gRNAde: <u>G</u>eometric Deep Learning for 3D <u>RNA</u> inverse <u>de</u>sign

Chaitanya K. Joshi¹ Arian R. Jamasb^{1,2} Ramon Viñas³ Charles Harris¹ Simon V. Mathis¹ Alex Morehead⁴ Pietro Liò¹ ¹University of Cambridge, UK ²Prescient Design, Genentech, Roche ³EPFL, Switzerland ⁴University of Missouri, USA

Open source code: github.com/chaitjo/geometric-rna-design

Abstract

Computational RNA design tasks are often posed as inverse problems, where sequences are designed based on adopting a single desired secondary structure without considering 3D geometry and conformational diversity. We introduce **gRNAde**, a **g**eometric **RNA de**sign pipeline operating on 3D RNA backbones to design sequences that explicitly account for structure and dynamics. Under the hood, gRNAde is a multi-state Graph Neural Network that generates candidate RNA sequences conditioned on one or more 3D backbone structures where the identities of the bases are unknown. On a single-state fixed backbone re-design benchmark of 14 RNA structures from the PDB identified by Das et al. [2010], gRNAde obtains higher native sequence recovery rates (56% on average) compared to Rosetta (45% on average), taking under a second to produce designs compared to the reported hours for Rosetta. We further demonstrate the utility of gRNAde on a new benchmark of multi-state design for structurally flexible RNAs, as well as zero-shot ranking of mutational fitness landscapes in a retrospective analysis of a recent RNA polymerase ribozyme structure.



Figure 1: **The gRNAde pipeline for 3D RNA inverse design.** gRNAde is a generative model for RNA sequence design conditioned on backbone 3D structure(s). gRNAde processes one or more RNA backbone graphs (a conformational ensemble) via a multi-state GNN encoder which is equivariant to 3D roto-translation of coordinates as well as conformer order, followed by conformer order-invariant pooling and autoregressive sequence decoding.

1 Introduction

Why RNA design? Historical efforts in computational drug discovery have focussed on designing small molecule or protein-based medicines that either treat symptoms or counter the end stages of disease processes. In recent years, there is a growing interest in designing new RNA-based therapeutics that intervene earlier in disease processes to cut off disease-causing information flow in the cell [Damase et al., 2021, Zhu et al., 2022]. Notable examples of RNA molecules at the forefront of biotechnology today include mRNA vaccines [Metkar et al., 2024] and CRISPR-based genomic medicine [Doudna and Charpentier, 2014]. Of particular interest for structure-based design are ribozymes and riboswitches in the untranslated regions of mRNAs [Mandal and Breaker, 2004, Leppek et al., 2018]. In addition to coding for proteins (such as the spike protein in the Covid vaccine), naturally occurring mRNAs contain riboswitches that are responsible for cell-state dependent protein expression of the mRNA. Riboswitches act by 'switching' their 3D structure from an unbound conformation to a bound one in the presence of specific metabolites or small molecules. Rational design of riboswitches will enable translation to be dependent on the presence or absence of partner molecules, essentially acting as 'on-off' switches for highly targeted mRNA therapies in the future [Felletti et al., 2016, Mustafina et al., 2019, Mohsen et al., 2023].

Challenges of RNA modelling. Despite the promises of RNA therapeutics, proteins have instead been the primary focus in the 3D biomolecular modelling community. Availability of large-scale protein structures from the PDB combined with advances in deep learning for structured data [Bronstein et al., 2021, Duval et al., 2023] have revolutionized protein 3D structure prediction [Jumper et al., 2021] and rational design [Dauparas et al., 2022, Watson et al., 2023]. Applications of deep learning for computational RNA design are underexplored compared to proteins due to paucity of 3D structure without considering 3D geometry [Churkin et al., 2018] and use non-learnt algorithms for aligning 3D RNA fragments [Han et al., 2017, Yesselman et al., 2019], which can be restrictive due to the hand-crafted nature of the heuristics used.

In addition to limited 3D data for training deep learning models, the key technical challenge is that RNA is more dynamic than proteins. The same RNA can adopt multiple distinct conformational states to create and regulate complex biological functions [Ganser et al., 2019, Hoetzel and Suess, 2022, Ken et al., 2023]. Computational RNA design pipelines must account for both the 3D geometric structure and conformational flexibility of RNA to engineer new biological functions.

Our contributions. This paper introduces **gRNAde**, a **g**eometric deep learning-based pipeline for **RNA** inverse **de**sign conditioned on 3D structure, analogous to ProteinMPNN for proteins [Dauparas et al., 2022]. As illustrated in Figure 1, gRNAde generates candidate RNA sequences conditioned on one or more backbone 3D conformations, enabling both single- and multi-state fixed-backbone sequence design. The model is trained on RNA structures from the PDB at 4.0Å or better resolution (12K 3D structures from 4.2K unique RNAs) [Adamczyk et al., 2022], ranging from short RNAs such as riboswitches, aptamers and ribozymes to larger ribosomal RNAs.

We demonstrate the utility of gRNAde for the following design scenarios:

- Improved performance and speed over Rosetta. We compare gRNAde to Rosetta [Leman et al., 2020], the state-of-the-art physically based tool for 3D RNA inverse design, for single-state fixed backbone design of 14 RNA structures of interest from the PDB identified by Das et al. [2010]. We obtain higher native sequence recovery rates with gRNAde (56% on average) compared to Rosetta (45% on average). Additionally, gRNAde is significantly faster than Rosetta for inference; e.g. sampling 100+ designs in 1 second for an RNA of 60 nucleotides on an A100 GPU, compared to the reported hours for Rosetta.
- Enables multi-state and partial RNA design, which were previously not possible with Rosetta. gRNAde with multi-state GNNs moderately improves sequence recovery over an equivalent single-state model on a new benchmark of structurally flexible RNAs, especially for surface nucleotides which undergo positional or secondary structural changes.
- Zero-shot learning of RNA fitness landscape. In a retrospective analysis of mutational fitness landscape data for an RNA polymerase ribozyme [McRae et al., 2024], we show how gRNAde's perplexity, the likelihood of a sequence folding into a backbone structure, can be used to rank mutants based on fitness in a zero-shot/unsupervised manner and outperforms random mutagenesis for improving fitness over the wild type in low throughput scenarios.



Figure 2: **gRNAde featurizes RNA backbone structures as 3D geometric graphs.** Each RNA nucleotide is a node in the graph, consisting of 3 coarse-grained beads for the coordinates for P, C4', N1 (pyrimidines) or N9 (purines) which are used to compute initial geometric features and edges to nearest neighbours in 3D space. Backbone chain figure adapted from Ingraham et al. [2019].

2 The gRNAde pipeline

2.1 The 3D RNA inverse folding problem

Figure 1 illustrates the RNA inverse folding problem: the task of designing new RNA sequences conditioned on a structural backbone. Given the 3D coordinates of a backbone structure, machine learning models must generate sequences that are likely to fold into that shape. The underlying assumption behind inverse folding (and rational biomolecule design) is that structure determines function [Huang et al., 2016]. To the best of our knowledge, gRNAde is the first explicitly multi-state inverse folding pipeline, allowing users to design sequences for backbone conformational ensembles (a set of 3D backbone structures) as opposed to a single structure. Our multi-state design framework aims to better capture RNA conformational dynamics which is often important for functionality in structured RNAs [Ken et al., 2023].

2.2 RNA conformational ensembles as geometric multi-graphs

Featurization. The input to gRNAde is a set of 3D RNA backbone structures (a conformational ensemble) and the corresponding sequence of *n* nucleotides, for instance via PDB files. As shown in Figure 2, gRNAde builds a geometric graph representation for each input structure:

- We start with a 3-bead coarse-grained representation of the RNA backbone, retaining the coordinates for P, C4', N1 (pyrimidine) or N9 (purine) for each nucleotide [Dawson et al., 2016]. This 'pseudotorsional' representation can describe RNA backbone conformations completely in most cases while reducing the size of the torsional space [Wadley et al., 2007].
- 2. Each nucleotide *i* is assigned a node in the geometric graph with the 3D coordinate $\vec{x}_i \in \mathbb{R}^3$ corresponding to the centroid of the 3 bead atoms. Random Gaussian noise with standard deviation 0.1Å is added to coordinates during training to prevent overfitting on crystallisation artifacts, following Dauparas et al. [2022]. Each node is connected by edges to its 32 nearest neighbours as measured by the pairwise distance in 3D space, $\|\vec{x}_i \vec{x}_i\|_2$.
- 3. Nodes are initialized with geometric features analogous to the featurization used in protein inverse folding [Ingraham et al., 2019, Jing et al., 2020]: (a) forward and reverse unit vectors along the backbone backbone from the 5' end to the 3' end, $(\vec{x}_{i+1} \vec{x}_i \text{ and } \vec{x}_i \vec{x}_{i-1})$; and (b) unit vectors, distances, angles, and torsions from each C4' to the corresponding P and N1/N9.
- 4. Edge features for each edge from node j to i are initialized as: (a) the unit vector from the source to destination node, x
 _j − x
 _i; (b) the distance in 3D space, ||x
 _j − x
 _i||₂, encoded by 32 radial basis functions; and (c) the distance along the backbone, j − i, encoded by 32 sinusoidal positional encodings.



Figure 3: **gRNAde model architecture.** One or more RNA backbone geometric graphs are encoded via a series of SE(3)-equivariant Graph Neural Network layers [Jing et al., 2020] to build latent representations of the local 3D geometric neighbourhood of each nucleotide within each state. Representations from multiple states for each nucleotide are then pooled together via permutation invariant Deep Sets [Zaheer et al., 2017], and fed to an autoregressive decoder to predict a probabilities over the four possible bases (A, G, C, U). The probability distribution can be sampled to design a set of candidate sequences. During training, the model is trained end-to-end by minimising a cross-entropy loss between the predicted probability distribution and the true sequence identity.

Multi-graph representation. As described in the previous section, given a set of k structures in the input conformational ensemble, each RNA backbone is featurized as a separate geometric graph $\mathcal{G}^{(k)} = (\mathbf{A}^{(k)}, \mathbf{S}^{(k)}, \vec{\mathbf{V}}^{(k)})$ with the scalar features $\mathbf{S}^{(k)} \in \mathbb{R}^{n \times f}$, vector features $\vec{\mathbf{V}}^{(k)} \in \mathbb{R}^{n \times f' \times 3}$, and $\mathbf{A}^{(k)}$, an $n \times n$ adjacency matrix. For clear presentation and without loss of generality, we omit edge features and use f, f' to denote scalar/vector feature channels.

The input to gRNAde is thus a set of geometric graphs $\{\mathcal{G}^{(1)}, \ldots, \mathcal{G}^{(k)}\}\$ which is merged into what we term a 'multi-graph' representation of the conformational ensemble, $\mathcal{M} = (\mathbf{A}, \mathbf{S}, \mathbf{V})$, by stacking the set of scalar features $\{\mathbf{S}^{(1)}, \ldots, \mathbf{S}^{(k)}\}\$ into one tensor $\mathbf{S} \in \mathbb{R}^{n \times k \times f}$ along a new axis for the set size k. Similarly, the set of vector features $\{\mathbf{V}^{(1)}, \ldots, \mathbf{V}^{(k)}\}\$ is stacked into one tensor $\mathbf{V} \in \mathbb{R}^{n \times k \times f' \times 3}$. Lastly, the set of adjacency matrices $\{\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(k)}\}\$ are merged via a union \cup into one single joint adjacency matrix \mathbf{A} .

2.3 Multi-state GNN for representation learning on conformational ensembles

The gRNAde model, illustrated in Figure 3, processes one or more RNA backbone graphs (a conformational ensemble) via a multi-state GNN encoder which is equivariant to 3D roto-translation of coordinates as well as to the ordering of conformers, followed by conformer order-invariant pooling and sequence decoding. We describe each component in the following sections.

Multi-state GNN encoder. When representing conformational ensembles as a multi-graph, each node feature tensor contains three axes: (no. of nodes, no. of conformations, feature channels). We perform message passing on the multi-graph adjacency to *independently* process each conformer, while maintaining permutation equivariance of the updated feature tensors along both the first (no. of nodes) and second (no. of conformations) axes. This works by operating on only the feature channels axis and generalising the PyTorch Geometric [Fey and Lenssen, 2019] message passing class to account for the extra conformations axis; see Figure 10 and the pseudocode in the Appendix.

We use multiple rotation-equivariant GVP-GNN [Jing et al., 2020] layers to update scalar features $s_i \in \mathbb{R}^{k \times f}$ and vector features $\vec{v}_i \in \mathbb{R}^{k \times f' \times 3}$ for each node *i*:

$$\boldsymbol{m}_{i}, \boldsymbol{\vec{m}}_{i} \coloneqq \sum_{j \in \mathcal{N}_{i}} \mathrm{MSG}\big(\left(\boldsymbol{s}_{i}, \boldsymbol{\vec{v}}_{i}\right), \left(\boldsymbol{s}_{j}, \boldsymbol{\vec{v}}_{j}\right), \boldsymbol{e}_{ij}\big),$$
(1)

$$\boldsymbol{s}_{i}^{\prime}, \boldsymbol{\vec{v}}_{i}^{\prime} \coloneqq \mathrm{UPD}\big(\left(\boldsymbol{s}_{i}, \boldsymbol{\vec{v}}_{i}\right), \left(\boldsymbol{m}_{i}, \boldsymbol{\vec{m}}_{i}\right)\big),$$
(2)

where MSG, UPD are Geometric Vector Perceptrons, a generalization of MLPs to take tuples of scalar and vector features as input and apply O(3)-equivariant non-linear updates. The overall GNN encoder is SO(3)-equivariant due to the use of reflection-sensitive input features (dihedral angles) combined with O(3)-equivariant GVP-GNN layers.

Our multi-state GNN encoder is easy to implement in any message passing framework and can be used as a *plug-and-play* extension for any geometric GNN pipeline to incorporate the multi-state inductive bias. It serves as an elegant alternative to batching all the conformations, which we found required major alterations to message passing and pooling depending on downstream tasks.

Conformation order-invariant pooling. The final encoder representations in gRNAde account for multi-state information while being invariant to the permutation of the conformational ensemble. To achieve this, we perform a Deep Set pooling [Zaheer et al., 2017] over the conformations axis after the final encoder layer to reduce $S \in \mathbb{R}^{n \times k \times f}$ and $\vec{V} \in \mathbb{R}^{n \times k \times f' \times 3}$ to $S' \in \mathbb{R}^{n \times f}$ and $\vec{V}' \in \mathbb{R}^{n \times f' \times 3}$:

$$\boldsymbol{S}', \vec{\boldsymbol{V}}' := \frac{1}{k} \sum_{i=1}^{k} \left(\boldsymbol{S}[:, i], \vec{\boldsymbol{V}}[:, i] \right).$$
(3)

A simple sum or average pooling does not introduce any new learnable parameters to the pipeline and is flexible to handle a variable number of conformations, enabling both single-state and multi-state design. It may be interesting to ablate the impact of more expressive geometric GNN layers [Joshi et al., 2023] as well as set pooling functions [Maron et al., 2020] in future iterations of gRNAde.

Sequence decoding and loss function. We feed the final encoder representations after pooling, S', \vec{V}' , to autoregressive GVP-GNN decoder layers to predict the probability of the four possible base identities (A, G, C, U) for each node/nucleotide. Decoding proceeds according to the RNA sequence order from the 5' end to 3' end. gRNAde is trained in a self-supervised manner by minimising a cross-entropy loss (with label smoothing value of 0.05) between the predicted probability distribution and the ground truth identity for each base. During training, we use autoregressive teacher forcing [Williams and Zipser, 1989] where the ground truth base identity is fed as input to the decoder at each step, encouraging the model to stay close to the ground-truth sequence.

Sampling. When using gRNAde for inference and designing new sequences, we iteratively sample the base identity for a given nucleotide from the predicted conditional probability distribution, given the partially designed sequence up till that nucleotide/decoding step. We can modulate the smoothness or sharpness of the probability distribution by using a temperature parameter. At lower temperatures, for instance ≤ 1.0 , we expect higher native sequence recovery and lower diversity in gRNAde's designs. At higher temperatures, the model produces more diverse designs by sampling from a smoothed probability distribution. We can also consider unordered decoding [Dauparas et al., 2022] and masking or logit biasing during sampling, depending on the design scenario at hand. This enables gRNAde to perform partial re-design of RNA sequences, retaining specified nucleotide identities while designing the rest of the sequence. Similar approaches for functional protein design have been shown to be successful in the wet lab [Sumida et al., 2024].

2.4 Evaluation metrics for designed sequences

In principle, inverse folding models can be sampled from to obtain a large number of designed sequences for a given backbone structure. Thus, in-silico metrics to determine which sequences are useful and which ones to prioritise in wet lab experiments are a critical part of the overall pipeline. We currently use the following metrics to evaluate gRNAde's designs, visualised in Appendix Figure 8:

- **Native sequence recovery**, which is the average percentage of native (ground truth) nucleotides correctly recovered in the sampled sequences. Recovery is the most widely used metric for biomolecule inverse design [Dauparas et al., 2022].
- Secondary structure self-consistency score, where we 'forward fold' the sampled sequences using a secondary structure prediction tool (we used EternaFold [Wayment-Steele et al., 2022]) and measure the average Mathew's Correlation Coefficient (MCC) to the ground truth secondary structure, represented as a binary adjacency matrix. MCC values range between -1 and +1, with +1 represents a perfect prediction, 0 an average random prediction and -1 an inverse prediction. Similar forward folding metrics, sometimes termed 'designability', have been found to correlate with experimental success in protein design [Watson et al., 2023]¹.

¹It would be ideal to consider a 3D structure self-consistency score, instead. While RNA 3D structure prediction is currently not at the same level of reliability as that for proteins, our framework is flexible enough to include new self-consistency scores based on improved RNA structure prediction tools in the future.

• **Perplexity**, which can be thought of as the average number of bases that the model is selecting from for each nucleotide. Formally, perplexity is the average exponential of the negative log-likelihood of the sampled sequences. A perfect model would have perplexity of 1, while a perplexity of 4 means that the model is making random predictions (the model outputs a uniform probability over 4 possible bases). Perplexity does not require a ground truth structure to calculate, and can also be used for ranking sequences as it is the model's estimate of the compatibility of a sequence with the input backbone structure.

3 Experimental Setup

3D RNA structure dataset. We create a machine learning-ready dataset for RNA inverse design using RNASolo [Adamczyk et al., 2022], a novel repository of RNA 3D structures extracted from solo RNAs, protein-RNA complexes, and DNA-RNA hybrids in the PDB. We used structures at resolution \leq 4.0Å resulting in 4,223 unique RNA sequences for which a total of 12,011 structures are available (downloaded from RNASolo on 31 October 2023). Dataset statistics are available in Figure 9 in the Appendix, illustrating the diversity of our dataset in terms of sequence length, number of structures per sequence, as well as structural variations among conformations per sequence.

Structural clustering. In order to ensure that we evaluate gRNAde's generalization ability to novel RNAs, we cluster the 4,223 unique RNAs into groups based on structural similarity. We use qTMclust for efficiently applying US-align [Zhang et al., 2022] with a similarity threshold of TM-score > 0.45, and ensure that we train, validate and test gRNAde on structurally dissimilar clusters (see next paragraph). We also provide utilities for clustering based on sequence homology using CD-HIT [Fu et al., 2012], which leads to splits containing biologically dissimilar clusters of RNAs.

Splits to evaluate generalization. After clustering, we split the RNAs into training (\sim 4000 samples), validation and test sets (100 samples each) to evaluate two different design scenarios:

- 1. **Single-state split.** This split is used to fairly evaluate gRNAde for single-state design on a set of RNA structures of interest from the PDB identified by Das et al. [2010], which mainly includes riboswitches, aptamers, and ribozymes (full listing in Table 1). We identify the structural clusters belonging to the RNAs identified in Das et al. [2010] and add all the RNAs in these clusters to the test set (100 samples). The remaining clusters are randomly added to the training and validation splits.
- 2. **Multi-state split.** This split is used to tests gRNAde's ability to design RNA with multiple distinct conformational states. We order the structural clusters based on median intra-sequence RMSD among available structures within the cluster². The top 100 samples from clusters with the highest median intra-sequence RMSD are added to the test set. The next 100 samples are added to the validation set and all remaining samples are used for training.

Validation and test samples come from clusters with at most 5 unique sequences, in order to ensure diversity. Any samples that were not assigned clusters are directly appended to the training set. We also directly add very large RNAs (> 1000 nts) to the training set, as it is unlikely that we want to design very large RNAs. We exclude very short RNA strands (< 10 nts).

Evaluation metrics. For a given data split, we evaluate models on the held-out test set by designing 16 sequences (sampled at temperature 0.1) for each test data point and computing averages for each of the three metrics described in Section 2.4: native sequence recovery, secondary structure MCC self-consistency score, and perplexity. We employ early stopping by reporting test set performance for the model checkpoint for the epoch with the best validation set recovery. Standard deviations are reported across 3 consistent random seeds for all models.

Hyperparameters. All models use 4 encoder and 4 decoder GVP-GNN layers, with 128 scalar/16 vector node features, 64 scalar/4 vector edge features, and drop out probability 0.5, resulting in 2,147,944 trainable parameters. All models are trained for a maximum of 50 epochs using the Adam optimiser with an initial learning rate of 0.0001, which is reduced by a factor 0.9 when validation performance plateaus with patience of 5 epochs. Ablation studies of key modelling decisions are available in Appendix Table 2.

²For each RNA sequence, we compute the pairwise C4' RMSD among all available structures. We then compute the median RMSD across all sequences within each structural cluster.



Figure 4: **gRNAde compared to Rosetta for single-state design.** (a) We benchmark native sequence recovery of gRNAde, Rosetta, and FARNA on 14 RNA structures of interest identified by Das et al. [2010]. gRNAde obtains higher native sequence recovery rates (56% on average) compared to Rosetta (45%) and FARNA (32%). (b) Sequence recovery per sample for Rosetta and gRNAde, shaded by gRNAde's perplexity for each sample. gRNAde's perplexity is correlated with native sequence recovery for designed sequences. Full results available in Appendix Table 1.

4 Results

4.1 Single-state RNA design benchmark

We set out to compare gRNAde to Rosetta, a state-of-the-art physically based toolkit for biomolecular modelling and design [Leman et al., 2020]. We reproduced the benchmark setup from Das et al. [2010] for Rosetta's fixed backbone RNA sequence design workflow on 14 RNA structures of interest from the PDB, which mainly includes riboswitches, aptamers, and ribozymes (full listing in Table 1). We trained gRNAde on the single-state split detailed in Section 3, explicitly excluding the 14 RNAs as well as any structurally similar RNAs in order to ensure that we fairly evaluate gRNAde's generalization abilities vs. Rosetta.

gRNAde improves sequence recovery over Rosetta. In Figure 4, we compare gRNAde's native sequence recovery for single-state design with numbers taken from Das et al. [2010] for Rosetta and FARNA (a predecessor of Rosetta). gRNAde has higher recovery of 56% on average compared to 45% for Rosetta and 32% for FARNA. See Appendix Table 1 for full results.

gRNAde is significantly faster than Rosetta. In addition to superior sequence recovery, gRNAde is significantly faster than Rosetta for high-throughout design pipelines. Training gRNAde from scratch takes roughly 2–6 hours on a single A100 GPU, depending on the exact hyperparameters. Once trained, gRNAde can design hundreds of sequences for backbones with hundreds of nucleotides in ~1 second with GPU acceleration. On the other hand, Rosetta takes order of hours to produce a single design due performing expensive Monte Carlo optimisations³. Deep learning methods like gRNAde are arguably easier to use since no expert customization is required and setup is easier compared to Rosetta [Dauparas et al., 2022], potentially making RNA design more broadly accessible.

gRNAde's perplexity correlates with recovery. In Figure 4b, we plot native sequence recovery per sample for Rosetta vs. gRNAde, shaded by gRNAde's average perplexity for each sample. Perplexity is an indicator of the model's confidence in its own prediction (lower perplexity implies higher confidence) and appears to be correlated with native sequence recovery. In Section 4.3, we further demonstrate the utility of gRNAde's perplexity for unsupervised learning of RNA fitness landscapes.

4.2 Multi-state RNA design benchmark

Structured RNAs often adopt multiple distinct conformational states to perform biological functions [Ken et al., 2023]. For instance, riboswitches adopt at least two distinct functional conformation: a

³While we have not run Rosetta ourselves, we note that its documentation states that "runs on RNA backbones longer than ten nucleotides take many minutes or hours".



Figure 5: **Multi-state design benchmark.** Multi-state gRNAde show marginal improvement over an equivalent single-state model in terms of average native sequence recovery and secondary structure self-consistency score. We plot performance for one consistent random seed across all models; collated results are available in Appendix Table 2.

ligand bound (holo) and unbound (apo) state, which helps them regulate and control gene expression [Stagno et al., 2017]. If we were to attempt single-state inverse design for such RNAs, each backbone structure may lead to a different set of sampled sequences. It is not obvious how to select the input backbone as well as designed sequence when using single-state models for multi-state design⁴. gRNAde's multi-state GNN, descibed in Section 2.3, directly 'bakes in' the multi-state nature of RNA as an architectural inductive bias and designs sequences explicitly conditioned on multiple states.

In order to evaluate gRNAde's multi-state design capabilities, we trained equivalent single-state and multi-state gRNAde models on the multi-state split detailed in Section 3, where the validation and test sets contain progressively more structurally flexible RNAs as measured by median RMSD among multiple available states for an RNA.

Multi-state gRNAde boosts sequence recovery. In Figure 5, we compared a single-state variant of gRNAde with otherwise equivalent multi-state models (with up to 3 and 5 states, respectively) in terms of native sequence recovery and secondary structure self-consistency score⁵. Multi-state variants show marginal improvements for both metrics. As a caveat, it is worth noting that multi-state models consume more GPU memory than an equivalent single-state model during mini-batch training (approximate peak GPU usage for max. number of states set to 1: 12GB, 3: 28GB, 5: 50GB on a single A100 with otherwise equivalent hyperparameters).

Improved recovery in structurally flexible regions. In Figure 6, we evaluated gRNAde's multi-state sequence recovery at a fine-grained, per-nucleotide level. gRNAde with multi-state GNN architectures improves sequence recovery over its single-state variant on structurally flexible nucleotides, as characterised by undergoing changes in base pairing/secondary structure, higher average RMSD between 3D coordinates across states, and larger solvent accessible surface area.

4.3 Zero-shot ranking of RNA fitness landscape

Lastly, we explored the use of gRNAde as an zero-shot ranker of mutants in RNA engineering campaigns. Given the backbone structure of a wild type RNA of interest as well as a candidate set of mutant sequences, we can compute gRNAde's perplexity of whether a given sequence folds into the

⁴ProteinMPNN [Dauparas et al., 2022] proposes to average logits from multiple backbones for multi-state protein design. Here is a simple example to highlight issues with such an approach: Consider two states A and B, and choice of labels X, Y, and Z. For state A: X, Y, Z are assigned probabilities 75%, 20%, 5%. For state B: X, Y, Z are assigned probabilities 5%, 20%, 75%. Logically, label Y is the only one that is compatible with both states. However, averaging the probabilities would lead to label X or Z being more likely to be sampled in designs.

⁵Self-consistency scores for multi-state design are perhaps less reliable than single-state design due to the lack of multi-state structure prediction tools. At present, we compute the average MCC score between the multiple ground truth secondary structures and the predicted structure for designed sequences.



Figure 6: **Per-nucleotide sequence recovery for multi-state design.** Multi-state gRNAde shows improved sequence recovery over a single-state model for structurally flexible regions of RNAs, as characterised by nucleotides that tend to undergo changes in base pairing (left), nucleotides with greater average solvent accessible surface area (centre), and nucleotides with higher average RMSD (right) across multiple states. Marginal histograms in blue show the distribution of values.

backbone structure. Perplexity is inversely related to the likelihood of a sequence conditioned on a structure, as described in Section 2.4. We can then rank sequences based on how 'compatible' they are with the backbone structure in order to select a subset to be experimentally validated in wet labs.

Retrospective analysis on ribozyme fitness landscape. A recent study by McRae et al. [2024] determined a cryo-EM structure of an RNA polymerase ribozyme at 5Å resolution⁶, along with a fitness landscape of \sim 75,000 mutant sequences for the catalytic subunit 5TU. We design a retrospective study using this data of (sequence, fitness value) pairs where we simulate an RNA engineering campaign with the aim of improving fitness over the wild type sequence.

We consider various design budgets ranging from hundreds to thousands of sequences selected for experimental validation, and compare 4 unsupervised approaches for ranking/selecting variants: (1) random choice from all \sim 75,000 sequences; (2) random choice from all 449 single mutant sequences; (3) random choice from all single and double mutant sequences (as sequences with higher mutation order tend to be less fit); and (4) negative gRNAde perplexity (lower perplexity is better). For each design budget and ranking approach, we compute the expected maximum change in fitness over the wild type that could be achieved by screening as many variants as allowed in the given design budget. We run 10,000 simulations to compute confidence intervals for the 3 random baselines.

gRNAde outperforms random baselines in low design budget scenarios. Figure 7 illustrates the results of our retrospective study. At low design budgets of up to hundreds of sequences, which are relevant in the case of a low throughput fitness screening assay, gRNAde outperforms all random baselines in terms of the maximum change in fitness over the wild type. The top 10 mutants as ranked by gRNAde contains a sequence with 4-fold improved fitness, while the top 200 leads to a 5-fold improvement⁷. Beyond design budgets of thousands, random selection from double mutants starts outperforming gRNAde.

Perspective. Overall, it is promising that gRNAde's perplexity is somewhat correlated with experimental fitness measurements out-of-the-box (zero-shot) and can be a useful ranker of mutant fitness in our retrospective study. In realistic design scenarios, improvements could likely be obtained by fine-tuning gRNAde on a low amount of experimental fitness data. For example, latent features from gRNAde may be finetuned or used as input to a prediction head with supervised learning on fitness landscape data. This study acts as a sanity check before committing to wet lab validation of gRNAde designs. We see random mutagenesis and directed evolution-based approaches as complementary to de-novo design and inverse folding approaches like gRNAde. Random mutagenesis can be thought of as local exploration around a wild type sequence, optimising fitness within an 'island' of activity. Structure-based design approaches are akin to global jumps in sequence space, with the potential to find new islands further away from the wild type [Huang et al., 2016].

⁶This RNA was not present in gRNAde's training data, which contains structures at \leq 4.0Å resolution.

⁷As a caveat, the fitness assays from McRae et al. [2024] used for creating the landscape have inherent noise and cannot easily differentiate between mutants of similar activity.



Figure 7: **Retrospective study of gRNAde for ranking ribozyme mutant fitness.** Using the backbone structure and mutational fitness landscape data from an RNA polymerase ribozyme [McRae et al., 2024], we retrospectively analyse how well we can rank variants at multiple design budgets using random selection vs. gRNAde's perplexity for mutant sequences conditioned on the backbone structure. Note that gRNAde is used zero-shot here, i.e. it was not fine-tuned on any assay data. For stochastic strategies, bars indicate median values, and error bars indicate the interquartile range estimated from 10,000 simulations per strategy and design budget. At low throughput design budgets of up to \sim 500 sequences, selecting mutants using gRNAde outperforms random baselines in terms of the expected maximum improvement in fitness over the wild type. In particular, gRNAde performs better than single site saturation mutagenesis, even when all single mutants are explored (449 mutants for the RNA polymerase ribozyme in McRae et al. [2024]).

5 Conclusion

We introduce gRNAde, a geometric deep learning pipeline for RNA sequence design conditioned on one or more 3D backbone structures. gRNAde is superior to the physically based Rosetta for 3D RNA inverse folding in terms of performance, inference speed, and ease of use. Additionally, gRNAde enables explicit multi-state design for structurally flexible RNAs which was previously not possible with Rosetta. To the best of our knowledge, gRNAde is also the first geometric deep learning architecture for multi-state biomolecule representation learning; the model is generic and can be repurposed for other learning tasks on conformational ensembles, including multi-state protein design. Key limitations of gRNAde and avenues for future development include the lack of support for multiple interacting RNA chains, or accounting for biomolecular interactions of RNAs with proteins, small molecules, and other ligands. More advanced user configurations, such as negative design against undesired conformations or modulating conformational propensities are also not currently possible. Finally, we are hopeful that advances in RNA structure determination and computationally assisted cryo-EM [Kappel et al., 2020, Bonilla and Kieft, 2022] will further increase the amount of RNA structures available for training geometric deep learning models in the future.

Acknowledgements

We would like to thank Roger Foo, Phillip Holliger, Alex Borodavka, Rhiju Das, Janusz Bujnicki, Edoardo Gianni, Ben Porebski, Michael Mohsen, Christian Choe, and John Boom for helpful comments and discussions. CKJ was supported by the A*STAR Singapore National Science Scholarship (PhD). SVM was supported by the UKRI Centre for Doctoral Training in Application of Artificial Intelligence to the study of Environmental Risks (EP/S022961/1). AM was supported by a U.S. NSF grant (DBI2308699) and two U.S. NIH grants (R01GM093123 and R01GM146340). This research was partially supported by Google TPU Research Cloud and Cambridge Dawn Supercomputer Pioneer Project compute grants.

References

- B. Adamczyk, M. Antczak, and M. Szachniuk. Rnasolo: a repository of cleaned pdb-derived rna 3d structures. *Bioinformatics*, 2022. (Cited on page 2, 6)
- M. Baek, F. DiMaio, I. Anishchenko, J. Dauparas, S. Ovchinnikov, G. R. Lee, J. Wang, Q. Cong, L. N. Kinch, R. D. Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 2021. (Cited on page 14)
- M. Baek, R. McHugh, I. Anishchenko, H. Jiang, D. Baker, and F. DiMaio. Accurate prediction of protein–nucleic acid complexes using rosettafoldna. *Nature Methods*, 2024. (Cited on page 14)
- S. L. Bonilla and J. S. Kieft. The promise of cryo-em to explore rna structural dynamics. *Journal of Molecular Biology*, 2022. (Cited on page 10)
- E. Bonnet, P. Rzazewski, and F. Sikora. Designing rna secondary structures is hard. *Journal of Computational Biology*, 2020. (Cited on page 14)
- M. M. Bronstein, J. Bruna, T. Cohen, and P. Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint*, 2021. (Cited on page 2)
- J. Chen, Z. Hu, S. Sun, Q. Tan, Y. Wang, Q. Yu, L. Zong, L. Hong, J. Xiao, T. Shen, et al. Interpretable rna foundation model from unannotated data for highly accurate rna structure and function predictions. *arXiv preprint*, 2022. (Cited on page 14)
- A. Churkin, M. D. Retwitzer, V. Reinharz, Y. Ponty, J. Waldispühl, and D. Barash. Design of rnas: comparing programs for inverse rna folding. *Briefings in bioinformatics*, 2018. (Cited on page 2, 14)
- T. R. Damase, R. Sukhovershin, C. Boada, F. Taraballi, R. I. Pettigrew, and J. P. Cooke. The limitless future of rna therapeutics. *Frontiers in bioengineering and biotechnology*, 2021. (Cited on page 2)
- R. Das, J. Karanicolas, and D. Baker. Atomic accuracy in predicting and designing noncanonical rna structure. *Nature methods*, 2010. (Cited on page 1, 2, 6, 7, 14, 15)
- J. Dauparas, I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. Wicky, et al. Robust deep learning based protein sequence design using proteinmpnn. *Science*, 2022. (Cited on page 2, 3, 5, 7, 8)
- W. K. Dawson, M. Maciejczyk, E. J. Jankowska, and J. M. Bujnicki. Coarse-grained modeling of rna 3d structure. *Methods*, 2016. (Cited on page 3)
- K. Didi, F. Vargas, S. Mathis, V. Dutordoir, E. Mathieu, U. J. Komorowska, and P. Lio. A framework for conditional diffusion modelling with applications in motif scaffolding for protein design. In *NeurIPS 2023 Machine Learning for Structural Biology Workshop*, 2023. (Cited on page 14)
- J. A. Doudna and E. Charpentier. The new frontier of genome engineering with crispr-cas9. *Science*, 2014. (Cited on page 2)
- A. Duval, S. V. Mathis, C. K. Joshi, V. Schmidt, S. Miret, F. D. Malliaros, T. Cohen, P. Lio, Y. Bengio, and M. Bronstein. A hitchhiker's guide to geometric gnns for 3d atomic systems. *arXiv preprint*, 2023. (Cited on page 2)
- M. Felletti, J. Stifel, L. A. Wurmthaler, S. Geiger, and J. S. Hartig. Twister ribozymes as highly versatile expression platforms for artificial riboswitches. *Nature communications*, 2016. (Cited on page 2)
- M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. *ICLR 2019 Representation Learning on Graphs and Manifolds Workshop*, 2019. (Cited on page 4)
- L. Fu, B. Niu, Z. Zhu, S. Wu, and W. Li. Cd-hit: accelerated for clustering the next-generation sequencing data. *Bioinformatics*, 2012. (Cited on page 6)
- L. R. Ganser, M. L. Kelly, D. Herschlag, and H. M. Al-Hashimi. The roles of structural dynamics in the cellular functions of rnas. *Nature reviews Molecular cell biology*, 2019. (Cited on page 2)

- D. Han, X. Qi, C. Myhrvold, B. Wang, M. Dai, S. Jiang, M. Bates, Y. Liu, B. An, F. Zhang, et al. Single-stranded dna and rna origami. *Science*, 2017. (Cited on page 2, 14)
- S. He, R. Huang, J. Townley, R. C. Kretsch, T. G. Karagianes, D. B. Cox, H. Blair, D. Penzar, V. Vyaltsev, E. Aristova, et al. Ribonanza: deep learning of rna structure through dual crowdsourcing. *bioRxiv*, 2024. (Cited on page 14)
- J. Hoetzel and B. Suess. Structural changes in aptamers are essential for synthetic riboswitch engineering. *Journal of Molecular Biology*, 2022. (Cited on page 2)
- P.-S. Huang, S. E. Boyken, and D. Baker. The coming of age of de novo protein design. *Nature*, 2016. (Cited on page 3, 9)
- J. Ingraham, V. Garg, R. Barzilay, and T. Jaakkola. Generative models for graph-based protein design. *NeurIPS*, 2019. (Cited on page 3)
- J. B. Ingraham, M. Baranov, Z. Costello, K. W. Barber, W. Wang, A. Ismail, V. Frappier, D. M. Lord, C. Ng-Thow-Hing, E. R. Van Vlack, et al. Illuminating protein space with a programmable generative model. *Nature*, 2023. (Cited on page 14)
- B. Jing, S. Eismann, P. Suriana, R. J. L. Townshend, and R. Dror. Learning from protein structure with geometric vector perceptrons. In *International Conference on Learning Representations*, 2020. (Cited on page 3, 4)
- C. K. Joshi, C. Bodnar, S. V. Mathis, T. Cohen, and P. Lio. On the expressive power of geometric graph neural networks. In *International Conference on Machine Learning*, 2023. (Cited on page 5)
- J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Zidek, A. Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 2021. (Cited on page 2, 14)
- K. Kappel, K. Zhang, Z. Su, A. M. Watkins, W. Kladwang, S. Li, G. Pintilie, V. V. Topkar, R. Rangan, I. N. Zheludev, et al. Accelerated cryo-em-guided determination of three-dimensional rna-only structures. *Nature methods*, 2020. (Cited on page 10)
- M. L. Ken, R. Roy, A. Geng, L. R. Ganser, A. Manghrani, B. R. Cullen, U. Schulze-Gahmen, D. Herschlag, and H. M. Al-Hashimi. Rna conformational propensities determine cellular activity. *Nature*, 2023. (Cited on page 2, 3, 7)
- J. K. Leman, B. D. Weitzner, S. M. Lewis, J. Adolf-Bryfogle, N. Alam, R. F. Alford, M. Aprahamian, D. Baker, K. A. Barlow, P. Barth, et al. Macromolecular modeling and design in rosetta: recent methods and frameworks. *Nature methods*, 2020. (Cited on page 2, 7)
- K. Leppek, R. Das, and M. Barna. Functional 5' utr mrna structures in eukaryotic translation regulation and how to find them. *Nature reviews Molecular cell biology*, 2018. (Cited on page 2)
- S. Li, S. Moayedpour, R. Li, M. Bailey, S. Riahi, L. Kogler-Anele, M. Miladi, J. Miner, D. Zheng, J. Wang, et al. Codonbert: Large language models for mrna design and optimization. *bioRxiv*, 2023a. (Cited on page 14)
- Y. Li, C. Zhang, C. Feng, R. Pearce, P. Lydia Freddolino, and Y. Zhang. Integrating end-toend learning with deep geometrical potentials for ab initio rna structure prediction. *Nature Communications*, 2023b. (Cited on page 14)
- M. Mandal and R. R. Breaker. Gene regulation by riboswitches. *Nature reviews Molecular cell biology*, 2004. (Cited on page 2)
- H. Maron, O. Litany, G. Chechik, and E. Fetaya. On learning sets of symmetric elements. In *International conference on machine learning*, 2020. (Cited on page 5)
- E. K. McRae, C. J. Wan, E. L. Kristoffersen, K. Hansen, E. Gianni, I. Gallego, J. F. Curran, J. Attwater, P. Holliger, and E. S. Andersen. Cryo-em structure and functional landscape of an rna polymerase ribozyme. *Proceedings of the National Academy of Sciences*, 2024. (Cited on page 2, 9, 10)

- M. Metkar, C. S. Pepin, and M. J. Moore. Tailor made: the art of therapeutic mrna design. *Nature Reviews Drug Discovery*, 2024. (Cited on page 2)
- M. G. Mohsen, M. K. Midy, A. Balaji, and R. R. Breaker. Exploiting natural riboswitches for aptamer engineering and validation. *Nucleic Acids Research*, 2023. (Cited on page 2)
- K. Mustafina, K. Fukunaga, and Y. Yokobayashi. Design of mammalian on-riboswitches based on tandemly fused aptamer and ribozyme. *ACS Synthetic Biology*, 2019. (Cited on page 2)
- R. J. Penic, T. Vlasic, R. G. Huber, Y. Wan, and M. Sikic. Rinalmo: General-purpose rna language models can generalize well on structure prediction tasks. *arXiv preprint*, 2024. (Cited on page 14)
- F. Runge, D. Stoll, S. Falkner, and F. Hutter. Learning to design RNA. In *ICLR*, 2019. (Cited on page 14)
- B. Schneider, B. A. Sweeney, A. Bateman, J. Cerny, T. Zok, and M. Szachniuk. When will rna get its alphafold moment? *Nucleic Acids Research*, 2023. (Cited on page 2)
- J. Stagno, Y. Liu, Y. Bhandari, C. Conrad, S. Panja, M. Swain, L. Fan, G. Nelson, C. Li, D. Wendel, et al. Structures of riboswitch rna reaction states by mix-and-inject xfel serial crystallography. *Nature*, 2017. (Cited on page 8)
- K. H. Sumida, R. Nunez-Franco, I. Kalvet, S. J. Pellock, B. I. Wicky, L. F. Milles, J. Dauparas, J. Wang, Y. Kipnis, N. Jameson, et al. Improving protein expression, stability, and function with proteinmpnn. *Journal of the American Chemical Society*, 2024. (Cited on page 5)
- R. J. Townshend, S. Eismann, A. M. Watkins, R. Rangan, M. Karelina, R. Das, and R. O. Dror. Geometric deep learning of rna structure. *Science*, 2021. (Cited on page 14)
- L. M. Wadley, K. S. Keating, C. M. Duarte, and A. M. Pyle. Evaluating and learning from rna pseudotorsional space: quantitative validation of a reduced representation for rna structure. *Journal of molecular biology*, 2007. (Cited on page 3)
- W. Wang, C. Feng, R. Han, Z. Wang, L. Ye, Z. Du, H. Wei, F. Zhang, Z. Peng, and J. Yang. trrosettarna: automated prediction of rna 3d structure with transformer network. *Nature Communications*, 2023. (Cited on page 14)
- M. Ward, E. Courtney, and E. Rivas. Fitness functions for rna structure design. *Nucleic Acids Research*, 2023. (Cited on page 14)
- A. M. Watkins, R. Rangan, and R. Das. Farfar2: improved de novo rosetta prediction of complex global rna folds. *Structure*, 2020. (Cited on page 14)
- J. L. Watson, D. Juergens, N. R. Bennett, B. L. Trippe, J. Yim, H. E. Eisenach, W. Ahern, A. J. Borst, R. J. Ragotte, L. F. Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 2023. (Cited on page 2, 5, 14)
- H. K. Wayment-Steele, W. Kladwang, A. I. Strom, J. Lee, A. Treuille, A. Becka, E. Participants, and R. Das. Rna secondary structure packages evaluated and improved by high-throughput experiments. *Nature methods*, 2022. (Cited on page 5)
- R. J. Williams and D. Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1989. (Cited on page 5)
- J. D. Yesselman, D. Eiler, E. D. Carlson, M. R. Gotrik, A. E. d'Aquino, A. N. Ooms, W. Kladwang, P. D. Carlson, X. Shi, D. A. Costantino, et al. Computational design of three-dimensional rna structure and function. *Nature nanotechnology*, 2019. (Cited on page 2, 14)
- M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. R. Salakhutdinov, and A. J. Smola. Deep sets. *NeurIPS*, 2017. (Cited on page 4, 5)
- C. Zhang, M. Shine, A. M. Pyle, and Y. Zhang. Us-align: universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nature methods*, 2022. (Cited on page 6)
- Y. Zhu, L. Zhu, X. Wang, and H. Jin. Rna-based therapeutics: An overview and prospectus. *Cell Death & Disease*, 2022. (Cited on page 2)

A Related Work

We attempt to briefly summarise recent developments in RNA structure modelling and design, with an emphasis on deep learning-based approaches.

RNA inverse folding. Most tools for RNA inverse folding focus on secondary structure without considering 3D geometry [Churkin et al., 2018, Runge et al., 2019] and approach the problem from the lens of energy optimisation [Ward et al., 2023]. Rosetta fixed backbone re-design [Das et al., 2010] is the only energy optimisation-based approach that accounts for 3D structure. Deep neural networks such as gRNAde can incorporate 3D structural constraints and are orders of magnitude faster than optimisation-based approaches; this is particularly attractive for high-throughput design pipelines as solving the inverse folding optimisation problem is NP hard [Bonnet et al., 2020].

RNA structure design. Inverse folding models for protein design have often been coupled with backbone generation models which design structural backbones conditioned on various design constraints [Watson et al., 2023, Ingraham et al., 2023, Didi et al., 2023]. Current approaches for RNA backbone design use classical (non-learnt) algorithms for aligning 3D RNA motifs [Han et al., 2017, Yesselman et al., 2019], which are small modular pieces of RNA that are believed to fold independently. Such algorithms may be restricted by the use of hand-crafted heuristics and we plan to explore data-driven generative models for RNA backbone design in future work.

RNA structure prediction. There have been several recent efforts to adapt protein folding architectures such as AlphaFold2 [Jumper et al., 2021] and RosettaFold [Baek et al., 2021] for RNA structure prediction [Li et al., 2023b, Wang et al., 2023, Baek et al., 2024]. A previous generation of models used GNNs as ranking functions together with Rosetta energy optimisation [Watkins et al., 2020, Townshend et al., 2021]. None of these architectures aim at capturing conformational flexibility of RNAs, unlike gRNAde which represents RNAs as multi-state conformational ensembles. Neither can structure prediction tools be used for RNA design tasks as they are not generative models.

RNA language models. Self-supervised language models have been developed for predictive and generative tasks on RNA sequences, including general-purpose models such as RNA FM [Chen et al., 2022] and RiNaLMo [Penic et al., 2024] as well as mRNA-specific CodonBERT [Li et al., 2023a]. RNA sequence data repositories are orders of magnitude larger than those for RNA structure (eg. RiNaLMo is trained on 36 million sequences). However, standard language models can only implicitly capture RNA structure and dynamics through sequence co-occurence statistics, which can pose a chellenge for designing structured RNAs such as riboswitches, aptamers, and ribozymes. RibonanzaNet [He et al., 2024] represents a recent effort in developing structure-informed RNA language models by supervised training on experimental readouts from chemical mapping, although RibonanzaNet cannot be used for RNA design. Inverse folding methods like gRNAde are language models conditioned on 3D structure, making them a natural choice for structure-based design.

B FAQs on using gRNAde

How to chose the number of states to provide as input to gRNAde? In general, this would depend on the design objective. For instance, designing riboswitches may necessitate multi-state design, while a single-state pipeline may be more sensible for locking an aptamer into its bound conformation [Yesselman et al., 2019]. Note that it may be possible to benefit from multi-state gRNAde models even when performing single-state design by using slightly noised variations of the same backbone structure as an input conformational ensemble.

How to prioritise or chose amongst designed sequences? We have currently provided 3 evaluation metrics: native sequence recovery, secondary structure self-consistency score and perplexity, towards this end. We suspect that recovery may not be the ideal choice, except for design scenarios where we require certain regions of the RNA to be conserved or native-like. Self-consistency score may be a more holistic evaluation metric as it accounts for alternative base pairings leading to similar structures, but may inherit the limitations of the secondary structure prediction method used as part of its computation. In realistic design scenarios, we can pair gRNAde with another machine learning model (an 'oracle') for ranking or predicting the suitability of designed sequences for the objective (for instance, binding affinity or some other notion of fitness). We hope to conduct further experimental validation of gRNAde designs in the wet lab in order to better understand these tradeoffs.

C Additional Figures and Tables

Figure 8: In-silico evaluation metrics for gRNAde.

Table 1: Full results for Figure 4 comparing gRNAde to Rosetta and FARNA for single-state design.

Figure 9: RNASolo data statistics.

Table 2: Aggregated benchmark results for gRNAde.

Figure 10: Multi-graph tensor representation of RNA conformational ensembles.

Listing 1: Pseudocode for multi-state GNN encoder layer.



Figure 8: **In-silico evaluation metrics for gRNAde designed sequences.** We consider (1) the percentage of native nucleotides recovered in designed samples, (2) self-consistency of 'forward folded' secondary structures of designs with the native secondary structure, as well as (3) perplexity of the model's prediction (not shown in the figure).

Table 1: Full results for Figure 4 comparing gRNAde to Rosetta and FARNA for single-state desig
on 14 RNA structures of interest identified by Das et al. [2010]. Rosetta and FARNA recovery value
are taken from Das et al. [2010], Supplementary Table 2.

		FARNA	Rosetta	gRNAde (single-state)		
PDB ID	Description	Recovery	Recovery	Recovery	Perplexity	Self-cons.
1CSL	RRE high affinity site	0.20	0.44	0.5719	1.2812	0.8644
1ET4	Vitamin B12 binding RNA aptamer	0.34	0.44	0.6250	1.3457	-0.0135
1F27	Biotin-binding RNA pseudoknot	0.36	0.37	0.3437	1.6203	0.4523
1L2X	Viral RNA pseudoknot	0.45	0.48	0.4721	1.3181	0.5692
1LNT	RNA internal loop of SRP	0.27	0.53	0.5843	1.4337	0.1379
1Q9A	Sarcin/ricin domain from E.coli 23S rRNA	0.40	0.41	0.5044	1.3411	0.0597
4FE5	Guanine riboswitch aptamer	0.28	0.36	0.5300	1.3824	0.9116
1X9C	All-RNA hairpin ribozyme	0.31	0.50	0.5000	1.3905	0.6630
1XPE	HIV-1 B RNA dimerization initiation site	0.24	0.40	0.7037	1.2177	0.7768
2GCS	Pre-cleavage state of glmS ribozyme	0.26	0.44	0.5078	1.3053	0.4062
2GDI	Thiamine pyrophosphate-specific riboswitch	0.38	0.48	0.6500	1.2363	-0.0251
20EU	Junctionless hairpin ribozyme	0.30	0.37	0.9519	1.0913	0.7768
2R8S	Synthetic FAB bound to ribozyme domain	0.36	0.53	0.5689	1.1881	0.7281
354D	Loop E from E. coli 5S rRNA	0.35	0.55	0.4410	1.4938	0.0430
	Overall recovery:	0.32	0.45	0.5682		



(a) **Sequence length.** The dataset is long-tailed in terms of RNA sequence length, with many short sequences including aptamers, riboswitches, ribozymes, and tRNAs (fewer than 200 nucleotides). The dataset also includes several longer ribosomal RNAs (thousands of nucleotides).



(b) **Number of structures per sequence.** The dataset covers a wide range of RNA conformation ensembles, with on average 3 structures per sequence. There are multiple structures available for 1,547 sequences. The remaining 2,676 sequences have one corresponding structure.





(c) **Average pairwise RMSD per sequence.** For 1,547 sequences with multiple structures, there is significant structural diversity among conformations. On average, the pairwise C4' RMSD among the set of structures for a sequence is greater than 1Å.

(d) **Bivariate distribution for sequence length vs. avg. RMSD.** The joint plot illustrates how structural diversity (measured by avg. pairwise RMSD) varies across sequence lengths. We notice similar structural variations regardless of sequence length.

Figure 9: **RNASolo data statistics.** We plot histograms to visualise the diversity of RNAs available in terms of (a) sequence length, (b) number of structures available per sequence, as well as (c) structural variation among conformations for those RNA that have multiple structures. The bivariate distribution plot (d) for sequence length vs. average pairwise RMSD illustrates structural diversity regardless of sequence lengths.

Table 2: Aggregated benchmark results for gRNAde. **Split:** Single- and multi-state splits are described in Section 3; models trained on 'All data' use all RNASolo samples for training and are evaluated on the multi-state test set, solely for the purpose of releasing the best possible gRNAde checkpoints for real-world usage. **Model:** 'AR' implies autoregressive decoding (described in Section 2.3), while 'NAR' implies non-autoregressive, one-shot decoding using an MLP. AR models have significantly higher self-consistency scores than NAR, as autoregressive decoding can condition predictions at each iteration on past predictions while one-shot decoding cannot. **Max. #states:** Multi-state models marginally improve native sequence recovery over an equivalent single state variant, even for the single-state benchmark. **Max. train RNA length:** Limiting the maximum length of RNAs used for training often leads to lower sequence recovery and higher self-consistency scores. Results in the main paper are reported for models shaded in gray which are all trained on maximum length of 5000 nucleotides.

Split	Model	Max. #states	Max. train RNA length	Native seq. recovery	Sec. struct. MCC self-consistency
Single-state split	AR	1	500	$0.4364 {\pm} 0.0059$	0.6206 ± 0.0662
	AR	1	1000	$0.4534{\pm}0.0063$	$0.6481 {\pm} 0.0037$
	AR	1	2500	$0.4945 {\pm} 0.0077$	$0.6278 {\pm} 0.0215$
	AR	1	5000	$0.5271 {\pm} 0.0019$	$0.5706 {\pm} 0.0119$
	AR	3	5000	$0.5386 {\pm} 0.0134$	$0.6255 {\pm} 0.0057$
	AR	5	5000	$0.5400 {\pm} 0.0270$	0.6006 ± 0.0345
Multi-state split	AR	1	500	$0.4452{\pm}0.0067$	$0.6031 {\pm} 0.0288$
	AR	1	1000	$0.4472 {\pm} 0.0132$	$0.5799 {\pm} 0.0156$
	AR	1	2000	$0.4799 {\pm} 0.0180$	$0.5668 {\pm} 0.0120$
	AR	1	5000	$0.4549 {\pm} 0.0082$	$0.5689 {\pm} 0.0185$
	AR	3	5000	$0.4666 {\pm} 0.0288$	$0.5605 {\pm} 0.0305$
	AR	5	5000	$0.4722 {\pm} 0.0106$	$0.5488 {\pm} 0.0323$
	NAR	1	5000	$0.4324{\pm}0.0097$	$0.4229 {\pm} 0.0050$
	NAR	3	5000	$0.4327 {\pm} 0.0212$	$0.4185 {\pm} 0.0249$
	NAR	5	5000	$0.4450 {\pm} 0.0116$	$0.4130 {\pm} 0.0177$
All data	AR	1	5000	$0.7361 {\pm} 0.0039$	$0.6301 {\pm} 0.0160$
	AR	2	5000	$0.7845 {\pm} 0.0099$	$0.6215 {\pm} 0.0218$
	AR	3	5000	$0.7886 {\pm} 0.0106$	$0.6052 {\pm} 0.0128$
	AR	5	5000	$0.8110 {\pm} 0.0077$	$0.6135 {\pm} 0.0280$



Figure 10: Multi-graph tensor representation of RNA conformational ensembles, and the associated symmetry groups acting on each axis. We process a set of k RNA backbone conformations with n nodes each into a tensor representation. Each multi-state GNN layer updates the tensor while being equivariant to the underlying symmetries; pseudocode is available in Listing 1. Here, we show a tensor of 3D vector-type features with shape $n \times k \times 3$. As depicted in the equivariance diagram, the updated tensor must be equivariant to permutation S_n of n nodes for axis 1, permutation S_k of k conformers for axis 2, and rotation SO(3)/O(3) of the 3D features for axis 3.

```
1 class MultiGVPConv(MessagePassing):
       '''GVPConv for handling multiple conformations'''
2
3
       def __init__(self, ...):
4
5
           . . .
6
       def forward(self, x_s, x_v, edge_index, edge_attr):
7
8
           # stack scalar feats along axis 1:
9
           # [n_nodes, n_conf, d_s] -> [n_nodes, n_conf * d_s]
10
           x_s = x_s.view(x_s.shape[0], x_s.shape[1] * x_s.shape[2])
           # stack vector feat along axis 1:
13
           # [n_nodes, n_conf, d_v, 3] -> [n_nodes, n_conf * d_v*3]
x_v = x_v.view(x_v.shape[0], x_v.shape[1] * x_v.shape[2]*3)
14
15
16
17
           # message passing and aggregation
           message = self.propagate(
18
19
                edge_index, s=x_s, v=x_v, edge_attr=edge_attr)
20
21
           # split scalar and vector channels
           return _split_multi(message, d_s, d_v, n_conf)
22
23
24
      def message(self, s_i, v_i, s_j, v_j, edge_attr):
25
26
           # unstack scalar feats:
           # [n_nodes, n_conf * d] -> [n_nodes, n_conf, d_s]
27
           s_i = s_i.view(s_i.shape[0], s_i.shape[1]//d_s, d_s)
28
           s_j = s_j.view(s_j.shape[0], s_j.shape[1]//d_s, d_s)
29
30
           # unstack vector feats:
31
           # [n_nodes, n_conf * d_v*3] -> [n_nodes, n_conf, d_v, 3]
32
           v_i = v_i.view(v_i.shape[0], v_i.shape[1]//(d_v*3), d_v, 3)
v_j = v_j.view(v_j.shape[0], v_j.shape[1]//(d_v*3), d_v, 3)
33
34
35
           # message function for edge j-i
36
37
           message = tuple_cat((s_j, v_j), edge_attr, (s_i, v_i))
           message = self.message_func(message) # GVP
38
39
40
           # merge scalar and vector channels along axis 1
           return _merge_multi(*message)
41
42
43
  def _split_multi(x, d_s, d_v, n_conf):
       , , ,
44
45
       Splits a merged representation of (s, v) back into a tuple.
46
       , , ,
      s = x[..., :-3 * d_v * n_conf].view(x.shape[0], n_conf, d_s)
47
      v = x[..., -3 * d_v * n_conf:].view(x.shape[0], n_conf, d_v, 3)
48
      return s, v
49
50
51 def _merge_multi(s, v):
       , , ,
52
      Merges a tuple (s, v) into a single 'torch.Tensor',
53
       where the vector channels are flattened and
54
       appended to the scalar channels.
55
56
      , , ,
      # s: [n_nodes, n_conf, d] \rightarrow [n_nodes, n_conf * d_s]
57
      s = s.view(s.shape[0], s.shape[1] * s.shape[2])
58
      # v: [n_nodes, n_conf, d, 3] -> [n_nodes, n_conf * d_v*3]
v = v.view(v.shape[0], v.shape[1] * v.shape[2]*3)
59
60
      return torch.cat([s, v], -1)
```

Listing 1: **PyG-style pseudocode for a multi-state GVP-GNN layer.** We update node features for each conformer independently while maintaining permutation equivariance of the updated feature tensors along both the first (no. of nodes) and second (no. of conformations) axes.