PEPRIG: RECEPTOR-GUIDED SEQUENCE–STRUCTURE COUPLING FOR ALL-ATOM PEPTIDE GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Peptides have diverse biological activities that support applications in targeted therapy, functional foods, and nutritional health. These activities depend on molecular interactions shaped by peptide sequence and structure, making their joint design important for developing peptides with desired functions. Generative models support joint peptide sequence–structure design, yet receptor compatibility requires residue identities and spatial arrangements to jointly establish complementary interactions at the binding interface. The central challenge is to coordinate these interdependent choices throughout generation, so that sequence and structure adapt together to the target receptor. To address this challenge, we introduce PepRIG, a receptor-guided all-atom diffusion framework that couples sequence and structure generation at the representation and potential levels. Specifically, Interface-Guided Co-Adaptation (IGCA) integrates multiscale receptor-interface information into shared representations, guiding sequence and structure updates with a common description of peptide–receptor interactions. Complementing this shared representation, Reciprocal Binding Potential (RBP) derives coupled gradients from a receptor-conditioned scalar potential, allowing each modality's correction to account for the other modality. Extensive experiments on LNR and PepBDB demonstrate that PepRIG outperforms strong baselines on key evaluation criteria, achieving state-of-the-art results across both benchmarks. Ablation studies demonstrate the complementary benefits of the two mechanisms, supporting shared receptor guidance for coupled sequence–structure generation.

1 INTRODUCTION

Peptides are short chains of amino-acid residues that participate in diverse biological processes through interactions with other molecules, particularly target proteins known as receptors (Petsalaki & Russell, 2008; Kaspar & Reichert, 2013). Through these molecular interactions, peptides exert biological effects that support their use in targeted therapy, functional foods, and nutritional health (Muttenthaler et al., 2021; Udenigwe & Aluko, 2012). Translating this potential into practical applications requires peptides with the desired biological activities. These activities depend on how peptide sequences and three-dimensional structures determine interactions with biological targets (London et al., 2010). Designing sequences and structures jointly is therefore important for developing peptides with target-specific functions.

To support such joint design, receptor-conditioned generative models learn to generate peptide sequences and structures using information about the target binding site (Li et al., 2024). Recently, diffusion models have been used to formulate this receptor-conditioned generation as an iterative denoising process, using binding-site information to guide the refinement of peptide representations (Ho et al., 2020; Wang et al., 2024). Methods such as PepGLAD and PepTri apply this approach to jointly generate residue identities and all-atom structures while promoting sequence–structure consistency (Kong et al., 2024; Nguyen et al., 2026). However, internal sequence–structure consistency is insufficient to ensure compatibility with the target receptor, which requires sequence and structure to jointly support complementary interactions at the binding interface.

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.

Meeting this requirement is challenging because residue identity and spatial arrangement impose interdependent constraints on these interactions. A residue substitution changes the available chemical groups and can alter peptide conformations and interactions within the binding site (Zhao et al., 2014). Conversely, moving a peptide residue relative to the receptor changes its interaction partners and the chemical properties favored at that position. Effective coordination therefore requires more than access to a common receptor context. Sequence updates need to account for the current geometry of the binding interface, while structural updates need to accommodate the interaction requirements of the peptide sequence. This motivates using shared information about peptide–receptor interactions to guide both updates throughout generation.

Building on this insight, we propose Peptide Generation with Receptor Interface Guidance (PepRIG), an all-atom diffusion framework for receptor-conditioned peptide design. PepRIG couples sequence and structure updates at two complementary levels. Specifically, at the representation level, Interface-Guided Co-Adaptation (IGCA) integrates local interaction environments and broader pocket context into shared multiscale representations that guide both updates. This grounds sequence and structure refinement in a common description of peptide–receptor interactions. To complement this shared information with coordinated correction directions, Reciprocal Binding Potential (RBP) defines a receptor-conditioned scalar potential over the joint semantic and geometric state. Gradients of this shared potential provide coupled corrections, allowing sequence refinement to account for the current geometry and structural refinement to reflect sequence-dependent interaction requirements. Together, these mechanisms treat receptor compatibility as a joint requirement on sequence and structure, shaping how the two modalities co-evolve throughout denoising.

The main contributions of this work are summarized as follows.

- We introduce PepRIG, an all-atom diffusion framework that treats receptor compatibility as a joint requirement on peptide sequence and structure. The framework uses shared receptor-interface guidance to coordinate their co-evolution throughout generation.
- We develop two complementary coupling mechanisms, IGCA and RBP, that coordinate the interface information used by both modalities and their correction directions. IGCA constructs shared multiscale interface representations, while RBP derives coupled corrections from a common receptor-conditioned scalar potential.
- Extensive experiments on LNR and PepBDB demonstrate the strong performance of PepRIG across both benchmarks. Ablation studies highlight the complementary contributions of IGCA and RBP, supporting receptor-guided sequence–structure coupling as an effective approach to all-atom peptide generation.

2 Related Work

Peptide Design Computational peptide design seeks to identify sequences and bound structures compatible with a target receptor (Bryant & Elofsson, 2023; Chen et al., 2024). Target-conditioned approaches such as PepMLM and EvoBind2 use protein sequence information to design peptide binders (Chen et al., 2026; Li et al., 2025). Complementary structure-aware methods exploit binding-site geometry or known interaction patterns, as exemplified by DiffPepBuilder and PepMimic (Wang et al., 2024; Kong et al., 2026). Related approaches extend structure-prediction-guided optimization to cyclic peptides (Rettie et al., 2025a) and diffusion-based design to protein-binding macrocycles (Rettie et al., 2025b). Peptide docking and complex-scoring models complement design by predicting receptor-bound poses and assessing their interaction quality (Zhao et al., 2025; Tao et al., 2025). Within this landscape, joint generative approaches directly model peptide sequences and all-atom structures in the receptor context. PepGLAD introduces geometric latent diffusion, compressing variable-size residues into fixed-dimensional representations and normalizing binding-site geometry through receptor-specific affine transformations (Kong et al., 2024). PepFlow instead uses multimodal flow matching to jointly generate residue identities, backbone geometry and side-chain angles (Li et al., 2024). More recently, PepTri integrates physical, evolutionary and mutual-information guidance into latent diffusion to promote structural plausibility and sequence–structure coherence (Nguyen et al., 2026). These developments establish a foundation for receptor-conditioned all-atom peptide generation. Building on this foundation, PepRIG focuses on how shared information about peptide–receptor interactions coordinates sequence and structure updates throughout denoising, making receptor compatibility a joint requirement of their evolution.

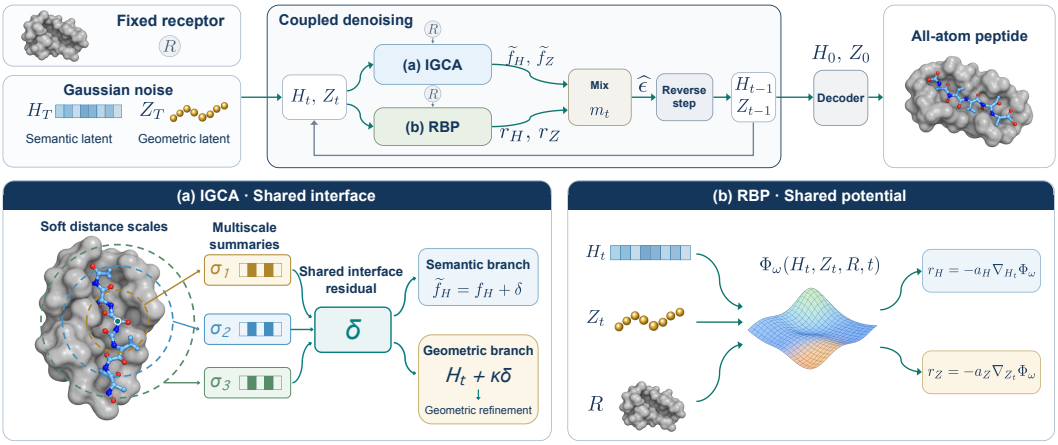


Figure 1: **Overview of PepRIG.** H_t, Z_t are semantic and geometric latents at step t, with receptor $\mathcal{R}$ fixed. In (a), Gaussian scales σ_k (dashed circles) yield a shared interface residual δ that guides both branches, producing noise predictions $\widetilde{f}_H, \widetilde{f}_Z$. In (b), negative gradients of a shared scalar potential Φ_ω produce coupled corrections r_H, r_Z. The weight m_t combines these predictions and corrections into the reverse-step noise estimate $\widehat{\epsilon}$.

Geometric Generative Models for Biomolecular Design Geometric generative models provide a foundation for molecular design by incorporating spatial symmetries into the generation of three-dimensional structures (Satorras et al., 2021; Hoogeboom et al., 2022). Genie and FrameDiff develop equivariant diffusion models for protein backbones, while RFdiffusion and Chroma extend diffusion-based generation to diverse protein design tasks (Lin & AlQuraishi, 2023; Yim et al., 2023; Watson et al., 2023; Ingraham et al., 2023). Backbone-first pipelines can subsequently assign amino-acid sequences through inverse-folding models (Hsu et al., 2022), including ProteinMPNN (Dauparas et al., 2022). To capture dependencies between residue identity and geometry during generation, joint models instead couple sequence and structure through equivariant transformations, multimodal flows or sequence-space diffusion (Shi et al., 2023; Campbell et al., 2024; Lisanza et al., 2025). All-atom approaches, including AbDiffuser and Protpardelle, further incorporate the side-chain geometry needed to describe molecular interactions (Martinkus et al., 2023; Chu et al., 2024). Alongside richer molecular representations, gradient guidance steers diffusion using differentiable objectives (Guo et al., 2024), and studies of score-field conservativeness distinguish scalar-derived gradients from unconstrained vector predictions (Chao et al., 2023). PepRIG brings these two aspects of generative modeling together in receptor-conditioned peptide design. Shared interface representations provide common information for both modalities, while a receptor-conditioned scalar potential couples their correction fields.

3 Method

Given a receptor pocket $\mathcal{R}$ and peptide length L, PepRIG generates an amino-acid sequence $S = (s_1, \ldots, s_L)$ and its all-atom bound conformation X. A frozen autoencoder represents the peptide through semantic latents H and geometric latents Z, following geometric latent peptide design (Kong et al., 2024; Nguyen et al., 2026). Generation starts from Gaussian peptide latents and iteratively denoises both blocks while keeping the receptor fixed. As illustrated in Figure 1, PepRIG coordinates these updates through two complementary mechanisms. IGCA supplies a common interface representation to both prediction branches, and RBP links their correction fields through a shared scalar potential. The final latent pair is decoded into sequence and structure.

3.1 Interface-Guided Co-Adaptation

The receptor interface provides both local interaction information and broader pocket context, whose relevance changes with the peptide's position. A single pooled summary can blur this distance-dependent organization of the interface. IGCA therefore computes a multiscale interface residual from

the current noisy geometry and shares it across semantic and geometric predictions. Recomputing this residual at each step lets both branches adapt to the evolving peptide–receptor interaction context while the receptor remains fixed.

For peptide residue i, let q_i contain its semantic, time and positional features, and let q_j denote the receptor context feature. We use all peptide–receptor residue pairs within a complex and measure distance d_{ij} in the backbone's receptor-normalized frame. At scale σ_k, the receptor feature summary, weighted occupancy and mean distance are

$$w_{ij}^{(k)} = \exp\left(-\frac{d_{ij}^2}{2\sigma_k^2}\right), \qquad m_i^{(k)} = \sum_j w_{ij}^{(k)}, \tag{1}$$

$$c_i^{(k)} = \frac{\sum_j w_{ij}^{(k)} q_j}{\max(m_i^{(k)}, \varepsilon)}, \qquad \bar{d}_i^{(k)} = \frac{\sum_j w_{ij}^{(k)} d_{ij}}{\max(m_i^{(k)}, \varepsilon)}. \tag{2}$$

Here ε is a small positive constant that stabilizes normalization. The summaries distinguish receptor features from neighborhood density and proximity, with smaller scales emphasizing immediate surroundings and larger scales preserving wider pocket context. We combine pooled and scale-specific summaries into a single interface residual

$$\begin{aligned}\delta_i = {} & g_{\text{pool}}\big(q_i, \bar{c}_i, \log(1+\bar{m}_i), \log(1+\bar{d}_i)\big) \\ & + g_{\text{scale}}\left(q_i \,\|\, \Big\|_{k=1}^{K_s} [c_i^{(k)}, \log(1+m_i^{(k)}), \log(1+\bar{d}_i^{(k)})]\right).\end{aligned} \tag{3}$$

Here K_s is the number of distance scales, and $\|$ and $\|$ denote feature concatenation. The pooled quantities $\bar{c}_i$, $\bar{m}_i$ and $\bar{d}_i$ use the same summary operations with scale-averaged edge weights. Both g functions are multilayer perceptrons that produce semantic-space residuals. The pooled route captures overall context, while the scale-specific route preserves its distance-dependent organization for shared refinement.

Let f_H and f_Z denote the backbone noise predictions, B_Z its coordinate output before subtracting Z_t, and G_Z an equivariant residual coordinate encoder. IGCA adds δ directly to semantic prediction and injects it into the peptide features driving geometric refinement

$$\widetilde{f}_H = f_H + \delta, \tag{4}$$

$$\widetilde{f}_Z = f_Z + \rho\big[G_Z(q^\delta, Z_t, \mathcal{R}) - B_Z(q, Z_t, \mathcal{R})\big], \tag{5}$$

$$q_i^\delta = [H_{t,i} + \kappa\delta_i, q_i^{\text{other}}]. \tag{6}$$

The vector q_i^{other} contains the unchanged time and positional features. The collection q^δ updates peptide semantic features while retaining these features and all receptor entries. The factor ρ scales the coordinate residual $G_Z - B_Z$, and κ controls semantic conditioning of G_Z. We set $\kappa = \kappa_{\max}\tanh(\kappa_{\text{raw}})$, where $\kappa_{\max}$ is a prescribed bound and κ_{raw} is learned. Initializing G_Z from the backbone and the new multiscale head at zero lets shared refinement begin from the backbone prediction.

3.2 Reciprocal Binding Potential

IGCA coordinates the interface information used by both branches, while RBP links their correction directions through a scalar function of the joint latent state. A semantic change alters a residue's interaction preferences, and a displacement changes its spatial relationship to the receptor. Evaluating both changes through one function makes each correction depend on the other modality. This shared construction couples the responses of the two correction fields during denoising.

For peptide residue i, let $\mathcal{N}_R(i)$ index its receptor neighbors in the RBP interaction graph. Each pair combines peptide semantic latent $H_{t,i}$, receptor feature $H_{R,j}$, diffusion-time embedding $\tau(t)$ and radial basis features of d_{ij}. A shared network maps these features to scalar contributions, averaged over receptor neighbors and summed over peptide residues

$$\Phi_\omega(H_t, Z_t, \mathcal{R}, t) = \sum_{i=1}^{L} \frac{1}{|\mathcal{N}_R(i)|} \sum_{j\in\mathcal{N}_R(i)} \varphi_\omega([H_{t,i}, H_{R,j}, \tau(t), \text{RBF}(d_{ij})])\,. \tag{7}$$

Here φ_ω is the pairwise network with parameters ω, and RBF denotes radial basis encoding. The time embedding is $\tau(t) = [\beta_t, \sin\beta_t, \cos\beta_t]$, with β_t denoting the stepwise noise variance. Neighbor averaging normalizes each peptide residue's contribution by its receptor edge count. Differentiating this latent compatibility potential with respect to both latent blocks yields

$$r_H = -a_H \nabla_{H_t} \Phi_\omega, \qquad r_Z = -a_Z \nabla_{Z_t} \Phi_\omega. \tag{8}$$

The positive gains a_H and a_Z set semantic and geometric correction strengths, respectively, and masks restrict updates to valid generated peptide entries. Because both fields originate from one smooth scalar, their cross-modal sensitivities on a fixed interface graph satisfy

$$\frac{\partial(a_H^{-1} r_H)}{\partial Z_t} = \left[\frac{\partial(a_Z^{-1} r_Z)}{\partial H_t}\right]^\top. \tag{9}$$

This relation links the sensitivity of semantic correction to geometry with the sensitivity of geometric correction to semantics. RBP imposes this reciprocal structure on its auxiliary fields alongside the backbone's independently learned noise predictions. IGCA thus shares the interface representation, and RBP couples the cross-modal responses used to refine the joint state.

3.3 Coupled denoising and training

At each reverse step, we evaluate both modules on the current latent pair and combine their predictions with a common time-dependent coefficient

$$\widehat{\epsilon}_H = (1 - m_t)\widetilde{f}_H + m_t r_H, \qquad \widehat{\epsilon}_Z = (1 - m_t)\widetilde{f}_Z + m_t r_Z, \qquad m_t = m_0(1 - \beta_t)^p. \tag{10}$$

Here $\widehat{\epsilon}_H$ and $\widehat{\epsilon}_Z$ are the final noise predictions, and β_t follows the cosine noise schedule. The parameters m_0 and p control the mixing scale and its time dependence, attenuating the potential contribution at high-noise stages. These predictions drive the reverse transition to the next latent pair. The interface summaries and potential derivatives are then recomputed from the updated state, providing shared receptor guidance throughout generation.

Training corrupts each clean latent block as $V_t = \sqrt{\bar{\alpha}_t} V_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_V$ for $V \in \{H, Z\}$ and $\epsilon_V \sim \mathcal{N}(0, I)$, where $\bar{\alpha}_t$ is cumulative signal retention. The subscript V selects the corresponding semantic or geometric prediction. The overall training objective combines denoising supervision, direct potential-field supervision and the backbone's auxiliary calibration losses

$$\mathcal{L} = \mathcal{L}_{\text{diff}} + \lambda_{\text{pot}} \mathcal{L}_{\text{pot}} + \mathcal{L}_{\text{aux}}. \tag{11}$$

Here $\mathcal{L}_{\text{diff}}$ supervises the mixed noise predictions, while $\mathcal{L}_{\text{pot}}$ trains the scalar-derived correction fields against the injected noise. The term $\mathcal{L}_{\text{aux}}$ retains the backbone's decoded sequence and geometry calibration losses. The first two terms are defined as

$$\mathcal{L}_{\text{diff}} = \mathbb{E} \sum_{V \in \{H,Z\}} \lambda_V \left\| \widehat{\epsilon}_V - \epsilon_V \right\|_M^2,$$
$$\mathcal{L}_{\text{pot}} = \mathbb{E} \sum_{V \in \{H,Z\}} \lambda_V \left\| r_V - \epsilon_V \right\|_M^2.$$

The expectation averages over training complexes, sampled diffusion steps and injected noise. The mask M selects valid generated entries, and $\|\cdot\|_M^2$ denotes their mean squared error. The weights λ_H and λ_Z balance semantic and geometric errors, while λ_{pot} controls direct supervision of the potential fields. This supervision remains active even when the mixing coefficient m_t is small. With the autoencoder frozen, training differentiates through the potential's input gradients, whereas generation requires only first-order gradients.

4 Experiments

4.1 Experimental Setup

We evaluate generation quality on two complementary benchmarks and then examine the mechanisms underlying the observed gains. Cross-dataset evaluation trains on 4,157 PepBench complexes with

Table 1: Generation performance on LNR and PepBDB. Throughout the tables, bold marks the best value per metric within each comparison.

Dataset	Method	GDT↑	F1↑	DockQ↑	LRMSD↓	Div.↑
LNR	PepGLAD	0.7148	0.7936	0.6073	1.2147	**0.9138**
LNR	PepFlow	0.7368	0.7947	0.5364	1.0648	0.8265
LNR	PepTri	0.7401	0.8283	0.6282	1.1811	0.9062
LNR	PepRIG	**0.7544**	**0.8557**	**0.6531**	**1.0540**	0.9046
PepBDB	PepGLAD	0.5167	0.5148	0.4372	1.9136	0.9143
PepBDB	PepFlow	0.6368	0.7046	0.4167	**1.2648**	0.6659
PepBDB	PepTri	0.6373	0.7493	**0.5117**	1.4435	**0.9231**
PepBDB	PepRIG	**0.6578**	**0.7773**	0.5044	1.3539	0.9153

114 validation complexes and tests on 93 LNR targets (Tsaban et al., 2022; Nguyen et al., 2026). Within-dataset evaluation uses 7,014 training, 323 validation and 142 test complexes from PepBDB. We compare PepRIG with PepTri, PepGLAD and PepFlow using 40 candidates per target. The primary results are unrelaxed, with molecular relaxation assessed separately. Component ablations and independent-head replacements on LNR examine the two forms of coupling, followed by candidate-budget, joint-quality and randomness analyses.

The metrics separate sequence recovery, internal peptide structure and receptor-bound placement. Sliding-AAR (AAR) measures offset-aligned sequence recovery, and GDT-TS (GDT) measures globally aligned backbone agreement. Contact-F1 (F1) measures intra-peptide Cα contacts within 8 Å, excluding sequence separations of at most two residues. DockQ evaluates bound pose and interface quality. Local RMSD (LRMSD), reported in Å, averages independently aligned five-residue windows, using the whole chain if shorter. Sequence Diversity (Div.) is the sequence-cluster count divided by 40, using single linkage at distance 0.4 with normalized BLOSUM62 similarity. Best-of-K takes each target's maximum separately for recovery, GDT-TS, Contact-F1 and DockQ and its minimum for Local RMSD. Target scores are averaged, except that Local RMSD is reported as the median across targets for all methods. Mean-candidate scores average within targets before averaging across targets.

The backbone uses width 128, three geometric layers and 100 reverse steps. IGCA uses normalized scales $\{0.3, 0.5, 0.8\}$ over all interface pairs with denominator floor 10^{-6}. RBP uses two width-128 SiLU layers, four zero-initialized output channels and 16 evenly spaced Gaussian radial centers over $[0, 1]$. Channel outputs are summed and normalized by the square root of the channel count. Generator training freezes the autoencoder and uses AdamW with weight decay 0.01 and gradient clipping at 1. Learning rates are 5×10^{-7} on LNR and 6×10^{-6} on PepBDB. Generation uses $(m_0, p, a_H, a_Z) = (0.005, 32, 1, 20)$ and $(0.02, 8, 0.5, 8)$, respectively. LNR training uses constant mixing 0.005 and gains $(1, 16)$, while PepBDB retains its generation settings. Potential-loss weights are 0.05 and 0.1. Primary generation uses seed 12.

4.2 Comparison with Baseline Methods

In Table 1, PepRIG improved global backbone agreement, internal contact recovery and local geometry over PepTri on both benchmarks. On LNR, Contact-F1 increased from 0.8283 to 0.8557 and DockQ from 0.6282 to 0.6531, while Local RMSD decreased from 1.1811 to 1.0540 Å. These gains combine improved internal peptide organization with more accurate receptor-relative placement. Diversity remained close to PepTri at 0.9046 versus 0.9062, indicating that the structural gains did not require a large reduction in sequence variety.

PepBDB showed improved internal peptide organization without a corresponding gain in receptor-relative placement. GDT-TS increased from 0.6373 to 0.6578, Contact-F1 from 0.7493 to 0.7773, and Local RMSD decreased from 1.4435 to 1.3539 Å. However, DockQ decreased from 0.5117 to 0.5044 and diversity from 0.9231 to 0.9153. This distinction shows that improved local geometry does not necessarily translate into a more accurate bound pose.

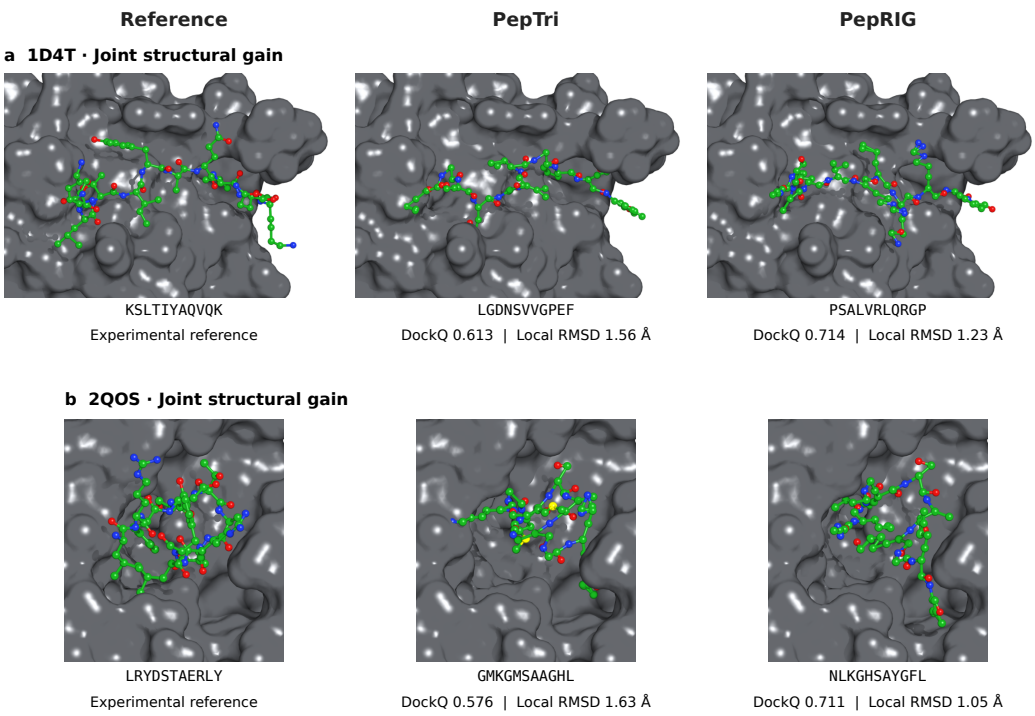


Figure 2: **Qualitative comparison on two selected LNR targets.** Gray surfaces show receptor pockets. Peptide C, N, O and S atoms are green, blue, red and yellow, respectively. Scores refer to the displayed candidates.

Figure 2 connects these metrics to the spatial organization of individual peptides. In 1D4T, both methods span the exposed receptor surface but produce different chain bends and residue placements. In the 2QOS view, PepTri occupies a more compact central region, whereas PepRIG extends along the visible receptor groove. Higher DockQ and lower Local RMSD in both PepRIG candidates link improved receptor-relative placement to closer local conformations, beyond a rigid change in peptide position. These cases illustrate the joint refinement targeted by receptor-guided coupling, with visualization details in Appendix A.1. The following ablations examine the contributions of the two mechanisms.

4.3 Ablation Studies

Table 2: Component ablations on LNR.

Configuration	AAR↑	GDT↑	F1↑	DockQ↑	LRMSD↓	Div.↑
Backbone	0.3580	0.7366	0.8326	0.6284	1.1921	0.9042
Backbone + IGCA	0.3642	0.7479	0.8482	0.6439	1.1462	0.9010
Backbone + RBP	0.3605	0.7462	0.8420	0.6394	1.1163	**0.9136**
PepRIG (full)	**0.3659**	**0.7544**	**0.8557**	**0.6531**	**1.0540**	0.9046

The ablations in Table 2 vary module inclusion while retaining the same backbone, training settings and candidate budget. IGCA alone gave larger Contact-F1 and DockQ gains than RBP alone, supporting shared interface representations for contact reconstruction and bound-pose refinement. RBP alone achieved lower Local RMSD than IGCA alone, at 1.1163 versus 1.1462 Å, consistent with scalar-derived corrections contributing to local conformational recovery. Combining the modules surpassed either alone on all five recovery, structure and interface metrics, reducing Local RMSD by 0.1381 Å over the backbone. This complementary pattern links shared interface information

Table 3: Shared coupling versus independent prediction heads on LNR.

Configuration	AAR↑	GDT↑	F1↑	DockQ↑	LRMSD↓	Div.↑
Independent interface head	0.3616	0.7471	0.8461	0.6431	1.1578	**0.9096**
Independent correction heads	0.3607	0.7488	0.8377	0.6469	1.1334	0.9093
PepRIG (full)	**0.3659**	**0.7544**	**0.8557**	**0.6531**	**1.0540**	0.9046

Table 4: Joint candidate quality on LNR at AAR ≥ 0.30. Rate differences and pointwise 95% confidence intervals are in percentage points.

LRMSD	DockQ	Candidates / 3,720		Difference	Targets / 93	
($\leq$Å)	($\geq$)	PepTri	PepRIG	[95% CI]	PepTri	PepRIG
2.0	0.23	229	235	$0.16\,[-0.48, 0.86]$	51	51
2.0	0.49	155	174	$0.51\,[0.03, 1.02]$	34	35
1.5	0.49	122	142	$0.54\,[0.11, 1.05]$	21	26

Table 5: Generation quality after FastRelax with shared settings.

	LNR		PepBDB	
Metric	PepTri	PepRIG	PepTri	PepRIG
Sliding-AAR↑	0.3609	0.3659	0.3431	0.3406
GDT-TS↑	0.7383	0.7565	0.6373	0.6578
Contact-F1↑	0.8321	0.8519	0.7493	0.7773
DockQ↑	0.6308	0.6498	0.5103	0.5037
Local RMSD↓	1.1964	1.0407	1.4435	1.3539
Diversity↑	0.9062	0.9046	0.9231	0.9153

with coupled correction directions. RBP alone had the highest diversity, while full-model diversity remained close to the backbone.

The independent-head replacements in Table 3 test whether the form of coupling matters beyond supplying interface information. Replacing IGCA with an independent interface head while retaining RBP reduced Contact-F1 by 0.0096 and increased Local RMSD by 0.1038 Å. Thus, interface-dependent prediction benefited from a shared adjustment across modalities. Replacing RBP with separate correction heads while retaining IGCA reduced Contact-F1 by 0.0180 and increased Local RMSD by 0.0794 Å. RBP explicitly enforces the reciprocal cross-modal relationship in Equation 9 through a shared potential. Its gains over separate correction heads support the value of this coupled design.

4.4 Further Analyses

The candidate-budget analysis in Appendix A.10 shows higher Contact-F1 and DockQ at every evaluated budget from $K = 1$ to $K = 40$. Gains in mean-candidate scores further show that improvement extends beyond the highest-scoring samples, although Local RMSD favors PepTri at $K = 5, 10$. Sections A.5 and A.8 examine variation across targets and peptide lengths.

To test whether the improvements coexist in individual peptides, the analysis in Table 4 requires each candidate to meet joint sequence-recovery, local-geometry and bound-pose thresholds. Target coverage counts targets with at least one successful candidate, with further analysis in Appendix A.6. At AAR ≥ 0.30, Local RMSD ≤ 1.5 Å and DockQ ≥ 0.49, successful candidates increased from 122 to 142 and target coverage from 21 to 26. Under the loosest thresholds, coverage remained unchanged and the rate-difference interval included zero. The stronger separation at stricter thresholds supports improved joint sequence, structure and interface quality in the same candidates. Intervals use 20,000 paired target-bootstrap resamples with seed 20260924, retaining all 40 candidates within each target, and are pointwise rather than multiplicity-adjusted.

Table 6: Training and generation variability on LNR, reported as mean ± sample standard deviation.

	Training repeats	Generation repeats with fixed weights		
Metric	PepTri	PepTri seed 12	PepTri seed 13	PepRIG
Sliding-AAR↑	0.3672 ± 0.0055	0.3640 ± 0.0074	0.3570 ± 0.0050	0.3655 ± 0.0058
GDT-TS↑	0.7332 ± 0.0129	0.7174 ± 0.0018	0.7356 ± 0.0043	0.7463 ± 0.0073
Contact-F1↑	0.7995 ± 0.0202	0.7851 ± 0.0012	0.8255 ± 0.0029	0.8417 ± 0.0146
DockQ↑	0.6088 ± 0.0099	0.6006 ± 0.0016	0.6032 ± 0.0052	0.6416 ± 0.0114
Local RMSD↓	1.2194 ± 0.0416	1.2865 ± 0.0180	1.2118 ± 0.0309	1.1462 ± 0.0878
Diversity↑	0.9045 ± 0.0102	0.9171 ± 0.0027	0.9004 ± 0.0061	0.9091 ± 0.0040

We next examine whether structural refinement preserves the generation advantage, complemented by unrelaxed geometry diagnostics in Appendix A.9. Both methods undergo FastRelax with identical settings and seed rules on all targets, with 40 candidates per target. Sequence identities remain fixed, leaving AAR and diversity unchanged. LNR uses backbone coordinate restraints to the generated structures with standard deviation 0.15 Å, while PepBDB fixes backbones and optimizes side chains. In Table 5, PepRIG retained its LNR advantage, with 13.02% lower Local RMSD and 3.03% higher DockQ. LNR coordinate restraints anchor refinement to each generated conformation, so this comparison evaluates quality after local refinement of the initial structures. PepBDB retained its backbone-metric gains and lower DockQ under side-chain relaxation, which leaves backbone placement unchanged.

Finally, Table 6 separates training variability from sampling variability on all 93 LNR targets. PepTri training seeds 12, 13 and 14 retrain both AE and LDM, with checkpoints selected by validation loss. The single trained PepRIG model exceeded all three runs in GDT-TS, Contact-F1 and DockQ, with lower Local RMSD but AAR below their mean. For sampling variability, generation seeds 12, 13 and 14 evaluate fixed PepTri checkpoints from training seeds 12 and 13 and the selected PepRIG checkpoint. PepRIG exceeded both baselines in mean Contact-F1 and DockQ, at 0.8417 and 0.6416, with lower mean Local RMSD of 1.1462 Å. Its structural and interface gains persisted across these sampling repeats, with greater variation among generated pools.

5 Conclusion

In this study, we introduced PepRIG, a receptor-guided diffusion framework for target-specific all-atom peptide design. PepRIG coordinates sequence and structure generation at both the representation and potential levels. IGCA constructs shared multiscale interface representations to guide semantic and geometric predictions, while RBP derives coupled corrections from a single receptor-conditioned scalar potential. By incorporating both mechanisms throughout denoising, PepRIG aligns sequence and structure refinement with a common receptor-interface context. Evaluations on LNR and PepBDB demonstrate improved peptide structural quality over PepTri, with additional gains in receptor-bound pose quality on LNR. Ablation studies further demonstrate the complementary benefits of the two coupling mechanisms. Future work will extend the framework to flexible receptor environments and experimentally assess the binding properties of generated peptides. Overall, PepRIG advances target-specific peptide design by learning how sequence and structure should co-adapt to the receptor interface.

Use of Generative AI

Generative AI tools assisted with manuscript drafting and language editing, code development and debugging, experimental workflow support, data analysis and visualization preparation. They also supported exploration of methodological options and discussion of experimental results. The authors are responsible for verifying AI-assisted outputs against the implemented methods and experimental evidence and for ensuring the accuracy of the manuscript. Responsibility for the final text, scientific claims and research contributions rests entirely with the authors.

Reproducibility Statement

We will provide the source code, training and inference configurations, evaluation scripts, and experiment logs as anonymous supplementary materials to facilitate reproduction of the reported results. These materials will include the hyperparameter settings and random seeds used in our experiments. Additional evaluation procedures and qualitative analyses are documented in Appendix A.

References

Patrick Bryant and Arne Elofsson. Peptide binder design with inverse folding and protein structure prediction. *Communications Chemistry*, 6:229, 2023.

Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 5453–5512, 2024.

Chen-Hao Chao, Wei-Fang Sun, Bo-Wun Cheng, and Chun-Yi Lee. On investigating the conservative property of score-based generative models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 4076–4095, 2023.

Leo Tianlai Chen, Zachary Quinn, Madeleine Dumas, Christina Peng, Lauren Hong, Moises Lopez-Gonzalez, Alexander Mestre, Rio Watson, Sophia Vincoff, Lin Zhao, Jianli Wu, Audrey Stavrand, Mayumi Schaepers-Cheu, Tian Zi Wang, Divya Srijay, Connor Monticello, Pranay Vure, Rishab Pulugurta, Sarah Pertsemlidis, Kseniia Kholina, Shrey Goel, Matthew P. DeLisa, Jen-Tsan Ashley Chi, Ray Truant, Hector C. Aguilar, and Pranam Chatterjee. Target sequence-conditioned design of peptide binders using masked language modeling. *Nature Biotechnology*, 44(6):1002–1010, 2026.

Sijie Chen, Tong Lin, Ruchira Basu, Jeremy Ritchey, Shen Wang, Yichuan Luo, Xingcan Li, Dehua Pei, Levent Burak Kara, and Xiaolin Cheng. Design of target specific peptide inhibitors using generative deep learning and molecular dynamics simulations. *Nature Communications*, 15:1611, 2024.

Alexander E. Chu, Jinho Kim, Lucy Cheng, Gina El Nesr, Minkai Xu, Richard W. Shuai, and Po-Ssu Huang. An all-atom protein generative model. *Proceedings of the National Academy of Sciences*, 121(27):e2311500121, 2024.

Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J. Ragotte, Lukas F. Milles, B. I. M. Wicky, A. Courbet, R. J. de Haas, N. Bethel, P. J. Y. Leung, T. F. Huddy, S. Pellock, D. Tischer, F. Chan, B. Koepnick, H. Nguyen, A. Kang, B. Sankaran, A. K. Bera, N. P. King, and D. Baker. Robust deep learning-based protein sequence design using ProteinMPNN. *Science*, 378 (6615):49–56, 2022.

Yingqing Guo, Hui Yuan, Yukang Yang, Minshuo Chen, and Mengdi Wang. Gradient guidance for diffusion models: An optimization perspective. In *Advances in Neural Information Processing Systems*, volume 37, pp. 90736–90770, 2024.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 8867–8887, 2022.

Chloe Hsu, Robert Verkuil, Jason Liu, Zeming Lin, Brian Hie, Tom Sercu, Adam Lerer, and Alexander Rives. Learning inverse folding from millions of predicted structures. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 8946–8970, 2022.

John B. Ingraham, Max Baranov, Zak Costello, Karl W. Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M. Lord, Christopher Ng-Thow-Hing, Erik R. Van Vlack, Shan Tie, Vincent Xue, Sarah C. Cowles, Alan Leung, João V. Rodrigues, Claudio L. Morales-Perez, Alex M. Ayoub, Robin Green, Katherine Puentes, Frank Oplinger, Nishant V. Panwar, Fritz Obermeyer, Adam R.

Root, Andrew L. Beam, Frank J. Poelwijk, and Gevorg Grigoryan. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.

Allan A. Kaspar and Janice M. Reichert. Future directions for peptide therapeutics development. *Drug Discovery Today*, 18(17–18):807–817, 2013.

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

Xiangzhe Kong, Rui Jiao, Haowei Lin, Ruihan Guo, Wenbing Huang, Wei-Ying Ma, Zihua Wang, Yang Liu, and Jianzhu Ma. Peptide design through binding interface mimicry with PepMimic. *Nature Biomedical Engineering*, 10(7):1295–1310, 2026.

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, pp. 27615–27640, 2024.

Qiuzhen Li, Efstathios Nikolaos Vlachos, and Patrick Bryant. Design of linear and cyclic peptide binders from protein sequence information. *Communications Chemistry*, 8(1):211, 2025.

Yeqing Lin and Mohammed AlQuraishi. Generating novel, designable, and diverse protein structures by equivariantly diffusing oriented residue clouds. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 20978–21002, 2023.

Sidney Lyayuga Lisanza, Jacob Merle Gershon, Samuel W. K. Tipps, Jeremiah Nelson Sims, Lucas Arnoldt, Samuel J. Hendel, Miriam K. Simma, Ge Liu, Muna Yase, Hongwei Wu, Claire D. Tharp, Xinting Li, Alex Kang, Evans Brackenbrough, Asim K. Bera, Stacey Gerben, Bruce J. Wittmann, Andrew C. McShan, and David Baker. Multistate and functional protein design using RoseTTAFold sequence space diffusion. *Nature Biotechnology*, 43(8):1288–1298, 2025.

Nir London, Dana Movshovitz-Attias, and Ora Schueler-Furman. The structural basis of peptide-protein binding strategies. *Structure*, 18(2):188–199, 2010.

Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, and Andreas Loukas. AbDiffuser: Full-atom generation of in-vitro functioning antibodies. In *Advances in Neural Information Processing Systems*, volume 36, pp. 40729–40759, 2023.

Markus Muttenthaler, Glenn F. King, David J. Adams, and Paul F. Alewood. Trends in peptide drug discovery. *Nature Reviews Drug Discovery*, 20(4):309–325, 2021.

Ngoc-Quang Nguyen, Jaeyoon Jung, Seijung Kim, Sunkyu Kim, and Jaewoo Kang. PepTri: Tri-guided all-atom diffusion for peptide design via physics, evolution, and mutual information. In *The Fourteenth International Conference on Learning Representations*, 2026.

Evangelia Petsalaki and Robert B. Russell. Peptide-mediated interactions in biological systems: new discoveries and applications. *Current Opinion in Biotechnology*, 19(4):344–350, 2008.

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, Mila Lamb, Analisa Murray, Victor Adebomi, Guangfeng Zhou, Frank DiMaio, Sergey Ovchinnikov, and Gaurav Bhardwaj. 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, Alex Kang, Evans Brackenbrough, Mila Lamb, Stacey R. Gerben, Analisa Murray, Paul M. Levine, Maika Schneider, Vibha Vasireddy, Sergey Ovchinnikov, Oliver H. Weiergräber, Dieter Willbold, Joshua A. Kritzer, Joseph D. Mougous, David Baker, Frank DiMaio, and Gaurav Bhardwaj. Accurate de novo design of high-affinity protein-binding macrocycles using deep learning. *Nature Chemical Biology*, 21 (12):1948–1956, 2025b.

Víctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 9323–9332, 2021.

Chence Shi, Chuanrui Wang, Jiarui Lu, Bozitao Zhong, and Jian Tang. Protein sequence and structure co-design with equivariant translation. In *The Eleventh International Conference on Learning Representations*, 2023.

Huanyu Tao, Xiaoyu Wang, and Sheng-You Huang. An interaction-derived graph learning framework for scoring protein–peptide complexes. *Nature Machine Intelligence*, 7(11):1858–1869, 2025.

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*, 13:176, 2022.

Chibuike C. Udenigwe and Rotimi E. Aluko. Food protein-derived bioactive peptides: Production, processing, and potential health benefits. *Journal of Food Science*, 77(1):R11–R24, 2012.

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.

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, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100, 2023.

Jason Yim, Brian L. Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. SE(3) diffusion model with application to protein backbone generation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 40001–40039, 2023.

Baoyu Zhao, Peng Xu, Longguang Jiang, Berit Paaske, Tobias Kromann-Hansen, Jan K. Jensen, Hans Peter Sørensen, Zhuo Liu, Jakob T. Nielsen, Anni Christensen, Masood Hosseini, Kasper K. Sørensen, Niels Christian Nielsen, Knud J. Jensen, Mingdong Huang, and Peter A. Andreasen. A cyclic peptidic serine protease inhibitor: Increasing affinity by increasing peptide flexibility. *PLoS ONE*, 9(12):e115872, 2014.

Huifeng Zhao, Odin Zhang, Dejun Jiang, Zhenxing Wu, Hongyan Du, Xiaorui Wang, Yihao Zhao, Yuansheng Huang, Jingxuan Ge, Tingjun Hou, and Yu Kang. Protein–peptide docking with a rational and accurate diffusion generative model. *Nature Machine Intelligence*, 7(8):1308–1321, 2025.

A Appendix

A.1 Qualitative comparison protocol

The following analyses use unrelaxed LNR test samples from 93 targets, with 40 candidates per target from one generation run per method. PepRIG denotes the complete model. Figures 2 and 3 show three examples selected for joint gains in bound pose and local structure. Each method contributes its highest-DockQ candidate per target. We align entire complexes using receptor Cα atoms and share the camera and crop within each comparison. The reference structure determines camera orientation. Renderings omit reference overlays and hydrogens and use standard residue connectivity without coordinate optimization. Gray surfaces contain the same receptor residues within 12 Å of any displayed peptide. Scores describe the displayed candidates, with Local RMSD computed over independently aligned five-residue windows.

The candidates in Figure 2 combine improved bound-pose recovery with lower Local RMSD. Chemical geometry shows a mixed pattern, with increased backbone bond-length error for 1D4T and remaining interface heavy-atom overlaps in 2QOS. These observations motivate the separate geometry diagnostics in Section A.9.

A.2 Pose recovery with reduced interface overlap

Figure 3 presents 5N85, where improved reference agreement accompanies reduced receptor overlap. DockQ increases from 0.507 to 0.555 and Local RMSD decreases from 1.802 to 1.520 Å. Under the screen in Section A.9, overlapping peptide heavy atoms decrease from 16.67% to 2.94%. Neither candidate has Cα clashes, so the heavy-atom analysis captures a difference beyond the backbone-level screen. This example combines gains in pose, local structure and interface packing.

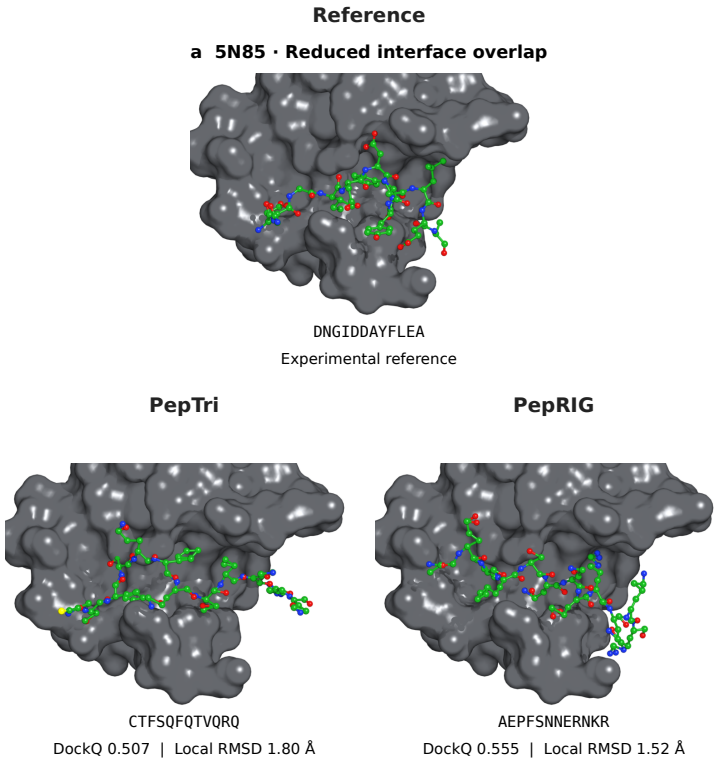


Figure 3: **Receptor-interface packing for the 12-residue peptide in 5N85.** Reference and generated complexes use the selection, alignment, colors and annotations of Figure 2.

A.3 SEQUENCE SUBSTITUTIONS AND RESIDUE-LEVEL POSE ERRORS

Figures 4 and 5 examine an outcome-stratified set comprising 2CNZ, 4W50 and 6MLC, with the largest DockQ gain, median positive gain and largest loss, respectively. Each method contributes its highest-DockQ candidate for each target. In 2CNZ, PepRIG reduces receptor-aligned displacement over much of the central peptide segment, while the N-terminal residues remain displaced. In 4W50, the generated sequences differ at positions 7 and 11, and the structural differences extend beyond these substitutions. In 6MLC, larger C-terminal errors accompany the lower DockQ. Together, the sequence alignments and residue-level profiles locate the changes underlying both improved and degraded bound poses. These receptor-aligned errors measure bound-pose displacement, unlike Local RMSD, which independently aligns local windows.

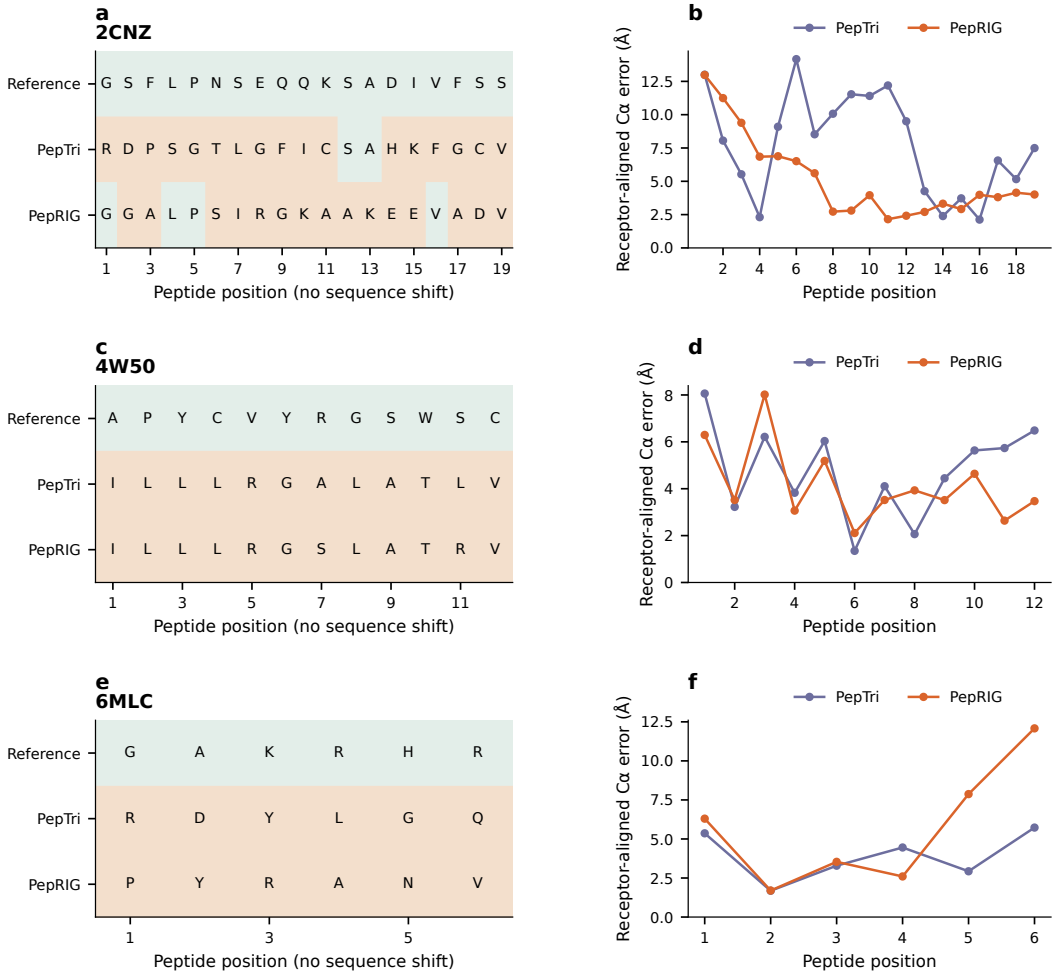


Figure 4: **Sequence and pose changes for 2CNZ, 4W50 and 6MLC.** Left panels align original residue indices, with matches in green and mismatches in beige. Right panels show peptide Cα Euclidean displacements after receptor alignment, with one trace per candidate.

A.4 RECEPTOR–PEPTIDE CONTACT PATTERNS

Figure 5 relates pose differences to receptor interactions through residue-pair contacts defined by a heavy-atom distance within 4.5 Å. Each row retains all receptor residues contacting any displayed structure and uses fixed peptide positions, making recovered, missing and additional contacts directly comparable. Interface-contact F1 is the harmonic mean of precision and recall against reference contacts, whereas benchmark Contact-F1 measures intra-peptide Cα contacts.

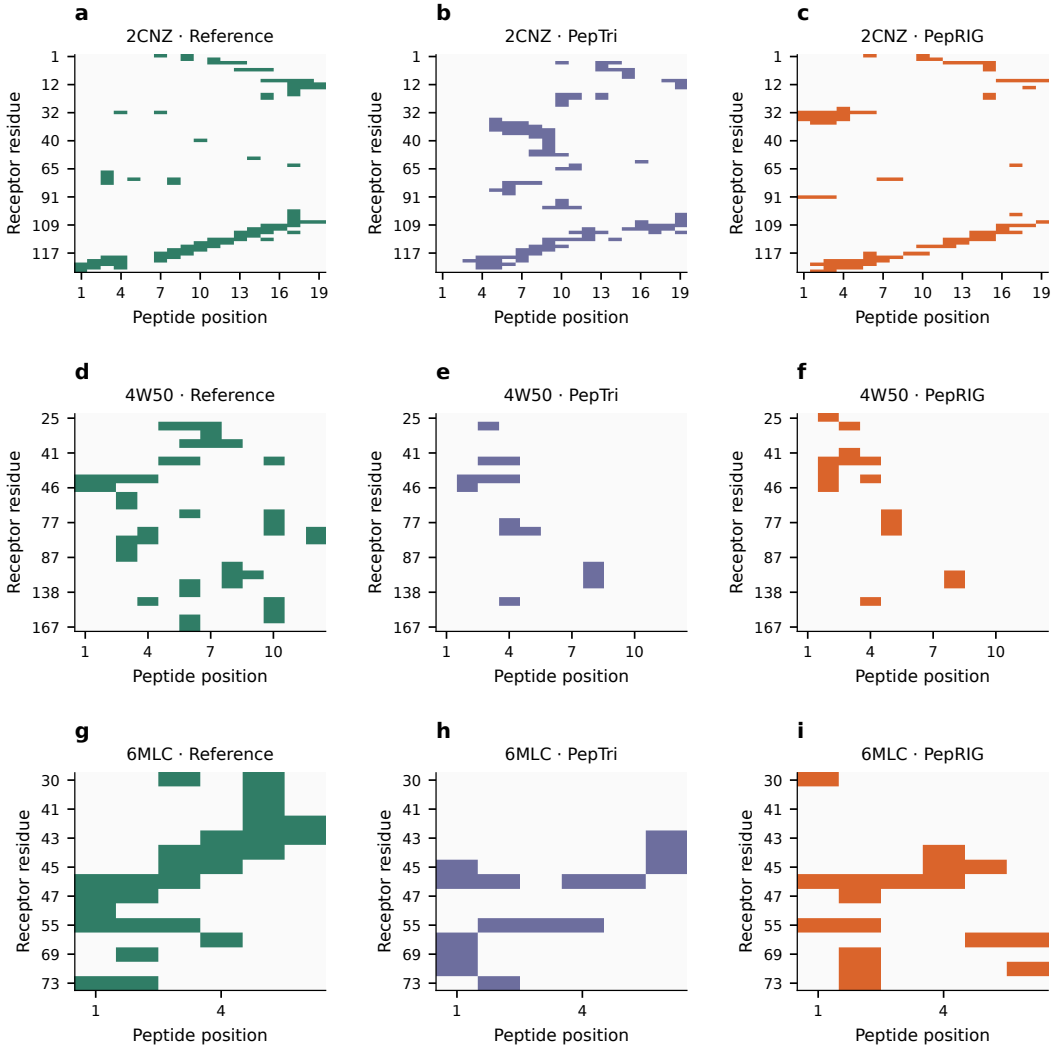


Figure 5: **Receptor–peptide contact maps.** Colored cells mark heavy-atom contacts within 4.5 Å, and white cells mark their absence. Rows and columns identify receptor residues and peptide positions, respectively.

Interface-contact F1 changes from 0.306 to 0.452 for 2CNZ, from 0.327 to 0.211 for 4W50, and from 0.286 to 0.455 for 6MLC. In 2CNZ, contact recovery and DockQ improve together, while their directions differ in the other two examples. Binary contact recovery can change as distances cross the threshold, whereas DockQ also reflects overall geometric agreement. Together, the maps and pose-error profiles locate the receptor interactions accompanying each structural outcome.

A.5 TARGET-LEVEL HETEROGENEITY

Figure 6 places the selected examples within the full distribution of paired target outcomes. PepRIG improves DockQ on 64 of 93 targets, ties on one and worsens on 28. Local RMSD improves on 52 targets and worsens on 41, showing that the bound-pose advantage extends to a larger fraction of targets than the local-geometry gain. Each target score independently selects the highest Sliding-AAR, GDT-TS, Contact-F1 or DockQ, or the lowest Local RMSD, from 40 candidates. Diversity uses the complete pool, and ties use a numerical tolerance of 10^{-9}. The scatter plots also expose differences in effect magnitude across targets. Diversity, for example, improves on more targets than it worsens, yet its overall mean is slightly lower because the balance of magnitudes differs from the balance of counts.

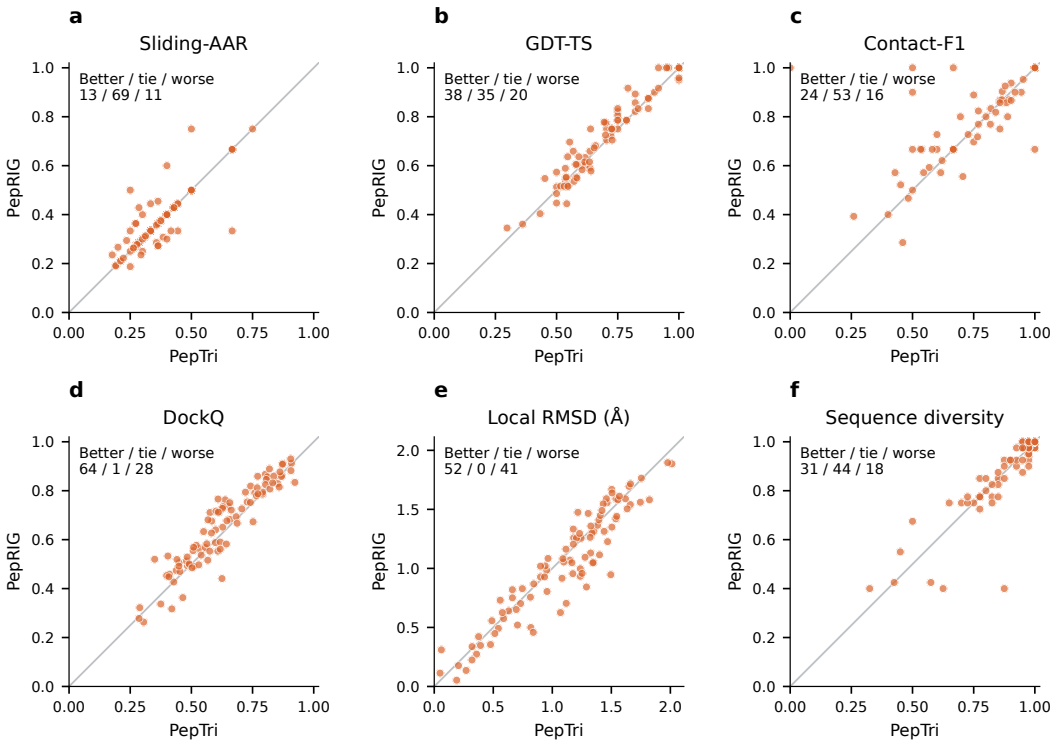


Figure 6: **Paired comparisons across 93 LNR targets.** Points plot PepTri horizontally and PepRIG vertically. Above-diagonal points favor PepRIG except for Local RMSD, where below-diagonal points favor PepRIG. Counts indicate better, tied and worse outcomes.

A.6 JOINT QUALITY WITHIN INDIVIDUAL CANDIDATES

The joint-quality analysis in Figure 7 tests whether sequence recovery, local geometry and bound-pose quality improve together in individual peptides. A successful candidate must satisfy all three thresholds simultaneously, linking the measurements to one generated sequence and structure. We retain the complete candidate pools and fix the Sliding-AAR lower bound at 0.30 throughout the grid. Varying the geometric and pose thresholds then reveals how the joint advantage changes with the required level of reference agreement.

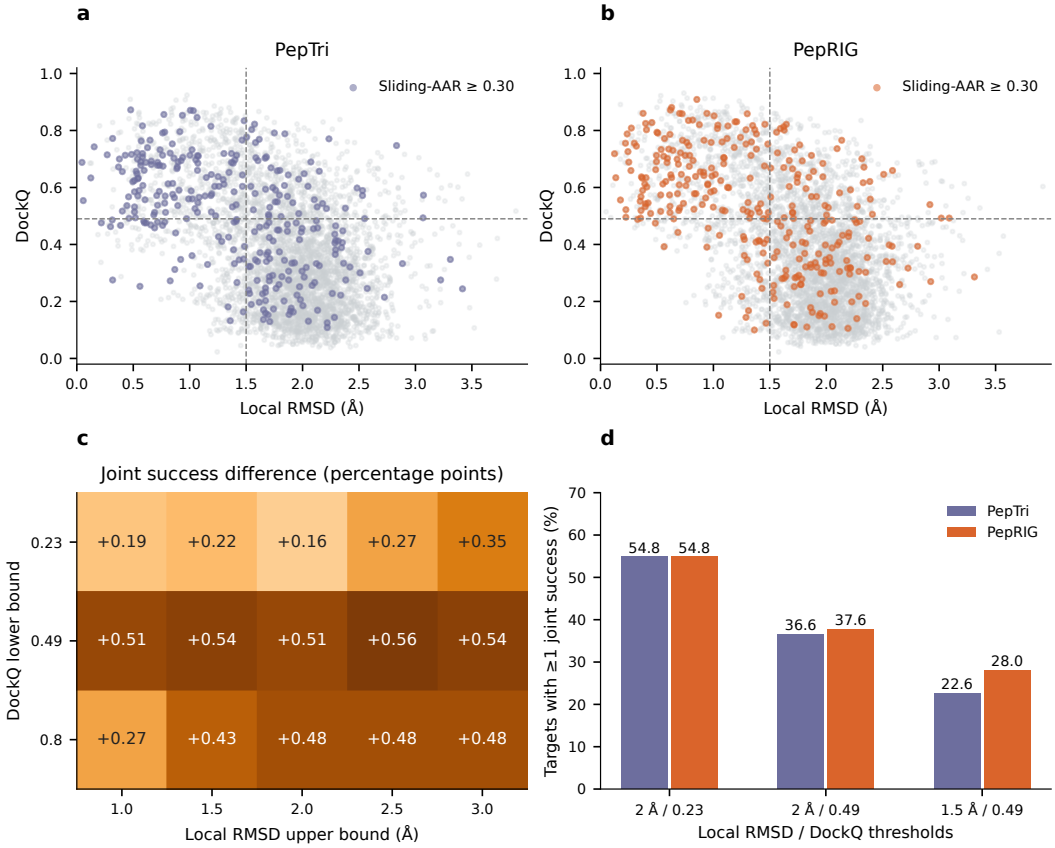


Figure 7: **Joint quality and threshold sensitivity. a,b**, All 3,720 candidates per method, highlighting Sliding-AAR ≥ 0.30. Dashed lines mark Local RMSD $= 1.5$ Å and DockQ $= 0.49$. **c**, PepRIG minus PepTri joint-success rates in percentage points. **d**, Fractions of targets with at least one jointly successful candidate. All joint thresholds require Sliding-AAR ≥ 0.30.

PepRIG has a higher joint-success rate at all 15 displayed settings. At Local RMSD ≤ 1.5 Å and DockQ ≥ 0.49, successful candidates cover 26 targets, compared with 21 for PepTri. The candidate success-rate difference is $+0.54$ percentage points, with a 95% interval from $+0.11$ to $+1.05$. The gain therefore extends to additional targets, while the broadest setting yields a smaller difference whose interval includes zero. Intervals use 20,000 paired resamples of the 93 targets with seed 20260924, retaining all candidates within each target. These pointwise intervals quantify target-sampling uncertainty for the evaluated runs, without multiplicity adjustment or independent-training variation.

A.7 Relationships between candidate-quality metrics

Joint success identifies peptides meeting simultaneous quality requirements, whereas correlation characterizes agreement in candidate rankings. To examine this complementary property, we compute Spearman correlations across the 40 candidates within each target and summarize each metric pair by its median across targets. Pairs with a constant metric within a target are undefined and omitted. Contact-F1 correlations therefore use 86 targets for PepTri and 87 for PepRIG, while other pairs use all 93. This target-wise calculation separates agreement among metrics from variation in target difficulty. We negate Local RMSD so that larger values consistently indicate better quality, allowing the signs of the correlations in Figure 8 to be read uniformly.

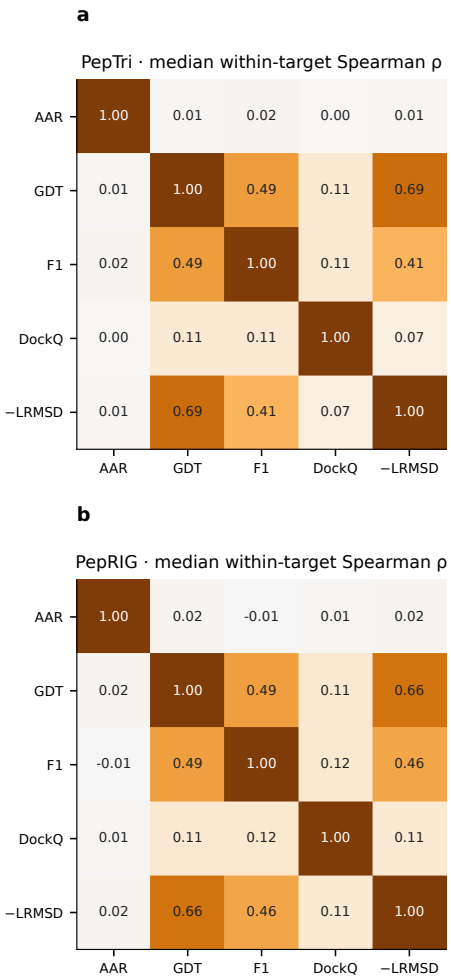


Figure 8: **Within-target metric relationships.** Cells show median target-wise Spearman correlations on a shared scale from −1 to 1. Local RMSD is negated so that higher values indicate better quality.

GDT-TS is associated with lower Local RMSD within targets, with median correlations of 0.69 for PepTri and 0.66 for PepRIG. By contrast, Sliding-AAR has weak median associations with geometric and interface scores in both methods. Sequence recovery therefore provides information distinct from direct measurements of geometry and receptor compatibility. Joint success measures threshold attainment rather than ranking agreement, so a larger successful subset need not produce stronger correlations across the full distribution. Together, the two analyses describe both the frequency of reference-like candidates and the agreement among their quality rankings.

A.8 SEQUENCE COMPOSITION AND LENGTH-STRATIFIED BEHAVIOR

Figure 9 examines whether the structural gains accompany changes in sequence composition and whether their magnitude varies with peptide length. Reference composition uses one native sequence per target, while generated composition includes all 40 sequences per target. Frequencies are normalized by the corresponding total residue count, giving longer peptides proportionally greater weight in the composition plot. The length-stratified analysis compares paired target-level improvements for peptides with at most 10 residues ($n = 50$), 11–15 residues ($n = 25$), and at least 16 residues ($n = 18$). Intervals use 10,000 within-group bootstrap resamples with seed 20260924.

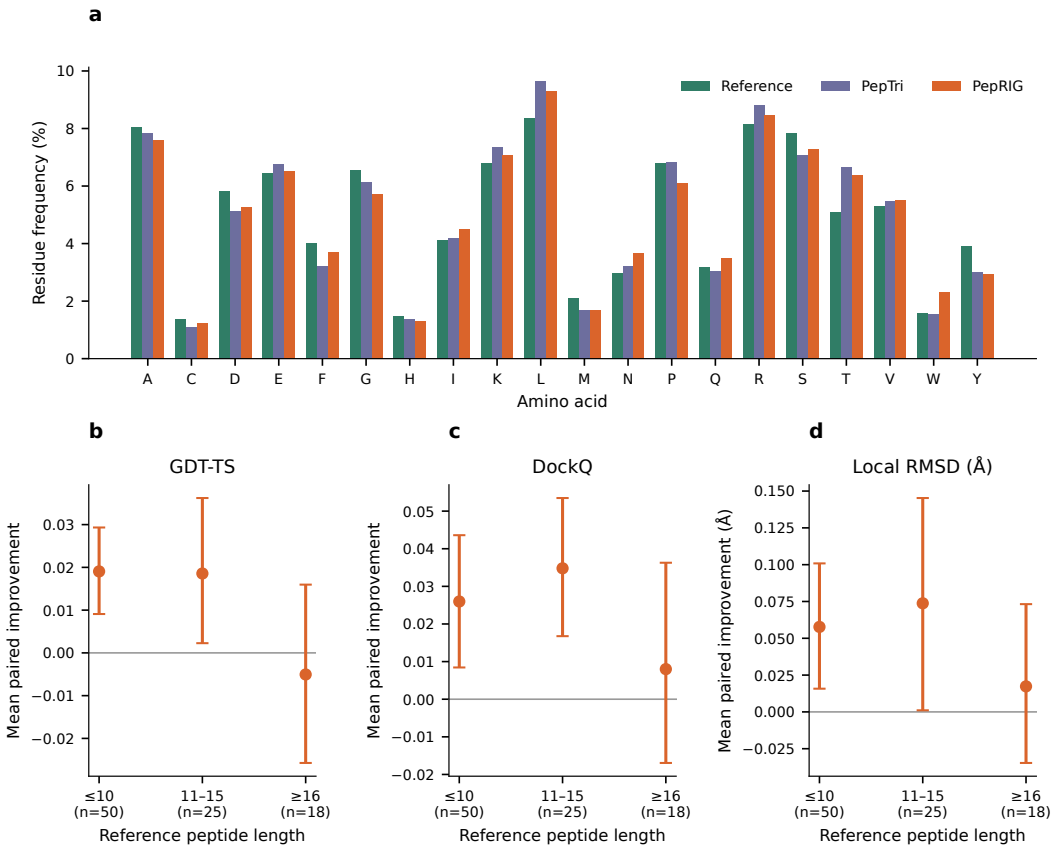


Figure 9: **Sequence composition and length-dependent gains. a**, Amino-acid frequencies. **b–d**, Mean paired gains in GDT-TS, DockQ and Local RMSD. Positive values favor PepRIG, with Local RMSD computed as PepTri minus PepRIG. Error bars show pointwise 95% percentile bootstrap intervals over targets.

Both models generate more leucine and threonine, and less tyrosine, than occur in the reference sequences. Relative to PepTri, PepRIG reduces the leucine and threonine excess but increases tryptophan frequency, producing residue-specific shifts in composition. Across the three length groups, mean DockQ gains are 0.0260, 0.0348 and 0.0080, respectively, with mean Local RMSD reductions of 0.0577, 0.0738 and 0.0173 Å. GDT-TS improves in the first two groups but decreases slightly in the longest group. All three intervals include zero for the longest group, which contains 18 targets. The structural gains are therefore clearest for short and medium-length peptides, while longer peptides show smaller average improvements and greater uncertainty.

A.9 BACKBONE GEOMETRY AND INTERFACE CROWDING

To complement reference-recovery metrics, we evaluate backbone crowding, bond-length errors and interface overlap across all 3,720 unrelaxed candidates per method. All candidates have finite coordinates and complete peptide backbone atoms, and the 93 experimental reference structures provide a geometric calibration. Following Nguyen et al. (2026), an internal Cα clash occurs below 3.6574 Å, excluding the same residue and immediate sequence neighbors. An external clash occurs when a peptide Cα lies within this distance of any receptor Cα. Each clash rate measures the fraction of peptide residues affected, rather than the number of atom pairs.

Bond-length diagnostics assess local covalent geometry independently of backbone crowding. We measure N–Cα, Cα–C, C–O and consecutive-residue C–N bonds in sequence order against reference lengths of 1.460, 1.525, 1.231 and 1.329 Å, respectively. For each candidate, the mean absolute error averages deviations across these bonds, and outlier fractions record deviations exceeding 0.10, 0.20 or 0.50 Å. The combination measures both average distortion and the prevalence of larger local errors.

The heavy-atom screen extends the interface assessment beyond Cα proximity by counting peptide atoms whose overlap with a receptor atom exceeds 0.4 Å. Overlap is the summed van der Waals radii minus the interatomic distance. Fixed radii for C, N, O, S and Se are 1.70, 1.55, 1.52, 1.80 and 1.90 Å, respectively. The rate uses all peptide heavy atoms as its denominator and excludes hydrogen atoms. This distance-based diagnostic applies no bond-specific exemptions and measures heavy-atom crowding rather than a hydrogen-aware, chemistry-validated clashscore. Candidate measurements are averaged within each target, followed by equal-weight averaging over 93 targets. Reference values use one experimental structure per target. Pointwise 95% percentile intervals use 20,000 paired target-bootstrap resamples.

Table 7: **Geometry of the complete candidate pools.** Lower is better. Differences are PepRIG minus PepTri, in percentage points for rates and Å for bond error, with pointwise 95% confidence intervals.

Diagnostic	Reference	PepTri	PepRIG	Difference [95% CI]
Internal peptide geometry				
Cα internal clash (%)	0.31	6.65	4.29	-2.36 [-3.01, -1.73]
Bond error (Å)	0.0087	0.2062	0.1926	-0.0136 [-0.0175, -0.0098]
Bond outliers > 0.10 Å (%)	0.09	66.43	63.89	-2.53 [-3.31, -1.78]
Bond outliers > 0.20 Å (%)	0.09	38.90	35.54	-3.37 [-4.35, -2.44]
Bond outliers > 0.50 Å (%)	0.04	7.09	5.81	-1.28 [-1.77, -0.86]
Receptor-interface crowding				
Cα external clash (%)	0.87	6.30	7.68	$+1.38$ [$+0.94$, $+1.85$]
Interface overlap atoms (%)	2.21	24.62	28.33	$+3.70$ [$+3.16$, $+4.26$]

Table 7 shows that PepRIG improves internal peptide geometry, with fewer internal Cα clashes, lower bond-length error and fewer outliers at every reported tolerance. Bond geometry remains substantially less regular than in the experimental references. Both interface-crowding measures increase relative to PepTri in the unrelaxed pools, so the internal gains do not extend to reduced receptor-interface conflicts.

A.10 CANDIDATE-BUDGET ANALYSIS

Using the first K samples from each frozen pool, Table 8 compares per-metric best-of-K scores. PepRIG achieved higher Contact-F1 and DockQ at every reported budget from $K = 1$ to $K = 40$. Across all candidates, mean Contact-F1 increased from 0.3902 to 0.4183, DockQ from 0.3750 to 0.3947, and GDT-TS from 0.5496 to 0.5603. Sliding-AAR changed little, indicating that the main gains concern structural organization and receptor-relative placement. These mean-candidate gains show that the improvement extends beyond the highest-scoring samples. Local RMSD favored PepRIG at $K = 1, 20, 40$ but PepTri at $K = 5, 10$. Sections A.5 and A.8 provide target-level and length-stratified analyses of these gains.

Table 8: Effect of candidate budget K on LNR. Local RMSD is reported in Å.

	Contact-F1↑		DockQ↑		Local RMSD↓	
K	PepTri	PepRIG	PepTri	PepRIG	PepTri	PepRIG
1	0.3583	0.4098	0.3863	0.3894	1.8085	1.7718
5	0.6290	0.6693	0.5177	0.5271	1.5103	1.5572
10	0.7266	0.7410	0.5562	0.5700	1.3652	1.3815
20	0.7769	0.8156	0.5954	0.6224	1.2714	1.1799
40	0.8283	0.8557	0.6282	0.6531	1.1811	1.0540