



A comparative study of protein structure prediction tools for challenging targets

Snake venom toxins

Kalogeropoulos, Konstantinos; Bohn, Markus-Frederik; Jenkins, David E.; Ledergerber, Jann; Sørensen, Christoffer V.; Hofmann, Nils; Wade, Jack; Fryer, Thomas; Nguyen, Giang Thi Tuyet; auf dem Keller, Ullrich

Total number of authors:
12

Published in:
Toxicon

Link to article, DOI:
[10.1016/j.toxicon.2023.107559](https://doi.org/10.1016/j.toxicon.2023.107559)

Publication date:
2024

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Kalogeropoulos, K., Bohn, M-F., Jenkins, D. E., Ledergerber, J., Sørensen, C. V., Hofmann, N., Wade, J., Fryer, T., Nguyen, G. T. T., auf dem Keller, U., Laustsen, A. H., & Jenkins, T. P. (2024). A comparative study of protein structure prediction tools for challenging targets: Snake venom toxins. *Toxicon*, 238, Article 107559.
<https://doi.org/10.1016/j.toxicon.2023.107559>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



A comparative study of protein structure prediction tools for challenging targets: Snake venom toxins

Konstantinos Kalogeropoulos^a, Markus-Frederik Bohn^a, David E. Jenkins^b, Jann Ledergerber^{a,c}, Christoffer V. Sørensen^a, Nils Hofmann^a, Jack Wade^a, Thomas Fryer^a, Giang Thi Tuyet Nguyen^a, Ulrich auf dem Keller^a, Andreas H. Laustsen^a, Timothy P. Jenkins^{a,*}

^a Department of Biotechnology and Biomedicine, Technical University of Denmark, Kongens Lyngby, Denmark

^b BettercallPaul, Munich, Germany

^c Department of Chemistry and Applied Bioscience, ETH Zurich, Zurich, Switzerland

ARTICLE INFO

Handling editor: Ray Norton

Keywords:

Protein structure prediction
Structural modelling
AlphaFold2
ColabFold
Modeller
Toxins
Challenging targets

ABSTRACT

Protein structure determination is a critical aspect of biological research, enabling us to understand protein function and potential applications. Recent advances in deep learning and artificial intelligence have led to the development of several protein structure prediction tools, such as AlphaFold2 and ColabFold. However, their performance has primarily been evaluated on well-characterised proteins and their ability to predict structures of proteins lacking experimental structures, such as many snake venom toxins, has been less scrutinised. In this study, we evaluated three modelling tools on their prediction of over 1000 snake venom toxin structures for which no experimental structures exist. Our findings show that AlphaFold2 (AF2) performed the best across all assessed parameters. We also observed that ColabFold (CF) only scored slightly worse than AF2, while being computationally less intensive. All tools struggled with regions of intrinsic disorder, such as loops and propeptide regions, and performed well in predicting the structure of functional domains. Overall, our study highlights the importance of exercising caution when working with proteins with no experimental structures available, particularly those that are large and contain flexible regions. Nonetheless, leveraging computational structure prediction tools can provide valuable insights into the modelling of protein interactions with different targets and reveal potential binding sites, active sites, and conformational changes, as well as into the design of potential molecular binders for reagent, diagnostic, or therapeutic purposes.

1. Introduction

Understanding 3-dimensional (3D) protein structures is key to many research questions, ranging from fundamental topics concerning how a given protein functions to translational hurdles involving using such information to manipulate protein function for industrial, therapeutic or other purposes. The most reliable approach towards resolving protein structures has been the use of experimental technologies, such as X-ray crystallography or nuclear magnetic resonance (NMR), and more recently, cryogenic electron microscopy. Yet, such approaches are time-consuming, low-throughput, and sometimes impossible for difficult to crystallise targets such as membrane proteins (Kermani, 2021). In an attempt to increase throughput and allow for 3D characterisation of large structural datasets and produce structures of proteins where

conventional approaches fail, scientists have explored computational methods for the prediction of structures based on evolutionary history. Homology modelling aims to predict the 3D structure of a given protein (target) sequence based on its homology to solved structures (templates) (Thompson et al., 2020; Bai et al., 2015) and pairwise evolutionary correlations (Jones et al., 2012; Marks et al., 2011; Weigt et al., 2009; Shindyalov et al., 1994; Altschuh et al., 1987). Homology modelling has been used to produce structure models for at least one domain in more than half of all known sequences and a total of over 38 million models deposited on ModBase (Pieper et al., 2014). The majority of these deposited structures were generated using the software program MODELLER (MDLR), an established and excellent protein modelling structure tool (Bitencourt-Ferreira and de Azevedo, 2019; Webb and Sali, 2016). This approach has aided in our understanding of protein structures and

* Corresponding author. Soltofts Plads, 2800, Kgs. Lyngby, Denmark.

E-mail address: tpaje@dtu.dk (T.P. Jenkins).

<https://doi.org/10.1016/j.toxicon.2023.107559>

Received 11 October 2023; Received in revised form 6 December 2023; Accepted 8 December 2023

Available online 17 December 2023

0041-0101/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

has grown in accuracy with the increasing number of experimental structures deposited in the Protein Data Bank (PDB) (Protein Data Bank, 2019), the rise of genomic sequencing, and the availability of deep learning techniques allowing rapid interpretation of these data. Nevertheless, more often than not, current evolutionary-history-based approaches fall short of generating predictions comparable to experimental accuracy, especially for proteins lacking closely related and experimentally resolved homologues. This has curtailed the utility of homology protein structure modelling for many biological applications to date.

Recently, protein structure prediction has, however, undergone a renaissance with the application of sophisticated machine learning approaches. Whilst many of these algorithms are still reliant on template proteins and multiple sequence alignment (MSA), the shift from a rational decision tree to transformer-based neural networks has seen a substantial improvement in prediction accuracy (Jumper et al., 2021). Indeed, the first protein structure prediction tool using transformers, i.e. AlphaFold2 (AF2) (Jumper et al., 2021), achieved the highest accuracy prediction at the Critical Assessment of protein Structure Prediction (CASP) competition (CASP14) (Pereira et al., 2021). Its accuracy was comparable to experimental protein structure determination with 36% of their submitted protein targets having a root-mean-square deviation (RMSD) under 2 Å (generally considered to be a solved structure), and 86% under 5 Å, with a total mean of 3.8 Å, which presented an impressive performance comparable to experimental accuracies. Since then, AF2 predicted structures for the near-whole proteome of 48 species with over 200 million entries (<https://alphafold.ebi.ac.uk/>) have been made publicly available, and the transformer-based approach of AF2 has also been independently reproduced in another tool, RoseTTAFold (Baek et al., 2021).

Whilst the prediction precision of these tools is unprecedented, the computational power required to leverage them is substantial. Primarily, building the MSAs is computationally intensive (multiple hours and >2 TB of storage per protein) and involves sensitive homology searches via HMMER (Eddy, 2011) and HHblits (Steinegger et al., 2019). Further, running the deep neural networks for the modelling itself also requires computational power and memory, albeit negligible compared to building the MSAs. With this need for substantial computation in mind, researchers developed ColabFold (CF) (Mirdita et al., 2021), which harnesses Google Colaboratory and thus provides free access to powerful graphics processing units (GPUs). CF also accelerates predictions (20–30 times faster than AF2) by using MMseqs2 search (Mirdita et al., 2019; Steinegger et al., 2017) instead of AF2's native input feature generation. Further, CF also leverages optimisation strategies for predictions of multiple structures by avoiding recompilation and adding early stop criteria. Still, the question remains of what approach to rely on, if any, for protein structure prediction. The reliance of MDLR, AF2, and CF on homology alignments suggests that whilst all of these approaches excel in well-characterised areas of biology, they might struggle when few high quality templates exist. This is particularly relevant for snake venom toxins. Here, 19,000–25,000 toxins are predicted to exist (Laustsen et al., 2016), but only around 2,000 different proteins have been described, and an even smaller percentage (<10%) of their structures have been experimentally resolved. Yet, understanding venom toxin structures could carry many benefits in either harnessing their beneficial potential as therapeutics (Mohamed Abd El-Aziz et al., 2019; Li et al., 2018; Ferraz et al., 2019) or for developing better treatments for snakebite envenomings (Gutiérrez et al., 2017; Knudsen et al., 2018; Jenkins et al., 2019).

Thus, to investigate the performance of three commonly used protein structure prediction tools on less characterised and potentially challenging protein targets, we predicted the structures of 1,062 snake venom toxins using MDLR and CF and compared them to each other and to AF2 predicted structures.

2. Materials and methods

2.1. Retrieval of snake venom toxin sequences

To retrieve the sequences of all published snake venom toxins (Nov. 2020) belonging to potentially medically relevant toxin families, we used the utilities offered by VenomZone (VenomZone). We selected C-type lectins (CTLs), disintegrins (DISs), kunitz-type serine protease inhibitors (KUNs), phospholipase A₂s (PLA₂s), snake venom metalloproteases (SVMPs), snake venom serine proteases (SVSPs), and three-finger toxins (3FTxs), and searched for entries on UniProt (e.g., taxonomy:serpentes family:"phospholipase a2 family" (annotation: (type:"tissue specificity" venom) OR locations: (location:nematocyst)) AND reviewed:yes). The toxin UniProt entry information was retrieved with custom Python scripts using Biopython's (Cock et al., 2009) ExPASy package (Gasteiger et al., 2003).

2.2. Identification of experimentally resolved toxin structures

The protein database (PDB) information on those entries was examined by cross-reference documentation in the obtained UniProt files. Where there were multiple annotated PDB entries for a UniProt accession, the entry with the highest resolution and chains corresponding to the toxin was used. This was achieved simply by parsing UniProt files using custom scripts. Preference was given to X-ray structures. Thereafter, for a given toxin, the selected PDB file containing structural information for that entry was retrieved using pdb-tools (Rodrigues et al., 2018). Using the same tool, files obtained were filtered, selecting for chains corresponding to the toxin, and removing ligands and hydrogens in the structure. This process was also automated via bash scripting.

2.3. Generating modeller structures

For the remaining toxins that did not contain an annotated structure, the theoretical structure was predicted using homology modelling algorithms. To that purpose, we employed the Python MODELLER software (Webb and Sali, 2016). Toxins below the threshold length of 30 amino acids were discarded. Thereafter, the "multiple template modelling" approach and "loop refinement" were used with the PDB_95 database for template search. These methods significantly improve modelling quality, especially of complex regions, such as loops (Webb and Sali, 2016). Predictions were automated and run using multiple threading with subprocesses. For each toxin sequence, a total of 10 models were generated, and the best model, based on the lowest discrete optimised protein energy (DOPE) score, was selected as the representative model. The DOPE score (Shen and Sali, 2006) is based on an improved reference state that corresponds to non-interacting atoms in a homogeneous sphere with the radius dependent on a sample native structure; it thus accounts for the finite and spherical shape of the native structures. It is used to assess the energy of the protein model generated through many iterations by MODELLER. Finally, we evaluated the DOPE score of all of our toxin models to ensure the overall quality was sufficient to merit their usage for further research.

2.4. Generating ColabFold structures

A version of ColabFold/MMseqs2 (Steinegger et al., 2019) was modified to allow batch generation of structures. The program, originally designed to process single proteins via an online interface, was restructured in order to process multiple sequences from file input. Functions were rewritten to eliminate use of global variables, in order to permit loop processing. A mechanism was added to allow processing to be interrupted and restarted without repeating previously generated structures. This was essential to allow use of cheap compute facilities such as Google Colab and Colab Pro, which reserve the right to interrupt

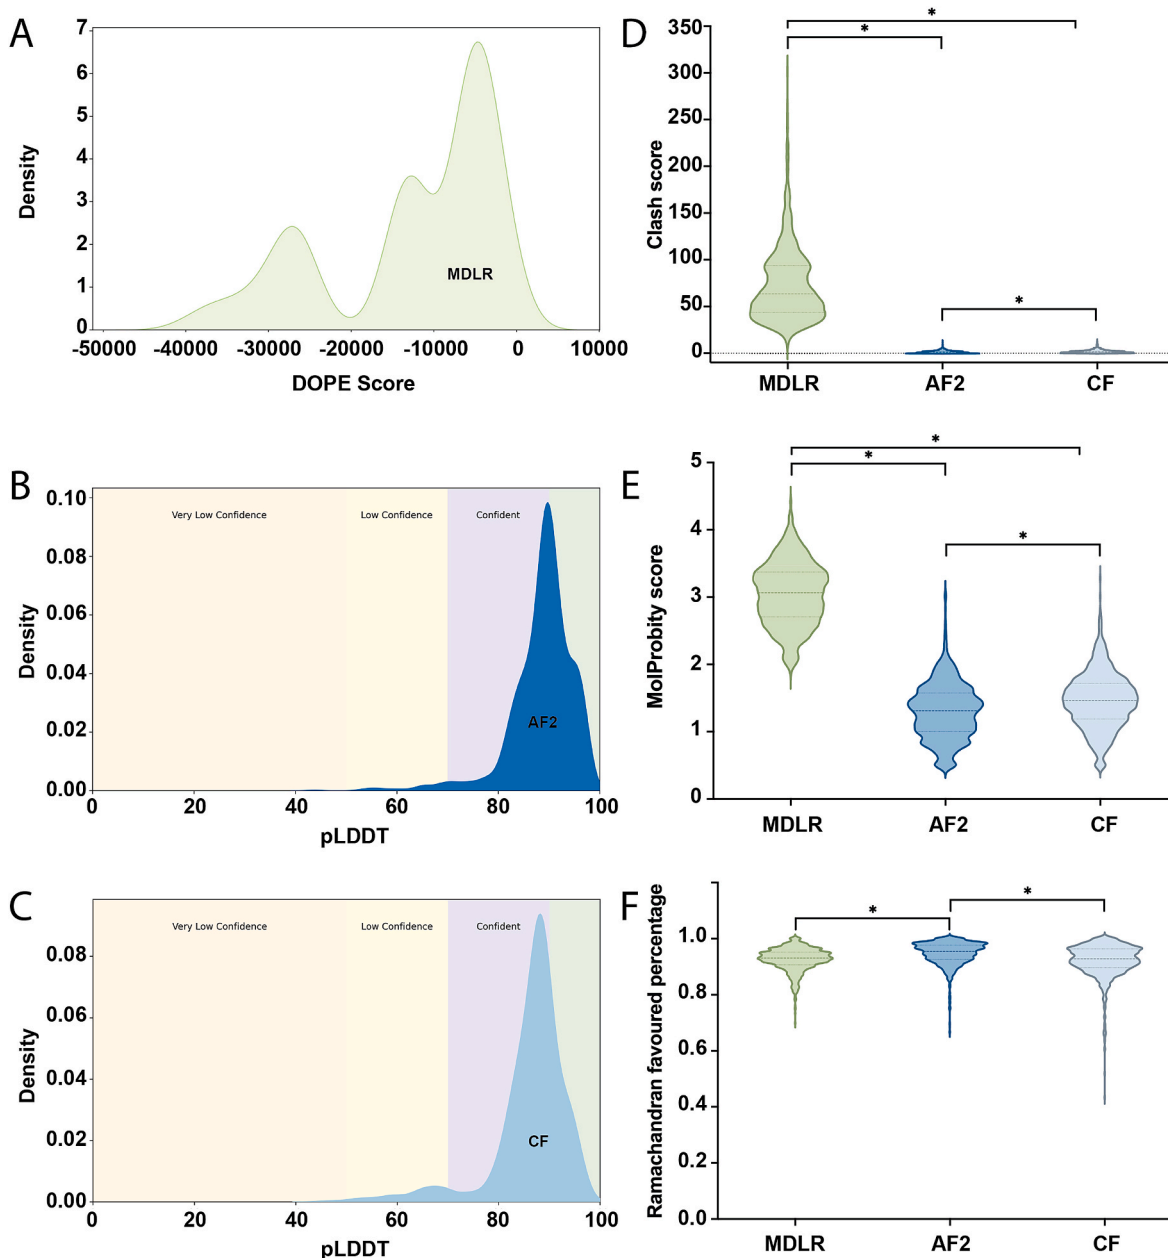


Fig. 1. Quality evaluation of the toxin structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF). Kernel Density Estimate of the distribution of DOPE scores for Modeller (Panel A), per-residue pLDDT values for AF2 (Panel B), and CF (Panel C). For Panels B and C, the pLDDT values are averaged for each model and categorized as Very Low Confidence (pLDDT < 50, orange), Low Confidence (50 ≤ pLDDT < 70, yellow), Confident (70 ≤ pLDDT < 90, blue), and Very Confident (pLDDT ≥ 90, green). The shading under the curves in all panels indicates the relative density of models falling within each quality range or score. D) Clash scores, i.e. the number of serious clashes per 1000 atoms, defined as all non-donor-acceptor atoms overlapping by more than 0.4 Å. E) MolProbity scores (scores measured in percentiles; percentile ≥ 66 being the best); Ramachandran outliers (scores from 0 to 1; 1 being the best); F) Ramachandran favoured percentage (scores from 0 to 1; 1 being the best). Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

long-running jobs. This mechanism also allowed processing to be shared across multiple compute instances. The feature dictionary and all other parameters for reproducing the structure generation of any particular structure (and generation errors, if any) were written to sequence-specific files during batch processing. The number of CF models calculated for each residue was set to the maximum, five; multi-sequence alignment was set; environmental processing was set; Amber relaxation was enabled; use of templates was enabled; and homooligomer processing was disabled. Query sequences of less than 20 residues or more than 1400 residues were skipped. Amber relaxation does not work for incomplete sequences with undetermined residues, so changes to the program were made to recognise these cases

automatically and to invoke ColabFold without relaxation. The FASTA file processing was based on code from brentp (github) and Minkyung Baek's modification to allow Colabfold to process complexes, as used in MMSeqs2, was retained (Baek and Baker, 2022).

2.5. Retrieving AlphaFold2 structures

To retrieve relevant toxin structures predicted by AF2, we downloaded the latest database (<https://alphafold.ebi.ac.uk/download>) and selected the same list of UniProtIDs as were used in our MDLR and CF predictions.

2.6. Trimming of propeptides

To enhance the accuracy and efficiency of our protein structure prediction we developed a code to detect and trim the signal peptides since these are not part of the mature protein or its function. The code is available via GitHub (https://github.com/nilshof01/signal_cleaver). Before starting the main script it is necessary to download the SignalP6 package. For the processing of the UniProtIDs and predicting potential signal peptides, we developed a Python3.9 script to retrieve the sequences by giving a list of IDs as input. Next, the protein sequences are analysed using SignalP6, a state-of-the-art software capable of predicting the presence and location of all five signal peptide cleavage sites in protein sequences. The software generates a GFF (General Feature Format) file for each ID in a separate folder, which contains the start-, end position and score of the predicted signal peptide regions (Teufel et al., 2022). A further Python script was developed to parse over the GFF files and extract this information for each ID in a pandas dataframe which was subsequently used to trim the sequences subsequently. Finally, the actual N-terminal trimming of the PDB file was performed using phenix.pdbtools (Liebschner et al., 2019). All structures have been deposited in Mendeley data (<https://doi.org/10.17632/gjk47cjm26.1>).

2.7. Quality control and comparisons of structures

To assess the quality and validate the toxin structures predicted by MDLR, AF2, and CF, we used several different approaches. Both AlphaFold2 and ColabFold provide their own measure of predicted model confidence as a per-residue predicted local distance difference test (pLDDT) score. MDLR assesses model quality using the discrete optimised protein energy (DOPE) method and provides a score for each model. These intrinsic accuracy measures were also considered in our evaluation of the models produced by each tool. The distribution of model-averaged pLDDT and DOPE scores was visualised using Kernel Density Estimate (KDE) plots. Plots were generated using seaborn and matplotlib in Python, with specific color-coded regions to denote varying confidence levels of the pLDDT values. We evaluated the Ramachandran scores (outliers and favoured percentage), which indicate the number of amino acid residues with poor or favoured ϕ/ψ (Phi/Psi) angles. Phi/Psi angles are the dihedral angles in the protein backbone. Only certain angles are typically found in proteins. Finally, we compared our models' MolProbity scores (Davis et al., 2007). MolProbity is a structure-validation web service which uses a weighted function of clashes, Ramachandran favoured, and rotamer outliers, scaled and normalised so that its value approximates the resolution at which that score would be average. To assess global and local differences between models generated by each of the three tools, we performed pairwise structural alignment of each toxin model using the superimpose method (Shindyalov and Bourne, 1998) from the Bio.PDB module in the Biopython toolbox and ranked aligned pairs by RMSD, separated by toxin family. Cases where RMSD was highest/lowest for each toxin family were inspected manually to identify areas of structural disagreement. Statistical analysis was performed as outlined above. To establish statistical significance the dataset was analysed via Prism (v.9.5.1). Once a lack of Gaussian distribution was established via both the Shapiro-Wilk and Kolmogorov-Smirnov test, the Wilcoxon matched-pairs signed rank test was selected as appropriate non-parametric analysis to conduct pairwise comparisons of AF2, CF, and MDLR.

3. Results

3.1. Generation, retrieval, and validation of toxin structures

For this study, we retrieved 1062 snake venom toxin sequences from UniProt, including 220 C-type lectins (CTLs), 82 disintegrins (DISs), 145 kunitz-type serine protease inhibitors (KUNs), 190 phospholipase A₂s (PLA₂s), 135 snake venom metalloproteases (SVMPs), 147 snake venom

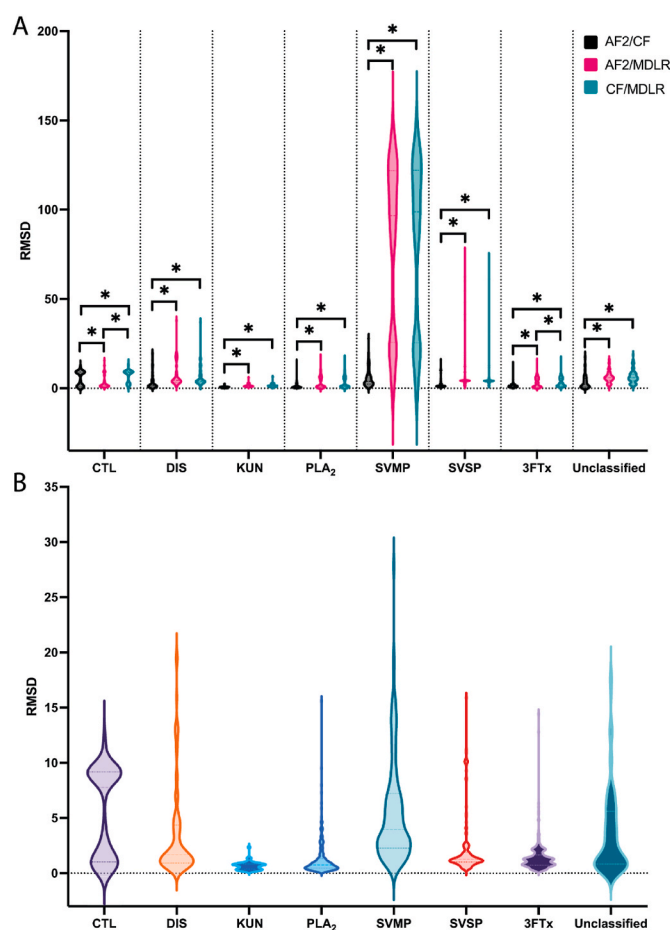


Fig. 2. Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF) and across the following toxin families: C-type lectins (CTLs), disintegrins (DISs), kunitz-type serine protease inhibitors (KUNs), phospholipase A₂s (PLA₂s), snake venom metalloproteases (SVMPs), snake venom serine proteases (SVSPs), and three-finger toxins (3FTxs). A) This resulted in three comparisons, i.e. AF2/CF, AF2/MDLR, and CF/MDLR. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$). B) Differences between AF2 and CF only.

serine proteases (SVSPs), and 274 three-finger toxins (3FTxs). Structures were generated for all of these sequences using MDLR, as well as CF. The respective AF2 structures were retrieved from the database. For any structures that had a propeptide, it was trimmed to allow for an equal comparison across tools. Notably, all tools appeared to model these propeptides very poorly. The 3186 trimmed structures were evaluated based on their individual confidence scoring, as well as via a series of parameters, including Clash scores, MolProbity scores, and Ramachandran favoured percentage. AF2 was significantly better than both MDLR and CF across all scores. MDLR performed the worst, though it was only slightly worse than CF in both Ramachandran evaluations (Fig. 1; Table S1). Notably, MDLR also performed worse on Clash and MolProbity scores, the more amino acid residues a given structure had (Fig. S2).

3.2. Comparing all modeller, ColabFold, and AlphaFold2 structures

To understand the differences in the structure prediction of the three modelling tools, all of the toxin structures underwent a three-way comparison. The largest variation in RMSD between models was observed in SVMPs, whereas the smallest was found to be for KUNs (Fig. 2; Tables S3 and S4). RMSD was significantly different across all toxin families between AF2/MDLR and CF/MDLR, when compared to




Fig. 3. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (P60774) and (B) least (Q10749) overlap across the entire dataset. Peptidase domain used for alignment (B) highlighted (red).




Fig. 4. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (D8VNS6) and (B) least (Q6X5S5) overlap across C-type lectins. (C) The largest difference between AF2 and CF predictions (A7X3W1). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

AF2/CF. Meanwhile, differences between AF2/MDLR and CF/MDLR were insignificant for all toxin families besides CTLs and 3FTxs. This indicated that AF2 and CF models exhibited a greater degree of similarity to each other than to structures generated by MDLR. The largest differences between AF2 and CF models were found within the CTL and SVMP families.

Across all tools, it was found that the mean difference in RMSD was

14.80 Å with a standard deviation of 30.94 Å between MDLR and CF, 14.44 ± 30.97 Å between MDLR and AF2, and 2.77 ± 0.08 Å between AF2 and CF (Table S3). The highest similarity across all three was observed between models of toxins sharing the classic three-finger toxin fold, such as in P60774 from *Naja samarensis*, with an RMSD of 0.4 Å between MDLR and AF2 and no appreciable differences (Fig. 3A). On the other hand, the largest difference observed by RMSD (145 Å) was




Fig. 5. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (P81631) and (B) least (P0DJ43) overlap across disintegrins. (C) The largest difference between AF2 and CF predictions (P0DJ43). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

between the metalloprotease Q10749 model generated by MDLR and either model generated by CF or AF2 (*Naja mossambica*). Modeller and AlphaFold agree only on the core of the peptidase domain, with AF2/CF identifying additional N- and C-terminal domains, which remain unstructured in the MDLR model (RMSD 145 Å between MDLR and both AF2 and CF; 3 Å between AF2 and CF). Restricting the alignment to the peptidase domain leads to a much better fit between models (RMSD 0.4/0.42 Å between MDLR and both AF2 and CF; 0.29 Å between AF2 and CF; Fig. 3B).

3.3. Toxin family specific patterns

Snake venom toxins comprise a plethora of protein families substantially varying in size and structural complexity (Tasoulis and Isbister, 2017). Therefore, we also explored the predictions across all represented protein families.

3.3.1. C-type lectins (CTLs)

Snake venom CTLs are 10–30 kDa glycoproteins that contain conserved carbohydrate recognition domains and can bind to specific sugar residues, resulting in various biological effects (Oliveira et al., 2022). Overall, we observed a mean RMSD difference of 5.2 ± 2.6 Å between all CTL models. The highest similarity was found to be between D8VNS6 from *Cerberus rynchops* (RMSD 0.4 Å between AF2 and CF, 0.9/08 Å between Modeller and AF2/CF) (Fig. 4A). The largest difference was between Q6X5S5 from *Echis ocellatus* (8.2 Å between AF2 and CF, as well as 15.7 Å and 13.6 Å differences between Modeller and AF2/CF, respectively; Fig. 4B). The largest difference between AF2 and

CF predictions was found to be for A7X3W1 from *Pseudoferania polylepsis* (12.6 Å; Fig. 4C). RMSD was significantly different between all three comparisons, with AF2/MDLR having the lowest average RMSD (3.3 Å) and CF/MDLR the highest (6.9 Å; Fig. 4D). Though not assessed in this study, we would expect to observe similar structural patterns in C-type lectin-like (SNACLECS) toxins; specifically that they are more on of the toxin families with higher levels of inaccuracies in their computational structure predictions.

3.3.2. Disintegrins (DISs)

Snake venom disintegrins, with few exceptions, derive from the proteolytic processing of PII SVMPs (Oliveira et al., 2022). Overall, we observed a mean RMSD difference of 7.0 ± 6.3 Å between MDLR and AF2 models. For example, the disintegrin fold of EC3B from *Echis carinatus* is predicted under good agreement (RMSD 4.9 Å between AF2 and CF, as well as 1.1/1.0 Å between MDLR and AF2/CF Fig. 5A), with all solutions showing the same conserved four disulfide bonds expected from the heterodimeric disintegrin fold family. Here, most of the differences arise from the terminal peptidic regions that are not part of the disulfide-linked core. The largest differences were found between P0DJ43 models (RMSD 38/37 Å between MDLR and AF2/CF, 20 Å between AF2 and CF) from *Micropechis ikaheka* (Fig. 5B). There is very little agreement across the three different models. Notably, the only section scored by AlphaFold to have “very high confidence” (pLDDT >90) is a <20 amino acid residue segment within the disintegrin domain. All three tools were able to model the predicted disulfide bonding pattern. The largest difference between AF2 and CF predictions was found to be for P0DJ43 (19.7 Å; Fig. 5C). RMSD was significantly different between




Fig. 6. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (C11C51) and (B) least (H6VC05) overlap across kunitz-type serine protease inhibitors. (C) The largest difference between AF2 and CF predictions (POCAR0). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

AF2/CF and both AF2/MDLR and CF/MDLR, with AF2/CF having the lowest average RMSD (3.6 Å) and AF2/MDLR the highest (7 Å; Fig. 5D).

3.3.3. Kunitz-type serine protease inhibitors (KUNs)

KUNs are 6–7 kDa small proteins containing three unique disulfide bonds that can inhibit the proteolytic activities of serine proteases and also include the highly neurotoxic dendrotoxins (Oliveira et al., 2022). Overall, we observed a mean RMSD difference of 2.8 ± 3.3 Å between all KUN MDLR and AF2 models. The highest similarity is between C11C51 from *Walterinnesia aegyptia* (RMSD 0.2 Å between AF2 and CF, 0.6 Å between Modeller and AF2/CF) (Fig. 4A). The largest difference driven by the terminal loop regions is between H6VC05 from *Daboia russelii* (2.0 Å between AF2 and CF, as well as 6.0 Å and 6.4 Å differences between Modeller and AF2/CF, respectively; Fig. 4B). The largest difference between AF2 and CF predictions was found to be for POCAR0 (2.5 Å; Fig. 4C). RMSD was significantly different between AF2/CF and both AF2/MDLR and CF/MDLR, with AF2/CF having the lowest average RMSD (0.7 Å) and AF2/MDLR as well as CF/MDLR shared highest (1.6 Å; Fig. 6D).

3.3.4. Phospholipase A₂s (PLA₂s)

Snake venom PLA₂s are 13–19 kDa proteins and one of the main components of animal venoms (Oliveira et al., 2022). Overall, we observed a mean RMSD difference of 3.2 ± 3.4 Å between MDLR and AF2 models (Table S3). The highest similarity is detected for P04417 from *Gloydius blomhoffii* (with an RMSD of 0.9 Å between AF2 and CF, 0.4/0.9 Å between MDLR and AF2/CF) (Fig. 5A). The PLA₂ domain itself is modelled to great agreement in all cases. The largest difference across predictions was found between P14411 from *Bungarus fasciatus* (5.2 Å between AF2 and CF, 5.8 Å between MDLR and AF2/CF), with differences mainly arising in positioning of the terminal loop regions

(Fig. 5B). The largest difference between AF2 and CF predictions was found to be for Q8AY47 (15.6 Å; Fig. 5C). RMSD was significantly different between AF2/CF and both AF2/MDLR and CF/MDLR, with AF2/CF having the lowest average RMSD (1.4 Å) and CF/MDLR the highest (3.3 Å; Fig. 7D).

3.3.5. Snake venom metalloproteases (SVMs)

SVMs are a class of enzymes found in the venom of snakes, with sizes ranging from 20 to 100 kDa (Oliveira et al., 2022). Out of 137 SVMs (average RMSD 80 ± 46 Å between MDLR and AF2 models), the highest similarity between models can be found in P20897 from *Crotalus ruber* (3.2 Å between AF2 and CF, 145 Å between Modeller and AF2/CF; Fig. 6A). Meanwhile, the largest difference was observed between models of Q10749 from *Naja mossambica* (6.8 Å between AF2 and CF, 8.7 Å/7.8 Å between Modeller and AF2/CF; Fig. 6B) mentioned earlier. The largest difference between AF2 and CF predictions was found to be for Q3HTN1 (27.8 Å; Fig. 6C). RMSD was significantly different between AF2/CF and both AF2/MDLR and CF/MDLR, with AF2/CF having the lowest average RMSD (6.6 Å) and CF/MDLR the highest (81 Å; Fig. 8D). Further analysis revealed that when separately analysing the smaller P-I and the larger P-III SVMs, there are larger discrepancies within the P-III models.

3.3.6. Snake venom serine proteases (SVSPs)

SVSPs are another class of enzymes found in snake venom, with sizes ranging from 26 to 67 kDa (Oliveira et al., 2022). Out of the 147 serine proteases (average RMSD 8 ± 13 Å between MDLR and AF2 models), the largest similarity between models can be found for Q7SZE2, from *Gloydius ussuriensis* (0.7 Å between AF2 and CF, 0.5 Å and 0.9 Å between MDLR and AF2/CF respectively; Fig. 7A). The largest disagreement between tools can be found for Q58L94 from *Notechis scutatus* (10 Å




Fig. 7. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (P04417) and (B) least (P14411) overlap across phospholipase A₂s. (C) The largest difference between AF2 and CF predictions (Q8AY47). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

between AF2 and CF, 78/75 Å between Modeller and AF2/CF respectively; Fig. 7B). Overall, the spread of RMSD between highest and lowest similarity models between MDLR and AF2 is 8.3 ± 14 Å. The biggest difference between AF2 and CF in RMSD for a given SVSP was found to be 16 Å (A6MFK8; Fig. 7C). RMSD was significantly different between AF2/CF and both AF2/MDLR and CF/MDLR, with AF2/CF having the lowest average RMSD (2.2 Å) and AF2/MDLR as well as CF/MDLR with the shared highest (8.3 Å; Fig. 9D).

3.3.7. Three-finger toxins (3FTxs)

3FTxs comprise three major subfamilies of toxins, i.e. cytotoxins, long-chain neurotoxins, and short-chain neurotoxins (Casewell et al., 2013). Out of the 275 3FTxs (average RMSD 3 ± 3 Å between MDLR and AF2 models), the greatest overlap between models was found to be for P60774 (0.37/0.92 Å between MDLR and AF2/CF, as well as 0.97 Å between AF2 and CF respectively; Fig. 8A). The largest difference in model RMSD was found to be for Q9W7K1 (12.8/12.7 Å between MDLR and AF2/CF, as well as 0.95 Å between AF2 and CF respectively; Fig. 8B). The protein found to have the largest differences in RMSD between AF2 and CF was P34074 (14.4 Å; Fig. 8C). RMSD was significantly different between all three comparisons, with AF2/MDLR having the lowest average RMSD (1.4 Å) and CF/MDLR the highest (3.5 Å; Fig. 10D).

4. Discussion

Prediction of protein structures is a critical aspect of biological research, particularly in understanding the function and potential therapeutic applications of proteins or how protein-based disease targets can be targeted with therapeutic agents. With the advent of deep

learning and AI, several protein structure prediction tools have emerged, such as Modeller (MDLR), ColabFold (CF), and AlphaFold2 (AF2). Each of these tools utilise different algorithms and approaches, resulting in different predictions of protein structures. Therefore, in this study we aimed to provide insight into the reliability and accuracy of these three modelling tools when dealing with proteins that lack experimentally solved structures, using snake venom toxins as model systems. A total of 1,062 snake venom toxin sequences, representing seven protein families, were retrieved, and structures were generated for each sequence using MDLR and CF. The respective AF2 structures were retrieved from the database, and all 3,186 structures (1,062 from each tool) were evaluated using various parameters, including Clash score, MolProbity score, as well as Ramachandran favoured and outlier percentage to ensure that no inherent quality bias was introduced by any of the tools. These metrics mostly capture stereochemical quality and geometric realism, but they may not always correlate with the accuracy of the model in reproducing the true, experimentally determined structure. However, they allow assessment of the relative quality and reliability of models generated by different prediction tools. The results of the analysis revealed significant differences in performance across all three models, with the two non-homology-based approaches AF2 and CF exhibiting superior performance in all four evaluated parameters compared to the homology-based method Modeller. As an additional observation, AF2 performed better than ColabFold across all evaluated parameters and significantly so in both Clash and MolProbity scores, which is in line with prior findings (Mirdita et al., 2021). These results indicate that AF2 and CF, which rely less on the availability of structurally resolved homologues, are more suited for predicting structures of diverse protein families, especially those without experimentally solved structures. The performance of Modeller, being a homology-based method, may be

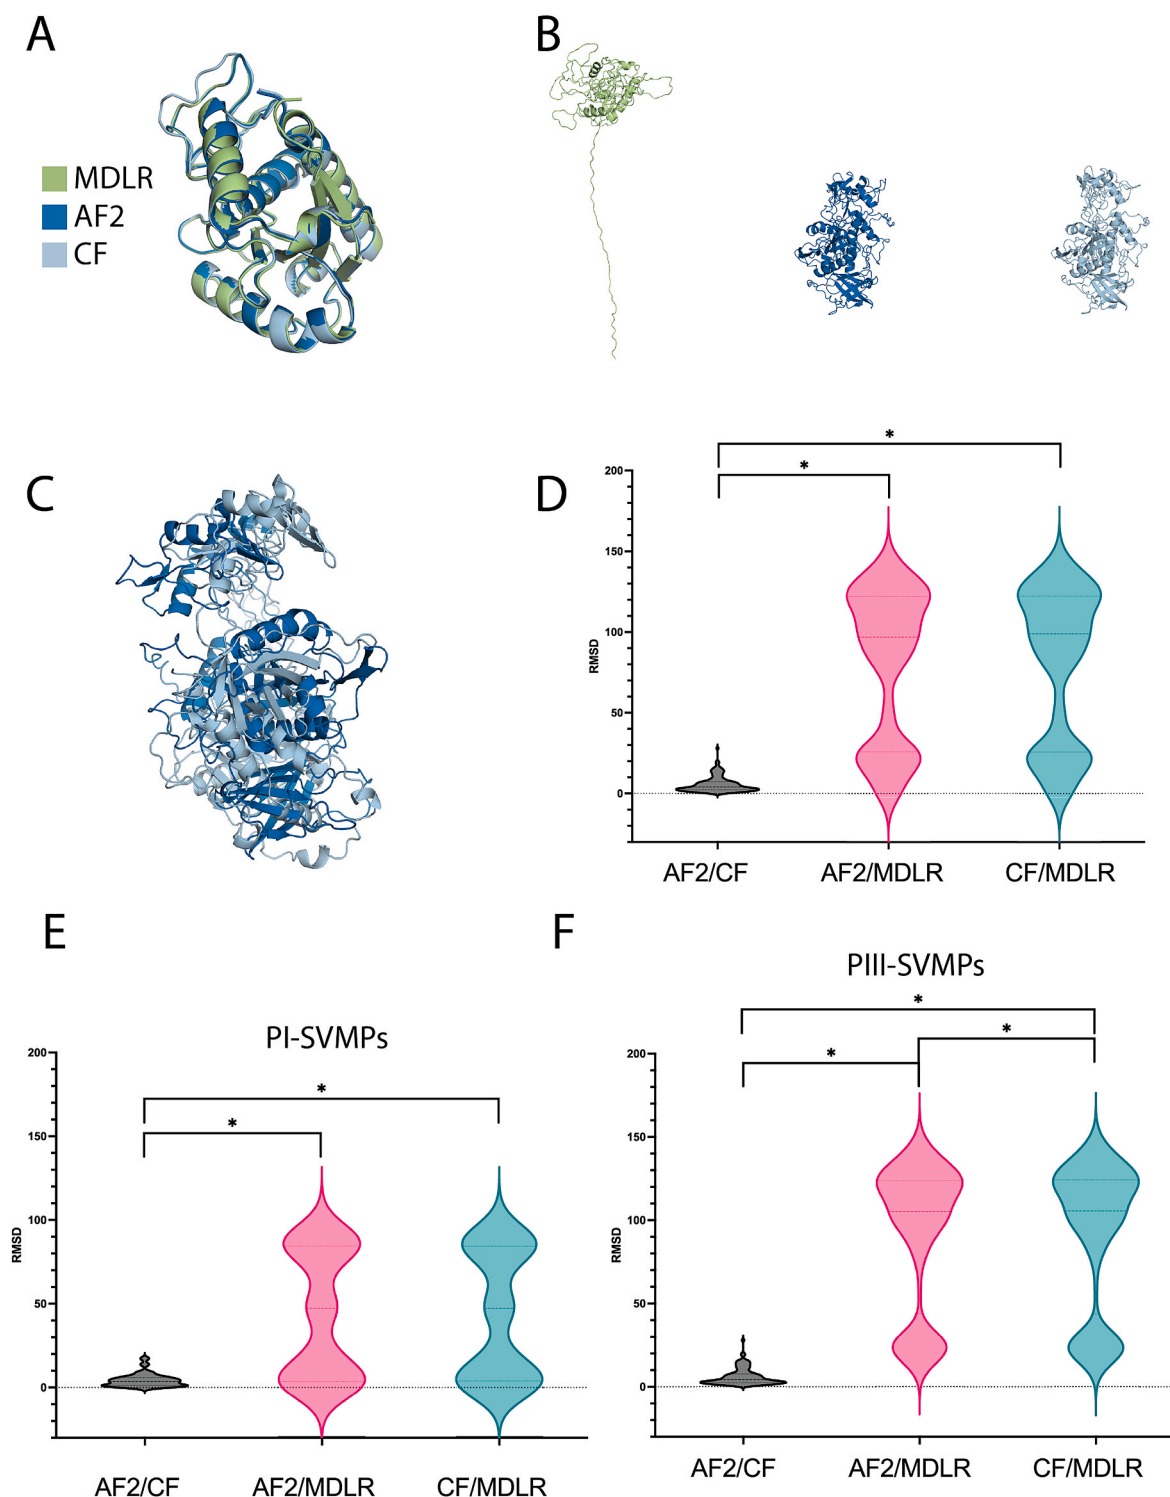


Fig. 8. The structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (P20897) and (B) least (Q10749) overlap across snake venom metalloproteases. (C) The largest difference between AF2 and CF predictions (Q3HTN1). (D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$). Differences are also shown for both P-I (E) and P-III (F) SVMs separately.

indicative of the challenges faced with poorly structurally characterised proteins, such as those found among snake venom toxins.

To have a better understanding of the differences in structure prediction quality of the three modelling tools, a three-way comparison was conducted over all assessed 1,062 snake toxins. RMSD analysis clearly indicated that a substantial spread existed in model overlap across the different toxins, ranging from 0.14 to 145 Å. The differences between

tools were again greatest when MDLR was involved, whereas differences between AF2 and CF were rarely significant. Notably, when considering different snake toxin families, MDLR exhibited the smallest performance differences in short and highly conserved proteins, such as 3FTxs that share the classic three-finger toxin fold. MDLR is typically not suitable for domains that do not belong to a fold family because it is designed to work with targets where it can successfully assign a fold (Webb and Sali,




Fig. 9. The snake venom serine protease structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (Q7SZE2) and (B) least (Q58L94) overlap across snake venom serine proteases. (C) The largest difference between AF2 and CF predictions (A6MFK8). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

2016), and differences may arise in the twilight zone of short sequences with relatively low similarity (Khor et al., 2015). Consequently, the largest performance differences were observed for long and complex proteins, such as SVMPs – particularly PIII SVMPs. Nevertheless, MDLR was able to successfully model conserved functional cores in SVMPs, such as the peptidase domain. This is likely due to the high level of structural conservation of such sites due to the need to also conserve function (Alberts et al., 2002). Another area where disagreement was observed between the generated protein structures was in loop regions. Whilst loop regions were typically identified across all three models we tested, their conformations often differed across MDLR, AF2, and CF. This was unsurprising and stems from proteins being dynamic molecules with a large conformational plasticity, allowing them to perform complex biological functions (Mukhopadhyay, 2022). Yet, these features are not uniformly distributed across the molecule, but are usually localised to parts with larger degrees of kinematic freedom (Papaleo et al., 2016). Thus, modelling conformations of loop regions remains challenging in computational biology and is usually inversely related to loop length (Barozet et al., 2021); these dynamics of flexibility, but also lack of reference structures, explains why we observed poor modelling of propeptides. It also highlights the need for a better structural understanding due to key roles of the propeptides in chaperoning and to provide insights into how the interactions of the propeptide region inhibits enzymatic function for the design of inhibitor molecules. Finally, it is notable that the overall mean difference between AF2 and CF was close to 3 Å, which is large enough to have a substantial impact when used for generative ML approaches for protein design (Wang et al., 2021; Watson et al., 2022). For some SVMPs, these differences even exceed 27 Å, highlighting the current limitations of the explored protein structure

prediction tools (Bryant et al., 2022).

Nevertheless, several of the toxin families assessed in this article have structural models that are of sufficiently high quality for further analysis, and thus a multitude of use cases; specifically, the highly conserved 3FTxs, KUNs, and PLA₂s can be used for computational simulations to model their interactions with different targets. This can help reveal the details of their binding sites, active sites, and conformational changes, as well as for the discovery and design of potential molecular binders for reagent, diagnostic, or therapeutic purposes (Norman et al., 2020). Notably, even toxin families with higher model variability and uncertainty (i.e. DISs, SVSPs, and SVMPs) could be used for similar purposes, as long as the focus falls on their conserved functional domains, such as their active sites. Importantly, these considerations regarding the limitations of computational structural modelling are not only relevant for toxin researchers, but can also be transferred to any other research area that is using computational predictions for protein structures and particularly ones with poor experimental coverage, such as rare diseases (Rossi Sebastiano et al., 2022). Overall, we hope the unravelled dynamics of computational structure prediction and the provision of all models and their comparisons constitute a key resource that can help de-risk future analyses.

5. Conclusion

The availability and accessibility of a range of powerful computational models are a game changer for structural biology. Whilst extremely powerful, the plethora of available tools each come with their own set of advantages and disadvantages; these are often somewhat understood within their application in model organisms, but little data




Fig. 10. The three-finger toxin structures predicted by Modeller (MDLR), AlphaFold2 (AF2), and ColabFold (CF), with the (A) most (P60774) and (B) least (Q9W7K1) overlap. (C) The largest difference between AF2 and CF predictions (P34074). D) Differences in root-mean-square deviation (RMSD) between the toxin structures predicted by MDLR, AF2, and CF. Significant differences were established via Wilcoxon matched-pairs signed rank test and indicated by an asterisk ($P < 0.05$).

exists on their performance on poorly characterised protein targets. Here, we studied the performance of three different computational structural biology tools on their predictions of over 2000 snake toxin structures that have few experimental reference structures available. We generated these structures using Modeller and ColabFold and compared them to each other as well as AlphaFold2 designs. We found that predictions were often closely aligned between CF and AF2, whereas MDLR often offered differing predictions. Nevertheless, differences between AF and CF were common, highlighting the need for cross-model validation of predicted structures. Notably, all tools performed well in predicting functional domains, while struggling with elements that are intrinsically disordered, such as loop regions. We further identified toxin families and structural features within these, as well as specific toxins, that were associated with substantially differing predictions across models. We therefore conclude that it is important to consider the complexity of the modelling task and use orthogonal modelling methods, such as AF2 and CF in combination with each other, to improve the reliability of structural assumptions. This will not only help future research quickly identify potential discrepancies and de-risk their use of these models, but also highlights key protein families, such as SVMPs, SVSPs, and DISs that require further experimental validation.

Ethical statement

No experimentation on human or animal subjects was involved in this study.

CRedit authorship contribution statement

Konstantinos Kalogeropoulos: Conceptualization, Methodology,

Investigation, Visualization, Resources, Writing - original draft, Writing - review & editing. **Markus-Frederik Bohn:** Conceptualization, Methodology, Investigation, Visualization, Resources, Writing - original draft, Writing - review & editing. **David E. Jenkins:** Methodology, Investigation, Writing - review & editing. **Jann Ledergerber:** Investigation, Visualization, Writing - review & editing. **Christoffer V. Sørensen:** Investigation, Writing - review & editing. **Nils Hofmann:** Investigation, Writing - review & editing. **Jack Wade:** Investigation, Writing - review & editing. **Thomas Fryer:** Investigation, Writing - review & editing. **Giang Thi Tuyet Nguyen:** Investigation, Writing - review & editing. **Ulrich auf dem Keller:** Resources, Supervision, Writing - review & editing. **Andreas H. Laustsen:** Resources, Supervision, Writing - review & editing. **Timothy P. Jenkins:** Conceptualization, Methodology, Investigation, Visualization, Funding acquisition, Project administration, Resources, Supervision, Writing - original draft, Writing - review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All structures have been deposited in Mendeley data (DOI: 10.17632/gjk47cjm26.1)

Acknowledgements

TPJ has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement no. 713683 (COFUNDfellowsDTU). UadK acknowledges support by a Young Investigator Award from the Novo Nordisk Foundation [NNF160C0020670]. AHL is supported by a grant from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program [850974], by a grant from the Villum Foundation [00025302], and a grant from Wellcome [221702/Z/20/Z].

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.toxicon.2023.107559>.

References

- Alberts, B., et al., 2002. Protein Function. Mol. Biol. Cell, 4th Ed.
- Altschuh, D., Lesk, A.M., Bloomer, A.C., Klug, A., 1987. Correlation of co-ordinated amino acid substitutions with function in viruses related to tobacco mosaic virus. J. Mol. Biol. 193, 693–707.
- Baek, M., Baker, D., 2022. Deep learning and protein structure modeling. Nat. Methods 19, 13–14.
- Baek, M., et al., 2021. Accurate prediction of protein structures and interactions using a three-track neural network. Science 373, 871–876.
- Bai, X.-C., McMullan, G., Scheres, S.H., 2015. How cryo-EM is revolutionizing structural biology. Trends Biochem. Sci. 40, 49–57.
- Barozet, A., Chacón, P., Cortés, J., 2021. Current approaches to flexible loop modeling. Curr. Res. Struct. Biol 3, 187–191.
- Bitencourt-Ferreira, G., de Azevedo, W.F., 2019. Homology modeling of protein targets with MODELLER. In: de Azevedo Jr, W.F. (Ed.), Docking Screens for Drug Discovery. Springer, New York, pp. 231–249. https://doi.org/10.1007/978-1-4939-9752-7_15.
- Bryant, P., et al., 2022. Predicting the structure of large protein complexes using AlphaFold and Monte Carlo tree search. Nat. Commun. 13, 6028.
- Casewell, N.R., Wüster, W., Vonk, F.J., Harrison, R.A., Fry, B.G., 2013. Complex cocktails: the evolutionary novelty of venoms. Trends Ecol. Evol. 28, 219–229.
- Cock, P.J.A., et al., 2009. Biopython: freely available Python tools for computational molecular biology and bioinformatics. Bioinformatics 25, 1422–1423.
- Davis, I.W., et al., 2007. MolProbity: all-atom contacts and structure validation for proteins and nucleic acids. Nucleic Acids Res. 35, W375–W383.
- Eddy, S.R., 2011. Accelerated profile HMM searches. PLoS Comput. Biol. 7, e1002195.
- Ferraz, C.R., et al., 2019. Multifunctional toxins in snake venoms and therapeutic implications: from pain to hemorrhage and necrosis. Front. Ecol. Evol. 7.
- Gasteiger, E., et al., 2003. ExPASy: the proteomics server for in-depth protein knowledge and analysis. Nucleic Acids Res. 31, 3784–3788.
- Gutiérrez, J.M., et al., 2017. Snakebite envenoming. Nat. Rev. Dis. Prim. 3, 17063.
- Jenkins, T.P., et al., 2019. Toxin neutralization using alternative binding proteins. Toxins 11 (53).
- Jones, D.T., Buchan, D.W., Cozzetto, D., Pontil, M., 2012. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. Bioinformatics 28, 184–190.
- Jumper, J., et al., 2021. Highly accurate protein structure prediction with AlphaFold. Nature 596, 583–589.
- Kermani, A.A., 2021. A guide to membrane protein X-ray crystallography. FEBS J. 288, 5788–5804.
- Khor, B.Y., Tye, G.J., Lim, T.S., Choong, Y.S., 2015. General overview on structure prediction of twilight-zone proteins. Theor. Biol. Med. Model. 12 (15).
- Knudsen, C., Laustsen, A., Knudsen, C., Laustsen, A.H., 2018. Recent advances in next generation snakebite antivenoms. Trav. Med. Infect. Dis. 3 (42).
- Laustsen, A.H., et al., 2016. From fangs to pharmacology: the future of snakebite envenoming therapy. Curr. Pharmaceut. Des. <http://www.eurekaselect.com/143508/article>.
- Li, L., Huang, J., Lin, Y., 2018. Snake venoms in cancer therapy: past, present and future. Toxins 10 (346).
- Liebschner, D., et al., 2019. Macromolecular structure determination using X-rays, neutrons and electrons: recent developments in Phenix. Acta Crystallogr. Sect. Struct. Biol 75, 861–877.
- Marks, D.S., et al., 2011. Protein 3D structure computed from evolutionary sequence variation. PLoS One 6, e28766.
- Mirdita, M., Steinegger, M., Söding, J., 2019. MMseqs2 desktop and local web server app for fast, interactive sequence searches. Bioinformatics 35, 2856–2858.
- Mirdita, M., et al., 2021. ColabFold-Making Protein Folding Accessible to All.
- Mohamed Abd El-Aziz, T., Soares, A.G., Stockand, J.D., 2019. Snake venoms in drug discovery: valuable therapeutic tools for life saving. Toxins 11 (564).
- Mukhopadhyay, S., 2022. Shapeshifting proteins: the role of structural disorder and conformational plasticity in physiology and disease. Essays Biochem. 66, 817–819.
- Norman, R.A., et al., 2020. Computational approaches to therapeutic antibody design: established methods and emerging trends. Brief. Bioinform 21, 1549–1567.
- Oliveira, A.L., et al., 2022. The chemistry of snake venom and its medicinal potential. Nat. Rev. Chem 6, 451–469.
- Papaleo, E., et al., 2016. The role of protein loops and linkers in conformational dynamics and allostery. Chem. Rev. 116, 6391–6423.
- Pereira, J., et al., 2021. High-accuracy protein structure prediction in CASP14. Proteins: Struct., Funct., Bioinf. 89, 1687–1699.
- Pieper, U., et al., 2014. ModBase, a database of annotated comparative protein structure models and associated resources. Nucleic Acids Res. 42, D336–D346.
- Protein Data Bank: the single global archive for 3D macromolecular structure data. Nucleic Acids Res. 47, 2019, D520–D528.
- Rodrigues, J.P., Teixeira, J.M., Trellet, M., Bonvin, A.M., 2018. Pdb-tools: a swiss army knife for molecular structures. F1000Research 7.
- Rossi Sebastiano, M., Ermondi, G., Hadano, S., Caron, G., 2022. AI-based protein structure databases have the potential to accelerate rare diseases research: AlphaFoldDB and the case of IAHSF/Alsin. Drug Discov. Today 27, 1652–1660.
- Shen, M., Sali, A., 2006. Statistical potential for assessment and prediction of protein structures. Protein Sci. 15, 2507–2524.
- Shindyalov, I.N., Bourne, P.E., 1998. Protein structure alignment by incremental combinatorial extension (CE) of the optimal path. Protein Eng. 11, 739–747.
- Shindyalov, I.N., Kolchanov, N.A., Sander, C., 1994. Can three-dimensional contacts in protein structures be predicted by analysis of correlated mutations? Protein Eng. Des. Sel. 7, 349–358.
- Steinegger, M., Söding, J., 2017. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. Nat. Biotechnol. 35, 1026–1028.
- Steinegger, M., et al., 2019. HH-suite3 for fast remote homology detection and deep protein annotation. BMC Bioinf. 20, 1–15.
- Tasoulis, T., Isbister, G.K., 2017. A review and database of snake venom proteomes. Toxins 9 (290).
- Teufel, F., et al., 2022. SignalP 6.0 predicts all five types of signal peptides using protein language models. Nat. Biotechnol. 40, 1023–1025.
- Thompson, M.C., Yeates, T.O., Rodriguez, J.A., 2020. Advances in methods for atomic resolution macromolecular structure determination. F1000Research 9.
- VenomZone. <https://venomzone.expasy.org/>.
- Wang, J., et al., 2021. Deep learning methods for designing proteins scaffolding functional sites, 2021.11.10.468128 Preprint at. <https://doi.org/10.1101/2021.11.10.468128>.
- Watson, J.L., et al., 2022. Broadly applicable and accurate protein design by integrating structure prediction networks and diffusion generative models, 2022.12.09.519842 Preprint at. <https://doi.org/10.1101/2022.12.09.519842>.
- Webb, B., Sali, A., 2016. Comparative protein structure modeling using MODELLER. Curr. Protoc. Bioinforma 54, 5.6.1–5.6.37.
- Weigt, M., White, R.A., Szurmant, H., Hoch, J.A., Hwa, T., 2009. Identification of direct residue contacts in protein–protein interaction by message passing. Proc. Natl. Acad. Sci. USA 106, 67–72.