TANGO: RNA Topology and Geometry Co-Design

Tianmeng Hu[1,2,3] **Biao Luo**[4*] **Ke Li**[1,2,3,5*]
[1]Earlham Institute, UK. [2]Living Systems Institute, University of Exeter, UK.
[3]Department of Computer Science, University of Exeter, UK.
[4]School of Automation, Central South University.
[5]School of Computing Sciences, University of East Anglia, UK.
tianmeng0824@gmail.com, biao.luo@hotmail.com, Ke.Li@earlham.ac.uk

Abstract

Designing RNA sequences that satisfy target tertiary-structure objectives is an important problem in RNA engineering, with broad applications in synthetic biology and therapeutics. Existing 3D RNA design methods have made substantial progress, but they largely treat the target fold as a geometry-only objective. This is incomplete for RNA, whose functional folds are specified not only by tertiary geometry but also by secondary base-pairing topology. We therefore formulate RNA inverse design as a *topology-and-geometry co-design* problem and propose TANGO, a *multi-objective multi-agent reinforcement learning* framework for this task. TANGO adopts a staged curriculum: it first learns a topology-aware policy from secondary-structure targets, then transfers this policy to tertiary-structure targets to train a general co-design model, and finally applies frontier-guided fine-tuning to refine target-specific trade-offs among secondary-structure fidelity, tertiary-structure fidelity, and sequence diversity. Experiments show that TANGO generates diverse and novel designs, improving tertiary- and secondary-structure performance by 18.3% and 52.7%, respectively, over the strongest baseline.

1 Introduction

RNA plays a central role in molecular biology. It participates in every stage of gene regulation, from transcription and processing to export, translation, and degradation [1–3]. Unlike DNA, RNA molecules are not static lines of code. They fold into complex secondary (2D) and tertiary (3D) structures in vivo. These structures are not merely passive geometries; they actively dictate biological traits. Computational RNA inverse design aims to generate sequences that satisfy specified structural objectives [4], and recent progress in deep learning has made tertiary-structure-based RNA design increasingly practical [5–10].

RNA design, however, differs from standard structure-conditioned protein design in a central way: an RNA fold is organized not only by 3D geometry but also by an explicit base-pairing topology. Local Watson-Crick and wobble interactions form stems, loops, junctions, and pseudoknots, providing a secondary-structure scaffold for tertiary assembly [11]. Long-range and non-canonical interactions then stabilize particular 3D conformations within the topologically allowed ensemble [12–14]. Thus, neither secondary topology nor tertiary geometry alone fully specifies an RNA design objective: topology constrains the feasible fold space, while geometry determines the final spatial arrangement.

Existing 3D RNA design methods have made important progress by conditioning sequence generation on tertiary structures. Some also incorporate secondary-structure information as priors and constraints [7] or repair heuristics [5]. Nevertheless, topology fidelity is rarely treated as a first-class optimization and evaluation target jointly with tertiary-structure fidelity. This leaves a gap between

*Corresponding authors.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

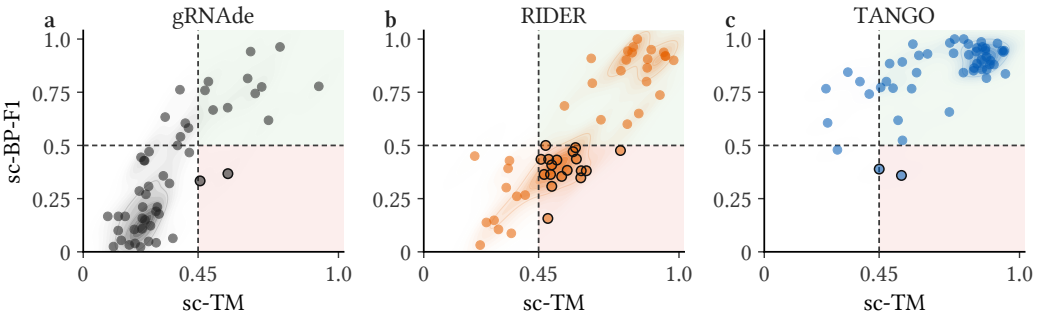

Figure 1: **RNA** 3D **inverse design requires joint topology and geometry fidelity.** Each point denotes a designed sequence evaluated by fold-back self-consistency, with tertiary-structure fidelity measured by sc-TM and secondary-topology fidelity measured by sc-BP-F1. The upper-right region contains designs that satisfy both objectives. Existing 3D RNA design methods can produce geometry-consistent but topology-inconsistent designs, showing that tertiary similarity alone does not guarantee recovery of the target base-pairing topology. Conversely, topology preservation alone is insufficient for specifying the final tertiary fold.

what is optimized and what functional RNA design often requires: a generated sequence should not merely fold back to a geometrically similar structure, but should also recover the base-pairing organization that scaffolds that fold. A detailed discussion of related work is provided in Appendix A.

This gap is not merely conceptual. As shown in Figure 1, strong tertiary-structure-based baselines can generate candidates with reasonable fold-back geometry but poor recovery of the target base pairs. Conversely, candidates that preserve much of the secondary topology may still fail to realize the desired 3D arrangement. These failure modes indicate that topology fidelity and geometry fidelity are complementary objectives rather than redundant metrics. They also suggest that RNA 3D inverse design should be evaluated as a *co-design problem*: a successful sequence should recover both the target tertiary geometry and the target base-pairing topology.

We therefore formulate RNA 3D inverse design as *topology-and-geometry co-design*. Given a target tertiary structure and its associated base-pairing topology, the goal is to generate sequences whose predicted folds match both. This objective is challenging for two reasons. First, base-paired positions are combinatorially coupled: changing one nucleotide may require a coordinated change at its pairing partner. Second, tertiary-structure rewards are delayed and sparse, because candidate sequences must be completed and folded before their 3D quality can be assessed. Although the sequence space contains 4^L possibilities for length L, only a small subset is compatible with a target topology, and only a smaller subset of those sequences realizes the target tertiary geometry.

We propose `TANGO`, a *multi-objective multi-agent reinforcement learning* framework for RNA topology-and-geometry co-design. `TANGO` uses the target topology to decompose sequence generation into structural units: unpaired residues are designed individually, while paired positions are designed jointly. Multiple agents fill these units in parallel under a shared structural reward, reducing the effective decision horizon relative to nucleotide-by-nucleotide generation and enabling coordinated decisions over coupled base pairs. Training proceeds in three stages. First, topology-aware policy initialization learns sequence-construction patterns that satisfy secondary-structure scaffolds. Second, geometry-conditioned training transfers this policy to tertiary-structure targets, allowing 3D information to adjust topology-compatible sequence decisions. Third, frontier-guided fine-tuning adapts the general policy to each target by directly improving the Pareto frontier over secondary-topology fidelity, tertiary-structure fidelity, and sequence diversity. Experiments on a curated PDB-derived benchmark show that `TANGO` improves tertiary fold-back fidelity by 18.3% and secondary-topology fidelity by 52.7% over the strongest baseline, while generating diverse non-native designs.

2 Problem Setup

RNA sequence, topology, and geometry. An RNA sequence is a length-L string $x = (x_1, \ldots, x_L)$ over the alphabet $\mathcal{V} = \{\mathtt{A}, \mathtt{U}, \mathtt{G}, \mathtt{C}\}$. We represent the target secondary topology as $\mathcal{T}^\star = (\mathcal{P}^\star, \mathcal{S}^\star)$, where $\mathcal{P}^\star \subseteq \{(p, q) : 1 \le p < q \le L\}$ is the set of target base pairs and $\mathcal{S}^\star$ is the set of unpaired positions. The same topology can be written in dot-bracket notation; pseudoknots correspond to crossing pairs in $\mathcal{P}^\star$ and are represented with additional bracket types. Unless otherwise stated, $\mathcal{P}^\star$ denotes the canonical and wobble base-pairing scaffold of the target. Non-canonical tertiary contacts are not imposed as discrete pair actions, but are evaluated through the tertiary geometry objective.

The target tertiary structure is denoted by $\mathcal{G}^\star$. We use a 3-bead coarse-grained representation [15], where each nucleotide is described by the Cartesian coordinates of the backbone atoms P and $\mathrm{C4}'$ and an unlabeled glycosidic nitrogen bead (N1 for pyrimidines and N9 for purines in the source structure). The model receives geometric coordinates but not native nucleotide identities. This representation compactly captures backbone conformation and approximate base placement for tertiary-structure-conditioned design.

Forward oracles. Given a candidate sequence x, a tertiary-structure oracle F_{3D} predicts a fold $\widehat{\mathcal{G}}(x) = F_{\mathrm{3D}}(x)$. We distinguish two secondary-structure oracles. During topology-only initialization, we use a secondary folding oracle $\widehat{\mathcal{T}}^{\mathrm{fold}}(x) = F_{\mathrm{2D}}^{\mathrm{fold}}(x)$, which predicts secondary structure directly from sequence. During geometry-conditioned training and evaluation, we derive secondary topology from the predicted tertiary structure: $\widehat{\mathcal{T}}^{\mathrm{ann}}(x) = A_{\mathrm{2D}}(F_{\mathrm{3D}}(x))$, where A_{2D} is a 3D-to-2D annotation procedure such as DSSR [16]. This distinction lets the first stage learn from a cheaper topology objective, while later stages measure topology and geometry under the same fold-back structure.

Topology-and-geometry co-design. RNA 3D inverse design is commonly evaluated by tertiary fold-back fidelity. We instead formulate a co-design objective that requires both tertiary-structure fidelity and secondary-topology fidelity. For a completed sequence x, we define

$$R_{\mathrm{3D}}(x; \mathcal{G}^\star) = \mathrm{TM}\Big(\widehat{\mathcal{G}}(x), \mathcal{G}^\star\Big), \tag{1}$$

and

$$R_{\mathrm{2D}}(x; \mathcal{T}^\star) = \text{BP-F1}\Big(\widehat{\mathcal{T}}^{\mathrm{ann}}(x), \mathcal{T}^\star\Big), \tag{2}$$

where TM-score measures global tertiary alignment and BP-F1 measures recovery of the target base pairs. A sequence is considered a co-design success when it satisfies both

$$R_{\mathrm{3D}}(x; \mathcal{G}^\star) \ge \tau_{\mathrm{3D}}, \qquad R_{\mathrm{2D}}(x; \mathcal{T}^\star) \ge \tau_{\mathrm{2D}}. \tag{3}$$

Unlike prior sequence–structure co-design [17], which jointly generates sequences and structures, our formulation designs RNA sequences to recover prescribed topology and geometry targets. Beyond finding a single high-scoring sequence, we seek a diverse set of candidates that achieve favorable trade-offs among topology fidelity, geometry fidelity, and sequence novelty.

3 Method

TANGO is a cooperative reinforcement learning framework for RNA topology-and-geometry co-design. Given a target pair $(\mathcal{G}^\star, \mathcal{T}^\star)$, TANGO generates RNA sequences by filling topology-derived structural units rather than individual nucleotides. Unpaired residues are designed as single-position units, while paired residues are designed jointly as paired units. This unit-level factorization exposes the base-pair couplings in the target topology and shortens the decision horizon.

Training proceeds in three stages. First, topology-aware initialization learns a policy for satisfying secondary-structure scaffolds. Second, geometry-conditioned co-design transfers this policy to tertiary-structure targets by adding 3D geometric conditioning and optimizing a joint topology-and-geometry reward. Third, frontier-guided fine-tuning adapts the general policy to each target by improving the Pareto frontier over topology fidelity, geometry fidelity, and sequence diversity. Figure 2 illustrates the overall pipeline. Pseudocode is provided in Appendix B.

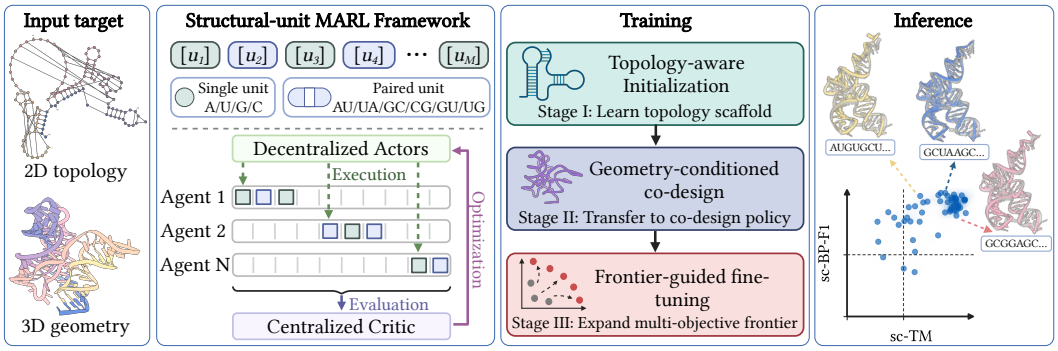

Figure 2: Overview of TANGO. TANGO formulates RNA design as cooperative sequence generation over structural units and follows a three-stage training pipeline.

3.1 Structural-unit Cooperative Generation

Structural units. Given the target topology $\mathcal{T}^\star = (\mathcal{P}^\star, \mathcal{S}^\star)$, we construct an ordered list of structural units $\mathcal{U}^\star = (u_1, \ldots, u_M)$. Each unpaired position $p \in \mathcal{S}^\star$ forms a single-position unit $u = \{p\}$, and each target base pair $(p, q) \in \mathcal{P}^\star$ forms a paired unit $u = \{p, q\}$. Crossing pairs from pseudoknots are handled identically because the representation is set-based rather than nested. The units cover the full sequence and assign each residue to exactly one unit.

The action space depends on the unit type:

$$\mathcal{A}_{\text{single}} = \{\texttt{A}, \texttt{U}, \texttt{G}, \texttt{C}\}, \qquad \mathcal{A}_{\text{pair}} = \{\texttt{AU}, \texttt{UA}, \texttt{GC}, \texttt{CG}, \texttt{GU}, \texttt{UG}\}. \tag{4}$$

Thus, paired units are filled by coordinated canonical or wobble base-pair assignments. This action factorization encodes the explicit secondary scaffold; non-canonical tertiary interactions are left to the 3D structural reward.

Multi-agent generation. The ordered unit list is split into N contiguous segments, one for each agent. Agent i maintains a cursor over its segment. At environment step t, every active agent observes its current unit, samples an action for that unit, writes the corresponding nucleotide assignment into the partial sequence, and advances its cursor. The episode terminates when all units have been assigned. Since up to N units are filled per step, the effective horizon is approximately $\lceil M/N \rceil$ rather than M.

All agents share a single actor network with unit-type-specific output heads. A centralized critic observes a pooled representation of the full partial design state during training, while each agent acts from its local observation at inference time. We optimize the policy under centralized training and decentralized execution. A formal MARL formulation and the training objective are provided in Appendix C.

Topology-aware unit embeddings. For each residue position p, we embed the dot-bracket token and structural metadata $\boldsymbol{r}_p = E_{\text{res}}(b_p, \boldsymbol{m}_p)$, where $\boldsymbol{m}_p$ includes pairing status, normalized position, pairing distance, and pseudoknot indicators. A single-position unit $u_j = \{p\}$ receives embedding

$$\boldsymbol{e}_j = E_{\text{single}}\big([\boldsymbol{r}_p, \boldsymbol{m}_{u_j}]\big)\,, \tag{5}$$

and a paired unit $u_j = \{p, q\}$ receives embedding

$$\boldsymbol{e}_j = E_{\text{pair}}\big([\boldsymbol{r}_p, \boldsymbol{r}_q, \boldsymbol{r}_p \odot \boldsymbol{r}_q, |\boldsymbol{r}_p - \boldsymbol{r}_q|, \boldsymbol{m}_{u_j}]\big)\,. \tag{6}$$

The actor observation for agent i at step t is $o_t^i = [\boldsymbol{e}_{j_t^i}, \boldsymbol{q}_{t,j_t^i}, \boldsymbol{c}_{j_t^i}]$, where j_t^i is the current unit index, $\boldsymbol{q}_{t,j_t^i}$ encodes the current assignment state of the unit, and $\boldsymbol{c}_{j_t^i}$ is a local context vector obtained by mean pooling embeddings in a fixed window around $u_{j_t^i}$.

3.2 Stage I: Topology-aware Initialization

The first stage trains the actor to construct sequences that satisfy target secondary topologies. After a complete sequence x_T is generated, the topology-only oracle predicts $\widehat{\mathcal{T}}^{\text{fold}}(x_T) = F_{\text{2D}}^{\text{fold}}(x_T)$. The

terminal reward is

$$R_{\text{topo}}(x_T; \mathcal{T}^\star) = \text{BP-F1}\left(\widehat{\mathcal{T}}^{\text{fold}}(x_T), \mathcal{T}^\star\right) + \beta_{\text{ex}} \mathbb{I}\left[\widehat{\mathcal{T}}^{\text{fold}}(x_T) = \mathcal{T}^\star\right]. \tag{7}$$

All agents receive this shared reward. The policy is trained by maximizing

$$\max_{\theta_{\text{topo}}} \mathbb{E}_{x_T \sim \pi_{\theta_{\text{topo}}}} \left[R_{\text{topo}}(x_T; \mathcal{T}^\star)\right]. \tag{8}$$

This stage learns coordinated nucleotide assignments for single and paired structural units, providing a scaffolded initialization for subsequent 3D co-design.

3.3 Stage II: Geometry-conditioned Co-design

The second stage transfers the topology-aware policy to tertiary-structure targets. The structural units and action spaces remain unchanged, but each unit is augmented with geometric information in $\mathcal{G}^\star$.

Geometry encoder. We encode the target tertiary structure with a `GVP-GNN` [18] over the 3-bead representation:

$$\boldsymbol{h}_1^{\mathcal{G}}, \ldots, \boldsymbol{h}_L^{\mathcal{G}} = E_{\text{GVP}}(\mathcal{G}^\star). \tag{9}$$

Residue-level embeddings are pooled into unit-level geometric embeddings. For a single-position unit $u_j = \{p\}$, we set $\boldsymbol{g}_j = \boldsymbol{h}_p^{\mathcal{G}}$. For a paired unit $u_j = \{p, q\}$, we use

$$\boldsymbol{g}_j = \frac{1}{2}\left(\boldsymbol{h}_p^{\mathcal{G}} + \boldsymbol{h}_q^{\mathcal{G}}\right). \tag{10}$$

A local 3D context vector is computed by pooling spatial neighbors of each unit:

$$\boldsymbol{c}_j^{\text{3D}} = \text{MeanPool}\left(\{\boldsymbol{g}_k : k \in \mathcal{W}_{\text{3D}}(j)\}\right). \tag{11}$$

Geometry-conditioned policy. For agent i at unit j_t^i, the geometry-conditioned observation is $o_{t,\text{geo}}^i = [o_t^i, \boldsymbol{g}_{j_t^i}, \boldsymbol{c}_{j_t^i}^{\text{3D}}]$. The actor is initialized from the topology-aware policy and augmented with two zero-initialized 3D-conditioned adapters. The final logits are

$$\ell_i = \ell_i^{\text{topo}} + \Delta\ell_i^{\text{FiLM}} + \Delta\ell_i^{\text{3D}}, \tag{12}$$

where ℓ_i^{topo} is the inherited topology-aware logit vector, $\Delta\ell_i^{\text{FiLM}}$ is produced by FiLM modulation of the hidden state using $\boldsymbol{g}_{j_t^i}$, and $\Delta\ell_i^{\text{3D}}$ is a direct geometric correction. Zero initialization makes the transferred policy start near the topology-aware policy, so 3D conditioning initially acts as a small residual adjustment rather than overwriting the learned topology scaffold.

Co-design reward. At the end of an episode, the completed sequence x_T is folded by the 3D oracle: $\widehat{\mathcal{G}}(x_T) = F_{\text{3D}}(x_T)$ and $\widehat{\mathcal{T}}^{\text{ann}}(x_T) = A_{\text{2D}}\left(\widehat{\mathcal{G}}(x_T)\right)$. Based on Eqs. (1) and (2), we compute the Stage-II terminal reward as

$$R_{\text{co}}(x_T) = R_{\text{3D}}(x_T; \mathcal{G}^\star) + \lambda_{\text{2D}} R_{\text{2D}}(x_T; \mathcal{T}^\star) - \lambda_{\text{viol}}\left[\tau_{\text{2D}} - R_{\text{2D}}(x_T; \mathcal{T}^\star)\right]_+. \tag{13}$$

The positive R_{2D} term encourages topology recovery, while the hinge penalty discourages geometry-improving updates that fall below the desired topology-fidelity threshold. This stage produces a general topology-and-geometry co-design policy.

3.4 Stage III: Frontier-guided Fine-tuning

The final stage performs target-wise policy adaptation. This shifts costly trial-and-error exploration from the wet lab to *in silico* optimization, prioritizing promising candidates for subsequent experimental validation. The trade-offs among topology fidelity, geometric fidelity, and sequence diversity motivate a multi-objective reinforcement learning formulation [19–24], rather than optimization under a single fixed scalarization. `TANGO` therefore fine-tunes the policy to expand the target-specific Pareto frontier.

Archive and objective vectors. For each target $(\mathcal{G}^\star, \mathcal{T}^\star)$, we maintain an archive of evaluated sequences, $\mathcal{H}_k^\star = \{x^{(1)}, \ldots, x^{(K)}\}$, at fine-tuning iteration k. The archive is warm-started with candidates sampled from the Stage-II co-design policy. Each sequence is mapped to an objective vector. The default bi-objective setting uses

$$\boldsymbol{r}(x) = [R_{\mathrm{2D}}(x; \mathcal{T}^\star), R_{\mathrm{3D}}(x; \mathcal{G}^\star)]. \tag{14}$$

We also consider a diversity-aware tri-objective setting:

$$\boldsymbol{r}_{\mathrm{div}}(x) = [R_{\mathrm{2D}}(x; \mathcal{T}^\star), R_{\mathrm{3D}}(x; \mathcal{G}^\star), R_{\mathrm{div}}(x; \mathcal{H}_k^\star)], \tag{15}$$

where

$$R_{\mathrm{div}}(x; \mathcal{H}_k^\star) = \min_{x' \in \mathcal{H}_k^\star} \frac{1}{L} \sum_{p=1}^{L} \mathbb{I}[x_p \neq x'_p]. \tag{16}$$

All objectives are normalized to $[0, 1]$ with larger values preferred. Duplicate sequences receive zero diversity and are inserted into the archive only once.

A sequence x dominates y if it is no worse in every selected objective and strictly better in at least one. The non-dominated archive is $\mathcal{F}_k^\star = \mathrm{ND}(\mathcal{H}_k^\star)$.

Batch-level hypervolume reward. At iteration k, the current policy samples a batch of candidates $\mathcal{C}_k$. We evaluate their objective vectors and compute the post-batch frontier $\mathcal{P}_k = \mathrm{ND}(\mathcal{H}_k^\star \cup \mathcal{C}_k)$. Let $\mathcal{R}(\mathcal{A}) = \{\boldsymbol{r}(x) : x \in \mathcal{A}\}$ denote the set of objective vectors for a sequence set $\mathcal{A}$. With reference point $\boldsymbol{r}_{\mathrm{ref}}$ in the selected objective space, the frontier reward for candidate $x \in \mathcal{C}_k$ is its exclusive hypervolume contribution:

$$R_{\mathrm{front}}(x) = \mathrm{HV}\left(\mathcal{R}(\mathcal{P}_k); \boldsymbol{r}_{\mathrm{ref}}\right) - \mathrm{HV}\left(\mathcal{R}\left(\mathrm{ND}\left(\mathcal{H}_k^\star \cup (\mathcal{C}_k \setminus \{x\})\right)\right); \boldsymbol{r}_{\mathrm{ref}}\right). \tag{17}$$

The batch-level definition makes the reward independent of the order in which candidates are inserted into the archive. A candidate receives positive reward only if it expands the target-wise frontier. After reward computation, all evaluated candidates are added to the archive to form $\mathcal{H}_{k+1}^\star$.

Fine-tuning objective. Starting from the Stage-II co-design policy, we optimize

$$\max_{\theta_{\mathrm{front}}} \mathbb{E}_{x_T \sim \pi_{\theta_{\mathrm{front}}}} \left[R_{\mathrm{front}}(x_T)\right]. \tag{18}$$

The reward is terminal and shared across agents. The resulting policy is target-adapted: it is rewarded not for improving a fixed scalarized score, but for producing sequences that expand the achievable frontier of topology fidelity, geometry fidelity, and diversity.

4 Experimental Setup

Benchmark construction. We construct a topology-and-geometry co-design benchmark from experimentally resolved RNA structures in the PDB [25]. We retain single-chain RNA structures with resolution better than 4 Å, remove structures with chain breaks, and deduplicate targets by secondary topology. To make predictor-based fold-back evaluation meaningful, we refold each native sequence with both `RhoFold` [26] and `AlphaFold3` [27], and retain only targets whose native sequence recovers the experimental fold with $\mathrm{TM\text{-}score} \geq 0.45$ under both predictors, a threshold commonly used to indicate shared global RNA topology [28]. This filtering yields $1{,}213$ targets. It reduces ambiguity in fold-back evaluation, but also biases the benchmark toward RNAs that current structure predictors can fold reasonably well.

We split targets by structural dissimilarity rather than random sequence-level partitioning. Specifically, we cluster targets by pairwise tertiary-structure similarity using k-medoids and choose held-out targets with the lowest nearest-neighbor similarity to the training set. The final test set contains 50 structurally remote targets, with mean nearest-neighbor similarity 0.34 and maximum nearest-neighbor similarity 0.42 to the training set. Among these targets, 44% contain pseudoknotted base pairs.

Table 1: Structural design performance on the test set. Each method generates $K = 10$ candidates per target. Continuous metrics report mean $\pm$ standard deviation over target-level aggregated scores. Success denotes *co-design success*, defined as sc-TM > 0.45 and sc-BP-F1 > 0.50. Solved targets denotes the fraction of targets for which at least one candidate satisfies both thresholds. Threshold sensitivity is reported in Appendix D.1.

Method	sc-TM ↑	sc-BP-F1 ↑	sc-RMSD ↓	**Success rate ↑**	**Solved targets ↑**
`gRNAde`	0.32 ± 0.12	0.36 ± 0.26	3.54 ± 1.00	0.17	38.00%
`RDesign`	0.31 ± 0.10	0.31 ± 0.21	3.61 ± 0.91	0.13	42.00%
`R3Design`	0.35 ± 0.14	0.33 ± 0.27	3.47 ± 1.03	0.21	50.00%
`RiboDiffusion`	0.32 ± 0.13	0.30 ± 0.25	3.58 ± 1.14	0.17	42.00%
`RIDER`	0.60 ± 0.13	0.55 ± 0.22	2.94 ± 1.26	0.49	60.00%
`TANGO`	$\mathbf{0.71 \pm 0.17}$	$\mathbf{0.84 \pm 0.15}$	$\mathbf{2.16 \pm 1.07}$	**0.86**	**88.00%**

Evaluation protocol and metrics. We instantiate the co-design evaluation protocol defined in Section 2. Unless otherwise stated, main results use `RhoFold` as the fold-back oracle; cross-oracle experiments evaluate the same designed sequences with `AlphaFold3`. For each target, we first fold the native sequence with the selected oracle to obtain a predictor-specific reference structure. We then fold each designed sequence with the same oracle and compare the two predicted structures. This self-consistency protocol avoids mixing experimental structures and predictor-specific folds in the primary metric.

We report tertiary fidelity by self-consistency TM-score (sc-TM) and self-consistency RMSD (sc-RMSD). We report secondary-topology fidelity by extracting base pairs from the predicted designed structure with DSSR [16] and computing self-consistency base-pair F1 (sc-BP-F1) against the target topology. Our primary task-level metric is co-design success: a candidate is successful only if it satisfies both

$$\text{sc-TM} > \tau_{3D} \quad \text{and} \quad \text{sc-BP-F1} > \tau_{2D}. \tag{19}$$

We use $\tau_{3D} = 0.45$ and $\tau_{2D} = 0.50$ in the main experiments, and report threshold sensitivity in Appendix D.1. The co-design solved target rate is the fraction of test targets for which at least one generated candidate satisfies both thresholds.

Baselines and inference budget. We compare against recent 3D RNA inverse-design methods: `gRNAde` [6], `RDesign` [5], `R3Design` [29], `RiboDiffusion` [8], and `RIDER` [9]. All baselines are retrained on the same training split and evaluated on the same held-out targets. Each method generates $K = 10$ candidate sequences per target under its default sampling procedure.

Metrics and statistics. For continuous metrics, we first aggregate the K candidates for each target using mean-of-K, and then report mean $\pm$ standard deviation across targets. For success metrics, we report candidate-level success rate and target-level solved rate separately. For sc-TM, sc-BP-F1, and success rate, we compute p-values using a two-sided paired Wilcoxon signed-rank test over target-level scores, followed by Holm correction for multiple comparisons. For diversity metrics and success-conditioned novelty, where the set of valid targets can differ across methods, we use a two-sided Wilcoxon rank-sum test. For solved target rates, we use a two-proportion z-test.

5 Results

5.1 `TANGO` improves topology and geometry fidelity

We first evaluate whether `TANGO` improves RNA 3D inverse design under the co-design criterion defined in Section 4. Table 1 compares `TANGO` with five recent 3D RNA inverse-design methods: `gRNAde`, `RDesign`, `R3Design`, `RiboDiffusion`, and `RIDER`. All methods are trained on the same split and evaluated on the same structurally remote test targets.

`TANGO` achieves the strongest performance on both topology and geometry metrics. Relative to the strongest baseline, `RIDER`, `TANGO` improves sc-TM from 0.60 to 0.71 and sc-BP-F1 from 0.55 to 0.84, corresponding to relative gains of 18.3% and 52.7%, respectively. The improvement in

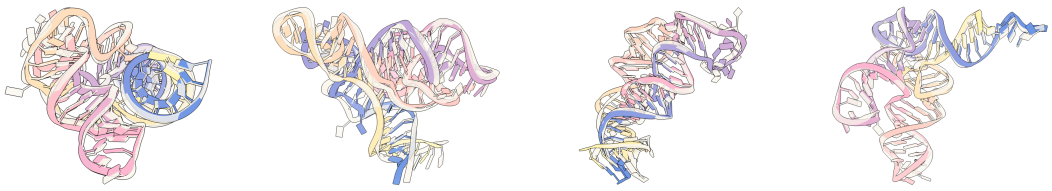

Figure 3: Representative co-design successes for PDB targets `3IQP`, `6WJS`, `3UCU`, and `3ADB`. Gray denotes the target reference structure, and color denotes the fold-back structure of the designed sequence. All shown designs satisfy the joint co-design success criterion.

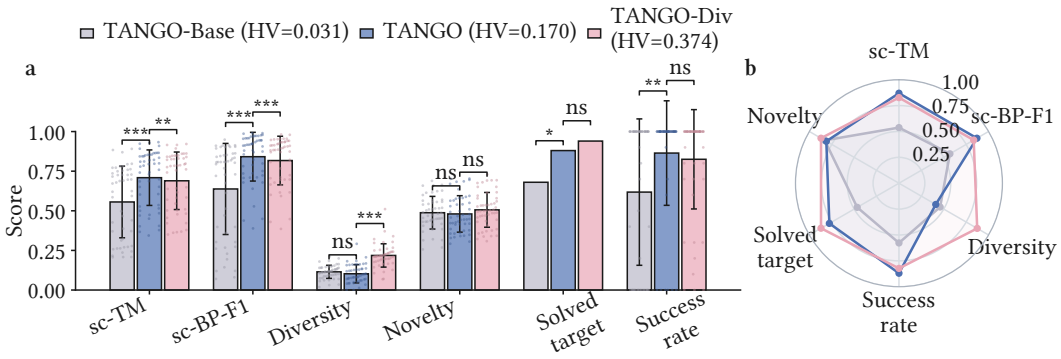

Figure 4: Target-wise frontier-guided refinement. (a) Structural quality and sequence diversity for the stage-II policy (`TANGO-Base`), the bi-objective frontier-refined model (`TANGO`), and the diversity-aware tri-objective variant (`TANGO-Div`). (b) Multi-objective trade-off visualization between design quality and diversity. HV denotes the dominated hypervolume of each method's Pareto frontier in the normalized objective space. Significance markers denote $*$ for $p < 0.05$, $**$ for $p < 0.01$, $***$ for $p < 0.001$, and ns for not significant.

sc-BP-F1 is central to our setting: the gain in tertiary fold-back fidelity is accompanied by substantially better recovery of the target base-pairing topology. `TANGO` also obtains the lowest sc-RMSD, indicating lower structural deviation after alignment.

The joint success metrics provide the primary evidence for topology-and-geometry co-design. `TANGO` improves the co-design success rate from 0.49 to 0.86 and the solved target rate from 60.00% to 88.00% compared with `RIDER`. Because a successful candidate must satisfy both the geometry and topology thresholds, these results support our formulation of RNA 3D inverse design as a joint co-design problem rather than a geometry-only fold-back task. Representative co-design successes are shown in Figure 3.

5.2 Target-wise frontier-guided refinement expands the co-design frontier

We next isolate the contribution of the target-wise refinement stage. Figure 4(a) compares the stage-II policy before target-wise fine-tuning (`TANGO-Base`), the default bi-objective frontier-refined model (`TANGO`), and the diversity-aware tri-objective variant (`TANGO-Div`). Compared with `TANGO-Base`, `TANGO` improves sc-TM by 27.6% and sc-BP-F1 by 31.9%, indicating that the target-wise refinement stage contributes to both tertiary fold-back fidelity and secondary-topology recovery. The improvements in success rate and solved target rate further show that the refined policies place more candidates in the high-sc-BP-F1, high-sc-TM region, rather than improving only one objective.

Adding sequence diversity as an explicit objective increases diversity by 90.1% relative to `TANGO-Base`, while preserving competitive topology and geometry fidelity. Figure 4(b) summarizes the resulting multi-objective trade-offs. The hypervolume increases from 0.031 for `TANGO-Base` to

0.170 for TANGO and 0.374 for TANGO-Div, indicating that frontier-guided refinement expands the set of useful target-specific candidates.

To test whether this gain is merely due to additional test-time sampling, Appendix D.2 compares against sampling-based re-ranking under the same candidate budget, and Appendix D.3 compares against fixed weighted-reward fine-tuning. These controls suggest that the improvement comes from optimizing the target-wise frontier, rather than only from selecting better samples from an unchanged policy or using a fixed scalarization of the objectives.

5.3 Topology-aware initialization provides a useful scaffold for 3D design

Here we ablate the Stage-I topology-aware initialization by using the same structural-unit action space and geometry-conditioned architecture, but starting 3D co-design training without first learning to satisfy secondary-structure scaffolds. Figure 5 shows that removing this initialization substantially degrades all primary co-design metrics. With topology-aware initialization, TANGO-Base achieves relative improvements of 38.9% in sc-TM, 93.3% in sc-BP-F1, and 106.0% in co-design success rate.

This ablation supports the curriculum design. Although Stage I optimizes only a secondary-topology objective, it improves both topology recovery and tertiary fold-back fidelity after 3D training.

This is consistent with the intuition that a topology-compatible initialization places the policy in a better exploration region for subsequent tertiary optimization, where fold-back rewards are less sparse and base-paired positions are already coordinated.

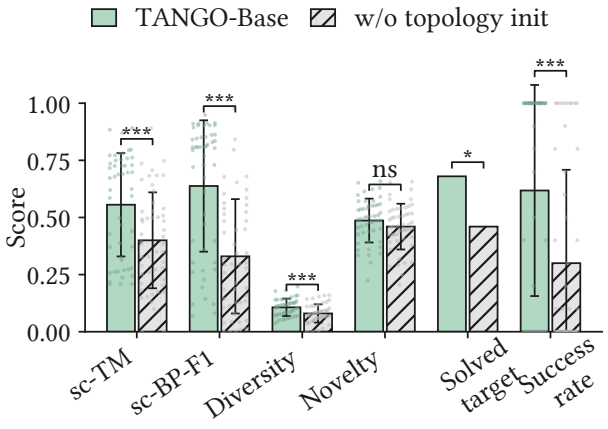

Figure 5: Effect of topology-aware initialization. The ablated variant removes Stage-I topology pretraining and trains the geometry-conditioned policy from scratch, while keeping the structural-unit action space and model architecture unchanged.

5.4 TANGO generates non-native high-quality designs

We next examine whether high fold-back scores are achieved by recovering native-like sequences. Figure 6(a) reports sequence novelty, defined as the normalized Hamming distance between a designed sequence and the native sequence of the target. Figure 6(b) and Figure 6(c) report the corresponding sc-TM and co-design success rate, respectively. TANGO achieves the strongest combination of novelty and design quality among the compared methods: its generated sequences remain far from the native target sequences while maintaining the highest co-design success rate.

This result suggests that TANGO can identify non-native sequence solutions that preserve both target topology and tertiary geometry. However, distance-to-native alone does not rule out all forms of memorization or predictor-specific sequence bias; we therefore interpret this analysis as evidence of non-native design rather than as definitive proof of memorization-free generation.

5.5 Robustness across oracles and architectures

We finally stress-test whether the main conclusions depend on a single fold-back oracle or architecture choice.

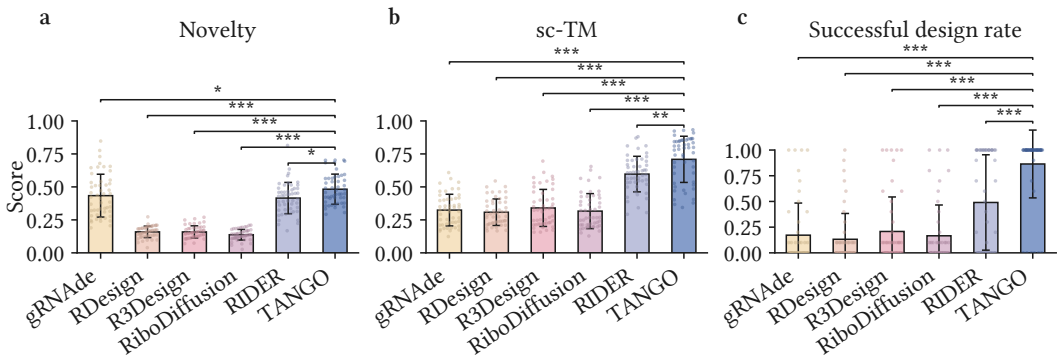

Figure 6: Sequence novelty and design quality across methods. Novelty is measured as normalized Hamming distance to the native target sequence. Design quality is reported by sc-TM and co-design success rate, where success requires both topology and geometry thresholds.

First, Figure 7 evaluates the same TANGO-generated sequences under both RhoFold and AlphaFold3. The sc-TM and sc-BP-F1 scores remain comparable across the two predictors, suggesting that TANGO designs are not obvious artifacts of one fold-back oracle. This remains predictor-based validation rather than experimental folding validation.

Second, Appendix D.4 ablates the multi-agent architecture. Compared with a single-agent variant, the multi-agent policy improves both topology and geometry metrics, supporting the use of topology-derived structural units and parallel cooperative generation for RNA co-design.

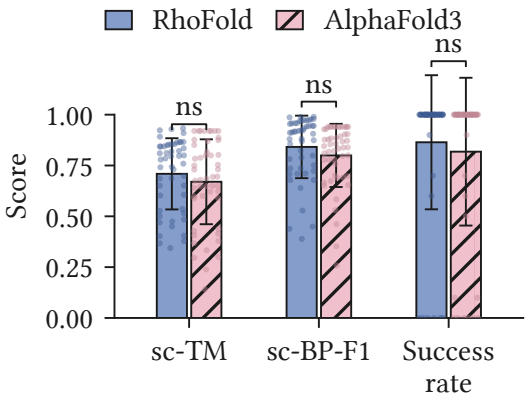

Figure 7: Cross-oracle evaluation of the same TANGO designs using RhoFold and AlphaFold3. No AlphaFold3-based refinement is performed.

6 Conclusion

This work argues that RNA 3D inverse design should not be treated as a geometry-only fold-back problem. Because RNA tertiary structure is scaffolded by base-pairing topology, useful designs should recover both the target 3D fold and the target secondary topology. We operationalize this view through TANGO, a curriculum MARL framework that designs topology-derived structural units, transfers a secondary-structure policy to geometry-conditioned co-design, and uses frontier-guided fine-tuning to improve target-wise trade-offs among topology fidelity, geometry fidelity, and sequence diversity. Experiments demonstrated that TANGO improves fold-back geometry, topology recovery, and joint co-design success over recent 3D RNA design baselines. Future work could extend TANGO to more complex biological settings and experimentally validate selected designs.

Acknowledgments

This work was supported by the UKRI Future Leaders Fellowship under grants MR/S017062/1 and MR/X011135/1; the National Natural Science Foundation of China under grants 62376056, 62076056 and 62625312; the Science and Technology Innovation Program of Hunan Province under Grant 2024RC1011; the Major Basic Research Projects in Hunan Province under Grant 2026JC0001; the Isambard AI Open Access (AIOA - 01 - 2026-06-04 - EvoRNA: Predictive Scaling of Evolution-Aware and Multimodal RNA Foundation Models for Crop Resilience); the Modal compute grant; the Royal Society Faraday Discovery Fellowship (FDF/S2/251014); BBSRC Transformative Research Technologies (UKRI1875); the Royal Society International Exchanges Award (IES/R3/243136); the

Kan Tong Po Fellowship (KTP/R1/231017); the Amazon Research Award; and the Alan Turing Fellowship.

References

[1] Scott A. Strobel and Jesse C. Cochrane. RNA catalysis: Ribozymes, ribosomes, and riboswitches. *Curr. Opin. Chem. Biol.*, 11(6):636–643, 2007.

[2] Alexander Serganov and Evgeny Nudler. A decade of riboswitches. *Cell*, 152(1–2):17–24, 2013.

[3] Jinsong Zhang, Yuhan Fei, Lei Sun, and Qiangfeng Cliff Zhang. Advances and opportunities in RNA structure experimental determination and computational modeling. *Nat. Methods*, 19(10): 1193–1207, 2022.

[4] Alexander Churkin, Matan Drory Retwitzer, Vladimir Reinharz, Yann Ponty, Jérôme Waldispühl, and Danny Barash. Design of RNAs: Comparing programs for inverse RNA folding. *Brief. Bioinform.*, 19(2):350–358, 2018.

[5] Cheng Tan, Yijie Zhang, Zhangyang Gao, Bozhen Hu, Siyuan Li, Zicheng Liu, and Stan Z. Li. RDesign: Hierarchical data-efficient representation learning for tertiary structure-based RNA design. In *Proc. International Conference on Learning Representations*, pages 10304–10333, 2024.

[6] Chaitanya K. Joshi, Arian R. Jamasb, Ramon Viñas, Charles Harris, Simon V. Mathis, Alex Morehead, Rishabh Anand, and Pietro Liò. gRNAde: Geometric deep learning for 3D RNA inverse design. In *Proc. International Conference on Learning Representations*, pages 10166–10191, 2025.

[7] Felix Wong, Dongchen He, Aarti Krishnan, Liang Hong, Alexander Z. Wang, Jiuming Wang, Zhihang Hu, Satotaka Omori, Alicia Li, Jiahua Rao, Qinze Yu, Wengong Jin, Tianqing Zhang, Katherine Ilia, Jack X. Chen, Shuangjia Zheng, Irwin King, Yu Li, and James J. Collins. Deep generative design of RNA aptamers using structural predictions. *Nat. Comput. Sci.*, 4(11): 829–839, 2024.

[8] Han Huang, Ziqian Lin, Dongchen He, Liang Hong, and Yu Li. RiboDiffusion: Tertiary structure-based RNA inverse folding with generative diffusion models. *Bioinformatics*, 40 (Supplement 1):i347–i356, 2024.

[9] Tianmeng Hu, Yongzheng Cui, Biao Luo, and Ke Li. RIDER: 3D RNA inverse design with reinforcement learning-guided diffusion. In *Proc. International Conference on Learning Representations*, 2026.

[10] Wanli Ma, Tianmeng Hu, and Ke Li. REFLEX: RNA ensemble generation via flexibility-calibrated stochastic bridge. *arXiv preprint arXiv:2610.26261*, 2026.

[11] Philippe Brion and Eric Westhof. Hierarchy and dynamics of RNA folding. *Annu. Rev. Biophys. Biomol. Struct.*, 26(1):113–137, 1997.

[12] Samuel E. Butcher and Anna Marie Pyle. The molecular interactions that stabilize RNA tertiary structure: RNA motifs, patterns, and networks. *Acc. Chem. Res.*, 44(12):1302–1311, 2011.

[13] Maximillian H. Bailor, Xiaoyan Sun, and Hashim M. Al-Hashimi. Topology links RNA secondary structure with global conformation, dynamics, and adaptation. *Science*, 327(5962): 202–206, 2010.

[14] Anthony M. Mustoe, Charles L. Brooks III, and Hashim M. Al-Hashimi. Topological constraints are major determinants of tRNA tertiary structure and dynamics and provide basis for tertiary folding cooperativity. *Nucleic Acids Res.*, 42(18):11792–11804, 2014.

[15] Wayne K. Dawson, Maciej Maciejczyk, Elzbieta J. Jankowska, and Janusz M. Bujnicki. Coarse-grained modeling of RNA 3D structure. *Methods*, 103:138–156, 2016.

[16] Xiang-Jun Lu, Harmen J. Bussemaker, and Wilma K. Olson. DSSR: An integrated software tool for dissecting the spatial structure of RNA. *Nucleic Acids Res.*, 43(21):e142, 2015.

[17] Runze Ma, Zhongyue Zhang, Zichen Wang, Chenqing Hua, Jiahua Rao, Zhuomin Zhou, and Shuangjia Zheng. RiboFlow: Conditional de novo RNA co-design via synergistic flow matching. In *Advances in Neural Information Processing Systems*, volume 38, pages 69799–69831, 2025.

[18] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron O. Dror. Learning from protein structure with geometric vector perceptrons. In *Proc. International Conference on Learning Representations*, 2021.

[19] D. M. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A survey of multi-objective sequential decision-making. *J. Artif. Intell. Res.*, 48:67–113, 2013.

[20] Tianmeng Hu, Biao Luo, Chunhua Yang, and Tingwen Huang. MO-MIX: Multi-objective multi-agent cooperative decision-making with deep reinforcement learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(10):12098–12112, 2023.

[21] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[22] Tianmeng Hu and Biao Luo. PA2D-MORL: Pareto ascent directional decomposition based multi-objective reinforcement learning. In *Proc. AAAI Conf. Artif. Intell.*, volume 38, pages 12547–12555, 2024.

[23] Conor F. Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M. Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai A. Irissappane, Patrick Mannion, Ann Nowé, Gabriel Ramos, Marcello Restelli, Peter Vamplew, and Diederik M. Roijers. A practical guide to multi-objective reinforcement learning and planning. *Auton. Agents Multi-Agent Syst.*, 36(1):26, 2022.

[24] Tianmeng Hu, Biao Luo, and Ke Li. Human-in-the-loop policy optimization for preference-based multi-objective reinforcement learning. In *Proc. International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 45520–45549, 2026.

[25] 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 Res.*, 28 (1):235–242, 2000.

[26] Tao Shen, Zhihang Hu, Siqi Sun, Di Liu, Felix Wong, Jiuming Wang, Jiayang Chen, Yixuan Wang, Liang Hong, Jin Xiao, Liangzhen Zheng, Tejas Krishnamoorthi, Irwin King, Sheng Wang, Peng Yin, James J. Collins, and Yu Li. Accurate RNA 3D structure prediction using a language model-based deep learning approach. *Nat. Methods*, 21(12):2287–2298, 2024.

[27] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, 2024.

[28] Chengxin Zhang, Morgan Shine, Anna Marie Pyle, and Yang Zhang. US-align: Universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nat. Methods*, 19(9):1109–1115, 2022.

[29] Cheng Tan, Yijie Zhang, Zhangyang Gao, Hanqun Cao, Siyuan Li, Siqi Ma, Mathieu Blanchette, and Stan Z. Li. R3Design: Deep tertiary structure-based RNA sequence design and beyond. *Brief. Bioinform.*, 26(1):bbae682, 2025.

[30] I. L. Hofacker, W. Fontana, P. F. Stadler, L. S. Bonhoeffer, M. Tacker, and P. Schuster. Fast folding and comparison of RNA secondary structures. *Monatsh. Chem.*, 125(2):167–188, 1994.

[31] Mirela Andronescu, Anthony P. Fejes, Frank Hutter, Holger H. Hoos, and Anne Condon. A new algorithm for RNA secondary structure design. *J. Mol. Biol.*, 336(3):607–624, 2004.

[32] Anke Busch and Rolf Backofen. INFO-RNA—a fast approach to inverse RNA folding. *Bioinformatics*, 22(15):1823–1831, 2006.

[33] Robert Kleinkauf, Martin Mann, and Rolf Backofen. antaRNA: Ant colony-based RNA sequence design. *Bioinformatics*, 31(19):3114–3121, 2015.

[34] Akito Taneda. Multi-objective genetic algorithm for pseudoknotted RNA sequence design. *Front. Genet.*, 3:36, 2012.

[35] Ali Esmaili-Taheri and Mohammad Ganjtabesh. ERD: A fast and reliable tool for RNA design including constraints. *BMC Bioinform.*, 16(1):20, 2015.

[36] Nono S. C. Merleau and Matteo Smerlak. aRNAque: An evolutionary algorithm for inverse pseudoknotted RNA folding inspired by Lévy flights. *BMC Bioinform.*, 23(1):335, 2022.

[37] Tianshuo Zhou, Ning Dai, Sizhen Li, Max Ward, David H. Mathews, and Liang Huang. RNA design via structure-aware multifrontier ensemble optimization. *Bioinformatics*, 39(Supplement 1):i563–i571, 2023.

[38] Peter Eastman, Jade Shi, Bharath Ramsundar, and Vijay S. Pande. Solving the RNA design problem with reinforcement learning. *PLOS Comput. Biol.*, 14(6):e1006176, 2018.

[39] Frederic Runge, Danny Stoll, Stefan Falkner, and Frank Hutter. Learning to design RNA. In *Proc. International Conference on Learning Representations*, 2019.

[40] Frederic Runge, Jörg Franke, Daniel Fertmann, Rolf Backofen, and Frank Hutter. Partial RNA design. *Bioinformatics*, 40(Supplement 1):i437–i445, 2024.

[41] Tianmeng Hu, Biao Luo, and Ke Li. CocoRNA: Collective RNA design with cooperative multi-agent reinforcement learning. In *Proc. International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 45495–45519, 2026.

[42] Julia Koehler Leman, Brian D. Weitzner, Steven M. Lewis, Jared Adolf-Bryfogle, Nawsad Alam, Rebecca F. Alford, Melanie Aprahamian, David Baker, Kyle A. Barlow, Patrick Barth, et al. Macromolecular modeling and design in Rosetta: Recent methods and frameworks. *Nat. Methods*, 17(7):665–680, 2020.

[43] Rubo Wang, Xingyu Gao, and Peilin Zhao. Geometric algebra-enhanced bayesian flow network for RNA inverse design. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025.

[44] Ronny Lorenz, Stephan H. Bernhart, Christian Höner zu Siederdissen, Hakim Tafer, Christoph Flamm, Peter F. Stadler, and Ivo L. Hofacker. ViennaRNA package 2.0. *Algorithms Mol. Biol.*, 6(1):26, 2011.

[45] Stanislav Bellaousov, Jessica S. Reuter, Matthew G. Seetin, and David H. Mathews. RNAstructure: Web servers for RNA secondary structure prediction and analysis. *Nucleic Acids Res.*, 41 (W1):W471–W474, 2013.

[46] Jaswinder Singh, Jack Hanson, Kuldip Paliwal, and Yaoqi Zhou. RNA secondary structure prediction using an ensemble of two-dimensional deep neural networks and transfer learning. *Nat. Commun.*, 10(1):5407, 2019.

[47] Laiyi Fu, Yingxin Cao, Jie Wu, Qinke Peng, Qing Nie, and Xiaohui Xie. UFold: Fast and accurate RNA secondary structure prediction with deep learning. *Nucleic Acids Res.*, 50(3):e14, 2022.

[48] Shujun He, Rui Huang, Jill Townley, Rachael C. Kretsch, Thomas G. Karagianes, David B. T. Cox, Hamish Blair, Dmitry Penzar, Valeriy Vyaltsev, Elizaveta Aristova, et al. Ribonanza: Deep learning of RNA structure through dual crowdsourcing. *bioRxiv*, 2024. Preprint 2024.02.24.581671.

[49] Heng Yang, Renzhi Chen, and Ke Li. Bridging sequence-structure alignment in RNA foundation models. In *Proc. AAAI Conf. Artif. Intell.*, pages 21929–21937, 2025.

[50] Ning Wang, Jiang Bian, Yuchen Li, Xuhong Li, Shahid Mumtaz, Linghe Kong, and Haoyi Xiong. Multi-purpose RNA language modelling with motif-aware pretraining and type-guided fine-tuning. *Nat. Mach. Intell.*, 6(5):548–557, 2024.

[51] Rafael Josip Penić, Tin Vlašić, Roland G. Huber, Yue Wan, and Mile Šikić. RiNALMo: General-purpose RNA language models can generalize well on structure prediction tasks. *Nat. Commun.*, 16(1):5671, 2025.

[52] Jiayang Chen, Zhihang Hu, Siqi Sun, Qingxiong Tan, Yixuan Wang, Qinze Yu, Licheng Zong, Liang Hong, Jin Xiao, Tao Shen, Irwin King, and Yu Li. Interpretable RNA foundation model from unannotated data for highly accurate RNA structure and function predictions. *CoRR*, abs/2204.00300, 2022.

[53] Yuan Li, Heng Yang, Renzhi Chen, and Ke Li. Understanding transfer learning of RNA foundation models on downstream tasks. In *Proc. International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 71724–71744, 2026.

[54] Mariusz Popenda, 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 Res.*, 40(14):e112, 2012.

[55] 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. *Nat. Commun.*, 14(1):7266, 2023.

[56] 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. *Nat. Commun.*, 14(1):5745, 2023.

[57] Sumit Tarafder and Debswapna Bhattacharya. RNAbpFlow: Base pair-augmented SE(3) flow matching for conditional RNA 3D structure generation. *Nat. Methods*, 23(7):1349–1358, 2026.

[58] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27 (4):819–840, 2002.

[59] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, volume 30, pages 6379–6390, 2017.

[60] Tianmeng Hu, Yongzheng Cui, Rui Tang, Biao Luo, and Ke Li. Beyond monotonicity: Revisiting factorization principles in multi-agent Q-learning. In *Proc. AAAI Conf. Artif. Intell.*, volume 40, pages 21876–21884, 2026.

[61] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proc. International Conference on Autonomous Agents and Multiagent Systems*, pages 2085–2087, 2018.

[62] Yongzheng Cui, Biao Luo, Rui Tang, Chunhua Yang, and Weihua Gui. QMIP: Enhancing multiagent reinforcement learning coordination through macro intention prediction. *IEEE Trans. Artif. Intell.*, 7(6):3346–3360, 2026.

[63] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proc. International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4295–4304, 2018.

[64] Rui Tang, Biao Luo, Yongzheng Cui, Bing-Chuan Wang, Wenfeng Hu, Wu Song, and Chunhua Yang. HMU: Heuristic memory utilization for multi-agent reinforcement learning. *IEEE Trans. Cogn. Develop. Syst.*, pages 1–14, 2026. Early access.

[65] Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proc. AAAI Conf. Artif. Intell.*, volume 32, pages 2974–2982, 2018.

[66] Rui Tang, Biao Luo, and Yongzheng Cui. Autonomous partner selection for cooperative multi-agent reinforcement learning. In *Proc. AAAI Conf. Artif. Intell.*, volume 40, pages 25840–25848, 2026.

[67] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of PPO in cooperative multi-agent games. In *Advances in Neural Information Processing Systems*, volume 35, pages 24611–24624, 2022.

[68] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *Proc. International Conference on Learning Representations*, 2016.

[69] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.

[70] Zhangyang Gao, Cheng Tan, Yijie Zhang, Xingran Chen, Lirong Wu, and Stan Z. Li. ProteinInvBench: Benchmarking protein inverse folding on diverse tasks, models, and metrics. In *Advances in Neural Information Processing Systems*, volume 36, pages 68207–68220, 2023.

[71] Shiting Wang, Yingjie Zou, and Ke Li. Separating common and unique directions for model merging. *arXiv preprint arXiv:2611.04108*, 2026.

[72] Mingyu Huang, Shasha Zhou, and Ke Li. Augmenting biological fitness prediction benchmarks with landscapes features from GraphFLA. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

[73] Mandy Yu Theng Lim, Chermaine Tan, Charannya Sozheesvari Subhramanyam, Shi Jie Teo, Louis DeFalco, Samuel Kevin Pasaribu, Chong Hui Koh, Dheeraj Rayamajhi, Jieying Chi, Shengnan Li, et al. A programmable ribozyme for RNA signal transduction. *Nature Communications*, 17:1428, 2026.

[74] James M Robson and Alexander A Green. Toehold-VISTA: a machine learning approach to decipher programmable RNA sensor-target interactions. *Nucleic Acids Res.*, 54(4):gkag097, 2026.

[75] Shasha Zhou, Mingyu Huang, and Ke Li. Position: Genomic model research must move beyond anecdotal evaluation of interpretability methods. In *Proc. International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 172974–172999, 2026.

A Related Work and Positioning

RNA inverse design. Computational RNA inverse design seeks sequences that satisfy specified structural objectives. Early work primarily studied secondary-structure inverse folding, where the target is a dot-bracket topology. Local-search methods [30–32] iteratively modify sequences using folding feedback, metaheuristic approaches such as ant-colony optimization [33] and evolutionary algorithms [34–37] explore large combinatorial spaces, and reinforcement-learning methods [38–40] learn sequence-construction policies. `CocoRNA` [41] applies cooperative multi-agent reinforcement learning to secondary-structure inverse design by distributing sequence construction among agents trained with a centralized critic. Many widely used secondary-design methods are built for pseudoknot-free structures, although methods such as `antaRNA` [33] and `aRNAque` [36] support broader topology classes. These methods directly optimize secondary topology, but do not by themselves enforce the tertiary geometry required by many functional RNA folds.

Recent work has shifted toward tertiary-structure-based RNA design. Rosetta provides physics-based fixed-backbone RNA design through energy modeling [42]. Deep learning methods further condition sequence generation on RNA 3D structure: graph neural networks [5, 6, 29] encode geometric structure, diffusion models such as `RiboDiffusion` [8] learn conditional sequence distributions over 3D backbones, and Bayesian flow networks [43] provide another generative formulation. `RIDER` [9] is closest to our use of structural optimization, as it fine-tunes a generative model with reinforcement learning to improve 3D fold-back similarity.

Our work is complementary to these approaches. Existing 3D RNA design methods primarily condition on tertiary geometry and are often evaluated by native sequence recovery or tertiary fold-back fidelity. Some methods incorporate secondary-structure information as priors, constraints, or repair mechanisms; for example, `RDesign` uses extracted secondary structures with base-pair information to guide design [5]. However, topology fidelity has not been systematically formulated, optimized, and evaluated as a first-class objective jointly with tertiary-structure fidelity. `TANGO` addresses this gap by using the target topology to define structural-unit actions, initializing the policy on secondary-structure design, and refining candidates with multi-objective topology-and-geometry rewards.

RNA structure prediction as design oracles. RNA structure prediction is the forward problem complementary to inverse design. Secondary-structure predictors include thermodynamic models such as `ViennaRNA` [44] and `RNAstructure` [45], as well as deep learning models [46–48]. RNA foundation models [49–52] provide pretrained representations for structure-related tasks [53]. Tertiary-structure prediction has advanced from fragment assembly methods such as `RNAComposer` [54] to deep geometric predictors including `trRosettaRNA` [55], `DRFold` [56], `RhoFold` [26], and `AlphaFold3` [27]. Related generative models such as `RNAbpFlow` [57] incorporate base-pairing conditions for RNA 3D structure generation. `REFLEX` [10] generates sequence-conditioned RNA conformational ensembles using a flexibility-calibrated stochastic bridge.

In this work, structure predictors are used as fold-back oracles for evaluating designed sequences, not as the proposed design model itself. Because predictor-based evaluation can introduce oracle bias, our experiments report cross-oracle validation with both `RhoFold` and `AlphaFold3`.

B Algorithm Pseudocode

This section provides pseudocode for the main algorithmic components of `TANGO`. We include three procedures: Algorithm 1 constructs structural units from the target topology, Algorithm 2 performs multi-agent sequence generation over these units, and Algorithm 3 summarizes the shared MARL training framework used across topology-aware initialization, topology-and-geometry co-design, and frontier-guided fine-tuning.

B.1 Structural-unit Cooperative Generation

TANGO first converts the target topology into an ordered list of structural units, and then assigns these units to multiple agents for cooperative sequence generation. Algorithm 1 gives the pseudocode for constructing structural units from a target pair map.

Algorithm 1 TANGO: Constructing Structural Units

Require: Pair map $m^\star$ for target topology $\mathcal{T}^\star$, with $m^\star_p = q$ for paired residues and $m^\star_p = 0$ for unpaired residues

Ensure: Ordered structural units $\mathcal{U}^\star = (u_1, \ldots, u_M)$

```
Initialize empty unit list U* and empty assigned set V
for p = 1 to L do
    if p ∈ V then
        Continue
    end if
    if m*_p = q and q > 0 then
        if p < q then
            Create paired unit u = (p, q)
            Store paired-unit metadata
            Append u to U* and add p, q to V
        end if
    else
        Create single-position unit u = (p)
        Store single-unit metadata
        Append u to U* and add p to V
    end if
end for

return U*
```

The resulting unit list is used as the action schedule for multi-agent rollout. As shown in Algorithm 2, each agent is responsible for generating the sequence for its assigned units, and the joint action at each step consists of the nucleotide choices made by all active agents. The rollout proceeds until all units are assigned, at which point a complete sequence is returned for reward evaluation.

B.2 Multi-Agent Reinforcement Learning Framework

Let $s \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}\}$ denote the training stage. We use $\mathcal{R}_s$ to denote the stage-specific terminal reward: the topology-training reward for Stage I, the topology-and-geometry co-design reward for Stage II, and the frontier-guided reward for Stage III. All stages share the same structural-unit rollout interface and MAPPO update, while differing in the conditioning information and terminal reward construction. Algorithm 3 summarizes the overall training framework for TANGO.

In Stage I, $\mathcal{R}_s$ is the topology-training reward. In Stage II, $\mathcal{R}_s$ is the topology-and-geometry co-design reward. In Stage III, $\mathcal{R}_s$ is the frontier-guided reward computed from the target-wise archive.

C Reinforcement Learning Formulation

C.1 Decentralized partially observable Markov decision process

We model structural-unit generation as a fully cooperative decentralized partially observable Markov decision process (Dec-POMDP) [58]. The process is denoted by

$$\mathcal{M} = \langle \mathcal{N}, \mathcal{S}, \{\mathcal{O}_i\}_{i=1}^N, \{\mathcal{A}_i\}_{i=1}^N, P, O, r, \gamma \rangle. \tag{20}$$

The tuple components are defined as follows:

- $\mathcal{N} = \{1, \ldots, N\}$ is the set of cooperative agents;
- $\mathcal{S}$ is the global state space;
- $\mathcal{O}_i$ is the local observation space of agent i;
- $\mathcal{A}_i$ is the action space of agent i;
- $P(s' \mid s, \boldsymbol{a})$ is the transition function under joint action $\boldsymbol{a} = (a^1, \ldots, a^N)$;

Algorithm 2 TANGO: Multi-Agent Structural-Unit Rollout

Require: Structural units $\mathcal{U}^\star$, actor π_θ, number of agents N

Ensure: Completed sequence x_T and rollout trajectory τ

Split $\mathcal{U}^\star$ into N contiguous agent segments
Initialize partial sequence x_0, assignment mask, and one cursor per agent
Initialize empty trajectory buffer τ
for $t = 0$ **to** $T - 1$ **do**
  Initialize empty joint action map $\boldsymbol{a}_t$
  for each agent $i = 1, \ldots, N$ **do**
    if agent i has no remaining unit **then**
      Mark agent i inactive
    else
      Build local observation o_t^i for the unit at agent i's cursor
      if 3D conditions are available **then**
        Augment o_t^i with unit-level 3D information
      end if
      Sample $a_t^i \sim \pi_\theta(\cdot \mid o_t^i)$ using the corresponding unit-type head
      Store a_t^i in $\boldsymbol{a}_t$ and record rollout information in τ
    end if
  end for
  Write active unit assignments from $\boldsymbol{a}_t$ into the partial sequence
  Advance cursors for active agents and update the assignment mask
  if all units are assigned **then**
    Break
  end if
end for

return completed sequence x_T and trajectory τ

Algorithm 3 TANGO: Multi-Agent Reinforcement Learning Framework

Require: Training stage s, actor π_θ, centralized critic V_ψ, MAPPO hyperparameters

Require: Training targets $\mathcal{D}$, structure oracles when required, and stage-specific reward $\mathcal{R}_s$

Ensure: Updated policy parameters θ

for each training update **do**
  Initialize rollout buffer $\mathcal{B}$
  for each sampled target in $\mathcal{D}$ **do**
    Obtain target topology $\mathcal{T}^\star$ and construct $\mathcal{U}^\star$ using Algorithm 1
    Build stage-specific conditioning information, including 3D conditions when required
    Generate (x_T, τ) using Algorithm 2
    Compute terminal reward $r_T = \mathcal{R}_s(x_T; \mathcal{G}^\star, \mathcal{T}^\star)$
    Assign shared terminal reward r_T to active agent transitions in τ
    Store τ in $\mathcal{B}$
  end for
  Compute returns and GAE advantages using centralized critic V_ψ
  Update π_θ and V_ψ with MAPPO
end for

return θ

- $O(o^1, \ldots, o^N \mid s', \boldsymbol{a})$ is the observation function;
- $r(s, \boldsymbol{a})$ is the shared reward function;
- $\gamma \in [0, 1)$ is the discount factor.

At step t, let $\mathcal{I}_t \subseteq \mathcal{N}$ denote the active agents whose assigned unit segments have not been completed. Each active agent $i \in \mathcal{I}_t$ receives a local observation o_t^i and samples an action from the shared actor:

$$a_t^i \sim \pi_\theta(\cdot \mid o_t^i). \tag{21}$$

The joint policy factorizes over active agents as

$$\pi_\theta(\boldsymbol{a}_t \mid \boldsymbol{o}_t) = \prod_{i \in \mathcal{I}_t} \pi_\theta(a_t^i \mid o_t^i), \tag{22}$$

where $\boldsymbol{a}_t = \{a_t^i\}_{i \in \mathcal{I}_t}$. The environment writes the selected single-position or paired-unit assignments into the partial sequence and then advances the active cursors. The transition is deterministic given the sampled joint action. All active agents receive the same reward r_t.

The objective is to learn a shared policy that maximizes the expected discounted return:

$$J(\pi_\theta) = \mathbb{E}_{\pi_\theta}\left[\sum_{t=0}^{T-1} \gamma^t r_t\right]. \tag{23}$$

In TANGO, rewards are terminal at the end of sequence generation. Stage I uses the topology-training reward, Stage II uses the topology-and-geometry co-design reward, and Stage III uses the frontier-guided reward. The rollout dynamics remain the same across stages; only the conditioning information and terminal reward construction change.

The global design state includes the target topology, the ordered structural-unit list, the current partial sequence, the assignment mask, all agent cursors, and the rollout progress. In Stage II and Stage III, the state also includes target 3D information. Each agent observes the unit at its cursor, the unit type, the current assignment state of that unit, and a local topology context. When 3D conditioning is used, the observation is augmented with the current unit's 3D embedding and a pooled local 3D context. The episode ends when all structural units are assigned.

C.2 Centralized training and decentralized execution

TANGO uses centralized training and decentralized execution (CTDE), a standard paradigm in cooperative MARL [59–67]. The actor is shared across agents and uses unit-type-specific output heads for single-position and paired units. During execution, each agent acts only from its local observation. During training, a centralized critic receives a global critic state s_t^{crit} constructed from the partial design state and outputs $V_\phi(s_t^{\mathrm{crit}})$. This separation matches the RNA design problem: actions are local to structural units, but sequence quality is evaluated by a global structural objective.

For a rollout, we compute temporal-difference residuals and generalized advantage estimates (GAE) [68] as

$$\delta_t = r_t + \gamma V_\phi(s_{t+1}^{\mathrm{crit}}) - V_\phi(s_t^{\mathrm{crit}}), \qquad A_t = \sum_{\ell=0}^{T-t-1} (\gamma\lambda_{\mathrm{GAE}})^\ell \delta_{t+\ell}. \tag{24}$$

Because the reward is shared, the same advantage A_t is assigned to all active agent transitions at step t.

C.3 MAPPO objective

Policy updates use the clipped MAPPO objective [67, 69]. Let θ_{old} denote the policy parameters before the update. For an active agent transition, the probability ratio is

$$\rho_t^i(\theta) = \frac{\pi_\theta(a_t^i \mid o_t^i)}{\pi_{\theta_{\mathrm{old}}}(a_t^i \mid o_t^i)}. \tag{25}$$

The actor objective is

$$\mathcal{L}_{\mathrm{actor}}(\theta) = \mathbb{E}_{t, i \in \mathcal{I}_t}\left[\min\left(\rho_t^i(\theta)A_t, \mathrm{clip}(\rho_t^i(\theta), 1-\epsilon, 1+\epsilon)A_t\right)\right]. \tag{26}$$

The critic is trained to predict the target return $\hat{R}_t$:

$$\mathcal{L}_{\mathrm{value}}(\phi) = \mathbb{E}_t\left[\left(V_\phi(s_t^{\mathrm{crit}}) - \hat{R}_t\right)^2\right]. \tag{27}$$

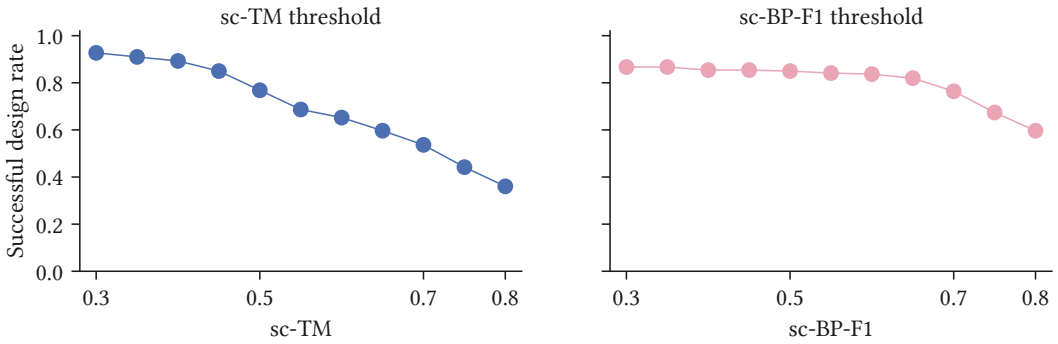

Figure A1: Sensitivity of TANGO to the joint success criterion. Left: successful design rate when varying the sc-TM threshold while fixing sc-BP-F1 > 0.50. Right: successful design rate when varying the sc-BP-F1 threshold while fixing sc-TM > 0.45.

The total loss minimized during training is

$$\mathcal{L}(\theta, \phi) = -\mathcal{L}_{\text{actor}}(\theta) + c_v \mathcal{L}_{\text{value}}(\phi) - c_H \mathcal{H}(\pi_\theta). \tag{28}$$

The same MAPPO objective is used across topology-aware initialization, topology-and-geometry co-design, and frontier-guided fine-tuning.

C.4 Effective horizon

Let M be the number of structural units in the target topology. A single-agent unit-by-unit policy requires M environment steps to complete the sequence. In contrast, TANGO splits the ordered unit list into N contiguous segments and lets agents act jointly. The rollout length is therefore the maximum segment length, approximately $\lceil M/N \rceil$ when the split is balanced. This shorter effective horizon reduces delayed rewards and improves temporal credit assignment, which is especially useful because structural rewards are obtained only after a complete sequence is generated.

D Additional Results

In this section, we present additional experiments under various settings to further demonstrate the effectiveness and robustness of TANGO. Specifically, we include:

- sensitivity analysis of the success thresholds;
- comparison between Top-10 sampling and frontier-guided fine-tuning;
- comparison between frontier-guided fine-tuning and fixed weighted rewards;
- ablation of the cooperative multi-agent formulation.

D.1 Threshold Sensitivity

A TM-score above 0.45 is commonly used to indicate that two RNA structures share the same global topology [28]. We therefore use sc-TM > 0.45 as the tertiary-structure threshold in the main experiments, together with sc-BP-F1 > 0.50 to require secondary-structure consistency. Figure A1 further evaluates how sensitive the successful design rate is to these thresholds. As expected, the success rate decreases as either requirement becomes stricter. When the sc-TM threshold is increased from 0.30 to 0.80 with sc-BP-F1 fixed at 0.50, the successful design rate decreases by 61.3% relative to the loosest threshold setting. When the sc-BP-F1 threshold is increased from 0.30 to 0.80 with sc-TM fixed at 0.45, the rate decreases by 31.0%. These results show that joint success is sensitive to the required levels of both tertiary-structure and secondary-structure fidelity. Nevertheless, TANGO retains a substantial successful design rate under stricter criteria, supporting its robustness in generating high-quality RNA designs.

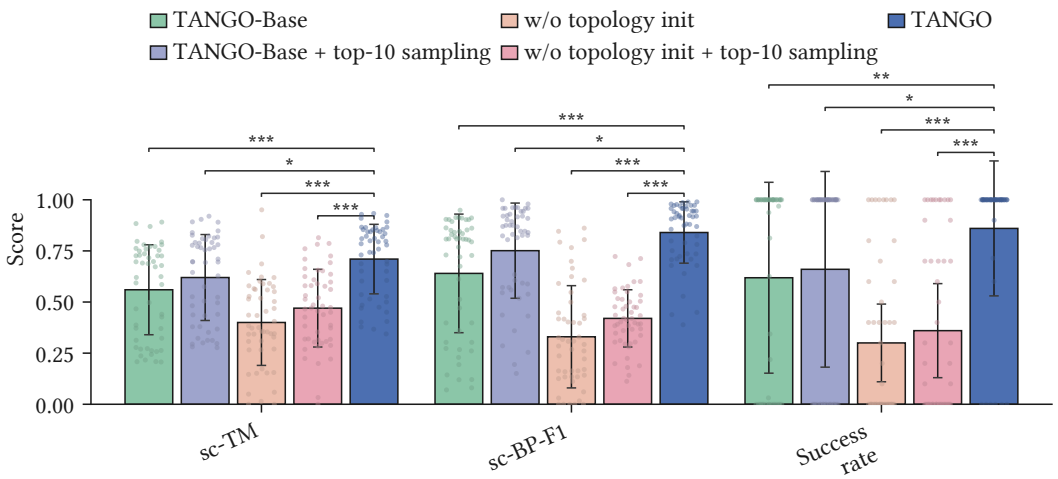

Figure A2: Comparison between Top-10 sampling and frontier-guided fine-tuning. For each target, we sample sequences from 30 random seeds and select the 10 designs with the highest sc-TM for evaluation. We compare raw policies, their Top-10 sampling variants, and the full TANGO model. Bars show mean ± standard deviation across targets.

D.2 Top 10 Sampling versus Frontier-guided Fine-tuning

Figure A2 compares frontier-guided fine-tuning with a Top-10 sampling strategy. We include TANGO-Base, an ablated variant without frontier-guided fine-tuning, and evaluate both its raw sampling performance and its Top-10 sampling performance. For each target, Top-10 sampling selects the 10 designs with the highest sc-TM from 30 sampled sequences. We also evaluate Top-10 sampling for the variant without topology-aware initialization.

The full TANGO model clearly outperforms all Top-10 sampling variants. Compared with TANGO-Base with Top-10 sampling, TANGO achieves relative improvements of 14.5% in sc-TM and 12.0% in sc-BP-F1. Compared with the Top-10 sampling variant without topology-aware initialization, TANGO achieves relative improvements of 51.1% in sc-TM and 100.0% in sc-BP-F1. This shows that the gains from TANGO cannot be explained by selecting better samples from a fixed policy.

Without policy optimization, Top-10 sampling still improves the raw policies, indicating that the sampled designs remain diverse enough for post-hoc selection to find stronger candidates. For TANGO-Base, Top-10 sampling improves sc-TM by 10.7% and sc-BP-F1 by 17.2% over raw sampling. For the variant without topology-aware initialization, it improves sc-TM by 17.5% and sc-BP-F1 by 27.3%. These results show that sampling and reranking can improve candidate selection, but frontier-guided fine-tuning is more effective because it improves the policy to generate higher-quality designs rather than only selecting from existing samples.

D.3 Frontier-guided Fine-tuning versus Weighted Rewards

Figures A3 and A4 compare frontier-guided fine-tuning with fixed weighted reward alternatives. The 2D-heavy setting matches TANGO in sc-BP-F1 (0.84), but achieves lower sc-TM (0.61) and diversity (0.08). Conversely, the 3D-heavy setting achieves slightly higher sc-TM (0.72), but reduces sc-BP-F1 to 0.62. These results show that fixed scalarization can bias optimization toward one structural objective at the expense of others. In contrast, TANGO maintains a stronger balance between topology and geometry, while TANGO-Div further increases diversity from 0.10 to 0.22 with only small decreases in sc-TM and sc-BP-F1. This supports the use of frontier-guided fine-tuning for target-wise multi-objective refinement without relying on manually chosen scalarization weights.

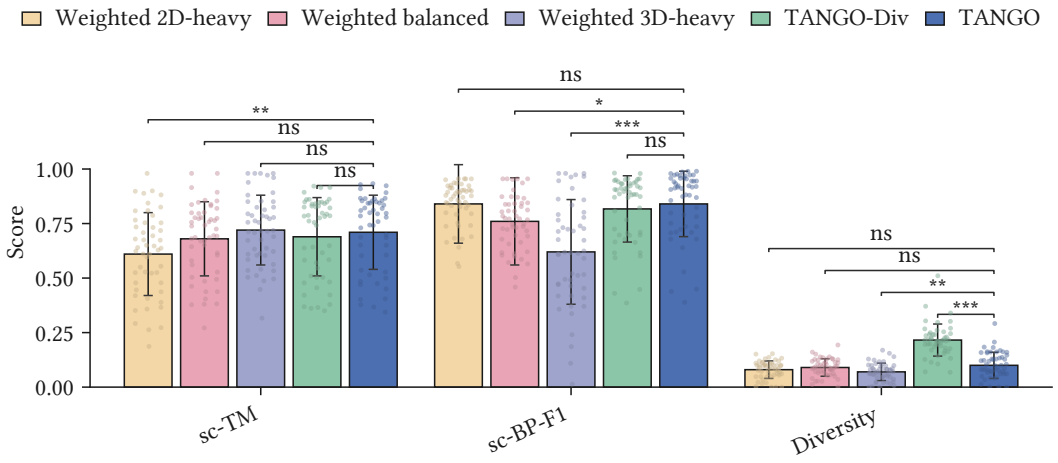

Figure A3: Comparison between frontier-guided fine-tuning and fixed weighted reward alternatives. Three scalarized reward settings with different 2D/3D weights are compared with `TANGO` and `TANGO-Div` on sc-TM, sc-BP-F1, and sequence diversity. Bars show mean ± standard deviation across targets. Significance markers denote * for $p < 0.05$, ** for $p < 0.01$, *** for $p < 0.001$, and `ns` for not significant.

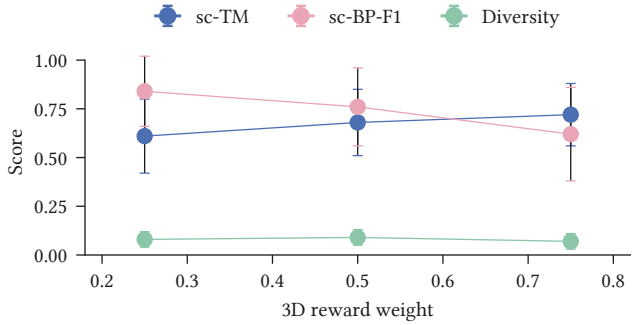

Figure A4: Trade-offs among multiple objectives under fixed weighted rewards. Changing the scalarization weights shifts the balance among 2D fidelity, 3D fidelity, and diversity, whereas frontier-guided fine-tuning directly optimizes the target-wise trade-off frontier.

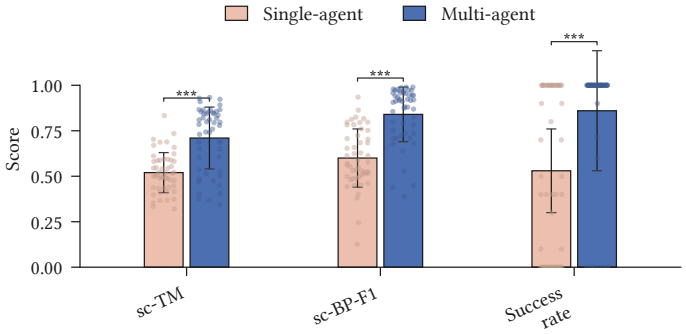

Figure A5: Ablation of the cooperative multi-agent formulation. We compare the default multi-agent TANGO with a single-agent variant under the same training and evaluation protocol. Bars show mean ± standard deviation across held-out targets. Significance markers denote * for $p < 0.05$, ** for $p < 0.01$, *** for $p < 0.001$, and `ns` for not significant.

D.4 Effect of the Cooperative Multi-agent Formulation

Figure A5 evaluates the effect of the cooperative multi-agent formulation by comparing TANGO with a single-agent variant. The multi-agent formulation consistently improves all metrics, increasing sc-TM from 0.52 to 0.71, sc-BP-F1 from 0.60 to 0.84, and success rate from 0.53 to 0.86. These correspond to relative improvements of 36.5%, 40.0%, and 62.3%, respectively. The results show that decomposing RNA sequence construction into coordinated local decisions benefits both tertiary-structure recovery and secondary-structure consistency.

E Model Architecture

E.1 Topology and Unit Encoders

The policy operates on structural units, while the first representation is constructed at the residue level. For residue p, the topology encoder embeds the dot-bracket token b_p and concatenates it with structural metadata $\boldsymbol{m}_p$, including pairing status, normalized position, pairing distance, and pseudoknot-related indicators. The token vocabulary covers unpaired positions, standard parenthesis brackets, and additional bracket symbols used for pseudoknotted base pairs. The concatenated representation is mapped by an MLP to a residue-level topology embedding $\boldsymbol{h}_p^{\mathcal{T}} \in \mathbb{R}^{256}$.

The unit encoder then converts residue embeddings into the structural-unit representation used by all agents. For a single-position unit $u_j = \{p\}$, the input is $\boldsymbol{h}_p^{\mathcal{T}}$ together with unit metadata $\boldsymbol{m}_{u_j}$, such as normalized position and local topology context. For a paired unit $u_j = \{p, q\}$, the input contains $\boldsymbol{h}_p^{\mathcal{T}}$, $\boldsymbol{h}_q^{\mathcal{T}}$, their elementwise product, their absolute difference, and pair-level metadata such as normalized pair distance, normalized endpoints, and pseudoknot/context indicators. Separate MLPs encode single and paired units into a shared 256-dimensional unit space. This design aligns the representation with the action granularity: paired residues are represented and assigned as one coupled decision rather than two independent nucleotide decisions.

A local topology context vector is computed for the current unit by mean pooling unit embeddings in a fixed local window around that unit. This context is used as local information for the actor. The centralized critic receives a separate global summary of the partial design state.

E.2 Shared Actor with Unit-specific Heads

All agents share one actor. At a decision step, the actor input for an active agent is the current unit embedding, the assignment state of that unit, and the local topology context, corresponding to the observation $o_t^i = [\boldsymbol{e}_{j_t^i}, \boldsymbol{q}_{t,j_t^i}, \boldsymbol{c}_{j_t^i}]$ used in the main method. A shared MLP trunk maps this observation to a hidden representation.

The final action distribution is unit-specific. The single-unit head produces logits over $\{\texttt{A}, \texttt{U}, \texttt{G}, \texttt{C}\}$, and the paired-unit head produces logits over $\{\texttt{AU}, \texttt{UA}, \texttt{GC}, \texttt{CG}, \texttt{GU}, \texttt{UG}\}$. Agents with no remaining units are inactive at that step and do not contribute policy transitions.

E.3 Centralized Critic

The critic is centralized during training. In the topology-aware initialization stage, it receives a pooled representation of the global partial design state, summarizing the target topology, current partial sequence, assignment mask, and rollout progress. In the topology-and-geometry co-design and frontier-guided fine-tuning stages, the critic additionally receives a pooled 3D target condition. The actor remains decentralized: at inference time, each active agent only uses its local observation and the shared actor parameters.

E.4 3D Target Encoder

The target tertiary structure is represented with three beads per nucleotide: P, $\mathrm{C4}'$, and the glycosidic nitrogen atom, N1 for pyrimidines and N9 for purines [15]. This coarse-grained representation preserves backbone geometry and approximate base placement while avoiding all-atom inputs.

TANGO uses a GVP-GNN encoder to map the target geometry to residue-level embeddings [18, 6]. The encoder constructs a spatial graph whose nodes are nucleotides and whose edges connect nearby residues in 3D space. Node and edge features contain scalar channels, such as distances and positional encodings, and vector channels, such as backbone directions and relative displacement vectors. Native nucleotide identities are not used as input features, preventing leakage from the target sequence. A GVP-GNN layer updates scalar-vector node features by aggregating messages from spatial neighbors:

$$\begin{aligned}(\boldsymbol{m}_i^{(l)}, \vec{\boldsymbol{m}}_i^{(l)}) &= \mathrm{AGG}_{j\in\mathcal{N}(i)}\, \phi_{\mathrm{msg}}^{(l)}\left((\boldsymbol{s}_i^{(l-1)}, \vec{\boldsymbol{v}}_i^{(l-1)}), (\boldsymbol{s}_j^{(l-1)}, \vec{\boldsymbol{v}}_j^{(l-1)}), (\boldsymbol{e}_{ij}, \vec{\boldsymbol{e}}_{ij})\right), \\ (\boldsymbol{s}_i^{(l)}, \vec{\boldsymbol{v}}_i^{(l)}) &= \phi_{\mathrm{upd}}^{(l)}\left((\boldsymbol{s}_i^{(l-1)}, \vec{\boldsymbol{v}}_i^{(l-1)}), (\boldsymbol{m}_i^{(l)}, \vec{\boldsymbol{m}}_i^{(l)})\right).\end{aligned} \tag{29}$$

Here $\phi_{\mathrm{msg}}^{(l)}$ and $\phi_{\mathrm{upd}}^{(l)}$ are GVP-based message and update functions. Scalar channels provide rotation-invariant quantities for the downstream policy, while vector channels allow the encoder to propagate orientation-dependent geometric information during message passing. The resulting residue-level geometric embeddings are used as conditioning signals for the structural-unit actor.

Residue-level geometric embeddings are pooled to match the structural units. A single-position unit inherits the embedding of its residue. A paired unit uses the mean of the two residue embeddings. We also compute a unit center from the target $\mathrm{C4}'$ coordinates and build a unit-level 3D nearest-neighbor list. The local 3D context for a unit is the mean of the embeddings of its nearest neighboring units, using 15 nearest units in the reported configuration.

E.5 Zero-initialized Geometry Adapters

The topology-and-geometry co-design stage initializes from the topology-aware actor and adds two 3D-conditioned paths. For unit u_j, let $\boldsymbol{g}_j$ denote the unit-level 3D embedding and $\boldsymbol{c}_j^{\mathrm{3D}}$ denote the local 3D context. We define the geometry condition as $\boldsymbol{d}_j^{\mathrm{3D}} = [\boldsymbol{g}_j, \boldsymbol{c}_j^{\mathrm{3D}}]$.

The first 3D-conditioned path applies FiLM modulation to the inherited hidden state:

$$\widetilde{\boldsymbol{z}}_j = \left(1 + \gamma(\boldsymbol{d}_j^{\mathrm{3D}})\right) \odot \boldsymbol{z}_j + \beta(\boldsymbol{d}_j^{\mathrm{3D}}), \tag{30}$$

where $\boldsymbol{z}_j$ is the hidden representation produced by the topology-aware base path. The second path directly maps the topology-aware unit state together with the geometry condition through a small 3D-conditioned network. The final action logits are

$$\ell_j = \ell_j^{\mathrm{base}} + \Delta\ell_j^{\mathrm{FiLM}} + \Delta\ell_j^{\mathrm{3D}}. \tag{31}$$

The last linear layers of both 3D-conditioned paths are initialized to zero. Therefore, at the beginning of topology-and-geometry co-design, the transferred policy is close to the topology-aware policy. Geometry conditioning can then gradually adjust topology-compatible sequence decisions toward the target tertiary structure.

F Training Details

The three learning stages share the structural-unit rollout interface and the MAPPO update described in Appendix C. At the beginning of each episode, the target is parsed into ordered structural units, and cooperative agents fill these units until a complete candidate sequence is generated. The learning signal is assigned at the end of the episode. Stage I uses a topology-training reward, Stage II uses a topology-and-geometry co-design reward, and Stage III uses a frontier-guided reward based on target-wise multi-objective improvement.

Across stages, MAPPO uses discount factor $\gamma = 0.99$, GAE parameter $\lambda_{\mathrm{GAE}} = 0.95$, PPO clip range 0.2, value coefficient 0.5, and gradient clipping at norm 5.0. Stage-specific rollout, optimization, and reward settings are summarized below. Protocol choices such as the fold-back oracle, structure-alignment backend, and frontier objective set are described in prose when they define the training interface rather than scalar hyperparameters.

F.1 Stage I Hyperparameters

Stage I trains the topology-aware policy on secondary-structure targets. After the agents generate a complete sequence, the secondary-structure oracle evaluates its agreement with the target secondary

Table A1: Stage I rollout and optimization settings.

Setting	Value
Number of agents	4
Actor learning rate	1×10^{-4}
Critic learning rate	1×10^{-4}
Encoder learning-rate scale	0.3
Optimizer	AdamW
Weight decay	1×10^{-4}
Learning-rate decay floor	0.3
Buffer size	4096
Epochs per update	4
Minibatches per update	4
Entropy coefficient	0.01

Table A2: Stage II rollout, optimization, and co-design reward settings.

Setting	Value
Number of agents	4
Actor learning rate	5×10^{-4}
Critic learning rate	5×10^{-4}
Learning-rate decay floor	0.3
Buffer size	1024
Epochs per update	10
Minibatches per update	2
Entropy coefficient	0.01
λ_{2D}	1.0
λ_{viol}	2.0
τ_{2D}	0.4

structure, and the resulting topology-training reward is shared by all agents. This stage teaches the policy to make base-pair-compatible structural-unit assignments before tertiary-structure conditioning is introduced. Training is performed on a server with one `NVIDIA H100` GPU. The training takes approximately 5 hours, and inference for a single sample takes only a few seconds. The rollout and optimization settings are shown in Table A1.

F.2 Stage II Hyperparameters

Stage II starts from the topology-aware policy and trains a topology-and-geometry co-design policy. For each completed sequence, we predict its tertiary structure with the fold-back oracle, align it to the target structure, and compute the TM-score as the 3D reward. The predicted tertiary structure is also annotated into a secondary structure, which is used to compute the BP-F1 topology reward. Following the main text, the Stage II co-design reward contains three terms: a 3D fidelity term, a 2D fidelity term, and a hinge penalty when the 2D fidelity falls below the target threshold:

$$R_{\mathrm{co}}(x_T; \mathcal{G}^\star, \mathcal{T}^\star) = R_{\mathrm{3D}}(x_T; \mathcal{G}^\star) + \lambda_{\mathrm{2D}} R_{\mathrm{2D}}(x_T; \mathcal{T}^\star) - \lambda_{\mathrm{viol}} \max\left(0, \tau_{\mathrm{2D}} - R_{\mathrm{2D}}(x_T; \mathcal{T}^\star)\right). \quad (32)$$

Here R_{3D} is measured by TM-score, and R_{2D} is measured by BP-F1. RMSD is recorded for evaluation but is not used as a reward term. The rollout, optimization, and reward settings are shown in Table A2.

F.3 Stage III Hyperparameters

Stage III applies frontier-guided fine-tuning. For each target, we warm-start an archive with candidate sequences sampled from the co-design policy. During fine-tuning, each update samples complete

Table A3: Stage III frontier-guided fine-tuning settings.

Setting	Value
Warm-start archive size	64
Actor learning rate	1×10^{-4}
Critic learning rate	1×10^{-4}
Buffer size	1024
Epochs per update	4
Minibatches per update	2
Entropy coefficient	0.02

candidate sequences, evaluates them with the same fold-back protocol, and assigns each trajectory a terminal reward given by its exclusive hypervolume contribution to the target-wise frontier. The default bi-objective frontier uses sc-BP-F1 and sc-TM, while the diversity-aware variant additionally includes sequence diversity. The fine-tuning settings are shown in Table A3.

G Evaluation Metrics

G.1 Fold-back Self-consistency Protocol

The main evaluation follows a fold-back self-consistency protocol. Let F_{3D} denote the selected tertiary-structure prediction oracle, and let A_{2D} denote the 3D-to-2D annotation procedure. For a target with native sequence $x^\star$, we define the reference tertiary structure and its associated secondary topology as

$$\mathcal{G}^\star = F_{\mathrm{3D}}(x^\star), \qquad \mathcal{T}^\star = A_{\mathrm{2D}}(\mathcal{G}^\star). \tag{33}$$

For a designed sequence x, we compute its fold-back tertiary structure and annotated secondary topology as

$$\widehat{\mathcal{G}}(x) = F_{\mathrm{3D}}(x), \qquad \widehat{\mathcal{T}}^{\mathrm{ann}}(x) = A_{\mathrm{2D}}\left(\widehat{\mathcal{G}}(x)\right). \tag{34}$$

The designed sequence is evaluated by comparing $\widehat{\mathcal{G}}(x)$ with $\mathcal{G}^\star$ and $\widehat{\mathcal{T}}^{\mathrm{ann}}(x)$ with $\mathcal{T}^\star$. This differs from directly comparing a predicted design structure with experimental coordinates. Since the reference and designed sequences are folded by the same oracle, the shared predictor-specific inductive bias is expected to cancel out in the structural similarity calculation [70]. We additionally report cross-oracle evaluation of `TANGO` designs in Figure 7 to examine whether the observed trends remain consistent under an independent fold-back predictor.

G.2 3D Self-consistency Metrics

We report tertiary-structure fidelity using sc-TM and sc-RMSD. The sc-TM score is the TM-score computed by USalign between the designed fold-back structure and the reference fold-back structure:

$$\text{sc-TM}(x, x^\star) = \mathrm{TM}\left(\widehat{\mathcal{G}}(x), \mathcal{G}^\star\right). \tag{35}$$

TM-score is length-normalized and evaluates global fold similarity [28]. Given an aligned set of positions $\mathcal{A}$ with pairwise distances d_i, USalign computes TM-score using its standard length normalization:

$$\mathrm{TM} = \max\left[\frac{1}{L_{\mathrm{ref}}} \sum_{i \in \mathcal{A}} \frac{1}{1 + (d_i/d_0(L_{\mathrm{ref}}))^2}\right]. \tag{36}$$

The sc-RMSD score is the RMSD over the aligned region after structural alignment:

$$\text{sc-RMSD} = \sqrt{\frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} d_i^2}. \tag{37}$$

Thus, sc-TM measures global fold agreement, whereas sc-RMSD reflects geometric deviation within the aligned region.

G.3 2D Self-consistency Metrics

Secondary-topology fidelity is measured by comparing the target topology $\mathcal{T}^\star$ with the annotated topology $\widehat{\mathcal{T}}^{\mathrm{ann}}(x)$ extracted from the designed fold-back structure. Let $\mathcal{P}^\star$ denote the target base-pair set and let $\widehat{\mathcal{P}}(x)$ denote the base-pair set extracted from $\widehat{\mathcal{T}}^{\mathrm{ann}}(x)$. The self-consistent base-pair F1 score is

$$\text{sc-BP-F1}(x, x^\star) = \frac{2|\widehat{\mathcal{P}}(x) \cap \mathcal{P}^\star|}{|\widehat{\mathcal{P}}(x)| + |\mathcal{P}^\star|}. \tag{38}$$

When both pair sets are empty, we define $\text{sc-BP-F1} = 1$; when only one pair set is empty, the score is 0.

G.4 Target-level Aggregation

Each method generates multiple designs for each target. We first aggregate metrics within each target and then report averages across test targets. For sc-TM, sc-RMSD, and sc-BP-F1, we compute the target-level mean over sampled designs and then average over targets. A design is considered successful if it satisfies

$$\text{sc-TM}(x, x^\star) \geq \tau_{3\mathrm{D}}, \qquad \text{sc-BP-F1}(x, x^\star) \geq \tau_{2\mathrm{D}}. \tag{39}$$

The success rate reports the average fraction of sampled designs satisfying both thresholds. The solved target rate reports the fraction of test targets for which at least one sampled design satisfies both thresholds. Appendix D.1 evaluates how TANGO changes under alternative sc-TM and sc-BP-F1 thresholds. This target-level aggregation prevents targets with more valid sampled outputs from dominating the benchmark average.

G.5 Novelty and Diversity

Sequence novelty is defined as the normalized Hamming distance between a designed sequence and the native sequence of the target:

$$\text{Novelty}(x, x^\star) = \frac{1}{L} \sum_{p=1}^{L} \mathbb{I}[x_p \neq x_p^\star]. \tag{40}$$

For a target-level set of designed sequences $\mathcal{X}_t = \{x^{(1)}, \ldots, x^{(m)}\}$ with common length L, sequence diversity is the average pairwise normalized Hamming distance:

$$\text{Diversity}(\mathcal{X}_t) = \frac{2}{m(m-1)} \sum_{1 \leq i < j \leq m} \frac{1}{L} \sum_{p=1}^{L} \mathbb{I}\left[x_p^{(i)} \neq x_p^{(j)}\right], \qquad m \geq 2. \tag{41}$$

When diversity is analyzed under structural success, $\mathcal{X}_t$ is restricted to successful designs for that target, so that random failed sequences do not artificially inflate diversity.

H Limitations and Broader Impact

Future Work and Limitations Future work could investigate model merging [71] across policies adapted from the same Stage-II initialization, assessing whether a merged policy can retain useful design capabilities across RNA targets. Future evaluation could also consider dimensions beyond structural fidelity, including measures of biological function [72].

Beyond the design method itself, structure predictors are a key component of the RNA design loop. The design framework is not tied to a specific structure-prediction oracle, but more accurate predictors can improve the translation of computational designs into experimentally useful molecules. Recent progress in RNA and biomolecular structure prediction has made these predictors increasingly accurate, and they have already been used as practical components in experimental RNA engineering pipelines [73, 74]. Future advances in RNA structure prediction are expected to further improve the reliability and practical utility of TANGO.

Broader Impact TANGO may contribute to synthetic biology, RNA-based therapeutics, and studies of gene regulation by improving the computational design of RNAs with desired structural properties. Such methods could support the engineering of functional RNA circuits, aptamers, and ribozymes. However, AI-driven biological sequence design also raises concerns about reliability, interpretability, and experimental validation. Recent work on genomic models advocates systematic evaluation of explanations for consistency, faithfulness, and biological validity [75]. Such assessment will also be important when using model interpretations to guide RNA engineering.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction accurately reflect the paper's contributions and scope.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: A discussion of the limitations is provided in the Conclusion and Appendix H.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Algorithmic details are provided in Section 3, Appendices B, E and F.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The code will be made publicly available upon acceptance.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Experimental setting/details are provided in Section 4, Section 5 and Appendix F.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The results include information on statistical significance (error bars, statistical significance tests). Experimental results are provided in Section 5 and Appendix D.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Information about the computer resources is included in Appendix F.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research fully complies with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Broader impacts are discussed in Appendix H.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All datasets and algorithms used in this paper are properly cited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.