

000 001 CHIFLOW: TORSIONAL ASYMMETRY FLOW MATCH- 002 ING FOR CHIRALITY-AWARE PROTEIN BACKBONE 003 GENERATION 004 005

006 **Anonymous authors**
007 Paper under double-blind review
008
009
010
011

ABSTRACT

013 Protein backbone generation is critical for de novo protein design, yet existing
014 methods suffer from two key limitations: over-reliance on $SE(3)$ modeling, which
015 introduces unnecessary complexity for cyclic dihedral angles, and lack of explicit
016 chirality control, leading to nonfunctional D-chiral outputs. We present **ChiFlow**,
017 a chirality-aware backbone generator based on flow matching on **toroidal Rie-**
018 **mannian manifolds**. ChiFlow models backbone dihedrals ϕ, ψ, ω as points on
019 \mathbb{T}^3 , extending PPFlow to backbone variables and using periodicity to avoid bound-
020 ary artifacts. Unlike the previous $SE(3)$ -based flows such as Frameflow and Fold-
021 flow2, ChiFlow operates directly on the hypertorus, simplifying computations for
022 angles. We also introduce a Riemannian mirroring operator and impose asymme-
023 try on the learned vector field to enforce L-chirality. And we extended the meth-
024 ods in Foldingdiff by reconstructing the 3D atomic coordinates using fixed bond
025 lengths and trigonometric calculations. To increase the diversity that was low-
026 ered by the chirality constraint, we added Stochastic Flow Matching to ChiFlow,
027 resulting in an increase in diversity of the generated backbone. With extensive
028 experiments on real-world protein datasets, ChiFlow approaches the leading flow
029 models in the benchmark while maintaining absolute chirality purity. Our imple-
030 ment detail is at <https://anonymous.4open.science/r/anonym1>.
031

1 INTRODUCTION

032 Proteins play a pivotal role in sustaining life, as their structural diversity underlies a wide spectrum
033 of biological functions. The ability to design proteins with predetermined architectures and activities
034 represents a major advance in biotechnology, providing systematic approaches that extend beyond
035 natural evolutionary processes. Progresses made by methods like AlphaFold3 (Abramson et al.,
036 2024) and ESM3 (Hayes et al., 2024) carries profound significance for medicine and public health,
037 offering the potential to generate novel therapeutic strategies and address long-standing challenges
038 in disease prevention and treatment. At the core of protein design lies the generation of accurate
039 backbone structures (Tang et al., 2024a), since the backbone determines the overall fold and strongly
040 constrains the positioning of side chains. Figure 1 shows the comparasion between the whole protein
041 and the protein backbone.
042

043 **Limitations of Existing Approaches:** Current $SE(3)$ -based frameworks (Fuchs et al., 2020) repre-
044 sent protein backbones as sequences of $SE(3)$ rigid frames (Yim et al., 2023a;b; 2024; Bose et al.,
045 2024; Huguet et al., 2024), which are effective for capturing global structure but impose unnecessary
046 complexity when modeling cyclic dihedral angles ϕ, ψ, ω —the primary determinants of backbone
047 conformation—by embedding them into $SE(3)$ ’s rotation component, thereby creating periodicity
048 mismatches and boundary artifacts (Zhang et al., 2024). Moreover, natural proteins are exclusively
049 L-chiral, yet existing approaches (Yim et al., 2023a;b; 2024; Bose et al., 2024; Fu et al., 2023; Song
050 et al., 2023) lack explicit chirality enforcement, and they mostly base on $SE(3)$ which can not model
051 chirality (Dumitrescu et al., 2025b; Childs et al., 2025), which risks generating non-physical right-
052 handed backbones (Childs et al., 2025). While torus geometry has been shown to effectively capture
053 periodicity for sidechain torsion angles (Huguet et al., 2024), it has not been extended to backbone
dihedrals or coupled with chirality constraints. Therefore, there remains a critical gap for a backbone
generation framework that jointly leverages torus geometry and chirality-aware modeling.

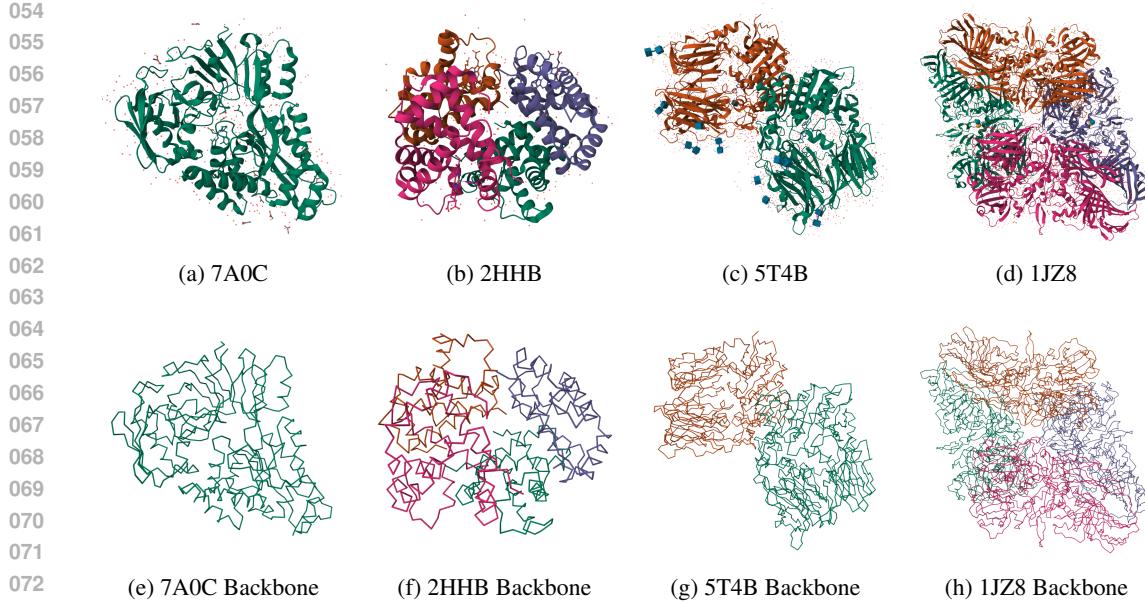


Figure 1: The comparasion between the original protein and the protein backbone

Proposed Method: To address these limitations, we introduce **ChiFlow**, a chirality-aware flow-matching framework that operates directly on toroidal Riemannian manifolds (Chen & Lipman, 2024) rather than $SE(3)$ flow (Bose et al., 2024). Unlike $SE(3)$ FrameFlow (Yim et al., 2023a), ChiFlow removes the burden of embedding periodic angles into rigid-body rotations; and unlike sidechain-focused torus approach (Lin et al., 2024a), it directly targets the backbone, the core structural determinant of protein folds. Through this integration, we achieve both competitive designability and sampling efficiency, while uniquely delivering **one hundred percent chirality purity**, bridging a fundamental gap in backbone generation.

Contributions: ChiFlow introduces three key advances in protein backbone generation: (1) **Toroidal Riemannian Manifold** framework that models backbone dihedrals natively on \mathbb{T}^{3N} , eliminating periodicity artifacts inherent in $SE(3)$ -based methods; (2) **Chirality-Aware Flow Matching** approach via a Riemannian mirroring operator and asymmetric vector field constraint, guaranteeing absolute L-chirality without sacrificing designability; and (3) **Stochastic Flow Matching** on the torus that enhances conformational diversity while preserving geometric constraints.

2 BACKGROUND & PRELIMINARIES

2.1 PROTEIN BACKBONE DIHEDRALS & CHIRALITY

A protein backbone consists of N residues linked by peptide bonds (Abdin & Kim, 2023), with each residue's conformation defined by three dihedral angles : ϕ around the $N-C\alpha$ bond, ψ around the $C\alpha-C$ bond, and ω around the $C-N$ peptide bond, which is nearly fixed at 0 in the trans state or π in the cis state but remains cyclic. These angles determine secondary structures such as α -helices, with $\phi \approx -57^\circ$ and $\psi \approx -47^\circ$, and β -sheets, with $\phi \approx -130^\circ$ and $\psi \approx 120^\circ$ (Ramachandran et al., 1963). And the dihedral angles will determine the chiral of the protein backbone, note that natural proteins are all inherently L-chiral (Cintas, 2002). A protein backbone is L-chiral if its dihedral angles, when mapped to 3D $C\alpha$ coordinates, cannot be superimposed on their mirror image, that is, the D-chiral configuration, via translations alone. To model this attribute, let $X = \{x_1, \dots, x_N\} \in \mathbb{R}^{3N}$ be $C\alpha$ coordinates derived from dihedrals $\chi \in \mathbb{T}^{3N}$, and

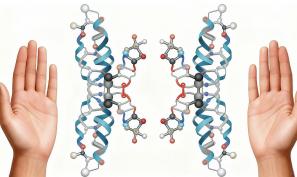


Figure 2: Chirality

108 $M : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be a mirroring operator, for example $M(x) = (-x_1, x_2, x_3)$. X is L-chiral if
 109 $\nexists b \in \mathbb{R}^3$ such that $M(X) + b = X$.
 110

111 2.2 TOROIDAL RIEMANNIAN MANIFOLDS 112

113 The 1-dimensional torus \mathbb{T} can be defined as the quotient space $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$ (Lipman et al., 2023),
 114 where two real numbers a and b are equivalent if $a - b = 2\pi k$ for some $k \in \mathbb{Z}$, making it a
 115 natural model for cyclic variables (Lin et al., 2024b) like dihedral angles. For protein backbone
 116 dihedrals, we work on the hypertorus $\mathbb{T}^3 = \mathbb{T} \times \mathbb{T} \times \mathbb{T}$, where each point $\chi = (\chi_\phi, \chi_\psi, \chi_\omega) \in \mathbb{T}^3$
 117 represents the three dihedral angles of a residue. At a point χ , the tangent space $T_\chi \mathbb{T}^3$ is endowed
 118 with a conformal product metric $g_{\mathbb{T}^3}|_\chi(u, v) = \sum_{k \in \{\phi, \psi, \omega\}} w_k(\chi_k) u_k v_k$, where the smooth, 2π -
 119 periodic weights are given by $w_k(\chi_k) = \frac{1}{1 + \sin^2(\chi_k/2)}$. Extending this to N residues that make up
 120 the whole protein backbone yields the product metric $g = \bigsqcup_{n=1}^N g_{\mathbb{T}^3}$, which induces the pointwise
 121 norm $\|u\|_g^2 = g(u, u)$ and the volume element $d\text{vol}_g = \sqrt{\det g(\chi)} d\chi$. Under this metric, the
 122 Riemannian gradient, divergence, and Laplace–Beltrami operators are respectively $\nabla_g f = g^{-1} \cdot \partial f$,
 123 $\text{div}_g X = \frac{1}{\sqrt{\det g}} \partial_i (\sqrt{\det g} X^i)$, and $\Delta_g f = \text{div}_g (\nabla_g f)$, all computed componentwise due to
 124 the product structure. For interpolation, we follow (Zhang et al., 2024) and define the geodesic
 125 $\chi_t : [0, 1] \rightarrow \mathbb{T}^3$ between χ_0 and χ_1 as $\chi_t = \exp_{\chi_0}(t \log_{\chi_0}(\chi_1))$ Mathieu & Nickel (2020), which
 126 in the flat product metric reduces to shortest-path interpolation on each circle, meaning that each
 127 angular coordinate evolves linearly along the minimal arc between its start and end values, with
 128 wrapping modulo 2π to preserve periodicity.
 129

130 2.3 FLOW MATCHING ON RIEMANNIAN MANIFOLDS 131

132 Riemannian Flow Matching (RFM) provides a simulation-free
 133 framework for learning continuous normalizing flows on general
 134 Riemannian manifolds. Given a complete, connected smooth Rie-
 135 mannnian manifold \mathcal{M} with metric g , the goal is to learn a time-
 136 dependent vector field $v_t(x)$ that transports an initial distribution
 137 p_0 to a target distribution p_1 by satisfying the continuity equation
 138 $\partial_t p_t + \text{div}_g(p_t v_t) = 0$ with boundary conditions $p_{t=0} = p_0$ and
 139 $p_{t=1} = p_1$ Chen & Lipman (2024).

140 The key idea is to construct conditional probability paths $p_t(x|x_1)$
 141 that interpolate between p_0 and Dirac distributions δ_{x_1} centered at
 142 data points $x_1 \sim p_1$. The marginal probability path is then defined
 143 as:

$$144 \quad 145 \quad p_t(x) = \int_{\mathcal{M}} p_t(x|x_1) p_1(x_1) d\text{vol}_{x_1}. \quad (1)$$

146 and the target vector field is obtained by marginalizing conditional vector fields:
 147

$$148 \quad 149 \quad u_t(x) = \int_{\mathcal{M}} u_t(x|x_1) \frac{p_t(x|x_1) p_1(x_1)}{p_t(x)} d\text{vol}_{x_1}. \quad (2)$$

150 where $u_t(x|x_1)$ generates $p_t(x|x_1)$ from p_0 .
 151

152 For simple geometries where closed-form geodesics are available, the conditional flow can be de-
 153 fined using the geodesic distance d_g as a premetric. With the linear scheduler $\kappa(t) = 1 - t$, the
 154 conditional flow is:

$$155 \quad x_t = \psi_t(x_0|x_1) = \exp_{x_0}(t \log_{x_0}(x_1)). \quad (3)$$

156 which yields the conditional vector field:
 157

$$158 \quad 159 \quad u_t(x_t|x_1) = \frac{d}{dt} x_t = \frac{\log_{x_t}(x_1)}{1 - t}. \quad (4)$$

160 This ensures that the interpolation path follows the geodesic of the manifold, respecting its intrinsic
 161 geometry (Bose et al., 2024; Chen & Lipman, 2024).

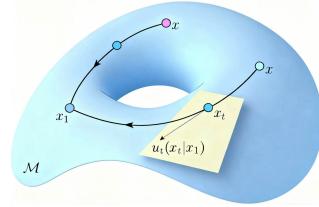


Figure 3: Torus Manifold
 (Chen & Lipman, 2024)

162 The parametric vector field v_θ is learned by minimizing the Riemannian Conditional Flow Matching
 163 (RCFM) objective:

$$165 \quad \mathcal{L}_{\text{RCFM}}(\theta) = \mathbb{E}_{t \sim \rho(t), x_1 \sim p_1, x_0 \sim p_0} \left[\|v_\theta(t, x_t) - u_t(x_t | x_1)\|_g^2 \right]. \quad (5)$$

166 where $\rho(t)$ is a time weighting distribution, often chosen as $\rho(t) \propto (1-t)^\alpha$ for $\alpha \geq 0$, and
 167 expectations are taken over the conditional coupling induced by the geodesic interpolation (Chen &
 168 Lipman, 2024). This objective is equivalent to the marginal Riemannian Flow Matching objective
 169 up to an additive constant, ensuring that the learned vector field generates flows that satisfy the
 170 continuity equation in expectation.
 171

172 RFM is simulation-free on simple geometries and does not require divergence computation during
 173 training, and provides exact target vector fields without approximation errors (Chen & Lipman,
 174 2024). For general manifolds, spectral distances can be used as premetrics to maintain tractability,
 175 ensuring broad applicability across diverse geometries.

176 2.4 MIRRORING OPERATOR FOR CHIRALITY

178 To enforce strict L-chirality in generated backbones, we introduce a Riemannian mirroring operator
 179 $M : \mathbb{T}^{3N} \rightarrow \mathbb{T}^{3N}$ that acts residue-wise on the dihedral angles. For each residue, the operator is
 180 defined as:

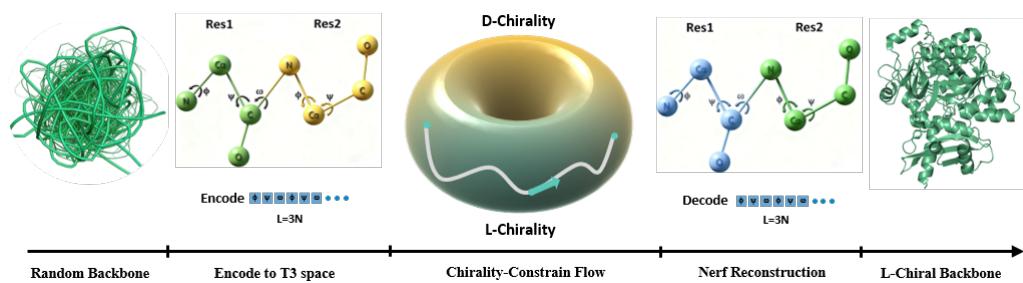
$$181 \quad M(\phi_n, \psi_n, \omega_n) = (2\pi - \phi_n \bmod 2\pi, 2\pi - \psi_n \bmod 2\pi, \omega_n). \quad (6)$$

182 This transformation corresponds to the molecular mirroring operation that converts an L-chiral back-
 183 bone to its D-chiral enantiomer while preserving the cyclic nature of the dihedral angles. Crucially,
 184 M is an isometric diffeomorphism on \mathbb{T}^{3N} equipped with the product metric g . The pushforward
 185 $M_* : T_\chi \mathbb{T}^{3N} \rightarrow T_{M(\chi)} \mathbb{T}^{3N}$ has a constant Jacobian $J_M = \text{diag}(-1, -1, 1)$ per residue, satisfying
 186 $M_*^2 = I$. The conformal weights $w_k(\chi_k) = \frac{1}{1 + \sin^2(\chi_k/2)}$ satisfy $w_k(\chi_k) = w_k(2\pi - \chi_k)$ for
 187 $k \in \{\phi, \psi\}$, and $w_\omega(\chi_\omega) = w_\omega(\chi_\omega)$, ensuring that M preserves the metric:
 188

$$189 \quad g|_{M(\chi)}(M_* u, M_* v) = g|_\chi(u, v). \quad (7)$$

190 Through this, we explicitly encodes chirality awareness into the generative process, ensuring that all
 191 sampled backbones maintain correct L-chirality while preserving the natural periodicity of dihedral
 192 angles.

194 3 METHODOLOGY: CHIFLOW



207 Figure 4: ChiFlow represents protein backbone dihedrals (ϕ, ψ, ω) on a toroidal manifold T^3 to
 208 inherently preserve angular periodicity. And employs Riemannian flow matching to learn a vector field
 209 that transports noise distributions to data distributions on this manifold. To ensure strict L-chirality,
 210 a mirroring operator M is introduced, constraining the vector field to be asymmetric under M (i.e.,
 211 $\hat{v}(M\chi, t) = -M_*\hat{v}(\chi, t)$). Finally, 3D atomic coordinates are reconstructed deterministically from
 212 dihedrals using fixed bond lengths and trigonometric operations via the SN-NeRF algorithm.

213 **ChiFlow** is a generative model based on flow-matching model for protein backbone generation, ad-
 214 dressing critical flaws of $\text{SE}(3)$ -based methods by grounding dihedral angle (ϕ, ψ, ω) modeling in
 215 a toroidal manifold \mathbb{T}^{3N} (one \mathbb{T}^3 per residue). Geodesic interpolation on \mathbb{T}^{3N} enables continuous,

shortest-path transitions between conformations, and Riemannian operators $(\nabla_g, \text{div}_g, \Delta_g)$ ensure consistent calculus on the manifold. ChiFlow first uses deterministic Riemannian flow matching to learn a vector field $\hat{v}(\chi, t)$ that transports a prior distribution p_0 to a native-like target p_1 , minimizing a L^2 loss $\mathcal{L}_{\mathbb{T}^3}$. To enhance conformational diversity and model robustness, we extended it to Stochastic Flow Matching(SFM): a torus Brownian motion (with generator $\sigma(t)^2/2\Delta_g$, where $\sigma(t)$ is a decreasing noise schedule) defines a reference stochastic process, with noise projected onto the manifold’s tangent spaces to preserve geometric validity. Framed as a Schrödinger bridge (Tang et al., 2024b), SFM learns the optimal vector field v_θ via a flow matching loss \mathcal{L}_{SFM} , avoiding the computational cost of simulating the entire stochastic process. Critically, to guarantee biologically essential L-chirality, a **chiral asymmetry loss** $\mathcal{L}_{\text{chiral}}$ enforces $v_\theta(M(\chi_t), t) = -M_*v_\theta(\chi_t, t)$, where M is a chirality-flipping mirroring operator. To reconstruct the protein backbone using the dihedral angles generated by the flow model, we employed the SN-NeRF algorithm (Parsons et al., 2005), paired with forward kinematics using standard protein backbone bond lengths ($\text{C}\alpha\text{-N} = 1.45$ Å, $\text{C}\alpha\text{-C} = 1.54$ Å) for efficiency and stability.

3.1 BACKBONE MODELING AND TOROIDAL FLOW MATCHING

Proteins fold into specific three-dimensional structures to perform their biological functions (Pollock, 2007). The backbone dihedral angles—phi (ϕ), psi (ψ), and omega (ω)—are crucial degrees of freedom that define the protein’s conformational space. Unlike linear Euclidean spaces, these angles are cyclic by nature, with values wrapping around every 2π radians (Atavin & Vilkov, 2003). This periodicity must be respected in any meaningful mathematical model of protein structure and dynamics. To address this, we turn to the concept of the Torus (Wiemeler, 2015). We represent the full protein backbone as a point on a hypertorus, $\chi = \{\chi_1, \dots, \chi_N\} \in \mathbb{T}^{3N}$, where each residue n has its own set of dihedral angles $\chi_n = (\phi_n, \psi_n, \omega_n)$. The 1-dimensional torus \mathbb{T} is defined as the quotient space $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$, which identifies points a and b whenever $a - b = 2\pi k$ for any integer k . This identification effectively gather together the ends of the interval $[0, 2\pi]$, forming a circle and eliminating artificial boundaries at 0 and 2π . The 3-dimensional product torus $\mathbb{T}^3 = \mathbb{T} \times \mathbb{T} \times \mathbb{T}$ thus naturally encodes the three dihedral angles of a single residue without introducing discontinuities. The space \mathbb{T}^{3N} is not just a set of points, it is a Riemannian manifold. At a point $\chi \in \mathbb{T}^3$, the tangent space $T_\chi \mathbb{T}^3$ is a vector space containing all possible directions of change for the dihedral angles. We equip this space with a conformal product metric:

$$g_{\mathbb{T}^3}|_\chi(u, v) = \sum_{k \in \{\phi, \psi, \omega\}} w_k(\chi_k) u_k v_k, \quad \text{where} \quad w_k(\chi_k) = \frac{1}{1 + \sin^2(\chi_k/2)}. \quad (8)$$

This metric is smooth and 2π -periodic in each coordinate, meaning it respects the circular nature of the angles. The weight function $w_k(\chi_k)$ is chosen to reflect known physical constraints; for instance, it assigns higher cost to changes in angles when they are near regions of steric clash (Ramachandran et al., 2011). Extending this metric to N residues yields the product metric $g = \bigsqcup_{n=1}^N g_{\mathbb{T}^3}$ for the entire backbone. This metric induces a norm $\|u\|_g^2 = g(u, u)$ on tangent vectors and a volume element $d\text{vol}_g = \sqrt{\det g(\chi)} d\chi$ for integration on the manifold. On a Riemannian manifold, the familiar calculus operations from Euclidean space generalize. The Riemannian gradient, divergence, and Laplace–Beltrami (Zhou & Lähner, 2025) operators become:

$$\nabla_g f = g^{-1} \cdot \partial f, \quad \text{div}_g X = \frac{1}{\sqrt{\det g}} \partial_i \left(\sqrt{\det g} X^i \right), \quad \Delta_g f = \text{div}_g (\nabla_g f). \quad (9)$$

Due to the product structure of our torus, these operators can be computed componentwise, which significantly simplifies calculations. A key geometric concept is the *geodesic*—the shortest path between two points on the manifold. For interpolation between two backbone configurations χ_0 and χ_1 on \mathbb{T}^{3N} , we follow the shortest-path geodesic (Wu et al., 2025) on each individual angular circle:

$$\chi_t = \exp_{\chi_0} (t \log_{\chi_0} (\chi_1)). \quad (10)$$

This operation reduces to linear interpolation along the minimal arc in each angular coordinate, with values wrapped modulo 2π . This provides a natural and continuous path between the desired protein backbone and the origin protein backbone. *Flow matching* is a powerful framework for learning probability paths between distributions on manifolds. On \mathbb{T}^{3N} , we seek a time-dependent vector field $\hat{v}(\chi, t)$ that transports an initial probability distribution p_0 (e.g., a simple prior distribution over

structures) to a target distribution p_1 (e.g., a distribution of native-like structures) by satisfying the continuity equation:

$$\partial_t p_t + \operatorname{div}_g(p_t \hat{v}(\cdot, t)) = 0. \quad (11)$$

with boundary conditions $p_{t=0} = p_0$ and $p_{t=1} = p_1$. Using the geodesic interpolation $(\chi_t)_{t \in [0,1]}$ defined between samples $\chi_0 \sim p_0$ and $\chi_1 \sim p_1$, the exact conditional vector field that transports χ_0 to χ_1 along the geodesic is known in closed form: $u(\chi_t, t \mid \chi_1) = \frac{\log_{\chi_t}(\chi_1)}{1-t}$. We estimate the vector field \hat{v} by minimizing a time-weighted Riemannian L^2 mismatch loss:

$$\mathcal{L}_{\mathbb{T}^3} = \mathbb{E}_{t \sim \rho} \mathbb{E}_{p(\chi_0), p(\chi_1)} [\|u(\chi_t, t \mid \chi_1) - \hat{v}(\chi_t, t)\|_g^2]. \quad (12)$$

where $\rho(t) \propto (1-t)^\alpha$ is a weighting function that often puts more emphasis on later times (e.g., $\alpha = 2$), and the inner expectation is over the conditional coupling induced by the geodesic interpolation. This geometric formulation provides a rigorous foundation for modeling protein dynamics, folding pathways (Englander & Mayne, 2014), and for generating novel structures, as it inherently avoids the coordinate singularities and artifacts that plague methods treating dihedral angles as linear Euclidean variables.

3.2 STOCHASTIC FLOW MATCHING AND CHIRALITY CONTROL

To address the limitations of deterministic flow and significantly improve the sample diversity and robustness of the ChiFlow model, we extend its framework by incorporating a **Stochastic Flow Matching (SFM)** (Albergo et al., 2023) formulation on the toroidal manifold \mathbb{T}^{3N} . This extension transforms the deterministic flow into a stochastic bridge, leveraging controlled noise and optimal transport theory.

3.2.1 STOCHASTIC DYNAMICS AND THE SCHRÖDINGER BRIDGE

The core of the SFM involves defining a reference stochastic process on \mathbb{T}^{3N} . Let Δ_g be the Laplace–Beltrami operator under the product metric. We consider a torus Brownian motion with generator $\frac{\sigma(t)^2}{2} \Delta_g$ as the reference process, where $\sigma(t)$ is a carefully chosen decreasing noise schedule. The dynamics of the system are then governed by the following Stratonovich stochastic differential equation (SDE) (Anderson, 1982):

$$d\chi_t = v_\theta(\chi_t, t) dt + \sigma(t) \Pi_{T_{\chi_t}} \circ dW_t. \quad (13)$$

Here, $\Pi_{T_{\chi_t}}$ projects the standard Wiener noise dW_t onto the tangent space of the manifold at χ_t , ensuring that the added noise respects the geometric constraints of \mathbb{T}^{3N} . The vector field v_θ , parameterized by a neural network, learns to guide this stochastic process. We seek the *most efficient* stochastic path connecting the initial and target distributions relative to this reference Brownian motion. This is formulated as the Schrödinger bridge problem, an entropy-regularized optimal transport problem on the path space:

$$\min_{u_t} \mathbb{E} \int_0^1 \frac{\|u_t(\chi_t)\|_g^2}{2\sigma(t)^2} dt \quad \text{subject to} \quad \partial_t p_t = -\operatorname{div}_g(p_t u_t) + \frac{\sigma(t)^2}{2} \Delta_g p_t. \quad (14)$$

The solution to this problem yields the control u_t that minimizes the expected energy required to drive the system from p_0 to p_1 amidst the stochastic noise.

3.2.2 CHIRALITY PRESERVATION

Directly solving the Schrödinger bridge can be computationally challenging (Jing et al., 2025). Instead, we adopt a **flow matching** objective to efficiently learn the optimal vector field v_θ that approximates the solution. We define a conditional probability path (Mathieu & Nickel, 2020) between samples. The flow matching loss is constructed to regress the learned vector field v_θ towards a target vector field u_σ that defines the desired flow:

$$\mathcal{L}_{\text{SFM}} = \mathbb{E}_{\chi_1 \sim p_1, t \sim U(0,1)} [\|u_\sigma(\chi_t, t \mid \chi_1) - v_\theta(\chi_t, t)\|_{g_{\mathbb{T}^3}}^2]. \quad (15)$$

This objective allows us to train the model without simulating the entire stochastic process, greatly enhancing computational efficiency. A critical concern when introducing stochasticity is the potential violation of fundamental physical symmetries. To enforce the conservation of chirality—a

324 crucial property of biomolecules—we introduce the chirality penalty term into the loss function
 325 (Ciampiconi et al., 2024). Let M denote the operator that flips the chirality of a structure. We
 326 require the learned vector field to exhibit asymmetry under this transformation: $v_\theta(M(\chi_t), t) =$
 327 $-M_*(v_\theta(\chi_t, t))$. This is enforced via the chiral asymmetry loss:

$$\mathcal{L}_{\text{chiral}} = \mathbb{E}_{\chi_t, t} [\|\|v_\theta(M(\chi_t), t) + M_*(v_\theta(\chi_t, t))\|_{g_{\mathbb{T}^3}}^2]. \quad (16)$$

3.2.3 TOTAL OBJECTIVE AND SUMMARY OF IMPROVEMENTS

332 The overall training objective for the Stochastic Flow Matching part of ChiFlow combines the flow
 333 matching loss and the chiral symmetry loss:

$$\mathcal{L}_{\text{total}}^{\text{SFM}} = \mathcal{L}_{\text{SFM}} + \lambda \mathcal{L}_{\text{chiral}}. \quad (17)$$

336 The hyperparameter λ balances the trade-off between sample diversity (achieved through SFM) and
 337 strict geometric consistency (enforced by the chiral loss) and is defined in the config file during training
 338 between 0.1 and 2. The controlled stochastic noise enables the generation of a wider variety of
 339 structurally plausible protein backbone. The SFM extension transforms ChiFlow into a more powerful
 340 and versatile tool for generating diverse, robust, and physically accurate molecular structures on the toroidal manifold.

3.3 $C\alpha$ COORDINATE RECONSTRUCTION

344 We employ forward kinematics with standard protein backbone parameters (e.g., $C\alpha\text{-N} =$
 345 1.45 \AA , $C\alpha\text{-C} = 1.54 \text{ \AA}$) to reconstruct $C\alpha$ coordinates from χ_1 torsional angles. To
 346 maximize computational efficiency and numerical stability, we adopt the Self-Normalizing
 347 Natural Extension Reference Frame (SN-NeRF) algorithm Parsons et al. (2005). This
 348 method reduces the Cartesian coordinate conversion to 66 floating-point operations per
 349 atom, which is fewer than classical rigid-based reconstruction methods (Yim et al., 2023b).
 350 Angle-dependent placement: Compute initial atom position in a local
 351 frame using bond lengths and angles.

$$\vec{D}_2 = (R \cos \theta, R \cos \phi \sin \theta, R \sin \phi \sin \theta) \quad (18)$$

352 where $R = \text{bond}_{CD}$, $\theta = \text{angle}_{BCD}$, $\phi = \text{torsion}_{BC}$.

355 Reference frame alignment: Transform to global coordinates via an orthonormal basis derived from prior atoms:

$$\vec{D} = \hat{M} \vec{D}_2 + \vec{C} \quad (19)$$

$$\hat{M} \equiv [\hat{bc}, \hat{n} \times \hat{bc}, \hat{n}] \quad (20)$$

360 with $\hat{bc} = \overrightarrow{BC} / |\overrightarrow{BC}|$ and \hat{n} as the ABC plane normal.

4 EXPERIMENTS SETUP

4.1 DATASET

366 We train **ChiFlow** on a filtered PDB (Burley et al., 2022) dataset. The raw data was downloaded
 367 from the RCSB (Burley et al., 2022) in mmcif format and we followed method proposed in
 368 FrameDiff (Yim et al., 2023b) to preprocess the dataset to pdb format for visualization and generated
 369 metadata.csv and clusters-by-entity-30.txt to help process the dataset more efficiently. During
 370 training, we use different length of protein backbone such as [60,256] and [100,512], as shown in
 371 Figure 7. Also, pair geometry is encoded via the features of protein backbone heavy atoms. What's
 372 more, during the preprocess process of the datasets, full in-memory caching is disabled by default
 373 to keep preprocessing lightweight.

4.2 METRIC

375 Designability, defined as the fraction of designs with $\text{scRMSD} < 2.0 \text{ \AA}$; Novelty, evaluated by
 376 the maximum TM-score to known PDB structures and the fraction of designs with an averaged

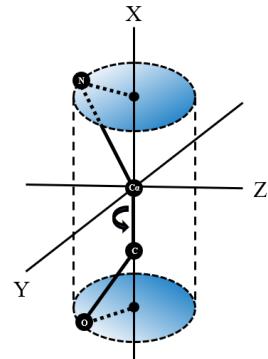


Figure 5: SN-NeRF

378 maximum TM-score < 0.3 and scRMSD $< 2.0 \text{ \AA}$, and Diversity, assessed by the average pairwise
 379 TM-score and the MaxCluster fraction. Standard errors are included for both Designability and
 380 Novelty metrics (Bose et al., 2024). And the Chirality was calculated through the Cahn-Ingold-
 381 Prelog (CIP) rule algorithm from RDKit (Parsons et al., 2005).

382

383

384

4.3 BASELINE

385

386 We compare our results to RFDiffusion (Watson et al., 2022), FrameDiff (Yim et al., 2023b), Frame-
 387 Flow (Yim et al., 2023a) and FoldFlow (Bose et al., 2024) for protein backbone generation. Unlike
 388 RFDiffusion which leverages a pre-trained folding network, both Genie and FrameDiff are diffu-
 389 sion models that generate protein backbones without such dependency. For the flow-based models,
 390 FrameFlow and FoldFlow, we utilize their official implementations and pre-trained weights, these
 391 baselines are selected to illustrate the inherent trade-offs between computational speed and design
 392 quality in state-of-the-art protein structure generation. We can tell from Table 1 that while there
 393 still remains a gap between the golden standard(RFDiffusion) and ChiFlow, ChiFlow approaches
 394 previous methods like FrameFlow, FrameDiff and FoldFlow in benchmark and shows no big dif-
 395 ference. Note that the SFM greatly improved the Diversity of the protein backbone designed by
 396 ChiFlow, even outperform the golden standard. And Figure 7 shows some of the backbone gener-
 397 ated by ChiFlow. The ablation study 2 shows that the Riemannian mirroring operator is crucial to
 398 the control of the Chirality of the protein backbone. And the Riemannian HyperTorus ensures that
 399 the Flow Matching process on dihedral angles is fluent and smooth. Overperform SE(3)-only Flow
 400 Matching method. **Note:** Best results are in bold. Standard errors represent variability across mul-
 401 tiple runs. Chirality purity measures the percentage of generated backbones with correct L-chirality
 402 configuration.

403

404

Table 1: Comprehensive experiments of Baseline over the PDB Datasets Under Various Metrics

Model	Designability		Novelty		Diversity		Chirality
	Fraction \uparrow	scRMSD \downarrow	Fraction \uparrow	avg.max TM \downarrow	pairwise TM \downarrow	MaxCluster \uparrow	Score \uparrow
RFDiffusion	0.969 ± 0.034	0.650 ± 0.136	0.708 ± 0.060	0.449 ± 0.012	0.256 ± 0.010	0.172 ± 0.015	0.94 ± 0.143
FrameDiff	0.414 ± 0.064	3.970 ± 0.436	0.181 ± 0.128	0.556 ± 0.047	0.244 ± 0.011	0.320 ± 0.021	0.88 ± 0.094
FrameFlow	0.798 ± 0.044	—	—	0.698 ± 0.033	—	0.292 ± 0.020	0.93 ± 0.039
FoldFlow	0.671 ± 0.044	3.090 ± 0.285	0.449 ± 0.078	0.469 ± 0.025	0.274 ± 0.014	—	0.95 ± 0.047
ChiFlow	0.674 ± 0.042	2.309 ± 0.271	0.383 ± 0.074	0.502 ± 0.024	0.267 ± 0.013	0.295 ± 0.019	1.00
ChiFlowpro	0.743 ± 0.040	2.192 ± 0.249	0.497 ± 0.068	0.409 ± 0.020	0.197 ± 0.010	0.302 ± 0.020	1.00

411

412

413

414

Table 2: Comprehensive ablation study of ChiFlow components. The table evaluates the contribution
 415 of explicit chirality constraints and Stochastic Flow Matching (SFM) to the overall performance.
 416

Method	Designability		Novelty		Performance	
	Fraction \uparrow	scRMSD (\AA) \downarrow	Fraction \uparrow	Avg. TM \downarrow	Diversity \downarrow	Chirality (%) \uparrow
ChiFlow-pro (Full)	0.743 ± 0.040	2.192 ± 0.249	0.497 ± 0.068	0.409 ± 0.020	0.197 ± 0.010	1.00
SE(3)	0.671 ± 0.044	3.090 ± 0.285	0.449 ± 0.078	0.469 ± 0.025	0.274 ± 0.014	0.95 ± 0.047
ChiFlow(w/o Chirality)	0.735 ± 0.039	2.205 ± 0.251	0.505 ± 0.069	0.425 ± 0.021	0.275 ± 0.010	0.623 ± 0.057

422

423

424

Configuration: The model learns a continuous vector field over the backbone torsion angles
 425 (ϕ, ψ, ω) in the 3-torus space $(-\pi, \pi]^3$, and reconstructs Cartesian coordinates (N, C_α, C) via a
 426 deterministic NeRF decoder. The system was trained for 500 epochs using Adam (lr = 1×10^{-4} ,
 427 β_{default} , batch size 512) with gradient norm clipping at 1.0. Distributed Data Parallel (DDP) (Li
 428 et al., 2020) on two A6000 GPUs was used for training, synchronizing gradients via an All-Reduce
 429 algorithm. Sampling integrates the learned field using 100 steps of explicit Euler reverse diffusion
 430 with torus wrapping, augmented by stochastic paths (SFM) to promote diversity. All runs were con-
 431 figured and versioned via Hydra. And inference is executed on a single RTX 4090 GPU (24 GB
 432 GDDR6X, 330 TFLOPS FP16).

432

5 RELATED WORKS

433
 434
 435 **Flow-based models.** Flow models learn a time-dependent vector field and generate structures via
 436 ODE or controlled SDE integration. SE(3)-equivariant frame flows introduced by FrameFlow target
 437 backbone generation with fast sampling and strong designability (Yim et al., 2023a), followed by im-
 438 provements in conditional motif-scaffolding (Yim et al., 2024) and added stochasticity for diversity
 439 (Bose et al., 2024). FoldFlow-base/OT/SFM refines flow-matching objectives and training, reporting
 440 improved designability and throughput across variants (Bose et al., 2024). In parallel, FoldFlow2
 441 advances training stability and sampling efficiency through refined flow objectives and architectural
 442 updates, achieving further gains in both accuracy and designability (Huguet et al., 2024). More
 443 recently, PROTEINA extends flow-based generation with hierarchical representations and geometry-
 444 aware conditioning, enabling scalable modeling of complex protein families (Geffner et al., 2025).
 445 In general, flows sample deterministically, support direct conditioning in the vector field, and offer
 446 faster sampling than diffusion in comparable setups.

447
 448 **Diffusion-based models.** Diffusion models construct a stochastic forward noising process and
 449 learn a reverse denoising dynamics to sample structures. RFDiffusion popularized diffusion in
 450 protein backbone design, enabling flexible conditional generation and motif scaffolding (Watson et al.,
 451 2022). Subsequent motif-conditioned diffusion systems extended conditioning mechanisms for joint
 452 sequence-structure tasks (Song et al., 2023). Latent diffusion accelerates sampling by operating
 453 in an autoencoded Euclidean latent space (Fu et al., 2023), trading some designability for speed.
 454 Large-scale scaffold diffusion emphasizes dataset scale and diversity, improving novelty and cover-
 455 age while retaining conditional control (Lin et al., 2024b). Diffusion models are versatile and robust
 456 but typically require more sampling steps and, when operating in SE(3) or Euclidean latents, do not
 457 align natively with the periodic geometry of dihedral angles.

458 **Chirality Constraints.** Chirality, the property of a molecule being non-superimposable on its
 459 mirror image, is a critical factor in chemistry and pharmacology. Enantiomers, or pairs of chiral
 460 molecules, can exhibit vastly different biological activities; for instance, one enantiomer of a drug
 461 can be therapeutic while the other is toxic (Dumitrescu et al., 2025a; Gaiński et al., 2023). Con-
 462 sequently, the ability of computational models to accurately represent and predict chirality is of
 463 paramount importance, particularly in drug discovery (Gaiński et al., 2023). Historically viewed as
 464 a binary property, recent work has also focused on developing mathematical methods to quantify
 465 chirality as a continuous variable, with primary approaches based on either geometric overlap, such
 466 as the Continuous Chirality Measure (Grieder et al., 2025), or scalar triple products, like the Chirality
 467 Characteristic (χ) (Abramson et al., 2024). Despite its importance, modeling chirality presents
 468 a significant challenge for many previous models such as SE(3)-based models (Dumitrescu et al.,
 469 2025a; Childs et al., 2025).

470
 471

6 CONCLUSION

472
 473 We present **ChiFlow**, a novel generative model for protein backbones that directly addresses the
 474 critical challenge of protein backbone chirality preservation—a fundamental constraint overlooked
 475 by existing SE(3)-based methods. By reformulating structure generation on a toroidal riemannian
 476 manifold and incorporating an asymmetry constraint, ChiFlow guarantees perfect chiral consistency
 477 without compromising performance. Furthermore, by operating natively on the space of torsion
 478 angles, ChiFlow avoids the computational overparameterization associated with SE(3) models and
 479 well models the dihedral angles of the protein backbone, resulting in a simpler architecture and faster
 480 sampling. Also, by utilizing the SFM, we improve the diversity of the protein backbone generated by
 481 ChiFlow. Despite these advantages, the torsional representation is inherently sensitive to local angu-
 482 lar deviations, which can accumulate over very long proteins. Future work could focus on improving
 483 the stability of the generative process for extended chains, as well as exploring conditional genera-
 484 tion for functional motifs and multimeric assemblies. ChiFlow establishes the toroidal manifold as
 485 an efficient way for protein backbone generation, offering a promising path toward chirality-aware
 486 design in computational structural biology.

486 REFERENCES
487488 Osama Abdin and Philip M. Kim. Pepflow: Direct conformational sampling from peptide energy
489 landscapes through hypernetwork-conditioned diffusion, June 2023.490 Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf
491 Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein,
492 David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool,
493 Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridg-
494 land, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael
495 Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishabh Jain, Yousuf A. Khan, Caroline M. R. Low,
496 Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisun-
497 daram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Vic-
498 tor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate
499 structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500,
500 June 2024. ISSN 1476-4687.501 Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unify-
502 ing framework for flows and diffusions, November 2023.
503504 Brian D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their*
505 *Applications*, 12(3):313–326, May 1982. ISSN 0304-4149.506 E. G. Atavin and L. V. Vilkov. Dihedral angles in cyclic molecules. *Journal of Structural Chemistry*,
507 44(5):846–851, September 2003. ISSN 1573-8779.
508509 Avishek Joey Bose, Tara Akhoud-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks,
510 Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander
511 Tong. Se(3)-stochastic flow matching for protein backbone generation, April 2024.512 Stephen K. Burley, Helen M. Berman, Jose M. Duarte, Zukang Feng, Justin W. Flatt, Brian P.
513 Hudson, Robert Lowe, Ezra Peisach, Dennis W. Piehl, Yana Rose, Andrej Sali, Monica Sekharan,
514 Chenghua Shao, Brinda Vallat, Maria Voigt, John D. Westbrook, Jasmine Y. Young, and Christine
515 Zardecki. Protein data bank: A comprehensive review of 3d structure holdings and worldwide
516 utilization by researchers, educators, and students. *Biomolecules*, 12(10):1425, October 2022.
517 ISSN 2218-273X.
518

519 Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries, February 2024.

520 Henry Childs, Pei Zhou, and Bruce R. Donald. Has alphafold 3 solved the protein folding problem
521 for d-peptides?, March 2025.
522523 Lorenzo Ciampiconi, Adam Elwood, Marco Leonardi, Ashraf Mohamed, and Alessandro Rozza. A
524 survey and taxonomy of loss functions in machine learning, November 2024.525 Pedro Cintas. Chirality of living systems: A helping hand from crystals and oligopeptides. *Ange-
526 wandte Chemie (International Ed. in English)*, 41(7):1139–1145, April 2002. ISSN 1433-7851.
527528 Alexandru Dumitrescu, Dani Korpela, Markus Heinonen, Yogesh Verma, Valerii Iakovlev, Vikas
529 Garg, and Harri Lähdesmäki. E(3)-equivariant models cannot learn chirality: Field-based molec-
530 ular generation, April 2025a.
531532 Alexandru Dumitrescu, Dani Korpela, Markus Heinonen, Yogesh Verma, Valerii Iakovlev, Vikas
533 Garg, and Harri Lähdesmäki. E(3)-equivariant models cannot learn chirality: Field-based molec-
534 ular generation, April 2025b.
535Stephan Eckstein and Aziz Lakhal. Exponential convergence of general iterative proportional fitting
536 procedures, April 2025.
537538 S. Walter Englander and Leland Mayne. The nature of protein folding pathways. *Proceedings of*
539 *the National Academy of Sciences*, 111(45):15873–15880, 2014. doi: 10.1073/pnas.1411798111.
URL <https://www.pnas.org/doi/abs/10.1073/pnas.1411798111>.

540 Cong Fu, Keqiang Yan, Limei Wang, Wing Yee Au, Michael McThrow, Tao Komikado, Koji
 541 Maruhashi, Kanji Uchino, Xiaoning Qian, and Shuiwang Ji. A latent diffusion model for pro-
 542 tein structure generation, December 2023.

543

544 Fabian B. Fuchs, Daniel E. Worrall, Volker Fischer, and Max Welling. Se(3)-transformers: 3d roto-
 545 translation equivariant attention networks, November 2020.

546

547 Piotr Gaiński, Michał Koziarski, Jacek Tabor, and Marek Śmieja. Chienn: Embracing molecular
 548 chirality with graph neural networks, July 2023.

549

550 Tomas Geffner, Kieran Didi, Zuobai Zhang, Danny Reidenbach, Zhonglin Cao, Jason Yim, Mario
 551 Geiger, Christian Dallago, Emine Kucukbenli, Arash Vahdat, and Karsten Kreis. Proteina: Scal-
 552 ing flow-based protein structure generative models, March 2025.

553

554 Andrew Grieder, Shihao Tu, and Yuan Ping. Relation of continuous chirality measure to spin and
 555 orbital polarization, and chiroptical properties in solids, September 2025.

556

557 Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J. Sofroniew, Deniz Oktay, Zeming Lin, Robert
 558 Verkuil, Vincent Q. Tran, Jonathan Deaton, Marius Wiggert, Rohil Badkundri, Irhum Shafkat,
 559 Jun Gong, Alexander Derry, Raul S. Molina, Neil Thomas, Yousuf A. Khan, Chetan Mishra, Car-
 560 olyn Kim, Liam J. Bartie, Matthew Nemeth, Patrick D. Hsu, Tom Sercu, Salvatore Candido, and
 561 Alexander Rives. Simulating 500 million years of evolution with a language model, December
 562 2024.

563

564 Guillaume Huguet, James Vuckovic, Kilian Fatras, Eric Thibodeau-Laufer, Pablo Lemos, Riashat
 565 Islam, Cheng-Hao Liu, Jarrid Rector-Brooks, Tara Akhoun-Sadegh, Michael Bronstein, Alexan-
 566 der Tong, and Avishek Joey Bose. Sequence-augmented se(3)-flow matching for conditional
 567 protein backbone generation, December 2024.

568

569 Yang Jing, Lei Li, and Jingtong Zhang. Solving schrödinger bridge problem via continuous normal-
 570 izing flow, March 2025.

571

572 Gen Li, Yanxi Chen, Yu Huang, Yuejie Chi, H. Vincent Poor, and Yuxin Chen. Fast computation of
 573 optimal transport via entropy-regularized extragradient methods, June 2024.

574

575 Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff
 576 Smith, Brian Vaughan, Pritam Damania, and Soumith Chintala. Pytorch distributed: Experiences
 577 on accelerating data parallel training, June 2020.

578

579 Haitao Lin, Odin Zhang, Huifeng Zhao, Dejun Jiang, Lirong Wu, Zicheng Liu, Yufei Huang, and
 580 Stan Z. Li. Ppflow: Target-aware peptide design with torsional flow matching, December 2024a.

581

582 Yeqing Lin, Minji Lee, Zhao Zhang, and Mohammed AlQuraishi. Out of many, one: Designing and
 583 scaffolding proteins at the scale of the structural universe with genie 2, May 2024b.

584

585 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 586 for generative modeling, February 2023.

587

588 Emile Mathieu and Maximilian Nickel. Riemannian continuous normalizing flows, December 2020.

589

590 Jerod Parsons, J. Bradley Holmes, J. Maurice Rojas, Jerry Tsai, and Charlie E. M. Strauss. Prac-
 591 tical conversion from torsion space to cartesian space for in silico protein synthesis. *Journal of*
 592 *Computational Chemistry*, 26(10):1063–1068, 2005. ISSN 1096-987X.

593

594 Veronica Pollock. Proteins. In S. J. Enna and David B. Bylund (eds.), *xPharm: The Compre-
 595 hensive Pharmacology Reference*, pp. 1–11. Elsevier, New York, January 2007. ISBN 978-0-08-055232-
 596 3.

597

598 G.N. Ramachandran, C. Ramakrishnan, and V. Sasisekharan. Stereochemistry of polypeptide chain
 599 configurations. *Journal of Molecular Biology*, 7(1):95–99, July 1963. ISSN 00222836.

600

601 Srinivas Ramachandran, Pradeep Kota, Feng Ding, and Nikolay V. Dokholyan. Automated mini-
 602 mization of steric clashes in protein structures. *Proteins*, 79(1):261–270, January 2011. ISSN
 603 0887-3585.

594 Mauro Rosestolato. A note on stochastic fubini's theorem and stochastic convolution, June 2018.
 595

596 Yujie Shi, Alex Jie Yang, and Sanhong Deng. Entropy in science of science, March 2025.

597 Zhenqiao Song, Yunlong Zhao, Yufei Song, Wenxian Shi, Yang Yang, and Lei Li. Joint design of
 598 protein sequence and structure based on motifs, October 2023.

599

600 Xiangru Tang, Howard Dai, Elizabeth Knight, Fang Wu, Yunyang Li, Tianxiao Li, and Mark Ger-
 601 stein. A survey of generative ai for *de novo* drug design: New frontiers in molecule and protein
 602 generation. *Briefings in Bioinformatics*, 25(4):bbae338, May 2024a. ISSN 1467-5463, 1477-
 603 4054.

604 Zhicong Tang, Tiankai Hang, Shuyang Gu, Dong Chen, and Baining Guo. Simplified diffusion
 605 schrödinger bridge, October 2024b.

606

607 Alexis M. H. Teter, Wenqing Wang, and Abhishek Halder. Schrödinger bridge with quadratic state
 608 cost is exactly solvable, October 2024.

609

610 Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E.
 611 Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M.
 612 Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham
 613 Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile
 614 Mathieu, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker.
 615 Broadly applicable and accurate protein design by integrating structure prediction networks and
 616 diffusion generative models. December 2022.

617 Michael Wiemeler. Torus manifolds and non-negative curvature. *Journal of the London Mathe-
 618 matical Society*, 91(3):667–692, June 2015. ISSN 00246107.

619 Jiawen Wu, Bingguang Chen, Yuyi Zhou, Qi Meng, Rongchan Zhu, and Zhi-Ming Ma. Riemannian
 620 neural geodesic interpolant, April 2025.

621 Jason Yim, Andrew Campbell, Andrew Y. K. Foong, Michael Gastegger, José Jiménez-Luna, Sarah
 622 Lewis, Victor Garcia Satorras, Bastiaan S. Veeling, Regina Barzilay, Tommi Jaakkola, and Frank
 623 Noé. Fast protein backbone generation with $se(3)$ flow matching, October 2023a.

624

625 Jason Yim, Brian L. Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay,
 626 and Tommi Jaakkola. $Se(3)$ diffusion model with application to protein backbone generation,
 627 May 2023b.

628

629 Jason Yim, Andrew Campbell, Emile Mathieu, Andrew Y. K. Foong, Michael Gastegger, José
 630 Jiménez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S. Veeling, Frank Noé, Regina
 631 Barzilay, and Tommi S. Jaakkola. Improved motif-scaffolding with $se(3)$ flow matching, July
 632 2024.

633 Yang Zhang, Zhewei Wei, Ye Yuan, Chongxuan Li, and Wenbing Huang. Equipocket: An $e(3)$ -
 634 equivariant geometric graph neural network for ligand binding site prediction, July 2024.

635 Hongyu Zhou and Zorah Lähner. Laplace-beltrami operator for gaussian splatting, February 2025.

636

637

638

639

640

641

642

643

644

645

646

647

648 A SUPPLEMENTARY PROOFS AND TECHNICAL DETAILS

649 A.1 SCHRÖDINGER BRIDGE THEORY

650 The Schrödinger bridge problem (SBP) finds the most likely stochastic evolution between two prob-
 651 ability distributions given a prior stochastic evolution. It is the dynamical version of the entropy-
 652 regularized optimal transport (OT) problem (Li et al., 2024), where the mean square distance is
 653 replaced by the relative entropy (Shi et al., 2025). Specifically, given a prior diffusion process Q^γ
 654 (e.g., a Wiener measure with volatility γ) and two endpoint marginals π_0 and π_1 , the dynamic SBP
 655 is defined as:

$$656 \inf_{Q \in \mathcal{D}(\pi_0, \pi_1)} D_{\text{KL}}(Q \| Q^\gamma), \quad (21)$$

657 where $\mathcal{D}(\pi_0, \pi_1)$ denotes the set of measures on path space with marginals π_0 at time 0 and π_1 at
 658 time 1. The minimizing Q^* is called the *Schrödinger bridge*, and its optimal value is the *entropic*
 659 *transportation cost*.

660 The SBP admits a dual formulation through the *Schrödinger system*, a pair of coupled PDEs (or
 661 integral equations in discrete settings) for potentials φ and $\hat{\varphi}$. For a diffusion prior with drift b_0 and
 662 volatility $\sqrt{\gamma}$, the optimal drifts of the forward and backward Schrödinger bridges are given by:

$$663 b^+(t) = \gamma \nabla_x \varphi(x^+(t), t), \quad b^-(t) = \gamma \nabla_x \hat{\varphi}(x^-(t), t), \quad (22)$$

664 where x^+ and x^- are the forward and backward processes, respectively. The potentials satisfy the
 665 boundary conditions:

$$666 \varphi(x, 0) \propto \frac{d\pi_0}{dQ_0^\gamma}(x), \quad \hat{\varphi}(x, 1) \propto \frac{d\pi_1}{dQ_1^\gamma}(x), \quad (23)$$

667 with Q_0^γ and Q_1^γ being the marginals of the prior at $t = 0$ and $t = 1$.

668 A key numerical approach for solving the SBP is the *iterative proportional fitting procedure* (IPFP)
 669 (Eckstein & Lakhai, 2025), also known as the Sinkhorn algorithm in the static OT setting. For
 670 dynamic SBP, IPFP alternates between solving forward and half-bridge problems:

$$671 P_i = \arg \inf_{P \in \mathcal{D}(\cdot, \pi_1)} D_{\text{KL}}(P \| Q_{i-1}), \quad (24)$$

$$672 Q_i = \arg \inf_{Q \in \mathcal{D}(\pi_0, \cdot)} D_{\text{KL}}(Q \| P_i), \quad (25)$$

673 which converges to the Schrödinger bridge Q^* . This procedure is equivalent to an iterative scaling
 674 algorithm in the discrete case.

675 Recently, (Teter et al., 2024) proposed a *regularized* SBP variant with a quadratic state cost-to-go,
 676 which incentivizes paths to stay close to a nominal level. Unlike the conventional SBP, this regu-
 677 larization induces a state-dependent rate of killing and creation of probability mass. Remarkably,
 678 they showed that this regularized SBP is *exactly solvable* even for non-Gaussian endpoints, by de-
 679 riving the Markov kernel of the associated reaction-diffusion PDE in closed form. Their solution
 680 recovers the heat kernel (and hence the conventional Schrödinger bridge) as a special case when the
 681 regularization vanishes.

682 The SBP also extends to quantum channels, where the goal is to scale a completely positive operator
 683 Q to a quantum channel R such that $R(\alpha) = \beta$ for given density matrices α and β . This is the
 684 quantum analog of the classical SBP and can be solved using fixed-point methods.

695 A.2 SE(3)-INVARIANT FEATURE DISTRIBUTIONS CANNOT ENCODE CHIRALITY

696 We formalize why standard E(3) / SE(3)-invariant (or equivariant) architectures based purely on Eu-
 697 clidean distances and inner products cannot distinguish mirror-image molecular or protein backbone
 698 configurations. This limitation motivates our explicit asymmetry constraint on dihedral torus flows.

699 **Proposition A.1** (Inability of E(3)-invariant point-cloud distributions to distinguish enantiomers).
 700 Let p_ϕ be a distribution over point clouds $M = (m_1, \dots, m_N) \in (\mathbb{R}^3)^N$ parameterized only
 701 through features that are invariant under the action of the Euclidean group E(3) (translations +

702 rotations) and depend exclusively on pairwise distances $\|m_i - m_j\|$ and/or inner products of cen-
 703 tered coordinates. Then for any enantiomeric pair (M, M') related by an improper orthogonal
 704 transformation (a reflection), we have

$$705 \quad 706 \quad p_\phi(M) = p_\phi(M'). \quad (26)$$

707 Consequently, p_ϕ cannot enforce or prefer a single molecular chirality.

709 **Sketch of proof.** Any collection of features constructed from (i) pairwise Euclidean distances
 710 $d_{ij} = \|m_i - m_j\|$, (ii) dot products of centered vectors $\langle m_i - \bar{m}, m_j - \bar{m} \rangle$, or (iii) higher-order
 711 tensor contractions thereof is invariant under the full orthogonal group $O(3)$, because both distances
 712 and inner products are preserved by any orthogonal matrix Q with $Q^T Q = I$, including reflections
 713 ($\det Q = -1$). A reflection that maps M to its mirror image M' lies in $O(3) \setminus SO(3)$. Thus all such
 714 features take identical values on M and M' (Dumitrescu et al., 2025b). strengthen this by showing
 715 that for functions of at most n input vectors, $SE(n)$ -invariance already implies reflection invariance
 716 via centering arguments and Householder reflections; hence $SE(3)$ -invariant parameterizations can-
 717 not break mirror symmetry. Therefore any likelihood or density model p_ϕ expressed purely in those
 718 invariants assigns equal probability to both enantiomers.

719 **Implication for ChiFlow.** To obtain chirality selectivity, one must either (a) augment invariant
 720 features with orientation-sensitive pseudoscalars (e.g., triple products) or (b) impose explicit asym-
 721 metric constraints under a mirroring isometry, as we do via the torsional mirroring operator M and
 722 the chiral penalty (Section 2.4).

723 A.3 PROOF OF PROPOSITION 3.1 (CONDITIONAL DRIFT UNDER TORUS INTERPOLATION)

726 We give a concise proof for the closed form of the conditional target drift on the torus under the
 727 time-dependent affine interpolation used in ChiFlow. Let $\tau \in \mathbb{T}^d$ denote a generic angle vector
 728 (componentwise on the circle), and consider an interpolation of the form

$$729 \quad \tau_t = (\mu_t(\tau_1) + \sigma_t(\tau_1) \varepsilon) \bmod 2\pi, \quad t \in [0, 1], \quad (27)$$

730 where μ_t and $\sigma_t > 0$ are smooth in t (and may depend on the endpoint τ_1), and ε is time-invariant
 731 along the conditional path. Differentiating τ_t w.r.t. t in the tangent yields

$$732 \quad 733 \quad u_t(\tau | \tau_1) = \frac{d}{dt} \tau_t = \dot{\sigma}_t \varepsilon + \dot{\mu}_t. \quad (28)$$

735 By rearranging the interpolation, we have the (componentwise, principal-branch) identity on the
 736 circle

$$737 \quad \varepsilon = \frac{\text{wrap}(\tau_t - \mu_t)}{\sigma_t}, \quad (29)$$

739 so substituting eliminates ε and expresses the drift as a function of (τ, t) :

$$740 \quad 741 \quad u_t(\tau | \tau_1) = \dot{\sigma}_t \frac{\text{wrap}(\tau_t - \mu_t)}{\sigma_t} + \dot{\mu}_t = \frac{\dot{\sigma}_t(\tau_1)}{\sigma_t(\tau_1)} \text{wrap}(\tau - \mu_t(\tau_1)) + \dot{\mu}_t(\tau_1), \quad (30)$$

742 where all operations are taken componentwise on \mathbb{T}^d using the shortest-arc difference $\text{wrap}(\cdot) \in$
 743 $(-\pi, \pi]$. This is the claimed expression.

745 A.4 PROOF OF PROPOSITION 3.2 (UNBIASEDNESS VIA DISINTEGRATION ON \mathbb{T}^d)

747 We show the equivalence between the conditional objective and its marginal counterpart by dis-
 748 integration of measures on the product torus and Fubini's theorem (Rosestolato, 2018). Let
 749 $\mathbb{T}^d = S^1 \times \dots \times S^1$ and assume the endpoint distributions $p_0, p_1 \in \mathcal{P}(\mathbb{T}^d)$ factorize across angles
 750 (orthogonality/independence of torsions), hence the interpolant p_t admits the same disintegration.
 751 Define

$$752 \quad 753 \quad u_t(\tau) = \mathbb{E}_{\tau_0 \sim p_0, \tau_1 \sim p_1} \left[\frac{p_t(\tau | \tau_0, \tau_1)}{p_t(\tau)} u_t(\tau | \tau_0, \tau_1) \right], \quad (31)$$

754 and consider the objective difference

$$755 \quad \nabla_\theta \left(\mathbb{E}_{\tau_0, \tau_1, \tau \sim p_t(\cdot | \tau_0, \tau_1)} \left[\|v_t(\tau) - u_t(\tau | \tau_0, \tau_1)\|_g^2 \right] - \mathbb{E}_{\tau \sim p_t} \left[\|v_t(\tau) - u_t(\tau)\|_g^2 \right] \right), \quad (32)$$

756 where $\langle \cdot, \cdot \rangle_g$ and $\|\cdot\|_g$ denote the Riemannian inner product and norm induced by the product metric
 757 g on \mathbb{T}^d . Expanding squares and canceling like terms reduces this to
 758

$$759 -2 \nabla_\theta \left(\mathbb{E}_{\tau_0, \tau_1, \tau \sim p_t(\cdot | \tau_0, \tau_1)} \langle v_t(\tau), u_t(\tau | \tau_0, \tau_1) \rangle_g - \mathbb{E}_{\tau \sim p_t} \langle v_t(\tau), u_t(\tau) \rangle_g \right). \quad (33)$$

761 Using the definition of $u_t(\tau)$ and changing the order of integration,
 762

$$763 \mathbb{E}_{\tau \sim p_t} \langle v_t(\tau), u_t(\tau) \rangle_g = \int \langle v_t(\tau), u_t(\tau) \rangle_g p_t(\tau) d\tau \quad (34)$$

$$765 = \int \left\langle v_t(\tau), \mathbb{E}_{\tau_0, \tau_1} \left[\frac{p_t(\tau | \tau_0, \tau_1)}{p_t(\tau)} u_t(\tau | \tau_0, \tau_1) \right] \right\rangle_g p_t(\tau) d\tau \quad (35)$$

$$767 = \iint \left\langle v_t(\tau), u_t(\tau | \tau_0, \tau_1) \right\rangle_g p_t(\tau | \tau_0, \tau_1) p_0(\tau_0) p_1(\tau_1) d\tau d\tau_0 d\tau_1 \quad (36)$$

$$769 = \mathbb{E}_{\tau_0 \sim p_0, \tau_1 \sim p_1, \tau \sim p_t(\cdot | \tau_0, \tau_1)} \langle v_t(\tau), u_t(\tau | \tau_0, \tau_1) \rangle_g. \quad (37)$$

771 Hence the two expectations are equal, and the gradient of their difference is zero. This establishes
 772 that optimizing the conditional objective is an unbiased surrogate for the marginal objective on \mathbb{T}^d
 773 under the stated disintegration assumptions.

774 A.5 PROOF THAT THE MIRRORING OPERATOR M IS AN ISOMETRIC DIFFEOMORPHISM ON \mathbb{T}^3

775 We show that the mirroring map $M : \mathbb{T}^3 \rightarrow \mathbb{T}^3$ defined componentwise by
 776

$$778 M(\phi, \psi, \omega) = (2\pi - \phi \bmod 2\pi, 2\pi - \psi \bmod 2\pi, \omega) \quad (38)$$

779 is an isometric diffeomorphism under the conformal product metric introduced in Section 2 (the
 780 extension to \mathbb{T}^{3N} follows by the product structure, since M acts residuewise).
 781

782 **1) M is a diffeomorphism.** Smoothness: Each component of M is the composition of a smooth
 783 linear map on the circle (e.g., $\phi \mapsto 2\pi - \phi$) with the canonical identification modulo 2π that defines
 784 the smooth structure of the torus. Hence M is smooth. Invertibility: M is an involution, i.e.,
 785 $M \circ M = \text{id}$, so $M^{-1} = M$. Since M is smooth, so is M^{-1} . Therefore M is a diffeomorphism on
 786 \mathbb{T}^3 ; by independence across residues, the product map on \mathbb{T}^{3N} is also a diffeomorphism.
 787

788 **2) M is an isometry for the metric $g_{\mathbb{T}^3}$.** Recall the metric
 789

$$790 g_{\mathbb{T}^3}|_\chi(u, v) = \sum_{k \in \{\phi, \psi, \omega\}} w_k(\chi_k) u_k v_k, \quad w_k(\chi_k) = \frac{1}{1 + \sin^2(\chi_k/2)}. \quad (39)$$

793 Let $\chi' = M(\chi)$ and let $M_* : T_\chi \mathbb{T}^3 \rightarrow T_{\chi'} \mathbb{T}^3$ be the pushforward of M . By differentiation of
 794 the component maps, the Jacobian of M is $J_M = \text{diag}(-1, -1, 1)$, hence for any $u \in T_\chi \mathbb{T}^3$ we
 795 have $M_* u = J_M u$ with components $(M_* u)_\phi = -u_\phi$, $(M_* u)_\psi = -u_\psi$, $(M_* u)_\omega = u_\omega$. Using the
 796 trigonometric identity $\sin(\pi - \theta) = \sin \theta$, we obtain for $k \in \{\phi, \psi\}$

$$797 w_k(\chi'_k) = \frac{1}{1 + \sin^2((2\pi - \chi_k)/2)} = \frac{1}{1 + \sin^2(\chi_k/2)} = w_k(\chi_k), \quad \text{and} \quad w_\omega(\chi'_\omega) = w_\omega(\chi_\omega). \quad (40)$$

800 Therefore, for any $u, v \in T_\chi \mathbb{T}^3$,

$$802 g_{\mathbb{T}^3}|_{\chi'}(M_* u, M_* v) = \sum_{k \in \{\phi, \psi, \omega\}} w_k(\chi'_k) (M_* u)_k (M_* v)_k \quad (41)$$

$$804 = w_\phi(\chi'_\phi)(-u_\phi)(-v_\phi) + w_\psi(\chi'_\psi)(-u_\psi)(-v_\psi) + w_\omega(\chi'_\omega)u_\omega v_\omega \quad (42)$$

$$806 = w_\phi(\chi_\phi)u_\phi v_\phi + w_\psi(\chi_\psi)u_\psi v_\psi + w_\omega(\chi_\omega)u_\omega v_\omega \quad (43)$$

$$808 = g_{\mathbb{T}^3}|_\chi(u, v). \quad (44)$$

809 which proves that M preserves the metric, hence is an isometry. The product argument yields the
 810 same conclusion on \mathbb{T}^{3N} .

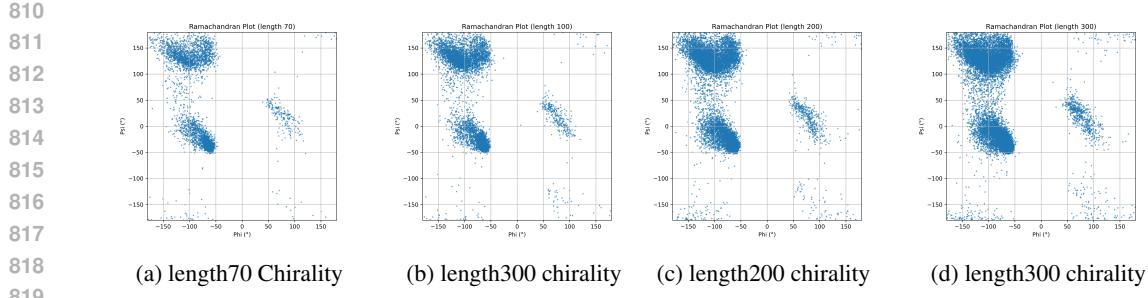


Figure 6: Ramachandran distribution image of the generated backbone of different length.

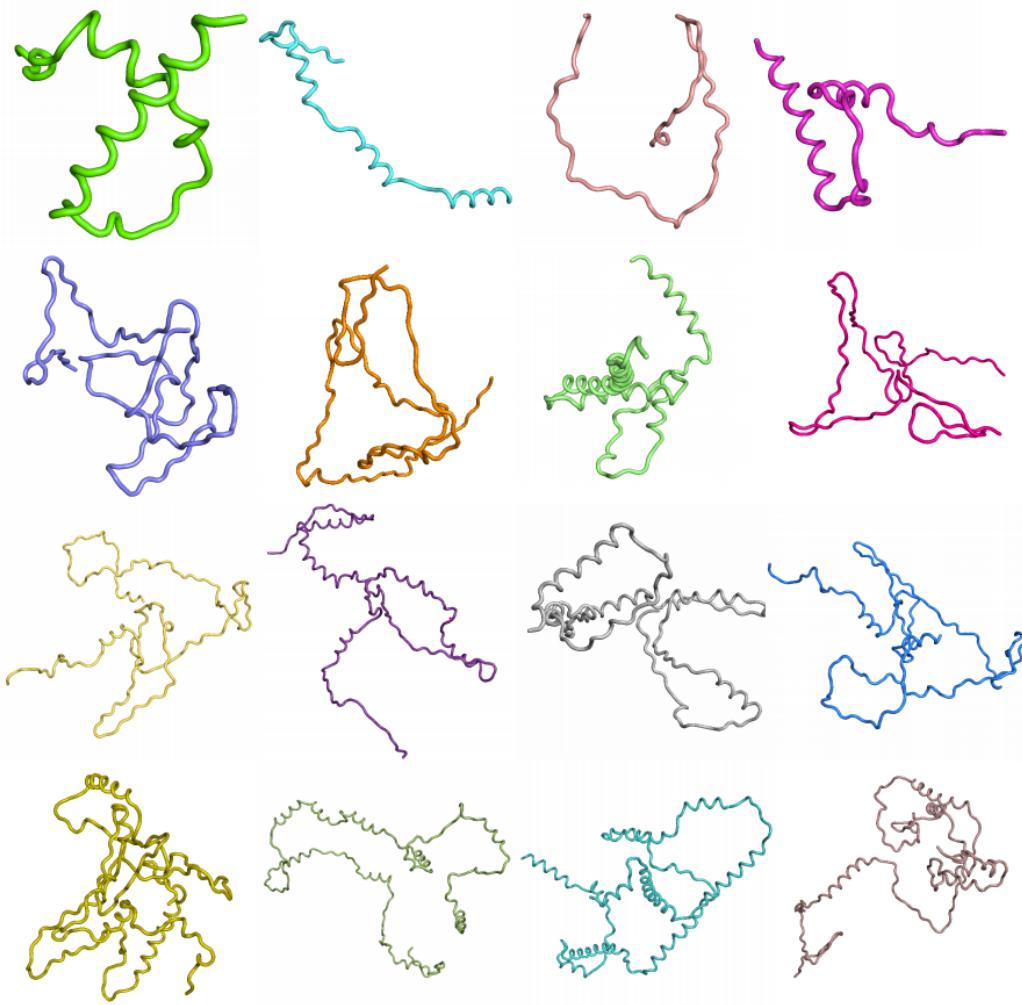


Figure 7: Different length of Protein Backbone Designed by ChiFlow. Varing from 60 to 256.

858
859
860
861
862
863

Consequence. Since M is an isometric diffeomorphism and an involution, the asymmetry constraint $\hat{v}(M\chi, t) = -M_*\hat{v}(\chi, t)$ is metric-consistent and well-defined everywhere on \mathbb{T}^{3N} .

864 A.6 THE USE OF LARGE LANGUAGE MODELS
865866 In this paper, we used Deepseek to polish our writing and made use of copilot to help debug our
867 code.
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917