COMPOSITIONAL REPRESENTATION LEARNING FOR RNA MODIFICATIONS VIA CHEMICAL TRANSFORMATION EQUIVARIANCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting nanopore responses for RNA modifications with limited or no response labels requires transferring knowledge across chemical types and sequence contexts. Recurring chemical transformations provide a basis for this sharing: they act across nucleoside substrates and compose along valid chemical paths. We formulate **Chemical Transformation Equivariance (CTE)**, under which valid chemical transformations are approximately preserved by corresponding actions on molecular representations. **RiboCTE** learns these actions in a shared latent subspace from precursor–product relations and valid paths, using frozen molecular embeddings as prediction targets. The same actions transform sequence-conditioned states, and a shared decoder predicts changes in nanopore substitution-plus-deletion rates relative to canonical controls. We introduce **RiboModBench** to evaluate cross-modification response prediction alongside chemical-action transfer and composition. RiboCTE achieves the highest Spearman correlation among the compared methods across modification/context holdouts and external evaluation, reaching 0.184 when both modifications and sequence contexts are held out and 0.268 on external modifications without source response labels. Shared chemical actions thus provide an effective inductive bias for RNA modification representation learning. Code and RiboModBench are available at `https://anonymous.4open.science/r/cte-7B11`.

1 INTRODUCTION

Chemical modifications expand RNA's chemical diversity and influence its folding, stability, and translation (Sordyl et al., 2026). Available modification annotations remain uneven across chemical types (Xuan et al., 2024). Nanopore sequencing provides a modification-sensitive readout through changes in basecalling behavior (Liu et al., 2019; Jenjaroenpun et al., 2021; Begik et al., 2021). We study differential response prediction: given a modification's chemistry and canonical local RNA sequence, predict the change in substitution-plus-deletion rate relative to a matched canonical control. We transfer chemical knowledge to modifications with limited or no response labels.

RNA language models and molecular encoders represent sequence context and supplied chemistry (Penić et al., 2025; Burns et al., 2025). Chemical structure has enabled current prediction for held-out nucleotide chemistries (Ding et al., 2021; Stephenson et al., 2025). Learning the changes linking molecular states offers a shared basis for transfer across modifications. The central question is how to represent these changes as reusable actions for context-dependent RNA response prediction.

RNA modification chemistry provides a concrete basis for sharing transformations. MODOMICS links nucleosides through precursor–product relations annotated with reaction types and chemical sites (Sordyl et al., 2026). For example, ribose 2′-O-methylation converts A, C, G, and U into Am, Cm, Gm, and Um (Figure 1A). These recurring edits provide a natural unit of sharing across substrates, while valid paths specify how actions compose.

We formulate **Chemical Transformation Equivariance (CTE)**: approximate consistency between valid chemical transformations and learned actions on molecular representations. Building on latent-action and action-conditioned predictive learning (Keurti et al., 2023; Assran et al., 2025), **RiboCTE** learns shared actions from precursor–product relations and valid paths. Chemical training predicts

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.

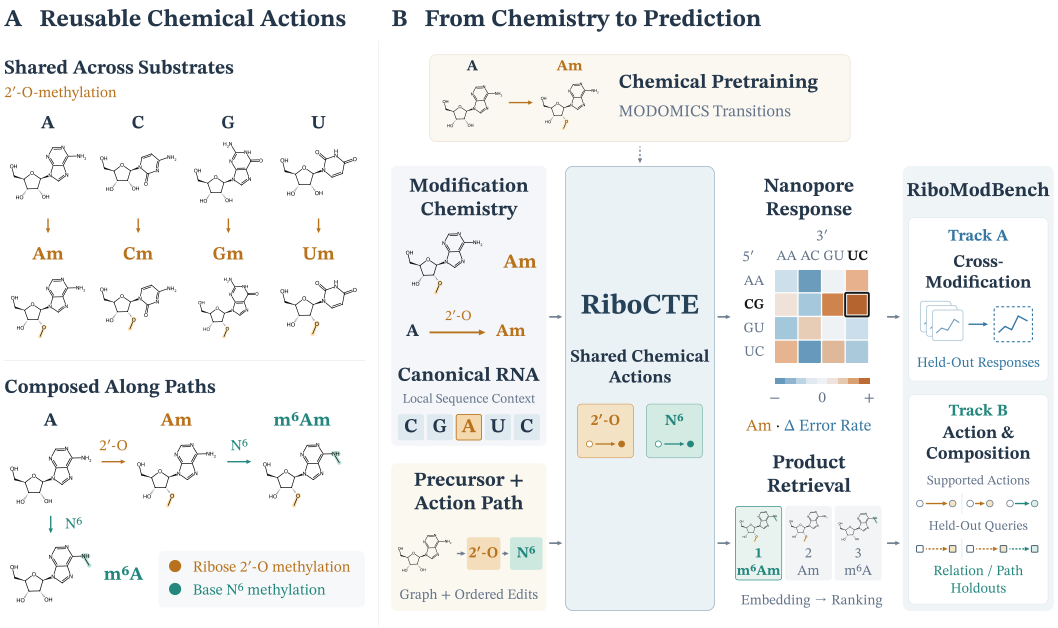


Figure 1: **Shared chemical actions and prediction tasks. (A)** Ribose 2′-O-methylation is shared across nucleosides and composes with N^6-methylation. **(B)** MODOMICS transitions train reusable actions without response labels (Sordyl et al., 2026). RiboModBench evaluates differential nanopore responses for held-out modifications (Track A) and product ranking for held-out chemical relations and compositions (Track B).

frozen molecular embeddings through actions in a learned subspace. The same actions are then reused on sequence-conditioned canonical states, where a shared nonlinear decoder predicts the transformed-minus-canonical response.

We evaluate this sharing principle with **RiboModBench**, which pairs response generalization with chemical-action transfer and composition tests (Figure 1B). ONT all-5-mer data provide modification/context holdouts and few-shot adaptation (Stoiber, 2026), while Curlcake RNA004 tests external transfer (Cruciani et al., 2025). Separate chemical splits test reuse across substrates and composition with observed or withheld constituent relations. Results show improved response ranking under modification and context shifts, with controlled comparisons demonstrating gains from learned actions and path supervision.

Contributions. **(1) CTE and RiboCTE:** reusable, composable chemical actions connecting molecular transformations to contextual RNA response prediction, with molecular rollout analysis. **(2) Systematic evaluation with RiboModBench:** response-transfer and chemical-action tests demonstrating improved ranking, chemical-action fidelity, and few-shot and external transfer.

2 Related Work

RNA modification and sequence representations. MultiRM, TransRNAm, MRM-BERT, RNA-ModX, and EvoRMD predict modification occurrence or type from sequence and biological context (Song et al., 2021; Chen et al., 2023; Wang & Zhou, 2024; Yuge et al., 2025; Wang et al., 2026). TandemMod, NanoSPA, and ORCA study transfer and multi-modification analysis in nanopore sequencing (Wu et al., 2024; Huang et al., 2024; Dong et al., 2026). ELIGOS and DRUMMER use comparative basecalling errors to detect modifications, while nanoRMS combines modification-sensitive read features for stoichiometry estimation (Jenjaroenpun et al., 2021; Abebe et al., 2022; Begik et al., 2021). Ding et al. (2021) predict nanopore ionic currents from chemical k-mer graphs, including transfer from canonical and m^6A-containing RNA to m^2G-containing RNA. Stephenson et al. (2025) combine molecular graph and sequence features for current prediction in expanded

DNA alphabets, evaluating ten-XNA base-dropout and few-shot transfer. RiboCTE learns chemical actions for predicting differential nanopore responses across RNA contexts. RNA-FM, RNAErnie, RiNALMo, and ERNIE-RNA provide contextual sequence representations complementary to this chemical input (Chen et al., 2022; Wang et al., 2024; Penić et al., 2025; Yin et al., 2025).

Modified biomolecules and molecular representations. PTM-Mamba, COMPASS-PTM, and ProtSyntax incorporate modification identities or biochemical priors into protein representations (Peng et al., 2024; 2025; Zhang et al., 2026; Lin et al., 2026). Bi-TEAM combines local noncanonical-residue chemistry with protein context through residual fusion (Gu et al., 2026). chemCPA predicts cellular responses to unseen drugs from molecular representations (Hetzel et al., 2022). Reaction-aware approaches, including MolR, ReaKE, RXGL, and KAGT, exploit reaction relations or atom correspondence (Wang et al., 2021; Xie et al., 2024; Li et al., 2024; Qiao et al., 2026); graph-edit models explicitly represent reaction changes (Sacha et al., 2021; Zhong et al., 2023). We test whether action structure adds value beyond direct molecular endpoints, including CheMeleon and Mol-JEPA representations (Burns et al., 2025; Rottach et al., 2026).

Latent actions, composition, and predictive learning. Equivariant representation learning relates input transformations to latent transformations, including noninvertible semigroup settings (Cohen & Welling, 2016; Bronstein et al., 2021; Worrall & Welling, 2019). The Homomorphism Autoencoder (HAE) learns latent group actions from transitions (Keurti et al., 2023). Attributes as Operators models visual attributes with shared linear transformations for unseen attribute–object combinations (Nagarajan & Grauman, 2018). Relation embeddings provide precedents for learned relational operators, shared bases, and path composition (Nickel et al., 2011; Bordes et al., 2013; Lin et al., 2015; Sun et al., 2019; Liu et al., 2017; Schlichtkrull et al., 2018; Guu et al., 2015). Composing learned transformations enables novel combinations (Chang et al., 2019); compositional splits (Keysers et al., 2020) and kernel analyses (Lippl & Stachenfeld, 2025) further characterize generalization to such combinations. Joint-embedding predictive architectures, including V-JEPA 2 and Mol-JEPA, motivate latent-target prediction (Assran et al., 2025; Rottach et al., 2026). RiboCTE learns descriptor-conditioned chemical actions from transitions and paths against frozen molecular targets, then reuses them in sequence-conditioned response prediction.

3 Problem Formulation

Typed chemical actions. Let $\mathcal{G}$ be the space of standardized nucleoside graphs, including stereochemistry, formal charge, and attachment-site information. A typed action $a \in \mathcal{A}$ specifies an edit and its site, and induces a partial map

$$\tau_a : \mathcal{D}_a \to \mathcal{G}, \qquad G_m = \tau_a(G_p). \tag{1}$$

The domain $\mathcal{D}_a$ describes chemical applicability, not enzymatic or cellular feasibility. A path $\pi = (a_1, \ldots, a_L)$ is valid when each intermediate state belongs to the next action's domain. Chemical actions need not be invertible or globally composable; we assume neither a symmetry group nor commutativity of chemical transformations.

Transformation preservation. For a molecular encoder $E : \mathcal{G} \to \mathbb{R}^d$, CTE requires

$$E(\tau_a(G)) \approx T_a(E(G)), \qquad G \in \mathcal{D}_a. \tag{2}$$

For a valid path, we seek $T_\pi(E(G_0)) \approx E(G_L)$, where $T_\pi = T_{a_L} \circ \cdots \circ T_{a_1}$. CTE denotes approximate transformation consistency on valid chemical paths. RiboCTE learns the shared subspace and action operators while keeping the molecular encoder fixed.

Downstream generalization. Given a canonical RNA context x centered at position i, a canonical parent G_0, and a modification m reachable by a specified valid path π_m, the target is

$$\Delta y_\nu(x, m) = y_\nu(x, G_m) - y_\nu(x, G_0). \tag{3}$$

Here ν fixes the experimental cohort, processing protocol, and incorporation pattern; we suppress it within a dataset. The target averages designated-site responses under that condition. Leave-one-modification-out (LOMO) evaluation requires $\mathcal{M}_{\text{train}} \cap \mathcal{M}_{\text{test}} = \varnothing$ for response labels; test structures and chemical descriptions remain available during chemical pretraining. Separate chemical splits withhold relations or compositions. Inputs specify the designated modification and canonical 5-mer; other modified sites follow the recorded cohort design.

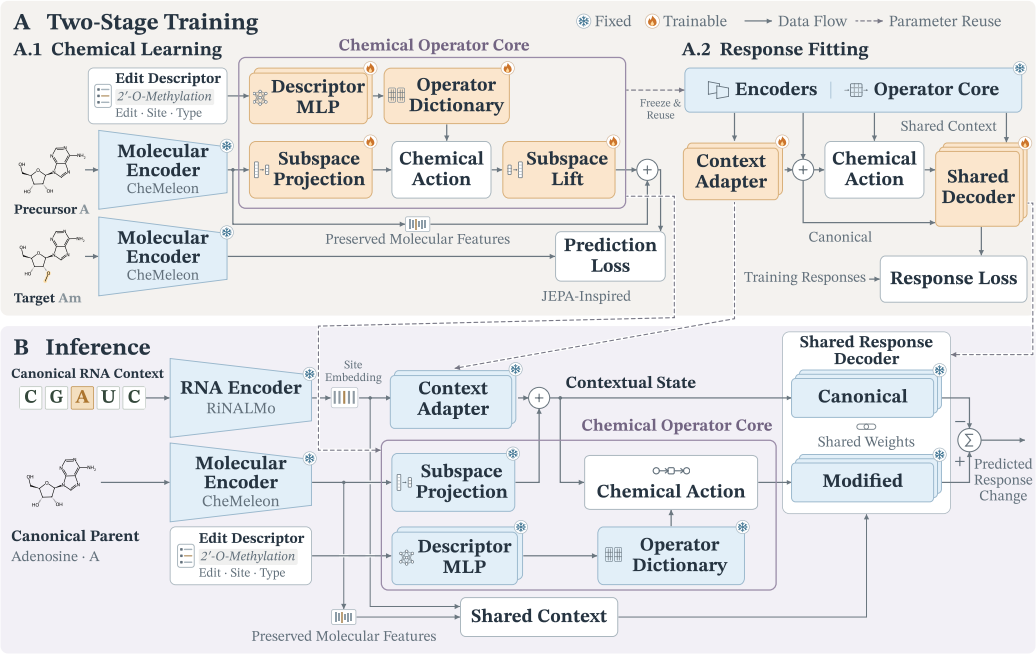


Figure 2: **RiboCTE architecture. (A)** Chemical actions predict single-step and path endpoints; the chemical core remains fixed during response fitting. **(B)** RNA context offsets the canonical-parent state before the same action is applied. A shared decoder predicts the transformed-minus-canonical response, conditioned on RNA context and preserved molecular features.

4 Learning and Transferring Shared Chemical Actions

RiboCTE first learns shared chemical actions from molecular transitions, then freezes them for contextual response prediction (Figure 2). Both stages use the same molecular coordinate system.

4.1 Shared, Composition-Closed Operators

States and action descriptors. A frozen molecular encoder supplies $z_G = E_0(G) \in \mathbb{R}^d$ with fixed, chemical-training-fitted normalization. The descriptor d_a encodes an atom-aligned edit template, attachment site, and reaction type. Atom correspondence supports curation; the operator generator receives neither the complete product graph nor product/edge identities. The precursor environment is represented in z_G; templates and applicability rules appear in Appendix E.

Operator parameterization. Let $Q \in \mathbb{R}^{d\times r}$ have orthonormal columns, $r < d$, and write $u = Q^\top z$. An action applies a shared displacement in this subspace:

$$\widetilde{T}_a(u) = u + b_a, \qquad T_a(z) = (I - QQ^\top)z + Q\widetilde{T}_a(Q^\top z). \tag{4}$$

The complement is unchanged. To share statistical strength across sparse actions, a small descriptor network predicts coefficients $\eta(a) = g_\phi(d_a)$ over a shared dictionary:

$$b_a = \sum_{j=1}^{J} \eta_j(a) t_j. \tag{5}$$

The same typed action uses the same b_a on all compatible substrates, and the identity action has $b_a = 0$. Related edit types share dictionary parameters without forcing every methylation site to use the same operator.

Proposition 1 (Composition closure) *Fix Q with $Q^\top Q = I_r$. The family in Eq. (4), allowing arbitrary offsets, is closed under composition. Applying a and then b gives*

$$b_{b\circ a} = b_a + b_b. \tag{6}$$

For a valid path, $b_\pi = \sum_{t=1}^{L} b_{a_t}$.

Paths compose by summing primitive offsets in the full translation family of Proposition 1. Paths with the same action multiset share a latent displacement, while chemical applicability determines valid orders. Proofs and approximation limits appear in Appendix A.

4.2 Action-Conditioned Predictive Learning

We learn reusable actions by predicting molecular endpoints across substrates and along valid paths. Given allowed training edges $\mathcal{E}_{\text{tr}}$ and valid paths $\mathcal{P}_{\text{tr}}$, we minimize

$$\mathcal{L}_{\text{pred}} = \mathbb{E}_{(G_p,a,G_m)\in\mathcal{E}_{\text{tr}}} \|T_a(z_p) - \text{sg}(z_m)\|_2^2, \tag{7}$$

$$\mathcal{L}_{\text{path}} = \mathbb{E}_{(G_0,\pi,G_L)\in\mathcal{P}_{\text{tr}}} \|T_\pi(z_0) - \text{sg}(z_L)\|_2^2, \tag{8}$$

$$\mathcal{L}_{\text{chem}} = \mathcal{L}_{\text{pred}} + \lambda_{\text{path}}\mathcal{L}_{\text{path}}. \tag{9}$$

Here sg denotes a fixed target: E_0 and its normalization are frozen while Q, the dictionary, and g_ϕ are learned. Path supervision jointly constrains primitive actions through their cumulative effects along valid chemical paths, using no nanopore outcomes.

Proposition 2 (Rollout error) *Let* $z_t = E_0(G_t)$ *follow a valid length-L path. If* $\|T_{a_t}(z_{t-1}) - z_t\|_2 \le \epsilon_t$ *and* T_{a_t} *is* ℓ_t*-Lipschitz, then, starting at* z_0*,*

$$\|T_\pi(z_0) - z_L\|_2 \le \sum_{t=1}^{L} \epsilon_t \prod_{s=t+1}^{L} \ell_s. \tag{10}$$

For Eq. (4), $\ell_t = 1$*, so the bound reduces to* $\sum_{t=1}^{L} \epsilon_t$.

4.3 Contextual Prediction with a Shared Action Core

To transfer the learned chemistry, we apply the same actions to sequence-conditioned canonical states (Figure 2B). Let $h_x = F_{\text{RNA}}(x)_i$ be a frozen RNA representation at the designated site, and $v_x = B_\psi(h_x) \in \mathbb{R}^r$ a learned contextual offset. Starting from the canonical parent,

$$u_0 = Q^\top z_0, \qquad s_0(x) = u_0 + v_x, \qquad \hat{s}_m(x) = s_0(x) + b_{\pi_m}. \tag{11}$$

The chemical action core is reused directly, with the context fixed throughout a path. Define $c_x = [h_x; (I - QQ^\top)z_0]$. A shared nonlinear decoder predicts the response difference:

$$\widehat{\Delta y}(x,m) = D_\omega(c_x, \hat{s}_m(x)) - D_\omega(c_x, s_0(x)). \tag{12}$$

The base predictor therefore assigns exactly zero effect to the identity action. The shared chemical displacement is

$$\hat{s}_m(x) - s_0(x) = b_{\pi_m}. \tag{13}$$

The decoder combines the shared displacement with the canonical-parent state and RNA context. The same action can thus yield different response profiles across substrates and contexts.

Training and adaptation. After chemical training, we freeze $E_0, Q, g_\phi, \{t_j\}$ and F_{RNA}, and fit B_ψ, D_ω with modification-macro squared error plus a RankNet term (Burges et al., 2005) comparing response orderings between contexts of the same source modification. The squared-error term learns response values, while the ranking term targets contextual ordering. In the K-shot setting, a ridge residual head is fitted on the query modification's support contexts; all representation modules remain frozen. Appendix F specifies the head and tuning protocol.

5 RiboModBench: Cross-Modification Generalization

RiboModBench separates response generalization (Track A) from chemical-action transfer and composition (Track B), with model-independent splits and separately recorded supervision.

5.1 Track A: Cross-Modification Response Generalization

For modification m and canonicalized centered context c, Track A predicts

$$Y_{m,c} = R(m, c) - R(0_m, c), \tag{14}$$

where 0_m is the matched canonical parent and R the center-site substitution-plus-deletion rate against the canonical reference. Inputs comprise chemistry and the canonical 5-mer; modified/control reads construct targets only (Appendix C).

Resource roles. MODOMICS supplies outcome-free chemistry (Sordyl et al., 2026); ONT all-5-mer supports development and nested selection (Stoiber, 2026); Curlcake RNA004 evaluates external transfer with ONT-selected configurations (Cruciani et al., 2025). Versions and accessions appear in Appendix B.

Splits and supervision. LOMO withholds a target modification's entire response profile. Double-OOD additionally excludes its query flank groups from every training modification, using common center-masked groups across parent bases. K-shot supports and queries are disjoint; Double-OOD supports remain within permitted training-context groups. Grouping, nested selection, and sampling follow Appendix D.

Metrics. For query contexts $\mathcal{Q}_m$, we report per-modification Spearman ρ_m and

$$\mathrm{nMAE}_m = \frac{|\mathcal{Q}_m|^{-1} \sum_{c \in \mathcal{Q}_m} |\hat{Y}_{m,c} - Y_{m,c}|}{\mathrm{IQR}(Y_{m,\mathcal{Q}_m}) + \epsilon}. \tag{15}$$

IQR is used only for scoring. Scores macro-average modifications; Appendix D specifies eligibility and aggregation rules.

5.2 Track B: Action Transfer and Composition Diagnostics

Single-step folds withhold an action–substrate relation while retaining the action on other substrates. For eligible length-2–4 paths, **edge-seen composition** retains constituent edges but withholds the ordered combination's path and aggregate supervision. **Edge-unseen composition** additionally removes path-specific constituent edges while retaining each primitive on another substrate. Neither setting requires unseen molecular nodes. Appendix E specifies folds, shortcut exclusions, and duplicate-route exclusions.

Given G_0 and an action specification, predictors roll out to $\hat{z}_L$ and rank a common candidate bank against frozen molecular embeddings. Candidate products cannot construct query descriptors or supervise excluded relations. MRR and Recall@5 macro-average eligible actions or composition groups; Appendix E specifies eligibility and ranking rules.

Symbolic execution provides a reference for fully specified edits. Retrieval measures latent-computation fidelity to known chemistry.

6 Experiments

We evaluate response generalization (RQ1), chemical-action transfer and composition (RQ2), contributions of action learning (RQ3), and few-shot and external transfer (RQ4).

6.1 Experimental Setup and Baselines

RiboCTE uses frozen CheMeleon embeddings; all response methods use the frozen RiNALMo-Giga sequence backbone (Burns et al., 2025; Penić et al., 2025). Evaluation covers 2,048 common-eligible ONT cells across eight modifications and a 106-state chemical release. Base predictors share folds and source-only model selection. Unless stated otherwise, learned-model results average three independent training seeds (0, 1, 2). Appendix F specifies training, calibration, and adaptation; the anonymous repository provides evaluation records and reproduction instructions.

Table 1: Differential response prediction on ONT. Scores macro-average eight modifications after averaging training seeds and context folds.

	LOMO		Double-OOD	
Method	$\rho \uparrow$	nMAE$\downarrow$	$\rho \uparrow$	nMAE$\downarrow$
Context-only	0.132	1.276	0.085	1.324
Morgan	0.110	1.232	0.098	1.279
CheMeleon	0.079	1.007	0.082	1.116
Mol-JEPA	0.120	1.166	0.090	1.221
HAE-E	0.172	1.171	0.096	1.234
HAE-C	0.077	0.987	0.062	**1.006**
RiboCTE	**0.216**	**0.983**	**0.184**	1.029

Table 2: Chemical-action transfer and composition. ES/EU denote edge-seen/edge-unseen composition. Scores macro-average action or ordered-action groups; bold marks the best learned scores, including ties. Symbolic execution provides a reference for fully specified edits, with 100% executable query coverage in each tier.

	Action–substrate		Composition (ES)		Composition (EU)	
Predictor	MRR$\uparrow$	R@5$\uparrow$	MRR$\uparrow$	R@5$\uparrow$	MRR$\uparrow$	R@5$\uparrow$
Action-conditioned MLP	0.927	**1.000**	**1.000**	**1.000**	0.962	**1.000**
HAE	0.877	**1.000**	**1.000**	**1.000**	0.850	0.993
RiboCTE	**0.943**	**1.000**	**1.000**	**1.000**	**0.981**	**1.000**
Symbolic edit reference	1.000	1.000	1.000	1.000	1.000	1.000

CheMeleon and **Mol-JEPA** combine known precursor/product embeddings, their difference, chemical descriptors, and RNA context (Burns et al., 2025; Rottach et al., 2026). We adapt **HAE** to chemical transitions with two response interfaces: **HAE-E** uses endpoint representations in the common readout, while **HAE-C** applies operators to contextual states before shared difference decoding (Keurti et al., 2023). Context-only and Morgan fingerprints provide controls. Fusion baselines share projection width, readout, and tuning budget; HAE-C uses the same adapter/decoder design in its own latent coordinates. Appendix F.1 details baselines and RQ3 controls.

6.2 RQ1: Generalization Across Modifications and Contexts

RiboCTE leads response ranking in both holdout settings, with the larger margin under joint modification and context shift (Table 1). Its correlations of 0.216 and 0.184 exceed the strongest LOMO and Double-OOD baselines by 0.044 and 0.086, respectively. Double-OOD excludes query context groups from every source modification, so this advantage demonstrates transfer across both held-out identities and unseen context combinations. RiboCTE also achieves the lowest LOMO nMAE.

Case study: a shared chemical action for context-dependent response ranking. Ribose $2'$-O-methylation provides a shared action across A, C, G and U, linking molecular transformations to context-dependent response ranking (Figure 3). Under LOMO, RiboCTE attains the highest mean Spearman correlation among the seven matched predictors for each of Am, Cm, Gm and Um.

6.3 RQ2: Transferable Actions and Held-Out Compositions

Table 2 evaluates chemical-action fidelity using common frozen molecular targets and a 106-state candidate bank. The action-conditioned MLP receives the same descriptors and precursor representations, and HAE predictions are decoded into the common target space. We evaluate 62 action–substrate queries, 194 edge-seen paths in 184 ordered-action groups, and 46 edge-unseen queries.

Edge-unseen composition tests whether actions learned on other substrates can recover paths whose constituent relations are withheld. RiboCTE reaches MRR 0.981 versus 0.962 for the action-conditioned MLP, demonstrating reuse beyond directly observed relations.

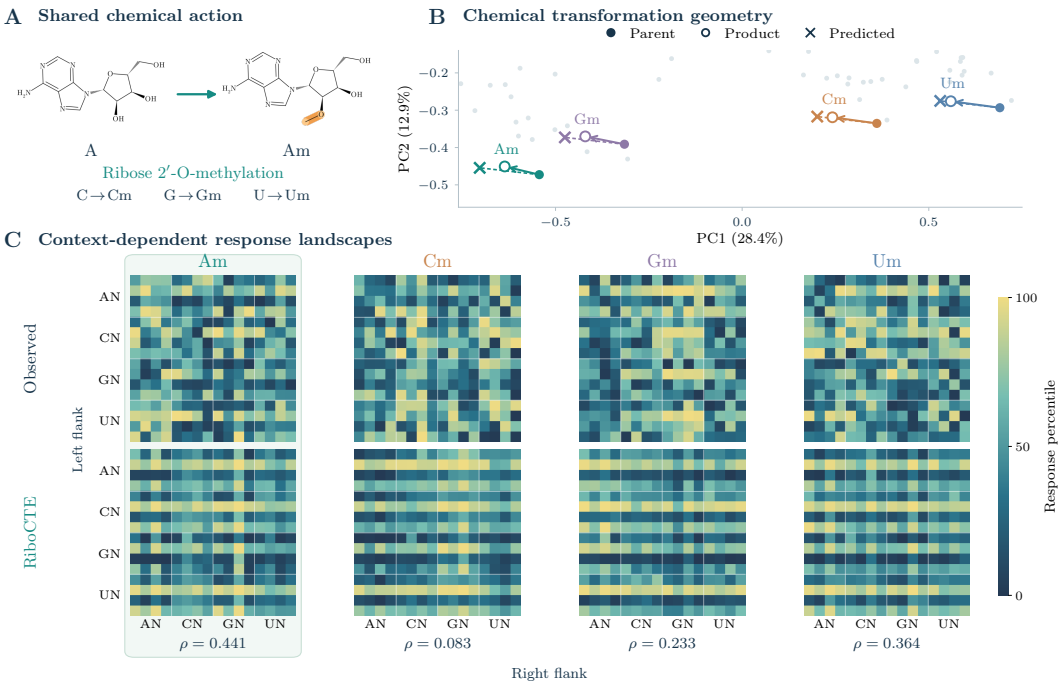


Figure 3: **(A)** Ribose 2′-O-methylation converts A to Am and applies to C, G and U. **(B)** A common PCA of normalized molecular embeddings shows parents (filled circles), products (open circles), and action-predicted endpoints (crosses). **(C)** Observed (top) and predicted (bottom) within-profile response percentiles across 256 contexts per modification under LOMO, with ρ averages 3 seeds.

Table 3: Chemical-action controls under matched response refitting. Entries are macro Spearman ρ.

Variant	LOMO $\rho \uparrow$	Double-OOD $\rho \uparrow$
Exact endpoints	0.162	0.168
Generic bilinear fusion	0.084	0.061
No path supervision ($\lambda_{\text{path}} = 0$)	0.063	0.055
Rotated action offsets	0.122	0.099
RiboCTE	**0.216**	**0.184**

6.4 RQ3: Contributions of Chemical-Action Learning

Table 3 compares learned actions, path supervision, and the molecular frame with all response modules refitted under common folds and tuning budgets. Exact endpoints retain the learned projection and contextual decoder but replace the action prediction with the known product embedding; no-path removes path supervision, and rotation changes offset orientation in the retained frame. Appendices F.1 and G detail these and supplementary controls.

Learned actions improve contextual transfer. With the same projection and decoder architecture, RiboCTE improves over exact endpoints from 0.162 to 0.216 under LOMO and from 0.168 to 0.184 under Double-OOD. It also outperforms generic bilinear fusion, which combines endpoint chemistry with RNA context. These comparisons show that learning shared transformations adds predictive value even when exact product structures are available.

Path supervision and the molecular frame contribute to generalization. Chemical path supervision improves Double-OOD correlation from 0.055 to 0.184, showing that learning the cumulative effects of primitive actions benefits downstream transfer. The learned offsets also outperform the rotated control (0.184 versus 0.099), with a positive difference for all eight modification means.

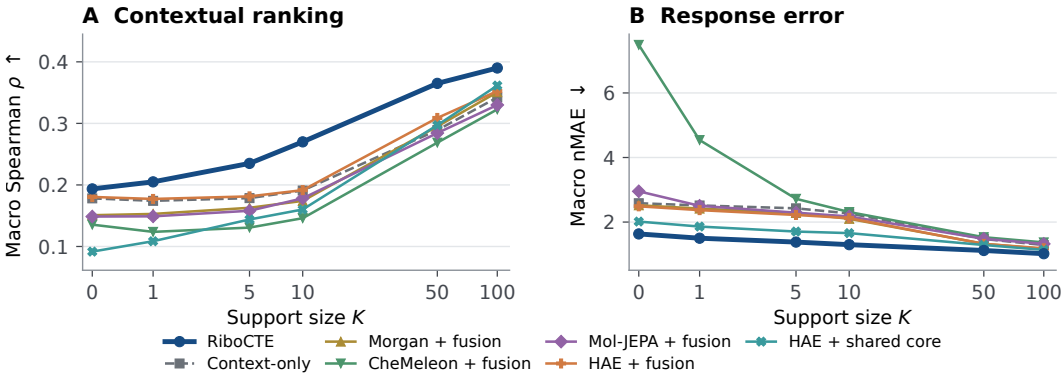


Figure 4: Few-shot adaptation on held-out modifications. Curves macro-average eight modifications, three seeds, and twenty paired support draws, with disjoint support/query sets and fixed queries across budgets within each draw. Positions use $\log(1 + K)$; see Appendix F.

6.5 RQ4: Few-Shot Adaptation and External Transfer

Adaptation with limited response labels. RiboCTE retains the highest mean ranking across support budgets, extending its transfer advantage as target response labels become available (Figure 4). We evaluate $K \in \{0, 1, 5, 10, 50, 100\}$ with nested supports and fixed, disjoint query sets within each draw. All methods use the common residual adaptation protocol in Appendix F. With 100 support contexts, RiboCTE leads both metrics at $\rho = 0.389$ and nMAE 1.021.

Transfer to an external experimental dataset. We select configurations using ONT-only inner validation, refit on all eligible ONT outcomes, and predict Curlcake at $K = 0$. The two resources differ in construct and modification-incorporation design (Appendix B). Table 4 separately reports the three source-shared and four source-unseen modifications.

Table 4: Zero-shot external transfer: three shared and four unseen ONT response-training identities. Chemical-pretraining exposure is separate. Scores average 3 seeds, then identities.

	Shared modifications		Unseen modifications	
Method	$\rho \uparrow$	nMAE↓	$\rho \uparrow$	nMAE↓
CheMeleon	0.397	0.845	0.054	1.504
Mol-JEPA	0.303	0.852	0.129	1.527
HAE-E	0.361	0.855	0.077	1.537
HAE-C	0.377	0.848	−0.009	1.530
RiboCTE	**0.412**	**0.827**	**0.268**	**1.346**

RiboCTE leads both external groups in mean correlation and nMAE, with the larger ranking margin on source-unseen modifications ($\rho = 0.268$ versus 0.129 for the strongest baseline; nMAE 1.346). The gain on identities absent from source response training shows that shared actions remain useful under simultaneous changes in modification identity and experimental dataset.

7 Limitations and Future Work

The sparse chemical graph limits coverage of actions and paths, while low-dimensional sharing restricts representable molecular changes. Enzymatic feasibility and cellular regulation additionally depend on cofactors, RNA structure, and readers; PCIF1-mediated m^6Am formation, for example, depends on cap context (Sendinc et al., 2019). Our evaluation addresses predictive transfer of controlled nanopore responses; native modification mixtures, multi-site interactions, and other RNA functions remain directions for extension.

AI USE STATEMENT

In this work, we used generative AI tools to implement parts of the method and experiment code, i.e., AI-assisted vibe coding, and to design and provide feedback on the research methodology and experiments. We have not used generative AI tools to generate synthetic datasets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses, clean and reformat datasets, or interpret results; assisting with translation and supporting qualitative or thematic data analysis are not applicable to this work. Additionally, we used generative AI tools to draft and edit parts of this paper, i.e., to polish the writing. We have reviewed all AI-assisted work: AI-assisted code was tested and verified for correctness by the authors, and AI-assisted text was reviewed and edited by the authors for accuracy. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work uses only publicly available resources, namely the MODOMICS chemical database and public nanopore datasets of synthetic in vitro transcribed RNA constructs, for computational evaluation of response predictions. It involves no human subjects, personal or clinical data, or new wet-lab experiments. RiboModBench redistributes only data derived from these public sources under their original terms, with per-resource provenance recorded in Appendix B.

REPRODUCIBILITY STATEMENT

Section 4 and Appendix F specify the model, action parameterization, training objectives, optimization, and baseline and adaptation details. Appendix A states the assumptions and gives complete proofs of the propositions in Sections 4.1 and 4.2. Appendices B, C, D, and E, together with Table 5, document data provenance, response construction, split definitions, and evaluation protocols. Code and RiboModBench are available at the anonymous repository linked in the abstract.

REFERENCES

Jonathan S. Abebe, Alexander M. Price, Katharina E. Hayer, Ian Mohr, Matthew D. Weitzman, Angus C. Wilson, and Daniel P. Depledge. DRUMMER—rapid detection of RNA modifications through comparative nanopore sequencing. *Bioinformatics*, 38(11):3113–3115, 2022. doi: 10.1093/bioinformatics/btac274.

Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Mojtaba Komeili, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, Sergio Arnaud, Abha Gejji, Ada Martin, Francois Robert Hogan, Daniel Dugas, Piotr Bojanowski, Vasil Khalidov, Patrick Labatut, Francisco Massa, Marc Szafraniec, Kapil Krishnakumar, Yong Li, Xiaodong Ma, Sarath Chandar, Franziska Meier, Yann LeCun, Michael Rabbat, and Nicolas Ballas. V-JEPA 2: Self-supervised video models enable understanding, prediction and planning, 2025. URL https://arxiv.org/abs/2506.09985.

Oguzhan Begik, Morghan C. Lucas, Leszek P. Pryszcz, Jose Miguel Ramirez, Rebeca Medina, Ivan Milenkovic, Sonia Cruciani, Huanle Liu, Helaine Graziele Santos Vieira, Aldema Sas-Chen, John S. Mattick, Schraga Schwartz, and Eva Maria Novoa. Quantitative profiling of pseudouridylation dynamics in native RNAs with nanopore sequencing. *Nature Biotechnology*, 39(10):1278–1291, 2021. doi: 10.1038/s41587-021-00915-6.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, volume 26, 2013. URL https://papers.nips.cc/paper/5071-translating-embeddings-for-modeling-multi-relational-data.pdf.

Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021. URL https://arxiv.org/abs/2104.13478.

Chris J. C. Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. Technical Report MSR-TR-2005-06, Microsoft Research, August 2005. URL https://www.microsoft.com/en-us/research/publication/learning-to-rank-using-gradient-descent/.

Jackson W. Burns, Akshat Shirish Zalte, Charlles R. A. Abreu, Jochen Sieg, Christian Feldmann, Miriam Mathea, and William H. Green. Deep learning foundation models from classical molecular descriptors, 2025. URL https://arxiv.org/abs/2506.15792.

Michael Chang, Abhishek Gupta, Sergey Levine, and Thomas L. Griffiths. Automatically composing representation transformations as a means for generalization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=B1ffQnRcKX.

Jiayang Chen, Zhihang Hu, Siqi Sun, Qingxiong Tan, Yixuan Wang, Qinze Yu, Licheng Zong, Liang Hong, Jin Xiao, Tao Shen, Irwin King, and Yu Li. Interpretable rna foundation model from unannotated data for highly accurate rna structure and function predictions, 2022. URL https://arxiv.org/abs/2204.00300.

Taoning Chen, Tingfang Wu, Deng Pan, Jinxing Xie, Jin Zhi, Xuejiao Wang, Lijun Quan, and Qiang Lyu. Transrnam: identifying twelve types of rna modifications by an interpretable multi-label deep learning model based on transformer. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 20(6):3623–3634, 2023.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.

Sonia Cruciani, Anna Delgado-Tejedor, Leszek P Pryszcz, Rebeca Medina, Laia Llovera, and Eva Maria Novoa. De novo basecalling of rna modifications at single molecule and nucleotide resolution. *Genome Biology*, 26(1):38, 2025.

Hongxu Ding, Ioannis Anastopoulos, Andrew D. Bailey, IV, Joshua Stuart, and Benedict Paten. Towards inferring nanopore sequencing ionic currents from nucleotide chemical structures. *Nature Communications*, 12(1):6545, 2021. doi: 10.1038/s41467-021-26929-x. URL https://www.nature.com/articles/s41467-021-26929-x.

Han Dong, Yongsheng Gao, Zhengyi Cai, Yi Li, Xing Li, Fangqing Zhao, and Jinyang Zhang. Comprehensive mapping of rna modification dynamics and crosstalk via deep learning and nanopore direct rna-sequencing. *Nature Communications*, 17(1):1722, 2026.

Chunbin Gu, Zijun Gao, Mutian He, Jingjie Zhang, Haipeng Wen, Zihao Luo, Xiaorui Wang, Hanqun Cao, Jiajun Bu, Chang-Yu Hsieh, and Pheng Ann Heng. Bi-team: A unified cross-scale representation learning framework for chemically modified biomolecules, 2026. URL https://arxiv.org/abs/2603.01873.

Kelvin Guu, John Miller, and Percy Liang. Traversing knowledge graphs in vector space. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 318–327. Association for Computational Linguistics, 2015. doi: 10.18653/v1/D15-1038. URL https://aclanthology.org/D15-1038/.

Leon Hetzel, Simon Boehm, Niki Kilbertus, Stephan Günnemann, Mohammad Lotfollahi, and Fabian J. Theis. Predicting cellular responses to novel drug perturbations at a single-cell resolution. In *Advances in Neural Information Processing Systems*, volume 35, 2022. doi: 10.52202/068431-1937.

Sihao Huang, Adam C Wylder, and Tao Pan. Simultaneous nanopore profiling of mrna m6a and pseudouridine reveals translation coordination. *Nature biotechnology*, 42(12):1831–1835, 2024.

Piroon Jenjaroenpun, Thidathip Wongsurawat, Taylor D. Wadley, Trudy M. Wassenaar, Jun Liu, Qing Dai, Visanu Wanchai, Nisreen S. Akel, Azemat Jamshidi-Parsian, Aime T. Franco, Gunnar Boysen, Michael L. Jennings, David W. Ussery, Chuan He, and Intawat Nookaew. Decoding the epitranscriptional landscape from native RNA sequences. *Nucleic Acids Research*, 49(2):e7, 2021. doi: 10.1093/nar/gkaa620.

Hamza Keurti, Hsiao-Ru Pan, Michel Besserve, Benjamin F Grewe, and Bernhard Schölkopf. Homomorphism AutoEncoder – learning group structured representations from observed transitions. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 16190–16215. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/keurti23a.html.

Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. Measuring compositional generalization: A comprehensive method on realistic data. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SygcCnNKwr.

Anchen Li, Elena Casiraghi, and Juho Rousu. Chemical reaction enhanced graph learning for molecule representation. *Bioinformatics*, 40(10):btae558, 2024.

Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015. doi: 10.1609/aaai.v29i1.9491.

Yiyu Lin, Jiahui Wu, You Zhou, Xinye Ni, Shan Chang, Jun Ding, Yan Wang, Xin Gao, and Sen Yang. Protsyntax: a protein large language model for decoding post-translational modification syntax and function. *bioRxiv*, pp. 2026–07, 2026.

Samuel Lippl and Kim Stachenfeld. When does compositional structure yield compositional generalization? a kernel theory. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=FPBce2P1er.

Hanxiao Liu, Yuexin Wu, and Yiming Yang. Analogical inference for multi-relational embeddings. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 2168–2178. PMLR, 2017. URL https://proceedings.mlr.press/v70/liu17d.html.

Huanle Liu, Oguzhan Begik, Morghan C. Lucas, Jose Miguel Ramirez, Christopher E. Mason, David Wiener, Schraga Schwartz, John S. Mattick, Martin A. Smith, and Eva Maria Novoa. Accurate detection of m6A RNA modifications in native RNA sequences. *Nature Communications*, 10(1): 4079, 2019. doi: 10.1038/s41467-019-11713-9. URL https://pubmed.ncbi.nlm.nih.gov/31501426/.

Tushar Nagarajan and Kristen Grauman. Attributes as Operators: Factorizing unseen attribute-object compositions. In *Computer Vision – ECCV 2018*, volume 11205 of *Lecture Notes in Computer Science*, pp. 172–190. Springer, 2018. doi: 10.1007/978-3-030-01246-5_11. URL https://www.cs.utexas.edu/~grauman/papers/attributes-operators-eccv2018.pdf.

Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning*, 2011. URL https://icml.cc/2011/papers/438_icmlpaper.pdf.

Fred Zhangzhi Peng, Chentong Wang, Tong Chen, Benjamin Schussheim, Sophia Vincoff, and Pranam Chatterjee. Ptm-mamba: a ptm-aware protein language model with bidirectional gated mamba blocks. In *Proceedings of the 33rd ACM international conference on information and knowledge management*, pp. 5475–5478, 2024.

Fred Zhangzhi Peng, Chentong Wang, Tong Chen, Benjamin Schussheim, Sophia Vincoff, and Pranam Chatterjee. PTM-Mamba: a PTM-aware protein language model with bidirectional gated mamba blocks. *Nature Methods*, 22(5):945–949, 2025. doi: 10.1038/s41592-025-02656-9.

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.

Jianbo Qiao, Kefei Li, Junru Jin, Ding Wang, Wenjia Gao, Shaoye Zhang, Shuan Liu, and Leyi Wei. Relation-aware pretraining and reaction center modeling for chemical reaction graph representation learning. *Journal of Chemical Theory and Computation*, 2026.

Florian Rottach, Sebastian Schieferdecker, William Rudman, Randall Balestriero, and Carsten Eickhoff. Mol-JEPA: A multimodal joint embedding predictive architecture for molecules, 2026. URL https://arxiv.org/abs/2608.22642.

Mikołaj Sacha, Mikołaj Błaż, Piotr Byrski, Paweł Dąbrowski-Tumański, Mikołaj Chromiński, Rafał Loska, Paweł Włodarczyk-Pruszyński, and Stanisław Jastrzębski. Molecule edit graph attention network: Modeling chemical reactions as sequences of graph edits, 2021. URL https://arxiv.org/abs/2006.15426.

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC 2018*, pp. 593–607. Springer, 2018. doi: 10.1007/978-3-319-93417-4_38.

Erdem Sendinc, David Valle-Garcia, Abhinav Dhall, Hao Chen, Telmo Henriques, Jose Navarrete-Perea, Wanqiang Sheng, Steven P. Gygi, Karen Adelman, and Yang Shi. PCIF1 catalyzes m6Am mRNA methylation to regulate gene expression. *Molecular Cell*, 75(3):620–630.e9, 2019. doi: 10.1016/j.molcel.2019.05.030. URL https://pubmed.ncbi.nlm.nih.gov/31279659/.

Zitao Song, Daiyun Huang, Bowen Song, Kunqi Chen, Yiyou Song, Gang Liu, Jionglong Su, João Pedro de Magalhães, Daniel J Rigden, and Jia Meng. Attention-based multi-label neural networks for integrated prediction and interpretation of twelve widely occurring rna modifications. *Nature communications*, 12(1):4011, 2021.

Dominik Sordyl, Etienne Boileau, Agata Bernat, Satyabrata Maiti, Sunandan Mukherjee, S Naeim Moafinejad, Masoud Amiri Farsani, Anastasiya Shavina, Andrea Cappannini, Giada Agostini, et al. Modomics: a database of rna modifications and related information. 2025 update and 20th anniversary. *Nucleic Acids Research*, 54(D1):D219–D225, 2026.

Ashley Stephenson, Jayson Sumabat, Hinako Kawabe, Sidharth Lakshmanan, Hyo Joong Kim, Yubing Li, Jorge A. Marchand, and Jeff Nivala. Structure-informed models for ionic current prediction in nanopore sequencing of expanded DNA alphabets. *Nucleic Acids Research*, 53(22): gkaf1389, 2025. doi: 10.1093/nar/gkaf1389.

Marcus Stoiber. RNA modified-base validation with all-5-mer ground truth. Oxford Nanopore Technologies, EPI2ME Data Releases, July 2026. URL https://epi2me.nanoporetech.com/rna-mod-validation-all5mer-2026.07/. Released July 21, 2026. Dataset version: 2026.07. Accessed September 12, 2026.

Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. RotatE: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=HkgEQnRqYQ.

Bo Wang, Hao Zhang, Taoyong Cui, Xiaoyu Wang, Jiangning Song, and Hao Xu. Evormd: integrating biological context and evolutionary rna language models for interpretable prediction of rna modifications. *Genome Biology*, 27(1):204, 2026.

Hongwei Wang, Weijiang Li, Xiaomeng Jin, Kyunghyun Cho, Heng Ji, Jiawei Han, and Martin D. Burke. Chemical-reaction-aware molecule representation learning, 2021. URL https://arxiv.org/abs/2109.09888.

Linshu Wang and Yuan Zhou. Mrm-bert: a novel deep neural network predictor of multiple rna modifications by fusing bert representation and sequence features. *RNA biology*, 21(1):340–349, 2024.

Ning Wang, Jiang Bian, Yuchen Li, Xuhong Li, Shahid Mumtaz, Linghe Kong, and Haoyi Xiong. Multi-purpose rna language modelling with motif-aware pretraining and type-guided fine-tuning. *Nature Machine Intelligence*, 6(5):548–557, 2024.

Daniel Worrall and Max Welling. Deep scale-spaces: Equivariance over scale. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://papers.nips.cc/paper/8956-deep-scale-spaces-equivariance-over-scale.pdf.

You Wu, Wenna Shao, Mengxiao Yan, Yuqin Wang, Pengfei Xu, Guoqiang Huang, Xiaofei Li, Brian D Gregory, Jun Yang, Hongxia Wang, et al. Transfer learning enables identification of multiple types of rna modifications using nanopore direct rna sequencing. *Nature communications*, 15(1):4049, 2024.

Jiancong Xie, Yi Wang, Jiahua Rao, Shuangjia Zheng, and Yuedong Yang. Self-supervised contrastive molecular representation learning with a chemical synthesis knowledge graph. *Journal of Chemical Information and Modeling*, 64(6):1945–1954, 2024.

Jiajia Xuan, Lifan Chen, Zhirong Chen, Junjie Pang, Junhong Huang, Jinran Lin, Lingling Zheng, Bin Li, Lianghu Qu, and Jianhua Yang. RMBase v3.0: decode the landscape, mechanisms and functions of RNA modifications. *Nucleic Acids Research*, 52(D1):D273–D284, 2024. doi: 10.1093/nar/gkad1070. URL https://academic.oup.com/nar/article/52/D1/D273/7416807.

Weijie Yin, Zhaoyu Zhang, Shuo Zhang, Liang He, Ruiyang Zhang, Rui Jiang, Gan Liu, Jingyi Wang, Xuegong Zhang, Tao Qin, et al. Ernie-rna: an rna language model with structure-enhanced representations. *Nature Communications*, 16(1):10076, 2025.

Chelsea Chen Yuge, Ee Soon Hang, Madasamy Ravi Nadar Mamtha, Shashikant Vishwakarma, Sijia Wang, Cheng Wang, and Nguyen Quoc Khanh Le. Rna-modx: a multilabel prediction and interpretation framework for rna modifications. *Briefings in Bioinformatics*, 26(1):bbae688, 2025.

Jingjie Zhang, Hanqun Cao, Zijun Gao, Yu Wang, Shaoning Li, Jun Xu, Cheng Tan, Jun Zhu, Chang-Yu Hsieh, Chunbin Gu, et al. Learning the ptm code through a coarse-to-fine mechanism-aware framework. *Nature Communications*, 2026.

Weihe Zhong, Ziduo Yang, and Calvin Yu-Chian Chen. Retrosynthesis prediction using an end-to-end graph generative architecture for molecular graph editing. *Nature Communications*, 14(1): 3009, 2023.

APPENDIX

A OPERATOR PROPERTIES AND APPROXIMATION LIMITS

A.1 PROOF OF PROPOSITION 1

Let $P = QQ^\top$ and $P_\perp = I - P$. Since $P_\perp z + QQ^\top z = z$, Eq. (4) gives $T_a(z) = z + Qb_a$. Hence

$$T_b(T_a(z)) = z + Q(b_a + b_b). \tag{16}$$

This is Eq. (4) with the offset in Eq. (6). Associativity follows from function composition, the empty path has zero offset, and $b_\pi = \sum_{t=1}^{L} b_{a_t}$. Valid paths satisfy the chemical applicability rule at every intermediate state.

A.2 PROOF OF PROPOSITION 2

Let $\hat{z}_0 = z_0$, $\hat{z}_t = T_{a_t}(\hat{z}_{t-1})$, and $e_t = \|\hat{z}_t - z_t\|_2$. The triangle inequality gives

$$e_t \le \|T_{a_t}(\hat{z}_{t-1}) - T_{a_t}(z_{t-1})\|_2 + \|T_{a_t}(z_{t-1}) - z_t\|_2 \tag{17}$$

$$\le \ell_t e_{t-1} + \epsilon_t. \tag{18}$$

Unrolling from $e_0 = 0$ yields Eq. (10). For $T_a(z) = z + Qb_a$, differences between two inputs are unchanged, so the Lipschitz constant is one.

A.3 SUBSPACE APPROXIMATION

The model preserves $P_\perp z$, so for a true precursor–product pair,

$$\|T_a(z_p) - z_m\|_2^2 \ge \|P_\perp(z_m - z_p)\|_2^2. \tag{19}$$

This is an irreducible error at fixed Q, obtained by orthogonally decomposing the residual. The bound isolates the component of a molecular change outside the learned subspace.

B BENCHMARK DATA PROVENANCE AND COVERAGE

Chemical source and standardization. The MODOMICS 2025 update supplies nucleoside structures and precursor–product relations, supplemented by explicitly curated primary literature (Sordyl et al., 2026). Retained structures have a defined canonical parent and stereochemistry. We sanitize valence and aromaticity, remove disconnected counterions, and deduplicate canonical isomeric SMILES while preserving formal charge, protonation, tautomers, and covalent phosphates. Ambiguous or inconsistent records are excluded; sources and exclusions are recorded by identifier.

ONT development resource. The ONT all-5-mer 2026.07 release provides raw signals and canonical reference sequences with designated-site annotations (Stoiber, 2026). We use the release's `full` population, containing up to the first 20 million reads per run and eight modified conditions: m^6A, inosine, m^5C, Ψ, Am, Cm, Gm, and Um, with canonical A/C/G/U controls. Each canonical-parent design contains 32 constructs, 256 designated sites, and 256 distinct centered 5-mers before read-level quality filtering.

Curlcake external resource. Curlcake RNA004 (ENA `PRJEB82528`) contains an unmodified control and seven modified conditions: m^6A, m^5C, hm^5C, ac^4C, $m^1\Psi$, Ψ, and m^5U (Cruciani et al., 2025). We use its RNA004 raw signals and canonical construct references. The incorporation of a modified nucleotide throughout a construct is retained as part of the cohort definition. Curlcake contributed to task development; configurations for the reported external evaluation are selected using ONT alone.

Action exposure and path coverage. All eight ONT targets have single-step canonical-parent paths. Table 6 records same-action source response exposure; all five action types occur in Track A chemical pretraining. The ONT path ablation evaluates chemical path supervision for single-step response transfer, while Track B evaluates multi-step composition.

Table 5: Resources and measured coverage of the primary benchmark.

Resource	Source and role	Coverage measures
MODOMICS	2025 update and curated experimental literature	106 states; 84 relations; 41 actions; 194 paths of lengths 2–4
ONT	2026.07 full population; 8 modified conditions and 4 controls	230,964,676 reads; 2,048 retained cells
Curlcake	PRJEB82528 RNA004; 7 modified conditions (3 shared, 4 unseen)	1,789 retained cells: 768 shared and 1,021 unseen

Table 6: Measured ONT cohort and action exposure. The source-identity count gives the number of other ONT modifications bearing the query action after excluding the target identity, in both LOMO and Double-OOD. Chemical graph exposure is separate from response-label exposure. All rows retain 256 cells and a canonical-parent path of length one.

Identity	Parent	Primitive action	Source identities
Am	A	Ribose 2′-O-methylation	3
Cm	C	Ribose 2′-O-methylation	3
Gm	G	Ribose 2′-O-methylation	3
I	A	Base N6 deamination	0
Ψ	U	N1-to-C5 glycosidic isomerization	0
Um	U	Ribose 2′-O-methylation	3
m^5C	C	Base C5 methylation	0
m^6A	A	Base N6 methylation	0

Counts refer to the frozen primary release. The chemical splits retain 62 action–substrate, 194 edge-seen, and 46 edge-unseen queries. All eight ONT modifications retain 256 contexts and support every declared few-shot budget.

C Nanopore Response Construction

Primary statistic and processing. The response is the probability of a substitution or deletion at the designated reference position, following modification-sensitive basecalling measurements (Liu et al., 2019). All RNA004 cohorts are basecalled with Dorado 2.0.1 and the canonical `rna004_sup@v6.0.0` model using A/C/G/T outputs from native POD5 signals. The Dorado aligner uses canonical references (U represented as T), with `sr` for 100-nt ONT constructs and `map-ont` for 2.2–2.7-kb Curlcake constructs.

Retain primary alignments with MAPQ $\geq$ 20 and aligned bases on both sides of the centered 5-mer. Exclude reference skips and ambiguous reference bases in the window. Deduplicate read identifiers within each run by MAPQ, then aligned length. Site selection uses neither base-quality nor modification-confidence thresholds.

For site j, condition m, and library b, let $N^{=}_{mjb}$, $N^{\neq}_{mjb}$, and N^{D}_{mjb} count canonical matches, substitutions (including ambiguous called bases), and CIGAR deletions at that position. Define

$$R_{jb}(m) = \frac{N^{\neq}_{mjb} + N^{D}_{mjb}}{N^{=}_{mjb} + N^{\neq}_{mjb} + N^{D}_{mjb}}. \tag{20}$$

Each deletion contributes once per covered position; insertions are excluded. Sites require at least 100 eligible observations in both conditions. Missing pairs are masked, and observed rates and differences are used directly.

Control matching and aggregation. Match modified and control measurements by canonical construct, position, parent base, and processing protocol. For a documented paired-library design, average the within-pair differences equally. Otherwise, first average site rates equally across eligible libraries within each condition and subtract those means. Let $\bar{R}_j$ denote the resulting condition

mean. For the common eligible positions $\mathcal{J}_{m,c}$ with canonical centered context c,

$$Y_{m,c} = \frac{1}{|\mathcal{J}_{m,c}|} \sum_{j \in \mathcal{J}_{m,c}} \left[\bar{R}_j(m) - \bar{R}_j(0_m) \right]. \tag{21}$$

Sites receive equal weight. ONT uses designated BED sites; Curlcake uses occurrences of the incorporated parent base. Additional modified sites follow each resource's incorporation design, defining cohort-conditioned responses.

Observations are aggregated into response cells before splitting. Inputs are canonical centered 5-mers; the four flanks define context groups, and the central parent remains visible.

D Response Splits and Evaluation

Response holdouts. LOMO holds out one modification's complete response profile. Double-OOD additionally partitions the 256 four-flank strings into five fixed groups of sizes 52, 51, 51, 51, and 51, shared across parent bases and modifications. Each trial excludes the query flank group from every source modification. Inner folds hold out a remaining modification and, for Double-OOD, a remaining context group. Outcome-fitted preprocessing and checkpoint selection are fitted within these inner folds. Response cells remain intact throughout splitting.

Few-shot episodes. The LOMO curves in Figure 4 use 20 draws paired across methods and all eight ONT identities. Each draw fixes 128 query contexts and selects nested supports from the remaining 128 contexts for $K \in \{0, 1, 5, 10, 50, 100\}$. The query pool is unchanged across budgets. Double-OOD episodes instead use the held-out flank group for queries and the other four groups for supports. Representations and base predictors remain frozen; only the residual head is fitted to support labels using Appendix F. Table 1 uses the full eligible zero-shot query sets.

Scoring and aggregation. Spearman uses average ranks for ties and a common observed-profile cohort with at least 20 query contexts and nonzero variance. Constant predictions receive $\rho = 0$; constant observed profiles remain eligible for error metrics. In Eq. (15), $\epsilon = 10^{-3}$ and query IQR uses linearly interpolated quartiles for scoring. Results average three independent training seeds (0, 1, 2), then the applicable context folds and support draws within each modification, followed by equal-weight modification means.

E Chemical Split and Query Construction

Templates and descriptors. Templates specify atom/bond edits, formal-charge changes, attachment sites, and reactant-side applicability. Atom correspondence and descriptors are curated from reaction annotations, excluding ambiguous maps. The outcome-independent descriptor vocabulary comprises reaction-type levels 1–3, input/output groups, introduced-group name/type, atom address, site class, and the local edit signature, with normalized aliases and an explicit missing-field token. Product identifiers and unchanged precursor environments are excluded. One descriptor is shared across compatible substrates; template matching and valence/stereochemistry checks determine applicability.

Multi-step training paths contain two to four edges without repeated states. Track A uses the shortest supported canonical-parent path, resolving ties by stable template IDs; the identity path is empty. The fusion descriptor d_π concatenates up to four ordered action descriptors, zero padded with an occupancy mask.

Action–substrate folds. Each fold withholds one deduplicated (G_p, a) relation group and every training path containing it. The action must remain supported on another precursor. Endpoint nodes may occur in other allowed relations.

Composition folds. Each fold withholds an ordered action-sequence group across starting states, including path supervision containing it as a contiguous subsequence, direct start–endpoint shortcuts, and alternate routes for the held-out start–endpoint pairs. Paths containing these excluded

shortcuts or routes are also removed. Edge-seen queries retain all constituent edges; edge-unseen queries additionally exclude their constituent edges and paths containing them, while retaining each primitive on another substrate. Exclusions precede normalization, training, and model selection; the two tiers use their resulting eligible cohorts.

Up to three inner folds follow the outer relation or composition rule, with fixed random group assignments. Selection uses macro endpoint error, falling back to the fixed configuration in Appendix F when no inner fold is eligible. Each variant is trained independently. Table 3 uses Track A chemical fits; Table 2 uses the corresponding Track B exclusions.

Candidate bank and ties. All methods rank the same 106-state bank by squared Euclidean distance in frozen molecular coordinates. Each fold fits normalization and reconstruction on its allowed training states; HAE rollouts are decoded into these coordinates.

For a target tied at positions $s, \dots, t$, use the expectation under uniform ordering within that tie:

$$\mathrm{RR} = \frac{1}{t-s+1} \sum_{j=s}^{t} \frac{1}{j}, \qquad \mathrm{R@5} = \frac{|[s,t] \cap [1,5]|}{t-s+1}. \tag{22}$$

Float64 distances within 10^{-10} of a tie block's minimum form a tie. Scores average queries within groups, then groups. The symbolic reference executes the curated edits and encodes the result with E_0; all learned predictors use the common eligible queries.

F Model, Baseline, and Adaptation Details

Frozen representations. CheMeleon supplies graph embeddings after bond message passing and mean aggregation, using the released checkpoint and Chemprop featurizer (Burns et al., 2025). For output $f(G) \in \mathbb{R}^d$, set $z_G = (f(G) - \mu)/\sigma$, where μ is the mean over unique allowed chemical training states and σ^2 their mean squared distance to μ, floored at 10^{-8}. RiNALMo-Giga (650M parameters, 33 blocks) encodes the canonical 5-mer with its official tokenizer; the frozen model's center-nucleotide representation after the final LayerNorm supplies 1280 features (Penić et al., 2025).

Architecture and initialization. The descriptor MLP is $d_a \to 64 \to J$ and the context adapter is $1280 \to 128 \to r$, both using GELU and final $\tanh$. The decoder has two width-w GELU hidden layers and a scalar output, shared between its two evaluations. Dictionary offsets start from $\mathcal{N}(0, 0.01^2)$. A Gaussian-initialized Q is orthonormalized by differentiable thin QR with positive diagonal at each forward pass.

The chemical grid is $r \in \{8, 16\}$, $J \in \{4, 8\}$, and $\lambda_{\mathrm{path}} \in \{0.25, 1\}$, with fallback $(8, 4, 0.25)$. All response models use $w \in \{128, 256\}$, learning rate $\{3 \times 10^{-4}, 10^{-3}\}$, and weight decay $\{10^{-4}, 10^{-2}\}$. Each grid has eight configurations; ablated components remain fixed as specified.

Optimization and selection. AdamW uses $(\beta_1, \beta_2) = (0.9, 0.999)$, epsilon 10^{-8}, and gradient clipping at 1. Chemical training uses learning rate 10^{-3}, weight decay 10^{-4}, and batches of 64 edges and 64 paths sampled with replacement, uniformly over actions or composition groups before their members. The path term is zero when no training path is available.

Chemical training runs for at most 5,000 steps at a fixed rate, validating every 50 steps with patience 10. Response training uses a 3,000-step cosine schedule, validating every ten steps with patience 50. Selection minimizes inner macro endpoint error for chemistry and maximizes inner macro Spearman for responses; ties favor fewer parameters, then grid order. Final fits use three initializations and the median selected inner stopping step. Without eligible inner chemical validation, the fallback runs for 5,000 steps. Track A chemical selection uses outcome-free relation folds independently of response selection.

Freeze the chemical modules before fitting B_ψ, D_ω. The value loss is

$$\mathcal{L}_{\mathrm{MSE}} = \frac{1}{|\mathcal{M}_{\mathrm{tr}}|} \sum_{m \in \mathcal{M}_{\mathrm{tr}}} \frac{1}{|\mathcal{C}_{m,\mathrm{tr}}|} \sum_{c \in \mathcal{C}_{m,\mathrm{tr}}} (\widehat{\Delta y}(c,m) - Y_{m,c})^2. \tag{23}$$

For modification m, let $\mathcal{P}_m$ contain unordered pairs of distinct training contexts with unequal observed responses. Write $\hat{y}_i = \widehat{\Delta y}(c_i, m)$, $\sigma_{ij} = \operatorname{sign}(Y_{m,c_i} - Y_{m,c_j})$, and let v be the macro mean of within-modification sample variances in the current training fold. The pairwise objective is

$$\mathcal{L}_{\text{rank}} = \frac{v}{|\mathcal{M}_{\text{tr}}|} \sum_{m \in \mathcal{M}_{\text{tr}}} \frac{1}{|\mathcal{P}_m|} \sum_{(i,j) \in \mathcal{P}_m} \log\left(1 + \exp\left[-\frac{\sigma_{ij}(\hat{y}_i - \hat{y}_j)}{\sqrt{v}}\right]\right), \tag{24}$$
$$\mathcal{L}_{\text{response}} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{rank}} \mathcal{L}_{\text{rank}}, \qquad \lambda_{\text{rank}} = 1.$$

Pair labels and scales use current-fold training outcomes; identities without unequal pairs contribute zero to the ranking term. Use full-batch response updates on the eligible training cells and predict in probability-difference units. For Curlcake external evaluation, configurations and stopping steps are selected by ONT-only inner folds before refitting on all ONT training responses.

Source-only calibration. For Tables 1 and 4, fit a nonnegative, zero-intercept scalar to source out-of-fold predictions by minimizing modification-macro MAE. Freeze it across the three final query predictions and compute both metrics from these calibrated predictions. Positive scaling preserves ranks; zero scaling follows the constant-prediction convention.

Few-shot residual features. Write the decoder as $D_\omega(c, s) = a^\top H_\omega(c, s) + b$. For RiboCTE, HAE-C, and the two-branch controls, use the feature difference $q(x, m) = H_\omega(c_x, \hat{s}_m) - H_\omega(c_x, s_0)$, evaluated with each model's decoder and contextual states. For direct-fusion models, use the penultimate readout feature. Fit PCA separately for each method on its centered response-training features. Retain at most 32 components whose training standard deviation exceeds 10^{-6} times that of the leading component, then scale each retained coordinate by its training standard deviation. Discard lower-variance directions and zero-pad to 32 dimensions; if the training features are constant, use 32 zeros. Append a constant to obtain $\phi(x, m) \in \mathbb{R}^{33}$. The projection and scales are fitted on response-training observations.

For $K > 0$, fit the ridge correction

$$w_{m,K} = \arg\min_w \frac{1}{K} \sum_{(x,y) \in \mathcal{S}_{m,K}} \left[y - \hat{y}_0(x, m) - w^\top \phi(x, m)\right]^2 + \lambda_{\text{fs}} \|w\|_2^2. \tag{25}$$

Regularize all coordinates, including the intercept. Select $\lambda_{\text{fs}} \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$ by average inner-episode MSE across nonzero budgets, favoring larger penalties on ties. Predict $\hat{y}_K = \hat{y}_0 + w_{m,K}^\top \phi$, with $w_{m,0} = 0$.

Few-shot learning curves. Figure 4 uses MSE-trained bases without scalar calibration. Model features q and RiNALMo features h are separately centered and scaled by their source RMS norms, concatenated, and transformed by source-fitted 32-component PCA and ridge coordinates. For $K \geq 5$, target-support leave-one-out validation selects from the same six penalties, favoring larger values on ties; $K = 1$ uses per-budget source selection and $K = 0$ the frozen base. All methods share this protocol and the fixed half-context query pools in Appendix D.

F.1 Baselines and Controlled Comparisons

Molecular endpoint fusion. CheMeleon uses the pretrained encoder described above. Mol-JEPA uses the authors' pretrained structure encoder and its 512-dimensional graph representation (Rottach et al., 2026). Both backbones are frozen in evaluation mode. For an encoder E_b, define

$$e_m = [E_b(G_0); E_b(G_m); E_b(G_m) - E_b(G_0); d_{\pi_m}].$$

Keep binary descriptors, path-occupancy masks, and Morgan bits at their original $0/1$ values, and retain fingerprint differences in $[-1, 1]$. For continuous embedding blocks, subtract response-training feature means and divide by response-training standard deviations; use scale 1 for columns with standard deviation below 10^{-6}. Apply the same fitted transformation to validation, support, and query inputs. Project e_m and h_x separately to 128 dimensions with affine/GELU layers, concatenate, and apply two width-w GELU layers followed by a scalar linear output. The same architecture and response grid are used for CheMeleon, Mol-JEPA, HAE-E, and Morgan radius-2, 2048-bit fingerprints with chirality. Context-only uses only the context branch and the same readout.

HAE chemical model. HAE uses continuous descriptor-conditioned actions parameterized by a matrix exponential. Let $A_H : \mathbb{R}^d \to \mathbb{R}^{r_H}$ and $C_H : \mathbb{R}^{r_H} \to \mathbb{R}^d$ be two-hidden-layer MLPs of width w_H. An action MLP $d_a \to 64 \to r_H^2$ produces a matrix $S_H(d_a)$, and $H_a = \exp(S_H(d_a))$ acts linearly on the latent state. Reconstruct continuous E_0 vectors with squared error, without an output sigmoid or latent unit-norm normalization.

For each allowed state/path, optimize

$$\begin{aligned}\mathcal{L}_H = \mathbb{E}_z\|C_H(A_H(z)) - z\|_2^2 \\ + \lambda_H \mathbb{E}_{(z_0,\pi,z_L)}\|H_\pi A_H(z_0) - A_H(z_L)\|_2^2 \\ + \mathbb{E}_{(z_0,\pi,z_L)}\|C_H(H_\pi A_H(z_0)) - z_L\|_2^2,\end{aligned} \tag{26}$$

where $H_\pi = H_{a_L} \cdots H_{a_1}$. Use one-step relations and the same permitted multi-step paths, with unit weight for each available path-length regime; gradients reach both HAE endpoint encodings. The eight chemical configurations are $r_H \in \{8, 16\}$, $w_H \in \{128, 256\}$, and $\lambda_H \in \{0.1, 1\}$, with the smaller setting as fallback. Use the common chemical optimization/stopping budget and select by decoded endpoint error. HAE-E supplies $A_H(z_0), A_H(z_m)$ to the common readout; HAE-C reuses the same frozen chemical fit. Track B decodes $H_\pi A_H(z_0)$ with C_H and therefore has a single HAE entry.

HAE-C response decoder. Let $v_x^H = B_H(h_x) \in \mathbb{R}^{r_H}$ and $c_x^H = [h_x; z_0]$. Define

$$\begin{aligned}s_0^H(x) = A_H(z_0) + v_x^H, \qquad \hat{s}_m^H(x) = H_{\pi_m} s_0^H(x), \\ \widehat{\Delta y}_H(x, m) = D_H(c_x^H, \hat{s}_m^H(x)) - D_H(c_x^H, s_0^H(x)).\end{aligned} \tag{27}$$

Use $B_H : 1280 \to 128 \to r_H$ with GELU and final $\tanh$, and a shared decoder D_H with two width-w GELU hidden layers and a scalar linear output. Freeze the chemically fitted HAE modules and RNA encoder, and fit only B_H, D_H on seen-modification responses with the common response objective in Eq. (24), response grid, and stopping protocol. Apply this interface to both response OOD settings, the common few-shot residual protocol using its penultimate feature difference, and ONT-selected external evaluation.

Chemical controls. The action-conditioned MLP uses $\hat{z}' = z + \mathrm{MLP}([z; d_a])$ with two GELU hidden layers, recursively applied along a path. Its eight chemical settings combine width $\{128, 256\}$, weight decay $\{10^{-4}, 10^{-2}\}$, and path weight $\{0.25, 1\}$; it uses the same edge/path targets and stopping budget. The no-path variant fits the RiboCTE offset dictionary from initialization with $\lambda_{\text{path}} = 0$; it is not warm-started from the full predictor.

Response controls. The exact-endpoint control retains the chemical model's frozen Q and refits B_ψ, D_ω with $\hat{s}_m = Q^\top z_m + v_x$. Offset controls retain Q and parent coordinates: rotation applies a shared orthogonal matrix; norm-preserving reassignment deranges primitive-action offsets and rescales each to the receiving action's original norm. Each uses one outcome-independent assignment shared across three refitted seeds, with a common backend.

For generic bilinear fusion, use the same e_m and preprocessing as CheMeleon and define a scalar predictor

$$\hat{y}_{\text{bil}} = f_\theta([p_e(e_m); p_h(h_x); (Ue_m) \odot (Vh_x)]). \tag{28}$$

Here p_e, p_h are 128-dimensional affine/GELU projections, U, V produce 32 interaction coordinates, and f_θ has two width-w GELU layers and a scalar output. Use the common response grid.

G SUPPLEMENTARY RESULTS

Action assignments. Table 7 compares native offsets, rotation, and norm-preserving reassignment with independently refitted response modules. Perturbations preserve action sharing and use one outcome-independent draw shared across training seeds.

Table 7: Action-assignment comparisons with the common response objective. Entries are macro Spearman ρ, averaged over three training seeds; bold marks column-wise best values, including ties. Native and rotated results are shared with Table 3.

Action offsets	LOMO	Double-OOD
Native	**0.216**	**0.184**
Rotated	0.122	0.099
Permuted, original norms	0.204	0.149