DC-FLOW: TARGET-CONDITIONED CYCLIC PEPTIDE DESIGN VIA DELAYED-COUPLING FLOW MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Target-conditioned cyclic peptide design requires coordinating interfacial side chains with macrocyclic closure. Existing end-to-end generators concurrently update geometric and chemical variables from noise, assigning residues to disordered coordinates. Meanwhile, predictor-assisted workflows often get stuck in local traps because they lack feedback in the propose-and-filter stage and ignore 3D structural proposals during sequence search. To overcome the drawbacks of these two paradigms, we propose DC-FLOW, a staged generative framework built around **D**elayed-**C**oupling Flow Matching. In principle, DC-FLOW postpones sequence and side-chain activation until a target-conditioned backbone takes shape, resolving structural ambiguity before chemical assignment. Then, it feeds delayed activations back into the backbone, achieving co-evolution of backbone geometry and emerging residue identities by coupling their flows. This mechanism enables more reliable macrocyclic closure and favorable interface packing. For predictor-assisted workflows, DC-FLOW supplies 3D proposals with in-trajectory guidance, enabling cascaded test-time scaling to prune redundant backbones and terminate unpromising refinement early. Extensive experiments show that DC-FLOW improves cyclization fidelity, diversity, and interface energetics, substantially increasing the yield of confident cyclic peptide binders under a fixed final-candidate budget.

1 INTRODUCTION

As head-to-tail macrocycles, cyclic peptides can engage broad, shallow, polar protein-protein interaction (PPI) interfaces that often lack pockets suitable for small molecules (Dougherty et al., 2019; Vinogradov et al., 2019). Well-designed macrocycles can protect against protease attack, reduce access to protease-susceptible backbone states (Malde et al., 2019), and exhibit molecular chameleonicity.[1] However, target-conditioned design of cyclic peptides remains challenging due to data scarcity. High-resolution complex structures are severely limited compared with the vast space of peptide sequences, lengths, and cyclization modes (Yang et al., 2026; Rettie et al., 2025a;b). Moreover, minor changes near the cyclization site can shift the entire backbone, yielding multiple low-energy states and binding poses. Therefore, a target-conditioned design model must jointly solve target docking, backbone closure, sequence design, and side-chain packing under sparse, heterogeneous supervision (Zhou et al., 2025; Wang et al., 2026).

For cyclic peptides, existing target-conditioned design methods generally fall into two paradigms, as shown in Figure 1 (a) and (b). On one hand, full-atom generators such as PepFlow (Li et al., 2024), PepGLAD (Kong et al., 2024), DiffPepBuilder (Wang et al., 2024), and cyclic-specific models like APCyc (Zhao et al., 2026) (adapted from dyMEAN (Kong et al., 2023)) offer high throughput by amortizing design into generative trajectories. However, the data scarcity issue limits their generalizability, and their outputs are often geometrically plausible but fail to recover confident interfaces under cofolding. On the other hand, predictor-assisted methods, e.g., RFPeptides (Rettie et al., 2025b), AfCycDesign (Rettie et al., 2025a), and HFGuidedDesign (Hu et al., 2026), integrate structure prediction into the proposal, optimization, and filtering pipeline, enriching candidates with high self-consistency. While effective, repeated and exhaustive cofolding incurs prohibitive computational costs, severely restricting candidate exploration under a fixed compute budget.

[1] For some cyclic peptides, their polar groups remain exposed in water but fold to hide their polarity in apolar media, facilitating membrane permeation (Wang et al., 2021; Linker et al., 2022).

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.

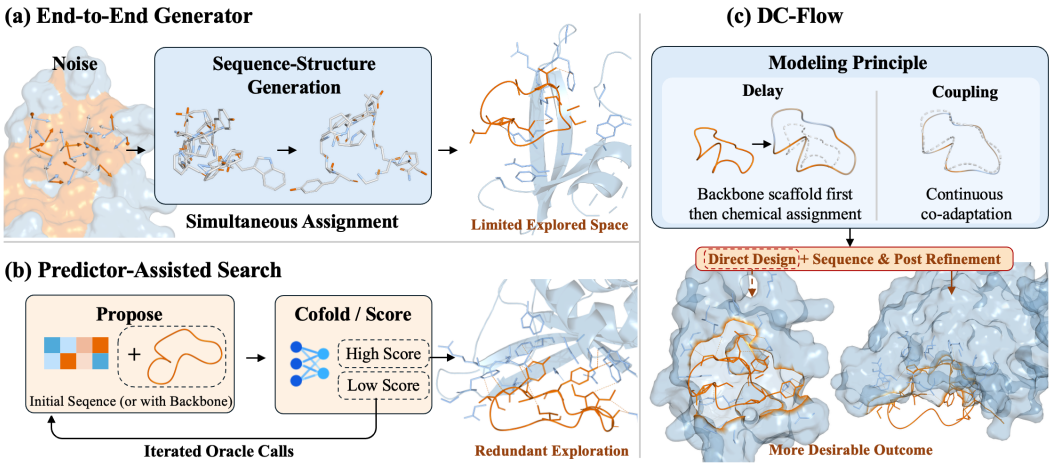

Figure 1: **DC-FLOW and Two Design Paradigms.** **(a)** End-to-end generators simultaneously update sequence and structure from noise, resulting in early disorder and limited explore space. **(b)** Predictor-assisted workflows rely on iterative oracle evaluations without dynamic structure-sequence feedback, resulting in redundant exploration. **(c)** DC-FLOW resolves early structural ambiguity by delaying chemical assignment, then couples their flows to achieve continuous co-adaptation for both end-to-end design and predictor-assisted refinement.

To bridge the gap between direct generation and predictor-assisted design and overcome their drawbacks, we propose DC-FLOW, a staged generative framework built on a delayed-coupling flow matching strategy. In principle, DC-FLOW is based on a key fact: **while side-chain interactions supply critical binding affinity at the PPI interface, the backbone scaffold dictates their optimal spatial presentation** (Clackson & Wells, 1995). Motivated by this, delayed-coupling flow matching postpones sequence and side-chain flows until the backbone scaffold takes shape, improving cyclic peptide design by enriching essential geometric context. As shown in Figure 1 (c), once each of the flows is activated, it couples with the backbone in a shared network, allowing backbone geometry and residue identities to co-evolve over time. To complement cyclic-complex supervision with general structure-conditioned sequence statistics, we augment a frozen ProteinMPNN prior (Dauparas et al., 2022) with trainable residual modules. Crucially, this staged formulation naturally unlocks *cascaded test-time scaling*, an adaptive search policy that selects diverse backbones mid-generation, guides sequence sampling along the trajectory, and prunes unpromising cofolding optimization early.

We comprehensively evaluate DC-FLOW on the LNR benchmark (Tsaban et al., 2022) as a full-atom generator. In addition, we test it on the design tasks in the published HFGuidedDesign targets (Hu et al., 2026) and a challenging TL1A trimer case (Zwolak et al., 2022) as a cofolding proposal generator. We also provide systematic ablations of our asynchronous staging, prior integration, and test-time search strategy. Experiments show that DC-FLOW can reliably discover cyclic peptide conformations with high prediction confidence and superior binding interfaces. It achieves the highest cyclization success on LNR, leads cofolding baselines in strict-hit rate and mean Rosetta ΔG_{sep}, and further boosts these metrics through cascaded search.

2 RELATED WORK

Target-conditioned generative peptide design. Generative approaches aim to generate protein sequences and structures conditioned on receptor pockets. For example, APM (Chen et al., 2025) models all-atom multi-chain complexes for conditional generation and inverse folding, and Genie 3 (Lin et al., 2026) enhances atomistic SE(3)-equivariant diffusion for efficient binder design. Focusing on peptides, PepGLAD (Kong et al., 2024) embeds variable-size residues for receptor-conditioned diffusion. PepFlow (Li et al., 2024) defines multiple flows over rigid frames, discrete residues, and side-chain torsions to support full-atom design. DiffPepBuilder (Wang et al., 2024) utilizes an SE(3)-equivariant diffusion model to design peptide sequences and conformations with disulfide stabilization. Focusing on cyclic peptides, APCyc (Zhao et al., 2026) extends dyMEAN (Kong et al., 2023) by incorporating an expanded residue vocabulary for cyclization sites and applying Bayesian

optimization. CPSea (Yang et al., 2026) benchmarks these methods on cyclic peptide design, but it relies on pseudo-complexes rather than real ones. Notably, these generators simultaneously update conformations and sequences from scratch along a single trajectory, while DC-FLOW stages the generation of backbone, sequence, and side chains in a delayed-coupling way.

Predictor-guided design and cofolding pipelines. A complementary paradigm incorporates structure predictors and cofolding models to guide candidate search, framing design as heuristic optimization over an energy landscape. At one end, AfCycDesign (Rettie et al., 2025a) introduces cyclic relative positional encodings into AlphaFold2 (Jumper et al., 2021) to establish a cyclic cofolding oracle, but its de novo design relies on time-consuming sequence hallucination and is prone to getting trapped in local optima. At the other extreme, RFPeptides (Rettie et al., 2025b) adopts a propose-and-filter pipeline, where RFdiffusion (Watson et al., 2023) proposes backbones, ProteinMPNN predicts sequences, and cofolding models (Rettie et al., 2025a) filter sequences. HFGuidedDesign (Hu et al., 2026) represents a middle ground by steering discrete sequence diffusion with online HighFold evaluations (Zhang et al., 2024), but it operates purely in sequence space and relies on the external conformation predictors. DC-FLOW bridges these approaches, generating coupled backbone-sequence candidates with in-trajectory guidance.

Test-time scaling in biomolecular design. Recently, Proteina-Complexa (Didi et al., 2026) combines a pretrained atomistic flow model with various searching and steering strategies, demonstrating that generative initializations benefit predictor-assisted sequence hallucination. In protein inverse folding, ProtInvTree (Liu et al., 2026) formulates sequence design as reward-guided tree search to increase search breadth and depth without retraining. MuCO (Wang et al., 2026) addresses a sequence-conditioned cyclic peptide setting by decomposing cyclization into backbone generation, side-chain packing, and all-atom refinement. Building on this concept, DC-FLOW proposes cascaded test-time scaling to concentrate expensive refinement compute on the most promising candidates.

3 PROPOSED METHOD

3.1 PROBLEM FORMULATION AND MODELING PRINCIPLE

In principle, we formulate the target-conditioned cyclic peptide design problem through the lens of conditional generative modeling, learning flow-based generators to capture the evolution of the *design state* over time. In particular, given a fixed receptor pocket $R = (S^R, \boldsymbol{X}^R)$ comprising its residue sequence S^R and 3D coordinates $\boldsymbol{X}^R$, we aim to generate a cyclic peptide of length L conforming to cyclization specification Γ (e.g., head-to-tail or disulfide closure), summarized as conditions $\xi = (L, \Gamma)$. At time $t \in [0, 1]$, the peptide design state is represented by $Y_t = (B_t, S_t, C_t)$, where B_t denotes backbone scaffold, S_t is a continuous carrier for residue identities over alphabet $\mathcal{A}$ yielding discrete sequence $S_1 \in \mathcal{A}^L$, and C_t describes side-chain conformations. Starting from an initial state sampled from an uninformative prior distribution, i.e., $Y_0 \sim p_0$, the model with trainable parameters θ learns the conditional distribution $p_\theta(Y_1 \mid R, \xi)$.

As aforementioned, at PPI interfaces, side-chain interactions determine critical binding affinity, but their spatial presentation is fundamentally dependent on backbone scaffold (Clackson & Wells, 1995). Therefore, generating all modalities (backbones, side chains, and sequences) simultaneously forces sequence and side chains to develop over disordered, unformed backbone scaffolds. In contrast, decoupling backbone generation with side-chain packing and sequence design restricts the exploration of the design space and thus leads to reduced search efficiency.

Motivated by the above analysis, DC-FLOW decomposes the generative trajectory into three sequential but overlapping stages demarcated by activation thresholds $0 = \tau_\mathrm{b} < \tau_\mathrm{s} < \tau_\mathrm{c} < 1$:

$$p_\theta(Y_1|R,\xi) = \int \underbrace{p_\theta(B_{\tau_\mathrm{s}}|R,\xi)}_{\text{Backbone sampler}} \underbrace{p_\theta(B_{\tau_\mathrm{c}}, S_{\tau_\mathrm{c}}|B_{\tau_\mathrm{s}}, R,\xi)}_{\text{Sequence generator}} \underbrace{p_\theta(Y_1|B_{\tau_\mathrm{s}}, B_{\tau_\mathrm{c}}, S_{\tau_\mathrm{c}}, R,\xi)}_{\text{Full-atom packer}} \mathrm{d}B_{\tau_\mathrm{s}}\mathrm{d}B_{\tau_\mathrm{c}}\mathrm{d}S_{\tau_\mathrm{c}}. \quad (1)$$

Here, τ's mark the modality activation times, B_τ and S_τ denote the backbone and sequence states at activation time $t = \tau$. This formulation first establishes a receptor-conditioned backbone to resolve structural ambiguity before sequence assignment. It activates sequence and side chains over the formed scaffold while keeping the backbone flexible, co-evolving backbone and residue identities.

Figure 2 provides an overview of DC-FLOW across both standalone full-atom generation and cofolding pipelines, which contains two cohesive mechanisms: (i) delayed-coupling flow matching

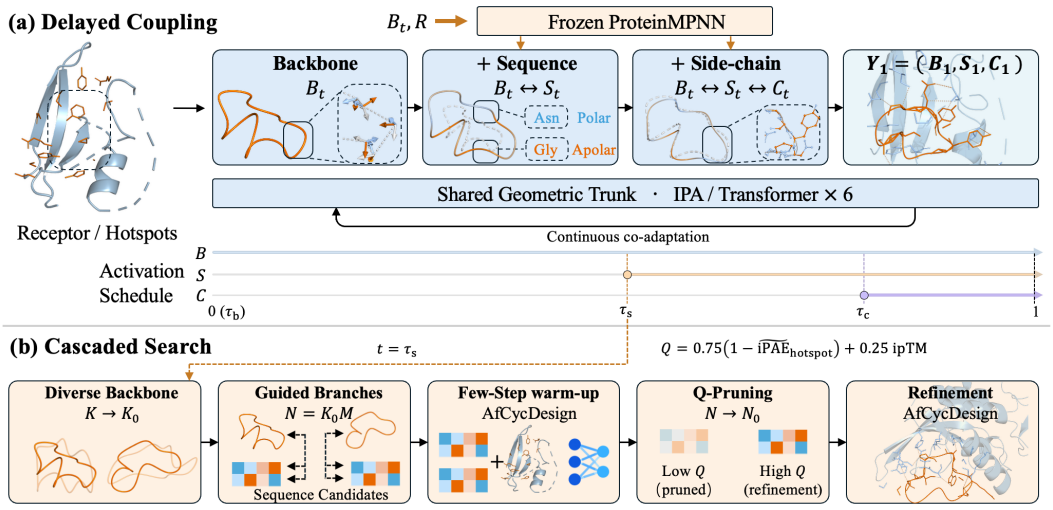

Figure 2: **The overview of DC-FLOW. (a) Delayed coupling:** Modality clocks activate backbone, sequence (augmented by ProteinMPNN prior), and side chains sequentially, enabling continuous co-adaptation in a shared network. **(b) Cascaded search:** At τ_s, backbones are pruned to diverse roots, each branching into M oracle-guided continuations ($N = K_0 M$). After a few-step warm-up, candidates are screened by score Q to restrict expensive refinement.

for staged backbone, sequence, and side-chain generation, conditioned on an external ProteinMPNN prior; and (ii) cascaded test-time scaling for efficient cofolding refinement.

3.2 DELAYED-COUPLING FLOW MATCHING

DC-FLOW learns a flow-based model to generate backbones, sequences, and side chains asynchronously. For modality $m \in \{\text{b}, \text{s}, \text{c}\}$, let τ_m denote the threshold for activating its generative flow. Accordingly, independent modality progress clocks are defined as $s_m(t) = \text{clip}(\frac{t-\tau_m}{1-\tau_m}, 0, 1)$, with activation boundaries $0 = \tau_\text{b} < \tau_\text{s} < \tau_\text{c} < 1$ (by default, $0, 0.5, 0.8$), which leads to the proposed delayed-coupling mechanism that corresponds to two complementary design principles:

Delaying: Resolving Structural Ambiguity. Before τ_s, the flow-based generator evolves only the backbone geometry B_t, while sequence inputs are replaced by uninformative mask tokens and side-chain inputs use a fixed zero-angle encoding. Under the reference-path assumptions in Appendix B.1, earlier backbone states are statistically less informative observations of the target structure B_1. Along this reference path, the residual sequence ambiguity, measured by conditional mutual information $D(t) \triangleq I(S_1; B_1 \mid B_t, R, \xi)$, is non-increasing as structural noise decreases, with $I(a; b \mid c) \triangleq H(a \mid c) - H(a \mid b, c)$ measuring the excess Bayes sequence risk. Delaying sequence activation to τ_s avoids conditioning amino-acid assignment in highly ambiguous states of the early trajectory.

Coupling: Continuous Mutual Co-Adaptation. Crucially, delaying sequence does *not* imply a hard, decoupled two-stage pipeline (e.g., generating a rigid backbone and subsequently fixing it for inverse folding). Once sequence activates at τ_s, it enters the shared geometric trunk to jointly update the remaining trajectory alongside B_t. Formally, letting R_B and R_{BS} denote Bayes risks of predicting backbone target without and with sequence S_t, we have $R_B - R_{BS} \triangleq \Delta_t^{B \leftarrow S} \geq 0$ (Proposition B.2, proof in Appendix B.2), where the error reduction is strict whenever sequence alters the conditional mean of the target. This continuous coupling allows the backbone scaffold to remain flexible, actively adjusting its local geometry to accommodate specific side-chain packing requirements and emerging residue interactions. At τ_c, side-chain torsions C_t activate to finalize all-atom packing.

Staged Training Objective. The generator is optimized end-to-end using flow matching and cross-entropy objectives. Let ℓ_B train backbone rigid frames and coordinates, ℓ_S denote sequence cross-entropy, ℓ_C supervise side-chain torsion angles, and ℓ_cyc penalize cyclic closure errors (head-to-tail and disulfide constraints). With nonnegative loss weights $\lambda_B, \lambda_S, \lambda_C$ and indicator function

$\mathbf{1}[\cdot]$, modality-specific loss terms activate in accordance with their respective time thresholds. The grouped objective is:

$$\mathcal{L} = \mathbb{E}_{t\sim\mathcal{U}[0.01,0.99]}[\lambda_B \ell_B + \mathbf{1}[t \geq \tau_s]\lambda_S \ell_S + \mathbf{1}[t \geq \tau_c]\lambda_C \ell_C + \ell_{cyc}]\,. \tag{2}$$

The disulfide component of ℓ_{cyc} is active only after side-chain activation; detailed mathematical definitions and decompositions of all individual loss terms are provided in Appendix C.5.

Leveraging ProteinMPNN as a Prior of Sequence. In both training and inference, DC-FLOW incorporates a frozen ProteinMPNN (Dauparas et al., 2022) F_ϕ as a sequence prior, which successfully compensates for the scarcity of cyclic peptide complex structures. In particular, for $t > \tau_s$, the geometric trunk outputs residue features $h_{i,t}$, from which a sequence head predicts base logits $\ell^B_{i,t} = f_{AA}(h_{i,t})$. Concurrently, F_ϕ evaluates the evolving intermediate backbone B_t and receptor R, outputting log-probabilities $\ell^P_{i,t} = F_\phi(B_t; R)_i$. We center $\ell^P_{i,t}$ across alphabet $\mathcal{A}$ to form zero-mean residual logits $\widetilde{\ell}^P_{i,t}$. With a trainable scalar gate $g \in \mathbb{R}$ (initialized to 0.1), the fused logits are

$$\ell_{i,t} = \ell^B_{i,t} + g\,\widetilde{\ell}^P_{i,t}, \qquad \text{where} \quad \widetilde{\ell}^P_{i,t,a} = \ell^P_{i,t,a} - \frac{1}{|\mathcal{A}|}\sum\nolimits_{a'\in\mathcal{A}} \ell^P_{i,t,a'}. \tag{3}$$

Disallowed residues (e.g., non-cysteines at disulfide sites) are masked before softmax normalization.

- **Training: Structural Transfer without Gradient Leakage.** During training, the sequence cross-entropy loss $\ell_S = -\frac{1}{L}\sum_{i=1}^{L} \log \pi_{i,t,a^*_i}$ supervises the fused distribution $\pi_{i,t} = \mathrm{softmax}(\ell_{i,t})$ against native residues a^*_i. Crucially, evaluating ProteinMPNN on the noisy intermediate backbone B_t rather than clean coordinates teaches the network sequence–structure compatibility throughout generation. We apply a stop-gradient operator to B_t at the input of F_ϕ, ensuring that ℓ_S optimizes only DC-FLOW's native trunk, sequence head, and gate g without differentiating through the external prior.
- **Inference: Dynamic In-Trajectory Guidance.** At inference time, residual fusion induces a product-of-experts distribution over admissible amino acids. Suppressing indices i, t, let $\pi^B = \mathrm{softmax}(\ell^B)$ and $q^P = \mathrm{softmax}(\ell^P)$. As shown in Appendix B.3, the fused categorical distribution satisfies
$$\pi(a) \propto \pi^B(a)[q^P(a)]^g, \tag{4}$$
directly scaling the injected structural log-odds by g.

Unlike conventional pipelines (e.g., RFPeptides (Rettie et al., 2025b)) that apply ProteinMPNN as an isolated post-hoc step on fixed coordinates, DC-FLOW embeds the prior *inside the generative trajectory*. Discrete sequences sampled from π feed back into the trunk at subsequent flow steps, steering continuous backbone relaxation and side-chain packing. This closes an active sequence–structure feedback loop during generation.

3.3 CASCADED TEST-TIME SCALING

Although DC-FLOW supports direct full-atom generation, cofolding (Rettie et al., 2025a) with generated proposals is still bottlenecked by expensive oracle evaluations and iterative refinement. We therefore introduce *cascaded test-time scaling*, dynamically allocating inference compute across structural exploration, sequence branching, and selective refinement (Appendix B.4).

- **Backbone Selection ($t = \tau_s$).** We initiate K parallel backbone trajectories up to τ_s. Before activating sequence generation, we downsample to $K_0 < K$ diverse roots via kernel-based structural crowding (Appendix B.4.2), pruning geometrically redundant backbones.
- **Guided Sequence Branching ($t \in [\tau_s, 1]$).** Each retained root branches into M independent sequence continuations ($N = K_0 M$ trajectories total) while backbones continue evolving (Appendix B.4.1). Concurrently, an external cyclic cofolding oracle periodically evaluates single-residue substitutions, biasing logits $\ell_{i,t}$ toward favorable hotspot interactions (Hu et al., 2026). Through sequence–structure coupling, this guidance simultaneously steers subsequent backbone updates (Appendices B.4.3 and B.4.4).
- **Early-Pruned Cofolding Refinement.** At $t = 1$, the N generated designs undergo a brief 5-step AfCycDesign warm-up. We then probabilistically retain $N_0 < N$ candidates weighted by an interface quality score
$$Q = 0.75\big(1 - \widetilde{\mathrm{iPAE}}_{\mathrm{hotspot}}\big) + 0.25\,\mathrm{ipTM}, \tag{5}$$

where $\widetilde{\text{iPAE}}_{\text{hotspot}}$ and ipTM denote the normalized hotspot interface PAE and interface TM-score, respectively. Only these N_0 survivors proceed to full refinement, subject to early stopping if the score plateaus for five consecutive soft-stage updates (Appendix B.4.5). This funnel concentrates expensive refinement on high-potential candidates, improving discovery efficiency under fixed final-candidate budget.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Data Construction. We train DC-FLOW and all direct full-atom baselines on the same stereochemically filtered CPSea-derived corpus. The complete data sources, quality-control criteria, and filtering statistics are provided in Appendix D. Raw full-atom generation is evaluated on our processed Large Non-Redundant (LNR) benchmark (Tsaban et al., 2022), which contains 93 receptor-peptide complexes.

To quantify generalization with respect to the training receptors, we use MMseqs2 (Steinegger & Söding, 2017) to compare every LNR receptor against all receptors in the training set. For a test receptor $\mathcal{R}$, let

$$\rho(\mathcal{R}) \triangleq \max_{\mathcal{R}' \in \mathcal{D}_{\text{train}}} \text{SeqId}_{\text{MMseqs2}}\left(\text{seq}(\mathcal{R}), \text{seq}(\mathcal{R}')\right), \tag{6}$$

where $\mathcal{D}_{\text{train}}$ denotes the set of training receptors, $\text{seq}(\mathcal{R})$ is the amino-acid sequence of $\mathcal{R}$, and $\text{SeqId}_{\text{MMseqs2}}$ is the sequence identity returned by MMseqs2. We define 37 receptors with $\rho(\mathcal{R}) < 0.50$ as the out-of-distribution subset (OOD) and the remaining 56 receptors with $0.50 \leq \rho(\mathcal{R}) \leq 0.70$ as the in-distribution subset (ID). No LNR receptor has greater than 70% sequence identity to a training receptor. For every target, all methods receive the same receptor, peptide-length condition, and generation budget.

The cofolding comparison uses the 12 published HFGuidedDesign targets, which cover peptide lengths and receptor interfaces of varying complexity (Hu et al., 2026). Each method contributes the same number of final candidates per target, and all candidates are evaluated through a common cofolding and scoring pipeline.

Baselines and Implementation. The direct full-atom baselines comprise PepFlow (Li et al., 2024), PepGLAD (Kong et al., 2024), DiffPepBuilder (Wang et al., 2024), APM (Chen et al., 2025), Genie 3 (Lin et al., 2026), and APCyc (Zhao et al., 2026). All six methods are retrained using the same training corpus and data split as DC-FLOW, with method-specific architectures and sampling procedures otherwise retained. For the predictor-guided and coarse-to-refine comparison, we include AfCycDesign (Rettie et al., 2025a) as cofolding predictor, RFPeptides (Rettie et al., 2025b), and HFGuidedDesign (Hu et al., 2026) as proposal generator baselines. We follow their released inference configurations and evaluate their final designs with the same candidate budget and post-generation scoring pipeline used for DC-FLOW. On LNR, each baseline generates 32 candidates per target; DC-FLOW generates 32 candidates separately for head-to-tail (HT) and terminal disulfide (SS) cyclization. Generated complexes undergo a common closure-constrained Rosetta FastRelax protocol. On the 12-target cofolding benchmark, each configuration contributes 32 final candidates per target. Generation-time guidance, subsequent AfCycDesign optimization, and cascaded search are identified separately. Implementation details and evaluation protocols are provided in Appendix E.

Evaluation Metrics. As a full-atom generator, we report five metrics: cyclization success (Cyc.), Rosetta total energy (Stability), sequence diversity (D_{seq}), structural diversity (D_{struct}), and Rosetta interface energy (dG) (Zhao et al., 2026; Alford et al., 2017). The diversity scores measure the fractions of sequence and structure clusters among valid candidates. Higher success and diversity, and lower energies, are preferred. Cyclization criteria, validity filtering, clustering thresholds, and receptor-wise aggregation are specified in Appendix E.4. As a cofolding proposal generator, we report ipTM, peptide pLDDT, interface and hotspot iPAE, composite confidence scores, and strict-hit rates (Hu et al., 2026). The main comparison uses the interface-based composite S, whereas the search comparison uses its hotspot-based counterpart S_{hotspot}. Rosetta evaluation additionally reports interface energy (ΔG_{sep}), buried surface area (ΔSASA), area-normalized energy ($100\Delta G_{\text{sep}}/\Delta\text{SASA}$), interface hydrogen bonds, buried unsatisfied polar groups, shape complementarity (SC), and packing quality (Packstat). Definitions, normalization, hit thresholds, and Mean/Best aggregation are given in Appendix E.5.

4.2 MAIN RESULTS

Direct full-atom generation on LNR. DC-FLOW combines high cyclization success with complementary sequence and conformational diversity under the common relaxation protocol. The HT configuration achieves the most favorable interface energy and highest sequence diversity on OOD receptors, supporting its applicability beyond training-related targets. The SS configuration achieves the highest structural diversity on both splits, although its sequence diversity is lower. The All/OOD/ID results in Appendix F confirm this topology-dependent balance: HT favors sequence exploration and interface energetics, while SS retains broader conformational variation. These results support direct generation as a source of diverse cyclic-complex hypotheses for downstream evaluation.

Table 1: **Direct full-atom cyclic peptide generation on LNR.** Evaluated across in-distribution (ID) and out-of-distribution (OOD) targets. Both DC-FLOW variants attain 100% cyclization success on ID targets. DC-FLOW(SS) achieves the highest closure fidelity and structural diversity (D_{struct}) across splits, whereas DC-FLOW(HT) provides superior binding energetics (Rosetta dG) and sequence exploration (D_{seq}), highlighting their complementary topological profiles.

Method	Cyc. ↑		Stability ↓		D_{seq} ↑		D_{struct} ↑		Rosetta dG ↓	
	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID
PepFlow	92.9	94.6	**-961.10**	-590.27	0.686	0.619	0.948	0.670	-24.78	-17.24
PepGLAD	97.5	87.4	-720.20	-579.45	0.895	**0.872**	0.932	0.817	-15.59	-13.50
DiffPepBuilder†	92.2	94.6	-939.25	-574.79	0.446	0.286	0.979	0.732	-20.66	-18.05
APM	97.3	**100.0**	-869.51	-497.36	0.173	0.120	0.938	0.754	-25.06	-16.50
Genie 3†	83.6	81.6	-726.29	-399.28	0.253	0.242	0.823	0.796	-14.07	-10.16
APCyc	95.4	73.9	-713.21	**-616.07**	0.867	0.831	0.849	0.705	-22.79	**-19.91**
DC-FLOW(HT)	97.5	**100.0**	-940.05	-604.28	**0.940**	0.845	0.934	0.459	**-28.37**	-19.14
DC-FLOW(SS)	**97.6**	**100.0**	-807.36	-509.06	0.668	0.205	**0.991**	**0.943**	-21.17	-15.39

† DiffPepBuilder and Genie 3 require the Rosetta side-chain packer.

Cofolding-based comparison. In the predictor-assisted workflow, DC-FLOW achieves the strongest Mean and Best confidence scores among the compared pipelines. Its strict-hit rate reaches 36.7%, compared with 33.3% for HFGuidedDesign and 29.2% for RFPeptides. The guidance and refinement configurations in Appendix F clarify this result: unrefined proposals remain substantially less confident, and the combination of generation-time guidance and subsequent optimization performs best. This supports the complementary role of predictor assistance in turning the generator's geometry-conditioned sequence proposals into confident predicted complexes.

Table 2: **Cofolding performance across 12 targets.** All methods use AFCyc optimization; HFGuidedDesign and DC-FLOW also use generation-time guidance. Mean summarizes all 384 candidates; Best averages the target-wise Top-1 candidates selected by composite score.

Method	Summary	ipTM ↑	Peptide pLDDT ↑	Interface iPAE ↓	Hotspot iPAE ↓	Composite S ↑	Strict hits (%) ↑
RFPeptides	Mean	0.666	61.21	0.334	0.234	1.944	29.2
	Best	0.877	84.90	0.170	0.115	2.556	
HFGuidedDesign	Mean	0.674	61.60	0.334	0.234	1.957	33.3
	Best	0.889	86.57	0.161	0.104	2.593	
DC-FLOW	Mean	**0.683**	**62.50**	**0.326**	**0.228**	**1.982**	**36.7**
	Best	**0.903**	**87.26**	**0.146**	**0.091**	**2.629**	

Note: Both iPAE metrics are normalized. Strict-hit rates use all 384 candidates (12 targets × 32 samples); full settings and evaluation details are in Table 10.

Figure 3 provides a physical assessment of these outputs. DC-FLOW achieves more favorable energy per buried interface area, more interface hydrogen bonds, and fewer buried unsatisfied polar groups in both the Mean and AF Top-1 summaries (Table 11). The improvements therefore extend beyond confidence scores to complementary measures of interface quality.

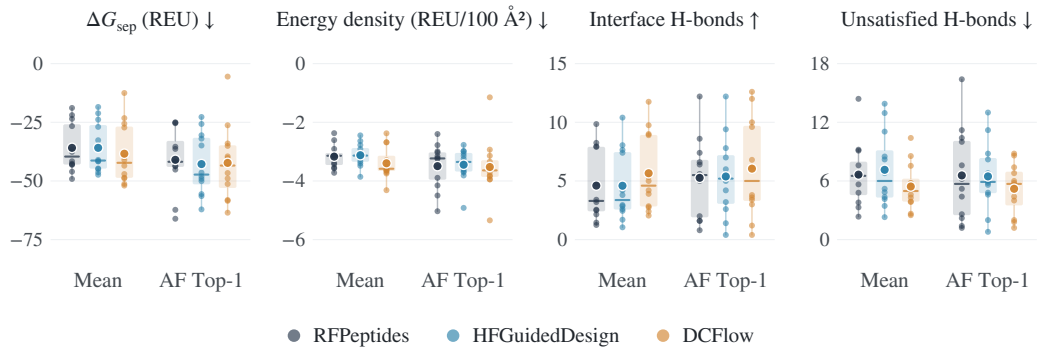

Figure 3: **Target-wise Rosetta interface quality across 12 targets.** Mean averages four sampled candidates per target; AF Top-1 uses the AF-selected candidate. Each candidate's scores are averaged over five models. Rounded bands show interquartile ranges, with median lines and 1.5-IQR whiskers; dots denote targets and larger circles denote across-target means. Energy density is $100\Delta G_{\text{sep}}/\Delta\text{SASA}$. Complete metrics are reported in Table 11.

4.3 Ablation Study and Analysis of Test-Time Scaling

We examine the three generative mechanisms through controlled ablations, with a separately trained model for each setting, retaining the receptor conditions, sampling-step budget, and candidate counts:

- **Sequence activation schedule** varies $\tau_s \in \{0, 0.2, 0.5\}$ while retaining $\tau_c = 0.8$, assessing how the timing of chemical assignment balances early backbone disambiguation against conformational exploration; $\tau_s = 0$ corresponds to fully synchronous generation from noise.
- **w/o coupling** freezes the backbone scaffold at B_{τ_s} for all $t \geq \tau_s$, mimicking conventional decoupled two-stage pipelines (i.e., fixed-scaffold inverse folding) to quantify the benefit of continuous mutual co-adaptation.
- **w/o ProteinMPNN prior** disables the external prior by setting $g = 0$ in equation 3, isolating the generator's native sequence trunk to evaluate how effectively the structural prior compensates for complex-data scarcity.

Continued coupling and the structural prior. Freezing the backbone causes the largest energetic deterioration: for OOD HT designs, interface energy worsens from -28.37 to -4.54 REU while sequence and structural diversity remain nearly unchanged. Backbone adaptability therefore contributes to chemical compatibility beyond maintaining candidate diversity. Removing the ProteinMPNN prior also weakens overall interface energies for both topologies, although its diversity effects vary (Appendix G). This is consistent with a structural prior that favors compatible residue choices without uniformly increasing the number of sequence modes.

Delayed activation balances exploration and compatibility. Simultaneous backbone–sequence activation weakens HT interface energetics, particularly on OOD targets, whereas the SS response is mixed. Earlier activation at $\tau_s = 0.2$ increases overall sequence and structural diversity relative to the default (Appendix G). The default $\tau_s = 0.5$ instead gives the highest overall cyclization success for both topologies and the best HT energy scores. Thus the schedule controls a trade-off between exploration and geometric compatibility; the benefit of delayed coupling cannot be reduced to maximizing diversity alone.

Cascaded search enriches the final candidate pool. We next evaluate the complete search policy with the generator and guidance/refinement backends unchanged. With 32 final candidates per target, search raises the strict-hit rate from 32.6% to 75.5% (Figure 4). Mean and Best both improve, with larger changes in Mean, showing that search enriches the returned candidate population as well as improving the strongest designs. The additional results in Appendix G show more favorable Rosetta energies and packing, alongside weaker hydrogen-bond metrics. These findings support the cascade as an integrated policy for allocating exploration and refinement, with its clearest benefit in the yield of confident predicted interfaces.

Table 3: **Ablations on LNR.** The default $\tau_s = 0.5$ achieves the highest overall cyclization success within each topology and the best HT Stability and Rosetta dG scores. Freezing backbone motion substantially worsens Stability for both topologies, while removing the MPNN prior weakens interface scores and removing the delay produces topology-dependent effects. Earlier activation at $\tau_s = 0.2$ improves overall sequence and structure diversity, indicating a trade-off between diversity and interface scores. OOD/ID correspond to 37/56 receptors.

Inference setting	Cyc. ↑		Stability ↓		D_{seq} ↑		D_{struct} ↑		Rosetta dG ↓	
	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID
w/o coupling (HT)	97.1	99.5	112.60	-255.79	0.937	0.824	0.933	0.439	-4.54	-12.92
w/o coupling (SS)	97.1	99.6	-167.87	-249.78	0.639	0.213	0.989	0.936	-7.97	-12.95
w/o MPNN prior (HT)	97.3	99.8	-914.12	-583.86	0.933	0.836	0.948	0.427	-27.56	-17.36
w/o MPNN prior (SS)	97.1	99.7	-808.71	-503.13	0.639	0.271	0.948	0.936	-18.88	-16.37
w/o delay, $\tau_s = 0$ (HT)	97.1	99.3	-905.72	-591.57	0.934	0.829	0.940	0.453	-20.40	-17.17
w/o delay, $\tau_s = 0$ (SS)	97.2	99.7	-838.82	-537.96	0.668	0.219	0.982	0.924	-19.61	-16.92
$\tau_s = 0.2$ (HT)	**97.6**	99.8	-913.84	-592.83	**0.940**	**0.903**	0.934	0.590	-23.38	-16.30
$\tau_s = 0.2$ (SS)	97.5	99.9	-812.85	-513.08	0.672	0.205	**1.000**	**0.943**	-19.59	-16.36
$\tau_s = 0.5$ (HT)	97.5	**100.0**	**-940.05**	**-604.28**	**0.940**	0.845	0.934	0.459	**-28.37**	**-19.14**
$\tau_s = 0.5$ (SS)	**97.6**	**100.0**	-807.36	-509.06	0.668	0.205	0.991	**0.943**	-21.17	-15.39

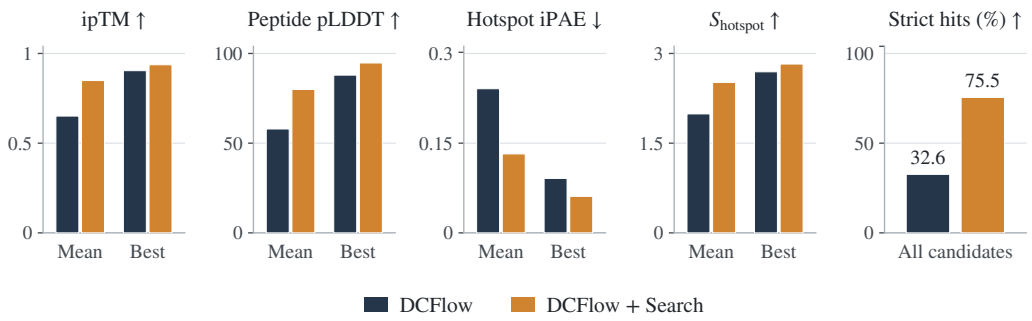

Figure 4: **Effect of inference-time search on cofolding.** Both DC-FLOW configurations use guidance and AfCycDesign optimization. Search improves Mean and Best on all four confidence metrics and increases the strict-hit rate from 32.6% to 75.5%. Best averages target-wise Top-1 candidates selected by S_{hotspot}; strict hits are measured over all 384 candidates. Detailed results appear in Table 13.

5 Conclusion and Future Work

We presented DC-FLOW, a target-conditioned cyclic peptide design framework combining delayed coupling, a frozen ProteinMPNN residual prior, and cascaded test-time search. Its staged trajectory introduces sequence and side-chain information after initial backbone formation while preserving continuous co-adaptation, supporting both direct full-atom generation and predictor-assisted design. Experiments show a favorable balance of cyclization success, diversity, and interface energetics, while integrated search enriches the final candidate pool with confident predicted complexes.

Limitations and Future Work. Despite geometric filtering, CPSea-derived supervision remains largely synthetic and may bias learned conformational and interface preferences. Reliance on AF2-Multimer confidence for guidance and evaluation may also favor predictor-specific solutions without corresponding improvements in experimental performance. Moreover, the current cofolding pipeline does not enforce terminal disulfide closure, limiting the use of DC-FLOW's SS design capability in this workflow. We plan to pair generated complex hypotheses with wet-lab measurements in an iterative design–experiment loop, using experimental feedback to guide dataset curation and model updates. This aims to improve the reliability of direct full-atom design and reduce dependence on costly cofolding refinement, accelerating the workflow as experimental evidence accumulates.

AI USE STATEMENT

In this work, generative AI tools were used to assist with language polishing, phrasing improvements, and the LaTeX formatting and typesetting of mathematical notations and proofs. Additionally, generative AI was utilized to assist with code organization and formatting; all code changes were manually inspected, executed, and verified by the authors. All scientific ideation, methodological designs, theoretical formulations, and mathematical proofs (Appendix B) were derived and completed by the authors. Model architectures, experimental implementations, and data analyses were likewise executed exclusively by the authors. The authors manually inspected and verified all AI-formatted text revisions against their original ones, and take full responsibility for the content, technical claims, code, and findings presented in this paper.

ETHICS STATEMENT

This work introduces DC-FLOW, a computational generative framework for target-conditioned cyclic peptide binder design aimed at facilitating therapeutic peptide discovery against challenging protein-protein interaction (PPI) interfaces. The research was conducted entirely *in silico* and relies exclusively on publicly accessible, anonymized structural repositories, including the Protein Data Bank (PDB), the CPSea database, and the open-source LNR benchmark. No human participants, personally identifiable health information, or vertebrate animals were involved in this study; therefore, no Institutional Review Board (IRB) or animal care approval was required. The methodology is designed as a foundational tool for computational biophysics and molecular design and is not intended for, nor tested on, the generation of biological toxins, dual-use hazardous agents, or pathogen engineering. The authors foresee no direct negative societal, discriminatory, or biosecurity implications from this work.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of all findings, models, and algorithms in this paper, we have provided:

- **Code and Checkpoints:** An anonymized codebase implementing DC-FLOW—including delayed-coupling flow matching, the ProteinMPNN prior integration, and the cascaded search pipeline—is provided in the supplementary material.
- **Architectures and Hyperparameters:** Detailed network structures (geometric trunk, IPA modules), loss weights, and training schedules are fully specified in Appendix C and Appendix E.
- **Data Curation and Benchmarks:** Complete stereochemical filtering scripts for the CPSea corpus (amide/disulfide geometry audits) and the MMseqs2-based sequence-identity splits for the LNR benchmark are documented in Appendix D.
- **Evaluation and Seeds:** All sampling trajectories use fixed 200-step schedules. Evaluation protocols for AlphaFold-Multimer-v3 and Rosetta FastRelax are detailed in Appendix E.
- **Theoretical Proofs:** Rigorous proofs and regularity assumptions for all theoretical propositions (delayed activation, coupling benefit, and search coverage) are provided in Appendix B.

REFERENCES

Rebecca F Alford, Andrew Leaver-Fay, Jeliazko R Jeliazkov, Matthew J O'Meara, Frank P DiMaio, Hahnbeom Park, Maxim V Shapovalov, P Douglas Renfrew, Vikram K Mulligan, Kalli Kappel, et al. The Rosetta all-atom energy function for macromolecular modeling and design. *Journal of chemical theory and computation*, 13(6):3031, 2017.

Donald S Berkholz, Camden M Driggers, Maxim V Shapovalov, Roland L Dunbrack Jr, and P Andrew Karplus. Nonplanar peptide bonds in proteins are common and conserved but not biased toward active sites. *Proceedings of the National Academy of Sciences*, 109(2):449–453, 2012.

Ruizhe Chen, Dongyu Xue, Xiangxin Zhou, Zaixiang Zheng, Xiangxiang Zeng, and Quanquan Gu. An all-atom generative model for designing protein complexes. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 9516–9545. PMLR, 2025.

Tim Clackson and James A Wells. A hot spot of binding energy in a hormone-receptor interface. *Science*, 267(5196):383–386, 1995.

Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J Ragotte, Lukas F Milles, Basile IM Wicky, Alexis Courbet, Rob J de Haas, Neville Bethel, et al. Robust deep learning–based protein sequence design using ProteinMPNN. *Science*, 378(6615):49–56, 2022.

Kieran Didi, Zuobai Zhang, Guoqing Zhou, Danny Reidenbach, Zhonglin Cao, Sooyoung Cha, Tomas Geffner, Christian Dallago, Jian Tang, Michael Bronstein, et al. Scaling atomistic protein binder design with generative pretraining and test-time compute. In *International Conference on Learning Representations*, volume 2026, pp. 61056–61109, 2026.

Alan A Dombkowski. Disulfide by design™: a computational method for the rational design of disulfide bonds in proteins. *Bioinformatics*, 19(14):1852–1853, 2003.

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

Roland L Dunbrack Jr and Martin Karplus. Backbone-dependent rotamer library for proteins application to side-chain prediction. *Journal of molecular biology*, 230(2):543–574, 1993.

Pengyuan Fang, Wing-Ka Pang, Shouhu Xuan, Wai-Lun Chan, and Ken Cham-Fai Leung. Recent advances in peptide macrocyclization strategies. *Chemical Society Reviews*, 53(24):11725–11771, 2024.

Heather C Hayes, Louis YP Luk, and Yu-Hsuan Tsai. Approaches for peptide and protein cyclisation. *Organic & biomolecular chemistry*, 19(18):3983–4001, 2021.

Haomeng Hu, Renjie Zhu, Ning Zhu, Chengyun Zhang, Tianfeng Shang, Chongyang Li, Jingjing Guo, Xudong Wang, and Hongliang Duan. HFGuidedDesign: de novo design of cyclic peptide binders via structure-guided discrete diffusion. *Chemical Science*, 2026.

Xinjian Ji, Alexander L Nielsen, and Christian Heinis. Cyclic peptides for drug development. *Angewandte Chemie International Edition*, 63(3):e202308251, 2024.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *nature*, 596(7873):583–589, 2021.

Bradley A Katz and Anthony Kossiakoff. The crystallographically determined structures of atypical strained disulfides engineered into subtilisin. *Journal of Biological Chemistry*, 261(33):15480–15485, 1986.

Horst Kessler. Conformation and biological activity of cyclic peptides. *Angewandte Chemie International Edition in English*, 21(7):512–523, 1982.

Xiangzhe Kong, Wenbing Huang, and Yang Liu. End-to-end full-atom antibody design. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 17409–17429. PMLR, 2023.

Xiangzhe Kong, Yinjun Jia, Wenbing Huang, and Yang Liu. Full-atom peptide design with geometric latent diffusion. *Advances in Neural Information Processing Systems*, 37:74808–74839, 2024.

Jiahan Li, Chaoran Cheng, Zuofan Wu, Ruihan Guo, Shitong Luo, Zhizhou Ren, Jian Peng, and Jianzhu Ma. Full-atom peptide design based on multi-modal flow matching. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 27615–27640. PMLR, 2024.

Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of machine learning research*, 18(185):1–52, 2018.

Yeqing Lin, Minji Lee, Aakarsh Vermani, Ellena Jiang, Sebastiaan De Cooman, Matej Špetko, and Mohammed AlQuraishi. Fast and ultra-capable protein design: Advancing the frontier through atomistic se (3)-equivariance with genie 3. *bioRxiv*, 2026.

Stephanie M Linker, Christian Schellhaas, Benjamin Ries, Hans-Jörg Roth, Marianne Fouché, Stephane Rodde, and Sereina Riniker. Polar/apolar interfaces modulate the conformational behavior of cyclic peptides with impact on their passive membrane permeability. *RSC advances*, 12(10): 5782–5796, 2022.

Mengdi Liu, Xiaoxue Cheng, Zhangyang Gao, Hong Chang, Cheng Tan, Shiguang Shan, and Xilin Chen. Protinvtree: Deliberate protein inverse folding with reward-guided tree search. *Advances in Neural Information Processing Systems*, 38:28337–28360, 2026.

Lucia Lombardi, Valentina Del Genio, Fernando Albericio, and Daryl R Williams. Advances in peptidomimetics for next-generation therapeutics: strategies, modifications, and applications. *Chemical Reviews*, 125(15):7099–7166, 2025.

Simon C Lovell, J Michael Word, Jane S Richardson, and David C Richardson. The penultimate rotamer library. *Proteins: Structure, Function, and Bioinformatics*, 40(3):389–408, 2000.

Lundquist and Jeffrey C Pelletier. A new tri-orthogonal strategy for peptide cyclization. *Organic Letters*, 4(19):3219–3221, 2002.

Malcolm W MacArthur and Janet M Thornton. Influence of proline residues on protein conformation. *Journal of molecular biology*, 218(2):397–412, 1991.

Alpeshkumar K Malde, Timothy A Hill, Abishek Iyer, and David P Fairlie. Crystal structures of protein-bound cyclic peptides. *Chemical Reviews*, 119(17):9861–9914, 2019.

Thi-Sau Migone, Jun Zhang, Xia Luo, Li Zhuang, Cecil Chen, Bugen Hu, June S Hong, James W Perry, Su-Fang Chen, Joe XH Zhou, et al. Tl1a is a tnf-like ligand for dr3 and tr6/dcr3 and functions as a t cell costimulator. *Immunity*, 16(3):479–492, 2002.

Stephen A Rettie, Katelyn V Campbell, Asim K Bera, Alex Kang, Simon Kozlov, Yensi Flores Bueso, Joshmyn De La Cruz, Maggie Ahlrichs, Suna Cheng, Stacey R Gerben, et al. Cyclic peptide structure prediction and design using alphafold2. *Nature Communications*, 16(1):4730, 2025a.

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, et al. Accurate de novo design of high-affinity protein-binding macrocycles using deep learning. *Nature Chemical Biology*, 21(12): 1948–1956, 2025b.

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion models. *arXiv preprint arXiv:2501.06848*, 2025.

Martin Steinegger and Johannes Söding. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature biotechnology*, 35(11):1026–1028, 2017.

Janet M Thornton. Disulphide bridges in globular proteins. *Journal of molecular biology*, 151(2): 261–287, 1981.

Tomer Tsaban, Julia K Varga, Orly Avraham, Ziv Ben-Aharon, Alisa Khramushin, and Ora Schueler-Furman. Harnessing protein folding neural networks for peptide–protein docking. *Nature communications*, 13(1):176, 2022.

Alexander A Vinogradov, Yizhen Yin, and Hiroaki Suga. Macrocyclic peptides as drug candidates: recent progress and remaining challenges. *Journal of the American Chemical Society*, 141(10): 4167–4181, 2019.

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.

Shuzhe Wang, Gerhard Konig, Hans-Jorg Roth, Marianne Fouché, Stephane Rodde, and Sereina Riniker. Effect of flexibility, lipophilicity, and the location of polar residues on the passive membrane permeability of a series of cyclic decapeptides. *Journal of Medicinal Chemistry*, 64 (17):12761–12773, 2021.

Yitian Wang, Fanmeng Wang, Angxiao Yue, Wentao Guo, Yaning Cui, and Hongteng Xu. MuCO: Generative peptide cyclization empowered by multi-stage conformation optimization. In *International Conference on Machine Learning (ICML)*, 2026. Accepted regular paper; arXiv:2602.11189.

Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100, 2023.

William J Wedemeyer, Ervin Welker, and Harold A Scheraga. Proline cis- trans isomerization and protein folding. *Biochemistry*, 41(50):14637–14644, 2002.

Manfred S Weiss, Andreas Jabs, and Rolf Hilgenfeld. Peptide bonds revisited. *Nature structural biology*, 5(8):676–676, 1998.

Christopher J White and Andrei K Yudin. Contemporary strategies for peptide macrocyclization. *Nature chemistry*, 3(7):509–524, 2011.

Ziyi Yang, Hanyuan Xie, Yinjun Jia, Xiangzhe Kong, Jiqing Zheng, Ziting Zhang, Yang Liu, Lei Liu, and Yanyan Lan. Cpsea: large-scale cyclic peptide-protein complex dataset for machine learning in cyclic peptide design. *Advances in Neural Information Processing Systems*, 38, 2026.

Jian Yao, Victoria A Feher, B Fabiola Espejo, Marine T Reymond, Peter E Wright, and H Jane Dyson. Stabilization of a type vi turn in a family of linear peptides in water solution. *Journal of molecular biology*, 243(4):736–753, 1994.

Chenyang Zhan, Yury Patskovsky, Qingrong Yan, Zhenhong Li, Udupi Ramagopal, Huiyong Cheng, Michael Brenowitz, Xiao Hui, Stanley G Nathenson, and Steven C Almo. Decoy strategies: the structure of tl1a: Dcr3 complex. *Structure*, 19(2):162–171, 2011.

Chenhao Zhang, Chengyun Zhang, Tianfeng Shang, Ning Zhu, Xinyi Wu, and Hongliang Duan. HighFold: accurately predicting structures of cyclic peptides and complexes with head-to-tail and disulfide bridge constraints. *Briefings in bioinformatics*, 25(3):bbae215, 2024.

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, 2026.

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.

Adam Zwolak, Szeman Ruby Chan, Paul Harvilla, Sally Mahady, Anthony A Armstrong, Leopoldo Luistro, Ninkka Tamot, Douglas Yamada, Mehabaw Derebe, Steven Pomerantz, et al. A stable, engineered tl1a ligand co-stimulates t cells via specific binding to dr3. *Scientific Reports*, 12(1): 20538, 2022.

APPENDIX CONTENTS

A Notation and Modeling Assumptions **15**

B Theory of Delayed-Coupling Flow Matching **16**

B.1 Delayed activation: structural information and sequence risk 16

B.2 Continued coupling: sequence-conditioned backbone prediction 18

B.3 ProteinMPNN prior: transfer mechanism and prediction benefit 18

B.4 Cascaded search: structural selection, sequence coverage, and refinement 20

C DC-Flow Formulation and Architecture **26**

C.1 State Representation and Modality-Specific Paths 27

C.2 Receptor and Cyclization Conditioning 27

C.3 Delayed Visibility and the Shared Geometric Trunk 28

C.4 Dynamic ProteinMPNN Residual Prior 28

C.5 Supervision, Numerical Updates, and Reconstruction 29

C.6 Interfaces to Guided Cofolding and Cascaded Search 30

D CPSea Quality Filtering **31**

D.1 Overview 31

D.2 Sequence Length Constraints 31

D.3 Head-to-Tail Amide Bond Geometries 31

D.4 Disulfide Bond Geometry and Quality Audit 32

D.5 Exclusion of Side-Chain Isopeptide Cyclizations 33

D.6 Dataset Statistics and Filtering Outcomes 33

D.7 Ablation Study 34

E Detailed Implementation and Evaluation Metrics **35**

E.1 Model Training and Sampling 35

E.2 Benchmark Conditions and Candidate Budgets 36

E.3 Physical Relaxation and Topology Handling 36

E.4 Direct Full-Atom Generation Metrics 37

E.5 Cofolding Evaluation and Reported Metrics 38

E.6 Generation-Time Guidance and AfCycDesign Refinement 39

F Additional Experiment Results **41**

G Additional Ablation Results **44**

H Challenging TL1A-Trimer Case **46**

I Visualization **47**

A Notation and Modeling Assumptions

Table 4 summarizes the principal symbols used in the main text and theoretical analysis. Additional architectural variables are defined locally in Appendix C. Time-indexed variables denote intermediate generative states, with $t = 1$ indicating the endpoint. The receptor remains fixed throughout each trajectory, while peptide generation is conditioned on the specified length and cyclization mode.

Table 4: Principal notation for generation, prior integration, and cascaded search.

Symbol	Meaning
R	Receptor structure and its representation
L, Γ, ξ	Peptide length, cyclization specification, and joint conditions $\xi = (L, \Gamma)$
$\mathcal{A}$	Amino-acid alphabet
$t \in [0, 1]$	Global generation time, distinct from the local modality clocks
B_t	Peptide backbone frames and carbonyl-orientation state at time t
S_t, S_1	Continuous sequence carrier during generation and decoded endpoint sequence $S_1 \in \mathcal{A}^L$
C_t	Side-chain torsion state; full heavy-atom coordinates are reconstructed from Y_1
Y_t	Joint peptide state (B_t, S_t, C_t)
p_0, p_θ	Initial sampling distribution and learned conditional generative distribution
$\tau_\mathrm{b}, \tau_\mathrm{s}, \tau_\mathrm{c}$	Backbone, sequence, and side-chain activation times; defaults 0, 0.5, and 0.8
$m, s_m(t)$	Modality index $m \in \{\mathrm{b}, \mathrm{s}, \mathrm{c}\}$ and its local progress clock
θ, ϕ	Trainable generator parameters and frozen ProteinMPNN parameters
F_ϕ	Frozen ProteinMPNN structural sequence prior
$\ell^B_{i,t}, \ell^P_{i,t}, \widetilde{\ell}^P_{i,t}$	Native sequence logits, prior log-probabilities, and centered prior residuals at residue i
$\ell_{i,t}, \pi_{i,t}$	Fused sequence logits and the resulting categorical distribution
π^B, q^P	Native and prior distributions over admissible amino acids; position and time indices suppressed
g	Trainable scalar gate controlling the prior contribution; initialized to 0.1
$\ell_B, \ell_S, \ell_C, \ell_\mathrm{cyc}$	Backbone, sequence, side-chain, and cyclic-closure loss terms
$\lambda_B, \lambda_S, \lambda_C$	Nonnegative weights of the modality-specific losses
K, K_0	Number of initial backbone hypotheses and retained roots, with $K_0 < K$
M, N	Continuations per retained root and total proposals $N = K_0 M$
N_0	Number of candidates retained after refinement warm-up, with $N_0 < N$
Q	Interface quality score for probabilistic early screening, defined in equation 5
$Z = (R, \xi)$	Receptor and design conditions grouped for the theoretical analysis
$\overline{B}_t$	Reference backbone observation used in the information-theoretic analysis
H_t	Complete sampler checkpoint, including active states and memory needed for continuation

The temporal factorization in equation 1 specifies a computational sampling schedule, not a physical backbone-first binding pathway or independence among modalities. The information decomposition in equation 10 is derived under the reference-path assumptions in Appendix B.1. More generally, the theoretical results state conditions under which delayed activation, continued coupling, prior transfer, and structural diversification can be beneficial; their practical effects are evaluated empirically. In cascaded refinement, Q is used for early candidate selection, whereas early stopping uses the composite score reported by AfCycDesign.

B THEORY OF DELAYED-COUPLING FLOW MATCHING

This appendix analyzes the four mechanisms used in Sections 3.2–3.3: delayed activation, continued coupling, residual-prior transfer, and cascaded inference-time allocation. Each result specifies its probability law, assumptions, and implication for the implemented method. We write $Z = (\mathcal{R}, \xi)$ for the receptor and design conditions and use natural logarithms. The finite sequence alphabet ensures that the sequence entropies below are finite.

The first two parts concern a joint endpoint distribution $p^\star$ and a specified interpolation or training coupling. The third concerns prediction risk on a fixed distribution of intermediate contexts. The fourth concerns the learned sampler and its conditional continuation distributions. Transferring a statement between these settings requires the stated additional assumptions.

Staged generation as composition of conditional kernels. Equation 1 gives a compact stage-wise decomposition in terms of the modality states exposed at the activation boundaries. For the implemented sampler, continuation can additionally depend on persistent source variables and other sampling memory; we therefore use the complete checkpoint H_t below to state the corresponding transition law exactly. Let H_t be the complete sampler checkpoint at time t. It contains the visible active state and any source variables or other memory needed to continue sampling. This distinction matters when updates depend on a fixed source such as x_0. Write $\mu_0(\mathrm{d}h_0 \mid Z)$ for the initialized checkpoint law and $P^\theta_{a:b}$ for the conditional transition kernel between times a and b, including intermediate categorical draws. If $D(h_1) = Y_1$ denotes terminal decoding, then for any measurable output set E,

$$p_\theta(Y_1 \in E \mid Z) = \int \mu_0(\mathrm{d}h_0 \mid Z)\, P^\theta_{0:\tau_s}(\mathrm{d}h_S \mid h_0, Z) \times P^\theta_{\tau_s:\tau_c}(\mathrm{d}h_C \mid h_S, Z)\, P^\theta_{\tau_c:1}(\mathrm{d}h_1 \mid h_C, Z)\, \mathbf{1}\{D(h_1) \in E\}. \tag{7}$$

This is the marginal law induced by the three stages. A product of the three kernels describes their joint path law before integrating out intermediate states. Replacing complete checkpoints by $\mathcal{B}_{\tau_s}$ or $(\mathcal{B}_{\tau_c}, \mathcal{S}_{\tau_c})$ alone additionally requires that those variables are sufficient for the future sampler. Temporal staging does not assume that the final backbone, sequence, and side chains are independent.

B.1 DELAYED ACTIVATION: STRUCTURAL INFORMATION AND SEQUENCE RISK

B.1.1 ADDITIONAL INFORMATION BEYOND RECEPTOR CONDITIONING

Let

$$J_B = I(\mathcal{S}_1; \mathcal{B}_1 \mid Z), \qquad D(t) = I(\mathcal{S}_1; \mathcal{B}_1 \mid \overline{\mathcal{B}}_t, Z), \qquad G(t) = I(\mathcal{S}_1; \overline{\mathcal{B}}_t \mid Z). \tag{8}$$

The bar distinguishes a reference backbone observation from the complete learned sampler state. In particular, the analysis does not assume that the receptor is uninformative. The residual quantity satisfies

$$J_B = \mathbb{E}[D_{\mathrm{KL}}(p^\star(\mathcal{S}_1 \mid \mathcal{B}_1, Z) \,\|\, p^\star(\mathcal{S}_1 \mid Z))]\,. \tag{9}$$

Thus $J_B > 0$ precisely when backbone conditioning changes sequence preferences on a set of positive probability, beyond the information already contained in Z.

Assumption D (ordered reference observations). For every t, assume $\mathcal{S}_1 \perp \overline{\mathcal{B}}_t \mid (\mathcal{B}_1, Z)$. For $t_1 < t_2$, assume there is a stochastic kernel that maps the observation at t_2 to one with the conditional law of the observation at t_1, without access to $\mathcal{B}_1$ or $\mathcal{S}_1$. Equivalently, the two observation laws admit a coupling with conditional Markov chain

$$\mathcal{S}_1 \longrightarrow \mathcal{B}_1 \longrightarrow \overline{\mathcal{B}}_{t_2} \longrightarrow \overline{\mathcal{B}}_{t_1} \quad \text{given } Z.$$

Assume the observation at zero is independent of the endpoints given Z, and $\overline{\mathcal{B}}_1 = \mathcal{B}_1$.

Proposition B.1 (Information ordering and optimal sequence loss). Under Assumption D, $G(t)$ is non-decreasing and $D(t)$ is non-increasing, with

$$J_B = G(t) + D(t), \qquad D(0) = J_B, \quad D(1) = 0. \tag{10}$$

Define the optimal sequence log-loss using the intermediate observation by

$$\mathcal{L}^\star_{\mathrm{seq}}(t) = \inf_q \mathbb{E}[-\log q(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z)]. \tag{11}$$

Then

$$\begin{aligned} \mathcal{L}^{\star}_{\mathrm{seq}}(t) &= H(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z), \\ \mathcal{L}^{\star}_{\mathrm{seq}}(t) - \mathcal{L}^{\star}_{\mathrm{seq}}(1) &= D(t). \end{aligned} \tag{12}$$

For $t_1 < t_2$, the loss reduction equals $G(t_2) - G(t_1)$ and is strict exactly when this increment is positive.

Proof. Conditional independence and the chain rule give

$$\begin{aligned} I(\mathcal{S}_1; \mathcal{B}_1, \overline{\mathcal{B}}_t \mid Z) &= I(\mathcal{S}_1; \mathcal{B}_1 \mid Z) \\ &= I(\mathcal{S}_1; \overline{\mathcal{B}}_t \mid Z) + I(\mathcal{S}_1; \mathcal{B}_1 \mid \overline{\mathcal{B}}_t, Z). \end{aligned} \tag{13}$$

This proves the decomposition. Conditional data processing under the degradation coupling gives $G(t_1) \le G(t_2)$. The endpoint identities follow from the endpoint assumptions.

For any conditional predictor q, its log-loss decomposes as

$$\begin{aligned} \mathbb{E}[-\log q(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z)] &= H(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z) \\ &\quad + \mathbb{E} D_{\mathrm{KL}}\big(p^{\star}(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z) \,\|\, q(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z)\big). \end{aligned} \tag{14}$$

Minimizing the nonnegative KL term proves the first risk identity. Furthermore, $H(\mathcal{S}_1 \mid \overline{\mathcal{B}}_t, Z) = H(\mathcal{S}_1 \mid Z) - G(t)$, and the analogous endpoint expression subtracts J_B. Their difference is $D(t)$. The strictness statement follows by subtracting risks at the two times. □

B.1.2 A REFERENCE PATH SATISFYING THE ASSUMPTION

Consider a Euclidean backbone representation with

$$\overline{\mathcal{B}}_t = (1-t)\mathcal{B}_0 + t\mathcal{B}_1, \qquad \mathcal{B}_0 \sim \mathcal{N}(0, I), \qquad \mathcal{B}_0 \perp (\mathcal{B}_1, \mathcal{S}_1, Z). \tag{15}$$

For $t > 0$, the invertible rescaling $V_t = \overline{\mathcal{B}}_t / t = \mathcal{B}_1 + \sigma_t \mathcal{B}_0$ has $\sigma_t = (1-t)/t$. For $0 < t_1 < t_2$, generate

$$V_2 = \mathcal{B}_1 + \sigma_{t_2} \varepsilon_2, \qquad V_1 = V_2 + \sqrt{\sigma_{t_1}^2 - \sigma_{t_2}^2}\, \varepsilon', \tag{16}$$

where $\varepsilon_2, \varepsilon'$ are independent standard Gaussian variables independent of $(\mathcal{B}_1, \mathcal{S}_1, Z)$. The conditional marginal laws of V_1, V_2 equal those of V_{t_1}, V_{t_2}, and V_1 is degraded from V_2. Consequently, this path satisfies Assumption D. Under this constructed coupling,

$$D(t_1) - D(t_2) = I(\mathcal{S}_1; V_2 \mid V_1, Z) \ge 0. \tag{17}$$

The actual shared-source interpolation need not be Markov across time; the argument compares its observation channels through an equivalent coupling.

Implication for delayed activation. The quantity $D(t)$ is an exact excess Bayes log-loss from using an intermediate backbone instead of the endpoint backbone. If $G(t)$ grows before τ_{s}, later sequence assignment has a lower information-limited risk. For finite predictors, write their risk as $\widehat{\mathcal{L}}(t) = \mathcal{L}^{\star}_{\mathrm{seq}}(t) + e(t)$, where $e(t) \ge 0$ is prediction excess risk. Then

$$\widehat{\mathcal{L}}(t_1) - \widehat{\mathcal{L}}(t_2) = G(t_2) - G(t_1) + e(t_1) - e(t_2). \tag{18}$$

Thus the gain persists whenever the information increment exceeds any increase in excess risk. For a factorized residue head, the same argument applies to each residue and the losses sum; it need not model the full sequence joint distribution.

Scope. The Gaussian construction applies to Euclidean coordinates, or whitened coordinates on the relevant translation subspace. It does not prove information ordering for rotations, the complete nonlinear learned trajectory, or a guided trajectory that already uses sequence feedback. Those cases require their own version of Assumption D. Nor does a fixed τ_{s} guarantee a confident or physically valid backbone. If $J_B = 0$, this information-based benefit vanishes. An unrestricted synchronous network can in principle learn to ignore early uninformative variables; the result supports explicit timing as an inductive bias, not universal dominance over all synchronous architectures.

B.2 Continued coupling: sequence-conditioned backbone prediction

Fix an active time $t \geq \tau_s$ and a joint law of $(U_t^B, \mathcal{B}_t, \mathcal{S}_t, Z)$. Here U_t^B is the backbone regression target in a common Euclidean or tangent-space representation. Assume $\mathbb{E}\|U_t^B\|^2 < \infty$. Let

$$\begin{aligned} m_B &= \mathbb{E}[U_t^B \mid \mathcal{B}_t, Z], & \mathcal{R}_B &= \inf_f \mathbb{E}\|U_t^B - f(\mathcal{B}_t, Z)\|^2, \\ m_{BS} &= \mathbb{E}[U_t^B \mid \mathcal{B}_t, \mathcal{S}_t, Z], & \mathcal{R}_{BS} &= \inf_f \mathbb{E}\|U_t^B - f(\mathcal{B}_t, \mathcal{S}_t, Z)\|^2. \end{aligned} \tag{19}$$

The infima range over square-integrable measurable predictors under the same target law.

Proposition B.2 (Available benefit from continued coupling). The optimal risks satisfy

$$\mathcal{R}_B - \mathcal{R}_{BS} = \mathbb{E}\|m_{BS} - m_B\|^2 \triangleq \Delta_t^{B \leftarrow S} \geq 0. \tag{20}$$

The decrease is strict if and only if $\mathbb{P}(m_{BS} \neq m_B) > 0$.

Proof. Conditional expectations attain the two optimal squared risks. Decompose $U_t^B - m_B = (U_t^B - m_{BS}) + (m_{BS} - m_B)$. The second summand is measurable with respect to $\sigma(\mathcal{B}_t, \mathcal{S}_t, Z)$, whereas the first has zero conditional mean. Their expected inner product is zero. Expanding the squared norm yields equation 20; a squared norm has strictly positive expectation exactly under the stated condition. □

When is the condition substantive? For the linear reference path and $t < 1$, the velocity target is $U_t^B = (\mathcal{B}_1 - \mathcal{B}_t)/(1-t)$, giving

$$\Delta_t^{B \leftarrow S} = \frac{1}{(1-t)^2} \mathbb{E}\left\| \mathbb{E}[\mathcal{B}_1 \mid \mathcal{B}_t, \mathcal{S}_t, Z] - \mathbb{E}[\mathcal{B}_1 \mid \mathcal{B}_t, Z] \right\|^2. \tag{21}$$

Thus strict gain requires sequence information to change the preferred endpoint geometry, not merely to correlate with it. For example, if

$$U_t^B = m_B + A_t \zeta_t + \varepsilon_t, \qquad \mathbb{E}[\zeta_t \mid \mathcal{B}_t, Z] = 0, \qquad \mathbb{E}[\varepsilon_t \mid \mathcal{B}_t, \mathcal{S}_t, Z] = 0,$$

where A_t is measurable given $(\mathcal{B}_t, Z)$ and $\zeta_t = \zeta(\mathcal{S}_t, \mathcal{B}_t, Z)$, then $\Delta_t^{B \leftarrow S} = \mathbb{E}\|A_t \zeta_t\|^2$. Nonzero sequence-dependent geometric corrections therefore give a strict benefit. Positive terminal mutual information alone does not establish this conditional-mean condition.

Finite models and late side-chain information. If fitted predictors have excess risks $e_B, e_{BS} \geq 0$ above their respective Bayes risks, then

$$\widehat{\mathcal{R}}_B - \widehat{\mathcal{R}}_{BS} = \Delta_t^{B \leftarrow S} + e_B - e_{BS}. \tag{22}$$

A fitted coupled predictor is better when its additional approximation or estimation error is smaller than the available conditioning gain. At $t \geq \tau_c$, the same orthogonal decomposition applies to the nested information sets $(\mathcal{B}_t, \mathcal{S}_t, Z)$ and $(\mathcal{B}_t, \mathcal{S}_t, C_t, Z)$.

Implication for the architecture. Freezing the backbone after τ_s prevents it from realizing further sequence-dependent corrections. Continued evolution retains that possibility whenever $\Delta_t^{B \leftarrow S} > 0$. The proposition compares regression information under a fixed law; it does not by itself prove lower terminal docking error or a lower force-field energy for the trained sampler. Together with the previous part, it supports delaying access to immature modalities while preserving their subsequent influence.

B.3 ProteinMPNN prior: transfer mechanism and prediction benefit

B.3.1 Residual fusion and admissible support

Fix an active prediction context, omitting residue and time subscripts. The prior below is supplied by the frozen ProteinMPNN model used in the Method (Dauparas et al., 2022). Let $\mathcal{A}_i \subseteq \mathcal{A}$ be the allowed amino acids at residue i. On this support, let $p(a) = \pi^B(a) > 0$ be the native categorical distribution and $q(a) = q^P(a) > 0$ the normalized frozen-prior distribution. Restricting and renormalizing a full-alphabet prior to $\mathcal{A}_i$ does not change the fused law below. Let $z(a) = \log q(a)$ and $\widetilde{\ell}_a^P = \ell_a^P - |\mathcal{A}|^{-1} \sum_{a' \in \mathcal{A}} \ell_{a'}^P$.

Proposition B.3 (Exact transfer of prior log odds). For every finite scalar gate g,

$$p_g(a) = \frac{p(a)q(a)^g}{\sum_{b\in\mathcal{A}_i} p(b)q(b)^g}, \qquad \log\frac{p_g(a)}{p_g(b)} = \log\frac{p(a)}{p(b)} + g\log\frac{q(a)}{q(b)}. \tag{23}$$

Moreover, $p_0 = p$. At a fixed context with unrestricted base logits, every strictly positive distribution on $\mathcal{A}_i$ can be represented for any fixed finite g.

Proof. In the softmax of $\ell^B + g\widetilde{\ell}^P$, centering contributes the same scalar to every category and cancels between numerator and denominator. Writing each uncentered exponential in terms of its normalized distribution cancels the remaining normalizers, yielding the first identity. Taking a ratio gives the second. The case $g = 0$ is immediate. For a desired distribution r, choose $\ell^B_a = \log r(a) - g\widetilde{\ell}^P_a$ on the allowed support. Its fused softmax is r. □

Centering is therefore a softmax-invariant parameterization; it does not itself improve categorical probabilities. When $g > 0$, the prior reinforces its preferred categories. A negative unconstrained gate reverses these preferences. The pointwise construction is not a claim that a finite shared network realizes arbitrary corrections at all contexts. The architectural nesting at $g = 0$, however, is exact.

B.3.2 A SUFFICIENT CONDITION FOR STRICTLY LOWER CROSS-ENTROPY

The product identity explains how transfer occurs, but does not alone show that it helps. We next give a prediction-risk criterion.

Let X denote an active residue-time context sampled from a fixed distribution ν, with true conditional residue law r_X supported on the admissible alphabet, native prediction p_X, and frozen prior q_X. Hold the base predictor and context distribution fixed while varying the global gate. Set $z_X(a) = \log q_X(a)$ and

$$\mathcal{E}(g) = \mathbb{E}_{X\sim\nu}\mathbb{E}_{a\sim r_X}[-\log p_{g,X}(a)]. \tag{24}$$

Assumption P (bounded, aligned structural preferences). Assume p_X, q_X have full support on $\mathcal{A}_i$, finite base risk, and

$$\max_a z_X(a) - \min_a z_X(a) \le D_P < \infty \quad \text{for } \nu\text{-almost every } X. \tag{25}$$

Define the average prior-alignment gap

$$\delta_P = \mathbb{E}_X[\mathbb{E}_{a\sim r_X} z_X(a) - \mathbb{E}_{a\sim p_X} z_X(a)]\,. \tag{26}$$

For positive-prior transfer, assume $\delta_P > 0$. This means that the true residue distribution receives a larger expected prior log score than the current native prediction does.

Proposition B.4 (Cross-entropy benefit of an aligned prior). Under Assumption P,

$$\begin{aligned} \mathcal{E}(g) - \mathcal{E}(0) &= \mathbb{E}_X\left[\log \mathbb{E}_{a\sim p_X} e^{g z_X(a)} - g\mathbb{E}_{a\sim r_X} z_X(a)\right], \\ \mathcal{E}'(0) &= -\delta_P, \qquad \mathcal{E}''(g) = \mathbb{E}_X \operatorname{Var}_{a\sim p_{g,X}}[z_X(a)] \le D_P^2/4. \end{aligned} \tag{27}$$

Consequently,

$$\mathcal{E}(g) \le \mathcal{E}(0) - g\delta_P + \frac{D_P^2 g^2}{8}. \tag{28}$$

If $D_P > 0$ and $\delta_P > 0$, every $0 < g < 8\delta_P/D_P^2$ gives a strict reduction, and the choice $g = 4\delta_P/D_P^2$ yields a reduction of at least $2\delta_P^2/D_P^2$.

Proof. Take the negative logarithm of the normalized product in equation 23, and average under r_X. Differentiation of the log normalizer gives its tilted mean and variance. To justify differentiation even if the absolute log scores vary across contexts, subtract $\min_a z_X(a)$ separately in each context; the risk difference and all derivatives are unchanged, while the centered scores are bounded in $[0, D_P]$.

For a random variable $W \in [a, b]$, the inequality $(W - a)(b - W) \ge 0$ gives $\operatorname{Var}(W) \le (\mathbb{E}W - a)(b - \mathbb{E}W) \le (b-a)^2/4$. Apply this to the tilted prior score to obtain the curvature bound. Taylor's theorem then gives equation 28; the asserted interval and quantitative reduction follow by completing the quadratic. □

Interpretation and training. The assumption is a calibration condition on current intermediate backbones, not on clean structures alone. The result supplies an explicit way in which externally learned preferences can reduce sequence loss: they must correct directions in which the native predictor underweights compatible residues. It also explains why the gate should be learned. For the same native predictor, $g = 0$ remains available; if $\delta_P < 0$, sufficiently small negative gates decrease the fixed-context risk instead. If the prior is uniform on the support, it has no effect. If $p_X = r_X$, then $\delta_P = 0$ and no strict improvement over the Bayes predictor is possible.

During joint training, the parameter class includes the native model and, under Assumption P at a given native predictor, a strictly better positive-gate prediction rule. This is a population-risk existence statement, not a guarantee that finite-data optimization finds that rule. The numerical gate above is an analytical witness, not a recommended replacement for the trained gate.

Gradient paths and dynamic use. The implementation evaluates $\mathrm{stopgrad}(F_\phi(\mathcal{B}_t; \mathcal{R}))$. The stop-gradient operator blocks that computational gradient path, while ϕ stays frozen. It does not make the mathematical function F_ϕ constant in its coordinates. Its values can change as $\mathcal{B}_t$ changes, and fused categorical samples can influence later backbone updates through the shared network.

The risk result fixes ν and the base predictions. It directly describes prediction on sampled training contexts. During inference, changing g can also change future contexts, so terminal cofolding confidence requires a separate trajectory-level argument or empirical evaluation. Frozen pretrained parameters transfer learned preferences; they do not create information beyond the input backbone in the information-theoretic sense.

B.4 Cascaded search: structural selection, sequence coverage, and refinement

This part analyzes the cascade as a sequence of conditional decisions. The structural selector reduces redundant conditioning contexts before sequence expansion. Quality-based selection occurs later, after short AfCycDesign trials. These operations use different surrogates and have different possible failure modes.

B.4.1 Branchable checkpoints and conditional continuation laws

Suppose the backbone, sequence, and side-chain sources are independent given Z, and inactive peptide sequence and side-chain inputs are masked before τ_{s}. Assume also that no external sequence-dependent guidance is applied during this prefix. Then the prefix checkpoint relevant to backbone evolution can be constructed without using the unrevealed sequence source. Replacing that source at τ_{s} leaves the generated prefix unchanged. This follows by induction over the prefix updates: each update uses only the previous prefix state, the conditions, and permitted prefix randomness.

Generate a root pool $\mathcal{H} = (H_1, \ldots, H_K)$ at τ_{s}. Each H_i contains $\mathcal{B}_{\tau_{\mathrm{s}}}^{(i)}$ and all memory needed for a valid continuation, including source-dependent sampler variables. Any stored unused sequence variables are replaced by the child's fresh source; they are not held fixed as conditioning variables in P_i below. For a retained index set A, let

$$P_i(s) = \mathbb{P}_\theta(\mathcal{S}_1 = s \mid H_i, Z; \text{fixed continuation and guidance protocol}). \tag{29}$$

Conditional on $\mathcal{H}$ and A, assume distinct children use independent fresh sequence sources and continuation randomness. The external oracle is deterministic or uses independent randomness for each child, with no cross-child adaptation. Then the M children of root i are independent draws from P_i. Their backbone states continue to evolve; no side-chain branching factor is added to the cofolding policy.

The selected roots are generally dependent because selection is history-dependent and without replacement. The sequence identities below condition on the realized root pool and selected set; they do not assume independent selected roots.

B.4.2 Structural potentials and reduced redundancy

Let d_{ij} be the declared distance between visible backbones of roots i, j, preserving the task's residue correspondence. Distances in a common receptor frame measure pose as well as shape; independent

peptide alignment measures shape alone. The proof applies to either explicitly chosen distance but does not identify them.

Let $\kappa_{ij} = \kappa(d_{ij}) \in [0, 1]$ be symmetric, with κ non-increasing. No positive-semidefinite-kernel or determinantal-process assumption is required. For a nonempty selected set A, define

$$c_i(A) = \frac{1}{|A|} \sum_{j \in A} \kappa_{ij}, \qquad i \in \mathcal{U}_A := \{1, \dots, K\} \setminus A. \tag{30}$$

Let $w(c) > 0$ be bounded and non-increasing on $[0, 1]$. Select

$$\mu_A(i) = \frac{w(c_i(A))}{\sum_{j \in \mathcal{U}_A} w(c_j(A))}, \qquad i \in \mathcal{U}_A. \tag{31}$$

The first root is selected uniformly; subsequent draws are without replacement until $|A| = K_0$. This formulation covers a family of probability-based crowding rules. The evaluated implementation must specify its distance, similarity, and weighting function.

Potential-based reweighting at an intermediate generation checkpoint is related to Feynman–Kac steering (Singhal et al., 2025). Here the potential also depends on the roots already selected. Equation 31 defines the conditional choice law being analyzed; it is not asserted to sample exactly from a prescribed terminal reward distribution.

Proposition B.5 (One-step reduction in structural crowding). Fix a root pool and a nonempty selected set A. Let u_A be uniform on $\mathcal{U}_A$. Then

$$\mathbb{E}_{\mu_A} c_i(A) - \mathbb{E}_{u_A} c_i(A) = \frac{\mathrm{Cov}_{u_A}(c_i(A), w(c_i(A)))}{\mathbb{E}_{u_A} w(c_i(A))} \le 0. \tag{32}$$

The inequality is strict when two remaining crowding values differ and the weighting function strictly distinguishes them. Every remaining root retains positive selection probability.

Proof. The equality follows by writing the reweighted expectation as $\mathbb{E}_{u_A}[cw]/\mathbb{E}_{u_A} w$. For independent indices $I, J \sim u_A$,

$$2\,\mathrm{Cov}_{u_A}(c, w(c)) = \mathbb{E}[(c_I - c_J)(w(c_I) - w(c_J))] \le 0, \tag{33}$$

because the differences have opposite signs. Strictness follows if a pair gives a strictly negative product; positivity follows from $w > 0$. □

For example, $w(c) = \exp(-\beta c)$ with $\beta > 0$ has the regularization interpretation

$$\mu_A = \arg\min_{\mu} \left\{ \mathbb{E}_{\mu} c_i(A) + \beta^{-1} D_{\mathrm{KL}}(\mu \| u_A) \right\}. \tag{34}$$

A Lagrange multiplier gives this form; strict convexity gives uniqueness. It tends to uniform selection as $\beta \downarrow 0$ and concentrates on minimum-crowding roots as $\beta \to \infty$. The covariance result is more general and does not require using this particular family in the implementation.

Scope and cost. The comparison uses the same root pool and selection history for both choices. It proves a local reduction in expected redundancy, not global optimality of the selected set or domination of a separate adaptive run. Low crowding alone establishes neither physical validity nor sequence diversity. With a cached similarity matrix, crowding sums can be updated in $O(K)$ work per root, or $O(KK_0)$ total. If a distance costs $O(L)$, the matrix costs $O(K^2 L)$ time and $O(K^2)$ storage; on-demand distances to selected roots cost $O(KK_0L)$ time and $O(K)$ auxiliary storage.

B.4.3 From geometric redundancy to conditional sequence coverage

Fix in advance a map $\psi : \mathcal{A}^L \to \{1, \dots, J\}$ describing sequence families. It may be the identity on exact sequences or a declared partition by sequence similarity or another fixed criterion. Let

$$\pi_i(v) = P_i(\psi(\mathcal{S}_1) = v), \qquad o_{ij} = \sum_{v=1}^{J} \pi_i(v)\pi_j(v). \tag{35}$$

Thus o_{ij} is the probability that independent continuations from roots i, j fall in the same sequence family. For a selected set A of size k, write $N = kM$. Let D_A count distinct families among the N children, and C_A count unordered pairs of children in the same family.

Proposition B.6 (Exact coverage and a collision-based certificate). Under conditional independence of the children,

$$\mathbb{E}[D_A \mid \mathcal{H}, A] = \sum_{v=1}^{J} \left[1 - \prod_{i \in A} (1 - \pi_i(v))^M \right], \tag{36}$$

and

$$\mathbb{E}[C_A \mid \mathcal{H}, A] = \binom{M}{2} \sum_{i \in A} o_{ii} + M^2 \sum_{\substack{i<j \\ i,j \in A}} o_{ij}. \tag{37}$$

Consequently, $\mathbb{E}D_A \geq N - \mathbb{E}C_A$. If structural similarity controls between-root sequence overlap as

$$o_{ij} \leq a_S \kappa_{ij} + \varepsilon_S \quad (i \neq j), \qquad a_S, \varepsilon_S \geq 0, \tag{38}$$

then

$$\begin{aligned} \mathbb{E}[D_A \mid \mathcal{H}, A] \geq N &- \binom{M}{2} \sum_{i \in A} o_{ii} \\ &- M^2 \left(a_S \sum_{\substack{i<j \\ i,j \in A}} \kappa_{ij} + \varepsilon_S \binom{k}{2} \right). \end{aligned} \tag{39}$$

Proof. A family v is absent exactly when every child misses it. Conditional independence gives the product in equation 36; sum the inclusion indicators. Within root i, $\binom{M}{2}$ pairs have collision probability o_{ii}. Between roots i, j, M^2 pairs have collision probability o_{ij}. Linearity of expectation proves the collision identity. If family v occurs n_v times, then $(n_v - 1)_+ \leq \binom{n_v}{2}$. Summing gives $N - D_A \leq C_A$. Substitute equation 38 for the final bound. □

Equation 38 is a cross-modal calibration assumption: backbone similarity must be informative about overlap between conditional sequence distributions. Coupling and the backbone-conditioned prior provide a mechanism for such dependence, but do not prove this bound. Arbitrarily large a_S or ε_S can make the certificate vacuous and do not establish useful cross-modal calibration. With the within-root term held fixed or controlled, reducing the kernel sum improves the coverage lower bound.

A stronger local calibration yields a reduction in actual expected sequence collisions. Fix A with $r = |A| > 0$, and suppose

$$\begin{aligned} o_{ii} &= d_A \quad \text{for all } i \in \mathcal{U}_A, \\ o_{ij} &= a_S \kappa_{ij} + b_A + e_{ij}, \quad |e_{ij}| \leq \varepsilon_S, \quad i \in \mathcal{U}_A,\ j \in A. \end{aligned} \tag{40}$$

Let $\Delta_c(A) = \mathbb{E}_{u_A} c_i(A) - \mathbb{E}_{\mu_A} c_i(A) \geq 0$.

Corollary B.7 (Reduced sequence redundancy under calibration). Under equation 40, the expected number of new colliding child pairs after adding one root satisfies

$$\mathbb{E}_{\mu_A}[\Delta C_i] - \mathbb{E}_{u_A}[\Delta C_i] \leq -M^2 r \big(a_S \Delta_c(A) - 2\varepsilon_S \big). \tag{41}$$

It is strictly smaller under the crowding selector if $a_S \Delta_c(A) > 2\varepsilon_S$.

Proof. For candidate i, $\mathbb{E}[\Delta C_i] = \binom{M}{2} o_{ii} + M^2 \sum_{j \in A} o_{ij}$. The diagonal and b_A terms cancel between selection laws. The kernel term is $-M^2 r a_S \Delta_c(A)$. The sum of errors lies in $[-r\varepsilon_S, r\varepsilon_S]$ for each i, so its two expectations differ by at most $2r\varepsilon_S$. □

An exact coverage case. Suppose roots belong to structural classes, κ_{ij} is one for roots in the same class and zero otherwise, and each class corresponds to a distinct sequence family with probability one. Roots from a class not yet selected have crowding zero and maximal weight $w(0)$. If n_{new} of the n remaining roots belong to unseen classes, then

$$\mathbb{P}_{\mu_A}(\text{new sequence family}) = \frac{n_{\text{new}} w(0)}{n_{\text{new}} w(0) + \sum_{i:\,\text{seen class}} w(c_i(A))} \geq \frac{n_{\text{new}}}{n}. \tag{42}$$

The right-hand side is the uniform-choice probability. Thus this separated conditional-family model directly converts geometric selection into a greater chance of adding a new family at the current decision. It is an example, not a universal model of the trained generator.

What breadth and branching each contribute. Let I be uniform over A and $V \sim \pi_I$ a random child's family. For $\bar{\pi}_A = k^{-1} \sum_{i \in A} \pi_i$,

$$H(\bar{\pi}_A) = \frac{1}{k} \sum_{i \in A} H(\pi_i) + \frac{1}{k} \sum_{i \in A} D_{\mathrm{KL}}(\pi_i \| \bar{\pi}_A). \tag{43}$$

This is $H(V) = H(V \mid I) + I(V; I)$. Distinct root-conditioned laws contribute between-root information; M children explore variability within each law. If all P_i are identical, structural selection cannot change the sequence sampling law. Also, sequence distributions may already have full support. Broader exploration here means finite-budget coverage and less repetition, not necessarily a larger mathematical support.

B.4.4 When broader exploration improves candidate discovery

For a fixed root set and a fixed per-child continuation/refinement protocol, let $p_i(q)$ be the probability that a child of root i achieves quality at least q. Conditional independence gives

$$\mathbb{P}\left(\max_{i \in A,\, m \le M} V_{i,m} \ge q \;\middle|\; \mathcal{H}, A \right) = 1 - \prod_{i \in A} (1 - p_i(q))^M. \tag{44}$$

Here V is a declared final cofolding score or hit indicator. Diversity is useful only if it changes the success opportunities being sampled; this identity alone does not favor diverse roots.

To make that requirement explicit, partition conditional sequence exploration into regions v. For a chosen quality threshold, let $\Theta_v \in [0, 1]$ be region v's latent success probability under a fixed child protocol.

Assumption Q (distinct, unresolved quality opportunities). The Θ_v are independent and identically distributed across regions, and independent of the geometry-only allocation. Conditional on them, child successes are independent Bernoulli draws with parameter Θ_v in region v. Allocate n_v children to region v, with $\sum_v n_v = N$ and fixed per-child computation, without observing the latent Θ_v.

Proposition B.8 (Diversification reduces an averaged miss probability). Under Assumption Q, let $m(n) = \mathbb{E}[(1 - \Theta)^n]$, with $m(0) = 1$. Then

$$P_{\mathrm{miss}}(n_1, \ldots, n_J) = \prod_{v=1}^{J} m(n_v). \tag{45}$$

For integers $a, b \ge 1$,

$$m(a)m(b) \le m(a + b). \tag{46}$$

Splitting $a + b$ trials from one region into a and b trials in two independent regions therefore cannot increase the averaged miss probability. The split factor decreases strictly when Θ is non-degenerate. For fixed N and a fixed number of available exchangeable regions, a balanced allocation minimizes the miss probability.

Proof. Conditional on the latent probabilities, all trials fail with probability $\prod_v (1 - \Theta_v)^{n_v}$. Independence gives the product of moments. Let $X = 1 - \Theta$ and let X' be an independent copy. Since both powers are increasing on $[0, 1]$,

$$m(a + b) - m(a)m(b) = \tfrac{1}{2} \mathbb{E}[(X^a - X'^a)(X^b - X'^b)] \ge 0. \tag{47}$$

For non-degenerate X, the integrand is strictly positive with positive probability. For balancing, Cauchy–Schwarz gives $m(n)^2 \le m(n - 1)m(n + 1)$ for $n \ge 1$. Thus the increments of $\log m(n)$ are non-decreasing when the moments are positive. If $n_u \ge n_v + 2$, moving one trial from u to v cannot increase $\log m(n_u) + \log m(n_v)$. Repeated moves yield a balanced allocation. If a positive moment vanishes, $\Theta = 1$ almost surely and the miss probability is already zero for any nonempty allocation. □

Relation to structural selection. The overlap assumptions link lower crowding to less redundant conditional sequence exploration. Assumption Q then gives a regime where allocating trials across distinct sequence regions improves discovery probability at the same leaf budget. The assumptions play different roles; neither follows from geometry alone. The comparison also fixes per-child computation and does not yet include quality pruning.

This is an average over unresolved region quality, not a guarantee for every fixed assignment of success rates. If all regions have the same deterministic rate, spreading trials leaves success probability unchanged. If high-quality regions are already known, concentration can be preferable. Quality-based refinement addresses that later information regime. No monotonic score gain from increasing K or reducing crowding alone is asserted.

B.4.5 Quality selection, early stopping, and total compute

After $b = 5$ initial AfCycDesign updates on each of $N = K_0 M$ leaves, let

$$Q_j = 0.75\big(1 - \widetilde{\mathrm{iPAE}}_{\mathrm{hotspot},j}\big) + 0.25\,\mathrm{ipTM}_j \tag{48}$$

be the screening score. Let positive weights $v_j = v(Q_j)$ increase with Q_j. The actual weighting rule and replacement policy must be reported. The result below considers sequential selection of N_0 distinct leaves using these fixed weights.

To separate screening error from generation, couple each candidate to its potential outcome under the same full-budget refinement protocol. Let $\mathcal{G}$ contain candidates that would reach the specified quality threshold under that protocol; future optimizer randomness can be fixed in advance for this definition. The selector's random choices are independent of those potential outcomes conditional on its observed scores.

Proposition B.9 (Retention of potentially successful candidates). Condition on screening scores and potential outcomes, with $\mathcal{G} \neq \varnothing$. Define

$$\omega_{\mathcal{G}} = \frac{\sum_{j \in \mathcal{G}} v_j}{\sum_{j=1}^{N} v_j}. \tag{49}$$

Then

$$\mathbb{P}(A_Q \cap \mathcal{G} = \varnothing \mid Q_1, \ldots, Q_N, \mathcal{G}) \leq (1 - \omega_{\mathcal{G}})^{N_0}. \tag{50}$$

Proof. Conditioned on having selected no member of $\mathcal{G}$, its total remaining weight is unchanged, while the denominator can only decrease. Hence the conditional probability of selecting a member at the next draw is at least $\omega_{\mathcal{G}}$. Multiply the conditional probabilities of missing it at all N_0 draws. □

The bound is useful when early scores assign appreciable weight to candidates that benefit from further optimization. For example, if all members of $\mathcal{G}$ score at least as highly as all nonmembers, monotonicity of v gives $\omega_{\mathcal{G}} \geq |\mathcal{G}|/N$. A high early score does not establish this calibration automatically. Positive weights preserve a chance of retaining each candidate but do not eliminate false negatives.

Let δ_{stop} be the probability that a retained potentially successful candidate exists but the actual stopping and discrete-completion policy returns no success. Combining retention with this definition gives

$$P_{\text{cascade hit}} \geq \mathbb{E}\big[\mathbf{1}\{\mathcal{G} \neq \varnothing\}\,\big(1 - (1 - \omega_{\mathcal{G}})^{N_0}\big)\big] - \delta_{\mathrm{stop}}. \tag{51}$$

The integrand is defined as zero if $\mathcal{G}$ is empty. This separates useful-candidate generation, retention, and stopping losses. Five soft-stage steps without improvement do not by themselves bound δ_{stop}; later improvement remains possible. The patience rule is therefore an empirical budget policy, not a proven convergence test.

Compute accounting. Let $T_{\max}$ be the full per-leaf optimizer cap including the b screening updates. If survivors resume their optimizer states and receive $\ell_j \leq T_{\max} - b$ additional updates, then

$$U_{\mathrm{cascade}} = bN + \sum_{j \in A_Q} \ell_j \leq bN + N_0(T_{\max} - b). \tag{52}$$

Relative to fully optimizing all N leaves,

$$NT_{\max} - U_{\text{cascade}} = (N - N_0)(T_{\max} - b) + \sum_{j \in A_Q} (T_{\max} - b - \ell_j). \tag{53}$$

The first term saves refinement on rejected leaves; the second saves work from stopping survivors early. If optimization is restarted, repeated work must be added. Different step types or recycle counts require measured costs rather than treating every update as equal.

With prefix cost c_{pre} and continuation cost c_{cont}, the end-to-end budget is

$$\begin{aligned} C_{\text{total}} = K c_{\text{pre}} + K_0 M c_{\text{cont}} + C_{\text{guide}} + C_{\text{select}} \\ + C_{\text{screen}} + C_{\text{refine}} + C_{\text{final evaluation}}. \end{aligned} \tag{54}$$

The resource-allocation principle is related to early-stopping methods such as Hyperband (Li et al., 2018), but this optimizer, its scores, and probabilistic selection do not inherit Hyperband's guarantees. AfCycDesign refinement is not a flow or diffusion trajectory.

Implications for the integrated policy. Structural selection reduces conditional geometric redundancy. When that redundancy tracks overlap between conditional sequence laws, it can reduce repeated exploration and improve a finite-budget coverage certificate. Under distinct unresolved quality opportunities, broader coverage can improve discovery probability. As quality evidence becomes available, weighted retention assigns additional refinement to promising leaves, with the stated retention bound. The policy trades exploration against optimization cost; this analysis neither attributes every empirical gain to a separate stage nor proves optimality of the chosen hyperparameters.

The results concern computational prediction and search. They do not make AlphaFold confidence a guarantee of experimental binding, or imply that full-atom generation cannot benefit from a different scaling policy. The evaluated scope is the integrated cofolding cascade in Section 3.3.

C DC-FLOW FORMULATION AND ARCHITECTURE

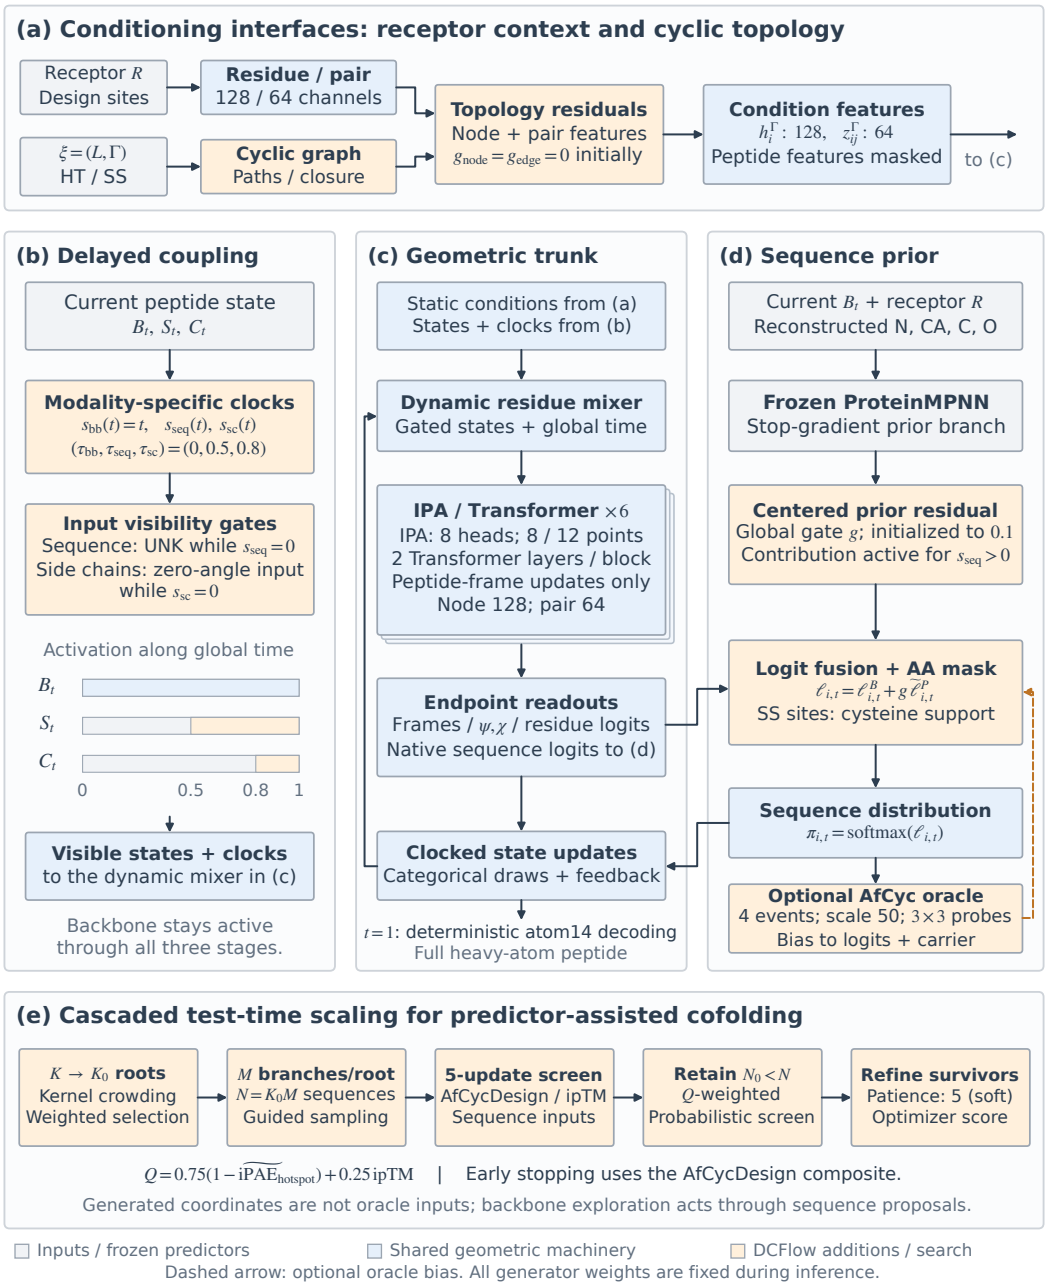

Figure 5: **Detailed architecture and inference interfaces of DC-FLOW.** (a) Receptor features and cyclic graph features form topology-conditioned residue and pair representations. (b) Modality clocks control both state progression and input visibility. (c) A shared geometric trunk predicts backbone, angular, and sequence endpoints; sampled sequences re-enter subsequent predictions while the backbone remains active. (d) Frozen ProteinMPNN predictions from the current intermediate backbone enter sequence logits through a trainable residual gate. Optional oracle feedback supplies persistent sequence biases. (e) The cofolding cascade selects backbone roots, branches guided continuations, and allocates further refinement after a short screen. Gray denotes inputs or frozen predictors, blue the shared geometric machinery, and orange the added conditioning and allocation mechanisms. All generator parameters are fixed during inference.

This appendix specifies how the three mechanisms in Section 3 are instantiated within a shared generative network. Figure 5 shows the conditioning interfaces, delayed modality activation, geometric trunk, and residual sequence prior, together with their connection to the cofolding cascade. The theoretical analysis is given in Appendix B; training settings, refinement protocols, and evaluation metrics are reported in Appendix E.

C.1 State Representation and Modality-Specific Paths

Geometric and chemical variables. We retain the main-text state $Y_t = (B_t, S_t, C_t)$. For peptide residue $i \in \{1, \dots, L\}$, let $x_{i,t} \in \mathbb{R}^3$ be its Cα translation and $O_{i,t} \in \mathrm{SO}(3)$ its local orientation, defined using N–Cα–C. We use $O_{i,t}$ to distinguish a residue rotation from the receptor R and the screening score Q. The backbone state B_t also includes a carbonyl-orientation channel $\psi_{i,t}$ computed from N–Cα–C–O. This channel controls the local carbonyl oxygen and is distinct from the conventional backbone dihedral defined with the next residue's nitrogen. The side-chain state C_t consists of up to four torsions $\chi_{i,t} = (\chi^1_{i,t}, \dots, \chi^4_{i,t})$ per residue, with residue-specific validity masks.

The implementation represents S_t by a continuous carrier $u_t \in \mathbb{R}^{L \times |\mathcal{A}|}$ and a sampled categorical sequence $a_t \in \mathcal{A}^L$. The carrier is mapped to probabilities by softmax; it is not itself constrained to the probability simplex. For a residue identity a, define

$$\mathrm{Encode}(a) = 10\ \mathrm{onehot}(a) - 5\mathbf{1}, \tag{55}$$

where $\mathbf{1}$ is the all-ones vector over the alphabet. The encoded category therefore has value $+5$ and the remaining categories have value -5. The terminal discrete sequence is denoted S_1 in the main text.

Training paths. Only peptide states are perturbed. Translation sources are centered Gaussian samples, orientations are sampled uniformly on $\mathrm{SO}(3)$, angular sources are uniform on the circle, and u_0 is Gaussian with standard deviation 5. Native endpoint states carry subscript 1. Using the modality clocks $s_m(t)$ defined in Section 3.2, the interpolation paths are

$$\begin{aligned} x_t &= (1 - s_{\mathrm{bb}}(t))x_0 + s_{\mathrm{bb}}(t)x_1, \\ u_t &= (1 - s_{\mathrm{seq}}(t))u_0 + s_{\mathrm{seq}}(t)u_1, \\ O_t &= \mathrm{Geo}_{\mathrm{SO}(3)}(O_0, O_1; s_{\mathrm{bb}}(t)), \\ \psi_t &= \mathrm{Geo}_{\mathbb{S}^1}(\psi_0, \psi_1; s_{\mathrm{bb}}(t)), \\ \chi_t &= \mathrm{Geo}_{\mathbb{S}^1}(\chi_0, \chi_1; s_{\mathrm{sc}}(t)). \end{aligned} \tag{56}$$

Here $\mathrm{Geo}_{\mathrm{SO}(3)}$ denotes geodesic interpolation between rotations, and $\mathrm{Geo}_{\mathbb{S}^1}$ denotes circular interpolation applied componentwise to the valid angular channels. Thus the carbonyl channel follows the backbone clock, whereas side-chain torsions follow the delayed side-chain clock. Intermediate residue identities are sampled from $\mathrm{softmax}(u_{i,t})$.

C.2 Receptor and Cyclization Conditioning

Static residue and pair features. The input contains the designated receptor pocket and peptide design positions. A residue encoder combines observed identities, local heavy-atom coordinates, and periodic backbone-angle features into 128-dimensional representations. A pair encoder combines residue-pair identities, same-chain relative positions clipped to ± 32, atom-distance features, and orientation features into 64-dimensional representations. The native peptide's identities and geometric features are masked in these static encoders during training. Observed pair geometry is retained only when both residues belong to the visible structural context. Peptide–receptor interactions are subsequently modeled using the evolving peptide frames in geometric attention.

The receptor and peptide participate jointly in attention. Receptor hidden features can therefore change, but its coordinates, identities, and angular states remain fixed. Static condition features are computed once per sampling trajectory, while the current peptide state is incorporated at each network evaluation.

Cyclic graph residuals. The cyclization specification Γ augments sequential peptide edges with the prescribed closing edge: an N–C linkage for head-to-tail (HT) closure or an SG–SG linkage

for disulfide (SS) closure. For each peptide residue pair, the graph encodes shortest-path distance, whether the selected shortest path crosses a closing edge, and the type of any direct closing edge. Equal-length paths prefer a path without the closing edge. These features supplement the original linear relative positions, allowing the network to distinguish sequence separation from chemical adjacency.

Let h_i^0 and z_{ij}^0 denote the static residue and pair features. Write c_Γ for the topology category, v_i^Γ for a closure-site indicator, and $d_{ij}^\Gamma, r_{ij}^\Gamma, e_{ij}^\Gamma$ for graph distance, path type, and direct closing-edge type. With learned embeddings $E_{\text{mode}}, E_{\text{site}}, E_d, E_r, E_e$, conditioning is injected as

$$\begin{aligned} h_i^\Gamma &= h_i^0 + g_{\text{node}}\big(E_{\text{mode}}(c_\Gamma) + E_{\text{site}}(v_i^\Gamma)\big), \\ z_{ij}^\Gamma &= z_{ij}^0 + g_{\text{edge}} m_{ij}^{\text{pep}}\big(E_d(d_{ij}^\Gamma) + E_r(r_{ij}^\Gamma) + E_e(e_{ij}^\Gamma)\big). \end{aligned} \tag{57}$$

Here m_{ij}^{pep} is one for peptide–peptide pairs and zero otherwise. The node residual is applied across residue tokens, with closure-site indicators identifying the relevant peptide positions. The unconstrained scalar gates g_{node} and g_{edge} are initialized to zero, preserving the pretrained condition features at initialization. At specified SS sites, the admissible sequence categories are restricted to cysteine before prediction probabilities are normalized.

C.3 Delayed Visibility and the Shared Geometric Trunk

Input visibility and clock conditioning. The delayed schedule affects interpolation, network inputs, supervision, and numerical updates. When $s_{\text{seq}}(t) = 0$, peptide residue identities supplied to the trunk are replaced by an unknown-residue token. When $s_{\text{sc}}(t) = 0$, the side-chain inputs are replaced by fixed zero angles before periodic encoding. The latter is a zero-angle representation rather than a zeroed feature vector. Consequently, inactive random sequence and side-chain states do not enter the trunk as sample-dependent chemical context.

In addition to the global-time embedding, the network receives a residual clock feature

$$c_t^{\text{clock}} = W_{\text{clock}}\left[E_{\text{seq}}(s_{\text{seq}}(t)); E_{\text{sc}}(s_{\text{sc}}(t))\right], \tag{58}$$

where each E produces a 128-dimensional time embedding, semicolons denote concatenation, and $W_{\text{clock}} \in \mathbb{R}^{128\times 256}$ is a bias-free, zero-initialized projection. The dynamic feature mixer combines these clock features with the static conditions, current residue identities, angular encodings, and global time.

Geometric updates and endpoint readouts. Six shared blocks process the joint receptor–peptide representation. Each block applies invariant point attention (IPA), residual normalization, two Transformer layers, a residue transition, and a local rigid-frame update. Pair representations are updated between blocks, except after the final block. IPA combines scalar feature similarity, pair-feature biases, and distances between points transformed by the current frames. Its geometric interactions therefore use the evolving peptide pose in the fixed receptor coordinate system.

Frame updates are restricted to valid peptide residues within every block. This constraint preserves receptor geometry throughout the forward pass, rather than restoring it only after the pass. The final frames provide backbone endpoint predictions; separate multilayer perceptrons produce 20 residue logits and five angular endpoints (one carbonyl and four side-chain channels). These heads share the same trunk. Once sequence or side chains become visible, their features can affect subsequent backbone predictions without requiring a separate backbone-refinement network.

C.4 Dynamic ProteinMPNN Residual Prior

The ProteinMPNN branch reconstructs peptide N, Cα, C, and O coordinates from the *current* frames and carbonyl angles, and combines them with the observed receptor backbone. It evaluates sequence-independent, backbone-conditioned residue probabilities using the frozen model F_ϕ (Dauparas et al., 2022). Neither the native peptide coordinates nor the predicted clean endpoint is supplied to this branch.

The 20 standard-amino-acid log-probabilities are centered across categories and fused with the native sequence head according to equation 3. Subtracting the category mean removes a softmax-invariant

Table 5: Architecture settings of the shared generator. The frozen ProteinMPNN branch is external to the six geometric blocks.

Component	Configuration
Residue / pair channels	128 / 64
Geometric blocks	6
IPA heads	8
IPA query/key / value points	8 / 12 per head
Transformer layers / heads	2 per block / 4
Transformer feed-forward width	128; dropout 0
Sequence head	$128 \to 128 \to 128 \to 20$
Angular head	$128 \to 128 \to 128 \to 5$
Additional clock projection	$256 \to 128$; zero initialization
Topology residual gates	Two trainable scalars; zero initialization
ProteinMPNN residual gate	Global scalar g; initialized to 0.1

offset; it does not change relative residue preferences. The gate g is a learned global scalar, and the prior contribution is enabled only when $s_{\text{seq}}(t) > 0$. This activation controls the contribution: the implementation can still evaluate the frozen branch before activation and multiply its residual by zero.

ProteinMPNN parameters and its coordinate-input path are excluded from backpropagation. The fused sequence loss trains the native trunk, sequence head, and gate g. After fusion, disallowed residue categories are masked before softmax normalization. Sampling the resulting distribution changes the sequence context provided at subsequent steps, allowing transferred structural preferences to influence later backbone and side-chain predictions through the shared trunk.

C.5 Supervision, Numerical Updates, and Reconstruction

Supervision follows modality activation. The grouped objective in equation 2 combines translation endpoint error, rotational matching, reconstructed-backbone error, sequence cross-entropy, angular matching, and cyclic-closure supervision. Given the predicted endpoints $\widehat{Y}_1 = (\widehat{B}_1, \widehat{S}_1, \widehat{C}_1)$ and native structures $Y_1 = (B_1, S_1, C_1)$, the individual loss terms are formally defined as follows:

- **Backbone Loss (ℓ_B):** Combines translation, rotation, reconstructed atoms, and carbonyl orientation via $\lambda_B \ell_B = \lambda_{\text{trans}}\ell_{\text{trans}} + \lambda_{\text{rot}}\ell_{\text{rot}} + \lambda_{\text{bb}}\ell_{\text{bb}} + \lambda_\psi \ell_\psi$, where

$$\begin{aligned} \ell_{\text{trans}} &= \frac{1}{L}\sum_{i=1}^{L} \|\widehat{x}_{i,1} - x_{i,1}\|_2^2, & \ell_{\text{rot}} &= \frac{1}{L}\sum_{i=1}^{L} \|\text{Log}(R_{i,1}^T \widehat{R}_{i,1})\|_2^2, \\ \ell_{\text{bb}} &= \frac{1}{3L}\sum_{i=1}^{L}\sum_{a\in\{\text{N,C}\alpha,\text{C}\}} \|\widehat{x}_{i,a} - x_{i,a}\|_2^2, & \ell_\psi &= \frac{1}{L}\sum_{i=1}^{L}\Big(1 - \cos(\widehat{\psi}_{i,1} - \psi_{i,1})\Big). \end{aligned} \tag{59}$$

  $\text{Log}(\cdot)$ denotes the $\text{SO}(3)$ logarithmic map penalizing Lie-algebra tangent deviations; ℓ_ψ supervises the N–Cα–C–O dihedral throughout generation.
- **Sequence Loss (ℓ_S):** Supervised via cross-entropy for $t \geq \tau_{\text{s}}$ against native residues a_i^*:

$$\ell_S = -\frac{1}{L}\sum_{i=1}^{L} \log \pi_{i,t,a_i^*}, \quad \text{where } \pi_{i,t} = \text{softmax}(\ell_{i,t}). \tag{60}$$

  Input visibility and ProteinMPNN prior fusion require $t > \tau_{\text{s}}$.
- **Side-Chain Angular Loss (ℓ_C):** Supervises valid channels $\chi = (\chi^1, \ldots, \chi^4)$ under mask $m_{i,k} \in \{0, 1\}$:

$$\lambda_C \ell_C = \lambda_{\text{angle}}\ell_{\text{angle}} + \lambda_{\text{torsion}}\ell_{\text{torsion}}, \quad \text{with } \ell_{\text{angle}} = \frac{\sum_{i=1}^{L}\sum_{k=1}^{4} m_{i,k}\left(1 - \cos(\widehat{\chi}_{i,1}^k - \chi_{i,1}^k)\right)}{\sum_{i=1}^{L}\sum_{k=1}^{4} m_{i,k}}. \tag{61}$$

  ℓ_{torsion} provides auxiliary endpoint supervision. Angular terms activate at $t \geq \tau_{\text{c}}$ using global time scaling.

- **Cyclic Closure Loss (ℓ_{cyc}):** Penalizes closing-atom distance errors on reconstructed heavy atoms:

$$\ell_{\text{cyc}} = \lambda_{\text{HT}} \mathbf{1}[\Gamma = \text{HT}] \ell_{\text{cyc}}^{\text{HT}} + \mathbf{1}[t \geq \tau_{\text{c}}, \Gamma = \text{SS}] \lambda_{\text{SS}} \ell_{\text{cyc}}^{\text{SS}},$$
$$\text{where} \quad \begin{cases} \ell_{\text{cyc}}^{\text{HT}} = \left(\|\widehat{x}_{1,\text{N}} - \widehat{x}_{L,\text{C}}\|_2 - 1.33\,\text{Å} \right)^2, \\ \ell_{\text{cyc}}^{\text{SS}} = \left(\|\widehat{x}_{1,\text{SG}} - \widehat{x}_{L,\text{SG}}\|_2 - 2.03\,\text{Å} \right)^2. \end{cases} \tag{62}$$

  The disulfide term is active only when $t \geq \tau_{\text{c}}$, as sulfur coordinates require side-chain packing. These terms supervise closure distances rather than bond-angle or steric constraints.

Following the loss configuration in Appendix E.1, the loss weights are assigned as: $\lambda_{\text{trans}} = 0.5$, $\lambda_{\text{rot}} = 0.5$, $\lambda_{\text{bb}} = 0.25$, $\lambda_{\psi} = 1.0$, $\lambda_S = 1.0$, $\lambda_{\text{angle}} = 1.0$, $\lambda_{\text{torsion}} = 0.5$, $\lambda_{\text{HT}} = 0.25$, and $\lambda_{\text{SS}} = 0.25$.

Source-dependent sampling and feedback. At each sampling step, the network predicts clean endpoints and the sampler advances the active states. Translation and sequence-carrier increments use the fixed source states, as specified in equation 69; they are not obtained by replacing the source with the current intermediate state. Rotations use endpoint-directed geodesic updates, and angular updates follow the backbone and side-chain clocks. Predicted sequence probabilities are sampled and encoded before the carrier update; a categorical sequence is then drawn from the updated carrier for the next trunk evaluation. The evaluated schedule uses 200 time points on $[0.01, 1]$, corresponding to 199 update intervals and one final endpoint prediction.

The complete checkpoint H_t introduced in Appendix B.4.1 therefore includes the current state and source-dependent sampling memory. Branching from a backbone prefix retains this context while providing independent randomness for the child continuations. At the endpoint, residue-specific templates and predicted frames and angles deterministically reconstruct up to 14 heavy atoms per residue. This atom14 reconstruction is a geometric decoder, not an additional atom-coordinate diffusion model. The full-atom designation refers to heavy atoms; hydrogens are not explicitly generated.

C.6 Interfaces to Guided Cofolding and Cascaded Search

Direct and predictor-assisted outputs. Direct generation returns the reconstructed peptide–receptor complex, which is subsequently assessed using the common physical protocol. In the cofolding workflow, the generated *sequence* is passed to the external predictor and AfCycDesign optimizer. The generated peptide coordinates are not provided as templates or initial poses. Structural exploration therefore affects this workflow through the sequence distributions induced by different backbone hypotheses.

In-trajectory sequence feedback. When guidance is enabled, the external cyclic oracle evaluates the current sequence and single-residue alternatives at four events during generation. Each event probes at most three positions and three alternative identities per position, using a base scale of 50. The resulting score-dependent biases modify later sequence logits, carrier updates, and the final categorical prediction. Oracle coordinates do not replace the peptide state, and generation-time guidance does not differentiate through the predictor. Sequence feedback can nevertheless influence later backbone updates through the shared trunk. The event schedule and bias rule are specified in Appendix E.6.

Cascade checkpoints and selection criteria. At τ_{s}, K backbone prefixes are probabilistically reduced to K_0 roots using structural crowding. Each retained root produces M guided continuations, yielding $N = K_0 M$ generated sequences. The backbones continue evolving during these continuations. All N sequences receive five initial AfCycDesign updates targeting ipTM, after which $N_0 < N$ candidates are retained probabilistically using Q from equation 5. This empirical score allocates further refinement; early stopping instead uses the composite score reported by AfCycDesign, with patience of five soft-stage updates. The final sequences undergo the common independent cofolding evaluation in Appendix E.5. The geometric root selection, sequence branching, and predictor-based retention thus operate at distinct checkpoints of one design workflow.

D CPSea Quality Filtering

D.1 Overview

CPSea provides a large source of cyclic peptide-receptor complexes (Yang et al., 2026). Our training corpus combines the CPTrans, CPBind, CPCore, and CPSea-PDB source pools and applies an additional strict quality-control pipeline designed to remove obvious geometric failures while preserving useful structural diversity.

The curation pipeline has two objectives: (i) to remove obvious geometric artifacts and conformations that are inconsistent with the intended chemical scope of the dataset, and (ii) to retain sufficient structural diversity for downstream learning tasks. To balance stereochemical conservatism against dataset coverage, we define two filtering regimes: a *Strict Pipeline*, which favors conformations close to commonly observed low-strain geometries, and a *Relaxed Pipeline*, which tolerates moderate ring-induced distortion while still excluding explicit geometric failures.

The qualitative locations of the preferred backbone, disulfide, and side-chain conformational regions are motivated by established peptide and protein structural chemistry. The exact numerical boundaries used below, however, are operational curation thresholds adopted for the present dataset and should not be interpreted as universal physical constants. The filtering strategy comprises four components: sequence-length constraints, head-to-tail amide-bond geometry, disulfide-bond quality audits, and the exclusion of side-chain isopeptide topologies.

D.2 Sequence Length Constraints

We restrict the working dataset to cyclic peptide chains with sequence length $\ell \geq 6$. This lower bound is used as an operational dataset-scope criterion rather than as a universal physical discontinuity at six residues. Classical conformational analyses show that very small cyclic peptides frequently require highly constrained backbone arrangements, unusual amide geometries, or sequence-specific turn-inducing features to achieve ring closure, whereas larger macrocycles more readily accommodate diverse turn-like and drug-relevant conformations (Kessler, 1982; White & Yudin, 2011; Ji et al., 2024). We therefore adopt $\ell \geq 6$ to reduce the representation of highly specialized, strongly strained microcycles and to focus the present dataset on the macrocyclic regime targeted by the downstream modeling tasks.

D.3 Head-to-Tail Amide Bond Geometries

The local geometry of the peptide backbone is governed in part by the amide torsion angle ω. Because peptide bonds strongly favor approximately planar geometries, we use ω as a primary criterion for backbone plausibility. Non-proline and proline-containing peptide bonds are treated separately.

Non-Proline Peptide Bonds. For ordinary secondary amide bonds, the *trans* state is strongly favored, although measurable nonplanarity and occasional substantial deviations are observed in high-resolution protein structures (Weiss et al., 1998; Berkholz et al., 2012). Non-proline *cis* peptide bonds can occur in exceptional structural contexts, but they lie outside the canonical backbone regime targeted by this dataset. We therefore discard any conformation containing a non-proline *cis* peptide bond.

For the remaining *trans* non-proline peptide bonds, we measure the deviation from ideal trans geometry according to

$$\big|\,|\omega| - 180^\circ\,\big| \leq \Delta\omega. \tag{63}$$

The values of $\Delta\omega$ are operational curation tolerances centered on the literature-supported preference for near-planar peptide bonds, rather than crystallographic pass/fail constants. We set $\Delta\omega = 10.0^\circ$ in the Strict Pipeline to favor near-planar amides, and relax the tolerance to $\Delta\omega = 15.0^\circ$ in the Relaxed Pipeline to preserve topologically closed conformers that remain chemically plausible but exhibit moderate ring-induced distortion (Weiss et al., 1998; Berkholz et al., 2012).

Proline-Containing Peptide Bonds. Peptide bonds preceding proline are a well-established exception to the behavior of ordinary secondary amides. The pyrrolidine ring reduces the free-energy

difference between the *cis* and *trans* states of the preceding X-Pro bond, such that both isomers can be populated under biologically relevant conditions (MacArthur & Thornton, 1991; Wedemeyer et al., 2002; Yao et al., 1994). A *cis* X-Pro bond can also support a tight local reversal, including type-VI-turn-like conformations, a feature that can be compatible with the conformational demands of cyclic peptide backbones (Kessler, 1982; Yao et al., 1994). We therefore retain both *cis* and *trans* states for proline-containing peptide bonds.

To reduce conformational sampling bias in the curated dataset, we balance the *cis* and *trans* pools of proline-containing structures using the following length-conditioned target fractions:

$$f_{\text{cis}}(\ell) = \begin{cases} 0.75 & \text{if } \ell = 6, \\ 0.30 & \text{if } 7 \le \ell \le 10, \\ 0.15 & \text{if } 11 \le \ell \le 16. \end{cases} \tag{64}$$

These values are dataset-balancing targets adopted for the present curation workflow; they are not intended to represent universal equilibrium *cis* populations. For each balanced pool of *cis* and *trans* structures, redundant or energetically unfavorable states are pruned by ranking conformations according to their computed Rosetta all-atom final energy and retaining only the lowest-energy subset (Alford et al., 2017). The energy function is used here as a physics-motivated ranking criterion for within-pool pruning rather than as an experimentally calibrated free energy.

D.4 Disulfide Bond Geometry and Quality Audit

Disulfide linkages (Cys-S-S-Cys) impose strong geometric constraints on peptide backbones and cysteine side chains. Protein disulfides occupy characteristic regions of bond-length, bond-angle, and dihedral space, although strained but physically realizable disulfides may deviate from the centers of these distributions (Thornton, 1981; Katz & Kossiakoff, 1986; Dombkowski, 2003). We therefore apply a four-part disulfide audit based on S-S distance, the two adjacent bond angles, the χ_3 dihedral, and the remaining cystine torsions.

The literature motivates the centers and qualitative shapes of the accepted conformational regions. The precise `PASS`, `WARN`, and `FAIL` boundaries below are operational quality-control envelopes adopted for this dataset. Under the *Strict Pipeline*, a disulfide bond is retained only if all four criteria receive `PASS`. Under the *Relaxed Pipeline*, geometries in the broader but still bounded `WARN` region are retained, whereas explicit geometric failures are discarded.

S-S Distance. Protein disulfides cluster around an S-S bond length of approximately 2.03 Å (Thornton, 1981; Katz & Kossiakoff, 1986). We classify the S-S distance $d_{\text{S}-\text{S}}$ as

$$\text{Status}(d_{\text{S}-\text{S}}) = \begin{cases} \text{PASS} & \text{if } 1.95\,\text{Å} \le d_{\text{S}-\text{S}} \le 2.12\,\text{Å}, \\ \text{WARN} & \text{if } 1.85\,\text{Å} \le d_{\text{S}-\text{S}} \le 2.20\,\text{Å}, \\ \text{FAIL} & \text{otherwise.} \end{cases} \tag{65}$$

Bond Angles. The C_β-S-S and S-S-C_β angles of protein disulfides are typically centered near 103°-104° (Thornton, 1981; Dombkowski, 2003). We evaluate each such angle θ according to

$$\text{Status}(\theta) = \begin{cases} \text{PASS} & \text{if } 95^\circ \le \theta \le 112^\circ, \\ \text{WARN} & \text{if } 90^\circ \le \theta \le 120^\circ, \\ \text{FAIL} & \text{otherwise.} \end{cases} \tag{66}$$

Dihedral Angle χ_3. The central disulfide dihedral χ_3 preferentially occupies near-perpendicular orientations, typically close to $\pm 90^\circ$, consistent with the stereoelectronic preferences of cystine linkages (Thornton, 1981; Katz & Kossiakoff, 1986). Ring closure can broaden the accessible χ_3 range, particularly in smaller macrocycles. To avoid disproportionately rejecting smaller rings, we use sequence-length-dependent operational PASS windows:

$$\text{Status}(\chi_3) = \begin{cases} \text{PASS} & \text{if } \begin{cases} |\chi_3| \in [50^\circ, 110^\circ] & \text{for } \ell \le 8, \\ |\chi_3| \in [65^\circ, 110^\circ] & \text{for } 9 \le \ell \le 10, \\ |\chi_3| \in [75^\circ, 105^\circ] & \text{for } \ell > 10, \end{cases} \\ \text{WARN} & \text{if outside PASS but within expanded boundaries,} \\ \text{FAIL} & \text{otherwise.} \end{cases} \tag{67}$$

The length dependence and the numerical limits in equation 67 are curation choices centered on the preferred disulfide conformational region; they are not claimed to be universal boundaries for all naturally occurring or synthetically constrained disulfides.

Cystine Torsional States. The remaining cystine torsions $(\chi_1, \chi_2, \chi_2', \chi_1')$ are biased toward staggered conformational basins characteristic of protein side chains and disulfide linkages (Thornton, 1981; Dunbrack Jr & Karplus, 1993; Lovell et al., 2000). For each angle $\chi_{\text{rot}} \in \{\chi_1, \chi_2, \chi_2', \chi_1'\}$, we calculate the absolute difference to the nearest canonical staggered state $\theta_r \in \{g^+ \approx +60°, g^- \approx -60°, t \approx 180°\}$:

$$\text{Status}(\chi_{\text{rot}}) = \begin{cases} \text{PASS} & \text{if } \min_r |\chi_{\text{rot}} - \theta_r| \leq 30.0°, \\ \text{WARN} & \text{if } \min_r |\chi_{\text{rot}} - \theta_r| \leq 45.0°, \\ \text{FAIL} & \text{otherwise.} \end{cases} \tag{68}$$

The $30°$ and $45°$ windows are operational tolerances used to distinguish core staggered states from progressively less typical torsional geometries. Any structure that triggers a `FAIL` status in any disulfide parameter is excluded.

D.5 Exclusion of Side-Chain Isopeptide Cyclizations

Cyclic peptides can be formed not only through head-to-tail backbone closure or disulfide formation, but also through side-chain-to-side-chain or side-chain-to-terminus isopeptide linkages, including Lys-to-Asp/Glu lactam bridges (Lundquist & Pelletier, 2002; White & Yudin, 2011; Hayes et al., 2021). We exclude side-chain-to-side-chain lactam topologies as a dataset-scope decision intended to keep the present corpus focused on the two cyclization classes emphasized here: canonical head-to-tail macrocycles and disulfide-bridged systems. This restriction should not be interpreted as evidence that side-chain lactams lack chemical, biological, or therapeutic value.

The scope restriction is motivated by practical consistency. Head-to-tail and disulfide-linked peptides are major and widely used classes in cyclic-peptide discovery and development (Ji et al., 2024; Hayes et al., 2021). By contrast, selective side-chain lactam formation commonly requires topology-specific protecting-group and deprotection strategies, often using orthogonal protection schemes, and may require additional route optimization to control competing coupling sites and cyclization outcomes (Lundquist & Pelletier, 2002; White & Yudin, 2011; Fang et al., 2024). These synthetic considerations can complicate uniform, high-throughput treatment across a heterogeneous structural database (Hayes et al., 2021; Fang et al., 2024; Lombardi et al., 2025). We therefore restrict the present dataset to head-to-tail and disulfide-bridged systems.

D.6 Dataset Statistics and Filtering Outcomes

To quantify the effect of the two quality regimes, we compare the retained dataset sizes under the Strict and Relaxed pipelines across all source subsets (Table 6).

Strict Filtering Outcomes. Under the Strict Quality Pipeline ($\Delta\omega = 10.0°$, disulfide `PASS` only), we obtain a high-fidelity dataset of $73,519$ structurally optimized conformations. This population comprises $58,107$ head-to-tail cyclized peptides and $15,412$ disulfide-bridged structures. Within this regime, the `CPBind` subset contributes the largest share with $39,539$ conformations, followed by `CPTrans` with $30,380$, while `CPCore` and `CPSea_PDB` provide $2,637$ and 963 structures, respectively. This subset is intended for modeling settings that benefit from conservative stereochemical quality control and geometries close to commonly populated conformational regions.

Relaxed Filtering Outcomes. Under the Relaxed Quality Pipeline ($\Delta\omega = 15.0°$, disulfide `PASS` or `WARN`), the dataset scale expands substantially to $129,198$ valid conformations, representing an approximate 75.7% increase in overall data coverage. This expanded dataset contains $76,929$ head-to-tail and $52,269$ disulfide-bridged structures. Across source subsets, `CPBind` and `CPTrans` expand to $74,212$ and $48,541$ conformations, respectively, whereas `CPCore` and `CPSea_PDB` scale to $4,724$ and $1,721$ conformations. These statistics show that allowing moderate, explicitly bounded backbone and side-chain strain recovers a substantial number of otherwise excluded topologically closed conformers, thereby increasing conformational coverage for data-intensive learning applications.

Table 6: Dataset statistics under the Strict and Relaxed quality filtering regimes.

	Strict Quality Pipeline			Relaxed Quality Pipeline		
Sub-Dataset	**Head-to-Tail**	**Disulfide**	**Total**	**Head-to-Tail**	**Disulfide**	**Total**
CPTrans	26,189	4,191	30,380	34,391	14,150	48,541
CPBind	29,149	10,390	39,539	38,897	35,315	74,212
CPCore	1,997	640	2,637	2,583	2,141	4,724
CPSea_PDB	772	191	963	1,058	663	1,721
Total	**58,107**	**15,412**	**73,519**	**76,929**	**52,269**	**129,198**

D.7 Ablation Study

We compare the geometric QC outcomes of samples generated by seven models trained on Strict CPSea or CPCore, using the audit described above (Figure 6). Training on Strict CPSea is associated with lower `FAIL` rates for all seven generators and a reduction of 7.53–12.54 percentage points in the combined `WARN+FAIL` rate. For DC-FLOW, the combined rate decreases from 36.72% to 26.31%. The `WARN` rate also decreases for six generators; PepGLAD is the exception, increasing from 6.42% to 7.98% despite a lower combined non-`PASS` rate.

Here, `WARN` and `FAIL` serve solely as reference indicators of geometric data quality under the adopted audit, not as measures of overall model quality. We therefore interpret these rates within each model across training corpora rather than use them to rank different generators. Since the corpora also differ in size and composition, this comparison does not isolate the effect of filtering alone.

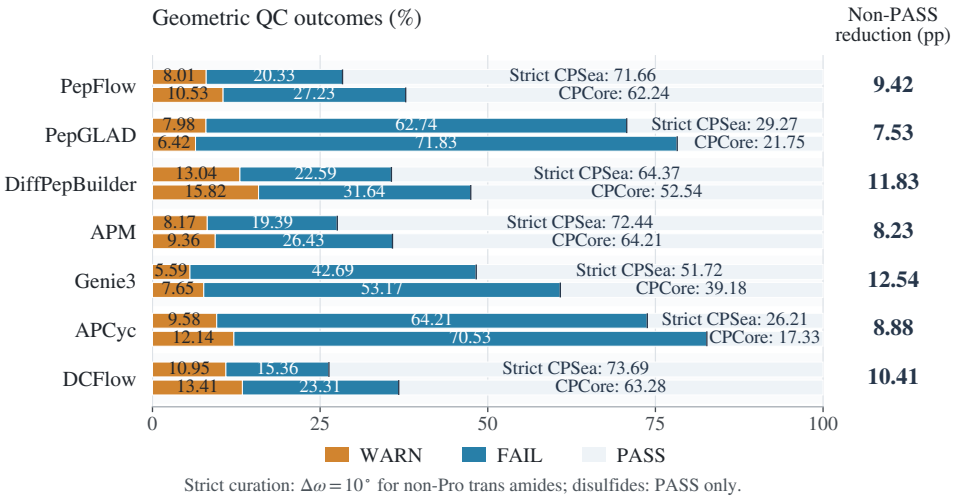

Figure 6: **Geometric QC by training corpus.** Models trained on Strict CPSea show lower FAIL and combined WARN+FAIL rates across all seven generators under the operational geometric audit described in the appendix. Right-hand values report WARN+FAIL reductions relative to CPCore in percentage points; WARN alone increases for PepGLAD.

E Detailed Implementation and Evaluation Metrics

This appendix specifies the model implementation, generation and refinement protocols, and metrics used in the experimental tables. Direct full-atom evaluation assesses generated complexes after a common physical relaxation procedure. Cofolding evaluation instead assesses the sequences produced by each design pipeline through a common structure-prediction protocol.

E.1 Model Training and Sampling

Geometric representation and conditioning. DC-FLOW represents each peptide residue by a rigid frame, a continuous sequence carrier, and torsional variables. Local frames are constructed from the N–Cα–C atoms. The angular representation contains one N–Cα–C–O channel controlling carbonyl orientation and up to four side-chain torsions. Residue templates reconstruct the full heavy-atom structure from the predicted residue identities, frames, and angles. Receptor coordinates and identities remain fixed throughout generation, although receptor features participate in the shared attention network. During training, the native peptide's identities and geometric features are masked in the static conditioning encoder.

The shared trunk contains six geometric update blocks. Each block combines invariant point attention, two Transformer layers, a residue transition, and rigid-frame updates; pair features are updated between blocks. Residue and pair representations have dimensions 128 and 64, respectively. The attention module uses eight heads, eight query/key points and twelve value points per head. The Transformer uses four heads, a feed-forward dimension of 128, and zero dropout. Separate prediction heads produce residue logits and angular endpoints.

Topology and modality activation. Cyclization conditions enter through residual node and pair features encoding the linkage type, endpoint identities, graph distances, and closure edges. These additions retain the original linear positional features. Topology gates and the added clock projection are initialized to zero, allowing initialization from a pretrained PepFlow model (Li et al., 2024). The default activation thresholds are $(\tau_{\mathrm{b}}, \tau_{\mathrm{s}}, \tau_{\mathrm{c}}) = (0, 0.5, 0.8)$. Before activation, peptide sequence inputs use an unknown-residue token and side-chain inputs use fixed zero-angle encodings. Backbone updates continue after both later modalities become active. For terminal disulfide generation, the designated endpoint identities are restricted to cysteine.

A frozen ProteinMPNN (Dauparas et al., 2022) evaluates the current intermediate backbone together with the observed receptor. Its centered log-probabilities enter the native sequence logits through the global gate g, initialized to 0.1. The prior is evaluated without gradients through either its parameters or coordinate inputs. Its contribution is enabled when the sequence clock is positive; sequence supervision starts at $t \geq \tau_{\mathrm{s}}$. The analogous side-chain loss threshold is $t \geq \tau_{\mathrm{c}}$.

Training configuration. Training uses the curated CPSea-derived corpus described in the dataset appendix. The shared generative network is initialized from a PepFlow checkpoint trained on the same corpus and optimized for 30 additional epochs. We use four GPUs with 32 examples per GPU, giving a nominal global batch size of 128. Adam uses coefficients $(0.9, 0.999)$, zero weight decay, and fixed learning rates of 10^{-5} for inherited parameters and 10^{-4} for the added clock, topology, and prior-gate parameters. Global gradient norms are clipped at 100. ProteinMPNN remains frozen. The reported configuration uses the epoch-30 checkpoint.

For each example, the interpolation time is sampled uniformly from $[0.01, 0.99)$. Translation sources are centered Gaussian samples, rotations are sampled uniformly, torsions use uniform circular sources, and sequence carriers use Gaussian noise with scale 5. The endpoint sequence encoding assigns $+5$ to the native category and -5 to other categories. The loss weights for translation, rotation, reconstructed backbone atoms, sequence cross-entropy, angular matching, auxiliary endpoint torsions, HT closure, and SS closure are, respectively, $0.5, 0.5, 0.25, 1.0, 1.0, 0.5, 0.25, 0.25$. Sequence cross-entropy is averaged over peptide residues and then over sequence-active examples. Angular supervision uses the observed and residue-specific torsion masks. SS closure supervision is activated with the side-chain modality.

Numerical generation. The evaluated sampler uses 200 time points spanning $[0.01, 1]$, with categorical sequence draws during the trajectory and a final endpoint prediction. To specify the

Table 7: Settings of the common Rosetta relaxation and interface analysis protocol. Closure restraints act during relaxation and are excluded from the reported energy scores.

Component	Setting
Energy function	`ref2015`
FastRelax repeats	5
Minimizer iteration cap	1,000 per minimization stage
Outputs	One relaxed structure per input
Receptor mobility	Backbone and side chains permitted to relax
Closure-restraint weight	1.0; constraint ramping disabled
Separated-state repacking	Disabled
Interface geometry	Hydrogen bonds, shape complementarity, and Packstat

source-dependent updates, let x_t denote peptide Cα translations and u_t the continuous sequence carrier, with fixed sources x_0, u_0 and predicted endpoints $\widehat{x}_1, \widehat{u}_1$. For a step of length Δt, define $\Delta s_{\text{seq}} = s_{\text{seq}}(t + \Delta t) - s_{\text{seq}}(t)$. The updates are

$$\begin{aligned} x_{t+\Delta t} &= x_t + (\widehat{x}_1 - x_0)\Delta t, \\ u_{t+\Delta t} &= u_t + (\widehat{u}_1 - u_0)\Delta s_{\text{seq}}. \end{aligned} \tag{69}$$

The sequence endpoint is obtained by sampling the predicted categorical distribution and encoding the sampled identity. Rotations and angles use their respective geometric updates, with modality clocks controlling the angular increments. The updated sequence carrier is sampled again to provide categorical input to the next network evaluation. Thus sequence discretization occurs throughout sampling, rather than only at its end.

E.2 Benchmark Conditions and Candidate Budgets

Direct full-atom generation. The LNR benchmark contains 93 receptors: 37 OOD and 56 ID targets under the receptor-sequence-identity split in Section 4.1. All methods receive the same target-specific receptor and peptide length. Each baseline generates 32 candidates per target. DC-FLOW generates 32 HT candidates and 32 terminal-SS candidates per target, reported as separate configurations. Each configuration therefore contains 2,976 candidates in total, including 1,184 OOD and 1,792 ID candidates. HT and SS outputs are not pooled to construct a single DC-FLOW result. DiffPepBuilder and Genie 3 outputs use Rosetta side-chain packing before the common full-atom evaluation, as indicated in the tables. The topology assigned during post-processing is described below; baseline models are not assumed to have identical native topology controls.

Predictor-assisted design. The cofolding benchmark comprises the 12 published HFGuidedDesign targets (Hu et al., 2026). Each reported configuration contributes 32 final candidates per target, or 384 candidates in total. The *Guide* flag denotes generation-time oracle feedback; *AFCyc optim.* denotes subsequent sequence-initialized optimization; and *Search* denotes the complete cascaded policy. Equal final-candidate counts do not imply equal oracle costs: guided proposals, intermediate branches, and refinement updates are separate contributions to inference compute. For DC-FLOW with inference-time search, the search procedure is configured with parameters $K = 16$, $K_0 = 8$, $M = 8$ (yielding $N = 64$), and $N_0 = 32$.

E.3 Physical Relaxation and Topology Handling

Common relaxation protocol. Generated or independently predicted complexes are processed with Rosetta FastRelax using `ref2015` (Alford et al., 2017). For scoring, the peptide and receptor are standardized as chains A and B, respectively, with receptor numbering and coordinates mapped to the corresponding target. The same preparation is applied across methods. FastRelax performs five repeats of repacking and minimization and returns one relaxed output per input. Both peptide and receptor backbone and side-chain degrees of freedom are permitted to relax: the protocol applies neither receptor-specific movement restrictions nor receptor coordinate restraints. The receptor is therefore fixed during neural generation but flexible during this shared physical post-processing step.

Chemical closure. HT structures use a declared bond between the first residue's nitrogen and the last residue's carbonyl carbon. Terminal SS structures use a declared disulfide between the terminal cysteine sulfur atoms, with side-chain repacking enabled. Let d be the closing-atom distance, d_0 its target length, and σ the restraint width. The harmonic contribution is

$$E_{\text{closure}} = \left(\frac{d - d_0}{\sigma} \right)^2 . \tag{70}$$

For HT closure, $(d_0, \sigma) = (1.33, 0.02)$ Å; for SS closure, $(d_0, \sigma) = (2.03, 0.03)$ Å. The atom-pair restraint has weight 1.0 throughout relaxation. No additional closure-specific angle, planarity, or dihedral restraints are introduced; the remaining geometry is governed by the declared chemical topology and the standard score function.

Topology fidelity and unsuccessful outputs. For DC-FLOW, relaxation follows the requested HT or SS topology. An output that changes the requested linkage type is counted as a failure. For baseline configurations processed without a prescribed linkage type, terminal-cysteine sequences are first attempted as SS structures; unsuccessful SS attempts can fall back to HT closure. Other sequences are processed as HT structures. The SS-to-HT fallback is triggered when SS closure cannot be completed, including a relaxed sulfur distance greater than $d_0 + 3\sigma$. A baseline output is accepted only if a permitted closure attempt succeeds; failed outputs remain in the generation denominator. This fallback policy is part of the evaluation adapter, rather than a claim about each baseline's intrinsic capabilities.

Unrestrained scoring. After relaxation, closure-restraint contributions are excluded from the reported `ref2015` scores. Total-complex energy supplies the Stability metric. InterfaceAnalyzer evaluates the peptide–receptor interface with separated-state repacking disabled; hydrogen-bond, surface-area, shape-complementarity, and packing calculations are enabled as required by the reported tables. Its cross-interface score, `dG_cross`, and separated-state score, `dG_separated`, are distinct quantities. The former sums cross-interface energy contributions; the latter compares the bound and separated states. These empirical energies are expressed in Rosetta energy units (REU).

E.4 DIRECT FULL-ATOM GENERATION METRICS

The full-atom tables report Cyc., Stability, D_{seq}, D_{struct}, and Rosetta dG. Following the benchmark convention adapted from APCyc (Zhao et al., 2026), all quantities are evaluated within each receptor before aggregation over the specified All, OOD, or ID split.

Cyclization success candidates. A candidate is considered successfully cyclized if the prescribed cyclization bond is successfully formed under the topology-handling protocol described above in equation 70. Cyc. is defined as the percentage of all generated candidates satisfying this bond-length criterion, with all unsuccessful outputs included in the denominator.

Energetic metrics. Stability is the total `ref2015` score of the relaxed complex without closure-restraint contributions. The Rosetta dG column uses the interface `dG_cross` score. Lower values are preferred for both metrics. Each is averaged over valid candidates for a receptor, and the resulting receptor-wise means are averaged across the split. These are conditional energy summaries of valid outputs, rather than unconditional estimates of binding affinity.

Sequence and structural diversity. Valid candidates for a receptor are clustered separately in sequence and structure space using single-linkage clustering. Sequence distance is one minus sequence similarity, with threshold 0.4; structural distance is peptide Cα RMSD, with threshold 4.0 Å. For a receptor with $n_{\text{valid}} > 0$ valid candidates, let c_{seq} and c_{struct} be the respective numbers of clusters. The reported diversities are

$$D_{\text{seq}} = \frac{c_{\text{seq}}}{n_{\text{valid}}}, \qquad D_{\text{struct}} = \frac{c_{\text{struct}}}{n_{\text{valid}}}. \tag{71}$$

Higher values indicate less redundancy within the evaluated candidate set. These cluster fractions are computed independently for each receptor and then averaged over the reported split. Cyc. retains all attempted samples in its denominator, whereas these diversity and energy summaries use the valid subset.

E.5 Cofolding Evaluation and Reported Metrics

Independent structure prediction. Final discrete sequences are independently evaluated with five AlphaFold-Multimer-v3 models and three recycles per model using cyclic peptide positional encoding. The input comprises the receptor, its MSA, the peptide sequence, and the designated cyclization configuration. Generated peptide coordinates are not supplied as structural templates. Prediction metrics are first averaged over the five models for each candidate. These independent evaluations, rather than scores from intermediate optimization states, determine the confidence tables.

Prediction confidence. We report ipTM, the predicted interface TM-score, and peptide pLDDT, the average per-residue confidence over the peptide chain on a 0–100 scale. Interface iPAE summarizes predicted alignment error across the peptide–receptor interface; hotspot iPAE restricts the receptor-side evaluation to the designated hotspot residues. Both raw iPAE values are divided by 31 before reporting or inclusion in composite scores. We denote these normalized quantities by $\widetilde{\mathrm{iPAE}}_{\mathrm{interface}}$ and $\widetilde{\mathrm{iPAE}}_{\mathrm{hotspot}}$. Higher ipTM and peptide pLDDT, and lower iPAE, indicate greater predicted confidence.

Composite scores. Let S denote the interface-based score used in the main comparison, and S_{hotspot} the hotspot-based score used in the search comparison. They are defined by

$$\begin{aligned} S &= \mathrm{ipTM} + \frac{\mathrm{pLDDT}_{\mathrm{pep}}}{100} + 1 - \widetilde{\mathrm{iPAE}}_{\mathrm{interface}}, \\ S_{\mathrm{hotspot}} &= \mathrm{ipTM} + \frac{\mathrm{pLDDT}_{\mathrm{pep}}}{100} + 1 - \widetilde{\mathrm{iPAE}}_{\mathrm{hotspot}}. \end{aligned} \tag{72}$$

Here $\mathrm{pLDDT}_{\mathrm{pep}}$ is peptide pLDDT. The tables use these sums without an additional factor of $1/3$. Hotspot iPAE is reported alongside the other metrics in the main comparison but does not enter its interface-based score S.

Strict hits. A candidate is a strict hit if its model-averaged metrics satisfy all three conditions:

$$\mathrm{ipTM} \geq 0.7, \qquad \mathrm{pLDDT}_{\mathrm{pep}} \geq 70, \qquad \widetilde{\mathrm{iPAE}}_{\mathrm{interface}} \leq 0.3. \tag{73}$$

The count and percentage are computed over all 384 final candidates. The same interface-iPAE threshold is used in the search comparison, even when its ranking score is S_{hotspot}.

Mean and Best summaries. Mean averages the candidate-level metrics over all 384 sequences. For Best, one candidate is selected from the 32 candidates for each target using the relevant composite score: S in the main comparison and S_{hotspot} in the search comparison. All Best columns then summarize these same 12 selected candidates. In particular, Best does not select a different sequence for each metric or a different AlphaFold model for each candidate. The main and search comparisons use separate random-seed sets; their baseline values need not coincide.

Rosetta interface metrics. Independent cofolded structures are relaxed with the common protocol in Appendix E.3. Let E_{bound} be the relaxed complex score and $E_{\mathrm{separated}}$ the score after separating the peptide and receptor without repacking. The cofolding energy column reports

$$\Delta G_{\mathrm{sep}} = E_{\mathrm{bound}} - E_{\mathrm{separated}} = \texttt{dG_separated}. \tag{74}$$

Buried solvent-accessible surface area is denoted ΔSASA (the `dSASA_int` output, in Å^2). Area-normalized energy is $100\Delta G_{\mathrm{sep}}/\Delta\mathrm{SASA}$, in REU per 100 Å^2. Lower raw and normalized energies are preferred. The remaining table columns are interface hydrogen-bond count (`hbonds_int`), buried unsatisfied polar groups (`delta_unsatHbonds`, labeled *Unsat. H-bonds*, measuring the change upon interface formation), surface shape complementarity (SC), and interface packing quality (Packstat). Fewer buried unsatisfied polar groups, and higher hydrogen-bond counts, SC, and Packstat are preferred in the reported comparisons. Buried surface area measures interface extent and is reported separately from its energy density.

Table 8: Generation-time AfCycDesign guidance settings for DC-FLOW. Position and alternative limits apply per candidate and per event.

Setting	Value
Base guidance scale	50.0
Maximum probed positions	3
Alternatives per position	3
Number of events	4
Event steps	125, 150, 175, 199 of the 200-point schedule
Maximum evaluations per event	9 single-residue alternatives plus 1 unchanged sequence
Online predictor	One AlphaFold-Multimer-v3 model, one recycle
Bias retention at each event	0.9

Physical-score aggregation. The Rosetta analysis uses four sampled candidates per target for the Mean summaries and the AF-selected Top-1 candidate for Best. Each candidate's Rosetta metrics are averaged across its five independently predicted and relaxed models, followed by aggregation within and across targets. AF Top-1 selection uses the corresponding confidence composite, rather than selecting the lowest Rosetta energy retrospectively. Thus the Rosetta summaries and the confidence tables have different evaluation sample counts, as stated in the interface-analysis figure. Predictor confidence and Rosetta scores serve as computational design criteria, not measurements of experimental affinity.

E.6 GENERATION-TIME GUIDANCE AND AFCYCDESIGN REFINEMENT

Oracle-guided sequence proposals. The external guide used by DC-FLOW and HFGuidedDesign (Hu et al., 2026) evaluates current peptide sequences and their single-residue alternatives with a cyclic AfCycDesign oracle (Rettie et al., 2025a), following the sequence-guidance principle of HFGuidedDesign. For efficiency, this online oracle uses `model_1_multimer_v3`, one recycle, cyclic positional offsets of type 2, the receptor MSA, and the designated hotspots. It receives sequences rather than the generated peptide coordinates, and its predicted coordinates do not overwrite the generative state.

Guidance is applied at four events during the 200-point trajectory: sampling steps 125, 150, 175, and 199, corresponding approximately to global times $0.627, 0.751, 0.876, 0.995$. At each event and for each candidate, the three positions with highest categorical entropy are considered. At each selected position, up to three alternative amino acids are selected from the model's high-probability categories, excluding the current identity. Every proposal changes only one residue. Thus an event evaluates at most nine alternatives, together with the unchanged sequence; the setting does not denote three simultaneous mutations. The guidance base scale is $\lambda_{\text{guide}} = 50.0$.

Score-dependent probability bias. For the online guide, S_{hotspot} in equation 72 is computed from the online oracle. Let $e_{\text{base}} = |3 - S^{\text{base}}_{\text{hotspot}}|$ and $e_{\text{mut}} = |3 - S^{\text{mut}}_{\text{hotspot}}|$ be the score gaps for the current sequence and a single-residue alternative. At global time t, the proposed residue receives the additive bias

$$\Delta b = \begin{cases} \lambda_{\text{guide}} t \dfrac{e_{\text{base}} - e_{\text{mut}}}{\max(e_{\text{base}}, 0.1)}, & e_{\text{mut}} < e_{\text{base}}, \\ -0.3\lambda_{\text{guide}} t \dfrac{e_{\text{mut}} - e_{\text{base}}}{\max(e_{\text{mut}}, 0.1)}, & \text{otherwise.} \end{cases} \tag{75}$$

Here Δb is accumulated at the tested position and amino-acid category. Existing accumulated biases are multiplied by 0.9 at each guidance event before new contributions are added. The base scale therefore controls score-dependent logit adjustments; it is neither an optimizer learning rate nor a direct multiplier of backbone coordinates.

The accumulated bias modifies subsequent sequence logits, the updated continuous sequence carrier, and the final sequence prediction. Its addition to the sequence carrier is not multiplied by the local time increment. The biased categorical sequence re-enters the shared trunk at subsequent steps and can thereby affect later backbone and side-chain updates. This is a discrete proposal-and-evaluation guide, without differentiation through the online folding predictor.

Post-generation optimization. In the fixed-budget comparisons, configurations marked *AFCyc optim.=Yes* initialize AfCycDesign with the generated sequence and run the specified 160-update refinement protocol using ipTM and hotspot-iPAE objectives. Configurations marked *AFCyc optim.=No* proceed directly to independent final evaluation. This optimization stage is distinct from the four generation-time guidance events, and the generated coordinates are not used as pose initializations. Final discrete sequences are reevaluated with the five-model, three-recycle protocol above before confidence aggregation.

Cascaded allocation and early stopping. The search configuration uses the integrated policy of Section 3.3: K backbone prefixes are probabilistically reduced to K_0 roots, each root produces M guided continuations, and $N = K_0 M$ proposals receive five initial AfCycDesign updates targeting ipTM. Of these, $N_0 < N$ are probabilistically retained using the empirical screening score $Q = 0.75(1 - \widetilde{\mathrm{iPAE}}_{\mathrm{hotspot}}) + 0.25\,\mathrm{ipTM}$. The retained candidates receive further optimization, with early stopping after five consecutive soft-stage updates without improvement in the composite score reported by AfCycDesign. Thus early retention uses Q, whereas stopping uses the optimizer's reported composite score. The resulting discrete sequences undergo the same independent final evaluation as other configurations. Search does not require every proposal to exhaust the fixed-budget refinement schedule. Oracle calls during guidance, the five-update screen, subsequent refinement, and independent evaluation all contribute to its total computational cost.

F Additional Experiment Results

Table 9: **Direct full-atom cyclic peptide generation on LNR.** DC-FLOW(SS) achieves the highest cyclization success across All/OOD/ID, reaching 99.1% overall, while both DC-FLOW variants reach 100% on ID receptors. DC-FLOW(HT) achieves the best overall Stability and Rosetta dG scores, together with the highest sequence diversity on All and OOD. DC-FLOW(SS) leads structural diversity across all three splits. Bold and underlined values indicate the best and second-best results, respectively, within each split.

Method	Cyc. (%) ↑			Stability ↓			Rosetta dG ↓		
	All	OOD	ID	All	OOD	ID	All	OOD	ID
PepFlow	94.0	92.9	94.6	-737.76	**-961.10**	-590.27	-20.23	-24.78	-17.24
PepGLAD	91.4	97.5	87.4	-635.44	-720.20	-579.45	-14.33	-15.59	-13.50
DiffPepBuilder†	93.7	92.2	94.6	-719.75	-939.25	-574.79	-19.09	-20.66	-18.05
APM	98.9	97.3	**100.0**	-645.42	-869.51	-497.36	-19.90	-25.06	-16.50
Genie 3†	82.4	83.6	81.6	-530.93	-726.29	-399.28	-11.73	-14.07	-10.16
APCyc	82.5	95.4	73.9	-654.07	-713.21	**-616.07**	-21.07	-22.79	**-19.91**
DC-FLOW(HT)	99.0	97.5	**100.0**	**-737.87**	-940.05	-604.28	**-22.82**	**-28.37**	-19.14
DC-FLOW(SS)	**99.1**	**97.6**	**100.0**	-627.74	-807.36	-509.06	-17.69	-21.17	-15.39

Method	D_{seq} ↑			D_{struct} ↑		
	All	OOD	ID	All	OOD	ID
PepFlow	0.646	0.686	0.619	0.781	0.948	0.670
PepGLAD	0.881	0.895	**0.872**	0.862	0.932	0.817
DiffPepBuilder†	0.349	0.446	0.286	0.830	0.979	0.732
APM	0.141	0.173	0.120	0.827	0.938	0.754
Genie 3†	0.247	0.253	0.242	0.807	0.823	0.796
APCyc	0.845	0.867	0.831	0.762	0.849	0.705
DC-FLOW(HT)	**0.883**	**0.940**	0.845	0.648	0.934	0.459
DC-FLOW(SS)	0.389	0.668	0.205	**0.962**	**0.991**	**0.943**

† DiffPepBuilder and Genie 3 require the Rosetta side-chain packer.

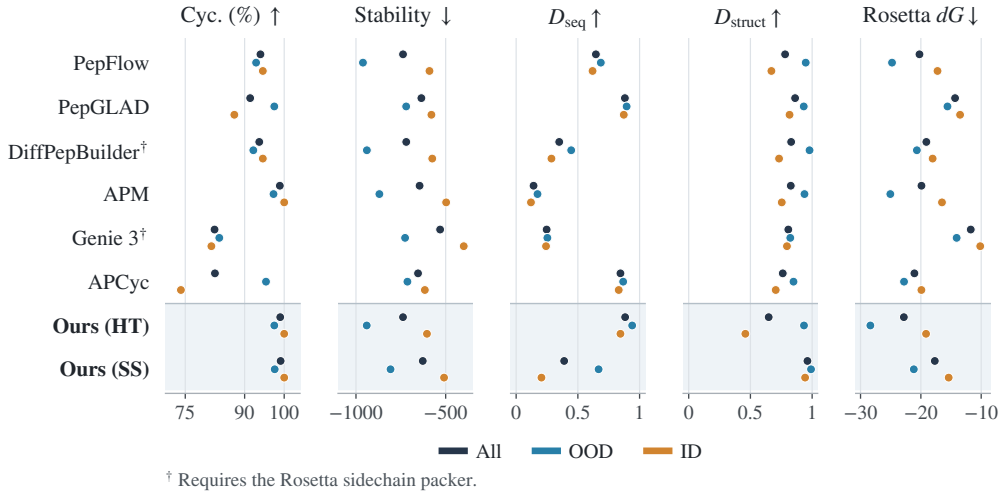

† Requires the Rosetta sidechain packer.

Table 10: **Detailed cofolding performance across 12 targets.** Each configuration contains 384 candidates (32 per target), independently evaluated with five AlphaFold-Multimer-v3 models and three recycles. Mean aggregates all candidates; Best averages the target-wise Top-1 candidates selected by composite score.

(a) Predicted complex quality

Method	Guide	AFCyc optim.	ipTM ↑		Peptide pLDDT ↑		Interface iPAE ↓	
			Mean	Best	Mean	Best	Mean	Best
RFPeptides	–	No	0.516	0.763	47.97	71.42	0.434	0.262
	–	Yes	0.666	0.877	61.21	84.90	<u>0.334</u>	0.170
HFGuidedDesign	No	No	0.470	0.627	41.65	57.39	0.477	0.365
	No	Yes	0.662	0.886	60.27	85.57	0.341	0.167
	Yes	No	0.519	0.720	46.38	67.36	0.442	0.295
	Yes	Yes	<u>0.674</u>	<u>0.889</u>	<u>61.60</u>	<u>86.57</u>	<u>0.334</u>	<u>0.161</u>
DC-FLOW	No	No	0.448	0.634	39.76	58.29	0.493	0.367
	No	Yes	0.648	0.886	58.81	85.99	0.352	0.165
	Yes	No	0.502	0.705	44.69	65.60	0.454	0.306
	Yes	Yes	**0.683**	**0.903**	**62.50**	**87.26**	**0.326**	**0.146**

(b) Hotspot alignment, composite score, and strict hits

Method	Guide	AFCyc optim.	Hotspot iPAE ↓		Composite S ↑		Strict hits ↑	
			Mean	Best	Mean	Best	Count	Rate (%)
RFPeptides	–	No	0.318	0.182	1.562	2.215	17/384	4.4
	–	Yes	<u>0.234</u>	0.115	1.944	2.556	112/384	29.2
HFGuidedDesign	No	No	0.369	0.283	1.410	1.836	2/384	0.5
	No	Yes	0.238	0.105	1.923	2.575	<u>128</u>/384	<u>33.3</u>
	Yes	No	0.336	0.223	1.541	2.099	9/384	2.3
	Yes	Yes	<u>0.234</u>	0.104	<u>1.957</u>	<u>2.593</u>	<u>128</u>/384	<u>33.3</u>
DC-FLOW	No	No	0.376	0.268	1.353	1.851	4/384	1.0
	No	Yes	0.248	<u>0.103</u>	1.885	2.581	93/384	24.2
	Yes	No	0.332	0.212	1.495	2.055	14/384	3.6
	Yes	Yes	**0.228**	**0.091**	**1.982**	**2.629**	**141/384**	**36.7**

Notes. Bold and underlined values indicate the best and second-best results across the ten configurations, respectively, separately for each numeric column; ties share the same rank. Best uses the same 12 candidates for every metric: for each target, select the candidate with the highest composite score among 32, then average across targets. These are not independently selected metric-wise extrema or a global maximum. Each candidate's metrics are averaged over five models. Guide denotes generation-time guidance; AFCyc optim. denotes subsequent sequence-initialized gradient optimization. DC-FLOW uses hotspot persistent-transition guidance. Both iPAE metrics are normalized by 31; hotspot iPAE is not included in S. Strict hits are counted over all 384 candidates and require $\mathrm{ipTM} \geq 0.7$, peptide pLDDT ≥ 70, and interface iPAE ≤ 0.3. Optimization-trajectory scores and pose-initialized arms are excluded.

Table 11: **Detailed Rosetta interface metrics for cofolding across 12 targets.** Mean and Best are shown separately for each of the ten generation and optimization configurations.

(a) Interface energetics and buried surface area

Method	Guide	AFCyc optim.	$\Delta G_{\text{sep}} \downarrow$ (REU)		$100\Delta G_{\text{sep}}/\Delta\text{SASA} \downarrow$ (REU/100 Å^2)		$\Delta\text{SASA} \uparrow$ (Å^2)	
			Mean	Best	Mean	Best	Mean	Best
RFPeptides	–	No	-30.289	-32.745	-2.824	-3.079	1062.6	1071.8
	–	Yes	-35.891	-40.984	-3.170	-3.494	1123.4	1174.1
HFGuidedDesign	No	No	-31.437	-34.868	-2.758	-2.942	1128.1	1166.5
	No	Yes	-34.938	-40.634	-3.037	-3.412	1138.1	1186.4
	Yes	No	-34.007	-37.760	-2.952	-3.150	**1149.0**	1198.8
	Yes	Yes	-35.882	**-42.806**	-3.121	-3.447	1145.7	**1249.5**
DC-FLOW	No	No	-29.865	-37.771	-2.831	-3.435	1045.1	1101.8
	No	Yes	-35.083	-40.063	-3.078	-3.365	1133.3	1202.0
	Yes	No	-32.730	-36.181	-2.982	-3.338	1087.8	1082.9
	Yes	Yes	**-38.386**	-42.298	**-3.396**	**-3.524**	1091.5	1154.6

(b) Hydrogen bonding and interface geometry

Method	Guide	AFCyc optim.	Interface H-bonds ↑		Unsat. H-bonds ↓		SC ↑		Packstat ↑	
			Mean	Best	Mean	Best	Mean	Best	Mean	Best
RFPeptides	–	No	4.083	5.100	7.367	6.183	0.653	0.685	0.635	0.640
	–	Yes	4.604	5.267	6.650	6.567	0.684	0.689	0.644	0.654
HFGuidedDesign	No	No	3.488	4.550	7.779	6.950	0.663	0.668	0.638	0.642
	No	Yes	4.729	5.550	7.250	8.683	0.685	0.701	0.646	**0.666**
	Yes	No	3.992	4.767	7.742	8.267	0.661	0.683	0.641	0.658
	Yes	Yes	4.592	5.383	7.142	6.467	0.686	**0.709**	0.648	0.661
DC-FLOW	No	No	3.358	4.667	7.875	7.550	0.664	0.698	0.644	0.655
	No	Yes	4.354	5.200	7.454	7.483	0.681	0.702	0.641	0.657
	Yes	No	3.946	4.700	7.650	6.550	0.672	0.698	0.635	0.645
	Yes	Yes	**5.658**	**6.050**	**5.433**	**5.200**	**0.702**	0.699	**0.649**	**0.666**

Notes. Bold and underlined values indicate the best and second-best results across the ten configurations, respectively, separately for each numeric column; ties share the same rank. Guide denotes generation-time guidance; AFCyc optim. denotes subsequent sequence-initialized gradient optimization.

G Additional Ablation Results

Table 12: **Ablations on LNR.** The default $\tau_s = 0.5$ achieves the highest overall cyclization success within each topology and the best HT Stability and Rosetta dG scores. Freezing backbone motion substantially worsens Stability for both topologies, while removing the MPNN prior weakens interface scores and removing the delay produces topology-dependent effects. Earlier activation at $\tau_s = 0.2$ improves overall sequence and structure diversity, indicating a trade-off between diversity and interface scores. All/OOD/ID correspond to 93/37/56 receptors.

Inference setting	Topo.	Cyc. (%) ↑			Stability ↓			Rosetta dG ↓		
		All	OOD	ID	All	OOD	ID	All	OOD	ID
w/o coupling	HT	98.5	97.1	99.5	-109.23	112.60	-255.79	-9.59	-4.54	-12.92
	SS	98.6	97.1	99.6	-217.19	-167.87	-249.78	-10.97	-7.97	-12.95
w/o MPNN prior	HT	98.8	97.3	99.8	-715.25	-914.12	-583.86	-21.42	-27.56	-17.36
	SS	98.7	97.1	99.7	-624.71	-808.71	-503.13	-17.37	-18.88	-16.37
w/o delay, $\tau_s = 0$	HT	98.4	97.1	99.3	-716.55	-905.72	-591.57	-18.46	-20.40	-17.17
	SS	98.7	97.2	99.7	-657.66	-838.82	-537.96	-17.99	-19.61	-16.92
$\tau_s = 0.2$	HT	98.9	**97.6**	99.8	-720.54	-913.84	-592.83	-19.12	-23.38	-16.30
	SS	98.9	97.5	99.9	-632.35	-812.85	-513.08	-17.64	-19.59	-16.36
$\tau_s = 0.5$	HT	99.0	97.5	**100.0**	**-737.87**	**-940.05**	**-604.28**	**-22.82**	**-28.37**	**-19.14**
	SS	**99.1**	**97.6**	**100.0**	-627.74	-807.36	-509.06	-17.69	-21.17	-15.39

Inference setting	Topo.	D_{seq} ↑			D_{struct} ↑		
		All	OOD	ID	All	OOD	ID
w/o coupling	HT	0.869	0.937	0.824	0.636	0.933	0.439
	SS	0.382	0.639	0.213	0.957	0.989	0.936
w/o MPNN prior	HT	0.875	0.933	0.836	0.634	0.948	0.427
	SS	0.417	0.639	0.271	0.941	0.948	0.936
w/o delay, $\tau_s = 0$	HT	0.871	0.934	0.829	0.647	0.940	0.453
	SS	0.398	0.668	0.219	0.947	0.982	0.924
$\tau_s = 0.2$	HT	**0.918**	**0.940**	**0.903**	0.727	0.934	0.590
	SS	0.391	0.672	0.205	**0.966**	**1.000**	**0.943**
$\tau_s = 0.5$	HT	0.883	**0.940**	0.845	0.648	0.934	0.459
	SS	0.389	0.668	0.205	0.962	0.991	**0.943**

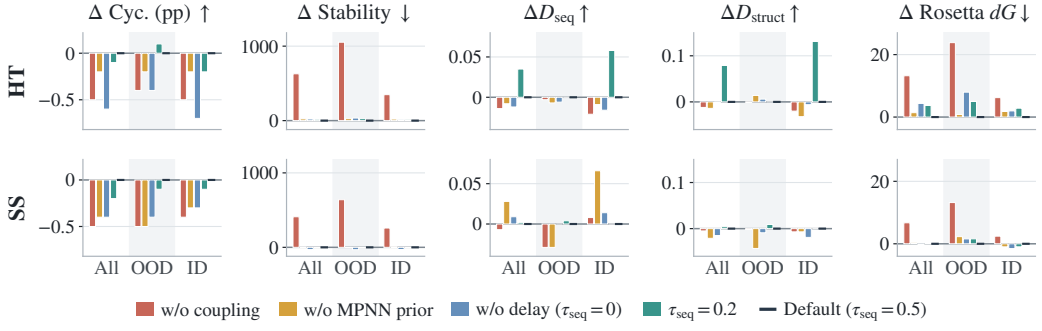

Table 13: **Detailed performance of DC-FLOW with and without search across 12 targets.** Both configurations use generation-time guidance and AFCyc optimization. This experiment reports the hotspot-based composite score S_{hotspot}.

(a) Prediction confidence

Method	Guide	AFCyc optim.	Search	ipTM ↑		Peptide pLDDT ↑	
				Mean	Best	Mean	Best
DC-FLOW	Yes	Yes	No	<u>0.651</u>	<u>0.904</u>	<u>57.90</u>	<u>87.92</u>
	Yes	Yes	Yes	**0.849**	**0.937**	**79.91**	**94.69**

(b) Hotspot score and strict hits

Method	Guide	AFCyc optim.	Search	Hotspot iPAE ↓		S_{hotspot} ↑		Strict hits ↑	
				Mean	Best	Mean	Best	Count	Rate (%)
DC-FLOW	Yes	Yes	No	<u>0.241</u>	<u>0.091</u>	<u>1.989</u>	<u>2.693</u>	<u>125/384</u>	<u>32.6</u>
	Yes	Yes	Yes	**0.132**	**0.061**	**2.516**	**2.823**	**290/384**	**75.5**

Notes. Bold and underlined values indicate the best and second-best results across the two configurations, respectively, separately for each numeric column; ties share the same rank. Best averages 12 candidates: select the highest-composite candidate among 32 for each target, then average its metrics across targets. Strict hits are counted over all 384 candidates and require $\text{ipTM} \geq 0.7$, peptide pLDDT ≥ 70, and interface iPAE ≤ 0.3.

Table 14: **Detailed Rosetta interface metrics for DC-FLOW with and without Search across 12 targets.** Both configurations use generation-time guidance and AFCyc optimization.

(a) Interface energetics and buried surface area

Method	Guide	AFCyc optim.	Search	ΔG_{sep} ↓ (REU)		$100\Delta G_{\text{sep}}/\Delta\text{SASA}$ ↓ (REU/100 Å^2)		ΔSASA ↑ (Å^2)	
				Mean	Best	Mean	Best	Mean	Best
DC-FLOW	Yes	Yes	No	<u>-38.386</u>	<u>-42.298</u>	<u>-3.396</u>	<u>-3.524</u>	<u>1091.5</u>	<u>1154.6</u>
	Yes	Yes	Yes	**-38.841**	**-43.514**	**-3.567**	**-3.724**	**1107.4**	**1180.5**

(b) Hydrogen bonding and interface geometry

Method	Guide	AFCyc optim.	Search	Interface H-bonds ↑		Unsat. H-bonds ↓		SC ↑		Packstat ↑	
				Mean	Best	Mean	Best	Mean	Best	Mean	Best
DC-FLOW	Yes	Yes	No	**5.658**	**6.050**	**5.433**	**5.200**	<u>0.702</u>	<u>0.699</u>	<u>0.649</u>	<u>0.666</u>
	Yes	Yes	Yes	<u>5.156</u>	<u>5.500</u>	<u>5.807</u>	<u>5.583</u>	**0.729**	**0.730**	**0.664**	**0.678**

Notes. Bold and underlined values indicate the best and second-best results across the two configurations, respectively, separately for each numeric column; ties share the same rank.

H CHALLENGING TL1A-TRIMER CASE

We study TL1A as a challenging target for cyclic-peptide design against an oligomeric protein–protein interface. TL1A (TNFSF15) is a TNF-superfamily ligand recognized by the signaling receptor DR3 and the soluble decoy receptor DcR3 (Migone et al., 2002). The TL1A–DcR3 crystal structure reveals a trimeric TL1A core bound by three DcR3 molecules (Zhan et al., 2011). Our objective is to design cyclic peptides that occupy the DcR3-binding surface of trimeric TL1A, with the intended function of competitively blocking TL1A–DcR3 association. Given the trimeric target and designated interface hotspots, the task is to generate a peptide sequence and cyclic conformation that can engage this surface in the assembled complex.

This case combines several geometric and modeling challenges. The designated hotspots span two TL1A chains, so recognizing the target requires a peptide pose compatible with a composite surface whose geometry depends on the relative arrangement of neighboring subunits. The peptide must establish contacts across this surface while avoiding steric clashes with the surrounding trimer. Meanwhile, the long target chains produce a large total residue context, requiring precise peptide localization within a substantially larger assembly. Cyclization adds a further constraint: the peptide must adopt a closed conformation while positioning its interacting residues to engage hotspots on both chains. Together, these requirements make the case a demanding test of whether cofolding can jointly resolve peptide conformation, interface placement, and compatibility with an oligomeric target.

For all evaluated configurations, we design cyclic peptides of length 12 residues and generate 12 candidate cyclic peptides per model. For DC-FLOW with inference-time search, the search procedure is configured with parameters $K = 8$, $K_0 = 4$, $M = 6$ (yielding $N = 24$), and $N_0 = 12$.

Table 15 compares RFPeptides, HFGuidedDesign, DC-FLOW, and DC-FLOW with inference-time search. Across all methods, ipTM values remain closely clustered at a high level. However, this uniformity obscures substantial uncertainty in peptide folding: all three baseline configurations without search yield consistently low peptide pLDDT, indicating that a favorable target-wide ipTM alone does not guarantee a well-resolved peptide structure. In contrast, incorporating inference-time search leads to pronounced improvements across the board, yielding the highest peptide pLDDT, the lowest interface and hotspot iPAE, and the strongest hotspot scores under both Mean and AF Top-1 summaries.

These trends demonstrate that inference-time search effectively guides the model toward more confident and geometrically consistent solutions on this multi-chain interface. Nevertheless, the absolute peptide confidence remains modest even after search, reflecting the inherent difficulty of the oligomeric target. While these computational gains indicate improved structural plausibility, experimental binding and competition assays will be essential to verify functional TL1A–DcR3 blockade.

Table 15: **Cofolding results on the challenging TL1A-trimer case.** Each cell reports Mean / AF Top-1. Search improves peptide confidence, interface iPAE, and hotspot-specific metrics, while mean ipTM remains similar across methods.

Method	ipTM ↑	Peptide pLDDT ↑	Interface iPAE ↓	Hotspot iPAE ↓	Hotspot S_{hotspot} ↑
RFPeptides	**0.868** / 0.873	26.82 / 30.38	0.808 / 0.725	0.751 / 0.623	1.385 / 1.555
HFGuidedDesign	0.867 / 0.868	25.52 / 27.44	0.818 / 0.790	0.749 / 0.705	1.374 / 1.438
DC-FLOW	**0.868** / 0.873	26.96 / 33.42	0.810 / 0.689	0.753 / 0.583	1.384 / 1.624
DC-FLOW (Search)	0.864 / **0.877**	**38.41 / 49.79**	**0.677 / 0.632**	**0.576 / 0.524**	**1.673 / 1.851**

Bold and underline indicate the best and second-best values, respectively, separately for each metric and summary; ties share the same rank. AF Top-1 follows the AF composite-score selection used above and is not selected independently for each metric.

I VISUALIZATION

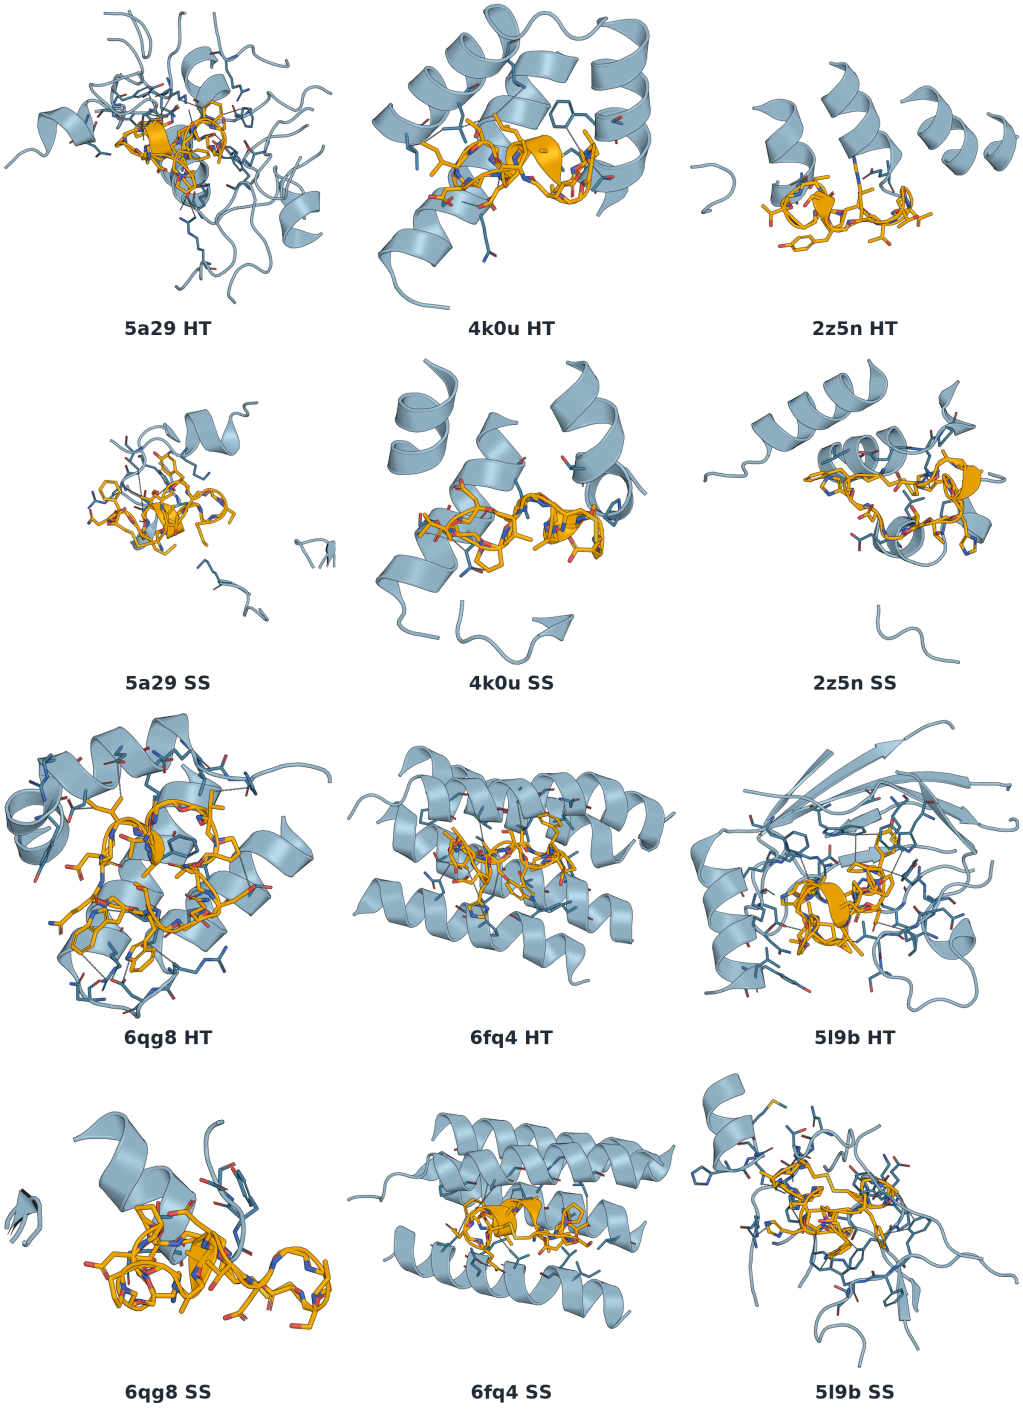

Figure 7: **Cyclic peptide designs generated directly by DC-FLOW across 6 targets and 2 cyclization modes in LNR.** Local interface views show the target in light blue and the designed peptide in orange. Dashed lines indicate hydrogen-bond/polar contacts (green), salt bridges (magenta), and hydrophobic contacts (gray), inferred from relaxed model geometry.

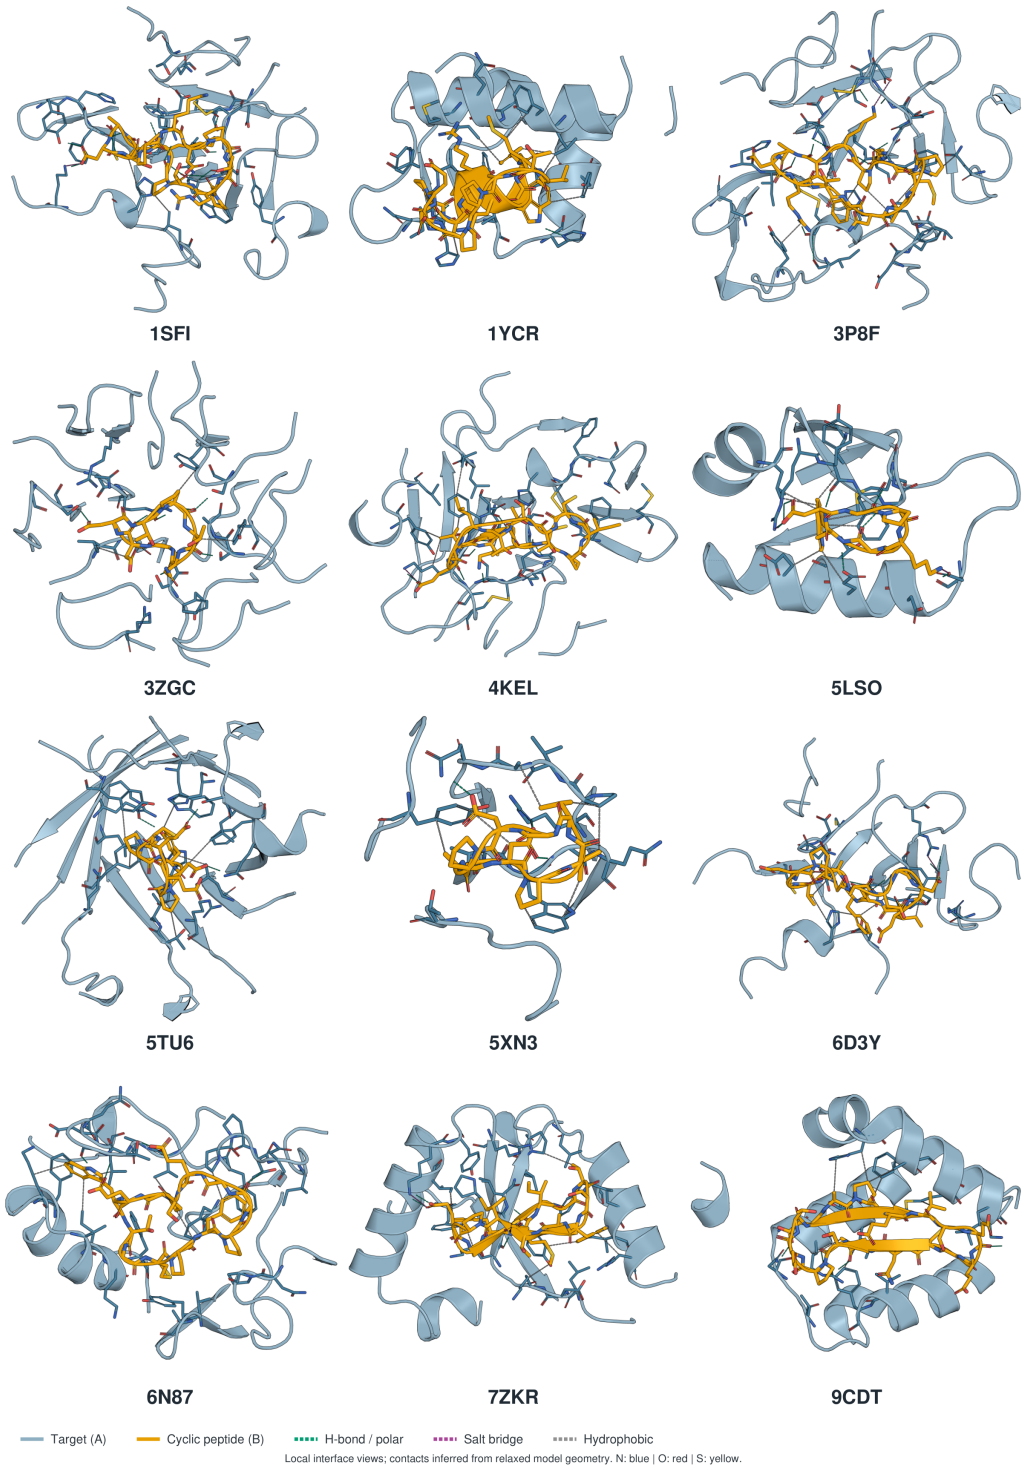

Figure 8: **Cyclic peptide designs generated by DC-FLOW (Search) across 12 targets in HFGuidedDesign test set.** Local interface views show the target in light blue and the designed peptide in orange. Dashed lines indicate hydrogen-bond/polar contacts (green), salt bridges (magenta), and hydrophobic contacts (gray), inferred from relaxed model geometry.