

RNAbpFlow: Base pair-augmented SE(3)-flow matching for conditional RNA 3D structure generation

Sumit Tarafder¹ and Debswapna Bhattacharya^{1,*}

¹Department of Computer Science, Virginia Tech, Blacksburg, Virginia, 24061, USA

*Corresponding author. dbhattacharya@vt.edu

Abstract

Motivation: Despite the groundbreaking advances in deep learning-enabled methods for bimolecular modeling, predicting accurate three-dimensional (3D) structures of RNA remains challenging due to the highly flexible nature of RNA molecules combined with the limited availability of evolutionary sequences or structural homology.

Results: We introduce RNAbpFlow, a novel sequence- and base-pair-conditioned SE(3)-equivariant flow matching model for generating RNA 3D structural ensemble. Leveraging a nucleobase center representation, RNAbpFlow enables end-to-end generation of all-atom RNA structures without the explicit or implicit use of evolutionary information or homologous structural templates. Experimental results show that base pairing conditioning leads to broadly generalizable performance improvements over current approaches for RNA topology sampling and predictive modeling in large-scale benchmarking.

Availability: RNAbpFlow is freely available at <https://github.com/Bhattacharya-Lab/RNAbpFlow>.

Key words: RNA 3D structure modeling, generative modeling, deep learning, flow matching

1 Introduction

The determination of three-dimensional (3D) structures of RNA has become a crucial challenge in structural biology, driven by the growing interest in RNA-based therapeutics [1, 2]. High-resolution characterization of 3D RNA structure is essential for the design and understanding of RNA molecules with specific therapeutic functions [3], thus expanding the scope of RNA-mediated drug discovery [4]. However, the intrinsic conformational flexibility of RNA presents significant challenges for experimental structure determination methods such as X-ray crystallography, nuclear magnetic resonance (NMR) spectroscopy, and cryo-electron microscopy (Cryo-EM). Computational RNA structure prediction, therefore, is emerging as an attractive alternative to fill the gap in RNA structure space and to elucidate RNA conformational dynamics that underpins diverse cellular processes [5].

Traditional RNA 3D structure prediction methods include template-based approaches like ModeRNA [6] and RNAbuilder [7], which rely on homologous structural information, as well as physics- and/or knowledge-based methods such as FARFAR2 [8], 3dRNA [9], RNAComposer [10], and Vfold3D [11], which exploit biophysical potentials and pre-built fragment libraries to assemble full-length RNA structures. However, these approaches are constrained by the scarcity of RNA structural data in the PDB [12] and are often computationally prohibitive, making them less suitable for the prediction of large RNAs with complex topologies [13, 14]. Although physics-based methods combined with expert human intervention have demonstrated success in community-wide blind RNA-Puzzles [15] and CASP (Critical Assessment of Structure Prediction) challenges [16], there remains a critical need for fully automated, fast and accurate methods for computational modeling of RNA structures.

Inspired by the transformative impact of AlphaFold2 [17] on protein structure prediction [18], a growing number of deep learning-based methods have recently been developed for modeling RNA structures including DRfold [13], trRosettaRNA [14], RoseTTAFoldNA [19], RhoFold+ [20], and NuFold [21] by leveraging attention-powered transformer architectures [22]. However, except for DRfold, most of these methods are highly dependent on explicit evolutionary sequence information derived from multiple sequence alignments (MSA) or implicitly make use of homologous information learned by biological language models [20]. Obtaining reliable MSAs for RNA sequences poses significant challenges due to the isosteric nature of base pair interactions, hindering sequence alignment efforts [23]. Furthermore, many existing methods fail to fully leverage the RNA base pair (2D) information, including canonical and non-canonical base pairing interactions, key determinants of the final 3D conformation of RNA [24, 25]. Finally, the static structure predictions made by these approaches might be inadequate to capture the inherent conformational flexibility of an RNA molecule that often adopts a distribution of conformational states instead of folding into a static structure [5, 26]. Thus, there is an urgent need to develop improved computational methods that can generate a conformational ensemble of all-atom RNA 3D structures directly from the nucleotide sequence by making use of base pairing information without explicitly or implicitly using any evolutionary information.

Diffusion-based generative modeling has achieved remarkable success in the image domain and is now attracting significant attention in structural bioinformatics, most notably in AlphaFold3 [27] for biomolecular interaction prediction, as well as 3D protein backbone generation with approaches such as RFdiffusion [28] and FrameDiff [29]. More recently, methods such as FrameFlow [30] have showcased the power of SE(3)-equivariant flow matching, achieving better precision




Fig. 1. Overview of RNAbpFlow. RNAbpFlow is a sequence- and base-pair-conditioned SE(3)-equivariant flow matching model for generating RNA 3D structural ensemble. Using the nucleotide sequence and base pairing information as conditions, our end-to-end framework enables efficient sampling of all-atom RNA 3D structures based on a nucleobase center representation without explicitly or implicitly using any evolutionary information.

in designability than their diffusion-based counterpart but with reduced sampling costs. RNA-FrameFlow [31], a recent adaptation of FrameFlow for RNA, is the first generative model that specifically targets 3D RNA backbone generation. However, the method employs naïve unconditional flow matching, generating backbones without using sequence or base pairing information. Such a limitation highlights the opportunity to develop improved deep generative models for RNA that can efficiently sample a large conformational ensemble of all-atom RNA 3D structures by explicitly conditioning on the nucleotide sequence and base pairing information through conditional flow matching, thereby offering an elegant combination of principles-based and data-driven approach to RNA structural ensemble generation free from sequence- and structural-level homology.

Here we present RNAbpFlow, a novel sequence- and base-pair-conditioned SE(3)-equivariant flow matching model for generating all-atom RNA conformational ensemble. The major contributions of this work lie in the methodology development of the first conditional flow matching-based method for RNA 3D structure generation. First, RNAbpFlow incorporates conditions on the nucleotide sequence and base pairing information from three complementary base pair annotation methods to comprehensively capture canonical and non-canonical interactions. Second, by incorporating a nucleobase center representation that enables the optimization of angles of all rotatable bonds of nucleobases, RNAbpFlow directly outputs all-atom RNA structures in an end-to-end fashion, bypassing the need for a post hoc geometry optimization module which is impractical in the context of large-scale sample generation. Third, we introduce base pair-centric auxiliary loss functions to enable maximal realization of the canonical and non-canonical base pairing interactions. We empirically observe performance improvements when base pairing condition is introduced across a wide range of evaluation metrics, outperforming a recent Molecular Dynamics (MD) simulation-based global topology sampling method RNAJP [32] which explicitly considers base pairing and base stacking interactions. Additionally, RNAbpFlow generalizes well for sequence- and base-pair-conditioned RNA 3D structure prediction compared to several state-of-the-art methods. RNAbpFlow is freely available at <https://github.com/Bhattacharya-Lab/RNAbpFlow>.

2 Materials & Methods

2.1 Overview of the RNAbpFlow framework

An overview of our method, RNAbpFlow, is shown in **Figure 1**. Our framework is built on the foundations of FrameFlow [30], a flow matching formulation tailored for fast protein backbone generation on the SE(3) frame representation. To represent each nucleotide in an RNA sequence as a rigid body frame defined by a translation from the global origin and a rotation matrix, as well as for the full atomic RNA 3D structure generation in an end-to-end manner, we follow the nucleotide representation presented in NuFold [21]. The rotation matrix is constructed using the Cartesian coordinates of C1' as the origin of the local frame and O4' – C1' – C2' for orientation. Given an RNA sequence of length N , we begin with N such frames sampled from a Gaussian distribution as the starting point for our iterative sampling process. Our method incorporates conditions on the nucleotide sequence and base pairing information to guide the conformational sampling process to generate the all-atom RNA structure. The three base-frame atoms: O4', C1', and C2' are derived from the learned frame representation, while the first nitrogen of the base (N1 for pyrimidines, N9 for purines) is imputed using tetrahedral geometry. The remaining atoms are partitioned into ten frames, and the corresponding atomic coordinates are generated based on the iterative update of the frames using the 9 predicted torsion angles lying on the bonds that connect these frames, leading to an all-atom RNA structure.

2.2 Model Input

Our method operates independently of any MSA or template information, relying solely on sequence and base pairing (2D) information for conditional generation of RNA 3D structure. Specifically, given an RNA sequence of length N as input, we encode the nucleotide sequence using one-hot encoding, represented as a binary vector of 4 entries corresponding to the four types of nucleotides (A, U, C, G). For base pairing information during the training process, we extract 2D annotations from experimental (native) 3D structures using three different software tools: RNAView [33], MC-Annotate [34], and DSSR [35]. These methods capture diverse base pairing information, including canonical and non-canonical base pairings identified from the 3D coordinates, which we represent as three separate 2D binary maps. We use all three maps as input for the

bias term in our denoiser architecture as edge features. During sampling, in the absence of native base pairing information, we use sequence-based predicted base pairs from three different RNA 2D structure predictors, namely MXfold2 [36], SimFold [37], and ContextFold [38].

2.3 Network Architecture

2.3.1 RNA Frame representation

We represent each nucleotide in a geometric abstraction using the concept of rigid body frames. Each nucleotide frame in the form of a tuple is defined as a Euclidean transform $T = (r, x)$, where $r \in SO(3)$ is a rotation matrix and $x \in \mathbb{R}^3$ is the translation vector that can be applied to transform a position in local coordinates to a position in global coordinates.

2.3.2 SE(3) flow matching on frames

Flow matching (FM) [39] is a class of deep generative models in which the goal is to learn a velocity field (or flow) that matches the probability flow of the data distribution to transform a simple distribution such as Gaussian to the desired complex data distribution in high-dimensional space. Flow matching directly learns this velocity field that describes how points should move from the simple distribution to the target distribution without completely destroying the data distribution. By integrating an ordinary differential equation (ODE) over this learned vector field, FM offers simpler trajectories towards achieving target distribution offering huge computational speed-up for large-scale sample generation.

In the context of generative modeling, the geodesic path describes a smooth transformation of one probability distribution into another while minimizing distortion. The concept of geodesics generalizes the notion of shortest paths in non-Euclidean spaces, enabling efficient computation and interpretation. Given a noisy frame T_0 sampled from a simple prior density $p_0(T_0)$ and the experimental frame T_1 sampled from a target distribution $p_1(T_1)$, the geodesic path connecting two points $\mathbf{T}_0$ and $\mathbf{T}_1$ in a combination of *simple manifolds* such as $\mathbb{R}^3$ and $SO(3)$ can be expressed using exponential and logarithmic maps following the generalization of flow matching to Riemannian manifolds [40] in the following way:

$$\mathbf{T}_t = \exp_{\mathbf{T}_0}(t \cdot \log_{\mathbf{T}_0}(\mathbf{T}_1)), \quad (1)$$

where $\log_{\mathbf{T}_0}(\mathbf{T}_1)$ is a vector in the tangent space of $\mathbf{T}_0$ pointing toward $\mathbf{T}_1$, and $t \in [0, 1]$ parameterizes the sequence of probability distributions a.k.a *probability path* p_t between the two distributions p_0 and p_1 . The conditional flow T_t constructed in Eq. 1 can be decomposed into separate individual flows for the simplification of training procedure in the following way:

$$\begin{aligned} \text{Translations } (\mathbb{R}^3) : \quad \mathbf{x}_t &= (1-t)\mathbf{x}_0 + t\mathbf{x}_1, \\ \text{Rotations } (SO(3)) : \quad \mathbf{r}_t &= \exp_{\mathbf{r}_0}(t \log_{\mathbf{r}_0}(\mathbf{r}_1)). \end{aligned} \quad (2)$$

where the prior distribution $p_0(T_0)$ during training takes the form of $p_0(T_0) = IGSO3(\sigma = 1.5) \otimes \mathcal{N}(0, I_3)$ where random translation x_0 is sampled from the unit Gaussian distribution $\mathcal{N}(0, I_3)$ and random rotation is sampled from $IGSO3(\sigma = 1.5)$ following [41, 29, 30] for better performance. During training, the parameter t is sampled from $\mathcal{U}([0, 1 - \epsilon])$ where $\epsilon = 0.1$ is chosen for training stability.

2.3.3 Denoiser model

The objective of our flow matching method is to learn the parametrized vector field $\mathbf{u}_t$, which represents a smooth, time-dependent (t) map that generates an ordinary differential equation (ODE) that describes the transformation between two distributions: $p_0(T_0)$ (noisy frames) and $p_1(T_1)$ (ground truth frames). To learn this mapping from noisy samples to cleaner ones, we train a parameterized denoiser model $\mathbf{v}_\theta(\mathbf{T}_t, t)$ which will predict clean frames given corrupted ground truth frames $\mathbf{T}_t$ at time t . Following FrameFlow [30], we use the structure module from AlphaFold2 [17] as the neural architecture of our denoiser model.

2.3.4 Sampling strategy

During conditional sampling of RNA 3D structures, our framework takes a random initialization of the backbone frames $T = (r, x)$, where translation x is sampled from a unit Gaussian distribution $\mathcal{N}(0, I_3)$ in $\mathbb{R}^3$ and rotation r is sampled from a uniform distribution in $SO(3)$. During inference, instead of the linear scheduler for rotation matrices, we use the exponential scheduler e^{-ct} with $c = 10$ for better performance. Thus, our $SO(3)$ flow for rotation in Eq. 2 changes according to the following equation:

$$r_t = \exp_{r_0}\left(\left(1 - e^{-ct}\right) \log_{r_0}(r_1)\right) \quad (3)$$

Based on the specified number of timesteps, we choose a set of values for t where $t \sim \mathcal{U}([0, 1])$. Starting from the random set of frames, we iteratively update the frame representations using the predictions from previous steps using our learned denoiser model with the specified condition at each timestep t in the following way:

$$\begin{aligned} \text{Translations :} \quad x_t &= x_{t-1} + \frac{\Delta t}{1-t} \cdot (x_t - x_{t-1}), \\ \text{Rotations :} \quad r_t &= \exp_{r_{t-1}}\left(c \Delta t \cdot \log_{r_{t-1}}(r_t)\right). \end{aligned} \quad (4)$$

2.4 Training of RNAbpFlow

Our training objective contains multiple loss function components related to base pair conditions and all-atom structure generation. The primary loss function for this framework is the same SE(3) loss formulated in FrameFlow [30] namely, the vector field loss in SE(3). To train our denoiser model $\mathbf{v}_\theta$, two separate components of this loss are calculated for the predicted rotation $\hat{\mathbf{r}}_t \in SO(3)$ and translation $\hat{\mathbf{x}}_t \in \mathbb{R}^3$ given the corrupted frames $\mathbf{T}_t$ at time t as shown below.

$$L_{trans} = \mathbb{E}_{t, p_0(\mathbf{T}_0), p_1(\mathbf{T}_1)} \left[\frac{1}{(1-t)^2} \sum_{n=1}^N \left\| \hat{\mathbf{x}}_t^{(n)} - \mathbf{x}_1^{(n)} \right\|_{\mathbb{R}^3}^2 \right] \quad (5)$$

$$\begin{aligned} L_{rot} = \mathbb{E}_{t, p_0(\mathbf{T}_0), p_1(\mathbf{T}_1)} \left[\frac{1}{(1-t)^2} \sum_{n=1}^N \left\| \log_{\mathbf{r}_t^{(n)}}(\hat{\mathbf{r}}_t^{(n)}) \right. \right. \\ \left. \left. - \log_{\mathbf{r}_t^{(n)}}(\mathbf{r}_1^{(n)}) \right\|_{SO(3)}^2 \right] \quad (6) \end{aligned}$$

where N is the total number of frames. To predict the 9 torsion angles ϕ_n , we use an additional head and calculate the torsional loss in the following way by comparing with the experimental angles ϕ :

$$L_{tors} = \frac{1}{9N} \sum_{n=1}^N \sum_{\phi \in \Phi_n} \|\phi - \hat{\phi}\|^2 \quad (7)$$

Our base pair-augmented auxiliary loss function components are described below. At the 3D level, we optimize the predicted distance between the annotated base pairs (m, n) extracted from all three base pair annotation methods, termed the *bp3D* loss, as follows:

$$L_{bp3D} = \frac{1}{3} \sum_{i=1}^3 \frac{1}{N_i^{nbpairs}} \sum_{n,m=1}^{N_i^{nbpairs}} \|D_i^{(nm)} - \hat{D}_i^{(nm)}\|^2 \quad (8)$$

where $D_i^{(nm)} = \|C1'(m) - C1'(n)\|^2$ denotes the actual distance between the $C1'$ atoms of the nucleotide pair (m, n) denoted by the i^{th} 2D structure annotation method and $\hat{D}_i^{(nm)}$ is the predicted distance. At the 2D level, we utilize an additional head to reduce the dimensionality of the predicted pair features and compute the BCEWithLogitsLoss against the experimental 2D maps, termed the *bp2D* loss, as described below:

$$L_{bp2D} = - \sum_{i=1}^3 \left[SS^{(i)} \odot \hat{S}^{(i)} - \log(1 + e^{\hat{S}^{(i)}}) \right] \quad (9)$$

Our final combined weighted loss function is shown below:

$$L_{total} = 2 \times L_{trans} + L_{rot} + L_{tors} + L_{bp3D} + L_{bp2D} \quad (10)$$

We train our model in PyTorch-Lightning using the Adam optimizer with a learning rate of 0.0001. The distributed training process runs on 4 80-GB NVIDIA A100 GPUs for 1000 epochs, taking approximately 8 hours.

2.5 Training and benchmark datasets

To develop our method RNAbpFlow, we harness a recently published RNA 3D structural dataset from RNA3DB [42], which is both sequentially and structurally non-redundant, making it highly suitable for training and benchmarking deep learning models. We use the version of RNA3DB parsed from the PDB [12] on April 26, 2024, selecting a representative set of RNA sequences from the provided train-test splits for our purposes. To ensure that the dataset contains only high-quality native structures, we apply several filtering steps such as excluding structures with only one atom per nucleotide, removing protein residues from RNA structures, extracting contiguous sequences from experimental structures to address mismatches between the provided FASTA sequences and the experimental 3D structures for preserving base pairing integrity, and excluding sequences lacking any base pairs in their corresponding native structures. Finally, we only retain sequences with a minimum length of 30 and a maximum length of 200 to ensure efficient training. This results in a clean training set consisting of 573 RNA sequences and 52 non-redundant test sequences (excluding component #0) for benchmarking our method development. To compare RNAbpFlow against RNAJP [32], a recent MD simulation-based method for RNA 3D structure sampling based on given 2D structure, we use the dataset of 22 RNAs containing three-way junctions as used in the original RNAJP study. We exclude multimeric structures from the dataset and apply CD-HIT-est-2D between the remaining sequences and our training set

of 573 RNA sequences from RNA3DB to avoid redundancy. This reduces the benchmark set to 12 RNAs, for which, we download the predicted all-atom decoy structures generated by RNAJP from their publicly available repository at <https://rna.physics.missouri.edu/RNAJP/index.html>. To benchmark our method RNAbpFlow on the CASP15 targets, we curate and filter a completely separate training set sourced from trRosettaRNA, adhering to the same filtering criteria used previously. We additionally perform a sequence level non-redundancy check by applying CD-HIT-est-2D [43] using 80% sequence similarity threshold between the naturally occurring RNA sequences from CASP15 and our training sequences, leading to 874 sequences in the training set and four CASP15 test targets (R1107, R1108, R1149, and R1156) without having chain breaks in the native 3D structure to preserve base pairing integrity.

2.6 Performance evaluation and competing methods

To evaluate the performance of our method, RNAbpFlow, compared to other state-of-the-art 3D structure sampling and prediction methods for RNA, we use a wide range of evaluation metrics. To evaluate global fold similarity, we calculate the TM-score (based on C3') using US-align [44] and GDT-TS (based on C4') using the LGA program for RNA [45]. For the assessment of local environment fitness, we compute IDDT [46] using OpenStructure [47] version 2.8 and the clash score metric using MolProbity package [48]. Furthermore, we utilize the CASP-RNA pipeline [49] to (1) evaluate the full atomic structural accuracy by calculating the all-atom RMSD and (2) compute the INF-All score [50], which quantifies how well a structural model reproduces the base interactions of the reference (usually an experimentally resolved native structure) by considering canonical, non-canonical, and stacking interactions.

Since our method RNAbpFlow relies purely on sequence- and base-pair-conditioned SE(3)-equivariant flow matching model for RNA 3D structure generation without utilizing MSAs or template information, we choose two deep learning-based methods that can only take sequence and base pair information for RNA 3D structure prediction for the sake of a fair performance comparison. These include DRfold that integrates end-to-end and geometric potentials and RhoFold+ which leverages language model. Both DRfold and RhoFold+ are installed and run locally with their default parameters, with RhoFold+ run in single-sequence mode. In addition to these deep learning-based approaches, we compare against three physics- and/or knowledge-based methods: RNAComposer [10], 3dRNA [51], and Vfold-Pipeline [52]. RNAComposer and 3dRNA are accessed via their respective web servers, while Vfold-Pipeline is installed and run locally. For these methods, base pairing information is provided in dot-bracket notation format, as predicted by MXfold2 software. To maintain a fair comparison and accurately assess predictive performance, we do not employ any post-prediction optimization or refinement procedures for structure prediction across all methods, except in the cases of Vfold-Pipeline and RNAComposer, where the option to switch off the post-prediction refinement functionality is not available.

3 Results

3.1 Impact of base pairs on RNA 3D generation

To assess the significance of the base pairing information incorporated in our conditional flow matching formulation for RNA 3D structure generation, we individually condition on

base pairs extracted from the native 3D structures using each of the three different 2D annotation tools described in Section 2.2 as well as their combination. We independently train each variant on the RNA3DB training set (detailed in Section 2.5) and evaluate on the 52 targets from the corresponding non-redundant test set. We also train a baseline model by conditioning only on the nucleotide sequence but not on the base pairing information. For each target sequence, we generate 1000 3D structural samples and calculate the maximum and mean of both TM-score and IDDT. The distribution of the average of these scores across all 52 targets is shown in **Figure 2** based on various base pair conditioning (or lack thereof).




Fig. 2. Distributions of the maximum and mean of 1000 3D structural samples per target in terms of (a) TM-score and (b) IDDT across 52 RNA3DB test targets for various base pairing conditioning (or lack thereof). The green triangles indicate averages.

As shown in **Figure 2**, RNAbpFlow achieves the best performance when all three base pair maps are incorporated as condition, resulting in an average over the maximum per-target TM-score distribution of 0.51 and IDDT of 0.71. This represents an average increase of 34.2% in TM-score and 44.9% in IDDT compared to the baseline sequence-conditioned variant of our flow matching formulation, which achieves an average over the maximum per-target TM-score distribution of 0.38 and

IDDT of 0.49. These improvements underscore the critical role of incorporating base pair information in RNA 3D generative modeling. The relative contribution of individual base pair maps is also evident from **Figure 2**, which demonstrates that none of the individual maps can alone outperform their combination. Furthermore, the distributions reveal the consistency of the generated sample qualities when all three maps are used, which achieves a relatively smaller interquartile range while maintaining the highest mean values for both TM-score and IDDT distributions, indicating both high-quality samples and reduced variability. This consistency is also reflected in sampling performance where RNAbpFlow successfully generates at least one 3D structure with a TM-score exceeding 0.45 (a threshold for assessing RNA global fold correctness [49]) for 37 out of 52 targets with 71.2% correctly folded targets from the test set, demonstrating the robustness of RNAbpFlow in RNA 3D structure generation when accurate base pairing information is available.

3.2 Structural ensemble generation performance

To evaluate the sampling performance of RNAbpFlow, we compare against RNAJP, a recent coarse-grained MD simulation-based method for RNA 3D structure sampling with explicit consideration of base pair information, including non-canonical base pairing and base stacking interactions as well as long-range loop-loop interactions. For the benchmark set of 12 RNA targets containing three-way junctions described in Section 2.5, we generate 1000 3D structural samples per target and compute the mean and maximum scores (TM-score and IDDT) for each target, which are then averaged across the 12 RNAs. To ensure a fair comparison, we use the first 1000 decoy structures for each target provided by RNAJP.

Table 1. Sampling results in terms of maximum and mean scores of 1000 3D structural samples for 12 RNAJP test RNAs. Values in bold indicate the best performance.

Method	TM-score		IDDT	
	Max	Mean	Max	Mean
RNAbpFlow	0.5	0.38	0.72	0.67
RNAJP	0.44	0.32	0.65	0.59

The results, summarized in **Table 1**, demonstrate that RNAbpFlow consistently outperforms RNAJP in both metrics. For example, RNAbpFlow achieves an average over the mean IDDT score of 0.67, surpassing that of 0.59 achieved by RNAJP. Similarly, in terms of global topology sampling, RNAbpFlow generates higher-quality structures, achieving an average over the mean TM-score of 0.38 compared to RNAJP's 0.32. The trend suggests that RNAbpFlow produces better structural ensembles by capturing accurate global and local structural features.

Figure 3 showcases two representative examples from the RNAJP benchmark set, illustrating the predictive performance of RNAbpFlow compared to the RNAJP method. For each target, the best structure from the 3D structural samples generated by both methods is selected based on the highest TM-score. For the first case study target RNA 2HGH of 55 nucleotides containing one three-way junction, the best structure predicted by RNAbpFlow shown in **Figure 3(a)** achieved a TM-score of 0.58, indicating a correct global fold and a precise

prediction of the three-way junction, with an all-atom RMSD of 2.58 Å. By contrast, the RNAJP method achieves a lower TM-score of 0.42, partly due to the misorientation of the three-way junction and the hairpin loop. A visual inspection of the predicted base pairs extracted from the 3D structures in **Figure 4(a)** demonstrates that RNAbpFlow more effectively preserves the overall base pairing interactions, particularly the non-canonical ones in the three-way junction region than RNAJP. By capturing the non-canonical base pair interactions with high fidelity, RNAbpFlow achieves a higher INF-NWC score of 0.78, which is much higher than that of RNAJP (0.52).




Fig. 3. Two representative RNA targets with PDB IDs (a) 2HGH and (b) 3PDR, both containing three-way junctions shown with the predicted structural models colored in blue superimposed on the experimental structures in green and the corresponding evaluation metrics displayed on the bottom right side each superposition.

For the second case study target 3PDR, a 160-nucleotide target and the largest in the set containing two three-way junctions, our method outperforms RNAJP by achieving an all-atom RMSD of 5.44 Å and a TM-score of 0.7, primarily due to the better packing of the helices and both of the three-way junction regions which is visually evident. By contrast, the best prediction by RNAJP results in an RMSD of 7.84 Å, primarily due to the deviations in one of the three-way junction regions, as shown in **Figure 3**. The extracted base pairs annotated in **Figure 4(b)**, further highlight the trend that RNAbpFlow more accurately captures non-canonical and long-range loop-loop interactions, with fewer false-positives compared to RNAJP.

3.3 Performance on CASP15 targets

3.3.1 Sampling performance

Figure 5 presents a comparison of the sampling performance of RNAbpFlow on four natural CASP15 RNA targets, conditioned on both native and noisy (predicted) base pairs as input. When provided with native base pairs, RNAbpFlow achieves an impressive average over the maximum per-target TM-score of 0.62 and successfully generates at least one 3D structure with a TM-score greater than 0.45 for 100% of the cases. By contrast, when noisy predicted base pairs are used, the performance declines significantly, with the average over the




Fig. 4. Canonical and non-canonical base pair annotations extracted using RNApdbee 2.0 [53] from the experimental and predicted 3D structures of the two case study targets with PDB IDs (a) 2HGH and (b) 3PDR. Base pair fidelity for the prediction is evaluated and annotated in terms of INF-All (all interactions) and INF-NWC (non-canonical interactions).

per-target maximum TM-score dropping to 0.48. Despite this performance drop, RNAbpFlow still demonstrates its robust sampling capabilities by generating at least one accurate global fold (TM-score > 0.45) for 3 out of 4 targets, showcasing its resilience to noisy 2D information. However, the average sampling performance declines when the input consists entirely of predicted base pair maps, with the average over the mean TM-score per target dropping from 0.42 to 0.32. Such a performance decline highlights the importance of accurate base pairing information in guiding the sample generation process. Additionally, the total number of unique non-canonical base pairs across the three maps for the four targets decreases sharply from 26 to only 3 (with no non-canonical pairs present in 3 out of the 4 targets) when predicted base pairs are used, which could contribute to the observed drop in performance.




Fig. 5. Distributions of the maximum and mean of 1000 3D structural samples per target in terms of TM-score and IDDT across 4 CASP15 test targets for both native and predicted base pair conditioning.

3.3.2 Predictive modeling performance

To directly compare our conditional flow matching-based RNA 3D structure generation method RNAbpFlow against the state-of-the-art sequence- and base-pair-conditioned RNA 3D structure prediction approaches, we leverage our recently published RNA 3D structure scoring method lociPARSE [54], to select the top structure from the RNAbpFlow generated structural ensemble based on the estimated IDDT score (pMoL). The detailed performance comparison is summarized in **Table 2**, showcasing the ability of RNAbpFlow to predict RNA 3D structures across a wide range of evaluation metrics. RNAbpFlow consistently shows better accuracies than the state-of-the-art deep learning-based method DRfold, achieving a well rounded performance across all the metrics and consistently outperforms all physics- and/or knowledge-based methods across almost all metrics except for clash score where it falls short of the physics-based method Vfold-Pipeline, which employs a post-refinement strategy specifically designed to resolve structural clashes. However, the overall predictive performance of RNAbpFlow can be improved, as none of the 3 accurate global folds generated during sampling is selected as the top structure, highlighting the need for further improvement of the scoring function to identify accurate 3D structure in the ensemble.

Table 2. Benchmarking the predictive modeling performance on four natural RNAs from CASP15 with predicted base pairs. Values in bold indicate the best performance.

Predictors	TM-score	IDDT	RMSD	GDT-TS	INF-All	Clash score
RNAbpFlow*	0.34	0.61	13.95	32.90	0.80	73.57
DRfold*	0.31	0.55	17.10	31.48	0.73	149.44
RhoFold+* ‡	0.22	0.37	20.29	22.27	0.46	684.04
Vfold-Pipeline †	0.28	0.47	21.23	26.02	0.65	2.46
RNAComposer †	0.27	0.50	18.51	26.68	0.73	16.22
3dRNA †	0.27	0.50	24.16	24.90	0.69	141.11

* Deep learning methods.

† Physics and/or knowledge guided methods.

‡ Language model based method, run in single sequence mode.

3.3.3 Impact of native base pairing information

Table 3. Comparison of RNAbpFlow with other predictors on CASP15 targets with native base pairing information. Values in bold indicate the best performance.

Predictors	TM-score	IDDT	RMSD	GDT-TS	INF-All	Clash score
RNAbpFlow	0.51	0.72	7.79	49.66	0.89	46.97
DRfold	0.36	0.67	14.47	36.82	0.85	131.68
RNAComposer	0.32	0.59	16.81	31.01	0.84	20.73
3dRNA	0.30	0.54	17.9	28.08	0.72	144.51
VFold	0.28	0.58	19.69	26.44	0.8	0.64

Table 3 shows how the performance of our method RNAbpFlow varies in the presence of accurate native base pairing information compared to all other competing methods. For this benchmarking, all three base pair maps extracted from native 3D structure are provided to DRfold for prediction while the rest of the methods are supplied with DSSR-extracted base pair map since they can accept only one map at a time. The predictive modeling performance of our method RNAbpFlow significantly improves in the presence of accurate base pairs, achieving an average TM-score of 0.51 and average RMSD of 7.79, compared to 0.34 and 13.95 TM-score and RMSD respectively, when

predicted base pairs are used, showing a 50% improvement in TM-score and a 44.1% reduction in RMSD, which is the highest performance gain compared to all other methods. This highlights the versatility of RNAbpFlow in its ability to achieve a higher performance ceiling by effectively leveraging accurate base pairing information as a key condition in deep generative modeling.

3.4 Ablation study and hyperparameter selection




Fig. 6. Distributions of the maximum and mean of 1000 3D structural samples per target in terms of (a) TM-score and (b) IDDT distribution across 52 RNA3DB test targets for five different loss combinations. The green triangles indicate averages.

To evaluate the importance of various loss components and their contributions to sampling performance, along with the architectural hyperparameters used during RNAbpFlow's training, we perform ablation study using the RNA3DB dataset described in Section 2.5. **Figure 6** and **Table 4** report the sampling performance in terms of TM-score and IDDT on the RNA3DB test set, calculated as the average over the per-target maximum and mean scores across 52 RNA3DB test targets for different ablated variants. **Figure 6** shows that RNAbpFlow achieves the best performance when all three loss components are used, with a consistent performance decline when any loss component is removed. For instance, excluding either or both 2D and 3D base pair loss components (i.e., *bp2D* and/or *bp3D*) leads to a noticeable drop in average TM-score

and IDDT, underscoring the importance of incorporating base pairs conditioning in the training. Furthermore, the average IDDT achieved using only the SE(3)-flow matching loss is 0.44, significantly lower than the 0.66 average IDDT obtained when incorporating torsion loss and 2D and 3D base pair losses altogether. This demonstrates the effectiveness of all the auxiliary loss terms used in RNAbpFlow.

Table 4. Hyperparameter selection based on number of structure-module blocks, batch size, number of timesteps used to generate a sample RNA 3D structure and number of samples generated per target on RNA3DB test set. Values in bold indicate the best performance.

# of blocks	TM-score		IDDT		Batch size	TM-score		IDDT	
	Max	Mean	Max	Mean		Max	Mean	Max	Mean
2	0.38	0.28	0.63	0.58	1	0.47	0.36	0.69	0.65
4	0.44	0.34	0.67	0.63	4	0.48	0.38	0.7	0.66
6	0.48	0.38	0.7	0.66	8	0.47	0.37	0.69	0.65
8	0.47	0.38	0.7	0.66	16	0.47	0.37	0.69	0.66

# of timesteps	TM-score		IDDT		# of samples	TM-score		IDDT	
	Max	Mean	Max	Mean		Max	Mean	Max	Mean
5	0.25	0.19	0.36	0.34	10	0.44	0.38	0.68	0.66
10	0.47	0.37	0.7	0.66	100	0.48	0.38	0.7	0.66
50	0.48	0.38	0.7	0.66	500	0.5	0.38	0.71	0.66
100	0.48	0.38	0.7	0.66	1000	0.51	0.38	0.71	0.66

Table 4 highlights the impact of hyperparameter selection on sampling performance. The number of structural module blocks significantly influences TM-score and IDDT, with 6 blocks providing the best balance between model complexity and computational cost. Similarly, a batch size of 4 is sufficient for training efficiency without compromising accuracy. The efficiency of the flow matching formulation is also evident from **Table 4**, as our method RNAbpFlow achieves near-optimal results with just 10 timesteps, with marginal gains observed at 50 timesteps, demonstrating the practical advantage of our approach for large-scale RNA 3D structure sampling with minimal computational overhead. Additionally, increasing the number of generated samples consistently improves performance, with 1000 samples per target yielding the best performance. These findings justify our hyperparameter choices used in this work: 6 structure-module blocks, a batch size of 4, 50 timesteps, and 1000 samples generated per target.

4 Conclusion

In this work, we developed RNAbpFlow, the first sequence- and base-pair-conditioned all-atom RNA 3D structure generation method based on SE(3)-equivariant flow matching model. Experimental results demonstrate that the introduction of base pairing conditioning leads to performance improvements, and the accuracy gain is connected to the quality of the base pairs. Free from the confines of sequence- and structural-level homology, RNAbpFlow enables direct generation of all-atom RNA 3D structural models in an end-to-end manner, thereby opening promising avenues for RNA conformational dynamics through large-scale structural ensemble generation in atomic detail.

5 Acknowledgements

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

- Jayesh A Kulkarni, Dominik Witzigmann, Sarah B Thomson, Sam Chen, Blair R Leavitt, Pieter R Cullis, and Roy van der

- Meel. The current landscape of nucleic acid therapeutics. *Nature nanotechnology*, 16(6):630–643, 2021.
- Cormac Sheridan. First small-molecule drug targeting RNA gains momentum. *Nature biotechnology*, 39(1):6–9, 2021.
- Di Liu, François A Th  lot, Joseph A Piccirilli, Maofu Liao, and Peng Yin. Sub-3-   cryo-EM structure of RNA enabled by engineered homomeric self-assembly. *Nature Methods*, 19(5):576–585, 2022.
- Katherine Deigan Warner, Christine E Hajdin, and Kevin M Weeks. Principles for targeting RNA with drug-like small molecules. *Nature reviews Drug discovery*, 17(8):547–558, 2018.
- Yueyi Li, Anibal Arce, Tyler Lucci, Rebecca A Rasmussen, and Julius B Lucks. Dynamic RNA synthetic biology: new principles, practices and potential. *RNA biology*, 20(1):817–829, 2023.
- Magdalena Rother, Kristian Rother, Tomasz Puton, and Janusz M Bujnicki. ModeRNA: a tool for comparative modeling of RNA 3D structure. *Nucleic acids research*, 39(10):4007–4022, 2011.
- Samuel C Flores, Yaqi Wan, Rick Russell, and Russ B Altman. Predicting RNA structure by multiple template homology modeling. In *Biocomputing 2010*, pages 216–227. World Scientific, 2010.
- Andrew Martin Watkins, Ramya Rangan, and Rhiju Das. FARFAR2: improved de novo rosetta prediction of complex global RNA folds. *Structure*, 28(8):963–976, 2020.
- Yi Zhang, Jun Wang, and Yi Xiao. 3dRNA: Building RNA 3D structure with improved template library. *Computational and structural biotechnology journal*, 18:2416–2423, 2020.
- Mariusz Popena, Marta Szachniuk, Maciej Antczak, Katarzyna J Purzycka, Piotr Lukasiak, Natalia Bartol, Jacek Blazewicz, and Ryszard W Adamiak. Automated 3D structure composition for large RNAs. *Nucleic acids research*, 40(14):e112–e112, 2012.
- Song Cao and Shi-Jie Chen. Physics-based de novo prediction of RNA 3D structures. *The journal of physical chemistry B*, 115(14):4216–4226, 2011.
- Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic Acids Research*, 28(1):235–242, 2000.
- Yang Li, Chengxin Zhang, Chenjie Feng, Robin Pearce, P Lydia Freddolino, and Yang Zhang. Integrating end-to-end learning with deep geometrical potentials for ab initio RNA structure prediction. *Nature Communications*, 14(1):5745, 2023.
- Wenkai Wang, Chenjie Feng, Renmin Han, Ziyi Wang, Lisha Ye, Zongyang Du, Hong Wei, Fa Zhang, Zhenling Peng, and Jianyi Yang. trRosettaRNA: automated prediction of RNA 3D structure with transformer network. *Nature Communications*, 14(1):7266, 2023.
- Zhichao Miao, Ryszard W Adamiak, Maciej Antczak, Micha   J Boniecki, Janusz Bujnicki, Shi-Jie Chen, Clarence Yu Cheng, Yi Cheng, Fang-Chieh Chou, Rhiju Das, et al. RNA-Puzzles Round IV: 3D structure predictions of four ribozymes and two aptamers. *Rna*, 26(8):982–995, 2020.
- Joanna Sarzynska, Mariusz Popena, Maciej Antczak, and Marta Szachniuk. RNA tertiary structure prediction using RNAComposer in CASP15. *Proteins: Structure, Function, and Bioinformatics*, 91(12):1790–1799, 2023.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin   dek, Anna Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.
- Bernard Moussad, Rahmatullah Roche, and Debswapna Bhattacharya. The transformative power of transformers in protein structure prediction. *Proceedings of the National Academy of Sciences*, 120(32):e2303499120, 2023.
- Minkyung Baek, Ryan McHugh, Ivan Anishchenko, Hanlun Jiang, David Baker, and Frank DiMaio. Accurate prediction of protein–nucleic acid complexes using RoseTTAFoldNA. *Nature methods*, 21(1):117–121, 2024.

20. Tao Shen, Zhihang Hu, Siqi Sun, Di Liu, Felix Wong, Jiuming Wang, Jiayang Chen, Yixuan Wang, Liang Hong, Jin Xiao, et al. Accurate RNA 3D structure prediction using a language model-based deep learning approach. *Nature Methods*, pages 1–12, 2024.
21. Yuki Kagaya, Zicong Zhang, Nabil Ibtehaz, Xiao Wang, Tsukasa Nakamura, Pranav Deep Punuru, and Daisuke Kihara. NuFold: end-to-end approach for RNA tertiary structure prediction with flexible nucleobase center representation. *Nature communications*, 16(1):881, 2025.
22. Sumit Tarafder, Rahmatullah Roche, and Debswapna Bhattacharya. The landscape of rna 3d structure modeling with transformer networks. *Biology Methods and Protocols*, 9(1):bpae047, 2024.
23. Yikun Zhang, Mei Lang, JiuHong Jiang, Zhiqiang Gao, Fan Xu, Thomas Litfin, Ke Chen, Jaswinder Singh, Xiansong Huang, Guoli Song, et al. Multiple sequence alignment-based RNA language model and its application to structural inference. *Nucleic Acids Research*, 52(1):e3–e3, 2024.
24. Quentin Vicens and Jeffrey S Kieft. Thoughts on how to think (and talk) about RNA structure. *Proceedings of the National Academy of Sciences*, 119(17):e2112677119, 2022.
25. Andriy Kryshchak, Maciej Antczak, Marta Szachniuk, Tomasz Zok, Rachael C Kretsch, Ramya Rangan, Phillip Pham, Rhiju Das, Xavier Robin, Gabriel Studer, et al. New prediction categories in CASP15. *Proteins: Structure, Function, and Bioinformatics*, 91(12):1550–1557, 2023.
26. Laura R Ganser, Megan L Kelly, Daniel Herschlag, and Hashim M Al-Hashimi. The roles of structural dynamics in the cellular functions of RNAs. *Nature reviews Molecular cell biology*, 20(8):474–489, 2019.
27. Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, pages 1–3, 2024.
28. Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100, 2023.
29. Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. SE (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023.
30. Jason Yim, Andrew Campbell, Andrew YK Foong, Michael Gastegger, José Jiménez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S Veeling, Regina Barzilay, Tommi Jaakkola, et al. Fast protein backbone generation with SE (3) flow matching. *arXiv preprint arXiv:2310.05297*, 2023.
31. Rishabh Anand, Chaitanya K. Joshi, Alex Morehead, Arian Rokkum Jamasb, Charles Harris, Simon V Mathis, Kieran Didi, Bryan Hooi, and Pietro Lio. RNA-frameflow for de novo 3d RNA backbone design. In *ICML 2024 AI for Science Workshop*, 2024.
32. Jun Li and Shi-Jie Chen. Rnarp: enhanced rna 3d structure predictions with non-canonical interactions and global topology sampling. *Nucleic acids research*, 51(7):3341–3356, 2023.
33. Patrick Gendron, Sébastien Lemieux, and François Major. Quantitative analysis of nucleic acid three-dimensional structures. *Journal of molecular biology*, 308(5):919–936, 2001.
34. Huanwang Yang, Fabrice Jossinet, Neocles Leontis, Li Chen, John Westbrook, Helen Berman, and Eric Westhof. Tools for the automatic identification and classification of RNA base pairs. *Nucleic acids research*, 31(13):3450–3460, 2003.
35. Xiang-Jun Lu, Harmen J Bussemaker, and Wilma K Olson. DSSR: an integrated software tool for dissecting the spatial structure of RNA. *Nucleic Acids Research*, 43(21):e142–e142, 2015.
36. Kengo Sato, Manato Akiyama, and Yasubumi Sakakibara. Rna secondary structure prediction using deep learning with thermodynamic integration. *Nature communications*, 12(1):941, 2021.
37. Mirela Andronescu, Anne Condon, Holger H Hoos, David H Mathews, and Kevin P Murphy. Computational approaches for rna energy parameter estimation. *RNA*, 16(12):2304–2318, 2010.
38. Shay Zakov, Yoav Goldberg, Michael Elhadad, and Michal Ziv-Ukelson. Rich parameterization improves rna structure prediction. *Journal of Computational Biology*, 18(11):1525–1542, 2011.
39. Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
40. Ricky TQ Chen and Yaron Lipman. Riemannian flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2023.
41. Dmitry I Nikolayev and Tatjana I Savolov. Normal distribution on the rotation group SO (3). *Texture, Stress, and Microstructure*, 29(3-4):201–233, 1997.
42. Marcell Szikszai, Marcin Magnus, Siddhant Sanghi, Sachin Kadyan, Nazim Bouatta, and Elena Rivas. RNA3DB: A structurally-dissimilar dataset split for training and benchmarking deep learning models for RNA structure prediction. *Journal of Molecular Biology*, page 168552, 2024.
43. Limin Fu, Beifang Niu, Zhengwei Zhu, Sitao Wu, and Weizhong Li. CD-HIT: accelerated for clustering the next-generation sequencing data. *Bioinformatics*, 28(23):3150–3152, 2012.
44. Chengxin Zhang, Morgan Shine, Anna Marie Pyle, and Yang Zhang. US-align: universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nature methods*, 19(9):1109–1115, 2022.
45. Adam Zemla. LGA: a method for finding 3D similarities in protein structures. *Nucleic acids research*, 31(13):3370–3374, 2003.
46. Valerio Mariani, Marco Biasini, Alessandro Barbato, and Torsten Schwede. IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics*, 29(21):2722–2728, 2013.
47. Marco Biasini, Tobias Schmidt, Stefan Bienert, Valerio Mariani, Gabriel Studer, Jürgen Haas, Niklaus Johner, Andreas Daniel Schenk, Ansgar Philippsen, and Torsten Schwede. OpenStructure: an integrated software framework for computational structural biology. *Acta Crystallographica Section D: Biological Crystallography*, 69(5):701–709, 2013.
48. Vincent B Chen, W Bryan Arendall, Jeffrey J Headd, Daniel A Keedy, Robert M Immormino, Gary J Kapral, Laura W Murray, Jane S Richardson, and David C Richardson. Molprobity: all-atom structure validation for macromolecular crystallography. *Acta Crystallographica Section D: Biological Crystallography*, 66(1):12–21, 2010.
49. Rhiju Das, Rachael C Kretsch, Adam J Simpkin, Thomas Mulvaney, Phillip Pham, Ramya Rangan, Fan Bu, Ronan M Keegan, Maya Topf, Daniel J Rigden, et al. Assessment of three-dimensional RNA structure prediction in CASP15. *Proteins: Structure, Function, and Bioinformatics*, 91(12):1747–1770, 2023.
50. Marc Parisien, José Almeida Cruz, Éric Westhof, and François Major. New metrics for comparing and assessing discrepancies between RNA 3D structures and models. *RNA*, 15(10):1875–1885, 2009.
51. Jun Wang, Jian Wang, Yanzhao Huang, and Yi Xiao. 3dRNA v2. 0: an updated web server for RNA 3D structure prediction. *International journal of molecular sciences*, 20(17):4116, 2019.
52. Jun Li, Sicheng Zhang, Dong Zhang, and Shi-Jie Chen. Vfold-Pipeline: a web server for RNA 3D structure prediction from sequences. *Bioinformatics*, 38(16):4042–4043, 2022.
53. Tomasz Zok, Maciej Antczak, Michal Zurkowski, Mariusz Popenda, Jacek Blazewicz, Ryszard W Adamiak, and Marta Szachniuk. Rnapdb 2.0: multifunctional tool for rna structure annotation. *Nucleic acids research*, 46(W1):W30–W35, 2018.
54. Sumit Tarafder and Debswapna Bhattacharya. lociparse: a locality-aware invariant point attention model for scoring rna 3d structures. *Journal of Chemical Information and Modeling*, 64(22):8655–8664, 2024.