
Simple End-to-end Deep Learning Model for CDR-H3 Loop Structure Prediction

Natalia Zenkova

JetBrains Research

natalia.zenkova@jetbrains.com

Ekaterina Sedykh

University of Tartu

ekaterina.sedykh@ut.ee

Tatiana Shugaeva

BIOCAD

shugaeva@biocad.ru

Vladislav Strashko

BIOCAD

strashko@biocad.ru

Timofei Ermak

BIOCAD

ermak@biocad.ru

Aleksei Shpilman

HSE University

JetBrains Research

alexey@shpilman.com

Abstract

Predicting a structure of an antibody from its sequence is important since it allows for a better design process of synthetic antibodies that play a vital role in the health industry. Most of the structure of an antibody is conservative. The most variable and hard-to-predict part is the *third complementarity-determining region of the antibody heavy chain* (CDR H3). Lately, deep learning has been employed to solve the task of CDR H3 prediction. However, current state-of-the-art methods are not end-to-end, but rather they output inter-residue distances and orientations to the RosettaAntibody package that uses this additional information alongside statistical and physics-based methods to predict the 3D structure. This does not allow a fast screening process and, therefore, inhibits the development of targeted synthetic antibodies. In this work, we present an end-to-end model to predict CDR H3 loop structure, that performs on par with state-of-the-art methods in terms of accuracy but an order of magnitude faster. We also raise an issue with a commonly used RosettaAntibody benchmark that leads to data leaks, i.e., the presence of identical sequences in the train and test datasets.

1 Introduction

Antibodies play a key role in adaptive immunity. Their ability to specifically bind target molecules makes antibodies effective fighters against pathogens [18] and one of the major classes of therapeutic molecules in the field of anticancer therapy, autoimmune diseases, and more [5] [13]. Experimental determination of the tertiary structure is expensive, therefore, computational approaches are widely used nowadays [18]. Antibodies are composed of two types of protein chains: heavy and light, presented in two copies. The light chain contains variable (VL) and constant domains (CL), the heavy chain contains one variable (VH) and three constant domains (CH1, CH2, CH3). The VH-VL pair forms a variable antigen-binding Fv fragment, which in turn is subdivided into conserved framework regions (FRs) and hypervariable complementarity-determining regions (CDRs). CDRs are the most important regions for antigen binding; they form six variable loops on the surface of the Fv fragment.

Three of them are located on the light chain and are named L1, L2, L3, respectively, three others are located on the surface of the heavy chain and are designated as H1, H2, and H3 [10].

Loops L1-L3 and H1-H2 typically fall into a limited set of conformations, therefore, it is relatively easy to determine their 3D configuration [2]. In contrast, the third loop of the heavy chain is the most variable and cannot be described using the limited set of structural clusters [11]. Its high variability does not allow the application of modern protein folding techniques based on multiple sequence alignment (MSA) of homologous sequences, such as AlphaFold2 [9] and RoseTTAFold [1]. Thus, CDR H3 structure prediction requires special approaches.

Recently, in addition to statistical and physics-based [16] prediction methods, deep learning models have been employed for the task of H3 structure prediction. Current state-of-the-art solution DeepH3 uses deep learning to calculate the distribution of inter-residue distances and orientations, which are then passed to the RosettaAntibody package that employs both physics-based and statistical methods. Because of that second step and the causing time constraints, this method can not be feasibly used to perform a wide *in silico* screening of possible antibody structures.

In this paper, we present a simple end-to-end deep learning method based on ELMo embedding and forward and backward LSTM passes that directly output coordinates of backbone atoms of amino acids in the H3 loop. This method, that we named **SimpleDH3**, performs on on par with DeepH3 in terms of root mean square deviation (RMSD) but is at least an order of magnitude faster.

2 Model

Our SimpleDH3 model pipeline consists of the following steps:

1. Convert each amino acid from the input sequence into a 1024-dimensional embedding representation using a context-dependent model ELMo (Embeddings from Language Models) [12] pretrained on the UniRef50 dataset [17] from [6];
2. Independent forward and backward LSTM [7] pass. LSTM cell with 256 neurons in the hidden and the output layer receives an amino acid embedding from ELMo model concatenated with the prediction coordinates for the previous amino acid.
3. Independent feedforward neural networks with 9-dimensional output after LSTM cell. 9 dimensions correspond to the coordinates of C , C_α and N atoms in a single amino acid that completely determine its position and orientation.
4. Averaging the output from forward and backward passes.




Figure 1: Architecture of our simple model (SimpleDH3). It takes the sequence for CDR H3 and 2 additional amino acids: one on the each end. Each amino acid from a sequence is transformed with pretrained ELMo model from [6]. Afterwards, the embeddings are concatenated with the prediction coordinates for the previous amino acid are passed to the backward LSTM cell pass combined with linear layer. Then independently forward LSTM cell pass combined with its own linear layer. As a result, the model outputs backbone coordinates, i.e., for C , C_α , and N atoms, averaged from both passes.

It is important to note that SimpleDH3 model is trained only on CDR H3 loops with two amino acids at the ends. In contrast to DeepH3 [14] and other methods [15] which take Fv fragment, our model use only loop sequences.

The model was trained with the following loss function:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{coords} + \lambda_2 \mathcal{L}_{dist} + \lambda_3 \mathcal{L}_{torsion},$$

where

$$\begin{aligned} \mathcal{L}_{coords} &= MSE_{coords}, \quad \mathcal{L}_{dist} = MSE_{dist}, \\ \mathcal{L}_{torsion} &= MSE_{torsion}, \quad \lambda_1 = 0.2, \quad \lambda_2 = 3.0, \quad \lambda_3 = 10.0. \end{aligned}$$

\mathcal{L}_{coords} is a mean-square deviation of C , C_α and N atoms coordinates; \mathcal{L}_{dist} is a mean-square deviation of distance between two consecutive atoms, such as $(C; C_\alpha)$, $(C_\alpha; N)$ and $(N; C)$; $\mathcal{L}_{torsion}$ is similar to other components of the loss function, but considers torsion angles cosines instead.

3 Experiments and Results

3.1 Data collection and preprocessing

The training and testing datasets were retrieved from the SABDab database [4]. Only the structures with resolution better than 3.0\AA threshold are selected and all sequences are annotated by Chothia numbering scheme [3]. Then Fv fragments are extracted using the annotation information. To make the dataset non-redundant, agglomerative clustering is applied [8] (with the complete linkage rule) with $1 - \textit{sequence identity}$ as a distance metric (for two chain antibodies, the distance metric is the average mean of $1 - \textit{sequence identity}$ for heavy and light chains). We break down two-chain (VH-VL) and single-chain (VHH) structures into different bins and clusterize them separately. Next, we determine 100 farthest clusters, i.e., the ones with the maximal sum of the distances to all other clusters, from the two-chain and 50 farthest ones from the single-chain bin with a greedy approach and marked them as the test dataset. The rest of the clusters are the training set. Then, from each cluster we take only the center, which is determined by RMSD among the cluster members, considering the center as the most representative structure of the cluster. Based on the annotation, CDR H3 loop structures are extracted and then only those structures are used for training and testing. Afterwards, we filter out CDR H3 structures containing gaps.

We take +1 additional amino acids before the loop and +1 after. We then shift the coordinates so that the first atom’s coordinates are $(0, 0, 0)$ and the last one is placed on the Y-axis. It is important to note that DeepH3 is trained on a full Fv fragment and our algorithm, in that way, requires less data to train.

3.2 Train and test split

The model was trained on 1970 structures from data described in subsection 3.1.

To compare our results with other methods, we have tested the network on the commonly used RosettaAntibody benchmark [19], containing 49 structures with CDR H3 loop lengths ranging from 7 to 17 residues. First, we have deleted structures with these 49 IDs from our dataset. However, we have quickly identified identical CDR H3 loop sequences with different IDs in our training data. The structures are different if compared by the whole Fv-fragment sequences, but if we only measure RMSD on the CDR H3 loop, RMSD on the test dataset will be wrongfully better due to a problem, known as data leaks. Therefore, we eliminated them so there are only unique sequences in all train and test sets. After this step, the size of our train dataset has decreased to 1945 structures. To make a fair comparison of our method to others we do it on the original RosettaAntibody benchmark. However, to make a fair evaluation, we also present the result on the data without duplicate sequences.

3.3 Training setup

SimpleDH3 model was trained using NVIDIA Tesla K80 GPU for 500 epochs (2 hours) until convergence on the antibody dataset, described in Section 3.1 with batch size of 64 and a learning rate of 0.001.

3.4 Results

Table 1 shows the results of our model in terms of mean RMSD and mean inference time per sequence in seconds.

As we can see from the results, our simple model performs on par with current state-of-the-art DeepH3 model, but dramatically outperforms it in terms of inference time. Moreover, the time results for DeepH3 [14] only show the time it takes to run the network part of their method, but, as Rosetta package is needed to get final results, the speed benefit will only increase.

Methods	Mean RMSD	Mean time per sequence
DeepH3 [14]	2.2	6.65
SimpleDH3 (ours)	2.04	0.59
SimpleDH3 (ours, unique)	2.77	0.59

Table 1: Performance of CDR H3 loops prediction on RosettaAntibody benchmark for both non-unique (original) and unique data.

We also show the dependence of RMSD on the sequence length for both non-unique (original) and unique data (Figure 2). RMSD is measured only on target H3 loop sequence, not taking into account additional amino acids used in training and inference.




Figure 2: SimpleDH3 RMSD performance across various H3 loop sequence lengths.

4 Conclusion

In this paper, we present a novel end-to-end deep learning method for the prediction of CDR H3 loop structure. By utilizing pretrained ELMo model and independent forward and backward LSTM passes, we can directly output the coordinates of atoms in the backbone of the amino acid sequence by only training on loop sequences and not whole Fv regions.

We demonstrate that SimpleDH3 approach performs on par with state-of-the-art methods in terms of RMSD, but is at least an order of magnitude faster. This brings us closer to the possibility of large-scale *in silico* screening for a specific antibody for a targeted therapeutic target.

We also point out a flaw in the current benchmark for this task regarding identical sequences with different PDB-IDs and present a method for creating a better benchmark.

Our code is available at <https://github.com/jbr-ai-labs/SimpleDH3>.

References

- [1] M. Baek, F. DiMaio, I. Anishchenko, J. Dauparas, S. Ovchinnikov, G. R. Lee, J. Wang, Q. Cong, L. N. Kinch, R. D. Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.
- [2] C. Chothia and A. M. Lesk. Canonical structures for the hypervariable regions of immunoglobulins. *Journal of molecular biology*, 196(4):901–917, 1987.
- [3] C. Chothia, A. M. Lesk, A. Tramontano, M. Levitt, S. J. Smith-Gill, G. Air, S. Sheriff, E. A. Padlan, D. Davies, W. R. Tulip, et al. Conformations of immunoglobulin hypervariable regions. *Nature*, 342(6252):877–883, 1989.
- [4] 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–D1146, 11 2013. ISSN 0305-1048. doi: 10.1093/nar/gkt1043. URL <https://doi.org/10.1093/nar/gkt1043>.
- [5] D. M. Ecker, S. D. Jones, and H. L. Levine. The therapeutic monoclonal antibody market. In *MAbs*, volume 7, pages 9–14. Taylor & Francis, 2015.
- [6] M. Heinzinger, A. Elnaggar, Y. Wang, C. Dallago, D. Nechaev, F. Matthes, and B. Rost. Modeling aspects of the language of life through transfer-learning protein sequences. *BMC bioinformatics*, 20(1):1–17, 2019.
- [7] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [8] A. K. Jain and R. C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., 1988.
- [9] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [10] C. Marks and C. Deane. Antibody h3 structure prediction. *Computational and Structural Biotechnology Journal*, 15:222–231, 2017.
- [11] B. North, A. Lehmann, and R. L. Dunbrack Jr. A new clustering of antibody cdr loop conformations. *Journal of molecular biology*, 406(2):228–256, 2011.
- [12] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1202. URL <https://aclanthology.org/N18-1202>.
- [13] J. M. Reichert. Antibodies to watch in 2017. In *MAbs*, volume 9, pages 167–181. Taylor & Francis, 2017.
- [14] J. A. Ruffolo, C. Guerra, S. P. Mahajan, J. Sulam, and J. J. Gray. Geometric potentials from deep learning improve prediction of cdr h3 loop structures. *Bioinformatics*, 36(Supplement_1): i268–i275, 2020.
- [15] J. A. Ruffolo, J. Sulam, and J. J. Gray. Antibody structure prediction using interpretable deep learning. *bioRxiv*, 2021.
- [16] A. Sivasubramanian, A. Sircar, S. Chaudhury, and J. J. Gray. Toward high-resolution homology modeling of antibody fv regions and application to antibody–antigen docking. *Proteins: Structure, Function, and Bioinformatics*, 74(2):497–514, 2009.
- [17] B. E. Suzek, Y. Wang, H. Huang, P. B. McGarvey, C. H. Wu, and U. Consortium. Uniref clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics*, 31(6):926–932, 2015.

- [18] K. E. Tiller and P. M. Tessier. Advances in antibody design. *Annual review of biomedical engineering*, 17:191–216, 2015.
- [19] B. D. Weitzner, D. Kuroda, N. Marze, J. Xu, and J. J. Gray. Blind prediction performance of rosettaantibody 3.0: grafting, relaxation, kinematic loop modeling, and full cdr optimization. *Proteins: Structure, Function, and Bioinformatics*, 82(8):1611–1623, 2014.