PIPER: MULTI-CONFORMER REASONING WITH PHYSICS-INFORMED REGULARIZATION FOR RNA–LIGAND INTERACTION PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

RNA–ligand interaction prediction is an important task in RNA-targeted drug discovery, yet it remains challenging because RNAs can adopt multiple conformational states while high-resolution tertiary-structure supervision is scarce. Most existing deep learning approaches adopt a single-state formulation, compressing each RNA–ligand pair into a static representation before prediction. Collapsing multiple latent conformational states into a single pair-level representation can blur state-specific interaction signals, making it difficult to compare alternative conformational hypotheses or impose state-specific geometric and physical constraints. To address these limitations, we introduce PIPER, a physics-informed framework for explicit reasoning over a compact set of surrogate conformers without requiring atomistic structures. Given RNA sequence, secondary structure, and ligand graphs, PIPER represents each RNA–ligand pair as a compact set of surrogate conformers rather than a single static state. It performs conformer-wise competition by integrating prior-guided conformer ranking with sparse RNA-residue–ligand-atom interactions, retains the two highest-utility surrogate conformers through a sparse Top-2 posterior, and aggregates their conformer-specific scores. A source-aware screened residual regularizes the local RNA–ligand interaction regions for all generated surrogate conformers. Extensive evaluations show that PIPER outperforms strong baselines across standard benchmarks and achieves particularly strong generalization under RNA cold-start, scaffold-hopping, and dual-cold settings.

1 INTRODUCTION

RNA is an increasingly important therapeutic target, expanding small-molecule drug discovery beyond proteins to non-coding transcripts implicated in oncogenesis, viral replication, and genetic disorders (Warner et al., 2018; Childs-Disney et al., 2022; Connelly et al., 2016; Falese et al., 2021; Morgan et al., 2019; Donlic et al., 2022). RNA–ligand interaction prediction is nevertheless difficult in the prevalent structure-scarce regime: functional RNAs can occupy alternative interaction-relevant configurations, whereas high-resolution tertiary structures are available for only a limited subset of RNA–ligand systems (Fei et al., 2026; Berman et al., 2000; Burley et al., 2021; Ganser et al., 2019). Ligand recognition can consequently depend on conformational selection, in which a molecule preferentially stabilizes one state from an underlying ensemble (Boehr et al., 2009; Stelzer et al., 2010; Ganser et al., 2019; Umuhire Juru et al., 2019). The resulting learning problem is not only to map sequence and ligand features to a binding label, but also to represent the unobserved state variation through which recognition may occur.

Most current RNA–ligand predictors instead use a single-state representation paradigm: they compress an RNA–ligand pair into one fused embedding and directly predict a binding score. Representative methods such as SMRTnet (Fei et al., 2026), DeepRSMA (Huang et al., 2024), and RNAsmol (Ma et al., 2025) span sequence-, fusion-, and augmentation-based modeling strategies and have advanced prediction under standard supervised settings. However, representing multiple latent conformational states with a single pair-level embedding can obscure state-specific interaction signals and prevents explicit reasoning over competing conformational hypotheses. This limitation becomes particularly important under distribution shifts in ligand chemistry or RNA sequence patterns.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To address these limitations, we propose PIPER, a multi-conformer reasoning framework with physics-informed regularization for RNA–ligand interaction prediction. Rather than relying on direct single-state fusion, PIPER formulates prediction as explicit finite-state inference over a compact set of surrogate conformers. Starting from RNA sequence, secondary structure, and ligand graphs, it constructs a residue-level surrogate RNA geometry and generates multiple surrogate conformers. These conformers are evaluated using prior-guided conformer ranking and sparse RNA-residue–ligand-atom interaction evidence, followed by conformer-wise competition through an explicit utility function. A sparse Top-2 posterior retains the two highest-utility surrogate conformers for conformer-specific score aggregation, while a source-aware screened residual regularizes the local RNA–ligand interaction regions for all generated surrogate conformers.

Across three RNA–ligand benchmarks and four controlled distribution-shift protocols, PIPER improves on matched baselines under standard splits and shows clear ranking advantages under RNA cold-start, scaffold-hopping, and dual-cold settings. Mechanistic analyses further support the key design choices of PIPER. Top-2 inference outperforms both Top-1 selection and All-3 aggregation, the utility margins between surrogate conformers increase during training, and physics-specific controls show that the full source-aware screened formulation is more effective than generic regularization alternatives. Our contributions are:

• **Explicit sparse latent-state inference.** To our knowledge, PIPER is the first RNA–ligand framework to perform explicit finite-state inference over surrogate conformational hypotheses, using competition between prior-guided conformer ranking and complementary ligand-conditioned evidence to retain sparse Top-2 support rather than selecting a single conformer or averaging over all candidate conformers. This preserves conformer-specific interaction signals and focuses prediction on the highest-utility binding-relevant hypotheses.

• **Physics-informed conformer-specific regularization.** We introduce a source-aware screened residual that regularizes local RNA–ligand interaction regions for each surrogate conformer, providing a conformer-specific physics-informed inductive bias. Matched controls using generic L2, a homogeneous partial differential equation (PDE) residual, and no-residual-loss variants further support the contribution of the complete source-aware screened formulation.

• **Evaluation and mechanistic validation.** We evaluate PIPER across three benchmarks and four controlled distribution-shift protocols, complemented by analyses of posterior support, physics specificity, perturbation robustness, and failure modes to characterize both its predictive gains and modeling behavior.

2 Related Work

2.1 RNA–Ligand Prediction under Structural Scarcity

Recent RNA–ligand prediction methods learn pair-level interaction scores from RNA sequence, secondary structure, and ligand representations, with progress driven by multimodal learning, cross-modal interaction modeling, and data augmentation (Fei et al., 2026; Huang et al., 2024; Ma et al., 2025; Oliver et al., 2020; Carvajal-Patiño et al., 2025; Krishnan et al., 2024; Sun & Gao, 2024; Zhuo et al., 2024; Xia et al., 2025). These methods demonstrate that effective prediction is possible without experimentally resolved RNA tertiary structures. RNA foundation models further provide transferable sequence representations for downstream structural and functional tasks (Chen et al., 2022; Penić et al., 2025; Zhou et al., 2025; Li et al., 2025; Shen et al., 2024). However, existing predictors generally compress each RNA–ligand pair into a single representation, leaving alternative conformational hypotheses implicit rather than explicitly compared during prediction.

2.2 Conformational and Latent-State Reasoning

Conformational ensembles and conformational selection provide a fundamental view of molecular recognition (Boehr et al., 2009; Stelzer et al., 2010; Ganser et al., 2019; Al-Hashimi & Walter, 2008; Cruz & Westhof, 2009; Dethoff et al., 2012; Bose et al., 2024). RNA molecules can populate multiple interaction-relevant states, whose relative populations may shift upon ligand binding. Although atomistic ensemble characterization typically requires structural experiments or detailed simulation, predictive latent states need not correspond to resolved native conformers. Finite latent-state models

can instead retain and compare multiple candidate hypotheses during prediction. However, explicit latent conformational reasoning remains underexplored for RNA–ligand interaction prediction.

2.3 Geometry- and Physics-Informed Molecular Learning

Geometry-aware molecular learning incorporates coordinate information and geometric symmetries into biomolecular models. Equivariant architectures such as the E(n)-equivariant graph neural network (EGNN) (Satorras et al., 2021) and protein–ligand methods including EquiBind, DiffDock, TankBind, and FABind (Stärk et al., 2022; Corso et al., 2022; Lu et al., 2022; Pei et al., 2023; Lu et al., 2024) exploit explicit three-dimensional geometry for docking and interaction modeling, often supported by large structural resources (Berman et al., 2000; Burley et al., 2021). These advances motivate geometric inductive biases for RNA–ligand prediction, where comparable co-resolved structural data are less abundant.

Physics-informed learning complements geometric modeling by incorporating differential-equation constraints and physically motivated regularization (Raissi et al., 2019; Karniadakis et al., 2021). Screened electrostatic formulations further motivate coarse field-based constraints in biomolecular systems (Ren et al., 2012), which can act as inductive regularization without implying calibrated atomistic simulation. Their integration with explicit latent conformational reasoning remains relatively unexplored in RNA–ligand interaction prediction.

3 Method

3.1 Problem Formulation and Overview

PIPER formulates RNA–ligand interaction prediction under structural scarcity as sparse inference over a finite surrogate conformational space. Given an RNA sequence R, secondary structure S, and ligand L, PIPER constructs a compact set of surrogate conformers,

$$\mathcal{Z} = \{z_k\}_{k=1}^{K}, \tag{1}$$

from a shared residue-level geometric scaffold. Each surrogate conformer represents a latent residue-level geometric hypothesis for interaction reasoning rather than an atomistically resolved RNA structure.

PIPER combines prior-guided conformer ranking with ligand-conditioned residue–atom evidence through an explicit conformer utility, and a sparse Top-2 posterior retains the two highest-utility conformers. Their conformer-specific energy-style scores are produced by a multilayer perceptron (MLP) together with field summaries and then aggregated for interaction prediction, while source-aware screened regularization is applied across the full conformer set during training. Training combines binary cross-entropy (BCE), Kullback–Leibler (KL) divergence for prior alignment, and this screened regularization.

The resulting formulation separates surrogate conformer generation, ligand-conditioned conformer evaluation, and sparse multi-conformer aggregation, enabling explicit competition among alternative interaction hypotheses rather than collapsing them into a single pair-level representation. Figure 1 illustrates the overall architecture of PIPER.

3.2 Molecular Encoding and Surrogate Conformational Space

To support multi-conformer reasoning without relying on resolved RNA tertiary structures, PIPER represents RNA and ligand inputs in complementary semantic and surrogate geometric spaces. The RNA encoder maps sequence and secondary-structure features to residue embeddings $\{h_i^R\}_{i=1}^{N_R}$ and a graph-level representation. A deterministic residue-level scaffold defines the base surrogate coordinates

$$x_i^{(0)} = \left[\frac{i}{N_R},\ \sin\!\left(\frac{i}{2}\right),\ \cos\!\left(\frac{i}{3}\right)\right]^{\top}, \qquad i = 1, \ldots, N_R. \tag{2}$$

The ligand encoder processes the molecular graph to obtain atom embeddings $\{h_a^L\}_{a=1}^{N_L}$ and a graph-level representation h^L.

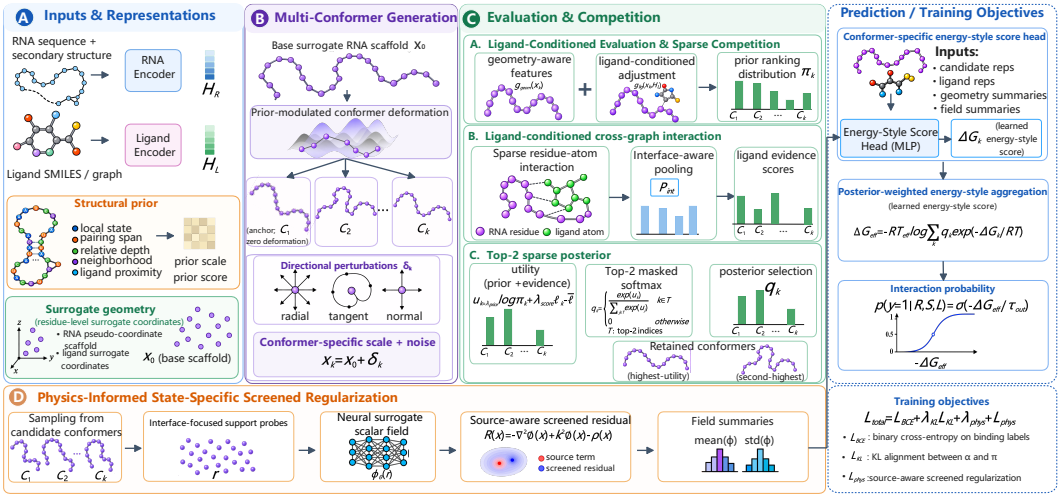


Figure 1: Overview of PIPER. RNA and ligand inputs are encoded into semantic and surrogate geometric representations. Multiple surrogate conformers are ranked using prior-guided and ligand-conditioned evidence, followed by sparse Top-2 prediction and source-aware screened-residual regularization.

For coarse geometric operations, ligand atoms are assigned a deterministic reference chart,

$$r_a^L = \left[\cos\left(\frac{2\pi(a-1)}{N_L}\right),\ \sin\left(\frac{2\pi(a-1)}{N_L}\right),\ 0\right]^\top, \qquad a = 1, \ldots, N_L. \tag{3}$$

The learned representations encode RNA sequence and structural context together with ligand chemistry, while surrogate coordinates provide computational geometry for conformer construction and interaction reasoning rather than atomistically resolved RNA or ligand structures.

3.3 Surrogate Conformer Generation and Geometry-Aware Prior

PIPER constructs a compact set of surrogate conformers by deforming the base RNA scaffold along locally defined radial, tangential, and normal directions. Let

$$c_L = \frac{1}{N_L}\sum_{a=1}^{N_L} r_a^L, \qquad d_i = \min_{a=1,\ldots,N_L} \|x_i^{(0)} - r_a^L\|_2, \qquad w_i^{\text{int}} = \exp(-d_i/T_I), \tag{4}$$

where c_L is the centroid of the ligand reference chart and $T_I > 0$ controls the interface scale. A structured residue score s_i combines secondary-structure information, sequence-position centrality, and a coarse chart-based interface cue. The local interface weight and structural score are integrated as

$$g_i = (1-\omega_g)w_i^{\text{int}} + \omega_g s_i, \qquad 0 \le \omega_g \le 1, \tag{5}$$

with residue deformation scale $a_{k,i} = \gamma_k a_i$, where a_i is derived from structural context and γ_k controls conformer-specific perturbation magnitude. For residue i, r_i denotes the radial direction toward c_L, t_i the local scaffold tangent, and n_i the corresponding normal direction. Using learned conformer coefficients $b_{k,1:4}$ and perturbation scale η_k, the displacement is defined as

$$\delta_{k,i} = \eta_k a_{k,i} g_i \left(b_{k,1} r_i + b_{k,2} t_i + b_{k,3} n_i\right) + \eta_k a_{k,i}\, \text{sigmoid}(b_{k,4})\epsilon_{k,i}, \qquad x_{k,i} = x_i^{(0)} + \delta_{k,i}. \tag{6}$$

The undeformed scaffold serves as an anchor conformer, while the remaining conformers are generated through conformer-specific perturbations around this reference. Stochastic perturbations during training promote local variation among candidate hypotheses, whereas evaluation uses deterministic conformer construction. This procedure yields multiple residue-level geometric hypotheses anchored to a common surrogate scaffold. Surrogate-conformer implementation details are provided in Appendix B.2.

Geometry-aware conformational prior. Coarse structural plausibility is evaluated using deformation strain E_k^{def}, contact-oriented geometric support B_k^{ct}, and ligand-approach support B_k^{app}:

$$E_{\mathrm{rank},k}^{(0)} = E_k^{\mathrm{def}} - \lambda_c B_k^{\mathrm{ct}} - \lambda_a B_k^{\mathrm{app}}. \tag{7}$$

The resulting score defines the geometry-derived component of conformer ranking, while structural context shapes conformer generation through $a_{k,i}$ and g_i.

3.4 Ligand-Conditioned Interaction and Sparse Conformer Competition

Geometric plausibility alone does not determine whether a conformer is favored by a given ligand. PIPER therefore evaluates each surrogate conformer using ligand-conditioned residue–atom interactions before posterior selection. RNA residue queries attend to ligand-atom keys and values, with the highest-scoring interactions retained for each residue. The resulting ligand context is gated into the RNA residue features and used for conformer discrimination rather than conformer-coordinate construction.

For conformer k, its geometry defines normalized interface-pooling weights

$$\omega_{k,i} = \mathrm{softmax}_i\left(-\frac{\|x_{k,i} - c_L\|_2}{T_I}\right), \qquad r_k = \sum_i \omega_{k,i} \tilde{h}_i^R, \tag{8}$$

where $\tilde{h}_i^R$ denotes the ligand-conditioned residue representation. The conformer representation h_k integrates coordinate summaries, r_k, the ligand representation h^L, and a learned conformer token c_k. The token-augmented interface representation $\hat{r}_k = r_k + c_k$ is then used to compute a ligand-conditioned adjustment,

$$b_k^{\mathrm{cond}} = \lambda_{\mathrm{cond}} \tanh f_{\mathrm{cond}}([c_k; \hat{r}_k; h^L]), \qquad E_{\mathrm{rank},k} = E_{\mathrm{rank},k}^{(0)} - b_k^{\mathrm{cond}}. \tag{9}$$

This adjustment refines the geometry-derived ranking score with ligand-conditioned information, yielding a conformer-specific ranking energy.

The corresponding prior-guided conformer-ranking distribution is

$$\pi_k = \frac{\exp(-E_{\mathrm{rank},k}/\tau_{\mathrm{prior}})}{\sum_j \exp(-E_{\mathrm{rank},j}/\tau_{\mathrm{prior}})}, \qquad \tau_{\mathrm{prior}} > 0, \tag{10}$$

while prior-guided attention is defined as

$$e_k = \frac{(W_q h^L)^\top (W_k h_k)}{\sqrt{d}}, \qquad \alpha_k = \mathrm{softmax}_k(e_k + \log \pi_k). \tag{11}$$

We define the ligand-conditioned conformer evidence as $\ell_k = e_k + b_k^{\mathrm{cond}}$. Here, π_k represents the prior-guided ranking distribution after ligand-conditioned adjustment, ℓ_k provides complementary ligand-conditioned evidence, and α serves as the soft distribution used for prior alignment.

Sparse conformer competition. Conformer competition is governed by the utility

$$u_k = \lambda_\pi \log \pi_k + \lambda_\ell(\ell_k - \bar{\ell}), \qquad \tilde{q}_k = \mathrm{softmax}_k(u_k/\tau_{\mathrm{post}}), \qquad \bar{\ell} = \frac{1}{K}\sum_k \ell_k, \tag{12}$$

where λ_π and λ_ℓ control the relative contributions of ranking and ligand-conditioned evidence, and $\tau_{\mathrm{post}} > 0$ controls posterior sharpness. Formal properties of this finite-state competition rule are characterized in Appendix A. Appendix A.3 shows that pairwise utility differences decompose exactly into prior-guided ranking and ligand-conditioned evidence contributions. The corresponding empirical decomposition and representative conformer-competition cases are provided in Appendices B.3 and B.4. Rather than selecting a single conformer or averaging over the full hypothesis set, PIPER retains the two highest-utility conformers and renormalizes their posterior mass:

$$q_k = \frac{\tilde{q}_k \mathbf{1}[k \in \mathrm{Top2}(u)]}{\sum_j \tilde{q}_j \mathbf{1}[j \in \mathrm{Top2}(u)]}. \tag{13}$$

The resulting sparse posterior preserves ambiguity between the strongest competing hypotheses while excluding low-utility conformers from the final prediction. Appendix A.4 characterizes the local stability of Top-2 selection under bounded utility perturbations determined by the support margin. Here, $\tilde{q}$ denotes the dense pre-sparsification support, whereas q is the normalized Top-2 posterior used for prediction.

3.5 PHYSICS-INFORMED STATE-SPECIFIC SCREENED REGULARIZATION

To introduce a state-specific physics-informed inductive bias into the surrogate conformational space, PIPER associates each conformer with a neural scalar field $\phi_k(p)$ and constrains it through a source-aware screened residual:

$$\mathrm{Res}_k(p) = -\nabla^2\phi_k(p) + \kappa^2\phi_k(p) - \rho_k(p), \qquad \mathcal{L}_{\mathrm{res},k} = \frac{1}{M}\sum_{m=1}^{M}\mathrm{Res}_k(p_m)^2. \tag{14}$$

The source term $\rho_k(p)$ is constructed from RNA-centered and ligand-centered surrogate sources, making the residual dependent on the conformer-specific surrogate geometry. The resulting constraint serves as a physics-informed regularizer in surrogate space rather than a calibrated atomistic electrostatic model.

For each conformer, probe points emphasize interface-focused regions in surrogate space while maintaining broader support across RNA residues. Small stochastic perturbations are introduced during training to regularize field evaluation. The screened-residual regularization is applied to all generated conformers. Sampling and perturbation details are provided in Appendix B.5.

The field branch serves two complementary roles: summaries of ϕ_k contribute to conformer-specific interaction scoring, while $\mathcal{L}_{\mathrm{res},k}$ imposes screened regularization on each conformer.

3.6 SPARSE MULTI-CONFORMER PREDICTION AND JOINT LEARNING

Each surrogate conformer yields a learned energy-style interaction score conditioned on its conformer representation, ligand representation, geometric summaries, and field statistics:

$$\Delta G_{\mathrm{bind},k} = \left(\overline{\phi}_k - \tfrac{1}{2}\,\mathrm{sd}(\phi_k)\right) + f_{\mathrm{bind}}([h_k; h^L; h_k \odot h^L; d_k]), \tag{15}$$

where d_k collects conformer-specific distance and field-summary features. Here, $\Delta G_{\mathrm{bind},k}$ denotes a learned energy-style score rather than a calibrated physical binding free energy.

The sparse posterior aggregates conformer-specific scores through posterior-weighted log-sum-exp aggregation:

$$\Delta G_{\mathrm{eff}} = -\tau_{\mathrm{agg}}\log\sum_{k=1}^{K} q_k \exp\left(-\frac{\Delta G_{\mathrm{bind},k}}{\tau_{\mathrm{agg}}}\right), \qquad p(y=1 \mid R,S,L) = \mathrm{sigmoid}\left(-\frac{\Delta G_{\mathrm{eff}}}{\tau_{\mathrm{out}}}\right). \tag{16}$$

The temperatures τ_{agg} and τ_{out} control score aggregation and probability mapping, respectively.

The model is trained jointly with supervised prediction, prior alignment, and screened regularization:

$$\mathcal{L} = \mathcal{L}_{\mathrm{BCE}} + \lambda_{\mathrm{KL}} D_{\mathrm{KL}}(\pi\|\alpha) + \lambda_{\mathrm{PDE}}\frac{1}{K}\sum_{k=1}^{K}\mathcal{L}_{\mathrm{res},k}. \tag{17}$$

where λ_{KL} and λ_{PDE} weight prior alignment and screened regularization, respectively, with $D_{\mathrm{KL}}(\pi\|\alpha)$ aligning the prior-guided conformer-ranking distribution π with the soft attention distribution α. Algorithm 1 is provided in Appendix B.1.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Datasets and protocols. We evaluate PIPER on three RNA–ligand benchmarks with distinct dataset scales and ligand coverage: SMRTnet (Fei et al., 2026), Robin (Yazdani et al., 2023), and Biosensor (Townshend et al., 2022). SMRTnet contains 24,687 labeled RNA–ligand pairs involving 4,838 RNAs and 341 ligands, with 8,229 positive and 16,458 negative samples. Robin contains 46,557 pairs involving 26 RNAs and 1,908 ligands, whereas Biosensor contains 16,999 pairs involving 191 RNAs and 89 ligands. Each dataset is partitioned into training, validation, and test sets at an 8:1:1 ratio, with all methods evaluated on matched partitions under the same metric definitions. We further

Table 1: Benchmark performance across SMRTnet, Robin, and Biosensor. Bold and underlined values mark the best and second-best results.

Dataset	Metric	PIPER	DeepRSMA	SMRTnet	SDNE	BINDTI	RNAsmol	DeepConv-DTI
SMRTnet	AUC	**0.8781±0.0031**	0.8152±0.0078	0.8343±0.0051	0.6041±0.1211	0.7060±0.0038	0.7327±0.0005	<u>0.8484±0.0053</u>
	AUPR	**0.7753±0.0052**	0.6698±0.0091	0.6991±0.0064	0.4405±0.0705	0.5513±0.0274	0.5587±0.0039	<u>0.7177±0.0095</u>
	ACC	**0.7983±0.0052**	0.7154±0.0207	0.7211±0.0115	0.6021±0.0939	0.6532±0.0222	0.6155±0.0058	<u>0.7716±0.0085</u>
	F1	**0.7285±0.0067**	0.6578±0.0111	0.6665±0.0091	0.5891±0.0367	0.5526±0.0030	0.5866±0.0058	<u>0.6854±0.0055</u>
Robin	AUC	**0.7241±0.0085**	0.6266±0.0022	<u>0.7122±0.0067</u>	0.5138±0.0209	0.7014±0.0118	0.5311±0.0052	0.6037±0.0146
	AUPR	**0.2310±0.0015**	0.1165±0.0019	0.2184±0.0134	0.1295±0.0182	<u>0.2206±0.0260</u>	0.0858±0.0009	0.1026±0.0069
	ACC	**0.8830±0.0093**	0.6314±0.1094	<u>0.8712±0.0107</u>	0.8678±0.0164	0.8605±0.0151	0.7503±0.0388	0.7068±0.0884
	F1	**0.2764±0.0133**	0.1789±0.0063	0.2541±0.0073	0.1783±0.0086	<u>0.2603±0.0071</u>	0.1511±0.0052	0.1693±0.0028
Biosensor	AUC	**0.9812±0.0011**	0.9487±0.0034	0.9598±0.0096	0.8599±0.0119	<u>0.9628±0.0086</u>	0.9158±0.0028	0.8464±0.0572
	AUPR	**0.8779±0.0137**	0.7709±0.0154	0.8131±0.0330	0.5918±0.0039	<u>0.8305±0.0364</u>	0.6511±0.0114	0.5294±0.1050
	ACC	**0.9645±0.0041**	0.9490±0.0025	0.9547±0.0075	0.9202±0.0003	<u>0.9598±0.0070</u>	0.9322±0.0019	0.8865±0.0575
	F1	**0.8013±0.0230**	0.6933±0.0190	0.7491±0.0350	0.5608±0.0069	<u>0.7785±0.0343</u>	0.6029±0.0166	0.4669±0.1106

evaluate ligand cold-start, RNA cold-start, scaffold-hopping, and dual cold-start settings; detailed split construction and leakage controls are provided in Appendix B.6.

Baselines and metrics. We compare PIPER with six representative baselines: DeepRSMA (Huang et al., 2024), SMRTnet (Fei et al., 2026), SDNE (Wang et al., 2016), BINDTI (Peng et al., 2024), RNAsmol (Ma et al., 2025), and DeepConv-DTI (Lee et al., 2019). We report the area under the receiver operating characteristic curve (AUC) and the area under the precision–recall curve (AUPR) for ranking performance, together with accuracy (ACC) and F1 score (F1) for thresholded classification. For ACC and F1, the decision threshold is selected by maximizing F1 on the validation set and then fixed for test evaluation. Benchmark results are reported as mean $\pm$ standard deviation over three independent runs.

Implementation details. PIPER uses a hidden dimension of 96 and is trained for 80 epochs with AdamW (Loshchilov & Hutter, 2019), a learning rate of 7×10^{-4}, and cosine annealing (Loshchilov & Hutter, 2017). The model uses three surrogate conformers with sparse Top-2 posterior support during both training and evaluation. The KL-alignment and screened-residual loss weights are set to $\lambda_{\mathrm{KL}} = \lambda_{\mathrm{PDE}} = 0.02$. Additional training and implementation details are provided in Appendix B.7.

4.2 Benchmark Results Across Datasets

We evaluate PIPER across RNA–ligand datasets with substantially different RNA and ligand compositions. This comparison evaluates both ranking quality and thresholded classification performance, providing a cross-dataset assessment of the model beyond a single benchmark. The results are reported in Table 1. On SMRTnet, the primary benchmark, PIPER achieves an AUC of 0.8781 and an AUPR of 0.7753, improving over the strongest competing results by 0.0297 and 0.0576, respectively. The gains extend to thresholded classification, where PIPER also obtains the highest ACC and F1. A similar pattern is observed on Robin despite its markedly different RNA–ligand composition: PIPER achieves the best performance across all four metrics, including an AUC of 0.7241 and an AUPR of 0.2310. On Biosensor, PIPER further reaches 0.9812 AUC and 0.8779 AUPR while maintaining the highest ACC and F1 among the compared methods. Taken together, the improvements across all three benchmarks indicate that the predictive advantage of PIPER is not confined to a particular dataset composition or evaluation metric.

4.3 Generalization under Molecular Distribution Shifts

We assess PIPER under out-of-distribution (OOD) shifts in ligand chemistry, RNA sequence, and their combination. Ligand and RNA cold-start separate molecular entities by Morgan-fingerprint and 5-mer Jaccard similarity, respectively, while scaffold-hopping enforces separation between Bemis–Murcko scaffold groups (Bemis & Murcko, 1996; Wu et al., 2018; Yang et al., 2019; Wallach & Heifets, 2018); dual cold-start jointly holds out ligand and RNA groups. Detailed split construction and leakage controls are provided in Appendix B.6.

Figure 2 summarizes OOD performance. PIPER achieves the best AUC/AUPR under RNA cold-start (0.8702/0.7584), scaffold-hopping (0.7322/0.5584), and dual cold-start (0.6931/0.4864), while remaining competitive under ligand cold-start (0.7377/0.5879), where DeepRSMA performs best. The

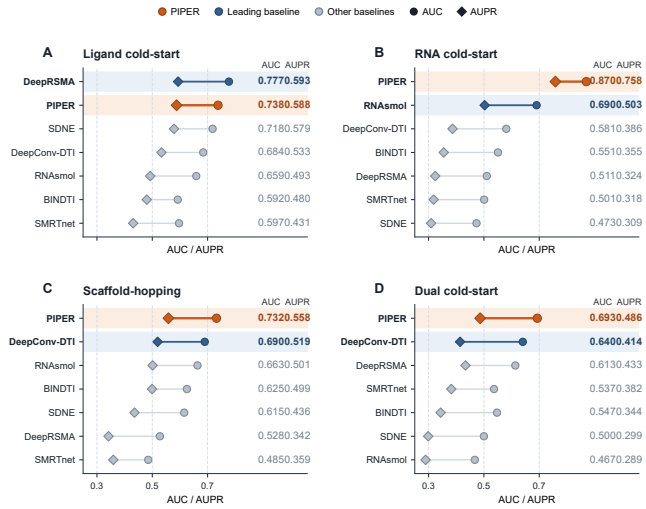


Figure 2: OOD ranking performance across four SMRTnet-derived distribution-shift protocols.

Table 2: Ablation and physics-specificity controls on SMRTnet. ACC and F1 use validation-selected thresholds.

Variant	AUC	AUPR	ACC	F1
PIPER	**0.8744**	**0.7691**	**0.7914**	**0.7203**
Multi-state reasoning				
Without conformer conditioning	0.8695	0.7601	0.7643	0.7043
Single conformer	0.8676	0.7459	0.7797	0.7106
Prior / interaction				
Uniform prior	0.8647	0.7314	0.7817	0.7030
Without graph cross-interaction	0.8657	0.7472	0.7882	0.7099
Physics				
Generic L2 field regularizer	0.8518	0.7368	0.7768	0.6913
Homogeneous PDE residual	0.8460	0.7242	0.7663	0.6716
Shuffled source assignments	0.8685	0.7486	0.7906	0.7133
No screened-residual loss ($\lambda_{\mathrm{PDE}} = 0$)	0.8382	0.7186	0.7461	0.6642

largest margin appears under RNA cold-start, where the strongest baseline reaches only 0.6904/0.5027. These results suggest stronger generalization to RNA-side, scaffold-level, and joint RNA–ligand shifts, with ligand-only cold-start remaining the most challenging. Matched OOD ablations further show that PIPER attains the highest AUPR across all four protocols over independently trained single-conformer and no-screened-residual-loss variants, supporting the contributions of sparse multi-conformer reasoning and screened regularization (Appendix B.6). Complete thresholded results, target-conditioned ligand prioritization, scaffold-level subgroup analysis, and representative scaffold-hopping cases are provided in Appendices B.6, B.8, B.9, and B.10.

4.4 Ablation and Physics-Specificity Analysis

We evaluate the contributions of multi-conformer reasoning, prior-guided interaction modeling, and source-aware screened regularization on the SMRTnet benchmark. As shown in Table 2, every targeted ablation degrades predictive performance. Replacing multi-conformer inference with a single conformer reduces AUPR from 0.7691 to 0.7459, while using a uniform prior or removing cross-graph interaction further lowers AUPR to 0.7314 and 0.7472, respectively. These results support the complementary roles of conformer competition, prior-guided ranking, and ligand-conditioned interaction modeling.

The screened-residual objective provides the largest individual contribution among the evaluated components: removing it reduces AUC/AUPR from 0.8744/0.7691 to 0.8382/0.7186, while retaining the neural-field features for interaction scoring. Replacing the source-aware screened residual with either a generic L2 penalty or a homogeneous PDE residual also degrades performance, indicating that the improvement is not attributable merely to the presence of an auxiliary field regularizer. A matched shuffled-source control, which preserves the operator, source counts, probe count, and

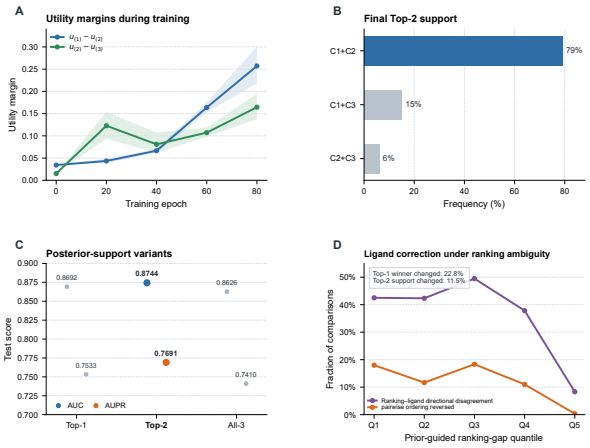


Figure 3: Posterior-support dynamics and ligand-conditioned conformer competition.

coordinate pool while permuting RNA and ligand source assignments, achieves 0.8685/0.7486 and remains below PIPER. Together, these controls support the importance of the aligned source-aware formulation for predictive performance. A complementary residual-consistency analysis is provided in Appendix B.11.

4.5 Mechanistic Analysis of Sparse Conformer Competition

We characterize sparse conformer competition and posterior support dynamics. Figure 3 shows that utility margins increase during training, while the final Top-2 posterior remains distributed across C1+C2, C1+C3, and C2+C3, accounting for 79%, 15%, and 6% of retained supports, respectively. Thus, training separates candidate utilities without collapsing support to a single conformer pair. The increasing $u_{(2)} - u_{(3)}$ margin also implies greater local stability: by Proposition A.4, a larger support margin permits larger bounded utility perturbations while preserving the selected Top-2 support.

The posterior-support comparison further shows that sparse Top-2 inference provides the strongest predictive trade-off. Independently trained Top-1, Top-2, and All-3 variants achieve AUC/AUPR values of 0.8692/0.7533, 0.8744/0.7691, and 0.8626/0.7410, respectively. Together with the margin analysis above, these results support sparse Top-2 inference as a balance between premature commitment to a single hypothesis and unrestricted aggregation over all candidate states.

Ligand-conditioned evidence changes conformer ordering most often when the prior-guided ranking gap is small. It changes the Top-1 conformer in 22.8% of cases and the retained Top-2 support in 11.5%. This suggests that ligand-specific evidence mainly refines uncertain prior-guided rankings rather than broadly overriding them. Additional analyses characterize high-confidence failure cases (Appendix B.12), robustness to incomplete secondary-structure inputs (Appendix B.13), and posterior-support and conformer-count sensitivity (Appendix B.14).

5 Conclusion

We present PIPER, a framework that replaces single-state fusion with explicit multi-conformer reasoning for RNA–ligand interaction prediction. By integrating prior-guided conformer ranking with lightweight physics-informed regularization, PIPER enables explicit latent-state reasoning over multiple structured interaction hypotheses without requiring tertiary-structure supervision. Across standard benchmarks and distribution shifts, it achieves strong performance, with particularly favorable generalization under stringent scaffold-level shifts, suggesting that multi-conformer reasoning provides a useful inductive bias beyond static fusion.

PIPER uses coarse surrogate geometries rather than atomistic RNA conformational ensembles, and its neural fields are designed to provide physics-informed regularization rather than quantitatively accurate electrostatic predictions. Future work will incorporate richer tertiary-structure supervision, more expressive conformational dynamics, and prospective experimental validation for RNA-targeted ligand discovery.

AI USE STATEMENT

We used a large language model (LLM) for grammar checking and improving the clarity of this manuscript. All technical content, experimental design, and analysis are original human work. The LLM suggestions were manually reviewed and modified to ensure that they align with the paper's objectives and maintain technical accuracy.

ETHICS STATEMENT

All authors adhere to the ICLR Code of Ethics. This work focuses on computational modeling of RNA–ligand interactions for biomolecular prediction. The experiments use existing benchmark datasets, including SMRTnet, Robin, and Biosensor, as described in Section 4.1; no new human-subject data or personally identifiable information are collected in this study. PIPER is intended as a computational research tool for RNA–ligand interaction prediction. Any downstream use in RNA-targeted drug discovery requires appropriate experimental validation and domain-specific assessment.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the PIPER architecture and training objective in Section 3 and summarize the end-to-end procedure in Algorithm 1. Dataset construction, distribution-shift protocols, partition statistics, training hyperparameters, software and hardware environments, and evaluation procedures are documented in Appendix B. Complete formal statements and proofs for the sparse conformer-competition mechanism are provided in Appendix A. Benchmark results are reported over independent runs where specified, and additional ablations, sensitivity analyses, and statistical tests are included in Appendix B. The supplementary materials contain the necessary code and example configuration files to reproduce PIPER training and evaluation. Our experimental code and data preprocessing pipeline will be made publicly available upon acceptance. The reported results can be reproduced using the datasets, split protocols, model configurations, training procedures, and evaluation settings described in the paper and supplementary materials.

REFERENCES

Hashim M Al-Hashimi and Nils G Walter. Rna dynamics: it is about time. *Current opinion in structural biology*, 18(3):321–329, 2008.

Guy W Bemis and Mark A Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of medicinal chemistry*, 39(15):2887–2893, 1996.

Helen M Berman, Talapady N Bhat, Philip E Bourne, Zukang Feng, Gary Gilliland, Helge Weissig, and John Westbrook. The protein data bank and the challenge of structural genomics. *Nature structural biology*, 7(11): 957–959, 2000.

David D Boehr, Ruth Nussinov, and Peter E Wright. The role of dynamic conformational ensembles in biomolecular recognition. *Nature chemical biology*, 5(11):789–796, 2009.

Ritwika Bose, Irfana Saleem, and Anthony M Mustoe. Causes, functions, and therapeutic possibilities of rna secondary structure ensembles and alternative states. *Cell chemical biology*, 31(1):17–35, 2024.

Stephen K Burley, Charmi Bhikadiya, Chunxiao Bi, Sebastian Bittrich, Li Chen, Gregg V Crichlow, Cole H Christie, Kenneth Dalenberg, Luigi Di Costanzo, Jose M Duarte, et al. Rcsb protein data bank: powerful new tools for exploring 3d structures of biological macromolecules for basic and applied research and education in fundamental biology, biomedicine, biotechnology, bioengineering and energy sciences. *Nucleic acids research*, 49(D1):D437–D451, 2021.

Juan G Carvajal-Patiño, Vincent Mallet, David Becerra, Luis Fernando Niño Vasquez, Carlos Oliver, and Jérôme Waldispühl. Rnamigos2: accelerated structure-based rna virtual screening with deep graph learning. *Nature Communications*, 16(1):2799, 2025.

Jiayang Chen, Zhihang Hu, Siqi Sun, Qingxiong Tan, Yixuan Wang, Qinze Yu, Licheng Zong, Liang Hong, Jin Xiao, Tao Shen, et al. Interpretable rna foundation model from unannotated data for highly accurate rna structure and function predictions. *BioRxiv*, pp. 2022–08, 2022.

Jessica L Childs-Disney, Xueyi Yang, Quentin MR Gibaut, Yuquan Tong, Robert T Batey, and Matthew D Disney. Targeting rna structures with small molecules. *Nature Reviews Drug Discovery*, 21(10):736–762, 2022.

Colleen M Connelly, Michelle H Moon, and John S Schneekloth. The emerging role of rna as a therapeutic target for small molecules. *Cell chemical biology*, 23(9):1077–1090, 2016.

Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.

José Almeida Cruz and Eric Westhof. The dynamic landscapes of rna architecture. *Cell*, 136(4):604–609, 2009.

Elizabeth A Dethoff, Jeetender Chugh, Anthony M Mustoe, and Hashim M Al-Hashimi. Functional complexity and regulation through rna dynamics. *Nature*, 482(7385):322–330, 2012.

Anita Donlic, Emily G Swanson, Liang-Yuan Chiu, Sarah L Wicks, Aline Umuhire Juru, Zhengguo Cai, Kamillah Kassam, Chris Laudeman, Bilva G Sanaba, Andrew Sugarman, et al. R-bind 2.0: an updated database of bioactive rna-targeting small molecules and associated rna secondary structures. *ACS chemical biology*, 17(6):1556, 2022.

James P Falese, Anita Donlic, and Amanda E Hargrove. Targeting rna with small molecules: from fundamental principles towards the clinic. *Chemical Society Reviews*, 50(4):2224–2243, 2021.

Yuhan Fei, Pengfei Wang, Jiasheng Zhang, Xinyue Shan, Zilin Cai, Jianbo Ma, Yangming Wang, and Qiangfeng Cliff Zhang. Predicting small molecule–rna interactions without rna tertiary structures. *Nature Biotechnology*, pp. 1–12, 2026.

Laura R Ganser, Megan L Kelly, Daniel Herschlag, and Hashim M Al-Hashimi. The roles of structural dynamics in the cellular functions of rnas. *Nature reviews Molecular cell biology*, 20(8):474–489, 2019.

Zhijian Huang, Yucheng Wang, Song Chen, Yaw Sing Tan, Lei Deng, and Min Wu. Deeprsma: a cross-fusion-based deep learning method for rna–small molecule binding affinity prediction. *Bioinformatics*, 40(12): btae678, 2024.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Sowmya R Krishnan, Arijit Roy, and M Michael Gromiha. Reliable method for predicting the binding affinity of rna-small molecule interactions using machine learning. *Briefings in Bioinformatics*, 25(2):bbae002, 2024.

Ingoo Lee, Jongsoo Keum, and Hojung Nam. Deepconv-dti: Prediction of drug-target interactions via deep learning with convolution on protein sequences. *PLoS computational biology*, 15(6):e1007129, 2019.

Guipeng Li, Feifei Jiang, Junhao Zhu, Huanhuan Cui, Zefeng Wang, and Wei Chen. Hydrarna: a hybrid architecture based full-length rna language model. *Genome Biology*, 26(1):383, 2025.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Wei Lu, Qifeng Wu, Jixian Zhang, Jiahua Rao, Chengtao Li, and Shuangjia Zheng. Tankbind: Trigonometry-aware neural networks for drug-protein binding structure prediction. *Advances in neural information processing systems*, 35:7236–7249, 2022.

Wei Lu, Jixian Zhang, Weifeng Huang, Ziqiao Zhang, Xiangyu Jia, Zhenyu Wang, Leilei Shi, Chengtao Li, Peter G Wolynes, and Shuangjia Zheng. Dynamicbind: predicting ligand-specific protein-ligand complex structure with a deep equivariant generative model. *Nature Communications*, 15(1):1071, 2024.

Hongli Ma, Letian Gao, Yunfan Jin, Jianwei Ma, Yilan Bai, Xiaofan Liu, Pengfei Bao, Ke Liu, Zhenjiang Zech Xu, and Zhi John Lu. Rna–ligand interaction scoring via data perturbation and augmentation modeling. *Nature computational science*, 5(8):648–660, 2025.

Brittany S Morgan, Bilva G Sanaba, Anita Donlic, Diane B Karloff, Jordan E Forte, Yuqi Zhang, and Amanda E Hargrove. R-bind: an interactive database for exploring and developing rna-targeted chemical probes. *ACS chemical biology*, 14(12):2691, 2019.

Carlos Oliver, Vincent Mallet, Roman Sarrazin Gendron, Vladimir Reinharz, William L Hamilton, Nicolas Moitessier, and Jérôme Waldispühl. Augmented base pairing networks encode rna-small molecule binding preferences. *Nucleic acids research*, 48(14):7690–7699, 2020.

Qizhi Pei, Kaiyuan Gao, Lijun Wu, Jinhua Zhu, Yingce Xia, Shufang Xie, Tao Qin, Kun He, Tie-Yan Liu, and Rui Yan. Fabind: Fast and accurate protein-ligand binding. *Advances in Neural Information Processing Systems*, 36:55963–55980, 2023.

Lihong Peng, Xin Liu, Long Yang, Longlong Liu, Zongzheng Bai, Min Chen, Xu Lu, and Libo Nie. Bindti: a bi-directional intention network for drug-target interaction identification based on attention mechanisms. *IEEE Journal of Biomedical and Health Informatics*, 29(3):1602–1612, 2024.

Rafael Josip Penić, Tin Vlašić, Roland G Huber, Yue Wan, and Mile Šikić. Rinalmo: general-purpose rna language models can generalize well on structure prediction tasks. *Nature Communications*, 16(1):5671, 2025.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

Pengyu Ren, Jaehun Chun, Dennis G Thomas, Michael J Schnieders, Marcelo Marucho, Jiajing Zhang, and Nathan A Baker. Biomolecular electrostatics and solvation: a computational perspective. *Quarterly reviews of biophysics*, 45(4):427–491, 2012.

Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

Tao Shen, Zhihang Hu, Siqi Sun, Di Liu, Felix Wong, Jiuming Wang, Jiayang Chen, Yixuan Wang, Liang Hong, Jin Xiao, et al. Accurate rna 3d structure prediction using a language model-based deep learning approach. *Nature Methods*, 21(12):2287–2298, 2024.

Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *International conference on machine learning*, pp. 20503–20521. PMLR, 2022.

Andrew C Stelzer, Jeremy D Kratz, Qi Zhang, and Hashim M Al-Hashimi. Rna dynamics by design: biasing ensembles towards the ligand-bound state. *Angewandte Chemie (International ed. in English)*, 49(33):5731, 2010.

Saisai Sun and Lin Gao. Contrastive pre-training and 3d convolution neural network for rna and small molecule binding affinity prediction. *Bioinformatics*, 40(4):btae155, 2024.

Brent Townshend, Matias Kaplan, and Christina D Smolke. Highly multiplexed selection of rna aptamers against a small molecule library. *Plos one*, 17(9):e0273381, 2022.

Aline Umuhire Juru, Neeraj N Patwardhan, and Amanda E Hargrove. Understanding the contributions of conformational changes, thermodynamics, and kinetics of rna–small molecule interactions. *ACS chemical biology*, 14(5):824–838, 2019.

Izhar Wallach and Abraham Heifets. Most ligand-based classification benchmarks reward memorization rather than generalization. *Journal of chemical information and modeling*, 58(5):916–932, 2018.

Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1225–1234. ACM, 2016.

Katherine Deigan Warner, Christine E Hajdin, and Kevin M Weeks. Principles for targeting rna with drug-like small molecules. *Nature reviews Drug discovery*, 17(8):547–558, 2018.

Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9 (2):513–530, 2018.

Wentao Xia, Jiasai Shu, Chunjiang Sang, Kang Wang, Yan Wang, Tingting Sun, and Xiaojun Xu. The prediction of rna-small-molecule ligand binding affinity based on geometric deep learning. *Computational Biology and Chemistry*, 115:108367, 2025.

Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, et al. Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling*, 59(8):3370, 2019.

Kamyar Yazdani, Deondre Jordan, Mo Yang, Christopher R Fullenkamp, David R Calabrese, Robert Boer, Thomas Hilimire, Timothy EH Allen, Rabia T Khan, and John S Schneekloth Jr. Machine learning informs rna-binding chemical space. *Angewandte Chemie*, 135(11):e202211358, 2023.

Hanwen Zhou, Yue Hu, Yulong Zheng, Jiefu Li, Jielong Peng, Jiang Hu, Yun Yang, Wei Chen, Guoqing Zhang, and Zefeng Wang. A foundation language model to decipher diverse regulation of rnas. *Genome Biology*, 26 (1):301, 2025.

Chen Zhuo, Jiaming Gao, Anbang Li, Xuefeng Liu, and Yunjie Zhao. A machine learning method for rna–small molecule binding preference prediction. *Journal of Chemical Information and Modeling*, 64(19):7386–7397, 2024.

A Formal Analysis of Sparse Conformer Competition

This appendix formalizes sparse conformer competition under fixed learned conformer scores. We characterize how pairwise utility differences and ligand-conditioned evidence determine conformer ranking, sparse Top-2 support selection, and posterior-weighted score aggregation. Empirical analyses of conformer competition and screened regularization are provided separately in Appendix B.

A.1 Finite-State Formulation of Conformer Selection

Consider a fixed RNA–ligand pair and a finite set of surrogate conformers $\mathcal{Z} = \{z_1, \dots, z_K\}$. We condition on the learned model outputs for this pair. Each surrogate conformer z_k is associated with a ranking energy $E_{\mathrm{rank},k}$ defined in Section 3.4, ligand-conditioned evidence ℓ_k, and an energy-style interaction score $\Delta G_{\mathrm{bind},k}$.

The corresponding prior-guided conformer-ranking distribution is

$$\pi_k = \frac{\exp(-E_{\mathrm{rank},k}/\tau_{\mathrm{prior}})}{\sum_{j=1}^{K} \exp(-E_{\mathrm{rank},j}/\tau_{\mathrm{prior}})}.$$

PIPER then forms the centered conformer utility

$$u_k = \lambda_\pi \log \pi_k + \lambda_\ell(\ell_k - \bar{\ell}), \qquad \bar{\ell} = \frac{1}{K}\sum_{j=1}^{K} \ell_j,$$

followed by the dense pre-sparsification posterior

$$\tilde{q}_k = \frac{\exp(u_k/\tau_{\mathrm{post}})}{\sum_{j=1}^{K} \exp(u_j/\tau_{\mathrm{post}})}.$$

Let $\mathcal{S}_2$ denote the indices of the two largest utilities. The sparse Top-2 posterior is obtained by retaining and renormalizing the corresponding mass,

$$q_k = \begin{cases} \dfrac{\tilde{q}_k}{\sum_{j\in\mathcal{S}_2} \tilde{q}_j}, & k \in \mathcal{S}_2, \\ 0, & k \notin \mathcal{S}_2. \end{cases}$$

The retained conformer-specific scores are then aggregated as

$$\Delta G_{\mathrm{eff}} = -\tau_{\mathrm{agg}} \log \sum_{k=1}^{K} q_k \exp\left(-\frac{\Delta G_{\mathrm{bind},k}}{\tau_{\mathrm{agg}}}\right).$$

This finite-state formulation isolates the conformer-selection and aggregation mechanism under fixed learned scores, independently of the neural parameterization and training dynamics. Details of the surrogate field, screened residual, and probe construction are provided separately in Appendix B.5.

A.2 Coherence of the Surrogate RNA Scaffold

Proposition A.1 (Local coherence of the surrogate RNA scaffold)**.** *Let*

$$\mathbf{x}_i^{(0)} = \left[\frac{i}{N_R}, \sin\left(\frac{i}{2}\right), \cos\left(\frac{i}{3}\right)\right]^\top, \qquad i = 1, \dots, N_R.$$

Then, for $i = 1, \dots, N_R - 1$,

$$\frac{1}{N_R} \le \left\|\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_i^{(0)}\right\| \le \sqrt{\frac{1}{N_R^2} + \frac{1}{4} + \frac{1}{9}},$$

and for every interior index $i = 2, \dots, N_R - 1$,

$$\left\|\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_{i-1}^{(0)}\right\| \ge \frac{2}{N_R}.$$

Moreover, the second finite difference is uniformly bounded:

$$\left\|\mathbf{x}_{i+1}^{(0)} - 2\mathbf{x}_i^{(0)} + \mathbf{x}_{i-1}^{(0)}\right\| \le \sqrt{\frac{1}{16} + \frac{1}{81}}.$$

Consequently, the surrogate scaffold is deterministic and order-preserving, with a non-degenerate centered tangent and a uniformly bounded second-difference descriptor.

Proof of Proposition A.1. For adjacent residues,

$$\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_i^{(0)} = \left[\frac{1}{N_R}, \sin\left(\frac{i+1}{2}\right) - \sin\left(\frac{i}{2}\right), \cos\left(\frac{i+1}{3}\right) - \cos\left(\frac{i}{3}\right)\right]^\top.$$

The first coordinate immediately gives

$$\left\|\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_i^{(0)}\right\| \ge \frac{1}{N_R}.$$

By the mean value theorem,

$$\left|\sin\left(\frac{i+1}{2}\right) - \sin\left(\frac{i}{2}\right)\right| \le \frac{1}{2}, \qquad \left|\cos\left(\frac{i+1}{3}\right) - \cos\left(\frac{i}{3}\right)\right| \le \frac{1}{3},$$

since the derivatives of $\sin$ and $\cos$ are bounded by one in magnitude. Therefore,

$$\left\|\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_i^{(0)}\right\| \le \sqrt{\frac{1}{N_R^2} + \frac{1}{4} + \frac{1}{9}}.$$

For an interior residue,

$$\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_{i-1}^{(0)} = \left[\frac{2}{N_R}, \sin\left(\frac{i+1}{2}\right) - \sin\left(\frac{i-1}{2}\right), \cos\left(\frac{i+1}{3}\right) - \cos\left(\frac{i-1}{3}\right)\right]^\top,$$

and hence

$$\left\|\mathbf{x}_{i+1}^{(0)} - \mathbf{x}_{i-1}^{(0)}\right\| \ge \frac{2}{N_R}.$$

Thus, the centered finite-difference tangent is nonzero at every interior residue.

For the second finite difference, the first coordinate is linear in i and therefore cancels:

$$\mathbf{x}_{i+1}^{(0)} - 2\mathbf{x}_i^{(0)} + \mathbf{x}_{i-1}^{(0)} = \left[0, \Delta^2 \sin\left(\frac{i}{2}\right), \Delta^2 \cos\left(\frac{i}{3}\right)\right]^\top,$$

where $\Delta^2 f(i) = f(i+1) - 2f(i) + f(i-1)$. For a twice differentiable scalar function sampled with step size h,

$$|f(t+h) - 2f(t) + f(t-h)| \le h^2 \sup_\xi |f''(\xi)|.$$

Applying this bound to $\sin t$ with $h = \frac{1}{2}$ and to $\cos t$ with $h = \frac{1}{3}$ gives

$$\left|\Delta^2 \sin\left(\frac{i}{2}\right)\right| \le \frac{1}{4}, \qquad \left|\Delta^2 \cos\left(\frac{i}{3}\right)\right| \le \frac{1}{9}.$$

Therefore,

$$\left\|\mathbf{x}_{i+1}^{(0)} - 2\mathbf{x}_i^{(0)} + \mathbf{x}_{i-1}^{(0)}\right\| \le \sqrt{\frac{1}{16} + \frac{1}{81}}.$$

Together, these bounds establish a deterministic, order-preserving local geometry with non-degenerate tangent directions and bounded second-difference features, providing a well-defined scaffold for the conformer perturbations used by PIPER.

A.3 Utility Decomposition and Ranking Effects

Proposition A.2 (Utility-gap decomposition). *For any two surrogate conformers i and j,*

$$u_i - u_j = -\frac{\lambda_\pi}{\tau_{\text{prior}}} \left(E_{\text{rank},i} - E_{\text{rank},j}\right) + \lambda_\ell(\ell_i - \ell_j).$$

Proof. Substituting the definition of π_k into u_k gives

$$u_k = \lambda_\pi \left(-\frac{E_{\text{rank},k}}{\tau_{\text{prior}}} - \log Z_\pi\right) + \lambda_\ell(\ell_k - \bar{\ell}),$$

where

$$Z_\pi = \sum_j \exp\left(-\frac{E_{\text{rank},j}}{\tau_{\text{prior}}}\right).$$

Subtracting u_j from u_i cancels both the normalization term $\log Z_\pi$ and the common centering term $\bar{\ell}$, yielding the stated identity. Thus, pairwise conformer competition decomposes exactly into a prior-guided ranking contribution and a ligand-conditioned evidence contribution.

Corollary A.3 (Ranking-driven separation). *Suppose $\ell_i = \ell_j$, $E_{\text{rank},i} < E_{\text{rank},j}$, $\lambda_\pi > 0$, and $\tau_{\text{prior}} > 0$. Then $u_i > u_j$. Holding the learned scores fixed, the utility gap $u_i - u_j$ increases linearly with λ_π and with $1/\tau_{\text{prior}}$.*

Proof. With $\ell_i = \ell_j$, Proposition A.2 reduces to

$$u_i - u_j = \frac{\lambda_\pi}{\tau_{\text{prior}}} \left(E_{\text{rank},j} - E_{\text{rank},i}\right).$$

Because $E_{\text{rank},j} - E_{\text{rank},i} > 0$, the utility difference is positive. For fixed ranking energies, its magnitude is proportional to both λ_π and $1/\tau_{\text{prior}}$, so increasing the ranking weight or decreasing the ranking temperature strengthens the separation between the two conformers.

A.4 Stability of Sparse Conformer Selection

Proposition A.4 (Local stability of Top-2 support). *Let*

$$u_{(1)} \geq u_{(2)} \geq u_{(3)}$$

denote the three largest conformer utilities, and define the support margin

$$m = u_{(2)} - u_{(3)} > 0.$$

If the utility vector is perturbed by δ with

$$\|\delta\|_\infty < \frac{m}{2},$$

then the retained Top-2 support remains unchanged.

Proof of Proposition A.4. Let i_2 and j_3 denote the indices of the second- and third-largest utilities, respectively. Under the perturbation δ, their separating margin satisfies

$$(u_{i_2} + \delta_{i_2}) - (u_{j_3} + \delta_{j_3}) \geq m - |\delta_{i_2}| - |\delta_{j_3}| \geq m - 2\|\delta\|_\infty.$$

When $\|\delta\|_\infty < m/2$, the right-hand side is strictly positive. Thus, the smallest retained utility remains larger than the largest excluded utility, and the Top-2 support is unchanged. The ordering of the two retained conformers may change, but their support set is preserved.

Corollary A.5 (Bounds for sparse energy-style aggregation). *Let $\mathcal{S}_2$ denote the retained Top-2 support, with $q_k \geq 0$, $\sum_{k\in\mathcal{S}_2} q_k = 1$, and $\tau_{\text{agg}} > 0$. Then*

$$\min_{k\in\mathcal{S}_2} \Delta G_{\text{bind},k} \leq \Delta G_{\text{eff}} \leq \sum_{k\in\mathcal{S}_2} q_k \Delta G_{\text{bind},k}.$$

Proof. Let

$$g_k = \Delta G_{\mathrm{bind},k}, \qquad k_{\min} \in \arg\min_{k \in \mathcal{S}_2} g_k.$$

Since $g_k \geq g_{k_{\min}}$ for every $k \in \mathcal{S}_2$,

$$\sum_{k \in \mathcal{S}_2} q_k e^{-g_k/\tau_{\mathrm{agg}}} \leq \sum_{k \in \mathcal{S}_2} q_k e^{-g_{k_{\min}}/\tau_{\mathrm{agg}}} = e^{-g_{k_{\min}}/\tau_{\mathrm{agg}}},$$

where the last equality follows from $\sum_{k \in \mathcal{S}_2} q_k = 1$. Applying $-\tau_{\mathrm{agg}} \log$ gives

$$\Delta G_{\mathrm{eff}} \geq g_{k_{\min}} = \min_{k \in \mathcal{S}_2} \Delta G_{\mathrm{bind},k}.$$

For the upper bound,

$$\exp\left(-\frac{\Delta G_{\mathrm{eff}}}{\tau_{\mathrm{agg}}}\right) = \sum_{k \in \mathcal{S}_2} q_k e^{-g_k/\tau_{\mathrm{agg}}}.$$

Because the exponential function is convex, Jensen's inequality yields

$$\sum_{k \in \mathcal{S}_2} q_k e^{-g_k/\tau_{\mathrm{agg}}} \geq \exp\left(-\frac{1}{\tau_{\mathrm{agg}}} \sum_{k \in \mathcal{S}_2} q_k g_k\right).$$

Applying $-\tau_{\mathrm{agg}} \log$ to both sides therefore gives

$$\Delta G_{\mathrm{eff}} \leq \sum_{k \in \mathcal{S}_2} q_k g_k,$$

which establishes the claimed upper bound.

A.5 Further Properties of Conformer Competition

The preceding results characterize sparse support selection and aggregation. Two additional properties follow directly from the utility decomposition. When prior-guided ranking and ligand-conditioned evidence favor the same conformer, their contributions reinforce the same utility ordering. When they favor different conformers, the utility-gap identity gives an explicit threshold determining which contribution dominates.

Proposition A.6 (Agreement regime)**.** *Suppose a surrogate conformer $k^\star$ satisfies*

$$E_{\mathrm{rank},k^\star} \leq E_{\mathrm{rank},j}, \qquad \ell_{k^\star} \geq \ell_j$$

for every competitor $j \neq k^\star$, with at least one of the two inequalities strict for each j. If $\lambda_\pi > 0$, $\lambda_\ell > 0$, and $\tau_{\mathrm{prior}} > 0$, then

$$u_{k^\star} > u_j \qquad \text{for all } j \neq k^\star.$$

Consequently, $k^\star$ has the largest dense-posterior mass and is retained by the Top-2 posterior for every $\tau_{\mathrm{post}} > 0$.

Proof. For any competitor $j \neq k^\star$, Proposition A.2 gives

$$u_{k^\star} - u_j = -\frac{\lambda_\pi}{\tau_{\mathrm{prior}}} \left(E_{\mathrm{rank},k^\star} - E_{\mathrm{rank},j}\right) + \lambda_\ell(\ell_{k^\star} - \ell_j).$$

The first term is nonnegative because $E_{\mathrm{rank},k^\star} \leq E_{\mathrm{rank},j}$, and the second is nonnegative because $\ell_{k^\star} \geq \ell_j$. At least one inequality is strict for each competitor, and all coefficients are positive, so

$$u_{k^\star} - u_j > 0.$$

Hence $k^\star$ has the largest utility. Since the dense posterior is a strictly increasing function of utility for $\tau_{\mathrm{post}} > 0$, it preserves this ordering, and $k^\star$ is necessarily included in the retained Top-2 support.

Proposition A.7 (Conflict threshold between ranking and ligand-conditioned evidence)**.** *For any two surrogate conformers i and j,*

$$u_i \geq u_j \quad \Longleftrightarrow \quad \lambda_\ell(\ell_i - \ell_j) \geq \frac{\lambda_\pi}{\tau_{\mathrm{prior}}} \left(E_{\mathrm{rank},i} - E_{\mathrm{rank},j}\right).$$

Thus, when prior-guided ranking and ligand-conditioned evidence favor different conformers, their pairwise ordering is determined by a linear trade-off between the two score gaps.

Proof. By Proposition A.2,

$$u_i - u_j = -\frac{\lambda_\pi}{\tau_{\text{prior}}} \left(E_{\text{rank},i} - E_{\text{rank},j}\right) + \lambda_\ell(\ell_i - \ell_j).$$

Therefore,

$$u_i - u_j \geq 0$$

if and only if

$$\lambda_\ell(\ell_i - \ell_j) \geq \frac{\lambda_\pi}{\tau_{\text{prior}}} \left(E_{\text{rank},i} - E_{\text{rank},j}\right).$$

This relation makes explicit the pairwise trade-off between the prior-guided ranking contribution and ligand-conditioned evidence in the conformer-selection rule.

A.6 Scope of the Analysis

The analysis establishes utility-gap decomposition, ranking-driven separation, local stability of Top-2 support, bounds on sparse energy-style aggregation, and the agreement and conflict regimes governing conformer competition. Together, these results characterize the finite-state selection and aggregation behavior of PIPER under fixed learned scores.

B Additional Method and Experimental Details

B.1 Overall Training and Inference Procedure

The end-to-end training and inference procedure of PIPER is summarized in Algorithm 1.

Algorithm 1 PIPER Multi-Conformer Reasoning and Joint Learning

Require: RNA sequence R, secondary structure S, ligand graph L, and label y
Ensure: Interaction probability $\hat{y}$ and training objective $\mathcal{L}$

```
Encode RNA residues and ligand atoms to obtain {h_i^R}_{i=1}^{N_R}, {h_a^L}_{a=1}^{N_L}, and h^L
Construct the shared RNA scaffold {x_i^(0)} and ligand reference chart {r_a^L}
Generate surrogate conformers {z_k}_{k=1}^K through prior-modulated deformation
Compute sparse RNA-residue–ligand-atom cross-graph interactions
for k = 1, ..., K do
    Pool ligand-conditioned residue features under z_k to obtain h_k
    Compute the ligand-conditioned adjustment b_k^cond and ranking energy E_{rank,k}
    Compute conformer evidence ℓ_k
end for
Form the prior-guided conformer-ranking distribution π and attention distribution α
Compute utilities {u_k}_{k=1}^K and dense posterior q̃
Retain the two highest-utility conformers and renormalize their mass to obtain sparse posterior q
for k = 1, ..., K do
    Evaluate the surrogate field φ_k and screened residual loss L_{res,k}
    Compute the conformer-specific energy-style score ΔG_{bind,k}
end for
Aggregate the retained conformers using q to obtain ΔG_eff
Set ŷ ← sigmoid(−ΔG_eff/τ_out)
Set L ← L_BCE + λ_KL D_KL(π‖α) + λ_PDE K^{-1} Σ_k L_{res,k}
return ŷ, L
```

B.2 Surrogate-Conformer Construction

The surrogate-conformer generator starts from the deterministic RNA coordinate chart defined in Eq. 2, with coordinates $\mathbf{x}_i^{(0)}$ for $i = 1, \ldots, N_R$. The ligand unit-circle chart is used only as a coarse geometric reference during candidate construction. No learned ligand atom features, ligand graph representation, or RNA–ligand cross-graph features are used to construct the surrogate

coordinates; ligand-specific chemical information enters subsequently during conformer evaluation and competition.

The reported configuration uses three surrogate conformers during both training and evaluation. Their deformation amplitudes are linearly spaced between 0 and 0.08, with the first conformer retained as the undeformed anchor state. A trainable basis matrix $B \in \mathbb{R}^{K_{\max} \times 4}$ is initialized entrywise from $\mathcal{N}(0, 0.05^2)$. Its first three columns weight the radial, tangential, and normal deformation directions, while the fourth modulates the stochastic perturbation through a sigmoid gate.

For residue i in conformer k, the displacement is

$$\boldsymbol{\delta}_{k,i} = \eta_k a_{k,i} g_i \left(B_{k1} \mathbf{r}_i + B_{k2} \mathbf{t}_i + B_{k3} \mathbf{n}_i \right) + \eta_k a_{k,i} \, \mathrm{sigmoid}(B_{k4}) \boldsymbol{\epsilon}_{k,i},$$

where η_k controls the conformer-specific deformation amplitude, $a_{k,i}$ is the structure-dependent residue scale, and g_i combines a coarse interface-distance term with the structured residue score. The conformer-specific multiplier within $a_{k,i}$ is linearly spaced from 0.92 to 1.08 across the active conformers.

During training, $\boldsymbol{\epsilon}_{k,i} \sim \mathcal{N}(\mathbf{0}, I)$ is applied only to valid RNA residues, introducing local stochastic variation around the shared scaffold. At evaluation time, the stochastic term is removed, yielding a deterministic set of surrogate conformers. The structured residue score therefore shapes conformer construction, whereas ligand-specific discrimination is introduced later through cross-graph interaction, conformer-specific pooling, and utility-based competition.

B.3 Utility Decomposition and Ligand Correction

To examine how prior-guided ranking and ligand-conditioned evidence jointly shape conformer competition, we decompose each pairwise utility difference into its ranking and ligand-evidence contributions. The analysis is performed on 1,000 validation pairs, yielding 3,000 pairwise conformer comparisons, with each comparison oriented from the higher-utility to the lower-utility conformer.

The two contributions favor the same ordering in 63.9% of comparisons. Across all comparisons, ligand-conditioned evidence contributes 9.2% of the total absolute utility contribution, computed as

$$\frac{\sum |\Delta u_{\mathrm{lig}}|}{\sum |\Delta u_{\mathrm{rank}}| + \sum |\Delta u_{\mathrm{lig}}|}.$$

Figure 4 visualizes the agreement and disagreement regimes together with the pairwise equal-utility boundary.

B.4 Representative Cases of Ligand-Conditioned Conformer Competition

To illustrate how ligand-conditioned evidence affects conformer competition, we select representative cases from the same 1,000 validation pairs considered in the utility-decomposition analysis. Three cases are examined: ranking–evidence agreement, a Top-1 change after ligand-evidence integration, and a Top-2 support change with the Top-1 conformer unchanged. Within each category, the representative case is selected using a pre-specified criterion that maximizes the smaller of the relevant prior-guided-ranking and final-utility margins.

In the agreement case, ligand-conditioned evidence preserves both the original conformer ordering and the retained Top-2 support. In the Top-1 change case, it promotes $C3$ from the lowest-ranked conformer to the highest-utility conformer, changing the support from $\{C1, C2\}$ to $\{C1, C3\}$. In the Top-2 support-change case, $C1$ remains the highest-utility conformer, but $C3$ replaces $C2$ as the second retained state. Together, these cases illustrate how ligand-conditioned evidence can either preserve the prior-guided ordering, change the leading conformer, or refine the second retained state, supporting the role of sparse Top-2 reasoning in resolving conformer-level ambiguity. The selected cases and their sparse posterior allocations are summarized in Table 3 and visualized in Figure 5.

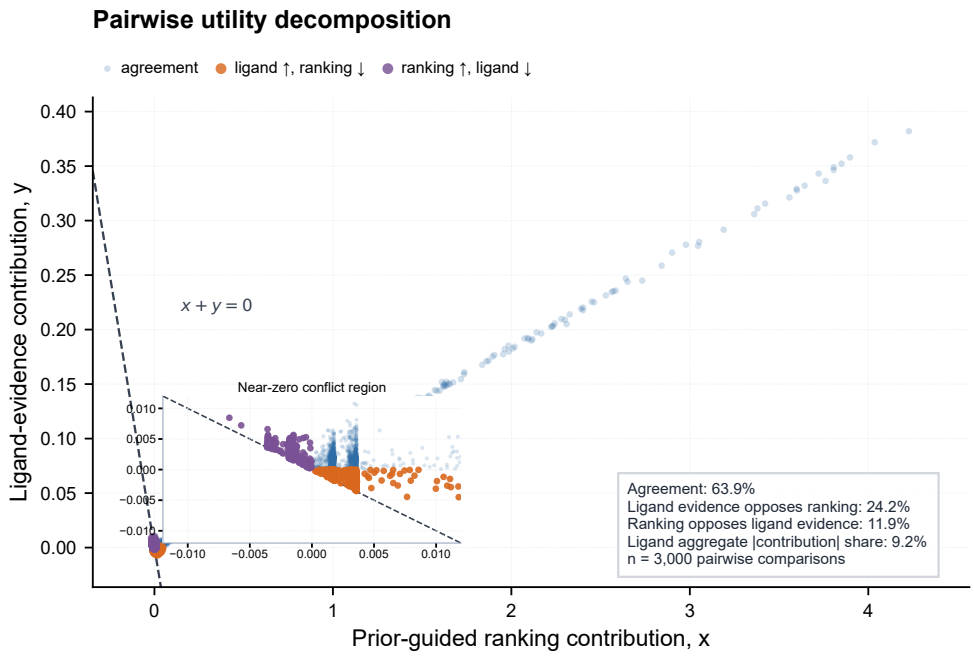


Figure 4: Pairwise utility decomposition on the SMRTnet validation set. Points compare higher- and lower-utility candidate conformers; colors indicate agreement or opposition between prior-guided ranking and ligand-conditioned evidence. The dashed line marks $x + y = 0$.

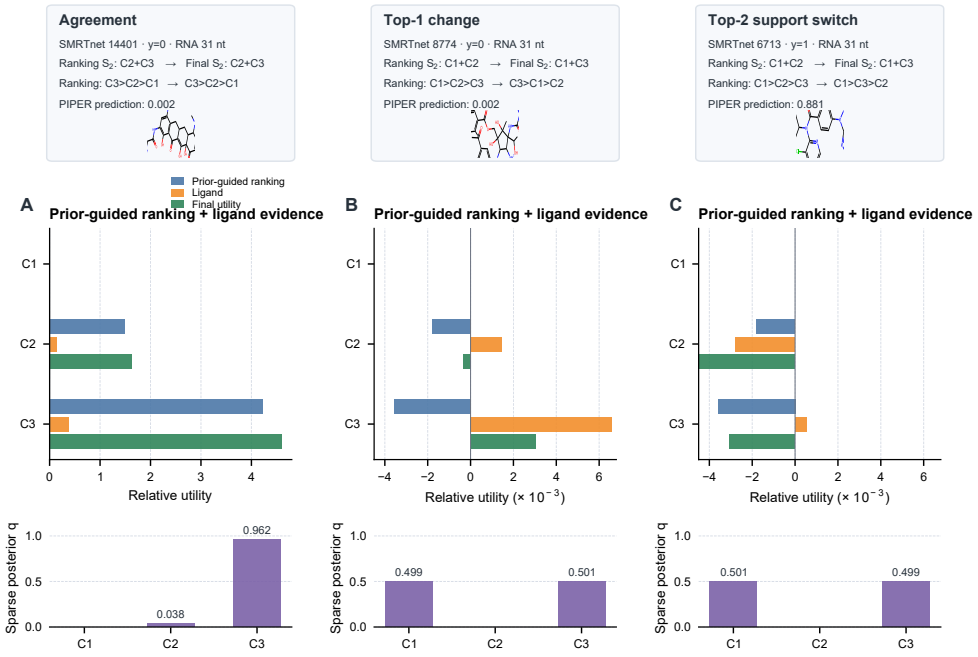


Figure 5: Representative cases of ligand-conditioned conformer competition on the SMRTnet validation set. Panels show ranking–evidence agreement, a Top-1 change after ligand-evidence integration, and a Top-2 support change with the Top-1 conformer unchanged. Each case reports the ligand structure, utility decomposition, sparse posterior, and prediction.

Table 3: Representative ligand-conditioned support cases from 1,000 SMRTnet validation pairs. Relative utilities are anchored at $C1 = 0$; q denotes the sparse posterior.

Case	ID	y	Ranking order	Final utility order	Ranking Top-2	Final Top-2
Agreement	14401	0	C3>C2>C1	C3>C2>C1	C2+C3	C2+C3
Posterior $q(C1, C2, C3)$						(0.000, 0.038, 0.962)
Top-1 change	8774	0	C1>C2>C3	C3>C1>C2	C1+C2	C1+C3
Posterior $q(C1, C2, C3)$						(0.499, 0.000, 0.501)
Top-2 switch	6713	1	C1>C2>C3	C1>C3>C2	C1+C2	C1+C3
Posterior $q(C1, C2, C3)$						(0.501, 0.000, 0.499)

B.5 Physics-Informed Field Parameterization

For each generated surrogate conformer, PIPER defines a conformer-specific scalar field conditioned on the probe coordinate, conformer representation, and ligand representation:

$$\phi_k(\mathbf{p}) = \mathrm{MLP}_\theta\left([\mathbf{p}; h_k; h_L]\right).$$

The field MLP has hidden dimension 96 and consists of two linear layers with Gaussian error linear unit (GELU) activations followed by a scalar output layer.

Each conformer is evaluated at six probe points, comprising four interface-focused probes and two support probes. The interface-focused probes are constructed from valid RNA residues nearest to the ligand-chart centroid, with their coordinates interpolated toward the centroid using coefficients spanning 0.2 to 0.8. During training, the two support probes are sampled uniformly without replacement from valid residues, whereas evaluation uses the next interface-ranked residues in a deterministic order. If padding is required, the final available residue index is repeated. A Gaussian perturbation with standard deviation 0.05 is applied to all probe coordinates during training and removed at evaluation time.

For conformer k, the source term is constructed from unit signed surrogate sources centered on the RNA residues and ligand reference-chart coordinates:

$$\rho_k(\mathbf{p}) = \sum_i m_i \exp\left(-\frac{\|\mathbf{p} - \mathbf{x}_{k,i}\|^2}{2\sigma^2}\right) - \sum_a n_a \exp\left(-\frac{\|\mathbf{p} - \mathbf{r}_a^{\mathrm{L}}\|^2}{2\sigma^2}\right),$$

where m_i and n_a denote validity indicators. The corresponding source-aware screened residual is

$$\mathrm{Res}_k(\mathbf{p}) = -\nabla^2\phi_k(\mathbf{p}) + \kappa^2\phi_k(\mathbf{p}) - \rho_k(\mathbf{p}),$$

with $\sigma = 1.2$ and $\kappa = 0.15$ in the surrogate coordinate system. The Laplacian is computed by automatic differentiation with respect to the probe coordinates, and the conformer-specific residual loss is obtained by averaging the squared residual over its probe points. This loss enters the training objective with $\lambda_{\mathrm{PDE}} = 0.02$ and is averaged across all three generated surrogate conformers, independently of the Top-2 support used for prediction.

The source construction uses unit surrogate sources rather than atom-specific partial charges. Accordingly, the field branch provides a physics-informed constraint in the surrogate coordinate space rather than a calibrated atomistic electrostatic model.

B.6 Evaluation Protocols under Distribution Shift

Table 4 summarizes the construction of the SMRTnet distribution-shift protocols. The fixed split uses a row-level partition and therefore permits RNA identities, ligand identities, and ligand scaffolds to overlap across training, validation, and test sets. Ligand cold-start groups ligands using Morgan fingerprints (radius 2, 2,048 bits) with a Tanimoto cutoff of 0.70, whereas RNA cold-start groups RNA sequences using 5-mer Jaccard similarity at a threshold of 0.80. Scaffold-hopping groups Bemis–Murcko scaffolds after Tanimoto-based merging at 0.50. Dual cold-start jointly enforces ligand-cluster and RNA-sequence separation: an RNA–ligand pair is retained only when its ligand and RNA groups are assigned to the same train/validation/test partition, and mixed assignments are excluded. For acyclic ligands, a deterministic RDKit ElementGraph hash of the complete molecular graph is used as a grouping identifier to ensure consistent partitioning.

Table 5 reports the resulting pair counts, entity coverage, and positive-class prevalence for each partition. The dual cold-start protocol excludes 8,353 mixed-assignment pairs as a buffer; these pairs are not used for training, validation, or test evaluation.

Table 4: Evaluation protocols and split constraints on SMRTnet. Partitions use existing labeled pairs only; no additional negative sampling is introduced.

Protocol	RNA grouping constraint	Ligand grouping constraint	Scaffold constraint	Held-out partition rule
Fixed split	No disjointness constraint	No disjointness constraint	No disjointness constraint	Row-level partition
Ligand cold-start	No disjointness constraint	Disjoint by Morgan-fingerprint ligand clusters (Tanimoto 0.70)	Not additionally constrained beyond ligand clustering	Ligand-cluster-disjoint partition
RNA cold-start	Disjoint by 5-mer Jaccard RNA clusters (0.80)	No disjointness constraint	No disjointness constraint	RNA-cluster-disjoint partition
Scaffold-hopping	No disjointness constraint	Grouped through scaffold partitioning	Bemis–Murcko scaffold groups with Tanimoto merging (0.50)	Scaffold-group-disjoint partition
Dual cold-start	Disjoint by 5-mer Jaccard RNA clusters (0.80)	Disjoint by Morgan-fingerprint ligand clusters (Tanimoto 0.70)	No additional scaffold constraint	Pairs retained only when RNA and ligand groups share the same split label; mixed assignments excluded

Table 5: Partition statistics for the SMRTnet evaluation protocols. Pair counts, positive-label prevalence, and the numbers of unique RNAs, ligands, and scaffolds are reported for each train/validation/test partition.

Protocol	Partition	Pairs	Pos. (%)	#RNA	#Ligand	#Scaffold
Fixed split	Train	19,749	33.3	4,770	341	252
	Val	2,469	33.3	1,811	277	217
	Test	2,469	33.3	1,836	278	216
Ligand cold-start	Train	19,740	33.3	4,788	245	194
	Val	2,478	33.3	1,770	51	38
	Test	2,469	33.3	1,817	45	38
RNA cold-start	Train	19,712	33.3	3,913	341	252
	Val	2,476	34.5	469	282	217
	Test	2,499	32.3	456	294	224
Scaffold-hopping	Train	19,410	33.3	4,761	260	191
	Val	2,784	33.3	2,111	37	27
	Test	2,493	33.3	1,756	44	34
Dual cold-start	Train	15,805	33.0	3,812	245	194
	Val	274	31.0	199	40	31
	Test	255	31.8	190	35	32

Each distribution-shift protocol enforces its corresponding disjointness constraint for both validation and test data. Thus, ligand clusters are unseen during training under ligand cold-start, RNA-sequence clusters are unseen under RNA cold-start, and scaffold groups are unseen under scaffold-hopping after the prescribed scaffold-merging step. Because SMRTnet dataset already contains labeled positive and negative RNA–ligand pairs, all partitions are constructed directly from the observed pairs without additional negative sampling. Model selection is performed exclusively on the validation partition defined by each protocol, with the held-out test partition reserved for final evaluation. The scaffold-hopping and ligand cold-start protocols apply the specified scaffold- and ligand-group separation rules.

Table 6 reports the complete numerical results corresponding to the four distribution-shift settings summarized in Figure 2. All methods are evaluated on matched held-out partitions. AUC and AUPR are computed directly from test scores, whereas ACC and F1 use the validation-selected threshold for each model and protocol.

Table 6: OOD performance under four SMRTnet-derived protocols: ligand cold-start, RNA cold-start, scaffold-hopping, and dual cold-start. ACC and F1 use validation-selected thresholds.

Protocol	Model	AUC	AUPR	ACC	F1
Ligand cold-start	PIPER	0.7377	0.5879	0.6784	0.5623
	DeepRSMA	**0.7771**	**0.5934**	0.6748	0.6149
	SMRTnet	0.5967	0.4309	0.6772	0.1512
	SDNE	0.7180	0.5786	**0.7691**	**0.6388**
	BINDTI	0.5921	0.4802	0.6865	0.2964
	RNAsmol	0.6589	0.4926	0.6877	0.1876
	DeepConv-DTI	0.6842	0.5332	0.7072	0.3421
RNA cold-start	PIPER	**0.8702**	**0.7584**	**0.7899**	**0.7022**
	DeepRSMA	0.5105	0.3241	0.6791	0.0000
	SMRTnet	0.5009	0.3183	0.5750	0.2872
	SDNE	0.4732	0.3090	0.5586	0.2852
	BINDTI	0.5508	0.3548	0.5042	0.4451
	RNAsmol	0.6904	0.5027	0.5594	0.5559
	DeepConv-DTI	0.5808	0.3865	0.6635	0.4090
Scaffold-hopping	PIPER	**0.7322**	**0.5584**	0.4954	0.5526
	DeepRSMA	0.5276	0.3419	0.5933	0.2929
	SMRTnet	0.4854	0.3591	0.5828	0.3131
	SDNE	0.6154	0.4358	0.6398	**0.5955**
	BINDTI	0.6252	0.4992	0.5961	0.4635
	RNAsmol	0.6634	0.5012	**0.6827**	0.4313
	DeepConv-DTI	0.6896	0.5194	0.3402	0.5017
Dual cold-start	PIPER	**0.6931**	**0.4864**	**0.5931**	**0.5204**
	DeepRSMA	0.6135	0.4335	0.4978	0.4579
	SMRTnet	0.5366	0.3821	0.4762	0.4211
	SDNE	0.5000	0.2987	0.2987	0.4600
	BINDTI	0.5474	0.3436	0.3810	0.4348
	RNAsmol	0.4672	0.2891	0.2987	0.4600
	DeepConv-DTI	0.6405	0.4139	0.4416	0.4861

Table 7: Matched OOD ablations across four SMRTnet distribution-shift protocols. The single-conformer variant uses one surrogate conformer during training and evaluation with conformer conditioning disabled, whereas the no-screened-residual-loss variant retains the neural-field features but removes screened-residual regularization. Boldface indicates the best value for each metric within each protocol.

Protocol	PIPER	Single conformer	No screened-residual loss
Ligand cold-start	0.7377 / **0.5879**	0.7265 / 0.5861	**0.7385** / 0.5803
RNA cold-start	**0.8702** / **0.7584**	0.8693 / 0.7466	0.8494 / 0.7249
Scaffold-hopping	**0.7322** / **0.5584**	0.7318 / 0.5491	0.6985 / 0.5395
Dual cold-start	0.6931 / **0.4864**	0.6923 / 0.4748	**0.7030** / 0.4765

To isolate the contribution of PIPER's core components from architectural differences among heterogeneous baselines, Table 7 compares independently trained matched variants evaluated on the same OOD partitions. PIPER achieves the highest AUPR across all four protocols. Compared with the single-conformer variant, it improves AUPR by 0.0018, 0.0118, 0.0093, and 0.0116 under ligand cold-start, RNA cold-start, scaffold-hopping, and dual cold-start, respectively. Removing the screened-residual loss also decreases AUPR consistently across all four settings, with the largest reduction observed under RNA cold-start (0.0335). The corresponding AUC differences are less uniform: the no-screened-residual-loss variant yields marginally higher AUC under ligand cold-start and dual cold-start, while remaining lower in AUPR in both cases.

B.7 Reproducibility and Baseline Protocols

To ensure controlled comparisons, all methods are evaluated on matched train/validation/test partitions using the same metric definitions. The distribution-shift experiments likewise use common held-out partitions defined by the corresponding ligand, RNA, scaffold, or dual exclusion criteria. Baseline methods retain their method-specific configurations while following the same data partitions and evaluation protocol as PIPER.

Table 8 summarizes the evaluation protocol and the main PIPER configuration. AUC and AUPR measure ranking performance, whereas ACC and F1 are computed using a decision threshold selected

Table 8: Experimental protocol and PIPER training configuration.

Item	Protocol
Baseline data partitions	All comparison models are evaluated under the same data partitions as PIPER. The ligand cold-start, RNA cold-start, and scaffold-hopping experiments use matched held-out sets generated under the same split criteria.
Evaluation criteria	All models use the same metric definitions. AUC and AUPR evaluate ranking quality, whereas ACC and F1 use a threshold selected by maximum validation F1 and then fixed on the test partition.
Model selection	Validation performance is used for model selection before held-out test evaluation.
PIPER configuration	The same model configuration is used across the reported experiments unless otherwise specified.
Epochs / batch size	80 epochs with a batch size of 32.
Optimizer / schedule	AdamW optimizer with learning rate 7×10^{-4}, weight decay 10^{-4}, and cosine annealing.
Weight averaging / precision	Exponential moving-average decay 0.995; automatic mixed precision disabled.
Hidden dimension	96.
Cross-graph sparsity	Each RNA residue attends to at most 8 highest-scoring ligand atoms in the residue–atom interaction module.
Reference score scale	$\tau_0 = 298.15 \times 1.987 \times 10^{-3}$. This fixed numerical reference initializes learned score scales and is not a physical complex temperature.
Prior scale	$\tau_{\text{prior}} = \tau_0 \exp(\theta_{\text{prior}})$, with $\exp(\theta_{\text{prior}})$ initialized to 0.7 and learned end-to-end.
Posterior temperature	$\tau_{\text{post}} = \exp(\theta_{\text{post}})$, initialized to 0.75 and learned end-to-end.
Posterior weights	Prior and ligand-score weights initialized to 3.0 and 0.6, respectively, and learned end-to-end.
Aggregation and interface scales	$\tau_{\text{agg}} = \tau_0 \exp(\theta_{\text{agg}})$ and $\tau_{\text{out}} = \tau_{\text{agg}} \exp(\theta_{\text{out}})$ are learned; the interface-distance scale is initialized to 1.5 and lower-bounded by 0.25.
Objective weights	The KL and PDE terms use weights $\lambda_{\text{KL}} = 0.02$ and $\lambda_{\text{PDE}} = 0.02$.
Surrogate conformers	3 during training and 3 during evaluation.
Posterior support	The posterior retains the Top-2 highest-utility surrogate conformers and renormalizes their weights.
Conformer conditioning	Conformer-conditioning weight $\lambda_{\text{cond}} = 1.5$.
Screened-residual probes	6 during training and 6 during evaluation: 4 interface-focused probes and 2 support probes.
Hardware	Single NVIDIA A100-SXM4-40GB GPU.
Software environment	Rocky Linux 9.7, Python 3.8.20, PyTorch 2.4.1, and CUDA 11.8.

by maximizing F1 on the validation set and then fixed for test evaluation. Table 1 reports benchmark performance as mean $\pm$ standard deviation over three independent runs. Detailed construction of the distribution-shift partitions is provided in Appendix B.6.

All experiments were conducted on a single NVIDIA A100-SXM4-40GB graphics processing unit (GPU) using Rocky Linux 9.7, Python 3.8.20, PyTorch 2.4.1, and CUDA 11.8.

Table 9 compares the computational cost of the single-conformer variant, PIPER ($K = 3$), the no-screened-residual-loss variant, and the $K = 5$ configuration under the same hardware setting. PIPER requires 0.812 ± 0.034 ms per pair for inference, compared with 0.775 ± 0.016 ms for the single-conformer variant and 0.806 ± 0.019 ms for $K = 5$. Relative to the single-conformer model, PIPER increases training time per epoch by 2.1% and peak GPU memory by 19.8%. All evaluated variants contain approximately 0.448 million trainable parameters.

Table 9: Efficiency of PIPER variants on one NVIDIA A100-SXM4-40GB GPU. Training time is per epoch and inference time is end-to-end on 2,469 validation pairs. Values are mean ± standard deviation over five timing rounds.

Variant	**Params (M)**	**Train time/epoch (s)**	**Inference (ms/pair)**	**Peak GPU memory (MB)**
Single conformer	0.448071	41.88	0.775 ± 0.016	77.9
PIPER ($K = 3$)	0.448079	42.77	0.812 ± 0.034	93.3
No screened-residual loss	0.448079	35.11	0.748 ± 0.012	73.2
$K = 5$	0.448087	46.20	0.806 ± 0.019	108.6

To assess statistical significance, we compare PIPER with each baseline across the three main datasets using paired DeLong tests on AUC. The resulting p-values are jointly adjusted over all 18 comparisons using the Holm procedure and are reported in Table 10.

Table 10: Global Holm-adjusted p-values for paired DeLong tests comparing PIPER with each baseline on AUC.

Dataset	DeepRSMA	SMRTnet	SDNE	BINDTI	RNAsmol	DeepConv-DTI
SMRTnet main-fixed	7.97×10^{-13}	7.68×10^{-6}	1.02×10^{-37}	2.16×10^{-44}	1.98×10^{-37}	1.73×10^{-4}
Robin	1.74×10^{-10}	0.5316	3.81×10^{-35}	0.0300	2.29×10^{-19}	2.53×10^{-13}
Biosensor	1.59×10^{-4}	1.48×10^{-4}	1.01×10^{-8}	0.0784	8.91×10^{-9}	2.15×10^{-11}

B.8 Held-Out Target-Conditioned Ligand Ranking

To complement pair-level evaluation under ligand cold-start, we examine whether the predicted scores support ligand prioritization for individual RNA targets. For each RNA, all observed held-out ligands are ranked by their continuous model scores, without applying a classification threshold. Each candidate pool contains the complete set of test ligands associated with that RNA, without synthetic decoys or resampled negatives.

RNAs are included if they have at least five held-out ligand candidates and contain both positive and negative labels. For visualization, we select the RNA with the largest eligible candidate pool, resolving ties lexicographically by RNA sequence. The resulting pool contains seven held-out ligands, including five positives and two negatives.

We evaluate ranking quality using Recall@K, mean reciprocal rank (MRR), and enrichment factor at K (EF@K). Recall@K measures the fraction of positive ligands recovered within the top K positions, MRR is the reciprocal rank of the highest-ranked positive ligand, and EF@K compares the positive fraction among the top K ligands with that of the complete candidate pool. Table 11 reports both the selected-RNA results and macro-averaged performance across all eligible RNAs, while Figure 6 visualizes the corresponding held-out ligand rankings. These results complement the global ligand cold-start metrics by showing that PIPER can retain useful ligand-prioritization ability within held-out RNA-specific candidate pools.

Table 11: Target-conditioned ranking of held-out ligands under the ligand cold-start protocol. The selected RNA has the largest eligible candidate pool, with lexicographic tie-breaking.

Model	Best rank	Recall@1	Recall@3	Recall@5	MRR	EF@3
PIPER	1	0.200	0.600	1.000	1.000	1.400
Single conformer	1	0.200	0.600	0.800	1.000	1.400
DeepConv-DTI	1	0.200	0.400	0.800	1.000	0.933
Macro-average across all eligible RNAs ($n = 15$)						
PIPER	–	0.243	0.752	0.978	0.933	1.353
Single conformer	–	0.243	0.736	0.964	0.933	1.340
DeepConv-DTI	–	0.132	0.478	0.953	0.711	0.879

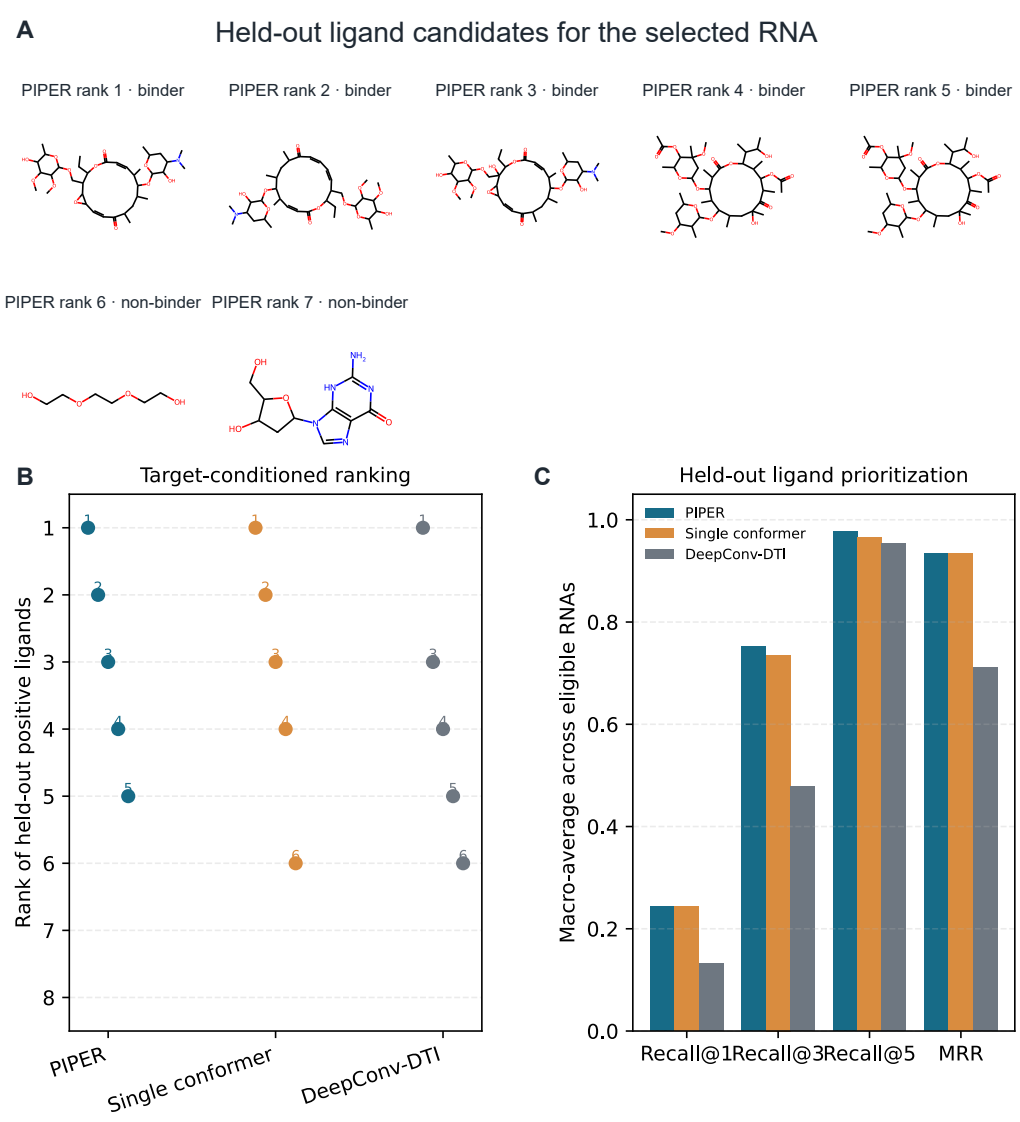


Figure 6: Target-conditioned ligand prioritization under ligand cold-start. (A) Held-out ligand candidates for one RNA. (B) Positive-ligand ranks for PIPER, the single-conformer variant, and DeepConv-DTI. (C) Macro Recall@K and MRR across eligible RNAs.

Table 12: Scaffold-hopping performance across ligand ring-topology strata.

Model	Acyclic AUC	Acyclic AUPR	Single-ring AUC	Single-ring AUPR	Multi-ring AUC	Multi-ring AUPR
PIPER	0.6324	0.4382	0.8112	0.6575	0.7506	0.5670
DeepRSMA	0.4568	0.3021	0.6031	0.3946	0.5500	0.3613
SMRTnet	0.5586	0.4268	0.4315	0.3074	0.4513	0.3122
SDNE	0.6205	0.4552	0.5735	0.4172	0.6364	0.4540
BINDTI	0.4898	0.3414	0.7645	0.6330	0.6312	0.5298
RNAsmol	0.5296	0.3725	0.7638	0.5673	0.6743	0.5378
DeepConv-DTI	0.4910	0.4364	0.5297	0.3147	0.4900	0.3262

Test subgroup counts: Acyclic: $n = 939$, $n_+ = 313$, Single-ring: $n = 636$, $n_+ = 212$, Multi-ring: $n = 918$, $n_+ = 306$.

Table 13: PIPER scaffold-hopping performance stratified by ligand-chemistry complexity.

Stratification	Group	n	n_+	AUC	AUPR
Heavy atoms	Low	831	277	0.6183	0.4266
	Mid	831	291	0.8181	0.6660
	High	831	263	0.7376	0.5523
Hetero atoms	Low	831	382	0.6883	0.6928
	Mid	831	206	0.7727	0.5212
	High	831	243	0.7303	0.4801

B.9 Chemical Subgroup Analysis under Scaffold Hopping

To examine whether scaffold-hopping performance is concentrated in particular regions of ligand chemical space, we further stratify the held-out test set by ring topology and molecular complexity. The scaffold-hopping protocol uses Bemis–Murcko framework partitioning with scaffold-level Tanimoto merging. Figure 7 and Table 12 report results across acyclic, single-ring, and multi-ring ligands, while Table 13 further characterizes PIPER across heavy-atom and hetero-atom tertiles.

The held-out test set contains 939 pairs involving acyclic ligands, 636 pairs involving single-ring ligands, and 918 pairs involving multi-ring ligands. PIPER achieves the highest AUC across all three ring-topology strata and the highest AUPR on the single-ring and multi-ring subsets. Its AUPR on the acyclic subset remains competitive. Together, these results indicate that PIPER's scaffold-hopping advantage in ranking performance is not restricted to a single ring-topology regime, while the strongest AUPR gains occur for ring-containing ligands.

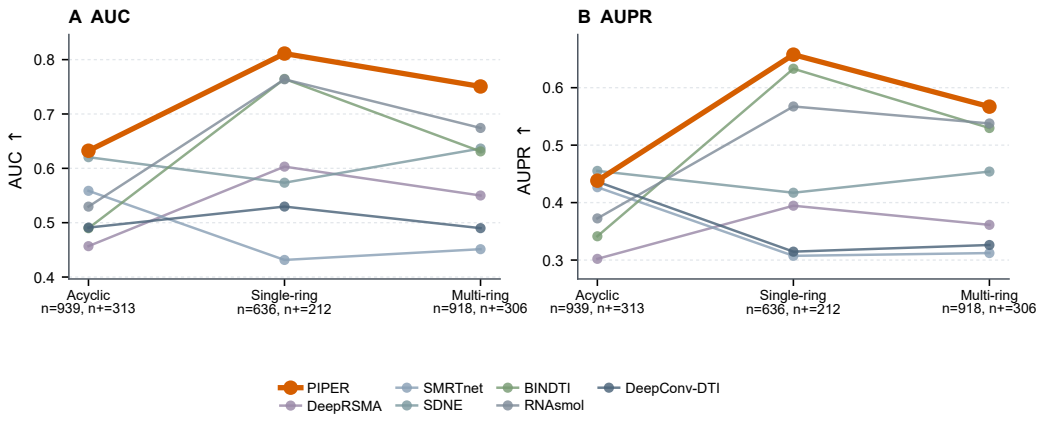


Figure 7: Scaffold-hopping performance stratified by ligand ring topology. Panels show AUC and AUPR on the held-out test set.

B.10 Representative Scaffold-Hopping Cases

To examine how PIPER's conformer competition operates on individual scaffold-hopping examples, we select representative test pairs for which PIPER predicts correctly but the single-conformer variant does not. Each selected pair contains a ligand scaffold absent from the training set, and at least one of the six comparison baselines also makes an incorrect prediction.

For model m with score s_m, validation-selected threshold t_m, and binary label y, we define the signed classification margin as

$$d_y(s_m - t_m), \qquad d_y = 2y - 1.$$

A positive signed margin denotes a correct prediction, whereas a negative margin denotes an error. We therefore rank eligible pairs according to how clearly PIPER predicts them correctly while the single-conformer variant and at least one baseline predict them incorrectly. Specifically, the ranking score is the minimum of the PIPER correct margin, the absolute single-conformer error margin, and the largest baseline error margin. We retain the two highest-ranked positive and two highest-ranked negative cases. This procedure selects examples with clear prediction differences between PIPER and the comparison models.

The four selected cases are summarized in Table 14 and visualized in Figure 8. For each case, the conformer-level bars decompose the utility into prior-guided ranking and ligand-evidence contributions relative to $C1$, while the overlaid points show the resulting sparse posterior allocation. The figure also compares the prediction probabilities of PIPER and the single-conformer variant and reports how many of the six baselines make an incorrect prediction. Together, these cases complement the aggregate scaffold-hopping results by providing a concrete view of how ligand-conditioned conformer competition and sparse posterior selection operate on previously unseen ligand scaffolds.

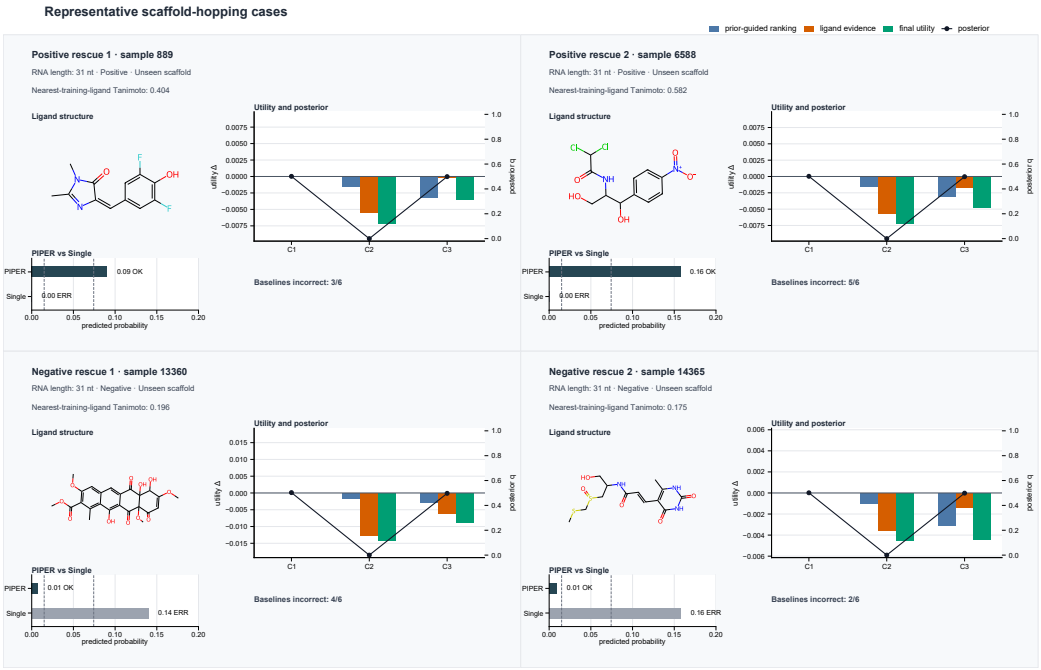


Figure 8: Representative scaffold-hopping cases selected by a pre-specified criterion. Cards summarize ligand similarity, utility contributions, sparse posterior support, and PIPER versus single-conformer predictions; baseline calls are listed in Table 14.

Table 14: Representative scaffold-hopping cases selected by a pre-specified criterion. Each case has a correct PIPER prediction, an incorrect single-conformer prediction, an unseen scaffold, and an incorrect baseline prediction. In the two p columns, *pred.* denotes the thresholded predicted label.

Case	y	RNA length	nearest-train ligand Tanimoto	PIPER p (pred.)	Single p (pred.)	Incorrect fixed baselines
positive_1	1	31	0.404	0.090 (1)	0.000 (0)	DeepRSMA, SDNE, BINDTI
positive_2	1	31	0.582	0.157 (1)	0.000 (0)	DeepRSMA, SMRTnet, SDNE, BINDTI, RNAsmol
negative_1	0	31	0.196	0.007 (0)	0.140 (1)	DeepRSMA, SDNE, BINDTI, DeepConv-DTI
negative_2	0	31	0.175	0.009 (0)	0.158 (1)	RNAsmol, DeepConv-DTI

B.11 Physics Specificity Controls and Sensitivity

To examine whether the predictive gains are associated with the specific source-aware screened constraint rather than with the presence of an auxiliary field regularizer alone, we compare PIPER with two matched controls. Both controls use the same data partitions and training protocol as PIPER, with all non-residual components held fixed.

The generic field-magnitude control removes the differential operator and source term and penalizes only the magnitude of the learned field:

$$\mathcal{L}_{\text{L2-field}} = \frac{1}{n}\sum_{r=1}^{n} |\phi_\theta(x_r)|^2 .$$

The homogeneous PDE control retains the Laplacian operator but removes both the source term and screening component:

$$\mathcal{R}_{\text{hom}}(x) = -\nabla^2\phi_\theta(x), \qquad \mathcal{L}_{\text{hom}} = \frac{1}{n}\sum_{r=1}^{n} |\mathcal{R}_{\text{hom}}(x_r)|^2 .$$

Neither control therefore imposes the complete source-aware screened constraint, $-\nabla^2\phi + \kappa^2\phi - \rho$.

Screened-residual consistency. To directly assess satisfaction of the complete screened constraint, we evaluate PIPER and the three physics-related variants under the same source-aware residual. Let $\mathbf{p}_{k,r}$ denote the rth probe point of conformer k:

$$\text{Res}_k(\mathbf{p}_{k,r}) = -\nabla^2\phi_k(\mathbf{p}_{k,r}) + \kappa^2\phi_k(\mathbf{p}_{k,r}) - \rho_k(\mathbf{p}_{k,r}).$$

All variants are evaluated on the same 2,469 validation RNA–ligand pairs, using three surrogate conformers and six fixed probe points per conformer. To normalize for differences in source magnitude, we use the zero-field reference obtained by setting $\phi = 0$:

$$\mathcal{R}_{\text{scr}} = \frac{1}{N}\sum_{k,r} |\text{Res}_k(\mathbf{p}_{k,r})|^2 , \qquad \mathcal{R}_0 = \frac{1}{N}\sum_{k,r} |\rho_k(\mathbf{p}_{k,r})|^2 , \qquad \widetilde{\mathcal{R}} = \frac{\mathcal{R}_{\text{scr}}}{\mathcal{R}_0},$$

where N is the total number of probe evaluations.

PIPER achieves a normalized residual of $\widetilde{\mathcal{R}} = 0.001233$, whereas the generic L2, homogeneous PDE, and no-screened-residual-loss variants yield 0.999964, 1.000097, and 1.000051, respectively. Under the same aligned-source audit, the shuffled-source variant attains $\widetilde{\mathcal{R}} = 0.009739$, with a screened residual of 1.349397 relative to a zero-field reference of 138.550295. Although source shuffling preserves substantial consistency with the screened constraint, it remains notably less consistent than PIPER and is accompanied by lower predictive performance. The lower performance of the shuffled-source variant despite its low residual further supports the importance of preserving source–geometry alignment. Table 15 summarizes the residual statistics together with the corresponding predictive metrics. Accordingly, $\widetilde{\mathcal{R}}$ provides a measure of consistency with the imposed surrogate screened constraint.

Table 15: Consistency with the source-aware screened constraint. Lower residual values indicate stronger constraint satisfaction; normalized residuals are scaled by the corresponding zero-field reference. The shuffled-source variant is evaluated with the common aligned source assignment for comparability.

Variant	Screened residual ↓	Zero-field reference	Normalized residual ↓	AUC	AUPR
PIPER (screened residual)	0.170858	138.603478	0.001233	0.8744	0.7691
Generic L2	138.363852	138.368782	0.999964	0.8518	0.7368
Homogeneous PDE	138.758695	138.745302	1.000097	0.8460	0.7242
Shuffled source assignments	1.349397	138.550295	0.009739	0.8685	0.7486
No screened-residual loss	138.706774	138.699747	1.000051	0.8382	0.7186

Table 16: Sensitivity to screened PDE weight and probe count on the SMRTnet fixed split. Variants are independently trained under matched settings; ACC and F1 use validation-selected thresholds.

Variant	λ_{PDE}	N_{probe}	AUC	AUPR	ACC	F1
λ_{PDE} sweep	0.01	6	0.8688	0.7449	0.7955	0.7171
λ_{PDE} sweep	0.02	6	0.8744	0.7691	0.7914	0.7203
λ_{PDE} sweep	0.05	6	0.8689	0.7587	0.7785	0.7132
N_{probe} sweep	0.02	4	0.8577	0.7336	0.7772	0.7027
N_{probe} sweep	0.02	8	0.8575	0.7396	0.7598	0.6847

Physics-parameter sensitivity. We further examine sensitivity to the screened-residual loss weight and the number of probe points on the SMRTnet fixed split. All sensitivity variants are trained independently under matched data partitions and training settings. The results are reported in Table 16.

To test whether the residual benefit depends on source–geometry alignment, we also trained a matched shuffled-source control. For each example, the RNA and ligand source assignments were permuted within the same source-location pool, while the screened operator, source counts, probe count, loss weight, and all other settings were kept fixed. This control reaches 0.8685 AUC and 0.7486 AUPR, below PIPER but above the generic L2 control in Table 2. The result supports a contribution from aligned surrogate source geometry, without implying that the surrogate coordinates are physically accurate.

With six probes, varying $\lambda_{\mathrm{PDE}} \in \{0.01, 0.02, 0.05\}$ yields AUC values between 0.8688 and 0.8744 and AUPR values between 0.7449 and 0.7691. Changing the probe count to four or eight reduces AUC to 0.8577 and 0.8575 and AUPR to 0.7336 and 0.7396, respectively. Among the tested probe counts, six probes provide the strongest performance, while the results remain relatively stable across the tested residual-loss weights, particularly in terms of AUC.

B.12 High-confidence Failure Case Analysis

To characterize failure modes beyond aggregate performance, we examine the highest-confidence false-negative (FN) and false-positive (FP) predictions made by PIPER on the fixed SMRTnet split. Cases are ranked by their distance from the validation-selected decision threshold, with ties resolved by the number of correct baseline predictions and then by sample ID. We retain the two highest-confidence FNs and two highest-confidence FPs. PIPER produces 159 FNs and 357 FPs among 2,469 test pairs at a decision threshold of 0.3169. Table 17 summarizes the selected cases, while Figure 9 reports model-wise test error rates using validation-selected F1 thresholds together with the predicted probabilities of the selected cases. These single-split error rates complement the three-run benchmark results reported in Table 1.

All four selected cases are highly similar to examples observed during training: each has a maximum train-ligand Morgan Tanimoto similarity of 1.000, a maximum train-RNA 5-mer Jaccard similarity of 1.000, and a ligand scaffold already seen during training. Both FNs are also misclassified by the single-conformer variant, whereas the two FPs receive near-unit predicted probabilities from both PIPER and the single-conformer variant. The posterior effective number, $K_{\mathrm{eff}} = \exp[-\sum_k q_k \log q_k]$, remains close to two across all selected cases, indicating that these errors are not associated with posterior collapse. Overall, the analysis shows that high-confidence errors are not confined to molecular distribution shift and can also occur for RNA–ligand pairs that closely resemble training examples.

Table 17: Pre-specified hard false-negative and false-positive cases on the SMRTnet fixed test split. Binary calls use validation-selected F1 thresholds; baseline correctness counts are computed across DeepRSMA, SMRTnet, SDNE, BINDTI, RNAsmol, and DeepConv-DTI.

Case	ID	True	PIPER p	Single p	No screened-residual loss p	Base. correct	Lig. sim.	RNA sim.	Scaffold	K_{eff}
Hard FN-1	1717	1	1.0e-06	1.0e-06	0.006	3/6	1.000	1.000	seen	2.00
Hard FN-2	1718	1	1.0e-06	2.3e-04	0.224	3/6	1.000	1.000	seen	2.00
Hard FP-1	20003	0	1.000	0.997	0.965	1/6	1.000	1.000	seen	2.00
Hard FP-2	21102	0	1.000	1.000	0.131	1/6	1.000	1.000	seen	1.98

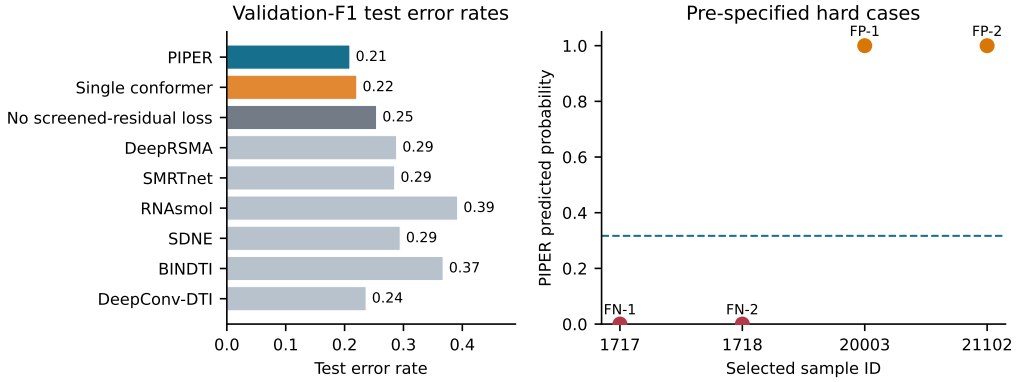


Figure 9: Validation-F1 error analysis on the SMRTnet fixed split. Left: test error rates obtained using each model's validation-selected F1 threshold. Right: predicted probabilities for selected false-negative and false-positive cases; the dashed line denotes PIPER's validation-selected threshold.

B.13 Secondary-Structure Perturbation Robustness

We evaluate PIPER's robustness to incomplete RNA secondary structures through controlled perturbations of the input dot-bracket representation. Because the structured prior incorporates pairing state, pair span, and local topology, we remove 5%, 10%, or 20% of the existing matched base pairs by replacing both bracket symbols with dots and deleting the corresponding base-pair graph edges. No new pairs are introduced. For a given RNA–structure input, the same perturbation mask is applied to all of its associated ligand pairs, and five deterministic masks are evaluated at each intermediate perturbation level. Table 18 reports AUC and AUPR on the SMRTnet fixed split.

Performance degrades progressively as secondary-structure information is removed. Relative to the unperturbed setting, deleting 5%, 10%, and 20% of matched pairs reduces AUC by 0.0045, 0.0100, and 0.0215, respectively, with corresponding AUPR decreases of 0.0075, 0.0152, and 0.0315. Removing all matched pairs reduces performance substantially to 0.5087 AUC and 0.3414 AUPR. These results indicate that PIPER remains relatively stable under moderate loss of pairing information, while the pronounced degradation after complete removal highlights the importance of secondary-structure signals to the model. This analysis specifically considers missing pair information through pair deletion.

B.14 Posterior-Support and Conformer-Count Sensitivity

We examine the effect of posterior support by comparing independently trained Top-1, Top-2, and All-3 variants, all constructed with the same three surrogate conformers. The Top-1 variant retains only the highest-utility conformer from the three-conformer model and therefore differs from the single-conformer ablation in Table 2, which is trained with $K = 1$ and conformer conditioning disabled. AUC and AUPR are computed directly from test scores, whereas ACC and F1 use the validation-selected threshold for each variant. Table 19 reports the resulting comparison.

Table 18: Inference-time robustness to missing RNA secondary-structure pairs on the SMRTnet fixed test split. Perturbations delete matched pairs without rewiring; intermediate doses average over five deterministic masks.

Perturbation	Realized drop (%)	AUC	ΔAUC	AUPR	ΔAUPR
0% pair deletion	0.0	0.8744	+0.0000	0.7691	+0.0000
5% pair deletion	5.0±0.4	0.8699±0.0013	-0.0045	0.7617±0.0023	-0.0075
10% pair deletion	9.9±0.2	0.8644±0.0026	-0.0100	0.7540±0.0053	-0.0152
20% pair deletion	19.4±0.5	0.8529±0.0041	-0.0215	0.7376±0.0050	-0.0315
All pairs removed	100.0	0.5087	-0.3658	0.3414	-0.4277

Table 19: Independently trained Top-1, Top-2, and All-3 posterior-support variants on the SMRTnet fixed split. All use three candidate conformers; Top-1 is distinct from the single-conformer ablation because it retains one state from the three-candidate model. ACC and F1 use validation-selected thresholds.

Retained support	AUC	AUPR	ACC	F1
Top-1	0.8692	0.7533	0.7695	0.7095
Top-2 (PIPER)	**0.8744**	**0.7691**	**0.7914**	**0.7203**
All-3	0.8626	0.7410	0.7744	0.6914

Utility-perturbation stability. To relate the local support-stability result to the learned utility landscape, we perturb the final three-conformer utility vectors on the fixed SMRTnet validation set. For each sample, let $m = u_{(2)} - u_{(3)}$, and scale each fixed-seed random perturbation direction such that $\|\delta\|_\infty = rm$. At $r = 0.10$, 0.25, and 0.40, all $128{,}000$ perturbations at each magnitude preserve the original Top-2 support, consistent with Proposition A.4. Support retention also remains 100% at $r = 0.50$ for the sampled directions, although this boundary case is not covered by the strict stability condition. Retention then decreases to 85.7% and 76.4% at $r = 0.75$ and 1.00, respectively. Figure 10 summarizes this behavior, showing that the learned support remains stable within the theoretically guaranteed regime and degrades progressively as perturbations move beyond this regime.

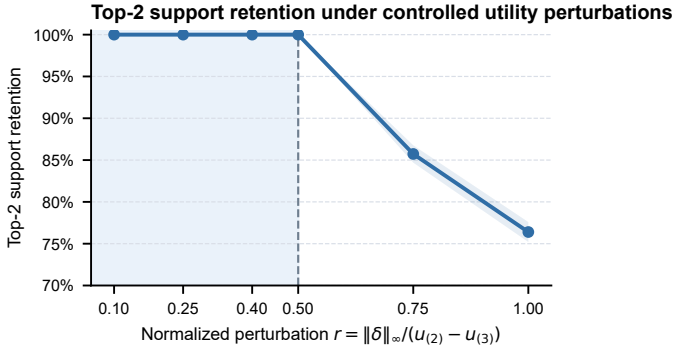


Figure 10: Top-2 support retention under controlled utility perturbations. Perturbations are normalized by the support margin $m = u_{(2)} - u_{(3)}$. The dashed line marks the formal stability boundary $r = 0.5$. Results are estimated from 1,000 validation samples with 128 seeded perturbation directions per sample; the shaded band denotes the standard deviation across directions.

We further evaluate sensitivity to the number of surrogate conformers by varying K from one to five, using the same value during training and evaluation. For $K = 1$, the posterior is necessarily concentrated on the sole available conformer, whereas for $K \geq 2$, it retains the two highest-utility conformers. All variants use matched data partitions and training settings. The $K = 1$ variant retains the conformer-conditioning pathway used throughout the conformer-count sweep and therefore differs from the single-conformer ablation in Table 2, which uses $K = 1$ with conformer conditioning disabled. Table 20 reports the corresponding conformer-count sensitivity. Among the tested settings, the Top-2 posterior with $K = 3$ achieves the highest AUC and AUPR.

Table 20: Sensitivity to the number of surrogate conformers on SMRTnet. The $K = 1$ variant retains conformer conditioning, whereas the single-conformer ablation in Table 2 disables it. For $K = 1$, the sole conformer is retained; for $K \geq 2$, the posterior retains the two highest-utility conformers. ACC and F1 use validation-selected thresholds.

K (train/eval)	AUC	AUPR	ACC	F1
1	0.8665	0.7510	0.7890	0.7120
2	0.8677	0.7513	0.7874	0.7036
3	0.8744	0.7691	0.7914	0.7203
4	0.8695	0.7531	0.7849	0.7137
5	0.8714	0.7638	0.7890	0.6987