HYBRIDGE: RECEPTOR-GUIDED CYCLIC-PEPTIDE GENERATION ON COMPILED COVALENT SUPPORT

Anonymous authors
Paper under double-blind review

ABSTRACT

Cyclic peptides offers attractive scaffolds for protein-targeted therapeutics by coupling rich amino-acid chemistry with a constrained conformational landscape. However, existing generative methods using geometric conditioning or cyclic encodings leave ring constraints unenforced during the sampling process. These geometric inconsistencies can compromise both the physical validity and receptor compatibility of their generated designs. Here, we propose HYBRIDGE, a hybrid sampler that preserves compiled covalent constraints throughout joint sequence–structure generation. HYBRIDGE leverages a receptor-free reference model to capture intrinsic peptide diversity. Following geometric placement, a shared receptor potential guides backbone motion and residue substitutions, with accepted updates preserving the compiled constraints. To learn pairing preferences, we incorporate crossed peptide–receptor panels alongside chemical supervision, canceling additive single-chain scoring offsets. Evaluated on held-out receptor targets under a head-to-tail (H2T) benchmark, HYBRIDGE achieves 61.2% structural delivery under a common closure, contact, and severe-overlap criterion. Specifically, suffix guidance improves delivery by 15.1 percentage points over execution without the learned potential and by 5.6 over terminal-only use. Experiments demonstrate receptor-responsive cyclic-peptide generation with persistent ring constraints and improved structural delivery, complemented by construction examples spanning disulfide, lactam, and bicyclic topologies.

1 INTRODUCTION

Cyclic peptides offer diverse amino-acid chemistry within a constrained molecular shape, making them attractive scaffolds for recognizing protein targets and developing therapeutics (Rettie et al., 2025). Generally, designing cyclic peptides for a receptor involves adapting sequence and conformation while preserving the ring that couples them.

Recent peptide models jointly generate sequence and structure (Lin et al., 2024; Wang et al., 2024), while cyclic systems incorporate explicit bonds, topology-aware dynamics, or macrocycle-specific architectures (Zhou et al., 2025; Jiang et al., 2025; Zhang et al., 2026; Rettie et al., 2025; Zhao et al., 2026). Scarce paired cyclic peptide–receptor data motivate reliance on geometric conditioning or cyclic encodings. However, these mechanisms alone do not guarantee ring-constraint preservation throughout sampling. Consequently, their guided updates may violate the intended covalent geometry, and subsequent ring closure could disrupt receptor contacts established during sampling. These inconsistencies could compromise the physical validity and receptor compatibility of their final designs, limiting successful delivery.

To close this gap, we introduce HYBRIDGE, a hybrid sampler that lets sequence and conformation adapt to a receptor while keeping the requested ring structures intact (Figure 1). A covalent graph specifies the topology, including head-to-tail (H2T), disulfide, lactam, and bicyclic programs, and defines the constraints that each accepted update must preserve. To address paired-data scarcity, we decouple receptor-free learning of constrained peptide variation from receptor guidance. After geometric placement in the receptor frame, one learned potential guides both backbone motion and residue substitutions. Receptor preferences thus shape the cyclic candidate itself as it develops, changing its chemistry and conformation while its ring structures remain intact.

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.

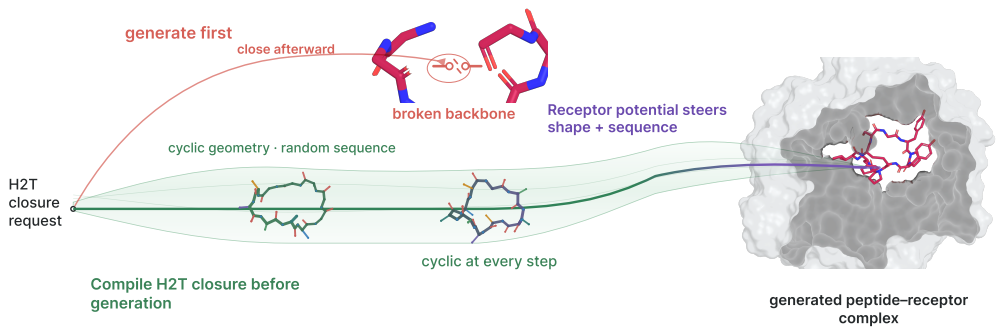


Figure 1: **Compile closure first, then steer.** HYBRIDGE stays cyclic as one receptor potential steers shape and sequence; the upper branch shows post-hoc failure.

The potential must distinguish pairing preferences from additive single-partner score offsets. Crossed comparisons of two peptides and two receptors cancel these offsets. Design and chemistry supervision complement this crossed signal to obtain the potential that guides shape and residue changes. Experiments examine how these components work together. Path controls test whether ring constraints persist during generation; receptor interventions measure responses to pairing and chemistry changes. We then compare three uses of the same potential within a fixed sampling procedure: guiding the suffix, scoring only terminal candidates, and remaining inactive. This isolates guidance timing and tests whether adaptation before selection improves structural delivery.

Specific contributions of this work are as follows:

- **Continuous Ring Geometry Preservation in Guided Generation.** To address the failure of endpoint closure in constraining intermediate states, our hybrid sampler preserves the covalent graph through accepted sequence and coordinate updates, ensuring that receptor guidance operates continuously on structurally valid cyclic candidates.
- **Pair-Specific Guidance via Cross-Conditioned Scoring.** Crossed score comparisons cancel additive peptide-only and receptor-only offsets. Combined with chemical supervision, they train a shared potential for conformational updates and residue substitutions.
- **Early-Stage Guidance for Active Candidate Shaping.** Terminal scoring only selects completed candidates. Using the same potential during generation reaches 61.2% H2T delivery on 103 held-out receptor components, gaining 15.1 percentage points over no learned potential and 5.6 over terminal-only use in matched schedules.

2 RELATED WORK

Peptide generators differ in how they combine receptor conditioning with molecular constraints. PPFlow conditions linear-peptide generation on a receptor (Lin et al., 2024). DiffPepBuilder jointly denoises sequence and structure with optional disulfide construction (Wang et al., 2024), while CpSDE brings explicit bonds into all-atom cyclic generation (Zhou et al., 2025). CP-Composer conditions a latent denoising network on residue-type and distance constraints, using classifier-free guidance to generate sequences and structures (Jiang et al., 2025). RFpeptides adapts RFdiffusion and ProteinMPNN to binding macrocycles (Rettie et al., 2025). GeoCycler aligns generation with cyclic constraints (Zhang et al., 2026), and APCyc learns a hierarchical latent model for cyclic peptides (Zhao et al., 2026).

HYBRIDGE puts covalent constraints in the execution rule: accepted sequence and coordinate updates remain on compiled support. Within this space, a learned receptor potential supplies coordinate derivatives and residue differences, connecting pairing preferences to the evolving cyclic candidate. CP-Composer and RFpeptides provide released pipelines for the same receptor-plus-H2T task and form the direct comparison. Appendix F details the generation interfaces and release status of the methods above.

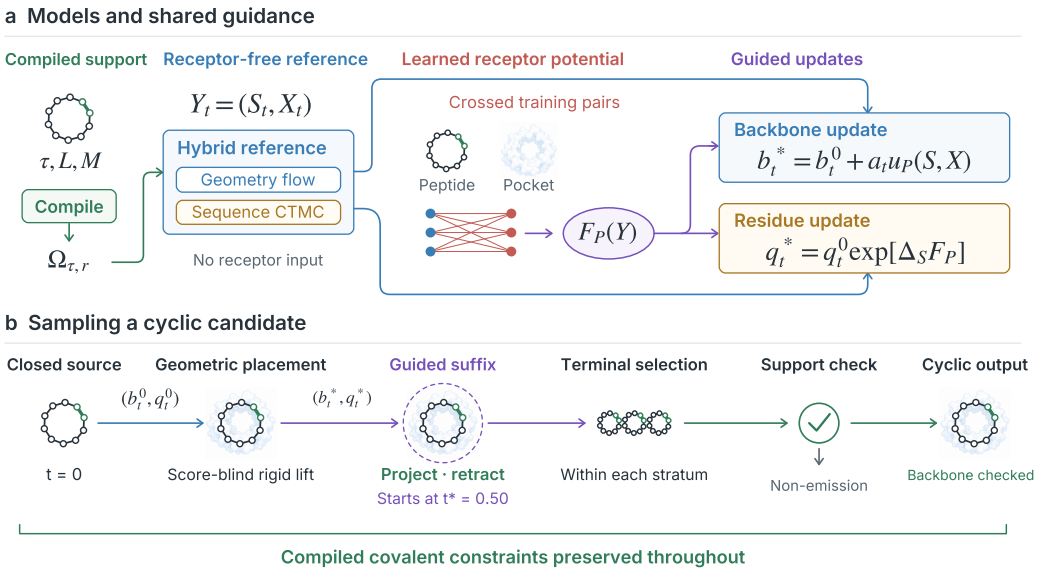


Figure 2: **Learned components and sampling procedure of HyBridge. a,** A receptor-free reference supplies the coordinate drift and residue rates. The learned receptor potential guides both through the local coordinate direction $u_P(S, X)$ and the residue difference $\Delta_S F_P = F_P(S', X) - F_P(S, X)$, giving the updates in Eq. (5). **b,** The prefix follows the reference with score-blind proper-rigid placement. From $t^\star = 0.50$, the numerical coordinate and token controls drive the suffix while projection and retraction preserve support; within-stratum selection returns a backbone-checked ring or typed non-emission.

3 Method: HyBridge

HyBridge organizes generation around two changes to a cyclic candidate: backbone motion adjusts the arrangement of residues, and substitutions change their chemistry. Compiled constraints keep both updates on admissible molecular states. A receptor-free reference proposes variation, and one receptor potential guides both updates after geometric placement. Figure 2 shows the models (panel a) and alternating updates from a closed source to a backbone-checked output (panel b).

3.1 Compiled molecular support

Compilation fixes the covalent structure while allowing the peptide to change conformation. For H2T, the graph includes the closing bond $C_L - N_1$, and local distance targets are inherited from the placed source backbone. These targets remain fixed along each particle path; the scaffold can change shape while satisfying them.

For a peptide of length L, $S \in \{1, \dots, 20\}^L$ denotes residue tokens and $X \in \mathbb{R}^{L\times 4\times 3}$ stores N, Cα, C, and O backbone atoms. A topology program τ compiles to an immutable molecular multigraph G_τ, attachment locks A_τ, and an execution adapter. The locks restrict residue identities at attachment sites; $A_\tau(S) = 1$ means that the sequence satisfies them. Internal disulfide, lactam, and composed topologies likewise enter through the graph and these locks.

The adapter's coordinate constraints are collected in the residual $c_{\tau,r}$. Here $r \in \mathcal{R}_{\tau,L}$ is the target-free constraint reference fixing the numerical targets along one particle path. The admissible molecular states satisfy both the sequence locks and coordinate constraints:

$$\Omega_{\tau,r} = \{(S, X) : A_\tau(S) = 1,\ c_{\tau,r}(S, X) = 0\}, \qquad \Omega_\tau = \bigcup_{r\in\mathcal{R}_{\tau,L}} \Omega_{\tau,r}. \tag{1}$$

At fixed S, the coordinate states form the active leaf $\mathcal{M}_{\tau,S,r}$. Coordinate derivatives and retractions act on this leaf while G_τ stays fixed.

Accepted coordinate updates are retracted to the active leaf, and sequence updates preserve the graph and attachment locks. Failed support checks retain the previous state. Exact retraction preserves support; the numerical solver certifies it within the declared residual tolerance. The support proof and solver details are in Appendices B.2 and B.1.

3.2 Learning the receptor-free reference

The receptor-free reference learns $\bar{\mu}_\tau$ by projected flow matching along an aligned interpolation from locally corrupted backbones to observed structures. Retraction keeps this path on its active constraint leaf. A symmetric 20-state continuous-time Markov chain corrupts sequences; the model learns reverse substitution rates and a residue posterior. It returns coordinate velocity and categorical rates, with a_t setting inference noise and receptor-correction scale. Appendix C gives the objectives, architecture, and checkpoints. H2T inference begins with locally perturbed training-bank backbones and corrupted donor sequences. The admitted source coordinates fix each path's constraint reference; reference training uses no receptor coordinates, pocket features, or binding labels.

3.3 Receptor placement and learned preferences

A finite score-blind search places the peptide using a proper-rigid lift $g_P = (R, t)$, retaining poses that pass receptor-heavy exclusion and pocket-contact tests. The lift preserves internal distances, covalent constraints, and chirality. The potential scores the resulting pairing; placement itself uses no IBC or binding score. The target model reads a bounded receptor pocket together with the peptide interface. Its cross-edges encode distance, direction, residue chemistry, steric clash, and uncertainty. This representation is the interface boundary condition (IBC). Training compares factual pairs with cross-pocket, peptide-chemistry-shuffled, pocket-chemistry-shuffled, role-intervened, and exact-null counterparts. The null graph fixes the zero-interaction reference; the crossed panels compare preferences across pairings. The scalar used during generation is $F_{\mathrm{IBC}} = F_{\mathrm{ctx}} + F_{\mathrm{chem}} + F_{\mathrm{syn}} + F_{\mathrm{prior}}$, where the context branch also depends on receptor–peptide cross-edges.

Learning pairing preferences. Crossed panels isolate pairing preferences by exchanging both partners. With $F_{ij} = F(Y_i, P_j)$, the rectangular contrast for peptide rows i, k and receptor columns j, ℓ is formulated as below.

$$\Delta_{ik;j\ell} = F_{ij} + F_{k\ell} - F_{i\ell} - F_{kj}. \tag{2}$$

Peptide-only and receptor-only shifts cancel (Appendix B.5). Auxiliary supervision learns the peptide-dependent terms retained for generation.

Training the receptor signal. The training objective combines fixed-pocket retrieval, a double-centred crossed-panel loss, and peptide, pocket, and role interventions. The crossed term learns pairing-dependent contrasts; the other terms select the scalar representative used for generation. Appendix C gives the full objective and Table 5 its fixed weights.

To avoid learning a coordinate-frame convention, candidate backbones are aligned to a fixed bank of eight proper-rigid anchors. Let $m_k \in \{0, 1\}$ denote the detached support mask and F_k the bounded IBC score at anchor k. With a fixed positive score scale β, the scalar used within each chamber is the following log-mean-exp:

$$F_P(Y) := \log \phi_P(Y) = \log \sum_{k=1}^{8} m_k \exp(\beta F_k) - \log 8. \tag{3}$$

This value is finite whenever a chamber contains a supported anchor. For the coordinate action, backpropagation holds the selected Kabsch rotations fixed and differentiates through the candidate coordinates and their Cα centering. Removing global translation and rotation and projecting onto the constraint tangent gives the local direction $u_P(S, X)$. At the same coordinates, re-evaluating F_P on each allowed single-site mutant gives the residue preferences. Both actions use a common local scalar at the current state, with alignment renewed after coordinates change (Appendix B.3).

3.4 EXECUTING THE GUIDED SAMPLER

A receptor-guided path target. Each seed screens 64 source states to form an anchor bank and $R = 8$ immutable output strata. Let H record the bank and admitted lineages, and let $Q^{\text{guide}}_{P,H,r}$ denote the supported path proposal in stratum r. For a path $Y_{0:J}$, set $B_{P,H}(Y) = F_{P,H}(Y)$ on admitted support and $-\infty$ otherwise, where $F_{P,H}$ is Eq. (3) conditional on H. For strata with positive terminal mass, the population path target is

$$\Gamma^{\text{FK}}_{P,H}(dY_{0:J}) = \frac{1}{R}\sum_{r=1}^{R} Z^{-1}_{P,H,r} Q^{\text{guide}}_{P,H,r}(dY_{0:J}) \exp[B_{P,H}(Y_J)], \tag{4}$$

where J is the final checkpoint and $Z_{P,H,r}$ normalizes each stratum. The receptor potential acts both inside the path proposal, shaping sequence and backbone, and at its endpoint, weighting completed candidates. Compiled support constrains the entire path. The prefix uses receptor-free actions with score-blind placement; potential-guided actions begin at $t^\star = 0.50$.

Sequence and backbone actions. The proposal combines coordinate guidance with sequence preferences. For reference drift b^0_t, token rates q^0_t, coordinate variance a_t, and an allowed fixed-coordinate jump $S \to S'$, its guided actions are

$$b^*_t = b^0_t + a_t u_P(S, X), \qquad q^*_t(S, S') = q^0_t(S, S') \exp[F_P(S', X) - F_P(S, X)], \tag{5}$$

where allowed token states preserve attachment locks and remain supported at the current coordinates. A substitution with a positive potential difference receives a higher rate than under the reference alone. Appendix B.6 derives the continuous gradient-and-rate construction; the local differentiation rule above specifies its numerical use here. Within each guided clock interval, the sampler first updates the backbone using the corrected drift in Eq. (5) and coordinate noise. Tangent projection and Newton retraction return this proposal to the active constraint leaf. At the resulting coordinates, a Gillespie simulation draws residue substitutions from the corrected rates, re-evaluating the reference and receptor scores as the sequence changes. Failed support checks retain the previous state. H2T constraints are sequence-independent; typed adapters additionally preserve attachment locks. The graph stays fixed, and coordinate residuals are recorded for certification.

Terminal branching and output. At $t_b = 0.76$, each lineage branches into eight continuations; one is drawn in proportion to its exponentiated terminal scalar. Finite execution retains each sampled prefix and normalizes over its children, whereas the population target in Eq. (4) reweights prefixes by expected terminal mass. The eight outputs retain stratum order, with unsupported or failed requests recorded as typed non-emission. Algorithm S1 and Appendix B.4 specify the finite law.

4 EXPERIMENTS

We first test whether receptor guidance during generation improves structural delivery over terminal selection alone. We then compare the complete native procedures, examine ring preservation and receptor response, and inspect structural examples across pockets and covalent programs.

4.1 EXPERIMENTAL SETUP

Datasets and competing baselines. The receptor-free data pool contains 39,128 eligible structures: 35,478 linear structures from PepBench and ProtFrag, and 3,650 H2T structures from CPCore and CPSea-PDB (lengths 6–16). Training excludes receptor coordinates, pocket features, and binding labels. The evaluation panel draws on experimentally observed complexes from 56 PepBench family components and 47 external LNR components. The native complex defines the receptor pocket and requested peptide length during preprocessing. Generation uses this frozen receptor-side input with the native peptide sequence and coordinates withheld. Appendix D.1 details source grouping, processing, and exclusions. The receptor-family panel comprises 103 components shared by HYBRIDGE, CP-Composer, and RFpeptides, all held out from HYBRIDGE training and model selection. Receptor sequence families and shared structure or source identifiers define components before allocation; Appendix D.1 gives both tracks, their sizes, and exclusion rules.

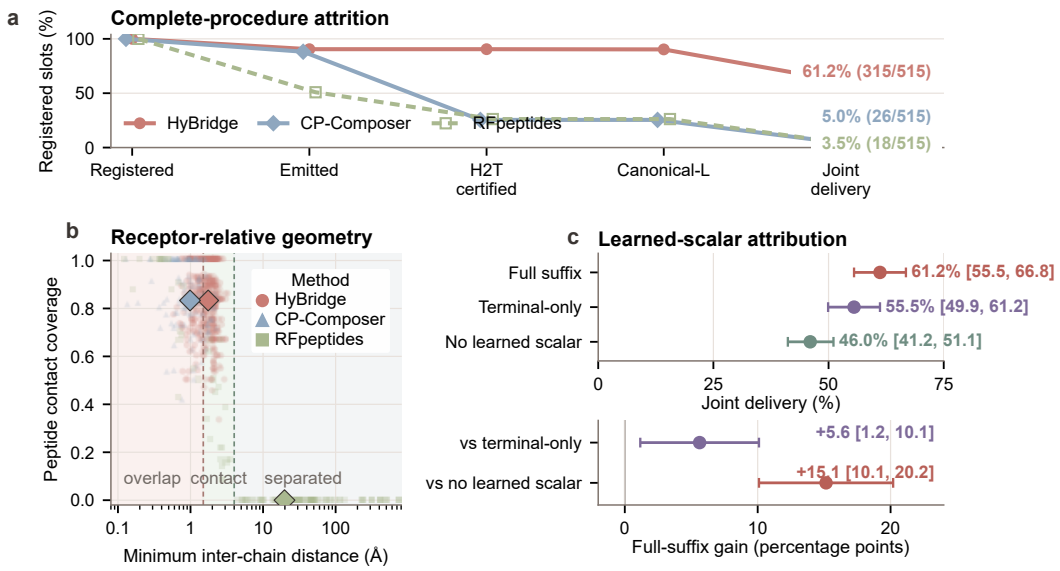


Figure 3: **Complete-procedure delivery and guidance-schedule comparison on the shared 103-component panel.** Five registered slots are evaluated per method and receptor. **a,** Retention from registration to joint delivery. **b,** Minimum interchain distance versus contact coverage; shaded regions distinguish overlap, contact, and separation, and diamonds mark medians. **c,** Joint delivery under full-suffix, terminal-only, and no-learned-scalar execution, followed by paired full-suffix gains. Checkpoint, support, placement, seeds, and slots are fixed. Error bars are 95% receptor-component-blocked confidence intervals.

Table 1: Guidance-schedule ablation study. Brackets show 95% receptor-component bootstrap intervals. The last column is the paired gain of full-suffix guidance over each control in percentage points. All schedules retain score-blind receptor placement.

Schedule	Joint delivery (%)	Full-suffix gain (pp)
No learned scalar	46.0 [41.2, 51.1]	+15.1 [10.1, 20.2]
Terminal-only	55.5 [49.9, 61.2]	+5.6 [1.2, 10.1]
Full suffix	**61.2** [55.5, 66.8]	—

Evaluation protocol. Each native procedure returns 24 indexed outputs per receptor. The first five, fixed before evaluation and used without endpoint reranking, form the common receptor-relative panel. The complete 24-output registration measures topology reliability as output opportunities accumulate. Missing and invalid slots stay in the denominator, and every export passes through the same materialization and interface criteria. Confidence intervals use a receptor-component-blocked bootstrap. Appendix D.1 gives the full protocol, and Algorithm S1 gives HyBridge execution.

4.2 Receptor guidance improves structural delivery

Structural delivery measures the yield of constructed receptor-relative candidates: an output must materialize as a canonical-L H2T peptide, contact the intended receptor, and contain no interchain heavy-atom pair below 1.5 Å. Each of the 103 components contributes five fixed slots. The schedule comparison tests when receptor information helps candidates reach this endpoint within the same generation procedure.

Guidance during generation gives the highest delivery rate in the matched comparison (Table 1). With the checkpoint, source procedure, support, placement, seeds, and output slots fixed, delivery rises from 46.0% without the learned scalar to 55.5% with terminal-only scoring and 61.2% with full-suffix guidance. Terminal scoring selects completed outputs; suffix guidance also changes the developing sequence and backbone. The gains are 15.1 percentage points over no learned scalar and

Table 2: Receptor coverage within the first five registered outputs under each native procedure. Counts indicate components with at least one evaluable or delivered structure; coverage is the delivered count divided by 103, expressed as a percentage. The structural endpoint is defined in Section 4.2.

Method	Evaluable components	Delivered components	Coverage (%)
CP-Composer	78/103	21/103	20.4
RFpeptides	76/103	11/103	10.7
HYBRIDGE	**100/103**	**96/103**	**93.2**

5.6 over terminal-only use, with both paired intervals above zero (Figure 3c). Higher materialization success and fewer severe overlaps after the common placement step account for the delivery gain. The same potential is thus useful both for selecting outputs and for guiding the changes that produce them. Figure 4e–f shows the associated receptor response and backbone changes. Appendix E.3 separates the outcomes by stage and source track.

Complete-procedure comparison. On the common set of 515 registered requests, HYBRIDGE delivers 315 structures (61.2%), compared with 26 (5.0%) for CP-Composer and 18 (3.5%) for RFpeptides (Figure 3a). Each method uses its native sampling budget. The common structural endpoint measures delivery from a registered request through emission and materialization to a receptor-contacting structure. Most of the cross-method difference appears before interface evaluation. HYBRIDGE supplies 465 canonical-L backbones, compared with 131 and 135 from the two baselines. Figure 3a–b follows these candidates through the subsequent contact and overlap checks. Appendix Tables 10 and 7, and Figures S13 and S9, resolve the individual failure stages and the accumulation of valid backbones over all 24 output opportunities.

Coverage across receptor components. The 315 delivered HYBRIDGE candidates span 96 of 103 receptor components, compared with 21 for CP-Composer and 11 for RFpeptides (Table 2). Delivery therefore extends across most of the panel. Appendix Figure S8 traces coverage from one to five outputs.

4.3 RECEPTOR SIGNAL ALONG THE SUPPORTED PATH

Path validity. Hard support preserves validity in all 480 paths, whereas every soft-penalty path leaves the fibre (Figure 4a). Yet both $\lambda = 10$ and 100 recover closure in the 24-path endpoint sentinel. Endpoint recovery and path preservation therefore describe different behavior. Compiled closure keeps the ring constraints active at the states used for scoring and resampling throughout generation.

Receptor-dependent response. The learned signal responds to the peptide–receptor pairing. On the held-out crossed panel, Eq. (2) reaches 0.928 rectangular-quartet accuracy. Removing peptide–receptor cross-edges sends the response exactly to zero. At the same checkpoint, removing the non-context readout branches lowers pairwise preference, Recall@1, and MRR. Pocket-counterfactual supervision strengthens spatial-chemistry recognition. Retraining three objectives gives a preference gain of 0.061 across five fit seeds, while ordinary matched-pair preference remains similar (Table 3). The intervention preserves pocket geometry and amino-acid composition, so the gain reflects sensitivity to the spatial arrangement of pocket chemistry. Together with the crossed-panel results, this shows that the signal distinguishes both the interacting partners and the chemical arrangement within a pocket. Appendix D.2, Table 9, and Figure S1 give estimators, intervals, and results.

Response during generation. The placement control translates the same closed cyclic source to the pocket centroid but performs no orientation or support search. The score-blind $\mathrm{SE}(3)$ lift raises valid placement by a mean 65 percentage points across two folds and leaves internal cyclic geometry unchanged. On matched source strata, suffix activation raises the frozen readout's $\Delta_\times$ over both controls (Figure 4e). Within each arm, median cyclic-aligned Cα RMSD between matched outputs for reciprocal receptors is 0.13 Å under either control and 2.08 Å under guidance (Figure 4f). Placement establishes the receptor frame; suffix guidance then changes both pairing preference and the geometry of matched outputs in response to the receptor. AF3 re-prediction additionally tests the response of generated sequences to intended and reciprocal receptors (Appendix H).

Table 3: Training-objective ablation on 60 held-out receptor components, evaluated over five fit seeds. Preference compares the factual pocket with geometry-fixed, composition-preserving chemistry shuffles. Brackets show 95% receptor-component bootstrap intervals.

Training objective	Preference	95% interval
Retrieval only	0.600	[0.519, 0.679]
Without pocket counterfactuals	0.611	[0.532, 0.689]
Complete objective	**0.672**	[0.593, 0.745]

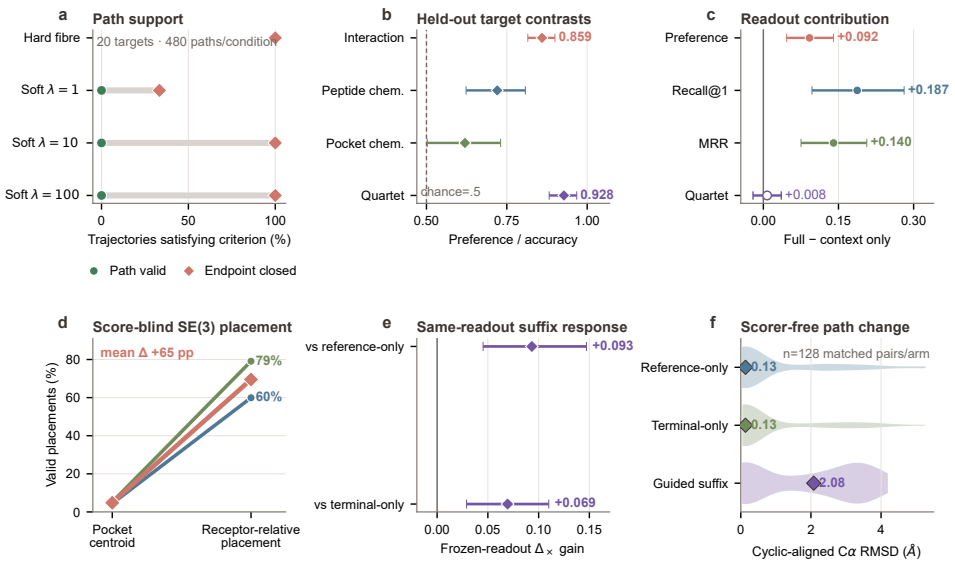


Figure 4: **Component-level controls and path-response diagnostics. a,** Hard and soft cyclic support: path validity over 480 paths (20 receptors), with endpoint closure from a 24-path sentinel control. **b,** Held-out interaction, chemistry-counterfactual, and rectangular-quartet controls; the no-cross-edge residual is exactly zero. **c,** Same-checkpoint removal of non-context readout branches. **d,** Receptor-relative placement in two independent folds. **e,** Paired changes in the frozen readout's $\Delta_\times$ when the scalar acts through the suffix. **f,** Scorer-free cyclic Cα RMSD between matched reciprocal-receptor outputs within each arm: reference-only, terminal-only, and guided suffix. Reciprocal directions give 128 records from 64 unique pairs per arm. The reference-only control retains score-blind receptor-relative placement. Error bars show 95% receptor-component-blocked confidence intervals; Appendix D details the controls.

4.4 Structural case studies across receptors and topologies

The structures in Figure 5 give the delivery results a molecular interpretation. Specifically, panels a–c show how the same H2T closure supports different receptor interfaces, while Figure 6a–c changes the covalent program itself.

Different interfaces under one ring topology. The H2T cases in panels a–c retain the same closure type while forming interfaces of different sizes. In all three, 70–80% of peptide residues contact the receptor, so contact extends across most of the ring. The 4TZM example has 190 interchain heavy-atom pairs within 4 Å and a BSA of 1422 Å^2, compared with 39 pairs and 818 Å^2 for 6SA8. The 1BJR interface falls between them at 118 pairs and 1030 Å^2. All pass canonical-L materialization and the common 1.5 Å interchain exclusion threshold.

Changing the covalent program. Figure 6a–c uses three construction routes: frozen-IBC selection of an internal-disulfide source, a terminal-disulfide rollout from a receptor-only pocket, and frozen-IBC sequence edits on an H2T-plus-disulfide source. The first two form one cycle; the last forms a cycle-rank-2 bicycle. All three programs produce receptor-relative structures with the specified

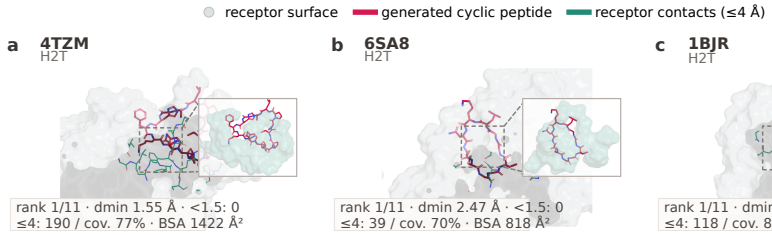


Figure 5: **Generated H2T complexes on 4TZM, 6SA8, and 1BJR (a–c).** Pink traces show the peptide, gray surfaces the receptor, and green traces contacts within 4 Å. Cards report fixed-pose ranks among 11 pockets, minimum interchain distance, close-pair counts, contact coverage, and BSA.

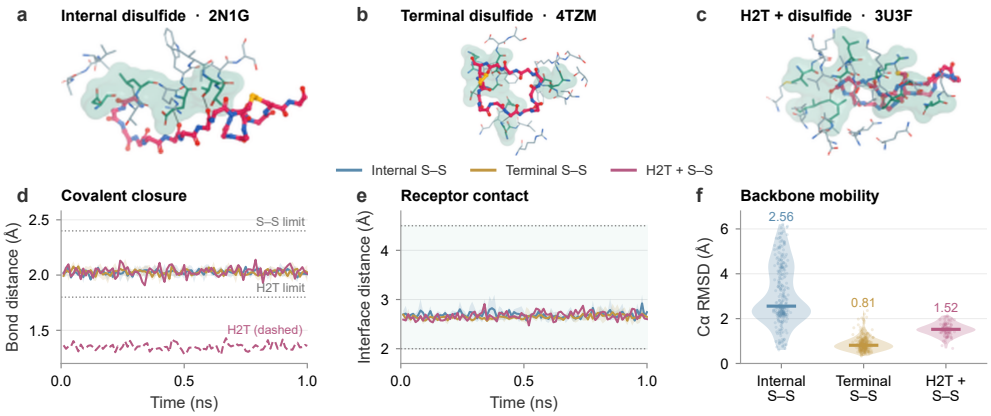


Figure 6: **Short-time persistence across additional topologies. a–c,** Internal S–S, terminal S–S, and H2T+S–S complexes. **d,** Closure distances over 1 ns; limits are 2.4 Å for S–S and 1.8 Å for H2T. **e,** Minimum interchain heavy-atom distance, with 2.0 Å overlap and 4.5 Å contact boundaries. **f,** Frame-level Cα RMSDs with medians. Curves and bands in **d,e** show means and ranges across three trajectories per monocycle; the bicycle has one. Frames are not independent replicates.

closures, receptor contacts, and no interchain pair below 1.5 Å. Appendix Figure S6 shows additional compiled products, and Table 16 summarizes the evidence for each topology. Across seven 1 ns trajectories, declared closure edges and canonical-L chirality persist in every stored frame. Receptor contact stays within the 4.5 Å criterion as backbone RMSD varies by topology (Figure 6d–f; Appendix Table 15). These trajectories distinguish covalent persistence from conformational rigidity.

5 Conclusion

HyBridge combines persistent covalent constraints with receptor guidance for joint sequence–structure generation. The compiled graph keeps accepted updates on cyclic states, while crossed-pair and chemistry supervision train a shared potential to guide backbone motion and residue substitutions. This construction achieves 61.2% structural delivery on 103 held-out receptor components, with at least one delivered complex for 96 components within five registered outputs. In matched schedules, guidance improves delivery by 15.1 percentage points over execution without the learned potential and by 5.6 over terminal-only use. The latter result captures the benefit of letting receptor preferences shape a candidate before selection. Path and receptor controls verify ring preservation and measure pairing and backbone responses. Structural examples span distinct H2T interfaces, disulfide rings, and bicyclic programs. Together, persistent covalent support and a shared receptor signal let sequence and shape adapt to the receptor while the ring remains intact.

AI USE STATEMENT

The authors conceived the research questions and mathematical formulation, designed the study, selected and froze every protocol, and made all scientific decisions. AI tools assisted feasibility exploration, repository navigation, implementation and analysis code, figure and prose drafting, mathematical exposition, preliminary result summaries, and consistency checks. The authors reviewed the analyses and interpretations, verified the reported evidence, and take full responsibility for the paper.

REPRODUCIBILITY STATEMENT

The full code and frozen-model package is available through the anonymous Figshare link `https://figshare.com/s/f90bcb408b99a21d3ae2`. It includes source data, tests, reproduction scripts, frozen checkpoints, and the GPU demo. Appendices B–E and H document sampling, training, and evaluation.

REFERENCES

Josh Abramson, Jonas Adler, Jack Dunger, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630:493–500, 2024. doi: 10.1038/s41586-024-07487-w.

Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 28266–28279. Curran Associates, Inc., 2022. doi: 10.52202/068431-2049. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/b5b528767aa35f5b1a60fe0aaeca0563-Abstract-Conference.html.

Vincent B. Chen, W. Bryan Arendall, III, Jeffrey J. Headd, Daniel A. Keedy, Robert M. Immormino, Gary J. Kapral, Laura W. Murray, Jane S. Richardson, and David C. Richardson. MolProbity: All-atom structure validation for macromolecular crystallography. *Acta Crystallographica Section D: Biological Crystallography*, 66(1):12–21, 2010. doi: 10.1107/S0907444909042073.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17695–17709. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/hash/940392f5f32a7ade1cc201767cf83e31-Abstract.html.

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. In *Advances in Neural Information Processing Systems*, volume 35, pp. 2406–2422. Curran Associates, Inc., 2022. doi: 10.52202/068431-0175. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/105112d52254f86d5854f3da734a52b4-Abstract-Conference.html.

Alexander Denker, Francisco Vargas, Shreyas Padhy, Kieran Didi, Simon Mathis, Vincent Dutordoir, Riccardo Barbano, Emile Mathieu, Urszula Julia Komorowska, and Pietro Lio. DEFT: Efficient fine-tuning of diffusion models by learning the generalised h-transform. In *Advances in Neural Information Processing Systems*, volume 37, pp. 19636–19682. Curran Associates, Inc., 2024. doi: 10.52202/079017-0620. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/22d258dfbdf840ccbf266bbc545dd95f-Abstract-Conference.html.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.

Marton L. Gyimesi, Victor Webersberger, Marco Del Giudice, Martin Voracek, and Ulrich S. Tran. One app to rule them all: A one-stop calculator and guide for 95 effect-size variants for two-group

comparisons of central tendency, variability, overlap, dominance, and distributional tails. *Advances in Methods and Practices in Psychological Science*, 8(2), 2025. doi: 10.1177/25152459251323186.

Jeremy Heng, Adrian N. Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential Monte Carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020. doi: 10.1214/19-AOS1914.

Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky T. Q. Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary Markov processes. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=RuP17cJtZo.

Dapeng Jiang, Xiangzhe Kong, Jiaqi Han, Mingyu Li, Rui Jiao, Wenbing Huang, Stefano Ermon, Jianzhu Ma, and Yang Liu. Zero-shot cyclic peptide design via composable geometric constraints. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 27553–27568. PMLR, 2025. URL https://proceedings.mlr.press/v267/jiang25d.html.

Haitao Lin, Odin Zhang, Huifeng Zhao, Dejun Jiang, Lirong Wu, Zicheng Liu, Yufei Huang, and Stan Z. Li. PPFLOW: Target-aware peptide design with torsional flow matching. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 30510–30528. PMLR, 2024. URL https://proceedings.mlr.press/v235/lin24z.html.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Stephen A. Rettie, David Juergens, Victor Adebomi, Yensi Flores Bueso, Qinqin Zhao, Alexandria N. Leveille, Andi Liu, Asim K. Bera, Joana A. Wilms, Alina Üffing, Alex Kang, Evans Brackenbrough, Mila Lamb, Stacey R. Gerben, Analisa Murray, Paul M. Levine, Maika Schneider, Vibha Vasireddy, Sergey Ovchinnikov, Oliver H. Weiergräber, Dieter Willbold, Joshua A. Kritzer, Joseph D. Mougous, David Baker, Frank DiMaio, and Gaurav Bhardwaj. Accurate de novo design of high-affinity protein-binding macrocycles using deep learning. *Nature Chemical Biology*, 21 (12):1948–1956, 2025. doi: 10.1038/s41589-025-01929-w.

Fanhao Wang, Yuzhe Wang, Laiyi Feng, Changsheng Zhang, and Luhua Lai. Target-specific de novo peptide binder design with DiffPepBuilder. *Journal of Chemical Information and Modeling*, 64 (24):9135–9149, 2024. doi: 10.1021/acs.jcim.4c00975.

Jingjie Zhang, Hanqun Cao, Haosen Shi, Mutian He, Yu Wang, Zijun Gao, Fang Wu, Xiaojun Yao, Chang-Yu Hsieh, Sinno Jialin Pan, Pranam Chatterjee, Chunbin Gu, and Pheng-Ann Heng. GeoCycler: Reward-aligned 3D diffusion for constraint-conditioned cyclic peptide design. *arXiv preprint arXiv:2605.23407*, 2026. URL https://arxiv.org/abs/2605.23407.

Yifan Zhao, Lang Qin, and Jintai Chen. APCyc: Property-informed design of cyclic peptides via automated cyclization. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, pp. 13020–13031. ACM, 2026. doi: 10.1145/3770855.3818908.

Xiangxin Zhou, Mingyu Li, Yi Xiao, Jiahan Li, Dongyu Xue, Zaixiang Zheng, Jianzhu Ma, and Quanquan Gu. Designing cyclic peptides via harmonic SDE with atom-bond modeling. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 78827–78858. PMLR, 2025. URL https://proceedings.mlr.press/v267/zhou25i.html.

A NOTATION AND SUPPORTED TOPOLOGIES

Table 4 collects the objects used throughout the paper. A rollout keeps one topology identifier fixed from source to output. H2T has matched receptor-conditioned supervision and a held-out receptor-family split. The remaining compiled graphs test the same support machinery in structural pilots.

Table 4: Core notation.

Symbol	Meaning
P	receptor pocket with coordinates, identities, frames, and chemistry
τ	discrete topology program, fixed for one rollout
S	amino-acid tokens in $\{1, \dots, 20\}^L$
$\bar{Y}, Y$	intrinsic state modulo rigid pose; receptor-frame posed state
X	N/Cα/C/O backbone coordinates in $\mathbb{R}^{L\times 4\times 3}$
G_τ, A_τ	compiled covalent graph and attachment-identity locks
r	registered target-free constraint reference fixed along a particle path
$\mathcal{M}_{\tau,S,r}$	smooth coordinate leaf for fixed τ, S, r
$c_{\tau,r}$	active coordinate-constraint residual on one leaf
Ω_τ	discrete union of admissible sequence–coordinate leaves
$\bar{\mu}_{\tau,r}$	intrinsic target-free measure on one registered fibre
$\bar{\mu}_\tau$	receptor-blind mixture over registered intrinsic measures
K_P^{lift}	score-blind proper-rigid receptor support kernel
$\mu^0_{\tau,P}$	posed base measure induced by $\bar{\mu}_\tau K_P^{\mathrm{lift}}$
F_{int}	selected cross-edge representative; rectangular contrasts supervise its quotient invariants
R_τ, C	peptide interface readout and receptor IBC encoding
(b_t, q_t, a_t, F_t)	hybrid drift, token rates, coordinate variance, and target potential
H	realized source-screen record, anchor bank, and admitted particle paths
$\Gamma^{\mathrm{FK}}_{P,H}$	population Feynman–Kac path target conditional on H
$\widehat{\Gamma}^{R,Q}_{P,H}$	exact finite request law over R returned paths or $\perp$

B SAMPLING DEFINITIONS AND SUPPORT PROPERTIES

B.1 FINITE SAMPLING PROCEDURE

The source-screen record H contains the anchor bank and admitted source lineages. Index its immutable output strata by r and write $B_{P,H}$ for the terminal log boundary, equal to $-\infty$ outside admitted support. Algorithm S1 uses $K = R$ strata and Q terminal children per stratum; Appendix B.4 defines the resulting request law.

At the nonterminal checkpoints $0 = t_0 < \cdots < t_J < T$, let ψ_j be the auxiliary lookahead and let $B(Y_T)$ be the terminal log boundary. Along each selected ancestry, the implemented increments are

$$\Delta\ell_0 = \psi_0(Y_0), \tag{6}$$

$$\Delta\ell_j = \psi_j(Y_j) - \psi_{j-1}(Y_{j-1}), \tag{7}$$

$$\Delta\ell_T = B(Y_T) - \psi_J(Y_J). \tag{8}$$

They telescope exactly to $B(Y_T)$ before floating-point error, including after particle cloning because the stored denominator follows the selected ancestry. This identity accounts for weights along a realized guided proposal. The released configuration uses eight immutable one-particle strata, making intermediate resampling the identity; the usual infinite-particle SMC limit is therefore not its operating regime. Each retained particle path produces $Q = 8$ continuations only at the final checkpoint. For continuations $C_{r,1:Q}$, let $w_{r,q} = \exp[B_{P,H}(C_{r,q})]$. A nonempty stratum draws

$$\Pr(I_r = q \mid H, C_{r,1:Q}) = \frac{w_{r,q}}{\sum_{m=1}^{Q} w_{r,m}}. \tag{9}$$

If any stratum has no supported terminal mass, the request returns $\perp$.

Algorithm S1 Supported target-conditioned H2T generation.

Input receptor P, request (τ, L, M), source bank $\mathcal{B}$, reference actions (b_t^0, q_t^0), and target scalar F_P; $\mathcal{T} = (0, .24, .50, .76, 1)$, $t^\star = .50$, $K = Q = 8$.

1 $(G_\tau, A_\tau) \leftarrow \mathrm{Compile}(\tau, L, M)$; build H from 64 source previews in frozen order without target-score ranking. The first K preview-supported endpoints form the anchor bank; using that bank, the first K quotient-supported source states initialize the particle paths. Either count below K yields FAIL.

2 Assign one immutable source stratum to each particle path and initialize its log weight.

3 **for** consecutive checkpoints $[t_j, t_{j+1}] \subset \mathcal{T}$:

4 If $t_{j+1} = 1$, branch each particle into Q children that inherit its reference and stratum.

5 On each fine step, use (b_t^0, q_t^0) for $t < t^\star$. For $t \geq t^\star$, construct the numerical guided actions from the masked anchor scalar's fixed-rotation coordinate derivative and allowed token differences.

6 Update coordinates by Euler–Maruyama, tangent projection, and retraction; update tokens by masked Gillespie simulation. A proposal self-loops if it leaves $\Omega^{\varepsilon}_{\tau,r_k}$, changes G_τ, or violates A_τ.

7 Apply the score-blind proper-rigid pose map; add $\Delta\Phi_j^k$ to each log weight ℓ^k. Resampling, when active, selects ancestry only within the same stratum.

8 At intermediate checkpoints, apply the ESS rule within strata. At the terminal checkpoint, draw I_r by Eq. (9) in every stratum. If any stratum has zero supported mass, return FAIL.

9 **end for**; serialize each selected state, reload it, and certify the explicit H2T closure. Any unrecoverable error returns FAIL.

10 **return** the ordered tuple $(Y_1^1, \ldots, Y_1^K)$.

The auxiliary twist may change stored weights before $t^\star$, but with one particle path per stratum it does not select among prefixes or alter their transition kernel; the local actions remain (b_t^0, q_t^0). Since F_P in Eq. (3) already includes the fixed score scale β, the released twist is $\psi_t(Y) = tF_P(Y) + \mathbf{1}\{t \in \{.50, .76\}\} \log m_P(\widetilde{Y})$, where one frozen reference shadow $\widetilde{Y}$ supplies the positive conformal margin $m_P \in [e^{-30}, 1]$ and is never emitted. The terminal potential is $B_P(Y) = F_P(Y) + \log \mathsf{H}_P(Y, 1)$, where $\mathsf{H}_P(Y, 1) \in \{0, 1\}$ is the terminal support indicator, and $\Delta\Phi_j = \psi_{t_{j+1}} - \psi_{t_j}$ for $t_{j+1} < 1$ and $\Delta\Phi_j = B_P - \psi_{t_j}$ at $t_{j+1} = 1$. The released coordinate variance is $a = .036$.

Coordinate increments are capped at 0.10 Å RMS; full-constraint solves use effective rcond 10^{-10} and residual tolerance 10^{-10}. At the terminal boundary, the emitted pose uses the supported anchor with the largest contribution to Eq. (3); its log-mean-exp supplies the terminal weight. Token rates are recomputed after each Gillespie event, with at most sixteen events per fine step. Proper-rigid placement follows every checkpoint and preserves intrinsic geometry. With one particle path per stratum, intermediate resampling is the identity; terminal selection acts within the $Q = 8$ child pool of each immutable stratum. Only F_P changes coordinate and token actions; m_P and H_P affect weights and support. If any stratum fails, the request is recorded as a typed non-emission; $\perp$ never enters the molecular state space.

B.2 Compiled support and numerical transport

Let $c_{\tau,r}(S, X) = 0$ collect the active coordinate constraints on one leaf and $J_X = D_X c_{\tau,r}(S, X)$. For a positive-definite metric W_X, define

$$\Pi_X^W v = v - W_X^{-1} J_X^\top (J_X W_X^{-1} J_X^\top)^+ J_X v. \tag{10}$$

Let $B = J_X W_X^{-1} J_X^\top$. Since W_X is positive definite, $\mathrm{range}(B) = \mathrm{range}(J_X)$, and BB^+ is the orthogonal projector onto this range. Therefore $J_X \Pi_X^W v = (I - BB^+) J_X v = 0$, so the projected action is tangent to first order. The implementation uses a tolerance-controlled Moore–Penrose pseudoinverse and records the resulting tangent residual. A bounded minimum-metric-norm Newton solver retracts each proposal and records its final constraint residual.

Let Ω_τ^ε denote the union of states that satisfy the attachment locks and whose active constraint residual has infinity norm at most ε.

Proposition 1 (finite-step fibre invariance). Suppose (i) $Y_0 \in \Omega_\tau$; (ii) every accepted coordinate proposal is retracted on its active leaf to residual at most ε; (iii) every token jump either keeps X fixed and verifies the destination leaf, or applies a topology-specific coordinate map followed by

destination-leaf retraction; and (iv) unsupported proposals are self-loops. Then every accepted state belongs to the numerical fibre Ω_τ^ε and the covalent graph never changes.

Proof. The initial state is supported. Each accepted coordinate update retracts to its active leaf; each token update verifies the destination leaf after any required coordinate map and retraction. Unsupported proposals retain the preceding state. Induction therefore preserves Ω_τ^ε, and every operation leaves G_τ fixed.

The certificate reports three nested levels. A compiled edge fixes graph topology, the closure residual tests its seam, and the full certificate checks every modeled backbone bond and required stereochemical predicate.

Proposition 2 (proper-rigid lift preserves intrinsic state). For $g = (R, t) \in \mathrm{SE}(3)$ with $R^\top R = I$ and $\det R = 1$, set $gX = XR^\top + \mathbf{1}t^\top$. All pairwise distances, bond angles, signed chiral volumes, G_τ, and A_τ are unchanged. Therefore a receptor-relative lift can alter interface placement without changing the intrinsic cyclic state.

Proof. Translation cancels in displacements, and orthogonality of R preserves their inner products, hence distances and angles. Signed chiral volumes are multiplied by $\det R = 1$. The lift leaves G_τ and A_τ fixed, preserving the intrinsic cyclic state.

The implementation first performs a finite score-blind proper-rigid receptor support search and rejects unsupported lifts. It then uses a finite bank of eight proper-rigid anchors for the IBC scalar and explicitly projects global translation and rotation from its coordinate action. Backpropagation fixes support masks and Kabsch rotations while retaining candidate-coordinate and centering derivatives within each supported evaluation. The eight anchors provide the finite quotient approximation used at runtime.

For a rigid pose with peptide–receptor heavy-atom distances d_a, peptide Cα–pocket Cα distances d_i, and translation u, the search minimizes

$$50\ \mathrm{mean}_a[2.2 - d_a]_+^2 + 1 - \mathrm{mean}_i\,\sigma\left(\frac{8 - d_i}{0.25}\right) + 10^{-3}\|u\|_2^2.$$

Eight quaternion restarts are optimized for 300 Adam steps at learning rate 0.03. A pose is admitted when its minimum heavy-atom distance is at least 2 Å and at least 80% of peptide Cα atoms lie within 8 Å of the pocket; the first supported restart is retained.

B.3 Local receptor actions in finite execution

The continuous identities in Appendix B.6 use the exact scalar gradient on each active leaf. The released sampler computes its coordinate action with the selected Kabsch rotations held fixed during differentiation, followed by rigid-component removal and constraint-tangent projection. Candidate coordinates and their Cα centering remain differentiable; token actions use re-evaluated scalar differences. These operations specify the numerical derivative used in the reported rollouts.

The common local score used by the two actions. At current coordinates X_0, hold the selected anchor rotations R_k^0 fixed and let $m_k^0(S)$ be the anchor support evaluated separately for each token state at X_0, then held fixed for coordinate differentiation. With differentiable candidate Cα centering, the resulting local score is

$$\widetilde{F}_{X_0}(S, Z) = \log \sum_{k=1}^{8} m_k^0(S) \exp\{\beta F_{\mathrm{IBC}}^{\mathrm{bounded}}(S, T_{k;X_0}(Z), P)\} - \log 8, \tag{11}$$

where $T_{k;X_0}(Z) = (Z - \mathbf{1}\bar{Z}_{C\alpha})R_k^0 + \mathbf{1}\bar{X}_{k,C\alpha}^{\mathrm{anchor}}$. At $Z = X_0$, this has the same value as the recomputed anchor score for each supported token state. The coordinate action is its derivative projected onto horizontal and constraint-tangent directions; allowed token corrections use $\widetilde{F}_{X_0}(S', X_0) - \widetilde{F}_{X_0}(S, X_0)$. Unsupported token transitions are masked out. Thus both actions use one local score at the current state, with alignment renewed after coordinates change. The continuous telescoping result concerns a fixed scalar and its exact derivative, including alignment dependence; the local score above specifies the finite sampler's differentiation convention.

The sampler combines this action with an eight-anchor chamber-wise pose quotient, finite receptor-support search, 50-step integration, pseudoinverse projection, nonlinear retraction, eight one-particle strata with eight terminal children each, a finite lookahead schedule, side-chain reconstruction, and constrained all-atom materialization. These choices define the finite law evaluated in the experiments. Terminal selection normalizes within each fixed prefix stratum, preserving its proposal weight; the population Feynman–Kac target instead reweights prefixes by their expected terminal mass (Appendix B.4).

B.4 Base measure and finite path law

This subsection collects the formal distributional definitions used by the generator in Section 3.

A score-blind base and endpoint tilt. For intrinsic state $\bar{Y}$, posed state $Y = (S, X)$, receptor P, length L, and optional motif M, let $\bar{\mu}_{\tau,r}$ be a target-free measure on the fibre indexed by r. The mixing weights $\pi_\tau(r \mid L, M)$ are receptor-blind, and K_P^{lift} is a score-blind proper-rigid placement kernel. They define

$$\bar{\mu}_\tau = \sum_{r \in \mathcal{R}_{\tau,L}} \pi_\tau(r \mid L, M)\bar{\mu}_{\tau,r}, \qquad \mu^0_{\tau,P}(dY) = \int K_P^{\mathrm{lift}}(dY \mid \bar{Y})\, \bar{\mu}_\tau(d\bar{Y}), \tag{12}$$

with $\operatorname{supp} \bar{\mu}_{\tau,r} \subseteq \Omega_{\tau,r}$. Writing $R_\tau(Y)$ and $C(P)$ for peptide-interface and receptor-pocket encodings, let $F_{\mathrm{int}}(R, C)$ denote the scalar representative selected by the cross-edge architecture and auxiliary objectives; its rectangular contrasts are supervised by Eq. (25). One useful endpoint diagnostic is

$$p^\star(dY \mid P, \tau, L, M) = \frac{\mu^0_{\tau,P}(dY \mid L, M)\, \exp[F_{\mathrm{int}}(R_\tau(Y), C(P))]}{Z^\star(P, \tau, L, M)}, \tag{13}$$

where $Z^\star$ normalizes the measure. The tilt combines a posed base built without receptor scoring with one interaction representative. Equation (13) specifies an endpoint tilt, and Eq. (4) gives the population target of the guided path proposal. Equation (9) and Algorithm S1 specify the finite procedure.

The population Feynman–Kac path target is given in Eq. (4). Its proposal $Q^{\mathrm{guide}}_{P,H,r}$ uses the anchor bank and admitted lineages recorded in H. Receptor-free molecular actions precede $t^\star = 0.50$; scalar guidance then acts on the suffix, with score-blind proper-rigid placement and compiled support active throughout. The rollout has length J. The finite sampler uses $R = 8$ one-particle strata and $Q = 8$ terminal continuations per lineage, defining the request law $\widehat{\Gamma}^{8,8}_{P,H}$ on $\Omega^8_\tau \sqcup \{\perp\}$. An empty stratum or a failed rollout, serialization, or certificate yields $\perp$. Equation (9) and Algorithm S1 specify this finite procedure. It retains the proposal weights of fixed prefixes; the population target in Eq. (4) reweights them by expected terminal mass. Increasing the number of continuations within one prefix leaves this distinction intact.

B.5 Interaction quotient and base-measure completion

Lemma 1 (rectangular-contrast gauge invariance). For a finite crossed panel, let $F_{ij} = F(Y_i, P_j)$ and $F'_{ij} = F_{ij} + a_i + b_j$. For arbitrary rows i, k and columns j, ℓ, every rectangular contrast is unchanged:

$$F'_{ij} + F'_{k\ell} - F'_{i\ell} - F'_{kj} = F_{ij} + F_{k\ell} - F_{i\ell} - F_{kj}. \tag{14}$$

Proof. Substituting $F'_{ij} = F_{ij} + a_i + b_j$ adds $(a_i + a_k - a_i - a_k) + (b_j + b_\ell - b_\ell - b_j) = 0$ to the rectangular contrast, leaving the value unchanged.

Lemma 2 (converse on a complete connected panel). If two real matrices F and G have the same rectangular cross differences, then $D = F - G$ can be written $D_{ij} = a_i + b_j$.

Proof. Fix row 1 and column 1. The zero cross difference for rows $i, 1$ and columns $j, 1$ gives $D_{ij} + D_{11} - D_{i1} - D_{1j} = 0$. Hence $D_{ij} = D_{i1} + (D_{1j} - D_{11})$. Setting $a_i = D_{i1}$ and $b_j = D_{1j} - D_{11}$ proves the claim for every i, j.

The complete numerical contrast table therefore determines the class

$$[F] = \{F + a(Y) + b(P) : a, b \text{ arbitrary}\}. \tag{15}$$

The matched-pair quantity $F_{ii} + F_{jj} - F_{ij} - F_{ji}$ is one rectangular contrast but does not alone provide the converse. Crossed supervision identifies this class; the base measure, exact-null normalization, and auxiliary losses select the representative used for generation, as described next.

Proposition 3 (base-measure completion). Let $\bar{\mu}_\tau = \sum_r \pi_\tau(r \mid L, M)\bar{\mu}_{\tau,r}$ be the intrinsic target-free mixture and let $\mu^0_{\tau,P} = \bar{\mu}_\tau K_P^{\mathrm{lift}}$ be its score-blind posed measure. Assume $0 < Z_P = \int \exp F(Y, P)\, \mu^0_{\tau,P}(dY) < \infty$. Then

$$p_F(dY \mid P, \tau) = Z_P^{-1} \exp F(Y, P)\, \mu^0_{\tau,P}(dY) \tag{16}$$

is a normalized conditional law. Replacing F by $F + b(P)$ leaves this law unchanged. Replacing F by $F + a(Y)$ generally changes the law and is equivalent to changing the posed base measure to $\exp a(Y)\mu^0_{\tau,P}(dY)$ followed by normalization.

Proof. Normalization gives total mass one. A receptor-only shift $b(P)$ multiplies both the numerator and Z_P by $\exp b(P)$, so it cancels. A peptide-only shift replaces the base measure by $\exp a(Y)\mu^0_{\tau,P}(dY)$ before normalization and generally changes relative molecular mass unless a is constant almost everywhere on its support.

Every intrinsic component $\bar{\mu}_{\tau,r}$ is fibre-supported. The proper-rigid lift preserves intrinsic topology and stereochemistry while its predeclared support tests constrain receptor-relative placement. This is equivalent to restricting an ambient intrinsic measure with a fibre indicator; we use the supported-measure notation throughout.

An exact-null graph, formed by removing all receptor–peptide cross-edges, sets the origin of the implemented representative through $F_{\mathrm{int}}(H^{\varnothing}) = 0$. Gauge-invariant and utility losses then select among representatives, including peptide-only terms that become active in the coupled graph (Eq. (25)). Proposition 3 shows how that representative combines with the base measure; the algebra identifies interaction contrasts rather than a unique physical energy.

B.6 One scalar and two hybrid actions

On the active coordinate leaf $\mathcal{M}_{\tau,S,r}$, let the target-free time-dependent generator act on a smooth test function as

$$\mathcal{L}^0_t f = b^0_t \cdot \nabla_X f + \tfrac{1}{2} a_t \Delta_{\mathcal{M}_{\tau,S,r}} f + \sum_{S' \neq S} q^0_t(S, S')\,[f(S', X) - f(S, X)]. \tag{17}$$

Here $a_t > 0$ is the inference variance; its stochastic marginals need not equal those of deterministic flow matching.

Proposition 4 (hybrid Doob generator). For positive h_t in the domain of $\mathcal{L}^0_t$, define

$$\mathcal{L}^h_t f = h_t^{-1}\mathcal{L}^0_t(h_t f) - h_t^{-1} f \mathcal{L}^0_t h_t. \tag{18}$$

Writing $F_t = \log h_t$ yields

$$b^h_t = b^0_t + a_t \operatorname{grad}_{\mathcal{M}_{\tau,S,r}} F_t, \tag{19}$$

$$q^h_t(S, S') = q^0_t(S, S') \exp\{F_t(S', X) - F_t(S, X)\}. \tag{20}$$

Proof. The manifold product rule gives $\Delta(h_t f) = h_t \Delta f + f \Delta h_t + 2 \operatorname{grad} h_t \cdot \operatorname{grad} f$. After division by h_t and subtraction of $h_t^{-1} f \mathcal{L}^0_t h_t$, the diffusion part is the reference diffusion plus $a_t \operatorname{grad} \log h_t \cdot \operatorname{grad} f$. For a jump $S \to S'$, the same subtraction gives

$$q^0_t(S, S') \frac{h_t(S', X)}{h_t(S, X)} [f(S', X) - f(S, X)].$$

Thus the drift and jump rates have the stated forms.

If h_t also satisfies $(\partial_t + \mathcal{L}^0_t) h_t = 0$, the transform is the standard local Doob conditioning of the reference process. These gradient and rate formulas define the continuous control construction. The numerical actions in Appendix B.3 specify the guided proposal used in the population reference, Eq. (4), and the finite execution law, Eq. (9).

Proposition 5 (mixed-path telescoping). For a time-independent scalar F, consider a path composed of smooth coordinate segments on one active leaf at fixed sequence and accepted fixed-coordinate token jumps $(S^-, X) \to (S^+, X)$ whose pre- and post-states belong to their respective leaves. The continuous potential contribution is the line integral of $\mathrm{grad}_{\mathcal{M}_{\tau,S,r}} F$, equivalently of $a_t^{-1}(b_t^h - b_t^0)$, and a jump contributes $F(S^+, X) - F(S^-, X)$. Summing both types gives

$$\sum_{\text{segments}} \int \mathrm{grad}_{\mathcal{M}_{\tau,S,r}} F \cdot dX + \sum_{\text{jumps}} \big(F(S^+, X) - F(S^-, X)\big) = F(Y_T) - F(Y_0). \tag{21}$$

Proof. On a smooth segment from (S, X_a) to (S, X_b), the fundamental theorem of line integrals gives $\int \mathrm{grad}_{\mathcal{M}_{\tau,S,r}} F \cdot dX = F(S, X_b) - F(S, X_a)$. At a fixed-coordinate token jump, the discrete term is the difference between the post- and pre-jump values of F. Ordering these increments along the path makes every interior endpoint appear once with each sign, leaving $F(Y_T) - F(Y_0)$.

For a time-dependent F_t, the corresponding identity is

$$\sum_{\text{segments}} \int \mathrm{grad}_{\mathcal{M}_{\tau,S,r}} F_t \cdot dX + \sum_{\text{jumps}} \Delta_S F_t + \int_0^T \partial_t F_t(Y_t)\, dt = F_T(Y_T) - F_0(Y_0). \tag{22}$$

On each smooth segment, the chain rule gives $dF_t(S, X_t)/dt = \mathrm{grad}\, F_t \cdot \dot{X}_t + \partial_t F_t$. Integrating this identity and adding the fixed-coordinate jump differences proves the time-dependent formula. Thus a closed fixed-time mixed loop has zero circulation for the exact scalar within a differentiable chamber. The switch at $t^\star$ separates two such chambers. Equation (21) applies within each phase; crossing the switch would introduce a formal boundary jump that is not a runtime action. The identity characterizes the exact continuous controls derived from one scalar. Appendix B.3 specifies the numerical actions used by the finite sampler.

C Training and inference details

Reference path and geometry objective. The reference learns receptor-free motion within the compiled constraint leaf. For an observed endpoint X_1, a proper Kabsch map removes global rigid pose from the coordinate target. A locally corrupted source is projected onto its active coordinate leaf, and every point along the interpolation is retracted to that leaf. For $t \in [0, 1]$, let T^* denote the proper Kabsch alignment and $\mathrm{Retr}_{\tau,r}$ the numerical retraction,

$$X_t = \mathrm{Retr}_{\tau,r}\big((1-t)X_0 + tT^*X_1\big), \qquad u_t = \frac{d}{dt}\,\mathrm{Retr}_{\tau,r}\big((1-t)X_0 + tT^*X_1\big). \tag{23}$$

The concurrent token state is S_t, and differentiating $c_{\tau,r}(S_t, X_t) = 0$ shows that u_t is tangent. Writing v_θ for the learned coordinate velocity, Π for tangent projection, and $\epsilon > 0$ for a fixed stabilizer, the geometry network is trained by projected flow matching,

$$\mathcal{L}_{\mathrm{geom}} = \mathbb{E}\left[\|\Pi_{T_{X_t}\mathcal{M}_{\tau,S_t,r}} v_\theta(X_t, S_t, t) - u_t\|^2 / (\|u_t\|^2 + \epsilon)\right]. \tag{24}$$

The expectation is over target-free training endpoints, their corruptions, and the sampled interpolation time. The velocity describes motion along the active constraint leaf.

The anonymous release includes a machine-readable `docs/CODE_MAP.json` linking the mathematical operations to their implementation. The accompanying source package also provides a readable index in `verification/method_code_map.md`.

Within the outer training split, component hashes define separate partitions: five folds for the reference and six for the target stack. Reference fitting uses folds 1–4; the target stages use the allocations described below.

The final target-free reference contains ten geometric transformer blocks, hidden width 384, and 12 attention heads. A frozen ESM-2 35M model provides token features. The released checkpoint is a 1,200-step posterior-rate refinement initialized from the preceding 1,200-step reference checkpoint; the optimizer is reinitialized for the final stage. The final stage uses batch size 8, geometry weight 1, rate weight 1, and auxiliary token-posterior weight 0.25. The reference receives no receptor feature or target score.

H2T inference uses 50 reference-clock steps, eight particles per registered seed, three registered seeds, 64 score-blind source proposals, and the frozen scalar scale $\beta = 7.17265462069118$ from the released temperature asset. The 50-step grid discretizes the hybrid reference. Coordinate actions are evaluated at clock steps, with additional sequence-action evaluations at Gillespie events.

The activation time comes from a score-blind support check on the grid $\{0, 0.24, 0.50, 0.76, 1\}$. We select the first point at which every declared stratum reaches support 0.75. Length-11 support is $0.625 < 0.75$ at 0.24, making 0.50 the first qualifying point and fixing $t^\star = 0.50$. Receptor guidance begins there; compiled topology projection and retraction remain active along the entire path.

Three seeds yield 24 ordered outputs for the emission, support, and diversity analyses. Before evaluation, a particle-major and then seed-major rule assigns axes 0–4 to positions 1–5 in the common endpoint. Later outputs do not replace these positions, and no endpoint reranking is applied.

Target-adapter training objective. For a complete pair panel, let F be the interaction matrix, $D = \alpha\ell_0\mathbf{1}^\top + F$ the completed design score, and $H = I - \mathbf{1}\mathbf{1}^\top/n$. Here ℓ_0 is the receptor-free design-logit vector, α is a learned global scale, $\mathbf{1}$ is the length-n all-ones vector, and n is the panel size. The target adapter was trained with fixed nonnegative weights $\lambda_\bullet$,

$$\mathcal{L}_{\text{target}} = \mathcal{L}_{\text{design}}(D) + \lambda_{\text{int}}\mathcal{L}_{\text{cross}}(HFH) + \lambda_p\mathcal{L}_{\text{cf}}^{\text{peptide}} + \lambda_P\mathcal{L}_{\text{cf}}^{\text{pocket}} + \lambda_r\mathcal{L}_{\text{cf}}^{\text{role}}. \tag{25}$$

$\mathcal{L}_{\text{design}}$ performs fixed-pocket multi-positive retrieval, and $\mathcal{L}_{\text{cross}}$ is the bidirectional multi-positive loss on the double-centred interaction matrix. The remaining three terms compare factual scores with peptide, pocket, and role interventions. The HFH term is invariant to arbitrary row and column additions and supervises the crossed contrasts; the other losses choose the score representative used during generation.

Table 5: Executed target-adapter objective. The design term has unit weight; all other values are fixed coefficients in the frozen training configuration.

Term	Weight	Role
$\mathcal{L}_{\text{design}}(D)$	1.00	fixed-pocket multi-positive retrieval
$\mathcal{L}_{\text{cross}}(HFH)$	0.25	gauge-invariant interaction contrast
$\mathcal{L}_{\text{cf}}^{\text{peptide}}$	0.10	peptide intervention
$\mathcal{L}_{\text{cf}}^{\text{pocket}}$	0.10	pocket intervention
$\mathcal{L}_{\text{cf}}^{\text{role}}$	0.20	typed-role intervention

Losses and intervention inputs. Rows of the complete score matrix index peptides and columns index pockets. Let $M_{ij} = 1$ when the two entries share a receptor-component identity. For matrix A, define the directional retrieval loss

$$\mathcal{R}(A, M) = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}}\left[\log\sum_j e^{A_{ij}} - \log\sum_{j:M_{ij}=1} e^{A_{ij}}\right], \qquad \mathcal{I} = \{i : 0 < \textstyle\sum_j M_{ij} < n\}. \tag{26}$$

The executed losses are

$$\begin{aligned}
\mathcal{L}_{\text{design}}(D) &= \mathcal{R}(D^\top/T_d, M^\top),\\
\mathcal{L}_{\text{cross}}(HFH) &= \tfrac{1}{2}\{\mathcal{R}(HFH/T_c, M) + \mathcal{R}((HFH)^\top/T_c, M^\top)\},\\
\mathcal{L}_{\text{cf}}^a &= n^{-1}\sum_i \log\left(1 + \exp[(F_i^a - F_{ii})/T_a]\right), \quad a \in \{\text{peptide}, \text{pocket}, \text{role}\}.
\end{aligned} \tag{27}$$

Here $T_d = T_c = T_{\text{peptide}} = T_{\text{pocket}} = 0.1$, while the role loss uses $T_{\text{role}} = 1$. The base-score scale is $\alpha = 0.1 + 1.9\,\text{sigmoid}(a_0)$, learned during adapter fitting and frozen at inference. The receptor-free sequence-prior score sums the observed-token log probabilities on the given backbone and divides by $\sqrt{L}$; the frozen prior is evaluated at clock 0.95 and noise scale 0.12.

Peptide counterfactuals cyclically shift residue identities by one position among valid peptide residues, holding coordinates fixed. Pocket counterfactuals apply the same shift to pocket identities and their aligned ESM features. Role counterfactuals instead shift the pocket's chemical-role vectors and uncertainties, retaining its identities, ESM features, and geometry. Off-diagonal panel entries place the donor receptor in the recipient peptide's Kabsch frame; diagonal entries keep their experimental frame. All three counterfactual losses compare the resulting bounded scores with the factual diagonal.

Four branches of the receptor score. Let c_{ij} and u_{ij} be the frozen context and typed-chemistry pair representations on an active peptide–pocket edge. Separate LayerNorm–linear–SiLU projections map them to width 192, giving h^c_{ij} and h^u_{ij}. The context and chemistry heads score these vectors; the synergy head scores their elementwise product $h^c_{ij} \odot h^u_{ij}$. Each head consists of LayerNorm, dropout 0.15, and a scalar linear output. For any edge score z, aggregation is

$$\mathcal{A}(z) = \frac{1}{|I_A|} \sum_{i \in I_A} \frac{1}{|A_i|} \sum_{j \in A_i} z_{ij}, \qquad I_A = \{i : |A_i| > 0\}, \tag{28}$$

where A_i contains active pocket neighbors of peptide residue i. This gives F_{ctx}, F_{chem}, and F_{syn}. The prior branch is $F_{\mathrm{prior}} = \eta F_{\mathrm{typed}}$, where $\eta = \mathrm{sigmoid}(a_1)$ is learned by the adapter and F_{typed} is the frozen typed model's aggregated chemistry score. The sum F_{IBC} is mapped to $4\tanh(F_{\mathrm{IBC}}/4)$ for the panel losses and anchor scores. An empty edge set gives zero in every branch. The context features encode residue context and interface geometry; the typed features additionally combine residue chemical roles and uncertainty with radial distances, local-frame directions, and relative frame orientation.

The deployed target scalar is a frozen three-stage stack. A three-layer, six-head context encoder of width 192 supplies receptor–peptide pair features; a target-free typed module of pair width 192 adds explicit residue-role features. The typed module is trained on interface records from CPCore, CPSea-PDB, PepBench, and ProtFrag in component folds 3–5, then selected at step 2,000 by minimum fold-2 completed-sequence NLL among checkpoints that pass the exact-null, rigid-motion, and node-permutation checks. Binding labels enter only the final 76,613-parameter adapter (hidden width 192, dropout 0.15).

Development fixes its 250-step endpoint, after which the adapter is refit on folds 2–5 with AdamW (learning rate 2×10^{-4}, weight decay 0.02), batches of eight distinct components, and gradient clipping at 1.0. Every batch scores the complete 8×8 peptide–pocket panel: diagonal pairs are factual, and off-diagonal assignments are crossed counterfactuals. This refit holds the context encoder, typed module, ESM-2 features, and loss coefficients fixed; fold 1 is reserved for confirmation, and OOD components are absent from model and checkpoint selection.

Sequence corruption and reverse rates. The reference uses a symmetric 20-state corruption kernel. On the training clock $s \in [s_0, s_1)$, $s_0 = 0.05$, $s_1 = 0.95$, let $u = (s - s_0)/(s_1 - s_0)$ and $\eta = 10^{-3}$. For a clean residue a,

$$K_s(b \mid a) = \alpha_s \mathbf{1}\{b = a\} + (1 - \alpha_s)/20, \qquad \alpha_s = \eta^{1-u}, \qquad \rho = -\frac{\log \eta}{20(s_1 - s_0)}. \tag{29}$$

The clean-token posterior combines the logarithm of this corruption likelihood, frozen structural sequence-prior logits, and a learned residual, followed by softmax normalization over the 20 residues. Writing this posterior as $\pi_\theta(a \mid Y_s)$, the off-diagonal reference rates are

$$q^0_s(y, z) = \rho \sum_{a=1}^{20} \pi_\theta(a \mid Y_s) \frac{K_s(z \mid a)}{K_s(y \mid a)}, \qquad z \neq y. \tag{30}$$

The posterior-averaged density ratio is capped at 10^6 for numerical safety; self-transition rates are zero. Training targets replace the posterior with the observed clean residue. Rate matching averages the generalized Poisson KL term $r \log(r/q) - r + q$ over valid residue-rate entries, where r and q are target and predicted rates. The loss is normalized by the detached loss of the constant off-diagonal baseline ρ, with a 10^{-6} denominator floor. The auxiliary clean-token cross-entropy is divided by $\log 20$ and weighted by 0.25. The reference checkpoint uses unit geometry and rate weights.

Source states and constraint references. During reference learning, X_1 is the current data example. A local corruption uses a Gaussian scale sampled by randomized log-uniform strata from 0.12 to 4.0 Å before tangent projection and retraction around that example, and the retracted path in Eq. (23) supplies the velocity target. Inference draws empirical scaffolds from an H2T bank built from the training split of CPCore and CPSea-PDB. Each seed selects 64 donors of the requested length, uniformly without replacement over eligible donor components, retaining one bank representative per component and excluding the requested receptor component. A random cyclic recut changes the residue indexing. Gaussian perturbation with scale 4.0 Å before tangent projection and full-constraint

retraction gives the source coordinates; symmetric token corruption initializes the sequence. This uses training scaffolds while withholding the native peptide sequence and coordinates of the requested target complex.

The H2T constraint reference r stores the initial admitted source coordinates and is inherited by the lineage and its terminal children. Each residual component is half the difference between a squared interatomic distance and its value in that reference. The constraint graph includes backbone bonds, the H2T seam, and virtual distances fixing the modeled bond angles and local peptide-bond dihedral branch. These virtual distances serve as numerical constraints. Proper-rigid receptor placement leaves all reference distances unchanged.

Explicit H2T distance constraints. Indices below are cyclic, so $i+1=1$ when $i=L$. The full constraint graph contains the following edges for every residue:

$$\begin{aligned} E_{\text{bond}} &: (N_i, C\alpha_i),\ (C\alpha_i, C_i),\ (C_i, O_i),\ (C_i, N_{i+1}), \\ E_{\text{angle}} &: (N_i, C_i),\ (C\alpha_i, O_i),\ (C\alpha_i, N_{i+1}),\ (O_i, N_{i+1}),\ (C_i, C\alpha_{i+1}), \\ E_{\omega} &: (C\alpha_i, C\alpha_{i+1}). \end{aligned} \tag{31}$$

Thus there are $4L$ bond, $5L$ angle-distance, and L peptide-plane-distance rows. For edge $e=(a,b)$, the residual is $c_e(X)=\frac{1}{2}(\|X_a-X_b\|^2-\|X_a^r-X_b^r\|^2)$. The coordinate projector uses the Euclidean metric and the float64 rectangular Jacobian, with pseudoinverse relative cutoff 10^{-10}. The local internal dimension at a regular non-collinear state is $12L-6-\operatorname{rank} J_X$; the pseudoinverse accounts for dependent rows. These distances preserve the source's local covalent geometry while allowing conformational motion tangent to the active leaf.

D Experimental protocol

The following sections specify the data, output registration, structural evaluation, and controls used in Section 4.

D.1 Data splits and registered outputs

The eligible reference pool contains 39,128 structures drawn from 5,381 global receptor components among 107,918 training observations (Table 6). Folds 1–4 fit the reference, and fold 0 is held out for calibration. The H2T branch contributes 3,650 structures from 1,912 components over lengths 6–16. Length-stratified component incidences are not additive because one component may occur at several lengths.

Table 6: Reference-training separation. Counts are structures/components after the declared component-hash split. H2T length columns report structures (top) and component incidences (bottom) on fit folds 1–4.

Reference subset	Structures	Components	Sources / support
Typed total	39,128	5,381	linear + H2T
Linear	35,478	3,492	PepBench, ProtFrag
H2T	3,650	1,912	CPCore, CPSea-PDB; $L=6$–16

L	6	7	8	9	10	11	12	13	14	15	16
Structures	219	219	286	357	296	280	328	523	481	344	317
Component incidences	94	142	168	202	178	184	211	294	260	246	244

The reference uses declared H2T structure sources and component-level splits. A receptor component is the transitive union induced by per-chain MMseqs2 clusters at 40% identity and 80% bilateral coverage, exact chain identity, shared PDB identity, and the source-provided receptor group. These rules define the sequence-family split. PepBench and LNR entries come from experimentally observed linear complexes, which define the binding site and length request during preprocessing. All methods use that frozen site through their respective interfaces: pocket features for HYBRIDGE, a pocket PDB

and epitope for CP-Composer, and receptor coordinates with deterministic pocket-derived hotspots for RFpeptides. At generation time, native peptide sequences, peptide coordinates, and binding labels are withheld.

The held-out receptor set contains 114 components, 104 of them within the declared H2T length range. One length-supported component requires noncanonical receptor chemistry unavailable to all three methods. This leaves a common-input panel of 56 PepBench family components and 47 external LNR components. The ten out-of-length components and the full 104-component length-supported set appear in sensitivity analyses. We set component assignments and exclusions before adaptation, threshold selection, and candidate evaluation.

Held-out status is defined relative to HYBRIDGE training. CP-Composer and RFpeptides use the checkpoints released by their authors, without project-specific retraining. Each method produces 24 ordered outputs per component. Positions 1–5 were fixed before evaluation and enter the 11-pocket endpoint, yielding $103 \times 3 \times 5 = 1{,}545$ registered slots. Missing early outputs are not replaced by later ones.

Executable baselines and structural outcomes. CP-Composer contributes the receptor-relative complex from its released checkpoint and native score-blind H2T closure. RFpeptides contributes a receptor-relative cyclic backbone from 50-step RFdiffusion and one ProteinMPNN sequence per backbone at temperature 0.1. Canonical-L materialization preserves every submitted N/Cα/C/O backbone and pose while building the same all-heavy representation.

The anonymous artifact records the upstream repositories, implementation commits, released checkpoints, and sampling arguments in `docs/BASELINE_PROTOCOL.md` and `source_data/baseline_configuration.json`. Table 7 accounts for every first-five slot before endpoint statistics are computed.

Table 7: Mutually exclusive endpoint outcomes in the first five outputs. Each column sums to 515; later candidates do not replace an unavailable slot.

Outcome	HYBRIDGE	CP-Composer	RFpeptides
Endpoint evaluable	465	131	135
No supported product or method-native ring	49	61	253
Post-reload coordinate certificate failed	0	319	127
Canonical-L numerical failure	1	4	0

The post-reload certificate evaluates each submitted N/Cα/C/O backbone in its original receptor-relative pose. It rejects 319 CP-Composer and 127 RFpeptides backbones on full-backbone geometry. Among these, 49 and 14 also fail at the seam, and 2 and 13 fail the explicit closure bond. These failures remain after CP-Composer's native post-hoc H2T minimization or RFpeptides' native cyclic generation. A terminal C–N bond can coexist with an invalid surrounding backbone. HyBridge records 49 support-related typed non-emissions. All 466 emitted first-five backbones pass the independent full-backbone audit; 465 then reach canonical-L, with one numerical failure during materialization. Receptor contact and overlap are evaluated for canonical-L structures.

Backbone validity is audited independently of endpoint construction (Table 17). The four CP-Composer numerical failures have invalid backbone geometry in that audit; the single HyBridge numerical failure has valid backbone geometry. Their topology-valid counts are therefore 131 and 466, respectively, while their evaluable canonical-L counts are 131 and 465. Figure S13 follows this progression, and Figure S14 shows representative coordinate examples.

Certificate sensitivity. The analysis in Table 8 recomputes the certificate on the same 1,177 coordinate-bearing requests over a grid chosen in advance. Coordinates, poses, slot identities, and the explicit H2T edge remain untouched. Each grid point changes only the bond, angle, and peptide-bond torsion limits. Interface metrics come from the existing deterministic canonical-L target complex, so the underlying structure never changes.

Wider limits recover many CP-Composer backbones, although most still contain a sub-1.5 Å pair; newly accepted RFpeptides products are usually non-contacting. The ordering is unchanged in both

Table 8: Certificate and interface sensitivity on unchanged coordinates. Each cell reports canonical-L-evaluable requests accepted by the stated geometry rule / requests that also contact the intended pocket without a sub-1.5 Å pair (percentage of 515). Moderate limits are 0.05 Å, 5°, and 40°; permissive limits are 0.10 Å, 10°, and 50°. The main benchmark uses the prespecified rule.

Geometry rule	HYBRIDGE	CP-Composer	RFpeptides
Frozen	465 / 315 (61.2%)	131 / 26 (5.0%)	135 / 18 (3.5%)
Bond margin 0.05 Å	465 / 315 (61.2%)	131 / 26 (5.0%)	165 / 19 (3.7%)
Angle margin 5°	465 / 315 (61.2%)	169 / 38 (7.4%)	136 / 18 (3.5%)
Torsion limit 40°	465 / 315 (61.2%)	206 / 45 (8.7%)	135 / 18 (3.5%)
Moderate joint	465 / 315 (61.2%)	271 / 64 (12.4%)	174 / 19 (3.7%)
Permissive joint	465 / 315 (61.2%)	339 / 79 (15.3%)	201 / 22 (4.3%)

OOD tracks at every boundary. Under the permissive stress test, the pooled HYBRIDGE margin is 45.8 percentage points over CP-Composer and 56.9 over RFpeptides.

The release includes identities for all 1,545 slots, coordinates for the 731 evaluable complexes, and submitted backbones with residuals for 446 failed certificates. The remaining 363 non-emissions and five numerical failures carry status records because no coordinate structure exists. Public wrappers reproduce baseline generation, closure, and the shared CPU endpoint.

For a candidate Y and 11 pockets $P_\star, P_1, \ldots, P_{10}$, the evaluator orders the intended pocket with ten component-distinct controls chosen from receptor-only information before generated outputs were inspected. A canonical frame built from pocket Cα coordinates and residue chemistry maps each control rigidly into the intended frame; peptide coordinates do not determine that map. Sequence, N/Cα/C/O backbone, receptor-relative pose, topology edges, candidate identity, and denominator remain unchanged. The evaluator then adds canonical L-Cβ, completes side chains and hydrogens at pH 7, fixes N/Cα/C/O/Cβ, minimizes only the added atoms, and reloads the complex.

Here, one indexed output is a method-native export after inference. Each procedure contributes 24 such outputs; materialization and evaluation are then shared. Table 11 records how each set is produced.

Fixed-pose receptor controls. For a candidate emitted against $P_\star$, the receptor-relative coordinates are part of the model output $Y_\star = (S, X_\star)$. The benchmark compares

$$\varepsilon(Y_\star, P_\star) \quad \text{with} \quad \varepsilon(Y_\star, T_j P_j), \tag{32}$$

where T_j is determined from receptor-only pocket frames. Holding Y fixed isolates receptor identity and records how the emitted pose orders $P_\star$ relative to the controls. Crossed and pocket-chemistry interventions provide the complementary learned interaction-response controls.

Let $U_{36}^{\mathrm{NC}}(Z; x)$ denote the CHARMM36 potential of atom set Z at coordinates x, using OpenMM `NoCutoff`, no explicit or implicit solvent, and the same hydrogen-bond constraints and particle parameters for all three systems. No coordinates are relaxed during the following single-point calculation. For each pocket,

$$\varepsilon(Y, P; x) = U_{36}^{\mathrm{NC}}(Y \cup P; x) - U_{36}^{\mathrm{NC}}(Y; x_Y) - U_{36}^{\mathrm{NC}}(P; x_P). \tag{33}$$

The peptide-only and receptor-only systems are exact subsets of the complex at the same coordinates, so Eq. (33) returns the fixed-coordinate cross-system force-field term. We use this term to order the intended pocket among the ten controls. With lower values preferred,

$$r(Y) = 1 + \sum_{j=1}^{10} \mathbf{1}\{\varepsilon(Y, P_j) \le \varepsilon(Y, P_\star)\}. \tag{34}$$

Ties are pessimistic. Rank@k is the fraction of the 1,545 registered first-five endpoints whose intended-pocket rank is at most k. With $R = 11$ pockets and unavailable slots assigned rank $R + 1$, normalized registered-denominator rank-AUC for $r \in \{1, \ldots, R + 1\}$ is $(R + 1 - r)/R = \frac{1}{R} \sum_{k=1}^{R} \mathbf{1}\{r \le k\}$. Thus the quantity called rank-AUC in the source data is the mean retrieval

CDF across all 11 cutoffs. An evaluated rank-11 candidate contributes $1/11$, whereas an unavailable rank-12 slot contributes zero.

The registered energy matrix has 1,545 candidates times 11 pockets, or 16,995 cells; ineligible cells remain NA and unavailable candidates retain worst rank. Confidence intervals are 10,000 receptor-component-blocked bootstrap resamples with base seed 20260817.

Rigid transport produces a strongly heavy-tailed energy scale. Across the 731 evaluable products, the intended-pocket median is 2.43×10^3 kcal/mol and 87.0% of values are positive; across 7,310 transported controls, the median is 1.33×10^9 kcal/mol and 74.6% exceed 10^6 kcal/mol. We therefore report ordinal rank together with contact and severe overlap. Each control contributes one comparison, and r is invariant under any strictly increasing transformation of the eleven energies. The resulting rank measures how well an emitted fixed pose favors its construction pocket over aligned receptor controls. The crossed-response and pocket-chemistry panels then examine the learned target signal itself.

D.2 Path controls and receptor response

Hard and soft path controls. The frozen fold-2 panel contains 20 receptors and 24 paths per receptor, giving 480 paths per condition. A paired hard replay and three soft controls start from the same admitted H2T coordinates and tokens. They use 50 Euler intervals on normalized time $[0, 1]$, with constraint metrics recorded after every interval. Coordinate Gaussian draws are shared, and the soft arms replay the hard arm's CTMC sequence states. Hard updates project and retract on the full H2T graph. Soft updates retain the full open-chain constraints and replace the terminal-link constraints with a harmonic seam velocity: the terminal C and initial N receive equal and opposite velocities along their connecting line, with magnitude λ times the absolute departure from the initial seam length, directed toward that length. The dimensionless strengths $\lambda \in \{1, 10, 100\}$ are fixed before execution.

Path validity requires the maximum relative distance error over every real and virtual H2T constraint edge to stay within 10^{-7} at all recorded steps; the same paths are also checked at 10^{-4} and 10^{-2}. All hard paths pass, and every soft path crosses each tolerance. Endpoint closure uses a separate terminal C–N distance threshold of 1.8 Å. Figure 4a reports that endpoint test on the frozen 24-path sentinel control. Its closure counts are 24, 8, 24, and 24 for hard support and the three increasing penalty strengths. This pairs a full-panel path audit with a concrete example of endpoint recovery.

Crossed receptor response. For a reciprocal factual/crossed pair, let $s(Y, P)$ denote the frozen target readout and let Y_F, Y_C be the factual and counterfactual output populations for receptors P_F, P_C. We compute

$$D_F = \frac{1}{|Y_F|} \sum_{Y \in Y_F} [s(Y, P_F) - s(Y, P_C)], \tag{35}$$

$$D_C = \frac{1}{|Y_C|} \sum_{Y \in Y_C} [s(Y, P_F) - s(Y, P_C)], \tag{36}$$

$$\Delta_\times = D_F - D_C. \tag{37}$$

This is the empirical population form of the rectangular contrast in Eq. (2); adding any peptide-only term $a(Y)$ or receptor-only term $b(P)$ leaves $\Delta_\times$ unchanged. Figure 4e compares the guided-suffix arm with terminal-only scoring and the reference-only geometry/support arm. The latter uses receptor-free molecular actions while retaining the common score-blind receptor-relative placement. The 11 reciprocal pairs are the statistical units.

These response diagnostics use the learned readout and generated geometry; U_{36} does not enter. On the held-out crossed panel, the readout gives factual-pair preference 0.859 [0.815, 0.900], rectangular-quartet accuracy 0.928 [0.882, 0.968], and an exact-null residual of zero.

Applying the same target correction during the guided suffix increases the frozen-readout $\Delta_\times$ diagnostic by 0.069 [0.029, 0.110] relative to terminal-only scoring and by 0.093 [0.045, 0.147] relative to the reference-only control. In the same-checkpoint removal, the complete readout also improves force-field-free Recall@1 by 0.187 [0.097, 0.281] and MRR by 0.140 [0.075, 0.207] over its

context-only component (Table 9). The pocket-chemistry experiment below gives a second fixed-state check of the learned target signal.

The geometry response compares matched factual and counterfactual output populations. Candidates are paired by common ancestry when the fixed seed yields a bijection; otherwise a minimum-cost bijection uses cyclic sequence edit distance. For each pair, the counterfactual ring is cyclically recut at the shift with the fewest token edits, with proper-Kabsch Cα RMSD as the deterministic tie-breaker. The plotted value is the proper-Kabsch Cα RMSD after that selected recut. This captures backbone change across matched output populations; the fixed-pose screen above reports receptor ordering.

D.3 Calibrating the fixed-coordinate screen

The ordering direction is calibrated on 47 held-out LNR co-crystal complexes with linear peptide ligands. Their native coordinates are scored against the cognate receptor and the same ten aligned, component-distinct controls used by the OOD endpoint; no generator or learned checkpoint is involved. The cognate receptor ranks first in all 47 cases. This uniform outcome makes the component bootstrap a point mass at one. An IID Bernoulli Wilson interval gives [0.924, 1.000] as a familiar finite-sample reference. The native poses establish the direction of the fixed-coordinate screen used for the generated cyclic peptides.

D.4 Removing non-context readout branches

The held-out fold-1 panel is replayed with the same pair tensors, candidates, geometry, base scores, and checkpoint. The complete readout is $F_{\mathrm{IBC}} = F_{\mathrm{ctx}} + F_{\mathrm{chem}} + F_{\mathrm{syn}} + F_{\mathrm{prior}}$, followed by the common bound $B\tanh(F_{\mathrm{IBC}}/B)$. We compare it with $B\tanh(F_{\mathrm{ctx}}/B)$, which removes the remaining three components. Because F_{ctx} is computed from receptor–peptide cross-edges, the context-only arm remains receptor-conditioned. Their difference isolates the contribution of the other three components at the same checkpoint. Table 9 reports component-balanced estimates from 10,000 paired component bootstrap resamples (seed 20260816). The complete arm exactly reproduces the registered confirmation metrics before comparison.

Table 9: Frozen complete-readout versus context-only removal. Positive Δ favors the complete readout; intervals are paired 95% confidence intervals.

Metric	Complete	Context only	Δ	95% CI
Interaction preference	0.859	0.767	0.092	[0.047, 0.140]
Recall@1	0.494	0.307	0.187	[0.097, 0.281]
MRR	0.649	0.509	0.140	[0.075, 0.207]
Target–best-decoy margin	-0.030	-0.075	0.045	[0.005, 0.087]
Quartet accuracy	0.928	0.920	0.008	[-0.021, 0.036]

The complete readout improves pairwise preference, Recall@1, MRR, and the target–best-decoy margin at this checkpoint. Quartet accuracy changes little from its already high value, and its paired interval crosses zero. Both arms retain an exact-null maximum error of zero. Thus, the added $F_{\mathrm{chem}} + F_{\mathrm{syn}} + F_{\mathrm{prior}}$ components improve individual-pair ordering while preserving the receptor contrast shared by both arms.

D.5 Held-out spatial-chemistry recognition

Three target adapters start from the same initialization and fit receptor components in folds 2–5: retrieval only, the complete objective without $\mathcal{L}_{\mathrm{cf}}^{\mathrm{pocket}}$, and the complete objective. Evaluation uses five fresh fit seeds and 60 held-out fold-1 components. For each peptide–pocket state, eight deterministic Sattolo derangements jointly permute the pocket amino-acid tokens and ESM rows over active residues while preserving pocket coordinates, the peptide state, and the pocket amino-acid multiset. The factual pocket is ranked among these geometry-fixed chemical counterfactuals with pessimistic ties. We average observations within component–length items and then across fit seeds and lengths within each receptor component; 95% intervals use 10,000 component-bootstrap replicates.

Pocket-counterfactual training improves recognition of the factual chemical arrangement. The complete objective reaches pairwise preference 0.672 [0.593, 0.745], compared with 0.611 [0.532, 0.689] without the pocket counterfactual and 0.600 [0.519, 0.679] for retrieval only. Relative to removal, the paired gains are 0.061 [0.041, 0.081] for pairwise preference and 0.046 [0.021, 0.071] for MRR; Rank@1 changes by 0.030 [-0.009, 0.069]. Rank@1 requires the factual pocket to beat all eight counterfactuals, whereas the fitted softplus objective sums factual-versus-counterfactual orderings. Pairwise preference and MRR are therefore the direct objective-aligned readouts. Ordinary matched-pair preference is 0.856 for the complete objective and 0.861 after removal, placing the gain specifically in spatial-chemistry recognition.

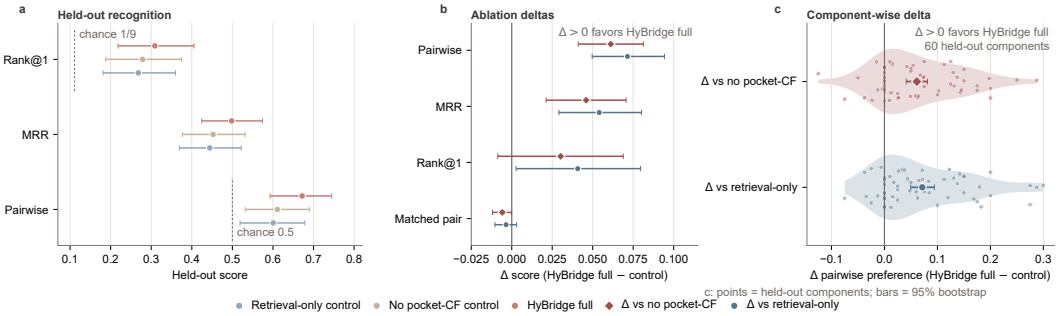


Figure S1: **Pocket-counterfactual supervision improves held-out spatial-chemistry ordering. a,** Factual Rank@1, MRR, and mean preference over eight geometry-fixed, composition-preserving pocket- chemistry shuffles on 60 held-out receptor components. **b,** Paired HyBridge-full-minus-control deltas. **c,** Component-wise deltas in the objective-aligned pairwise preference; markers and bars show the mean and 95% receptor-component bootstrap interval. Each arm uses five paired fit seeds and the same fixed molecular states.

Figure S2 complements this fixed-state recognition task with generated outputs: panel **a** measures reciprocal target response, and panel **b** follows matched changes in cyclic backbone geometry.

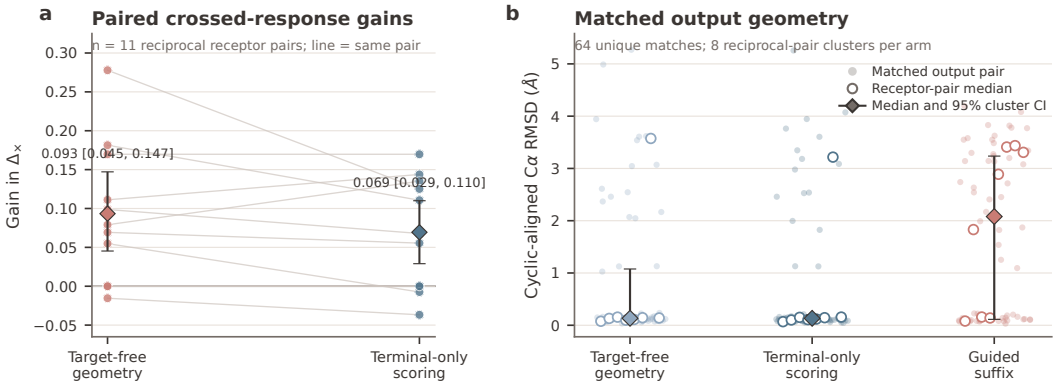


Figure S2: **Mechanism-level response with explicit pairing and uncertainty. a,** Guided-suffix gains in the gauge-invariant response $\Delta_\times$ relative to target-free geometry and terminal-only scoring. Each line joins the same one of 11 reciprocal receptor pairs; diamonds show means and bars the registered 95% pair-bootstrap intervals. **b,** Proper-Kabsch Cα RMSD after deterministic cyclic recutting of matched factual/counterfactual outputs. Faint points are 64 unique matched output pairs per arm after collapsing mirrored reciprocal directions; open circles are medians within eight reciprocal receptor-pair clusters. Diamonds show overall medians and bars descriptive 95% cluster-bootstrap intervals. Matching uses common-seed bijective ancestry when available and otherwise a minimum-cost bijection based on cyclic sequence edit distance.

Table 10: Complete-procedure comparison on 103 held-out receptor components. Counts use all 515 registered slots per method. Joint delivery requires canonical-L materialization, receptor contact within 4 Å, and no interchain heavy-atom pair below 1.5 Å. Percentages are per registered slot; native procedures are not matched for compute.

Method	Emitted	Canonical-L	Joint delivery
CP-Composer	454	131	26 (5.0%)
RFpeptides	262	135	18 (3.5%)
HYBRIDGE	466	465	**315 (61.2%)**

D.6 FROM EMITTED BACKBONES TO RECEPTOR INTERFACES

Each output passes through three structural stages (Table 12). Backbone availability requires a sequence with N/CA/C/O coordinates. Ring certification then checks the explicit H2T edge, closure seam, and full covalent backbone after reload. Canonical-L evaluation adds deterministic side chains and hydrogens with the declared stereochemistry. This progression separates missing output, invalid backbone geometry, and failure to construct the all-heavy representation.

The frozen backbone predicate accepts N–Cα, Cα–C, C–O, and successive C–N lengths in $[1.20, 1.70]$, $[1.30, 1.80]$, $[1.00, 1.50]$, and $[1.10, 1.70]$ Å, respectively. N–Cα–C, Cα–C–N, and C–N–Cα angles must lie in $[95, 125]^\circ$, $[105, 135]^\circ$, and $[105, 135]^\circ$, and each peptide ω must fall within 30° of its nearest cis or trans reference. The same rules apply across the serialized H2T seam; canonical-L certification also requires the declared chirality and an unchanged reloaded backbone.

Table 11 compares how each method produces the 24 candidate positions fixed before evaluation.

Table 11: Method-specific construction of 24 indexed outputs for one component. Each row describes the frozen inference routine that produces those outputs.

Method	Internal candidate construction	Indexed outputs
HYBRIDGE	Per seed: one shared 64-state source-preview pool forms an eight-state anchor bank and, in a separate support pass, initializes eight particle paths. Each path branches into eight continuations only over $t = .76 \to 1$, and normalized terminal-weight sampling returns one output	Three seeds, eight outputs each
CP-Composer	Released checkpoint and sampler on 24 fixed axes, followed by its method-native score-blind H2T closure	24 outputs
RFpeptides	24 RFdiffusion backbones (50 steps), with one ProteinMPNN sequence per backbone at temperature 0.1	24 outputs

Table 12: Stage-by-stage accounting for the common first-five denominator. Backbone is a generated sequence plus N/Cα/C/O coordinates; Ring adds the H2T graph and coordinate certificate; Canonical-L is the common all-heavy evaluation representation. Joint additionally requires direct 4 Å contact and no sub-1.5 Å inter-chain pair.

Track	Method	Registered	Backbone	Ring	Canonical-L	Joint
PepBench	HYBRIDGE	280	249	249	248	161
	CP-Composer	280	248	77	77	17
	RFpeptides	280	146	77	77	11
LNR	HYBRIDGE	235	217	217	217	154
	CP-Composer	235	206	54	54	9
	RFpeptides	235	116	58	58	7

The fixed-pose interface screen uses all non-hydrogen receptor and peptide atoms from the exact canonical-L PDB bytes used by the energy endpoint. It reports the minimum inter-chain heavy-atom distance and the fraction of peptide residues with a receptor heavy atom within 4 Å. The 1.5 Å cutoff

is a method-neutral severe-overlap sentinel; contact coverage and buried area provide complementary descriptions of the generated pose. Table 13 reports the corresponding fixed-pose receptor ordering.

Table 13: Generated-complex validity and a secondary rigid-transport fixed-pose compatibility screen. Entries are point estimate [95% interval] from 10,000 receptor-component bootstrap draws. Complete-denominator metrics retain unavailable candidates at rank 12; conditional AUC uses only certified canonical-L evaluable candidates.

PepBench family OOD (56 components)	HYBRIDGE	CP-Composer	RFpeptides
Rank@1	0.614 [0.532, 0.693]	0.071 [0.043, 0.104]	0.050 [0.021, 0.082]
Rank@2	0.800 [0.729, 0.864]	0.150 [0.111, 0.189]	0.068 [0.036, 0.107]
Rank@5	0.882 [0.818, 0.939]	0.232 [0.179, 0.286]	0.161 [0.111, 0.214]
AUC_{reg}	0.850 [0.789, 0.904]	0.221 [0.175, 0.271]	0.162 [0.120, 0.208]
AUC_{eval}	0.960 [0.948, 0.971]	0.805 [0.755, 0.854]	0.590 [0.498, 0.678]
Evaluable/registered	248/280	77/280	77/280
LNR external OOD (47 components)	**HYBRIDGE**	**CP-Composer**	**RFpeptides**
Rank@1	0.664 [0.579, 0.745]	0.094 [0.055, 0.136]	0.051 [0.021, 0.089]
Rank@2	0.821 [0.745, 0.889]	0.123 [0.081, 0.166]	0.060 [0.026, 0.098]
Rank@5	0.923 [0.855, 0.979]	0.191 [0.140, 0.243]	0.111 [0.072, 0.153]
AUC_{reg}	0.887 [0.823, 0.940]	0.183 [0.141, 0.228]	0.146 [0.110, 0.184]
AUC_{eval}	0.961 [0.946, 0.973]	0.798 [0.717, 0.875]	0.591 [0.511, 0.671]
Evaluable/registered	217/235	54/235	58/235

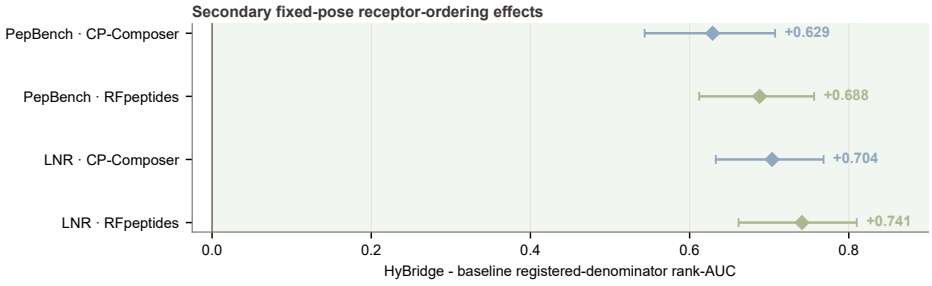


Figure S3: **Secondary rigid-transport compatibility-ordering effects.** Paired HYBRIDGE-minus-baseline differences in registered-denominator rank-AUC on PepBench and LNR. Markers show mean paired effects; error bars are 95% receptor-component-blocked bootstrap intervals. The ordinal endpoint uses the unchanged fixed poses and complete first-five denominator from Figure 3.

D.7 SENSITIVITY TO ALL-HEAVY MATERIALIZATION

Table 14 compares two views of the same fixed pose. The backbone-only distance is the minimum between receptor heavy atoms and peptide N/Cα/C/O atoms; the canonical-L distance uses every non-H atom in the receptor and evaluator-side peptide representation. Hydrogen atoms are excluded from both. Sequence, backbone, topology, receptor-relative pose, candidate identity, and registered denominator remain the same. The comparison therefore shows how side-chain completion changes the resolution of the interface diagnostic.

For context, MolProbity adds and optimizes hydrogens and defines a serious non-donor–acceptor clash through atom-specific van der Waals overlap greater than 0.4 Å (Chen et al., 2010). Among endpoints flagged by our simpler sentinel, the inter-chain pair-count median is 1 for HYBRIDGE and 3 for CP-Composer; across all certified endpoints the corresponding 90th percentiles are 1 and 12, and the maxima are 5 and 33. RFpeptides has only 22 flagged endpoints because 95/135 evaluable products are non-contacting; among the flagged cases, the median pair count is 19.5. Packing is therefore read from contact and overlap together, rather than from the pooled overlap rate alone.

Table 14: Backbone-only and canonical-L all-heavy interface sensitivity. Distance and contact-coverage entries are medians; the final column is the fraction with at least one sub-1.5 Å pair. All entries are computed only for certified canonical-L evaluation endpoints; unavailable registered slots remain in the retrieval denominator but are not imputed here.

Scope	Method	n	Median backbone $d_{\min}$ (Å)	Median all-heavy $d_{\min}$ (Å)	Median contact coverage	Fraction any < 1.5 Å
Pooled	HYBRIDGE	465	2.186	1.752	0.833	0.323
	CP-Composer	131	1.163	0.985	0.833	0.802
	RFpeptides	135	22.069	19.754	0.000	0.163
PepBench	HYBRIDGE	248	2.179	1.702	0.833	0.351
	CP-Composer	77	1.280	1.041	0.818	0.779
	RFpeptides	77	35.483	32.727	0.000	0.182
LNR	HYBRIDGE	217	2.194	1.849	0.833	0.290
	CP-Composer	54	0.990	0.895	0.909	0.833
	RFpeptides	58	17.424	15.165	0.000	0.138

D.8 COMPONENTS, TRAINING PROXIMITY, AND TOPOLOGY

The aggregate compatibility results vary across receptor components (Figures S3 and S4). We examine this variation alongside nearest-training sequence identity and backbone distance (Figure S5). The pocket-chemistry and matched-rollout controls in Figures S1 and S2 address receptor response, while the topology examples below extend the structural analysis beyond H2T. Source tables for Figure 3 give the track-specific interface summaries.

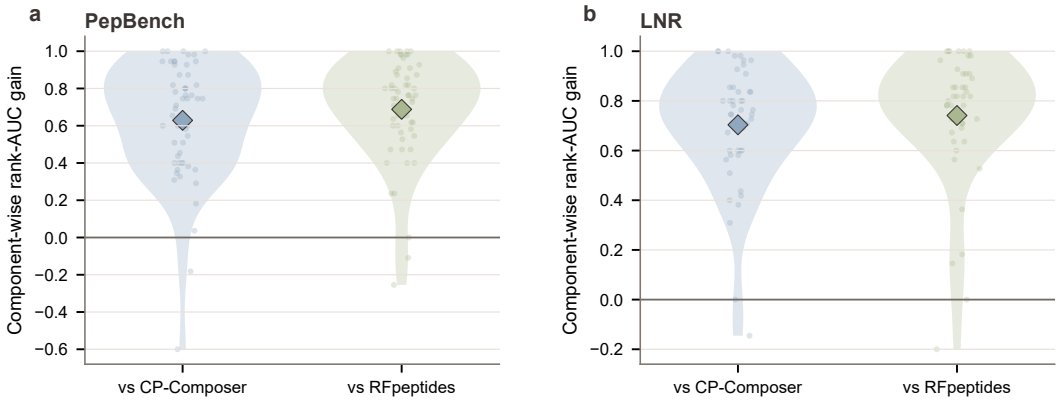


Figure S4: **Secondary rigid-transport compatibility-ordering effects.** Component-wise registered-denominator rank-AUC differences from CP-Composer and RFpeptides on both tracks; zero denotes no difference.

The component-level distributions show the spread of compatibility-ordering effects; source proximity is summarized below.

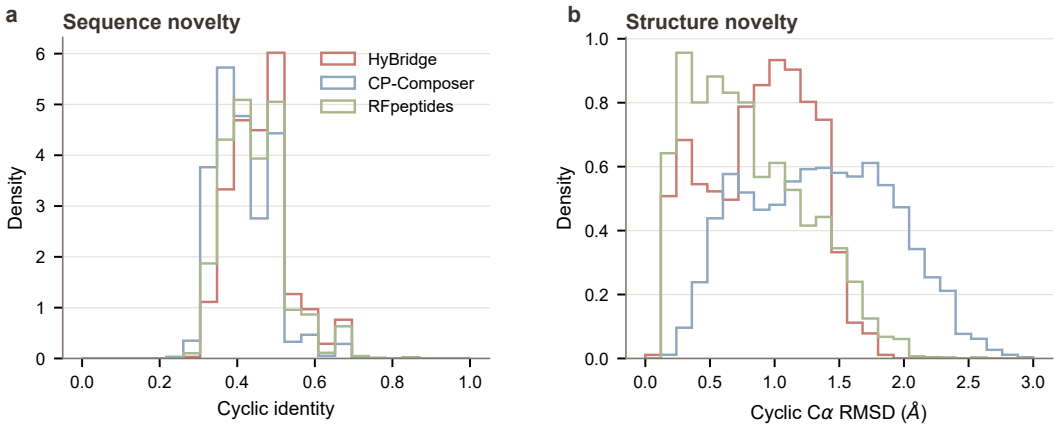


Figure S5: **Source proximity.** Nearest-training cyclic sequence identity and cyclic-recut-invariant Cα RMSD. The analysis identifies exact copying; distances beyond that are descriptive.

Topology execution is shown directly in Figure S6, which records coordinate-level products from six compiled covalent programs.

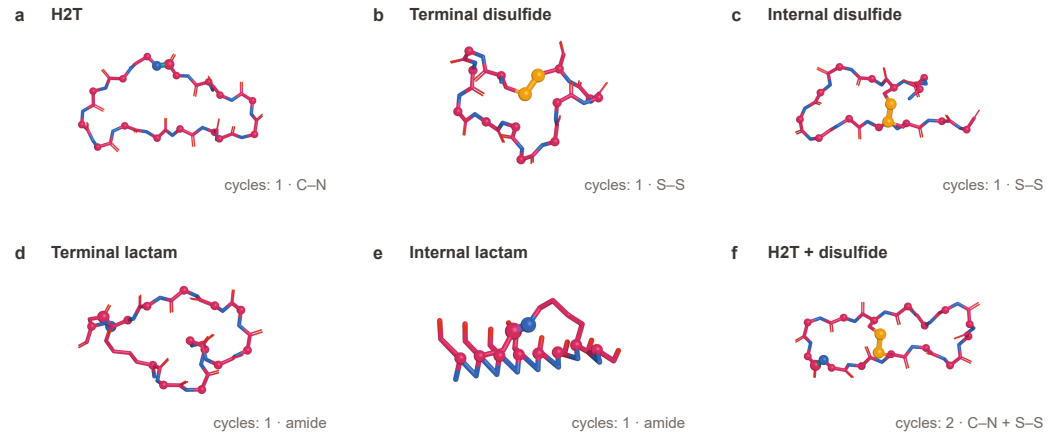


Figure S6: **Compiled topology atlas.** Six target-free coordinate products from the shared compiler. Panel **d** is an independently executed eight-step terminal-lactam path: every step passed the topology certificate, the final amide closure was 1.330 Å, and the maximum tangent and retraction residuals were 2.03×10^{-12} and 3.22×10^{-14}, respectively. The lactam examples demonstrate compiled graphs and coordinate materialization.

D.9 Short-time persistence of receptor-relative topology examples

To complement the static topology certificates, we ran an exploratory 1 ns explicit-solvent check of local geometric persistence at 300 K for the three disulfide topologies in Figure 6a–c. Figure S7 provides the full receptor views and interface metrics. Terminal and internal disulfides each use three fixed trajectories; the cycle-rank-2 H2T-plus-disulfide case uses one frozen trajectory and is descriptive. All seven trajectories retained the declared closure edges, showed no canonical-L chirality inversion, and remained within the 4.5 Å contact boundary at every stored frame (Table 15 and Figure 6). The different Cα RMSDs show that covalent persistence does not imply a rigid backbone.

The trajectories preserve closure, chirality, and receptor contact while allowing backbone motion over 1 ns. The topology-dependent RMSDs make that distinction visible in the same receptor-relative structures.

Table 15: Exploratory 1 ns explicit-solvent persistence. Contact counts and Cα RMSDs are frame averages; closure is the largest declared-bond distance observed. The first two rows summarize three trajectories each; the bicycle row is one trajectory.

Topology	Traj.	Closure (Å)	Inversions	CA RMSD (Å)	Contacts	Contact frames (%)
Internal S–S	3	2.111	0	3.063	80.1	100%
Terminal S–S	3	2.150	0	0.876	133.3	100%
H2T + S–S	1	1.429 / 2.139	0	1.514	79.8	100%

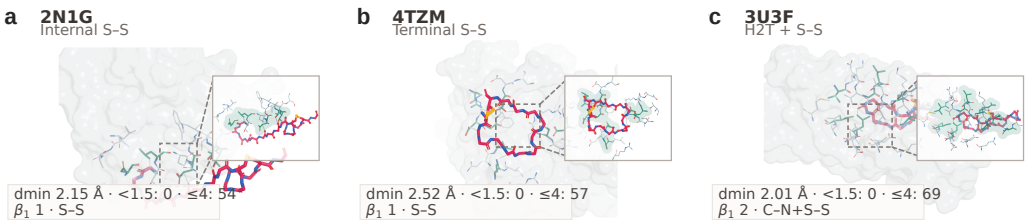


Figure S7: **Receptor-relative complexes for additional covalent programs. a–c,** Internal S–S on 2N1G, terminal S–S on 4TZM, and H2T+S–S on 3U3F, corresponding to Figure 6a–c. Pink traces show the peptide, gray surfaces the receptor, and green traces receptor contacts within 4 Å. Cards report minimum interchain distance, close-pair counts, and graph cycle rank, distinguishing single-ring and bicyclic products.

Table 16: Evidence available for each compiled topology.
Here $\beta_1 = |E| - |V| + c$ is the graph cycle rank, with c connected components.

Topology	β_1	Evidence	Demonstrated here
H2T	1	learned reference; receptor-family OOD	controlled crossed response; registered H2T output evaluation
Terminal disulfide	1	receptor-only pocket pilot; three 1-ns trajectories	topology-valid receptor-relative product
Internal disulfide	1	fixed-potential selection; three 1-ns trajectories	topology-valid receptor-relative product
Terminal lactam	1	compiled graph and coordinates	explicit closure product
Internal lactam	1	compiled graph and coordinates	explicit closure product
H2T + disulfide	2	fixed-potential edits; one 1-ns trajectory	topology-valid cycle-rank-2 product

E Robustness, Quality, and Failure Modes

E.1 Endpoint reach in the first five outputs

The frozen first-five order shows how quickly each procedure produces an evaluable structure for a receptor component. Evaluable@k records whether the first k outputs contain a common endpoint. Delivery@k further requires contact within 4 Å and no interchain heavy-atom pair below 1.5 Å. Unavailable outputs remain in their registered positions.

Across the five-slot prefix, HyBridge reaches a common endpoint for 100 of 103 components and meets the contact and exclusion criteria for 96. The corresponding delivery counts are 21 for CP-Composer and 11 for RFpeptides (Fig. S8).

E.2 Topology recovery across the 24-output registration

Each procedure has 24 ordered outputs. The full order shows how additional draws recover valid cyclic backbones. A counted output is emitted, passes the method-native ring check, and satisfies the common full-backbone geometry rule. Receptor-relative contact and overlap remain evaluated on the first-five panel. The horizontal axis follows registered output order and does not normalize compute.

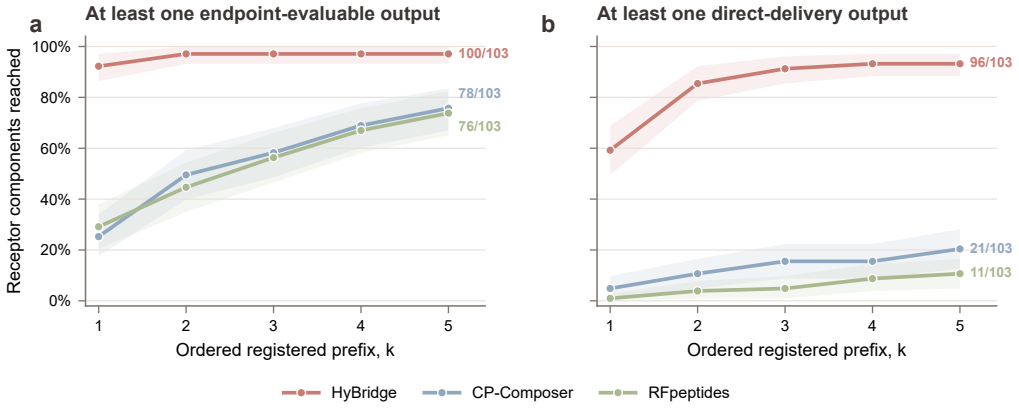


Figure S8: **Endpoint reach and direct delivery along the ordered first-five prefix. a,** Fraction of 103 receptor components with at least one common endpoint-evaluable output among the first k slots. **b,** Fraction with at least one endpoint that also contacts the intended pocket and contains no sub-1.5 Å inter-chain pair. Ribbons are percentile 95% intervals from 10,000 receptor-component bootstrap resamples (seed 20260820), with outputs nested within receptor component.

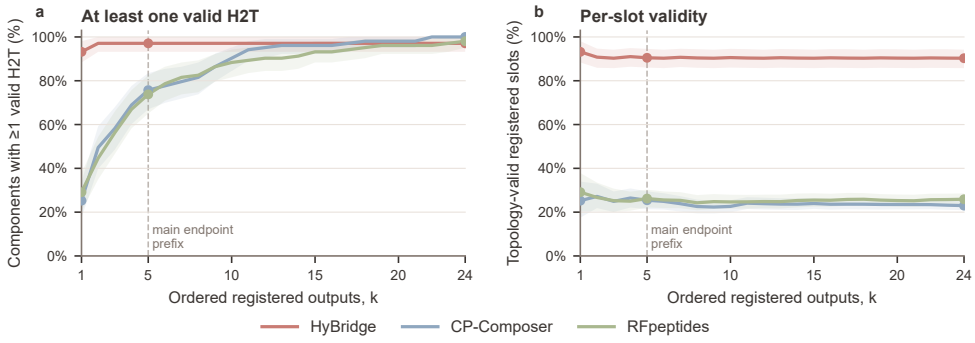


Figure S9: **Topology validity across the ordered 24-output registration. a,** Fraction of 103 receptor components with at least one topology-valid H2T backbone among the first k registered outputs. **b,** Complete-denominator fraction of topology-valid slots within the same prefix. Curves follow the frozen output order; bands are 95% receptor-component bootstrap intervals (10,000 resamples; seed 20260820). The dashed line marks the five-slot prefix used by the common interface endpoint.

Table 17: Receptor components with at least one topology-valid H2T backbone and per-slot topology validity at the five- and 24-output prefixes.

Method	≥ 1 valid H2T @5	≥ 1 valid H2T @24	Valid slots @5	Valid slots @24
HYBRIDGE	100/103	100/103	466/515	2,232/2,472
CP-Composer	78/103	103/103	131/515	568/2,472
RFpeptides	76/103	101/103	135/515	640/2,472

Both baselines reach more receptor components as outputs accumulate. Within five outputs, at least one topology-valid H2T backbone is available for 100/103 HYBRIDGE components, 78/103 CP-Composer components, and 76/103 RFpeptides components. At 24 outputs, the corresponding counts become 100/103, 103/103, and 101/103. Across all 24 outputs, the topology-valid fractions are 90.3%, 23.0%, and 25.9%. The larger output pool therefore broadens baseline coverage, while the per-slot fractions retain a substantial difference in topology reliability.

E.3 GUIDANCE SCHEDULES THROUGH STRUCTURAL OUTCOMES

The matched schedule comparison evaluates structural delivery on the same 103 receptor components under three schedules. All three share the checkpoint, H2T support, receptor-aware score-blind placement, seeds, and five output slots per component. Full guidance steers the entire suffix. Terminal-only guidance enters at the final choice, while the no-scalar control still sees receptor geometry without the learned receptor–peptide preference.

Delivery requires a canonical-L structure that contacts the intended receptor within 4 Å and has no interchain heavy-atom pair below 1.5 Å. Missing structures count as failures, so all 515 registered slots remain in the denominator. Confidence intervals come from 10,000 bootstrap resamples of the 103 receptor components (seed 20260820).

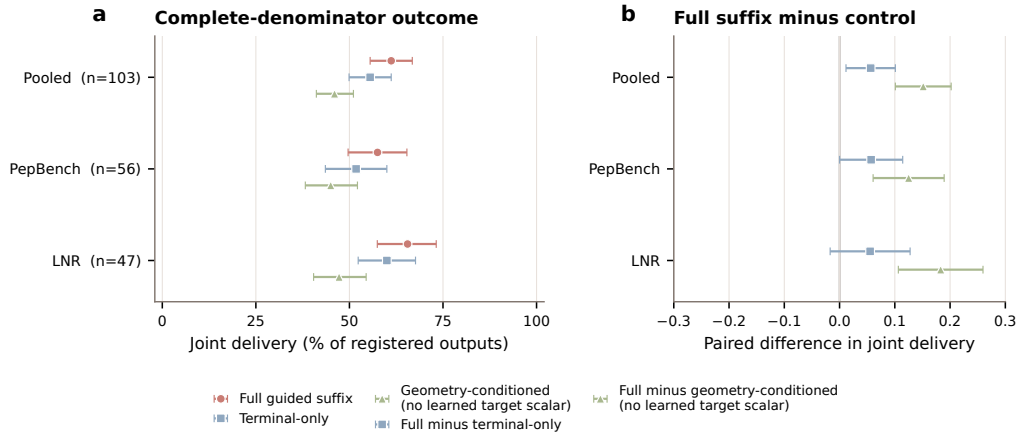


Figure S10: **Structural outcomes under three guidance schedules. a,** Joint-delivery estimates for the pooled panel and each receptor-family track. **b,** Paired full-guided-minus-control differences on the same registered component blocks. Points are estimates; lines are percentile 95% component-bootstrap intervals.

Full-suffix guidance gives the highest joint delivery across all 515 slots. The estimates are 61.2% (95% CI 55.5–66.8) for the full guided suffix, 55.5% (95% CI 49.9–61.2) for terminal-only, and 46.0% (95% CI 41.2–51.1) for geometry-conditioned execution without the learned target scalar. Relative to terminal-only and the geometry-conditioned control, the paired differences are 5.6 points (95% CI 1.2–10.1) and 15.1 points (95% CI 10.1–20.2), respectively. Figure S10 shows the same ordering in both source tracks. The main text and contact-conditioned overlap analysis use the same full-suffix estimate and interval. Each schedule contrast uses matched component resamples.

Table 18: **Where the three schedules succeed or fail.** Percentages and percentile 95% component-bootstrap intervals are shown for each stage. H2T support records a supported emission at the generator boundary; Canonical-L additionally requires the common post-reload construction. Overlap denotes an interchain heavy-atom pair below 1.5 Å; joint delivery combines canonical-L availability, 4 Å contact, and no such pair.

Path support and materialization

Pooled

Procedure	Emission	H2T support	Canonical-L
Full suffix	90.5 [86.0, 94.4]	90.5 [86.0, 94.4]	90.3 [85.8, 94.2]
Terminal-only	86.8 [82.7, 90.7]	86.8 [82.7, 90.7]	86.4 [82.3, 90.5]
No learned scalar	87.6 [83.5, 91.5]	87.6 [83.5, 91.5]	85.8 [81.6, 89.9]

PepBench

Procedure	Emission	H2T support	Canonical-L
Full suffix	88.9 [82.5, 94.3]	88.9 [82.5, 94.3]	88.6 [82.1, 93.9]
Terminal-only	86.1 [80.0, 91.4]	86.1 [80.0, 91.4]	85.7 [79.3, 91.1]
No learned scalar	86.8 [80.7, 92.1]	86.8 [80.7, 92.1]	86.8 [80.7, 92.1]

LNR

Procedure	Emission	H2T support	Canonical-L
Full suffix	92.3 [85.5, 97.9]	92.3 [85.5, 97.9]	92.3 [85.5, 97.9]
Terminal-only	87.7 [82.1, 92.8]	87.7 [82.1, 92.8]	87.2 [81.3, 92.8]
No learned scalar	88.5 [83.0, 93.6]	88.5 [83.0, 93.6]	84.7 [78.7, 90.2]

Interface outcome

Pooled

Procedure	Contact	Overlap	Joint
Full suffix	90.3 [85.8, 94.2]	29.1 [24.7, 33.8]	61.2 [55.5, 66.8]
Terminal-only	86.4 [82.3, 90.5]	30.9 [26.6, 35.1]	55.5 [49.9, 61.2]
No learned scalar	85.8 [81.6, 89.9]	39.8 [35.5, 44.3]	46.0 [41.2, 51.1]

PepBench

Procedure	Contact	Overlap	Joint
Full suffix	88.6 [82.1, 93.9]	31.1 [24.6, 37.5]	57.5 [49.6, 65.4]
Terminal-only	85.7 [79.3, 91.1]	33.9 [27.9, 40.4]	51.8 [43.6, 60.0]
No learned scalar	86.8 [80.7, 92.1]	41.8 [36.1, 47.5]	45.0 [38.2, 52.1]

LNR

Procedure	Contact	Overlap	Joint
Full suffix	92.3 [85.5, 97.9]	26.8 [20.9, 33.2]	65.5 [57.4, 73.2]
Terminal-only	87.2 [81.3, 92.8]	27.2 [21.3, 33.2]	60.0 [52.3, 67.7]
No learned scalar	84.7 [78.7, 90.2]	37.4 [31.1, 43.8]	47.2 [40.4, 54.5]

In Table 18, H2T support tracks emission across schedules, while the larger separation appears at export. Without the scalar, lower canonical-L availability and more severe overlap reduce joint delivery. Terminal-only recovers part of that gap, and full-suffix guidance performs best at both stages.

Table 19: **Receptor ordering among registered and available outputs.** Rank@k entries use all registered outputs and are percentages; AUC entries use the normalized 0–1 scale. Brackets give percentile 95% component-bootstrap intervals. Registered AUC assigns absent outputs rank 12, whereas conditional AUC uses canonical-L outputs only.

Rank thresholds

Pooled

Procedure	Rank@1	Rank@2	Rank@5
Full suffix	63.7 [57.9, 69.3]	81.0 [75.7, 85.6]	90.1 [85.6, 94.0]
Terminal-only	60.0 [54.0, 65.6]	75.3 [70.1, 80.4]	85.6 [81.2, 89.9]
No learned scalar	55.5 [50.1, 61.0]	71.8 [66.6, 76.9]	85.2 [80.8, 89.5]

PepBench

Procedure	Rank@1	Rank@2	Rank@5
Full suffix	61.4 [53.2, 69.3]	80.0 [72.9, 86.4]	88.2 [81.8, 93.6]
Terminal-only	55.0 [47.5, 62.5]	73.9 [66.8, 80.4]	85.0 [78.9, 90.4]
No learned scalar	55.4 [48.2, 62.5]	72.1 [65.4, 78.6]	86.4 [80.4, 91.8]

LNR

Procedure	Rank@1	Rank@2	Rank@5
Full suffix	66.4 [58.3, 74.5]	82.1 [74.5, 88.9]	92.3 [85.5, 97.9]
Terminal-only	66.0 [57.0, 74.9]	77.0 [68.5, 84.7]	86.4 [80.0, 92.3]
No learned scalar	55.7 [46.8, 64.7]	71.5 [63.0, 79.1]	83.8 [77.0, 89.8]

Rank AUC

Pooled

Procedure	Registered AUC	Conditional AUC
Full suffix	0.867 [0.823, 0.906]	0.960 [0.951, 0.969]
Terminal-only	0.824 [0.782, 0.864]	0.954 [0.941, 0.965]
No learned scalar	0.808 [0.766, 0.849]	0.942 [0.928, 0.954]

PepBench

Procedure	Registered AUC	Conditional AUC
Full suffix	0.850 [0.787, 0.904]	0.960 [0.948, 0.971]
Terminal-only	0.814 [0.754, 0.868]	0.949 [0.934, 0.963]
No learned scalar	0.819 [0.761, 0.871]	0.944 [0.928, 0.959]

LNR

Procedure	Registered AUC	Conditional AUC
Full suffix	0.887 [0.822, 0.940]	0.961 [0.947, 0.973]
Terminal-only	0.837 [0.775, 0.894]	0.960 [0.939, 0.976]
No learned scalar	0.795 [0.732, 0.854]	0.939 [0.915, 0.959]

Conditional AUCs are close once a canonical-L product is available (Table 19). Registered AUC also reflects whether a product reaches that stage, linking receptor ordering to end-to-end delivery. The source tables provide the underlying slot records and track-specific intervals.

E.4 CONTACT-CONDITIONED ATOMIC OVERLAP

A delivered complex has evaluable coordinates, contacts the receptor, and avoids severe overlap. Each yield averages five outputs within each of 103 receptor components; continuous coordinate summaries state their coverage. At the thresholds used in the main text, the direct-delivery rates are 315/515 (61.2%; 95% CI 55.5–66.8%) for HYBRIDGE, 26/515 (5.0%; 3.1–7.2%) for CP-Composer, and 18/515 (3.5%; 1.6–5.8%) for RFpeptides.

Among components with contacting coordinates ($n = 100$, 78, and 23), the mean of the component-median maximum overlaps is 1.58, 2.20, and 1.94 Å for HYBRIDGE, CP-Composer, and RFpeptides. Corresponding overlap-pair densities are 362, 932, and 1,422 pairs per 1,000 peptide heavy atoms. Conditioning this comparison on contact keeps a distant output from appearing well packed simply because it has no atom pairs near the receptor.

E.5 BACKBONE GEOMETRY AND CANONICAL-L AVAILABILITY

Recomputing the released certificate on all 1,177 coordinate-bearing outputs reproduces every stored full-backbone flag. Across all 515 requests, the Canonical-L endpoint is available for 465 HYBRIDGE outputs (90.3%; 85.8–94.4%), 131 CP-Composer outputs (25.4%; 21.6–29.5%), and 135 RFpeptides outputs (26.2%; 22.3–30.5%).

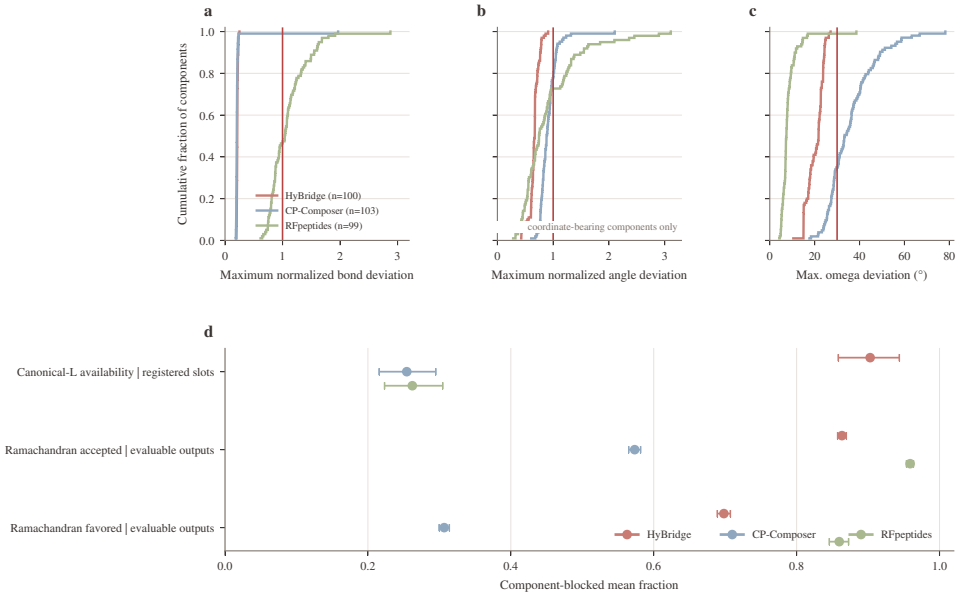


Figure S11: **Backbone geometry and Canonical-L availability. a–c,** Component-level distributions of maximum normalized bond and angle deviations and maximum ω deviation among coordinate-bearing components; dashed lines mark the certificate limits. **d,** Canonical-L availability among the first five registered outputs per component (515 requests per method), and Ramachandran accepted and favored fractions from the separate 24-output geometry audit. The latter uses 2,232, 2,169, and 2,472 available native backbones for HYBRIDGE, CP-Composer, and RFpeptides, respectively. Points are receptor-component means and bars are 95% bootstrap intervals.

The coordinate statistics separate backbone availability from the geometry of returned structures. Mean component-median maximum bond deviations are 0.214, 0.223, and 1.085; angle deviations are 0.658, 0.911, and 0.903; and maximum ω deviations are 20.4°, 35.9°, and 8.2°, respectively. In the 24-output geometry audit, RFpeptides has the strongest conditional Ramachandran distribution (95.9% accepted and 86.0% favored), whereas HYBRIDGE combines 86.3% acceptance and 69.8% favored torsions while reaching the Canonical-L endpoint for a much larger share of registered requests (Fig. S11d).

Among the 315 delivered HYBRIDGE complexes, the component-mean Ramachandran accepted and favored fractions are 86.0% and 69.0%. The corresponding fractions are 58.5% and 34.3% for the 26 delivered CP-Composer complexes, and 98.3% and 93.5% for the 18 RFpeptides complexes. These subsets use the same frozen preference grids and include every residue of the closed ring. Of the 315 delivered HYBRIDGE complexes, 96 have every residue in an accepted Ramachandran region.

E.6 SEQUENCE DIVERSITY AND LENGTH DEPENDENCE

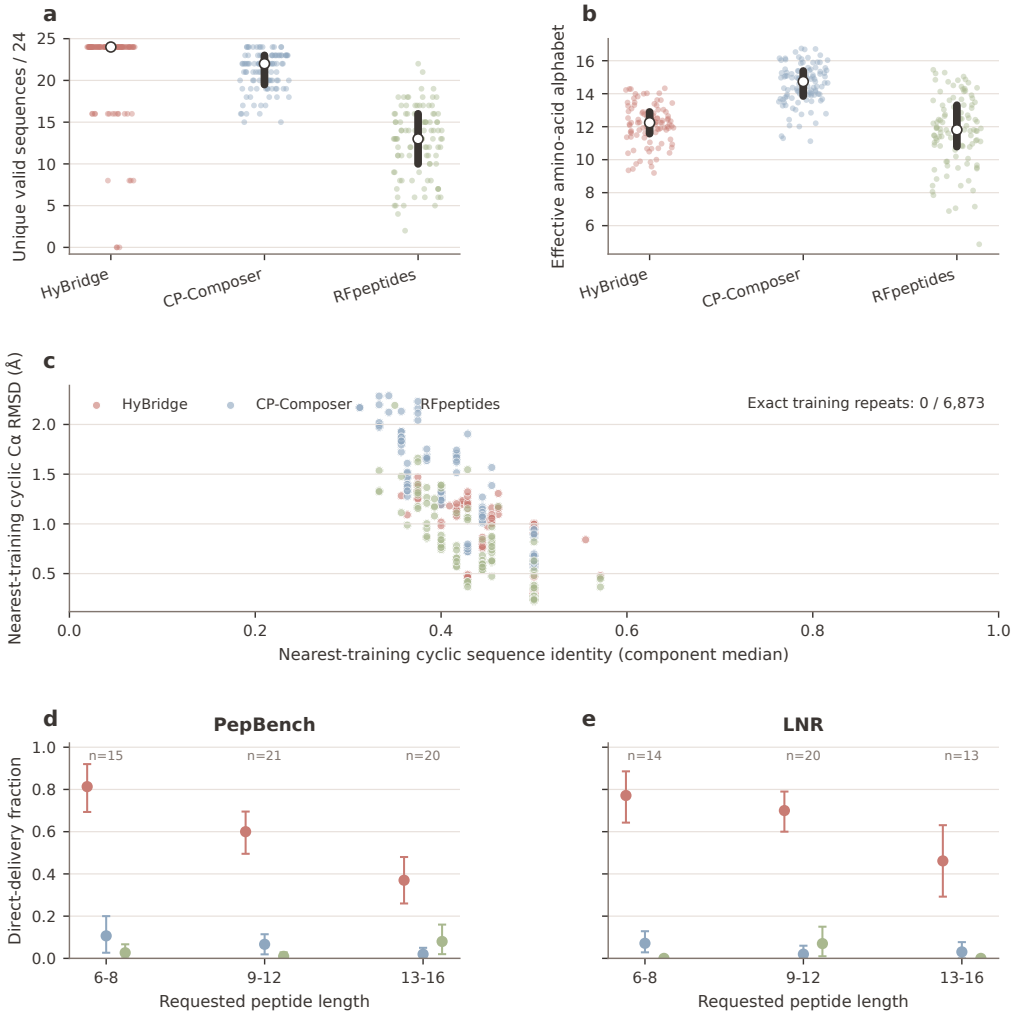


Figure S12: **Diversity, source proximity, and length-stratified delivery. a,b,** Component distributions of distinct cyclic sequences passing method-native ring checks among 24 registered outputs, and effective amino-acid alphabet size. Here "valid" refers to each method's native ring check; the common full-backbone certificate is evaluated in Table 17. **c,** Nearest-training cyclic sequence identity and cyclic Cα RMSD. **d,e,** Direct delivery by requested-length bin in PepBench and LNR; labels give the number of receptor components. Bars show 95% receptor-component bootstrap intervals.

Across 24 registered outputs per receptor component, HYBRIDGE yields a mean 21.7 distinct cyclic sequences passing its native ring check, compared with 21.1 for CP-Composer and 12.7 for RFpeptides under their respective native checks. Cyclic rotations count as the same sequence. Its effective amino-acid alphabet is 12.1 (14.6 and 11.7 for the baselines). Across the 6,873 axes with a training comparison, none exactly repeats a training sequence under cyclic rotation. Figure S12c adds the identity and RMSD of the nearest observed training neighbor.

Delivery declines with requested length but preserves the cross-method ordering in every track-by-length block. For HYBRIDGE, rates across lengths 6–8, 9–12, and 13–16 are 81.3%, 60.0%, and 37.0% in PepBench and 77.1%, 70.0%, and 46.2% in LNR. Both tracks show the same dependence on peptide length.

E.7 WHERE OUTPUTS FAIL

The procedures lose candidates at different stages (Figure S13). Before interface evaluation, HYBRIDGE records 49 source-support non-emissions and one numerical failure. CP-Composer has 61 native-closure non-emissions, 319 certificate failures, and four numerical failures; RFpeptides has 253 native-ring non-emissions and 127 certificate failures. Among the surviving structures, severe overlap accounts for 150 HYBRIDGE and 105 CP-Composer requests, while RFpeptides contributes 95 separated complexes and 22 with severe overlap. The certificate failures divide into full-backbone only, seam plus full-backbone, and closure-edge plus seam plus full-backbone counts of 270/47/2 for CP-Composer and 113/1/13 for RFpeptides.

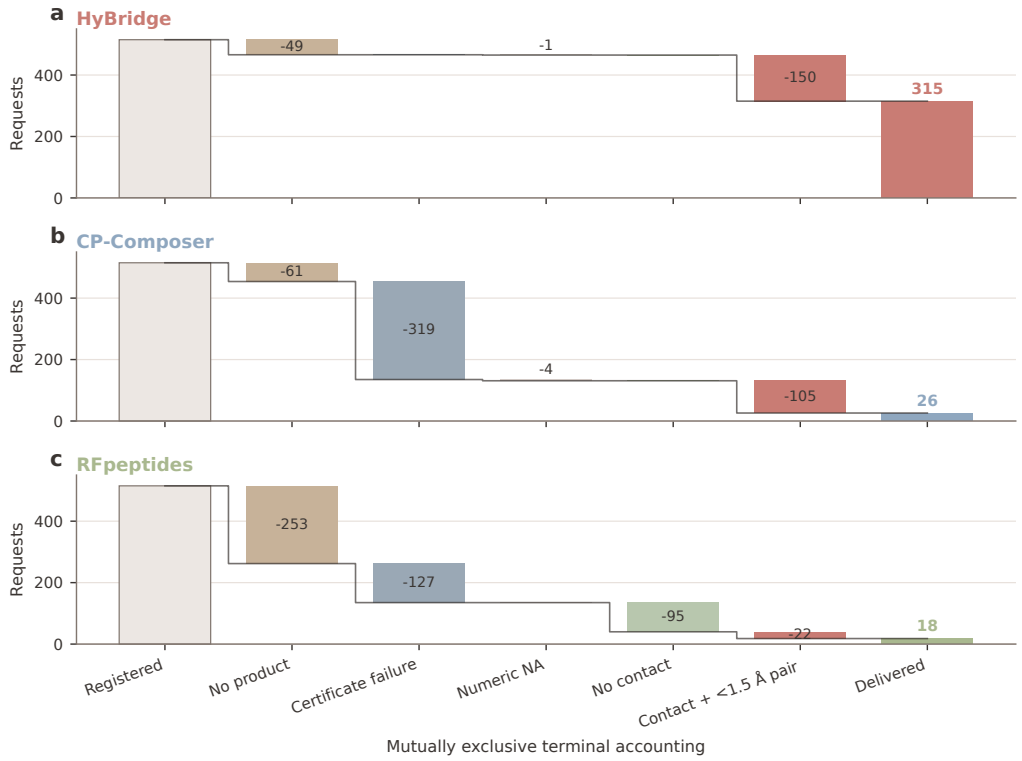


Figure S13: **Mutually exclusive outcomes for all requests.** Each panel begins with 515 requests and removes, in order, no method-supported product, post-reload certificate failure, numerical endpoint failure, lack of a 4 Å contact, and contact with a sub-1.5 Å interchain pair. The final bar gives direct delivery.

Figure S14 puts coordinates behind these categories. The selected HYBRIDGE overlap has $d_{\min} = 1.23$ Å and a maximum van der Waals overlap of 2.17 Å; the CP-Composer backbone has a 42.3° maximum ω deviation; and the nearest interchain heavy atoms in the RFpeptides separation case are 36.74 Å apart. The threshold grid and complete funnel show how often each type recurs across all 515 requests.

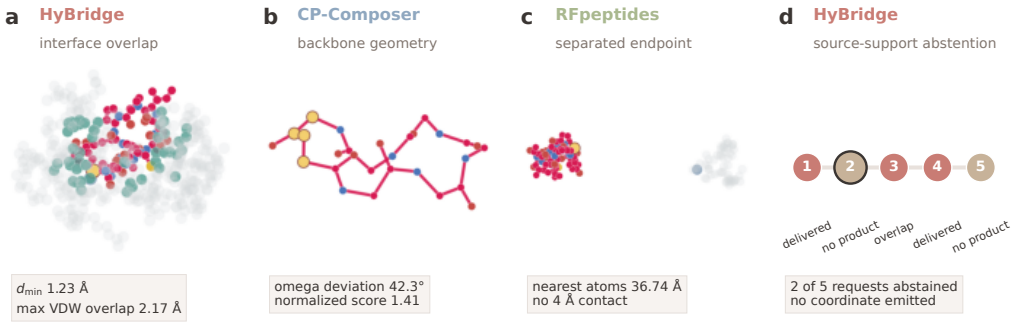


Figure S14: **Representative failures at median category severity. a–c,** Coordinate views closest to the median within the HYBRIDGE interface-overlap, CP-Composer backbone-geometry, and RFpeptides separated-endpoint categories. Receptor atoms are pale gray, contacting receptor atoms teal, peptide atoms magenta, and the diagnostic atoms gold. **d,** A HYBRIDGE source-support abstention component; the outlined slot has no coordinate and is shown as a status card.

F THE CYCLIC-PEPTIDE GENERATION LANDSCAPE

Peptide generators differ in the molecular graph they produce and in when the receptor enters generation. PPFlow models receptor-conditioned linear peptides in torsion space (Lin et al., 2024). DiffPepBuilder jointly generates backbone frames and sequence before optional disulfide construction (Wang et al., 2024). CpSDE couples all-atom structure and residue updates with explicit bond modeling (Zhou et al., 2025). Together, they establish much of the sequence, geometry, and receptor-conditioning machinery on which cyclic peptide generation builds.

CP-Composer jointly generates sequences and structures using type- and distance-conditioned latent diffusion with classifier-free guidance at each denoising step (Jiang et al., 2025). Its released pipeline decodes candidates and applies a cyclization-specific relaxer, including H2T closure. RFpeptides combines RFdiffusion with ProteinMPNN (Rettie et al., 2025). Both accept receptor-plus-H2T requests and return receptor-relative cyclic coordinates, matching the complete task evaluated here.

GeoCycler develops topology-aware cyclic generation (Zhang et al., 2026), and APCyc uses an autoencoder–diffusion hierarchy for cyclic peptides (Zhao et al., 2026). At the benchmark freeze, the public release status also differed. GeoCycler and CpSDE had no official runnable release, while APCyc provided code but no public checkpoint for its reported model. The direct numerical panel uses the native executable pipelines of CP-Composer and RFpeptides; the other methods are discussed through their published evidence and generation interfaces.

HYBRIDGE draws on flow matching (Lipman et al., 2023), continuous-time Markov chains (Campbell et al., 2022), generator matching (Holderrieth et al., 2025), and Riemannian diffusion (De Bortoli et al., 2022). Diffusion guidance (Dhariwal & Nichol, 2021), diffusion bridges (De Bortoli et al., 2021), controlled particle methods (Heng et al., 2020), and DEFT (Denker et al., 2024) supply complementary control ideas. HyBridge organizes these ingredients into an executable target-conditioning process with compiled support, crossed interaction contrasts, and coordinate and token controls constructed from one scalar.

G REPRODUCING THE RESULTS

The anonymous package includes a relative-path receptor-PDB-to-H2T-PDB demo, all five frozen checkpoints, the packaged ESM-2 asset, and their public configuration. Source data accompany every reported figure and table. Tests cover quotient invariance, scalar consistency, topology transport, rollout accounting, and PDB materialization, and scripts reproduce the public OOD tables. A target-free compile–project–retract–certify demonstration runs all six executable topology fibres; H2T additionally provides the full receptor-PDB-to-cyclic-PDB generation path. The compact mirror retains the code, source data, and CPU verification paths without the large numerical assets.

H SEQUENCE-LEVEL RECEPTOR RESPONSE WITH ALPHAFOLD 3

The crossed-response analysis in Appendix D.2 measures how generation responds to the learned receptor potential. Here we evaluate the generated sequences with AlphaFold 3 (AF3) (Abramson et al., 2024), which predicts a new complex for each peptide–receptor pairing. This provides a sequence-level view of receptor response using a separate structure predictor.

Matched sequences and receptor swaps. Of the eleven reciprocal receptor pairs in the parent matched-rollout panel, seven have completed outputs in both directions under all three arms. In the other four pairs, at least one run abstains because source or terminal-stratum support is insufficient. We retain all eight frozen output slots per direction for the common seven pairs, giving 112 outputs per arm and 336 in total. This analysis measures receptor response conditional on completed matched generation. The three arms share the score-blind receptor-relative placement protocol and use no learned potential, terminal-only scoring, or guided-suffix execution. This common-pair panel is separate from the 103-component delivery benchmark. Each sequence is paired with its intended receptor and with the reciprocal receptor, giving 672 AF3 tasks; all tasks completed.

AF3 v3.0.1 receives the cyclic sequence as a single multi-CCD ligand with explicit H2T covalent bonds, together with the receptor protein and its preprocessed MSA and templates. The peptide inputs contain neither generated coordinates nor the learned receptor score. We retain all five samples from model seed 20260826 for every task, yielding 3,360 predicted complexes.

Pair-level analysis. For each output, the target-swap margin is the median five-sample receptor–peptide chain-pair ipTM for the intended receptor minus the corresponding median for the reciprocal receptor. We average the 16 margins within each reciprocal pair, then give each of the seven pairs equal weight. Guided-control contrasts use the same pairs. Percentile 95% intervals resample intact pairs 10,000 times with seed 20260826, retaining both directions and all output slots together.

Table 20: **AF3 target-swap response of frozen generated sequences.** Higher margins indicate greater predicted interface confidence with the intended receptor. Intervals and positive counts use seven reciprocal pairs.

Execution or paired contrast	Mean	95% interval	Positive pairs
No learned potential	-0.0083	$[-0.0211,\ 0.0030]$	4/7
Terminal-only scoring	0.0039	$[-0.0072,\ 0.0141]$	5/7
Guided suffix	0.0170	$[\ 0.0074,\ 0.0271]$	6/7
Guided − no learned potential	0.0253	$[\ 0.0093,\ 0.0407]$	6/7
Guided − terminal-only	0.0130	$[-0.0030,\ 0.0270]$	5/7

The guided arm has a positive target-swap margin in six of seven pairs (Table 20). The improvement over the no-potential control has an interval above zero. Against terminal-only scoring, the mean improvement is positive in five pairs and its interval includes zero.

The absolute ipTM values show how the target-swap margins arise. Mean intended/reciprocal values are 0.6925/0.7008 without the learned potential, 0.6938/0.6898 with terminal-only scoring, and 0.6991/0.6821 with the guided suffix, using the same within-pair averaging as the margins.

For a scale-free summary, Cohen's paired effect is $d_z = \bar{d}/s_d$, where d_i is the guided-control margin difference for reciprocal pair i and s_d is its sample standard deviation. The effects are $d_z = 1.13$ against no learned potential and $d_z = 0.59$ against terminal-only scoring. The empirical paired probability of superiority is $\widehat{\mathrm{PS}} = \#\{d_i > 0\}/\#\{d_i \neq 0\}$ (Gyimesi et al., 2025, Table 2, entry 29). With no tied pair means, it equals 0.857 (6/7) and 0.714 (5/7), respectively. PS expresses the frequency of a positive paired effect on a 0–1 scale; the raw mean difference reports its magnitude.

All 3,360 predictions retain an H2T terminal C–N distance of at most 1.8 Å, and every task has at least three samples for which AF3's `has_clash` flag is false. Together, the results support a sequence-level receptor response relative to the no-potential control under independent structure prediction, measured through predicted interface confidence.