D-Flow: Multi-modality Flow Matching for D-peptide Design

Fang Wu [1] [*], Shuting Jin[2], Xiangru Tang[3], Junlin Xu[2], Mark Gerstein[3], James Zou[1]
[1] Stanford University [2] Wuhan University of Science and Technology [3] Yale University

Abstract

Proteins play crucial roles in biological processes, with therapeutic peptides emerging as promising pharmaceutical agents. They allow for new possibilities to leverage target binding sites that were previously undruggable. Although deep learning has advanced peptide discovery, generating D-proteins composed of D-amino acids remains challenging because of the scarcity of natural examples. This paper proposes D-Flow, a full-atom flow-based framework for de novo D-peptide design. D-Flow is conditioned on receptor binding and utilizes a comprehensive representation of peptide structure, incorporating backbone frames, side-chain angles, and discrete amino acid types. A mirror-image algorithm is implemented to address the lack of training data for D-proteins, which convert the chirality of L-receptors. Furthermore, we enhance D-Flow's capacity by integrating large protein language models with structural awareness through a lightweight structural adapter. A two-stage training pipeline and a controlling toolkit also enable D-Flow to transition from a general protein design to a targeted binder design while preserving pre-training knowledge. Extensive experimental results on the PepMerge benchmark demonstrate D-Flow's effectiveness, particularly in developing peptides with entire D-residues. This approach represents a significant advancement in computational D-peptide design, offering unique opportunities for bioorthogonal and stable molecular tools and diagnostics.

1 Introduction

Proteins are the building blocks of life and make essential contributions to nearly all fundamental biological processes in the cell. They fold into 3D conformations to perform distinct functionalities. Therapeutic peptides, comprising a limited number of well-ordered residues, are single-chain proteins and an irregular class of pharmaceutical agents [1–3]. Peptide drugs occupy a unique chemical and pharmacological space between small and large molecules [4]. Due to this specialty, increasing efforts have adopted deep learning (DL) algorithms to facilitate their discovery [5, 6]. Notably, unbound peptide chains typically have high free energy and entropy, resulting in unstable conformations. In contrast, they can elicit pharmacological effects upon binding to specific receptors, forming stable complexes where both the receptor and ligand adopt equilibrium structures.

Diffusion or flow models [7, 8], a recent generative family, are revolutionizing several domains. Notable initiatives have explored the potential of diffusion or flow models in the context of drug discovery. These models have been applied to various protein representations, including carbon-alpha only structures [9], torsion angles [10], and the SE(3) backbone frame representation [11], and diverse scenes such as molecular design [12], antibody engineering [13], *de novo* protein design [14], and peptide discovery [15] as well. For a more detailed description, see Appendix A.

[*]Corresponding Author, email: fangwu97@stanford.edu

The NeurIPS 2025 Workshop on AI Virtual Cells and Instruments: A New Era in Drug Discovery and Development (AI4D3 2025), San Diego, California, USA, 2025.

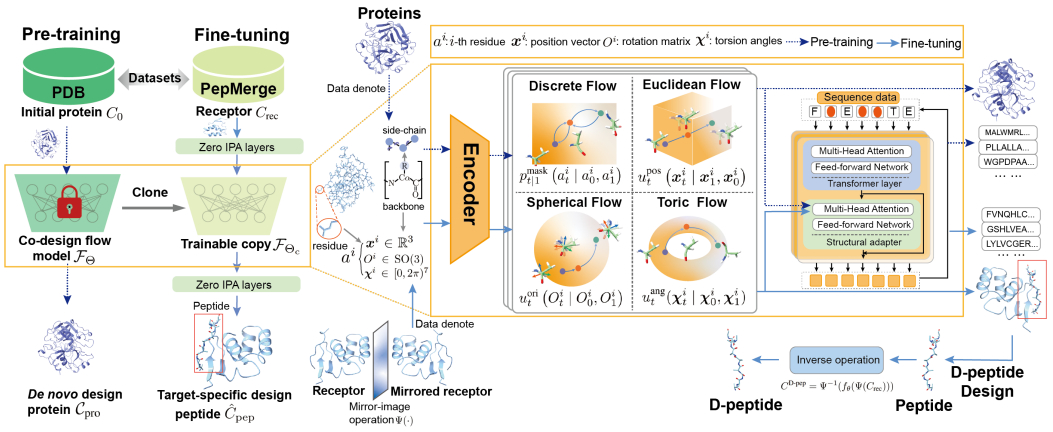

Figure 1: Illustrative overview of our D-Flow framework.

Despite emerging interest in developing 3D diffusion or flow to peptides, little attention has been paid to D-protein generation in machine learning (ML). It is worth mentioning that though every cell in the human body contains proteins, their cornerstone, 20 categories of amino acids, can exist in two stereoisomers: L (levo) and D (dextro), which are non-superimposable mirror images of each other [16], and this chirality characteristic is determined by the orientation of the functional groups around the alpha carbon. D-proteins are protein molecules whose polypeptide chains consist of D-amino acids and the achiral amino acid glycine. They can form specific heterochiral protein-protein interactions (PPIs) with natural L-protein targets and possess remarkable potential as therapeutics, and diagnostics because of their high bioorthogonality and stability [17]. However, in nature, the vast majority of peptides and proteins are made from L-amino acids, which poses a great challenge for DL mechanisms, as no training data is accessible.

To overcome this obstacle, we introduce a mirror-imaged multi-modality model, named D-Flow (see Fig. 1), to produce D-peptides. D-Flow characterizes the peptide sequence and structure using rigid backbone frames within the SE(3) manifold, the side-chain angles on the high-dimensional torus, and the discrete amino acid types on the categorical space [18]. Each modality has an analytical flow, and they jointly capture the distribution of full-atom peptide structure conditioned on the receptor. During the inference phase, a novel mirror-image algorithm is adopted to transform the L-receptor and convert its natural chirality. Moreover, as datasets containing complete 3D structures of peptide complexes are orders of magnitude smaller than sequence-only databases [19], the scarcity of high-quality peptide-receptor pairs inevitably constrains the performance of DL approaches. As a resolution, we propose to empower the flow model with large protein language models (PLMs) and a two-stage training pipeline. Specifically, a structural surgery is implemented on PLMs using a lightweight structural adapter [20]. Consequently, PLMs are endowed with structural awareness and can facilitate the regression of residue categories given intermediate-state conformation information. Secondly, D-Flow embraces a controlling toolkit [21] to achieve awareness of target proteins and transfer from the general protein design task to the binder design task with the minimum loss of the pretraining knowledge. Comprehensive experiments on the PepMerge benchmark [5] verify the effectiveness of D-Flow, particularly for its breakthrough in generating pure D-peptides. Besides, it outpaces existing DL baselines in generating L-peptides by a large margin. Our study opens up avenues for systematically exploring the DL-based mirror-image protein universe, paving the way for various design applications targeting L-proteins.

2 Methods

2.1 Flow for Full-atom Protein Generation

Discrete Flow for 1D Amino Acid Sequence. We provide the preliminary and background of *Proteins* and *Flow on Riemannian Manifolds* in Appendix B. Although flow models succeed in continuous spatial signals like images [22], they falter when applied to discrete sequential data. Two recent lines adapt diffusion or flow to the discrete setting: one embeds the discrete data in continuous

space and adopts the continuous diffusion or flow [5], and the other designs the transformation process immediately over categorical state spaces [18, 23]. In this work, we investigate both approaches and observe a significant advantage of discrete flow over the continuous version. Following Campbell et al. [18] and Gat et al. [23], we define FM in a discrete setting and consider a Continuous-Time discrete Markov Chain (CTMC) paradigm. The categorical variable a_t^i jumps between states in the amino acid type space $\mathcal{A}$ depending on a continuous time $t \in [0, 1]$. u_t, called the *probability velocity* as reminiscent of the velocity field in CNF, is defined as the rate of probability change of the current sample a_t^i in each of its $\|\mathcal{A}\| = 20$ categories. Thus, each category of the residue type $a_t^i \sim p_t$ is updated independently with the Euler step as $a_{t+\frac{1}{N}}^i \sim \delta_{a_t^i}(\cdot) + \frac{1}{N} u_t\left(a^i, \cdot\right)$. For a sufficiently large number of timesteps $N \to \infty$ and any potential state $z \in \mathcal{A}$, u_t satisfies

$$\sum_{a^i \in \mathcal{A}} u_t\left(a^i, z\right) = 0 \text{ and } u_t\left(a^i, z\right) \geq 0, \tag{1}$$

for $\forall i \in [n]$ and $a^i \neq z$. Then, the unconditional probability velocity u_t can be constructed as the marginalization of the conditional one $u_t\left(a^i, z \mid a_0^i, a_1^i\right)$. As a consequence, we attain the marginal velocity written as $u_t\left(a^i, z\right) = \sum_{a_0^i, a_1^i \in \mathcal{A}} u_t\left(a^i, z \mid a_0^i, a_1^i\right) p_t\left(a_0^i, a_1^i \mid z\right)$, which generates the probability path $p_t\left(a^i\right)$. Generally, the conditional probability paths can be represented as a convex sum of m conditional probabilities, namely, $p_t\left(a^i \mid a_0^i, a_1^i\right) = \sum_{j=1}^k \kappa_t^j p^j(a^i \mid a_0^i, a_1^i)$, where the schedulers κ_t^j are collectively non-negative and satisfy $\sum_j \kappa_t^j = 1$. A simple yet useful instance of conditional probability paths is $p_t\left(a^i \mid a_0^i, a_1^i\right) = \kappa_t \delta_{a_1^i}\left(a^i\right) + (1 - \kappa_t)\delta_{a_0^i}\left(a^i\right)$, where the scheduler κ_t satisfies $\kappa_0 = 0$, $\kappa_1 = 1$ and monotonically increases in t (*i.e.*, $\dot{\kappa}_t \geq 0$). This results in $p_0(a^i|a_0^i, a_1^i) = \delta_{a_0^i}(a^i)$ and $p_1(a^i|a_0^i, a_1^i) = \delta_{a_1^i}(a^i)$. Subsequently, we get:

$$u_t\left(a^i, z\right) = \frac{\dot{\kappa}_t}{1 - \kappa_t}\left[p_{1|t}\left(a^i \mid z\right) - \delta_z\left(a^i\right)\right], \tag{2}$$

where $p_{1|t}\left(a^i \mid z\right) = \sum_{a_0^i, a_1^i} \delta_{a_1^i}\left(a^i\right) p_t\left(a_0^i, a_1^i \mid z\right)$ is used as the notation for the probability denoiser. The training goal is to learn $p_{1|t}$ by minimizing the cross-entropy (CE) loss:

$$\mathcal{L}_{\text{aa}}(\theta) = \mathbb{E}_{i \in [n], t \sim \mathcal{U}(0,1), p_0\left(a^i\right), p_1\left(a^i\right), p_t\left(a^i \mid a_0^i, a_1^i\right)}\left[\log p_{1|t}\left(a_1^i \mid a_t^i\right)\right]. \tag{3}$$

where $\mathcal{U}(0, 1)$ is a uniform time distribution on $[0, 1]$, and $v^{\text{aa}}\left(a_t^i, t\right)$ is the neural network with parameters θ to approximate the true denoising distribution $p_{1|t}$. Given a noisy input $a_t^i \sim p_{t|1}\left(a^i \mid a_1^i, a_0^i\right)$, the model learns to predict the clean data point a_1^i. Here, rather than a linear interpolation towards a_1^i from a prior of uniform categorical distribution $p_{t|1}^{\text{unif}}\left(a_t^i \mid a_0^i, a_1^i\right) = \text{Cat}\left(t\delta_{a_1^i}(\cdot) + (1 - t)\frac{1}{\|\mathcal{A}\|}\right)$, we adopt an artificially introduced mask state M and the conditional path becomes [18]:

$$p_{t|1}^{\text{mask}}\left(a_t^i \mid a_0^i, a_1^i\right) = \text{Cat}\left(t\delta_{a_1^i}(\cdot) + (1 - t)\delta_{\text{M}}(\cdot)\right). \tag{4}$$

Multi-modality FM for 3D Protein Structures. Apart from the sequential amino acid type, we follow common practice [4, 5, 14, 18] and construct different probability flows containing Euclidean, spherical, and toric FMs for positions $\boldsymbol{x} \in \mathbb{R}^{n \times 3}$, orientations $O \subseteq \text{SO}(3)$, and torsion angles $\boldsymbol{\chi} \in [0, 2\pi)^{n \times 7}$, respectively. To be specific, a vanilla Gaussian FM on Euclidean manifolds with the prior $\mathcal{N}(0, \boldsymbol{I}_3)$ is leveraged to generate $\boldsymbol{x}^i$. As for the 3D rotation group $\text{SO}(3)$, it is also a smooth Riemannian manifold with its tangent space $\mathfrak{so}(3)$ forming a Lie algebra consisting of skew-symmetric matrice. We establish flows based on the geodesics in the context of $\text{SO}(3)$ and select the uniform distribution over $\text{SO}(3)$ as the prior $p\left(O_0^i\right)$. For torsion angles $\boldsymbol{\chi}^i \in [0, 2\pi)^7$, each can be represented as a point on the unit circle $\mathbb{S}^1$. Thus, $\boldsymbol{\chi}^i$ lies on the 7-dimensional hypertorus $\mathbb{T}^7 = (\mathbb{S}^1)^7$ as the Cartesian product of all seven unit circles. This flat torus $\mathbb{T}^7$ can be viewed as the quotient space $(\mathbb{R}^7 \mod 2\pi\mathbb{Z})^7$ that inherits the Riemannian metric from Euclidean space, where the uniform distribution on $[0, 2\pi)^7$ is utilized as the prior. The conditional flows are prescribed as:

$$\phi_t^{\text{pos}}\left(\boldsymbol{x}_0^i, \boldsymbol{x}_1^i\right) = t\boldsymbol{x}_1^i + (1 - t)\boldsymbol{x}_0^i, \quad \boldsymbol{x}_0^i \sim \mathcal{N}\left(0, \boldsymbol{I}_3\right), \tag{5}$$

$$\phi_t^{\text{ori}}\left(O_0^i, O_1^i\right) = \exp_{O_0^i}\left(t \log_{O_0^i}\left(O_1^i\right)\right), \quad O_0^i \sim \mathcal{U}(\text{SO}(3)), \tag{6}$$

$$\phi_t^{\text{ang}}(\boldsymbol{\chi}_0^i, \boldsymbol{\chi}_1^i) = \left[t\boldsymbol{\chi}_1^i + (1 - t)\boldsymbol{\chi}_0^i\right] \mod 2\pi, \quad \boldsymbol{\chi}_0^i \sim \mathcal{U}\left([0, 2\pi)^7\right), \tag{7}$$

where $\exp(\cdot)$ and $\log(\cdot)$ are the exponential and logarithm maps on $\mathrm{SO}(3)$ that can be computed efficiently using Rodrigues' formula [5, 14]. Subsequently, the conditional vector fields u_t^{pos}, u_t^{ori}, and u_t^{ang} can be obtained by taking the time derivative of linear flows ϕ_t^{pos}, ϕ_t^{ori}, and ϕ_t^{ang} using Independent Coupling techniques:

$$u_t^{\text{pos}}\left(\boldsymbol{x}_t^i \mid \boldsymbol{x}_1^i, \boldsymbol{x}_0^i\right) = \boldsymbol{x}_1^i - \boldsymbol{x}_0^i = \frac{\boldsymbol{x}_1^i - \boldsymbol{x}_t^i}{1-t}. \tag{8}$$

$$u_t^{\text{ori}}\left(O_t^i \mid O_0^i, O_1^i\right) = \frac{\log_{O_t^i}\left(O_1^i\right)}{1-t}, \tag{9}$$

$$u_t^{\text{ang}}(\boldsymbol{\chi}_t^i \mid \boldsymbol{\chi}_0^i, \boldsymbol{\chi}_1^i) = \left(\frac{\boldsymbol{\chi}_1^i - \boldsymbol{\chi}_t^i}{1-t} + \pi\right) \mod 2\pi - \pi, \tag{10}$$

Ultimately, time-dependent equivariant neural networks $v^{\text{pos}}(\cdot)$, $v^{\text{ori}}(\cdot)$, and $v^{\text{ang}}(\cdot)$, with the same backbone architecture as $v^{\text{aa}}(\cdot)$ but different head predictors, are employed to approximate the conditional gradient fields u_t^{pos}, u_t^{ori}, and u_t^{ang}. The FM objectives are:

$$\mathcal{L}_{\text{pos}}(\theta) = \mathbb{E}_{i\in[n], t\sim\mathcal{U}(0,1), p\left(\boldsymbol{x}_1^i\right), p\left(\boldsymbol{x}_0^i\right), p\left(\boldsymbol{x}_t^i \mid \boldsymbol{x}_0^i, \boldsymbol{x}_1^i\right)} \left\| v^{\text{pos}}\left(\boldsymbol{x}_t^i, t\right) - \left(\boldsymbol{x}_1^i - \boldsymbol{x}_0^i\right)\right\|_2^2, \tag{11}$$

$$\mathcal{L}_{\text{ori}}(\theta) = \mathbb{E}_{i\in[n], t\sim\mathcal{U}(0,1), p\left(O_1^i\right), p\left(O_0^i\right), p\left(O_t^i \mid O_0^i, O_1^i\right)} \left\| v^{\text{ori}}\left(O_t^i, t\right) - \frac{\log_{O_t^i}\left(O_1^i\right)}{1-t}\right\|_{\mathrm{SO}(3)}^2, \tag{12}$$

$$\mathcal{L}_{\text{ang}}(\theta) = \mathbb{E}_{i\in[n], t\sim\mathcal{U}(0,1), p\left(\boldsymbol{\chi}_1^i\right), p\left(\boldsymbol{\chi}_0^i\right), p\left(\boldsymbol{\chi}_t^i \mid \boldsymbol{\chi}_0^i, \boldsymbol{\chi}_1^i\right)} \left\| v^{\text{ang}}\left(\boldsymbol{\chi}_t^i, t\right) - \left(\boldsymbol{\chi}_1^i - \boldsymbol{\chi}_0^i\right)\right\|_2^2. \tag{13}$$

The overall FM objective is the loss sum of all different modalities as $\mathcal{L}_{\text{CFM}} = \lambda_{\text{pos}}\mathcal{L}_{\text{pos}} + \lambda_{\text{ori}}\mathcal{L}_{\text{ori}} + \lambda_{\text{ang}}\mathcal{L}_{\text{ang}} + \lambda_{\text{aa}}\mathcal{L}_{\text{aa}}$, where λ_* are balance weight hyperparameters. Two additional losses are also imposed concerning the backbone atoms and the distance matrix [14].

During the inference phase, we first sample from several distinct priors, *i.e.*, $\boldsymbol{x}_0^i \sim \mathcal{N}(0, \boldsymbol{I}_3)$, $O_0^i \sim \mathcal{U}(\mathrm{SO}(3))$, and $\boldsymbol{\chi}_0^i \sim [0, 2\pi)^7$. After that, we solve the probability flow with learned predictors $v^{\text{pos}}(\cdot)$, $v^{\text{ori}}(\cdot)$, and $v^{\text{ang}}(\cdot)$ using the N-step forward Euler method to get the position, orientations, and torsion angles of i-th residue with $t = \left\{0, \ldots, \frac{N-1}{N}\right\}$:

$$\boldsymbol{x}_{t+\frac{1}{N}}^i = \boldsymbol{x}_t^i + \frac{1}{N} v^{\text{pos}}\left(\boldsymbol{x}_t^i, t\right),\ O_{t+\frac{1}{N}}^i = \exp_{O_t^i}\left(\frac{1}{N} v^{\text{ori}}\left(O_t^i, t\right)\right),\ \boldsymbol{\chi}_{t+\frac{1}{N}}^i = \left[\boldsymbol{\chi}_t^i + \frac{1}{N} v^{\text{ang}}\left(\boldsymbol{\chi}_t^i, t\right)\right] \mod 2\pi. \tag{14}$$

2.2 Parameterization with Adapter-guided Protein Language Models

To efficiently learn $(a^i, \boldsymbol{x}^i, O^i, \boldsymbol{\chi}^i)$ for every residue, we adopt and improve the FramePred architecture [11, 14]. It incorporates Invariant Point Attention (IPA) [24] to encode spatial features and ensure equivariance. IPA consists of multiple Transformer layers [25] to capture sequence-level attributes. In addition, considering the periodicity, torsion angles $\boldsymbol{\chi} \in [0, 2\pi)^7$ are flexibly encoded by applying multi-frequency sine and cosine transformations [5], which are also fused with the timestep embedding and residue sequence embedding into IPAs.

PLMs capture the evolutionary patterns from large-scale sequence data, and this knowledge is proven to be supportive of protein folding [26] and inverse design [20], which is evidently beneficial for our co-design task as well. Some prior studies [19] integrate this knowledge by immediately appending the geometric networks to PLMs. Drawing inspiration from LM-Design [20], we employ a lightweight structural adapter to endow PLMs with structural awareness. In our approach, if the IPA output is denoted as $\boldsymbol{h}_{\text{IPA}} \in \mathbb{R}^{\psi_{\text{IPA}}}$, the l-th layer's attention is computed as: $\boldsymbol{o} = \mathrm{softmax}\left(\frac{\boldsymbol{h}_{\text{seq}}\boldsymbol{W}_Q \cdot \boldsymbol{h}_{\text{IPA}}^\top \boldsymbol{W}_K^\top}{\sqrt{\psi_{\text{seq}}}}\right)\boldsymbol{h}_{\text{IPA}}\boldsymbol{W}_V$, where $\boldsymbol{h}_{\text{seq}} \in \mathbb{R}^{\psi_{\text{seq}}}$ is the sequential embedding of the last $(l-1)$-th layer. $\boldsymbol{W}_K$, $\boldsymbol{W}_Q$, and $\boldsymbol{W}_V$ are trainable weights for key, query, and values, respectively. This module allows the usage of sequence- and structure-based information for protein understanding.

2.3 Controlling Flow Matching with Target Proteins

Existing *de novo* design algorithms [11, 14] are mainly targeting monomers or biomolecules without receptors. Some recent studies [4, 5, 27] have considered binder design but their performance is strictly

restricted by the number of available complex structures (*e.g.*, SabDAB [28] and PepMerge [5]). To bridge the gap and exploit all crystal structures, we pretrain our flow model on the vast amount of general proteins (*i.e.*, Protein Data Bank) and then fine-tune it on the target-specific design challenges.

To begin with, $\mathcal{F}_\Theta : \mathcal{P}_0 \to \mathcal{C}_{\text{pro}}$ is a co-design flow model without target awareness, which can *de novo* design a protein C_{pro} from any initial protein C_0 that is drawn from a prior distribution p_0. It is trained on general proteins to approximate $p(a, \boldsymbol{x}, O, \boldsymbol{\chi})$. To enable awareness of $\mathcal{F}_\Theta$ to the receptor C_{rec}, we borrow ideology from ControlNet [21], locking (freeze) the parameters Θ of the original block and simultaneously cloning the block to a trainable copy with parameters Θ_c. The trainable copy takes a receptor protein C_{rec} as input and is connected to the locked model with zero IPAs denoted as $\mathcal{Z}(\cdot)$. In particular, both the weight and the bias in $\mathcal{Z}(\cdot)$ are initialized to zeros. In practice, two instances of zero IPAs are used with parameters Θ_{z1} and Θ_{z2}, respectively:

$$\hat{C}_{\text{pep}} = \mathcal{F}_\Theta(C_0) + \mathcal{Z}_{\Theta_{\text{z2}}}\left(\mathcal{F}_{\Theta_{\text{c}}}\left(C_0 + \mathcal{Z}_{\Theta_{\text{z1}}}\left(C_{\text{rec}}\right)\right)\right), \tag{15}$$

where $\hat{C}_{\text{pep}}$ is the designed peptide that is expected to bind with C_{rec}. In the first training step, since both the weight and bias parameters of IPAs are initialized to zero, both $\mathcal{Z}(\cdot)$ terms evaluate to zero, and $\hat{C}_{\text{pep}} = \hat{C}_{\text{pro}}$. In this way, harmful noise cannot influence the hidden states of the network layers in the trainable copy when the training starts. Zero IPA layers protect this backbone by eliminating random noise as gradients in the initial training steps.

2.4 D-peptide Design

Mirror-image peptides and proteins, composed of D-amino acids and the achiral glycine, have garnered significant attention as potential therapeutic and enzymatic tools due to their remarkable resistance to enzymatic digestion by natural-chirality enzymes and their exceptional biostability [29]. However, the development of mirror-image biological systems and related applications faces challenges, primarily due to the lack of effective methods for designing mirror-image (D-) proteins [16].

In this work, we generate mini D-protein binders, specifically D-peptides, using the mirror-image algorithm [17]. The mirror-image operation, denoted as $\Psi(\cdot)$, transforms any given protein structure into its corresponding mirror image. $\Psi(\cdot)$ inverts the spatial configuration of the protein, effectively reflecting it across an imaginary plane while preserving the relative distances and angles between atoms. In our approach, we first apply $\Psi(\cdot)$ to the receptor protein C_{rec}, converting it into its mirror-image counterpart. Then, we pass this mirrored receptor through a flow model $f_\theta(\cdot) : \mathcal{C}_{\text{rec}} \to \mathcal{C}_{\text{pep}}$, which is trained on natural-chirality (L-) proteins, to yield the corresponding peptide. Finally, we apply the inverse operation $\Psi^{-1}(\cdot)$ to convert the resulting peptide back to the D-configuration. Formally, the generation process can be expressed as:

$$C^{\text{D-pep}} = \Psi^{-1}(f_\theta(\Psi(C_{\text{rec}}))). \tag{16}$$

Notably, this reflection $\Psi(\cdot)$ can be represented by a parity transformation, written as $\Psi(C) = MC$. where M is a reflection matrix and satisfies $M^T M = I$ and $\det(M) = -1$. Unlike rotation matrices (which have determinant +1 and preserve orientation), a reflection matrix M flips orientation and reverses chirality (*e.g.*, converting an L-protein to a D-protein). Since mirroring preserves pairwise distances and angles but reverses chirality, it commutes with translations but interacts with rotations depending on their handedness. Therefore, the following conclusion holds:

Theorem 1 *If f_θ is roto-translational equivariant, then the mirror-image algorithm in Equ. 16 maintains the equivariance property. That is $\Psi^{-1}\left(f_\theta\left(\Psi\left(RC_{\text{rec}} + t\right)\right)\right) = R\Psi^{-1}\left(f_\theta\left(\Psi\left(C_{\text{rec}}\right)\right)\right) + t$., where (R, t) is the rigid transformation.*

It indicates that applying any rigid transformation (rotation R and translation t) to the input receptor structure C_{rec} before running it through the mirror-image algorithm results in the same transformation applied to the output $C^{\text{D-pep}}$. The proof of the theorem is in C.

3 Results

We execute two types of experiments to validate the effectiveness of our D-Flow approach. The first, detailed in Sec. 3.1, involves the conventional co-design challenge for L-proteins, where models generate peptides' both sequences and structures based on a given receptor binding site. The second

Table 1: Evaluation of methods in the traditional L-peptide sequence-structure co-design task, where metrics are divided into three main categories. The **best** and suboptimal results are labeled boldly and underlined, respectively.

	Geometry				Energy		Design	
	AAR % ↑	RMSD Å ↓	SSR % ↑	BSR % ↑	Stb. % ↑	Aff. % ↑	Des. % ↑	Div. ↑
RFdiffusion [35]	40.14	4.17	63.86	26.71	26.82	16.53	**78.52**	0.38
ProteinGen [36]	45.82	4.35	29.15	24.62	23.48	13.47	71.82	0.54
Diffusion [27]	47.04	3.28	74.89	49.83	15.34	17.13	48.54	0.57
PPIFlow [4]	48.35	3.59	68.13	25.94	15.77	12.08	46.53	0.51
PepFlow-Bb [5]	50.46	2.30	82.17	82.17	14.04	18.10	50.03	**0.64**
PepFlow-Seq [5]	53.25	2.21	85.22	85.19	19.20	19.39	56.04	0.50
PepFlow-Ang [5]	51.25	2.07	83.46	86.89	18.15	21.37	65.22	0.42
D-Flow	**58.69**	**1.63**	**89.02**	**88.47**	**26.85**	**24.31**	75.14	0.60

Table 2: Ablation studies on the effects of each module. To be specific, DFM is the abbreviation of *discrete flow matching*, PLM is the usage of *protein language models*, Pretrain is the transfer learning from PDB to PepMerge, and CN is the employment of the *ControlNet* technique.

	DFM	PLM	Pretrain	CN	AAR% ↑	Δ	RMSD↓	Δ
1	-	-	-	-	50.43	–	2.34	–
2	✓	-	-	-	51.86	+1.43	2.09	- 0.25
3	✓	✓	-	-	52.44	+0.58	2.01	-0.08
4	✓	✓	✓	-	58.37	+5.93	1.78	-0.23
5	✓	✓	✓	✓	58.69	+0.32	1.63	-0.15

experiment, described in Sec. 3.2, focuses on generating D-proteins. For benchmarking, we use the **PepMerge** dataset [5], derived from PepBDB [30] and Q-BioLip [31]. To ensure a fair comparison with prior work [5], we cluster peptide-protein complexes based on 40% sequence identity using MMseqs2 [32], after removing duplicates and applying empirical filters (*e.g.*, resolution < 4Å, peptide length between 3 and 25). This process resulted in 8,365 non-redundant complexes across 292 clusters. For consistent comparison, we use the same test set as Li et al. [5], containing 10 clusters and 158 complexes. Pretraining data contain monomers between lengths 60 and 384 with resolution < 5Å from PDB [33] on August 8, 2021, ensuring no data leakage for peptides. The data is then filtered by including proteins with high secondary structure compositions only. Monomers with more than 50% loops are also removed using DSSP [34], resulting in 18,684 proteins.

3.1 Unconditioned Sequence-structure Co-design

Baselines. We select two categories of state-of-the-art protein design methods as baselines. The first category disregards side-chain conformations and includes approaches such as RFDiffusion [35] and ProteinGen [36]. RFDiffusion generates protein backbones, with sequences predicted afterward using ProteinMPNN [37], while ProteinGen improves on RFDiffusion by jointly sampling both backbones and sequences. The second category considers full-atom protein generation, containing Diffusion [27], PPIFlow [4], and PepFlow [5]. PepFlow has three variants based on whether backbones, sequences, and side-chain angles are sampled. For each test case, we sample 64 peptides simultaneously.

Main Results. Generated peptides are evaluated across three key aspects:*Geometry*, *Energy*, and *Design* (Detailed see Appendix D).As documented in Tab. 1, D-Flow generates peptide sequences with the closest resemblance to native ones, achieving a high amino acid identity of 58.69%, a 14.51% improvement over PepFlow. It excels in all geometry-related metrics, highlighting its close alignment with the binding sites of native peptides. Besides, D-Flow scores 75.14% in designability and 0.60 in diversity, maintaining a good balance between structural fidelity and variety. Moreover, D-Flow also demonstrates superior energy-based properties, achieving the best stability score at 26.85% and the best affinity score at 24.31%, which are critical for the formation of strong and stable complexes. In a nutshell, D-Flow outperforms all baselines across most key metrics, indicating its strength in consistently producing the most accurate peptide sequences and structures, along with the optimal stability and affinity for their targets.

We also navigate the contribution of D-Flow components, containing the discrete flow matching (DFM), structural-adapted PLM, the pretraining on PDB, and the ControlNet-style transfer learning

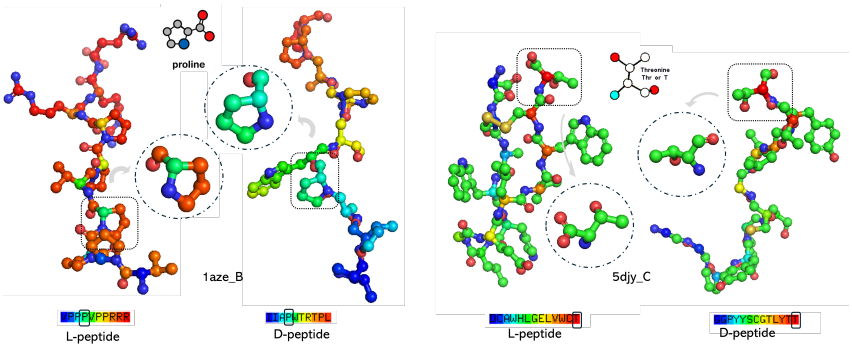

Figure 2: Perspective view of D-residues within two randomly selected D-peptides generated by our D-flow, where all stereogenic alpha carbons to the amino group in D-peptides have the D-configuration.

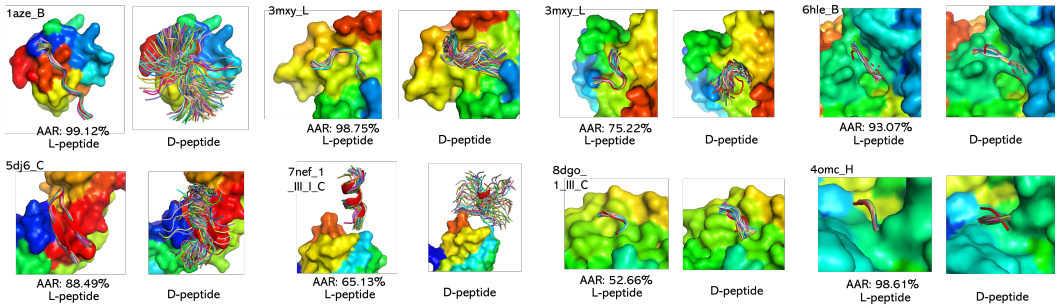

Figure 3: Visualization of L- and D-peptide examples generated by D-Flow with the same target receptor, where the red one is the ground truth L-peptide. Despite that D-flow achieves particularly high recovery of both amino acid sequences and structures for L-peptides, its sampling space of D-peptides is relatively diverse and unconstrained. This phenomenon of diversity can be explained by the distribution mismatch between the training and inference data, loss of evolutionary constraints, asymmetric energy landscape exploration, and stereochemical environment adaptation.

technique. Tab. 2 reports the ablation study results. It can be observed that the integration of additional unlabeled structural data yielded the most substantial improvement, increasing AAR by 5.93%. This supports our hypothesis that limited binding complex data significantly constrains generative models' *de novo* design capabilities. While incorporating 1D sequence information improved D-Flow's performance by 0.58%, this impact was less pronounced than in previous protein-related tasks [19, 38] like ligand docking [39, 40], ligand efficacy prediction [41], or model quality assessment. We attribute this to PLMs' inherent limitations with short sequences like peptides, compared to their effectiveness with longer proteins. This gap generally arises because PLMs learn contextual relationships by observing long-range dependencies and conserved motifs – key elements that characterize protein families and evolutionary relationships [42]. In short sequences, such as peptides, these relationships are limited due to their short length and lack of structure, making it harder for PLMs to capture evolutionary patterns effectively. Peptides often lack secondary and tertiary structure information present in longer sequences, which is critical for understanding function and evolutionary context. Attention-based models like transformers [25] distribute attention weights over the input length. Shorter sequences provide fewer tokens to which attention can be allocated, limiting the richness of learned representations. Regarding other components, DFM demonstrated superior performance over the probability simplex mechanism [5] for amino acid type representation, improving AAR by 1.43%. The control technique provided a modest 0.32% enhancement in AAR.

3.2 D-peptide Design

Generation of Pure D-peptides. Our analysis confirms that D-Flow successfully produces *pure D-peptides* in all cases (Fig. 2). Remarkably, this achievement occurs without any D-protein training

corpus, requiring only our specialized post-processing technique $\Psi(\cdot)$. To understand it, as D-Flow gradually moves the peptide's sequence and structure $(a, \boldsymbol{x}, O, \boldsymbol{\chi})$ from a prior distribution p_0 to the target data distribution p_1 at a speed of v_t during $t \in [0, 1]$, its moving trajectory is heavily dependent on the anchor fixed receptor $\mathcal{C}_{\text{rec}}$. The velocity vectors $v^{\text{pos}}(\cdot), v^{\text{ori}}(\cdot)$, and $v^{\text{ang}}(\cdot)$ for Cα locations, frame orientations, and angles are computed based on the interactions between the intermediate state of peptide $\mathcal{C}^t_{\text{pep}}$ and the target protein $\mathcal{C}_{\text{rec}}$. Once a D-chirality receptor $\Psi(\mathcal{C}_{\text{rec}})$ is provided as the condition, our D-Flow $f_\theta(\cdot)$ captures the chirality nuance of this given input. As a response, it adjusts $v^{\text{pos}}(\cdot), v^{\text{ori}}(\cdot)$, and $v^{\text{ang}}(\cdot)$ to accord with this D-target protein, which leads to a product with considerable D-residues.

Model Robustness Assessment. We examined the robustness of D-Flow by testing multiple mirror-image approaches for the input receptor. Despite there being infinite possible mirror transformations, we focused on the three fundamental axes (x, y, and z). Our results demonstrate that the chirality of generated peptides remains consistently D-configured regardless of the mirroring approach used.

Conformational Distribution Analysis. Fig. 3 visualizes a few examples of designed L- and D-peptides. While D-Flow successfully generates both configurations, our analysis reveals a striking contrast: L-peptides show remarkable consistency in their structural and sequence distributions, whereas D-peptides exhibit significantly higher variability, particularly in longer sequences. This phenomenon can be attributed to several fundamental factors.

Training Distribution Mismatch. Our model, trained exclusively on L-peptides, has learned to navigate L-peptide conformations' natural distribution and energy landscape. When presented with a D-receptor through the mirror-image algorithm, the model must operate in a conformational space absent from its training distribution. This distributional shift forces the model to explore the D-space with less confidence, increasing structural diversity. *Loss of Evolutionary Constraints.* L-peptides in our training data reflect millions of years of evolutionary optimization, containing implicit biases about energetically favorable conformations [43]. While the mirror-image transformation preserves theoretical symmetry, these evolutionary constraints don't translate perfectly to the D-space, potentially contributing to the observed variability. *Asymmetric Energy Landscape Exploration.* The FM process relies on smooth transitions across the conformational space [7]. However, the mirror-image transformation introduces subtle numerical asymmetries in the computational representation of D-peptides. These asymmetries can affect how the model explores the energy landscape, leading to more diverse intermediate states and final structures [44]. *Stereochemical Environment Adaptation.* The significant change in the stereochemical environment during mirroring poses a unique challenge. The flow model must adapt its learned representations of molecular interactions to account for the inverted chirality [45]. This adaptation process isn't perfect, resulting in broader conformational sampling and increased structural diversity in the generated D-peptides.

These challenges align with recent findings in protein design that highlight the challenges of transferring learned molecular representations across different stereochemical spaces [46]. The energy landscape learned for L-peptides doesn't directly translate to the D-space due to fundamental differences in molecular interactions and packing arrangements. This leads to broader structural exploration and increased diversity in the generated structures. This analysis suggests potential avenues for improvement, such as incorporating symmetry-aware constraints or developing specialized adaptation mechanisms for D-peptide generation, which could help stabilize the output distribution while maintaining the desired chirality.

4 Conclusion

This work introduces a multi-modality flow model dubbed D-Flow to generate D-peptides from scratch. D-Flow produces peptides containing a considerable proportion of D-residues through a novel mirror-image algorithm. In addition, it incorporates an adapter-guided protein language model with structural awareness and employs a ControlNet-style mechanism to bridge the gap between pretraining and fine-tuning stages. By utilizing a mirror-image adaptation mechanism, D-Flow successfully generates pure D-peptides, achieving comparable or superior stability and binding affinity relative to their L-peptide counterparts. This represents a significant advancement in generative peptide design and demonstrates strong performance across key metrics.

References

[1] Lei Wang et al. Therapeutic peptides: current applications and future directions. *Signal transduction and targeted therapy*, 7(1):48, 2022.

[2] Fang Wu, Bozhen Hu, and Stan Z Li. Generalized implicit neural representations for dynamic molecular surface modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 877–885, 2025.

[3] Fang Wu, Shuting Jin, Jianmin Wang, Zerui Xu, Jinbo Xu, Brian Hie, et al. Surfdesign: Effective protein design on molecular surfaces.

[4] Haitao Lin et al. Ppflow: Target-aware peptide design with torsional flow matching. *bioRxiv*, pages 2024–03, 2024.

[5] Jiahan Li et al. Full-atom peptide design based on multi-modal flow matching. *arXiv preprint arXiv:2406.00735*, 2024.

[6] Fang Wu. A semi-supervised molecular learning framework for activity cliff estimation. In *33rd International Joint Conference on Artificial Intelligence, IJCAI*, pages 6080–6088, 2024.

[7] Yaron Lipman et al. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

[8] Fang Wu and Stan Z Li. Diffmd: A geometric diffusion model for molecular dynamics simulations. In *AAAI Conference on Artificial Intelligence*, volume 37, pages 5321–5329, 2023.

[9] Brian L Trippe et al. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. *arXiv preprint arXiv:2206.04119*, 2022.

[10] Kevin E Wu et al. Protein structure generation via folding diffusion. *Nature communications*, 15(1):1059, 2024.

[11] Jason Yim et al. Se (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023.

[12] Jiaqi Guan, Wesley Wei Qian, Xingang Peng, Yufeng Su, Jian Peng, and Jianzhu Ma. 3d equivariant diffusion for target-aware molecule generation and affinity prediction. *arXiv preprint arXiv:2303.03543*, 2023.

[13] Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, et al. Abdiffuser: full-atom generation of in-vitro functioning antibodies. *Advances in Neural Information Processing Systems*, 36, 2024.

[14] Jason Yim et al. Fast protein backbone generation with se (3) flow matching. *arXiv preprint arXiv:2310.05297*, 2023.

[15] Vishva Saravanan Ramasubramanian et al. A hybrid diffusion model for stable, affinity-driven, receptor-aware peptide generation. *bioRxiv*, pages 2024–03, 2024.

[16] Alexander J Lander, Yi Jin, and Louis YP Luk. D-peptide and d-protein technology: recent advances, challenges, and opportunities. *ChemBioChem*, 24(4):e202200537, 2023.

[17] Ke Sun et al. Accurate de novo design of heterochiral protein-protein interactions. *Cell Research*, pages 1–13, 2024.

[18] 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. *arXiv preprint arXiv:2402.04997*, 2024.

[19] Fang Wu, Lirong Wu, Dragomir Radev, Jinbo Xu, and Stan Z Li. Integration of pre-trained protein language models into geometric deep learning networks. *Communications Biology*, 6 (1):876, 2023.

[20] Zaixiang Zheng et al. Structure-informed language models are protein designers. In *International conference on machine learning*, pages 42317–42338. PMLR, 2023.

[21] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.

[22] Robin Rombach et al. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

[23] Itai Gat et al. Discrete flow matching. *arXiv preprint arXiv:2407.15595*, 2024.

[24] John Jumper et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596 (7873):583–589, 2021.

[25] Ashish Vaswani et al. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[26] Zeming Lin et al. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022.

[27] Shitong Luo et al. Antigen-specific antibody design and optimization with diffusion-based generative models. *bioRxiv*, 2022.

[28] James Dunbar et al. Sabdab: the structural antibody database. *Nucleic acids research*, 42(D1): D1140–D1146, 2014.

[29] Guanwei Zhang and Ting F Zhu. Mirror-image trypsin digestion and sequencing of d-proteins. *Nature Chemistry*, 16(4):592–598, 2024.

[30] Zeyu Wen, Jiahua He, Huanyu Tao, and Sheng-You Huang. Pepbdb: a comprehensive structural database of biological peptide–protein interactions. *Bioinformatics*, 35(1):175–177, 2019.

[31] Hong Wei, Wenkai Wang, Zhenling Peng, and Jianyi Yang. Q-biolip: A comprehensive resource for quaternary structure-based protein–ligand interactions. *Genomics, Proteomics & Bioinformatics*, 22(1), 2024.

[32] 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.

[33] Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1):235–242, 2000.

[34] Wolfgang Kabsch and Christian Sander. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers: Original Research on Biomolecules*, 22(12):2577–2637, 1983.

[35] Joseph L Watson et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

[36] Sidney Lyayuga Lisanza et al. Joint generation of protein sequence and structure with rosettafold sequence space diffusion. *bioRxiv*, pages 2023–05, 2023.

[37] Justas Dauparas et al. Robust deep learning–based protein sequence design using proteinmpnn. *Science*, 378(6615):49–56, 2022.

[38] Fang Wu and Stan Z Li. Dynamics-inspired structure hallucination for protein-protein interaction modeling. *Transactions on Machine Learning Research*, 2025.

[39] Yize Jiang, Xinze Li, Yuanyuan Zhang, Jin Han, Youjun Xu, Ayush Pandit, Zaixi Zhang, Mengdi Wang, Mengyang Wang, Chong Liu, et al. Posex: Ai defeats physics approaches on protein-ligand cross docking. *arXiv preprint arXiv:2505.01700*, 2025.

[40] Houtim Lai, Longyue Wang, Ruiyuan Qian, Junhong Huang, Peng Zhou, Geyan Ye, Fandi Wu, Fang Wu, Xiangxiang Zeng, and Wei Liu. Interformer: an interaction-aware model for protein-ligand docking and affinity prediction. *Nature Communications*, 15(1):10223, 2024.

[41] Arthur Deng, Karsten Householder, Fang Wu, Sebastian Thrun, K Christopher Garcia, and Brian Trippe. Predicting mutational effects on protein binding from folding energy. *arXiv preprint arXiv:2507.05502*, 2025.

[42] Alexander Rives et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *National Academy of Sciences*, 118(15):e2016239118, 2021.

[43] Arnon Levy and Yair Levy. Evolutionary debunking arguments meet evolutionary science. *Philosophy and Phenomenological Research*, 100(3):491–509, 2020.

[44] Peter G Wolynes. Symmetry and the energy landscapes of biomolecules. *National Academy of Sciences*, 93(25):14249–14255, 1996.

[45] Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022.

[46] Alexey Strokach and Philip M Kim. Deep generative modeling for protein design. *Current opinion in structural biology*, 72:226–236, 2022.

[47] Xiangru Tang, Howard Dai, Elizabeth Knight, Fang Wu, Yunyang Li, Tianxiao Li, and Mark Gerstein. A survey of generative ai for de novo drug design: new frontiers in molecule and protein generation. *Briefings in Bioinformatics*, 25(4), 2024.

[48] Fang Wu and Stan Z Li. A hierarchical training paradigm for antibody structure-sequence co-design. *Advances in Neural Information Processing Systems*, 36, 2024.

[49] Fang Wu and Stan Z Li. Surface-vqmae: Vector-quantized masked auto-encoders on molecular surfaces. In *International Conference on Machine Learning*, pages 53619–53634. PMLR, 2024.

[50] Fang Wu. Diffantiseq: A controllable diffusion model for efficient antibody library design. In *LLM for Scientific Discovery: Reasoning, Assistance, and Collaboration*, 2025.

[51] Xiangru Tang, Xinwu Ye, Fang Wu, Daniel Shao, Dong Xu, and Mark Gerstein. Bc-design: A biochemistry-aware framework for highly accurate inverse protein folding. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*.

[52] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Neurips*, 32, 2019.

[53] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

[54] Alexander Tong, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.

[55] Ricky TQ Chen and Yaron Lipman. Riemannian flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2023.

[56] Aram-Alexandre Pooladian et al. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.

[57] Xiangzhe Kong, Wenbing Huang, and Yang Liu. Conditional antibody design as 3d equivariant graph translation. *arXiv preprint arXiv:2208.06073*, 2022.

[58] Shitong Luo, Yufeng Su, Zuofan Wu, Chenpeng Su, Jian Peng, and Jianzhu Ma. Rotamer density estimator is an unsupervised learner of the effect of mutations on protein-protein interaction. *bioRxiv*, pages 2023–02, 2023.

[59] Xiangzhe Kong, Wenbing Huang, and Yang Liu. Full-atom peptide design with geometric latent diffusion. *arXiv preprint arXiv:2402.13555*, 2024.

[60] Vadim A Davankov. Biological homochirality on the earth, or in the universe? a selective review. *Symmetry*, 10(12):749, 2018.

[61] Yuki Goto, Hiroshi Murakami, and Hiroaki Suga. Initiating translation with d-amino acids. *Rna*, 14(7):1390–1398, 2008.

[62] Jyumpei Kobayashi, Yasuhiro Shimizu, Yuta Mutaguchi, Katsumi Doi, and Toshihisa Ohshima. Characterization of d-amino acid aminotransferase from lactobacillus salivarius. *Journal of Molecular Catalysis B: Enzymatic*, 94:15–22, 2013.

[63] Yasushi Ogasawara and Tohru Dairi. Peptide epimerization machineries found in microorganisms. *Frontiers in microbiology*, 9:156, 2018.

[64] RC deL Milton, SCF Milton, and SBH Kent. Total chemical synthesis of a d-enzyme: the enantiomers of hiv-1 protease show reciprocal chiral substrate specificity. *Science*, 256(5062): 1445–1448, 1992.

[65] Suwei Dong, Ji-Shen Zheng, Yiming Li, Huan Wang, Gong Chen, Yongxiang Chen, Gemin Fang, Jun Guo, Chunmao He, Honggang Hu, et al. Recent advances in chemical protein synthesis: method developments and biological applications. *Science China Chemistry*, 67(4): 1060–1096, 2024.

[66] Ton NM Schumacher, Lorenz M Mayr, Daniel L Minor Jr, Michael A Milhollen, Michael W Burgess, and Peter S Kim. Identification of d-peptide ligands through mirror-image phage display. *Science*, 271(5257):1854–1857, 1996.

[67] Hao-Nan Chang, Bei-Yuan Liu, Yun-Kun Qi, Yang Zhou, Yan-Ping Chen, Kai-Mai Pan, Wen-Wen Li, Xiu-Man Zhou, Wei-Wei Ma, Cai-Yun Fu, et al. Blocking of the pd-1/pd-l1 interaction by ad-peptide antagonist for cancer immunotherapy. *Angewandte Chemie International Edition*, 54(40):11760–11764, 2015.

[68] Xiuman Zhou, Chao Zuo, Wanqiong Li, Weiwei Shi, Xiaowen Zhou, Hongfei Wang, Shaomeng Chen, et al. A novel d-peptide identified by mirror-image phage display blocks tigit/pvr for cancer immunotherapy. *Angewandte Chemie International Edition*, 59(35):15114–15118, 2020.

[69] Alex J Callahan, Satish Gandhesiri, Tara L Travaline, Rahi M Reja, Lia Lozano Salazar, Stephanie Hanna, Yen-Chun Lee, Kunhua Li, Olena S Tokareva, Jean-Marie Swiecicki, et al. Mirror-image ligand discovery enabled by single-shot fast-flow synthesis of d-proteins. *Nature communications*, 15(1):1813, 2024.

[70] Yun-Kun Qi, Ji-Shen Zheng, and Lei Liu. Mirror-image protein and peptide drug discovery through mirror-image phage display. *Chem*, 2024.

[71] Rebecca F Alford et al. The rosetta all-atom energy function for macromolecular modeling and design. *Journal of chemical theory and computation*, 13(6):3031–3048, 2017.

[72] Yang Zhang and Jeffrey Skolnick. Scoring function for automated assessment of protein structure template quality. *Proteins: Structure, Function, and Bioinformatics*, 57(4):702–710, 2004.

A Related Works

A.1 Diffusion and Flow for Protein Design

Generative models have displayed incredible promise in designing novel protein structures for custom functions [47–51], with much of this progress driven by advancements in diffusion models [52, 53]. Among distinct protein representations, the frame representation has achieved state-of-the-art performance in tasks like *de novo* backbone design and motif scaffolds [9], as exemplified by RFdiffusion [35]. Recently, flow methods [54] appear as an alternative to diffusion models, offering

a deterministic approach by eliminating stochasticity from the sampling process. They have been extended to Riemannian manifolds [55] and possess an attractive connection to optimal transport (OT). The linear interpolation schedule of flow results in more direct sampling trajectories, reducing the number of integration steps needed [7]. In the computer vision field, flow matching (FM) has already demonstrated comparable results to diffusion models with more computational efficiency [56].

The past few years have witnessed a growth in applying diffusion and flow models [13, 27, 57] for antibody design, a special protein family in the immune system to resist antigens, mainly focusing on inpainting complementarity-determining regions (CDRs) at the interface between the antigen and the framework. Lately, non-trivial efforts have been put into developing DL algorithms for peptides. Li et al. [5] encompasses the dynamics of side-chain angles. Contrarily, Lin et al. [4] relies on RDE [58] to complete the solution to full-atom design. Kong et al. [59] presents a full-atom diffusion on the latent geometric space learned by VAE. Nevertheless, none considers the extension of sequence models (*e.g.*, PLMs) and multi-stage training paradigms to alleviate the shortage of available binding structures.

A.2 D-peptide Technology

Proteins are built from chiral molecules, specifically amino acids, that exist in two mirror-image forms: L (Levorotatory) and D (Dextrorotatory) [60]. Natural ribosomes exclusively use L-amino acids to synthesize proteins, resulting in L-framework proteins throughout nature. While engineered ribosomes [61], post-translational modification systems [62], and non-ribosomal peptide synthetases [63] can incorporate some D-amino acids into L-protein chains, proteins made entirely of D-amino acids have never been found in nature [16] and must be chemically synthesized in the laboratory. Despite the synthetic challenges, D-proteins are valuable research tools. They fold into mirror images of their L-counterparts and offer unique opportunities for studying fundamental protein mechanisms, developing stable molecular binding agents, and even exploring the possibility of mirror-image biological systems [64, 65]. To identify D-proteins capable of binding to a target L-protein, mirror-image peptide phage display methods have been developed [66–70], which involves screening a phage library of L-peptides with a target D-protein. However, it remains challenging to precisely target a specific surface region of the target protein and confirm the presence of valid binders within the initial random library [17]. More importantly, no preceding works have explored the possibility of DL co-design methods [11, 27] for D-protein generation, not to mention the incorporation and verification of this mirror-image technique.

B Preliminary and Background

Proteins. A protein is a biomolecule consisting of n amino acid residues, each defined by its type, backbone frame, and side-chain torsion angles. The type of the i-th residue, denoted by $a^i \in \mathcal{A}$, is determined by its side-chain R group, where $\mathcal{A} = \{1 \dots 20\}$. The rigid frame of each residue is constructed from the coordinates of four backbone heavy atoms N, Cα, C, and O, with Cα positioned at the origin. This frame is represented by a position vector $\boldsymbol{x}^i \in \mathbb{R}^3$ and a rotation matrix $O^i \in \mathrm{SO}(3)$ [24]. Additionally, seven torsion angles $\boldsymbol{\chi}^i \in [0, 2\pi)^7$ are considered. Three of them belong to the backbone, where $\boldsymbol{\chi}^i[0]$ is the angle around $\mathrm{C} - \mathrm{N}$ bond, $\boldsymbol{\chi}^i[1]$ is the angle around $\mathrm{N} - \mathrm{C}\alpha$ bond, and $\boldsymbol{\chi}^i[2]$ affects the position of the oxygen atom. Unlike the backbone, the side-chain conformation is more flexible, involving up to four rotatable torsion angles between side-chain atoms, denoted by $\boldsymbol{\chi}^i[3:] \in [0, 2\pi)^4$.

Flow on Riemannian Manifolds. Flow matching (FM) [7] is a simulation-free method for learning a continuous normalizing flow (CNF). On a manifold $\mathcal{M}$ with a Riemannian metric g, the CNF $\Phi : \mathcal{M} \to \mathcal{M}$ is defined by a one-parameter diffeomorphism that integrates along a time-dependent vector field $u_t \in \mathcal{U}$, where $u_t(x) \in \mathcal{T}_x\mathcal{M}$ falls at the tangent space of the manifold at $x \in \mathcal{M}$.

With an initial condition of $\phi_0(x) = x$, the time-dependent flow $\phi_t : \mathcal{M} \to \mathcal{M}$ and the final diffeomorphismis are attained by setting $\Phi(x) = \phi_1(x)$ and solving the ordinary differential equation (ODE) $\frac{\mathrm{d}\phi_t}{\mathrm{d}t}(x) = u_t(\phi_t(x))$ on $\mathcal{M}$ over $t \in [0, 1]$. Remarkably, $\phi_t(x)$ transports the point x along the vector field $u_t(x)$ from time 0 up to time t, satisyfing another ODE: $\mathrm{d}x = u_t(x)\mathrm{d}t$. Denote $\mathbb{P}(\mathcal{M})$ as the space of probability distributions on $\mathcal{M}$. Then the flow can transform a simple prior density $p_0 \in \mathbb{P}(\mathcal{M})$ towards the data distribution $p_1 \in \mathbb{P}(\mathcal{M})$ via the push-forward equation $p_t = [\phi_t]_{\#}p_0$,

and its density is $p_t(x) = [\phi_t]_{\#} p_0(x) = p_0\left(\phi_t^{-1}(x)\right) \det\left[\frac{\partial \phi_t^{-1}}{\partial x}(x)\right]$. This sequence of probability distributions $\{p_t : t \in [0,1]\}$ is referred to the probability path, which interpolates $\mathbb{P}(\mathcal{M})$ between p_0 and p_1.

The density p_t is characterized by the Fokker-Planck equation: $\frac{\mathrm{d}p_t}{\mathrm{d}t} = -\mathrm{div}(u_t p_t)$, also known as the continuity equation. Under these conditions, u_t is said to be the probability flow for p_t, and p_t is said to be the marginal probability path generated by u_t. Although u_t is intractable in general, it can be learned efficiently by decomposing the target probability path p_t as a mixture of tractable conditional probability paths $p_t(x|x_1)$, which satisfy $p_0(x|x_1) = p_0(x)$ and $p_1(x|x_1) = \delta_{x_1}(\cdot)$. Consequently, the desired unconditional p_t is an average of the conditional probability paths with respect to the data distribution: $p(x) = \int p_t(x|x_1)p(x_1)dx_1$.

The key insight of conditional Riemannian flow matching (CRFM) is to fit a vector field $v_\theta \in \mathcal{U}$ with parameters θ to the conditional vector field $u_t(x|x_1)$, which produces $p_t(x|x_1)$. Its objective falls at the tangent space $\mathcal{T}_x\mathcal{M}$ as:

$$\mathcal{L}_{\text{CRFM}}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],p_1(x_1),p_t(x|x_1)}\|v_\theta(x,t) - u_t(x|x_1)\|_g^2, \tag{17}$$

This loss can be parameterized by defining the conditional flow $x_t = \phi_t(x_0|x_1)$, where ϕ_t is the solution to $\frac{\mathrm{d}\phi_t}{\mathrm{d}t}(x) = u_t(\phi_t(x_0|x_1)|x_1)$ with an initial condition of $\phi_0(x_0|x_1) = x_0$. Finally, the loss can be written as $\mathcal{L}_{\text{CRFM}}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],x_1\sim p_1(x_1),x_0\sim p_0(x_0)}\|v_\theta(x,t) - \dot{x}_t\|_g^2$. During the inference period, samples can be generated by solving the ODE related to the neural vector field v_θ to push $x_0 \in \mathcal{M}$ from the source density p_0 to the data distribution p_1 in time efficiently.

C Proof of Theorem

Applying a rigid transformation (R, t) to the receptor before mirroring gives:

$$\Psi\left(RC_{\text{rec}} + t\right) = M\left(RC_{\text{rec}} + t\right) = MRC_{\text{rec}} + Mt$$

Since M is a reflection, it interacts with rotation matrices. If R is an arbitrary rotation matrix, then MR is generally not equal to RM, meaning that mirroring and rotation do not necessarily commute.

Applying the flow model to the mirrored receptor:

$$f_\theta\left(\Psi\left(RC_{\text{rec}} + t\right)\right) = f_\theta\left(MRC_{\text{rec}} + Mt\right)$$

Since the flow model f_θ is trained on natural L-proteins and learns roto-translational equivariant mappings, it satisfies:

$$f_\theta\left(R'C\right) = R' f_\theta(C), \quad \forall R' \in SO(3), t' \in \mathbb{R}^3$$

Thus, applying f_θ to the mirrored input:

$$f_\theta\left(MRC_{\text{rec}} + Mt\right) = MRf_\theta\left(C_{\text{rec}}\right) + Mt.$$

Now, applying Ψ^{-1} to the generated peptide:

$$\Psi^{-1}\left(f_\theta\left(\Psi\left(RC_{\text{rec}} + t\right)\right)\right) = M^{-1}\left(MRf_\theta\left(C_{\text{rec}}\right) + Mt\right)$$

Since $M^{-1} = M$ (because mirroring is an involution),

$$M\left(MRf_\theta\left(C_{\text{rec}}\right) + Mt\right) = Rf_\theta\left(C_{\text{rec}}\right) + t$$

We have shown that:

$$\Psi^{-1}\left(f_\theta\left(\Psi\left(RC_{\text{rec}} + t\right)\right)\right) = R\Psi^{-1}\left(f_\theta\left(\Psi\left(C_{\text{rec}}\right)\right)\right) + t$$

Thus, the mirror-image algorithm is roto-translationally equivariant.

D Metrics.

Generated peptides are evaluated across three key aspects. (1) *Geometry*: Designed peptides should closely mirror native sequences and structures. We quantify sequence similarity through the amino acid recovery rate (**AAR**), which measures sequence identity between generated and ground truth

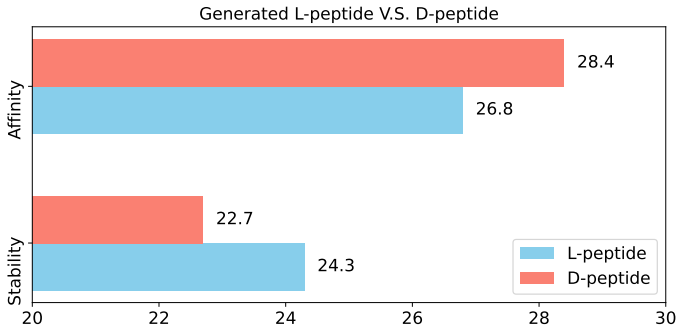

Figure 4: Energy-based comparison between generated L- and D-peptides, simulated by Rosetta. *Stability* represents the fraction of generated complexes exhibiting lower total energy than the native states, while *Affinity* indicates the percentage of generated peptides achieving higher binding affinity than the native ones.

peptides. Structural similarity is evaluated using the root-mean-square deviation (**RMSD**) of C_α atoms after complex alignment. The secondary-structure similarity ratio (**SSR**) captures the proportion of matching secondary structures, while the binding site ratio (**BSR**) evaluates the overlap between generated and native peptide binding sites on the target protein. (2) *Energy*: We aim to design high-affinity peptide binders that enhance protein-peptide complex stability. **Affinity** represents the percentage of generated peptides achieving higher binding affinities (lower binding energies) than the native peptide. **Stability** indicates the fraction of complexes exhibiting lower total energy than the native state. All energy calculations are performed using Rosetta [71]. (3) *Design*: **Designability** measures the consistency between designed sequences and structures, calculated as the proportion of sequences that fold into structures similar to their generated conformations (C_α RMSD < 2 Å). Sequence refolding is performed using ESMFold [26]. **Diversity**, computed as the average of one minus the pairwise TM-Score [72], quantifies structural variation among the designed peptides.

E Quantitative Analysis.

Though the ground truth D-peptides that bind to receptors are unavailable, we conduct a simulation via Rosetta to quantitatively analyze the stability and binding affinity of the system for the generated D-peptides. As the conformation distribution is divergent, we generate 1K samples for each receptor in the test set and select the most stable one for evaluation. These filtered D-peptides exhibit competitive energy metrics compared to their L-peptide counterparts (see Fig. 4). Specifically, they achieve an average improvement in binding affinity of **28. 4%** over native L-peptides, while maintaining a high stability improvement of 22.7%, significantly outperforming the state-of-the-art generated L-peptides (24.31%). These results suggest that D-Flow's ability to effectively leverage receptor-specific chirality information enhances its ability to optimize peptide-protein interactions.