

MERGE-RNA: a physics-based model to predict RNA secondary structure ensembles with chemical probing

Giuseppe Sacco ¹, Jianhui Li ², Redmond P. Smyth ², Guido Sanguinetti ¹, Giovanni Bussi ^{1,*}

¹Scuola Internazionale Superiore di Studi Avanzati, via Bonomea 265, Trieste, 34136, Italy


²Architecture et Réactivité de l'ARN, Université de Strasbourg, CNRS, Institute of Molecular and Cellular Biology (IBMC), 2 Allée Konrad Roentgen, 67000 Strasbourg, France

*To whom correspondence should be addressed. Email: bussi@sissa.it

Abstract

RNA function is tied to secondary structure, operating through dynamic and heterogeneous structural ensembles. While current analysis tools typically output single static structures or averaged contact maps, chemical probing methods like DMS capture nucleotide-resolution signals representing the full structural ensemble, which remain difficult to interpret structurally. To address this, we present MERGE-RNA, a framework that describes and outputs RNA as a structural ensemble. By modeling the physics of the experimental pipeline, MERGE-RNA learns a small set of transferable and interpretable parameters, enabling the integration of measurements across different molecules, probe concentrations, and replicates in a single optimization to improve robustness. Our model employs a maximum-entropy principle to predict thermodynamic populations, with the minimal adjustments necessary to align the ensemble with experimental data. We validate MERGE-RNA on diverse RNAs, showing that it achieves structural accuracy surpassing standard pseudo-free-energy methods and yields ensembles better recapitulating measured DMS reactivity. Applied to the *Vibrio vulnificus* adenine riboswitch, MERGE-RNA recovers the NMR-resolved conformations and their ligand-induced rearrangement, with population shifts matching the NMR-derived K_d . In a designed RNA construct for which we report new DMS data, MERGE-RNA deconvolves mixed states to reveal transient intermediate populations involved in strand displacement, dynamics invisible to methods based on enumerating a small number of structures.

Graphical abstract



Introduction

RNA molecules play central roles in cellular processes, ranging from catalysis (ribozymes and ribosomes) [1, 2] to gene regulation (including riboswitches and microRNAs) [3]. This functional versatility arises from their ability to fold into specific secondary and tertiary structures [4]. Importantly, RNA molecules exist as thermodynamic ensembles at both the secondary [5] and tertiary [6] levels, with function arising from

stable conformations, minority states, or dynamic equilibria. While tertiary structure modeling is computationally demanding [7], secondary structure thermodynamic models allow efficient computation of optimal structures [8] and base-pairing probabilities [9]. However, thermodynamic models face significant limitations from fundamental approximations: the nearest-neighbor approximation [10], the neglect of pseudo-knots [8], and the hierarchical folding assumption [11] may

Received: May 15, 2026. Revised: August 25, 2026. Accepted: August 26, 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

make it difficult to fully capture long-range interactions and context-dependent effects. To address these limitations, researchers have integrated co-evolutionary analysis or chemical probing data. Dimethyl sulfate (DMS) [12, 13] has become a powerful tool for structural analysis, providing single-nucleotide resolution data both *in vitro* and *in vivo*. DMS reacts preferentially with unpaired adenines and cytosines, with modifications detected as mutations during sequencing (mutational profiling, MaP) [13]. However, interpreting this data is challenging: measured reactivities report on structural ensembles rather than single conformations, and experimental processes introduce systematic biases from reverse transcription and sequencing artifacts. Current approaches typically convert measured reactivities into pseudo-free-energy restraints [14], but this modifies predicted ensembles even when baseline models already agree with data [15], and the coefficients used to compute pseudo-free-energies are based on heuristics rather than physical principles. Despite widespread practice, focusing on a single minimum free-energy (MFE) structure can be misleading as RNA molecules are known to adopt an ensemble of conformations, of which the MFE structure is often only a minor component. In fact, ignoring structural heterogeneity can pose problems both for interpreting RNA function and for properly modeling chemical probing experiments.

In this work we present MERGE-RNA (Multi-system Ensemble Refinement via Generalizable parameters Estimation), a unified physics-based model that describes chemical probing experiments with physically meaningful parameters and predicts RNA secondary structure ensembles. Our approach combines a standard thermodynamic folding model with explicit probe-RNA interactions characterized by two key parameters: a chemical potential (μ_+) setting the probe concentration scale, and an energetic penalty ($\Delta\mu_{\text{pairing}}$) for probe binding to paired nucleotides. We employ maximum entropy inference to determine sequence-specific soft constraints (λ_i) representing minimal adjustments for model-data consistency, avoiding artificial constraints when baseline predictions already match data. The model explicitly accounts for mutational profiling artifacts, including false positives, false negatives, and position-specific biases. Crucially, our framework enables the simultaneous analysis of data from multiple RNA sequences, probe concentrations, and experimental replicates within a unified model, leveraging the shared physical parameters to improve robustness and extract transferable experimental knowledge.

After outlining the model and inference procedure, we benchmark it on a set of well-characterized structured RNAs for which DMS probing data are available across multiple concentrations and replicates under consistent experimental conditions [16]. This homogeneity makes the dataset well suited to assess both the transferability of the inferred physical parameters and the structural accuracy of the predicted ensembles, which we benchmark against the standard pseudo-free-energy approach [14]. On these systems, we also show that no single static structure fully explains the probing data, indicating that conformational heterogeneity is required even for RNAs traditionally treated as well-folded, and that the ensembles predicted by MERGE-RNA significantly improve the agreement with probing data when compared to the traditional pseudo-free-energy approach. For a more stringent validation of the inferred heterogeneity, we then turn to the *Vibrio vulnificus* adenine riboswitch, for which NMR provides independent ground truth on the co-existing conformations and

their ligand-dependent populations [17] and which has benchmarked other deconvolution methods [18, 19]: MERGE-RNA recovers the structural details of each NMR-resolved conformer (except the pseudoknot, excluded by the thermodynamic baseline) and reproduces the ligand-dependent population shifts, matching the NMR-derived K_d . Finally, we demonstrate accurate deconvolution of mixed states in synthetic bistable constructs, and leverage experimental data to reveal intermediate, suboptimal populations involved in dynamic processes (e.g. strand displacement) that are missed by traditional approaches.

Methods

Physical model

Our model follows the experimental pipeline of chemical probing experiments, as illustrated in Fig. 1. The underlying physics remains consistent across experiments using the same reagent, probing time, and temperature, allowing us to simultaneously fit multiple experimental datasets within a unified framework. By maximizing the likelihood of the model to reproduce the observed mutation rates, our approach estimates physically meaningful parameters (detailed below) and identifies the ensemble of secondary structures most consistent with both the underlying physics and experimental evidence.

In this section we aim at introducing the physical model underlying our approach, while in the Supplementary data we provide a more in-depth description and rigorous discussion.

RNA folding and chemical probing

Folding

To model the RNA secondary structure ensemble, we utilize as baseline the ViennaRNA package [20]. The probability to observe a given secondary structure s at temperature T is defined as

$$\mathbb{P}(s) = \frac{\exp\left(-\frac{F(s)}{k_B T}\right)}{z}, \quad (1)$$

where z is the partition function, k_B is the Boltzmann constant, and the baseline $F(s) = F_0(s)$ is computed by the underlying ViennaRNA thermodynamic model, including the standard nearest-neighbor contributions (e.g. stacking interactions, hairpin, internal, bulge, and multiloop terms). To improve predictions, we refine the ensemble to align with experimental data while minimally perturbing the baseline model. Using the principle of maximum entropy [21], we optimize sequence-specific parameters (λ_i) that act as soft constraints for each nucleotide position [15]. These soft constraints modify the baseline ViennaRNA-predicted energy of pairing and remain consistent across experimental replicates and probe concentrations for a given RNA sequence. This results in a structure-dependent energy correction

$$\Delta F_{\text{opt}}(s) = \sum_{i \in \text{paired}(s)} \lambda_i, \quad (2)$$

where $\text{paired}(s)$ denotes the list of paired nucleobases in structure s .

Probe-RNA interaction

We separately model the probe reactivity in two steps [22]: a reversible physical binding, described here, followed by an irreversible chemical reaction, as detailed in the next para-




Figure 1. Schematic overview of the method. Our approach builds a physical model—illustrated on the left side of the figure—that represents the entire pipeline of chemical probing experiments. This includes RNA folding into an ensemble of secondary structures, probe binding, adduct formation, and the final mutational profiling readout. The model is trained on chemical probing data from multiple RNA sequences, probe concentrations, and experimental replicates. From these data, it learns a set of physically meaningful parameters shared across all experiments [μ_r , $\Delta\mu_{\text{pairing}}$, $p_{\text{bind}}(\text{A})$, $p_{\text{bind}}(\text{C})$, $p_{\text{bind}}(\text{G})$, $p_{\text{bind}}(\text{U})$, m_0 , and m_1 , as defined in the main text]. In addition, it estimates sequence-specific soft constraints λ_i (one per nucleotide position). Once trained, the model can predict the structural ensemble of each RNA at any probe concentration. Importantly, it can also extrapolate to zero probe concentration, providing the unperturbed structural ensemble.

graph. At non-zero probe concentrations, we describe physical probe–RNA interactions through a thermodynamic binding model with two trainable parameters (μ_r and $\Delta\mu_{\text{pairing}}$). The reference chemical potential μ_r captures the probability of an unpaired nucleotide to bind to a chemical probe at a reference DMS concentration $[\text{DMS}]_r = 1\text{M}$. The chemical potential is then corrected based on the actual probe concentration $[\text{DMS}]$ as $\mu = \mu_r + k_B T \ln \left(\frac{[\text{DMS}]}{[\text{DMS}]_r} \right)$. $\Delta\mu_{\text{pairing}}$ differentiates between paired and unpaired sites, so that the effective chemical potential for paired nucleotides becomes $\mu' \equiv \mu - \Delta\mu_{\text{pairing}}$. Since chemical probes typically exhibit reduced accessibility to base-paired regions, $\Delta\mu_{\text{pairing}}$ is expected to be positive, though the model could be straightforwardly generalized to cases where this is not true. Incorporating a penalty term for base-paired sites results in a modified partition function. This modification is implemented by applying a concentration-dependent energetic contribution to paired bases $\Delta F_{[\text{DMS}]}(s)$ (derivation is detailed in the SI Appendix).

Putting together all these terms, the corrected free energy of a given secondary structure s is

$$F(s) = \underbrace{F_0(s)}_{\text{baseline}} + \underbrace{\sum_{i \in \text{paired}(s)} \lambda_i}_{\text{site-specific corrections}} + \underbrace{\Delta F_{[\text{DMS}]}(s)}_{\text{concentration-dependent perturbation}} \quad (3)$$

from which we can compute the population of observing each structure s via Equation (1).

Chemical binding

Following the physical interaction between a chemical probe and RNA, adducts are formed with probabilities that depend on both nucleobase identity and local RNA structure, under fixed experimental conditions (e.g. temperature and probing time). In the most general framework, this yields eight possible configurations, defined by the combination of nucleotide type (A, U, C, or G) and structural state (paired or unpaired). For DMS probing, due to the position of the potentially reactive nitrogens, adenine and cytosine show negligible reactivity when base-paired, while guanine and uracil retain reduced but secondary structure-independent reactivity. Accordingly, the model is parameterized by four effective binding probabilities, corresponding to the reactive nucleotide-structure combinations: $p_{\text{bind}}(\text{A, unpaired})$, $p_{\text{bind}}(\text{C, unpaired})$, $p_{\text{bind}}(\text{G})$, and $p_{\text{bind}}(\text{U})$.

Mutational profiling and prediction of reactivity profiles

Mutational profiling

Following chemical probing, modifications are typically detected through reverse transcription. Here, we focus on mutational profiling (MaP), though only minor modifications to the model would be required for compatibility with other experimental techniques, such as RT-stop. During reverse transcription, we account for the probability of producing false positive or false negative signals through the parameters m_0 and m_1 , as detailed in the SI Appendix. Applying this framework to a given secondary structure s , we can compute the expected reactivity profile for that structure, then sum over all

secondary structures in the ensemble to obtain the expected reactivity profile for the experiment.

Handling systematic errors

Sequencing errors are known to introduce systematic biases in reactivity profiles [23], and reverse transcription can be affected by the presence of secondary structures [24]. To account for such effects, we make use of experimental controls (performed without reagent) to estimate position-dependent correction factors ϵ_i , which effectively act as a baseline adjustment to the predicted mutation rates at each nucleotide position and are computed so that the prediction at zero concentration of DMS matches the control data. These factors capture systematic biases that are sequence-specific but independent of probe concentration. This allows the model to focus on the concentration-dependent signal in reactivity, which is most informative about RNA structure and probe interactions, while ϵ_i absorbs background effects that remain constant across conditions. In other words, the sequence-dependent ϵ_i parameters model the mutation rate in absence of reagents, whereas the sequence-dependent soft constraints λ_i affect the pairing population which, in turn, controls how much the mutation rate increases when the reagent is added. As discussed in the SI Appendix, when control experiments are available the parameter m_0 is omitted, since background correction factors can be directly estimated from the control data. If controls are unavailable, m_0 is instead fitted to provide a single background correction factor shared across the sequence.

Parameter optimization

Upon computation of the expected reactivity profile for a given sequence and probe concentration, we optimize the model parameters by minimizing a loss function that quantifies the discrepancy between predicted and observed reactivities. We model per-position mutation counts as binomial: given read depth n_i and model-predicted mutation probability M_i , each read is an independent Bernoulli trial. Accordingly,

$$\mathbb{P}(\mathcal{M}_i^{\text{exp}} | n_i, M_i) = \text{Binomial}(\mathcal{M}_i^{\text{exp}}; n_i, M_i), \quad (4)$$

where $\mathcal{M}_i^{\text{exp}}$ is the experimentally observed number of mutations at position i . As discussed in the SI Appendix, the overall likelihood of the model across all nucleotide positions is the product of the individual position likelihoods. To facilitate optimization, we minimize the negative log-likelihood, yielding the following loss function:

$$\mathcal{L} = - \sum_i \ln \mathbb{P}(\mathcal{M}_i^{\text{exp}} | n_i, M_i). \quad (5)$$

The derivation and all definitions are expanded in the SI Appendix.

This optimization is implemented via a gradient-based algorithm that systematically adjusts model parameters to minimize the discrepancy between predicted and experimental reactivity profiles [25] as implemented in SciPy [26]. The parameter optimization is performed in two consecutive phases, as illustrated in Fig. 1. In the first phase, we optimize the global physical parameters that are shared across all sequences and concentrations. These include the reference chemical potential (μ_r), the binding penalty for structured regions ($\Delta\mu_{\text{pairing}}$), the nucleotide-specific adduct formation probabilities [$p_{\text{bind}}(\text{A, unpaired})$, $p_{\text{bind}}(\text{C, unpaired})$, $p_{\text{bind}}(\text{G})$, and

$p_{\text{bind}}(\text{U})$], the false negative rate parameter (m_1), and—when experimental controls are unavailable—the background correction factor (m_0). This phase optimizes a total of seven parameters (eight if no controls are available). In the second phase, we look to refine the secondary structure ensemble to find the one that best explains the experimental data. This is achieved by optimizing the sequence-specific soft constraints (λ_i) for each nucleotide position, while keeping the global physical parameters fixed.

Following parameter optimization, $\Delta F_{\text{DMS}}(s)$ is set to zero, extrapolating to zero probe concentration and thereby revealing the ensemble of secondary structures most consistent with the experimental data under native conditions.

Experimental procedures

DMS probing of a synthetic RNA construct was performed as described in the appendix. DMS-treated RNA was column-purified and subjected to reverse transcription using MarathonRT [16, 27–29]. Sequencing reads were processed with an RNAFramework-based workflow [30–32].

Results

Transferability and accuracy of the physical model

We extracted physical parameters from our unified physics-based framework after training on a set of well-characterized RNA structures previously studied using DMS chemical probing [16]. This set spans a range of structural complexity, comprising five distinct systems: (i) bacterial RNase P (type A), (ii) hc16 ligase, (iii) *Tetrahymena* ribozyme, (iv) *V. cholerae* glycine riboswitch, and (v) HCV IRES. These sequences range from ~100–400 nucleotides in length. Because all systems were probed under consistent conditions, with multiple DMS concentrations and independent replicates, this dataset is well suited for testing whether the physical parameters of the model transfer across sequences rather than being overfitted to any single system. Reference structures for hc16 ligase, *Tetrahymena* ribozyme, and *V. cholerae* glycine riboswitch were obtained from cryo-EM-guided structure determination [33], and secondary structure was annotated from the 3D structure with Barnaba (v0.1.8) [34]. HCV IRES reference structure was manually transcribed from extended data fig. 6 of Bohn et al. [16]. For bacterial RNase P (type A), although a crystal structure exists [35], sequence differences between the crystallized construct and the chemically probed sequence preclude its use as a reference.

These references are expected to be more accurate than predictions from a simple thermodynamic model, but since they represent single conformations, they cannot by themselves report on conformational heterogeneity.

Our first objective was to assess whether the physical parameters extracted from our model represent genuine physical properties of the experiment and are hence transferable to different systems, rather than being overfitted on a specific case. We performed comprehensive cross-validation fits, training only the physical parameters of the model (all the parameters except the λ_i , which are fixed to 0 for the moment) on triplets of RNA systems and testing on the remaining pairs (Fig. 2a). Loss values [as defined in Supplementary Equation (31)] are then divided by the number of data points to enable direct comparison. Obtained values of the physical parameters are reported in Supplementary Table S3 for each fit. The loss val-




Figure 2. a) Cross-validation of physical parameters across five RNA systems. Each bar represents the per-datapoint loss obtained for the indicated system after training on three systems and testing on the remaining two. Yellow bars refer to results obtained when the corresponding system was excluded from the training set. (b) Frobenius distance divided by the squared number of nucleotides between base pairing probability (bpp) matrices, quantifying structural similarity between the HCV IRES reference structure and MFE structure, bpp of the original thermodynamic ensemble, bpp obtained from standard pseudo-free-energy method (Deigan) and bpp from MERGE-RNA. Deigan results are reported using different DMS concentrations, and MERGE-RNA results are reported for different replicates (independent fits). Deigan and MERGE-RNA results are reported for different baseline thermodynamic models (Turner and Andronescu). Solid bars for Deigan method use (m , b) coefficients determined with a standard jackknife procedure, while the caps mark its theoretical best performance (see text for definition). Colors as in panel (d). (c) Mutation rates at 57 mM DMS concentration after background subtraction for adenine and cytosine in HCV IRES, grouped by pairing status in the reference structure. Each point is colored according to the pairing probability predicted by MERGE-RNA. Strong disagreements (positions where MERGE-RNA predicts pairing probability below 25% or above 75% while the reference indicates the opposite) are highlighted. (d) Squared Pearson correlation coefficient between predicted base pairing probabilities and observed mutation rates, computed separately for A and C and averaged. The caps show the theoretical best performance of the Deigan method. Labels as in panel (b).

ues obtained across all sets are consistent both in training and in test, demonstrating that our model generalizes on multiple systems and successfully captures the underlying principles governing chemical probing experiments with a small set of parameters.

Inference of optimal ensemble

Before turning to the full inference of the per-nucleotide soft constraints, we verified with a simple controlled experiment, reported in the SI Appendix (“Importance of conformational heterogeneity in structured RNAs,” [Supplementary Figs S1 and S3–S5](#)), that a proper balance between reference structural information and ensemble heterogeneity is needed to reproduce chemical-probing data: applying a uniform reference-derived soft constraint of increasing magnitude reveals a

“sweet spot” where the model partially incorporates the reference while preserving the conformational diversity captured by the experiment, beyond which performance degrades. Building on this insight, instead of borrowing from the reference structure, we inferred the best sequence-specific soft constraints (λ_i parameters) using our model. This resulted in substantially improved agreement with experimental data (blue line in [Supplementary Figs S1 and S3–S5](#)), as expected given that the experimental data are used to infer the soft constraints.

To validate our model independently, we computed the Frobenius distance between the reference structure and base-pairing probability matrices obtained by the baseline ViennaRNA MFE, ViennaRNA BPP, and by standard pseudo-free-energy approach as in Deigan *et al.* [14] (Fig. 2b and [Supplementary Figs S6–S8](#)). For the Deigan approach, rather

than using the literature coefficients (m, b) = (1.8, -0.6), for each DMS concentration we selected the best-fitting slope and intercept through a procedure analogous to `rfr-jackknife` [18, 30], using a leave-one-out approach to avoid fitting these coefficients on the same structural data we are optimizing, as detailed in the SI Appendix (Refitting the Deigan slope and intercept, see [Supplementary Tables S1 and S2](#)). Solid bars report results using these coefficients, while the dashed caps mark the best value the Deigan functional form can reach (its ceiling), i.e. the value obtained by choosing (m, b) to maximize the agreement with the metric itself. To test robustness with respect to the choice of thermodynamic baseline, we repeated the analysis using both the standard Turner2004 [36, 37], and the Andronescu2007 parameters [38]. For ensemble-based metrics, our model improves significantly over Deigan's in three of the four systems studied, and in general remains comparable to the theoretical best results achievable by Deigan's method, obtained when it is fitted to the reported test data ([Supplementary Figs S6–S8](#), in SI Appendix). This is remarkable, considering that MERGE-RNA is not optimized for structure prediction at any stage of its pipeline. We also report in the SI Appendix ([Supplementary Figs S9–S12](#)) metrics more commonly reported for single-structure prediction: positive predictive value, sensitivity, and the same Frobenius distance computed on the MFE structure rather than on the ensemble in order to isolate the effect of selecting only the MFE from the ensemble. When evaluating the MFE structure, the Deigan approach is favored. This is expected, since the Deigan pseudo-free-energy approach is explicitly designed to improve single-structure prediction, whereas MERGE-RNA is optimized to infer an ensemble consistent with the probing data rather than to maximize agreement of any single structure with the reference. Arc plots (Fig. 3) illustrate how MERGE-RNA's individualized fitting preserves ensemble diversity (blue) while recapitulating the reference structure (red) more accurately than the standard thermodynamic ensemble (green). The same comparison with the Deigan ensemble superimposed in place of the unconstrained thermodynamic one, for each of the four systems, is reported in the SI Appendix ([Supplementary Figs S17–S20](#)). Similar improvements were observed across all tested RNA systems, confirming the broad applicability of our approach.

MERGE-RNA suggests localized structural dynamics supported by chemical probing data

Having established structural accuracy by comparison with experimentally resolved reference structures, we now examine where our predictions diverge from the static reference. Since our model predicts a dynamic ensemble rather than a single conformation, divergence is expected and informative: chemical probing experiments inherently report on the dynamic ensemble of the RNA molecule, which a static reference structure alone cannot fully capture. In Fig. 2c, we report mutation rates after control subtraction for adenine and cytosine, grouped by pairing status in the reference structure. As expected, paired positions show, on average, lower mutation rates than unpaired ones. Each point is colored according to the pairing status predicted by MERGE-RNA. The model's predictions correlate well with mutation rates, demonstrating its capability to infer and model pairing dynamics from probing data.

It is instructive to examine positions where our predictions disagree with the reference structure. When mutation rates are consistent with the reference structure, MERGE-RNA's predictions align with it. For positions with intermediate mutation rates, MERGE-RNA predicts intermediate pairing probabilities, suggesting conformational heterogeneity in solution. This is particularly informative: a static reference structure classifies each position as either paired or unpaired, and thus cannot account for intermediate reactivity. MERGE-RNA provides an alternative explanation for this discrepancy by proposing alternative conformations that together explain the observed mutation rates. Where probing data and reference structure are in contrast, MERGE-RNA captures new pairings absent from the reference but supported by probing, or identifies reference-annotated pairs that show elevated mutation rates, suggesting they are frequently unpaired in the dynamic ensemble.

The number of “strong disagreement sites” (positions where MERGE-RNA predicts pairing probability below 25% or above 75%, while the reference indicates the opposite for A and C nucleotides, highlighted in Fig. 2c) is modest, ranging from 0 to 13 across the systems studied here. Detailed analysis (SI Appendix, [Supplementary Figs S13–S16](#)) reveals that most correspond to small corrections of existing structural elements or formation/dismantling of short motifs (e.g. 3–4 bp stems), indicating localized inconsistencies rather than wide structural disagreements. Some level of disagreement is expected, as the reference is a single structure obtained under specific conditions (e.g. in a crystal or in a different buffer), whereas probing reflects the ensemble populated under the experimental conditions of the probing itself, which in principle can be very different; MERGE-RNA adjusts the reference only where the probing data require it, a behavior that becomes valuable when the reference is partial or originates from conditions that differ from those of the probing experiment. The correspondence between base-pairing probability and observed mutation rates can be quantified by computing their Pearson correlation (Fig. 2d): MERGE-RNA shows a noticeably higher correlation than the one obtained by the traditional pseudo-free-energy approach across all systems. This gap is not a matter of how the (m, b) Deigan coefficients are chosen. In fact, MERGE-RNA exceeds the results obtained with the coefficients cherry-picked to maximize the displayed metric by $\Delta R^2 = +0.18$ in the worst case and $+0.28$ on average ($1.4\times$ and $1.9\times$, respectively) across all systems and concentrations. This indicates that the gap stems from the limited functional form and degrees of freedom of the two-parameter pseudo-energy scheme, rather than from a suboptimal choice of coefficients or the data.

In summary, with our model (i) the inferred physical parameters are generalisable across systems (Fig. 2a); (ii) we show that there is an optimal balance between structural accuracy and ensemble heterogeneity, and inference of sequence-specific corrections further improves agreement with experimental data ([Supplementary Figs S1 and S3–S5](#)); (iii) the resulting structural ensembles more accurately capture native conformations compared to standard thermodynamic approaches (Fig. 2b); (iv) the resulting ensemble correlates significantly better with the observed mutation rates than traditional methods (Fig. 2d); and (v) comparing MERGE-RNA's results to the reference structures reveals that agreement reflects probing data consistency (Fig. 2c). In cases where reference models are inconsistent with experiment, MERGE-RNA goes beyond




Figure 3. Arc plot visualization comparing the secondary structures of HCV IRES: reference experimental structure (red), ensemble obtained from the thermodynamic model (green), and ensemble obtained from MERGE-RNA (blue).

the static reference by identifying localized discrepancies and proposing alternative conformations that better explain the data. These results are consistent across all tested RNA systems (SI Appendix [Supplementary Figs S6–S8](#)). However, in the absence of independent experimental evidence for ensemble populations, we cannot quantitatively validate the inferred alternative conformations here; in particular, we cannot rule out that MERGE-RNA over-estimates ensemble heterogeneity by accommodating intermediate reactivities with intermediate pairing probabilities. For this reason, we next turn to the adenine riboswitch, where independent NMR measurements provide ground truth on both the co-existing conformers and their ligand-dependent populations.

Application on a known case of functional structural rearrangement: *V. vulnificus* add adenine riboswitch

Having established structural accuracy and parameter transferability on structured RNAs, we now turn to the validation of the model's ability to deconvolve co-existing conformational states. For this purpose we use the *V. vulnificus* add adenine riboswitch, which undergoes a well-characterized ligand-induced conformational change that regulates translation, and provides independent experimental ground truth unavailable for the structured RNAs above: its ligand-free and ligand-bound conformations and the binding dissociation constant have been resolved by NMR [17], and the system has served

as a benchmark for other ensemble-deconvolution methods including DRACO [18] and DANCE-MaP [19].

We analyzed previously deposited DMS-MaP data acquired across a range of adenine concentrations, from 0 to 500 μM [19]. Since the experiment provides only a single DMS concentration (170 mM) alongside the no-DMS control, we resorted to the linear version of the model, which assumes that the probe perturbs pairing probabilities negligibly and that the probe concentration is well below saturation, which corresponds to the regime of low chemical potential of probe–RNA interactions, avoiding fitting a nonlinear model to only two points. Because our model does not explicitly account for the reduction in mutation rate of unpaired bases caused by ligand contacts and does not model pseudoknots, we masked, respectively, the ligand-binding pocket (J) and the loops (L2 and L3), to prevent the model from fitting these trends. For the same reason, the physical parameters of the model were fit using only the ligand-free experiments; we then optimized the λ_i parameters separately for each condition to infer the structural ensemble at every adenine concentration.

We present in Fig. 4 the inferred ensembles at the two limiting conditions, ligand-free (0 μM adenine) and ligand-saturated (500 μM adenine). The full concentration series is provided in [Supplementary Fig. S21](#) and smoothly interpolates between these two points. Panel (a) shows the arc-plot of the base-pairing probabilities, and panels (c1–2, d1–2, e1–2, f1–2) provide zoomed-in views of the predicted BPP matrices




Figure 4. MERGE-RNA captures ligand-induced structural rearrangements in the *V. vulnificus* add adenine riboswitch. (a) Arc plot of the inferred base-pairing probabilities at 0 μM (top) and 500 μM (bottom) adenine; colors indicate base-pairing probability. (b) Estimation of the ligand-binding dissociation constant (K_d) from the inferred pairing probability of a representative base pair in the P1 helix [highlighted with a purple box in panels (c1–2)], which is stabilized upon ligand binding. (c1–2, d1–2, e1–2, f1–2) Zoomed-in views of the predicted BPP matrices for the P1, P2, P3, and SD domains in the absence (top row) and presence (bottom row) of adenine. Color scale as in panel (a).

for the P1, P2, P3, and SD domains in the absence (top row) and presence (bottom row) of adenine.

The model recovers the molecule's structural details and the expected ligand-induced rearrangement: the P1 helix is stabilized upon ligand binding; the P2 region transitions from a mixture of alternative pairings (P2/P2-alt) to a single stabilized P2 conformation; the P3 helix remains stable in both conditions; and the SD domain (which harbors the Shine–Dalgarno sequence) is destabilized upon ligand binding, consistent with the known mechanism of translational regulation by this riboswitch. To validate against the independent NMR data, similarly to what has previously been done by DANCE-MaP authors, we estimated the dissociation constant from the ligand-dependent pairing probability predicted by MERGE-RNA of a representative base pair in the P1 helix (17–79, highlighted with a purple box in c1–2) that is stabilized when ligand is present. We tracked this pairing probability across all adenine concentrations and fit it to a Hill curve,

$$y(x) = y_{\min} + (y_{\max} - y_{\min}) \frac{x^n}{K^n + x^n},$$

where x is the adenine concentration (in μM), the 0- μM point is pinned as the lower plateau, and K and n are estimated by bounded nonlinear least squares. The midpoint concentration K is $K_d = 1.7 \mu\text{M}$ for both the independent fits, in good agreement with the NMR-derived value of $1.6 \mu\text{M}$ [17] and the previous prediction of $1.5 \pm 0.3 \mu\text{M}$ reported in Olson *et al.* [19].

Taken together, MERGE-RNA reproduces essentially all features of the adenine riboswitch that are accessible in NMR experiments [17]: the three resolved structures, their ligand-dependent rearrangement, and the dissociation constant K_d . The only structural element that is, by construction, out of reach for our model is the pseudoknot involving L2 and L3, since the underlying thermodynamic baseline does not allow pseudoknotted pairings. For comparison, we repeated the same analysis using the Deigan pseudo-free-energy approach (Supplementary Figs S22 and S23). While the optimized Deigan model predicts a similar apparent K_d , it substantially over-stabilizes the P1 helix (predicted population always > 97%) and fails to reproduce the ligand-dependent destabilization of the Shine–Dalgarno sequence.

Beyond the conformations resolved by NMR, MERGE-RNA also predicts intermediate states along the P1 $\leftrightarrow$ SD competition (P4 in Reining *et al.* [17]), in the form of a gradient of pairing probabilities within the two competing helices rather than only the fully-formed alternatives (Supplementary Fig. S21). We cannot directly validate this prediction against the NMR data, for two reasons. First, the construct used in the probing experiment is longer than the one resolved by NMR: it carries additional upstream nucleotides that can pair with the 3' side of P1, stabilizing a longer version of the helix (“P1-extended” in Olson *et al.* [19]). Second, NMR structure determination tends to discretize the conformational space into a small number of representative structures, so its absence in the NMR ensemble cannot be taken as evidence against the existence of intermediate populations. That said, the exis-

tence of intermediates is consistent with the NMR conformers themselves, which resolve P1 with 0, 4, and 7 base pairs in apo-B, apo-A, and holo, respectively, three discrete points along what is plausibly a continuous distribution of partially formed P1 helices. Read-clustering methods such as DRACO [18] and DANCE-MaP [19] similarly discretize the ensemble into a small number of representative structures and therefore struggle to describe such a continuum, whereas MERGE-RNA naturally captures it.

We also note a methodological distinction: DRACO and DANCE-MaP exploit the co-occurrence of modifications within individual reads to assign each read to a discrete cluster, which requires deep coverage and effectively discards reads with zero or a single modification. MERGE-RNA instead operates on the cumulative mutation profile alone, so every read contributes to the inference and the per-position pairing probabilities are recovered without any per-read correlation analysis. The population gradient observed in the P1 $\leftrightarrow$ SD competition is reminiscent of a strand-displacement. Interestingly, strand-displacement has been suggested as a possible mechanism to control the expression-platform conformation in transcriptional riboswitches [39]. Strand-displacement is also observed in a more controlled setting using new DMS data below (see the “A putatively bistable sequence exhibits a heterogeneous ensemble with strand displacement” section). We also note that although the data were collected with four-base DMS [40], our framework currently treats only A and C as pairing-sensitive reactive bases. Although this reduces the information used relative to the original four-base DMS analysis, the A/C-sensitive model is sufficient to recover the main ligand-dependent rearrangement. Extending the model to all four bases is a natural direction for future work.

Overall, this case study shows that, from chemical probing data alone, MERGE-RNA resolves the structural detail of multiple competing conformers identified by NMR in a biologically relevant RNA. Additionally, it captures the ligand-induced shift in the conformational ensemble, quantitatively matching the NMR-derived dissociation constant. Importantly, conformational ensembles of this riboswitch were here controlled by ligand concentration alone, which allowed a series of experiments performed in similar conditions. The more challenging case of a thermoswitch, which requires combining experiments at different temperatures and acquisition times [41, 42], is reported in the Supplementary data (see [Supplementary Figs S2 and S24–S29](#)).

Deconvolution of mixed structural states on synthetic data

To rigorously test the model’s ability to quantitatively predict secondary-structure populations in heterogeneous ensembles, we generated synthetic datasets for an in-house designed bistable RNA carrying two mutually exclusive hairpins. We deliberately designed this construct so that we could also assay the very same sequence experimentally (discussed in the next section); accordingly, we appended primer binding sites at the 3′ and 5′ ends. Concretely, the final construct has the architecture

$$5′\text{--PBS1--A1--L1--B--L2--A2--PBS2--}3′, \quad (6)$$

where PBS1 and PBS2 are distinct primer binding sites, A1 and A2 are identical copies of the same sequence, and B is comple-

mentary to A1/A2. In this design, A1 and A2 compete to pair with B, producing two mutually exclusive hairpins (helix 1: A1:B; helix 2: A2:B). L1 and L2 are short linkers designed to act as tetraloops for the hairpins.

While the original design targeted roughly equal populations of the two helices, the thermodynamic model allows us to artificially bias the ensemble to favor either helix. We then simulated mutation profiles for the two cases—when either helix is always formed—and combined these profiles in known proportions to create synthetic datasets with varying helix populations, ranging from 0 to 1 in increments of 0.2. To generate realistic mutation profiles, we used the physical parameters inferred from our analysis on experimental data from biologically relevant RNAs (the “Transferability and accuracy of the physical model” section).

We then trained our model on these synthetic datasets. Following the procedure established earlier, we first performed a joint fit of the physical parameters across all datasets, starting from random initializations. Subsequently, we optimized the soft constraints λ_i for each dataset individually. The results, presented in Fig. 5, show that the model accurately recovers the expected base-pairing probabilities for both helices across all synthetic datasets. The method is robust to the presence of additional base pairs formed by the random primer binding sites. Overall, these synthetic data tests validate the model’s ability to recover structural populations consistent with the known ground truth (Fig. 5c).

A putatively bistable sequence exhibits a heterogeneous ensemble with strand displacement

To validate our model on experimental data for the same system, we collected DMS-MaP data on the putative bistable RNA construct described in the previous section. The data were collected at multiple DMS concentrations, 0–100 mM at steps of 20 mM, two replicates each. The obtained mutation profiles are visualized in Fig. 6a. We then trained our model on these experimental datasets, following the same procedure as in the previous section. Note that the sequence, as discussed earlier, was designed to adopt two mutually exclusive helices. We then applied our model to infer the structural ensembles from the data.

Similarly to what we have done in a previous paragraph, to challenge our model’s ability to recover the correct structural ensemble from an inaccurate thermodynamic baseline [F_0 term in Equation (3)], we performed the fitting procedure using two different baseline models. We first used the default Turner 2004 parameters [36], which were employed in the original sequence design. We then repeated the analysis with the Andronescu 2007 parameters [38], which predict a substantially different initial ensemble with populations of ~20% helix 1 and 80% helix 2, in stark contrast to the roughly 50–50 distribution predicted by the Turner 2004 parameters, as illustrated in Fig. 6b, dotted lines.

The results of our model fitting are presented in the same panel as solid lines. Remarkably, the two predictions of MERGE-RNA are consistent with each other, regardless of the initial thermodynamic parameter used, signaling robustness of the prediction with respect to the baseline parameters. In both cases, the model adjusts the ensemble of structures so that the resulting ensemble slightly favors helix 1, in accordance with the lower mutation rates observed in the experimental data (Fig. 6a).




Figure 5. Model accurately deconvolves mixed structural states from synthetic data. **(a)** Base-pairing probabilities inferred from fitting on synthetic data for a bistable RNA sequence. The data were generated from an RNA construct with two competing helices (helix 1 and helix 2), where the ground-truth population of helix 1 was systematically varied from 0 to 1 with steps of 0.2. The model's predicted pairing probabilities are shown for each case. **(b)** Arc plots of the inferred structural ensemble for two limit cases: when helix 1 is dominant (gold) and when helix 2 is dominant (black). **(c)** The helix population inferred by the model is plotted against the ground-truth population. The inferred population is quantified as the median of the base-pairing probabilities over the respective helical regions, highlighted in panels (a) and (b).

Crucially, MERGE-RNA does more than adjust the populations of the two competing helices. It reveals a gradient in base-pairing probabilities within regions A1 and A2, where nucleotides closer to the central domain (B) are more likely to be paired, reflecting the slope observed in the experimental mutation profiles (Fig. 6a). Beyond the two fully-formed helices, the model resolves intermediate states in which B is partially paired with both A1 and A2, as schematically illustrated in Fig. 6c. This mechanism, known as strand displacement (or strand invasion) [43], has been observed in biological systems and is consistent with the observed mutation profile: regions nearest the central domain are expected to pair first during displacement, and accordingly show reduced reactivity closer to the loops (Fig. 6a). The presence of such intermediate states is expected if the two loops are sufficiently stable. Analysis of loop co-occupancy (Supplementary Table S4) demonstrates capture of intermediate states missed by standard models. While baseline ensembles rarely allow both loops to form ($< 0.2\%$), the model predicts simultaneous pairing in 42%–57% of structures, consistent with the observed mutation profiles and reflecting the strand-displacement mechanism in which the central domain transiently pairs with both competing helices. We repeated the same analysis with the optimized Deigan pseudo-free-energy approach (Supplementary Fig. S30). Although the method partially captures the strand-displacement behavior, the inferred ensemble depends strongly on the prob-

ing concentration and on the baseline ensemble. Overall, these results demonstrate MERGE-RNA's ability to capture a continuous spectrum of low-population intermediates that standard thermodynamic models fail to predict, independently of the baseline thermodynamic model.

Discussion

In this work, we have introduced MERGE-RNA, a physics-based framework for the analysis of chemical probing data that addresses fundamental limitations of conventional approaches. Our approach moves beyond the use of experimental data as simple pseudo-energetic restraints [14] and instead models the entire experimental pipeline, from the physical interaction of the probe with the RNA to the final mutational readout. This enables the determination of transferable, physically meaningful parameters and the integration of multiple datasets for inferring the native structural ensemble. Like all current thermodynamics-based ensemble prediction methods, MERGE-RNA ultimately relies on the underlying nearest-neighbor thermodynamic model as its prior. The inferred soft constraints compensate for systematic discrepancies between this prior and the experimental data, but they cannot fully eliminate limitations arising from an incomplete physical description of RNA folding. Improvements in thermodynamic models are therefore expected to translate directly




Figure 6. Ensemble inference on experimental data for a putative bistable RNA exhibits evidence of strand displacement. **(a)** Per-nucleotide DMS-MaP mutation profiles for adenines and cytosines across 0–100 mM (20 mM steps; two replicates each) for a putative bistable RNA. We leave out of the plot the primer binding sites for better visualization. Two regions [A1 on the left and A2 on the right, as defined in Equation (6)] are competing to pair with the central region (B). They show a slope in the reactivity profile that is compatible with a strand displacement mechanism. **(b)** Pairing probabilities obtained from the ensemble predicted by MERGE-RNA (solid lines) after fitting on experimental data, with Turner 2004 [36] (crimson) or Andronescu 2007 [38] (steel) thermodynamic models as baselines. MERGE-RNA is able to fix both the baseline predictions (dotted lines) and retrieve pairing probabilities consistent with the observed mutation rates. **(c)** Schematic depiction of the ensemble of structures as predicted by the baseline thermodynamic model (above) and MERGE-RNA fitted on the experimental data (below). While the thermodynamic model predicts only the two configurations with fully-formed helices, MERGE-RNA captures a wealth of intermediate states where strand displacement occurs.

into improved ensemble predictions within our framework. Applied to a homogeneous set of structured RNAs probed across multiple DMS concentrations and replicates, the model yields physical parameters that transfer across sequences and produces ensembles that agree with deposited structures more closely than the standard pseudo-free-energy approach. On the same systems, the sweet-spot behavior under uniform soft constraints, the improved correlation with mutation rates, and the localized disagreements between deposited structures and probing signals all point to the same conclusion: our results suggest that even RNAs typically regarded as well-folded may not be fully described by a single static structure under the conditions of the probing experiment. A more direct test of the inferred heterogeneity is provided by the *V. vulnificus* adenine riboswitch, where NMR offers independent ground truth on the co-existing conformers and where other deconvolution methods have been benchmarked; here MERGE-RNA recovers the resolved conformations and reproduces the ligand-dependent population shifts, including a ligand-dependent structural midpoint in agreement with the NMR-derived K_d . MERGE-RNA predicts a somewhat smaller ligand-induced increase in the P1 population than reported from the NMR and DANCE-MaP analyses. However, we regard this discrepancy to be minor compared with the overall agreement on the biologically relevant conformational transition considering the following factors: (i) the probing experi-

ment was performed on a longer construct, which may have altered the relative populations of the two conformers; (ii) NMR and DANCE-MaP assume the existence of only a few discrete states, while MERGE-RNA allows for a more heterogeneous ensemble; and (iii) other differences in the experimental conditions and inherent uncertainties in both the NMR and MERGE-RNA populations. Interestingly, the Deigan pseudo-free-energy approach yields a similar apparent K_d , but substantially over-stabilizes the P1 helix and fails to reproduce the accompanying destabilization of the Shine–Dalgarno sequence required for the regulatory conformational switch. Together with the deconvolution results on synthetic and experimental bistable constructs, these findings show that the framework can quantitatively characterize conformational populations, including the intermediates that arise in processes such as strand displacement and co-transcriptional folding.

We note that our model explicitly includes all four nucleotides even though DMS exhibits structure-dependent reactivity primarily with A and C. Including the complete sequence leads to more robust estimates of the physical parameters, ultimately improving the accuracy of inferred structural ensembles. In principle, it could be extended to model pair-dependent reactivity in U and G nucleotides so as to fully exploit information reported by four-base DMS [40]. The separation of RNA thermodynamics (baseline model and λ_i) from probe-specific physics (all of the other parameters) enables

future multi-probe extensions. One could analyze data from CMCT [44], SHAPE [45–47], hydroxyl radical footprinting [48], and ETS [49] simultaneously. IPANEMAP [50] showed improvements from multi-probe integration but relies on empirical conversions. In our framework, each probe would be modeled with its own parameters, while each RNA molecule would have specific λ_i values, leveraging orthogonal chemical information.

A central strength is the ability to avoid over-constraining: standard pseudo-free-energy approaches [14] modify the energy landscape globally, often altering already-correct regions. Similarly to Washietl *et al.* [15], our framework infers minimal soft constraints, ensuring modifications arise only where data support deviation, and resulting in a lower loss (see [Supplementary Fig. S1](#)). In addition, our physics-based model naturally enables the combination of multiple datasets in a single fitting process. Among ensembles compatible with experiment, the maximum entropy principle chooses the one closest to the prior, and modifications of the ensemble are here regularized by bounding soft constraints. Alternative regularization approaches might be explored [51, 52].

Several existing methods also aim to extract information about RNA structural ensembles from chemical probing data, although they address related rather than identical inference problems. Existing ensemble-deconvolution pipelines such as DREEM [53], DRACO [18], and DANCE-MaP [19] cluster sequencing reads with multiple co-occurring modifications and assign each cluster to a discrete structure. By capitalizing on per-read correlation patterns, these methods recover information about alternative conformations but at the cost of discarding the abundant reads with zero or single modifications and demanding deep coverage to resolve low-probability states. Alternative ensemble fitting approaches include Rsample [54] and SLEQ [55]. Rsample employs a sample-and-cluster approach, reweighting the Boltzmann ensemble through pseudo-free-energy adjustments resembling our soft constraints, though without physics-grounded calibration, but rather adopting an empirical strategy in a non self-consistent fashion. SLEQ uses a sample-and-select strategy rooted in the maximum parsimony principle to identify the sparsest set of structures explaining modification patterns, differing from our maximum entropy approach designed to minimally perturb the thermodynamic model. Our framework calibrates against averaged mutational fractions, ensuring every read (even negatives) contributes signal in sparse-modification regimes. This distinction is illustrated by the adenine riboswitch analyzed in this work, for which we use exactly the DMS dataset originally introduced in the DANCE-MaP study [19]. Despite relying only on the averaged mutation profile rather than read-level co-occurrence information, MERGE-RNA recovers the same ligand-dependent structural rearrangement reported by DANCE-MaP. The trade-off is loss of explicit per-read co-variation, though reads could in principle be integrated to fine-tune populations. In fact, our framework enables the estimation of the likelihood of generating each read, which could be used to assign reads to structures and further refine population estimates, though we leave this as a future extension. Read-level clustering would struggle to distinguish multiple similar conformations and represent continuous ensembles with many intermediates (e.g. strand-displacement pathways). Furthermore, alternative methods assume each read originates from a single structure, which may not hold for RNAs undergoing rapid dy-

namics. Our framework models ensemble-averaged behavior directly, accommodating fast interconversions at equilibrium. Consequently, our framework captures both discrete alternative structures and continuous ensembles enabling interpretation of dynamic processes. Strikingly, even without individual reads, the model reconstructs correlations and measures minor substate populations, which would be missed in methods that enumerate a small number of representative structures.

A key feature is predicting the structural ensemble at any probe concentration, especially extrapolating to zero concentration for the native state. In our analyses, base-pairing probability dependence on concentration was modest, introducing partial parameter under-determination (e.g. increasing μ mimics decreasing p_{bind}). This does not compromise robustness or transferability (Fig. 2a) and can be mitigated by widening the concentration range; Bayesian treatment could be used to quantify identifiability. At the same time, the model is equipped to handle systems where probes induce significant non-linearities [22] from rearrangement, denaturation, or saturation. Standard protocols recommend avoiding these effects by reaching the single-hit kinetics [56], which might be difficult or impossible for large RNAs, and intrinsically eliminates correlation information. Explicitly modeling these concentration-dependent perturbations might enable a more faithful representation of native conformational landscapes.

MERGE-RNA is intended for detailed ensemble inference on selected RNAs rather than transcriptome-wide analysis. Its optimization procedure is naturally more computationally demanding than a single thermodynamic folding calculation because it repeatedly evaluates partition functions during likelihood optimization. In the current implementation, the systems studied here required from ~ 1 – 2 h for the smaller RNAs (e.g. the adenine riboswitch and cspA) to ~ 1 – 2 days for the largest benchmark systems on a single workstation. Nevertheless, the framework scales well to datasets comprising multiple measurements of the same RNA or collections of RNAs acquired under comparable experimental conditions, where the global physical parameters can be inferred jointly. Extending the approach to transcriptome-wide applications will require more efficient optimization strategies. Such strategies are currently under development and are expected to substantially improve the scalability of the framework. An additional extension concerns RNA–protein interactions that confound probing readouts. Known contacts could be excluded from the loss and constrained against internal pairing; unknown interaction sites could be flagged by localizing discrepancies between predicted and observed profiles.

We note that validation of RNA structure prediction methods typically relies on reference structures from crystallography, cryo-EM, or covariance analysis, each with inherent biases. Crystals often fail to form for conformationally heterogeneous molecules, selecting against dynamic ensembles. Cryo-EM reconstructions typically use only a small fraction of recorded particles, with difficult-to-estimate effects on captured dynamics. Covariance-based methods reflect evolutionarily conserved, biologically functional structures, which may represent only a subset of conformations observed *in vitro*. These considerations highlight the value of probing-based ensemble methods as a complementary approach to traditional structure determination.

We acknowledge important limitations. The underlying ViennaRNA framework does not account for pseudoknots, critical for many RNAs. More generally, secondary structure alone

may not suffice to predict DMS reactivities, which can depend on 3D features [57]. Possible factors neglected here include solvent accessibility, backbone geometry, stacking context, conformational dynamics, tertiary contacts, and the local electronic environment of the reactive nitrogen. In particular, low reactivity should not always be interpreted as evidence of canonical base pairing, since nucleotides may also be protected by tertiary contacts, non-canonical interactions, solvent accessibility, or other structural features not represented in the current model. Consequently, the present implementation may attribute part of this protection to secondary structure, potentially overestimating localized structural heterogeneity in some regions. Despite these limitations, we here show that MERGE-RNA can be used in a wide range of contexts, systematically increasing the correlation between base-pairing probability and observed DMS reactivity and recapitulating reference RNA conformational ensembles. Finally, while physical parameters transfer across RNAs under constant conditions, transferability between varying conditions (temperature, buffer, probing duration) is not guaranteed; we recommend fitting similar-condition data together. Broader applicability would require modeling parameter dependence on experimental variables.

Acknowledgements

We thank Catherine Eichhorn, Anthony Mustoe, and Joseph Yesselman for useful discussions. Anthony Mustoe is also acknowledged for reading a preliminary version of this manuscript and providing useful suggestions.

Author contributions: Gi.S., Gu.S., and G.B. conceived and designed the study. Gi.S., Gu.S., and G.B. developed the theoretical modeling and statistical framework, and interpreted the data and the results. Gi.S. implemented MERGE-RNA, led the primary data analysis, processed sequencing data, performed model benchmarking, and prepared the figures. G.B. performed exploratory analyses and contributed methodological input. J.L. and R.P.S. designed and performed the DMS-MaP experiments. G.B., Gu.S., and R.P.S. supervised the research. All authors wrote the manuscript, revised it critically, and approved the final version.

Supplementary data

Supplementary data is available at NAR online.

Conflict of interest

None declared.

Funding

Gi.S. acknowledges financial support from the European Molecular Biology Organization (EMBO) through the EMBO Scientific Exchange Grant 11403. R.P.S. acknowledges the interdisciplinary Thematic Institute IMCBio+, as part of the ITI 2021–2028 program of the University of Strasbourg, CNRS and Inserm, IdEx Unistra (ANR-10-IDEX-0002), SFR1-STRAT'US (ANR 20-SFR1-0012), and EUR IMCBio (ANR-17-EURE-0023) under the framework of the French Investments of the France 2030 Program. G.B. acknowledges support from the Italian Ministry of University and Research (MUR) through the Fondo Italiano per la Scienza (FIS 3),

Grant No. FIS-2024-01745, RNAScale. Funding to pay the Open Access publication charges for this article was provided by the MUR Grant FIS-2024-01745.

Data availability

The MERGE-RNA source code is available at <https://github.com/giusSacco/MERGE-RNA> and has been archived on Zenodo under DOI: 10.5281/zenodo.20059549. The repository includes scripts and instructions to reproduce the fits, run the analyses reported here, and generate the plots. Sequencing data for all experiments performed in this work have been deposited in the European Nucleotide Archive (ENA) under accession number PRJEB113043.

References

1. Doudna JA, Cech TR. The chemical repertoire of natural ribozymes. *Nature* 2002;418:222–8. <https://doi.org/10.1038/418222a>
2. Ramakrishnan V. Ribosome structure and the mechanism of translation. *Cell* 2002;108:557–72. [https://doi.org/10.1016/S0092-8674\(02\)00619-0](https://doi.org/10.1016/S0092-8674(02)00619-0)
3. Morris KV, Mattick JS. The rise of regulatory RNA. *Nat Rev Genet* 2014;15:423–37. <https://doi.org/10.1038/nrg3722>
4. Butcher SE, Pyle AM. The molecular interactions that stabilize RNA tertiary structure: RNA motifs, patterns, and networks. *Acc Chem Res* 2011;44:1302–11. <https://doi.org/10.1021/ar200098t>
5. Bose R, Saleem I, Mustoe AM. Causes, functions, and therapeutic possibilities of RNA secondary structure ensembles and alternative states. *Cell Chem Biol* 2024;31:17–35. <https://doi.org/10.1016/j.chembiol.2023.12.010>
6. Ken ML, Roy R, Geng A *et al.* RNA conformational propensities determine cellular activity. *Nature* 2023;617:835–41. <https://doi.org/10.1038/s41586-023-06080-x>
7. Bu F, Adam Y, Adamiak RW *et al.* RNA-puzzles round V: blind predictions of 23 RNA structures. *Nat Methods* 2025;22:399–411. <https://doi.org/10.1038/s41592-024-02543-9>
8. Nussinov R, Jacobson AB. Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proc Natl Acad Sci USA* 1980;77:6309–13. <https://doi.org/10.1073/pnas.77.11.6309>
9. McCaskill JS. The equilibrium partition function and base pair binding probabilities for RNA secondary structure. *Biopolymers* 1990;29:1105–19. <https://doi.org/10.1002/bip.360290621>
10. Tinoco I, Borer PN, Dengler B *et al.* Improved estimation of secondary structure in ribonucleic acids. *Nature New Biol* 1973;246:40–1. <https://doi.org/10.1038/newbio246040a0>
11. Tinoco I, Bustamante C. How RNA folds. *J Mol Biol* 1999;293:271–81. <https://doi.org/10.1006/jmbi.1999.3001>
12. Zaug AJ, Cech TR. Analysis of the structure of Tetrahymena nuclear RNAs *in vivo*: telomerase RNA, the self-splicing rRNA intron, and U2 snRNA. *RNA* 1995;1:363–74.
13. Kubota M, Tran C, Spitale RC. Progress and challenges for chemical probing of RNA structure inside living cells. *Nat Chem Biol* 2015;11:933–41. <https://doi.org/10.1038/nchembio.1958>
14. Deigan KE, Li TW, Mathews DH *et al.* Accurate SHAPE-directed RNA structure determination. *Proc Natl Acad Sci USA* 2009;106:97–102. <https://doi.org/10.1073/pnas.0806929106>
15. Washietl S, Hofacker IL, Stadler PF *et al.* RNA folding with soft constraints: reconciliation of probing data and thermodynamic secondary structure prediction. *Nucleic Acids Res* 2012;40:4261–72. <https://doi.org/10.1093/nar/gks009>
16. Bohn P, Gribbling-Burrer AS, Ambi UB *et al.* Nano-DMS-MaP allows isoform-specific RNA structure determination. *Nat Methods* 2023;20:849–59. <https://doi.org/10.1038/s41592-023-01862-7>

17. Reining A, Nozinovic S, Schlepckow K *et al.* Three-state mechanism couples ligand and temperature sensing in riboswitches. *Nature* 2013;499:355–9. <https://doi.org/10.1038/nature12378>
18. Morandi E, Manfredonia I, Simon LM *et al.* Genome-scale deconvolution of RNA structure ensembles. *Nat Methods* 2021;18:249–52. <https://doi.org/10.1038/s41592-021-01075-w>
19. Olson SW, Turner AMW, Arney JW *et al.* Discovery of a large-scale, cell-state-responsive allosteric switch in the 7SK RNA using DANCE-MaP. *Mol Cell* 2022;82:1708–23. <https://doi.org/10.1016/j.molcel.2022.02.009>
20. Lorenz R, Bernhart SH, Höner zu Siederdissen C *et al.* ViennaRNA Package 2.0. *Algorithms Mol Biol* 2011;6:1–14. <https://doi.org/10.1186/1748-7188-6-26>
21. De Martino A, De Martino D. An introduction to the maximum entropy approach and its application to inference problems in biology. *Heliyon* 2018;4:e00596. <https://doi.org/10.1016/j.heliyon.2018.e00596>
22. Calonaci N, Bernetti M, Jones A *et al.* Molecular dynamics simulations with grand-canonical reweighting suggest cooperativity effects in RNA structure probing experiments. *J Chem Theory Comput* 2023;19:3672–85. <https://doi.org/10.1021/acs.jctc.3c00084>
23. Liu-Wei W, van der Toorn W, Bohn P *et al.* Sequencing accuracy and systematic errors of nanopore direct RNA sequencing. *BMC Genomics* 2024;25:1–15. <https://doi.org/10.1186/s12864-024-10440-w>
24. Verwilt J, Mestdagh P, Vandesompele J. Artifacts and biases of the reverse transcription reaction in RNA sequencing. *RNA* 2023;29:889–97. <https://doi.org/10.1261/rna.079623.123>
25. Byrd RH, Lu P, Nocedal J *et al.* A limited memory algorithm for bound constrained optimization. *SIAM J Sci Comput* 1995;16:1190–208. <https://doi.org/10.1137/0916069>
26. Virtanen P, Gommers R, Oliphant TE *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* 2020;17:261–72. <https://doi.org/10.1038/s41592-019-0686-2>
27. Zhao C, Liu F, Pyle AM. An ultraprogressive, accurate reverse transcriptase encoded by a Metazoan Group II intron. *RNA* 2018;24:183–95. <https://doi.org/10.1261/rna.063479.117>
28. Guo LT, Adams RL, Wan H *et al.* Sequencing and structure probing of long RNAs using MarathonRT: a next-generation reverse transcriptase. *J Mol Biol* 2020;432:3338–52. <https://doi.org/10.1016/j.jmb.2020.03.022>
29. Gribbling-Burrer AS, Bohn P, Smyth RP. Isoform-specific RNA structure determination using Nano-DMS-MaP. *Nat Protoc* 2024;19:1835–65. <https://doi.org/10.1038/s41596-024-00959-3>
30. Incarnato D, Morandi E, Simon LM *et al.* RNA framework: an all-in-one toolkit for the analysis of RNA structures and post-transcriptional modifications. *Nucleic Acids Res* 2018;46:e97. <https://doi.org/10.1093/nar/gky486>
31. Langmead B, Salzberg SL. Fast gapped-read alignment with Bowtie 2. *Nat Methods* 2012;9:357–9. <https://doi.org/10.1038/nmeth.1923>
32. Danecek P, Bonfield JK, Liddle J *et al.* Twelve years of SAMtools and BCFtools. *GigaScience* 2021;10:giab008. <https://doi.org/10.1093/gigascience/giab008>
33. Kappel K, Zhang K, Su Z *et al.* Accelerated cryo-EM-guided determination of three-dimensional RNA-only structures. *Nat Methods* 2020;17:699–707. <https://doi.org/10.1038/s41592-020-0878-9>
34. Bottaro S, Bussi G, Pinamonti G *et al.* Barnaba: software for analysis of nucleic acid structures and trajectories. *RNA* 2019;25:219–31. <https://doi.org/10.1261/rna.067678.118>
35. Torres-Larios A, Swinger KK, Krasilnikov AS *et al.* Crystal structure of the RNA component of bacterial ribonuclease P. *Nature* 2005;437:584–7. <https://doi.org/10.1038/nature04074>
36. Mathews DH, Disney MD, Childs JL *et al.* Incorporating chemical modification constraints into a dynamic programming algorithm for prediction of RNA secondary structure. *Proc Natl Acad Sci USA* 2004;101:7287–92. <https://doi.org/10.1073/pnas.0401799101>
37. Turner DH, Mathews DH. NNDB: the nearest neighbor parameter database for predicting stability of nucleic acid secondary structure. *Nucleic Acids Res* 2010;38:D280–2. <https://doi.org/10.1093/nar/gkp892>
38. Andronescu M, Condon A, Hoos HH *et al.* Efficient parameter estimation for RNA secondary structure prediction. *Bioinformatics* 2007;23:i19–28. <https://doi.org/10.1093/bioinformatics/btm223>
39. Cheng L, White EN, Brandt NL *et al.* Cotranscriptional RNA strand exchange underlies the gene regulation mechanism in a purine-sensing transcriptional riboswitch. *Nucleic Acids Res* 2022;50:12001–18. <https://doi.org/10.1093/nar/gkac102>
40. Mustoe AM, Lama NN, Irving PS *et al.* RNA base-pairing complexity in living cells visualized by correlated chemical probing. *Proc Natl Acad Sci USA* 2019;116:24574–82. <https://doi.org/10.1073/pnas.1905491116>
41. Giuliodori A, Di Pietro F, Marzi S *et al.* The cspA mRNA is a thermosensor that modulates translation of the cold-shock protein CspA. *Mol Cell* 2010;37:21–33. <https://doi.org/10.1016/j.molcel.2009.11.033>
42. Zhang Y, Burkhardt DH, Rouskin S *et al.* A stress response that monitors and regulates mRNA structure is central to cold shock adaptation. *Mol Cell* 2018;70:274–86. <https://doi.org/10.1016/j.molcel.2018.02.035>
43. Ratajczyk EJ, Bath J, Šulc P *et al.* Controlling DNA–RNA strand displacement kinetics with base distribution. *Proc Natl Acad Sci USA* 2025;122:e2416988122. <https://doi.org/10.1073/pnas.2416988122>
44. Ehresmann C, Baudin F, Moulé M *et al.* Probing the structure of RNAs in solution. *Nucleic Acids Res* 1987;15:9109–28. <https://doi.org/10.1093/nar/15.22.9109>
45. Merino EJ, Wilkinson KA, Coughlan JL *et al.* RNA structure analysis at single nucleotide resolution by selective 2′-hydroxyl acylation and primer extension (SHAPE). *J Am Chem Soc* 2005;127:4223–31. <https://doi.org/10.1021/ja043822v>
46. Weeks KM. Advances in RNA structure analysis by chemical probing. *Curr Opin Struct Biol* 2010;20:295–304. <https://doi.org/10.1016/j.sbi.2010.04.001>
47. Marinus T, Fessler AB, Ogle CA *et al.* A novel SHAPE reagent enables the analysis of RNA structure in living cells with unprecedented accuracy. *Nucleic Acids Res* 2021;49:e34. <https://doi.org/10.1093/nar/gkaa1255>
48. Adilakshmi T, Lease RA, Woodson SA. Hydroxyl radical footprinting in vivo: mapping macromolecular structures with synchrotron radiation. *Nucleic Acids Res* 2006;34:e64. <https://doi.org/10.1093/nar/gkl1291>
49. Douds CA, Babinzke P, Bevilacqua P. A new reagent for *in vivo* structure probing of RNA G and U residues that improves RNA structure prediction alone and combined with DMS. *RNA* 2024;30:901–19. <https://doi.org/10.1261/rna.079974.124>
50. Saaidi A, Allouche D, Regnier M *et al.* IPANEMAP: integrative probing analysis of nucleic acids empowered by multiple accessibility profiles. *Nucleic Acids Res* 2020;48:8276–89. <https://doi.org/10.1093/nar/gkaa607>
51. Hummer G, Köfinger J. Bayesian ensemble refinement by replica simulations and reweighting. *J Chem Phys* 2015;143:243150. <https://doi.org/10.1063/1.4937786>
52. Cesari A, Reißer S, Bussi G. Using the maximum entropy principle to combine simulations and solution experiments. *Computation* 2018;6:15. <https://doi.org/10.3390/computation6010015>
53. Tomezsko PJ, Corbin VDA, Gupta P *et al.* Determination of RNA structural diversity and its role in HIV-1 RNA splicing. *Nature* 2020;582:438–42. <https://doi.org/10.1038/s41586-020-2253-5>
54. Spasic A, Assmann SM, Bevilacqua PC *et al.* Modeling RNA secondary structure folding ensembles using SHAPE mapping data. *Nucleic Acids Res* 2018;46:314–23. <https://doi.org/10.1093/nar/gkx1057>

55. Li H, Aviran S. Statistical modeling of RNA structure profiling experiments enables parsimonious reconstruction of structure landscapes. *Nat Commun* 2018;9:606. <https://doi.org/10.1038/s41467-018-02923-8>
56. Arnold EB, Cohn D, Bose E *et al*. Investigating the interplay between RNA structural dynamics and RNA chemical probing experiments. *Nucleic Acids Res* 2025;53:gkaf290. <https://doi.org/10.1093/nar/gkaf290>
57. Deenalattha DHS, Jurich CP, Lange B *et al*. Characterizing RNA 3D structural features from DMS reactivity. *Nucleic Acids Res* 2026;54:gkag672. <https://doi.org/10.1093/nar/gkag672>