REMATCH: LEARNING TO RANK ALTERNATIVE PAIRINGS FOR RNA–PROTEIN INTERACTION PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting RNA–protein interactions from sequence can help reveal how RNAs are processed, localized, and translated. Although pretrained sequence models and expressive interaction architectures provide rich molecular features, generalization to unfamiliar RNAs and proteins remains challenging. Training often relies on an uneven mix of evidence, that is, observed interactions have experimental support, but many negatives are untested pairs formed by re-pairing molecules. Treating these alternatives as confirmed negatives can suppress plausible interactions, especially when hard-example selection favors high-scoring pairs. These uncertain alternatives can still be useful, however, if they provide comparisons rather than additional negative labels. We introduce **ReMatch**, short for *Ranking-enhanced Matching*, which combines binary classification with bidirectional ranking. For each observed interaction, ranking teaches the model to prefer that pair over alternatives formed by replacing either molecule. The comparisons act on the classifier's interaction log-odds, allowing an alternative to remain above the interaction boundary while the observed pair receives a higher score. Across five validation folds, adding ranking improves accuracy by approximately 4.7 and 2.5 percentage points on two benchmarks with unseen RNA and protein identities, with AUROC gains on both. Additional classification supervision under the same hard-pair selection policy does not reproduce these gains, supporting the value of learning relative preferences beyond selecting difficult examples. Ranking also improves preference for observed molecular assignments in existing predictors, extending the benefit of this supervision across interaction architectures. Code and preprocessing pipelines are available at `https://anonymous.4open.science/r/rematch/`.

1 INTRODUCTION

RNA–protein interactions help determine when and where RNA molecules are processed, localized, and translated. Predicting these interactions from sequence can therefore guide experiments beyond the pairs cataloged in interaction databases (Yuan et al., 2014; Lin et al., 2020). Richer representations have driven progress in RNA–protein interaction (RPI) prediction, from composition and physicochemical features (Muppirala et al., 2011; Suresh et al., 2015) to representations learned by RNA and protein language models (Penić et al., 2025; Lin et al., 2023). Attention, feature fusion, and graph methods further model relationships between molecules (Zhou et al., 2023; Abir et al., 2025; Liu et al., 2025a; Catalano et al., 2026). Yet generalization also depends on what the training objective rewards. Random pair splits can preserve molecular identities across training and evaluation (Park & Marcotte, 2012), whereas componentwise splits require predictions for unfamiliar RNAs and proteins (Catalano et al., 2026).

The training evidence, however, is uneven. Observed interactions support positive labels, whereas many benchmark negatives are created by re-pairing molecules and excluding recorded positives (Pan et al., 2016; Catalano et al., 2026). Consequently, an absent record becomes a negative target, and binary cross-entropy (BCE) rewards lowering its score without experimental confirmation. This becomes particularly ambiguous during hard-example selection, because a high-scoring alternative may be a difficult negative or an interaction missing from the catalog. Several approaches address this incomplete evidence. Positive–unlabeled learning distinguishes missing observations from negative evidence (Elkan & Noto, 2008; Kiryo et al., 2017), and SURE screens for reliable RPI negatives (Li et al., 2023). RPIPLM organizes representations through supervised contrastive learning (Liu et al.,

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.

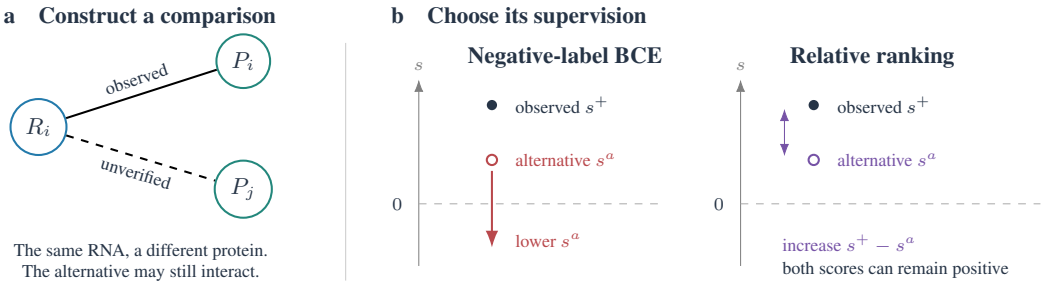


Figure 1: **Pointwise and pairwise supervision of alternative pairings.** Replacing one molecule in an observed interaction produces an unverified alternative. Negative-label BCE lowers its absolute interaction score; pairwise ranking increases the observed pair's relative advantage. The ranking constraint permits both scores to remain above the interaction boundary. ReMatch combines these comparisons with benchmark classification.

2025b), while NPI-DGCN learns relative interaction scores through pairwise constraints (Zhuo et al., 2022). These approaches motivate separating the choice of informative alternatives from the supervision assigned to them. To benefit prediction, the resulting preferences should shape the final interaction score.

Our premise is that *an alternative need not be a reliable negative to be a useful comparison*. An observed interaction provides a reference for alternatives formed by replacing either molecule while holding the other fixed. Experimental support gives the observed pair a basis for preference, while ranking allows both scores to remain above the interaction boundary (Figure 1). Building on this idea, we introduce **ReMatch**, short for *Ranking-enhanced Matching*. ReMatch combines benchmark BCE with bidirectional ranking at the classifier's interaction log-odds. Classification anchors individual predictions, while ranking teaches observed pairs to outrank hard alternatives sharing a molecule. Both objectives shape the prediction output, so existing RPI architectures can use this supervision without changing inference. The clearest gains occur when both molecular components are unseen. Across five validation folds, adding ranking to BCE improves accuracy by approximately 4.7 and 2.5 percentage points on the two componentwise benchmarks, with higher AUROC on both. Additional classification under the same hard-pair selection policy does not reproduce these gains, supporting the value of relative supervision beyond candidate selection. Ranking also improves assignment discrimination across predictors, while the combined objective retains more withheld interactions than additional negative-label supervision in controlled hidden-positive tests.

Our contributions are threefold. *(i)* We use unverified alternative pairings as comparisons, allowing them to guide learning without additional negative labels. *(ii)* We introduce ReMatch, which combines classification with bidirectional ranking at the interaction log-odds and can be applied to existing predictors without changing inference. *(iii)* We use controlled comparisons, score-location tests, and cross-model and hidden-positive analyses to connect the observed benefits to the design. We review related work in Section 2 and present the model and training objective in Section 3. Section 4 describes the evaluation protocol and examines predictive performance, design choices, and transfer across predictors. We conclude with the main findings and directions for future work.

2 Related Work

2.1 From molecular features to interaction models

Early sequence-based RPI predictors represented molecules through composition, physicochemical properties, and learned sequence features (Muppirala et al., 2011; Suresh et al., 2015; Yi et al., 2020; Armaos et al., 2021). RNA and protein language models now provide richer features learned from large sequence collections (Penić et al., 2025; Lin et al., 2023). With these representations in place, attention and feature-fusion methods learn relationships between the two molecular inputs (Zhou et al., 2023; Abir et al., 2025; Catalano et al., 2026), while graph-based approaches draw on interaction neighborhoods (Guan et al., 2025; Liu et al., 2025a; Zhao et al., 2026). These advances expand

the information available for recognizing molecular compatibility. Architectural expressiveness alone, however, does not require the prediction score to favor an observed interaction over plausible alternatives sharing either molecule, motivating explicit supervision of these pairing preferences.

2.2 Supervision from incomplete interaction records

Missing observations complicate learning because an unrecorded pair may be treated as negative without having been tested. Positive–unlabeled learning formalizes this distinction (Elkan & Noto, 2008; Kiryo et al., 2017), and SURE uses it to select reliable RPI negatives (Li et al., 2023). Comparisons offer another way to use the available evidence. Contrastive methods organize sequence, graph, and paired-complex representations (Khosla et al., 2020; Liu et al., 2025b; Gao et al., 2025; Han et al., 2025), while debiasing and hard-negative sampling address the effect of false negatives on these comparisons (Chuang et al., 2020; Robinson et al., 2021). Preference learning instead uses relative orderings: BPR learns from implicit observations (Rendle et al., 2009), and NPI-DGCN applies a pairwise AUC objective to graph-derived ncRNA–protein scores (Zhuo et al., 2022). These approaches establish the value of comparative evidence. However, representation-level contrastive losses and standalone link-ranking objectives do not jointly anchor absolute interaction predictions and supervise bidirectional pairing preferences at the classifier's decision score.

2.3 Evaluation beyond familiar molecular identities

These supervision choices become particularly relevant when the molecules themselves are unfamiliar. Random pair splits can reuse RNA and protein identities across training and evaluation, preserving entity-specific patterns (Park & Marcotte, 2012; Geirhos et al., 2020). Protein-diversity studies and CORAL's componentwise splits impose stronger separation between molecular components (Yu et al., 2026; Catalano et al., 2026), while PaRPI examines transfer across proteins, protocols, and batches in binding-site prediction (Peng et al., 2025). Such evaluations test whether learned interaction rules extend beyond familiar identities. Generalization scores alone, however, do not isolate dependence on the pairing itself, motivating complementary assignment tests that hold the molecules fixed and cancel additive molecule-specific score contributions.

3 ReMatch

ReMatch makes alternative pairings useful without requiring an additional binary label for each one. It retains benchmark classification and adds comparisons derived from observed interactions, with both objectives acting on the same interaction score. This simple coupling lets evidence about relative preference improve the function used for prediction (Figure 2).

3.1 A Shared Score for Classification and Ranking

For a sequence pair $x = (R, P)$, the predictor produces logits $\ell_0(x)$ and $\ell_1(x)$, with class 1 denoting interaction. We use their difference as the common score,

$$s_\theta(R, P) = \ell_1(R, P) - \ell_0(R, P), \qquad p_\theta(y = 1 \mid R, P) = \sigma(s_\theta(R, P)). \tag{1}$$

Using log-odds connects ranking directly to the interaction decision: increasing this score raises the predicted probability. It also removes the arbitrary common offset of the two raw logits.

Frozen RNA and protein language models provide pooled sequence features (Penić et al., 2025; Lin et al., 2023). Trainable projections map them to $z_R, z_P \in \mathbb{R}^d$. The interaction head combines a gated mixture, an elementwise product, and a factorized bilinear feature,

$$z_{RP} = \mathrm{LayerNorm}([\, g \odot z_R + (1 - g) \odot z_P \,;\; z_R \odot z_P \,;\; U(z_R) \odot V(z_P) \,]) \,. \tag{2}$$

Here $g \in (0, 1)^d$ is produced by a two-layer network conditioned on both molecules, and U, V are learned linear maps. An MLP maps the combined feature to the two logits. Observed pairs and alternatives share all hyperparamters; encoder configurations and the shared projection dimension appear in Appendix B.

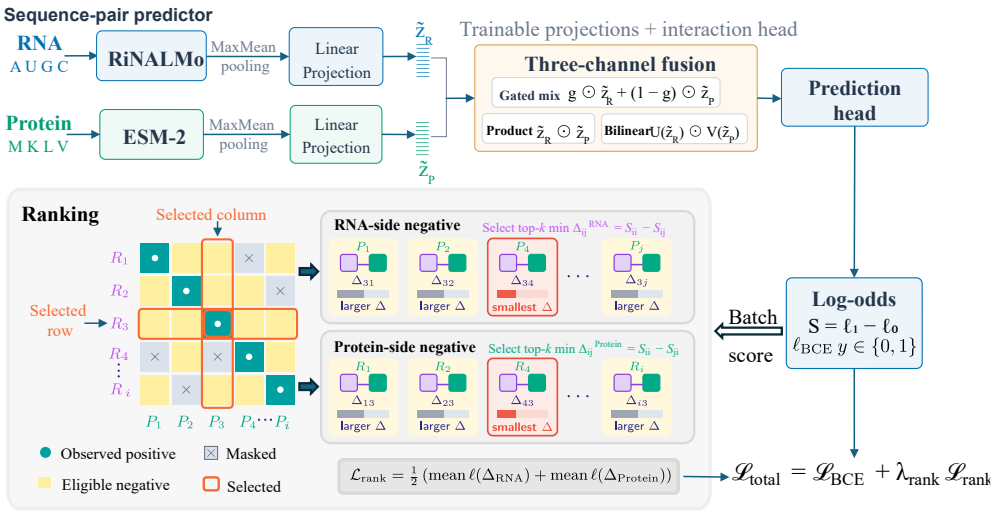


Figure 2: **Joint classification and ranking in the ReMatch framework.** Frozen sequence encoders feed a trainable interaction head. Benchmark pairs receive BCE supervision. Observed positive pairs form a cross-pair score matrix for ranking: known training interactions are masked, and hard alternatives are selected in both replacement directions. Both objectives supervise the classifier's interaction log-odds. Inference requires only the single-pair prediction path.

3.2 Alternative Pairings from Observed Interactions

Given positive minibatch pairs $\mathcal{B}^+ = \{(R_i, P_i)\}_{i=1}^n$, we compute $s_{ij} = s_\theta(R_i, P_j)$. Let K_{train} contain all positive pair identities in the training fold. Eligible alternatives replace the protein or RNA while excluding known interactions,

$$\mathcal{C}_i^P = \{j : j \neq i,\ (R_i, P_j) \notin K_{\text{train}}\}, \qquad \mathcal{C}_i^R = \{j : j \neq i,\ (R_j, P_i) \notin K_{\text{train}}\}. \tag{3}$$

The mask checks pair identities against the entire training fold and uses no held-out labels. An eligible cross-pair is therefore unobserved as positive in the training fold, rather than experimentally established as negative.

The corresponding score margins are

$$\Delta_{ij}^P = s_{ii} - s_{ij}, \qquad \Delta_{ij}^R = s_{ii} - s_{ji}. \tag{4}$$

We retain the k smallest eligible margins for each anchor and direction, or all candidates when fewer are available. These are the alternatives that the model finds most plausible, making them useful tests of its preference for the observed interaction. Selection is discrete, with gradients flowing through the selected scores. The two directions also make a score based on only one molecule insufficient: an RNA-only score ties under protein replacement, and a protein-only score ties under RNA replacement. To improve both comparisons, the predictor must use information from both inputs.

3.3 Joint Classification and Ranking

Let $\mathcal{H}_i^d$ contain the retained alternatives, $K_i^d = |\mathcal{H}_i^d|$, and $\mathcal{A}^d = \{i : K_i^d > 0\}$. The ranking objective averages comparisons within each anchor and then averages anchors within each direction,

$$\mathcal{L}_{\text{rank}} = \frac{1}{2} \sum_{d \in \{P,R\}} \frac{1}{|\mathcal{A}^d|} \sum_{i \in \mathcal{A}^d} \frac{1}{K_i^d} \sum_{j \in \mathcal{H}_i^d} \text{softplus}\left(-\frac{\Delta_{ij}^d}{\tau}\right), \tag{5}$$

where $\text{softplus}(u) = \log(1 + e^u)$ and $\tau > 0$ sets the score-difference scale. Averaging within each anchor prevents pairs with more eligible alternatives from receiving more weight. Each nonempty anchor contributes equally; an empty direction contributes zero.

Benchmark positives and negatives retain their binary classification loss,

$$\mathcal{L}_{\mathrm{BCE}} = \frac{1}{|\mathcal{B}|} \sum_{(x,y)\in\mathcal{B}} \big[\,\mathrm{softplus}(s_\theta(x)) - y\, s_\theta(x)\big], \qquad \mathcal{L}(t) = \mathcal{L}_{\mathrm{BCE}} + \lambda_{\mathrm{rank}}(t)\mathcal{L}_{\mathrm{rank}}. \quad (6)$$

Constructed alternatives enter the auxiliary objective without additional class labels. Benchmark pairs retain their classification targets. We use $k = 40$, $\tau = 0.25$, and unit ranking weight after a two-epoch warm-up, followed by linear annealing from epoch 15 (Appendix B). The extra computation is confined to training. Constructing a local score matrix requires up to n^2 interaction-head evaluations, while top-k selection limits the comparisons entering the loss. Frozen sequence features are reused, and inference evaluates each queried pair once.

3.4 Relative Constraints on Alternative Pairings

For observed and alternative scores s^+ and s^a, negative BCE and ranking impose different constraints,

$$\ell_{\mathrm{neg}}(s^a) = \mathrm{softplus}(s^a), \qquad \ell_{\mathrm{rank}}(s^+, s^a) = \mathrm{softplus}\left(\frac{s^a - s^+}{\tau}\right). \quad (7)$$

Negative BCE continues to reward lowering an alternative that remains above the interaction boundary. Ranking instead rewards the observed pair's advantage, so the same comparison can receive a low loss while both scores remain positive. This flexibility is useful when an alternative is an unrecorded interaction: the model can favor the observed pair without having to place the alternative below the interaction boundary.

This freedom comes from a simple invariance: shifting every score by a constant, $s(x) \mapsto s(x) + c$, preserves all ranking margins. It also explains why ReMatch keeps classification. The shift can move individual predictions across the decision boundary, so BCE anchors absolute scores to the benchmark labels while ranking resolves comparisons between related pairs. The two losses serve different purposes even though they share the same output. Our evaluation reflects both by reporting fixed-boundary accuracy and threshold-independent AUROC. The preference is based on experimental observation, rather than a claim that the alternative cannot interact. This suggests a testable benefit: relative supervision should retain more hidden positives above the interaction boundary than adding negative labels to those comparisons. Section 4.5 examines that behavior; Appendix A gives the gradient comparison.

Connection to preference and contrastive learning. Equation (5) uses the logistic preference loss associated with Bradley–Terry models, RankNet, and BPR (Bradley & Terry, 1952; Burges et al., 2005; Rendle et al., 2009). A single comparison is also a two-candidate softmax contrastive loss (van den Oord et al., 2018). With several alternatives, pairwise averaging applies a separate loss to each comparison, whereas multi-candidate normalization couples their scores. The choice central to ReMatch is to supervise the learned interaction function itself, so that the preferences shape the score used at inference. Section 4.3 examines both the supervised score and the aggregation rule, while Appendix G.1 describes the details of these controls and Appendix A shows the corresponding identities.

4 Experiments

We first evaluate prediction for familiar and unseen molecular identities, then examine the contributions of ranking and the score it supervises. Cross-model and hidden-positive analyses further connect these design choices to assignment discrimination and the treatment of unverified alternatives.

4.1 Evaluation Protocol

Benchmarks and model selection. We use three benchmarks with two generalization settings. The random-pair NPInter2.0 benchmark contains 10,412 positives and an equal number of re-paired negatives (Yuan et al., 2014). CORAL's componentwise CW30_50 and CW60_80 folds draw from the same pool of 24,012 positives and use re-paired negatives, with no exact RNA or protein identity shared between training and validation (Catalano et al., 2026). Their suffixes specify protein/RNA

Table 1: **Predictive performance on random-pair and componentwise RPI benchmarks.** Baselines reuse completed reproduction runs, with one run per reported setting; ReMatch reports five-fold validation means. Full metrics and the reproduction protocol appear in Appendix I.1.

	NPInter2.0		CW30_50		CW60_80	
Method	ACC	AUROC	ACC	AUROC	ACC	AUROC
CORAL	0.893	0.938	0.646	0.693	<u>0.684</u>	0.757
ZeRPI	0.946	0.983	0.657	<u>0.799</u>	0.663	<u>0.785</u>
PRInter	0.822	0.979	0.500	0.542	0.500	0.568
MHAM-NPI	0.893	0.936	0.634	0.605	0.638	0.654
ZHMolGraph	0.935	0.971	0.659	0.709	0.679	0.726
ZHMolTopoRPI	**0.958**	<u>0.985</u>	<u>0.676</u>	0.756	0.680	0.775
ReMatch (this work)	<u>0.954</u>	**0.986**	**0.770**	**0.844**	**0.729**	**0.799**

clustering thresholds. Compared with the random-pair setting of NPInter2.0, the CW benchmarks impose a more demanding generalization requirement, i.e., both molecular identities are unseen during training. They therefore provide our primary test of whether learned compatibility transfers to unfamiliar RNAs and proteins.

We train one run per fold across five folds for 40 epochs. The primary training configuration is selected on NPInter and fixed for both CW benchmarks; the score-location controls use fixed temperatures per score location. Within each fold, the checkpoint with highest validation accuracy is evaluated on that same validation split. Reported values are therefore validation means, without an independent selection split or untouched test set. We report accuracy at $s > 0$ and AUROC. Table values use three decimal places; bold and underlining indicate the first and second distinct ranks at source precision. Where used, blue shading follows the ordering within each comparison group.

Objective controls. BCE uses benchmark labels. BCE-H adds balanced cross-pair classification, pairing each selected hard alternative labeled 0 with a repeated observed anchor labeled 1. Rank uses only relative supervision, while BCE-H + Rank combines benchmark BCE with both auxiliary losses. BCE + Rank is ReMatch. The controls share features, architecture, training budget, and the top-k eligibility and hardness policy, although selected pairs can differ as the models evolve. Full training settings, data provenance and split audits, and complete results with fold dispersion are provided in Appendices B, C, and D.

4.2 Prediction for Unseen Molecular Identities

ReMatch's clearest advantage appears when both molecular components are unseen (Table 1). On CW30_50 and CW60_80, its accuracy exceeds the strongest reproduced baseline by approximately 9.4 and 4.5 percentage points, respectively, with higher AUROC on both. It also remains competitive on NPInter, achieving the highest AUROC while ZHMolTopoRPI has the highest accuracy. The gains thus extend beyond the random-pair setting in which molecular identities may recur.

The baselines cover attention, pretrained features, and molecular-network models (Catalano et al., 2026; Gao et al., 2025; Zhao et al., 2026; Zhou et al., 2023; Liu et al., 2025a; Chen et al., 2026). Their comparison shows the practical benefit of ReMatch across different generalization settings. Baseline rows reuse completed reproduction runs, with one run per reported setting, whereas ReMatch reports five-fold validation means. To examine what the ranking objective contributes within a shared architecture, we next turn to the controlled ablations.

4.3 Effects of the Ranking Objective and Supervised Score

Relative supervision contributes beyond candidate selection. Adding ranking to BCE raises accuracy by approximately 4.7 percentage points on CW30_50 and 2.5 points on CW60_80, with AUROC improving on both (Table 2). NPInter accuracy remains comparable and AUROC increases by 0.004. The largest accuracy gains therefore occur where the predictor must transfer beyond familiar molecular identities. This pattern is consistent with the purpose of ranking: it explicitly trains the interaction score to distinguish related pairings that differ in one molecule.

Table 2: **Ablation of training objectives and hard-pair sampling.** Five-fold validation means. BCE-H augments benchmark BCE with balanced classification of hard cross-pairs; ReMatch augments benchmark BCE with pairwise ranking. All variants use the same interaction head. Full metrics and fold dispersion appear in Appendix D.

	NPInter2.0		CW30_50		CW60_80	
Objective	ACC	AUROC	ACC	AUROC	ACC	AUROC
BCE	**0.955**	0.982	0.723	0.813	0.703	0.791
BCE-H	0.931	0.974	0.736	0.840	0.701	0.812
Rank	0.902	0.961	0.735	0.805	0.722	0.802
BCE-H + Rank	0.942	0.974	0.766	0.834	0.726	**0.813**
BCE + Rank (this work)	0.954	**0.986**	**0.770**	**0.844**	**0.729**	0.799

Table 3: **Ablation of supervision location and loss aggregation.** Five-fold validation means with benchmark BCE retained. Embedding cosine similarity and interaction log-odds use temperatures $\tau = 0.100$ and $\tau = 0.250$. Pairwise log-odds is the primary ReMatch configuration. Appendix G.1 provides the full control design.

		NPInter2.0		CW30_50		CW60_80	
Score	Aggregation	ACC	AUROC	ACC	AUROC	ACC	AUROC
Embedding cosine	Pairwise	0.955	0.986	0.711	0.802	0.694	0.766
Embedding cosine	Multi-candidate	**0.960**	**0.989**	0.717	0.816	0.704	0.772
Interaction log-odds	Pairwise (this work)	0.954	0.986	**0.770**	**0.844**	0.729	0.799
Interaction log-odds	Multi-candidate	0.959	0.988	0.766	0.834	**0.741**	**0.815**

Hard-pair classification alone gives a less consistent benefit. BCE-H improves CW30_50 accuracy modestly but lowers it on the other two benchmarks. Adding ranking to BCE-H then raises accuracy on all three and yields the highest CW60_80 AUROC. Thus, under the same candidate-selection policy, the way alternatives are supervised changes their value. Selecting difficult pairs alone does not reproduce the benefit of learning relative preferences.

The ranking-only control also explains why ReMatch retains classification. Rank exceeds BCE accuracy on both CW benchmarks, but loses substantially on NPInter and has lower CW30_50 AUROC. Combining the two objectives gives the strongest CW accuracy in this ablation: benchmark BCE anchors individual predictions, while ranking adds distinctions between related pairs.

Supervising the final interaction score. The objective controls establish the value of relative supervision. A separate 2×2 comparison examines where it should act: projected-embedding cosine similarity or final interaction log-odds, each under pairwise and multi-candidate aggregation. All variants retain benchmark BCE, with fixed temperatures per score location.

Direct log-odds supervision improves both accuracy and AUROC under both aggregation rules on the two CW benchmarks (Table 3). On CW30_50, the accuracy gains over embedding similarity are 5.9 percentage points for pairwise aggregation and approximately 5.0 points for multi-candidate aggregation. The corresponding gains on CW60_80 are 3.5 and 3.7 points. This consistent ordering supports placing comparisons after the interaction head, where they supervise the full computation used to judge a pair.

Aggregation has a more dataset-specific effect. Pairwise log-odds performs best on CW30_50, whereas multi-candidate log-odds leads on CW60_80; NPInter instead favors embedding similarity with multi-candidate aggregation. Across the two componentwise benchmarks, prediction-level supervision is therefore the more consistent design choice. A separate direction control also supports using both replacements. With the overall ranking coefficient held fixed, CW30_50 accuracy reaches 0.770, compared with 0.758 for protein replacement alone and 0.728 for RNA replacement alone.

4.4 TRANSFER OF RANKING SUPERVISION ACROSS PREDICTORS

Because ReMatch applies ranking to the prediction score, other interaction architectures can use the same supervision. We add it to ZHMolTopoRPI, ZeRPI, and BioLLMNet, alongside the ReMatch head. To test the behavior that ranking is intended to encourage, we compare the observed assignment of two interactions with their crossed assignment,

$$D_{ij} = s(R_i, P_i) + s(R_j, P_j) - s(R_i, P_j) - s(R_j, P_i). \tag{8}$$

Every molecule appears once on either side. Separate RNA and protein contributions in an additive score $s(R, P) = a(R) + b(P)$ cancel, so positive D identifies a preference for how the molecules are paired. All models use the same valid quartets within a benchmark, excluding known-positive crossed pairs. We report strict $P(D > 0)$, so exact ties do not count as successful preferences.

Table 4: **Transfer of ReMatch ranking supervision to existing RPI predictors.** Values are strict assignment-preference rates $P(D > 0)$ on shared quartets within each benchmark. Parentheses report relative improvements over Base, calculated from the displayed values. Ranking improves assignment preference for every predictor on both benchmarks. ZHMolTopoRPI and BioLLMNet use matched pre-sigmoid scores; the ReMatch head uses interaction log-odds. ZeRPI retains its reported score convention. Full statistics appear in Appendix E.

	NPInter2.0		CW30_50	
Predictor	Base	+ Rank	Base	+ Rank
ZHMolTopoRPI	0.829	0.866 (+ 4.5%)	0.687	0.723 (+5.2%)
ZeRPI	0.776	0.872 (+12.4%)	0.679	0.710 (+4.6%)
BioLLMNet	0.727	0.811 (+11.6%)	0.701	0.714 (+1.9%)
ReMatch head	0.764	0.854 (+11.8%)	0.708	0.721 (+1.8%)

Ranking improves assignment preference for every predictor in Table 4 on both benchmarks. On NPInter, the gains are approximately 3.6 percentage points for ZHMolTopoRPI and 8.4 points for BioLLMNet, with improvements on CW30_50 as well. Their Base and +Rank variants use matched pre-sigmoid scores, preserving the meaning of the diagnostic within each model family. ZeRPI follows the same trend under its retained evaluation convention.

The shared improvement is useful because it does not require adopting the ReMatch fusion module. Different predictors learn a stronger preference for observed molecular assignments through their own outputs, supporting ranking as reusable supervision. This diagnostic complements classification by holding the molecules fixed and testing their assignment, rather than the absolute decision on an individual pair.

4.5 HIDDEN-POSITIVE RETENTION UNDER RELATIVE SUPERVISION

The central motivation predicts that a useful alternative need not be driven below the interaction boundary. We examine this behavior in the Hub-H analysis on CW30_50. In each fold, a fixed set H containing 10% of training positives is selected from interactions involving high-degree proteins. These pairs are removed from positive sampling and the known-positive mask, then injected as alternatives at four representative realized rates from 0% to approximately 33%. Validation pairs remain unchanged. At zero injection, H is still hidden, providing a masking control.

The four objectives share the same hidden set, batches, and designated injections. BCE-H assigns negative labels to selected alternatives; Rank uses relative supervision without benchmark BCE; BCE-H + Rank includes both forms of auxiliary supervision; and ReMatch retains benchmark BCE with ranking alone. Ordinary alternatives are mined online under each objective. We measure the fraction of hidden positives above the interaction boundary at the selected checkpoint, alongside validation accuracy and AUROC (Table 5).

ReMatch retains the largest fraction of hidden positives above zero and achieves the highest accuracy and AUROC at every representative injection level. The comparison with BCE-H + Rank is particularly informative: both retain benchmark classification and ranking, but the latter additionally labels the selected alternatives negative. At approximately 11% injection, ReMatch retains 27.0% of hidden

Table 5: **Hidden-positive retention and predictive performance under relative supervision.** Five-fold CW30_50 means. Realized ρ is the fraction of comparison slots occupied by designated injections; $P(S_h > 0)$ is hidden-positive retention at the selected checkpoint. The zero-injection condition still hides these positives from supervision. Bold and underlining compare objectives within each setting. Complete score statistics appear in Appendix F.

Realized ρ (%)	Objective	$P(S_h > 0) \uparrow$	ACC	AUROC
0.000	BCE-H	<u>0.216</u>	0.688	0.803
	Rank	0.139	<u>0.716</u>	0.794
	ReMatch	**0.316**	**0.749**	**0.836**
	BCE-H + Rank	0.180	0.706	<u>0.820</u>
11.390	BCE-H	0.112	0.643	0.790
	Rank	<u>0.123</u>	<u>0.710</u>	0.785
	ReMatch	**0.270**	**0.739**	**0.836**
	BCE-H + Rank	0.108	0.689	<u>0.814</u>
21.720	BCE-H	<u>0.129</u>	0.673	0.814
	Rank	0.116	<u>0.709</u>	0.782
	ReMatch	**0.259**	**0.737**	**0.833**
	BCE-H + Rank	0.119	0.699	<u>0.824</u>
33.070	BCE-H	0.121	0.666	0.803
	Rank	0.105	<u>0.707</u>	0.778
	ReMatch	**0.168**	**0.728**	**0.828**
	BCE-H + Rank	<u>0.127</u>	0.693	<u>0.824</u>

positives above the boundary, compared with 10.8% for BCE-H + Rank. At approximately 22%, the corresponding fractions are 25.9% and 11.9%. These differences support the intended behavior of the relative constraint, which leaves more room for unrecorded interactions.

Better retention also accompanies better prediction. At approximately 11% injection, ReMatch achieves accuracy/AUROC of $0.739/0.836$, versus $0.689/0.814$ for BCE-H + Rank. Rank alone produces smaller average score shifts, yet has lower retention and predictive performance than ReMatch. The combination therefore matters: relative comparisons preserve flexibility for alternatives, while benchmark BCE provides an absolute anchor for prediction. The advantage is already present at zero injection and does not grow monotonically with the injection rate. These runs support the predicted score-retention behavior, without establishing false-negative tolerance as the sole explanation for the generalization gains.

5 Conclusion and Future Work

An uncertain pairing can guide learning without first being labeled negative. ReMatch combines benchmark classification with bidirectional ranking at the interaction log-odds, anchoring individual predictions while learning pairing preferences without changing inference. The strongest gains occur when both molecules are unseen. Controlled comparisons show benefits beyond hard-pair sampling and favor the final prediction score on componentwise benchmarks. Ranking also improves assignment discrimination across predictors, while hidden-positive tests show better retention of withheld interactions. Two directions follow from these findings. First, comparisons could reflect the strength of experimental evidence. An observed interaction provides a training reference, but it need not bind more strongly than an unobserved alternative. Measurements collected under matched conditions could support reliable orderings, ties, or unresolved comparisons, helping separate molecular compatibility from uneven experimental coverage. Second, context-aware ranking could distinguish molecular compatibility from the opportunity to interact. A model could combine sequence features with expression and localization, while accounting separately for assay detection. Paired in vitro and cellular measurements could then test whether it explains changes in pairing preference for the same molecules across conditions.

AI USE STATEMENT

Generative AI tools were used to assist with literature search and language polishing of the manuscript. The authors independently reviewed and verified the references, equations, source data, and interpretations, and take full responsibility for the final content.

REPRODUCIBILITY STATEMENT

The anonymous repository linked above provides the implementation and preprocessing pipelines. Section 3 defines the objectives; Appendix B records the training configuration, objective schedules, and model selection. Data sources and complete metrics appear in Appendices C and D. The supplementary snapshot includes processed folds and feature-cache manifests, configurations, selected-checkpoint identifiers, fold-level predictions, and scripts mapping reported tables to their source runs. Original reproduction results are separated from the current controlled experiments in Appendix I.

REFERENCES

Abrar Rahman Abir, Md Toki Tahmid, and Md Shamsuzzoha Bayzid. BioLLMNet: enhancing RNA-interaction prediction with a specialized cross-LLM transformation network. *Briefings in Bioinformatics*, 26(5):bbaf549, 2025. doi: 10.1093/bib/bbaf549. URL https://doi.org/10.1093/bib/bbaf549.

Alexandros Armaos, Alessio Colantoni, Gabriele Proietti, Jakob Rupert, and Gian Gaetano Tartaglia. catRAPID omics v2.0: going deeper and wider in the prediction of protein–RNA interactions. *Nucleic Acids Research*, 49(W1):W72–W79, 2021. doi: 10.1093/nar/gkab393. URL https://doi.org/10.1093/nar/gkab393.

Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual Information Neural Estimation. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 531–540, 2018. URL https://proceedings.mlr.press/v80/belghazi18a.html.

Ralph Allan Bradley and Milton E. Terry. Rank Analysis of Incomplete Block Designs: The Method of Paired Comparisons. *Biometrika*, 39:324–345, 1952. doi: 10.1093/biomet/39.3-4.324. URL https://doi.org/10.1093/biomet/39.3-4.324.

Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning - ICML '05*, pp. 89–96, 2005. doi: 10.1145/1102351.1102363. URL https://doi.org/10.1145/1102351.1102363.

Mario Catalano, Gerardo Pepe, Romina Appierdo, Gabriele Ausiello, Claire McWhite, Giorgio Gambosi, Manuela Helmer-Citterich, and Pier Federico Gherardini. Cross-Attention Over RNA And Protein Sequences Enables Generalizable Interaction Prediction. *bioRxiv*, 2026. doi: 10.64898/2026.04.22.720174. URL https://www.biorxiv.org/content/10.64898/2026.04.22.720174v2. Preprint. Version 2, July 31, 2026.

Long Chen, Haoquan Liu, and Yunjie Zhao. ZHMolTopoRPI: A Commutative Algebra-Driven Deep Learning Framework for Robust RNA-Protein Interaction Prediction. *Journal of Chemical Information and Modeling*, 66(13):7772–7788, 2026. doi: 10.1021/acs.jcim.6c01199. URL https://doi.org/10.1021/acs.jcim.6c01199.

Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debiased contrastive learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 8765–8775. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/63c3ddcc7b23daa1e42dc41f9a44a873-Paper.pdf.

Charles Elkan and Keith Noto. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 213–220, 2008. doi: 10.1145/1401890.1401920. URL https://doi.org/10.1145/1401890.1401920.

Yifei Gao, Runhan Shi, Gufeng Yu, Yuyang Huang, and Yang Yang. ZeRPI: A graph neural network model for zero-shot prediction of RNA-protein interactions. *Methods*, 235:45–52, 2025. doi: 10.1016/j.ymeth.2025.01.014. URL https://doi.org/10.1016/j.ymeth.2025.01.014.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2:665–673, 2020. doi: 10.1038/s42256-020-00257-z. URL https://doi.org/10.1038/s42256-020-00257-z.

Jiahui Guan, Lantian Yao, Peilin Xie, Zhihao Zhao, Dian Meng, Tzong-Yi Lee, Junwen Wang, and Ying-Chih Chiang. Graph-RPI: predicting RNA–protein interactions via graph autoencoder and self-supervised learning strategies. *Briefings in Bioinformatics*, 26(3):bbaf292, 2025. doi: 10.1093/bib/bbaf292. URL https://doi.org/10.1093/bib/bbaf292.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/guo17a.html.

Rong Han, Xiaohong Liu, Tong Pan, Jing Xu, Xiaoyu Wang, Wuyang Lan, Zhenyu Li, Zixuan Wang, Jiangning Song, Guangyu Wang, and Ting Chen. CoPRA: Bridging Cross-domain Pretrained Sequence Models with Complex Structures for Protein-RNA Binding Affinity Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 246–254, 2025. doi: 10.1609/aaai.v39i1.32001. URL https://doi.org/10.1609/aaai.v39i1.32001.

Xiaoqian Huang, Yi Shi, Jing Yan, Wenyan Qu, Xiaoyi Li, and Jianjun Tan. LPI-CSFFR: Combining serial fusion with feature reuse for predicting lncrna-protein interactions. *Computational Biology and Chemistry*, 99:107718, 2022. doi: 10.1016/j.compbiolchem.2022.107718. URL https://doi.org/10.1016/j.compbiolchem.2022.107718.

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18661–18673. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/d89a66c7c80a29b1bdbab0f2a1a94af8-Paper.pdf.

Ryuichi Kiryo, Gang Niu, Marthinus C du Plessis, and Masashi Sugiyama. Positive-unlabeled learning with non-negative risk estimator. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/7cce53cf90577442771720a370c3c723-Paper.pdf.

Ying Li, Hang Sun, Wensi Fang, Qin Ma, Siyu Han, Rui Wang-Sattler, Wei Du, and Qiong Yu. SURE: Screening unlabeled samples for reliable negative samples based on reinforcement learning. *Information Sciences*, 629:299–312, 2023. doi: 10.1016/j.ins.2023.01.112. URL https://doi.org/10.1016/j.ins.2023.01.112.

Yunqing Lin, Tianyuan Liu, Tianyu Cui, Zhao Wang, Yuncong Zhang, Puwen Tan, Yan Huang, Jia Yu, and Dong Wang. RNAInter in 2020: RNA interactome repository with increased coverage and annotation. *Nucleic Acids Research*, 48:D189–D197, 2020. doi: 10.1093/nar/gkz804. URL https://doi.org/10.1093/nar/gkz804.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574. URL https://doi.org/10.1126/science.ade2574.

Haoquan Liu, Yiren Jian, Chen Zeng, and Yunjie Zhao. RNA-protein interaction prediction using network-guided deep learning. *Communications Biology*, 8(1):247, 2025a. doi: 10.1038/s42003-025-07694-9. URL https://doi.org/10.1038/s42003-025-07694-9.

Yiwei Liu, Ting Bao, Peng Yin, Shumin Wang, and Yanbin Wang. RPIPLM: Prediction of ncRNA-protein interaction by post-training a dual-tower pretrained biological model with supervised contrastive learning. *PLOS One*, 20(8):e0329174, 2025b. doi: 10.1371/journal.pone.0329174. URL https://doi.org/10.1371/journal.pone.0329174.

Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Usha K. Muppirala, Vasant G. Honavar, and Drena Dobbs. Predicting RNA-protein interactions using only sequence information. *BMC Bioinformatics*, 12(1):489, 2011. doi: 10.1186/1471-2105-12-489. URL https://doi.org/10.1186/1471-2105-12-489.

Xiaoyong Pan, Yong-Xian Fan, Junchi Yan, and Hong-Bin Shen. IPMiner: Hidden ncrna-protein interaction sequential pattern mining with stacked autoencoder for accurate computational prediction. *BMC Genomics*, 17(1):582, 2016. doi: 10.1186/s12864-016-2931-8. URL https://doi.org/10.1186/s12864-016-2931-8.

Yungki Park and Edward M Marcotte. Flaws in evaluation schemes for pair-input computational predictions. *Nature Methods*, 9:1134–1136, 2012. doi: 10.1038/nmeth.2259. URL https://doi.org/10.1038/nmeth.2259.

Liangchen Peng, Lijun Quan, Lingkun Meng, Zhihong Zhang, Shengju Zhang, Zhijun Zhang, Yi Zhang, Qiufeng Chen, Bei Zhang, Lexin Cao, Tingfang Wu, and Qiang Lyu. PaRPI predicts RNA-Protein interactions from cross-protocol and cross-batch RNA-binding protein datasets. *Communications Biology*, 8(1):1396, 2025. doi: 10.1038/s42003-025-08807-0. URL https://doi.org/10.1038/s42003-025-08807-0.

Rafael Josip Penić, Tin Vlašić, Roland G. Huber, Yue Wan, and Mile Šikić. RiNALMo: general-purpose RNA language models can generalize well on structure prediction tasks. *Nature Communications*, 16(1):5671, 2025. doi: 10.1038/s41467-025-60872-5. URL https://doi.org/10.1038/s41467-025-60872-5.

Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 452–461, 2009. URL https://www.auai.org/uai2009/proceedings.html.

Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive Learning with Hard Negative Samples. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=CR1XOQ0UTh-.

Jeffrey A. Ruffolo, Stephen Nayfach, Joseph Gallagher, Aadyot Bhatnagar, Joel Beazer, Riffat Hussain, Jordan Russ, Jennifer Yip, Emily Hill, Martin Pacesa, Alexander J. Meeske, Peter Cameron, and Ali Madani. Design of highly functional genome editors by modelling CRISPR–Cas sequences. *Nature*, 645(8080):518–525, 2025. doi: 10.1038/s41586-025-09298-z. URL https://doi.org/10.1038/s41586-025-09298-z.

V. Suresh, Liang Liu, Donald Adjeroh, and Xiaobo Zhou. RPI-Pred: Predicting ncrna-protein interaction using sequence and structural information. *Nucleic Acids Research*, 43(3):1370–1379, 2015. doi: 10.1093/nar/gkv020. URL https://doi.org/10.1093/nar/gkv020.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding, 2018. URL https://arxiv.org/abs/1807.03748.

Hai-Cheng Yi, Zhu-Hong You, Li Cheng, Xi Zhou, Tong-Hai Jiang, Xiao Li, and Yan-Bin Wang. Learning distributed representations of RNA and protein sequences and its application for predicting lncRNA-protein interactions. *Computational and Structural Biotechnology Journal*, 18: 20–26, 2020. doi: 10.1016/j.csbj.2019.11.004. URL https://doi.org/10.1016/j.csbj.2019.11.004.

Yu-Huai Yu, Han-Ting Hong, and Tzu-Hsien Yang. Generalizable deep-learning-based mRNA-protein interaction prediction strongly depends on protein diversity. *Journal of Cheminformatics*, 18(1):74, 2026. doi: 10.1186/s13321-026-01197-3. URL https://doi.org/10.1186/s13321-026-01197-3.

Jiao Yuan, Wei Wu, Chaoyong Xie, Guoguang Zhao, Yi Zhao, and Runsheng Chen. NPInter v2.0: an updated database of ncRNA interactions. *Nucleic Acids Research*, 42:D104–D108, 2014. doi: 10.1093/nar/gkt1057. URL https://doi.org/10.1093/nar/gkt1057.

Guoqing Zhao, Baolu Shi, Pengpai Li, and Zhi-Ping Liu. PRInter: an inductive multimodal heterogeneous network embedding framework with contrastive learning for lncRNA-protein interaction prediction. *BMC Bioinformatics*, 2026. doi: 10.1186/s12859-026-06571-4. URL https://doi.org/10.1186/s12859-026-06571-4. Accepted article published online July 30, 2026; final volume and article number not yet assigned in the publisher record.

Zhecheng Zhou, Zhenya Du, Jinhang Wei, Linlin Zhuo, Shiyao Pan, Xiangzheng Fu, and Xinze Lian. MHAM-NPI: Predicting ncRNA-protein interactions based on multi-head attention mechanism. *Computers in Biology and Medicine*, 163:107143, 2023. doi: 10.1016/j.compbiomed.2023.107143. URL https://doi.org/10.1016/j.compbiomed.2023.107143.

Linlin Zhuo, Bosheng Song, Yuansheng Liu, Zejun Li, and Xiangzheng Fu. Predicting ncRNA–protein interactions based on dual graph convolutional network and pairwise learning. *Briefings in Bioinformatics*, 23(6):bbac339, 2022. doi: 10.1093/bib/bbac339. URL https://doi.org/10.1093/bib/bbac339.

A Score Conventions and Ranking Properties

Interaction log-odds. With class 1 denoting interaction,

$$\frac{p(y=1 \mid x)}{p(y=0 \mid x)} = \frac{e^{\ell_1(x)}}{e^{\ell_0(x)}} = e^{s(x)}, \qquad s(x) = \ell_1(x) - \ell_0(x). \tag{9}$$

The score is unchanged if an arbitrary input-dependent constant is added to both raw logits. Ranking ℓ_1 alone would not have this invariance. With the reverse convention $q = \ell_0 - \ell_1$, the observed interaction must have a smaller q than its alternative, giving the equivalent loss $\mathrm{softplus}((q^+ - q^a)/\tau)$.

Global shift and classification. Adding a common constant to every score leaves every ranking margin unchanged. The ranking-only objective therefore cannot identify the intercept needed for a global classification threshold. Benchmark BCE supplies an absolute score constraint; calibration under a different class prior or sampling scheme remains a separate empirical question (Guo et al., 2017). Ranking-only accuracy consequently depends on the threshold-selection protocol, whereas its AUROC does not depend on the intercept.

The exchange diagnostic. Substituting $s(R, P) = a(R) + b(P)$ into Equation (8) cancels each molecule-specific term. A positive rescaling $s' = cs$, $c > 0$, gives $D' = cD$ and preserves its sign. An increasing nonlinear transformation, such as a sigmoid, need not preserve this additive assignment difference. We therefore document each predictor's score convention and compare mean and median D within model families. The strict preference rate $P(D > 0)$ uses zero numerical tie tolerance, with exact ties reported separately.

Pairwise and multi-candidate normalization. For an observed score s_i^+ and selected alternative s_{ia}, the Bradley–Terry preference model gives

$$-\log \frac{e^{s_i^+/\tau}}{e^{s_i^+/\tau} + e^{s_{ia}/\tau}} = \log\left(1 + e^{(s_{ia} - s_i^+)/\tau}\right). \tag{10}$$

For the K_i retained alternatives in $\mathcal{H}_i$, pairwise averaging and multi-candidate normalization are

$$L_i^{\mathrm{pair}} = \frac{1}{K_i} \sum_{a \in \mathcal{H}_i} \log\left(1 + e^{(s_{ia} - s_i^+)/\tau}\right), \qquad L_i^{\mathrm{multi}} = \log\left(1 + \sum_{a \in \mathcal{H}_i} e^{(s_{ia} - s_i^+)/\tau}\right). \tag{11}$$

They coincide for $K_i = 1$. Otherwise, multi-candidate normalization couples the alternatives, while pairwise averaging applies a separate preference loss to each. The implementation averages comparisons within each anchor before averaging anchors and directions. Neither algebraic form alone establishes robustness to mislabeled alternatives.

A high-scoring ordered alternative. At $\tau = 1$, scores $(s^+, s^a) = (5, 2)$ give a ranking loss of approximately 0.049, although both predicted interaction probabilities exceed 0.5. Labeling the alternative negative instead gives BCE loss 2.127. Their gradients with respect to the alternative score are

$$\frac{\partial \ell_{\mathrm{neg}}}{\partial s^a} = \sigma(s^a), \qquad \frac{\partial \ell_{\mathrm{rank}}}{\partial s^a} = \frac{1}{\tau}\sigma\left(\frac{s^a - s^+}{\tau}\right). \tag{12}$$

The ranking gradient decreases as the observed pair's advantage grows. When the ordering is violated, a small temperature can instead make this gradient larger than that of negative BCE. The distinction is a relative constraint, rather than a uniformly weaker update. This numerical example illustrates the objectives and is not an experimental result.

B Training Setup and Objective Schedules

Model and training configuration. We use frozen RNA and protein language-model representations throughout training. NPInter2.0 uses RiNALMo-giga and ESM2-t36, whereas CW30_50 and CW60_80 use RiNALMo-mega and ESM2-t36. The pooled molecular representations are projected to a shared dimension $d = 1280$, followed by the interaction module defined in Section 3.1.

Unless otherwise stated, all ReMatch experiments are trained for 40 epochs using AdamW (Loshchilov & Hutter, 2019) with learning rate 1.5×10^{-4} and weight decay 0.02. Training uses a batch size of 128 per GPU under four-GPU distributed training. The primary ranking configuration uses $k = 40$, temperature $\tau = 0.25$, and an initial ranking coefficient $\lambda_{\text{rank}} = 1$. The same model architecture and optimization settings are used for the matched objective controls unless explicitly stated otherwise.

Ranking-weight schedule. Epochs are indexed from 0 to 39. For the primary ReMatch objective, the ranking term is disabled during epochs 0–1, activated with unit weight from epoch 2, and kept at full strength through epoch 15. It is then linearly annealed over the remaining training epochs:

$$\lambda_{\text{rank}}(e) = \begin{cases} 0, & 0 \le e < 2, \\ 1, & 2 \le e \le 15, \\ 1 - \dfrac{e-15}{25}, & 15 < e < 40. \end{cases} \tag{13}$$

Thus, $\lambda_{\text{rank}}(15) = 1$ and $\lambda_{\text{rank}}(39) = 0.04$. The ranking-only control uses the same two-epoch warm-up but does not apply the late-stage annealing. The BCE-H auxiliary classification term, when enabled, has fixed unit weight throughout training.

Model selection. Within each fold, the checkpoint with the highest validation accuracy is retained. Accuracy is evaluated using the fixed interaction boundary $s = 0$, equivalently $p(y = 1) = 0.5$. Reported ReMatch results are averaged over five folds.

C Data Sources and Split Audit

The NPInter2.0 RPI subset contains 10,412 experimentally supported positive pairs and 10,412 negatives obtained by random re-pairing with known positives excluded (Yuan et al., 2014; Pan et al., 2016). It is narrower than the complete NPInter database, which also records other ncRNA associations. CORAL pools strongly supported RNAInter associations (Lin et al., 2020) with positives from RPI488 (Pan et al., 2016), RPI369 and RPI2241 (Muppirala et al., 2011), and RPI1807 (Suresh et al., 2015), yielding 24,012 positives and an equal number of re-paired negatives before fold-specific filtering. These counts describe the source pool rather than individual training or evaluation folds.

RPI1460 is a derived collection constructed by LPI-CSFFR from RPI488, RPI2241, and RPI1807 after retaining RNAs longer than 200 nucleotides and removing duplicate interactions; it is not an additional independent source in the CORAL pool described above (Huang et al., 2022). RPI488 contains 243 positives and 245 negatives (Pan et al., 2016). The CW files are the released CORAL datasets obtained from the original CORAL repository (Apache-2.0) and its Zenodo release (concept DOI 10.5281/zenodo.19730809). The CORAL positive pool retains only RNAInter interactions marked by strong experimental evidence together with positive interactions from RPI488, RPI369, RPI2241, and RPI1807; proteins longer than 1,250 residues and RNAs longer than 6,000 nucleotides are removed, yielding 24,012 positive pairs. We use the released sequences as provided (uppercase; no additional U/T conversion). The exact upstream deduplication and conflicting-label resolution rules were not reimplemented in our pipeline. The NPInter2.0 benchmark uses the NPInter10412 release with 10,412 positives and an equal number of re-paired negatives (Yuan et al., 2014; Pan et al., 2016). An overlap audit found 120 of the 10,412 NPInter2.0 positives (1.15%) contained in the CORAL positive pool under the strict same-protein plus RNA $\geq 97\%$ identity criterion; the corresponding upper-bound match count was 188 (1.81%). A further 401 NPInter2.0 positives (3.85%) involve proteins exceeding CORAL's 1,250-residue filter. For source-specific licensing, the CORAL repository is distributed under Apache-2.0, while the experiment records do not contain independently verified redistribution-license metadata for NPInter2.0, RNAInter, or the legacy RPI subsets included through the CORAL release.

Split verification. The componentwise labels CW30_50 and CW60_80 correspond to the released 30/50 and 60/80 protein/RNA similarity controls; in CORAL's RNA/protein ordering these are 50/30 and 80/60, respectively. CORAL clusters RNA and protein sequences with the MMseqs2 `cluster` command, using protein identity thresholds of 30% and 60%, RNA identity thresholds of 50% and 80%, and a coverage parameter of 60%; the remaining MMseqs2 parameters are left at their defaults. For componentwise splitting, 20% of RNA clusters and 20% of protein clusters are independently

Table 6: Dataset source inventory. Counts precede pooling and fold-specific filtering. CORAL constructs its own re-paired negatives; the constituent datasets' native negative definitions are listed for provenance.

Collection	Positive	Negative	Label construction
NPInter2.0 (Yuan et al., 2014; Pan et al., 2016)	10,412	10,412	Experimental positives; random re-pairing excluding known positives.
CORAL pool	24,012	24,012	RNAInter and four RPI positive subsets; re-pairing within splits.
RPI1460 (Huang et al., 2022)	1,460	1,460	Pooled from RPI488, RPI2241, and RPI1807 after RNA-length and duplicate filtering; matched negatives from the released benchmark.
RPI369 (Muppirala et al., 2011)	369	369	Structural positives from PRIDB; matched negatives by random re-pairing.
RPI2241 (Muppirala et al., 2011)	2,241	2,241	Structural positives from PRIDB; matched negatives by random re-pairing.
RPI488 (Pan et al., 2016)	243	245	Structure-based labels using a 5 Å contact criterion after sequence-redundancy filtering.
RPI1807 (Suresh et al., 2015)	1,807	1,436	Structural interactions and weak/non-contact pairs using a 3.400 Å contact criterion.

assigned to the held-out partition and the remaining 80% to training. Pairs whose two components receive mixed train/held-out assignments are excluded. CORAL generates 100,000 random partitions and retains five splits satisfying at least 12,000 training positives while maximizing the number of held-out interactions. Negatives are generated independently within the training and held-out subsets by random re-pairing, excluding known positives, at a 1:1 class ratio. We use these released folds without reconstructing the clustering. Our fold audit confirms zero exact RNA-identity and zero exact protein-identity overlap between train and held-out data in every CW30_50 and CW60_80 fold. We did not independently recompute maximum cross-boundary sequence identities; the homology thresholds are those of the released CORAL splits. This audit concerns supervised fold overlap; exposure during encoder pretraining requires a separate provenance check.

D Complete Objective and Sampling Ablations

All objective controls follow the configuration in Appendix B and use the same five folds. G-mean is the geometric mean of sensitivity and specificity. Table 8 reports sample standard deviations across folds; Table 9 compares ReMatch with BCE within each fold. Checkpoints are selected by validation accuracy on the reported fold, and rounding is applied only for display.

Matching the two objectives within each fold removes differences in the folds' overall difficulty from the paired comparison. The intervals summarize variation over these five predefined splits, rather than uncertainty across independently repeated experiments.

Table 7: Complete objective and sampling ablation under the unified DDP protocol. Entries are five-fold means. All variants use four-GPU DDP with a batch size of 128 per GPU and the matched training configuration described in Appendix B. ReMatch combines benchmark BCE with pairwise ranking; ReMatch + MI additionally enables the optional representation regularizer. Accuracy and AUROC are shown at the same precision as in Table 2, and G-mean denotes the geometric mean of sensitivity and specificity.

Dataset	Objective	ACC	Macro-F1	G-mean	AUROC	MCC
NPInter2.0	BCE	0.955	0.955	0.955	0.982	0.910
NPInter2.0	BCE-H	0.931	0.931	0.931	0.974	0.863
NPInter2.0	Rank	0.902	0.902	0.901	0.961	0.805
NPInter2.0	BCE-H + Rank	0.942	0.942	0.942	0.974	0.885
NPInter2.0	ReMatch	0.954	0.954	0.954	0.986	0.909
NPInter2.0	ReMatch + MI	**0.960**	**0.960**	**0.960**	**0.987**	**0.921**
CW30_50	BCE	0.723	0.719	0.714	0.813	0.457
CW30_50	BCE-H	0.736	0.731	0.723	0.840	0.488
CW30_50	Rank	0.735	0.734	0.732	0.805	0.475
CW30_50	BCE-H + Rank	0.766	0.765	0.764	0.834	0.534
CW30_50	ReMatch	**0.770**	**0.769**	**0.768**	**0.844**	**0.542**
CW30_50	ReMatch + MI	0.766	0.764	0.761	0.838	**0.542**
CW60_80	BCE	0.703	0.696	0.687	0.791	0.421
CW60_80	BCE-H	0.701	0.690	0.675	0.812	0.432
CW60_80	Rank	0.722	0.719	0.716	0.802	0.451
CW60_80	BCE-H + Rank	0.726	0.723	0.718	**0.813**	**0.462**
CW60_80	ReMatch	**0.729**	**0.728**	**0.727**	0.799	0.458
CW60_80	ReMatch + MI	0.725	0.724	0.723	0.801	0.452

Table 8: **Fold-level dispersion of the objective ablations under the unified DDP protocol.** Entries are sample standard deviations across the five matched folds. The values describe variation across predefined data folds rather than repeated random-seed uncertainty; lower dispersion is therefore not interpreted as better predictive performance.

Dataset	Objective	ACC	Macro-F1	G-mean	AUROC	MCC
NPInter2.0	BCE	0.001	0.001	0.001	0.002	0.002
	BCE-H	0.004	0.004	0.004	0.006	0.008
	BCE-H + Rank	0.006	0.006	0.006	0.004	0.013
	Rank	0.013	0.013	0.013	0.013	0.025
	ReMatch + MI	0.002	0.002	0.002	0.002	0.003
	ReMatch (BCE + Rank)	0.002	0.002	0.002	0.002	0.003
CW30_50	BCE	0.026	0.027	0.029	0.042	0.051
	BCE-H	0.047	0.050	0.055	0.042	0.085
	BCE-H + Rank	0.026	0.026	0.026	0.029	0.054
	Rank	0.045	0.045	0.045	0.050	0.090
	ReMatch + MI	0.028	0.027	0.023	0.040	0.070
	ReMatch (BCE + Rank)	0.030	0.030	0.029	0.040	0.061
CW60_80	BCE	0.065	0.073	0.081	0.064	0.124
	BCE-H	0.050	0.057	0.064	0.058	0.096
	BCE-H + Rank	0.057	0.060	0.063	0.053	0.110
	Rank	0.052	0.054	0.055	0.045	0.104
	ReMatch + MI	0.064	0.064	0.064	0.060	0.127
	ReMatch (BCE + Rank)	0.057	0.058	0.058	0.066	0.114

Table 9: **Paired five-fold differences between ReMatch and BCE.** Differences are computed within matched folds. Entries report $100\times$ the mean difference $\pm$ the half-width of a two-sided 95% Student-t interval; divide by 100 to recover the metric scale. This scale preserves the small NPInter differences at three decimal places. Intervals describe variation across the five predefined folds, rather than independent test replications.

Dataset	ΔACC	ΔMacro-F1	ΔG-mean	ΔAUROC	ΔMCC
NPInter2.0	-0.050 ± 0.160	-0.040 ± 0.150	-0.050 ± 0.160	$+0.400 \pm 0.110$	-0.080 ± 0.320
CW30_50	$+4.690 \pm 5.270$	$+5.000 \pm 5.360$	$+5.420 \pm 5.520$	$+3.070 \pm 6.130$	$+8.480 \pm 10.470$
CW60_80	$+2.500 \pm 5.380$	$+3.180 \pm 6.080$	$+3.980 \pm 6.960$	$+0.820 \pm 5.580$	$+3.730 \pm 9.610$

E Pairing Diagnostic and Cross-Model Controls

We evaluate all predictors using a shared quartet-evaluation function. For two observed interactions (R_i, P_i) and (R_j, P_j), the assignment gap is

$$D_{ij} = s(R_i, P_i) + s(R_j, P_j) - s(R_i, P_j) - s(R_j, P_i).$$

Within each benchmark, every predictor is evaluated on the same set of valid quartets, with crossed assignments recorded as known positives excluded. This yields 9,822,511 valid quartets for NPInter2.0 and 7,662,966 for CW30_50.

The common evaluation function reports the mean and median assignment gap, the strict observed-assignment preference rate $P(D > 0)$, and the number of exact ties. Base and +Rank variants within each model family use the same scalar interaction-score definition. PRInter is reported separately because its ranking variant changes the prediction head in addition to the training objective. Because quartets reuse molecules and interactions extensively, these counts should not be interpreted as numbers of independent experimental replicates.

Table 10: **Pairing discrimination on shared quartets.** Mean and median assignment gap D, strict observed-assignment preference rate $P(D > 0)$, and the number of exact ties $D = 0$. All predictors within each benchmark are evaluated using the same quartet-evaluation function and the same retained quartets.

	NPInter2.0				CW30_50			
Method	mean D	median D	$P(D > 0)$	$D = 0$	mean D	median D	$P(D > 0)$	$D = 0$
ZHMolTopoRPI	7.471	5.100	0.829	0	0.473	0.159	0.687	3
ZHMolTopoRPI + Rank	4.549	4.065	<u>0.866</u>	1	1.187	0.321	**0.723**	5
ZeRPI	0.560	0.047	0.776	62	0.254	0.090	0.679	2
ZeRPI + Rank	0.911	0.977	**0.872**	10	0.334	0.121	0.710	8
BioLLMNet	19.839	0.748	0.727	4,427	10.304	1.740	0.701	10,497
BioLLMNet + Rank	24.424	0.998	0.811	76,114	7.241	1.247	0.714	64,647
PRInter (SVM)	4.300	4.146	0.845	0	0.113	0.064	0.515	5
PRInter (SVM) + Rank	0.000	0.000	0.460	1,030,901	0.000	0.000	0.484	227,322
BCE	14.799	1.981	0.764	2	4.537	1.458	0.708	1
ReMatch	19.098	5.140	0.854	0	2.667	0.594	<u>0.721</u>	0

Ranking increases assignment preference for both matched neural predictors on both benchmarks. ZHMolTopoRPI rises from 0.8292 to 0.8655 on NPInter2.0 and from 0.6872 to 0.7227 on CW30_50; BioLLMNet rises from 0.7271 to 0.8110 and from 0.7013 to 0.7138, respectively. ZeRPI follows the same trend, improving from 0.776 to 0.872 and from 0.679 to 0.710.

These gains concern how often the observed assignment is preferred, not just how widely the scores are separated. ZHMolTopoRPI on NPInter2.0 and BioLLMNet on CW30_50 both improve $P(D > 0)$ while their mean and median gaps decrease. BioLLMNet on NPInter2.0 improves all three statistics. The preference rate thus captures an improvement that average gap size alone would miss.

Exact ties are rare for ZHMolTopoRPI, with 0–5 across the pooled quartets. BioLLMNet has $76{,}114$ ties on NPInter2.0 and $64{,}647$ on CW30_50 after adding ranking, still below 1% on either benchmark. The strict $D > 0$ criterion counts none of these ties as successful preferences.

The ReMatch head shows the same distinction between preference frequency and gap size. Adding ranking to BCE increases $P(D > 0)$ from 0.7638 to 0.8535 on NPInter2.0 and from 0.7084 to

0.7206 on CW30_50. Mean and median gaps rise from $14.7991/1.9809$ to $19.0980/5.1400$ on NPInter2.0, but fall from $4.5366/1.4585$ to $2.6672/0.5940$ on CW30_50. Across both settings, the model chooses the observed assignment more often.

F Hidden-Positive Analyses

F.1 Construction and Objective Controls

We designate a subset of training-side positive interactions as hidden positives. These interactions are removed from ordinary positive supervision and from the known-positive mask, while the validation data remain unchanged. The experiment compares BCE-H, Rank, BCE-H + Rank, and ReMatch under matched hidden-positive conditions. All configurations use $k = 40$, $\tau = 0.25$, and the corresponding objective schedules defined in the main text.

At zero injection, the designated interactions remain hidden but are not actively introduced into the comparison set, providing a masking control.

Hub-H construction. For each fold, the designated hidden set H contains 10% of the training positives. Candidate interactions are restricted to proteins with degree at least 15, and the highest-degree interactions are retained. Members of H are excluded from ordinary positive supervision and from known-positive protection throughout the experiment.

We vary the number of active hidden interactions $h \in \{16, 32, 48, 64\}$ and the nominal injection cap ρ. These settings control how frequently designated hidden positives appear among the alternative comparisons while leaving the benchmark-supervised training set unchanged.

Because the number of available designated comparisons varies across folds, we report the realized contamination rate

$$\rho_{\text{real}} = \frac{N_{\text{inj}}}{2kn}, \tag{14}$$

where N_{inj} is the number of designated hidden-positive comparisons, $k = 40$, and $n = 64$ is the number of supervised positive anchors. The realized rates, rather than the nominal caps, are used in the reported analyses.

With $h = 16$, the number of available injectable cells saturates even when the nominal cap is increased. Higher contamination levels are therefore obtained mainly by increasing h. For example, $h = 32, \rho = 0.3$ yields 11.39% realized contamination, whereas $h = 32, \rho = 1.0$ and $h = 48, \rho = 1.0$ yield 21.72% and 33.07%, respectively.

Injection schedule and matched conditions. All four objectives share the same designated hidden set H and the same injection schedule, so the realized rate is identical across them; the reported series reaches 33.07% realized contamination at $h = 48$ and $\rho = 1.0$.

Table 11: **Fold-level realized Hub-H contamination on CW30_50.** Counts report the mean number of designated injected comparison cells per step, with percentages computed relative to $2kn = 5120$ selected comparison slots, i.e. $\rho_{\text{real}} = N_{\text{inj}}/(2kn)$ with $k = 40$ and $n = 64$. All four objectives share the designated hidden set and injection schedule and therefore have identical realized rates; the final column reports the five-fold mean realized ρ used in Tables 5 and 12.

Setting	43483	460	4708	55152	5608	Mean ρ (%)
$h = 16, \rho = 0.3$	452.800 (8.840%)	266.000 (5.200%)	266.000 (5.200%)	198.000 (3.870%)	291.700 (5.700%)	5.760
$h = 32, \rho = 0.3$	684.000 (13.360%)	603.000 (11.780%)	398.000 (7.770%)	552.000 (10.780%)	678.000 (13.240%)	11.390
$h = 32, \rho = 1.0$	1369.000 (26.740%)	1081.500 (21.120%)	1025.000 (20.020%)	1092.000 (21.330%)	994.000 (19.410%)	21.720
$h = 48, \rho = 1.0$	1928.000 (37.660%)	1777.500 (34.720%)	1322.000 (25.820%)	1812.000 (35.390%)	1626.000 (31.760%)	33.070
$h = 64, \rho = 1.0$	2237.500 (43.700%)	1821.000 (35.570%)	1322.000 (25.820%)	1812.000 (35.390%)	1784.000 (34.840%)	35.060

F.2 Score Retention and Predictive Performance

Table 12 expands the main-text summary in Table 5 with mean score changes, checkpoint medians, and training-time margins. The designated hidden set is fixed across settings; only the per-step active

injection size h and nominal cap ρ vary. We report realized injection rates to distinguish the number of injected comparisons from the requested cap.

ReMatch retains more hidden positives above the interaction boundary than BCE-H + Rank at all four settings, while also producing higher medians and smaller downward mean score changes. At approximately 11% injection, the retained fractions are 0.270 and 0.108; at approximately 22%, they are 0.259 and 0.119. ReMatch also achieves the highest accuracy and AUROC among the four objectives at every representative level.

Rank alone produces the smallest mean score changes and medians closest to zero. These statistics do not imply better retention: its fraction above zero is below ReMatch's at every setting. Without benchmark BCE, its score offset is unconstrained by the ranking objective (Appendix A). Reading retention together with classification therefore gives a more informative account than comparing score shifts alone.

The main contrast keeps benchmark BCE and ranking in both conditions and changes whether the selected alternatives also receive negative labels. This supports the proposed distinction between relative and absolute supervision. The performance ordering is already present in the masking control, however, and the gains do not increase monotonically with the realized injection rate. Online mining also allows the ordinary candidate sets to diverge. The analysis consequently documents score-retention behavior under the stated training procedures without attributing all predictive gains to false-negative tolerance.

Table 12: **Complete hidden-positive score statistics and predictive performance.** Five-fold CW30_50 means for the Hub-H settings summarized in Table 5. ΔS_h is the mean score change from epoch 0 to the selected checkpoint e^*; the median and $P(S_h(e^*) > 0)$ describe scores at that checkpoint. The observed–hidden margin $s^+ - S_h$ is averaged over logged training steps and is unavailable at zero injection. Emphasis compares objectives within each setting in the indicated direction. Mean shifts and medians must be read alongside retention because ranking alone does not anchor the score offset.

Setting	Objective	Real. ρ (%)	$\Delta S_h \uparrow$	Med. $S_h(e^*) \uparrow$	$P(S_h(e^*) > 0) \uparrow$	$s^+ - S_h \downarrow$	ACC↑	AUROC↑
Hub-H, $h = 16$, $\rho = 0$	BCE-H	0.000	−1.859	−2.472	0.216	–	0.688	0.803
	Rank		−0.803	−0.711	0.139	–	0.716	0.794
	ReMatch		−2.356	−2.766	**0.316**	–	**0.749**	**0.836**
	BCE-H + Rank		−2.680	−3.305	0.180	–	0.706	0.820
Hub-H, $h = 32$, $\rho = 0.3$	BCE-H	11.390	−3.535	−4.476	0.112	15.860	0.643	0.790
	Rank		−1.518	−1.468	0.123	**6.190**	0.710	0.785
	ReMatch		−2.301	−2.789	**0.270**	17.920	**0.739**	**0.836**
	BCE-H + Rank		−4.202	−5.651	0.108	22.680	0.689	0.814
Hub-H, $h = 32$, $\rho = 1.0$	BCE-H	21.720	−3.944	−4.891	0.129	20.840	0.673	0.814
	Rank		−1.596	−1.574	0.116	**6.310**	0.709	0.782
	ReMatch		−1.895	−2.296	**0.259**	17.660	**0.737**	**0.833**
	BCE-H + Rank		−4.394	−6.019	0.119	27.500	0.699	0.824
Hub-H, $h = 48$, $\rho = 1.0$	BCE-H	33.070	−3.957	−4.914	0.121	20.520	0.666	0.803
	Rank		−2.094	−2.193	0.105	**5.940**	0.707	0.778
	ReMatch		−2.097	−2.606	**0.168**	16.630	**0.728**	**0.828**
	BCE-H + Rank		−3.611	−5.178	0.127	26.040	0.693	0.824

G SCORE, AGGREGATION, AND DIRECTIONAL CONTROLS

G.1 SCORE LOCATION AND AGGREGATION

We compare two locations for relative supervision. The embedding-level score is the projected-embedding cosine similarity $S_{\text{emb}}(R, P) = \cos(z_R, z_P)$, whereas the prediction-level score is the final interaction log-odds $S_{\text{logit}}(R, P) = \ell_1(R, P) - \ell_0(R, P)$. Each score is evaluated with either pairwise logistic aggregation or multi-candidate aggregation. All four configurations retain the benchmark BCE objective, use $k = 40$, both replacement directions, the same known-positive masking rule, unit ranking weight, and the matched DDP training framework.

Table 3 summarizes the comparison in the main text. Table 14 reports the complete metrics and fold-level dispersion, while Table 13 reports the temperature grid.

Temperature settings. The recorded grids are $\tau_{\text{emb}} \in \{0.05, 0.10, 0.20\}$ and $\tau_{\text{logit}} \in \{0.15, 0.25, 0.50\}$ (Table 13). The primary ReMatch configuration uses $\tau = 0.25$ for interaction log-odds; on CW30_50 this is the best log-odds grid point, exceeding $\tau = 0.50$ on all five metrics, while the NPInter2.0 log-odds grid is flat, with $\tau = 0.50$ marginally higher in ACC, Macro-F1, G-mean, and MCC but lower in AUROC. The embedding-similarity rows serve as a reference for the supervised score location rather than as a candidate configuration; the reported value $\tau = 0.10$ is kept fixed across the pairwise and multi-candidate variants so that both aggregation rules share one reference setting. On CW30_50 the three embedding-similarity grid points differ by at most 0.6 accuracy points and 1.2 AUROC points ($\tau = 0.05$ and $\tau = 0.20$ are marginally higher than $\tau = 0.10$, while on NPInter2.0 $\tau = 0.10$ is marginally highest), so the location-level conclusions do not depend on this choice. Because the cosine score and the interaction log-odds have different numerical scales and functional roles, this analysis compares practical supervision designs rather than isolating score location under an otherwise identical scalar parameterization.

Candidate construction. All variants use the same eligibility rule, known-positive masking rule, $k = 40$, and hard-alternative selection strategy. Candidate selection follows the score associated with the corresponding supervision location.

G.2 REPLACEMENT DIRECTION

Protein replacement compares alternatives while holding the RNA fixed, whereas RNA replacement holds the protein fixed. For a single-direction condition, the corresponding ranking loss receives unit auxiliary weight. The bidirectional condition averages the two directional losses, preserving the overall auxiliary coefficient rather than doubling it. These directional controls use the same log-odds/pairwise targeted-analysis configuration as the ReMatch reference in Appendix G.1.

The two directions work best together on CW30_50. Bidirectional ranking improves accuracy by approximately 1.2 percentage points over protein replacement alone and 4.2 points over RNA replacement alone, with the same ordering in AUROC (Table 15). NPInter2.0 instead favors RNA replacement alone by a small margin. As in the score-location analysis, the benefit of the full design is clearest under componentwise generalization.

Directional margins. Directional score differences are evaluated at the checkpoint selected by validation accuracy and summarized across the five folds. These margins characterize the two replacement directions and are distinct from the held-out quartet diagnostic in Appendix E.

The margins reveal what the second comparison adds. With RNA replacement alone, the CW30_50 positive-margin rate reaches 0.850 in the trained direction but only 0.677 under protein replacement. Bidirectional training brings the two rates to 0.860 and 0.858 (Table 16). Protein-only training produces larger mean margins yet lower predictive accuracy. ReMatch learns a more balanced preference across both replacement types, accompanying its predictive gain.

Table 13: **Temperature grid for score-level and embedding-level supervision.** Five-fold means for pairwise aggregation. $\tau = 0.250$ is the interaction log-odds temperature of the primary ReMatch configuration; $\tau = 0.100$ is the reference temperature used for the embedding-similarity rows of Tables 3 and 14, shared across aggregation rules. NPInter2.0 and CW60_80 results are reported for completeness.

NPInter2.0

Score	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{logit}	0.150	0.953	0.953	0.952	0.985	0.906
S_{logit}	0.250	0.954	0.954	0.954	**0.986**	0.909
S_{logit}	0.500	**0.956**	**0.956**	**0.955**	0.985	**0.911**
S_{emb}	0.050	0.954	0.954	0.954	**0.986**	0.909
S_{emb}	0.100	**0.955**	**0.955**	**0.955**	**0.986**	**0.911**
S_{emb}	0.200	**0.955**	**0.955**	**0.955**	0.985	0.910

CW30_50

Score	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{logit}	0.150	0.761	0.761	0.760	0.828	0.525
S_{logit}	0.250	**0.770**	**0.769**	**0.768**	**0.844**	**0.542**
S_{logit}	0.500	0.761	0.761	0.760	0.835	0.524
S_{emb}	0.050	**0.717**	0.713	0.707	0.813	**0.445**
S_{emb}	0.100	0.711	0.708	0.703	0.802	0.430
S_{emb}	0.200	**0.717**	**0.714**	**0.709**	**0.814**	0.443

CW60_80

Score	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{logit}	0.150	0.723	0.722	0.720	**0.799**	0.451
S_{logit}	0.250	**0.728**	**0.728**	**0.727**	**0.799**	**0.458**
S_{logit}	0.500	0.724	0.722	0.718	0.798	0.454
S_{emb}	0.050	**0.698**	**0.695**	**0.690**	**0.769**	0.406
S_{emb}	0.100	0.693	0.689	0.683	0.766	0.398
S_{emb}	0.200	0.697	0.690	0.681	**0.769**	**0.409**

Table 14: **Score-location and aggregation controls.** Values are five-fold mean ± sample standard deviation. S_{emb} denotes projected-embedding cosine similarity and S_{logit} denotes final interaction log-odds. All variants retain benchmark BCE, use $k = 40$, both replacement directions, and the same matched training setup. Cosine and log-odds supervision use $\tau = 0.10$ and 0.25, respectively.

NPInter2.0

Score	Aggregation	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{emb}	Pairwise	0.100	0.955 ± 0.002	0.955 ± 0.002	0.955 ± 0.002	0.986 ± 0.001	0.911 ± 0.004
S_{emb}	Multi-candidate	0.100	$\mathbf{0.960} \pm 0.002$	$\mathbf{0.960} \pm 0.002$	$\mathbf{0.960} \pm 0.002$	$\mathbf{0.989} \pm 0.002$	$\mathbf{0.920} \pm 0.005$
S_{logit}	Pairwise	0.250	0.954 ± 0.002	0.954 ± 0.002	0.954 ± 0.002	0.986 ± 0.002	0.909 ± 0.003
S_{logit}	Multi-candidate	0.250	$\underline{0.959} \pm 0.002$	$\underline{0.959} \pm 0.002$	$\underline{0.959} \pm 0.002$	$\underline{0.988} \pm 0.002$	$\underline{0.918} \pm 0.004$

CW30_50

Score	Aggregation	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{emb}	Pairwise	0.100	0.711 ± 0.023	0.708 ± 0.024	0.703 ± 0.026	0.802 ± 0.022	0.430 ± 0.043
S_{emb}	Multi-candidate	0.100	0.717 ± 0.024	0.712 ± 0.027	0.707 ± 0.030	0.816 ± 0.023	0.446 ± 0.042
S_{logit}	Pairwise	0.250	$\mathbf{0.770} \pm 0.030$	$\mathbf{0.769} \pm 0.030$	$\mathbf{0.768} \pm 0.029$	$\mathbf{0.844} \pm 0.040$	$\mathbf{0.542} \pm 0.061$
S_{logit}	Multi-candidate	0.250	$\underline{0.766} \pm 0.035$	$\underline{0.765} \pm 0.035$	$\underline{0.763} \pm 0.034$	$\underline{0.834} \pm 0.044$	$\underline{0.538} \pm 0.073$

CW60_80

Score	Aggregation	τ	ACC	Macro-F1	G-mean	AUROC	MCC
S_{emb}	Pairwise	0.100	0.694 ± 0.057	0.689 ± 0.059	0.683 ± 0.062	0.766 ± 0.075	0.398 ± 0.114
S_{emb}	Multi-candidate	0.100	0.704 ± 0.055	0.699 ± 0.058	0.694 ± 0.060	0.772 ± 0.077	0.419 ± 0.110
S_{logit}	Pairwise	0.250	$\underline{0.729} \pm 0.057$	$\underline{0.728} \pm 0.058$	$\underline{0.727} \pm 0.058$	$\underline{0.799} \pm 0.066$	$\underline{0.458} \pm 0.114$
S_{logit}	Multi-candidate	0.250	$\mathbf{0.741} \pm 0.057$	$\mathbf{0.741} \pm 0.057$	$\mathbf{0.740} \pm 0.057$	$\mathbf{0.815} \pm 0.060$	$\mathbf{0.484} \pm 0.114$

Table 15: **Ablation of RNA and protein replacement directions.** Five-fold validation means with the overall ranking coefficient fixed. The bidirectional condition is the primary ReMatch configuration. Table 16 reports complete predictive metrics and directional-margin statistics.

	NPInter2.0		CW30_50	
Replacement direction	ACC	AUROC	ACC	AUROC
Protein only	0.950	0.982	<u>0.758</u>	<u>0.837</u>
RNA only	**0.956**	**0.987**	0.728	0.825
Both (ReMatch)	<u>0.954</u>	<u>0.986</u>	**0.770**	**0.844**

Table 16: **Directional ablation of RNA- and protein-side replacement.** Predictive metrics are five-fold mean ± sample standard deviation. Panel A reports predictive performance for protein-only, RNA-only, and bidirectional ranking. Panel B reports the corresponding directional-margin statistics at the validation-selected checkpoint.

Panel A: Predictive performance

Dataset	Direction	ACC	AUROC	Macro-F1	G-mean	MCC
NPInter2.0	Protein replacement only	0.950 ± 0.002	0.982 ± 0.002	0.950 ± 0.002	0.950 ± 0.002	0.902 ± 0.003
NPInter2.0	RNA replacement only	$\mathbf{0.956} \pm 0.001$	$\mathbf{0.987} \pm 0.002$	$\mathbf{0.956} \pm 0.001$	$\mathbf{0.956} \pm 0.001$	$\mathbf{0.912} \pm 0.003$
NPInter2.0	Bidirectional	$\underline{0.954} \pm 0.002$	$\underline{0.986} \pm 0.002$	$\underline{0.954} \pm 0.002$	$\underline{0.954} \pm 0.002$	$\underline{0.909} \pm 0.003$
CW30_50	Protein replacement only	$\underline{0.758} \pm 0.030$	$\underline{0.837} \pm 0.038$	$\underline{0.758} \pm 0.030$	$\underline{0.757} \pm 0.030$	$\underline{0.517} \pm 0.059$
CW30_50	RNA replacement only	0.728 ± 0.020	0.825 ± 0.014	0.725 ± 0.023	0.720 ± 0.028	0.467 ± 0.029
CW30_50	Bidirectional	$\mathbf{0.770} \pm 0.030$	$\mathbf{0.844} \pm 0.040$	$\mathbf{0.769} \pm 0.030$	$\mathbf{0.768} \pm 0.029$	$\mathbf{0.542} \pm 0.061$

Panel B: Directional training-construction margins at the selected checkpoint

Dataset	Direction	Mean Δ^P	Med. Δ^P	$P(\Delta^P > 0)$	Mean Δ^R	Med. Δ^R	$P(\Delta^R > 0)$
NPInter2.0	Protein replacement only	1.596 ± 0.108	0.208 ± 0.015	$\underline{0.698} \pm 0.013$	1.601 ± 0.105	0.537 ± 0.050	0.671 ± 0.008
NPInter2.0	RNA replacement only	$\mathbf{2.243} \pm 0.500$	$\mathbf{1.339} \pm 0.208$	0.655 ± 0.085	$\mathbf{2.218} \pm 0.514$	$\mathbf{0.972} \pm 0.302$	$\mathbf{0.792} \pm 0.010$
NPInter2.0	Bidirectional	$\underline{1.967} \pm 0.256$	$\underline{0.322} \pm 0.065$	$\mathbf{0.723} \pm 0.023$	$\underline{1.973} \pm 0.257$	$\underline{0.578} \pm 0.152$	$\underline{0.699} \pm 0.006$
CW30_50	Protein replacement only	$\mathbf{2.955} \pm 1.020$	$\mathbf{1.376} \pm 0.567$	$\mathbf{0.866} \pm 0.031$	$\mathbf{2.950} \pm 1.023$	$\mathbf{1.499} \pm 0.662$	0.824 ± 0.031
CW30_50	RNA replacement only	2.068 ± 1.306	1.002 ± 0.548	0.677 ± 0.081	2.025 ± 1.301	0.999 ± 0.537	$\underline{0.850} \pm 0.048$
CW30_50	Bidirectional	$\underline{2.694} \pm 0.872$	$\underline{1.333} \pm 0.417$	$\underline{0.858} \pm 0.028$	$\underline{2.689} \pm 0.873$	$\underline{1.339} \pm 0.468$	$\mathbf{0.860} \pm 0.036$

H Optional Representation Regularization

ReMatch was also evaluated with an optional representation-level compatibility regularizer. For projected RNA and protein representations z_{R_i} and z_{P_j}, define

$$T_{ij} = \tanh\left(\frac{z_{R_i}^{\mathsf{T}} z_{P_j}}{\sqrt{d}}\right), \tag{15}$$

and let Ω denote the set of eligible mismatched RNA–protein pairs excluding known training positives. The auxiliary regularizer is

$$\mathcal{L}_{\mathrm{MI}} = -\frac{1}{n}\sum_{i=1}^{n} T_{ii} + \log\left(\frac{1}{|\Omega|}\sum_{(i,j)\in\Omega} e^{T_{ij}}\right), \tag{16}$$

This objective has the form of a Donsker–Varadhan-style contrast between observed and mismatched molecular representations (Belghazi et al., 2018). We use it as an auxiliary compatibility regularizer rather than as a calibrated estimate of mutual information.

H.1 Matched Ablation

Matched ablation. We compare ReMatch with and without the auxiliary representation regularizer while keeping the interaction architecture, ranking objective, optimization settings, fold assignments, and model-selection procedure fixed. The two configurations differ only in whether $\mathcal{L}_{\mathrm{MI}}$ is included with unit weight.

Table 17: **Matched effect of the optional representation regularizer.** ReMatch + MI adds $\mathcal{L}_{\mathrm{MI}}$ with unit weight, while ReMatch omits the auxiliary regularizer. All remaining model and training settings are matched. Entries are five-fold means.

Dataset	Objective	ACC	Macro-F1	G-mean	MCC	AUROC
NPInter2.0	ReMatch + MI	**0.960**	**0.960**	**0.960**	**0.921**	**0.987**
NPInter2.0	ReMatch	0.954	0.954	0.954	0.909	0.986
CW30_50	ReMatch + MI	0.766	0.764	0.761	**0.542**	0.838
CW30_50	ReMatch	**0.770**	**0.769**	**0.768**	**0.542**	**0.844**
CW60_80	ReMatch + MI	0.725	0.724	0.723	0.452	**0.801**
CW60_80	ReMatch	**0.729**	**0.728**	**0.727**	**0.458**	0.799

The auxiliary regularizer improves the reported NPInter2.0 metrics but does not provide a consistent benefit on the two componentwise benchmarks. On CW30_50, removing the regularizer improves accuracy, Macro-F1, G-mean, and AUROC, while MCC is unchanged. On CW60_80, the MI-free model improves accuracy, Macro-F1, G-mean, and MCC. We therefore use BCE and ranking without the representation regularizer in the final ReMatch objective.

I Additional Results for the Original ReMatch + MI Configuration

This appendix retains supplementary analyses from the earlier ReMatch + MI experiment series. These experiments precede the final matched ReMatch configuration and are reported as complementary reproduction and sensitivity analyses rather than as matched comparisons with the final model. The contribution of the MI regularizer under the current training configuration is evaluated separately in Appendix H.

I.1 Complete Reproduction Metrics

The original ReMatch + MI model leads the reproduced baselines on both componentwise benchmarks, reaching 0.7855 accuracy on CW30_50 and 0.7398 on CW60_80. ZHMolTopoRPI leads on

NPInter2.0 accuracy and AUROC. Table 18 reports all five metrics, while Table 1 places the current ReMatch model alongside these baseline reproductions.

Table 18: **Complete reproduction comparison for the original ReMatch + MI configuration.** Entries consolidate the supplied baseline reproduction records. A dash denotes an unavailable metric. These results belong to the original reproduction configuration and are reported separately from the final matched DDP objective controls.

Dataset	Method	ACC	Macro-F1	G-mean	AUROC	MCC
CW30_50	**ReMatch + MI**	**0.786**	**0.785**	**0.785**	**0.857**	**0.572**
	CORAL	0.646	0.636	0.624	0.693	0.309
	ZeRPI	0.657	0.632	0.606	0.799	0.349
	PRInter	0.500	0.336	0.049	0.542	-0.005
	MHAM-NPI	0.634	0.608	0.582	0.605	0.304
	ZHMolGraph	0.659	0.620	0.547	0.709	0.328
	ZHMolTopoRPI	0.676	0.663	0.647	0.756	0.382
CW60_80	**ReMatch + MI**	**0.740**	**0.738**	**0.736**	**0.812**	**0.484**
	CORAL	0.684	–	–	0.757	–
	ZeRPI	0.663	0.638	0.609	0.785	0.369
	PRInter	0.500	0.336	0.054	0.568	0.004
	MHAM-NPI	0.638	0.624	0.609	0.654	0.298
	ZHMolGraph	0.679	0.669	0.657	0.726	0.383
	ZHMolTopoRPI	0.680	0.675	0.659	0.775	0.404
NPInter	ReMatch + MI	0.945	0.945	0.945	0.980	0.892
	CORAL	0.893	–	0.893	0.938	0.787
	ZeRPI	0.946	0.946	0.945	0.983	0.892
	PRInter	0.822	0.817	0.804	0.979	0.686
	MHAM-NPI	0.893	0.893	0.892	0.936	0.789
	ZHMolGraph	0.935	0.935	0.934	0.971	0.871
	ZHMolTopoRPI	**0.958**	**0.958**	**0.958**	**0.985**	**0.917**

The CW60_80 ZHMolGraph accuracy is 0.679, as reconciled with the original metric record.

I.2 Transfer Without Fine-Tuning

The Cas v2 evaluation applies the original NPInter-trained checkpoints without Cas-specific fine-tuning. We constructed this balanced transfer set from the Cas12 operon resource released with OpenCRISPR (Ruffolo et al., 2025), whose source collection contains 83,274 protein–RNA records and 29,074 unique Cas12 proteins. After filtering source sequences against CORAL at 60% protein and 80% RNA similarity, we retained 2,119 same-operon Cas12 protein–guide RNA pairs as positives. Negatives were formed by re-pairing these molecules while requiring each alternative protein to share less than 50% identity with the RNA's cognate protein and each alternative RNA to share less than 30% identity with the paired protein's cognate RNA. The resulting 4,238 examples are balanced, and every retained molecule appears in both classes. The record documents filtering against CORAL rather than the actual NPInter training folds, so Cas v2 serves as a transfer reproduction reference, not a fully audited leakage-free benchmark.

Table 19: **Original transfer evaluation on Cas v2.** Results use the original checkpoints, with ReMatch labeled ReMatch + MI. AUROC and AUPRC are reported without target-domain fine-tuning.

Method	AUROC	AUPRC
CORAL	0.499	0.505
ZeRPI	0.509	0.509
PRInter	0.511	0.508
MHAM-NPI	0.500	0.502
ZHMolGraph	0.531	0.527
ZHMolTopoRPI	0.523	0.522
ReMatch + MI	**0.543**	**0.529**

Transfer remains difficult for every model. ReMatch + MI reaches 0.5430 AUROC and 0.5288 AUPRC, compared with 0.5306 and 0.5274 for ZHMolGraph, the strongest baseline on both metrics. These modest gains leave absolute discrimination close to chance on this balanced set.

I.3 Feature Fusion

The fusion ablation examines how the two molecular representations enter the prediction head (Table 20). The gated mixture, elementwise product, and bilinear features give the highest CW30_50 scores. Adding a difference feature slightly improves NPInter2.0 accuracy but reduces it on CW30_50. The three retained operators already provide complementary ways to represent a pair; the additional channel offers no consistent gain across these benchmarks.

Table 20: **Fusion variants of the original ReMatch + MI configuration.** All metric values are retained. Different rows change the feature operators used before classification. These rows vary the feature operators within the original configuration and are not a re-evaluation of the final DDP ReMatch model.

Dataset	Fusion	ACC	Macro-F1	G-mean	AUROC	MCC
CW30_50	**Gated mix + product + bilinear**	**0.786**	**0.785**	**0.785**	**0.857**	**0.572**
	Concatenation + product + difference	0.763	0.762	0.762	0.843	0.526
	Gated mix + product + difference + bilinear	0.779	0.779	0.777	0.848	0.561
	Gated mix	0.778	0.778	0.777	0.851	0.558
NPInter	Gated mix + product + bilinear	0.945	0.945	0.945	0.980	0.892
	Concatenation + product + difference	0.943	0.943	0.943	**0.980**	0.888
	Gated mix + product + difference + bilinear	**0.946**	**0.946**	**0.946**	0.980	**0.893**
	Gated mix	0.938	0.938	0.937	0.978	0.877

I.4 Model Capacity and Encoder Sensitivity

Model capacity and encoder sensitivity. Table 21 summarizes capacity and encoder variants evaluated during the earlier ReMatch + MI study. These experiments are retained as sensitivity analyses rather than as the selection procedure for the final model. Rows corresponding to the final configurations used in Table 1 are aligned with the current matched ReMatch results; the remaining rows retain the archived sensitivity results.

Table 21: **Capacity and encoder sensitivity analyses.** Panel A varies the shared latent dimension and Panel B varies the encoder configuration. Rows corresponding to the final configurations used in Table 1 report the current matched ReMatch results; the remaining entries retain the archived sensitivity results from the earlier ReMatch + MI study.

Dataset	Configuration	ACC	Macro-F1	G-mean	AUROC	MCC
Panel A: Latent dimension						
CW30_50	d640	0.753	0.752	0.752	0.818	0.506
	d1280	0.770	0.769	0.768	0.844	0.542
	d2560	**0.773**	**0.773**	**0.772**	**0.847**	**0.548**
NPInter	d640	0.943	0.943	0.943	0.979	0.888
	d1280	**0.954**	**0.954**	**0.954**	**0.986**	**0.909**
	d2560	0.950	0.950	0.950	0.982	0.900
Panel B: FM/encoder size						
CW30_50	micro_t36	0.764	0.764	0.764	0.834	0.529
	mega_t12	0.733	0.732	0.731	0.800	0.467
	mega_t30	0.760	0.759	0.757	0.825	0.524
	mega_t33	0.731	0.729	0.726	0.803	0.467
	mega_t36	0.770	0.769	0.768	0.844	0.542
	mega_t48	0.764	0.763	0.762	0.837	0.530
	giga_t36	**0.790**	**0.789**	**0.788**	**0.861**	**0.583**
NPInter	micro_t36	0.945	0.945	0.945	0.978	0.891
	mega_t12	0.942	0.942	0.941	0.979	0.884
	mega_t30	0.943	0.943	0.942	0.979	0.886
	mega_t33	0.945	0.945	0.944	0.980	0.890
	mega_t36	0.945	0.945	0.945	0.980	0.892
	mega_t48	0.948	0.948	0.948	0.981	0.897
	giga_t36	**0.954**	**0.954**	**0.954**	**0.986**	**0.909**