

1	Appendix Contents	
2	A Backbone Representation	1
3	B Diffusion on the Toric Manifold	3
4	C Diffusion on SE(3)	3
5	D Architecture	4
6	E Experimental Details	4
7	E.1 Dataset	4
8	E.2 Baselines	4
9	E.3 Training Metrics Details	5
10	E.4 Hyperparameters	6
11	E.5 Computational Complexity	6
12	F Additional Experiments	6
13	F.1 Visualization	6
14	F.2 Training and Sampling Time Steps	6
15	G Limitations and Future Work	7

16 This appendix provides additional technical details and experimental insights. We first describe
17 the backbone representation and our diffusion processes on relevant geometric manifolds. We then
18 outline the model architecture and training setup, including datasets, baselines, and hyperparameters.
19 Further sections present additional experiments, followed by a discussion of limitations and directions
20 for future work.

21 **A Backbone Representation**

22 As introduced in 3.3 section, every frame is composed by four atomic group N^* , C_α^* , C^* , O^* , which
23 is idealized atom coordinates that assumes chemically idealized bond angles and lengths.

24 We use the tuple $T = (r, m)$ to denote the Euclidean transformations corresponding to frames, where
25 $r \in SO(3)$ for the rotation and $m \in \mathbb{R}^3$ for the translation components. We use the dot product
26 operator (\cdot) to denote application of a transformation to the position of frame $b \in \mathbb{R}^3$:

$$\begin{aligned}
\hat{b} &= T \cdot b \\
&= (r, m) \cdot b \\
&= rb + m.
\end{aligned}$$

27 The composition of Euclidean transformations denoted as:

$$\begin{aligned}
T &= T_1 \cdot T_2 \\
(r, m) &= (r_1, m_1) \cdot (r_2, m_2) \\
&= (r_1 r_2, r_1 m_2 + m_1).
\end{aligned}$$

28 The group inverse of the transform T is denoted as:

$$\begin{aligned} T^{-1} &= (\mathbf{r}, \mathbf{m})^{-1} \\ &= (\mathbf{r}^{-1}, -\mathbf{r}^{-1}\mathbf{m}) \end{aligned}$$

29 The tuple transforms a position in local coordinates $\mathbf{b}_{\text{local}} \in \mathbb{R}^3$ to a position in global coordinates
30 $\mathbf{b}_{\text{global}} \in \mathbb{R}^3$ as

$$\begin{aligned} \mathbf{b}_{\text{global}} &= T \cdot \mathbf{b}_{\text{local}} \\ &= \mathbf{r}\mathbf{b}_{\text{local}} + \mathbf{m}. \end{aligned}$$

31 In local position of frame, the bond angles and lengths values differ slightly per amino acid type.
32 Follow (Yim et al., 2023) and (Yang et al., 2023), we set the local coordinates as:

$$\begin{aligned} \mathbf{N}^* &= (-0.525, 1.363, 0.0) \\ \mathbf{C}_\alpha^* &= (0.0, 0.0, 0.0) \\ \mathbf{C}^* &= (1.526, 0.0, 0.0) \\ \mathbf{O}^* &= (0.627, 1.062, 0.0) \end{aligned} \tag{1}$$

33 where \mathbf{C}_α^* is central in protein backbones, connecting \mathbf{N}^* and \mathbf{C}^* groups. Using the transformation
34 T_n , we manipulate idealized coordinates to construct global coordinates of backbone atoms for
35 residue n via:

$$\begin{aligned} [\mathbf{N}_n, \mathbf{C}_n, (\mathbf{C}_\alpha)_n, \mathbf{O}_n] &= [T_n \cdot \mathbf{N}^*, T_n \cdot \mathbf{C}^*, \\ &\quad T_n \cdot \mathbf{C}_\alpha^*, T_n \cdot T_{\text{psi}}^*(\psi_n) \cdot \mathbf{O}^*]. \end{aligned} \tag{2}$$

36 Given the coordinates of three atoms $[\mathbf{N}_n, \mathbf{C}_n, (\mathbf{C}_\alpha)_n]$, we construct a local rigid frame using a
37 Gram-Schmidt process:

$$\begin{aligned} \boldsymbol{\omega}_1 &= \mathbf{C}_n - (\mathbf{C}_\alpha)_n, & \boldsymbol{\omega}_2 &= \mathbf{N}_n - (\mathbf{C}_\alpha)_n \\ \mathbf{e}_1 &= \boldsymbol{\omega}_1 / \|\boldsymbol{\omega}_1\|, & \mathbf{u}_2 &= \boldsymbol{\omega}_2 - \mathbf{e}_1(\mathbf{e}_1^T \boldsymbol{\omega}_2), \\ \mathbf{e}_2 &= \mathbf{u}_2 / \|\mathbf{u}_2\|, \\ \mathbf{e}_3 &= \mathbf{e}_1 \times \mathbf{e}_2, \\ \mathbf{r}_n &= (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3), \\ \mathbf{m}_n &= (\mathbf{C}_\alpha)_n, \\ T_n &= (\mathbf{r}_n, \mathbf{m}_n). \end{aligned} \tag{3}$$

38 In this construction, two directional vectors are first defined: from \mathbf{C}_α to \mathbf{C} and \mathbf{C}_α to \mathbf{N} . We
39 normalize the first direction $\boldsymbol{\omega}_1$ to define the local x-axis \mathbf{e}_1 , and orthogonalize and normalize $\boldsymbol{\omega}_2$
40 to define the y-axis \mathbf{e}_2 . The z-axis \mathbf{e}_3 is computed as the cross product of \mathbf{e}_1 and \mathbf{e}_2 , forming a
41 right-handed orthonormal basis. The resulting frame T_n placing the local frame at the \mathbf{C}_α of residue n .
42 To define the local frame for placing the oxygen atom, we begin with the residue's central frame T_n ,
43 then apply a secondary transformation: $T_{\text{psi}}^*(\psi_n) = (\mathbf{r}_x(\psi_n), \mathbf{m}_{\text{psi}})$, where ψ_n represents a rotation
44 angle along x-axis. The transformation funtions are defined as:

$$\mathbf{r}_x(\psi) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \psi & -\sin \psi \\ 0 & \sin \psi & \cos \psi \end{pmatrix}, \quad \mathbf{m}_{\text{psi}} = (1.526, 0.0, 0.0). \tag{4}$$

45 This transformation represents a rotation around the x-axis (aligned with the bond from \mathbf{C}_α to \mathbf{C}) by
46 an angle ψ_n , followed by a translation to the position of the carbon atom \mathbf{C}^* in the idealized frame
47 centered at \mathbf{C}_α^* . The combined transformation $T_n \cdot T_{\text{psi}}^*(\psi_n)$ thus defines the final frame in which the
48 idealized oxygen \mathbf{O}^* is placed to obtain its global coordinate.

49 B Diffusion on the Toric Manifold

50 A torsion vector $\chi \in [0, 2\pi)^d$ naturally resides on a flat d-dimensional torus, which can be represented
 51 as the quotient space \mathbb{R}^d/L , where $L = (2\pi\mathbb{Z}^d)$ denotes a discrete lattice subgroup of \mathbb{R}^d isomorphic
 52 to \mathbb{Z}^d . This space models periodic angular data, and inherits a flat metric from its covering Euclidean
 53 space. The tangent space of the torus at any point is identified with \mathbb{R}^d , and all operations are
 54 performed modulo 2π .

55 C Diffusion on SE(3)

56 Following previous work (Yim et al., 2023), we treat the group $\text{SE}(3)$ as the product space $\text{SO}(3) \times \mathbb{R}^3$,
 57 and endow it with a product Riemannian metric. Specifically, for tangent vectors $(a, b), (a', b') \in$
 58 $T_r \text{SO}(3) \times \mathbb{R}^3$, the metric is defined as:

$$\langle (a, b), (a', b') \rangle_{\text{SE}(3)} = \langle a, a' \rangle_{\text{SO}(3)} + \langle b, b' \rangle_{\mathbb{R}^3}.$$

59 This structure allows for a natural decomposition of differential geometric objects on $\text{SE}(3)$ into
 60 rotational and translational components. In particular, the gradient of a function $f : \text{SE}(3) \rightarrow \mathbb{R}$ at
 61 $\mathcal{T} = (\mathbf{r}, \mathbf{x})$ is given by:

$$\nabla_{\mathcal{T}} f(\mathcal{T}) = [\nabla_r f(\mathbf{r}, \mathbf{m}), \nabla_m f(\mathbf{r}, \mathbf{m})],$$

62 and the Laplace–Beltrami operator decomposes as:

$$\Delta_{\text{SE}(3)} f(\mathcal{T}) = \Delta_{\text{SO}(3)} f(\mathbf{r}, \mathbf{m}) + \Delta_{\mathbb{R}^3} f(\mathbf{r}, \mathbf{m}).$$

63 We define Brownian motion on $\text{SE}(3)$ as the product of independent Brownian motions on $\text{SO}(3)$
 64 and \mathbb{R}^3 :

$$\mathbf{B}_t^{\text{SE}(3)} = [\mathbf{B}_t^{\text{SO}(3)}, \mathbf{B}_t^{\mathbb{R}^3}]$$

65 where the rotational and translational components evolve independently. This product metric allows
 66 us to treat the rotational and translational components of the forward diffusion process independently,
 67 leading to the following decomposition of the conditional score:

$$\nabla_{\mathcal{T}_t} \log p_{t|0}(\mathcal{T}_t | \mathcal{T}_0) = [\nabla_{\mathbf{r}_t} \log p_{t|0}(\mathbf{r}_t | \mathbf{r}_0), \nabla_{\mathbf{m}_t} \log p_{t|0}(\mathbf{m}_t | \mathbf{m}_0)]$$

68 The forward process on $\text{SE}(3)$ is thus described by two independent SDEs. Let \mathcal{M} be a compact Lie
 69 group (e.g., $\text{SO}(3)$), and let χ_ℓ denote the character of the ℓ -th irreducible unitary representation of
 70 dimension d_ℓ . Then, the heat kernel (transition density of Brownian motion) on \mathcal{M} is given by:

$$p_{t|0}(x_t | x_0) = \sum_{\ell \in \mathbb{N}} d_\ell e^{-\lambda_\ell t/2} \chi_\ell((x_0)^{-1} x_t).$$

71 where λ_ℓ is the eigenvalue of the Laplace–Beltrami operator associated with χ_ℓ . Specializing to
 72 $\text{SO}(3)$, the heat kernel becomes the isotropic Gaussian on $\text{SO}(3)$:

$$f(\omega, t) = \sum_{\ell \in \mathbb{N}} (2\ell + 1) e^{-\ell(\ell+1)t/2} \frac{\sin((\ell+1/2)\omega)}{\sin(\omega/2)}. \quad (5)$$

73 where ω is the angle of the relative rotation. The corresponding score function is:

$$\nabla \log p_{t|0}(\mathbf{r}_t | \mathbf{r}_0) = \frac{\mathbf{r}_t}{\omega_t} \log(\mathbf{r}_0^\top \mathbf{r}_t) \frac{\partial_\omega f(\omega^{(t)}, t)}{f(\omega^{(t)}, t)}, \quad (6)$$

74 where ω_t is the angle of the relative rotation $\mathbf{r}_0^\top \mathbf{r}_t$, and the matrix logarithm term maps to the tangent
 75 space at \mathbf{r}_t .

76 For the translational component, we model the forward process using a Variance Preserving SDE
 77 (VP-SDE). The transition density is given by:

$$p_{t|0}(\mathbf{m}_t | \mathbf{m}_0) = \mathcal{N}(x_t; e^{-t/2} \mathbf{m}_0, (1 - e^{-t}) \mathbf{I}_3). \quad (7)$$

78 Then we can get the score as:

$$\nabla \log p_{t|0}(\mathbf{m}_t | \mathbf{m}_0) = \frac{1}{1 - e^{-t}} (e^{-t/2} \mathbf{m}_0 - \mathbf{m}_t). \quad (8)$$

79 We use a learned denoising network to approximate the conditional score of the full $\text{SE}(3)$ transfor-
 80 mation. The score is decomposed into rotational and translational components as follows:

$$\begin{aligned} \nabla_{\mathcal{T}_t} \log p_{t|0}(\mathcal{T}_t | \hat{\mathcal{T}}_0) &= (s_\theta^r(t, \mathcal{T}_t), s_\theta^m(t, \mathcal{T}_t)), \\ s_\theta^r(t, \mathcal{T}_t) &= \nabla_{\mathbf{r}_t} \log p_{t|0}(\mathbf{r}_t | \hat{\mathbf{r}}_0), \\ s_\theta^m(t, \mathcal{T}_t) &= \nabla_{\mathbf{m}_t} \log p_{t|0}(\mathbf{m}_t | \hat{\mathbf{m}}_0). \end{aligned} \quad (9)$$

81 D Architecture

82 Here we provide mathematical detail of PepBridge presented in method section. Let $\mathbf{h}^\ell =$
 83 $[h_1^\ell, \dots, h_N^\ell] \in \mathbb{R}^{N \times D_h}$ denote the node embeddings at layer ℓ , where h_n^ℓ represents the embedding
 84 for residue n . Similarly, let $\mathbf{z}^\ell \in \mathbb{R}^{N \times N \times D_z}$ represent the edge embeddings, where z_{ij}^ℓ encodes the
 85 interaction between residues i and j .

86 Node embeddings are initialized using residue indices, atom coordinates, backbone dihedral angles,
 87 side-chain angles h_B , and the diffusion timestep t . For edge (residue-pair) embeddings, we incorporate
 88 a combination of residue-type pairs, relative sequence positions, pairwise distances, and relative
 89 orientations. Each of these features is individually encoded using a dedicated multi-layer perceptron
 90 (MLP). The resulting feature vectors are concatenated and passed through another MLP to produce
 91 the final embeddings.

92 The initial layer-0 embeddings for residues i and residue pairs (i, j) are computed using MLPs
 93 applied to sinusoidal encodings $\phi(\cdot)$ of the input features:

$$\begin{aligned} h_i^0 &= \text{MLP}([\phi(h_{B_i}), \phi(t)]) \in \mathbb{R}^{D_h}, \\ z_{ij}^0 &= \text{MLP}([\phi(h_{B_i}), \phi(h_{B_j}), \phi(i-j), \phi(\text{dis}(i, j)), \phi(\text{ori}(i, j)), \phi(t)]) \in \mathbb{R}^{D_z}, \end{aligned} \quad (10)$$

94 where D_h, D_z denote the dimensions of the node and edge embeddings, respectively. The functions
 95 $\phi(\text{dis}(i, j))$ and $\phi(\text{ori}(i, j))$ represent sinusoidal encodings of the distance and relative orientation
 96 between residues i and j .

97 To encode the surface of the receptor protein, we extract node-level features from the surface points
 98 and apply an MLP to obtain embeddings. Each surface node is represented by its 3D position
 99 (surf_t), hydrogen bonding potential ($\text{surf}_{\text{hbond}}$), and hydrophobicity score (surf_{hp}). These features are
 100 concatenated and passed through an MLP to produce the surface node embeddings:

$$h_{\text{surf}} = \text{MLP}([\text{surf}_t, \text{surf}_{\text{hbond}}, \text{surf}_{\text{hp}}]). \quad (11)$$

101 For the peptide surface representation, we encode only the 3D positional coordinates using an MLP,
 102 omitting auxiliary features such as hydrogen bonding and hydrophobicity. At each diffusion timestep
 103 t , the model takes as input the receptor’s node and edge embeddings, the noised peptide descriptors,
 104 and a sinusoidal embedding of the timestep. It predicts a denoising score that guides the reverse
 105 diffusion process toward the clean peptide descriptors at $t = 0$. The model architecture is based
 106 on Invariant Point Attention (IPA), which employs SE(3)-invariant attention to capture interactions
 107 between the receptor and the peptide. The output of the IPA module is passed through separate MLP
 108 decoders to reconstruct various ground-truth peptide descriptors, such as atom coordinates, dihedral
 109 angles, and residue types. Notably, certain residue types may be partially inferred from the number
 110 of side-chain dihedral angles, due to structural constraints.

111 E Experimental Details

112 The experiments were conducted on a computing cluster with 2 NVIDIA RTX A6000, each with 48
 113 GB of memory. The total computation time for training was approximately 21 hours. We trained for
 114 900000 steps with batch size 8. We used the Adam optimizer with a start learning of $5e-4$. We also
 115 schedule to decay the learning rate exponentially with a factor of 0.6 and a minimum learning rate of
 116 $1e-6$. The learning rate is decayed if there is no improvement for the validation loss in 10 consecutive
 117 evaluations.

118 E.1 Dataset

119 The filtered dataset underwent sequence-based clustering using MMseqs2 Steinegger & Söding (2017),
 120 resulting in 9,816 protein-peptide complexes organized into 292 distinct clusters. For systematic
 121 evaluation, we designated 10 clusters encompassing 158 complexes as the test set, with the remaining
 122 complexes allocated to training and validation cohorts.

123 E.2 Baselines

124 We briefly summarize the baselines and tools used in our study, including generative approaches for
 125 protein and peptide design, as well as side-chain packing methods.

- **ProteinGenerator** (Lisanza et al., 2023) is a RoseTTAFold-based diffusion model that jointly generates protein sequences and structures, with flexible conditioning on target sequence and structural attributes.
- **RFDiffusion** (Watson et al., 2023) is a generative model fine-tuned on structure denoising tasks, enabling high-accuracy design of monomers, binders, and symmetric protein assemblies.
- **PPFLOW** (Lin et al., 2024) is a target-aware peptide designer that performs conditional flow matching on torus manifolds to model peptide torsion-angle geometry.
- **PepFlow** (Li et al., 2024) is a multimodal flow-matching model for full-atom peptide design targeting specific protein receptors. It jointly models backbone geometry, side-chain conformations, and residue identities over appropriate geometric manifolds.
- **PepGLAD** (Kong et al., 2024) combines geometric latent diffusion with receptor-specific transformations to generate full-atom peptides. The model operates in a learned latent space and adapts to diverse binding geometries for improved generalization.
- **Chroma** (Ingraham et al., 2023) is a unified generative framework for proteins and protein complexes that integrates a polymer-aware diffusion process with a scalable architecture, supporting constraint-driven design across sequence, structure, and function.
- **SCWRL4** (Krivov et al., 2009) is a widely used side-chain packing tool employing a backbone-dependent rotamer library and a statistical energy function.
- **DLPacker** (Misiura et al., 2022) is a 3D CNN-based model for residue side-chain conformation prediction. We utilize the official implementation along with the model weights.
- **AttnPacker** (McPartlon & Xu, 2022) utilizes equivariant attention mechanisms on backbone 3D geometry to predict all side-chain coordinates simultaneously.
- **DiffPack** (Zhang et al., 2023) is a diffusion-based generative model that autoregressively samples side-chain angles on a toric manifold.

E.3 Training Metrics Details

RMSD. Root-Mean-Square Deviation is a widely used metric for assessing structural similarity between proteins. In our evaluation, we align the generated peptide to the native peptide within the complex using the Kabsch algorithm, considering only the peptide portion for superposition. We then compute the RMSD based on normalized C_α atom distances between the generated and native peptides. Lower RMSD values indicate greater structural similarity.

BSR. Binding Site Recovery measures the similarity of interaction patterns between the generated and native peptide-protein complexes. Specifically, it evaluates whether the generated peptide engages target protein residues in a manner similar to the native peptide, potentially reflecting similar biological functions. A residue in the protein is considered part of the binding site if its C_β atom lies within 6 Å of any residue in the peptide. BSR is defined as the ratio of overlapping binding site residues between the generated and native complexes. Higher BSR values indicate greater similarity in binding interactions.

Consistency represents the statistical association between the clustering results of surface and structures. This metric quantifies how well a model captures the fundamental consistency between surfaces and their corresponding structures. A model that accurately represents the joint distribution should achieve a high score, while a low score indicates the model generates inconsistent surface-structure pairs. The evaluation process involves clustering both surfaces and structures, assigning discrete labels to each. These clustering labels can be interpreted as nominal variables. Given that similar surfaces should correspond to similar structures. We employ Cramér’s V association Cramér (1999) to measure this correlation, where a value of 1.0 indicates perfect association and 0.0 indicates no association. For surface representation, we first obtain molecular fingerprints using the methodology described in Song et al. (2024). These fingerprints serve as input for the clustering algorithm, which assigns labels to the generated peptide surfaces.

Diversity. To assess diversity, we compute all pairwise TM-scores among the generated peptides for a given target using the original TM-align program. TM-scores quantify structural similarity between peptide pairs. We define diversity as 1 minus the average TM-score, where higher values

178 indicate greater structural variability among the generated peptides. This metric reflects the breadth
 179 of structural exploration achieved during the design process.

180 E.4 Hyperparameters

181 Our proposed method incorporates several hyperparameters, including sample step, learning rate,
 182 batch size and feature dimensions. To validate these hyperparameters, we conducted a random search.
 183 The search space are presented in Table 1.

Table 1: Search space for all PepBridge. The parameters used in validation are marked in **bold**

Parameter	Search Space
Learning rate	0.0009, 0.0007, 0.0005 , 0.00001
Hidden dimension of residue feature	64, 128 , 256
Hidden dimension of edge feature	64 , 128, 256
Hidden dimension of surface feature	16 , 24, 32
Number of attention heads	8 , 16, 24
Loss weight of surface	0.1, 0.5 , 1
Loss weight of backbone position	0.1, 0.5, 1
Loss weight of backbone rotation	0.1, 0.5, 1
Sampling steps	500, 1000 , 1500
Training steps	500, 1000 , 1500
Batch size	4, 8 , 16

184 E.5 Computational Complexity

185 We compared PepBridge with several baseline methods in terms of training time, inference time per
 186 sample, GPU usage, and model size. The time cost is reported as the total time spent divided by the
 187 number of designed candidates. As summarized in Table 2, multi-modal processing does introduce
 188 additional complexity relative to uni-modal approaches, our analysis shows that PepBridge achieves
 189 a good balance between computational efficiency and performance.

Table 2: Computational cost and resource footprint across methods. Training and inference times are normalized per designed sample.

Method	Training time	Inference time (s/sample)	GPU(s) used	Params (M)
RFdiffusion	3 days	80–180	8×A100	~120
ProteinGenerator	4 weeks	152	64×V100	~120
PepFlow	20 hours	14–24	2×A6000	~7
Chroma	10 weeks	185–226	8×V100	~20
PepGLAD	29 hours	3	1×24 GB GPU	~2.5
PepBridge	21 hours	16–37	2×A6000	~10

190 F Additional Experiments

191 F.1 Visualization

192 Figure 1 provides additional examples of peptides generated by PepBridge. These visualizations
 193 include both surface and backbone structures of the generated peptides in a top-down view, further
 194 illustrating the model’s ability to produce geometrically coherent and interface-aware designs.

195 F.2 Training and Sampling Time Steps

196 We ablated the number of diffusion steps used during training while fixing the sampling procedure
 197 at 1000 steps. The results are summarized in Table 3. Increasing the number of steps consistently
 198 improved all quality metrics. The largest gains occurred up to 1000 steps, with smaller but still




Figure 1: Visualization of generated peptides by PepBridge. **Top:** Generated peptides (in orange) for receptors (in purple). The PDB ID of the receptors are 4CC2, 3AV9, 1B07, and 2XS1. **Bottom:** The generated backbone structure and surface. The ground-truth surface structure (in black) and generated surface (in orange) are shown to compare the ability of interface caption.

Table 3: Effect of training steps on performance (sampling fixed at 1000). Higher is better (\uparrow) except RMSD (\downarrow).

	Div _{stru} (\uparrow)	Aff. % (\uparrow)	Stab. % (\uparrow)	RMSD Å (\downarrow)	BSR (\uparrow)
PepBridge (time step =500)	0.57	18.96	24.79	2.96	82.84
PepBridge (time step =800)	0.61	19.07	26.31	2.36	85.57
PepBridge (time step =1000)	0.59	19.16	25.02	2.19	83.90
PepBridge (time step =1500)	0.62	23.28	26.68	2.11	86.92

measurable improvements between 1000 and 1500 steps. Given these trends, we adopt 1000 training steps as a favorable trade-off between overall accuracy and computational cost.

We further conducted experiments to assess how different inference time steps affect the final performance metrics using the model trained with 1000 steps. As shown in Table 4, the results indicate that longer sampling chains (e.g., 1000 steps) improve the quality of the generated structures. However, performance gains begin to plateau beyond 800 steps, and we observe a slight decrease in structural diversity, suggesting a trade-off between generation determinism and diversity. Based on this analysis, we selected 1000 steps as the default setting to achieve the best overall balance.

G Limitations and Future Work

While PepBridge presents a structured approach to joint protein surface and backbone design, several limitations remain that can be addressed in future work. One key limitation lies in the

Table 4: Effect of sampling steps on performance (model trained with 1000 steps). Higher is better (\uparrow) except RMSD (\downarrow).

	Div _{stru} (\uparrow)	Aff. % (\uparrow)	Stab. % (\uparrow)	RMSD Å (\downarrow)	BSR (\uparrow)
PepBridge (time step =500)	0.61	17.42	23.97	2.85	80.05
PepBridge (time step =800)	0.62	18.77	24.58	2.69	82.33
PepBridge (time step =1000)	0.59	19.16	25.02	2.19	83.90
PepBridge (time step =1500)	0.56	18.46	24.71	2.74	83.78

simplification of surface representations. The current model relies on solvent-accessible point clouds with biochemical annotations, which, while effective, may not fully capture finer molecular interactions such as electrostatic potential fields and solvent dynamics. These factors play crucial roles in protein-protein interactions and could enhance the accuracy of designed peptides if incorporated. Another limitation is the model’s reliance on receptor geometry. PepBridge assumes that receptor surface features sufficiently dictate the constraints on peptide binding. However, this does not account for receptor flexibility or conformational changes upon ligand binding, which are common in many biological systems. Addressing this aspect could make the model more applicable to highly dynamic binding sites. Computational efficiency also poses a challenge. The diffusion bridge model and SE(3) diffusion backbone generation require computationally intensive sampling. While the hierarchical generation process improves efficiency, generating high-quality peptides remains expensive, particularly for longer sequences. Further optimization is necessary to reduce the computational cost while maintaining or improving accuracy. Additionally, the current evaluation primarily focuses on geometric complementarity and binding affinity predictions. While these provide useful insights, they do not fully capture the functional stability of designed peptides. Wet-lab experiments and molecular dynamics simulations are necessary to assess real-world applicability, ensuring that the generated structures remain stable under physiological conditions.

Future work can address these limitations in several ways. Enhancing surface representations by incorporating higher-order biochemical features such as electrostatic potential fields, solvent effects, or graph-based molecular embeddings could improve the precision of surface-conditioned peptide generation. Furthermore, integrating receptor flexibility into the model by leveraging conformational ensembles or reinforcement learning-based refinement strategies would allow for more realistic modeling of dynamic binding sites. To improve computational efficiency, future research could explore accelerated sampling techniques, such as adaptive noise schedules, score distillation sampling (SDS), or flow-matching approaches. These methods could significantly reduce inference time while preserving or enhancing model accuracy. Finally, validating PepBridge through real-world applications, particularly in therapeutic peptide design, remains a crucial next step. Incorporating experimental validation through biochemical assays and integrating co-evolutionary signals into the design process could further enhance the biological relevance of the generated structures. By addressing these challenges, PepBridge can be refined to enable more accurate, efficient, and versatile protein design.

References

- Cramér, H. *Mathematical methods of statistics*, volume 26. Princeton university press, 1999.
- Ingraham, J. B., Baranov, M., Costello, Z., Barber, K. W., Wang, W., Ismail, A., Frappier, V., Lord, D. M., Ng-Thow-Hing, C., Van Vlack, E. R., et al. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.
- Kong, X., Jia, Y., Huang, W., and Liu, Y. Full-atom peptide design with geometric latent diffusion. *Advances in Neural Information Processing Systems*, 37:74808–74839, 2024.
- Krivov, G. G., Shapovalov, M. V., and Dunbrack Jr, R. L. Improved prediction of protein side-chain conformations with scwrl4. *Proteins: Structure, Function, and Bioinformatics*, 77(4):778–795, 2009.
- Li, J., Cheng, C., Wu, Z., Guo, R., Luo, S., Ren, Z., Peng, J., and Ma, J. Full-atom peptide design based on multi-modal flow matching. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 27615–27640, 2024.
- Lin, H., Zhang, O., Zhao, H., Jiang, D., Wu, L., Liu, Z., Huang, Y., and Li, S. Z. Ppflow: target-aware peptide design with torsional flow matching. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 30510–30528, 2024.
- Lisanza, S. L., Gershon, J. M., Tipps, S., Arnoldt, L., Hendel, S., Sims, J. N., Li, X., and Baker, D. Joint generation of protein sequence and structure with rosettafold sequence space diffusion. *bioRxiv*, pp. 2023–05, 2023.
- McPartlon, M. and Xu, J. An end-to-end deep learning method for rotamer-free protein side-chain packing. *bioRxiv*, pp. 2022–03, 2022.

- 262 Misiura, M., Shroff, R., Thyer, R., and Kolomeisky, A. B. Dlpacker: Deep learning for prediction of
263 amino acid side chain conformations in proteins. *Proteins: Structure, Function, and Bioinformatics*,
264 90(6):1278–1290, 2022.
- 265 Song, Z., Huang, T., Li, L., and Jin, W. Surfpro: Functional protein design based on continuous
266 surface. *arXiv preprint arXiv:2405.06693*, 2024.
- 267 Steinegger, M. and Söding, J. Mmseqs2 enables sensitive protein sequence searching for the analysis
268 of massive data sets. *Nature biotechnology*, 35(11):1026–1028, 2017.
- 269 Watson, J. L., Juergens, D., Bennett, N. R., Trippe, B. L., Yim, J., Eisenach, H. E., Ahern, W., Borst,
270 A. J., Ragotte, R. J., Milles, L. F., et al. De novo design of protein structure and function with
271 rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- 272 Yang, Z., Zeng, X., Zhao, Y., and Chen, R. Alphafold2 and its applications in the fields of biology
273 and medicine. *Signal Transduction and Targeted Therapy*, 8(1):115, 2023.
- 274 Yim, J., Trippe, B. L., De Bortoli, V., Mathieu, E., Doucet, A., Barzilay, R., and Jaakkola, T. Se (3)
275 diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*,
276 2023.
- 277 Zhang, Y. et al. Diffpack: A torsional diffusion model for autoregressive protein side-chain packing.
278 In *Advances in Neural Information Processing Systems*, volume 36, pp. 48150–48172, 2023.