SOURCE-ANCHORED FLOW MATCHING FOR LOCAL MOLECULAR EDITING

Anonymous authors
Paper under double-blind review

ABSTRACT

Lead optimisation requires local edits that improve a property while preserving the parent scaffold. Property-conditioned *de novo* generators steer samples toward a desired label without anchoring them to a given parent, so scaffold retention can only be enforced after sampling. We introduce a recipe for local property editing in the frozen latent space of a pretrained SMILES encoder. Sparse Cross-Improvement Neighbours (SCIN) mines directed source-to-target pairs whose observed property gain exceeds a margin within a latent-similarity neighbourhood, and Source-Anchored Flow Matching (SAFM) learns a velocity field conditioned on the source embedding and the property labels. Classifier-free guidance and an inference-time edit budget control the fidelity–improvement trade-off without retraining. We evaluate on ZINC250K and MOSES for QED maximisation, then transfer the same recipe to cyclic-peptide permeability on CycPeptMPDB across the PAMPA and Caco-2 assays. SAFM leads faithful editing on both modalities, at 36.3% on CycPeptMPDB against 0.3% for unconstrained generate-and-filter, and at 9.8% on ZINC250K. Descriptor analysis of the peptide edits recovers a consistent aromatic-to-aliphatic exchange that medicinal chemistry associates with improved passive permeation. The recipe transfers across chemical modalities with no domain-specific machinery beyond a frozen encoder and modality-appropriate validity gates. *Code will be released upon acceptance.*

1 INTRODUCTION

Molecular optimisation is predominantly an editing problem. Given a starting compound with an expected functionality and a measured property, the task is to propose a distinct but structurally similar analogue that improves that property while satisfying lightweight validity and scaffold-retention constraints, as in hit-to-lead optimisation (Bleicher et al., 2003; Nippa et al., 2025) and matched-molecular-pair analysis (Griffen et al., 2011). Existing approaches each cover part of this. Property-conditioned *de novo* generators explore chemical space effectively (Tang et al., 2025) without anchoring to a known parent, so scaffold similarity must be imposed by post-hoc filtering or nearest-neighbour selection. Latent gradient ascent on a learned proxy begins from the parent but has no explicit locality prior and can move off-manifold. Curated matched-pair datasets that supervise useful edits directly are rarely available at the scale of modern labelled corpora.

We argue that whenever a labelled property dataset exists, directed improvement trajectories can be mined from it and transport can be learned conditioned on the source molecule itself. Local editing is then source-anchored transport between real improvement neighbours, the fidelity–improvement trade-off becomes part of the model instead of a filter applied after *de novo* sampling, and a single inference-time knob, the edit budget, dials structural displacement without retraining.

The recipe is modality-agnostic by construction, and our experiments test whether it transfers across substantially different molecular regimes. We evaluate on ZINC250K (Sterling and Irwin, 2015) and MOSES (Polykovskiy et al., 2020) for QED (Bickerton et al., 2012) ↑, then transfer the same procedure to cyclic-peptide permeability on CycPeptMPDB (Li et al., 2023). Pair mining, flow-matching formulation, guidance mechanism, decoding and evaluation are held fixed; only the frozen encoder and the modality-specific validity gates change.

Contributions.

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.

1. **SCIN pair mining**, which builds supervised improvement pairs from any labelled dataset by combining a property-margin filter with a latent-similarity neighbourhood and optional structure constraints, requiring no additional experiments (Section 3.2).
2. **Source-Anchored Flow Matching (SAFM)**, a flow-matching editor whose velocity field is anchored to a real parent latent, admits classifier-free guidance on labels and on the source, and exposes the edit budget as a fidelity dial (Section 3.3).
3. A **budget-aware evaluation framework** reporting Pareto curves, faithful-improvement and improvement rates across edit budgets under an explicit similarity floor, comparable across modalities and oracle types (Section 4.2).
4. An **empirical evaluation** on small molecules and cyclic peptides, and a **chemical characterisation** for the cyclic peptides. SAFM leads faithful editing on both modalities, at 36.3% on CycPeptMPDB and at 9.8% on ZINC250K. Descriptor analysis of the peptide edits recovers a consistent aromatic-to-aliphatic exchange that medicinal chemistry associates with improved passive permeation.

2 Related Work

Molecular editing and latent-space optimisation. Constrained property optimisation on ZINC250K (Gómez-Bombarelli et al., 2018; Jin et al., 2018) is the standard benchmark for editing under explicit similarity constraints, and recent flow-based latent editors (Lobo et al., 2026) characterise the fidelity–improvement trade-off across guidance budgets. Related lines include latent-space optimisation (Eckmann et al., 2022), heuristic latent interpolation (Zhang et al., 2024), retrieval-augmented matched-pair transformations (Pan et al., 2026), Bayesian optimisation in molecular latent spaces (Griffiths and Hernández-Lobato, 2020), and graph-based methods (You et al., 2018) that act on discrete structures directly. SAFM differs in combining continuous latent flow matching, supervision from mined real improvement pairs, and an edit budget defined by partial traversal of the learned path.

Flow matching and guidance. Flow matching learns a time-dependent vector field transporting a base distribution toward a data distribution along a prescribed probability path (Lipman et al., 2022), and extends to conditional generation by conditioning that field on labels or properties. Classifier-free guidance (Lipman et al., 2024) adjusts conditioning strength at inference without a separate classifier. Discrete and mixed-state formulations (Chen et al., 2025) carry these ideas to sequences; SAFM instead operates in a continuous pretrained latent space and conditions on both the source embedding and the property labels.

Cyclic peptides. Permeability is a longstanding bottleneck for therapeutic cyclic peptides (Dougherty et al., 2019). CycPeptMPDB (Li et al., 2023) aggregates PAMPA, Caco-2, MDCK and RRCK measurements whose protocols and endpoints are not interchangeable, so assay identity must be accounted for in training and evaluation. Peptide-specialised encoders and permeability predictors (Feller et al., 2026; Zhang et al., 2026; Wang et al., 2025) model the endpoint, while editing-oriented systems include estimator-controlled proposers (Aerts et al., 2025) and preference-aligned geometric diffusion (Sun et al., 2026). Discrete generators with inference-time guidance (Tang et al., 2025), sequence flow matching (Chen et al., 2025) and equivariant 3D generators (Li et al., 2024) target permeability without requiring retention of a specified parent, and we evaluate both editing and *de novo* baselines under an explicit similarity floor.

3 Method

Figure 1 summarises the framework. SCIN mines directed, property-improving neighbour pairs that define a supervised local transport problem; SAFM learns source- and property-conditioned transport between their frozen embeddings; and a domain-specific decoder maps edited latents back to molecules. At inference, an edit budget truncates the transport path, after which one valid decoded sequence is selected and evaluated for source-relative feasibility and property change. We next define the task, representation, SCIN supervision, SAFM objective and inference, decoding, property

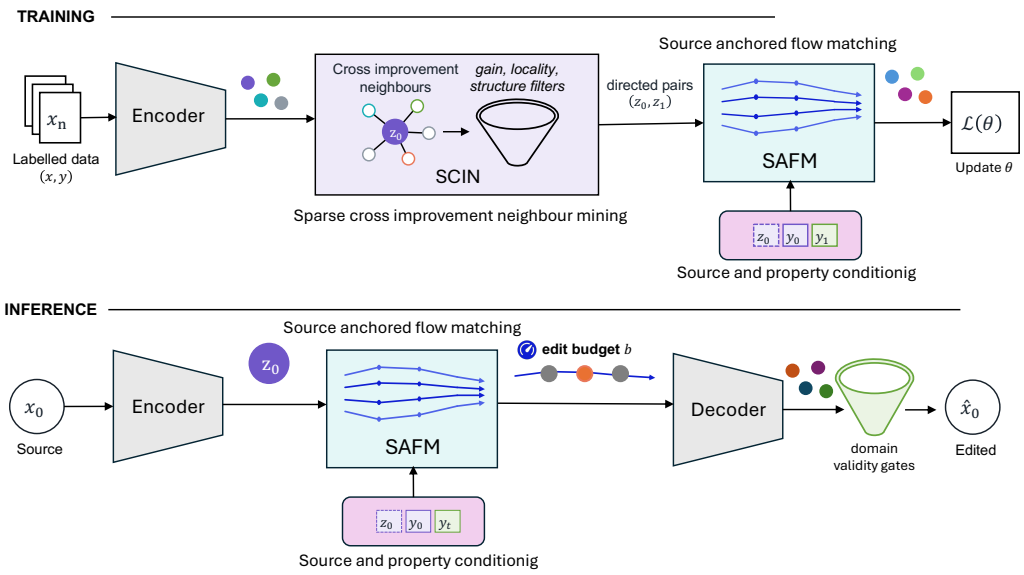

Figure 1: **SAFM framework for local molecular editing.** SCIN mines source-specific improvement pairs that establish the data-driven basis for anchoring. SAFM learns transport over these pairs, while the edit budget controls the trade-off between property improvement and structural fidelity.

scoring, and computational complexity; domain-specific configurations and reproducibility details appear in Experiments and Appendix K.

Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{n}$ be a labelled molecular dataset, where $x_i \in \mathcal{X}$ is a Simplified Molecular Input Line Entry System (SMILES) string and $y_i \in \mathbb{R}$ is its property value. We define the oriented utility $u(x_i) = y_i$, using QED for small molecules and PAMPA permeability for cyclic peptides; larger values are preferred in both settings. For an edit $\hat{x}$ of source x, the reported change is $\Delta u(x, \hat{x}) = u(\hat{x}) - u(x)$, obtained by post-hoc analytic rescoring or the learned proxy in Section 3.5.

For a held-out source x, the task is to generate a distinct, valid candidate $\hat{x}$ that maximises Δu while preserving source structure. Candidates must satisfy domain-specific validity gates and chemistry-light edit constraints, and structural fidelity is measured by Morgan-fingerprint Tanimoto similarity (radius 2; 2048 bits). At inference, the edit budget $b \in (0, 1]$ limits the fraction of the learned transport path traversed without retraining. We summarise the resulting fidelity–improvement trade-off with a faithfulness metric combining Δu, structural similarity, and constraint satisfaction.

3.1 Frozen Latent Representation

All learning takes place in the frozen embedding space of a domain-appropriate pretrained SMILES encoder, $E : \mathcal{X} \to \mathbb{R}^d$. We keep the parameters of E fixed during velocity-network, decoder, and oracle training. We use ChemBERTa-77M-MLM (Chithrananda et al., 2020) with mean pooling, yielding $d = 512$, for the small-molecule benchmarks, and PeptideCLM-2 MLM-large (Feller et al., 2026) with mean pooling, yielding $d = 1024$, for the cyclic-peptide setting. Encoder-scale ablations are reported in the appendix B.

3.2 Sparse Cross-Improvement Neighbours (SCIN)

Sparse Cross-Improvement Neighbours (SCIN) constructs supervised local improvement trajectories from labelled molecular data without paired edit annotations. Using training-fold embeddings only, it retrieves a cosine-similarity k-nearest-neighbour pool $\mathcal{N}_k(i)$ for each source i (default $k = 50$) and retains directed pairs $(i \to j)$ satisfying $y_j - y_i \geq \delta_{\min}$, $\cos(z_i, z_j) \geq \tau$ (default $\tau = 0.50$), and $g(x_i, x_j) = 1$. Here, $\delta_{\min}$ is the minimum gain on the target's natural scale, while g enforces equal monomer length for peptides and no additional small-molecule constraint. Retain-

ing at most three targets per source limits redundant supervision while preserving observed local improvement directions. The resulting pair set is:

$$\mathcal{P}_{\text{SCIN}} = \{(i,j) : j \in \mathcal{N}_k(i),\ y_j - y_i \geq \delta_{\min},\ \cos(z_i, z_j) \geq \tau,\ g(x_i, x_j) = 1\}\,. \tag{1}$$

SCIN assumes that useful molecular edits follow local latent directions associated with observed property gains. Pairs are mined independently within each training fold, excluding validation and test molecules, and 10% of the mined training pairs are held out for velocity-network early stopping. In the random-pair ablation, each source instead samples uniformly from a global candidate pool while retaining the positive improvement margin and applicable domain filter, but not the SCIN cosine threshold or exact pair count. Sensitivity to k and τ is reported in Appendix C.

3.3 Source-Anchored Flow Matching (SAFM)

Training objective. SAFM learns transport between observed SCIN improvement neighbours rather than from generic noise. For a pair (z_0, z_1) with normalised labels (y_0, y_{target}), we sample $t \sim \mathcal{U}(0,1)$, define $z_t = (1-t)z_0 + tz_1$, and use the constant target velocity $z_1 - z_0$. The source z_0 is both the path origin and an optional conditioning input, the labels specify the desired property transition, and SCIN supplies the supervised trajectories. Unlike the property-conditional de novo baseline, SAFM therefore remains anchored to the source.

A source-conditional velocity network $v_\theta(z_t, t, z_0, y_0, y_{\text{target}})$ is trained under the flow-matching loss

$$\mathcal{L}(\theta) \;=\; \mathbb{E}_{(z_0,\, z_1,\, y_0,\, y_{\text{target}}),\ t}\ \big\| v_\theta(z_t, t, z_0, y_0, y_{\text{target}}) \;-\; (z_1 - z_0) \big\|_2^2. \tag{2}$$

The velocity network is a three-layer MLP with hidden dimension 512, SiLU activations, and a sinusoidal time embedding of dimension 64. Its input concatenates the current state z_t, the source embedding z_0 (when source-conditional), time, and the normalised source and target labels. Classifier-free guidance is enabled at training time by independently dropping the label conditioning with probability $p_{\text{cfg}} = 0.1$ and, when source-conditional, dropping z_0 with probability $p_{\text{src}} = 0.2$; each dropped signal is replaced by a learned null token.

Guided inference. For a test source x_k, let z_k be its standardised embedding and let y_k be its stored property label. We set $y_{\text{target}} = y_0 + \Delta y_{\text{goal}}$. Starting from $z^{(0)} = z_k$, we integrate the learned velocity field with a first-order Euler solver for $N = 50$ steps. The size h of each step is controlled by the edit budget $b \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ as $h = b/N$. Because training independently drops property and source conditioning, inference forms guided velocities from the fully conditioned, null-label, and null-source predictions:

$$\widetilde{v}^{(n)} = v_{\varnothing}^{(n)} + \gamma \left(v_\theta^{(n)} - v_{\varnothing}^{(n)} \right) \tag{3}$$

The guidance scale γ (default 1.5) controls property steering and source anchoring; its sensitivity is reported in Appendix D. Larger budgets extend the integration horizon, typically trading structural fidelity for predicted improvement. The delta target over the source is set to 0.5 for QED, and 3 log units for permeability.

3.4 Autoregressive SMILES decoder

Because the encoder is not directly invertible, we learn a conditional autoregressive decoder D_ϕ that maps a latent representation back to a SMILES sequence. Let $x = (x_1, \ldots, x_L)$ denote the tokenised canonical SMILES of a molecule and let z be its standardised frozen representation. The decoder factorises the conditional sequence distribution as

$$p_\phi(x \mid z) = \prod_{\ell=1}^{L} p_\phi(x_\ell \mid x_{<\ell}, z)\,. \tag{4}$$

$$\mathcal{L}_{\text{dec}}(\phi) = -\mathbb{E}_{x \sim \mathcal{D}_{\text{train}}} \left[\sum_{\ell=1}^{L} \log p_\phi(x_\ell \mid x_{<\ell}, z) \right]. \tag{5}$$

The decoder is a Transformer with $d_{\text{model}} = 256$, four layers, four attention heads, maximum decode length 256, and dropout 0.1. Once trained, the decoder is frozen for all editing methods in a given experiment, so that decode quality affects every method equally. The reconstruction accuracy of the decoder is reported in the appendix E.

Given an edited latent $\hat{z}$, the decoder forms eight candidates $\mathcal{G}(\hat{z})$ from one greedy decode and $K = 7$ stochastic samples generated with temperature 0.8 and top-p 0.9. The earliest candidate satisfying the domain validity check is retained, falling back to the first candidate if none passes. Source-relative edit constraints, distinctness, Tanimoto similarity, and property change are computed after selection for evaluation.

Validity and edit constraints are domain-specific. Small molecules must yield RDKit-parseable SMILES and may additionally be subject to MOSES filters, Morgan-Tanimoto scaffold similarity. Peptides must be cyclic, contain at least 20 heavy atoms, have molecular weight at least 300, include at least three amide bonds and one backbone macrocycle, preserve the backbone-ring count, and equal amide-bond count; rings with fewer than eight atoms are treated as side-chain rings. Both domains enforce distinctness by canonical-SMILES inequality with the source.

3.5 Property oracle

We fit a regressor $R_\psi : \mathbb{R}^d \to \mathbb{R}$ on standardised frozen training-set embeddings to predict the target property. For a training fold $\mathcal{D}_{\text{train}}$, the predictor is learned by minimising

$$\mathcal{L}_{\text{oracle}}(\psi) = \mathbb{E}_{(x,y)\sim\mathcal{D}_{\text{train}}}\left[\left(R_\psi(E(x)) - y\right)^2\right]. \tag{6}$$

The regressor has two hidden layers of width 256 and dropout 0.1, with early stopping determined by validation error. The same fold-specific proxy provides differentiable scores for optimisation methods and post-decoding property changes, and is fitted without validation or test labels. For cyclic peptides, an independent public permeability predictor (Zhang et al., 2026) supplies an additional post-hoc robustness check. Held-out oracle performance is reported in Appendix F.

4 Experiments

The same SCIN construction, SAFM objective, classifier-free conditioning scheme, budget-controlled inference, decoding procedure, and evaluation definitions are used across small molecules and cyclic peptides. Only the frozen encoder, target-property interface, feasibility constraints, latent dimensionality, and property-scale thresholds vary by domain. We first describe the datasets and splits, then specify the property oracle, decoder, SAFM operating points, and candidate-generation settings before introducing baselines and evaluation metrics.

4.1 Datasets and baselines

Small-molecule benchmarks. ZINC250K (Sterling and Irwin, 2015) uses an $81/9/10$ scaffold-aware train/validation/test partition. QED↑ is the primary small-molecule benchmark and uses five scaffold folds by three seeds (CV15). MOSES uses its recommended split and the same encoder configuration, and is reported as directional replication evidence in appendix G.

Cyclic peptides (domain case study). CycPeptMPDB (Li et al., 2023) is evaluated separately for PAMPA and Caco-2 assays because their protocols and endpoints are not directly interchangeable. PAMPA is the primary peptide case study. Experiment on Caco-2 assay A.1 tests whether the same framework can be re-instantiated under different assay labels and data regimes. We use five-fold cross-validation with three seeds, fit every learned oracle and editor on the corresponding training fold only.

Controlled baselines. We compare SAFM with two alternative property-improvement strategies discussed in related work: latent optimisation of target utility (M1) and property-conditioned generate-and-filter (M2). On the matched QED and cyclic-peptide CV15 evaluations, SAFM, M1,

and M2 use the same held-out sources, frozen encoder and decoder within each domain, candidate-count limit, evaluation protocol, and post-generation validity and edit constraints. These controlled comparisons reduce confounding from differences in architectures and evaluation; they are not reproductions of specific published methods. We additionally report two diagnostic references: SCIN-target retrieval (M3) and interpolation towards a SCIN target (M4). M3 and M4 rely on training-derived targets and are not upper bounds. Finally, we ablate SCIN supervision using random latent-neighbour pairs (A1), to demonstrate the importance of the SCIN supervision. We also compare against a version without the optional source conditioning (SAFM without z_0 cond.). The details regarding the implementation of these baselines are provided in appendix A.2. External comparisons with MoltenFlow-style optimisation Lobo et al. (2026) on ZINC250K and fold-retrained C2PO Aerts et al. (2025) on CycPeptMPDB are reported separately in appendix A.3, as their model stacks and search procedures differ from the controlled in-stack baselines.

4.2 Evaluation

We present results from two main evaluations: (i) quantitative assessment of the proposed strategy to produce local edits, and (ii) downstream chemical characterisation of the edits in the peptide domain case study. For each seed–fold cell, test molecules were assigned to low, medium, and high property bins using tertile boundaries computed from the corresponding training fold. Within each bin, molecules were sorted by source property and the lowest-valued 100 were selected, yielding up to 300 held-out sources per cell.

4.2.1 Quantitative assessment

The primary endpoint is Faithful-Improve rate: the fraction of held-out sources for which the validation-fixed selected candidate is distinct, valid, constraint-passing, has non-negative target-property improvement, and has Morgan Tanimoto similarity at least 0.7. This endpoint operationalises the paper's central question: whether an editor can improve a molecule *locally*, rather than generate a different high-scoring molecule. We report the target-property lift, structural similarity to source and constraint pass rate breakdown in expanded results section in appendix A. For SAFM, we vary the ODE traversal budget and report Faithful-Improve rate and source similarity at each setting to demonstrate the fidelity–improvement trade-off. We also compare the target property improvement at different bins (low, medium, high) of property values between the source and the edited candidates for the SAFM method to analyse if the edits are applied where it matters the most.

4.2.2 Chemical characterisation of peptides

Topological descriptors. We characterise the edits with descriptors computed independently of the learned oracles. Topological descriptors are calculated with RDKit from canonical SMILES for every source and edited peptide and compared pairwise. Pairs whose edit decodes to the source are excluded, and the remainder are classified by comparing stereochemistry-stripped canonical SMILES and molecular formulae, which separates stereocentre inversions from constitutional isomers and from composition changes. Ring systems are enumerated per molecule excluding the backbone macrocycle, so that the rings removed and introduced by each edit can be counted across the set. Each macrocycle is additionally cleaved at its backbone amide bonds and the resulting residue multisets are compared, which separates substitutions that change the residue set from rearrangements that conserve it.

Medicinal chemistry rule score. Six descriptors are scored against a favourable direction taken from the literature before any edit is inspected, namely an increase in the fraction of sp3 carbon and in the N-methyl amide count, and a decrease in aromatic ring count, rotatable bonds, hydrogen bond donors and topological polar surface area. A pair scores $+1$ for every rule it satisfies and -1 for every rule it violates, with changes below a numerical tolerance held at zero. Per-rule counts are assessed with a two-sided exact sign test over the pairs in which the descriptor changes, and the composite sums the retained rules within each molecule.

Conformational descriptors. Conformational descriptors are computed on ensembles generated with ETKDGv3 using macrocycle torsion sampling and RMS pruning at $0.5\,\text{Å}$, minimised with MMFF94, restricted to conformers within 10 kcal mol^{-1} of the ensemble minimum and weighted

Table 1: **Faithful-Improve at $b = 1.0$.** Mean $\pm$ standard deviation over 15 matched seed–fold cells. Faithful-Improve requires a distinct, valid, constraint-passing edit with non-negative property change and Morgan Tanimoto ≥ 0.7. ZINC uses analytic QED; peptide edits are scored by both the in-house oracle and PeptiVerse.

Method	ZINC QED	PAMPA in-house	PAMPA PeptiVerse
Latent gradient (M1)	0.001 ± 0.003	0.061 ± 0.012	0.044 ± 0.014
Generate-and-filter (M2)	0.000 ± 0.000	0.003 ± 0.003	0.003 ± 0.003
SCIN target retrieval (M3)	0.000 ± 0.000	0.040 ± 0.013	0.023 ± 0.015
Latent interpolation (M4)	0.070 ± 0.016	0.042 ± 0.012	0.028 ± 0.013
SAFM, random SCIN (A1)	0.058 ± 0.021	0.130 ± 0.024	0.079 ± 0.019
SAFM, without z_0 cond.	0.055 ± 0.013	$\mathbf{0.363 \pm 0.036}$	$\mathbf{0.268 \pm 0.036}$
SAFM	$\mathbf{0.098 \pm 0.019}$	$\mathbf{0.363 \pm 0.036}$	0.244 ± 0.037

by Boltzmann populations at 298 K. Three-dimensional polar surface area is the solvent accessible surface of nitrogen, oxygen and polar hydrogen atoms computed with `rdFreeSASA`, and amide hydrogen exposure isolates the donor hydrogen contribution to that total.

5 Results

5.1 Faithful improvement at a common budget

Table 1 is deliberately restricted to matched in-stack methods and a single operating point to make a controlled baseline comparison. It therefore tests the local editing rule without selecting a different budget for each source or mixing external model stacks. On ZINC, source-anchored SAFM has the highest Faithful-Improve rate (0.098). Both Generate-and-filter and Latent gradient achieve high delta property improvement (0.292, 0.184) but at the cost of structural similarity to the source (0.113, 0.269), and constraint pass rate (0.461, 0.332). When compared to SAFM with random SCIN ($\Delta\text{QED} = 0.046$), SAFM without z_0 cond. optimises for property improvement ($\Delta\text{QED} = 0.130$), and achieves a lower structural similarity to source (0.462 vs 0.533). Therefore, in the ZINC case source conditioning acts as an extra regulariser for structural similarity. For cyclic peptides, both SAFM variants outperform the baselines and achieve the highest Faithful-Improve rate as per both the PAMPA scorers. The source conditioning here does not make a large difference because of the additional structural similarity constraints imposed during the SCIN construction when compared to ZINC. The expanded results are reported in appendix A.

Figure 2 illustrates the binned property improvement at budget $b = 1.0$ for SAFM to analyse if the edits are applied where it matters the most. Source and SAFM-edited property values are grouped by held-out source property bins and aggregated over CV15. The figure shows that SAFM balances the fidelity–improvement trade-off by making bigger jumps in the low property bins while showing modest improvement in the medium and high property value bins. Figure 3 further illustrates how the edit budget, a tunable knob provided by SAFM at inference, can be exploited to balance this trade-off. This is particularly useful to ground the editing process in the use cases where the domain specific validity gates are unknown.

5.2 Chemical characterisation of peptide edits

This section asks whether the improvement reported by the oracles has a chemical basis. Passive permeation in cyclic peptides is governed by a small and well characterised set of structural determinants, which supply favourable directions fixed independently of our data and make such a check possible. We restrict the analysis to this class, since the small-molecule benchmarks offer no equivalent reference against which an individual edit could be judged. A detailed analysis for the peptides characterisation is presented in Appendix H.

What the editor changes. Descriptors were computed for the 334 unique pairs accepted by the feasibility criteria, which leave 258 modifications after removing 76 identity mappings. Cleaving each macrocycle at its backbone amide bonds sorts these into 90 stereocenter inversions, 44 rear-

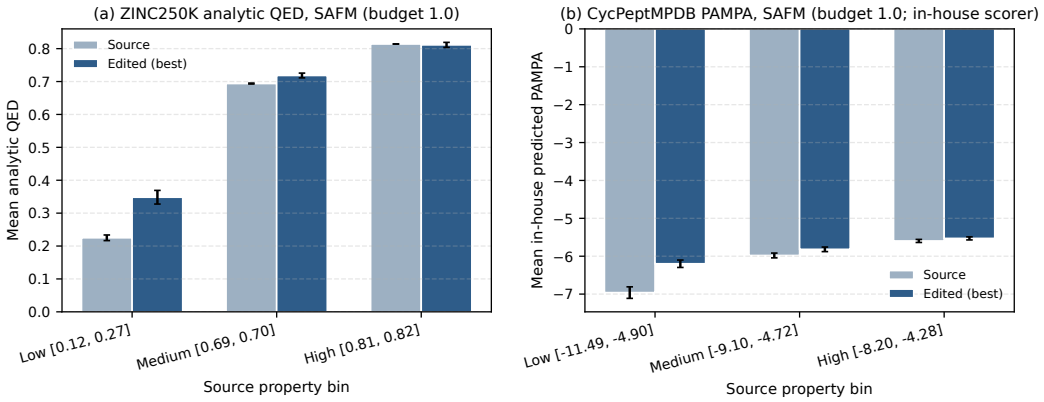

Figure 2: **Property improvement at the fixed operating point** $b = 1.0$**.** Bars show seed–fold mean and standard deviation.

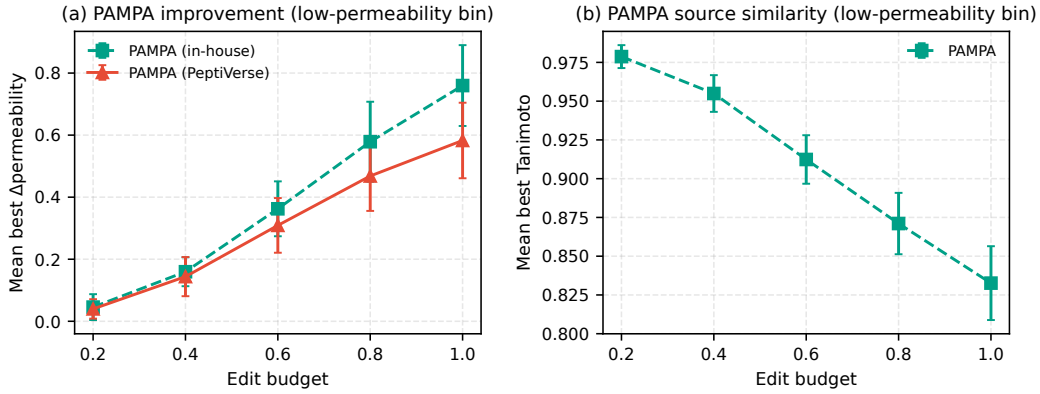

Figure 3: **SAFM budget sweep.** Predicted Δpermeability for low-permeability CycPeptMPDB sources under the in-house and PeptiVerse scorers, and CycPeptMPDB source similarity. Error bars are standard deviations over 15 seed–fold cells.

rangements of a conserved residue set, and 124 composition changes. The rearrangements relocate every residue in the ring at a higher mean Tanimoto than the single-residue substitutions. One transformation recurs, with 49 ring systems removed against 16 introduced, 36 of the removals phenyl and 9 of the introductions cyclohexane. Table 2 scores every edit against directions fixed in advance by the medicinal chemistry literature, and saturation and aromaticity move almost without exception, at 70 pairs against 7 and 33 against 1, while the determinants specific to peptides hold still. The composite is lopsided in the favourable direction at 75 edits against 6 and strengthens as the edit narrows, from 0.47 favourable rules per edit over the full set to 1.03 over the 80 single-residue substitutions. Of the 177 edits scoring zero, 134 conserve the residue set and are invisible to graph descriptors, and the rest cancel across rules.

Biophysical rationale. The reported changes support a specific biophysical mechanism. Transfer into the hydrocarbon core is dominated by desolvation of the backbone amide donors, which is why N-methylation and transannular hydrogen bonding are the established permeability routes in this class (Dougherty et al., 2019), and both are left intact by the editor. The Boltzmann-weighted ensembles confirm that conservation in three dimensions, with exposed polar surface moving from 194.9 to 192.6 Å^2, amide hydrogen exposure from 32.2 to 30.5 Å^2 and the radius of gyration unchanged. With polarity and shape fixed, the exchange acts on the apolar surface alone, replacing a rigid planar ring that disrupts acyl chain packing beyond its own volume with a puckered saturated one that accommodates itself between ordered chains. The edited peptides therefore pay less to form

Table 2: **Medicinal chemistry rule score over the 258 peptide edits.** A pair scores $+1$ for each rule it satisfies and -1 for each rule it violates, and the composite sums the first four rules. The p values are two-sided exact sign tests over the pairs in which the descriptor changes.

Descriptor	Favourable direction	Favourable	Unfavourable	Unchanged	p
General developability					
Fraction sp3 carbon	increase	70	7	181	3.5×10^{-14}
Aromatic rings	decrease	33	1	224	4.1×10^{-9}
Rotatable bonds	decrease	51	34	173	0.082
Macrocycle and peptide permeability					
H-bond donors	decrease	11	3	244	0.057
Reported, excluded from the composite					
N-methyl amides	increase	9	4	245	0.27
TPSA	decrease	19	15	224	0.61
Composite, all edits		75	6	177	2.9×10^{-16}
Composite, edits that change the residue set		75	6	43	2.9×10^{-16}
Composite, single residue substitutions		52	3	25	1.5×10^{-12}

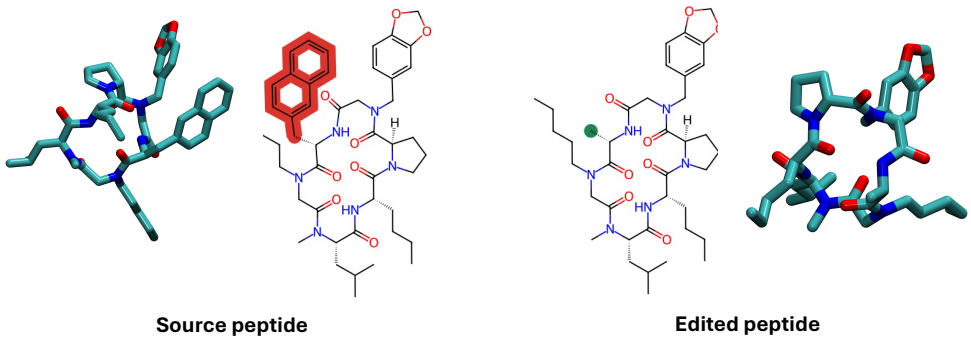

Figure 4: **Case study.** Representative example of the edit made by SAFM

a cavity in the bilayer while paying the same desolvation penalty, which places the editor on an axis independent of the two classical routes.

Case study. We took the edit with the strongest oracle agreement whose source lies inside the assay dynamic range, the top-ranked pair being reported at the PAMPA detection floor of -10. In the given example 4, measured at -8.89, the editor replaces a 2-naphthylalanine residue with alanine, so molecular weight falls from 853 to 727,Da, aromatic rings from three to one, sp3 fraction rises from 0.542 to 0.684 and lipophilicity falls from 5.20 to 2.82. Polar surface area is unchanged at 157.9 Å^2 and 151 of the 155 Å^2 of lost accessible surface are apolar, so the entire change falls on the apolar face. Both oracles score the edit as an improvement, by 3.01 and 2.40 log units, and a side chain of this size is a plausible pharmacophoric element, so the substitution warrants evaluation against a target, which the permeability objective does not cover.

6 Discussion and Limitations

Our results show that molecular lead optimisation can be formulated as source-anchored transport over mined improvement pairs. SCIN provides sparse, local supervision, while SAFM learns the corresponding trajectories and exposes the edit budget b as a direct fidelity–improvement control. Together, they support both small-molecule optimisation and cyclic-peptide permeability editing across distinct representations and constraints. Future work would include experimental validation using wet lab assays.

AI USE STATEMENT

Generative AI and large language model tools were utilized during the preparation of this manuscript solely for assisting with code maintenance and prose polishing. All conceptual framework designs, mathematical formulations, experimental pipelines, numerical analyses, and scientific conclusions were conceived, executed, verified, and curated by the authors.

REFERENCES

Roy Aerts, Joris Tavernier, Alan Kerstjens, Mazen Ahmad, Jose Carlos Gómez-Tamayo, Gary Tresadern, and Hans De Winter. C2po: an ml-powered optimizer of the membrane permeability of cyclic peptides through chemical modification. *Journal of Cheminformatics*, 17(1):168, 2025.

G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.

Konrad H Bleicher, Hans-Joachim Böhm, Klaus Müller, and Alexander I Alanine. Hit and lead generation: beyond high-throughput screening: a guide to drug discovery. *Nature reviews Drug discovery*, 2(5):369–378, 2003.

Tong Chen, Yinuo Zhang, Sophia Tang, and Pranam Chatterjee. Multi-objective-guided discrete flow matching for controllable biological sequence design. *arXiv preprint arXiv:2505.07086*, 2025.

Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*, 2020.

Patrick G Dougherty, Ashweta Sahni, and Dehua Pei. Understanding cell penetration of cyclic peptides. *Chemical reviews*, 119(17):10241, 2019.

Peter Eckmann, Kunyang Sun, Bo Zhao, Mudong Feng, Michael K Gilson, and Rose Yu. Limo: Latent inceptionism for targeted molecule generation. *Proceedings of machine learning research*, 162:5777, 2022.

Aaron L Feller, Maxim Secor, Sebastian Swanson, Claus O Wilke, and Kristine Deibler. Scaling smiles-based chemical language models for therapeutic peptide engineering. *Journal of Chemical Information and Modeling*, 2026.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268, 2018.

Ed Griffen, Andrew G Leach, Graeme R Robb, and Daniel J Warner. Matched molecular pairs as a medicinal chemistry tool: miniperspective. *Journal of medicinal chemistry*, 54(22):7739–7750, 2011.

Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science*, 11(2):577–586, 2020.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pages 2323–2332. PMLR, 2018.

Jiahan Li, Chaoran Cheng, Zuofan Wu, Ruihan Guo, Shitong Luo, Zhizhou Ren, Jian Peng, and Jianzhu Ma. Full-atom peptide design based on multi-modal flow matching. *arXiv preprint arXiv:2406.00735*, 2024.

Jianan Li, Keisuke Yanagisawa, Masatake Sugita, Takuya Fujie, Masahito Ohue, and Yutaka Akiyama. Cycpeptmpdb: a comprehensive database of membrane permeability of cyclic peptides. *Journal of chemical information and modeling*, 63(7):2240, 2023.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*, 2024.

Alexander Arjun Lobo, Urvi Awasthi, and Leonid Zhukov. Property-guided molecular generation and optimization via latent flows. *arXiv preprint arXiv:2603.26889*, 2026.

David F Nippa, Kenneth Atz, Yannick Stenzhorn, Alex T Müller, Andreas Tosstorff, Jörg Benz, Hayley Binch, Markus Bürkler, Achi Haider, Dominik Heer, et al. Expediting hit-to-lead progression in drug discovery through reaction prediction and multi-dimensional optimization. *Nature Communications*, 16(1):11646, 2025.

Bo Pan, Peter Zhiping Zhang, Hao-Wei Pang, Alex Zhu, Xiang Yu, Liying Zhang, and Liang Zhao. Retrieval-augmented foundation models for matched molecular pair transformations to recapitulate medicinal chemistry intuition. *arXiv preprint arXiv:2602.16684*, 2026.

Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, et al. Molecular sets (moses): a benchmarking platform for molecular generation models. *Frontiers in pharmacology*, 11:565644, 2020.

Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.

Heqi Sun, Hong Tan, Yanyi Chu, Jiayi Li, Ruixuan Wang, and Dong-Qing Wei. Design of permeability-optimized target-binding macrocycles via direct preference optimization. *Chemical Science*, 17(20):10223–10236, 2026.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *ArXiv*, pages arXiv–2412, 2025.

Zixu Wang, Yangyang Chen, Yifan Shang, Xiulong Yang, Wenqiong Pan, Xiucai Ye, Tetsuya Sakurai, and Xiangxiang Zeng. Multicycpermea: accurate and interpretable prediction of cyclic peptide permeability using a multimodal image-sequence model. *BMC biology*, 23(1):63, 2025.

Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018.

Yang Zhang, Jie Li, and Xu Chao. Chemnav: An interactive visual tool to navigate in the latent space for chemical molecules discovery. *Visual Informatics*, 8(4):60–70, 2024.

Yinuo Zhang, Sophia Tang, Tong Chen, Elizabeth Mahood, Sophia Vincoff, and Pranam Chatterjee. Peptiverse: A unified platform for therapeutic peptide property prediction. *Nature Communications*, 2026.

A DETAILED RESULTS FOR CYCPEPTMPDB AND ZINC250K

Table 3: **CycPeptMPDB PAMPA editing at budget $b = 1.0$.** Mean $\pm$ standard deviation is over fifteen matched CV cells. Faithful-Improve requires non-negative predicted improvement, distinctness, the joint edit gate, and Tanimoto ≥ 0.7.

Method	Oracle	Faithful ↑	Best $\Delta\hat{y}$ ↑	Best Tanimoto ↑	Constr. pass ↑
SAFM with z_0 cond.	In-house	0.363 ± 0.036	0.330 ± 0.056	0.870 ± 0.014	0.730 ± 0.034
SAFM with z_0 cond.	PeptiVerse	0.244 ± 0.037	0.259 ± 0.051	0.870 ± 0.014	0.730 ± 0.034
SAFM without z_0 cond.	In-house	0.363 ± 0.036	0.466 ± 0.052	0.843 ± 0.010	0.616 ± 0.032
SAFM without z_0 cond.	PeptiVerse	0.268 ± 0.036	0.341 ± 0.056	0.843 ± 0.010	0.616 ± 0.032
SAFM, random SCIN	In-house	0.130 ± 0.024	0.337 ± 0.115	0.718 ± 0.029	0.462 ± 0.048
SAFM, random SCIN	PeptiVerse	0.079 ± 0.019	0.277 ± 0.098	0.718 ± 0.029	0.462 ± 0.048
Latent gradient	In-house	0.061 ± 0.012	2.412 ± 0.104	0.727 ± 0.038	0.216 ± 0.021
Latent gradient	PeptiVerse	0.044 ± 0.014	-0.012 ± 0.261	0.727 ± 0.038	0.216 ± 0.021
Latent interpolation	In-house	0.042 ± 0.012	0.830 ± 0.126	0.556 ± 0.024	0.258 ± 0.041
Latent interpolation	PeptiVerse	0.028 ± 0.013	0.443 ± 0.104	0.556 ± 0.024	0.258 ± 0.041
SCIN retrieval	In-house	0.040 ± 0.013	0.830 ± 0.126	0.553 ± 0.023	0.299 ± 0.043
SCIN retrieval	PeptiVerse	0.023 ± 0.015	0.453 ± 0.108	0.553 ± 0.023	0.299 ± 0.043
Generate-and-filter	In-house	0.003 ± 0.003	1.810 ± 0.120	0.375 ± 0.012	0.040 ± 0.010
Generate-and-filter	PeptiVerse	0.003 ± 0.003	0.936 ± 0.087	0.375 ± 0.012	0.040 ± 0.010
Identity	In-house	0.000 ± 0.000	0.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Identity	PeptiVerse	0.000 ± 0.000	0.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000

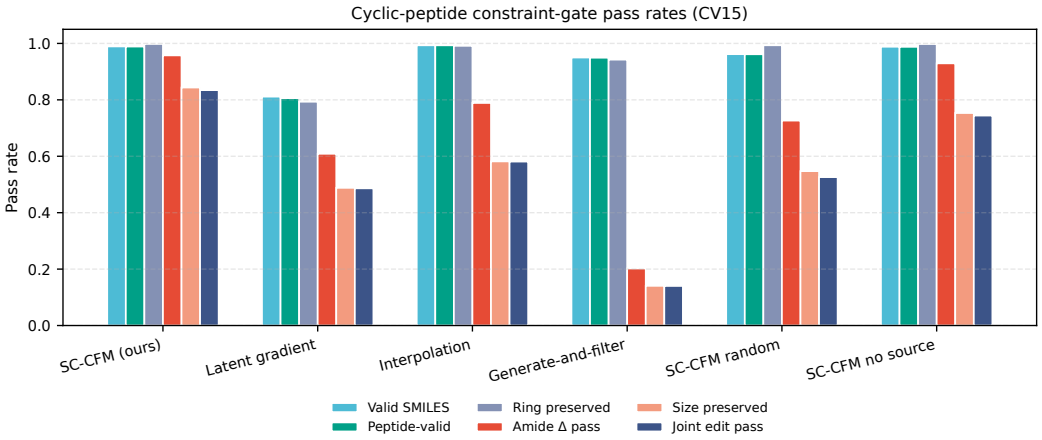

Figure 5: **Cyclic-peptide constraint-pass comparison at $b = 1.0$.** Individual and joint gate pass rates for SAFM, latent gradient, and generate-and-filter over the 15 matched seed–fold cells. Error bars show standard deviation across cells.

A.1 ASSAY ROBUSTNESS

In the zero-shot PAMPA-to-Caco-2 rescoring study, the ranking by Faithful-Improve remains unchanged between the PeptiVerse assay heads ($\rho = 1.00$). SAFM ranks first under both, attaining $40.9 \pm 4.6\%$ under the PAMPA head and $40.1 \pm 3.6\%$ under the Caco-2 head. Table 4 reports the full matched comparison.

The absolute values are not invariant across assay heads. For example, SAFM random and latent interpolation obtain higher Faithful-Improve under the Caco-2 head than under the PAMPA head, whereas latent gradient decreases. Generate-and-filter produces the largest mean predicted changes under both heads but almost never satisfies the joint fidelity endpoint. Thus, the cross-assay result

Table 4: **Zero-shot PAMPA-to-Caco-2 transfer of PAMPA-trained edits.** The same decoded candidates are rescored post hoc with the public PeptiVerse PAMPA and Caco-2 heads. Faithful-Improve requires non-negative assay-head change, Tanimoto similarity at least 0.7, distinctness, and all edit constraints. Values are mean ± standard deviation across 15 matched seed–fold cells at budget $b = 1.0$.

Method	PAMPA Faithful (%)	Caco-2 Faithful (%)
SAFM with z_0 cond.	$\mathbf{40.9 \pm 4.6}$	$\mathbf{40.1 \pm 3.6}$
SAFM without z_0 cond.	39.4 ± 4.4	38.8 ± 4.4
Latent interpolation	17.0 ± 2.9	25.1 ± 3.4
Latent gradient	16.2 ± 4.1	12.9 ± 2.9
SAFM, random SCIN	19.3 ± 3.0	27.4 ± 4.0
SCIN retrieval	2.6 ± 1.7	4.1 ± 1.4
Generate-and-filter	0.1 ± 0.2	0.2 ± 0.2

preserves the overall method ordering while retaining meaningful assay-specific differences. Because both columns are predictions from PeptiVerse heads applied to the same decoded candidates, this analysis is a robustness check under an alternative assay-specific predictor, not experimental evidence of improved Caco-2 permeability.

Table 5: **ZINC250K editing at budget $b = 1.0$.** Mean ± standard deviation is over fifteen matched CV cells. Faithful-Improve requires non-negative predicted improvement, distinctness, the joint edit gate, and Tanimoto ≥ 0.7.

Method	Faithful ↑	Best ΔQED ↑	Best Tanimoto ↑	Constr. pass ↑
SAFM with z_0 cond.	0.098 ± 0.019	0.038 ± 0.008	0.709 ± 0.027	0.592 ± 0.031
SAFM without z_0 cond.	0.055 ± 0.013	0.130 ± 0.007	0.462 ± 0.011	0.791 ± 0.027
SAFM, random SCIN	0.058 ± 0.021	0.046 ± 0.015	0.533 ± 0.041	0.591 ± 0.038
Latent interpolation	0.070 ± 0.016	0.226 ± 0.107	0.604 ± 0.115	0.662 ± 0.132
Latent gradient	0.001 ± 0.003	0.184 ± 0.022	0.269 ± 0.023	0.332 ± 0.031
Identity	0.000 ± 0.000	0.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
SCIN retrieval	0.000 ± 0.000	0.242 ± 0.002	0.696 ± 0.005	0.333 ± 0.000
De novo filter	0.000 ± 0.000	0.292 ± 0.032	0.113 ± 0.005	0.461 ± 0.054

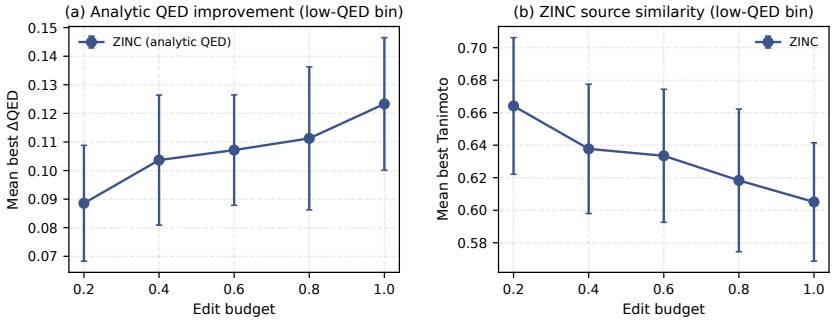

Figure 6: **SAFM budget sweep.** Predicted ΔQED for low-QED ZINC250K sources and ZINC250K source similarity. Error bars are standard deviations over 15 seed–fold cells.

A.2 Baseline method details

The following methods are evaluated with the same frozen encoder and decoder, held-out source set, candidate count, and feasibility checks, so differences reflect the editing rule rather than stack

changes: 1) **Identity (M0)**: The source latent is decoded without editing. This is the no-change reference for sanity check; 2) **Latent gradient (M1)**: Starting from the source latent, we apply iterative ascent on the predicted target utility and then decode with the same candidate-selection pipeline. It targets utility gain directly and serves as a local optimisation baseline in the shared latent space; 3) **Generate-and-filter (M2)**: We generate candidates from property-conditioned targets and retain outputs by post-hoc filtering with the same validity and constraint checks. Unlike local edit methods, it is not source-anchored and mainly tests whether unconstrained property-seeking generation can satisfy local-faithfulness criteria; 4) **SCIN retrieval (M3)**: For each source, we retrieve a mined SCIN neighbour that satisfies the improvement margin in the training fold, then decode that retrieved latent without integrating a learned flow field. This isolates neighbour retrieval as an editing primitive; 5) **Latent interpolation (M4)**: We interpolate between source and SCIN improvement target latents and decode the interpolated point. This provides a simple path-based baseline that uses the same neighbour information but no learned transport dynamics. Both M3 and M4 would require matched pairs during inference which is not the usual deployment scenario; 6) **SAFM, random SCIN (A1)**: This ablation keeps the SAFM architecture but replaces SCIN neighbourhood construction with random improving pairs. It tests whether learned transport benefits specifically from local SCIN supervision. 7) **SAFM without z_0 cond. (A2)**: This ablation keeps SCIN pairs but removes source conditioning during SAFM training and inference. It tests whether source anchoring is necessary once improving-pair supervision is provided.

A.3 External method comparisons

MoltenFlow-style optimisation. We implemented the optimisation-mode recipe of MoltenFlow Lobo et al. (2026) in the frozen ChemBERTa latent space. This is a method-level reimplementation and the comparison is reported for ZINC250K fold 0 across three seeds at $b = 1.0$.

Table 6: **MoltenFlow-style optimisation on ZINC250K.** Mean $\pm$ standard deviation over three fold-0 seeds at $b = 1.0$. Faithful-Improve requires non-negative analytic QED change, validity, distinctness, constraint passage, and Morgan Tanimoto ≥ 0.7.

Method	Faithful	Best ΔQED	Best Tanimoto	Constr. pass
MoltenFlow-style	0.000 ± 0.000	0.450 ± 0.001	0.054 ± 0.010	0.253 ± 0.031

MoltenFlow-style optimisation produces large QED increases but moves far from the source, resulting in no edits that satisfy the joint Faithful-Improve endpoint.

C2PO. C2PO Aerts et al. (2025) is retrained from scratch in each of the 15 peptide seed–fold cells using the same training and validation partitions as SAFM, without access to the held-out test fold. Both methods are evaluated on the same held-out sources per cell using the same fold-specific in-house oracle and Faithful-Improve definition. C2PO optimises toward a source-relative target of $+3$ PAMPA units and exports one final candidate per source.

Table 7: **SAFM and C2PO on CycPeptMPDB PAMPA.** Mean $\pm$ standard deviation over 15 matched seed–fold cells at $b = 1.0$. Both methods are evaluated on identical sources with the same in-house oracle and endpoint definition.

Method	Faithful	Best $\Delta\hat{y}$	Best Tanimoto	Constr. pass
SAFM	$\mathbf{0.363 \pm 0.036}$	$\mathbf{0.330 \pm 0.056}$	0.870 ± 0.014	$\mathbf{0.730 \pm 0.034}$
C2PO	0.086 ± 0.032	-0.063 ± 0.080	$\mathbf{0.904 \pm 0.036}$	0.640 ± 0.083

SAFM achieves a higher Faithful-Improve rate and mean predicted property change, whereas C2PO retains slightly higher source similarity.

B Encoder scale ablation on peptides

We evaluated downstream SAFM editing performance when using PeptideCLM-2 foundation models of varying capacities: MLM-small ($d = 512$) with $\sim 32M$, MLM-base ($d = 768$) with $\sim 114M$,

and MLM-large ($d = 1024$) with $\sim 337M$ parameters. Table 8 reports performance across the edit budgets across three seeds on the first fold. We use the same configuration as the cyclic peptide experiments reported in the main text.

Table 8: **PeptideCLM encoder capacity ablation on CycPeptMPDB (PAMPA).** Evaluated under both in-house and PeptiVerse oracles across three seeds on the first fold.

Encoder	Oracle	Faithful ↑	Best $\Delta\hat{y}$ ↑	Best Tanimoto ↑	Constr. pass ↑
MLM-large	In-house	$\mathbf{0.573 \pm 0.008}$	$\mathbf{0.602 \pm 0.000}$	0.844 ± 0.000	$\mathbf{0.854 \pm 0.008}$
MLM-large	PeptiVerse	$\mathbf{0.603 \pm 0.022}$	$\mathbf{0.416 \pm 0.046}$	0.853 ± 0.000	$\mathbf{0.854 \pm 0.008}$
MLM-base	In-house	0.429 ± 0.019	0.353 ± 0.062	0.874 ± 0.009	0.841 ± 0.047
MLM-base	PeptiVerse	0.392 ± 0.018	0.319 ± 0.061	0.874 ± 0.015	0.841 ± 0.047
MLM-small	In-house	0.406 ± 0.008	0.365 ± 0.065	$\mathbf{0.889 \pm 0.015}$	0.846 ± 0.014
MLM-small	PeptiVerse	0.407 ± 0.033	0.293 ± 0.052	$\mathbf{0.893 \pm 0.017}$	0.846 ± 0.014

MLM-large yields the highest faithful-improvement rate under both oracles in these experiments. MLM-small attains the highest best-of-source Tanimoto similarity, while MLM-base and MLM-small do not exhibit a monotonic capacity trend. The result therefore supports the selected MLM-large configuration under the joint fidelity–improvement criterion, but does not establish model capacity as the causal mechanism.

C SCIN HYPERPARAMETER SENSITIVITY

We report here the best-over-budget faithful-improvement performance for SAFM across different neighbour-pool size k and latent cosine threshold τ for the SCIN retrieval. The property margin is held fixed at $\delta_{\min} = 1.0$. The sweep reuses the frozen components from seed 0, fold 0 of cyclic peptide experiments and retrains only the SAFM method.

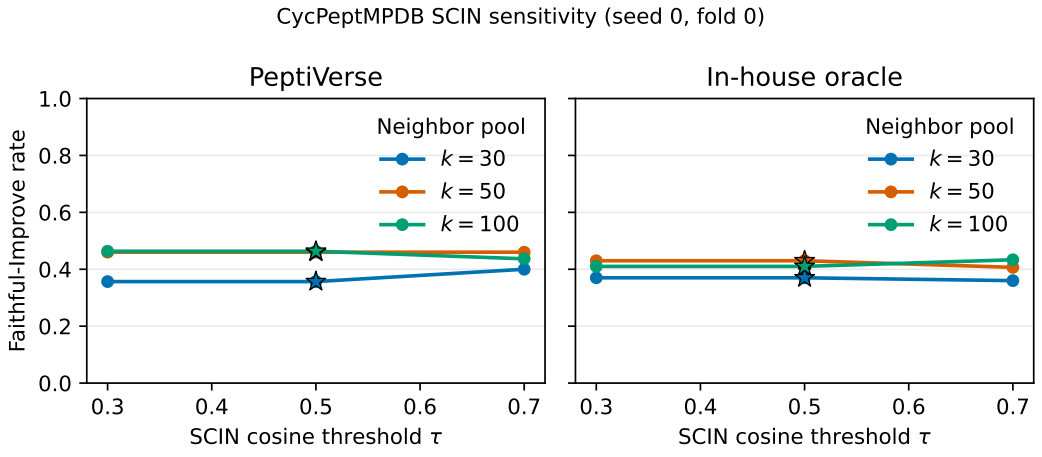

Figure 7: **SCIN hyperparameter sensitivity on CycPeptMPDB.** Faithful-Improve as a function of the latent cosine threshold τ and neighbour-pool size k. Stars identify the configured threshold $\tau = 0.50$.

Increasing k increases the number of mined pairs from 3,683 to 4,185 to 4,807, with performance generally improving from $k = 30$ to $k = 50$, then plateauing. PeptiVerse and the in-house oracle broadly agree that $k = 50 - 100$ is preferable. At $\tau = 0.3$ and 0.5, pair counts, pair statistics, and faithful-improvement are identical. This is because the selected pairs already have a high cosine similarity after meeting all the other criteria.

D SENSITIVITY TO GUIDANCE SCALE γ

Figure 8 reports a matched classifier-free guidance sweep on ZINC250K and CycPeptMPDB (seed0/fold0). For each guidance scale and source, the displayed point uses the candidate with the largest property improvement across the common edit-budget sweep. ZINC improvement is rescored with analytic QED; peptide improvement is rescored with PeptiVerse.

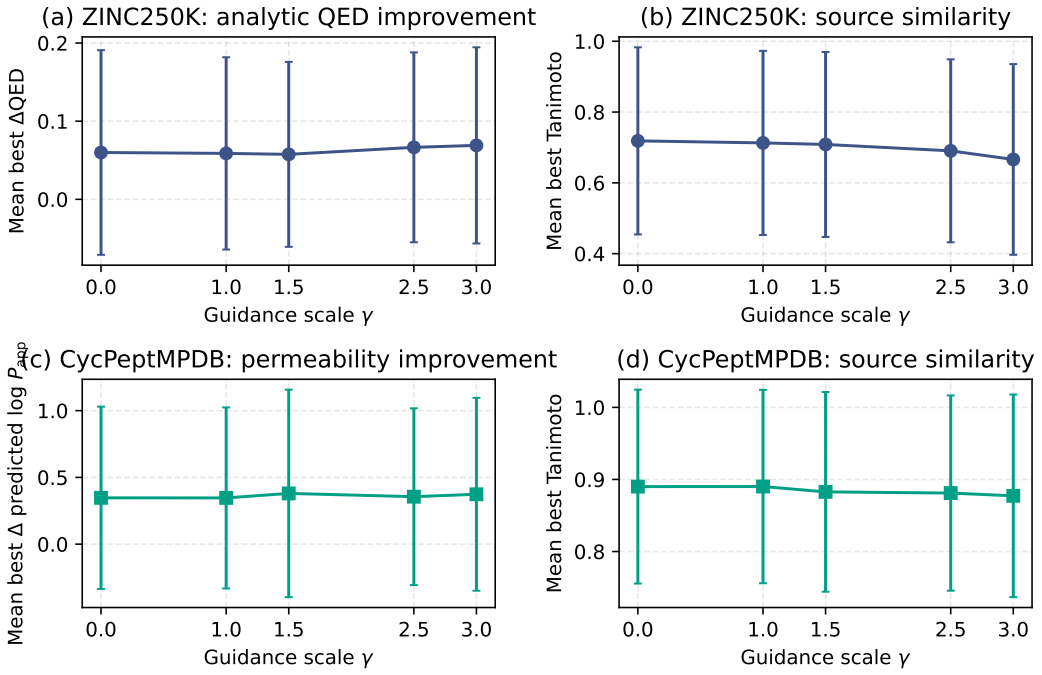

Figure 8: **Property-guidance sensitivity on ZINC250K and CycPeptMPDB.** Mean best property improvement versus mean best source Tanimoto across test sources for $\gamma \in \{0, 1, 1.5, 2.5, 3\}$ (seed0/fold0).

E PERFORMANCE OF THE DECODER

The edited outputs and editing-test property labels are not used for decoder training, and every matched editing method within an experiment cell uses the same frozen decoder. The decoder is trained and evaluated on an independent $80/10/10$ train/validation/test partition. The following table reports the Greedy latent-to-SMILES reconstruction on the test split. Validity is the fraction of outputs parsed by RDKit. Mean Tanimoto is computed over finite valid source–decode pairs using radius-2, 2048-bit Morgan fingerprints.

Table 9: **Decoder reconstruction quality.** Mean $\pm$ standard deviation across three unique seed-specific decoder splits repeated across the five editing folds.

Dataset	Valid (%)	Mean Tanimoto
ZINC250K	72.7 ± 0.1	0.824 ± 0.002
CycPeptMPDB	95.7 ± 0.3	0.958 ± 0.003

The cyclic-peptide decoder shows stronger direct reconstruction quality than the ZINC250K decoder: 95.7% versus 72.7% valid greedy outputs, and mean Tanimoto 0.958 versus 0.824. Thus, the peptide decoder more reliably maps in-distribution latent representations back to valid, source-faithful structures.

F PERFORMANCE OF THE PROPERTY ORACLES

Figure 9 reports correlation of the seed-0 out-of-fold oracle predictions with the true labels. It compares the learned ZINC250K QED proxy against analytic RDKit QED, the in-house cyclic-peptide scorer and PeptiVerse against measured PAMPA permeability.

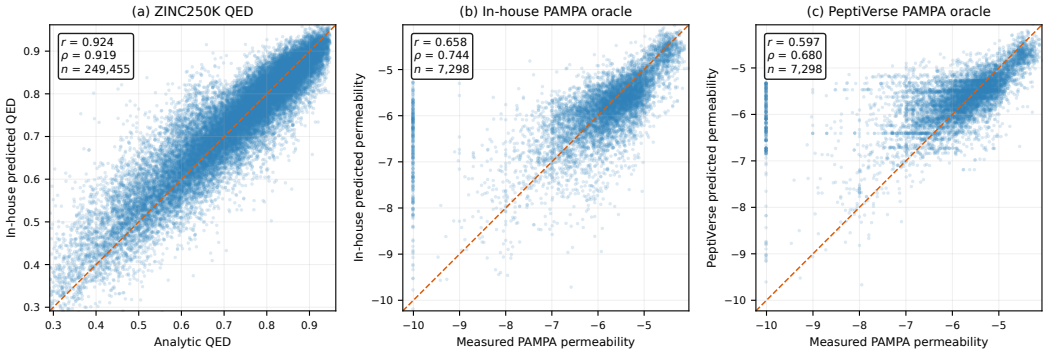

Figure 9: **Oracle–label correlations.** Each molecule appears in exactly one plotted CV test partition. Dashed lines mark identity; annotations report Pearson r, Spearman ρ, and the number of finite pairs.

G ADDITIONAL MOSES

To confirm that our findings are not idiosyncratic to ZINC250K, we report replication experiments on the official MOSES benchmark split. All MOSES experiments use the ChemBERTa-77M-MLM encoder ($d = 512$) with the official MOSES training, validation, and test splits. Table 10 reports the results for the experiments across the budgets. SAFM retains the source fidelity (best Tanimoto 0.862, mean Tanimoto 0.877, constraint pass 0.866).

Table 10: **MOSES editing evaluation (seed0/fold0).** Mean best metrics per source across the budgets.

Method	Best Tanimoto ↑	Best Δscore ↑	Mean Tanimoto ↑	Constr. pass ↑
SAFM random (A1)	0.710	0.037	0.747	0.730
SAFM without z_0 cond.	0.708	0.096	0.814	0.812
SAFM with z_0 cond.	**0.862**	0.013	**0.877**	**0.866**

H DETAILED CHARACTERISATION OF PEPTIDE EDIT

To identify the chemical modifications introduced by the editor, molecular descriptors were computed for every source and edited peptide that met the feasibility criteria on the PAMPA runs. The 699 accepted proposals are distributed over five edit budgets and reduce to 334 unique pairs once duplicate proposals are collapsed. Of these, 76 are identity mappings that return the source molecule, which leaves 258 actual modifications. Global fingerprint similarity provides limited spatial resolution regarding the modification site within a macrocycle. Each peptide ring was therefore cleaved at its backbone amide bonds to generate residue multisets, which quantifies the structural extent of every edit. The modifications comprise 90 stereocenter inversions with conserved residue composition, 44 rearrangements of the residue sequence within the ring, and 124 composition changes, 80 of which involve a single position. The rearrangements yield a mean Tanimoto similarity of 0.93 against 0.84 for single-residue substitutions, despite affecting every residue in the macrocycle. Morgan similarity therefore assigns its highest scores to the most structurally extensive class. The consequence extends beyond the present analysis, since any locality threshold expressed in fingerprint space admits edits that relocate an entire macrocyclic sequence.

Side-chain ring systems constitute the primary site of chemical modification, since the feasibility criteria hold the backbone macrocycle invariant. Ring composition also governs partitioning into the hydrocarbon core of a lipid bilayer. All ring systems external to the backbone macrocycle were therefore enumerated and mapped between source and edited counterparts, which converts every modification into an inventory of rings gained and lost. Among the modifications affecting the molecular graph, 49 ring systems are eliminated and 16 are introduced. Phenyl groups account for 36 removals and cyclohexane groups for 9 introductions. A further 33 pairs exhibit both halves of the exchange in one modification, losing an aromatic ring while increasing sp3 character. The editor consequently performs one conserved transformation, substituting planar aromatic side chains with saturated analogues of comparable steric bulk. It performs that transformation while the polar budget of the peptide remains fixed, since amide bond count is retained in 97.3% of pairs and the macrocycle in all cases.

The chemical significance of the exchange follows from what distinguishes the two ring types once size and lipophilicity are comparable. A benzene ring is planar and rigid, and its quadrupole moment and π electron density accept weak hydrogen bonds, which favours residence near the polar headgroup region. It also disrupts acyl chain packing over a lateral extent exceeding its van der Waals volume. A cyclohexane ring of identical carbon count is puckered and conformationally mobile, carries neither feature, and accommodates itself between ordered chains at lower cost. The editor therefore reaches permeability through the cavity formation term of the transfer free energy and leaves the desolvation term unaltered. This positions its modification as a third axis, independent of the two established routes in this molecular class and available in combination with them.

Prior medicinal chemistry literature establishes the directional effect of each molecular descriptor on permeability. Table 2 presents the six rules categorised by the population on which each direction was established, together with a composite score aggregating the four retained metrics.

Table 2 shows that saturation and aromaticity exhibit nearly invariant directional shifts, at 70 against 7 pairs and 33 against 1 pair. Both occur within the same modifications, which reproduces the ring exchange of the preceding paragraphs in descriptor space. Flexibility and hydrogen bond donor count display consistent directional tendencies with isolated counterexamples. The two parameters that govern permeability specifically in peptides hold still, since the editor preserves the methylation pattern in 245 of the 258 pairs and polar surface area in 224. The resulting profile indicates a specialised mechanism that operates on side chains while maintaining the chemistry of the backbone. That specialisation carries a practical consequence, because the handles medicinal chemists apply to this class remain unspent and available for a subsequent optimisation stage.

At the whole-molecule level the operation yields a favourable distribution, with 75 edits improving the composite score against 6 that decrease it. The outcome intensifies with the spatial locality of the edit. It rises from 0.47 favourable rules per edit across the entire dataset to 0.97 for the 124 edits altering the residue set, and reaches 1.03 for the 80 single-residue substitutions. That gradient indicates that the chemical signal is carried by localised substitutions and grows stronger as the edit narrows, which is the behaviour a locality-constrained generator is intended to produce. The 177 edits yielding a score of zero separate by cause. In 134 the residue set is conserved, so descriptors derived from the molecular graph are insensitive to the modification by construction. This bounds the reach of the present analysis and defines what the conformational treatment must address. The remaining 43 alter the residue set and produce rule effects that mutually cancel.

Conformational ensembles (Figure 10) establish which structural features the modification preserves, and that conservation is the condition under which the effect can be attributed to cavity formation. Partitioning from aqueous media into the hydrocarbon core is governed predominantly by desolvation, and backbone amide groups account for most of that penalty. A buried donor forfeits its aqueous hydrogen bonds without proximal replacement, which is why N-methylation and transannular hydrogen bonding are the established permeability pathways in this class (Dougherty et al., 2019). Both pathways survive the transformation. Donor count, methylation pattern and intramolecular hydrogen bond count remain constant, and topological polar surface area is conserved in 86.8% of pairs. The Boltzmann-weighted ensembles confirm that conservation in three dimensions. Exposed polar surface area shifts from 194.9 to $192.6\,\text{Å}^2$ and amide hydrogen exposure from 32.2 to $30.5\,\text{Å}^2$, variations of approximately 1% and 5%. Radius of gyration and conformational amplitude remain invariant at 5.31 against $5.32\,\text{Å}$ and 17.83 against 17.85. Molecular polarity, geometry and the classical permeability determinants therefore persist across the modified set. The

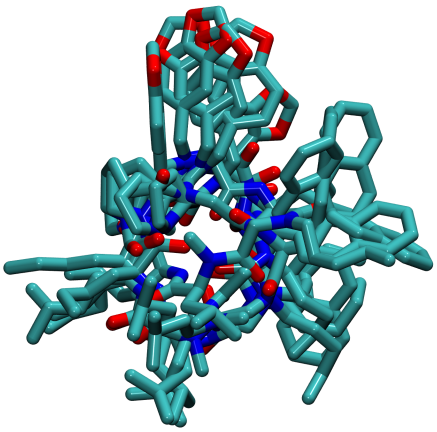

Figure 10: Superimposed conformational ensemble of a macrocyclic peptide generated using the MMFF94 force field in vacuum using ETKDG Method in RDKit, highlighting conformational flexibility and spatial variation of the macrocyclic ring and side chains (Shown N=5).

transformation operates on the apolar surface alone, which isolates cavity formation as the quantity the editor manipulates. These ensembles are generated in vacuum with a classical force field and without explicit solvent or membrane, an environment that favours conformations burying polar groups internally. The exposure magnitudes therefore reflect a relative ordering pending calibration against explicit solvent, and the directional evidence for the modification rests on the topological analysis above.

I Qualitative Case Studies and Review Protocol

To ensure chemical plausibility beyond scalar summary metrics, the evaluation suite includes an automated case-review tool that inspects edited molecules across multiple orthogonal dimensions:

1. **Physicochemical shifts:** Tracking changes in molecular weight ($\Delta\mathrm{MW}$), topological polar surface area ($\Delta\mathrm{TPSA}$), and lipophilicity ($\Delta\,\mathrm{logP}$). Edits that increase predicted permeability without causing unphysical lipophilicity surges ($\Delta\,\mathrm{logP} > 3$) are prioritized.

2. **Sequence and monomer integrity:** For cyclic peptides, parsed Hierarchical Editing Language for Macromolecules (HELM) representations verify that macrocycle length and backbone connectivity remain intact, identifying whether edits correspond to side-chain substitutions, stereocenter inversions, or cyclization modifications.

3. **Nearest-training neighbour similarity:** Candidate structures are compared against the entire training set using canonical Morgan fingerprints to distinguish true de novo local modifications from mere retrieval of known database entries.

J Computational Complexity

SCIN requires nearest-neighbour retrieval over the training embeddings and is constructed once per fold. With property guidance enabled, Euler inference with N steps performs two velocity-network evaluations per step, hence $2N = O(N)$ evaluations. Generating K sequences requires $O(K)$ autoregressive decoder passes, with sequence length and attention cost suppressed in this notation. Candidate deduplication and validity checks scale linearly with K.

K REPRODUCIBILITY DETAILS

Compute and hardware. All model training and inference sweeps were executed on Linux compute nodes equipped with NVIDIA RTX 5090 (32GB VRAM) and NVIDIA A100 (80GB VRAM) GPUs.

Environment and software. The repository is packaged with `uv` under Python 3.12. Core software versions include PyTorch 2.13.0 with CUDA 13.0, Hugging Face Transformers 4.49.0, RDKit 2024.03.5, and NumPy 1.26.4. Random seeds $\{0, 1, 2\}$ were set for Python `random`, NumPy, and PyTorch; CUDA deterministic mode was set to `True`.