with 85 additional complexes that were used for validation. The complexes were divided into train and test using 97% sequence identity cut-off for antibody sequences.

We have also validated our method using 32 unbound antibody-antigen complexes from the docking Benchmark 5.0 [13] by training a model that excludes all these complexes. This set contains unbound structures that enable us to test our method using more accurate antibody structures.

A.2 Assessment criteria

Each complex model is assessed for accuracy based on root mean square deviation (RMSD) from the correct structure, as used at CAPRI [20, 17]. A docking model is considered acceptable if the ligand C RMSD after superposition of the receptors is $< 10 \text{ \AA}$ or the interface C RMSD is $< 4 \text{ \AA}$. A docking model is of medium accuracy if ligand C RMSD is $< 5 \text{ \AA}$ or interface C RMSD is $< 2 \text{ \AA}$. The success rate is the percentage of benchmark cases with at least one medium or acceptable accuracy model in the TopN predictions.

A.3 Epitope prediction assessment

Antigen interface assessment is performed using Recall and Precision criteria from CAPRI [16]. Recall is the fraction of correctly predicted complex interface residues in the model and Precision is the fraction of residues in the model interface that are the actual complex interface residues. Interface residue is a residue whose solvent accessible area is lower in the complex (or model) than in the individual components. A predicted interface is considered correct if Recall and Precision are at least 0.5. Here, we calculated the Recall and Precision for the antigen only to estimate the accuracy of epitope identification.

A.4 Training

ContactNet was trained end-to-end to predict binary classification based on the assumption that the distribution of correct and incorrect complexes is separable while optimizing the cross-entropy loss. We trained the model with the WADAM optimizer with a learning rate of $1e-4$ that decays 100 times while training using the cosine decay as a learning rate scheduler with a weight of $5e-3$ for the decay factor. The batch size contained 52 docking models, randomly selected from different complexes for better generalization. One epoch was defined as 2,000 batches with 100,000 docking models per epoch. The network was trained for 160 epochs. The major challenge in the training process was the unbalanced nature of the data. Specifically, there were hundreds of thousands of negative complexes compared to a small number of positive ones. To cope with this challenge each batch is composed of 25% of the positive, acceptable accuracy docking models and 75% of the negative, incorrect ones. We used the W&B platform for experiment tracking [3] and trained the model on a single GPU 2080RTX for 20 hours.

A.5 Sampling antibody-antigen complexes

The antibodies were docked to antigens using the antibody-antigen docking protocol of PatchDock. PatchDock is an efficient geometric rigid docking method that maximizes shape complementarity [25]. Here we used a higher sampling precision to generate a higher fraction of "positives", eg. acceptable accuracy models, for training, typically generating 100,000 docking models. These models were scored and re-ranked using SOAP-PP statistical potential [7]. For training, we used

