RIBOSBDD AND RIBOSBDDBENCH: STRUCTURE BASED SMALL MOLECULE GENERATION FOR RNA

Anonymous authors
Paper under double-blind review

ABSTRACT

Structure based drug design (SBDD) for RNA is limited by scarce RNA ligand structures and flexible binding pockets. We introduce **RiboSBDD**, an $\mathrm{SE}(3)$ equivariant diffusion framework for RNA conditioned 3D ligand generation. Despite different backbones, interfaces between ligands and proteins or RNA share local packing principles. A unified all atom representation lets us transfer these geometric priors through protein pretraining and RNA specific fine tuning. To learn local pocket adjustment without receptor trajectories, we restrict motion to nucleobase rotations about glycosidic bonds with a fixed backbone. Synthetic rotations provide analytical inverse targets. This couples base correction with ligand denoising through a differentiable pocket transformation. To retain protein derived geometric guidance after RNA specialization, an adaptive sampler combines coordinate proposals from the frozen protein source and RNA target using local equal density composition and potential guided particle resampling. We leave atom types to the target. We also establish **RiboSBDDBench**, with 8,006 curated RNA pocket ligand instances, grouped structural splits, and evaluation across multiple metrics. On 713 test pockets, RiboSBDD improves contact recovery and steric compatibility over TargetDiff after RNA fine tuning. Together, RiboSBDD and RiboSBDDBench provide a generative framework that combines transfer with local flexibility and a standardized evaluation foundation for RNA targeted small molecule design.

1 INTRODUCTION

Ribonucleic acid (RNA) has regulatory, catalytic, and scaffolding roles and forms structures with pockets that bind small molecules (Serganov et al., 2004). The clinical use of risdiplam shows the potential of RNA targeting (Jaklevic, 2020). Geometric deep learning now enables designing ligands directly for RNA binding sites (Karroucha et al., 2026).

Structure based drug design (SBDD) uses target 3D pocket geometry to guide molecular generation. For proteins, equivariant models generate ligand atom types and coordinates within binding sites (Peng et al., 2022; Guan et al., 2023; Schneuing et al., 2024). Extending this to RNA faces two challenges (Figure 1a): *data scarcity*, because RNA ligand structures are much less abundant than protein ligand complexes; and *pocket flexibility*, as conformational changes alter accessible interactions (Stelzer et al., 2011). Few apo and holo structure pairs and receptor trajectories further complicate supervision of pocket motion.

These challenges motivate **RiboSBDD**, an $\mathrm{SE}(3)$ equivariant diffusion framework with three complementary design choices. First, ligand recognition by proteins and RNA shares steric packing, hydrogen bonding, and aromatic interactions despite different backbones. Shared all atom spatial graphs expose these local patterns to the same denoiser while retaining residue identities. Protein pretraining provides geometric priors, and RNA fine tuning adapts them to nucleotide specific environments (Figure 1b).

Second, we model a tractable subset of pocket flexibility instead of unrestricted receptor motion. Changes in the glycosidic torsion χ reorient nucleobase interaction edges and aromatic surfaces while preserving base geometry and keeping the backbone fixed. Bounded synthetic rotations provide analytical inverse targets ($-\Delta\chi$), enabling supervision without paired conformations or trajectories. A differentiable rotation operator couples an auxiliary base correction head to ligand denoising, so

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.

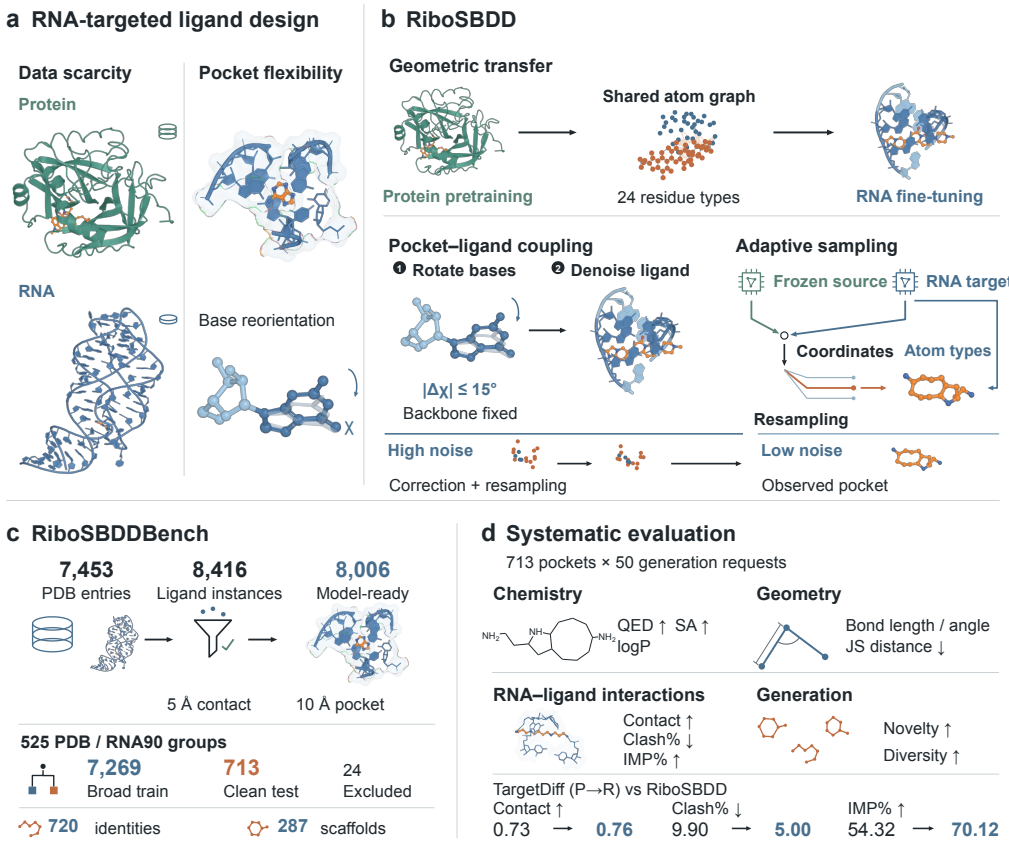


Figure 1: **Overview of RNA targeted ligand generation.** (a) Data scarcity and pocket flexibility. (b) RiboSBDD: protein to RNA transfer, local base correction, and adaptive sampling. (c) RiboSBDDBench: RNA-ligand complex curation and grouped splits. (d) Chemistry, geometry, RNA–ligand interactions, and generation. Molecular views are schematic.

ligand reconstruction also supervises local adjustment. During inference, corrections start from the observed pocket at each high noise step; low noise denoising uses the original pocket.

Third, transfer need not end at initialization. To retain protein derived geometric guidance after RNA specialization, we keep the frozen protein source during inference. Local equal density composition adapts coordinate mixing to each stochastic proposal. Potential guided particle resampling weights and selects candidate trajectories. Atom types are determined exclusively by the RNA target, allowing source guidance on spatial placement without replacing RNA adapted atom type predictions.

Our main contributions are:

1. **Method. RiboSBDD** integrates protein to RNA geometric transfer, jointly learned local base correction, and adaptive source and target sampling for RNA conditioned 3D ligand generation.
2. **Benchmark. RiboSBDDBench** provides 8,006 curated pocket ligand instances and evaluation across multiple metrics. Grouped splits prevent shared PDB identifiers or 90% RNA sequence identity clusters (RNA90) between training and test sets.
3. **Systematic findings.** Protein pretraining improves ligand geometry across all four evaluated architectures. Module comparisons show that local base correction improves contact recovery and reduces clashes, while adaptive source and target sampling further improves steric compatibility. On 713 test pockets, RiboSBDD achieves better contact recovery, fewer clashes, and a higher AnnapuRNA score improvement rate than TargetDiff after RNA fine tuning.

2 Related Work

Structure based drug design. Generators conditioned on binding pockets have primarily been developed for proteins. Autoregressive approaches construct ligands atom by atom (Luo et al., 2021; Peng et al., 2022) or fragment by fragment (Zhang et al., 2023). Diffusion models denoise coordinates and atom types (Guan et al., 2023; Schneuing et al., 2024; Guan et al., 2024). Recent work explores Bayesian flow networks (Qu et al., 2024), flow matching with interaction guidance (Zhou et al., 2025a), decoupled modeling of binding positions and conformations (Yang et al., 2026), and diffusion with synthesizability guidance and direct preference optimization (Liu et al., 2026). CBGBench provides a common evaluation framework (Lin et al., 2025). To model receptor flexibility, FlexSBDD jointly models pockets and ligands, with augmentation based on relaxation and side-chain repacking (Zhang et al., 2024). DynamicFlow learns from molecular dynamics ensembles (Zhou et al., 2025b), and Apo2Mol uses experimental pairs of apo and holo structures (Zheng et al., 2026). For RNA, we learn nucleobase correction under a fixed backbone from synthetic glycosidic rotations without requiring receptor trajectories or experimental apo and holo pairs.

RNA modeling and molecular recognition. RNA generative models address inverse folding, as in RiboDiffusion and gRNAde (Huang et al., 2024a; Joshi et al., 2025), and structure or ensemble generation, as in DynaRNA and RNAbpFlow (Li et al., 2025; Tarafder & Bhattacharya, 2026). RiboFlow jointly designs RNA sequence and structure conditioned on a ligand (Ma et al., 2025). Our task reverses this direction. We are given the RNA pocket, and we generate the ligand. Complementary predictive methods identify binding sites (Su et al., 2021), score poses (Stefaniak & Bujnicki, 2021), screen compounds (Carvajal-Patiño et al., 2025), or predict binding affinity (Huang et al., 2024b) and specificity (Xia et al., 2025). The Serna Bio GenAI platform generates analogs of seed compounds for chemical optimization targeting RNA (Allen et al., 2026). RiboSBDD focuses on *de novo* 3D ligand generation from an RNA pocket; RiboSBDDBench provides the corresponding structural data and evaluation protocol.

Model composition during inference. SuperDiff combines pretrained diffusion models during sampling (Skreta et al., 2025b), while Feynman–Kac correctors use particle weighting and sequential Monte Carlo resampling for controlled sampling (Skreta et al., 2025a). These ideas motivate us to retain the protein source after RNA adaptation. We compose its coordinate proposals with those of the RNA target and leave atom types to the target. Our proposal with noise dependent mixing and the particle potential define an empirical sampler, without an exact guarantee on the terminal distribution.

3 RiboSBDD

RiboSBDD predicts ligand atom types and Cartesian coordinates based on an RNA pocket (Figure 2). A unified all-atom diffusion backbone supports transfer from protein complexes. Differentiable local base correction links pocket flexibility to ligand denoising while keeping the backbone fixed. During inference, adaptive sampling integrates coordinate transitions from both the protein source and RNA target, using particle resampling for guidance.

3.1 All Atom Diffusion and Protein to RNA Transfer

Shared all atom spatial representation. We represent a biopolymer pocket as a point set $\mathcal{C} = \{(r_j, a_j, s_j)\}_{j=1}^{M}$, where r_j are atomic coordinates, a_j are atomic features, and s_j are residue types. Protein and RNA pockets share atomic features and a unified residue vocabulary of 20 amino acids and four canonical RNA nucleotides. Pocket inputs include all resolved atoms, with explicit hydrogens when available. A generated ligand consists of N heavy atoms, denoted by $\mathcal{L} = (X, V)$, where X are continuous Cartesian coordinates and V are one-hot categorical atom types across K element/aromaticity classes. Pocket and ligand atoms form a heterogeneous spatial graph that captures shared local packing geometry and preserves residue identities to enable learning nucleotide-specific interactions.

Pocket conditioned molecular diffusion. Following TargetDiff (Guan et al., 2023), we apply continuous Gaussian diffusion to ligand coordinates and a discrete uniform categorical process to atom types. For diffusion timesteps $t \in \{1, \ldots, T\}$, the forward diffusion marginals in pocket-

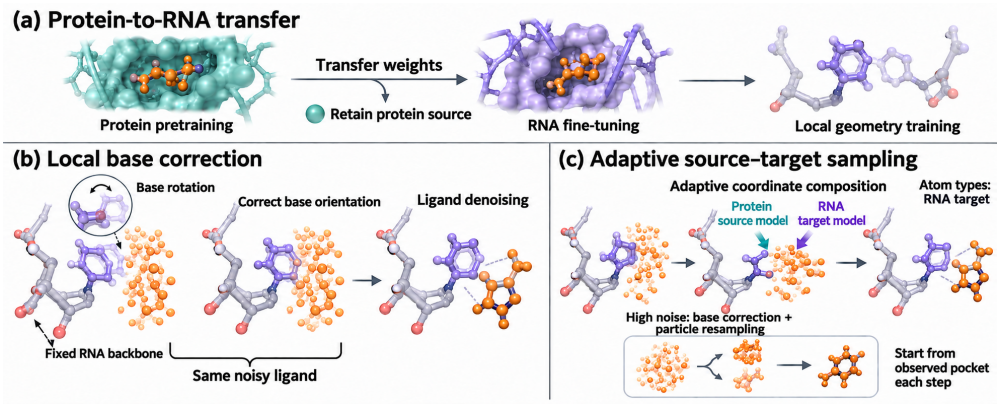


Figure 2: **RiboSBDD training and inference.** (a) Protein pretraining, RNA fine tuning, and local geometry training. (b) Differentiable base correction with a fixed backbone, then ligand denoising from the same noisy state. (c) Adaptive source–target coordinate composition and target-only atom types; base correction and particle resampling operate at high noise. Molecular views are schematic.

centered reference coordinates are parameterized as follows:

$$q(X_t \mid X_0) = \mathcal{N}\Big(\sqrt{\bar{\alpha}_t^x} X_0, (1-\bar{\alpha}_t^x) I\Big), \quad q(v_t \mid v_0) = \mathrm{Cat}(\bar{\alpha}_t^v v_0 + (1-\bar{\alpha}_t^v)\mathbf{1}/K), \tag{1}$$

where v_t represents an individual atom's categorical state, $\mathbf{1}$ is a vector of ones, and $\bar{\alpha}_t^z = \prod_{u=1}^{t}(1-\beta_u^z)$ is determined by separate coordinate ($z = x$) and atom type ($z = v$) noise schedules. A parameterized geometric denoiser $F_\theta(X_t, V_t, \mathcal{C}, t) = (\widehat{X}_0, \ell_0)$ predicts clean coordinates and unnormalized categorical logits. These outputs define the reverse transition kernels $p_\theta(X_{t-1} \mid X_t, V_t, \mathcal{C})$ and $p_\theta(V_{t-1} \mid X_t, V_t, \mathcal{C})$, which iteratively transform noise into a pocket-conditioned ligand. For a single corrupted complex, the diffusion loss is decomposed as follows:

$$\ell_{\mathrm{diff}}(\theta; \mathcal{C}) = \underbrace{\frac{1}{N}\|\widehat{X}_0 - X_0\|_F^2}_{\mathcal{L}_{\mathrm{pos}}} + \lambda_v \mathcal{L}_{\mathrm{type}}, \tag{2}$$

where $\mathcal{L}_{\mathrm{pos}}$ supervises coordinate reconstruction, and $\mathcal{L}_{\mathrm{type}}$ evaluates the atom averaged Kullback–Leibler divergence between the tractable true posterior $q(v_{t-1} \mid v_t, v_0)$ and the predicted reverse distribution $p_\theta(v_{t-1} \mid X_t, V_t, \mathcal{C})$, constructed from $\mathrm{softmax}(\ell_0)$ and the current noisy atom types through the categorical posterior for $t > 1$, and transitions to a reconstruction loss based on negative log likelihood at $t = 1$. We average losses over sampled diffusion timesteps, forward noise realizations, and complex samples. The ligand atom count N is sampled before diffusion using an invariant pocket-conditioned count estimator $p_\psi(N \mid \mathcal{C})$ and remains fixed for each trajectory (Appendix A.4).

Domain transfer through shared geometric generative priors. To address limited RNA ligand data, we pretrain F_{θ_s} on protein-ligand complexes using Eq. (2), then initialize F_{θ_r} with these weights and fine-tune on RNA complexes under the same objective. This approach reuses packing and covalent geometry priors while adapting to nucleotide-specific interactions. The frozen source F_{θ_s} is retained to guide sampling (Section 3.3).

3.2 Local Base Correction within Ligand Diffusion

We model local RNA pocket flexibility by allowing nucleobase rotations around glycosidic bonds (χ torsion). Treating nucleobases as rigid, planar aromatic units enables their interaction edges and stacking surfaces to reorient while preserving the sugar-phosphate backbone. We use an analytical rotation operator $\mathcal{R}(\mathcal{C}, \delta)$, which rotates eligible bases about the $C1'$–$N9$ axis for purines or the $C1'$–$N1$ axis for pyrimidines by angles δ defined for each residue.

Coupling base correction with ligand denoising. During training, we generate perturbed pocket conformations by sampling bounded angular offsets $a_r \sim \mathrm{TruncNorm}(0, \sigma^2)$ for each eligible

residue r, scaled with diffusion noise: $a_{t,r} = \sqrt{t/T}a_r$ and $\mathcal{C}_t = \mathcal{R}(\mathcal{C}, a_t)$. With probability 0.25, the pocket remains unperturbed, ensuring that training also includes ligand denoising conditioned on the observed pocket. We augment the RNA denoiser with an invariant base correction head for end-to-end pocket adaptation. In the first pass, the model processes $\mathcal{C}_t$ and (X_t, V_t), mean-pools node representations for each base, and predicts normalized angular adjustments $\widehat{u}_{t,r} \in [-1, 1]$. The scaled correction is:

$$\Delta\chi_{t,r} = b_t\widehat{u}_{t,r}, \quad b_t = \chi_{\max}\sqrt{t/T}, \quad \widetilde{\mathcal{C}}_t = \mathcal{R}(\mathcal{C}_t, \Delta\chi_t), \tag{3}$$

where $\chi_{\max} = 15^\circ$ sets the maximum correction magnitude. In the second forward pass, the denoiser processes the same noisy ligand state, now conditioned on the corrected pocket $\widetilde{\mathcal{C}}_t$:

$$(\widehat{X}_0, \ell_0) = F_{\theta_r}(X_t, V_t, \widetilde{\mathcal{C}}_t, t). \tag{4}$$

The rigid body rotation $\mathcal{R}$ is differentiable with respect to $\Delta\chi_t$, allowing the second pass loss ℓ_{diff} to backpropagate through the pocket transformation into the correction head. Ligand reconstruction therefore provides auxiliary supervision for base correction alongside generative denoising.

Multitask local geometry supervision. To promote recovery of observed base orientations, we use the inverse synthetic rotation $-a_{t,r}$ as the supervision target. For a complex with eligible base set $\mathcal{B}$, we define the normalized angular loss as follows:

$$\mathcal{L}_\chi = \frac{1}{|\mathcal{B}|}\sum_{r\in\mathcal{B}}\left(\frac{a_{t,r} + \Delta\chi_{t,r}}{\chi_{\max}}\right)^2. \tag{5}$$

To ensure steric and interaction compatibility between the pocket and ligand, we identify native contact pairs $\mathcal{S} = \{(i,j) : d^0_{ij} \le 6\,\text{Å}\}$, where $d^0_{ij} = \|X_{0,i} - r_j\|$ is the deposited interface distance, and $\widehat{d}_{ij} = \|\widehat{X}_{0,i} - \widetilde{r}_{t,j}\|$ is the reconstructed distance relative to corrected pocket atoms $\widetilde{r}_{t,j}$. Pocket compatibility is supervised through smooth distance recovery and a one-sided steric barrier:

$$\begin{aligned}\mathcal{L}_{\text{contact}} &= \frac{1}{|\mathcal{S}|}\sum_{(i,j)\in\mathcal{S}}\text{SmoothL1}(\widehat{d}_{ij} - d^0_{ij}),\\ \mathcal{L}_{\text{steric}} &= \frac{\mathbf{1}[t \le t_{\text{steric}}]}{N}\sum_{i=1}^{N}\left[2.2\,\text{Å} - \min_j \|\widehat{X}_{0,i} - r_j\|\right]_+^2.\end{aligned} \tag{6}$$

where $[\cdot]_+ = \max(0, \cdot)$, and $t_{\text{steric}} = 201$ restricts steric penalization to low noise steps, corresponding to network indices $\tau \le 200$. The contact loss promotes recovery of observed interface distances under the corrected pocket, while the steric loss penalizes overlaps with the original observed pocket. The total training objective minimizes the following joint expectation:

$$\mathcal{L}_{\text{local}} = \mathbb{E}_{t,q,a}\left[\ell_{\text{diff}}(\theta_r; \widetilde{\mathcal{C}}_t) + \lambda_\chi\mathcal{L}_\chi + \lambda_c\mathcal{L}_{\text{contact}} + \lambda_s\mathcal{L}_{\text{steric}}\right]. \tag{7}$$

Starting from the converged RNA fine-tuned checkpoint, we introduce these auxiliary objectives over three curriculum continuation stages (Algorithm 1). Stage schedules and loss weights are detailed in Appendix A.3.

3.3 Adaptive Source and Target Sampling

To integrate protein-derived geometric priors with RNA-specific adaptation, we retain both frozen models during inference and dynamically combine their reverse coordinate proposals. Discrete atom types are assigned exclusively by the RNA target. At each high noise step, the target applies local base correction to the observed pocket $\mathcal{C}$ (Eq. (3)) before denoising the ligand. Corrections are applied independently to $\mathcal{C}$ at each step, while low noise denoising uses the original observed pocket.

Product of experts coordinate composition. At a stochastic reverse step t, the source and target coordinate proposals are $P_i(x) = \mathcal{N}(x; \mu_i, \sigma_t^2 I)$ for $i \in \{s, r\}$, where $x \in \mathbb{R}^{3N}$ are vectorized ligand coordinates. The means μ_i are derived from the predicted clean coordinates $\widehat{X}_0^i$, and the shared diffusion schedule sets $\sigma_t^2 > 0$. Defining $d_t = \mu_r - \mu_s$, the weighted product satisfies the following identity for fixed κ:

$$P_r(x)^\kappa P_s(x)^{1-\kappa} = G_t(\kappa)\,\mathcal{N}(x; \mu_s + \kappa d_t, \sigma_t^2 I), \quad \text{with} \quad \log G_t(\kappa) = -\frac{\kappa(1-\kappa)\|d_t\|^2}{2\sigma_t^2}. \tag{8}$$

For fixed κ, Eq. (8) expresses the weighted product as a Gaussian with mean $\mu_s + \kappa d_t$ and a normalizing factor $G_t(\kappa)$, which motivates the particle potential described below.

Equal local density proposal. Instead of fixing the composition weight, we adapt it for each noise draw using the local equal density condition (Skreta et al., 2025b): $\log P_r(x) = \log P_s(x)$. For a proposal $x = \mu_s + \kappa d_t + \sigma_t \varepsilon$ with Gaussian noise $\varepsilon \sim \mathcal{N}(0, I)$, this condition reduces to the following linear constraint:

$$(2\kappa - 1)\|d_t\|^2 + 2\sigma_t \langle d_t, \varepsilon \rangle = 0. \tag{9}$$

For $d_t \neq \mathbf{0}$, solving Eq. (9) and clipping the resulting coefficient provides the adaptive coordinate transition to the next noise level t^-:

$$\kappa_t = \mathrm{clip}\left(\frac{1}{2} - \frac{\sigma_t \langle d_t, \varepsilon \rangle}{\|d_t\|^2}, -0.5, 1.5\right), \quad X_{t^-} = \mu_s + \kappa_t d_t + \sigma_t \varepsilon. \tag{10}$$

When unclipped, Eq. (10) projects the update onto the equal density hyperplane, removing the noise component along d_t and preserving orthogonal components. Clipping limits extrapolation but may break exact density equality. We set $\kappa_t = 1/2$ when the proposal means are numerically coincident, with $\|d_t\|^2 \leq 10^{-12}$. We apply coordinate composition throughout the reverse schedule, while limiting base correction to high-noise steps. The final coordinate update adds no noise and reuses the last composition coefficient (Eq. (17)).

Particle trajectory resampling via composition potential. To select among candidate trajectories, we simulate an ensemble of B ligand candidates per pocket and track the accumulated coordinate potential from Eq. (8). Particles are grouped by atom count into strata $\mathcal{G}_N = \{b : N_b = N\}$ of size B_N, ensuring that potentials are compared only among proposals of the same dimensionality. Within each stratum, cumulative weights and the effective sample size (ESS) are updated as follows:

$$\log w_b \leftarrow \log w_b + \log G_t(\kappa_{t,b}), \quad \widetilde{w}_b = \frac{w_b}{\sum_{a \in \mathcal{G}_N} w_a}, \quad \mathrm{ESS}_N = \left(\sum_{b \in \mathcal{G}_N} \widetilde{w}_b^2\right)^{-1}. \tag{11}$$

During high noise sampling, when $\mathrm{ESS}_N < B_N/2$, systematic resampling selects complete ligand states according to $\widetilde{w}_b$, followed by a weight reset (Algorithm 2) (Skreta et al., 2025a). For the noise-adaptive proposal, we use the factor derived for a fixed coefficient $G_t(\kappa_t)$ as an empirical particle weighting potential. Derivations and algorithmic details are provided in Appendices A.1–A.4.

4 Experiments

4.1 Experimental Setup

Dataset curation. We curated *RiboSBDDBench* from 7,453 RNA-containing experimental structures in the RCSB PDB (Burley et al., 2019). Applying chemical filters from HARIBOSS (Panei et al., 2022) and a 5 Å proximity criterion between ligand and RNA heavy atoms yields 8,416 ligand instances; model representability checks retain 8,006. Pockets include complete RNA residues with at least one heavy atom within 10 Å of the ligand. Instances sharing PDB identifiers or belonging to 90% RNA sequence identity clusters (RNA90) are assigned to the same split. The partition includes 7,269 Broad training instances and 713 Clean test instances; we exclude 24 non-Clean members of test groups. The Clean filter restricts nearby protein atoms. Details on curation and splitting are provided in Appendices C.1–C.2.

Baselines. We evaluate **TargetDiff** (Guan et al., 2023) (equivariant diffusion of coordinates and atom types), **DiffSBDD** (Schneuing et al., 2024) (SE(3) equivariant diffusion), **Pocket2Mol** (Peng et al., 2022) (autoregressive generation one atom at a time), and **GraphBP** (Liu et al., 2022) (autoregressive normalizing flows). Each baseline uses one of three regimes: protein training without RNA updates to the generator (P), training from scratch on RNA (R), or protein pretraining followed by RNA fine tuning (P→R). For TargetDiff and DiffSBDD, P uses an atom count histogram from RNA training ligands; the absence of RNA updates refers to the denoiser (Appendix B.5). RiboSBDD uses the P→R regime, incorporating local geometry training and adaptive sampling (Appendices A.3 and A.2).

Evaluation metrics. We assess four complementary dimensions: (1) *Chemistry*: QED (↑), normalized synthetic accessibility ($\mathrm{SA} = (10 - \mathrm{SA}_{\mathrm{raw}})/9$, ↑), and logP, with $[-0.4, 5.6]$ as a descriptive

Table 1: **RNA conditioned ligand generation.** All configurations use 713 test pockets and 50 generation requests per pocket. P/R: generator training on proteins/RNA; P→R: protein pretraining and RNA fine tuning. BL/BA: bond length/angle JS distances; Nov./Div.: novelty/diversity. Bold marks the best values among baselines and the full model; module controls are unranked.

Model	Train	Chemistry			Geometry		Interaction			Generation	
		QED ↑	SA ↑	logP	BL ↓	BA ↓	IMP% ↑	Contact ↑	Clash% ↓	Nov. ↑	Div. ↑
TargetDiff	P	0.35	0.64	0.21	0.57	0.50	63.86	0.70	5.73	**1.00**	0.85
	P→R	0.34	0.68	1.43	**0.54**	0.50	54.32	0.73	9.90	0.84	0.74
	R	0.34	0.68	1.32	0.56	0.52	53.78	0.71	9.48	0.83	0.73
DiffSBDD	P	0.32	0.65	0.67	0.63	0.62	63.90	0.65	60.70	0.99	**0.91**
	P→R	0.37	**0.75**	0.41	0.60	0.59	48.84	0.66	57.97	0.89	0.83
	R	0.37	0.71	0.27	0.63	0.64	45.13	0.65	57.10	0.91	0.84
Pocket2Mol	P	0.39	0.69	0.18	0.61	0.52	33.50	0.43	26.42	0.95	0.79
	P→R	0.41	0.64	0.32	0.63	0.53	29.08	0.72	16.74	0.95	0.78
	R	0.38	0.65	0.29	0.66	0.57	23.24	0.69	18.28	0.95	0.77
GraphBP	P	**0.49**	0.61	1.82	0.69	0.64	15.40	0.17	92.68	**1.00**	0.88
	P→R	0.48	0.67	0.72	0.63	0.59	26.52	0.25	79.97	**1.00**	0.86
	R	0.48	0.64	0.52	0.65	0.61	26.68	0.24	83.40	**1.00**	0.87
RiboSBDD	P→R	0.34	0.62	0.99	0.55	**0.48**	**70.12**	**0.76**	**5.00**	0.98	0.70
w/o Flexibility	P→R	0.35	0.62	0.87	0.55	0.49	65.27	0.74	5.80	0.98	0.71
w/o Sampling	P→R	0.30	0.59	1.24	0.55	0.50	60.31	0.78	9.20	0.98	0.67

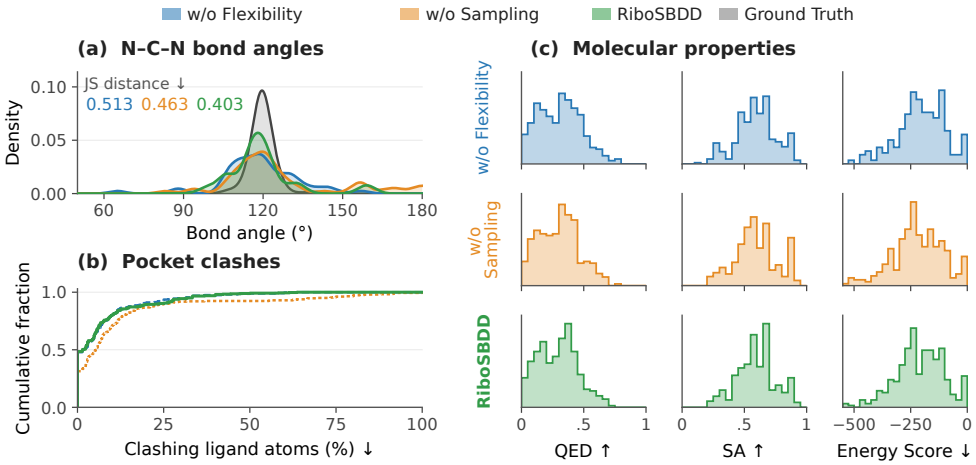


Figure 3: **Module distribution diagnostics.** (a) N–C–N angles (smoothed curves; JS distances from original histograms). (b) CDFs of per molecule clashing atom fractions. (c) QED, normalized SA, and AnnapuRNA Energy Score (↓). Panels (b,c) equally weight evaluable molecules; Table 1 pools atoms for Clash%. Details: Appendix D.2.

reference interval (Ghose et al., 1999). (2) *Geometry*: bond length and bond angle Jensen–Shannon *distances* (JSD_{BL}, JSD_{BA}; ↓), averaged over evaluated categories against reference ligands from CrossDocked2020 (Francoeur et al., 2020). (3) *Interaction*: the atom pooled receptor clash ratio (Clash%, ↓), mean recovery of reference contact residues within 4 Å over valid evaluated molecules pooled across pockets (Contact, ↑), and AnnapuRNA score improvement (IMP%, ↑), the percentage of successfully scored poses scoring below the reference ligand in the same pocket (Stefaniak & Bujnicki, 2021). (4) *Generation*: within-pocket mean pairwise fingerprint Tanimoto distance (Diversity, ↑), and the fraction of unique molecular identities not present in RNA training ligands (Novelty, ↑), using canonical isomeric SMILES. We evaluate all baselines, RiboSBDD, and both module controls on the same 713 test pockets, with 50 ligand generation requests per pocket. Reconstruction and metric definitions are provided in Appendix C.3.

4.2 Main Results

Protein pretraining supports geometric transfer. Compared with RNA-only training, protein pretraining followed by RNA fine tuning results in lower bond length and bond angle JS distances, and higher Contact across all four evaluated architectures (Table 1). For example, bond angle distance decreases from 0.64 to 0.59 for DiffSBDD, and from 0.57 to 0.53 for Pocket2Mol. This pattern across architectures supports initializing RNA generators using geometric representations learned from protein complexes. Additional geometry profiles appear in Appendix D.3.

RNA adaptation improves contact recovery, but not uniformly steric packing. Compared to protein transfer without RNA denoiser updates, RNA fine tuning increases Contact and decreases within-pocket Diversity across all four architectures. Novelty decreases for TargetDiff and DiffSBDD, but remains unchanged at the displayed precision for Pocket2Mol and GraphBP. For TargetDiff, Contact increases from 0.70 to 0.73, while Clash rises from 5.73% to 9.90%. Contact recovery measures overlap with reference interaction sites without penalizing steric penetration. Therefore, better site coverage does not necessarily imply improved atomic packing. This distinction motivates retaining coordinate guidance from the source in addition to RNA adaptation.

Chemical descriptors do not substitute for pocket compatibility. GraphBP has QED of 0.48–0.49 and Novelty of 1.00, but Clash of 79.97–92.68% and Contact of 0.17–0.25. Fine tuned DiffSBDD similarly combines the highest reported SA (0.75) with Clash of 57.97%. These comparisons support evaluating molecular properties and pocket-conditioned placement together, rather than inferring spatial compatibility from chemical descriptors alone.

RiboSBDD improves spatial compatibility with tradeoffs in chemical properties. Relative to TargetDiff after RNA fine tuning, RiboSBDD increases Contact from 0.73 to 0.76, reduces Clash from 9.90% to 5.00%, and increases IMP from 54.32% to 70.12% under AnnapuRNA. Bond angle distance improves from 0.50 to 0.48, whereas bond length distance changes from 0.54 to 0.55. SA decreases from 0.68 to 0.62, and within-pocket Diversity from 0.74 to 0.70, while QED remains unchanged at the displayed precision. Higher Novelty (0.98 vs. 0.84) with lower Diversity distinguishes generating unseen molecular identities from generating varied candidates within a pocket. The following comparisons examine the local correction and sampling modules.

4.3 Ablation Studies

We evaluate all three module configurations on the same 713 test pockets, with 50 generation requests per pocket (Table 1). *w/o Flexibility* uses a rigid continuation checkpoint and disables base correction, while retaining adaptive sampling. *w/o Sampling* retains the locally trained target, but removes the protein source, coordinate composition, and particle resampling. Both controls retain an RNA-adapted ligand generator, allowing us to examine the reported changes within the same generative framework. Figure 3 provides separate distributional diagnostics. Appendix D.2 details their scope and aggregation.

Local base correction improves interface compatibility. Compared with *w/o Flexibility*, the full model reduces Clash from 5.80% to 5.00%, increases Contact from 0.74 to 0.76, and increases IMP from 65.27% to 70.12%. Bond angle distance decreases from 0.49 to 0.48, SA remains at 0.62, and QED changes from 0.35 to 0.34. These results support coupling local correction with ligand denoising. Interface gains do not necessarily coincide with higher 2D property scores. The diagnostic distributions show a closer N–C–N angle profile and overlapping QED and SA distributions (Figure 3a,c). The control evaluates local geometry training and inference-time correction together.

Adaptive sampling improves packing beyond contact recovery. Compared with *w/o Sampling*, the full model reduces Clash from 9.20% to 5.00%, increases IMP from 60.31% to 70.12%, and decreases Contact from 0.78 to 0.76. Bond angle distance decreases from 0.50 to 0.48, while Diversity increases from 0.67 to 0.70. Thus, higher contact recovery alone does not necessarily yield better packing or pose scores. The diagnostic CDF also shows more candidates with low clash fractions (Figure 3b). This comparison supports retaining source guidance during sampling and jointly evaluates coordinate composition and particle resampling.

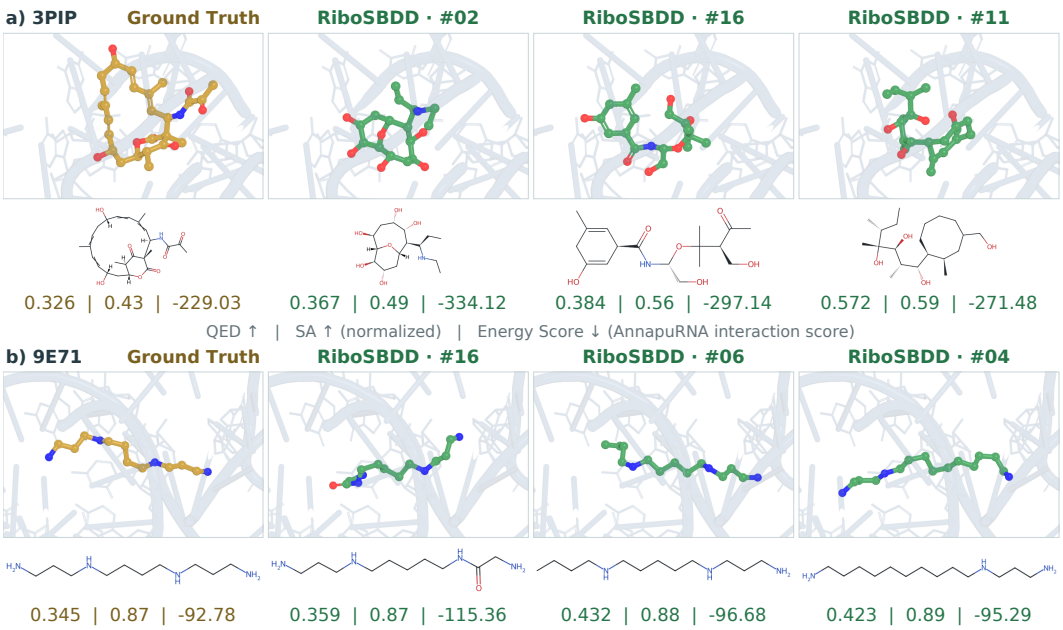


Figure 4: **Qualitative examples in (a) 3PIP and (b) 9E71.** Rows compare the deposited ligand and three RiboSBDD candidates in the same observed RNA pocket.

4.4 Qualitative Studies

Figure 4 shows compact, ring-containing candidates for 3PIP and extended, acyclic candidates for 9E71, each surrounding the corresponding deposited ligands. Their AnnapuRNA scores are lower, and their QED and normalized SA values are at least as high as those of the reference ligands. Displaying candidates alongside the deposited ligand in the same observed pocket allows score comparisons to be interpreted alongside their spatial placement. These selected cases illustrate shape variation, but do not indicate its prevalence across the test set. Appendix D.4 includes further examples with less favorable AnnapuRNA score comparisons.

5 Conclusion

We introduced RiboSBDD, a framework for RNA-conditioned *de novo* 3D ligand generation that integrates SE(3) equivariant diffusion, protein-to-RNA geometric transfer, fixed-backbone nucleobase correction, and adaptive coordinate sampling. RiboSBDDBench offers 8,006 curated pocket-ligand instances, grouped PDB/RNA90 splits, and standardized multi-metric evaluation. On 713 test pockets, protein pretraining consistently improves covalent geometry compared to RNA-only training across all four architectures. Module comparisons demonstrate that local base correction enhances contact recovery and reduces clashes, while adaptive sampling further reduces clashes and increases the AnnapuRNA score improvement rate. The complete framework improves pocket compatibility relative to fine-tuned TargetDiff. The benchmark also reveals that favorable 2D chemical properties do not guarantee compatible 3D placement. Overall, these results support the transfer of shared molecular geometry to RNA-targeted generation and provide a basis for evaluating candidate molecules.

Limitations and Future Work. Local flexibility is restricted to bounded glycosidic rotations with a fixed backbone, guided by synthetic perturbations rather than observed receptor dynamics. The adaptive sampler is empirical and does not ensure an exact terminal distribution. Spatial improvements over fine-tuned TargetDiff are accompanied by lower normalized SA and reduced within-pocket diversity. Evaluation depends on the available RNA ligand structures and computational metrics. Improved AnnapuRNA scores do not confirm experimental binding or biological activity. Future work includes modeling conformational ensembles, enhancing sampling efficiency and theoretical understanding, and balancing spatial compatibility with synthesis-aware and diversity objectives. Prospective synthesis and binding experiments are needed to further validate these candidates.

AI USE STATEMENT

We used generative AI tools for manuscript editing, literature checks, and assistance with code review and documentation. Generative AI also assisted the preparation of conceptual molecular illustrations. We take responsibility for the final manuscript, including its text, claims, code, and figures.

ETHICS STATEMENT

The experiments use publicly available biomolecular structures and computational models. Generated ligands are research candidates: geometric and chemical metrics do not establish experimental binding, therapeutic efficacy, or safety. The benchmark also reflects the coverage and selection biases of available RNA ligand structures.

REPRODUCIBILITY STATEMENT

Appendices A.2 and A.4 specify the sampling procedure and expanded algorithms. Appendix A.3 defines the local geometry objectives and continuation schedule, and Appendix B gives model configurations, optimization settings, atom count policies, and random seeds. Appendix C describes dataset curation, the grouped split, and metric definitions. We will open-source all datasets, benchmarks, and pretrained weights once accepted.

REFERENCES

Timothy EH Allen, Susan M Boyd, Maurinne Bonnet, and Rabia T Khan. Generative chemistry platform for small molecules targeting rna: A case study for chemical optimization. *Computational and Structural Biotechnology Journal*, 35(1):0218, 2026.

Stephen K Burley, Helen M Berman, Charmi Bhikadiya, Chunxiao Bi, Li Chen, Luigi Di Costanzo, Cole Christie, Ken Dalenberg, Jose M Duarte, Shuchismita Dutta, et al. Rcsb protein data bank: biological macromolecular structures enabling research and education in fundamental biology, biomedicine, biotechnology and energy. *Nucleic acids research*, 47(D1):D464–D474, 2019.

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*, 16:2799, 2025. doi: 10.1038/s41467-025-57852-0.

Paul G Francoeur, Tomohide Masuda, Jocelyn Sunseri, Andrew Jia, Richard B Iovanisci, Ian Snyder, and David R Koes. 3d convolutional neural networks and a crossdocked dataset for structure-based drug design. *Journal of chemical information and modeling*, 60(9):4200, 2020.

Arup K Ghose, Vellarkad N Viswanadhan, and John J Wendoloski. A knowledge-based approach in designing combinatorial or medicinal chemistry libraries for drug discovery. 1. a qualitative and quantitative characterization of known drug databases. *Journal of Combinatorial Chemistry*, 1(1): 55–68, 1999.

Jiaqi Guan, Wesley Wei Qian, Xingang Peng, Yufeng Su, Jian Peng, and Jianzhu Ma. 3d equivariant diffusion for target-aware molecule generation and affinity prediction. *arXiv preprint arXiv:2303.03543*, 2023.

Jiaqi Guan, Xiangxin Zhou, Yuwei Yang, Yu Bao, Jian Peng, Jianzhu Ma, Qiang Liu, Liang Wang, and Quanquan Gu. Decompdiff: diffusion models with decomposed priors for structure-based drug design. *arXiv preprint arXiv:2403.07902*, 2024.

Han Huang, Ziqian Lin, Dongchen He, Liang Hong, and Yu Li. RiboDiffusion: tertiary structure-based RNA inverse folding with generative diffusion models. *Bioinformatics*, 40(Supplement 1): i347–i356, 2024a. doi: 10.1093/bioinformatics/btae259.

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, 2024b. doi: 10.1093/bioinformatics/btae678.

Mary Chris Jaklevic. Oral drug approved for spinal muscular atrophy, 2020.

Chaitanya Joshi, Arian Jamasb, Ramon Viñas, Charles Harris, Simon Mathis, Alex Morehead, Rishabh Anand, and Pietro Liò. grnade: Geometric deep learning for 3d rna inverse design. In *International Conference on Learning Representations*, volume 2025, pp. 10166–10191, 2025.

Wissam Karroucha, Carlos Oliver, Veronique Stoven, and Vincent Mallet. Machine learning for rna-targeting drug design. *Journal of Chemical Information and Modeling*, 2026.

Zhengxin Li, Junjie Zhu, Xiaokun Hong, Junxi Mu, Zhuoqi Zheng, Taeyoung Cui, Yutong Sun, Ting Wei, and Hai-Feng Chen. DynaRNA: accurate dynamic RNA conformation ensemble generation with diffusion model. *Communications Biology*, 8:1472, 2025. doi: 10.1038/s42003-025-08875-2.

Haitao Lin, Guojiang Zhao, Odin Zhang, Yufei Huang, Lirong Wu, Cheng Tan, Zicheng Liu, Zhifeng Gao, and Stan Z Li. Cbgbench: fill in the blank of protein-molecule complex binding graph. In *International Conference on Learning Representations*, volume 2025, pp. 19861–19887, 2025.

Meng Liu, Youzhi Luo, Kanji Uchino, Koji Maruhashi, and Shuiwang Ji. Generating 3D molecules for target protein binding. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 13912–13924. PMLR, 2022. URL https://proceedings.mlr.press/v162/liu22m.html.

Yunfan Liu, Lirong Wu, Zhifeng Gao, Yufei Huang, Cheng Tan, Haitao Lin, Zicheng Liu, Changxi Chi, Chang Yu, and Stan Z. Li. SYNC: Measuring and advancing synthesizability in structure-based drug design. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=y1tPw4Uuzg.

Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. A 3d generative model for structure-based drug design. *Advances in neural information processing systems*, 34:6229–6239, 2021.

Runze Ma, Zhongyue Zhang, Zichen Wang, Chenqing Hua, Jiahua Rao, Zhuomin Zhou, and Shuangjia Zheng. Riboflow: conditional de novo rna co-design via synergistic flow matching. *Advances in Neural Information Processing Systems*, 38:69799–69831, 2025.

Francesco P Panei, Rachel Torchet, Herve Menager, Paraskevi Gkeka, and Massimiliano Bonomi. Hariboss: a curated database of rna-small molecules structures to aid rational drug design. *Bioinformatics*, 38(17):4185–4193, 2022.

Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2Mol: Efficient molecular sampling based on 3D protein pockets. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17644–17655. PMLR, 2022. URL https://proceedings.mlr.press/v162/peng22b.html.

Yanru Qu, Keyue Qiu, Yuxuan Song, Jingjing Gong, Jiawei Han, Mingyue Zheng, Hao Zhou, and Wei-Ying Ma. MolCRAFT: Structure-based drug design in continuous parameter space. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 41749–41768. PMLR, 2024. URL https://proceedings.mlr.press/v235/qu24a.html.

Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du, Carla Gomes, Tom L Blundell, Pietro Lio, et al. Structure-based drug design with equivariant diffusion models. *Nature Computational Science*, 4(12):899–909, 2024.

Alexander Serganov, Yu-Ren Yuan, Olga Pikovskaya, Anna Polonskaia, Lucy Malinina, Anh Tuân Phan, Claudia Hobartner, Ronald Micura, Ronald R Breaker, and Dinshaw J Patel. Structural basis for discriminative regulation of gene expression by adenine-and guanine-sensing mrnas. *Chemistry & biology*, 11(12):1729–1741, 2004.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correctors in diffusion: Annealing, guidance, and product of experts. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 55906–55949. PMLR, 2025a. URL https://proceedings.mlr.press/v267/skreta25a.html.

Marta Skreta, Lazar Atanackovic, Joey Bose, Alexander Tong, and Kirill Neklyudov. The superposition of diffusion models using the itô density estimator. In *International Conference on Learning Representations*, volume 2025, pp. 67004–67057, 2025b.

Filip Stefaniak and Janusz M. Bujnicki. AnnapuRNA: A scoring function for predicting RNA-small molecule binding poses. *PLOS Computational Biology*, 17(2):e1008309, 2021. doi: 10.1371/journal.pcbi.1008309.

Andrew C. Stelzer, Aaron T. Frank, Jeremy D. Kratz, Michael D. Swanson, Marta J. Gonzalez-Hernandez, Janghyun Lee, Ioan Andricioaei, David M. Markovitz, and Hashim M. Al-Hashimi. Discovery of selective bioactive small molecules by targeting an RNA dynamic ensemble. *Nature Chemical Biology*, 7(8):553–559, 2011. doi: 10.1038/nchembio.596.

Hong Su, Zhenling Peng, and Jianyi Yang. Recognition of small molecule–RNA binding sites using RNA sequence and structure. *Bioinformatics*, 37(1):36–42, 2021. doi: 10.1093/bioinformatics/btaa1092.

Sumit Tarafder and Debswapna Bhattacharya. RNAbpFlow: base pair-augmented SE(3) flow matching for conditional RNA 3d structure generation. *Nature Methods*, 23(7):1349–1358, 2026. doi: 10.1038/s41592-026-03128-4.

Yunpeng Xia, Jiayi Li, Yi-Ting Chu, Jiahua Rao, Jing Chen, Chenqing Hua, Dong-Jun Yu, Xiu-Cai Chen, and Shuangjia Zheng. Deciphering RNA–ligand binding specificity with GerNA-Bind. *Nature Machine Intelligence*, 7:1996–2008, 2025. doi: 10.1038/s42256-025-01154-z.

Julong Yang, Wen Huang, Junhui Chen, and Jian Peng. Decode: Decoupling binding position and molecular conformation in 3d ligand diffusion for structure-based drug design. In *Forty-third International Conference on Machine Learning*, 2026.

Zaixi Zhang, Yaosen Min, Shuxin Zheng, and Qi Liu. Molecule generation for target protein binding with structural motifs. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=Rq13idF0F73.

Zaixi Zhang, Mengdi Wang, and Qi Liu. FlexSBDD: Structure-based drug design with flexible protein modeling. In *Advances in Neural Information Processing Systems 37*, 2024. doi: 10.52202/079017-1707.

Xinzhe Zheng, Shiyu Jiang, Gustavo Seabra, Chenglong Li, and Yanjun Li. Apo2Mol: 3D molecule generation via dynamic pocket-aware diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 1614–1622, 2026. doi: 10.1609/aaai.v40i2.37138.

Jingyuan Zhou, Hao Qian, Shikui Tu, and Lei Xu. Prior-guided flow matching for target-aware molecule design with learnable atom number. *Advances in Neural Information Processing Systems*, 38:121513–121543, 2025a.

Xiangxin Zhou, Yi Xiao, Haowei Lin, Xinheng He, Jiaqi Guan, Yang Wang, Qiang Liu, Feng Zhou, Liang Wang, and Jianzhu Ma. Integrating protein dynamics into structure-based drug design via full-atom stochastic flows. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/2eea2b3ce7266e83785692cebbe1381d-Abstract-Conference.html.

A Theoretical Foundations and Algorithmic Details

A.1 Equal Density Coordinate Composition

We derive the equal density coordinate update used to combine the protein source and RNA target proposals (Skreta et al., 2025b). Let $x \in \mathbb{R}^{3N}$ denote the vectorized coordinates of an N-atom ligand. At a stochastic transition from t to the next selected reverse level t^-, the source and target proposals share covariance $\sigma_t^2 I$, with $\sigma_t^2 > 0$: $P_s(x) = \mathcal{N}(x; \mu_s, \sigma_t^2 I)$ and $P_r(x) = \mathcal{N}(x; \mu_r, \sigma_t^2 I)$. The transition coefficients are recomputed for the selected reverse levels specified in Appendix A.2. For a fixed coefficient κ, define $d_t = \mu_r - \mu_s$ and the affine composed mean $\mu_\kappa = \mu_s + \kappa d_t$. The weighted density product satisfies

$$\begin{aligned} P_r(x)^\kappa P_s(x)^{1-\kappa} &= G_t(\kappa)\, \mathcal{N}(x; \mu_\kappa, \sigma_t^2 I), \\ \log G_t(\kappa) &= -\frac{\kappa(1-\kappa)\|d_t\|^2}{2\sigma_t^2}. \end{aligned} \tag{12}$$

The scalar factor $G_t(\kappa)$ is the normalizing constant of this weighted density product and motivates the particle weighting potential used below.

Hyperplane projection. We choose κ for each noise draw by requiring the proposed update

$$X' = \mu_s + \kappa d_t + \sigma_t \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, I),$$

to satisfy $P_r(X') = P_s(X')$. Because the two proposals have the same covariance, equal density is equivalent to

$$\|X' - \mu_r\|^2 = \|X' - \mu_s\|^2.$$

For $d_t \neq 0$, substituting the proposed update and solving for κ yields

$$\kappa^*(\varepsilon) = \frac{1}{2} - \frac{\sigma_t \langle d_t, \varepsilon \rangle}{\|d_t\|^2}. \tag{13}$$

Let $m_t = (\mu_s + \mu_r)/2$ and let $\Pi_{d_t} = d_t d_t^\top / \|d_t\|^2$ denote the orthogonal projection onto the direction d_t. Substituting Eq. (13) into the update gives

$$X' = m_t + \sigma_t (I - \Pi_{d_t})\varepsilon, \qquad d_t^\top (X' - m_t) = 0. \tag{14}$$

Thus, the unclipped update lies on the perpendicular bisector hyperplane of μ_s and μ_r, retaining the Gaussian noise components orthogonal to d_t.

In implementation, we use $\kappa_t = \mathrm{clip}(\kappa^*(\varepsilon), -0.5, 1.5)$ in Eq. (10) to bound extrapolation. Clipping can break exact density equality. When $\|d_t\|^2 \leq 10^{-12}$, we treat the means as numerically coincident and use $\kappa_t = 1/2$. Because κ_t depends on the noise draw, the fixed coefficient identity in Eq. (12) does not directly specify the distribution of the adaptive update. We use $G_t(\kappa_t)$ as an empirical particle weighting potential. The zero variance terminal coordinate step follows the separate rule in Eq. (17).

A.2 Two Phase Reverse Sampling Dynamics

Algorithm 2 summarizes the sampling procedure. The full model uses 125 selected reverse levels from the original diffusion schedule with 1,000 levels and generates $B = 50$ particle trajectories per pocket. Coordinate updates combine the source and target predictions throughout sampling, while categorical atom type transitions use only the RNA target.

Reverse levels and timestep conventions. We select the original network indices as

$$\tau_j = \mathrm{rint}\left(\frac{999j}{124}\right), \qquad j = 0, \ldots, 124, \tag{15}$$

where rint rounds to the nearest integer, with half integer ties rounded to the nearest even integer. Sampling traverses these indices in descending order, from $j = 124$ to $j = 0$; consecutive indices differ by 8 or 9. The denoisers receive the original network index τ_j. The shortened loop index is

used only to retrieve transition coefficients recomputed for consecutive selected levels. The relation to the main text notation is

$$t = \tau + 1, \qquad T = 1000, \qquad \sqrt{t/T} = \sqrt{(\tau+1)/1000}. \tag{16}$$

Thus, the base rotation scaling in Eq. (3) uses the original diffusion timestep, not the index or length of the shortened loop. For $j > 0$, the reverse transition connects $t_j = \tau_j + 1$ to t_{j-1}; the last selected level $\tau_0 = 0$ corresponds to the terminal transition $t = 1 \to 0$. Each particle stores its most recent stochastic composition coefficient as κ_{last}.

Phase I: High noise base correction and particle resampling ($\tau \geq 500$). At each high noise level, the target model first predicts nucleobase rotations $\Delta\chi_t$ from the observed pocket $\mathcal{C}$ and the current noisy ligand. It then denoises the same ligand in a second forward pass conditioned on the corrected pocket $\widetilde{\mathcal{C}}_t$ (Eq. (3)). Each correction is applied independently to $\mathcal{C}$ at the current level, rather than accumulated across reverse steps. The source model supplies its coordinate proposal in a separate forward pass. Coordinates are updated using Eq. (10), and particle log weights accumulate $\log G_t(\kappa_t)$ from Eq. (12). Within each group of particles sharing the same atom count N, we perform systematic resampling when $\text{ESS}_N < B_N/2$, where B_N is the group size (Eq. (11)) (Skreta et al., 2025a). Coordinates, atom types, and the stored coefficient κ_{last} are copied together from the selected ancestor, followed by a reset of the group weights.

Phase II: Low noise ligand refinement ($0 \leq \tau < 500$). At low noise levels, the target model denoises the ligand in a single forward pass conditioned on the observed pocket $\mathcal{C}$. Base correction and particle resampling are disabled, while source and target coordinate composition remains active. For $\tau > 0$, coordinates follow the stochastic update in Eq. (10), and κ_{last} is updated to the current composition coefficient. The final level $\tau = 0$ uses the terminal rule below.

Terminal coordinate update ($\tau = 0$). For the final transition $t = 1 \to 0$, both models are evaluated again on the current ligand state to obtain clean coordinate predictions $\widehat{X}_0^s$ and $\widehat{X}_0^r$. We reuse the stored coefficient κ_{last} from the preceding stochastic level, which is $\tau = 8$ under Eq. (15), and set

$$X_0 = (1 - \kappa_{\text{last}})\widehat{X}_0^s + \kappa_{\text{last}}\widehat{X}_0^r. \tag{17}$$

No coordinate noise is added, and neither density ratios nor the Gaussian composition potential are recomputed. The terminal log weight increment is defined as

$$\Delta \log w_{\text{terminal}} = 0, \qquad G_{\text{terminal}} = 1. \tag{18}$$

This is a terminal step convention, not the $\sigma_t \to 0$ limit of Eq. (12). No terminal resampling is performed. Atom types are still sampled through the RNA target's categorical scheduler.

Denoiser evaluations. The selected schedule contains 63 high noise levels and 62 low noise levels, including the terminal level. Counting source and target forward passes separately, the full sampler uses 125 source evaluations and $125 + 63 = 188$ target evaluations, for a total of 313 denoiser evaluations over the reverse schedule.

A.3 Differentiable Fixed Backbone Base Correction

We model local pocket flexibility by rotating nucleobases about their glycosidic bonds while keeping the sugar phosphate backbone fixed. The operator $\mathcal{R}(\mathcal{C}, \Delta\chi)$ applies a rigid rotation to each eligible nucleobase about the C1′–N9 axis for purines or the C1′–N1 axis for pyrimidines. The corresponding torsion χ is defined by O4′–C1′–N9–C4 and O4′–C1′–N1–C2, respectively. We implement $\mathcal{R}$ using Rodrigues' rotation formula, which is differentiable with respect to $\Delta\chi$. The ligand denoising loss from the second forward pass therefore propagates through the corrected pocket coordinates to the base correction head. Training combines ligand denoising with angular, contact, and steric supervision:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pos}} + \lambda_v\mathcal{L}_{\text{type}} + \lambda_\chi\mathcal{L}_\chi + \lambda_c\mathcal{L}_{\text{contact}} + \lambda_s\mathcal{L}_{\text{steric}}. \tag{19}$$

The loss terms are defined in Eqs. (2), (5), and (6).

Starting from the RNA fine tuned checkpoint, we perform three continuation stages of 2,000 updates each (Table 2). Stage 1 uses angular supervision and a lower steric loss weight, with the contact loss disabled. Stage 2 increases the angular loss weight and enables the contact loss. Stage 3 increases the steric loss weight while retaining the other weights.

Table 2: **Loss weights for local geometry continuation.** Each stage uses 2,000 updates; the coordinate loss weight is fixed at one.

Stage	λ_{pos}	λ_v	λ_χ	λ_s	λ_c
Stage 1 (0–2k updates)	1.0	100.0	1.0	0.2	0.0
Stage 2 (2–4k updates)	1.0	100.0	10.0	0.2	1.0
Stage 3 (4–6k updates)	1.0	100.0	10.0	1.0	1.0

A.4 Expanded Training and Sampling Algorithms

Algorithm 1 Protein to RNA training with local base correction

Require: Protein ligand data $\mathcal{D}_s$; RNA ligand data $\mathcal{D}_r$; local rotation operator $\mathcal{R}$
Ensure: Protein source model F_{θ_s} and RNA target model F_{θ_r}
1: Train F_{θ_s} on $\mathcal{D}_s$; copy its weights to initialize F_{θ_r} and fine tune the target on $\mathcal{D}_r$.
2: Perturb local base orientations and the ligand in each RNA complex.
3: Predict a base correction, then denoise the ligand under the corrected pocket (Eq. (3)).
4: Update the target with Eq. (7); retain both frozen models for sampling.

Algorithm 2 Adaptive source and target sampling for an RNA pocket

Require: Observed RNA pocket $\mathcal{C}$; frozen source F_{θ_s} and target F_{θ_r}
Ensure: Candidate ligands $\{(X_0^b, V_0^b)\}_{b=1}^B$
1: **Initialize:** sample atom counts and $B = 50$ noisy ligand particles.
2: **for** each reverse level, from high to low noise **do**
3: **Predict:** obtain source and target transitions; apply local base correction to the target at high noise.
4: **Update:** use Eq. (10) at stochastic levels and Eq. (17) at the terminal level; sample atom types only from the RNA target.
5: **Resample:** at high noise, weight and select complete ligand states using Eq. (11).
6: **Decode:** reconstruct bonds and return the generated ligands.

Algorithms 1 and 2 emphasize the conceptual design. Algorithms 3 and 4 record the full schedule, indexing, and edge cases used by the final method. Each of the 50 particles corresponds to one generation request; molecular reconstruction can fail for individual outputs.

B Implementation Details

B.1 Model Configurations

We use the repository implementations of TargetDiff, DiffSBDD, Pocket2Mol, and GraphBP, keeping each architecture fixed across protein only (P), RNA training from scratch (R), and protein pretraining followed by RNA fine tuning (P→R). The TargetDiff and DiffSBDD implementations use UniTransformer backbones with nine layers, hidden dimension 128, and 16 attention heads. Both use 1,000 diffusion levels. TargetDiff uses a sigmoid coordinate schedule with endpoints 10^{-7} and 2×10^{-3}, a cosine atom type schedule with offset 0.01, and symmetric timestep sampling. DiffSBDD uses quadratic polynomial schedules and random timestep sampling. Pocket2Mol uses a GVP Transformer with six layers, 256 scalar channels, and 64 vector channels. GraphBP uses an encoder and a context embedder, each with six layers and hidden dimension 256.

RiboSBDD inherits the TargetDiff backbone and diffusion schedules. Its base correction head combines mean pooled nucleobase features with five time and angular features, followed by a $133 \to 128 \to 64 \to 1$ MLP with SiLU hidden activations and a $\tanh$ output. The output layer is initialized to zero, so the model predicts zero rotation at initialization.

B.2 Protein Pretraining

We train each source model from random initialization on CrossDocked2020 (Francoeur et al., 2020) using Adam. We select source checkpoints by the lowest loss on a fixed holdout of 100 protein

Algorithm 3 Expanded protein to RNA training procedure

Require: Protein ligand data $\mathcal{D}_s$; RNA training complexes $\mathcal{D}_r$; local rotation operator $\mathcal{R}$
Ensure: Retained protein source F_{θ_s}, RNA target F_{θ_r}, and independently trained count model p_ψ
Pretrain on $\mathcal{D}_s$; select F_{θ_s} by protein validation loss.
Initialize $F_{\theta_r} \leftarrow F_{\theta_s}$; fine tune it on $\mathcal{D}_r$; retain the source.
for continuation stage $j = 1, 2, 3$ **do**
    Keep the preceding target weights; reset the optimizer and scheduler; set stage-j loss weights (Appendix A.3).
    for each update in stage j **do**
        Sample an RNA complex and state t; corrupt its ligand and perturb base orientations to obtain $(X_t, V_t, \mathcal{C}_t)$, where the network level is $\tau = t - 1$.
        Run the target once and form $\widetilde{\mathcal{C}}_t$ from its predicted correction (Eq. (3)).
        Denoise the same ligand with the corrected pocket; update the target through both passes using the stage objective.
Fit p_ψ separately on RNA training pocket features and ligand counts.

Algorithm 4 Expanded RNA conditioned sampling procedure

Require: Observed RNA pocket $\mathcal{C}$; frozen source F_{θ_s}, target F_{θ_r}, and count model p_ψ; $B = 50$ particles
Ensure: Generated ligands $\{(X_0^b, V_0^b)\}_{b=1}^B$
for particle $b = 1, \ldots, B$ **do**
    Sample and fix $N_b \sim p_\psi(\cdot \mid \mathcal{C})$; initialize (X_T^b, V_T^b) and $\log w_b \leftarrow 0$.
Partition particles into equal count strata $\mathcal{G}_N = \{b : N_b = N\}$; write $B_N = |\mathcal{G}_N|$.
for each selected network level in Eq. (15), in descending order **do**
    Evaluate the source on the observed pocket $\mathcal{C}$.
    if $\tau \geq 500$ **then**
        Run the target twice: predict a local correction, then predict the ligand from the corrected current step pocket.
    else
        Run the target once on $\mathcal{C}$.
    Form source and target coordinate means μ_s, μ_r and the target only categorical reverse transition.
    if the coordinate transition is stochastic ($\sigma_t^2 > 0$) **then**
        Draw ε_b and compute $\kappa_{t,b}$ with Eq. (13), including clipping and the degenerate fallback.
        Sample coordinates with Eq. (10); sample atom types only from the target transition; save $\kappa_{t,b}$.
        Accumulate the coordinate log weight with Eq. (12).
        if $\tau \geq 500$ **then**
            for each stratum $\mathcal{G}_N$ with $\mathrm{ESS}_N < B_N/2$ **do**
                Systematically resample complete states and all particle indexed auxiliaries by ancestor; reset stratum weights.
    else
        Use Eq. (17) with the coefficient inherited from the selected ancestor; set the terminal log weight increment to zero and apply the target categorical update.
Return unweighted atom states for bond reconstruction; do not force terminal resampling.

pockets. Table 3 lists the selected checkpoints and batch sizes; optimizer settings appear in Table 4. RiboSBDD uses the selected TargetDiff checkpoint for both RNA initialization and the retained protein source. All reported training runs use random seed 2022.

Table 3: Selected protein checkpoints and RNA training budgets. Protein iterations identify selected checkpoints, not total source training budgets. The RNA budget applies separately to R and P→R.

Model	Protein batch	Selected iteration	RNA batch	RNA updates
TargetDiff	4	217,000	16	75,000
DiffSBDD	8	194,000	16	75,000
Pocket2Mol	8	88,000	16	75,000
GraphBP	16	56,000	16	75,000

Table 4: Optimizer settings for protein pretraining and baseline RNA training. LR denotes learning rate; clipping limits the gradient norm. Plateau schedulers use factor 0.6, with patience measured in monitoring evaluations.

Model	Initial LR	Adam betas	Decay	Clipping	Patience	Min. LR
TargetDiff	5×10^{-4}	$(0.95, 0.999)$	0	8	10	10^{-6}
DiffSBDD	5×10^{-4}	$(0.95, 0.999)$	0	8	10	10^{-6}
Pocket2Mol	2×10^{-4}	$(0.99, 0.999)$	0	100	8	10^{-5}
GraphBP	10^{-5}	$(0.99, 0.999)$	10^{-6}	100	100	10^{-5}

B.3 RNA Training

RNA training from scratch uses random model initialization, whereas fine tuning transfers the selected protein weights and initializes a new optimizer and scheduler. Both regimes update all model parameters on the same 7,269 RNA training complexes for 75,000 updates at batch size 16. The position and atom type loss weights are $(1, 100)$ for TargetDiff and $(1, 1)$ for DiffSBDD; Pocket2Mol and GraphBP use unit weights for their component losses. A fixed subset of 713 training complexes, disjoint from the test set, supplies loss measurements every 1,000 updates for the learning rate scheduler. Both regimes use the final checkpoint at 75,000 updates.

B.4 Local Geometry Training

RiboSBDD continues from the final TargetDiff checkpoint after RNA fine tuning through three stages of 2,000 updates each, using batch size 8. We train the target backbone and correction head jointly, while keeping the protein source frozen. Loss weights follow Table 2, and learning rates are listed in Table 5. We use Adam with betas $(0.95, 0.999)$, zero weight decay, gradient clipping at 8, and a plateau scheduler with factor 0.6, patience 10, and minimum learning rate 10^{-6}. Each stage transfers the preceding model weights and starts a new optimizer and scheduler; the final stage supplies the reported model.

Table 5: Learning rates during local geometry training. Each stage uses 2,000 updates; loss weights are given in Table 2.

Stage	Backbone learning rate	Correction head learning rate
Stage 1	10^{-5}	5×10^{-4}
Stage 2	10^{-5}	5×10^{-4}
Stage 3	10^{-5}	5×10^{-3}

Local augmentation draws bounded Gaussian angular perturbations with nominal standard deviation 10° and maximum magnitude 15°, scaled by $\sqrt{t/T}$. With probability 0.25, the observed pocket is left unperturbed. Ligand reconstruction is differentiated through the corrected pocket coordinates, coupling base correction to ligand denoising. The contact loss uses a 6 Å support cutoff and a SmoothL1 transition scale of 1 Å. The steric loss uses a 2.2 Å threshold and is active at $\tau \leq 200$, equivalent to $t \leq 201$. Training permits correction across sampled timesteps; the restriction to $\tau \geq 500$ applies only during generation.

Ligand atom count. An independent $11 \to 128 \to 256 \to 128 \to 1$ SiLU MLP predicts ligand size from invariant descriptors of pocket size and shape. The 11 inputs comprise atom count, radius of gyration, three covariance eigenvalue scales, the median of the ten largest interatomic distances, and five radial quantiles at 10, 25, 50, 75, and 90 percent. We fit it on the RNA training split for 200 epochs using Adam at 10^{-3} and a batch size of 256, and freeze its final checkpoint during generation. Atom counts are sampled by adding Gaussian noise with standard deviation 0.01 in normalized prediction units, followed by inverse normalization, rounding, and clipping to the range observed in RNA training. Native test ligand counts are not used.

B.5 Generation Protocol

We evaluate all configurations on the same 713 test pockets, with 50 generation requests per pocket. This budget counts requests, not successful molecular reconstructions. Baseline TargetDiff and DiffSBDD use their native samplers with 1,000 diffusion steps and an atom count histogram derived from the RNA training split. For these baselines, P denotes no RNA updates to the denoiser, while still using RNA statistics for ligand size. Pocket2Mol and GraphBP use their native autoregressive generation and stopping rules with batch size 1; GraphBP uses the `cbgbench_uniform` contact and focus policy.

RiboSBDD generates a batch of $B = 50$ particle trajectories per pocket using the 125 reverse levels defined in Appendix A.2. Each trajectory samples its atom count before diffusion and retains it throughout sampling. Stochastic coordinate updates, categorical atom type transitions, and particle resampling generate the candidate ensemble. Coordinates combine source and target proposals, whereas atom types use only the RNA target. At $\tau \geq 500$, the target applies local base correction and systematic resampling is triggered when the effective sample size falls below half the size of a group with equal atom count. The source always uses the observed pocket, and each target correction starts from that same pocket. Appendix A.2 specifies the composition coefficients and the terminal coordinate update; reconstruction and metric aggregation follow Appendix C.3.

B.6 Module Ablations

Without flexibility. Starting from the same checkpoint after 75,000 RNA updates, this control performs 6,000 continuation updates with matched batch size, backbone optimizer settings, and stage restarts. It uses only the position and atom type losses with weights 1 and 100, omitting the correction head and the angular, contact, and steric losses. Generation disables base correction but retains the learned atom count predictor, protein source, coordinate composition, and particle resampling. This comparison evaluates local geometry training and correction during inference together.

Without sampling. This control reuses the full model's final target checkpoint and atom count predictor. It removes the protein source, coordinate composition, and particle resampling, while retaining the 125 reverse levels and local correction at $\tau \geq 500$.

The full model and both controls use 81,000 nominal RNA denoiser updates, but require different numbers of denoiser evaluations during generation. Per particle batch, RiboSBDD uses 313 evaluations (125 source, 125 target, and 63 additional correction passes), compared with 250 without flexibility and 188 without sampling. Sampling budgets are specific to each configuration rather than matched by network evaluations or runtime.

C Dataset Construction and Evaluation Details

C.1 Structural Curation

Source and instance definition. We collected 7,453 RNA containing experimental structures from the PDB and enumerated 3,670,853 candidate component instances. Each instance is identified by its PDB entry, model, chain, and residue position. Structural and chemical curation removes polymers, water, ions, common crystallization additives, carbon free components, and unparseable molecules.

RNA association and pocket extraction. We retain a ligand instance if any ligand heavy atom lies within $5\,\text{Å}$ of an RNA heavy atom. Its pocket contains all deposited atoms of RNA residues with at least one heavy atom within $10\,\text{Å}$ of the ligand. Explicit hydrogens are retained when available, occurring in 1,221 of the 8,006 model ready pockets. We additionally store neighboring nonligand residues, including proteins, as auxiliary structural context for dataset analysis.

Broad and Clean views. The Broad set contains 8,416 curated instances. The Clean subset further limits protein atoms in the local ligand environment. Let n_P and n_R denote the numbers of distinct protein and RNA heavy atoms within $5\,\text{Å}$ of any ligand heavy atom. We retain instances satisfying $n_P/(n_P + n_R) \leq 2/3$, yielding 7,469 Clean instances before model conversion.

Model representability. Conversion checks the ligand atom and bond vocabulary, RNA residue mapping, and receptor element support. The union of these checks excludes 410 samples and produces 8,006 model ready Broad instances and their Clean subset of 7,113 instances, contributed by 2,258 PDB entries. Table 6 reports the individual checks and their overlapping exclusion counts.

Table 6: Model representability exclusions. Individual checks may overlap; the union counts each excluded instance once. Percentages are relative to the 8,416 curated Broad instances.

Conversion check	Samples	Raw Broad (%)
Ligand atom/bond vocabulary	80	0.95%
RNA residue mapping	320	3.80%
Receptor element support	30	0.36%
Unique excluded samples (union)	410	4.87%

C.2 Grouped Structural Split

We connect instances that share a PDB identifier or an RNA sequence cluster at 90% identity (an RNA90 cluster). Each instance is associated with every RNA90 cluster represented by its receptor chains. Taking the transitive closure produces 525 connected components (split groups) spanning 857 RNA90 clusters. Whole groups are assigned together, preventing a shared PDB identifier or RNA90 cluster across training and test sets.

The final partition contains 7,269 Broad training instances and 713 Clean test instances. The remaining 24 instances are non Clean members of test components and are excluded from both training and evaluation.

Table 7: Dataset construction and partition counts. Clean is a subset of Broad.

Quantity	Count
Source PDB entries	7,453
Curated Broad / Clean instances	8,416 / 7,469
Model ready Broad / Clean instances	8,006 / 7,113
Contributing PDB entries	2,258
RNA90 clusters / connected components	857 / 525
Broad training / Clean test instances	7,269 / 713
Non Clean members of test components, excluded	24

C.3 Evaluation Protocol and Metrics

Generation and evaluation protocol. We evaluate all twelve baseline settings, RiboSBDD, and both module controls on the same 713 RNA test pockets, with 50 ligand generation requests per pocket. We convert generated atom types and coordinates into molecular structures through bond reconstruction, using Open Babel bond recovery when the initial reconstruction raises an exception. Reconstructed molecules undergo RDKit sanitization. Following the CBGBench fragment selection protocol (Lin et al., 2025), we retain the largest connected fragment if it contains at least $\lfloor 0.8N \rfloor$ atoms, where N is the atom count of the reconstructed molecule before fragment selection. We save accepted structures as SDF files and parse, sanitize, and strip them of explicit hydrogens for evaluation.

We record the numbers of generated, reconstructed, sanitized, accepted, and chemistry evaluated outputs. Accepted fragment validity and evaluation coverage are the fractions of the 50 generation requests per pocket yielding an accepted fragment and a chemistry evaluated molecule, respectively. Chemical property summaries use chemistry evaluated molecules; geometry and interaction metrics additionally require the corresponding coordinate evaluations. IMP uses poses with valid generated and reference AnnapuRNA scores, as defined below. The main table shares the scope of 713 pockets and the generation budget across all configurations; the number of evaluable outputs can differ by metric.

Chemical properties. For each chemistry evaluated molecule m, we compute QED, synthetic accessibility (SA), and logP using RDKit. QED summarizes molecular properties associated with drug likeness on a scale from 0 to 1. SA estimates synthetic accessibility. Following the convention used in CBGBench (Lin et al., 2025), we transform the raw synthetic accessibility score as

$$\mathrm{SA}(m) = \frac{10 - \mathrm{SA}_{\mathrm{raw}}(m)}{9}, \tag{20}$$

so that higher values indicate greater estimated synthetic accessibility. Unless otherwise specified, chemical property summaries report the mean over chemistry evaluated molecules.

Bond length and bond angle distributions. We assess local molecular geometry by comparing generated bond length and bond angle distributions against empirical reference distributions from ligands in CrossDocked2020 (Francoeur et al., 2020). For each bond or angle category, we construct normalized histograms for generated and reference samples using identical bins. Let P and Q denote the resulting distributions and $M = (P+Q)/2$. Following the evaluator implementation, we compute the Jensen–Shannon distance:

$$d_{\mathrm{JS}}(P, Q) = \sqrt{\frac{1}{2} D_{\mathrm{KL}}(P \| M) + \frac{1}{2} D_{\mathrm{KL}}(Q \| M)}. \tag{21}$$

We report the category averaged distances as $\mathrm{JSD}_{\mathrm{BL}}$ and $\mathrm{JSD}_{\mathrm{BA}}$, where lower values indicate closer agreement with the empirical distributions of real ligands.

Receptor clashes. Receptor clashes are evaluated against the original observed RNA pocket stored in the frozen `pocket.pdb`. The RNA parser reads `ATOM` records for A, U, G, and C residues in the first model. We do not use transiently corrected receptor coordinates or protein atoms stored in the auxiliary context. Existing explicit receptor hydrogens are retained, but no receptor hydrogens are added; ligand hydrogens are removed before evaluation. We quantify steric clashes between ligand heavy atoms and the retained receptor atoms using their van der Waals radii. A ligand atom i is classified as clashing if there exists a receptor atom j such that

$$\|x_i - r_j\| < R_i^{\mathrm{vdW}} + R_j^{\mathrm{vdW}} - 0.4\,\text{Å}, \tag{22}$$

where x_i and r_j are the corresponding coordinates, and R^{vdW} denotes the van der Waals radius; receptor hydrogen uses a radius of 1.20 Å. Each ligand atom is counted once if it clashes with at least one receptor atom.

Let $\mathcal{G}_{\mathrm{geom}}$ denote the set of geometry evaluated molecules, $N(m)$ the number of ligand atoms in molecule m, and $N_{\mathrm{clash}}(m)$ the number of ligand atoms classified as clashing with the receptor. The atom level clash ratio is

$$\mathrm{Clash}_{\mathrm{atom}} = \frac{\sum_{m \in \mathcal{G}_{\mathrm{geom}}} N_{\mathrm{clash}}(m)}{\sum_{m \in \mathcal{G}_{\mathrm{geom}}} N(m)}. \tag{23}$$

The Clash% column in Table 1 reports $100\,\mathrm{Clash}_{\mathrm{atom}}$.

The CDF in Figure 3 instead uses the fraction $N_{\mathrm{clash}}(m)/N(m)$ for each evaluated molecule, giving molecules equal weight. This is distinct from both the pooled atom statistic in Table 1 and a binary indicator of whether a molecule has any clash. Lower values indicate fewer steric overlaps.

Contact recovery. We measure recovery of the receptor residues contacted by the deposited reference ligand. For pocket p, let $\mathcal{R}_p^{\mathrm{ref}}$ and $\mathcal{R}_p(m)$ denote the contact residue sets of the reference ligand and generated molecule m. A residue belongs to a contact set if one of its heavy atoms lies within 4 Å of a ligand heavy atom. The contact recall of a generated molecule is

$$c_p(m) = \frac{|\mathcal{R}_p(m) \cap \mathcal{R}_p^{\mathrm{ref}}|}{|\mathcal{R}_p^{\mathrm{ref}}|}. \tag{24}$$

Let $\mathcal{P}_{\mathrm{contact}}$ contain test pockets with valid, nonempty reference contact sets, and let $\mathcal{G}_p^{\mathrm{contact}}$ contain accepted generated molecules with valid contact evaluation for pocket p. We pool these molecules across pockets and report

$$\mathrm{Contact} = \frac{\sum_{p \in \mathcal{P}_{\mathrm{contact}}} \sum_{m \in \mathcal{G}_p^{\mathrm{contact}}} c_p(m)}{\sum_{p \in \mathcal{P}_{\mathrm{contact}}} |\mathcal{G}_p^{\mathrm{contact}}|}. \tag{25}$$

Each evaluated molecule receives equal weight, so pockets contribute in proportion to their numbers of valid molecules rather than equally. Failed generation or contact evaluation is excluded rather than assigned zero; valid molecules recovering no reference contacts remain in the average with $c_p(m) = 0$. Pockets without valid evaluated molecules contribute to neither the numerator nor the denominator. Higher Contact indicates greater recovery of reference contact residues but does not penalize steric overlaps.

AnnapuRNA score improvement. Let $s_p(m)$ and s_p^{ref} denote the AnnapuRNA scores of a generated pose and the deposited reference ligand in pocket p. Let $\mathcal{P}_{\mathrm{score}}$ contain pockets with a finite reference score and at least one successfully scored generated pose, and let $\mathcal{G}_p^{\mathrm{score}}$ contain those generated poses. We report

$$\mathrm{IMP}\% = 100\,\frac{\sum_{p\in\mathcal{P}_{\mathrm{score}}}\sum_{m\in\mathcal{G}_p^{\mathrm{score}}}\mathbf{1}\big[s_p(m) < s_p^{\mathrm{ref}}\big]}{\sum_{p\in\mathcal{P}_{\mathrm{score}}}|\mathcal{G}_p^{\mathrm{score}}|}. \tag{26}$$

Each successfully scored pose receives equal weight. Higher values indicate a larger fraction of generated poses scoring below the reference ligand in the same pocket. AnnapuRNA scores are interaction scores rather than physical binding energies.

Diversity. We measure chemical diversity within each pocket using RDKit fingerprints. For each pair of valid generated molecules m_i and m_j, we compute the Tanimoto distance between their RDKFingerprints,

$$d_{\mathrm{Tan}}(m_i, m_j) = 1 - \mathrm{Tanimoto}\big(\mathrm{FP}(m_i), \mathrm{FP}(m_j)\big). \tag{27}$$

For a pocket p with $n_p \geq 2$ valid molecules, diversity is the mean over all unordered molecular pairs,

$$\mathrm{Div}(p) = \frac{2}{n_p(n_p-1)}\sum_{1\leq i<j\leq n_p} d_{\mathrm{Tan}}(m_i, m_j). \tag{28}$$

Pockets with fewer than two valid molecules have undefined diversity. The reported dataset level diversity is the macro average over pockets with defined diversity.

Novelty. We assess novelty against RNA training ligands using canonical isomeric SMILES. Explicit hydrogens are removed before canonicalization, without additional neutralization or tautomer standardization. Let $\mathcal{U}_p$ denote the set of unique valid molecules generated for pocket p, and let $\mathcal{T}$ denote the set of canonicalized RNA training ligands. Pocket level novelty is

$$\mathrm{Novelty}(p) = \frac{|\{m \in \mathcal{U}_p : \mathrm{SMILES}(m) \notin \mathcal{T}\}|}{|\mathcal{U}_p|}. \tag{29}$$

The reported novelty is the macro average over pockets with at least one unique valid generated molecule.

D Additional Results

D.1 Aggregate Benchmark Comparisons

The radar plots and heatmap reproduce the values displayed in Table 1. For display, bond length and bond angle distances are transformed as $1 - d_{\mathrm{JS}}/\sqrt{\ln 2}$, IMP is divided by 100, and the clash axis uses one minus the clashing atom fraction. The logP axis maps $[-3, 3]$ linearly to $[0, 1]$ and is descriptive rather than a quality ordering. Radar area is not an aggregate performance score.

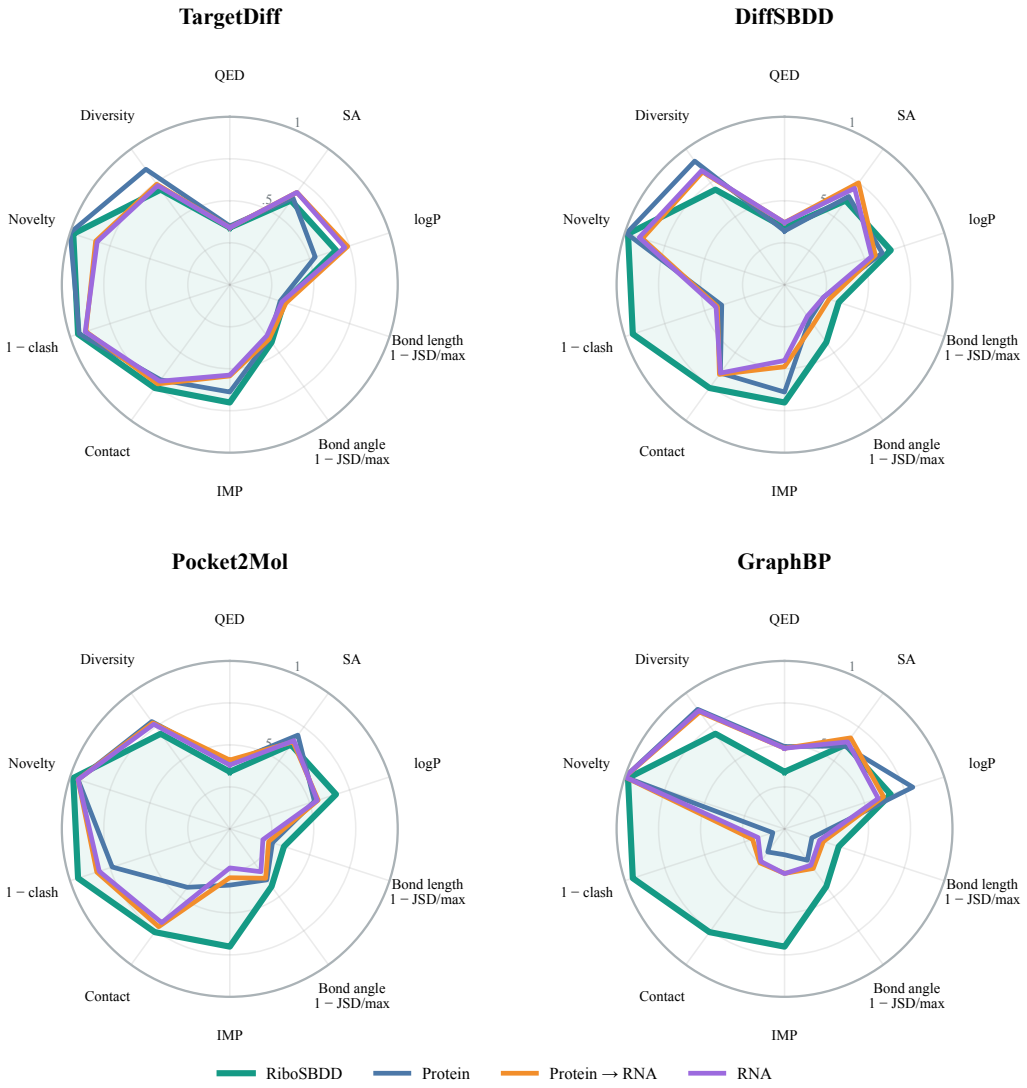


Figure 5: **RiboSBDD and twelve baseline configurations.** Each architecture is shown under P, P→R, and R using the corresponding values from Table 1.

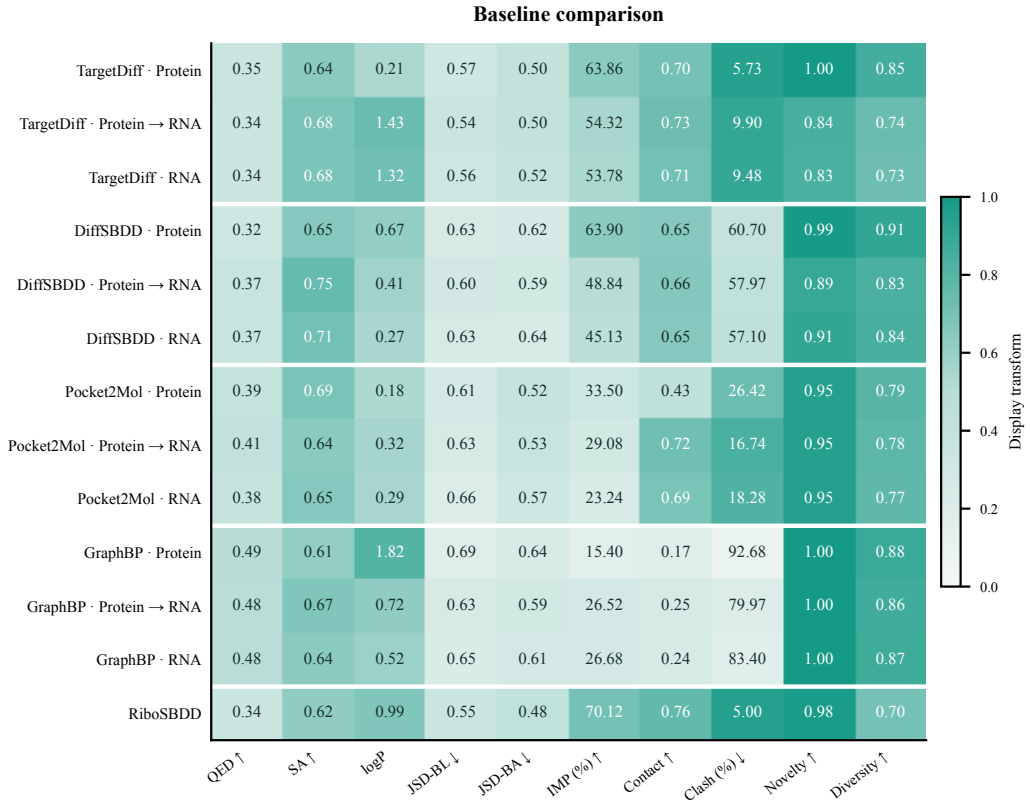


Figure 6: **Benchmark values.** Cell annotations reproduce Table 1; colors use the display transforms defined above.

D.2 Module Comparisons

Table 1 compares the full model and both module controls on 713 test pockets with 50 generation requests per pocket. Figure 7 visualizes these table entries.

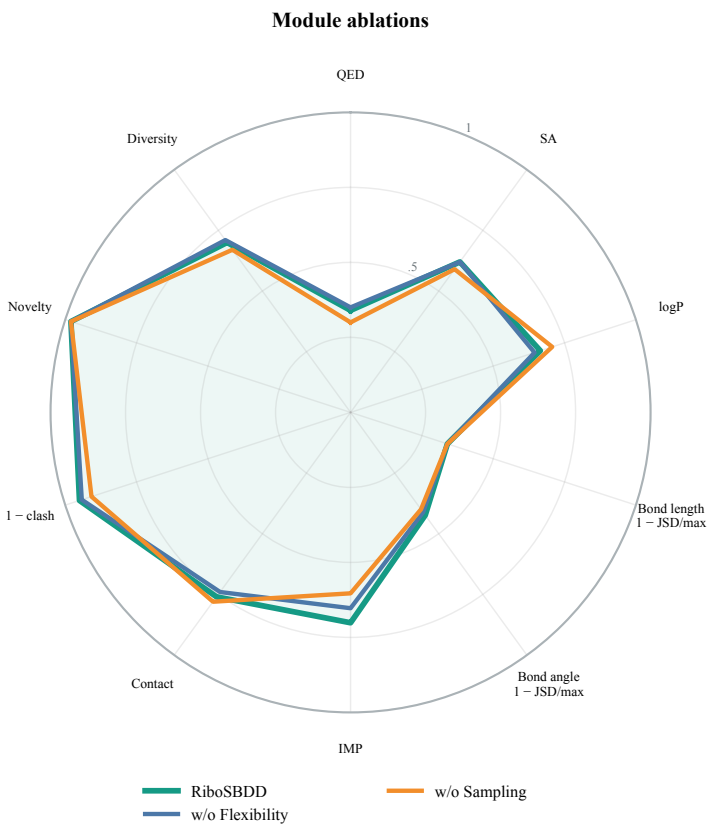


Figure 7: **Full model and module controls.** Values reproduce the three module rows of Table 1, with the same display transforms as the baseline radar plots.

Figure 3 and the module geometry profiles use the existing diagnostic records for 20 shared pockets, rather than the full set used for Table 1. The clash CDFs and property distributions give each evaluable molecule equal weight; the main table pools ligand atoms for Clash%. For Figure 3(a), angle profiles use bins of 2° and a common Gaussian smoothing bandwidth of 3°, with reflection at 0° and 180°. JS distances are calculated from the original complete histograms, not the smoothed curves. Energy Score denotes the AnnapuRNA score, with lower values preferred.

D.3 Geometric Distributions

Ground Truth in the geometry plots denotes the empirical reference distribution from CrossDocked2020 ligands, as defined in Appendix C.3. The module plots supplement the shared diagnostic comparison; the architecture plots compare RiboSBDD with the four baselines after protein pretraining and RNA fine tuning.

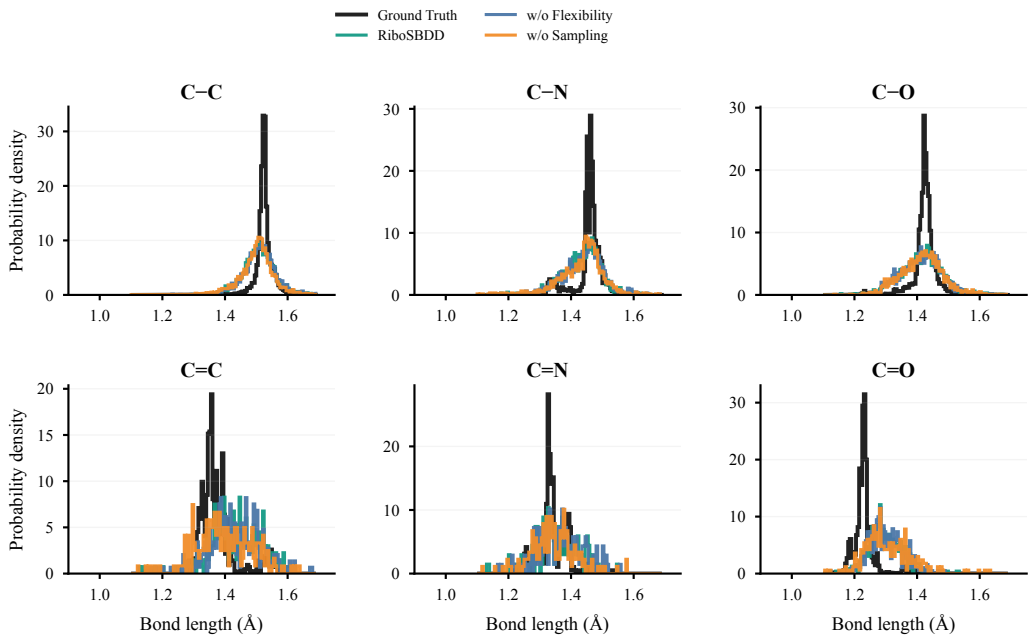


Figure 8: **Bond length profiles for the module comparisons.** Fixed interior bins retain their original normalization; JS distances include all tail mass.

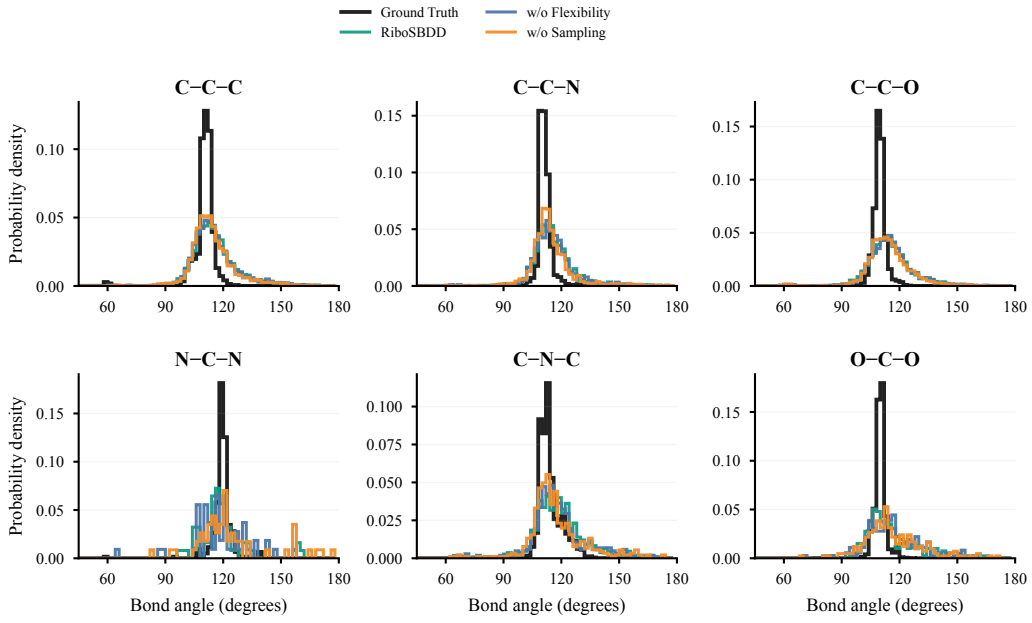


Figure 9: **Bond angle profiles for the module comparisons.**

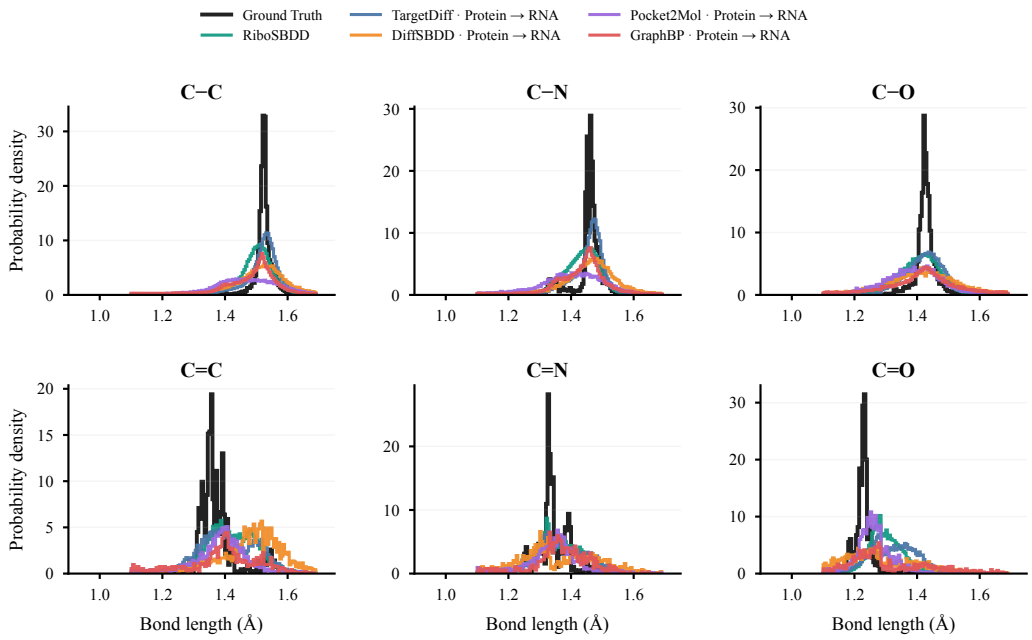


Figure 10: **Bond length profiles across architectures.**

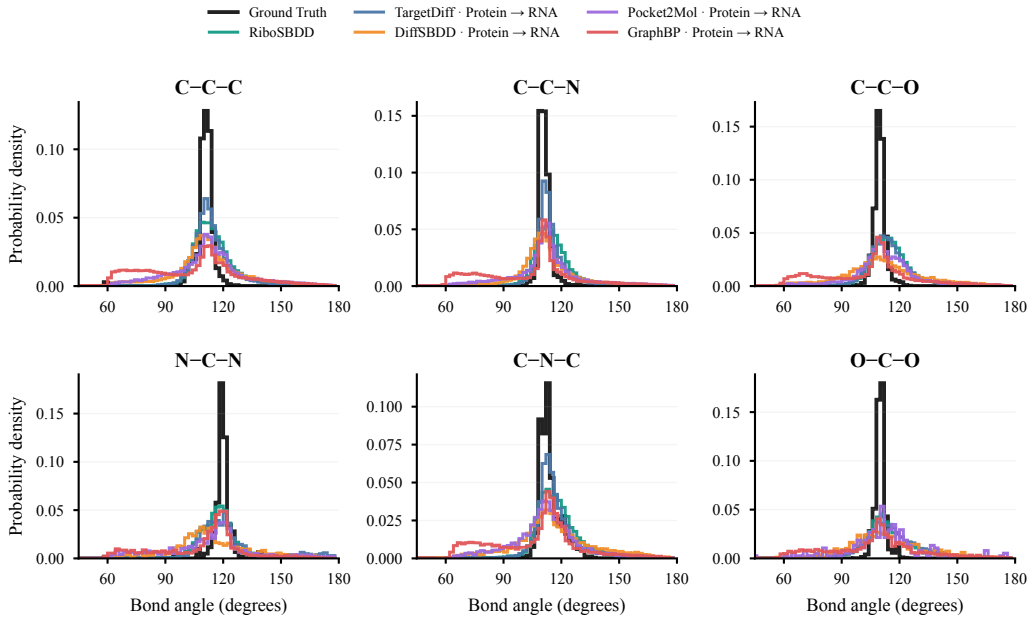
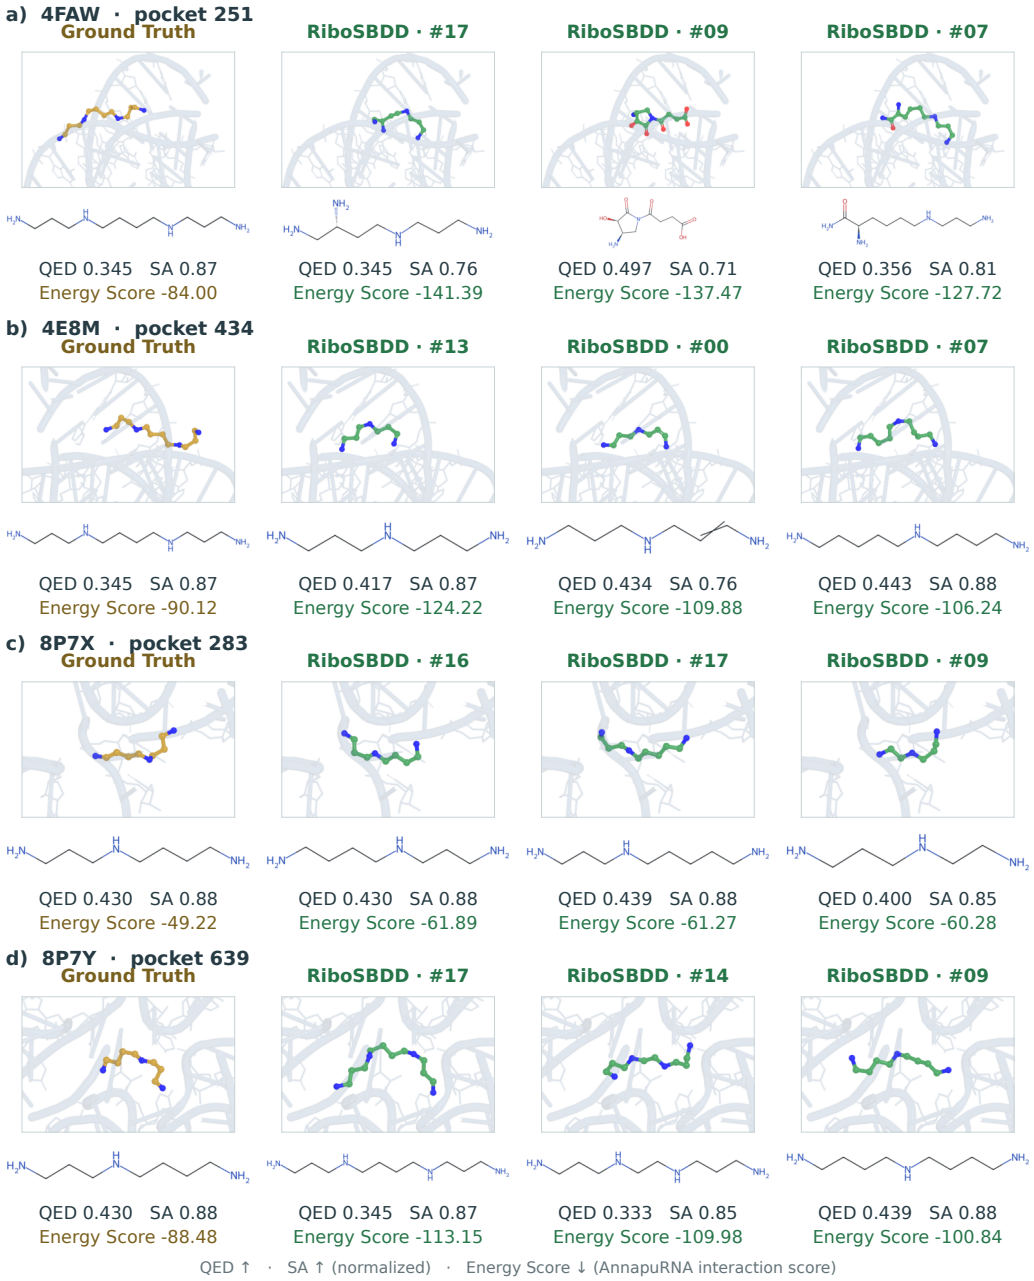
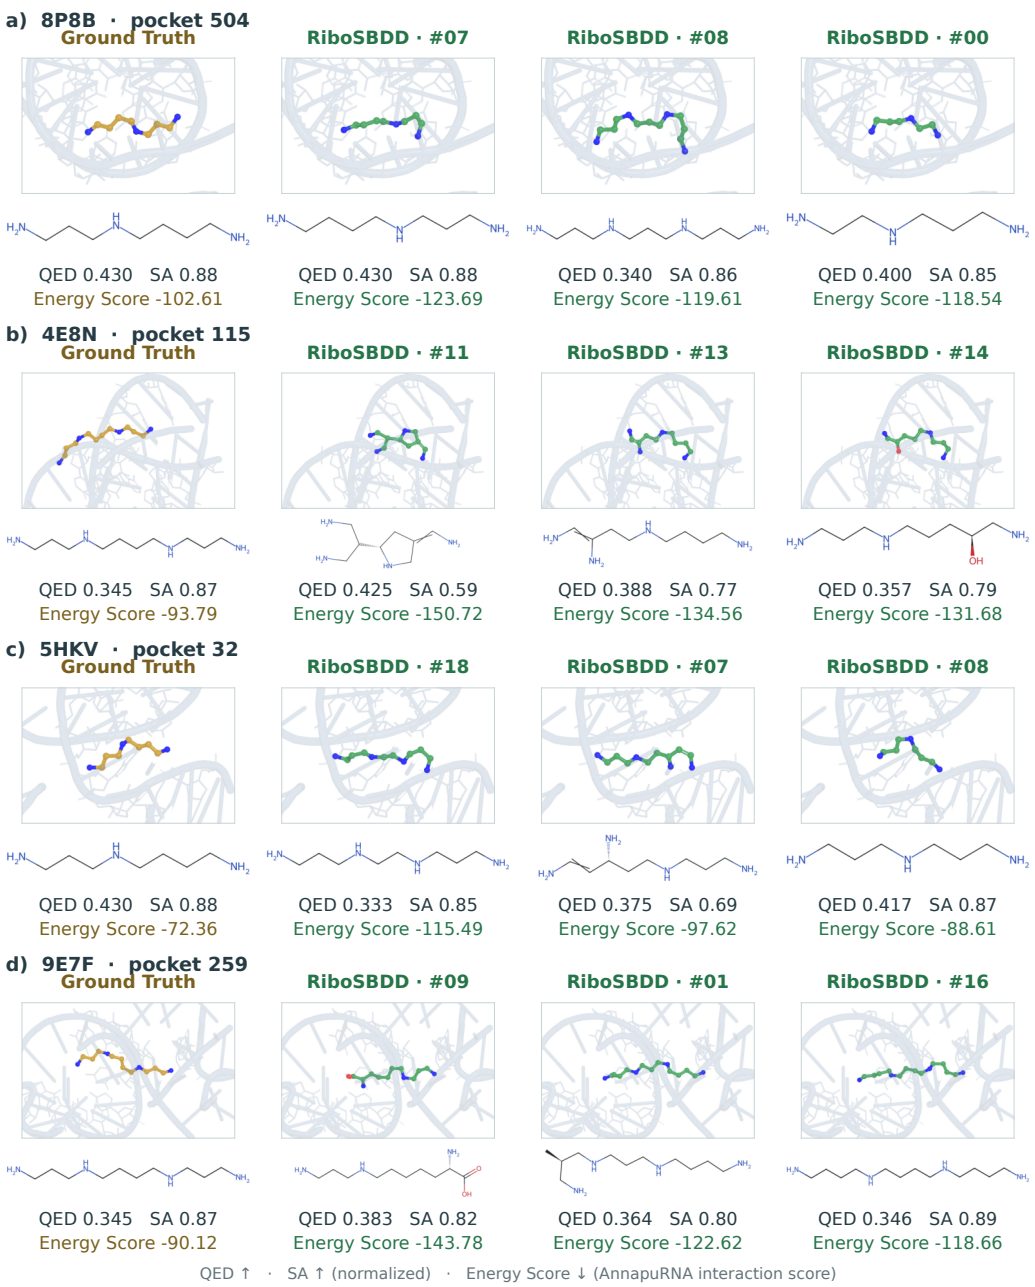


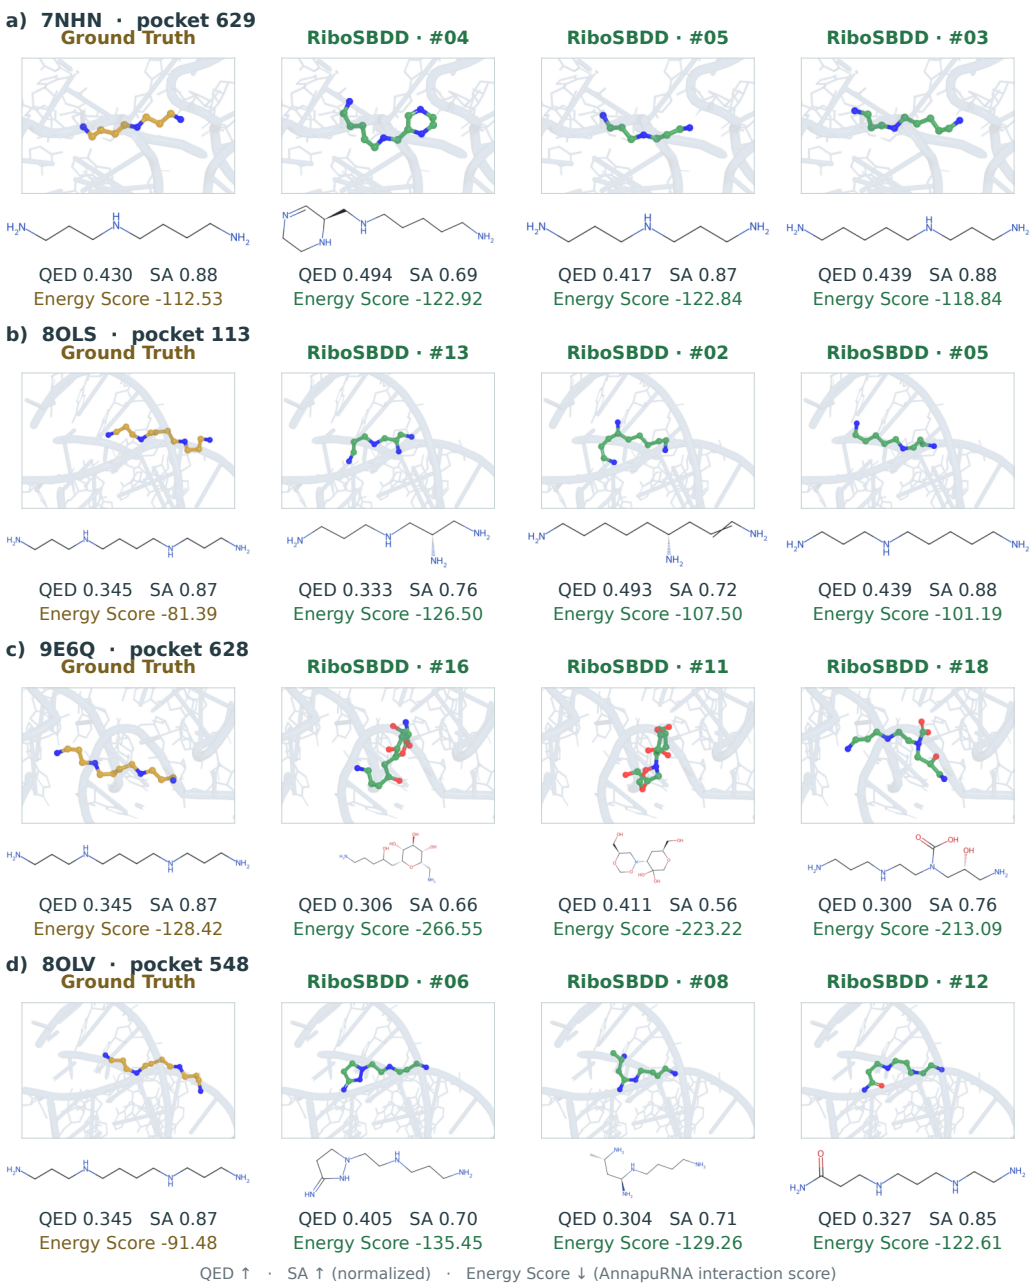
Figure 11: **Bond angle profiles across architectures.**

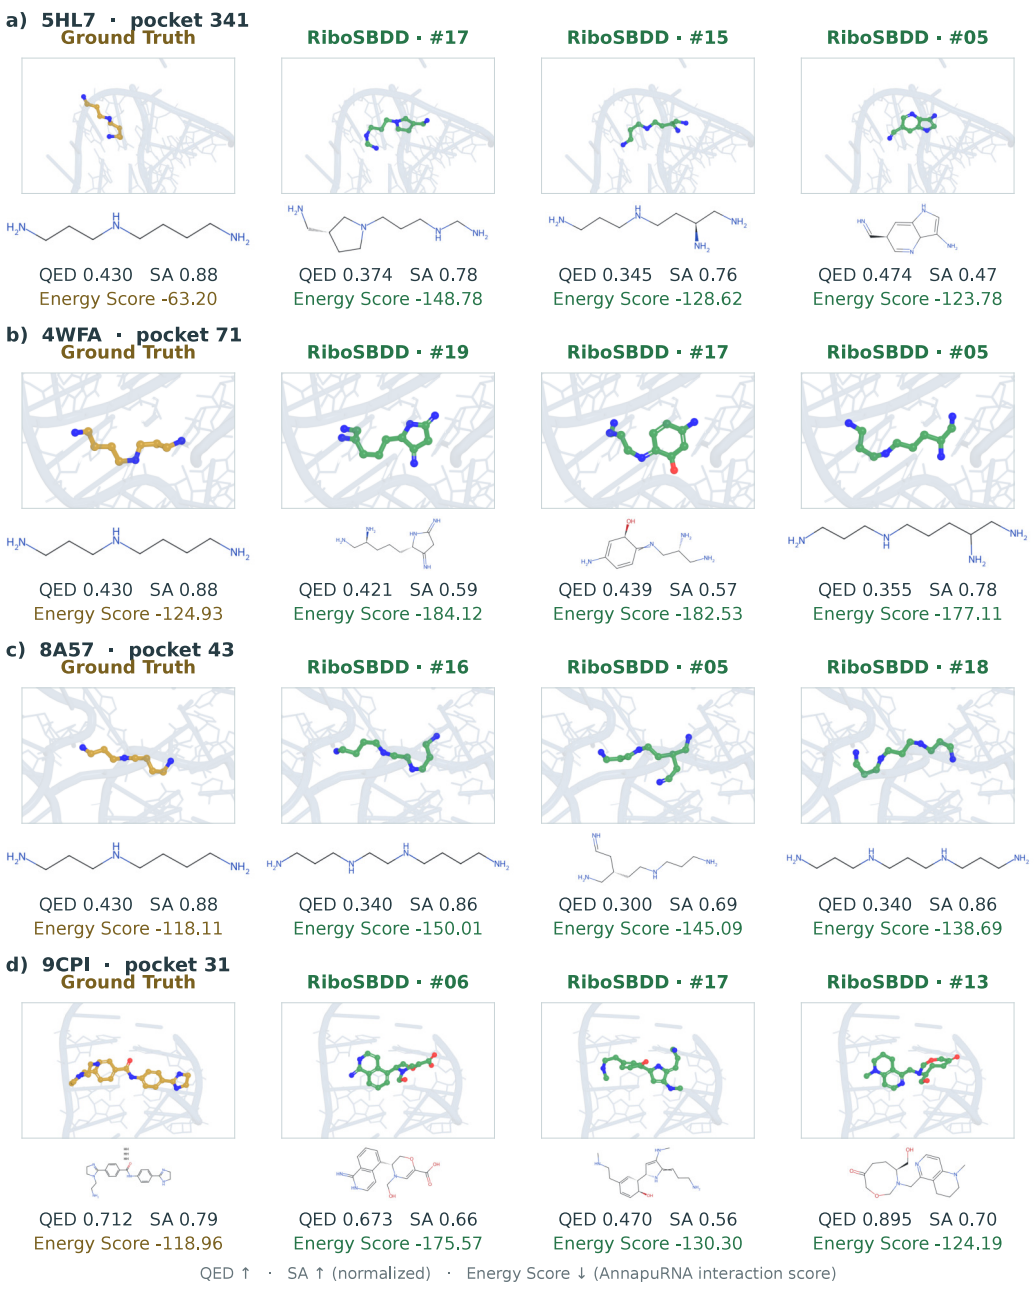
D.4 Additional Qualitative Examples

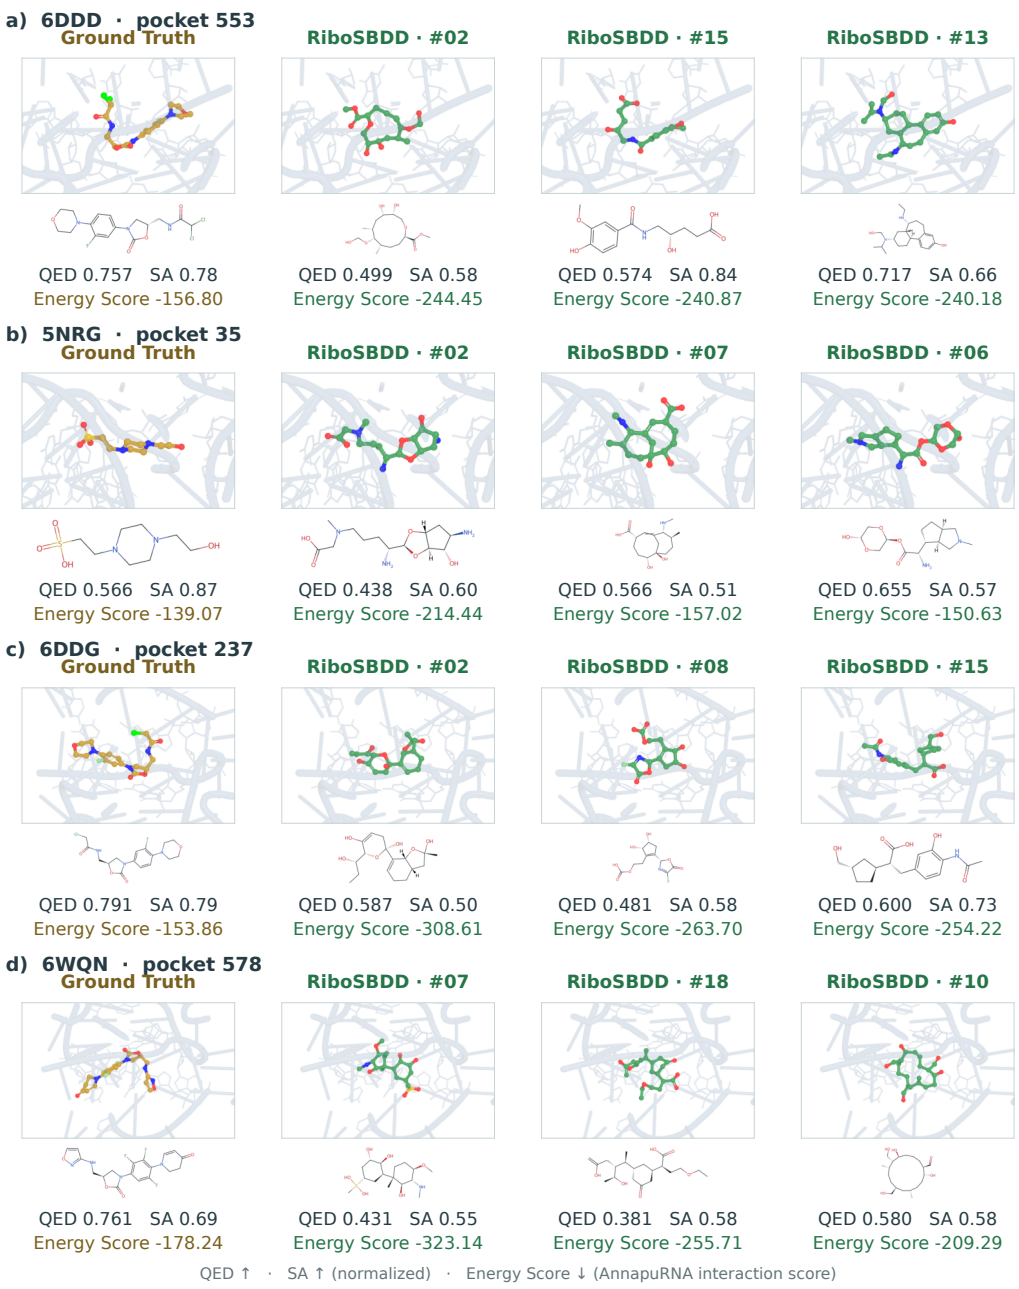
The following selected examples show generated candidates and deposited reference ligands in the same observed RNA pocket. The final panel includes candidates with less favorable AnnapuRNA scores than their references. These examples illustrate individual outcomes rather than overall success rates.

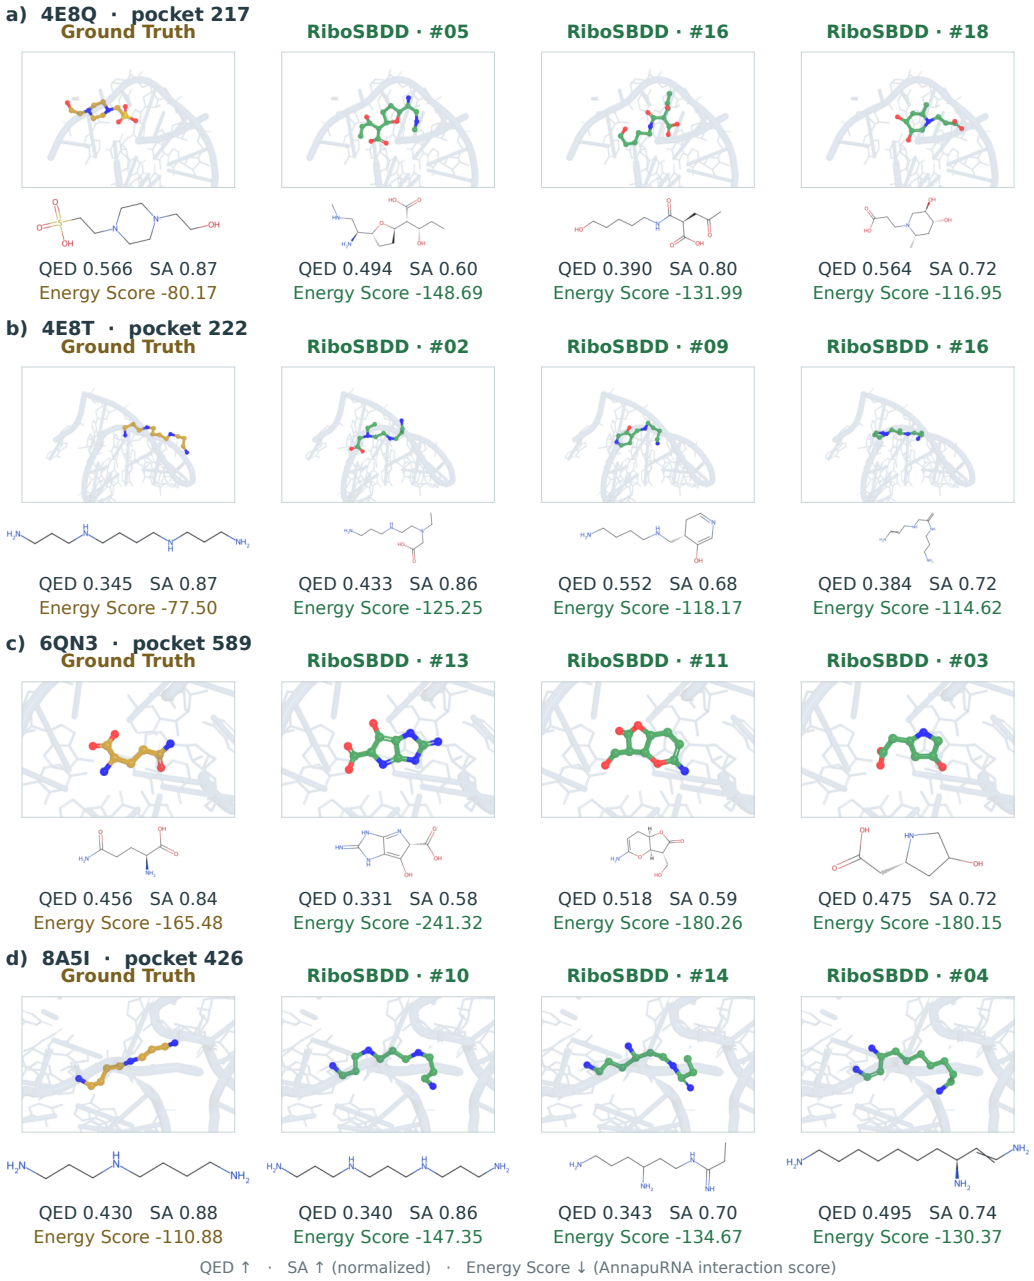

a) 4FAW · pocket 251
Ground Truth
RiboSBDD · #17
RiboSBDD · #09
RiboSBDD · #07
QED 0.345 SA 0.87
Energy Score -84.00
QED 0.345 SA 0.76
Energy Score -141.39
QED 0.497 SA 0.71
Energy Score -137.47
QED 0.356 SA 0.81
Energy Score -127.72
b) 4E8M · pocket 434
Ground Truth
RiboSBDD · #13
RiboSBDD · #00
RiboSBDD · #07
QED 0.345 SA 0.87
Energy Score -90.12
QED 0.417 SA 0.87
Energy Score -124.22
QED 0.434 SA 0.76
Energy Score -109.88
QED 0.443 SA 0.88
Energy Score -106.24
c) 8P7X · pocket 283
Ground Truth
RiboSBDD · #16
RiboSBDD · #17
RiboSBDD · #09
QED 0.430 SA 0.88
Energy Score -49.22
QED 0.430 SA 0.88
Energy Score -61.89
QED 0.439 SA 0.88
Energy Score -61.27
QED 0.400 SA 0.85
Energy Score -60.28
d) 8P7Y · pocket 639
Ground Truth
RiboSBDD · #17
RiboSBDD · #14
RiboSBDD · #09
QED 0.430 SA 0.88
Energy Score -88.48
QED 0.345 SA 0.87
Energy Score -113.15
QED 0.333 SA 0.85
Energy Score -109.98
QED 0.439 SA 0.88
Energy Score -100.84
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)

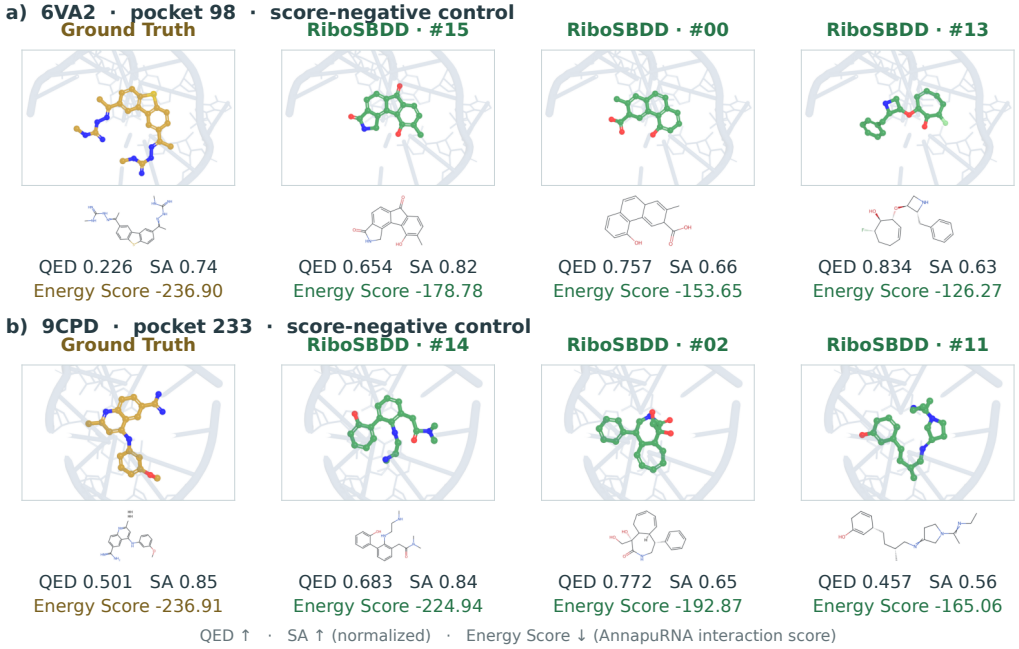

a) 8P8B · pocket 504
Ground Truth
QED 0.430 SA 0.88
Energy Score -102.61
RiboSBDD · #07
QED 0.430 SA 0.88
Energy Score -123.69
RiboSBDD · #08
QED 0.340 SA 0.86
Energy Score -119.61
RiboSBDD · #00
QED 0.400 SA 0.85
Energy Score -118.54
b) 4E8N · pocket 115
Ground Truth
QED 0.345 SA 0.87
Energy Score -93.79
RiboSBDD · #11
QED 0.425 SA 0.59
Energy Score -150.72
RiboSBDD · #13
QED 0.388 SA 0.77
Energy Score -134.56
RiboSBDD · #14
QED 0.357 SA 0.79
Energy Score -131.68
c) 5HKV · pocket 32
Ground Truth
QED 0.430 SA 0.88
Energy Score -72.36
RiboSBDD · #18
QED 0.333 SA 0.85
Energy Score -115.49
RiboSBDD · #07
QED 0.375 SA 0.69
Energy Score -97.62
RiboSBDD · #08
QED 0.417 SA 0.87
Energy Score -88.61
d) 9E7F · pocket 259
Ground Truth
QED 0.345 SA 0.87
Energy Score -90.12
RiboSBDD · #09
QED 0.383 SA 0.82
Energy Score -143.78
RiboSBDD · #01
QED 0.364 SA 0.80
Energy Score -122.62
RiboSBDD · #16
QED 0.346 SA 0.89
Energy Score -118.66
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)


a) 7NHN · pocket 629
Ground Truth
QED 0.430 SA 0.88
Energy Score -112.53
RiboSBDD · #04
QED 0.494 SA 0.69
Energy Score -122.92
RiboSBDD · #05
QED 0.417 SA 0.87
Energy Score -122.84
RiboSBDD · #03
QED 0.439 SA 0.88
Energy Score -118.84
b) 8OLS · pocket 113
Ground Truth
QED 0.345 SA 0.87
Energy Score -81.39
RiboSBDD · #13
QED 0.333 SA 0.76
Energy Score -126.50
RiboSBDD · #02
QED 0.493 SA 0.72
Energy Score -107.50
RiboSBDD · #05
QED 0.439 SA 0.88
Energy Score -101.19
c) 9E6Q · pocket 628
Ground Truth
QED 0.345 SA 0.87
Energy Score -128.42
RiboSBDD · #16
QED 0.306 SA 0.66
Energy Score -266.55
RiboSBDD · #11
QED 0.411 SA 0.56
Energy Score -223.22
RiboSBDD · #18
QED 0.300 SA 0.76
Energy Score -213.09
d) 8OLV · pocket 548
Ground Truth
QED 0.345 SA 0.87
Energy Score -91.48
RiboSBDD · #06
QED 0.405 SA 0.70
Energy Score -135.45
RiboSBDD · #08
QED 0.304 SA 0.71
Energy Score -129.26
RiboSBDD · #12
QED 0.327 SA 0.85
Energy Score -122.61
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)


a) 5HL7 · pocket 341
Ground Truth
QED 0.430 SA 0.88
Energy Score -63.20
RiboSBDD · #17
QED 0.374 SA 0.78
Energy Score -148.78
RiboSBDD · #15
QED 0.345 SA 0.76
Energy Score -128.62
RiboSBDD · #05
QED 0.474 SA 0.47
Energy Score -123.78
b) 4WFA · pocket 71
Ground Truth
QED 0.430 SA 0.88
Energy Score -124.93
RiboSBDD · #19
QED 0.421 SA 0.59
Energy Score -184.12
RiboSBDD · #17
QED 0.439 SA 0.57
Energy Score -182.53
RiboSBDD · #05
QED 0.355 SA 0.78
Energy Score -177.11
c) 8A57 · pocket 43
Ground Truth
QED 0.430 SA 0.88
Energy Score -118.11
RiboSBDD · #16
QED 0.340 SA 0.86
Energy Score -150.01
RiboSBDD · #05
QED 0.300 SA 0.69
Energy Score -145.09
RiboSBDD · #18
QED 0.340 SA 0.86
Energy Score -138.69
d) 9CPI · pocket 31
Ground Truth
QED 0.712 SA 0.79
Energy Score -118.96
RiboSBDD · #06
QED 0.673 SA 0.66
Energy Score -175.57
RiboSBDD · #17
QED 0.470 SA 0.56
Energy Score -130.30
RiboSBDD · #13
QED 0.895 SA 0.70
Energy Score -124.19
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)


a) 6DDD · pocket 553
Ground Truth
QED 0.757 SA 0.78
Energy Score -156.80
RiboSBDD · #02
QED 0.499 SA 0.58
Energy Score -244.45
RiboSBDD · #15
QED 0.574 SA 0.84
Energy Score -240.87
RiboSBDD · #13
QED 0.717 SA 0.66
Energy Score -240.18
b) 5NRG · pocket 35
Ground Truth
QED 0.566 SA 0.87
Energy Score -139.07
RiboSBDD · #02
QED 0.438 SA 0.60
Energy Score -214.44
RiboSBDD · #07
QED 0.566 SA 0.51
Energy Score -157.02
RiboSBDD · #06
QED 0.655 SA 0.57
Energy Score -150.63
c) 6DDG · pocket 237
Ground Truth
QED 0.791 SA 0.79
Energy Score -153.86
RiboSBDD · #02
QED 0.587 SA 0.50
Energy Score -308.61
RiboSBDD · #08
QED 0.481 SA 0.58
Energy Score -263.70
RiboSBDD · #15
QED 0.600 SA 0.73
Energy Score -254.22
d) 6WQN · pocket 578
Ground Truth
QED 0.761 SA 0.69
Energy Score -178.24
RiboSBDD · #07
QED 0.431 SA 0.55
Energy Score -323.14
RiboSBDD · #18
QED 0.381 SA 0.58
Energy Score -255.71
RiboSBDD · #10
QED 0.580 SA 0.58
Energy Score -209.29
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)


a) 4E8Q · pocket 217
Ground Truth
QED 0.566 SA 0.87
Energy Score -80.17
RiboSBDD · #05
QED 0.494 SA 0.60
Energy Score -148.69
RiboSBDD · #16
QED 0.390 SA 0.80
Energy Score -131.99
RiboSBDD · #18
QED 0.564 SA 0.72
Energy Score -116.95
b) 4E8T · pocket 222
Ground Truth
QED 0.345 SA 0.87
Energy Score -77.50
RiboSBDD · #02
QED 0.433 SA 0.86
Energy Score -125.25
RiboSBDD · #09
QED 0.552 SA 0.68
Energy Score -118.17
RiboSBDD · #16
QED 0.384 SA 0.72
Energy Score -114.62
c) 6QN3 · pocket 589
Ground Truth
QED 0.456 SA 0.84
Energy Score -165.48
RiboSBDD · #13
QED 0.331 SA 0.58
Energy Score -241.32
RiboSBDD · #11
QED 0.518 SA 0.59
Energy Score -180.26
RiboSBDD · #03
QED 0.475 SA 0.72
Energy Score -180.15
d) 8A5I · pocket 426
Ground Truth
QED 0.430 SA 0.88
Energy Score -110.88
RiboSBDD · #10
QED 0.340 SA 0.86
Energy Score -147.35
RiboSBDD · #14
QED 0.343 SA 0.70
Energy Score -134.67
RiboSBDD · #04
QED 0.495 SA 0.74
Energy Score -130.37
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)


a) 6VA2 · pocket 98 · score-negative control
Ground Truth
RiboSBDD · #15
RiboSBDD · #00
RiboSBDD · #13
QED 0.226 SA 0.74
Energy Score -236.90
QED 0.654 SA 0.82
Energy Score -178.78
QED 0.757 SA 0.66
Energy Score -153.65
QED 0.834 SA 0.63
Energy Score -126.27
b) 9CPD · pocket 233 · score-negative control
Ground Truth
RiboSBDD · #14
RiboSBDD · #02
RiboSBDD · #11
QED 0.501 SA 0.85
Energy Score -236.91
QED 0.683 SA 0.84
Energy Score -224.94
QED 0.772 SA 0.65
Energy Score -192.87
QED 0.457 SA 0.56
Energy Score -165.06
QED ↑ · SA ↑ (normalized) · Energy Score ↓ (AnnapuRNA interaction score)