

Article

IociPARSE: A Locality-aware Invariant Point Attention Model for Scoring RNA 3D Structures

Sumit Tarafder and Debswapna Bhattacharya*

Cite This: J. Chem. Inf. Model. 2024, 64, 8655–8664

ACCESS	III Metrics & More	E Article Recommendations	s Supporting Information
ABSTRACT: A	scoring function that can	reliably assess the	High Confidence

ABSTRACT: A scoring function that can reliably assess the accuracy of a 3D RNA structural model in the absence of experimental structure is not only important for model evaluation and selection but also useful for scoring-guided conformational sampling. However, high-fidelity RNA scoring has proven to be difficult using conventional knowledge-based statistical potentials and currently available machine learning-based approaches. Here, we present lociPARSE, a locality-aware invariant point attention architecture for scoring RNA 3D structures. Unlike existing machine learning methods that estimate superposition-based root-mean-square deviation (RMSD), lociPARSE estimates Local Distance Difference Test (IDDT) scores capturing the accuracy of each nucleotide and its surrounding local atomic environment in a



superposition-free manner, before aggregating information to predict global structural accuracy. Tested on multiple datasets including CASP15, lociPARSE significantly outperforms existing statistical potentials (rsRNASP, cgRNASP, DFIRE-RNA, and RASP) and machine learning methods (ARES and RNA3DCNN) across complementary assessment metrics. lociPARSE is freely available at https://github.com/Bhattacharya-Lab/lociPARSE.

1. INTRODUCTION

Computational prediction of RNA 3-dimensional structures from nucleotide sequence has garnered considerable research effort over the past decade,^{1–10} and deep learning-enabled RNA 3D modeling has gained significant attention in the recent past.^{11–15} To facilitate the practical applicability of predicted 3D models, it is critical to have a scoring function that can reliably assess their global topology and local quality in the absence of experimental structures.^{10–18} Moreover, the ability of a scoring function to distinguish accurate 3D models of previously unseen RNAs from misfolded alternatives plays an important role in guiding conformation sampling toward the native state.¹⁹

Existing methods for scoring RNA structures roughly belong to two categories: knowledge-based statistical potentials and supervised machine learning. Various knowledge-based statistical potentials have been developed, both at all-atom and coarse-grained levels,^{20–24} using different simulated reference states.^{25–30} However, reliably distinguishing accurate structural models of RNA from less accurate ones has proven to be difficult because the characteristics of energetically favorable RNA structures are not sufficiently well understood and thus the reference states may deviate largely from the ideal one. Machine learning-based methods^{17,31} aim to overcome such limitations by learning to predict the accuracy of an RNA structural model through supervised learning. Indeed, machine learning-based RNA scoring functions, trained to estimate the unfitness score either at the nucleotide level or at the structural level by learning to predict the root-mean-square deviation (RMSD) from the unknown true structure, have been shown to be effective in RNA-Puzzles blind structure prediction challenges.⁵

Despite the effectiveness, the existing machine learning methods do not consider some key factors that can significantly improve the sensitivity of RNA scoring functions. First, the global superposition-dependent RMSD metric is not length normalized, is affected by superposition, is dominated by outliers in poorly modeled structural regions, and does not take into account the accuracy of the local atomic environment. RNA is a flexible molecule in which irregular loops may affect RMSD measures and global superposition may not be optimal, leading to scoring anomalies. Yet, virtually all existing machine learning-based RNA scoring functions use RMSD as the ground truth during supervised training. Second, similar to other macromolecules, RNA structures have no natural canonical orientation i.e., the same RNA structure can be

Received:September 6, 2024Revised:October 17, 2024Accepted:October 29, 2024Published:November 11, 2024







Figure 1. Overview of lociPARSE. Given a 3D RNA structure, we estimate the local nucleotide-wise IDDT scores that are then aggregated over all nucleotides to predict global structural accuracy. (a) Input RNA 3D structure. (b) Architecture of our locality-aware invariant point attention (IPA) module to capture the accuracy of each nucleotide and the effect of its local atomic environment. (c) By aggregating information at the level of nucleotide-wise IDDT (pNuL) scores before making a prediction at the level of the entire RNA structure to output predicted molecular-level IDDT (pMoL).

rotated in space without affecting its biological function, thus allowing a structure to be represented in any orientation. As such, machine learning methods that are not invariant to global Euclidean transformations such as rotation must account for this aspect of variation by tweaking model architecture and/or parameters, which may affect their expressiveness and generalizability. Third, in consideration of RNA as a flexible molecule in which interplay between various local structural motifs defines the global topology, an effective scoring function should not be strongly influenced by the relative motions between the tertiary motifs. That is, the effects of relative movement between the motifs should not lead to artificially unfavorable scores.

Using the Local Distance Difference Test (lDDT)³² as the ground truth during supervised training is an attractive alternative to the popular RMSD metric. IDDT compares distances between atoms that are nearby (within 15 Å) in the experimental structure to the distances between those atoms in the predicted structure and offers several advantages over RMSD. First, being superposition-free and based on rotationinvariant properties of a structure, IDDT naturally preserves invariance with respect to the global Euclidean transformations of the input RNA structure such as global rotations and translations. Second, IDDT measures the accuracy of the local environment of the model in atomic detail, without being affected by superposition or dominated by outliers in poorly modeled structural regions. Third, IDDT exhibits robustness to movements between tertiary structural units such as domains in proteins that can generalize to RNA tertiary motifs, provided a way can be found that ensures rigid motion between a set of local structural units is invariant under global Euclidean transformations on the said units. A solution to this problem comes from Invariant Point Attention (IPA) proposed in AlphaFold2 as part of the structural module.³³ IPA is a form of attention that acts on a set of 3D point clouds and is invariant under global Euclidean transformations on said points, where 3D point clouds are represented using local frames.

How can we capture the aforementioned benefits of IDDT in a neural network architecture for RNA scoring, while maintaining invariance under global Euclidean transformations? Here, we provide such a solution by developing a new attention-based architecture, called lociPARSE (locality-aware

invariant Point Attention-based RNA ScorEr), for scoring RNA 3D structures. Different from previous supervised learning approaches that estimate the RMSD metric, our method estimates local nucleotide-wise IDDT scores that are then aggregated over all nucleotides to predict global structural accuracy. Inspired by AlphaFold2, we define nucleotide-wise frames parameterized by rotation matrices and translation vectors operating on predefined RNA conformation at the local level. To model the local atomic environment of each nucleotide as captured by IDDT, the IPA implementation used in the original AlphaFold2 has been modified to incorporate locality information derived from the RNA atomic coordinates. By so doing, we are able to effectively capture the accuracy of each nucleotide while considering the effect of its local atomic environment. It is worth mentioning here that, although an RNA 3D structure prediction method that uses AlphaFold2inspired IPA architecture can self-estimate the quality of its own predicted structure given an input RNA sequence, our method is the first general-purpose RNA scoring method that is capable of estimating the quality of any input RNA 3D structure using our modified implementation of AlphaFold2's IPA architecture.

Our method significantly outperforms traditional knowledge-based statistical potentials as well as state-of-the-art machine learning-based RNA scoring functions such as ARES³¹ on multiple independent test datasets including CASP15 blind test targets across a wide range of performance measures. In particular, lociPARSE exhibits superior ability to reproduce the ground truth IDDT scores both at the global and local levels, rank predictions for a given target with high fidelity, recognize the best predictions consistently, and better discriminate between "good" and "bad" predictions. An opensource software implementation of lociPARSE, licensed under the GNU General Public License v3, is freely available at https://github.com/Bhattacharya-Lab/lociPARSE.

2. MATERIALS AND METHODS

2.1. lociPARSE: Locality-aware Invariant Point Attention for RNA Scoring. An overview of our method, lociPARSE, is illustrated in Figure 1. The core component of our architecture, outlined in Figure 1b, is an invariant point attention (IPA) module that utilizes the geometry of the input RNA 3D structure to revise the nucleotide and pair features. This component is similar to the AlphaFold2's IPA formulation used in the structure module, but modified herein to incorporate locality information derived from the RNA atomic coordinates. To do this, we introduce locality-aware geometry and edge-biased attention (Section 2.3.2) to convert nucleotide pair features to edge adjacencies by considering a set of k-nearest neighbor nucleotides to capture the local atomic environment of each nucleotide based on the Euclidean distances of the C4'-C4' atoms between nucleotide pairs. In our setting, we define local nucleotide frames (Section 2.3.1) from the Cartesian coordinates of C4', P, and glycosidic N atoms. The IPA partitions the nucleotide query and value features into 3D vectors and transforms them from the target nucleotide's local frame into a global reference frame before computing both attention weights and the output of the attention mechanism. Further, we augment nucleotidenucleotide atomic distances between all pairs of 3 atoms P, C4' and N, encoded with Gaussian radial basis functions as pair features (Section 2.2), and make further use of the pair features to bias attention weights and update scalar features. Our network architecture consists of 4 IPA layers, with the IPA hyperparameters (N_{heads} , c, $N_{\text{query points}}$, $N_{\text{value points}}$) set to (4, 128, 8, 4) and we use 20 nearest neighbors for the locality computation, determined through ablation experiments using an independent validation set (Section 3.4). The output of the attention layer is invariant to the global Euclidean transformations such as global rotations and translations of the input RNA. Finally, a linear layer followed by a two-layer fully connected network is used to estimate the predicted nucleotide-wise lDDT (pNuL) scores. These nucleotide-level pNuL predictions are then aggregated over all nucleotides by taking an average to estimate the predicted molecular-level IDDT (pMoL), enabling our method to estimate both local and global qualities of the input RNA 3D structure.

2.2. Model Input. Our model uses only input features derived directly from a nucleotide sequence and RNA 3D structural coordinates. We use just the basic nucleotide-level encodings for our input. These include one-hot encoding of the nucleotide (i.e., a binary vector of 5 entries indicating each of the 4 nucleotide types and one for nonstandard nucleotides such as 'T' or modified nucleotides) and the relative position of the nucleotide in its sequence calculated as i/L (where *i* is the nucleotide index and L is the sequence length). For our pair features, we use the index i of a nucleotide's partner in sequence and the corresponding 3D coordinates, quantified by sequential separation and spatial proximity information. We do not consider the nucleotide type of the partner as a pair feature. The sequence separation i.e., the absolute difference between the two nucleotide indices is discretized into 5 bins and represented by one-hot encoding where the first two bins correspond to self-loops and adjacent bonds, respectively. The rest of the three bins are defined based on three types of interactions depending on the sequence separation: shortrange (2-5), medium-range (6-24) and long-range (>24), similar to.^{34,35} The other component of our pair features includes nucleotide-nucleotide atomic distances between all pairs of P, C4' and glycosidic N atoms, encoded with radial basis functions. The radial basis functions are defined based on the distance from a reference point, making them suitable for capturing distance-based similarities. We used Gaussian radial basis functions to encode the interatomic distances in the following manner. For a nucleotide pair (i, j), at first, k = 9

distances d_{ij}^k are calculated because of all possible combinations among the set of 3 atoms P, C4' and N. Then, the set of all d_{ij}^k values, D is encoded using Gaussian radial basis function as follows:

$$D_{\mu} = \text{linspace}(D_{\min}, D_{\max}, \text{RBF})$$

$$D_{\sigma} = \frac{D_{\max} - D_{\min}}{\text{RBF}}$$

$$D_{\text{RBF}} = \exp\left(-\left(\frac{D - D_{\mu}}{D_{\sigma}}\right)^{2}\right)$$
(1)

where, RBF = 16 is the number of radial basis functions, a value chosen based on the ablation experiment presented in Table 6 on an independent validation set. As such, D_{μ} is a set of 16 linearly spaced values between $D_{\min} = 0$ and $D_{\max} = 100$, which indicates the minimum and maximum interatomic distances, respectively. The output D_{RBF} contains 16 encoded distance pair features for each distance d_{ij}^k . It is important to note that all our nucleotide and pair features are invariant, consistent with the invariant layers of the IPA module.

2.3. Network Architecture. 2.3.1. Construction of Local Nucleotide Frames. To perform invariant point attention on a set of 3D points, we represent each nucleotide in a geometric abstraction using the concept of frames. Each nucleotide frame in the form of a tuple is defined as a Euclidean transform $T = (R, \vec{t})$, where $R \in \mathbb{R}^{3\times3}$ is a rotation matrix and $\vec{t} \in \mathbb{R}^3$ is the translation vector that can be applied to transform a position in local coordinates $(\vec{x}_{local} \in \mathbb{R}^3)$ to a position in global coordinates $(\vec{x}_{global} \in \mathbb{R}^3)$ as

$$\vec{\mathbf{x}}_{global} = T \circ \vec{\mathbf{x}}_{local}$$

$$= (R, \vec{\mathbf{t}}) \circ \vec{\mathbf{x}}_{local}$$

$$= R \vec{\mathbf{x}}_{local} + \vec{\mathbf{t}}$$
(2)

In our setting, we define local nucleotide frames from the Cartesian coordinates of P, C4', and glycosidic N atoms of the input RNA 3D structure and construct 3-bead coordinate frame using a Gram–Schmidt process specified in Alphafold2 (Algorithm 21) that takes the input coordinates (scaled by 0.1) of P as \vec{x}_1 , C4' as \vec{x}_2 , and N as \vec{x}_3 . Note that the translation vector \vec{t} is assigned to the center atom \vec{x}_2 .

2.3.2. Locality-aware Invariant Point Attention. The formulation of locality-aware IPA used in our work combines sequence representation, \mathbf{s}_{ij} from each nucleotide *i* of the input RNA, pair representation \mathbf{e}_{ij} of nucleotide *i* with other nucleotides *j* based on nucleotide pair adjacencies capturing the local atomic environment N of nucleotide *i*, where $j \in N_i$ is the locality information derived from the RNA atomic coordinates. Consequently, the update function of the IPA layer is as follows:

$$\mathbf{s}'_i = \mathbf{s}_i + \mathrm{IPA}(\mathbf{s}_i, \{\mathbf{s}_j\}_{j \in \mathcal{N}_i}, \{\mathbf{e}_{ij}\}_{j \in \mathcal{N}_i})$$
(3)

To perform attention on 3D point clouds, IPA derives query (\mathbf{q}_i^h) , key (\mathbf{k}_i^h) and value (\mathbf{v}_i^h) embeddings from a linear projection of \mathbf{s}_i to a latent representation of dimension *c* for each nucleotide *i*, where $\mathbf{q}_i^h, \mathbf{k}_i^h, \mathbf{v}_i^h \in \mathbb{R}^c$ and $h \in \{1, \dots, N_{\text{head}}\}$ which represents the number of attention heads in the IPA module. 3D query, key and value points are also generated considering the local frame T_i of each nucleotide *i*, where \mathbf{q}_i^{hp} ,

 $\vec{k}_i^{hp} \in \mathbb{R}^3$, $p \in \{1, \dots, N_{\text{query points}}\}$ and $\vec{v}_i^{hp} \in \mathbb{R}^3$, $p \in \{1, \dots, N_{\text{value points}}\}$. The IPA module acts on a set of frames (parameterized as Euclidean transforms of the local frame T_i) and is invariant under global Euclidean transformations T_{global} on said frames. By performing locality-aware geometry and edge-biased attention, the IPA module transforms the 3D points from the target nucleotide's local frame into a global reference frame for computing the attention weights as follows:

$$a_{ij}^{h} = \operatorname{softmax}_{j}(w_{\mathrm{L}}(\frac{1}{\sqrt{c}}\mathbf{q}_{i}^{h^{T}}\mathbf{k}_{j}^{h} + b_{ij}^{h}) - \frac{\gamma^{h}w_{\mathrm{C}}}{2}\sum_{p} \left\| T_{i}\circ\overline{\mathbf{q}}_{i}^{hp} - T_{j}\circ\overline{\mathbf{k}}_{j}^{hp} \right\|^{2})$$

$$(4)$$

where b_{ij}^h is the attention bias derived from the linear projection of \mathbf{e}_{ij} to hidden dimension c, weighting factors $w_{\rm L}$ and $w_{\rm C}$ are taken from the IPA formulation specified in AlphaFold2 and $\gamma^h \in \mathbb{R}$ is a learned scalar value. The attention mechanism acting on a set of local frames ensures invariance under global Euclidean transformations such as global rotations and translations of the input RNA due to the invariant nature of l_2 -norm of a vector under such rigid transformations.

The attention weights are used to compute the outputs of the attention mechanism while mapping them back to the local frame and preserving invariance, as follows:

$$\widetilde{\mathbf{o}}_{i}^{h} = \sum_{j} a_{ij}^{h} \mathbf{e}_{ij}$$
⁽⁵⁾

$$\mathbf{o}_{i}^{h} = \sum_{j} a_{ij}^{h} \mathbf{v}_{j}^{h}$$

$$\tag{6}$$

$$\overrightarrow{\mathbf{o}}_{i}^{hp} = T_{i}^{-1} \circ \sum_{j} a_{ij}^{h} (T_{j} \circ \overrightarrow{\mathbf{v}}_{j}^{hp})$$

$$\tag{7}$$

The outputs of the attention mechanism are then concatenated and passed through a linear layer to compute the updated sequence representations \mathbf{s}_i' of each nucleotide as follows:

$$\mathbf{s}_{i}' = \mathbf{s}_{i} + \operatorname{Linear}(\operatorname{concat}_{h,p}(\widetilde{\mathbf{o}}_{i}^{h}, \mathbf{o}_{i}^{h}, \overrightarrow{\mathbf{o}}_{i}^{hp}, \left\| \overrightarrow{\mathbf{o}}_{i}^{hp} \right\|))$$
(8)

Here, $\|\vec{o}_i^{hp}\|$ indicates the Euclidean norm of \vec{o}_i^{hp} . The updated sequence embeddings \mathbf{s}_i' for each nucleotide *i* are subsequently stacked together to obtain the embedding \mathbf{s}' for all nucleotides in the RNA. Finally, a linear layer followed by a 2 layer fully connected network implemented as a multilayer perceptron (MLP) is used to obtain the final representation \mathbf{s}^f before estimating nucleotide-wise lDDT scores as follows:

$$\mathbf{s}^{f} = \mathbf{s}' + \mathrm{MLP}(\mathbf{s}') \tag{9}$$

2.4. Training, Validation, and Test Datasets. To curate our training dataset, we first obtained the training dataset used in trRosettaRNA¹² containing 3, 632 RNA targets. We then filtered this set by removing duplicate chains and discontinuous structures, separating monomers from complexes, splitting multiple chains into single chains, and correcting formatting issues in the coordinates files. We removed sequences with length >200 nucleotides and ensured that our training and test sets were nonredundant by running CD-HIT-est³⁶ with default parameter settings, which reduced the training set to 1, 399 RNA targets. We generated 37 structural models per target with a total of 51, 763 structural models for

the 1, 399 targets using a combination of different RNA 3D structure prediction tools including recent deep learningenabled RNA structure prediction methods,11-15 physicsbased RNA folding,³⁷ and structure perturbation using PyRosetta.³⁸ For the deep learning-enabled RNA structure prediction methods, we used default parameter settings to generate the structural models. Additionally, we generated one model using SimRNA, selecting the first frame of the cluster with a 3.5 Å RMSD threshold. Finally, each of the 6 structures generated by DeepFoldRNA¹¹ was relaxed by FastRelax protocol³⁹ within the PyRosetta framework³⁸ with 5,000 and 10,000 steps aiming to introduce structural diversity in our training set. The number of structural models generated by each method is listed in Supplementary Table S1. Our test data includes 30 independent RNAs, also collected from trRosettaRNA following the train and test splits of the original work. We generated 3D structural models for each of these 30 RNAs using the deep learning-enabled RNA structure prediction methods.¹¹⁻¹⁵ We also use 12 RNA targets from CASP15 as an additional independent test set containing targets cleared for public access as of December 20, 2022, where the corresponding 3D structural models are collected directly from the CASP15 Web site https://predictioncenter.org/ casp15/ based on the blind predictions submitted by various participating groups in CASP15 RNA 3D structure prediction challenge. In addition, we separately curated a validation set of 60 RNA targets for the ablation study and hyperparameter selection from the Protein Data Bank (PDB)⁴⁰ with experimental structures released between January 1, 2022 and July 6, 2023. Such a date range was chosen to avoid any overlap with our training dataset collected from trRosettaRNA which used structures released before January 1, 2022. We generated 25 structural models for each of these 60 RNAs using the recent deep learning-based RNA structure prediction methods.¹¹⁻¹⁵ Supplementary Tables S2 and S3 list the number of structural models per target used for our two test sets Test30 and CASP15, respectively. The nucleotide-wise ground truth lDDT distributions of our training and test sets are shown in Supplementary Figures S3 and S4. We created a reduced training subset consisting of 6, 872 structural models for 1, 399 RNA targets through clustering⁴¹ for ablation study and hyperparameter selection.

2.5. Training Details. To train our model, lociPARSE, we obtained nucleotide-wise ground truth IDDT scores by comparing the predicted structural models in our training dataset against the corresponding experimental structures using the docker version of OpenStructure.⁴² During the ground truth lDDT computation, we enabled the option '-lddt-nostereochecks' to skip stereochemical quality checks in its calculation following the recent CASP assessment.⁴ lociPARSE was implemented in PyTorch⁴³ with the $\mathcal{L}1$ loss function to learn the mean absolute error between ground truth IDDT and predictions on nucleotide level, thereby formulating the local nucleotide-wise quality estimation as a regression task. We trained our model using the Adam optimizer⁴⁴ having parameters $\beta 1 = 0.9$ and $\beta 2 = 0.999$ with a learning rate of 0.0001 and dropout rate of 0.1. The training process consists of 50 epochs on an 80-GB NVIDIA A100 GPU.

2.6. Competing Methods. lociPARSE is compared against both traditional knowledge-based statistical potentials (rsRNASP,²⁰ RASP,²² DFIRE-RNA,²³ and cgRNASP²¹) and recent machine learning-based RNA scoring functions (RNA3DCNN¹⁷ and ARES³¹). rsRNASP is an all-atom

distance-dependent potential considering short and longranged interactions present in RNA based on sequence separation aiming to capture the hierarchical nature of RNA folding.⁴⁵ Ribonucleic Acids Statistical Potential (RASP) is another all-atom statistical potential based on the averaging reference state. Similar to rsRNASP, RASP also separates interaction pairs into local and nonlocal categories and takes into account the base stacking and base pairing interactions present in RNA. DFIRE-RNA is yet another distance-scaled statistical potential designed using a finite-ideal-gas reference state. Finally, cgRNASP, a coarse-grained counterpart of rsRNASP potential introduces three different variants of coarse-grained potentials for RNA scoring. We have used the 3-bead representation of cgRNASP in this work which takes into account P, C4' and N atoms. The Atomic Rotationally Equivariant Scorer (ARES) is an equivariant graph neural network that scores RNA structures by identifying complex structural motifs through equivariant convolutions. ARES employs E3NN⁴⁶ to predict the global RMSD of the structure. Finally, RNA3DCNN uses 3D convolutional neural network to predict the RMSD-like unfitness score of a nucleotide to its surroundings by considering RNA 3D structure as a 3D image and representing each nucleotide as an array of voxels. For prediction, we used the model that was trained on samples generated from both molecular dynamics (MD) and Monte Carlo (MC) simulations. It is worth noting that except lociPARSE, RNA3DCNN is the only other method that estimates both local and global quality.

2.7. Performance Assessment. To assess the accuracy of different aspects of quality estimation, we use a wide range of performance measures including the ability to reproduce the ground truth IDDT and all-atom RMSD scores both at the global and local levels, rank predictions for a given target, recognize the best predictions, and discriminate between "good" and "bad" predictions. Our performance assessment can be broadly grouped into two categories: global-level assessment and per-target average assessment. The global level assessment puts together the predicted scores of all the structural models across all targets during performance evaluation. In contrast, the per-target average assessment evaluates the predictions for each target's structures separately against their corresponding ground truths and then averages the results over all the targets. Since the global level assessment puts all the targets together to evaluate, it is important to length-normalize the estimated scores for the methods that predict RMSD unfitness scores. For the same reason, we length-normalize the ground truth all-atom RMSD metric during scoring performance evaluation. For the length normalization of RMSD in the range of (0-1], we use the formulation of US-align⁴⁷ as follows:

$$d_{0} = 1.24 \sqrt[3]{L} - 15 - 1.8$$

RMSD_{norm} = $\frac{1}{1 + \left(\frac{\text{RMSD}}{d_{0}}\right)^{2}}$ (10)

Except for lociPARSE, all the other six competing methods evaluated in this work estimate some form of structural unfitness score with a lower value representing better structural quality. ARES and RNA3DCNN estimate RMSD, whereas the four other knowledge-based statistical potentials estimate the potential energy of an RNA structural model. To ensure that the estimated scores from all competing methods are comparable, we length-normalize the predicted RMSD of ARES and RNA3DCNN using eq 10 that maps the predictions between (0-1] with higher values representing better structural quality. For the knowledge-based statistical potentials, we use min-max normalization to map the estimated potential energy to a normalized score *s* between (0-1] and use (1 - s) to recalibrate the score such that a higher value represents better structural quality. We also length-normalize the ground truth RMSD between (0-1] using eq 10. Our method lociPARSE estimates the lDDT score between (0-1] where higher values represent better structural quality. Meanwhile, the ground truth lDDT is between (0-1] by definition.

Our assessment metrics include both global and average pertarget Pearson's (r), Spearman's rank (ρ), and Kendall's Tau rank (τ) correlation coefficients between the estimated score of each method and the ground truth lDDT as well as ground truth all-atom RMSD where a higher correlation indicates better performance. While Pearson's r assumes a linear relationship between variables and is sensitive to outliers, Spearman's ρ and Kendall's τ account for nonlinear but monotonic trends and are less affected by outliers. Also, Kendall's τ is particularly robust for small sample sizes and many tied ranks. Therefore, reporting all three coefficients provides a comprehensive assessment of the predicted quality scores of different methods. Ground truth IDDT of the test set targets are calculated using the docker version of Open-Structure⁴² whereas the ground truth all-atom RMSD is calculated using casp-rna pipeline.⁴ All the predicted quality scores and the ground truth metrics are normalized between (0-1] during scoring performance evaluation. Diff, another assessment metric is calculated at the global level as the mean absolute difference between the structural level estimated scores (pMoL) of all the structures of all targets and their corresponding ground truth IDDT or RMSD metrics. Pertarget loss or top-1 loss is calculated as the absolute difference between the ground truth IDDT or RMSD of the structural model ranked at the top by each method and the ground truth IDDT or RMSD of the most accurate structural model for each target which is then averaged over all targets to get the loss value for each method over a test set. Lower values of diff and loss, therefore, indicate better performance. To calculate the local nucleotide-wise scoring performance in Table 5, we accumulated all the nucleotide-wise predicted scores (pNuL) and computed the coefficients and diff in the same manner as in our global assessment. For RNA3DCNN, we adhered to the method described in their work to scale the nucleotide-level predicted RMSD in the range (0-1] before the evaluation. We additionally performed receiver operating characteristics (ROC) analysis using a lDDT threshold of 0.75 to separate "good" and "bad" structural models, following.⁴ Consequently, the area under the ROC curve (AUC) quantifies the ability of a scoring function to distinguish good and bad models. Finally, an average of all assessment metrics is taken to combine the results of all the different metrics into a single composite quality score, called Q_o, defined as

$$Q_{c} = \frac{1}{5}(r_{g} + (1 - D) + r_{a} + (1 - L) + AUC)$$
(11)

where, $r_g = \text{global Pearson's } r$, D = global diff, $r_a = \text{per-target}$ average Pearson's r, L = average loss, and AUC = area under the ROC curve. We use the composite quality score for

Article

Table 1. Performance on 30 Independent RNA Targets Based on IDDT as the Ground Truth Metric, Sorted in Nonincreasing Order of Global Pearson's r^{a}

	global					per-target average				
method	$r\uparrow$	$\rho\uparrow$	τ \uparrow	diff ↓	AUC ↑	$r\uparrow$	$\rho\uparrow$	τ \uparrow	loss ↓	
lociPARSE	0.67	0.71	0.55	0.06	0.91	0.77	0.75	0.64	0.06	
rsRNASP	0.5	0.5	0.36	0.11	0.79	0.75	0.69	0.55	0.06	
RASP	0.46	0.5	0.36	0.16	0.78	0.72	0.68	0.58	0.08	
ARES	0.43	0.48	0.34	0.55	0.80	0.72	0.7	0.58	0.06	
DFIRE-RNA	0.33	0.32	0.22	0.19	0.69	0.75	0.7	0.56	0.05	
cgRNASP	0.27	0.19	0.13	0.15	0.60	0.12	0.07	0.07	0.07	
RNA3DCNN	0.19	0.17	0.12	0.64	0.58	0.47	0.36	0.27	0.06	
^{<i>a</i>} Values in bold indic	ate the best r	erformance.								

Table 2. Performance on 30 Independent RNA Targets Based on All-Atom RMSD as the Ground Truth Metric, Sorted in Nonincreasing Order of Global Pearson's r^a

		glob	al	per-target average				
method	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	diff ↓	$r\uparrow$	$\rho\uparrow$	τ \uparrow	loss ↓
lociPARSE	0.65	0.65	0.47	0.31	0.59	0.61	0.49	1.47
ARES	0.64	0.64	0.48	0.3	0.58	0.61	0.48	2.0
RASP	0.64	0.65	0.47	0.21	0.59	0.56	0.44	2.37
rsRNASP	0.63	0.65	0.47	0.26	0.52	0.52	0.37	1.93
DFIRE-RNA	0.54	0.53	0.39	0.22	0.49	0.49	0.35	1.88
cgRNASP	0.49	0.48	0.33	0.29	0.18	0.09	0.09	1.59
RNA3DCNN	-0.11	-0.12	-0.07	0.38	0.06	0.12	0.11	2.13
Values in hold indicate the best performance								

Values in bold indicate the best performance.

ablation study and hyperparameter selection, where higher values of Q_c indicate better performance.

3. RESULTS

3.1. Performance on 30 Independent RNA Targets. The performance of our new method lociPARSE and the other competing methods on 30 independent RNA targets in terms of ground truth as IDDT and RMSD is reported in Tables 1 and 2, respectively. Table 1 shows that lociPARSE consistently outperforms all other tested methods across almost all performance criteria when lDDT is used as the ground truth metric. For instance, lociPARSE attains the highest global Pearson's r of 0.67 which is much better than the second-best rsRNAsp (0.5). The same trend continues for global Spearman's ρ (lociPARSE: 0.71 vs the second-best score of 0.5 and global Kendall's τ (lociPARSE: 0.55 vs the second-best score of 0.36). Additionally, lociPARSE attains the lowest diff of 0.06, which is lower than the second-best rsRNASP (0.11). Furthermore, lociPARSE always delivers the highest per-target average correlations. In terms of average lDDT loss, however, DFIRE-RNA attains the lowest average loss (0.05). Meanwhile, lociPARSE, ARES, rsRNASP and RNA3DCNN are tied for the second spot with a comparably low loss of 0.06. To investigate the ability of lociPARSE to distinguish "good" and "bad" models in comparison with the other tested methods across all structural models for all 30 targets, we performed receiver operating characteristics (ROC) analysis using a cutoff of lDDT = 0.75 to differentiate "good" and "bad" models following the CASP15 official assessment.⁴ Meanwhile, the area under the ROC curve (AUC) quantifies the ability of a method to differentiate "good" and "bad" models. Table 1 shows that lociPARSE achieves the highest AUC of 0.91, which is noticeably better than the second-best AUC of 0.8 by ARES, demonstrating a better distinguishability aspect of our

method. The ROC curves for all competing methods across both test sets are provided in Supplementary Figure S6.

Since our method lociPARSE is trained to estimate the 1DDT scores, whereas methods such as ARES and RNA3DCNN are trained to estimate the RMSD scores, to ensure a fair performance evaluation, we perform an analogous set of assessments using the all-atom RMSD as the ground truth metric instead of lDDT. As reported in Table 2, lociPARSE exhibits remarkable robustness and performance resilience by achieving comparable correlations in both global and per-target levels with the lowest per-target loss, even when evaluated based on RMSD as the ground truth. To further investigate whether each method can effectively sort and rankorder the structures, we analyze the median lDDT and RMSD ground truth metrics of top-1 and best of top-10 ranked structural models from each method. As shown in Supplementary Figures S1 and S2, lociPARSE achieves the highest median top-1 lDDT score of 0.8 and the lowest median top-1 RMSD score of 1.91, demonstrating its effectiveness in identifying the optimal structural model irrespective of the choice of the ground truth metrics. Considering the best of top-10 median scores, lociPARSE is only 0.01 points lower than the highest median lDDT score and achieves the lowest median RMSD of 1.73. In summary, the results demonstrate that lociPARSE is effective in sorting and rank-ordering structural models while being robust and versatile in terms of the choice of the ground truth assessment metrics.

It is worth mentioning here that both the competing machine learning-based scoring functions ARES and RNA3DCNN exhibit inferior global diff values despite attaining comparable performance in terms of per-target averages. Meanwhile, DFIRE-RNA, the method attaining the lowest average per-target IDDT loss, does not deliver top performance in terms of global correlations in Table 1. That is, there are complementary aspects of scoring and model quality

global							per-targe	t average	
method	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	diff ↓	AUC ↑	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	loss ↓
lociPARSE	0.74	0.72	0.55	0.12	0.96	0.73	0.66	0.52	0.07
ARES	0.33	0.35	0.25	0.27	0.78	0.55	0.53	0.42	0.13
RASP	0.33	0.42	0.3	0.17	0.80	0.55	0.54	0.42	0.17
rsRNASP	0.31	0.41	0.3	0.11	0.83	0.66	0.61	0.47	0.11
DFIRE-RNA	0.29	0.3	0.22	0.39	0.77	0.69	0.6	0.46	0.16
cgRNASP	0.29	0.32	0.24	0.12	0.80	0.59	0.5	0.37	0.11
RNA3DCNN	0	-0.01	-0.01	0.52	0.35	0.5	0.42	0.31	0.09
^a Values in hold indic	cata tha hast i	porformanco							

Table 3. Performance on CASP15 RNA Targets Based on LDDT as the Ground Truth Metric, Sorted in Nonincreasing Order of Global Pearson's r^{a}

Values in bold indicate the best performance.

Table 4. Performance on CASP15 RNA Targets Based on All-Atom RMSD as the Ground Truth Metric, Sorted in Nonincreasing Order of Global Pearson's r^a

		glo	bal	per-target average				
method	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	diff \downarrow	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	loss ↓
lociPARSE	0.37	0.34	0.24	0.58	0.33	0.32	0.24	9.74
ARES	0.15	0.11	0.09	0.23	0.22	0.25	0.18	11.83
rsRNASP	0.13	0.21	0.14	0.49	0.34	0.37	0.27	14.37
cgRNASP	0.11	0.13	0.09	0.41	0.28	0.31	0.23	14.79
RASP	0.1	0.09	0.07	0.63	0.18	0.2	0.15	16.79
DFIRE-RNA	0.06	0.11	0.09	0.86	0.32	0.37	0.27	15.31
RNA3DCNN	0.01	0.04	0.03	0.05	0.19	0.15	0.11	13.39
4 Values in hold indicate the best performance.								

estimation that can lead to performance trade-offs. Our new method lociPARSE strikes an ideal balance to deliver a wellrounded RNA scoring performance across a wide range of assessment metrics. It is interesting to note that among the other tested methods, the two machine learning-based scoring functions ARES and RNA3DCNN show dramatically different performance. While ARES is comparable to the traditional knowledge-based statistical potentials in terms of global correlations and per-target average correlations, RNA3DCNN exhibits poor global and per-target average correlations, which are much lower than most knowledge-based statistical potentials. A similar trend can be observed between rsRNASP and its coarse-grained counterpart cgRNASP, where rsRNASP consistently attains good global and per-target average correlations but cgRNASP falls short. That is, subtle methodological differences such as the granularity of RNA conformational space representation or the choice of the neural network architecture can lead to dramatic differences in performance. Notably, lociPARSE exhibits remarkable robustness by being resilient to the choice of the ground truth assessment metrics or complementary aspects of scoring performance evaluation due to various factors including our novel use of the locality-aware IPA architecture as a generalpurpose RNA scoring function, invariant set of features, and the invariant nature of the IDDT metric that lociPARSE is trained to estimate, leading to high-fidelity RNA scoring performance.

3.2. Performance on CASP15 RNA Targets. Tables 3 and 4 report the performance of lociPARSE and the other competing methods on 12 CASP15 RNA targets based on IDDT and all-atom RMSD as the ground truth metrics, respectively. As shown in Table 3, the performance of lociPARSE for the global level assessment based on lDDT as the ground truth metric is considerably better, having the highest global Pearson's r of 0.74, which is much better than

the second-best method ARES (0.33). lociPARSE is the second-best method in terms of global diff, only slightly worse by 0.01 points than the lowest global diff. lociPARSE also achieves the highest AUC value of 0.96, outperforming the second-best method rsRNASP (0.83) by a large margin, demonstrating better distinguishability of lociPARSE in separating "good" and "bad" models over a diverse set of predicted structural models submitted by all CASP15 predictors. Furthermore, lociPARSE consistently attains higher per-target average correlations than the other competing methods and achieves the lowest average lDDT loss of 0.07, which is lower than the second-best RNA3DCNN (0.09).

Table 4 reports the full set of results based on all-atom RMSD as the ground truth metric that the two competing machine learning-based scoring functions, ARES and RNA3DCNN, are trained on. lociPARSE is better than both methods in terms of both global and per-target correlations with the lowest average RMSD loss but exhibits comparatively higher global diff. The global correlations of lociPARSE are noticeably better than all competing methods and comparable to most of the energy-based methods in terms of per-target assessment. It is interesting to note that DFIRE-RNA, the method attaining the lowest lDDT loss in 30 independent RNA targets, yields a poor IDDT loss (0.16) in CASP15. By contrast, lociPARSE consistently attains low loss in both test sets, indicating its ability to select the best model that generalizes across different datasets. Supplementary Figures S1 and S2 further demonstrate that the median lDDT score of the top-ranked structure on 12 CASP15 targets by lociPARSE is 0.69, higher than all other methods. In terms of best of top 10 predictions, lociPARSE is second-best with a median lDDT of 0.72 but jointly best in terms of median all-atom RMSD value of 8.29. In summary, the results underscore the ability of lociPARSE to consistently select high-quality structures from a diverse pool of structural models.

When nucleotide-wise quality is evaluated, lociPARSE is orders of magnitude better than RNA3DCNN, the only other method except for lociPARSE that can estimate per-nucleotide score. For example, lociPARSE attains more than three times higher global Pearson's r, Spearman's ρ , and Kendall's τ than RNA3DCNN, and achieves noticeably lower diff than that attained by RNA3DCNN (Table 5). While in principle, the

Table 5. Nucleotide-Wise Scoring Performance on CASP15 Set^a

	nucleotide-wise						
method	$r\uparrow$	$\rho\uparrow$	$\tau\uparrow$	diff ↓			
lociPARSE	0.72	0.73	0.53	0.15			
RNA3DCNN	0.17	0.19	0.13	0.24			
^a Values in bold indicate the best performance.							

nucleotide-wise scoring performance reported in Table 5 should be synchronized to the structural level performance reported in Table 3 as it is the case for our method lociPARSE, RNA3DCNN shows a discrepancy in this regard, possibly due to inconsistencies in nucleotide-wise scoring performance. Overall, lociPARSE delivers stable and consistent nucleotidewise scoring performance that translates to well-rounded structural level performance.

3.3. Case Study. Figure 2 shows a representative example of nucleotide-wise scoring performance using lociPARSE for a top-ranked structural model submitted by the winning group AIchemy_RNA2 (group 232) for the CASP15 target R1108 having a length of 69. The predicted nucleotide-wise IDDT (pNuL) scores are in close agreement with the ground truth IDDT with a high Pearson's r of 0.89 (Figure 2a). Two local problematic regions are estimated by lociPARSE in nucleotide positions (19–27) and (59–63). These two local problematic regions are visually noticeable when the predicted structural model is aligned with the experimental structure.

The poorly modeled structural regions around the hairpin loop in nucleotide positions (19-27) and part of the helix strand in positions (59-63) are obvious even with simple visual inspection (Figure 2b). By contrast, virtually all nucleotides with high pNuL values are structurally well modeled. A comparison with ARES on this target reveals the benefit of nucleotide-wise IDDT prediction as employed by our method over structural level RMSD prediction of ARES,

especially in the presence of irregular regions such as the loop in nucleotide positions (19-27). The predicted RMSD value by ARES is 7.8 Å which is noticeably higher than the ground truth all-atom RMSD of 4.63 Å. However, a structural level unfitness score alone fails to pinpoint the problematic regions in the model contributing to the higher RMSD. In contrast, lociPARSE demonstrates its strength by accurately estimating the quality of each nucleotide, effectively identifying the incorrectly modeled loop region in nucleotide positions (19-27) highlighted in blue. To further assess whether the learned representation in lociPARSE is consistent with the physical principles of RNA 3D structure, we analyze the predicted attention map extracted from one of the attention heads in the final IPA layer of our trained model for this same structural model. A side-by-side comparison against the base pairing information extracted from the RNA 3D structure presented in Supplementary Figure S5 shows a resemblance between the two maps, indicating the effectiveness of the learned representations by lociPARSE.

3.4. Ablation Study and Hyperparameter Selection. To examine the relative importance of the features and architectural hyperparameters adopted in lociPARSE, we conduct ablation experiments by systematically varying individual parameters during model training using the reduced training set and evaluating the accuracy of the independent validation set (Section 2.4). Table 6 reports the composite quality score (Q_c) defined in Section 2.7 of the full-fledged version of lociPARSE serving as a baseline and its ablated variants. The results demonstrate that all the parameters adopted in the full-fledged version of lociPARSE positively contribute to the overall accuracy achieved by lociPARSE. For example, we notice a performance decline when we vary the value of K used in the nearest neighbors for the locality computation from K = 20 used in the baseline to $K \in \{5, 10, 10\}$ 30, 40}. Furthermore, to bias the attention weights as well as to update scalar features, we make use of the pair features in the form of nucleotide-nucleotide atomic distances between all pairs of 3 atoms P, C4' and N, encoded with Gaussian radial basis functions (hereafter called $pair_{\varphi(d)}$). We notice a significant performance drop when $pair_{\varphi(d)}$ features are isolated as well as changing the number of radial basis functions from RBF = 16 used in the baseline to $RBF \in \{1, 2, 4, 8, 32\}$. Similarly, we notice consistent performance decline from the baseline configuration whenever we vary the network



Figure 2. Example of lociPARSE nucleotide-wise quality estimation for the CASP15 target R1108. (a) Predicted nucleotide-wise lDDT (pNuL) vs the ground truth lDDT for the top-ranked structural model submitted by AIchemy_RNA2 (group 232). (b) Predicted structural model in rainbow colored with color code ramping from blue to red for low to high pNuL values superimposed on the experimental structure in gray and two local problematic regions are highlighted.

Article

Table 6. Validation Set Performance in Terms of Composite Quality Score (Q_c) with Various Settings of Features (on the left) and Hyperparameters (on the right) Compared to the Full-Fledged Version of lociPARSE Serving as a Baseline^{*a*}

settings	Q _c	settings	RBF	hyperparamters	$Q_{\rm c}$	hyperparamters	$Q_{\rm c}$
baseline ($K = 20 \text{ w/pair}_{\varphi(d)}$)	0.764	baseline (RBF = 16)	0.764	baseline $(L = 4, H = 4)$	0.764	baseline $(Q = 8, V = 4)$	0.764
$K = 5 \text{ w/pair}_{\varphi(d)}$	0.744	RBF = 1	0.72	L = 2	0.75	Q = 4	0.762
$K = 10 \text{ w/pair}_{\varphi(d)}$	0.762	RBF = 2	0.742	L = 6	0.756	Q = 16	0.762
$K = 30 \text{ w/pair}_{\varphi(d)}$	0.754	RBF = 4	0.74	H = 2	0.748	V = 2	0.746
$K = 40 \text{ w/pair}_{\varphi(d)}$	0.756	RBF = 8	0.758	H = 8	0.758	V = 8	0.754
$K = 20$ w/o pair _{$\varphi(d)$}	0.728	RBF = 32	0.754	Q = 2	0.754	V = 16	0.758
^a Values in bold indicate the b	est perfor	mance.					

architecture such as the number of IPA layers $(N_{\text{layers}} = L)$ or various IPA hyperparameters $(N_{\text{heads}} = H, N_{\text{query points}} = Q, N_{\text{value points}} = V)$, justifying our choice of the parameters adopted in the full-fledged version of lociPARSE.

4. DISCUSSION

In this work, we developed lociPARSE, a locality-aware invariant point attention model for scoring RNA 3D structures. lociPARSE uses locality information derived from the RNA atomic coordinates to define nucleotide-wise frames together with its local atomic environment. This, coupled with the invariant point attention architecture, allows for the simultaneous estimation of local quality in the form of predicted nucleotide-wise lDDT (pNuL) scores which are then aggregated over all nucleotides by taking an average to estimate global structural correctness in the form of predicted molecular-level lDDT (pMoL). Our empirical results demonstrate the superiority of our method in scoring RNA 3D structures compared to existing approaches.

Our locality-aware attention-based architecture can be extended in several ways, including estimating other local quality measures such as the Interaction Network Fidelity (INF) score,⁴⁸ which is a local interaction metric that captures various types of base-base interactions in RNA. In fact, INF and IDDT have been shown to correlate well in a near-linear and size-independent relationship,⁴ suggesting that lDDT may capture the subset of interactions measured in INF whereas INF focuses on a selection of RNA-specific interactions. A model with a very similar architecture as lociPARSE would make an excellent candidate for jointly estimating INF and IDDT, thereby capturing complementary aspects of local quality. Further, a promising direction for future work is to investigate the potential benefits of capturing the multistate conformational landscape of RNA, since many RNA targets exhibit conformational flexibility.⁴ The lDDT score can be computed simultaneously against multiple reference structures of the same RNA at the same time, without arbitrarily selecting one reference structure for the target or removing parts that show variability. Training our model using multireference IDDT to capture different classes of conformations will allow our scoring function to account for conformational flexibility and pave the way to evaluate predictions of conformational ensembles instead of just a single structure. One limitation of our method is that it does not account for the stereochemical quality and physical plausibility of the model being evaluated. This is because, unlike proteins, the currently available implementation of lDDT for RNA does not penalize stereochemical violations. Using a customized version of IDDT that incorporates stereochemical quality checks in its calculation can address such limitations, and this aspect remains an important future direction.

ASSOCIATED CONTENT

Data Availability Statement

Source code, trained model and detailed instruction to use lociPARSE is available at https://github.com/Bhattacharya-Lab/lociPARSE. The training and test datasets are available at https://zenodo.org/records/12729654.

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jcim.4c01621.

List of the number of decoys per RNA target in training and test datasets; list of the number of decoys per RNA target in CASP15 test set, a demonstration of the ability to rank the structures in terms of IDDT and RMSD; nucleotide-wise ground truth IDDT distribution of training and test sets; interpretation of attention map; and distinguishability of the assessment methods on both test sets (PDF)

AUTHOR INFORMATION

Corresponding Author

Debswapna Bhattacharya – Department of Computer Science, Virginia Tech, Blacksburg, Virginia 24061, United States; orcid.org/0000-0002-9630-0141; Email: dbhattacharya@vt.edu

Author

Sumit Tarafder – Department of Computer Science, Virginia Tech, Blacksburg, Virginia 24061, United States

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jcim.4c01621

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was partially supported by the National Institute of General Medical Sciences (R35GM138146 to D.B.) and the National Science Foundation (DBI2208679 to D.B.).

REFERENCES

(1) Miao, Z.; Westhof, E. RNA structure: advances and assessment of 3D structure prediction. *Annual review of biophysics* **2017**, *46*, 483–503.

(2) Keating, K. S.; Humphris, E. L.; Pyle, A. M. A new way to see RNA. *Q. Rev. Biophys.* **2011**, *44*, 433–466.

(3) Zhang, J.; Lang, M.; Zhou, Y.; Zhang, Y. Predicting RNA structures and functions by artificial intelligence. *Trends Genet.* **2024**, 40, 94.

(4) Das, R.; et al. Assessment of three-dimensional RNA structure prediction in CASP15. *Proteins: Struct., Funct., Bioinf.* **2023**, *91*, 1747–1770.

(5) Miao, Z.; et al. RNA-Puzzles Round IV: 3D structure predictions of four ribozymes and two aptamers. *RNA* **2020**, *26*, 982–995.

(6) Kryshtafovych, A.; et al. New prediction categories in casp15. *Proteins: Struct., Funct., Bioinf.* **2023**, *91*, 1550–1557.

(7) Ou, X.; Zhang, Y.; Xiong, Y.; Xiao, Y. Advances in rna 3d structure prediction. J. Chem. Inf. Model. 2022, 62, 5862–5874.

(8) Capriotti, E.; Marti-Renom, M. A.; et al. Computational RNA structure prediction. *Curr. Bioinform* **2008**, *3*, 32–45.

(9) Laing, C.; Schlick, T. Computational approaches to 3D modeling of RNA. *J. Phys.: Condens. Matter* **2010**, *22*, No. 283101.

(10) Li, J.; Chen, S.-J. RNA 3D structure prediction using coarsegrained models. *Frontiers in Molecular Biosciences* **2021**, *8*, No. 720937.

(11) Pearce, R.; Omenn, G. S.; Zhang, Y. De novo RNA tertiary structure prediction at atomic resolution using geometric potentials from deep learning. *bioRxiv* 2022, 1755.

(12) Wang, W.; et al. trRosettaRNA: automated prediction of RNA 3D structure with transformer network. *Nat. Commun.* 2023, 14, 7266.

(13) Li, Y.; et al. Integrating end-to-end learning with deep geometrical potentials for ab initio RNA structure prediction. *Nat. Commun.* **2023**, *14*, 5745.

(14) Baek, M.; et al. Accurate prediction of protein-nucleic acid complexes using rosettafoldna. *Nat. Methods* **2024**, *21*, 117–121.

(15) Shen, T. E2Efold-3D: end-to-end deep learning method for accurate de novo RNA 3D structure prediction, arXiv preprint arXiv:2207.01586, 2022.

(16) Wang, X.; Yu, S.; Lou, E.; Tan, Y.-L.; Tan, Z.-J. RNA 3D Structure Prediction: Progress and Perspective. *Molecules* **2023**, *28*, 5532.

(17) Li, J.; et al. RNA3DCNN: Local and global quality assessments of RNA 3D structures using 3D deep convolutional neural networks. *PLoS computational biology* **2018**, *14*, No. e1006514.

(18) Micheletti, C.; Di Štefano, M.; Orland, H. Absence of knots in known RNA structures. *Proc. Natl. Acad. Sci. U. S. A.* **2015**, *112*, 2052–2057.

(19) Das, R.; Baker, D. Automated de novo prediction of native-like RNA tertiary structures. *Proc. Natl. Acad. Sci. U. S. A.* **2007**, *104*, 14664–14669.

(20) Tan, Y.-L.; Wang, X.; Shi, Y.-Z.; Zhang, W.; Tan, Z.-J. rsRNASP: A residue-separation-based statistical potential for RNA 3D structure evaluation. *Biophys. J.* **2022**, *121*, 142–156.

(21) Tan, Y.-L.; Wang, X.; Yu, S.; Zhang, B.; Tan, Z.-J. cgRNASP: coarse-grained statistical potentials with residue separation for RNA structure evaluation. *NAR Genom. Bioinf.* **2023**, *5*, lqad016.

(22) Capriotti, E.; Norambuena, T.; Marti-Renom, M. A.; Melo, F. All-atom knowledge-based potential for RNA structure prediction and assessment. *Bioinformatics* **2011**, *27*, 1086–1093.

(23) Zhang, T.; Hu, G.; Yang, Y.; Wang, J.; Zhou, Y. All-atom knowledge-based potential for RNA structure discrimination based on the distance-scaled finite ideal-gas reference state. *Journal of computational biology* **2020**, *27*, 856–867.

(24) Wang, J.; Zhao, Y.; Zhu, C.; Xiao, Y. 3dRNAscore: a distance and torsion angle dependent evaluation function of 3D RNA structures. *Nucleic acids research* **2015**, *43*, No. e63.

(25) Samudrala, R.; Moult, J. An all-atom distance-dependent conditional probability discriminatory function for protein structure prediction. *Journal of molecular biology* **1998**, *275*, 895–916.

(26) Lu, H.; Skolnick, J. A distance-dependent atomic knowledgebased potential for improved protein structure selection. *Proteins: Struct., Funct., Bioinf.* **2001**, *44*, 223–232.

(27) Rykunov, D.; Fiser, A. New statistical potential for quality assessment of protein models and a survey of energy functions. *BMC Bioinf.* **2010**, *11*, 128.

(28) Zhou, H.; Zhou, Y. Distance-scaled, finite ideal-gas reference state improves structure-derived potentials of mean force for structure selection and stability prediction. *Protein science* **2002**, *11*, 2714–2726.

(29) Shen, M.-Y.; Sali, A. Statistical potential for assessment and prediction of protein structures. *Protein science* **2006**, *15*, 2507–2524.

(30) Zhang, J.; Zhang, Y. A novel side-chain orientation dependent potential derived from random-walk reference state for protein fold selection and structure prediction. *PloS one* **2010**, *5*, No. e15386.

(31) Townshend, R. J.; et al. Geometric deep learning of RNA structure. *Science* 2021, 373, 1047–1051.

(32) Mariani, V.; Biasini, M.; Barbato, A.; Schwede, T. IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* **2013**, *29*, 2722–2728.

(33) Jumper, J.; et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **2021**, *596*, 583–589.

(34) Jian, Y.; et al. DIRECT: RNA contact predictions by integrating structural patterns. *BMC Bioinf.* **2019**, *20*, 497.

(35) Weinreb, C.; et al. 3D RNA and functional interactions from evolutionary couplings. *Cell* **2016**, *165*, 963–975.

(36) Fu, L.; Niu, B.; Zhu, Z.; Wu, S.; Li, W. CD-HIT: accelerated for clustering the next-generation sequencing data. *Bioinformatics* **2012**, 28, 3150–3152.

(37) Boniecki, M. J.; et al. SimRNA: a coarse-grained method for RNA folding simulations and 3D structure prediction. *Nucleic acids research* **2016**, *44*, No. e63.

(38) Chaudhury, S.; Lyskov, S.; Gray, J. J. PyRosetta: a script-based interface for implementing molecular modeling algorithms using Rosetta. *Bioinformatics* **2010**, *26*, 689–691.

(39) Conway, P.; Tyka, M. D.; DiMaio, F.; Konerding, D. E.; Baker, D. Relaxation of backbone bond geometry improves protein energy landscape modeling. *Protein science* **2014**, *23*, 47–55.

(40) Tarafder, S.; Roche, R.; Bhattacharya, D. The landscape of RNA 3D structure modeling with transformer networks. *Biol. Methods Protoc.* **2024**, *9*, bpae047.

(41) Zhang, Y.; Skolnick, J. SPICKER: a clustering approach to identify near-native protein folds. *Journal of computational chemistry* **2004**, *25*, 865–871.

(42) Biasini, M.; et al. OpenStructure: an integrated software framework for computational structural biology. *Acta Crystallographica Section D: Biological Crystallography* **2013**, *69*, 701–709.

(43) Paszke, A. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems;* ACM: 2019; Vol. 32.

(44) Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.

(45) Brion, P.; Westhof, E. Hierarchy and dynamics of RNA folding. *Annual review of biophysics and biomolecular structure* **1997**, *26*, 113–137.

(46) Geiger, M.; Smidt, T. e3nn: Euclidean neural networks. arXiv preprint arXiv:2207.09453, 2022.

(47) Zhang, C.; Shine, M.; Pyle, A. M.; Zhang, Y. Us-align: universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nat. Methods* **2022**, *19*, 1109–1115.

(48) Parisien, M.; Cruz, J. A.; Westhof, E.; Major, F. New metrics for comparing and assessing discrepancies between RNA 3D structures and models. *RNA* **2009**, *15*, 1875–1885.