

PAPER

Probabilistic RNA Designability via Interpretable Ensemble Approximation and Dynamic Decomposition

Tianshuo Zhou¹, David H. Mathews^{3,4,5} and Liang Huang^{1,2,*}¹School of EECS, ²Dept. of Biochemistry & Biophysics, Oregon State University, USA, ³Dept. of Biochemistry & Biophysics, ⁴Center for RNA Biology and ⁵Dept. of Biostatistics and Computational Biology, University of Rochester Medical Center, USA

*Corresponding author: liang.huang.sh@gmail.com

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

Abstract

Motivation: RNA design, also known as RNA inverse folding, aims to find RNA sequences that fold into a target secondary structure. However, recent work has shown that some target structures are provably *undesignable*, where no RNA sequence can fold into it as the *minimum free energy* (MFE) structure. In this paper, we go beyond this binary, MFE-based designability and explore a soft, probability-based designability that upperbounds the Boltzmann probability of any design and quantifies how easily or likely any design might possibly fold into the target structure. We introduce a theory of ensemble approximation and a probability decomposition framework for bounding the folding probabilities of RNA structures and motifs in an explainable way. We further develop a linear-time dynamic programming algorithm that efficiently searches over exponentially many decompositions. Combining ensemble approximation with dynamic decomposition search, our method efficiently identifies the optimal motif decomposition that yields the tightest probabilistic bound for a given structure. Our framework is applicable to any *factorizable* energy model or scoring function that decomposes onto loops.

Results: Applying our work, *LinearDecompose*, to both *native* and *artificial RNA structures in the ArchiveII and Eterna100 datasets*, we obtained much tighter probability bounds than baselines. Our work also provides anatomical tools for analyzing RNA structures and pinpointing the sources of design difficulty at the motif level.

Availability: Source code and data are available at <https://github.com/shanry/RNA-Undesign>.

Supplementary information: Supplementary text and data are available in a separate PDF.

Key words: RNA Design, RNA Designability, RNA Inverse Folding, Equilibrium Probability

1. Introduction

RNA secondary structures play crucial roles in the functions of non-coding RNAs such as rRNA (Doudna and Cech, 2002) and tRNA (Coller and Ignatova, 2024). RNA design, also known as RNA inverse folding, aims to find one or more RNA sequences that fold into a given target structure, typically using the default Turner energy model (Mathews and Turner, 2006; Turner and Mathews, 2010). However, not all RNA structures are designable (Aguirre-Hernández *et al.*, 2007; Yao, 2021; Zhou *et al.*, 2024, 2025, 2026; Malik *et al.*, 2026). Our recent works RIGEND (Zhou *et al.*, 2024) and FastMotif (Zhou *et al.*, 2025) proved that many native structures in the ArchiveII dataset and many artificial structures in the Eterna100 benchmark are undesignable, in the sense that no sequence can ever fold into them as MFE structures. However, these methods only study the MFE-based criteria and therefore only

provide a *binary* notion of designability. As a result, they cannot address the more important ensemble-level design objectives such as equilibrium probability which, by considering competing structures, is more appropriate for RNA design (Zadeh *et al.*, 2010; Ward *et al.*, 2023; Zhou *et al.*, 2023; Tang *et al.*, 2026).

This perspective motivates a shift from binary MFE-based criteria to a probabilistic characterization of designability. While MFE-based undesignability can certify that a structure is impossible to be a minimum free energy structure of any RNA sequence, it cannot quantify *how likely* or how easily a sequence might possibly fold into that structure. For example, consider two structures \mathbf{y}_1 and \mathbf{y}_2 that are both undesignable under the MFE criterion. It is nevertheless possible that there exists a sequence \mathbf{x}_1 where $p(\mathbf{y}_1 | \mathbf{x}_1) = 0.2$, whereas no sequence can make $p(\mathbf{y}_2 | \mathbf{x})$ exceed 0.01. This discrepancy is invisible to MFE-based analysis but is crucial for understanding practical designability of a target structure.

These observations motivate the study of *probabilistic designability*, defined as an upper bound on $p(\mathbf{y}^* | \mathbf{x})$ over all sequences \mathbf{x} for a given target structure \mathbf{y}^* . The tighter the upper bound, the better we understand the designability of that structure.

Despite its importance, probabilistic designability has received little attention, with CountingDesign (Yao et al., 2019; Yao, 2021) being the only existing method that addresses this question. However, it relies on exhaustive enumeration: for a given RNA structure or motif, it enumerates and folds all (partial) sequences that satisfy its base-pairing constraints and record the maximum equilibrium probability observed. This brute-force strategy is neither scalable nor interpretable, requiring weeks of computation for structures or motifs of length up to 14.

To bridge the gap between MFE-based and probabilistic designabilities, we propose two main ideas: *ensemble approximation* and *probability decomposition*. Ensemble approximation represents the full RNA folding ensemble using a small, explicitly identified set of rival structures (motifs) that compete with the target structure. This approximated ensemble enables us to derive rigorous upper bounds on folding probability while maintaining full explainability, as each bound is supported by concrete rival structures that thermodynamically dominate the target. Furthermore, we prove that the probability bound of a target structure is no greater than the *product of the local probability bounds* of its constituent motifs under *any* structural decomposition. In order to explore as many decompositions as possible, we develop a *linear-time* dynamic programming algorithm that explores *exponentially many* decompositions to find the one with the tightest upper bound. Our contributions are:

1. **Theory.** We propose a novel and elegant theory of ensemble approximation to interpretably characterize probabilistic RNA designability. We also prove that the probability bound of a structure can be factored into products of local bounds of structural motifs, revealing a more nuanced connection between local and global designabilities.
2. **Algorithms.** We develop efficient algorithms (Algs. 1 and 2) for approximating Boltzmann ensembles using rival structures or motifs, and introduce a linear-time dynamic programming approach (Algorithm 3) that efficiently explores exponentially many decompositions to obtain the tightest probability bound.
3. **Application.** Applying our methods to both native and artificial RNA structures in the ArchiveII and Eterna100 datasets, we obtained probability bounds that are much tighter than prior approaches. In addition, our methods further provide anatomical tools for analyzing RNA structures and understanding the sources of design difficulty at the motif level.

2. Preliminaries: RNA Structures and Motifs

An RNA sequence \mathbf{x} of length n is a string of nucleotides $\mathbf{x}_1\mathbf{x}_2\dots\mathbf{x}_n$, where $\mathbf{x}_i \in \{A, C, G, U\}$. A secondary structure \mathcal{P} for \mathbf{x} is a set of paired indices where each pair $(i, j) \in \mathcal{P}$ indicates two distinct bases $\mathbf{x}_i\mathbf{x}_j \in \{CG, GC, AU, UA, GU, UG\}$ and each base only be paired

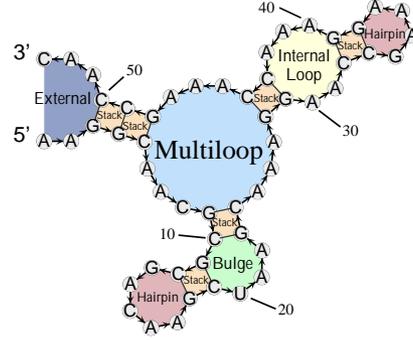


Fig. 1. A secondary structure and its loops (including stacks).

at most once. A secondary structure is pseudoknot-free (i.e., nested pairs only) if there are no two pairs $(i, j) \in \mathcal{P}$ and $(k, l) \in \mathcal{P}$ such that $i < k < j < l$. This work does not consider pseudoknots. \mathcal{P} can also be represented as a dot-bracket string $\mathbf{y} = \mathbf{y}_1\mathbf{y}_2\dots\mathbf{y}_n$, where a pair of indices $(i, j) \in \mathcal{P}$ corresponds to $\mathbf{y}_i = "("$ and $\mathbf{y}_j = ")"$ and any unpaired index k corresponds to $\mathbf{y}_k = "."$. The unpaired indices in \mathbf{y} are denoted as *unpaired*(\mathbf{y}) and the set of paired indices in \mathbf{y} is denoted as *pairs*(\mathbf{y}), which equals \mathcal{P} . For more details, see LinearFold (Huang et al., 2019) and LinearPartition (Zhang et al., 2020).

The *ensemble* of an RNA sequence \mathbf{x} is the set of all secondary structures that \mathbf{x} can possibly fold into, denoted as $\mathcal{Y}(\mathbf{x})$. The *free energy (change)* $\Delta G^\circ(\mathbf{x}, \mathbf{y})$ to characterizes the stability of $\mathbf{y} \in \mathcal{Y}(\mathbf{x})$. The lower $\Delta G^\circ(\mathbf{x}, \mathbf{y})$, the more stable the structure \mathbf{y} for \mathbf{x} .

A structure \mathbf{y} is composed of a set of loops denoted as *loops*(\mathbf{y}) where stacks of adjacent base pairs are also considered special loops, as illustrated in Fig. 1. The free energy change of a secondary structure \mathbf{y} is the sum of the free energy change of each loop, i.e.,

$$\Delta G^\circ(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{z} \in \text{loops}(\mathbf{y})} \Delta G^\circ(\mathbf{x}, \mathbf{z}), \quad (1)$$

where each term $\Delta G^\circ(\mathbf{x}, \mathbf{z})$ is the energy for loop \mathbf{z} . The energy of each loop is typically determined by nucleotides on the positions of enclosing pairs and their adjacent mismatch positions, which are named as *critical positions* and denoted as *critical*(\mathbf{z}). See Supplementary Section S1 for a detailed explanation of different loop types and their critical positions. The structure with the *minimum free energy* is the most stable structure in the ensemble $\mathcal{Y}(\mathbf{x})$. The minimum free energy of $\mathcal{Y}(\mathbf{x})$ is defined as

$$\text{MFE}(\mathbf{x}) = \min_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} \Delta G^\circ(\mathbf{x}, \mathbf{y}). \quad (2)$$

The partition function sums up the Boltzmann terms of all structures,

$$Q(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} e^{-\Delta G^\circ(\mathbf{x}, \mathbf{y})/RT}, \quad (3)$$

which defines the equilibrium probability or Boltzmann probability of a structure \mathbf{y} in the ensemble

$$p(\mathbf{y} | \mathbf{x}) = \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{y})/RT}}{Q(\mathbf{x})} = \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{y})/RT}}{\sum_{\mathbf{y}' \in \mathcal{Y}(\mathbf{x})} e^{-\Delta G^\circ(\mathbf{x}, \mathbf{y}')/RT}}. \quad (4)$$

A (structural) *motif* is defined as a contiguous set of loops in a structure, which generalizes RNA structures (See Supplementary Section S3 or FastMotif (Zhou *et al.*, 2024) for details). The thermodynamic definitions in Eqs. (1), (2) and (4) can be naturally extended to motifs. For example, the equilibrium probability of a motif \mathbf{m} in the ensemble of a (partial) sequence \mathbf{x} is

$$p(\mathbf{m} | \mathbf{x}) = \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m})/RT}}{Q(\mathbf{x})} = \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m})/RT}}{\sum_{\mathbf{m}' \in \mathcal{M}(\mathbf{x})} e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m}')/RT}}, \quad (5)$$

where $\mathcal{M}(\mathbf{x})$ is the motif ensemble of (partial) sequence \mathbf{x} .

3. RNA Designability: MFE vs. Probabilistic

We first review the MFE-based designability before moving on to probabilistic designability.

3.1. MFE-based Designability

Given a target structure \mathbf{y}^* , MFE-based RNA design aims to find a suitable RNA sequence \mathbf{x} such that \mathbf{y}^* is an MFE structure of \mathbf{x} . We define $\mathcal{X}(\mathbf{y})$ to be the design space of structure \mathbf{y} , which is the set of all RNA sequences that might fold into \mathbf{y} , i.e., $\mathcal{X}(\mathbf{y}) = \{\mathbf{x} | \mathbf{y} \in \mathcal{Y}(\mathbf{x})\}$. Structure \mathbf{y}^* is undesignable by the MFE criterion if and only if

$$\forall \mathbf{x} \in \mathcal{X}(\mathbf{y}), \exists \mathbf{y}' \neq \mathbf{y}^*, \Delta G^\circ(\mathbf{x}, \mathbf{y}') < \Delta G^\circ(\mathbf{x}, \mathbf{y}), \quad (6)$$

which means for any sequence \mathbf{x} there is always a rival structure \mathbf{y}' that is energetically more favored than \mathbf{y}^* .

Similarly, following the unique MFE (uMFE) criterion from previous studies (Haleš *et al.*, 2015; Yao *et al.*, 2019; Ward *et al.*, 2023; Zhou *et al.*, 2023), a structure \mathbf{y}^* is undesignable by the (uMFE) criterion if and only if

$$\forall \mathbf{x} \in \mathcal{X}(\mathbf{y}), \exists \mathbf{y}' \neq \mathbf{y}^*, \Delta G^\circ(\mathbf{x}, \mathbf{y}') \leq \Delta G^\circ(\mathbf{x}, \mathbf{y}), \quad (7)$$

which means for any sequence \mathbf{x} there is always a rival structure \mathbf{y}' that is energetically no worse than \mathbf{y}^* .

3.2. Probabilistic Designability

MFE-based designability cannot quantify how likely or how easily a target structure can possibly form in the ensemble of any RNA sequence. To address this problem, we aim to find an upper bound $pbound(\mathbf{y})$ such that

$$pbound(\mathbf{y}) \geq \max_{\mathbf{x}} p(\mathbf{y} | \mathbf{x}).$$

While it is hard to obtain the exact value of $\max_{\mathbf{x}} p(\mathbf{y} | \mathbf{x})$, the tighter $pbound(\mathbf{y})$ is, the better we know about the probabilistic designability of \mathbf{y} . Similarly, we can also find an upper bound $pbound(\mathbf{m})$ for a motif \mathbf{m} such that $pbound \geq (\mathbf{m}) \max_{\mathbf{x}} p(\mathbf{m} | \mathbf{x})$.

4. Ensemble Approximation via Rival Search

Rival structures (motifs) play a pivotal role in the previous works RIGEND (Zhou *et al.*, 2024) and FastMotif (Zhou *et al.*, 2025), which can prove RNA structures (motifs) undesignable under the uMFE criterion. The central idea



Fig. 2. Ensemble approximation with a single rival structure.

is to identify a small set of rival structures (motifs)

$$Y_r = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k\}, \mathbf{y}^* \notin Y_r,$$

such that

$$\forall \mathbf{x}, \exists \mathbf{y}' \in Y_r, \Delta G^\circ(\mathbf{x}, \mathbf{y}') \leq \Delta G^\circ(\mathbf{x}, \mathbf{y}^*). \quad (8)$$

In other words, for every RNA sequence \mathbf{x} , at least one explicitly identified rival structure is energetically no worse than the target structure. This condition guarantees that the target structure can never be the unique minimum free energy conformation of any sequence and is therefore uMFE-undesignable.

While effective, both RIGEND and FastMotif treat rival structures as *independent* competitors. Consequently, the collective thermodynamic effect of multiple rival structures is left unexploited. In this section, we introduce a new theory of Boltzmann ensemble approximation that leverages the collective contribution of rival structures to derive an upper bound on the equilibrium probability of the target structure.

To make the derivation intuitive, we begin with the simplest case, in which the ensemble is approximated by the target structure plus a single rival structure. We then extend this theory to the general case, where the ensemble approximation consists of multiple rival structures (motifs). Since motifs are generalizations of structures, we present the theory using structures for notational simplicity; all results apply to motifs directly.

4.1. Single Rival Structure (Motif)

We illustrate the idea using the Eterna100 structure “Simple Single Bond” (we cut off trailing unpaired bases to fit the page). When attempting to design this target structure \mathbf{y}^* (Fig. 2), the resulting sequence \mathbf{x} consistently folds into a different but structurally very similar conformation \mathbf{y}' .

As proven in RIGEND (Zhou *et al.*, 2024), it turns out that the free energy change difference between \mathbf{y}' and \mathbf{y}^*

$$\Delta \Delta G^\circ(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) \triangleq \Delta G^\circ(\mathbf{x}, \mathbf{y}') - \Delta G^\circ(\mathbf{x}, \mathbf{y}^*)$$

is only dependent on a number of *differential positions* $\Delta(\mathbf{y}', \mathbf{y}^*)$ (e.g., pairs and mismatches), defined as

$$\Delta(\mathbf{y}', \mathbf{y}^*) \triangleq \bigcup_{\mathbf{z} \in \text{loops}(\mathbf{y}^*) \ominus \text{loops}(\mathbf{y}')} \text{critical}(\mathbf{z}).^1$$

Fig. 2 annotates the differential positions. Consequently, to verify that $\Delta \Delta G^\circ(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) < 0$ holds for all sequences \mathbf{x} , it suffices to enumerate nucleotide assignments restricted to $\Delta(\mathbf{y}', \mathbf{y}^*)$.² This reduction

¹ \ominus denotes the symmetric difference between two sets.

² The maximum size of $\Delta(\mathbf{y}', \mathbf{y}^*)$ or enumeration is constrained to bound complexity or running time.

Algorithm 1 Ensemble Approximation with a Single Rival Structure (Motif)

function ENSEMBLEAPPROXIMATION(\mathbf{y}^* , \mathbf{y}')
 $\Delta\Delta G_{\max}^{\circ} \leftarrow \max_{\hat{\mathbf{x}} \in \{\mathbf{x} \vdash \Delta(\mathbf{y}', \mathbf{y}^*) \mid \mathbf{x} \in \mathcal{X}(\mathbf{y}^*)\}} \Delta\Delta G^{\circ}(\hat{\mathbf{x}}, \mathbf{y}', \mathbf{y}^*)$
return $\frac{1}{1 + e^{-\Delta\Delta G_{\max}^{\circ}/RT}}$ ▷ upper bound $pbound(\mathbf{y}^*)$

enables an efficient and interpretable proof that \mathbf{y}^* is undesignable under the uMFE criterion. Beyond the uMFE undesignability, however, we notice that this process yields additional thermodynamic insight. In particular, rather than merely establishing the sign of $\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*)$, we can compute the quantity,

$$\max_{\mathbf{x}} \Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*),$$

which represents the guaranteed energy advantage of the rival structure over the target across all sequences. This quantity directly leads to an upper bound on the equilibrium probability of the target structure:

$$\begin{aligned} p(\mathbf{y}^* \mid \mathbf{x}) &= \frac{e^{-\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)/RT}}{\sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} e^{-\Delta G^{\circ}(\mathbf{x}, \mathbf{y})/RT}} \\ &\leq \frac{e^{-\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)/RT}}{e^{-\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)/RT} + e^{-\Delta G^{\circ}(\mathbf{x}, \mathbf{y}')/RT}} \quad (9) \\ &= \frac{1}{1 + e^{-(\Delta G^{\circ}(\mathbf{x}, \mathbf{y}') - \Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)) / RT}}. \end{aligned}$$

Taking the maximum over all sequences \mathbf{x} , we obtain the following upper bound:

$$\max_{\mathbf{x}} p(\mathbf{y}^* \mid \mathbf{x}) \leq \frac{1}{1 + e^{-\max_{\mathbf{x}} \Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) / RT}}. \quad (10)$$

Algorithm 1 shows how to identify an upper bound for $\max_{\mathbf{x}} p(\mathbf{y}^* \mid \mathbf{x})$ using a single rival structure \mathbf{y}' , which can be obtained by first taking a sequence \mathbf{x} designed for \mathbf{y}^* and then folding \mathbf{x} (assume MFE(\mathbf{x}) is close to \mathbf{y}^*). Note that \vdash denotes the projection of a sequence \mathbf{x} onto a set of positions ($\Delta(\mathbf{y}', \mathbf{y}^*)$). (See Supp. Sec. S2)

4.2. Multiple Rival Structures (Motifs)

Algorithm 1 is effective and interpretable, but the resulting bound is limited by the coarse approximation of the ensemble using only one rival structure. Incorporating more rival structures will yield a tighter approximation, though enumerating the full ensemble is infeasible.

We therefore generalize the framework to include multiple rival structures. Figure 3 illustrates this idea using the Eterna100 structure “Zigzag Semicircle”. Our previous work (Zhou et al., 2024) proved this structure uMFE-undesignable by identifying nine rival structures

$$\forall \mathbf{x}, \exists \mathbf{y}' \in Y_r = \{\mathbf{y}'^1, \mathbf{y}'^2, \dots, \mathbf{y}'^9\}, \text{ s.t. } \Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) \leq 0.$$

Rather than selecting a single rival, we include all nine rivals in the ensemble approximation to obtain a tighter bound. Specifically,

$$\begin{aligned} p(\mathbf{y}^* \mid \mathbf{x}) &\leq \frac{e^{-\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)/RT}}{e^{-\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)/RT} + \sum_{\mathbf{y}' \in Y_r} e^{-\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}')/RT}} \\ &\leq \frac{1}{1 + \sum_{\mathbf{y}' \in Y_r} e^{-(\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}') - \Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}^*)) / RT}} \quad (11) \\ &\leq \frac{1}{1 + \sum_{\mathbf{y}' \in Y_r} e^{-\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) / RT}} \\ &\leq \frac{1}{1 + e^{-\Delta\Delta G^{\circ}(\mathbf{x}, Y_r, \mathbf{y}^*) / RT}}, \end{aligned}$$

where $e^{-\Delta\Delta G^{\circ}(\mathbf{x}, Y_r, \mathbf{y}^*) / RT} \triangleq \sum_{\mathbf{y}' \in Y_r} e^{-\Delta\Delta G^{\circ}(\mathbf{x}, \mathbf{y}', \mathbf{y}^*) / RT}$ summarizes the collective energetic dominance of the rival set over the target structure. Taking the maximum over all sequences yields

$$\max_{\mathbf{x}} p(\mathbf{y}^* \mid \mathbf{x}) \leq \frac{1}{1 + e^{-\max_{\mathbf{x}} \Delta\Delta G^{\circ}(\mathbf{x}, Y_r, \mathbf{y}^*) / RT}}. \quad (12)$$

As in the single-rival case, exhaustive enumeration over all sequences is unnecessary. Owing to the sparsity of loop energy contributions in the Turner model, it suffices to enumerate nucleotides on the *overall differential positions*

$$\Delta(Y_r, \mathbf{y}^*) = \cup_{\mathbf{y}' \in Y_r} \Delta(\mathbf{y}', \mathbf{y}^*). \quad (13)$$

Algorithm 2 outlines the resulting procedure to compute the upper bound using multiple rival structures.

4.3. Complexity Analysis

Algorithm 1 is a special case of Algorithm 2 with a single rival structure. The computational complexity of Algorithm 2 is determined by the size of the differential positions $\Delta(Y_r, \mathbf{y}^*)$ defined in Eq. 13, as the algorithm enumerates all nucleotide assignments restricted to these positions. Specifically, the time complexity is

$$O\left(6^{|\text{pairs}(\Delta(Y_r, \mathbf{y}^*))|} \cdot 4^{|\text{unpaired}(\Delta(Y_r, \mathbf{y}^*))|}\right), \quad (14)$$

where paired positions admit six canonical base-pair types and unpaired positions admit four nucleotides.

```

 $\mathbf{y}^*$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^1$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^2$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^3$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^4$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^5$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^6$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^7$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^8$ : ... ((((((.....)).).).)))...
 $\mathbf{y}'^9$ : ... ((((((.....)).).).)))...
 $\Delta$ : *****

```

Fig. 3. Ensemble approximation with nine rival structures.

Algorithm 2 Ensemble Approximation with Multiple Rival Structures (Motifs)

function ENSEMBLEAPPROXIMATION(\mathbf{y}^* , Y_r)
 $\Delta\Delta G_{\max}^{\circ} \leftarrow \max_{\hat{\mathbf{x}} \in \{\mathbf{x} \vdash \Delta(Y_r, \mathbf{y}^*) \mid \mathbf{x} \in \mathcal{X}(\mathbf{y}^*)\}} \Delta\Delta G^{\circ}(\hat{\mathbf{x}}, Y_r, \mathbf{y}^*)$
return $\frac{1}{1 + e^{-\Delta\Delta G_{\max}^{\circ}/RT}}$ ▷ upper bound $pbound(\mathbf{y}^*)$

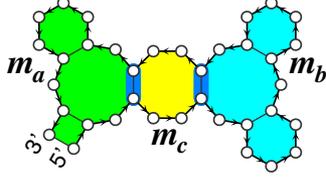


Fig. 4. Example of structure decomposition for the Eterna100 structure “multilooping fun”, which is decomposed into 3 motifs m_a , m_b , and m_c highlighted in different colors.

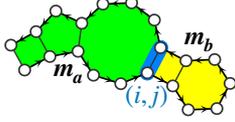


Fig. 5. A motif m is split into m_a and m_b at the base pair (i, j)

In practice, to control runtime, we sample rival structures (motifs) by folding sequences compatible with the target structure (motif) and impose an upper limit on the number of differential assignments enumerated. If this limit is exceeded and no suitable rival structures can be sampled, the corresponding structure or motif is skipped, trading completeness for efficiency.

5. Linear-time Dynamic Programming over Exponentially Many Decompositions

5.1. Structure and Probability Decomposition

FastMotif (Zhou *et al.*, 2025) shows that MFE-based structure undesignability can often be attributed to *local* (minimal) undesignable motifs. For example, Fig. 4 decomposes the Eterna100 structure *multilooping fun* into three motifs m_a , m_b , and m_c ; in this instance, either m_a or m_b alone can certify the uMFE-undesignability of the full structure, providing a more localized and interpretable explanation than relying solely on global rival structures.

The principle that global designability is constrained by local designability also extends to *probabilistic* designability. CountingDesign (Yao, 2021) proved that for any motif m contained in a target structure y^* , the motif probability upper-bounds the structure probability:

$$\forall m \in y^*, p(m | x) \geq p(y^* | x). \quad (15)$$

Applying this result to Fig. 4 yields, for example, $p(y^* | x) \leq p(m_a | x)$ and $p(y^* | x) \leq p(m_b | x)$.

However, CountingDesign bounds $p(y^* | x)$ by the probability of a *single* motif. The joint impact of multiple non-overlapping motifs remains unexplored. Here we show that when a structure can be decomposed into non-overlapping motifs, the structure probability is upper-bounded by the *product* of the motif probabilities.

Theorem 1 (Probability decomposition over non-overlapping motifs) *If a pseudoknot-free structure y^* can be decomposed into a set of non-overlapping motifs $M =$*

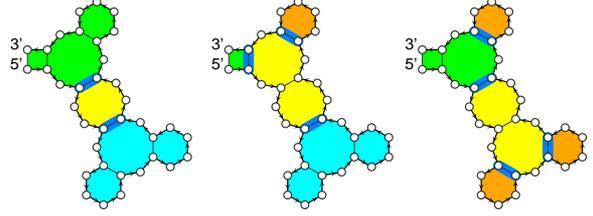


Fig. 6. Three different decompositions for the same structure for the same structure shown in Fig. 4. Within each decomposition, different motifs are highlighted in different colors.

$\{m_1, m_2, \dots, m_C\}$, then for any sequence $x \in \mathcal{X}(y^*)$,

$$p(y^* | x) \leq \prod_{m \in M} p(m | x). \quad (16)$$

Proof As a pseudoknot-free structure can be decomposed recursively, it suffices to prove the binary splitting step. Suppose a motif m is split into two submotifs m_a and m_b at a boundary base pair (i, j) , as illustrated in Fig. 5. Let $P_{i,j}$ denote the event that bases i and j are paired. By the chain rule of probability,

$$\begin{aligned} p(m | x) &= p(m_a + m_b | x) \\ &= p(P_{i,j} | x) \times p(m_a | x, P_{i,j}) \times p(m_b | x, P_{i,j}) \\ &= p(P_{i,j} | x) \times p(m_a | x) \times p(m_b | x) \\ &\leq p(m_a | x) \times p(m_b | x). \end{aligned} \quad (17)$$

In a pseudoknot-free structure, the boundary pair (i, j) ensures that the two submotifs m_a and m_b are independent (no interactions between them) in the Boltzmann ensemble, which justifies the third line. The fourth line holds because the boundary pairs are enforced in the motif-level folding. \square

Corollary 1

$$pbound(y^*) \leq \prod_{m \in M} pbound(m). \quad (18)$$

This corollary suggests that we can tighten the probability upper bound for a structure by exploiting probability bounds of smaller local motifs. For sufficiently small motifs, we can obtain an exact bound by enumerating all compatible (partial) sequences and folding them under the motif constraints. For motifs that are too large for exhaustive enumeration, we instead apply the motif ensemble approximation with rival motifs developed in the previous section.

Figure 6 shows three alternative decompositions of the same Eterna100 structure from Fig. 4. In fact, a given structure admits exponentially many decompositions, because each base pair may or may not be selected as a splitting boundary. This yields the following Theorem.

Theorem 2 *The total number of distinct decompositions of a secondary structure y is $2^{|\text{pairs}(y)|}$.*

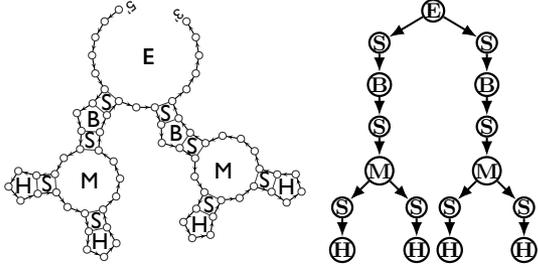


Fig. 7. Left: the Eternal100 structure *Chicken Feet*. Right: corresponding loop tree where each loop (including stacking) is represented as a node. E, S, B, H, M represent different loop types.

Let \mathcal{M} denote the set of decompositions that can be evaluated efficiently for a target structure \mathbf{y}^* . The best probability bound obtainable via decomposition is then

$$pbound(\mathbf{y}^*) \leq \min_{\mathbf{m} \in \mathcal{M}} \prod_{\mathbf{m} \in \mathcal{M}} pbound(\mathbf{m}). \quad (19)$$

5.2. Optimal Decomposition Search by Linear-Time Dynamic Programming

Although enumerating all decompositions is intractable, exploring more decompositions can only tighten the resulting upper bound. The key question is therefore how to search a large space of decompositions efficiently.

Inspired by dynamic programming techniques from compilers (Aho and Ullman, 1972; Aho and Johnson, 1976) and machine translation (Huang et al., 2006), we propose a linear-time memoized top-down algorithm (Algorithm 3) that explores many decompositions while remaining computationally efficient. We represent a target structure by a fixed loop tree τ , where each node corresponds to a loop in the structure (Fig. 7). We decompose the structure by traversing τ top-down, starting at the root. When visiting a loop node η , we are solving the subproblem for the subtree τ_η , and we must choose a motif \mathbf{m}_η that contains η and some of its descendants. To avoid enumerating all possibilities, we constrain candidate motifs using three parameters: *depth*, *width*, and *number of loops*. Figure 8 shows examples of motif candidates generated when visiting the external loop at the root; the candidate-generation procedure is described in Algorithm 5 (in Supplementary Section S4). Once \mathbf{m}_η is selected, we recursively process each adjacent descendant loop node η_i in the subtree τ_η . This procedure yields the following dynamic programming recurrence:

$$best(\tau_\eta) = \min_{\mathbf{m}_\eta \in \text{MOTIFGEN}(\eta)} pbound(\mathbf{m}_\eta) \prod_{\eta_i \in \text{descendants}(\mathbf{m}_\eta)} best(\tau_{\eta_i}), \quad (20)$$

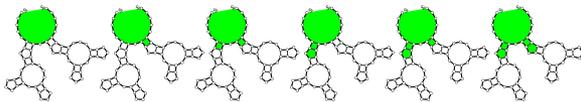


Fig. 8. Examples of motif candidates generated at the root node. For demonstration, here the maximum values for *depth*, *width*, and *number of loops* are set as 3, 2, and 5, respectively, and 6 out of 9 cases are shown. In experiments the maximum *depth* is set as 5.

Algorithm 3 Top-down Memoized Decomposition

```

1: function DECOMPOSE( $\eta$ )
2:   if cache[ $\eta$ ] defined then
3:     return cache[ $\eta$ ]
4:   best  $\leftarrow$  1
5:   candidates  $\leftarrow$  MOTIFGEN( $\eta$ , depth, width, max_loop)
6:   for  $\mathbf{m}_\eta \in$  candidates do  $\triangleright$  try each candidate  $\mathbf{m}_\eta$ 
7:     product  $\leftarrow$  pbound( $\mathbf{m}_\eta$ )  $\triangleright$  exact or approx.
8:     descendants  $\leftarrow$  loop nodes adjacent to  $\mathbf{m}$  in  $\tau_\eta$ 
9:     for  $\eta_i \in$  descendants do
10:      ( $p_i$ ,  $\mathbf{m}_{\eta_i}^{\text{best}}$ )  $\leftarrow$  DECOMPOSE( $\eta_i$ )  $\triangleright$  recursion
11:      product  $\leftarrow$  product  $\cdot$   $p_i$ 
12:     if product < best then
13:       best  $\leftarrow$  product
14:        $\mathbf{m}_\eta^{\text{best}}$   $\leftarrow$   $\mathbf{m}_\eta$   $\triangleright$  plug in the results
15:   cache[ $\eta$ ]  $\leftarrow$  (best,  $\mathbf{m}_\eta^{\text{best}}$ )  $\triangleright$  memoization
16:   return cache[ $\eta$ ]  $\triangleright$  optimal decomposition choice

```

where $best(\tau_\eta)$ represents the tightest upper bound for the substructure rooted at the loop node η . The base case occurs when η is a hairpin loop, for which $best(\tau_\eta) = 1$.

Each loop node may admit many candidate motifs; the number of possible decompositions can also be calculated by dynamic programming recurrence:

$$count(\tau_\eta) = \sum_{\mathbf{m}_\eta \in \text{MOTIFGEN}(\eta)} \prod_{\eta_i \in \text{descendants}(\mathbf{m}_\eta)} count(\tau_{\eta_i}), \quad (21)$$

with the base case value of 1 (leaf nodes). As a result, the decomposition possibilities grows exponentially with the size of the loop tree. We address this via memoization: each subtree rooted at a loop node is solved at most once and cached. We also store backpointers to recover the best motif choice at each node, and reconstruct the optimal decomposition by backtracking. The resulting dynamic program runs in $O(\alpha\beta)$ time, where α is the number of loop nodes in the loop tree and β is the maximum number of motif candidates considered at any node.

6. Evaluation Results

6.1. Implementation

Our work, *LinearDecompose*, consists of two core components: (i) Linear-Time Dynamic Decomposition (Alg. 3) and (ii) Ensemble Approximation (Alg. 2). The first module efficiently searches for the optimal motif decomposition by the recurrence formula Eq. 20 over exponentially many candidates. The second module estimates probability bounds $pbound(\mathbf{y}^*)$ or $pbound(\mathbf{m}^*)$ for a target structure or motif.

Our C++ code runs on Linux machines with 4.0 GHz CPUs and 32 GB memory. It calls ViennaRNA (v2.7) folding energy engine (Lorenz et al., 2011) with Turner 2004 parameters (Mathews et al., 2004). Ensemble approximation is parallelized with OpenMP over 90 cores.

During decomposition search, Algorithm 3 repeatedly queries motif probability bounds $pbound(\mathbf{m}_\eta)$ (line 7). To improve efficiency, we precompute and cache motif bounds using a hybrid strategy: (1) exact bounds via

Table 1. Average probability bounds on each dataset.

Dataset	ArchiveII	Eterna100
# of unique structures	1144	100
best achieved $p(\mathbf{y}^* \mathbf{x}) \uparrow$	0.766	0.599
$pbound(\mathbf{y}^*) \downarrow$		
CountingDesign	0.896 ± 0.242	0.898 ± 0.234
CountingDesign+	0.875 ± 0.255	0.876 ± 0.250
LinearDecompose	0.800 ± 0.288	0.741 ± 0.318
no ensemble approx.	0.866 ± 0.256	0.839 ± 0.275
no dynamic decomp.	0.819 ± 0.286	0.787 ± 0.295

exhaustive enumeration for motifs of length up to 14, following CountingDesign (Yao, 2021); and (2) non-exact bounds via ensemble approximation for larger motifs.

6.2. Datasets and Baselines

We evaluate on two popular RNA structure datasets:

1. **ArchiveII** (Sloma and Mathews, 2016; Cannone *et al.*, 2002) contains native RNA secondary structures spanning 10 families of naturally RNAs, including tRNA, rRNA, etc. We remove pseudoknotted structures and those incompatible with ViennaRNA’s default loop constraints. We only use structures longer than 200 *nt*, resulting in 1,144 structures.
2. **Eterna100** (Anderson-Lee *et al.*, 2016) consists of 100 secondary structures of varying design difficulty designed by human players of Eterna and is widely used for evaluating RNA design and designability.

We compare our method against CountingDesign and our enhanced version of it, CountingDesign+:

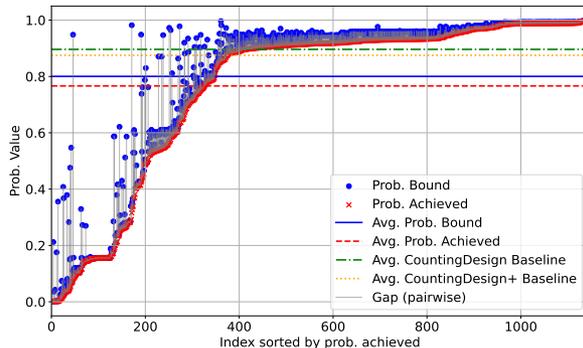
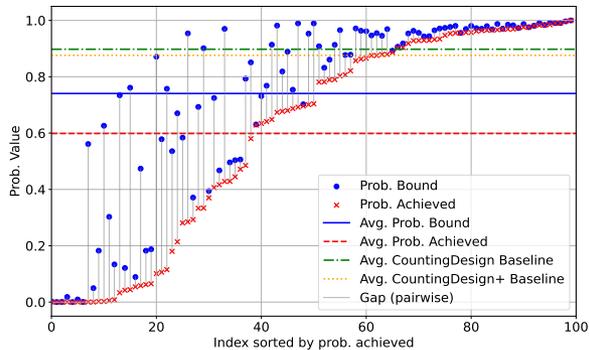
1. **CountingDesign** (Yao *et al.*, 2019; Yao, 2021) can identify exact probability bounds for very short motifs (up to 14 nucleotides) by exhaustively enumerating RNA sequences for each motif. However, the method is not scalable and lacks explainability.
2. **CountingDesign+**. We extend CountingDesign to include motifs with external loops, thus expanding their motif definition to align with ours.

6.3. Overall and Individual Results

Table 1 reports average probability bounds on both datasets. For reference, we also include the highest $p(\mathbf{y}^*|\mathbf{x})$ achieved by the state-of-the-art RNA design methods, taking the better of SAMFEO (Zhou *et al.*, 2023) and SamplingDesign (Tang *et al.*, 2026) for each puzzle.

LinearDecompose consistently yields the tightest bounds, improving over CountingDesign+ by 0.075 on ArchiveII and 0.135 on Eterna100. This improvement is expected, as LinearDecompose is theoretically guaranteed to produce bounds no worse than CountingDesign+.

We also observe a clear correlation between probability bounds and empirical designability. ArchiveII structures are generally more designable: the average bound (0.800) is close to the achieved probability (0.766). In contrast, Eterna100 structures are less designable, with a larger gap between the achieved probability (0.599) and the bound (0.741). This difference is statistically significant (Mann-Whitney $U = 72,687$, $p = 6.95 \times 10^{-6}$, $r = 0.27$).

**Fig. 9.** Probability bounds vs. achieved $p(\mathbf{y}^*|\mathbf{x})$ on ArchiveII.**Fig. 10.** Probability bounds vs. achieved $p(\mathbf{y}^*|\mathbf{x})$ on Eterna100.

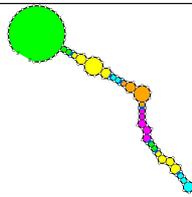
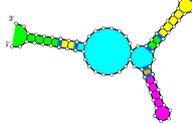
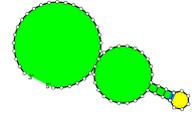
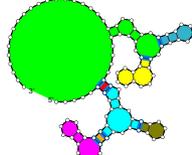
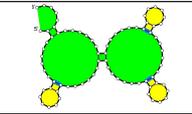
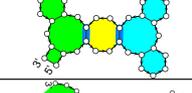
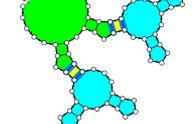
Figs. 9–10 show per-structure bounds and best achieved $p(\mathbf{y}^*|\mathbf{x})$. Bounds tend to be tighter for both highly designable and highly undesignable structures. Highly designable structures typically consist of motifs with uniformly high local probabilities, while highly undesignable structures often contain multiple motifs with very low bounds, often leading to near-zero global bounds.

These bounds also provide a useful lens for evaluating RNA design results. For example, a structure achieving probability 0.7 may be closer to its theoretical limit than another achieving 0.8 if its bound is substantially lower.

Table 2 presents a structural dissection of several Eterna100 puzzles previously shown to be uMFE-undesignable (Zhou *et al.*, 2025). For each structure \mathbf{y}^* , we report the number of decompositions explored by LinearDecompose (via dynamic programming Eq. 21), the total number of possible decompositions (via Theorem 2), the optimal decomposition achieving the tightest probability bound, the resulting global probability upper bound $pbound(\mathbf{y}^*)$, and the best $p(\mathbf{y}^*|\mathbf{x})$ achieved by SAMFEO and SamplingDesign (as a lowerbound).

LinearDecompose efficiently explores exponentially many decompositions per puzzle (e.g., 47M out of 67M total decompositions for the first puzzle), and consistently identifies a small set of local motifs whose combined bounds dominate the global probability bound. Notably, structures with complex multiloop or bulge configurations (e.g., *multilooping fun*) exhibit extremely low bounds, indicating that probabilistic undesignability can often be traced to a few highly constraining motifs. In addition, we also noticed that motifs with isolated base

Table 2. Structure Dissection with Eterna100 Puzzles.

Puzzle (y^*)	Decomposition		$pbound(y^*)$ $\max p(y^* x)$
	Explored	Optimal	
Mat - Elements & Sections	47, 012, 668 out of $2^{26} =$ 67, 108, 864		0.0889↓ 0.0426↑
Loop next to a Multiloop	2, 997, 104 out of $2^{22} =$ 4, 194, 304		0.4958↓ 0.4293↑
Simple Single Bond	32 out of $2^5 = 32$		0.3942↓ 0.3700↑
1, 2, 3 and 4 bulges	2, 734, 333 out of $2^{22} =$ 4, 194, 304		0.3026↓ 0.0054↑
Repetit. Seqs. 8/10	860 out of $2^{10} = 1, 024$		0.1874↓ 0.0651↑
multilooping fun	119 out of $2^7 = 128$		0.0006↓ 2×10^{-6} ↑
Chicken feet	52, 114 out of $2^{16} =$ 65, 536		0.1336↓ 0.0080↑

pairs tend to have low probability bounds, consistent with our prior systematic study (Zhou *et al.*, 2026) on undesignable motifs that isolated base pairs are energetically unfavorable and often avoided in natural RNA structures. These results highlight the value of decomposition-based analysis for pinpointing the structural origins of low RNA designability.

6.4. Ablation Studies

To assess the two core components of LinearDecompose, we perform two ablations: (1) disable motif ensemble approximation by setting approximate motif bounds to 1; and (2) disable dynamic decomposition by using only the best single-motif bound. As shown in the last two rows of Table 1, both ablations lead to looser bounds, confirming the importance of both components. Nevertheless, the ablated variants still outperform the baselines in Table 1, demonstrating the robustness of the framework.

6.5. Efficiency

The runtime of LinearDecompose consists of decomposition search and motif ensemble approximation. Decomposition search is fast, averaging 0.01 s per structure on ArchiveII and 0.9 s on Eterna100. Ensemble approximation estimates bounds for 15,800 and 8,000 unique motifs on ArchiveII and Eterna100, with an average cost of 7.6 s and 10.7 s per motif, respectively. Approximately 10% and 2% of motifs on ArchiveII and Eterna100, respectively, are skipped due to exceeding the enumeration time limit, for which the bound defaults to 1. As these fractions are small, we expect the impact on overall bound quality to be minimal. The average time costs for ensemble approximation per structure are 106 s (ArchiveII) and 1039 s (Eterna100). The total runtimes are 106 s and 1040 s per structure, respectively. Note that each motif bound can be cached once computed. As a result, the amortized runtime of LinearDecompose can be significantly reduced when analyzing multiple structures sharing common motifs, which is often the case in practice.

7. Discussion and Conclusion

We introduced a theoretical framework based on ensemble approximation and probability decomposition to quantify the probabilistic designability of RNA secondary structures under any *factorizable* energy model or scoring function that decomposes onto individual loops. We further develop a linear-time dynamic-programming algorithm that efficiently searches for optimal decompositions among exponentially many choices. Together, the resulting bounds offer a novel, interpretable characterization of probabilistic designability: rather than a black-box score, each bound is explicitly supported by rival motifs that thermodynamically dominate the target, and the optimal decomposition identifies the specific motifs responsible for low global designability.

Applied to the ArchiveII and Eterna100 datasets, our algorithm LinearDecompose consistently produces tighter probability bounds under the Turner energy model than existing methods. Moreover, the resulting optimal decompositions offer interpretable explanations of structural design difficulty by identifying the specific motifs that limit global designability.

LinearDecompose still has some limitations:

1. Unless a structure is designed with a probability close to the probability bound found by LinearDecompose, the bound may be optimistic but loose.
2. Ensemble approximation is not universally applicable; for some structures or motifs, suitable rival motifs are hard to sample.

Future work includes improving rival motif generation to tighten ensemble approximations (see Appendix Sec. S5 for more details) and leveraging the optimal decomposition given by LinearDecompose as a guide for RNA design.

Acknowledgement

This work was supported in part by NSF grant 2330737 (L.H. and D.H.M.). We thank the reviewers for comments.

References

- Aguirre-Hernández, R., Hoos, H. H. and Condon, A. (2007). Computational RNA secondary structure design: empirical complexity and improved methods. *BMC bioinformatics*, **8**(1), 1–16.
- Aho, A. V. and Johnson, S. C. (1976). Optimal code generation for expression trees. *Journal of the ACM (JACM)*, **23**(3), 488–501.
- Aho, A. V. and Ullman, J. D. (1972). *The theory of parsing, translation, and compiling*, volume 1. Prentice-Hall Englewood Cliffs, NJ.
- Anderson-Lee, J., Fisker, E., Kosaraju, V., Wu, M., Kong, J., Lee, J., Lee, M., Zada, M., Treuille, A. and Das, R. (2016). Principles for predicting RNA secondary structure design difficulty. *Journal of molecular biology*, **428**(5), 748–757.
- Cannone, J. J., Subramanian, S., Schnare, M. N., Collett, J. R., D’Souza, L. M., Du, Y., Feng, B., Lin, N., Madabusi, L. V., Müller, K. M., Pande, N., Shang, Z., Yu, N. and Gutell, R. R. (2002). The Comparative RNA Web (CRW) Site: An Online Database of Comparative Sequence and Structure Information for Ribosomal, Intron, and Other RNAs. *BioMed Central Bioinformatics*, **3**(2).
- Coller, J. and Ignatova, Z. (2024). tRNA therapeutics for genetic diseases. *Nature Reviews Drug Discovery*, **23**(2), 108–125.
- Doudna, J. A. and Cech, T. R. (2002). The chemical repertoire of natural ribozymes. *Nature*, **418**(6894), 222–228.
- Haleš, J., Maňuch, J., Ponty, Y. and Stacho, L. (2015). Combinatorial RNA design: designability and structure-approximating algorithm. In *Combinatorial Pattern Matching: 26th Annual Symposium, CPM 2015, Ischia Island, Italy, June 29–July 1, 2015, Proceedings*, pages 231–246. Springer.
- Huang, L., Knight, K. and Joshi, A. (2006). Statistical syntax-directed translation with extended domain of locality. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 66–73.
- Huang, L., Zhang, H., Deng, D., Zhao, K., Liu, K., Hendrix, D. A. and Mathews, D. H. (2019). LinearFold: linear-time approximate RNA folding by 5'-to-3'dynamic programming and beam search. *Bioinformatics*, **35**(14), i295–i304.
- Lorenz, R., Bernhart, S. H., Zu Siederdisen, C. H., Tafer, H., Flamm, C., Stadler, P. F. and Hofacker, I. L. (2011). ViennaRNA Package 2.0. *Algorithms for Molecular Biology*, **6**(1), 1.
- Malik, A., Zhou, T., Tang, W. Y., Mathews, D. H. and Huang, L. (2026). Motif Server: Web Server for Undesignable RNA Motifs and Structures. *Journal of Molecular Biology*.
- Mathews, D. H. and Turner, D. H. (2006). Prediction of RNA secondary structure by free energy minimization. *Current Opinion in Structural Biology*, **16**(3), 270–278.
- Mathews, D. H., Disney, M. D., Childs, J. L., Schroeder, S. J., Zuker, M. and Turner, D. H. (2004). Incorporating chemical modification constraints into a dynamic programming algorithm for prediction of RNA secondary structure. *Proceedings of the National Academy of Sciences U.S.A.*, **101**(19), 7287–7292.
- Sloma, M. and Mathews, D. H. (2016). Exact calculation of loop formation probability identifies folding motifs in RNA secondary structures. *RNA*, **22**, 1808–1818.
- Tang, W. Y., Dai, N., Zhou, T., Mathews, D. and Huang, L. (2026). SamplingDesign: RNA Design via Continuous Optimization with Coupled Variables and Monte-Carlo Sampling. *Nature Communications*.
- Turner, D. H. and Mathews, D. H. (2010). NNDB: the nearest neighbor parameter database for predicting stability of nucleic acid secondary structure. *Nucleic Acids Research*, **38**(suppl.1), D280–D282.
- Ward, M., Courtney, E. and Rivas, E. (2023). Fitness functions for RNA structure design. *Nucleic Acids Research*, **51**(7), e40–e40.
- Yao, H.-T. (2021). *Local decomposition in RNA structural design*. Ph.D. thesis, McGill University.
- Yao, H.-T., Chauve, C., Regnier, M. and Ponty, Y. (2019). Exponentially few RNA structures are designable. In *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, pages 289–298.
- Zadeh, J. N., Wolfe, B. R. and Pierce, N. A. (2010). Nucleic Acid Sequence Design via Efficient Ensemble Defect Optimization. *Journal of Computational Chemistry*, **32**(3), 439–452.
- Zhang, H., Zhang, L., Mathews, D. H. and Huang, L. (2020). LinearPartition: linear-time approximation of RNA folding partition function and base-pairing probabilities. *Bioinformatics*, **36**(Supplement_1), i258–i267.
- Zhou, T., Dai, N., Li, S., Ward, M., Mathews, D. H. and Huang, L. (2023). RNA design via structure-aware multifrontier ensemble optimization. *Bioinformatics*, **39**(Supplement_1), i563–i571.
- Zhou, T., Tang, W. Y., Mathews, D. H. and Huang, L. (2024). Undesignable RNA structure identification via rival structure generation and structure decomposition. In *International Conference on Research in Computational Molecular Biology*, pages 270–287. Springer.
- Zhou, T., Malik, A., Tang, W. Y., Mathews, D. H. and Huang, L. (2025). Scalable and interpretable identification of minimal undesignable RNA structure motifs with rotational invariance. In *International Conference on Research in Computational Molecular Biology*, pages 153–174. Springer.
- Zhou, T., Malik, A., Tang, W. Y., Mathews, D. H. and Huang, L. (2026). Theory, Algorithms, and Applications for Identification of Undesignable RNA Secondary Structures and Motifs. *Journal of Computational Biology*, **33**(1), 134–167.

Supplementary Information

S1. Structural Loops

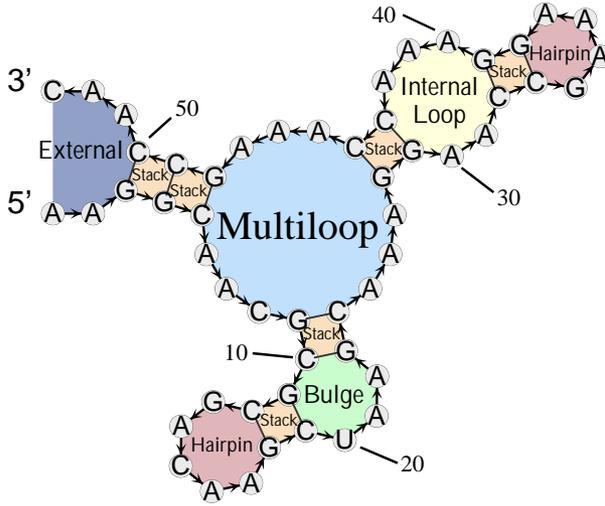


Fig. S1. An example of secondary structure and loops.

A secondary structure can be decomposed into a collection of loops, where each loop is usually a region enclosed by base pair(s). Depending on the number of pairs on the boundary, main types of loops include hairpin loop, internal loop and multiloop, which are bounded by 1, 2 and 3 or more base pairs, respectively. In particular, the external loop is the most outside loop and is bounded by two ends (5' and 3') and other base pair(s). Thus each loop can be identified by a set of pairs. Fig. S1 showcases an example of secondary structure with various types of loops, where some of the loops are notated as

1. Hairpin: $H\langle(12, 18)\rangle$.
2. Bulge: $B\langle(10, 23), (11, 19)\rangle$.
3. Stack: $S\langle(3, 50), (4, 49)\rangle$.
4. Internal Loop: $I\langle(29, 43), (32, 39)\rangle$.
5. Multiloop: $M\langle(5, 48), (9, 24), (28, 44)\rangle$.
6. External Loop: $E\langle(3, 50)\rangle$.

Table S1. Critical positions of loops in Fig. 1

Loop Type	Critical Positions	
	Closing Pairs	Mismatches (Unpaired)
External	(3, 50)	2, 51
Stack	(3, 50), (4, 49)	
Stack	(4, 49), (5, 48)	
Multi	(5, 48), (9, 24), (28, 44)	4, 49, 8, 25, 27, 45
Stack	(9, 24), (10, 23)	
Bulge	(10, 23), (11, 19)	
Stack	(11, 19), (12, 18)	
Hairpin	(12, 18)	13, 17
Stack	(28, 44), (29, 43)	
Internal	(29, 43), (32, 39)	30, 42, 31, 40
Stack	(32, 39), (33, 38)	
Hairpin	(33, 38)	34, 37

The function $loops(\mathbf{y})$ is used to denote the set of loops in a structure \mathbf{y} . The free energy of a secondary structure \mathbf{y} is the sum of the free energy of each loop,

$$\Delta G^\circ(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{z} \in loops(\mathbf{y})} \Delta G^\circ(\mathbf{x}, \mathbf{z}),$$

The energy of each loop is typically determined by nucleotides on the positions of enclosing pairs and their adjacent mismatch positions, which are named as *critical positions* in this article. Table S1 lists the critical positions for all the loops in Fig. 1 and Table S2 shows the indices of critical positions for each type of loops. Additionally, some special hairpins of unstable triloops and stable tetraloops and hexaloops in Turner model have a separate energy lookup table. When evaluating the energy of a loop, it suffices to input only the nucleotides on its critical positions, i.e.,

$$\Delta G^\circ(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{z} \in loops(\mathbf{y})} \Delta G^\circ(\mathbf{x} \vdash critical(\mathbf{z}), \mathbf{z}), \quad (22)$$

where $critical(\mathbf{z})$ denotes the critical positions of loop \mathbf{z} and $\mathbf{x} \vdash critical(\mathbf{z})$ denotes the nucleotides from \mathbf{x} that are "projected" onto $critical(\mathbf{z})$. The projection (\vdash) allows us to focus on the relevant nucleotides for energy evaluation. For instance,

$$critical(H\langle(12, 18)\rangle) = \{12, 13, 17, 18\}, \quad (23)$$

$$critical(I\langle(29, 43), (32, 39)\rangle) = \{29, 30, 31, 32, 39, 40, 42, 43\}. \quad (24)$$

For convenience, we also interchangeably put paired positions in brackets, i.e.,

$$critical(H\langle(12, 18)\rangle) = \{(12, 18), 13, 17\}, \quad (25)$$

$$critical(I\langle(29, 43), (32, 39)\rangle) = \{(29, 43), (32, 39), 30, 31, 40, 42\}. \quad (26)$$

the folding outcomes are unaffected by the nucleotides at the constrained positions. Thus, with slight notation abuse, we use \mathbf{x} to denote a partial sequence corresponding to a motif \mathbf{m} , where each position in \mathbf{x} matches a position in \mathbf{m} , and vice versa. By this definition, the motif ensemble of a partial sequence \mathbf{x} is denoted as $\mathcal{M}(\mathbf{x})$. Similarly, the notation $\mathcal{X}(\mathbf{m})$ generalizes $\mathcal{X}(\mathbf{y})$ and represents all (partial) RNA sequences whose motif ensembles contain \mathbf{m} . Motifs in $\mathcal{M}(\mathbf{x}, \text{bpairs}(\mathbf{m}))$ have the same boundary pairs, i.e.,

$$\forall \mathbf{m}', \mathbf{m}'' \in \mathcal{M}(\mathbf{x}), \text{bpairs}(\mathbf{m}') = \text{bpairs}(\mathbf{m}'') = \text{bpairs}(\mathbf{m}). \quad (28)$$

The free energy change of a motif \mathbf{m} is the sum of the free energy of the loops in \mathbf{m} ,

$$\Delta G^\circ(\mathbf{x}, \mathbf{m}) = \sum_{\mathbf{z} \in \text{loops}(\mathbf{m})} \Delta G^\circ(\mathbf{x}, \mathbf{z}). \quad (29)$$

The definitions of MFE and uMFE can also be generalized to motifs via constrained folding.

Definition 2 A motif $\mathbf{m}^* \subseteq \mathbf{y}$ is an MFE motif of folding \mathbf{x} under constraint $\text{bpairs}(\mathbf{m})$, i.e., $\text{MFE}(\mathbf{x}, \text{bpairs}(\mathbf{m}))$, if and only if

$$\forall \mathbf{m} \in \mathcal{M}(\mathbf{x}, \text{bpairs}(\mathbf{m})) \text{ and } \mathbf{m} \neq \mathbf{m}^*, \quad \Delta G^\circ(\mathbf{x}, \mathbf{m}^*) \leq \Delta G^\circ(\mathbf{x}, \mathbf{m}). \quad (30)$$

Definition 3 A motif $\mathbf{m}^* \subseteq \mathbf{y}$ is an uMFE motif of folding \mathbf{x} under constraint $\text{bpairs}(\mathbf{m})$, i.e., $\text{uMFE}(\mathbf{x}, \text{bpairs}(\mathbf{m}))$, if and only if

$$\forall \mathbf{m} \in \mathcal{M}(\mathbf{x}, \text{bpairs}(\mathbf{m})) \text{ and } \mathbf{m} \neq \mathbf{m}^*, \quad \Delta G^\circ(\mathbf{x}, \mathbf{m}^*) < \Delta G^\circ(\mathbf{x}, \mathbf{m}). \quad (31)$$

Similarly, the equilibrium probability of a sequence folding into the motif is defined as,

$$\begin{aligned} p(\mathbf{m} | \mathbf{x}) &= \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m})/RT}}{Q(\mathbf{x})} \\ &= \frac{e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m})/RT}}{\sum_{\mathbf{m}' \in \mathcal{M}(\mathbf{x}, \text{bpairs}(\mathbf{m}))} e^{-\Delta G^\circ(\mathbf{x}, \mathbf{m}')/RT}}. \end{aligned} \quad (32)$$

S4. Additional Pseudocode

The motif generation procedure used in LinearDecompose is shown in Algorithm 5, which generates candidate motifs for each loop node based on the specified constraints of depth, width, and number of loops.

S5. Discussion on the Tightness of Ensemble Approximation

The tightness of the probability bound produced by ensemble approximation (Algorithms 1 and 2) depends on two key factors: the *quality* and the *number* of rival structures (motifs).

Algorithm 5 Constrained Motif Generation from a Loop

```

1: function MOTIFGEN( $\eta, \text{max\_depth}, \text{max\_width}, \text{max\_loop}$ )
2:    $\text{candidates} \leftarrow \emptyset$ 
3:   for  $d \leftarrow 1$  to  $\text{max\_depth}$  do
4:      $\mathbf{M}_d \leftarrow \emptyset$   $\triangleright$  motifs of depth  $d$  to be generated
5:     if  $d = 1$  then
6:        $\mathbf{M}_d \leftarrow \{\eta\}$   $\triangleright$  the motif is a single node  $\eta$ 
7:     else
8:       for  $\mathbf{m} \in \mathbf{M}_{d-1}$  do
9:          $\mathbf{m}_{\text{new}} \leftarrow \text{GROW}(\mathbf{m})$ 
10:        if  $\text{WIDTH}(\mathbf{m}_{\text{new}}) > \text{max\_width}$  then
11:          continue
12:        if  $|\text{LOOPS}(\mathbf{m}_{\text{new}})| > \text{max\_loop}$  then
13:          continue
14:         $\mathbf{M}_d \leftarrow \mathbf{M}_d \cup \mathbf{m}_{\text{new}}$ 
15:       $\text{candidates} \leftarrow \text{candidates} \cup \mathbf{M}_d$ 
16:   return  $\text{candidates}$ 

```

S5.1. Quality of rival structures.

A rival structure is considered high quality if it thermodynamically dominates the target across a broad range of sequences, i.e., $\Delta \Delta G^\circ(\mathbf{x}, \mathbf{y}', \mathbf{y}^*)$ is large and positive for many sequence assignments. Such rivals push the estimated bound down substantially. In contrast, a rival that only marginally outcompetes the target in a narrow region of sequence space contributes little to tightening the bound. In practice, rival quality is limited by the sampling strategy: rivals are obtained by folding sequences designed for the target, so they tend to be structurally similar to the target with small differential position sets $\Delta(\mathbf{y}', \mathbf{y}^*)$. When the sampled rivals are structurally distant or energetically weak competitors, the resulting bound may be loose.

S5.2. Number of rival structures.

Including more rivals in the ensemble approximation directly tightens the bound, because each additional rival contributes an extra term to the denominator of the probability bound (Eq. 8). Intuitively, more competitors in the approximated ensemble means the target structure claims a smaller share of the Boltzmann weight. The bound can only decrease as rivals are added, so the method is conservative by construction. However, adding more rivals also increases the size of the overall differential positions $\Delta(Y_r, \mathbf{y}^*)$ (Eq. 13), which exponentially increases the enumeration cost. In practice, we impose an upper limit on the enumeration size, which may prevent the bound from fully benefiting from a large rival set.

S5.3. When bounds may be loose.

Bounds tend to be loose when: (1) the sampled rivals are of low quality, e.g., they only weakly dominate the target; (2) the number of sampled rivals is small; or (3) the enumeration limit is reached before the best sequence assignment is found. These conditions are more likely for large structures with complex loop configurations, where the differential position set grows quickly and sampling diverse, high-quality rivals is harder.

S5.4. Utility for evaluating RNA design results.

Beyond certifying undesignability, the probability bounds provide a practical tool for evaluating and comparing RNA design results. Table 2 illustrates varying degrees of bound tightness across different Eterna100 puzzles. For example, *Simple Single Bond* achieves a best design probability of 0.370 against a bound of 0.394 — a tight gap indicating the structure is close to its theoretical limit and the design is near-optimal. In contrast, *1, 2, 3 and 4 bulges* achieves only 0.005 against a bound of 0.303, suggesting substantial room for improvement. Similarly, *multilooping fun* has a bound of 0.0006, consistent with the negligible best design probability of 2×10^{-6} , confirming that this structure is highly constrained and essentially undesignable at the ensemble level. These examples demonstrate that probability bounds provide a richer, quantitative characterization of design difficulty than binary undesignability criteria alone.

S6. Detailed Plots

The following figures show the probability bounds vs. achieved probabilities for all 1144 ArchiveII structures and 100 Eterna100 structures. In addition, the per-structure bounds from baselines (CountingDesign and CountingDesign+) are also scattered for comparison.

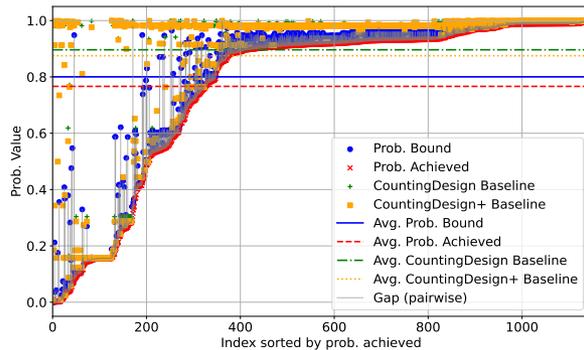


Fig. S3. Probability bounds vs. achieved $p(\mathbf{y}^* | \mathbf{x})$ on ArchiveII.

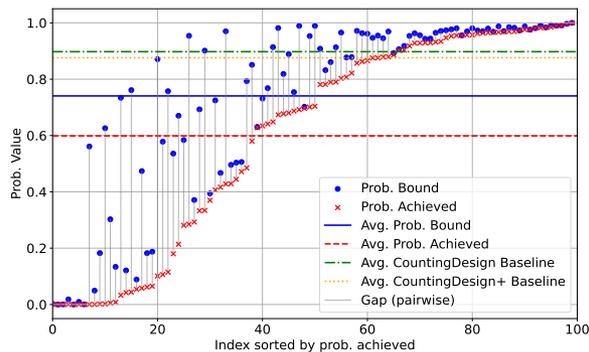


Fig. S4. Probability bounds vs. achieved $p(\mathbf{y}^* | \mathbf{x})$ on Eterna100.