nature machine intelligence

Supplementary information

https://doi.org/10.1038/s42256-026-01223-x

Predicting RNA 3D structure and conformers using a pre-trained secondary structure model and structure-aware attention

In the format provided by the
authors and unedited

Contents

Supplementary Results

1. Performance of trRNA2-SS for RNA secondary structure prediction

1.1. Benchmark results

The limited data of experimentally determined 3D RNA structures is a primary bottleneck hindering the generalization of deep learning methods in RNA structure prediction. To overcome this, we introduce an SS prior module (Fig. 1b), which was pre-trained on the extensive secondary structure data from bpRNA [1], to infuse 2D prior knowledge into the 3D structure prediction process. Meanwhile, this module also serves as an independent RNA secondary structure predictor, named trRNA2-SS. trRNA2-SS is trained using a transfer learning approach: initial pre-training on the bpRNA dataset, followed by fine-tuning on a smaller dataset derived from PDB (see Methods).

To rigorously evaluate the performance of trRNA2-SS, we collected two independent test sets. The first test set includes 50 RNAs from ArchiveII [2], a widely used SS benchmark set. The second set, PDB2D, includes 19 non-redundant RNA-only entries from PDB, all released after 2022-01. Both datasets were rigorously filtered to eliminate redundancy with our training data (see Methods for details).

We compare trRNA2-SS against a series of representative methods, including traditional energy- and/or evolution-based approaches (EternaFold [3], RNAfold [4], RNAstructure [5], PETfold [6]) and deep learning-based methods (RNA-FM [7], RNA-MSM [8], UFold [9], MXfold2 [10], SPOT-RNA [11], SPOT-RNA2 [12], KnotFold [13], BPfold [14]). Among them, PETfold, SPOT-RNA2, and RNA-MSM are based on MSA, whereas the other methods rely solely on the single sequence.

As shown in Supplementary Fig. 1 and Supplementary Tables 1-2, trRNA2-SS outperforms all compared methods on PDB2D and most compared methods on ArchiveII, in terms of various metrics. It is important to note that while we explicitly removed ArchiveII sequences redundant with our training set, we cannot guarantee their non-redundancy with the training sets of other deep learning-based methods. This potential data overlap may lead to performance inflation for those methods on ArchiveII. This is particularly evident in BPfold, where the per-target average AUPRC drops precipitously from a near-perfect 0.959 on ArchiveII to 0.723 on PDB2D. In contrast, trRNA2-SS maintains consistently robust performance across both datasets, achieving AUPRC values exceeding 0.8. Moreover, the compact distributions (relatively flat boxes) observed for trRNA2-SS in Supplementary Fig. 1c further illustrate its stability. These findings demonstrate the adaptability and robustness of trRNA2-SS, making it a state-of-the-art method for RNA secondary structure prediction.

The superiority of trRNA2-SS stems from its advanced network architecture, which comprises 8 RNAformer blocks. Unlike existing deep learning-based approaches that rely on simple ResNet, LSTM, or vanilla Transformer architectures, RNAformer utilizes a two-track mechanism to iteratively process and update bi-modal representations (Algorithm 4). Within each track, axial attention modules are employed to accommodate the multidimensional nature of these representations, while biased attention and outer-product expansion facilitate cross-modal interactions. Additionally, the pair track incorporates a multi-scale ResNet (Res2Net [15]) to enhance local detail modeling, which complements the global receptive field of the attention mechanism. Notably, this advanced architecture does not incur significant computational overhead or parameter growth (Supplementary Table 3). For instance, each trRNA2-SS model contains 5.9 million parameters, fewer than the Transformer-based KnotFold (6.2 M) and BPfold (7.6 M). For a typical RNA of ~200 nucleotides, trRNA2-SS generates predictions within ~5 seconds on average while consuming less than 2GB of GPU memory, making it a resource-efficient tool suitable for deployment on standard personal computers.

We further examine whether the improved secondary structure prediction could enhance the 3D structure prediction. In our previous method trRosettaRNA, SPOT-RNA was the default method for providing input secondary structures. Supplementary Fig. 1 has shown that trRNA2-SS can generate more accurate secondary structures compared to SPOT-RNA. Using these improved secondary structures as inputs, the average RMSD of trRosettaRNA on the TS28 set (the benchmark set for 3D structure prediction; see Methods) decreases by approximately 20% (from 11.35 Å to 9.07 Å). Notably, a significant majority of the RNAs (71.4%; 20 out of 28) exhibit improved 3D structure predictions (see Supplementary Fig. 1d). This result confirms that the enhanced secondary structure prediction contributes to the improved accuracy of RNA 3D structure prediction.

trRNA2-SS leverages pre-training on the large-scale bpRNA dataset to address data scarcity. Nevertheless, bpRNA suffers from redundancy and the noise inherent in predicted labels. We mitigate the first limitation by strictly filtering the dataset to ensure non-redundancy at an 80% sequence identity threshold. To assess the impact of label noise, we compare models trained on bpRNA or PDB alone (Supplementary Table 4). Consistent with the lower data quality, the model trained exclusively on bpRNA performs worse (AUPRC: 0.748; MCC: 0.731) than the PDB-only model (AUPRC: 0.787; MCC: 0.753). Crucially, however, trRNA2-SS (AUPRC: 0.822; MCC: 0.787), which employs a transfer learning strategy from bpRNA to PDB, outperforms the PDB-only baseline. This observation aligns with established findings in computer vision, where pre-training on large-scale noisy datasets enhances performance [16]. We therefore conclude that, despite its limitations, pre-training on bpRNA provides valuable inductive bias that is helpful for boosting final performance.

1.2. Application to RNAs without known secondary structures

The rapid inference powered by the deep learning model and GPUs enables the high-throughput prediction of secondary structures for a large volume of RNA sequences. As an application, we curate a dataset of 92,370 non-redundant RNAs (i.e., pairwise sequence identity < 80%) from RNAcentral without known secondary structures, strictly filtering out training set homologs (BLASTN E-value < 10 or CD-HIT identity > 80%). We then employed trRNA2-SS to generate predictions for this large-scale dataset. The prediction accuracy was estimated using the confidence score C-Score$_{\mathrm{SS}}$ defined in Eq. (8). A C-Score$_{\mathrm{SS}}$ >0.75 strongly suggests an accurate prediction (Supplementary Fig. 18).

As shown in Supplementary Fig. 2a, trRNA2-SS is more confident on the RNAs with well-defined functions, such as the regulatory pre-miRNA, the signal recognition particle RNA (SRP_RNA), and the hammerhead ribozyme, a well-known catalytic RNA. However, trRNA2-SS exhibits lower confidence for long non-coding RNAs (lncRNAs). This is consistent with the inherent characteristics of lncRNA, given that their structure and function are largely mediated by interactions with other biomolecules [17] and that they are exceedingly rare in SS-annotated databases (e.g., only 0.02% of RNAs in RNAcentral with annotated SSs are lncRNAs).

Despite this, trRNA2-SS can yield at least one confidence prediction for all related RNA types annotated by RNAcentral, including the challenging lncRNA. We present representative predictions for various RNA types in Supplementary Fig. 2b. The corresponding predicted 3D structures are shown in Supplementary Fig. 2c. These predictions can serve as valuable structural references for related research. For example, for a human lncRNA (RNAcentral ID URS00008B8B1D) sharing negligible homology with the training set (maximum sequence identity = 0.241), the trRNA2-SS prediction has a C-Score$_{\mathrm{SS}}$ of 0.756, indicating a high-confidence prediction. The secondary structure predicted by trRNA2-SS reveals a 70-nt loop, which may represent a potential binding site with other RNAs, thus providing valuable structural insight for investigating the function and mechanism of this lncRNA. To facilitate further research, we have made all confident trRNA2-SS predictions available on our website: https://yanglab.qd.sdu.edu.cn/trRosettaRNA/rnacentral/.

2. Ablation study of trRosettaRNA2

In this section, we systematically analyze the contributions of different components to the performance of trRosettaRNA2. The first part focuses on the network architecture and training strategies, while the second part evaluates the specific enhancements implemented within the structure module. For simplicity, all 3D structures were generated using the end-to-end (E2E) approach rather than the PyRosetta approach, unless otherwise specified.

2.1. Network architecture and training

This section analyzes the impact of input priors (the SS prior module and MSA embedding) and training strategies (violation fine-tuning and resampled fine-tuning; see Methods) on the performance of trRosettaRNA2. For the input priors, we performed an ablation study by selectively omitting each module and retraining the neural network while keeping all other configurations identical. For the two fine-tuning stages, we evaluated the performance using checkpoints saved immediately after each respective training phase. During inference, all ablation models were provided with identical input MSAs (with the exception of the "w/o MSA embedding" model, which utilized single sequences only).

Impact of SS prior module. As shown in Supplementary Fig. 3a, the most significant performance decline is observed when the SS prior module is removed (the SS-aware structure module also degrades to the vanilla one). The average RMSD on the TS28 set increases from 9.07 Å to 11.8 Å. Notably, for 85.7 % (24/28) of RNAs in TS28, the removal of the SS prior module leads to higher RMSD values. All other accuracy metrics, including lDDT, TM-score, and INF, also exhibit significant deterioration. While this comparison could also be impacted by the difference number of model parameters, our preliminary experiment confirms that merely increasing the blocks did not improve performance (Supplementary Table 13). This indicates the importance of the base-pairing prior knowledge embedded within the extensive secondary structure data.

Further investigation reveals weak correlations between the performance degradation and the quality of SSs generated by the SS prior module (PCC -0.266~-0.021; Supplementary Fig. 4). Notably, even for targets with seemingly poor SS predictions (e.g., PDB IDs: 8HBA and 8I3Z, whose SSs are heavily modulated by ligand binding[18]; AUPRC < 0.5), removing the SS prior module still reduced 3D modeling accuracy. Detailed analysis clarified this observation: while the predicted SSs for these targets were incomplete (recall: 0.1–0.3), they were perfectly accurate (precision: 1.0). This finding highlights that even sparse prior information remains highly valuable for modeling challenging targets, provided it is accurate.

Impact of MSA embedding. Removing the MSA embedding also exhibits a detrimental effect, although less pronounced than that of excluding the SS prior module. The average RMSD increases from 9.07 Å to 9.73 Å. For 70% of the 20 RNAs (14/20) possessing more than one effective homologous sequence (i.e., N_{eff}>1), the removal of MSA embedding worsens modeling accuracy. For the remaining 8 RNAs lacking effective homologous sequences (i.e., N_{eff}=1), the impact is less consistent: the performance degradation is observed in only half of the RNAs, with the other half exhibiting minimal improvements. This result illustrates the importance of the co-evolutionary signals implied in the MSA.

However, according to Supplementary Fig. 4, the correlations between the accuracy degradation and the MSA quality (measured by the logarithm of N_{eff}) are also weak (PCC -0.132~0.243). Considering the overall average RMSD, the introduction of MSA embedding proves beneficial for both the N_{eff} = 1 and N_{eff} > 1 groups. This finding suggests that, beyond providing useful co-evolutionary priors, MSA embedding may also contribute to training robustness, likely due to the data augmentation effect introduced by random MSA subsampling (see Methods).

Impact of violation fine-tuning. The first fine-tuning stage of trRosettaRNA2 aims to minimize steric clashes and bond violations in the predicted structures. As shown in Supplementary Fig. 3a, this stage significantly reduces clashes, with the average clashscore decreasing from 12.4 to 5.1, confirming the effectiveness of this process. The reduction of clash is more obvious for targets that initially exhibit higher clashscores (PCC = 0.944; Supplementary Fig. 4). Concurrently, the average RMSD remains largely unchanged, while both the lDDT and TM-score show slight improvements.

Remarkably, the improvement in the INF metric also shows a significant correlation with the initial clashscore (PCC=0.484). As depicted in Supplementary Fig. 4, the violation fine-tuning process results in improved interaction networks for all RNAs with an initial clashscore > 20. This may be attributed to the more realistic local structures generated by the fine-tuned model.

Impact of resampled fine-tuning. The final training stage of trRosettaRNA2 is designed to enhance performance on challenging RNAs by increasing the sampling weights for RNAs with low lDDT scores. According to Supplementary Fig. 3a, this fine-tuning stage results in improvements in RMSD, lDDT, and TM-score, while the INF and clashscore remain relatively stable. Among the 28 RNAs, 14 exhibited an RMSD > 6 Å before the resampled fine-tuning stage. After this stage, a majority (78.6%, or 11 out of 14) of these "hard" targets achieved an improvement in RMSD, a notably higher percentage compared to the remaining "easy" targets (64.3%, or 9 out of 14). This result confirms the effectiveness of our resampled fine-tuning strategy for improving the prediction accuracy of difficult RNA targets.

Case study. We reuse the example RNA (PDB ID:8YDC) to specifically illustrate the contribution of each component (Supplementary Fig. 3b). Apart from the pseudoknot, the triple helix formed by residues 3~7, 14~18, and 19~23 (indicated by gray boxes in Supplementary Fig. 3b) is also crucial for the overall 3D topology.

Without the SS prior module, the model successfully predicted the helix near the 5' terminus. However, the pseudoknot and the triple helix were not built well, resulting in an inaccurate 3D structure with an RMSD of 14.1 Å. For the "w/o MSA embedding" model, the pseudoknot was captured well, and the 3D structure was significantly more accurate with an RMSD of 6.4 Å, perhaps due to the contribution of the SS prior module. Nevertheless, the triple helix region remained incomplete, leading to an unnatural straight loop of 5 nucleotides. When both the SS prior module and MSA embedding were included, the model generated a more reasonable 3D structure that captured all key interactions, although the orientation of the triple helix was still incorrect. Finally, with the two-stage fine-tuning, the complete trRosettaRNA2 model predicts a more accurate 3D structure, achieving a lower RMSD of 6.1 Å. This example further demonstrates the effectiveness of the network architecture and training strategy employed by trRosettaRNA2.

2.2. Enhancements to the structure module and inference

This part evaluates the specific enhancements implemented within the structure module's inference protocols, including the integration of SS information (SS-aware), confidence score-guided model selection, and the recycling mechanism. All analyses were conducted using the final version of trRosettaRNA2.

SS-aware structure module. As shown in Supplementary Fig. 5a, the incorporation of SS information into the structure module improves the recovery of base pairs, particularly canonical ones. For example, in the case of a NAD-II riboswitch (PDB ID: 8HBA; Supplementary Fig. 5b), trRosettaRNA2 fails to recapitulate the native base pairs when using the vanilla structure module, yielding a median inter-base distance of >4 Å, significantly higher than the ideal value (~2.8 Å). In contrast, the SS-aware structure module accurately captures these local interactions, reducing the median distance of native base pairs to ~2.9 Å, which closely matches the ground truth.

This improvement is further visualized in a representative G-C pair (C15–G52) shown in Supplementary Fig. 5c. The model generated by the vanilla module (red) exhibits distorted geometry, characterized by an excessive distance of 5.8 Å. Conversely, the SS-aware module (blue) generates a much more physicochemically plausible base orientation, achieving an ideal distance of 2.8 Å. This observation aligns with our analysis in Section 2.1, where 8HBA was identified as a target with incomplete but accurate SS priors. This confirms that the SS-aware module can effectively leverage even partial prior information to refine 3D geometry.

Confidence score. In trRosettaRNA2, the confidence score of the predicted 3D structure is directly generated by the structure module, in the form of a predicted lDDT (pLDDT) score (predicted from the updated representation; see Algorithm 6). As illustrated in Supplementary Fig. 6a, the pLDDT score exhibits a strong correlation with the real RMSD values, achieving a Pearson correlation coefficient (PCC) of -0.812. This high correlation indicates that pLDDT serves as a reliable and reference-free metric for model selection. To demonstrate its practical utility, we employ the pLDDT score to select optimal models between trRosettaRNA2 predictions based on SPOT-RNA and trRNA2-SS. Specifically, for each target, we generated models using both SS sources and selected the one with the higher pLDDT score. The results reveal that this pLDDT-guided selection strategy effectively enhances prediction accuracy for both the E2E version and the PyRosetta version (Table 1 and Supplementary Fig. 6b). These findings highlight the potential for further improvements in trRosettaRNA2 performance through the systematic exploration of diverse inputs.

Recycling mechanism. We also observed that the recycling mechanism can help improve performance, especially when the input secondary structure is inaccurate. As shown in Supplementary Fig. 7, while recycling enhances the RMSD of predicted 3D structures across most input SS AUC levels, the most pronounced improvements occur at the low AUC levels (<0.6). This demonstrates the efficacy of the recycling mechanism in compensating for prior base-pairing information.

To illustrate the effect of the recycling mechanism more specifically, we examine an example RNA with PDB ID 8YDC, a pseudoknot-containing hammerhead ribozyme. As shown in Supplementary Fig. 8a, SPOT-RNA failed to identify several interactions crucial for the 3D topology of this RNA (e.g., the triple helix formed between residues 3~7, 14~18, and 19~23, highlighted in a dashed circle). Consequently, trRosettaRNA, using this limited SS input, could not accurately model these interactions (Supplementary Fig. 8b), resulting in a 3D structure with high RMSD (>20 Å; Supplementary Fig. 8c). In contrast, while trRosettaRNA2 initially struggled to capture SPOT-RNA-missed interactions within a single cycle, the final recycled prediction (cycle=4) successfully resolved this, resulting in a substantial RMSD reduction (Supplementary Fig. 9, RMSD dropping from ~15 Å to ~6 Å). Notably, the complete trRosettaRNA2 model, using trRNA2-SS input instead of SPOT-RNA, achieves acceptable accuracy even without recycling (~8 Å), while recycling provides only an incremental improvement (~2 Å). This observation illustrates that while enhanced prior knowledge is beneficial, trRosettaRNA2's recycling mechanism ensures model robustness when the base-pairing information from the input SS is limited or inaccurate.

3. Performance on CASP15 RNA targets

In 2022, the 15th CASP experiment (CASP15) introduced the evaluation of RNA structure prediction, comprising 8 natural and 4 synthetic RNA targets. Blind assessments in CASP15 demonstrated that trRosettaRNA achieved global accuracy (RMSD) competitive with top human groups (e.g., AIchemy_RNA2 [19]) on natural RNAs. However, its performance on local metrics (lDDT and INF) was suboptimal, revealing inaccuracies in local structural details. Furthermore, the automated modeling of synthetic RNAs, which are characterized by their intricate topologies and scarcity of homologous sequences, was a formidable challenge, evidenced by the substantial gap between automated servers and top human expert groups [20]. In this study, we benchmark trRosettaRNA2 on CASP15 targets to evaluate its capacity to overcome these limitations.

We first evaluate trRosettaRNA2 on 6 natural RNAs (excluding R1189 and R1190 due to their complex interactions with proteins that could skew the evaluation). As shown in Supplementary Table 6 and Supplementary Fig. 10a, trRosettaRNA2 consistently outperforms Yang-Server (powered by trRosettaRNA) and demonstrates competitive performance with AF3, aligning well with the results observed on the TS28 dataset. Notably, trRosettaRNA2 also surpasses the leading human group, AIchemy_RNA2, in terms of RMSD, while maintaining competitive performance across other metrics, which also represents an enhancement compared to trRosettaRNA. We highlight this enhancement using a challenging target R1116, an Enterovirus cloverleaf RNA (PDB ID: 8S95). Supplementary Fig. 10b reveals that the AIchemy_RNA2 model exhibits inaccuracies in the cloverleaf region and its orientation relative to the tRNA component, with an RMSD exceeding 20 Å. The Yang-Server model displays improved global topology (RMSD of 12 Å), but its local interaction details are inferior to AIchemy_RNA2, as indicated by lower lDDT and INF values. In contrast, the trRosettaRNA2 model demonstrates better accuracy than AIchemy_RNA2 in both global topology and local details, achieving a 49.8% lower RMSD, a 7.9% higher lDDT, and an 8.1% higher INF compared to AIchemy_RNA2. These findings further confirm the robust enhancements of trRosettaRNA2 over its predecessor, trRosettaRNA.

Regarding synthetic RNAs, trRosettaRNA2 demonstrates SOTA performance among the automated approaches in approximating global topology (indicated by its lower RMSD and higher TM-score; Supplementary Table 7). However, its overall accuracy remains inadequate and significantly lags behind the top human expert performance. Indeed, all evaluated automated methods, including AF3, consistently exhibit suboptimal performance on these RNAs. This persistent performance gap illustrates the inherent limitations of current automated approaches in capturing novel, synthetic RNA topologies, highlighting the necessity to integrate computational predictions with experimental/template priors to effectively resolve these challenging cases.

4. Comparison with two recently released methods

4.1. Comparison with DRfold2

For each target RNA, DRfold2 employs a composite language model (named RCLM; NA Composite Language Model) to embed the nucleotide sequence and uses 80 Transformer-based neural network models to predict 80 sets of backbone frames and 2D geometries. As detailed in Supplementary Table 8, the parameter count of DRfold2's prior module (RCLM) far exceeds that of trRosettaRNA2’s SS prior module (45.3 M vs. 17.7 M total). Furthermore, while the backbone of trRosettaRNA2 is slightly larger than a single DRfold2 model, DRfold2 requires aggregating 80 such models via a complex CSOR (Clustering, Scoring, Optimization, and Refinement) protocol to generate final predictions. trRosettaRNA2, in contrast, performs efficient single-model inference. To compare the predictive performance, we further filtered the TS28 set to exclude sequences redundant to the DRfold2 training set (based on a BLASTN E-value threshold of 10 and a CD-HIT-EST identity cutoff of 80%), resulting in a subset of 23 RNAs.

As shown in Supplementary Table 9, trRosettaRNA2 outperforms DRfold2 in terms of RMSD, lDDT, and INF, while being competitive on TM-score. Notably, the steric clashes in trRosettaRNA2 predictions are significantly fewer than those in DRfold2 predictions, with the average clashscores of 1.9 and 32.3, respectively. Overall, trRosettaRNA2 exhibits a more favorable balance between modeling accuracy and structural plausibility compared to DRfold2.

4.2. Comparison with RNAbpFlow

RNAbpFlow utilizes a flow-matching framework to generate 3D structures conditioned on sequence and SS information. While it benefits from the generative capabilities of flow matching, its accuracy heavily relies on ground-truth SS priors and extensive sampling. As shown in Supplementary Table 10, RNAbpFlow exhibits advantages only in a highly idealized scenario: specifically, when provided with three sets of native SSs (extracted via RNAview [21], MC-Annotate [22], and DSSR [23] from the experimental 3D structures) and evaluating the best-of-10 decoys. However, this setup is impractical for real-world applications, as neither the native SS nor the ground-truth structure (required for model selection) is available.

When benchmarked using the same predicted SS priors as trRosettaRNA2 (i.e., from trRNA2-SS) and evaluating only the top-ranked model, RNAbpFlow significantly underperforms trRosettaRNA2 (e.g., RMSD 10.37 Å vs. 8.66 Å). We attribute this performance gap to disparities in network architecture: unlike trRosettaRNA2, RNAbpFlow relies solely on a structure module where pair-wise embeddings are updated by simple MLPs, rather than employing a comprehensive encoder like RNAformer to systematically update representations. Consequently, RNAbpFlow fails to effectively capture global interactions beyond the provided SS priors, resulting in lower 3D accuracy (RMSD, lDDT, TM-score), despite achieving competitive INF. Moreover, the generated models exhibit significantly more steric clashes compared to trRosettaRNA2. These results confirm the superiority of trRosettaRNA2 over RNAbpFlow in practical predictive scenarios.

5. Generalizability analysis

To assess the generalizability of trRosettaRNA2 and trRNA2-SS, we evaluated the impact of training set similarity on model performance. Similarity was quantified across two dimensions: (1) maximum sequence identity, determined via an exhaustive BLASTN search (-evalue 1000000 -word_size 4); and (2) maximum secondary structure (SS) similarity, calculated using RNAdistance from the ViennaRNA package [4], defined as follows:

$$\mathrm{SS\_simiarity}(q, D) = \max_{t \in D} \left(1 - \frac{\left\| \mathrm{SS}_q - \mathrm{SS}_t \right\|}{L_q} \right) \tag{1}$$

where q represents the query (test sample), D denotes the training set, $\|.\|$ is the tree edit distance between the SSs of q and a training sample t calculated by RNAdistance, and L_q is the length of the query.

5.1 trRNA2-SS

As shown in Supplementary Fig. 13a-b, trRNA2-SS demonstrates robust generalizability, maintaining SOTA performance across all sequence and SS similarity levels. Notably, its performance advantage is most pronounced on 5 challenging targets that exhibit the lowest homology to the training sets (sequence identity < 30% and SS similarity < 50%), which primarily consist of synthetic RNAs such as RNA origami (4 out of 5). On these targets, trRNA2-SS achieves high AUPRC scores (0.868 on average), significantly outperforming the second-best method, SPOT-RNA2 (0.795). Furthermore, Mann-Whitney U tests confirm that, in most cases, the performance differences between similarity groups are statistically insignificant (Supplementary Fig.

13b). Regarding its contribution to 3D structure prediction, trRNA2-SS yields improvement over SPOT-RNA across all homology levels, even for the lowest homology group. However, accurately modeling the 3D topology of these low-homology RNAs remains challenging, even with improved SS priors (Supplementary Fig. 14). Overall, trRNA2-SS maintains robust SOTA performance regardless of similarity levels, confirming its commendable generalization capability.

5.2 trRosettaRNA2

A similar pattern is observed for trRosettaRNA2 in 3D structure prediction. As illustrated in Supplementary Fig. 15a, trRosettaRNA2 maintains SOTA performance across all similarity levels, with its advantage being most pronounced on the lowest homology groups. However, as noted in Supplementary Results 3 and 5.1, accurately modeling the 3D structures of the synthetic RNAs remains a formidable challenge shared by all mainstream automated methods, including both traditional and deep learning-based approaches (Supplementary Table 12).

Statistical analysis in Supplementary Fig. 15b reveals that performance differences between similarity groups are largely non-significant for trRosettaRNA2, with exceptions only in the lowest similarity group (discussed above) and the highest SS similarity group, where all methods exhibit comparable high performance. These results confirm that the competitive advantage of trRosettaRNA2 is robust across varying similarity levels, although the accurate modeling of novel, synthetic RNAs remains a common bottleneck for the field.

Supplementary Methods

1. Key contributions of trRosettaRNA2

(1) Internal pre-trained SS prior module: Unlike previous methods (trRosettaRNA and DeepFoldRNA) that merely leverage established third-party tools to provide static SS inputs, we pre-trained an RNAformer network (consisting of eight network blocks) on extensive SS datasets. This strategy aims to internalize fundamental RNA folding rules (inductive bias) and mitigate the scarcity of RNA 3D training data. Our ablation study confirms that this SS-pretrained approach is significantly more effective than simply increasing network depth (Supplementary Table 13 and Supplementary Fig. 3). This design enables trRosettaRNA2 to achieve competitive performance with AlphaFold3 using significantly fewer parameters and training resources (Fig. 2e-f). Meanwhile, this module, equipped with the sophisticated RNAformer architecture, achieves the state-of-the-art performance in SS prediction (Supplementary Fig. 1).

(2) SS-aware end-to-end prediction and efficient physical refinement: We transitioned from the "two-step" pipeline of trRosettaRNA (which predicted geometric restraints) to an end-to-end pipeline that directly predicts all-atom coordinates. Crucially, unlike other state-of-the-art end-to-end methods like AlphaFold 3 and RoseTTAFoldNA, we explicitly incorporate SS information into the structure module (i.e., the SS-aware structure module). This unique integration significantly improves the recovery of base pairs in the predicted 3D structures (see Supplementary Fig. 5). In addition, we introduced a fast relaxation protocol based on PyRosetta to resolve structural violations in the neural network outputs. We also retain an option to reconstruct 3D structures from predicted 2D restraints via energy minimization, which provides robustness for extreme cases. Consequently, trRosettaRNA2 generates models with significantly fewer steric clashes than other methods (Tables 1~2).

(3) Support for heterogeneous conformer sampling: A key feature of trRosettaRNA2 is its capability to generate heterogeneous conformations by leveraging diverse SS priors. This capability not only enabled our group (Yang-Server) to rank as the best automated server group in the CASP16 blind tests (Fig. 3) but also opened a new avenue for exploring RNA dynamics, a challenge not explored by existing methods. As demonstrated on the RNase P RNA (Fig. 4), trRosettaRNA2 effectively captures domain-specific dynamics consistent with AFM experimental observations, whereas the diffusion-based AlphaFold 3 tends to produce homogenous ensembles.

2. Pseudocode for trRosettaRNA2

As mentioned in the main text, the primary difference between the SS-aware structure module (line 8 in Algorithm 1) and the vanilla structure module (Algorithm 20 in the AlphaFold2 SI [24]) is the substitution of the standard softmax operation with an SS-weighted softmax (Eq. 2 in the main text). The distogram module (line 12 in Algorithm 1; which outputs the predicted geometric restraints) largely resembles that used in AlphaFold2 (specifically, Algorithm 29 and section 1.9.8 in its SI). Detailed pseudocode for the remaining modules is provided below.

Algorithm 1 Neural network Inference procedure of trRosettaRNA2

```
Input: multiple sequence alignment MSA
Output: predicted coordinates $\hat{\mathbf{X}}_{\mathbf{all}}$, confidence score $c_{\text{plddt}}$, predicted geometris {$\hat{\mathbf{g}}$}, predicted residue frames $\{\hat{\mathbf{T}}_j\}_{j=1}^{L}$    ▷ $\hat{\mathbf{X}}_{\text{all}}$: $(L, A, 3)$, $c_{\text{plddt}}$: scalar, $\hat{\mathbf{g}}$: $(L, L, N_{\text{bins}})$, $\hat{\mathbf{T}}_j$: $(L, 4, 3)$
1.  $\mathbf{z}_{\text{ss}}$ = SSPriorModule(MSA)    ▷ $\mathbf{z}_{\text{ss}}$: $(L, L, 1)$
2.  $\mathbf{m}_{\text{init}}$, $\mathbf{z}_{\text{msa}}$ = MSAEmbedder(MSA)    ▷ $\mathbf{m}_{\text{init}}$: $(N, L, 64)$, $\mathbf{z}_{\text{msa}}$: $(L, L, 46)$
3.  $\mathbf{z}_{\text{init}}$ = Linear(concat($\mathbf{z}_{\text{ss}}$, $\mathbf{z}_{\text{msa}}$; dim=-1))    ▷ $\mathbf{z}_{\text{init}}$: $(L, L, 64)$
4.  $\tilde{\mathbf{m}}_{\text{init}}$, $\tilde{\mathbf{z}}_{\text{init}}$ = $\mathbf{m}_{\text{init}}$, $\mathbf{z}_{\text{init}}$    ▷ initial 1D ($\tilde{\mathbf{m}}_{\text{init}}$) and 2D representations ($\tilde{\mathbf{z}}_{\text{init}}$)
5.  for $i$ in range(4):    ▷ recycle 3 times
6.      $\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$ = RNAformerStack($\tilde{\mathbf{m}}_{\text{init}}$, $\tilde{\mathbf{z}}_{\text{init}}$)    ▷ $\mathbf{m}_{\text{updated}}$: $(N, L, 64)$, $\mathbf{z}_{\text{updated}}$: $(L, L, 64)$
7.      $\mathbf{s}_{\text{init}}$ = $\mathbf{m}_{\text{updated}}[0]$    ▷ $\mathbf{s}_{\text{init}}$: $(L, 64)$
8.      $\hat{\mathbf{X}}_{\text{all}}$, $\hat{\mathbf{X}}_{\text{C1'}}$, $\{\hat{\mathbf{T}}_j\}_{j=1}^{L}$, $\mathbf{s}_{\text{updated}}$ = SSAwareStructureModule($\mathbf{s}_{\text{init}}$, $\mathbf{z}_{\text{updated}}$, $\mathbf{z}_{\text{ss}}$)    ▷ $\hat{\mathbf{X}}_{\text{all}}$: $(L, A, 3)$, $\hat{\mathbf{X}}_{\text{C1'}}$: $(L, 3)$, $\hat{\mathbf{T}}_j$: $(L, 4, 3)$, $\mathbf{s}_{\text{updated}}$: $(L, 64)$
9.      if $i$<3:
10.         ($\tilde{\mathbf{m}}_{\text{init}}[0]$, $\tilde{\mathbf{z}}_{\text{init}}$) = ($\mathbf{m}_{\text{init}}[0]$, $\mathbf{z}_{\text{init}}$) + RecycleEmbedder($\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$)
11. $c_{\text{plddt}}$ = ConfidenceModule($\mathbf{s}_{\text{updated}}$)    ▷ $c_{\text{plddt}}$: scalar
12. {$\hat{\mathbf{g}}$} = DistogramModule($\mathbf{z}_{\text{updated}}$)    ▷ $\hat{\mathbf{g}}$: $(L, L, N_{\text{bins}})$
```

Note: N=Number of sequences in MSA, L=Length of input sequence, A=number of atoms per nucleotide, N_{bins}=number of bins (38 for distances, 1 for contact).

Algorithm 2 MSAEmbedder

```
Input: multiple sequence alignment MSA
Output: initial MSA representations $\mathbf{m}_{\text{init}}$, MSA features $\mathbf{z}_{\text{msa}}$    ▷ $\mathbf{m}_{\text{init}}$: $(N, L, 64)$, $\mathbf{z}_{\text{msa}}$: $(L, L, 46)$
$\mathbf{m}_{\text{raw}}$ = Tokenization(MSA)    ▷ $\mathbf{m}_{\text{raw}}$: $(N, L)$
$\mathbf{m}_{\text{init}}$ = Embedding($\mathbf{m}_{\text{raw}}$)    ▷ $\mathbf{m}_{\text{init}}$: $(N, L, 64)$
$\mathbf{s}_{\text{1hot}}$ = one_hot($\mathbf{m}_{\text{raw}}[0]$)[:, :4]    ▷ $\mathbf{s}_{\text{1hot}}$: $(L, 4)$
$\mathbf{s}_{\text{pssm}}$ = PositionSpecificScoringMatrixWithIC(MSA)    ▷ $\mathbf{s}_{\text{pssm}}$: $(L, 6)$
$\mathbf{s}_{\text{all}}$ = concat($\mathbf{s}_{\text{1hot}}$, $\mathbf{s}_{\text{pssm}}$; dim=-1)    ▷ $\mathbf{s}_{\text{all}}$: $(L, 10)$
$\mathbf{z}_{\text{dca}}$ = DirectCouplingAnalysis(MSA)    ▷ $\mathbf{z}_{\text{dca}}$: $(L, L, 26)$
$\mathbf{z}_{\text{msa}}$=concat(outer_concat($\mathbf{s}_{\text{all}}$, $\mathbf{s}_{\text{all}}$), $\mathbf{z}_{\text{dca}}$)    ▷ $\mathbf{z}_{\text{msa}}$: $(L, L, 46)$
```

Note: MSAs are generated by searching for homologs in the RNAcentral database using a combined approach with BLASTN [25] and Infernal [26]. BLASTN is employed to rapidly generate an initial MSA via efficient heuristic searching. Subsequently, Infernal utilizes this initial alignment to perform a covariance model-based search, allowing for the detection of remote homologs that might be missed by sequence-only methods like BLASTN. Details regarding the Position Specific Scoring Matrix (PSSM) and Information Content (IC) are available at https://en.wikipedia.org/wiki/Position_weight_matrix, while the Direct Coupling Analysis (DCA) implementation follows the protocol outlined in [27].

Algorithm 3 SSPriorModule

```
Input: multiple sequence alignment MSA; number of models $N_{\text{model}}$
Output: predicted base-pairing probability map $\mathbf{z}_{\text{ss}}$        ▷ $\mathbf{z}_{\text{ss}}$: $(L, L, 1)$
all_preds = []
for _ in range($N_{\text{model}}$):
    $\mathbf{m}_{\text{init}}$, $\mathbf{z}_{\text{msa}}$ = MSAEmbedder(MSA)        ▷ $\mathbf{m}_{\text{init}}$: $(N, L, 64)$, $\mathbf{z}_{\text{msa}}$: $(L, L, 46)$
    $\mathbf{z}_{\text{init}}$ = Linear($\mathbf{z}_{\text{ss}}$)        ▷ $\mathbf{z}_{\text{init}}$: $(L, L, 64)$
    $\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$ = RNAformerStack($\tilde{\mathbf{m}}_{\text{init}}$, $\tilde{\mathbf{z}}_{\text{init}}$; $N_{\text{block}}$=8)        ▷ $\mathbf{m}_{\text{updated}}$: $(N, L, 64)$, $\mathbf{z}_{\text{updated}}$: $(L, L, 64)$
    $\mathbf{z}_{\text{symm}}$ = $\mathbf{z}_{\text{updated}}$ + permute($\mathbf{z}_{\text{updated}}$, [1,0,2])
    all_preds.append(sigmoid(Linear($\mathbf{z}_{\text{symm}}$)))
$\mathbf{z}_{\text{ss}}$ = mean(all_preds)        ▷ $\mathbf{z}_{\text{ss}}$: $(L, L, 1)$
```

Algorithm 4 RNAformerStack

```
Input: initial representations $\mathbf{m}_{\text{init}}$, $\mathbf{z}_{\text{init}}$; Number of blocks $N_{\text{block}}$        ▷ $\mathbf{m}_{\text{init}}$: $(N, L, 64)$, $\mathbf{z}_{\text{init}}$: $(L, L, 64)$
Output: updated representations $\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$        ▷ $\mathbf{m}_{\text{updated}}$: $(N, L, 64)$, $\mathbf{z}_{\text{updated}}$: $(L, L, 64)$
$\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$ = $\mathbf{m}_{\text{init}}$, $\mathbf{z}_{\text{init}}$
for _ in range($N_{\text{block}}$):
    $\mathbf{m}_{\text{updated}}$ += MSARowAttentionWithPairBias($\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$)        ▷ pair to MSA
    $\mathbf{m}_{\text{updated}}$ += MSAColumnAttention($\mathbf{m}_{\text{updated}}$)
    $\mathbf{m}_{\text{updated}}$ += MSATransition($\mathbf{m}_{\text{updated}}$)
    $\mathbf{z}_{\text{updated}}$ += OuterProductMean($\mathbf{m}_{\text{updated}}$)        ▷ MSA to pair
    $\mathbf{z}_{\text{updated}}$ += TriangleMultiplicationOutgoing($\mathbf{z}_{\text{updated}}$) + Res2NetBlock($\mathbf{z}_{\text{updated}}$)
    $\mathbf{z}_{\text{updated}}$ += TriangleMultiplicationIncoming($\mathbf{z}_{\text{updated}}$) + Res2NetBlock($\mathbf{z}_{\text{updated}}$)
    $\mathbf{z}_{\text{updated}}$ += TriangleAttentionStartingNode($\mathbf{z}_{\text{updated}}$) + Res2NetBlock($\mathbf{z}_{\text{updated}}$)
    $\mathbf{z}_{\text{updated}}$ += TriangleAttentionEndingNode($\mathbf{z}_{\text{updated}}$) + Res2NetBlock($\mathbf{z}_{\text{updated}}$)
    $\mathbf{z}_{\text{updated}}$ += PairTransition($\mathbf{z}_{\text{updated}}$)
```

Note: More details about the used functions in Algorithm 4 can be found in AlphaFold2 SI (Algorithms 7~15) and Res2Net paper[15].

Algorithm 5 RecycleEmbedder

```
Input: updated representations $\mathbf{m}_{\text{updated}}$, $\mathbf{z}_{\text{updated}}$, predicted C1' coordinates $\hat{\mathbf{X}}_{\text{C1'}}$        ▷ $\mathbf{m}_{\text{updated}}$: $(N, L, 64)$, $\mathbf{z}_{\text{updated}}$: $(L, L, 64)$, $\hat{\mathbf{X}}_{\text{C1'}}$: $(L, 3)$
Output: biases $\Delta\mathbf{m}$, $\Delta\mathbf{z}$        ▷ $\Delta\mathbf{m}$: $(L, 64)$, $\Delta\mathbf{z}$: $(L, L, 64)$
$\mathbf{d}$ = pairwise_distance($\hat{\mathbf{X}}_{\text{C1'}}$)        ▷ $\mathbf{d}$: $(L, L)$
$\mathbf{d}$ = Linear(one_hot($\mathbf{d}$, $\mathbf{v}_{\text{bins}}$=[3Å, 4Å,…,40Å]))        ▷ $\mathbf{d}$: $(L, L, 64)$
$\Delta\mathbf{z}$ = $\mathbf{d}$ + LayerNorm($\mathbf{z}_{\text{updated}}$)        ▷ $\Delta\mathbf{z}$: $(L, L, 64)$
$\Delta\mathbf{m}$ = LayerNorm($\mathbf{m}_{\text{updated}}$[0])        ▷ $\Delta\mathbf{m}$: $(L, 64)$
```

Algorithm 6 ConfidenceModule

```
Input: updated single representation $\mathbf{s}_{\text{updated}}$        ▷ $\mathbf{s}_{\text{updated}}$: $(L, 64)$
Output: confidence score (pLDDT) $c_{\text{plddt}}$        ▷ $c_{\text{plddt}}$: scalar
$\mathbf{a}$ = Linear(relu(Linear(relu(Linear(LayerNorm($\mathbf{s}_{\text{updated}}$))))))        ▷ $\mathbf{a}_{\text{updated}}$: $(L, 50)$
$\mathbf{p}_{\text{plddt}}$ = softmax($\mathbf{a}$, dim=-1)        ▷ $\mathbf{p}_{\text{plddt}}$: $(L, 50)$
$\mathbf{v}_{\text{bins}}$ = [0.01, 0.03, 0.05, …, 0.99]        ▷ $\mathbf{v}_{\text{bins}}$: $(50)$
$\mathbf{r}_{\text{plddt}}$ = $\mathbf{p}_{\text{plddt}}$ @ $\mathbf{v}_{\text{bins}}$        ▷ $\mathbf{r}_{\text{plddt}}$: $(L)$
$c_{\text{plddt}}$ = mean($\mathbf{r}_{\text{plddt}}$)
```

3. Loss functions

Frame-aligned-position-error (FAPE) loss $\mathcal{L}_{\mathbf{FAPE}}$: Adapted from AlphaFold2, this loss penalizes the deviation between predicted and ground-truth atom positions (denoted by $\hat{\mathbf{X}} = \{\hat{\mathbf{x}}_i\}_{i=1}^{A}$ and $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{A}$, respectively; A is the number of atoms) after aligned by the local residue frames (i.e., the orientation and position of residues). Let $\hat{\mathbf{T}}_j$ and $\mathbf{T}_j$ represent the predicted and ground-truth frames of residue j, then this loss is defined as:

$$\mathcal{L}_{\text{FAPE}} = \frac{1}{A \cdot L} \sum_{i=1}^{A} \sum_{j=1}^{L} \frac{1}{Z} \min\left(d_{\text{clamp}}, \sqrt{\left\| \mathbf{T}_j \circ \mathbf{x}_i - \hat{\mathbf{T}}_j \circ \hat{\mathbf{x}}_i \right\|^2 + \varepsilon}\right) \tag{2}$$

where L is sequence length (i.e., number of residues); Z =10 Å is a normalization constant; d_{clamp} is a clamping threshold set to 10 Å with a probability of 90% and $+\infty$ with a probability of 10%; ||.|| is the L2 norm; and ε=10^{-6} Å is a small constant added to prevent division by zero during gradient computation.

2D geometries loss $\mathcal{L}_{\mathbf{geo2D}}$: This loss item supervises the prediction of 2D geometries. In trRosettaRNA2, the predicted geometries encompass inter-nucleotide distances and contacts (defined as a minimum heavy-atom distance < 8 Å). Specifically, distances are calculated for six atom pairs: C3', P, C1', C4', N1 (for pyrimidines) / N9 (for purines), and C2 (for pyrimidines) / C4 (for purines).

For distance prediction, rather than directly regressing continuous values, we discretize the distances into 38 bins: 37 bins covering the range from 3 Å to 40 Å (with a bin size of 1 Å), plus an additional bin representing distances exceeding this range. Consequently, the cross-entropy loss is employed for both distance and contact predictions. The loss for each distance type is defined as:

$$\mathcal{L}_{\text{dist}} = -\frac{1}{L^2} \sum_{i=1}^{L} \sum_{j=1}^{L} \sum_{b=1}^{38} y_{ij}^{b} \log p_{ij}^{b} \tag{3}$$

The contact prediction is a binary classification task, whose loss can be written as:

$$\mathcal{L}_{\text{contact}} = -\frac{1}{L^2} \sum_{i=1}^{L} \sum_{j=1}^{L} \left[y_{ij} \log p_{ij} + \left(1 - y_{ij}\right) \log\left(1 - p_{ij}\right) \right] \tag{4}$$

Among Eqs. (3) and (4), y denotes the ground-truth labels (one-hot encoding for distance; binary label for contact), while p denotes the corresponding predicted probabilities. The final $\mathcal{L}_{\text{geo2D}}$ is calculated as the arithmetic mean of these seven loss terms (six distance losses and one contact loss).

Torsion angle loss $\mathcal{L}_{\mathbf{torsion}}$: This loss supervises the prediction of intra-nucleotide torsion angles. We adopt the definition from RhoFold+, which utilizes a set of four torsion angles. Although this constitutes a simplified representation of nucleotide geometry, our comprehensive structural optimization procedure ensures the generation of physicochemically plausible 3D structures (Tables 1~2 and Supplementary Table 5). Formally, this loss is defined as the Mean Squared Error (MSE) between the predicted and ground-truth trigonometric projections (sine and cosine values) of the torsion angles (following Algorithm 27 in the AlphaFold2 SI).

pLDDT loss $\mathcal{L}_{\mathbf{plddt}}$: This loss supervises the prediction of the per-residue confidence score (pLDDT). Specifically, we discretize the ground-truth lDDT values (ranging from 0 to 1) into 50 bins (bin width = 0.02) and predict the probability distribution over these bins. The objective is minimized using a multi-class cross-entropy loss.

Base pair loss $\mathcal{L}_{\mathbf{BP}}$: This term constrains the geometry of predicted base pairs based on native structural information. Specifically, for every native base pair, the inter-atom distances in the predicted structure are enforced to remain within the interval [1.5 Å, d+3 Å], where d is the corresponding inter-atom distance in the experimental structure. Any deviation outside this boundary incurs an L1 penalty.

Bond loss $\mathcal{L}_{\mathbf{bond}}$: To ensure stereochemical quality, this loss restricts backbone geometry, specifically the O3'–P bond lengths, P–P distances, and O3'–P–O5' angles (measured in radians) between neighboring residues, to statistically derived ideal ranges obtained from the training set. Values falling outside these valid ranges are penalized via an L1 loss.

Clash loss $\mathcal{L}_{\mathbf{clash}}$: This objective is designed to resolve steric clashes. It applies an L1 penalty to any inter-atom distance that falls below a dynamic threshold d_0 = min(d-0.1 Å,2 Å), where d is the corresponding distance in the experimental structure, thereby discouraging steric clashes.

Supplementary Tables

Supplementary Table 1. Benchmark results of RNA secondary structure (SS) prediction on the ArchiveII test set. Note that although this dataset was filtered to exclude sequences similar to our training set, it may still share varying degrees of overlap with the training sets of other methods. For trRNA2-SS, binary predictions were derived from the probability maps using a standard cutoff of 0.5 to compute MCC and F1-score.

Classification	Method	AUPRC	MCC	F1-score	precision	recall
Traditional	RNAfold	0.556	0.552	0.552	0.520	0.590
	RNAstructure	0.545	0.486	0.443	0.333	0.756
	PETfold	0.787	0.760	0.758	<u>0.734</u>	0.793
	EternaFold	0.677	0.595	0.595	0.561	0.637
Deep learning-based	SPOT-RNA	0.849	0.781	0.775	0.722	0.853
	SPOT-RNA2	0.772	0.755	0.753	0.699	0.820
	UFold	0.630	0.677	0.676	0.630	0.733
	MXfold2	0.674	0.659	0.659	0.642	0.681
	RNA-FM	0.752	0.744	0.734	0.635	0.876
	RNA-MSM	0.806	<u>0.792</u>	<u>0.784</u>	0.689	0.916
	KnotFold	0.852	0.767	0.765	0.709	0.835
	BPfold	**0.959**	**0.958**	**0.958**	**0.956**	**0.962**
	trRNA2-SS	<u>0.874</u>	0.791	0.780	0.670	<u>0.937</u>

Supplementary Table 2. Benchmark results of RNA secondary structure (SS) prediction on the PDB2D test set. For trRNA2-SS, binary predictions were derived from the probability maps using a standard cutoff of 0.5 to compute MCC and F1-score.

Classification	**Method**	**AUPRC**	**MCC**	**F1-score**	**precision**	**recall**
Traditional	RNAfold	0.766	0.745	0.733	0.889	0.637
	RNAstructure	0.731	0.709	0.697	0.805	0.651
	PETfold	0.604	0.701	0.678	0.904	0.561
	EternaFold	0.678	0.715	0.705	0.846	0.619
Deep learning-based	SPOT-RNA	0.731	0.731	0.716	0.886	0.621
	SPOT-RNA2	0.774	0.773	0.765	0.882	**0.692**
	UFold	0.561	0.599	0.591	0.716	0.516
	MXfold2	0.637	0.725	0.713	0.873	0.615
	RNA-FM	0.719	0.730	0.722	0.851	0.639
	RNA-MSM	0.532	0.491	0.442	0.814	0.350
	KnotFold	0.666	0.685	0.660	0.888	0.557
	BPfold	0.723	0.696	0.680	0.862	0.576
	trRNA2-SS	**0.822**	**0.788**	**0.771**	**0.944**	0.682

Supplementary Table 3. Comparison of model architecture, complexity and computational overhead for the evaluated deep learning SS predictors. All benchmarks were performed using an NVIDIA A800 GPU and an Intel(R) Xeon(R) Gold 6342 CPU. Metrics were measured using a 224-nucleotide RNA sequence and averaged over 10 independent runs. Inference times are presented as mean ± standard deviation.

Method	**Architecture**	**Parameters per model (M)**	**Inference time (s)**	**Peak GPU memory (GB)**	**Peak CPU memory (GB)**
SPOT-RNA*	BiLSTM+ResNet	1.5~4.4 (5 models)	47.2 ±1.16	\	1.88
SPOT-RNA2*	ResNet	3.4~5.0 (4 models)	45.3 ±0.68	\	1.74
MXfold2*	BiLSTM+ResNet	0.8 (1 model)	4.2 ±0.14	\	0.26
Ufold	U-Net	8.2 (1 model)	2.2 ±0.04	0.63	0.87
KnotFold	Transformer	6.2 (5 models)	1.8 ±0.08	0.79	0.84
BPfold	CNN-enhanced Transformer	7.6 (6 models)	12.5 ±0.31	0.90	1.07
RNA-FM	Language model: Transformer	94.9 (1 model; fine-tuned)	4.9 ±0.19	9.28	1.04
	Downstream head: ResNet	2.3 (1 model)			
RNA-MSM	Language model: Axial Transformer	9.15 (1 model; frozen)	3.1 ±0.08	3.50	1.01
	Downstream head: ResNet	1.3 (1 model)			
trRNA2-SS	RNAformer (two-track axial Transformer + Res2Net)	5.9 (3 models)	5.4 ±0.05	1.95	0.93

* These methods run exclusively on the CPU due to incompatibility with our GPU.

Supplementary Table 4. Impact of training datasets on trRNA2-SS performance evaluated on the PDB2D test set. To ensure consistency with the trRNA2-SS protocol, results for ablation models are derived from an ensemble of three independently trained replicates.

Training sets	AUPRC	MCC	F1-score	precision	recall
bpRNA only	0.748	0.731	0.711	0.925	0.596
PDB only	0.787	0.753	0.740	0.889	0.655
bpRNA+PDB	**0.822**	**0.788**	**0.771**	**0.944**	**0.682**

Supplementary Table 5. Geometric entanglement analysis of AlphaFold 3 and trRosettaRNA2 predictions (TS28 dataset), assessed via the RNAspider webserver [28].

PDB ID	AlphaFold 3	trRosettaRNA2	PDB ID	AlphaFold 3	trRosettaRNA2
7KUB	0	0	8ITS	0	0
7KVT	0	0	8TDY	0	0
7SHX	0	0	8TE2	0	0
7UGA	0	0	8TQX	0	0
7UMC	0	0	8TSV	0	0
7UR5	0	0	8UPT	0	0
7V06	0	0	8V1H	0	0
7WIE	0	0	8VFS	0	0
8EYU	0	0	8VT5	0	0
8HBA	0	0	8XZE	0	0
8I3Z	0	0	8YDC	0	0
8SCH	0	0	9BLM	0	0
8BU8	**1**	0	9BUN	**2**	0
8FER	0	0	9G7C	0	0

Supplementary Table 6. Comparisons on six natural RNAs from CASP15. For the evaluation of AIchemy_RNA, DF_RNA, and BAKER, we use the same configuration as described in the trRosettaRNA paper. Specifically, we prioritize the evaluation of their submitted models during CASP15, and in cases where a submitted model is unavailable, we replace it with a locally generated model via the corresponding algorithms.

Method	**RMSD (Å)** ↓	**lDDT** ↑	**TM-score** ↑	**INF** ↑	**clashscore** ↓
AIchemy_RNA ^{H} (RhoFold)	10.86	0.663	0.354	0.763	21.0
RhoFold+	11.65	0.633	0.321	0.687	32.3
DF_RNA ^{H} (DeepFoldRNA)	11.56	0.673	0.341	0.744	48.7
BAKER ^{H} (RoseTTAFoldNA)	13.81	0.613	0.302	0.734	28.9
AlphaFold 3	9.47	**0.754**	**0.428**	**0.834**	16.5
Yang-Server^{S} (trRosettaRNA)	12.30	0.599	0.319	0.699	26.9
AIchemy_RNA2 ^{H}	13.39	0.718	0.387	0.821	12.7
trRosettaRNA2	**9.19**	0.708	0.399	0.791	**2.6**

^{H}: Human group; ^{S}: Server group.

Supplementary Table 7. Comparisons on four synthetic RNAs from CASP15.

Group/Method	**RMSD (Å)** ↓	**lDDT** ↑	**TM-score** ↑	**INF** ↑	**clashscore** ↓
AIchemy_RNA^{H} (RhoFold)	40.3	0.492	0.236	0.788	10.0
RhoFold+	46.3	0.309	0.168	0.212	235.8
DF_RNA^{H} (DeepFoldRNA)	54.1	0.486	0.211	0.777	63.4
BAKER^{H} (RoseTTAFoldNA)	40.1	0.421	0.221	0.706	**0.2**
AlphaFold 3	36.5	0.520	0.252	0.834	8.6
Yang-Server^{S} (trRosettaRNA)	39.0	0.468	0.213	0.563	165.4
AIchemy_RNA2^{H}	**11.3**	**0.753**	**0.692**	**0.884**	8.7
trRosettaRNA2	29.1	0.492	0.278	0.689	14.9

Supplementary Table 8. Complexity comparison between trRosettaRNA2 and DRfold2. This table presents the number of parameters of the two methods.

Method	Prior module	Main module
trRosettaRNA2	5.9 M (3 trRNA2-SS models)	9.3 M (1 model)
DRfold2	45.3 M (1 RCLM model)	3.2~4.0 M (80 models)

Supplementary Table 9. Comparison between trRosettaRNA2 and DRfold2 on 23 RNAs from TS28 which are non-redundant to the training set of DRfold2.

Method	SS predictor	RMSD (Å) ↓	lDDT ↑	TM-score ↑	INF ↑	clashscore ↓
trRosettaRNA2 (PyRosetta)	SPOT-RNA	9.89	0.686	0.311	0.759	2.7
	trRNA2-SS	9.34	0.697	0.314	0.772	**1.9**
	pLDDT top1	**9.00**	**0.700**	0.328	**0.777**	2.3
DRfold2	\	10.57	0.693	**0.351**	0.769	32.3

Supplementary Table 10. Comparison between trRosettaRNA2 and RNAbpFlow on TS28. To ensure a fair comparison, the "CASP15" version of RNAbpFlow (trained on RNAs released before 2022) was used.

Method	SS source	RMSD (Å) ↓	lDDT ↑	TM-score ↑	INF ↑	clashscore ↓
trRosettaRNA2 (PyRosetta)	trRNA2-SS	8.66	0.695	0.364	0.765	2.0
RNAbpFlow (model 1)	trRNA2-SS	10.37	0.641	0.319	0.764	71.0
	native	8.56	0.717	0.346	0.823	91.0
RNAbpFlow (best-of-10)	trRNA2-SS	8.30	0.673	0.370	0.788	46.8
	native	6.92	0.747	0.402	0.845	65.5

Supplementary Table 11. Comparisons of Yang-Server, AF3-Server, and dNAfold on 23 CASP16 targets shorter than 400 nucleotides or with templates.

Group	Model	RMSD (Å) ↓	lDDT ↑	TM-score ↑	GDT-TS ↑	INF ↑	clashscore ↓
Yang-Server	First model	10.6	0.634	0.530	0.483	0.801	3.4
	Best model	**7.8**	0.663	**0.573**	**0.518**	0.826	**1.5**
AF3-server	First model	10.8	0.658	0.511	0.473	0.818	14.6
	Best model	9.8	**0.676**	0.535	0.494	**0.829**	10.6
dNAfold (DRfold2)	First model	11.2	0.369	0.528	0.483	0.749	35.7
	Best model	10.2	0.486	0.535	0.490	0.795	23.7

Supplementary Table 12. RMSD (Å) comparison for the six challenging RNA targets exhibiting the lowest similarity to the training set.

Methods	8BU8	9G7C	R1126	R1128	R1136	R1138
FarFar2	48.5	33.4	45.9	44.5	44.8	60.6
RoseTTAFoldNA	37.4	33.9	33.3	22.1	43.5	62.2
DeepFoldRNA	43.6	34.5	45.5	48.3	50.4	72.1
RhoFold+	53.1	35.5	46.7	**17.1**	52.4	69.2
AlphaFold 3	38.6	**29.5**	34.7	33.7	42.4	**35.3**
trRosettaRNA	45.6	41.6	38.1	22.3	46.2	50.5
trRosettaRNA2	**18.4**	36.8	**25.8**	26.7	**28.0**	36.1

Supplementary Table 13. Comparisons of trRosettaRNA (48 blocks) and its lightweight version (12 blocks) on the TS28 set.

Blocks	RMSD (Å) ↓	lDDT ↑	TM-score ↑	INF ↑	clashscore ↓
48	11.35	0.641	0.336	**0.723**	**3.4**
12	**11.05**	**0.649**	**0.338**	0.706	12.0

Supplementary Table 14. Training details of trRosettaRNA2.

Parameter	Stage 1: initial training	Stage 2: violation fine-tuning	Stage 3: resampled fine-tuning
Sequence crop size	256	384	
Initial parameters	Random	From stage 1	From stage 2
Initial learning rate	0.0002	0.00005	0.00007
Training sample*	3,665/4,992 clusters (8,598/10,699 RNAs)		8,598/10,699 RNAs
Training sampler	By sequence cluster (identical weights for all clusters)		8,000 samples (weight: reciprocal of the lDDT in the previous epoch)
Training times	~5 days	~3 days	~4 days
Training resource	1 A100 40GB GPU	1 A800 GPU	

* The numbers preceding the "/" indicate the subset used for benchmarking, which includes RNAs released before 2022, while the numbers following the "/" represent the total number of training RNAs.

Supplementary Table 15. Comparison of SS AUPRC on three CASP16 targets with >600 nucleotides. UFold is not included as it does not support inputs with >600 nucleotides.

Method	R1250 (744 nt)	R1251 (833 nt)	R1281 (718 nt)
RNAfold	0.472	0.522	0.848
RNAstructure	0.457	0.525	0.853
PETfold	0.129	0.113	0.071
EternaFold	0.407	0.468	0.794
SPOT-RNA	0.309	0.299	0.755
SPOT-RNA2	0.406	0.580	0.808
MXfold2	0.236	0.469	0.763
RNA-FM	0.306	0.328	0.678
RNA-MSM	0.283	0.270	0.190
KnotFold	0.289	0.343	0.575
BPfold	0.462	0.512	0.745
trRNA2-SS	**0.669**	**0.797**	**0.865**

Supplementary Table 16. Information about the datasets used in this work.

Type	Name	Size	Usage
Training sets	PDB	10,699	Training trRNA2-SS and trRosettaRNA2
	bpRNA	14,648	Pre-training trRNA2-SS
Test sets	TS28	28	Benchmarking trRosettaRNA2
	CASP15	10	
	ArchiveII	50	Benchmarking trRNA2-SS
	PDB2D	19	

Supplementary Table 17. Information of the RNAs for which RhoFold+ fails to perform relaxation.

Source	ID	Length	Clashscore of RhoFold+ unrelaxed model
CASP15	R1126	363	156.5
	R1136	374	186.5
	R1138	720	137.5
CASP16	R1248	407	579.0
	R1283	580	432.2
	R1286	526	222.2

Supplementary Figures

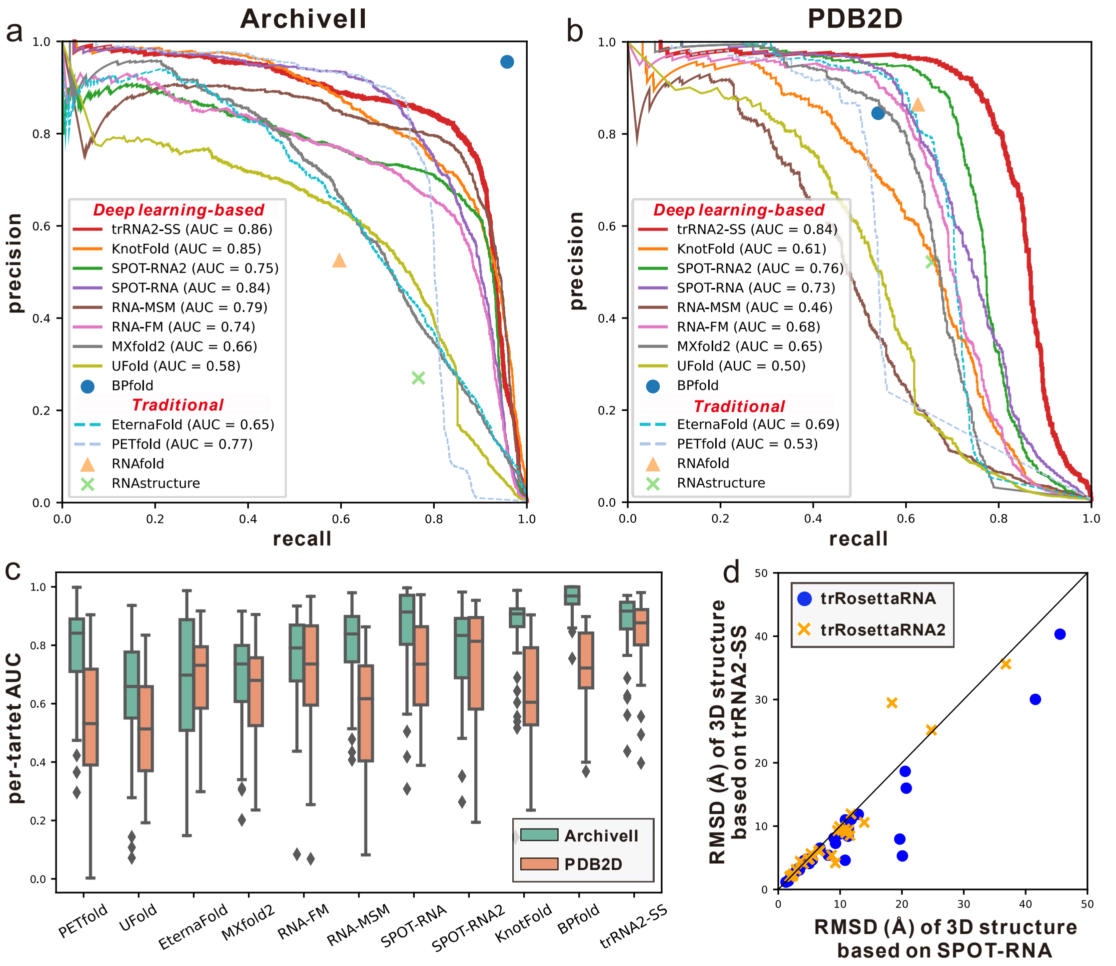

Supplementary Fig. 1. Performance of trRNA2-SS for RNA secondary structure prediction. (a, b) Precision-recall curves on ArchiveII (*n*=50) and PDB2D (*n*=19) test sets. (c) The distributions of per-target PR-AUC values. Each box encloses the middle 50% of the data (the Interquartile Range, IQR), with the bottom edge at the 25th percentile (Q1) and the top edge at the 75th percentile (Q3). The line inside the box indicates the median (50th percentile). The whiskers extend to the 1.5 * IQR from the box edges. Data points plotted individually outside the whiskers are outliers. (d) Head-to-head comparisons between the 3D modeling results based on SPOT-RNA and trRNA2-SS. Note that while the ArchiveII set has been filtered to guarantee non-redundancy with our training set, it inevitably overlaps with the training sets of other deep learning-based methods (especially BPfold), which may lead to performance inflation for these methods.

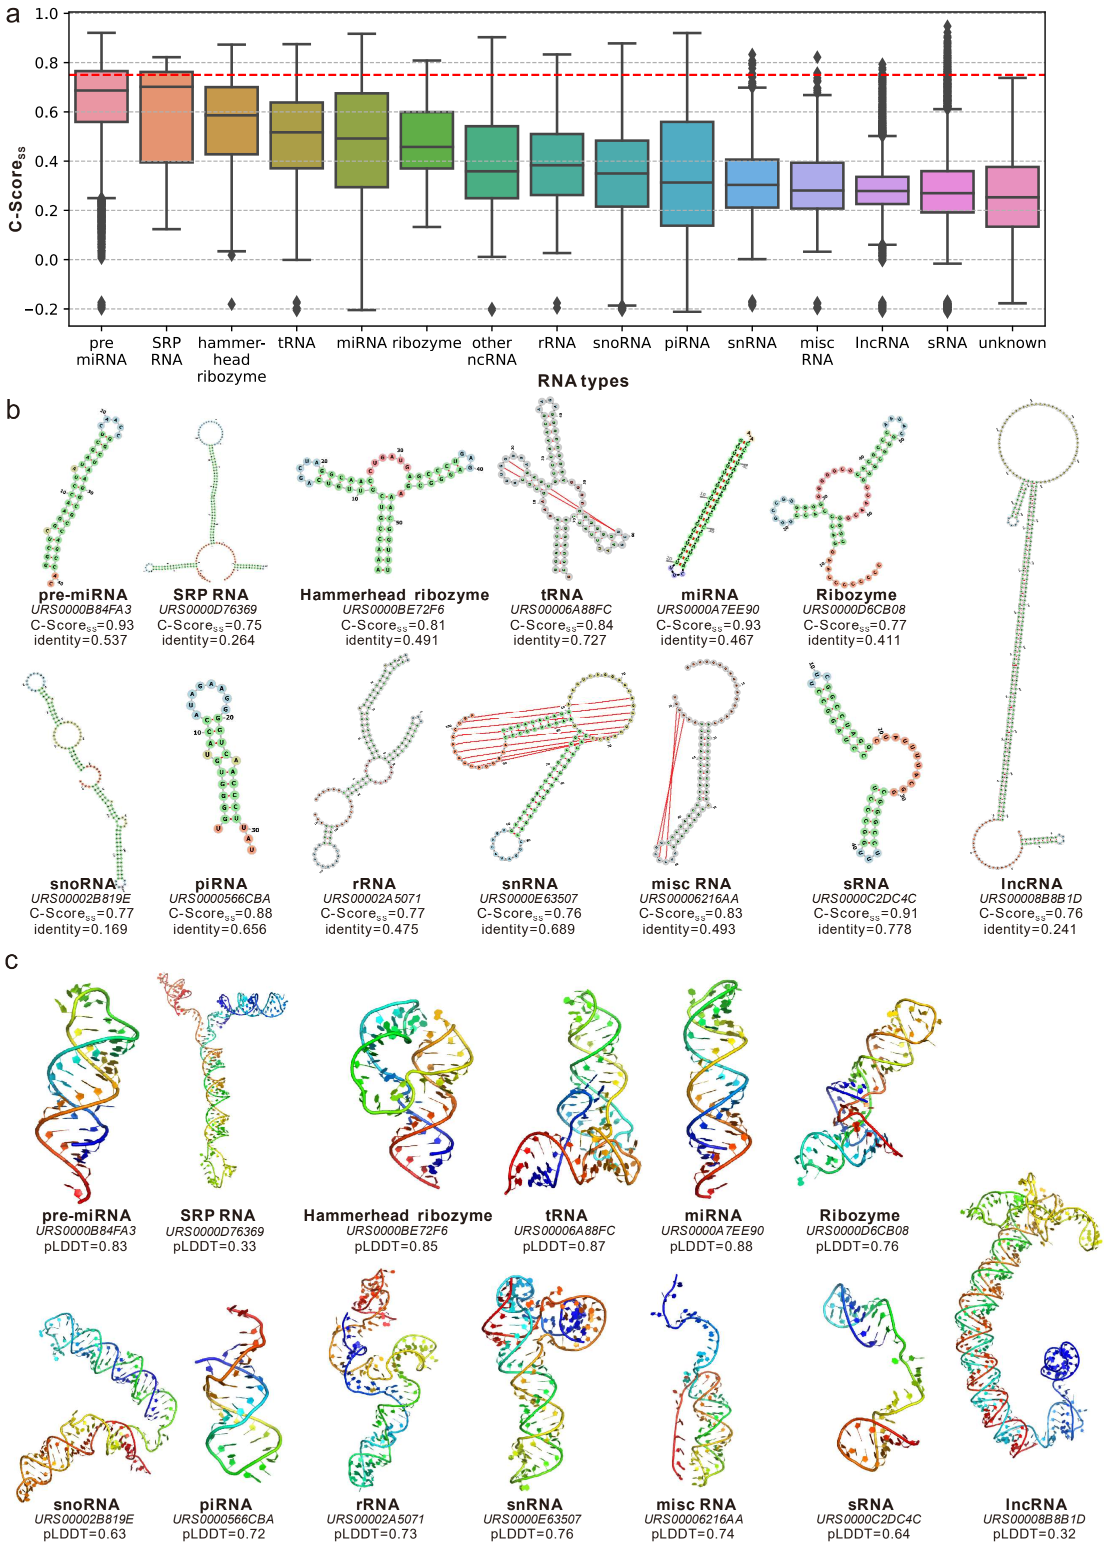

Supplementary Fig. 2. Application of trRNA2-SS to RNAcentral sequences without known secondary structures. (a) Distributions of the confidence score for predicted secondary structures (C-Score$_{SS}$) on different RNA types (*n*=92,370). The red dashed line indicates the threshold (0.75) for a "confident" prediction. Each box plot displays the median (center line), the IQR as the box, and whiskers extending 1.5 × IQR from the box's edges. Outliers are shown as individual points beyond the whiskers. (b) Representative examples of predicted secondary structures and (c) their trRosettaRNA2 3D models.

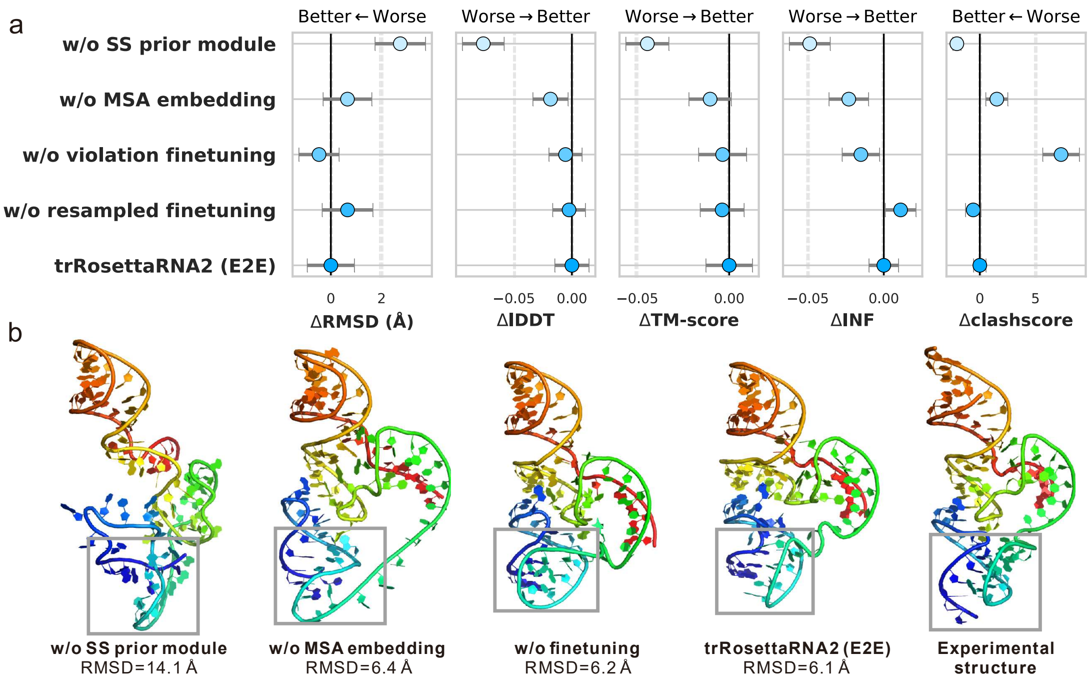

Supplementary Fig. 3. Ablation study of trRosettaRNA2. (a) Quantitative evaluation of component contributions on the TS28 dataset (*n*=28). Metrics are presented as the mean difference (Δ) relative to the full trRosettaRNA2 (E2E) model, which serves as the baseline (black dashed line at 0). Arrows at the top indicate the direction of the different metrics. Error bars denote one-tenth of the standard deviation of the metric. (b) 3D modeling results for the example RNA (PDB ID: 8YDC). The gray boxes highlight the triple helix, which is crucial for the 3D topology of this RNA. Note that the RMSD values are calculated for the entire structures rather than for the highlighted triple helix regions.

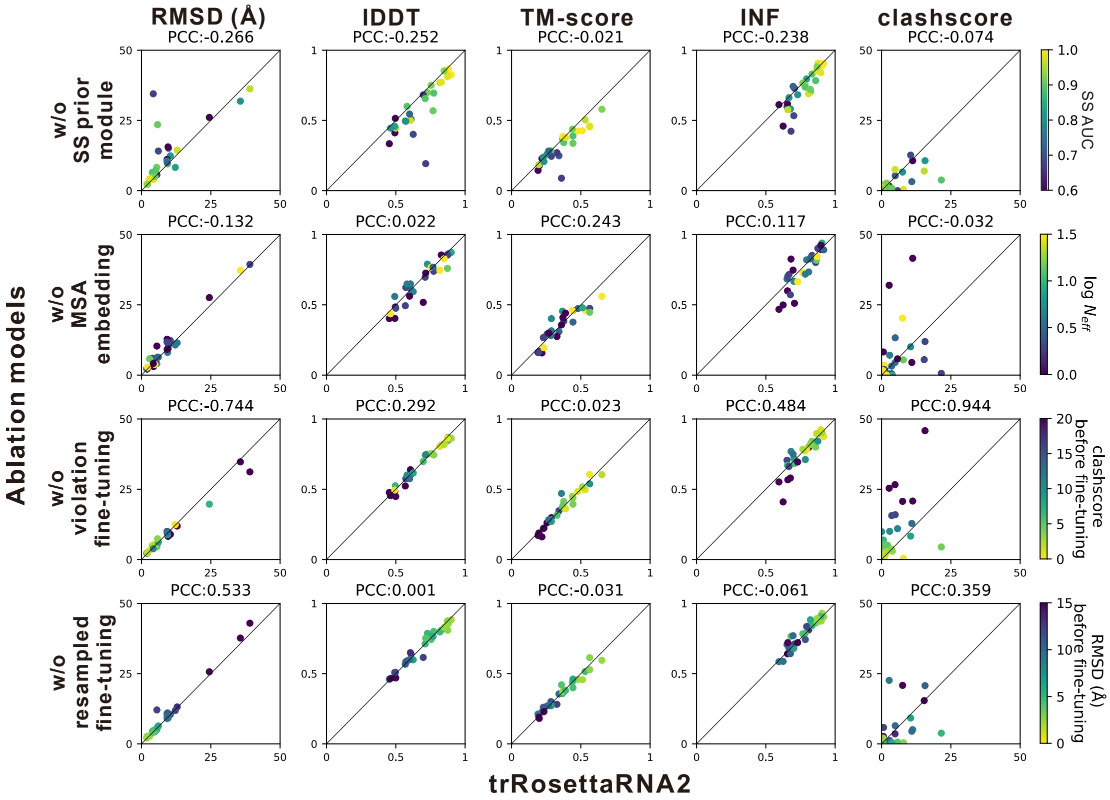

Supplementary Fig. 4. Head-to-head comparisons between trRosettaRNA2 and ablation models. These comparisons are performed on the TS28 dataset (*n*=28). The Pearson correlation coefficient (PCC; calculated using Python) value above each subplot indicates the correlation between the performance degradation and the quality metric associated with the ablated component. E.g., the second subplot in the first row shows a PCC of -0.252, indicating that the lDDT degradation after removing the SS prior module has a slight negative correlation with the quality (AUC) of SS generated by the SS prior module. This suggests that even for targets where the SS prior module generates low-quality prior information, the removal still tends to worsen the performance.

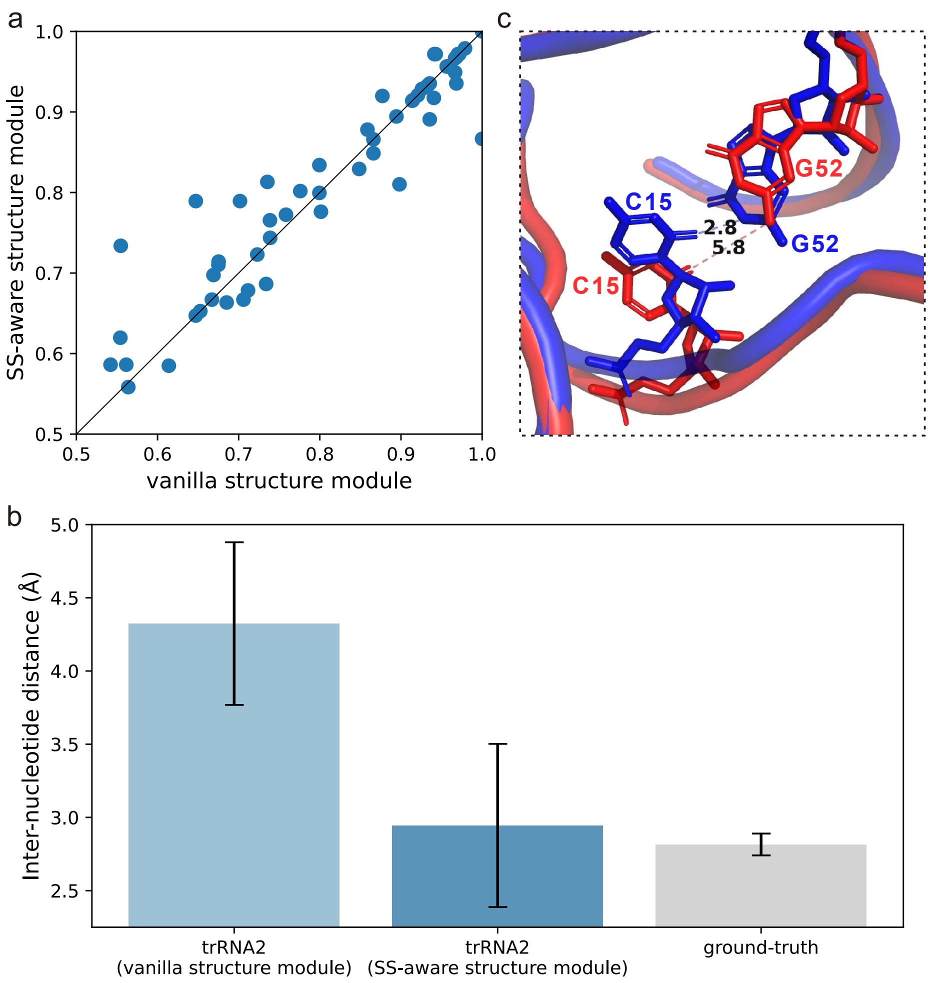

Supplementary Fig. 5. Impact of the SS-aware structure module. (a) Head-to-head comparison of INF for canonical base pairs between the trRNA2 results based on vanilla and SS-aware structure modules. (b) Comparison of inter-nucleotide distances for native base pairs in an example RNA (PDB ID: 8HBA). Bars represent median values, and error bars indicate half of the standard deviation. (c) Visualization of a representative nucleotide pair (C15–G52) that forms a base pair in the experimental structure. The models generated using the vanilla (red) and SS-aware (blue) structure modules are superimposed. Numbers indicate the minimum heavy-atom distances between C15 and G52 in the respective models.

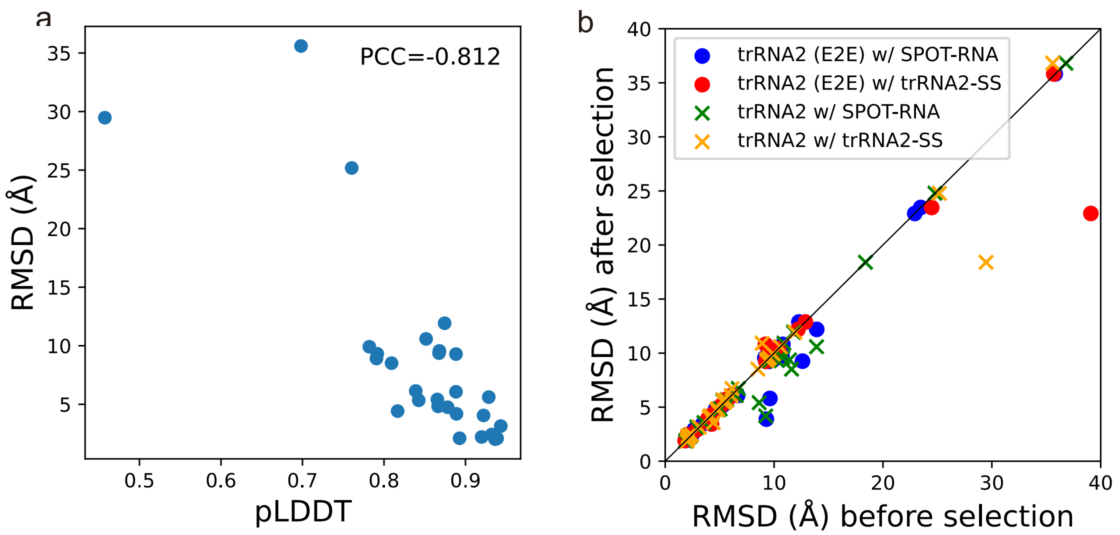

Supplementary Fig. 6. Relationship between pLDDT and RMSD on the TS28 test set (a) and its application to model selection (b).

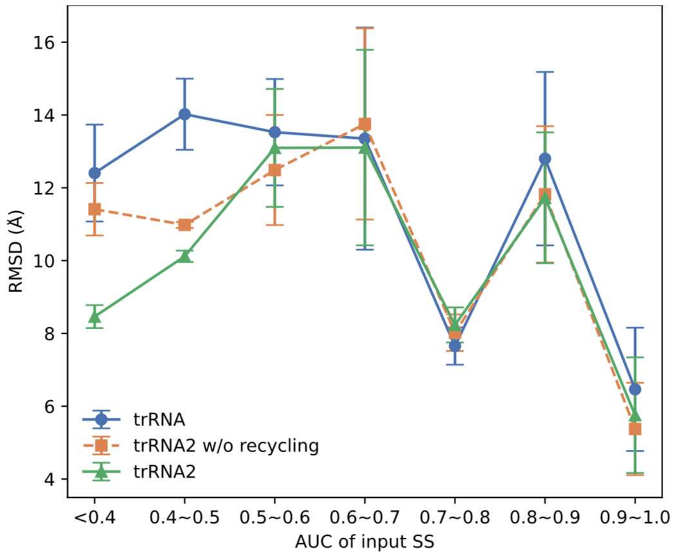

Supplementary Fig. 7. Relationship between RMSD of predicted 3D structure and AUC of input SS. The error bars indicate one-fifth of the standard deviation. "trRNA2" refers to trRosettaRNA2 (E2E) here.

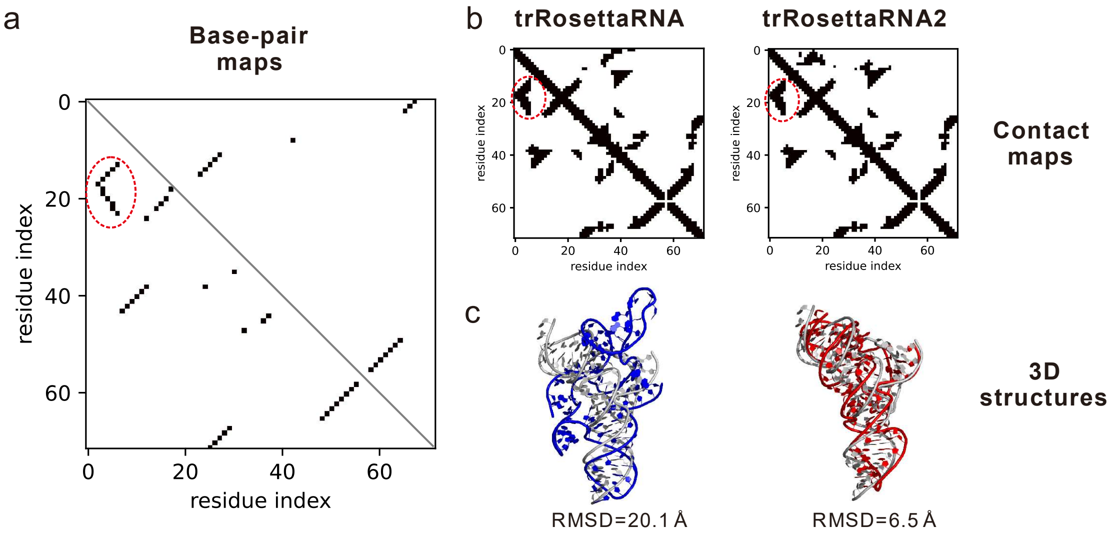

Supplementary Fig. 8. trRosettaRNA vs. trRosettaRNA2 on an example RNA (PDB ID: 8YDC) with SPOT-RNA input. (a) Comparison of the base-pair map predicted by SPOT-RNA (upper right) and that extracted from the experimental 3D structure (lower left). (b) Comparison of the inter-residue contact maps predicted by trRosettaRNA/trRosettaRNA2 (upper right) against the contact map derived from the experimental 3D structure (lower left). (c) Superposition of the 3D structures predicted by trRosettaRNA (blue) and trRosettaRNA2 (red) onto the experimental structure (gray). The red dashed circles in (a) and (b) highlight the interactions corresponding to the triple helix.

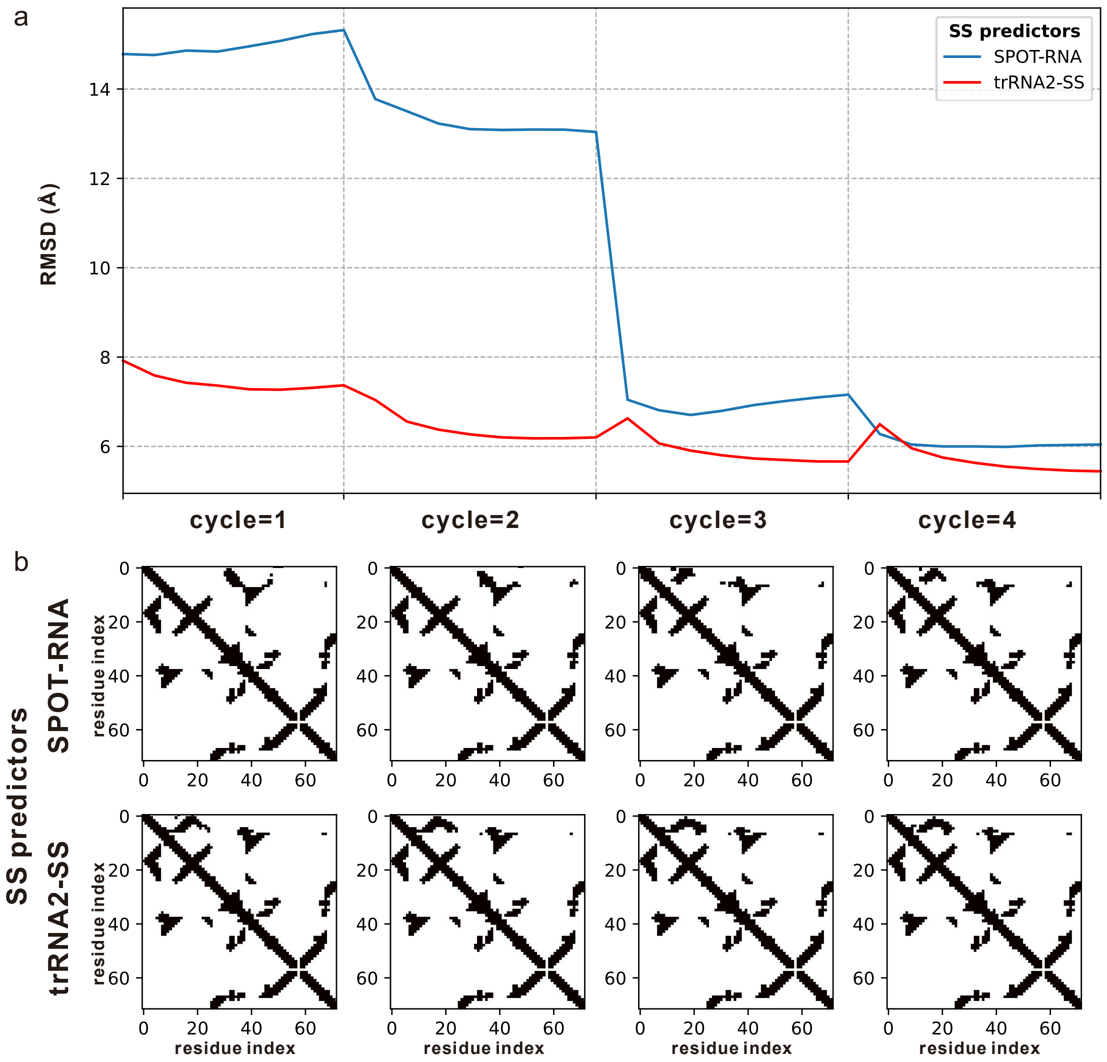

Supplementary Fig. 9. Impact of recycling on the prediction of the example RNA (PDB ID: 8YDC). (a) Relationship between RMSD and number of cycles for trRNA2 using different SS predictors. (b) Comparison of inter-residue contact maps generated at different cycles and using different SS predictors for input. Within each 2D plot in panel (b), the predicted contact map is displayed in the upper right corner, juxtaposed against the experimentally derived contact map shown in the lower left corner.

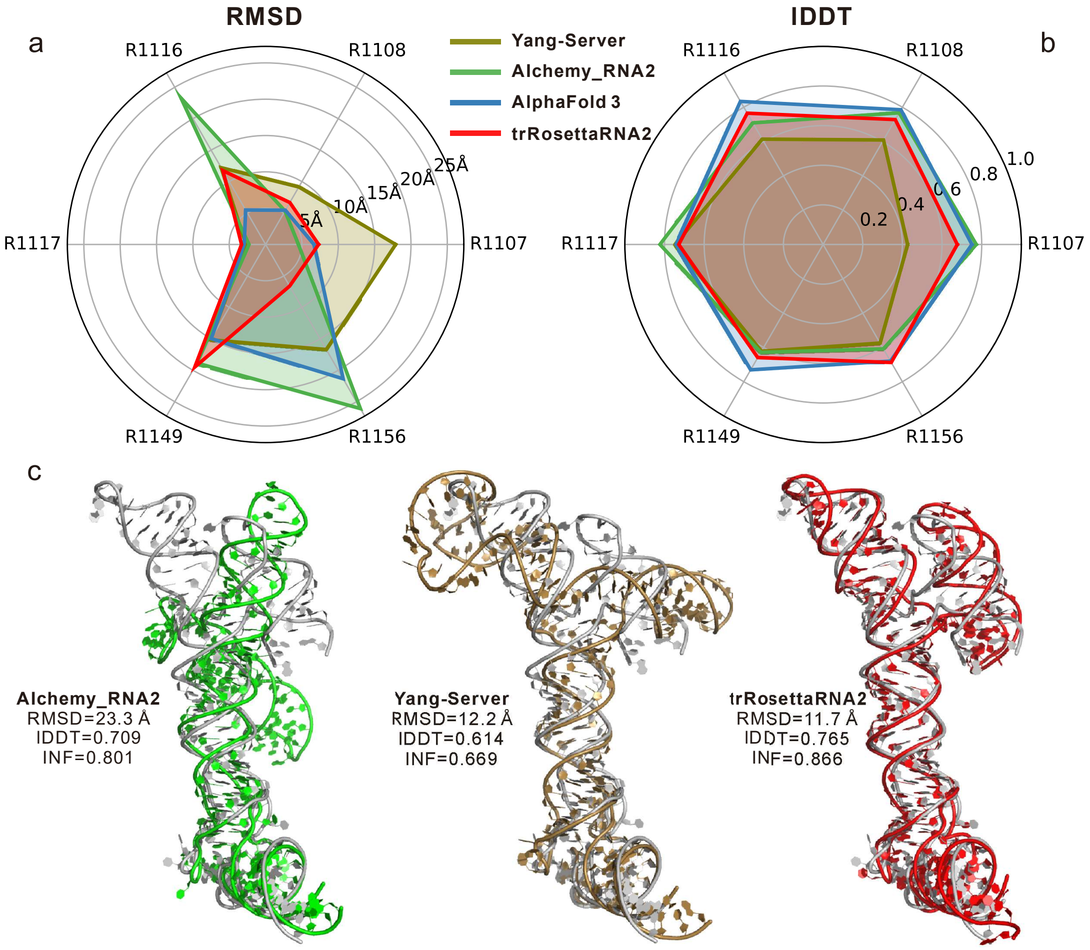

Supplementary Fig. 10. Performance on six natural RNAs from CASP15. (a) Comparison of RMSDs and lDDTs on six natural RNAs (*n*=6). (b) Comparison of AIchemy_RNA2, Yang-Server, and trRosettaRNA on R1116. Note that we evaluate the "first" models for fairness.

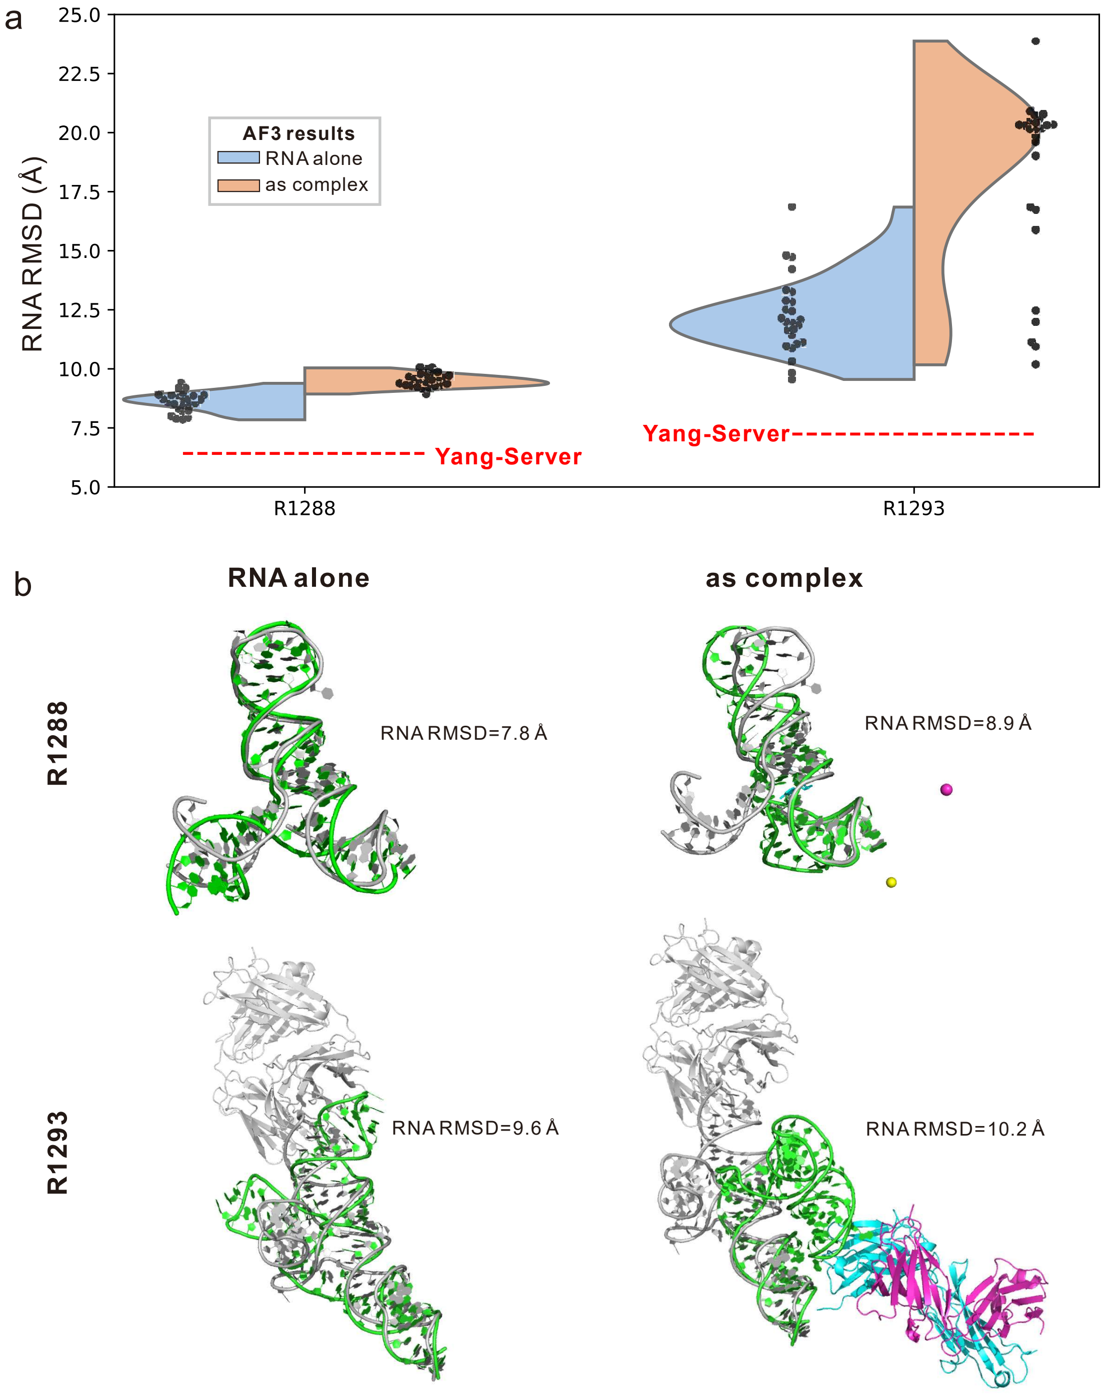

Supplementary Fig. 11. Evaluating AF3 predictions for CASP16 hybrid targets R1288 and R1293. AF3 predictions were generated using the local AlphaFold 3 package with a fixed seed set (0-4), producing 5 models per seed, totaling 25 models for each target. (a) RNA RMSD distribution of the AF3 models when modeling the RNA component alone versus modeling the entire complex (n=75 models). (b) Superimposition of the best AF3 models (RNA in green; proteins or ligands in other colors) onto the experimental structures (in gray).

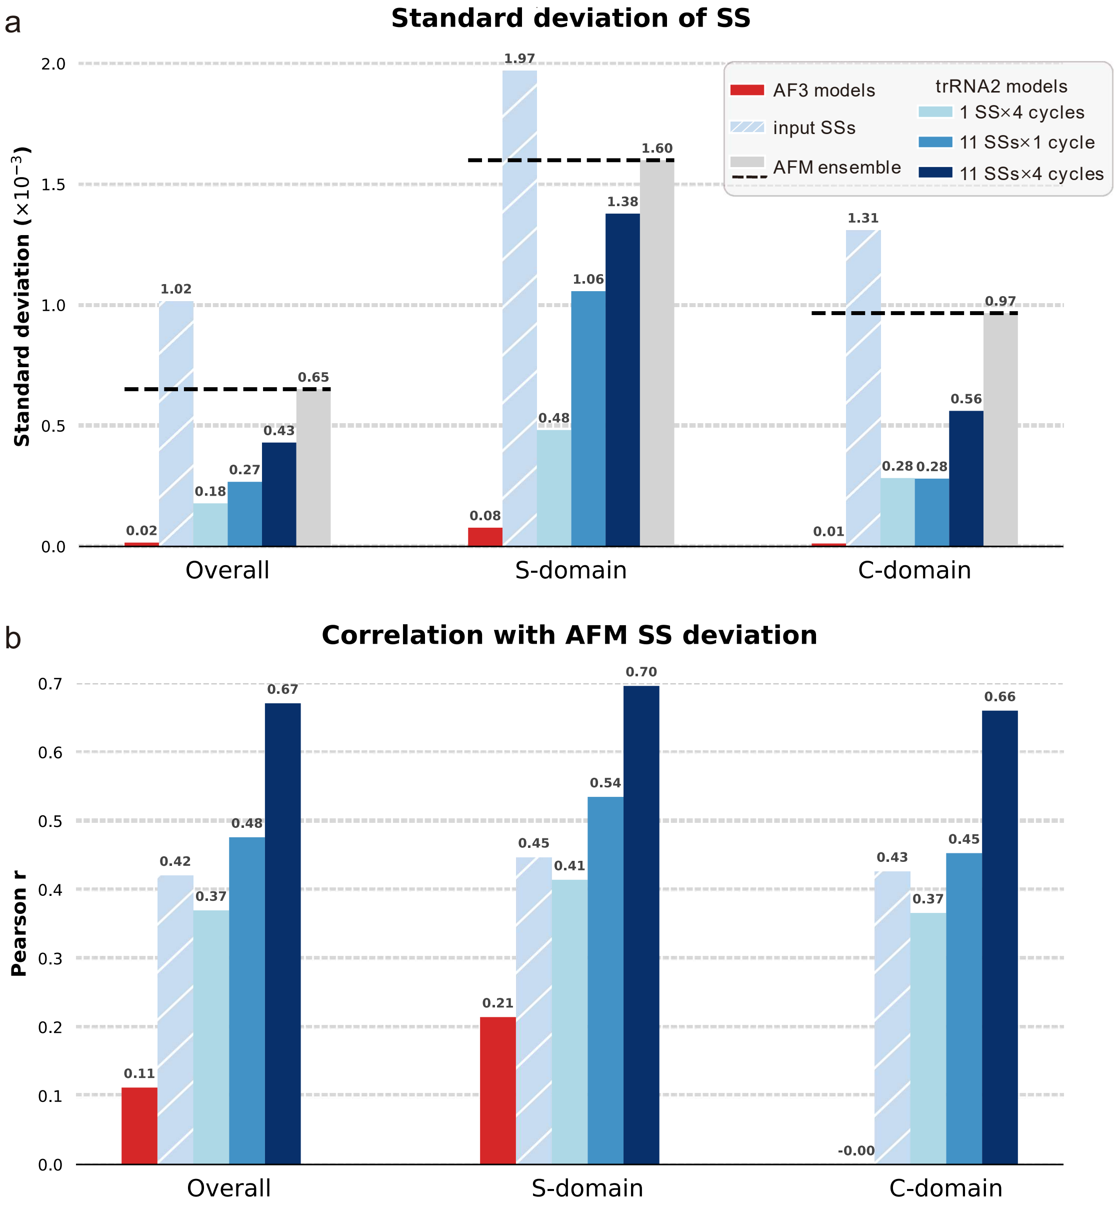

Supplementary Fig. 12. Analysis of SS dynamics for the RNase P RNA. (a) Standard deviations of secondary structures (SD_{SS}) across different ensembles (*n*=100 for AF3; *n*=11 for input SSs; *n*=4/11/44 for trRNA2 outputs; and *n*=158 for AFM), averaged over the whole structure, S-domain, and C-domain. (b) Pearson correlation coefficients (r) quantifying the agreement between the predicted SS standard deviations and the deviations derived from AFM conformations, calculated for the whole structure, S-domain, and C-domain.

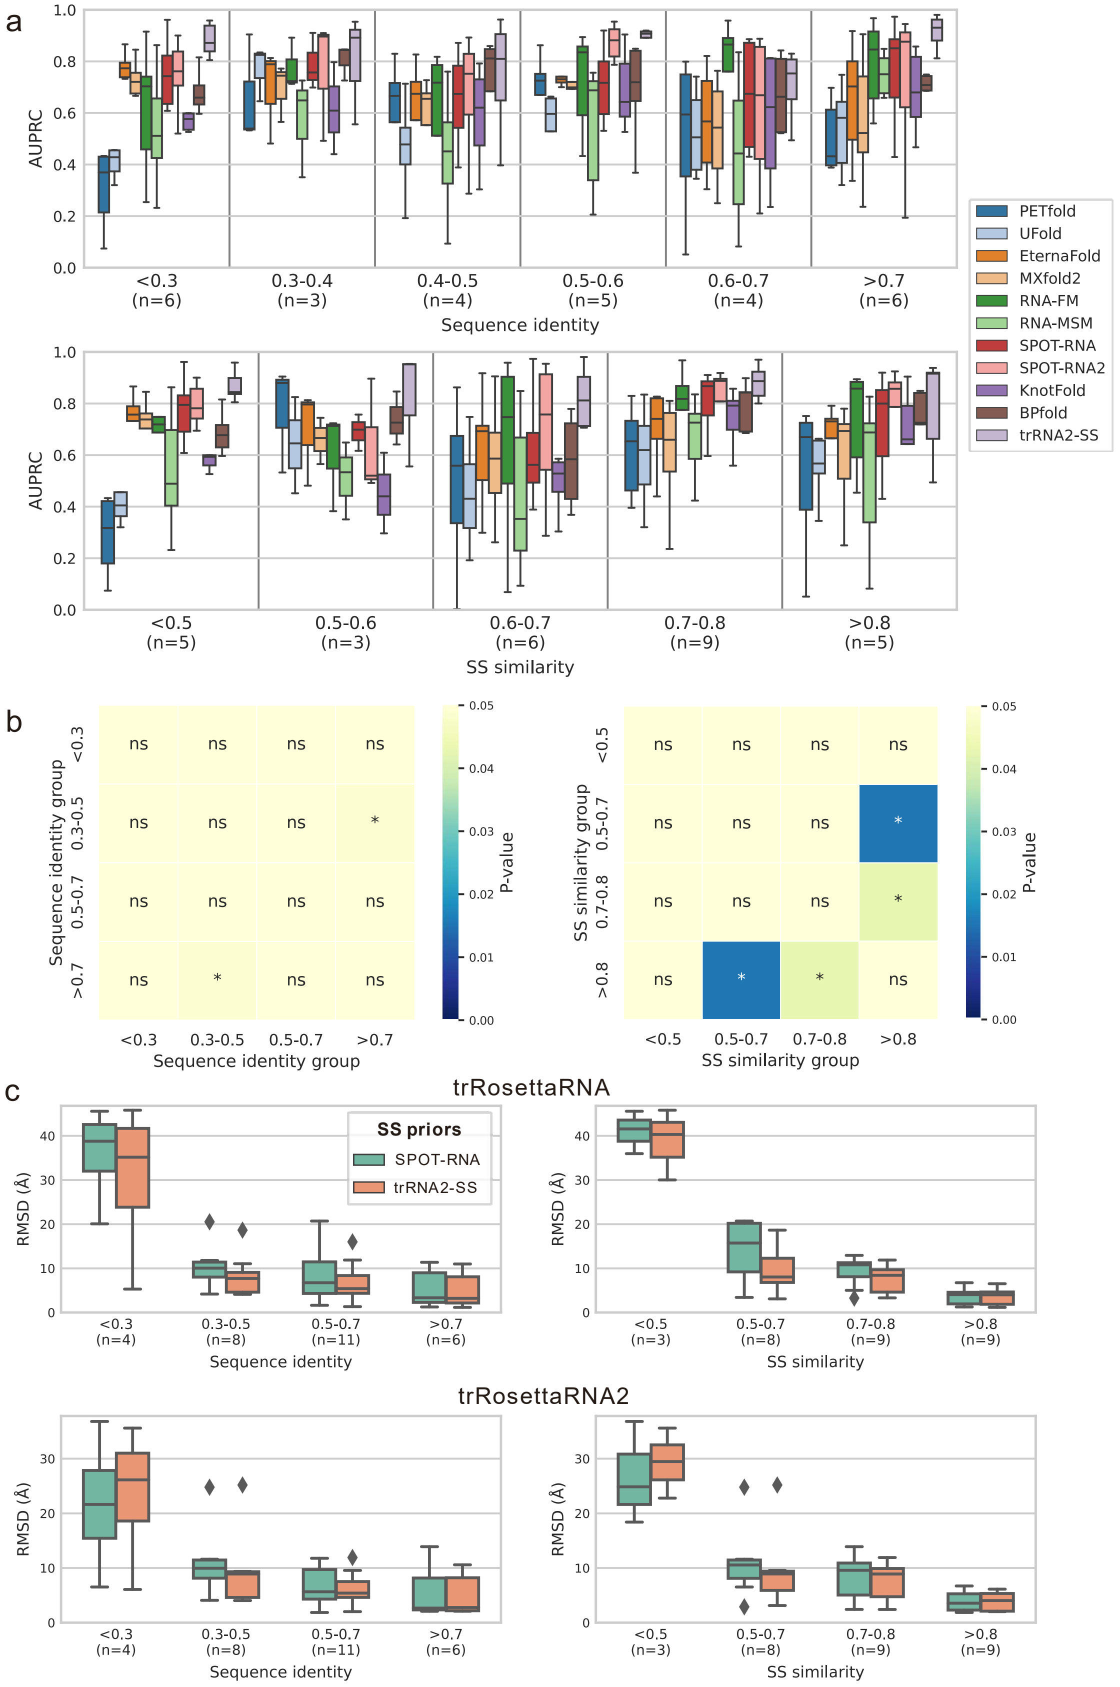

Supplementary Fig. 13. Assessment of trRNA2-SS generalizability and its contribution to 3D modeling accuracy. This analysis is performed on all benchmark RNAs (*n*=38 RNAs; 28 from

TS28 and 10 RNAs from CASP15). For the SS prediction benchmarking (panels a and b), the RNAs containing "N" nucleotide and >600 nt are further removed, resulting in 28 RNAs. (a) Performance comparison (AUPRC) of trRNA2-SS and other representative methods across different similarity levels relative to the training set. Data are grouped by maximum sequence identity (top) and maximum SS similarity (bottom). (b) Pairwise statistical significance of trRNA2-SS performance differences between similarity groups (top: sequence identity; bottom: SS similarity). P-values were calculated using the Mann-Whitney U test. (ns: not significant, $P \geq 0.05$; *: $0.01 \leq P < 0.05$). (c) Impact of SS priors on 3D modeling accuracy (RMSD). The performance of trRosettaRNA2 was evaluated using SS priors predicted by either SPOT-RNA (green) or trRNA2-SS (orange). Results are stratified by sequence identity (left column) and SS similarity (right column). In all boxplots, center lines represent medians, box limits indicate the interquartile range, and whiskers extend to 1.5×IQR.

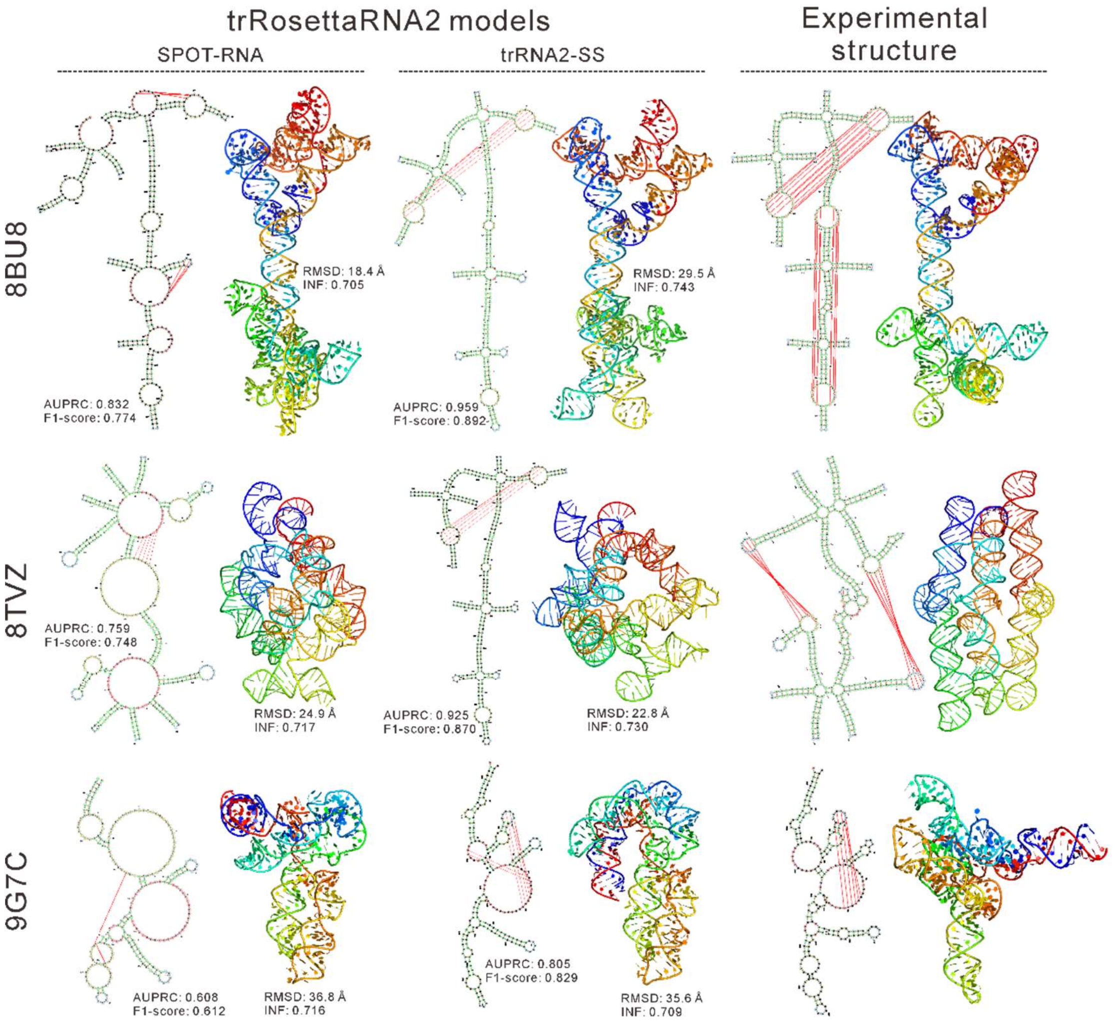

Supplementary Fig. 14. Comparison between trRosettaRNA2 models based on SPOT-RNA and trRNA2-SS priors on three example low-homology targets.

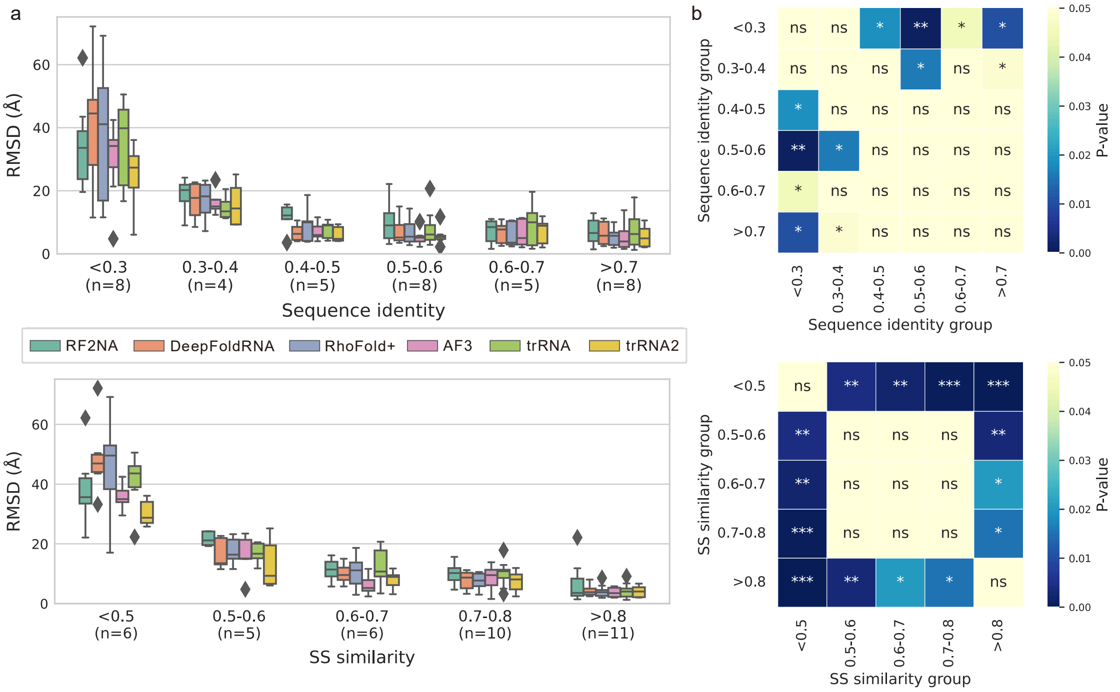

Supplementary Fig. 15. Generalizability analysis of trRosettaRNA2. This analysis is performed on all benchmark RNAs (*n*=38 RNAs; 28 from TS28 and 10 RNAs from CASP15). (a) Performance comparison (RMSD) of trRosettaRNA2 (trRNA2) and other representative methods across different similarity levels relative to the training set. Data are grouped by maximum sequence identity (top) and maximum SS similarity (bottom). The number of targets (*n*) in each group is indicated in parentheses. (b) Pairwise statistical significance of trRosettaRNA2 performance differences between similarity groups (top: sequence identity; bottom: SS similarity). P-values were calculated using the Mann-Whitney U test. (ns: not significant, $P \geq 0.05$; *: $0.01 \leq P < 0.05$; **: $0.001 \leq P < 0.01$; ***: $P < 0.001$).

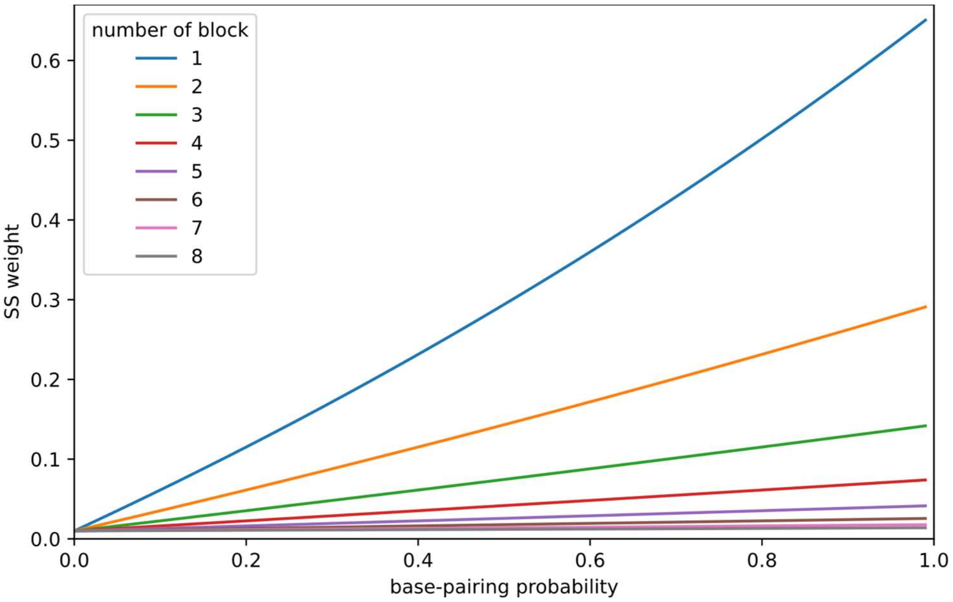

Supplementary Fig. 16. Relationship between the SS weight used in the structure module and the original base-pairing probability. For example, for the first block, the SS weight derived from the highest probability is around 0.6, which is 60 times greater than the weight from the lowest probability. For the last block, the weights are smoothed, with values around 0.01 across the entire range of probabilities.

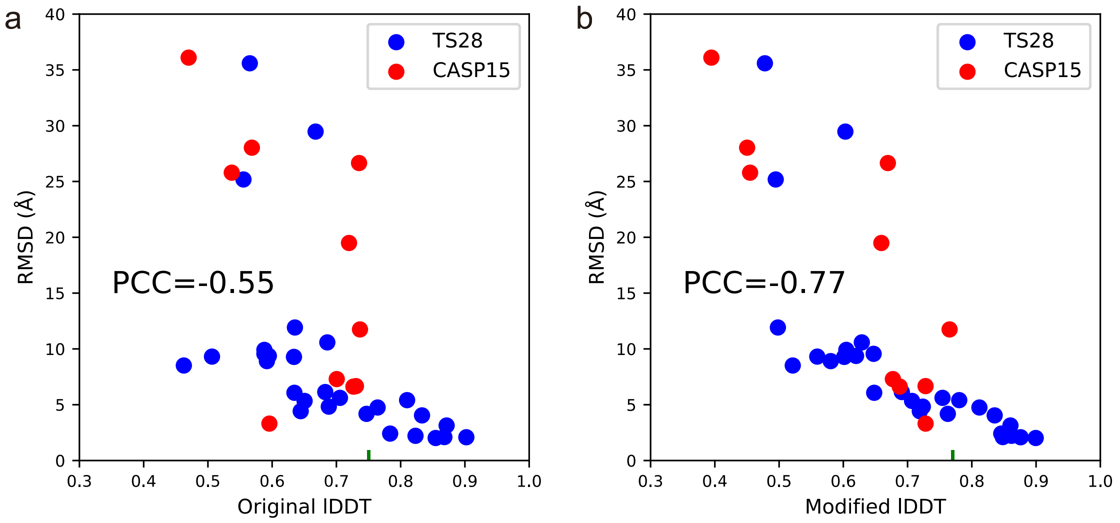

Supplementary Fig. 17. Relationship between RMSD and lDDT before and after modification (*n*=38). PCC values are calculated using Python.

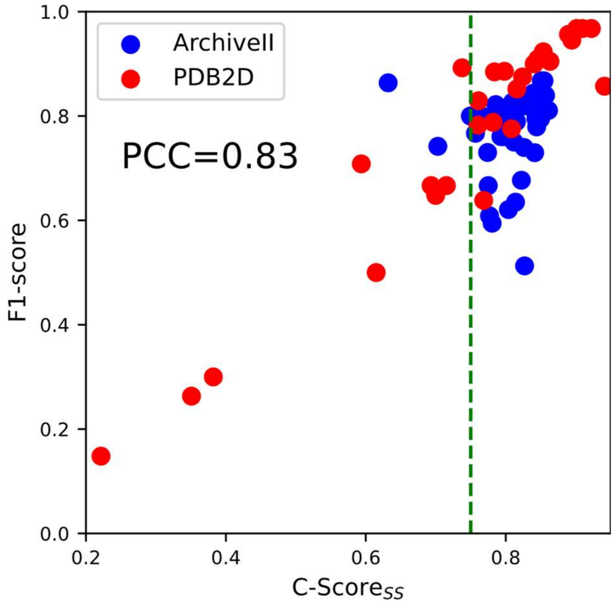

Supplementary Fig. 18. Relationship between F1-Score and C-Score$_{SS}$ on two test sets (*n*=69). For 83.1% of the targets exhibiting a C-Score$_{SS}$ greater than 0.75 (green dashed line), the corresponding F1-scores also exceeded 0.75.

Supplementary References

1. Danaee, P. *et al.* bpRNA: large-scale automated annotation and analysis of RNA secondary structure. *Nucleic Acids Research* **46**, 5381-5394 (2018).
2. Sloma, M.F. & Mathews, D.H. Exact calculation of loop formation probability identifies folding motifs in RNA secondary structures. *RNA* **22**, 1808-1818 (2016).
3. Wayment-Steele, H.K. *et al.* RNA secondary structure packages evaluated and improved by high-throughput experiments. *Nature Methods* **19**, 1234-1242 (2022).
4. Lorenz, R. *et al.* ViennaRNA package 2.0. *Algorithms for Molecular Biology* **6**, 26 (2011).
5. Reuter, J.S. & Mathews, D.H. RNAstructure: software for RNA secondary structure prediction and analysis. *BMC Bioinformatics* **11**, 129 (2010).
6. Seemann, S.E., Menzel, P., Backofen, R. & Gorodkin, J. The PETfold and PETcofold web servers for intra- and intermolecular structures of multiple RNA sequences. *Nucleic Acids Research* **39**, W107-W111 (2011).
7. Chen, J. *et al.* Interpretable RNA Foundation Model from Unannotated Data for Highly Accurate RNA Structure and Function Predictions. *arXiv e-prints*, arXiv:2204.00300 (2022).
8. Zhang, Y. *et al.* Multiple sequence alignment-based RNA language model and its application to structural inference. *Nucleic Acids Research* **52**, e3-e3 (2024).
9. Fu, L. *et al.* UFold: fast and accurate RNA secondary structure prediction with deep learning. *Nucleic Acids Research* **50**, e14-e14 (2022).
10. Sato, K., Akiyama, M. & Sakakibara, Y. RNA secondary structure prediction using deep learning with thermodynamic integration. *Nature Communications* **12**, 941 (2021).
11. Singh, J., Hanson, J., Paliwal, K. & Zhou, Y. RNA secondary structure prediction using an ensemble of two-dimensional deep neural networks and transfer learning. *Nature Communications* **10**, 5407 (2019).
12. Singh, J. *et al.* Improved RNA secondary structure and tertiary base-pairing prediction using evolutionary profile, mutational coupling and two-dimensional transfer learning. *Bioinformatics* **37**, 2589-2600 (2021).
13. Gong, T., Ju, F. & Bu, D. Accurate prediction of RNA secondary structure including pseudoknots through solving minimum-cost flow with learned potentials. *Communications Biology* **7**, 297 (2024).
14. Zhu, H. *et al.* Deep generalizable prediction of RNA secondary structure via base pair motif energy. *Nature Communications* **16**, 5856 (2025).
15. Gao, S.-H. *et al.* Res2Net: A New Multi-Scale Backbone Architecture. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 652–662 (2021).
16. Sun, C., Shrivastava, A., Singh, S. & Gupta, A. Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. *arXiv e-prints*, arXiv:1707.02968 (2017).
17. Statello, L., Guo, C.-J., Chen, L.-L. & Huarte, M. Gene regulation by long non-coding RNAs and its biological functions. *Nature Reviews Molecular Cell Biology* **22**, 96-118 (2021).
18. Peng, X., Liao, W., Lin, X., Lilley, D.M.J. & Huang, L. Crystal structures of the NAD+-II riboswitch reveal two distinct ligand-binding pockets. *Nucleic Acids Research* **51**, 2904-2914 (2023).
19. Chen, K., Zhou, Y., Wang, S. & Xiong, P. RNA tertiary structure modeling with BRiQ potential in CASP15. *Proteins: Structure, Function, and Bioinformatics* **91**, 1771-1778 (2023).
20. Das, R. *et al.* Assessment of three-dimensional RNA structure prediction in CASP15. *Proteins: Structure, Function, and Bioinformatics* **91**, 1747-1770 (2023).
21. Leontis, N.B. & Westhof, E. Geometric nomenclature and classification of RNA

base pairs. *RNA* **7**, 499-512 (2001).
22. Gendron, P., Lemieux, S. & Major, F. Quantitative analysis of nucleic acid three-dimensional structures11Edited by I. Tinoco. *Journal of Molecular Biology* **308**, 919-936 (2001).
23. Lu, X.-J., Bussemaker, H.J. & Olson, W.K. DSSR: an integrated software tool for dissecting the spatial structure of RNA. *Nucleic Acids Research* **43**, e142-e142 (2015).
24. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583-589 (2021).
25. Camacho, C. *et al.* BLAST+: architecture and applications. *BMC Bioinformatics* **10**, 421 (2009).
26. Nawrocki, E.P. & Eddy, S.R. Infernal 1.1: 100-fold faster RNA homology searches. *Bioinformatics* **29**, 2933-2935 (2013).
27. Morcos, F. *et al.* Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences* **108**, E1293-E1301 (2011).
28. Luwanski, K. *et al.* RNAspider: a webserver to analyze entanglements in RNA 3D structures. *Nucleic Acids Research* **50**, W663-W669 (2022).