Limits of deep-learning-based RNA prediction methods

Marko Ludaic[1], Arne Elofsson[1,*]

Science for Life Laboratory and Department of Biochemistry and Biophysics, Stockholm University, Solna 171 21, Sweden
*To whom correspondence should be addressed. Email: arne@bioinfo.se

Abstract

In recent years, tremendous advances have been made in predicting protein structures and protein-protein interactions. However, progress in predicting the structure of RNA, either alone or in complex with other macromolecules, has been less prominent, although some recent developments have been reported. It remains unclear whether the improved prediction accuracy is sustained for novel RNA structures. Here, we use an independent benchmark to evaluate the performance of the latest methods. First, we show that state-of-the-art methods can sometimes predict the structure of single-chain RNA strands, with accurate models observed for RNAs with well-defined or regular secondary structures. Next, our evaluation was extended to RNA complexes, where prediction accuracy was notably higher for those involving extensive canonical base pairing. Additionally, a structural similarity analysis revealed that prediction success strongly correlates with resemblance to known structures, indicating that current methods recognise recurring motifs rather than generalising to novel folds. Finally, we also noted that the estimation of accuracies for RNA models is far from accurate. Therefore, it is not possible to reliably identify the correctly predicted models with today's methods.

Keywords:

- RNA structure
- Deep Learning
- Structure Prediction
- RNA-Protein interactions

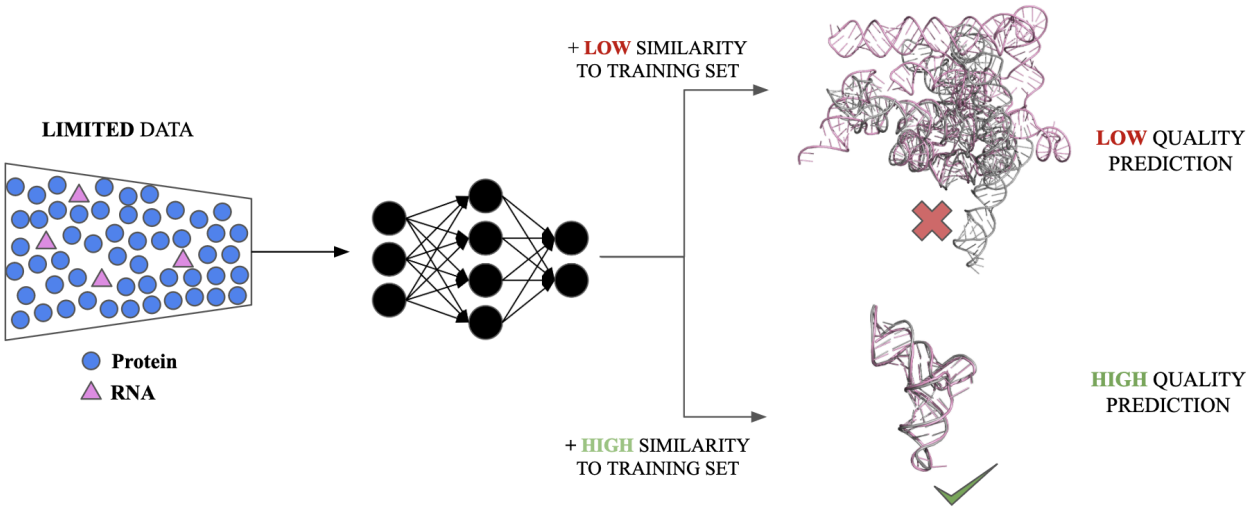

Graphical Abstract

Introduction

RNA molecules are highly dynamic macromolecules that adopt intricate three-dimensional (3D) structures essential for many cellular and viral processes, emergent RNA-based medicines, and biotechnologies. Different types of RNA, such as antisense oligonucleotides, small interfering RNAs, and microRNAs, have shown promising results as therapeutics (1). Even though RNAs were initially viewed as a mere intermediary in transferring genetic information from DNA to protein synthesis, they have been recently recognised for their functional versatility, extending far beyond their linear nucleotide sequence.

RNA structures are critical for elucidating structure-function relationships; therefore, modelling RNA structure from its nucleotide sequences remains an important and largely unsolved task. There are two main reasons for the limited success. First, the limited availability of determined RNA structures, at least one order of magnitude less than for proteins, poses a significant challenge for training AI-based models. Secondly, RNA's conformational dynamicity complicates the application of deep learning models. Many RNAs undergo drastic conformational changes to obtain a shape required for binding to a molecular partner (2). This is also one of the reasons why experimental methods such as X-ray crystallography, NMR, and cryo-EM have been primarily able to resolve regulatory and enzymatic RNAs such as tRNA, rRNA, and ribozymes (3). The structure of other RNA types like long non-coding RNAs (lncRNAs) remains less studied, possibly because they are indeed more disordered.

Prior to the release of RoseTTAFold2NA (RF2NA) in early 2024 (4) and AlphaFold3 later in 2024 (5), deep learning approaches were limited to predicting the structures of monomeric RNAs and could not handle RNA in a complex with other molecules. Recently, several deep learning approaches based on the AlphaFold3 architecture, such as Boltz-1 (6), Chai-1 (7), and HelixFold3 (8), were released. These methods can handle proteins, nucleic acids, small molecules, and ions, greatly expanding the scope and applicability of structure prediction. Furthermore, some methods for predicting single-chain RNA, such as RhoFold+ (Shen et al., 2024), utilize large language models (LLMs) to predict atom coordinates, while others are adaptations of AlphaFold. Table 1 provides an overview of the latest methods for predicting RNA structure.

Since advancements in RNA structure prediction are recent, only one benchmark study focused on a single-chain RNA has been reported previously (9). Some benchmarking results were

presented at the CASP16 meeting in December 2024, but those did not include the latest methods as they were released too late to participate in CASP16.

Table 1. A brief description of the most important characteristics of benchmarked methods for single-chain RNA and RNA-RNA/Protein prediction.

METHOD	DESCRIPTION	RNA
AlphaFold 3	A deep learning method consisting of a pairformer for residue interaction modeling and a diffusion module for denoising raw atomic coordinates, removing global rotational/translation constraints to enhance prediction flexibility and accuracy.	Single, complex
Boltz-1	AlphaFold 3-inspired architecture with modified order of operations in the MSAModule and transformer layer.	Single, complex
Chai-1	AlphaFold 3-inspired architecture incorporating residue-level embeddings from a large language model for proteins, also introducing experimental constraint features like pocket, contact, and docking constraints. Trained without RNA MSAs.	Single, complex
HelixFold3	AlphaFold 3-inspired architecture using PaddlePaddle library allowing for faster model training.	Single, complex
NuFold	A deep learning method built upon AlphaFold architecture design to predict RNA structures by incorporating nucleic acid sequences, secondary structure inputs, and RNA-specific inter-base angles, distances, and atom positions.	Single
RhoFold+	A deep learning model for RNA structure prediction that utilizes an RNA language model, alongside several deep learning modules, including an IPA module for modeling 3D positions.	Single
RoseTTAFoldNA	A three-track architectured model that simultaneously updates sequence (1D), residue-pair (2D), and structural (3D) representations to predict macromolecular structures..	Single, complex
trRosettaRNA	A deep learning model consisting of MSA and pair representations fed into a transformer network (RNAformer) predicting 1D and 2D geometries, refined in the final 3D structure folding step.	Single

In this study, we conducted a comprehensive benchmark study to evaluate the capabilities and limitations of these deep learning methods for both single-chain RNA and RNA complexes. The analysis compared the performance of nine methods for predicting monomeric RNA structures and four methods for predicting RNA complexes. By systematically assessing these models, we aim to provide valuable insights into their accuracy, strengths, and areas for further improvement.

Materials and Methods

Benchmark dataset

Nucleotide-containing entries released after September 30, 2021, were downloaded from the biological assemblies in the Protein Data Bank (PDB) along with their corresponding sequences. The first biological unit was used. In addition, we used CASP16 (Critical Assessment of Structure Prediction) targets containing RNA to supplement the benchmark dataset. Next, sequences were filtered using CD-HIT-EST (10) at 80% sequence identity to ensure dataset diversity. The identity was calculated using the longer sequence as a reference, which was critical, as many RNA sequences in the PDB are synthesized from short RNAs. Applying this

criterion avoids falsely reporting a sequence identity of 100% between a short and a long RNA sequence. Sequences containing ambiguous nucleotides, such as 'X' or 'N', were removed from the dataset. Furthermore, sequences shorter than 10 nucleotides that consisted predominantly of a single nucleotide type (>80% homogeneity) were also excluded, as they are likely to be linear and lack a modular structure, making them unsuitable for inclusion in the benchmark dataset.

MMseqs2 (11) was employed to assess sequence homology between proteins in the benchmark dataset and those in the AlphaFold training dataset (PDB structures released before September 30, 2021). For each RNA-protein complex, if at least one protein chain shared more than 40% sequence identity with any sequence in the AlphaFold training dataset, the corresponding RNA molecule in the same complex was also analyzed for sequence identity. Only if it had more than 80% sequence identity, the complex was excluded, since the primary focus of this benchmark was to evaluate the accuracy of RNA prediction methods, rather than protein prediction methods.

Finally, the dataset was reduced as we could not model some entries in the initial test set for some methods. Modelling complexes with more than five protein and/or RNA subunits with Boltz-1 version 0.4.1 failed due to out-of-memory errors. Consequently, such complexes were excluded from the analysis, reducing the RNA-RNA and RNA-protein complexes dataset to 42 examples. In the case of single-chain RNA structures, after filtering, our benchmark dataset consisted of 50 commonly generated single-chain RNA models.

Model generation

The benchmark evaluates the performance of nine prediction methods, AlphaFold3 (5), Boltz-1 (6), Chai-1 (7), HelixFold3 (8), NuFold (12), RhoFold+ (13), RoseTTAFoldNA (4), and trRosettaRNA (14), for single-chain RNA. Four of the methods. AlphaFold3, Boltz-1, HelixFold3, and RoseTTAFoldNA also predicted RNA-RNA and RNA-protein complexes. Chai-1, a multi-modal foundation model for molecular structure prediction, was excluded from the RNA complex analysis due to its inference limit of 2048 tokens.

For each method, the structures were generated using a single NVIDIA DGX-A100 Core GPU, with the latest version of each method (v1.0 for trRosettaRNA) as of December 1, 2024, and default settings. For all methods, the same set of MSAs was used. These were generated using rMSA (15) for RNA sequences and MMseqs2 for protein sequences. An rMSA search was performed against the RNACentral and NT databases using default settings. For proteins, MMseqs2 was run on the ColabFold databases as of December 1, 2024, which consist of UniRef30, BFD/Mgnify, and ColabFold DB (16). MSAs were converted into alternative formats depending on each model's specific requirements. The conversion scripts are available from the accompanying git repository. Using different seeds, twenty-five models, in the case of AlphaFold3, and five models for all other methods were generated.

The top-ranked model was used for analyzing each method. Several methods, including AlphaFold3, report the predicted template modeling (pTM) and the interface predicted template modeling (ipTM) scores, and rank models by combining pTM and ipTM scores calculated as follows:

```
0.8*ipTM + 0.2*pTM + 0.5*disorder - 100*clash.
```

Other methods, such as RF2NA or RhoFold+, report the predicted local distance difference test (pLDDT), which reflects confidence in the local structure (Mariani et al., 2013). We ranked the models based on the average pLDDT scores for these methods. If a method did not report any confidence or ranking score, we randomly selected one of the five models. We also noted that in almost all cases, the difference in quality between generated models is minimal (see Figures S1-S4).

Model evaluation

To evaluate the quality of predicted models, we used TM-score (17) for assessing single-chain RNA structures and the DockQ score (18, 19) in addition to the TM-score for evaluating RNA complexes.

When calculating TM-score using USalign's default parameters (C. Zhang et al., 2022), we observed unexpectedly low values, even when the structures appeared to be correct. For example, a double helical RNA (8SXL) modeled with AlphaFold3, which is almost perfectly predicted (RMSD = 0.47), had a TM-score of only 0.660 (Figure S5). The discrepancy was traced to the length normalization method used by TM-score. TM-score employs a scale factor (d_0) that is dependent on the shortest chain length, L_{min}, to calculate the similarity between two residues i and j, separated by a distance d_{ij}:

$$s(i,j) = \frac{1}{1+(\frac{d_{ij}^2}{d_0})} \text{ and } d_0 = 1.24\sqrt{L_{min} - 15} - 1.8$$

This formula is optimized for proteins, but we realized that d_0 was often too low for RNA chains, as they are usually shorter than proteins. To our knowledge, no one has tried to optimize TMscore for nucleotides. We observed that using a constant d_o of 5.0 produced realistic results and did not try to optimize that further. Therefore, we calculate the TM-score using the following command line options:

```
{model.pdb} {reference.pdb} -mm 1 -d 5 -ter 0
```

for RNA complexes and

```
{model.pdb} {reference.pdb} -mol RNA -d 5
```

for single-chain RNA molecules.

Similarly, we calculate the TM-score for protein chains using the following command line option:

```
{model.pdb} {reference.pdb} -mol prot
```

Changing the d_0 parameter to 5.0 provided a more realistic TM-score for the AlphaFold3 model of 8SXL of 0.965 (Figure S5).

DockQ provides an overall score, averaged across all interfaces, as well as individual scores, one for each interface. Therefore, we also assessed the accuracy of RNA-RNA and RNA-protein interfaces within each complex, using the per-interface DockQ score.

Structural similarity to the training set

We also conducted a structural similarity analysis between the benchmark dataset and the AlphaFold3 training dataset, which is identical to the training set of most methods, to identify structural redundancy between the training and benchmark sets. Specifically, we downloaded all RNA-containing entries, single chains and complexes, present in PDB before September 30, 2021, which corresponds to the cutoff date for the AlphaFold3 training set.

A structural search was performed for all entries in the test set against the downloaded PDB set using USalign with the same settings described above. The TM-scores were computed between each RNA structure in the training set and all the RNA-containing structures in the training set. For each entry in the benchmark dataset, the maximum TM-score in the training set represents the highest structural similarity to any structure in the training set. This number was then used to classify targets based on their structural similarity to the training set.

Results

Performance of structure prediction methods on single-chain RNAs

Figure 1 shows the performance of deep learning methods for single-chain RNA prediction. In addition to the benchmark set of single RNAs, we evaluated the prediction accuracy of individual chains in the RNA-RNA and RNA-protein complexes for the four methods that predict complexes.

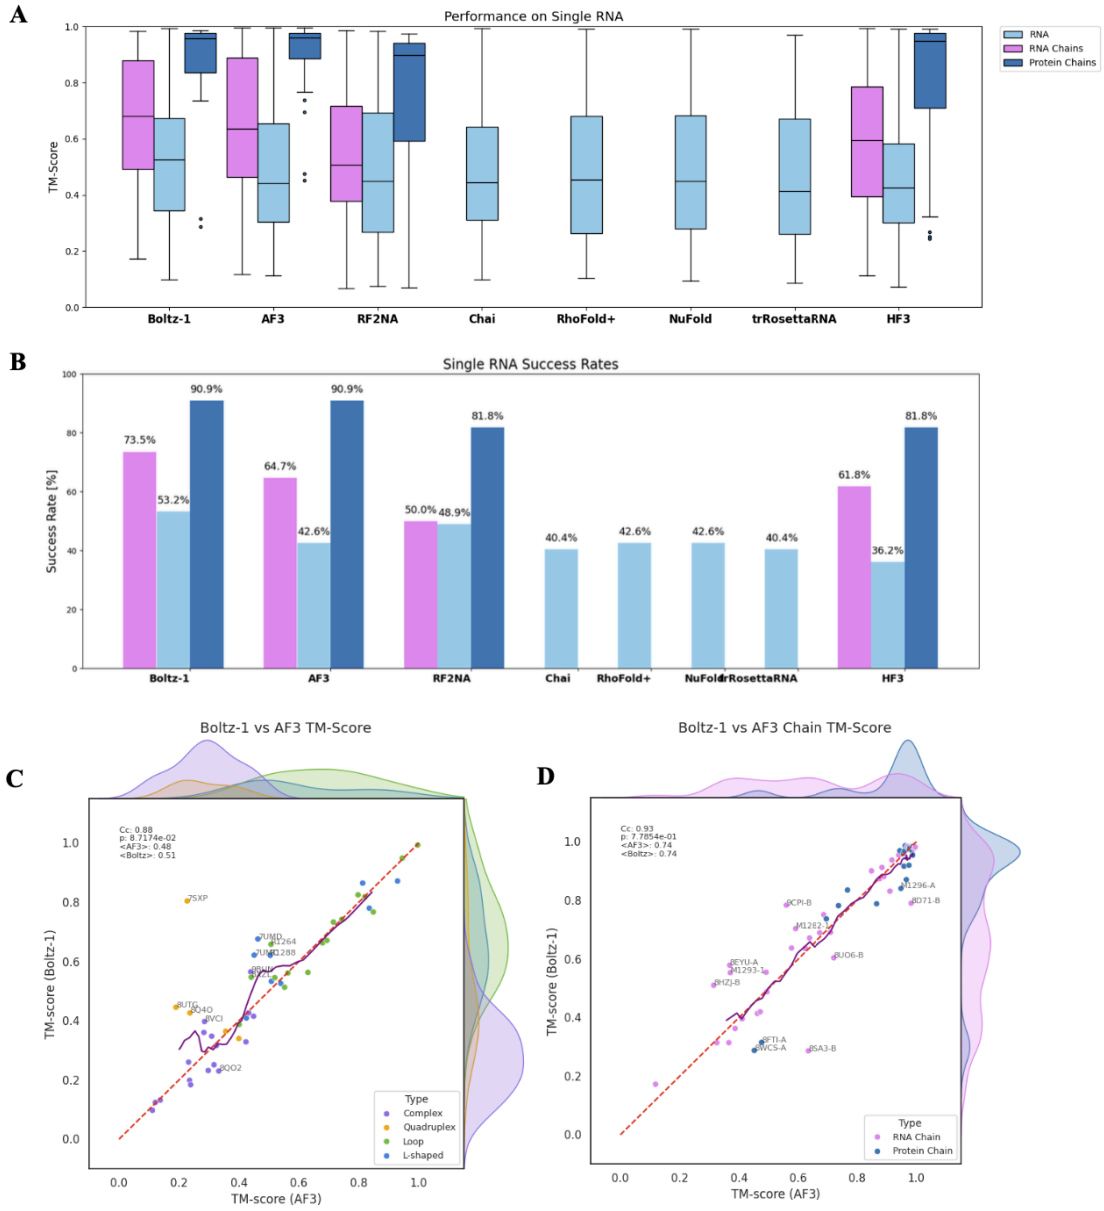

Figure 1. Performance comparison for different RNA prediction methods on single-chain RNA. (A) The performance of RNA prediction models was evaluated using the TM-score. The distribution of TM-scores for each method is presented as a box plot, with the median marked by a line and the interquartile ranges shown. The methods are arranged in descending order of the mean TM-score. Different colors correspond to different datasets: single-chain RNA (light blue), RNA chains derived from complexes (pink), and protein chains derived from complexes (dark blue). (B) Success rates achieved by all methods are defined as the fraction of acceptable models with a minimum TM score of 0.5. (C) Pairplot showing correlation between Boltz-1 (y-axis) and AlphaFold3 (x-axis) TM-scores for single-chain RNA. The labeled points show PDB IDs of predicted models whose difference in TM-score between the two methods is greater than 10%. Paired t-test reported no significant difference between

the two methods (Cc = 0.88, P = $8.7*10^{-2}$). (D) Pairplot showing correlation between Boltz-1 (y-axis) and AlphaFold3 (x-axis) TM-scores for single-chain RNA and protein chains (Cc = 0.93, P = $7.8*10^{-1}$).

Figure 1A shows the distribution of TM-scores for each method, and the success rates, i.e. the fraction of models with a TM-score ≥ 0.5, are shown in Figure 1B. TM-scores for all models are shown in Table S1. Overall, Boltz-1 demonstrates the best performance, achieving an average TM-score of 0.512 and a success rate of 53%, followed by AlphaFold3 (0.483/42%) and RF2NA (0.481/49%). A paired T-test reported no significant difference between the predictions from Boltz-1 and AlphaFold3, but NuFold (P = 3.09*10-3), HF3 (P = 6.75*10-3), and trRosettaRNA (P = 1*10-2) are statistically significantly less accurate than Boltz-1. A similar trend can be seen for the chains, predicted as a part of a complex, with the highest mean TM-score of 0.663 for RNA chains obtained by Boltz-1. As expected, the protein components are predicted better, 0.887/91% for AlphaFold3. However, no significant statistical difference was found between AlphaFold3 and Boltz-1 methods, see Table S2.

To better understand performance differences across different types of RNAs, we plotted the TM scores of the top-performing method, Boltz-1 (Figure 1C and 1D). We categorized the RNAs into short RNA loops, G-quadruplexes, L-shaped structure conformations such as those seen in tRNA, and Complex RNA structures, which are typically longer and composed of multiple distinct secondary structure motifs. First, the prediction quality is highly correlated across different methods, as shown in Figures 1C, 1D, and S6. That is, some targets are generally easier to predict than others.

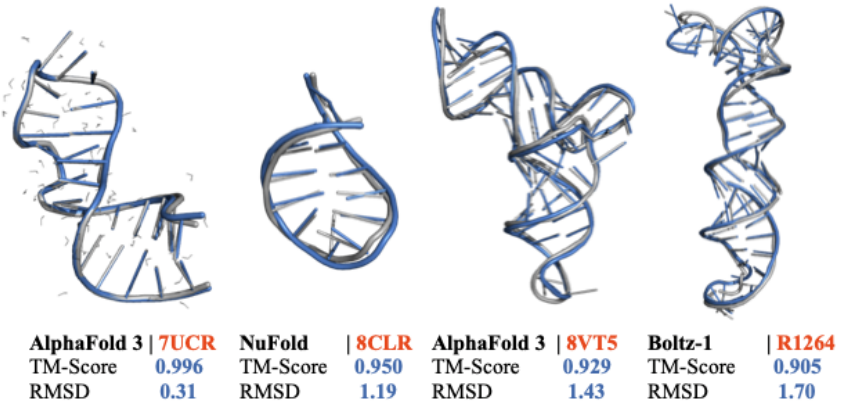

Figure 2. Top four predicted monomeric RNA structures, ranked by model quality. Predicted RNA is shown in blue, while the ground truth structure is colored grey. TM scores and RMSD are reported for each model, along with the name of the prediction method that achieved the highest score. Models are displayed in order from best to worst based on TM-score.

In general, short RNA loops and L-shaped structure conformations (Figure 2) are easier to predict, compared with longer and more complex RNA structures (Figure S7). Exception to this

trend is a couple of models, including one G-quadruplex structure, 7SXP, where Boltz-1 makes a significantly better model than AlphaFold3 and the other methods.

Performance of structure prediction for RNA-RNA and RNA-protein complexes

The performance of the benchmarked methods, evaluated using TM-score and DockQ score, is shown as violin plots in Figures 3A and 3B, respectively. Among the four methods, AlphaFold3 demonstrated the highest mean TM-score (0.642; Figure 3A) and DockQ score of (0.395; Figure 3B), followed by Boltz-1 (0.597/0.345), HF3 (0.496/0.315), and RF2NA (0.466/0.105). Using a DockQ score of ≥ 0.23 to define successful models, the success rates for benchmarked methods are 53.8%, 48.7%, 35.9%, and 17.9% for AlphaFold3, Boltz-1, HF3, and RF2NA, respectively. TM-scores and DockQ scores across all methods are shown in Tables S3 and S4, respectively. Mean TM-scores and success rates are shown in Table S5.

To gain more insights, we analyzed individual RNA-RNA and RNA-protein interfaces within each complex by calculating DockQ scores for each pair of chains, i.e., the interface DockQ. The individual DockQ scores for RNA-RNA interfaces, compared to those for RNA-protein interfaces, suggest that there is no significant difference in prediction accuracy between these two types (Tables S5 and S6). Here, AlphaFold3 shows the highest mean interface DockQ scores for both RNA-protein (0.382) and RNA-RNA interfaces (0.343), followed by Boltz-1 (0.309/0.330).

The relationships of the model qualities of the different methods are shown in Figure 3E, 3F, and S8. AlphaFold3 outperformed Boltz-1 in some cases. AlphaFold3 outperformed Boltz-1 in some cases. A statistically significant difference was found for TM-Score ($P = 1.95\times10^{-2}$; Figure 3D), but no statistically significant difference was found for interface DockQ score ($P = 8.53\times10^{-2}$; Figure 3E). Mean TM-scores and DockQ scores across all methods revealed that the same model consistently achieved the highest scores. Notably, this highly accurate model is a double helical RNA homodimer (8SXL), likely benefiting from the relative simplicity of base pair complementarity, which makes it easier to predict. The structure comparison across all methods for this model, as well as another high-scoring model, 8TQX, is shown in Figure S5.

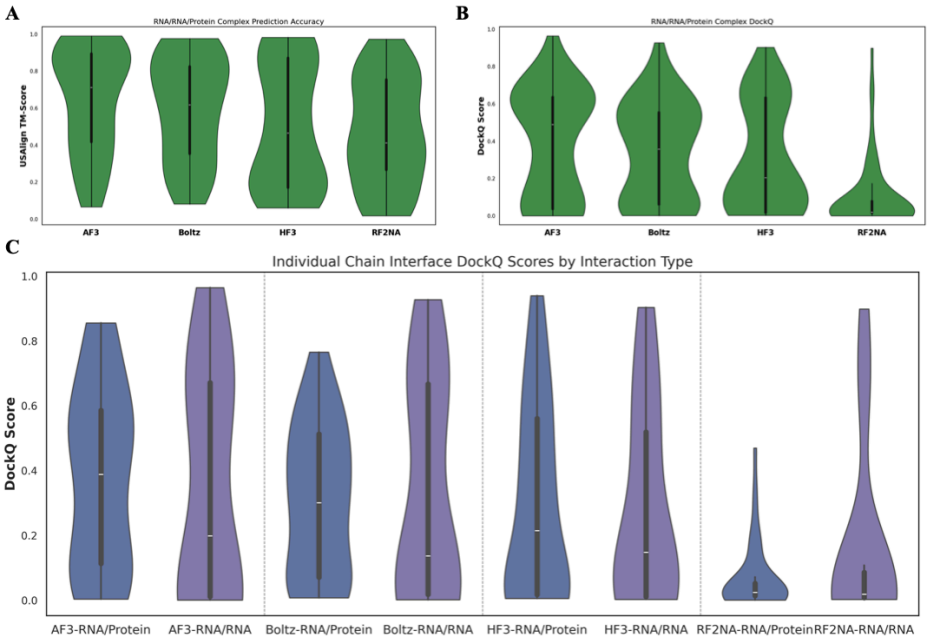

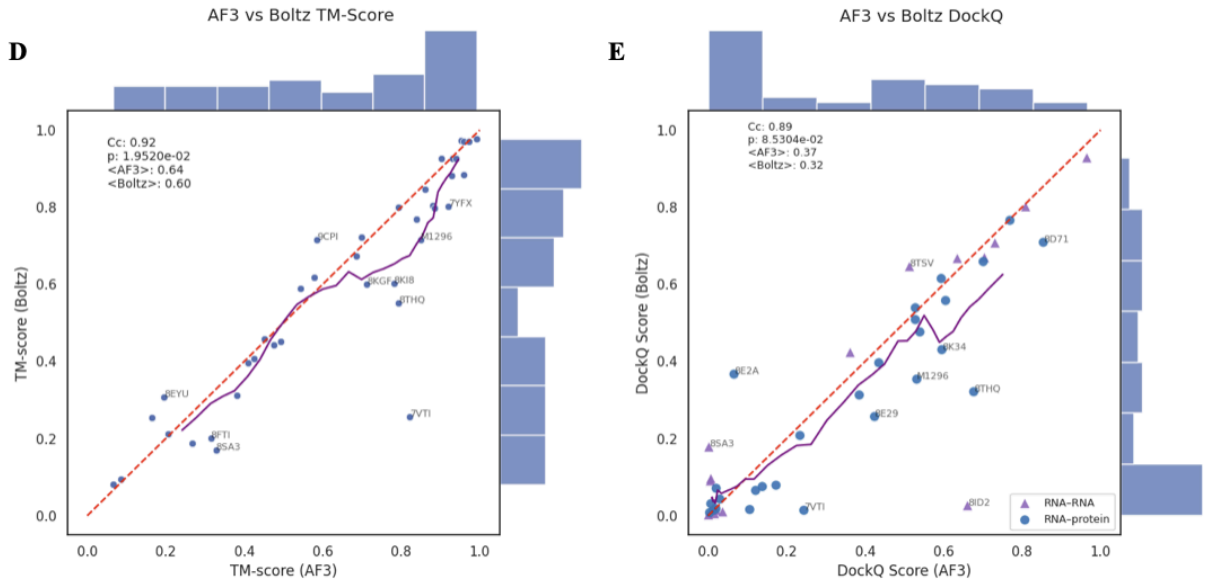

Figure 3. Performance comparison for different RNA prediction methods on RNA complexes. (A) Performance of RNA complex prediction models was evaluated using USAlign TM-score. Each method corresponds to a separate violin plot. The horizontal line inside each violin plot represents the median, while the thin bar in the centre represents the interquartile range. Violin plots are shown in a descending order of TM-score in both cases. (B) Performance of RNA complex prediction models evaluated using DockQ score, following the same layout as in TM-score analysis. (C) DockQ score for individual pairs of chains (Interface DockQ). DockQ scores were calculated separately for RNA-protein (blue) and RNA-RNA (purple) complexes and shown as violin plots arranged in descending order based on the RNA-protein interface DockQ scores. (D) Scatter plot showing correlation between AlphaFold3 and Boltz-1 TM-score for RNA complexes (Cc = 0.92, P = $1.95*10^{-2}$). A rolling average calculated between scores of the two methods is shown as a purple line. Bar plots on each side show the distribution of scores for each model separately (AlphaFold3 on the x-axis, Boltz on the y-axis). E) Scatter plot showing correlation between AlphaFold3 and Boltz-1 interface DockQ score for RNA complexes (Cc = 0.89, P = $8.53*10^{-2}$).

Models with low DockQ scores and high TM-scores

For most models, the DockQ score and TM-score agree well (Figure S9). However, for five RNA-protein interfaces (7VSJ, 7YEY, 8E2A, 8KI8, M1282) we observed a high TM score (>0.5), despite a very low interface DockQ score (<0.23). In these cases, the RNA molecule was misplaced, bound to the wrong region of the larger protein chain. In two of these cases (7VSJ, 7YEY), the single protein chain has an internal structural symmetry. The RNA chain was incorrectly placed on the opposite side of the symmetric fold, resulting in an incorrect prediction of the interface. Although the two sides are structurally similar, the sequences are not (Figure S10).

The reliability of pTM/ipTM confidence scores

Several methods, including AlphaFold3, Boltz-1, and HelixFold3, provide confidence metrics alongside predicted structures in the form of predicted Template Modelling (pTM) scores for both monomeric and complex structures, and the interface predicted Template Modelling (ipTM) scores for macromolecular complexes. To assess the predictive reliability of these metrics, we compared them with the calculated TM-score for single-chain RNA and RNA complexes, as well as the DockQ score for RNA complexes.

For single-chain RNAs, the reported pTM scores were low (pTM < 0.5), indicating limited confidence in the structural predictions across all methods (see Figure 4). Even more striking is the low correlation between pTM and TM scores. Often, the pTM scores are below 0.5, even for correctly predicted models, and the correlation is best for HelixFold3, at 0.61. In contrast, for some complexes, particularly protein-RNA complexes, the ipTM scores are relatively high (greater than 0.8), as shown in Figure 4. The correlation is also higher, with the best correlation between ipTM and DockQ score found for HF3 (0.81), followed by Boltz-1 (0.73) and AlphaFold3 (0.62). Notably, RNA-RNA complexes have lower ipTM and DockQ scores than RNA-protein complexes, indicating that these are harder to predict. We hypothesize that the increased confidence for RNA-protein complexes is due to the presence of protein chains in the complexes.

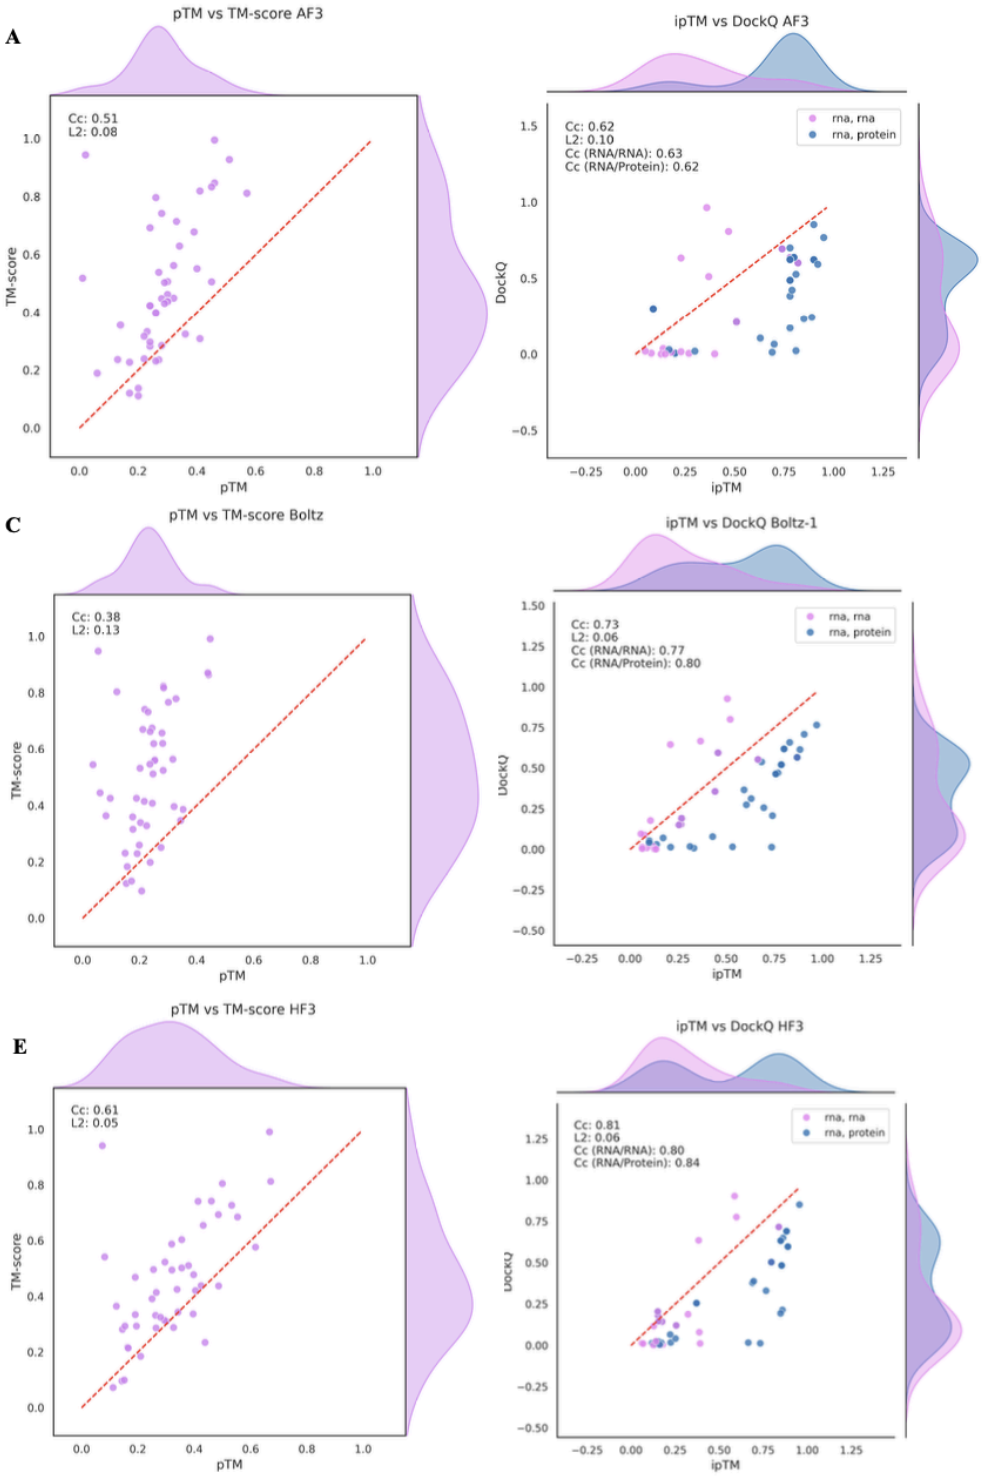

Figure 4. Evaluation of pTM/ipTM confidence scores. (A, C, E) Scatter plot between TM-score and pTM score for AlphaFold3/Boltz-1 and HelixFold3 models for single-chain models. (B, D, F) Scatter plot between DockQ score and ipTM score for RNA complexes, for AlphaFold3, Boltz-1, and HelixFold3. RNA-RNA complexes are colored in pink, and RNA-protein complexes in blue.

Prediction accuracy dependence on AlphaFold3 training set

The presence of commonly found RNA motifs in the PDB raises the possibility that the benchmarked methods may perform well by recognizing frequently occurring patterns, rather than generalizing to novel or less common RNA structures. In other words, the methods may primarily be able to detect structural similarity to the training set. As our results suggest, double helical and cloverleaf structures are usually predicted with higher accuracy, but these are also more common in PDB. To investigate the extent to which model accuracy is influenced by prior exposure to similar motifs during training, we assessed whether motifs of benchmarked structures were present in AlphaFold3's training data.

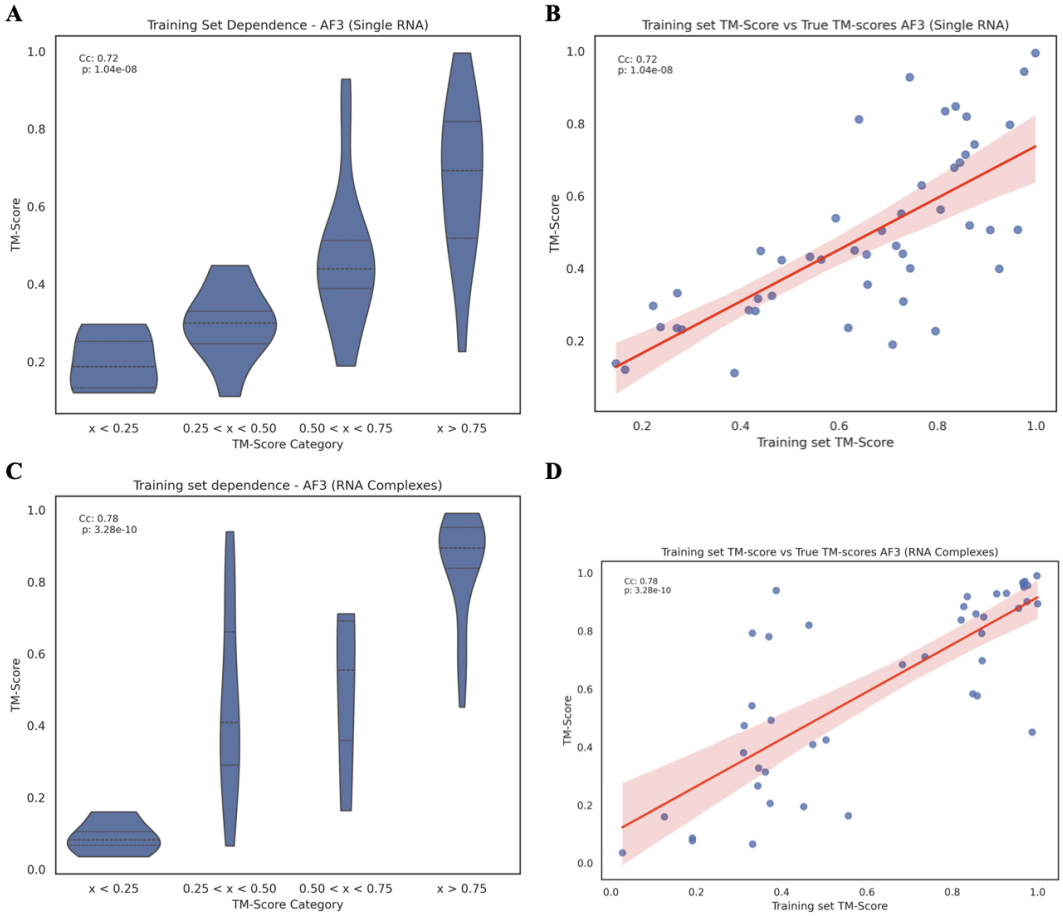

Figure 5. Similarity to the training set. (A) TM-score of benchmarked single RNA models dependent on the training set. TM-scores calculated between structures from the training set (training set TM-score) and benchmarked single RNA structures are divided into four categories. (B) Correlation between TM-scores for benchmarked structures (y-axis) and training set TM-score (x-axis), in the case of AlphaFold3 (Cc = 0.72, P = $1.04*10^{-8}$). (C) TM-score of benchmarked RNA complexes dependent on the training set. TM scores calculated between structures from the training set (training set TM score) and benchmarked RNA complexes are divided into four categories (TM score categories). (D) Correlation between TM-scores for benchmarked structures (y-axis) and training set TM-score (x-axis), in the case of AlphaFold3 (Cc = 0.78, P = $3.28*10^{-10}$).

For both single RNA (Figure 5A, 5B, and S11-S13) and RNA complexes (Figure 5C, 5D, and S14), we defined four similarity categories using the following TM-score thresholds: <0.25, 0.25–0.5, 0.5–0.75, and >0.75, We observed that prediction accuracy, as measured by TM-score, generally increased with higher structural similarity to the training set for all methods. The mean TM-score is below 0.1 for the category with the least similarity and increases gradually to over 0.6 for the category with the highest similarity to the training set. This suggests that AlphaFold3 and other methods tend to perform better when the target structure is more similar to motifs it encountered during training, highlighting the limitation of current methods in predicting unseen and structurally divergent RNAs.

Discussion

Advancements in deep learning have significantly impacted RNA structural biology, enabling the prediction of both monomeric RNA structures and RNA complexes. In this benchmark study, we evaluated several state-of-the-art deep learning methods for RNA structure prediction, revealing their strengths and limitations. Our results suggest that while current methods show promising performance, particularly for well-characterized RNA motifs, substantial challenges still remain.
Approximately one-third of the models achieve high accuracy.
The accurately predicted models correspond primarily to RNA molecules with structural similarity to already known RNA structures. Given that current methods are trained on existing experimentally solved structures, our findings suggest that the lack of RNA structural diversity in the PDB hinders generalization to novel and structurally complex RNAs, and that current state-of-the-art methods are unable to generalize outside of their training regime. Further, even within the training regime, the ability to reliably identify the correct models is a severe limitation for large-scale predictions. Therefore, our results also point to an urgent need for broader and more structurally diverse RNA training datasets.
RNA is increasingly recognized as a versatile macromolecule with essential catalytic and regulatory functions, many of which are unknown. We anticipate that growing interest in RNA biology will stimulate further experimental efforts towards solving complex RNA structures. Combined with innovations in techniques such as cryo-EM (2) and continued advances in machine learning (20), we are optimistic that RNA 3D structure prediction will evolve beyond simple motifs and toward accurately modelling the full complexity of RNA structural space. A possible first step would be to improve the estimates of model accuracy.

Acknowledgements

We thank the members of the Elofsson group for valuable input.

Supplementary data

Supplementary data are available at NAR online.

Conflict of interest

The authors declare no conflict of interest.

Author Contributions

ML did all the computational work, wrote the final draft and the final version, and perfomed the analysis.
AE obtained funding and helped with analysis and writing.

Funding

AE was funded by the Vetenskapsrådet Grant No. 2021-03979 and the Knut and Alice Wallenberg Foundation, Grant No. 2022.0032. The computations and data handling were enabled by the supercomputing resource Berzelius, provided by the National Supercomputer Centre at Linköping University, the Knut and Alice Wallenberg Foundation, and SNIC, with grant Numbers. SNIC 2021/5-297 and Berzelius-2021-29.

Data availability

Code used for running the methods and analyzing the main results is available from https://doi.org/10.5281/zenodo.15599802 and https://github.com/iammarcol/RNA-Benchmark

References

1. Damase,T.R., Sukhovershin,R., Boada,C., Taraballi,F., Pettigrew,R.I. and Cooke,J.P. (2021) The limitless future of RNA therapeutics. *Front. Bioeng. Biotechnol.*, **9**, 628137.

2. Ma,H., Jia,X., Zhang,K. and Su,Z. (2022) Cryo-EM advances in RNA structure determination. *Signal Transduct. Target. Ther.*, **7**, 58.

3. Zhang,J., Fei,Y., Sun,L. and Zhang,Q.C. (2022) Advances and opportunities in RNA structure experimental determination and computational modeling. *Nat. Methods*, **19**, 1193–1207.

4. Baek,M., McHugh,R., Anishchenko,I., Jiang,H., Baker,D. and DiMaio,F. (2024) Accurate prediction of protein-nucleic acid complexes using RoseTTAFoldNA. *Nat. Methods*, **21**, 117–121.

5. Abramson,J., Adler,J., Dunger,J., Evans,R., Green,T., Pritzel,A., Ronneberger,O., Willmore,L., Ballard,A.J., Bambrick,J., *et al.* (2024) Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, **630**, 493–500.

6. Wohlwend,J., Corso,G., Passaro,S., Reveiz,M., Leidal,K., Swiderski,W., Portnoi,T., Chinn,I., Silterra,J., Jaakkola,T., *et al.* (2024) Boltz-1 democratizing biomolecular interaction modeling. *bioRxivorg*, 10.1101/2024.11.19.624167.

7. Discovery,C., Boitreaud,J., Dent,J., McPartlon,M., Meier,J., Reis,V., Rogozhnikov,A. and Wu,K. (2024) Chai-1: Decoding the molecular interactions of life. *bioRxiv*, 10.1101/2024.10.10.615955.

8. Liu,L., Zhang,S., Xue,Y., Ye,X., Zhu,K., Li,Y., Liu,Y., Gao,J., Zhao,W., Yu,H., *et al.* (2024) Technical report of HelixFold3 for biomolecular structure prediction. *arXiv [q-bio.BM]*, 10.48550/ARXIV.2408.16975.

9. Bernard,C., Postic,G., Ghannay,S. and Tahi,F. (2024) State-of-the-RNArt: benchmarking current methods for RNA 3D structure prediction. *NAR Genom. Bioinform.*, **6**, lqae048.

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

11. Steinegger,M. and Söding,J. (2017) MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat. Biotechnol.*, **35**, 1026–1028.

12. Kagaya,Y., Zhang,Z., Ibtehaz,N., Wang,X., Nakamura,T., Punuru,P.D. and Kihara,D. (2025)

NuFold: end-to-end approach for RNA tertiary structure prediction with flexible nucleobase center representation. *Nat. Commun.*, **16**, 881.

13. Shen,T., Hu,Z., Sun,S., Liu,D., Wong,F., Wang,J., Chen,J., Wang,Y., Hong,L., Xiao,J., *et al.* (2024) Accurate RNA 3D structure prediction using a language model-based deep learning approach. *Nat. Methods*, **21**, 2287–2298.

14. Wang,W., Feng,C., Han,R., Wang,Z., Ye,L., Du,Z., Wei,H., Zhang,F., Peng,Z. and Yang,J. (2023) trRosettaRNA: automated prediction of RNA 3D structure with transformer network. *Nat. Commun.*, **14**, 7266.

15. Zhang,C., Zhang,Y. and Pyle,A.M. (2023) RMSA: A sequence search and alignment algorithm to improve RNA structure modeling. *J. Mol. Biol.*, **435**, 167904.

16. Mirdita,M., Schütze,K., Moriwaki,Y., Heo,L., Ovchinnikov,S. and Steinegger,M. (2022) ColabFold: making protein folding accessible to all. *Nat. Methods*, **19**, 679–682.

17. Zhang,Y. and Skolnick,J. (2005) TM-align: a protein structure alignment algorithm based on the TM-score. *Nucleic Acids Res*, **33**, 2302–2309.

18. Basu,S. and Wallner,B. (2016) DockQ: A quality measure for protein-protein docking models. *PLoS One*, **11**, e0161879.

19. Mirabello,C. and Wallner,B. (2024) DockQ v2: improved automatic quality measure for protein multimers, nucleic acids, and small molecules. *Bioinformatics*, **40**.

20. Kwon,D. (2025) RNA function follows form - why is it so hard to predict? *Nature*, **639**, 1106–1108.