MuCO2: Unlocking Multi-Stage Conformation Generation for Noncanonical Cyclic Peptides

Anonymous authors
Paper under double-blind review

Abstract

Cyclic peptides containing noncanonical amino acids (ncAAs) represent an indispensable modality in modern drug discovery, offering improved proteolytic stability and therapeutic profiles. However, efficiently modeling their conformational ensembles across complex energy landscapes remains challenging: existing methods like the HighFold series and MuCO mainly focus on canonical peptides and struggle with complicated structural perturbations. Building upon MuCO, we propose MuCO2, a multi-stage conformation generation framework for noncanonical cyclic peptides with α-backbone modifications, stereochemical transformations, and open-ended side-chain modifications. In particular, MuCO2 represents open-set ncAAs by aligning the atomic representations of Uni-Mol2 with the residue-level embeddings from ESM-2. Leveraging a curriculum learning paradigm driven by conformational energy and symmetries, MuCO2 learns a quaternion flow-based cyclic backbone generator coupled with a structure-guided Feynman-Kac steering mechanism, enabling efficient, generalizable cyclic peptide sampling conditioned on chirality parity and N-methylation. Extensive evaluations on canonical and noncanonical benchmarks show that MuCO2 achieves high-throughput cyclic peptide sampling with near-perfect stereochemical preservation, consistently outperforming state-of-the-art methods in physical stability, structural diversity, and potential-well coverage.

1 Introduction

Cyclic peptides occupy a privileged middle ground between small molecules and macromolecular biologics and have attracted considerable interest for drug discovery (Jing & Jin, 2020; Ji et al., 2024b). The ring-closed structure of cyclic peptides offers them greater proteolytic stability and better pharmaceutical properties than linear peptides (Vinogradov et al., 2019; Dougherty et al., 2019; Lee et al., 2019; Mannes et al., 2022) but also leads to highly dynamic conformations with complex energy landscapes (Schlick et al., 2021; Jiang et al., 2025). Physical simulation (Hollingsworth & Dror, 2018; Pracht et al., 2020) can provide high-precision conformations, whose computational complexity, however, is too high for large-scale exploration in drug development. Therefore, developing learning-based solutions to generate cyclic peptide conformations reliably and efficiently is critical.

Currently, HighFold and its variants (Zhu et al., 2025; Cao et al., 2025b;a; 2026) achieve end-to-end generation of cyclic peptide conformations, which are built on protein structure prediction models such as the AlphaFold series (Jumper et al., 2021; Abramson et al., 2024) and Boltz2 (Passaro et al., 2025). To improve efficiency and mitigate mode collapse, MuCO (Wang et al., 2026) establishes a multi-stage conformation generation paradigm for cyclic peptides, consisting of backbone generation, side-chain packing, and force-field refinement. By incorporating amortized conformation sampling into a hierarchical generative model, MuCO achieves high-throughput, diverse cyclic peptide sampling.

Notably, modern drug discovery routinely requires incorporating *noncanonical amino acids* (ncAAs) into cyclic peptides to fine-tune their properties, yet extending the above generative models to noncanonical chemistry remains challenging. In practice, noncanonical residues in macrocycle design are primarily based on three perturbations (Gfeller et al., 2012; Hickey et al., 2023; Katoh & Suga, 2022): α-backbone modifications (e.g., N-methylation[1]), stereochemical transformations, and

[1] N-methylation replaces the amide N–H with N–CH_3, eliminating a hydrogen-bond donor and shifting the *cis*/*trans* torsional preference of the preceding peptide bond (Slough et al., 2017).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

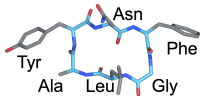
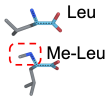
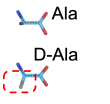
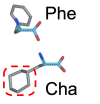
side-chain modifications (along with less common backbone extensions and non-amide linkages; see Appendix A for the full five-category taxonomy). In principle, HighFold models can only generate noncanonical cyclic peptides in one or two categories, and their performance is often undesirable. For the multi-stage method MUCO, its feasibility in the noncanonical regime is even worse because i) it applies ESM-2 (Lin et al., 2023) to represent amino acids, which cannot deal with ncAAs properly, and ii) it does not encode structural perturbations, e.g., chirality inversion, global ring reorganization, and localized steric crowding from N-modifications (Damjanovic et al., 2021; Slough et al., 2017).

Building on MUCO, we propose MUCO2 to generate noncanonical cyclic peptide conformations, a multi-stage generation method for this challenging task.

1. We align atomic embeddings of Uni-Mol2 (Ji et al., 2024a) with residue-level embeddings of ESM-2, enabling MUCO2 to encode both canonical and noncanonical amino acids.
2. In the aspect of model architecture, we replace the Riemannian SE(3)-Flow model (Bose et al., 2024; Yim et al., 2023b;a) used in MUCO with the Quaternion Flow model (Yue et al., 2025) for accelerated cyclic backbone generation and extend the model to support injecting noncanonical side-chain, sign of chirality, and N-modifications as conditions during generation.
3. We train our model with a data augmentation framework driven by conformational energy and symmetries, leading to a curriculum learning paradigm. The current model generates low-energy cyclic peptides with N-methylated and mixed-chirality macrocycles (Bayaraa et al., 2026) to support next-round training, gradually mitigating data scarcity and improving model generalizability.
4. We advance MUCO's sampling strategy by developing a *structure-guided Feynman-Kac steering* mechanism, which supports high-throughput conformation sampling with improved diversity.

As shown in Table 1, the above technical innovations ensure that MUCO2 can efficiently generate noncanonical cyclic peptides across the three most prevalent perturbations. We evaluate MUCO2 across canonical CPSea benchmarks (Yang et al., 2026), their noncanonical variants, and external noncanonical datasets (Zhu et al., 2025; Li et al., 2023). Extensive experiments demonstrate that MUCO2 achieves superior or at least competitive performance compared to state-of-the-art methods in physical stability, structural diversity, and potential well coverage.

Table 1: Illustrations of a canonical cyclic peptide and three typical perturbations. For each canonical cyclic peptide generator, "✓" and "✗" indicate whether it can generate noncanonical peptides with the corresponding perturbation or not; "✓/✗" indicates support under restricted conditions and with undesirable performance.

Canonical Cyclic Peptide	α-Backbone Modifications	Stereochemical Transformations	Side-chain Modifications
			
HighFold2	✗	✓	✓/✗
HighFold-MeD	✓/✗	✓	✗
HighFold-MeD2	✓/✗	✓	✗
HighFold3	✓	✓/✗	✓
MuCO	✗	✗	✗
MuCO2	✓	✓	✓

2 Related Work

Physics-Based Macrocycle Conformation Sampling. Molecular dynamics (MD) and enhanced-sampling methods remain the gold standard for mapping cyclic peptide solution landscapes (Damjanovic et al., 2021), but they suffer from severe kinetic trapping—most notably across the slow *cis/trans* isomerization of N-methylated backbones (Slough et al., 2017). Search-and-score protocols (e.g., Rosetta kinematic closure and CREST) and large-scale simulation efforts (e.g., CREMP and CycPeptMPDB-4D) require millions of CPU hours or extensive multi-solvent sampling to generate conformer libraries (Grambow et al., 2024; Liu et al., 2026). Moreover, simulating novel ncAAs incurs tedious parameterization bottlenecks for partial charges and bonded potentials (Wijesiri & Brown, 2025).

Deep Learning for Cyclic Peptide Structure and Ensembles. Recent machine learning methods for cyclic peptides broadly divide into two paradigms. The first focuses on reference-structure prediction: HighFold2 and NCPepFold extend protein folding models with atom-level features (Zhu et al., 2025; Mao et al., 2025); HighFold3 conditions all-atom diffusion on Chemical Component Dictionary (CCD) definitions (Cao et al., 2025a); and HighFold-MeD/MeD2 distill Rosetta structures for D- and N-methylated peptides (Cao et al., 2025b; 2026). However, these models target deterministic single structures rather than multi-basin energy landscapes, lacking compositional generalization

across dense modifications. The second paradigm directly generates conformational ensembles: StrEAMM (Miao et al., 2021; Hui et al., 2023) and RINGER (Grambow et al., 2023) amortize MD distributions; AGDIFF predicts all-atom graphs on CREMP (Wu & Zou, 2026); and recent diffusion models learn mixed-chirality ensembles from MD trajectories (Bayaraa et al., 2026). Nevertheless, the above methods remain restricted to fixed ring sizes, closed vocabularies, or dense trajectory data. Specifically, protein language models such as ESM-2 (Lin et al., 2023) capture sequence context but cannot encode noncanonical chemistry, which limits the methods like MuCO (Wang et al., 2026) in noncanonical scenarios. While molecular graphs (e.g., AGDIFF (Wu & Zou, 2026)) and CCD definitions (e.g., HighFold3 and HighFold-MeD2 (Cao et al., 2025a; 2026)) bypass vocabulary limits, they lack 3D stereoelectronic features and struggle with backbone-state shifts.

Building upon the hierarchical framework of MuCO, MuCO2 overcomes the above problems and enables scalable multi-basin exploration with open-ended noncanonical generalization.

3 Proposed Model

3.1 Data Representation and Modeling Principle

For canonical cyclic peptides, the conformation generation task aims at generating their conformations based on the amino acid sequences of the corresponding linear peptides. Typically, denote a linear peptide of length L as $\mathcal{S} = \{s_i\}_{i=1}^{L}$, where $s_i \in \mathcal{A}$ denotes the amino acid type of the i-th residue and $\mathcal{A}$ is the set of 20 canonical amino acids. The conformation of its corresponding canonical cyclic peptides is represented as a 3D coordinate matrix $\boldsymbol{X} \in \mathbb{R}^{N\times 3}$, where N is the number of atoms. The conformation can be equivalently represented as a backbone scaffold $\mathcal{B} = \{(\boldsymbol{T}_i, \psi_i)\}_{i=1}^{L}$ and side-chain torsions $\{\boldsymbol{\chi}_i\}_{i=1}^{L}$, where $\boldsymbol{T}_i \in \mathrm{SE}(3)$ denotes the local frame of the i-th residue (Jumper et al., 2021), $\psi_i \in \mathbb{T}$, where $\mathbb{T} = \mathbb{R} \setminus 2\pi\mathbb{Z}$, represents the backbone torsion angle for the carbonyl oxygen, and $\boldsymbol{\chi}_i \in \mathbb{T}^4$ contains up to four side-chain torsional angles.

For noncanonical cyclic peptides, however, their conformations are generated by first cyclizing linear peptides and then modifying residues to ncAAs by various perturbations. Therefore, MuCO2 **represents noncanonical cyclic peptides as canonical priors superposed by perturbations.** The sequence of canonical peptide $\mathcal{S}$ provides us with a prior for generating noncanonical cyclic peptides. The perturbations of its residues are represented as $\mathcal{R} = \{r_i\}_{i=1}^{L}$, where $r_i = (m_i, c_i, \eta_i, \mathcal{G}_i, \boldsymbol{Q}_i)$ indicates the three primary perturbations of the i-th residue:

- **α-backbone and side-chain modifications:** $m_i \in \{0, 1\}$ indicates the N-methylation of the i-th residue. $\mathcal{G}_i$ and $\boldsymbol{Q}_i \in \mathbb{R}^{N_i\times 3}$ are the side-chain chemical graph and its template conformation in Chemical Component Dictionary (Westbrook et al., 2015), describing more complex α-modifications, e.g., peptoid-like bulky N-modifications, and open-set side-chain modifications.
- **Stereochemical transformation:** $c_i \in \{1, -1\}$ indicates the stereochemical chirality (i.e., $c_i = 1$ for L-amino acid and -1 for D-amino acid), and $\eta_i \in \{0, 1\}$ indicates whether the residue is achiral Glycine ($\eta_i = 0$) or not ($\eta_i = 1$).

Notably, for the i-th residue, we represent its side-chain by the dihedral angles $\boldsymbol{\chi}_i$ if it is canonical. Otherwise, we represent its side-chain by atomic coordinates, denoted as $\boldsymbol{C}_i \in \mathbb{R}^{N_i\times 3}$. As a result, the set of side-chains $\mathcal{C} = \{\boldsymbol{\chi}_i\}_{i\in\mathcal{I}_c} \cup \{\boldsymbol{C}_i\}_{i\in\mathcal{I}_n}$, where $\mathcal{I}_c$ and $\mathcal{I}_n$ are sets of canonical and noncanonical residues, respectively. Given $(\mathcal{S}, \mathcal{R})$, we aim to efficiently sample physically valid, diverse, and low-energy conformational ensembles for the corresponding cyclic peptide. MuCO2 learns a hierarchical generative model to achieve this aim:

$$p(\boldsymbol{X} \mid \mathcal{S}, \mathcal{R}) = p(\mathcal{B}, \mathcal{C} \mid \mathcal{S}, \mathcal{R}) = \underbrace{p_\theta(\mathcal{B} \mid \mathcal{S}, \mathcal{R})}_{\text{Backbone Generator}} \cdot \underbrace{p_\tau(\mathcal{C} \mid \mathcal{B}, \mathcal{S}, \mathcal{R})}_{\text{Side-Chain Packer}}. \tag{1}$$

Here, p_θ is the backbone generator conditioned on $(\mathcal{S}, \mathcal{R})$, whose parameters are denoted as θ. p_τ is the side-chain packer conditioned on $(\mathcal{B}, \mathcal{S}, \mathcal{R})$, whose parameters are denoted as τ.

3.2 Multi-Stage Conformation Generation

As illustrated in Figure 1, MuCO2 implements a multi-stage pipeline for generating cyclic peptide conformations. Similar to MuCO (Wang et al., 2026), it first generates cyclic backbones, then packs

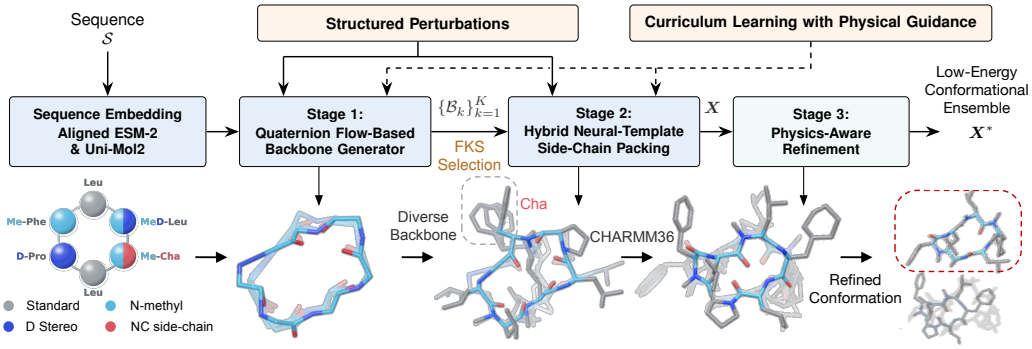


Figure 1: The scheme of MUCO2. MUCO2 reuses the physical-aware refinement module (green) of MUCO and improves the remaining modules (blue) in the multi-stage pipeline. The yellow frames indicate the new information and algorithms applied by MUCO2.

side-chains based on the generated backbones, and finally refines the all-atom conformations via force-field optimization. Notably, to make the pipeline applicable to noncanonical cyclic peptides, we improve the key modules as follows. A detailed model architecture is provided in Appendix B.

Embedding of Sequence and Perturbations. Following MUCO (Wang et al., 2026), we apply ESM-2 (Lin et al., 2023) followed by a trainable sequence-to-trunk projection to embed sequence $\mathcal{S}$, denoted as $\boldsymbol{H}_{\mathcal{S}} = \text{Seq2Trunk}(\text{ESM-2}(\mathcal{S})) = \{\boldsymbol{h}_i, \boldsymbol{h}_{ij}\}_{i,j=1}^L$, where $\boldsymbol{h}_i$ and $\boldsymbol{h}_{ij}$ represent the residue and residue-pair embeddings, respectively. To accommodate ncAAs across open-ended chemical space, we employ pre-trained Uni-Mol2 (Ji et al., 2024a) on residue graphs and corresponding template conformations to obtain atomic embeddings: $\boldsymbol{Z}_{\mathcal{R}} = \text{Uni-Mol2}(\{\mathcal{G}_i, \boldsymbol{Q}_i\}_{i=1}^L) = \{\boldsymbol{z}_m, \boldsymbol{z}_{mn}\}_{m,n=1}^N$, where $\boldsymbol{z}_m$ and $\boldsymbol{z}_{mn}$ are representations of the m-th atom and the atom-pair (m, n), respectively.

To reconcile discrepancies in granularity and latent space, a lightweight adapter g_ϕ maps $\boldsymbol{Z}_{\mathcal{R}}$ into residue representations: $g_\phi(\boldsymbol{Z}_{\mathcal{R}}) = \{\boldsymbol{h}_{i,\phi}, \boldsymbol{h}_{ij,\phi}\}_{i,j=1}^L$. Specifically, g_ϕ aggregates atom embeddings $\{\boldsymbol{z}_m\}_{m=1}^N$ via residue-query cross-attention, followed by residue-level updates and an output projection, to obtain $\boldsymbol{h}_{i,\phi}$. In parallel, g_ϕ aggregates atom-pair embeddings $\{\boldsymbol{z}_{mn}\}_{m,n=1}^N$ over each residue pair and passes them through an encoder to construct $\boldsymbol{h}_{ij,\phi}$. Given a set of sequences, we train g_ϕ by penalizing the alignment errors and encouraging the cosine similarities:

$$\phi^* = \arg\min_\phi \mathbb{E}_{\mathcal{S}}\Big[\sum\nolimits_{i=1}^L \Big(\|\boldsymbol{h}_{i,\phi} - \boldsymbol{h}_i\|_2^2 - \lambda \cos(\boldsymbol{h}_{i,\phi}, \boldsymbol{h}_i)\Big) + \beta \sum\nolimits_{i,j=1}^L \|\boldsymbol{h}_{ij,\phi} - \boldsymbol{h}_{ij}\|_2^2\Big], \quad (2)$$

where $\lambda \geq 0$ balances MSE and cosine similarity, and $\beta \geq 0$ balances the residue-level term and the pairwise one. Given the trained adapter g_{ϕ^*}, we apply a hybrid embedding strategy to make MUCO2 applicable for canonical peptides and those with complex N and side-chain modifications:

$$\{\hat{\boldsymbol{h}}_i, \hat{\boldsymbol{h}}_{ij}\}_{i,j=1}^L = \boldsymbol{H}_{\mathcal{S}} \text{ if all residues are canonical, and } g_{\phi^*}(\boldsymbol{Z}_{\mathcal{R}}), \text{ otherwise.} \quad (3)$$

Besides side-chain modifications, we further embed the remaining indicators in $\mathcal{R}$ and, accordingly, inject the information of α-backbone modifications and stereochemical transformations as follows:

$$\tilde{\boldsymbol{h}}_i = \hat{\boldsymbol{h}}_i + \boldsymbol{e}_s^{\text{Me}}(m_i) + \boldsymbol{e}_s^{St}(c_i), \qquad \tilde{\boldsymbol{h}}_{ij} = \hat{\boldsymbol{h}}_{ij} + \boldsymbol{e}_p^{\text{Me}}(m_i, m_j) + \boldsymbol{e}_p^{St}(c_i c_j), \quad (4)$$

where $\boldsymbol{e}_s^{\text{Me}}(m_i)$ and $\boldsymbol{e}_p^{\text{Me}}(m_i, m_j)$ are methylation embeddings, and $\boldsymbol{e}_s^{\text{St}}(c_i)$ and $\boldsymbol{e}_p^{\text{St}}(c_i c_j)$[2] are chirality embeddings. Equation 4 provides us with perturbation-aware peptide representations, denoted as $\widetilde{\boldsymbol{H}} = \{\tilde{\boldsymbol{h}}_i, \tilde{\boldsymbol{h}}_{ij}\}_{i,j=1}^L$. Conditioned on these unified representations, our hierarchical generative model generates cyclic peptide conformations accordingly.

Quaternion Flow-Based Conditional Backbone Generation. Given $\widetilde{\boldsymbol{H}}$, we apply a flow-based model to generate cyclic backbones. For the backbone generator, MUCO2 replaces the Riemannian SE(3) flow matching (Bose et al., 2024) used in MUCO with the Quaternion Flow (QFlow) matching (Yue et al., 2025). In particular, MUCO represents each local frame $\boldsymbol{T}_i$ as a tuple $(\boldsymbol{R}_i, \boldsymbol{t}_i)$, where $\boldsymbol{R}_i \in \text{SO}(3)$ is a rotation matrix and $\boldsymbol{t}_i \in \mathbb{R}^3$ is a translation vector. Instead, MUCO2 represents

[2]By parameterizing stereochemistry as $c_i \in \{+1, -1\}$, the scalar product $c_i c_j \in \{+1, -1\}$ naturally functions as a relational parity indicator (matching v.s. inverted chirality), so we set the input of $\boldsymbol{e}_p^{\text{St}}$ as $c_i c_j$.

each rotation by a unit quaternion $\boldsymbol{q}_i = [\cos\frac{\gamma_i}{2}, \boldsymbol{u}_i \sin\frac{\gamma_i}{2}] \in \mathbb{S}^3$, where $\boldsymbol{u}_i \in \mathbb{S}^2$ denotes the rotation axis and γ_i is the rotation angle. Let $\mathcal{B}_t = \{(\boldsymbol{q}_{i,t}, \boldsymbol{t}_{i,t}, \psi_{i,t})\}_{i=1}^{L}$ denote the intermediate backbone state at time $t \in [0,1]$. Conditioned on $\widetilde{\boldsymbol{H}}$, we learn the backbone generator as a time-dependent velocity field $v_\theta(\mathcal{B}_t \mid t, \widetilde{\boldsymbol{H}})$, which corresponds to solving the following flow matching problem:

$$\theta^* = \arg\min_\theta \mathbb{E}_{t,\mathcal{B}_0,\mathcal{B}_1}\left[\|v_\theta(\mathcal{B}_t \mid t, \widetilde{\boldsymbol{H}}) - u(\mathcal{B}_0, \mathcal{B}_1, t)\|_2^2\right], \tag{5}$$

where $\mathcal{B}_0$ denotes initial local frame set sampled from random noise, and $\mathcal{B}_1$ denotes the local frame set of a cyclic backbone sampled from the training data. The interpolation between $\mathcal{B}_0$ and $\mathcal{B}_1$ at time t results in the vector field $u(\mathcal{B}_0, \mathcal{B}_1, t)$ defined on $\mathbb{R}^3 \times \mathbb{S}^3 \times \mathbb{T}$. Compared to Riemannian $\mathrm{SE}(3)$ flow matching, which relies on computationally demanding geodesic projections and logarithmic maps on $\mathrm{SO}(3)$, QFlow parameterizes rotations on $\mathbb{S}^3$ and integrates trajectories via quaternion-based spherical linear interpolation (SLERP) (Yue et al., 2025). As a result, instead of using the Riemannian metric (Wang et al., 2026), we can apply the MSE loss in equation 5 and avoid numerical singularities during training and inference. In the following experiments, we will show that this improvement helps enhance stability for closed-loop geometries and accelerate backbone sampling.

Hybrid Neural-Template Side-Chain Packing. Given the generated backbone $\mathcal{B}$ together with the representations $\widetilde{\boldsymbol{H}}$, MuCO2 further packs side-chains to obtain full-atom conformations of cyclic peptides. To handle open-set side-chain chemistry, we apply a hybrid packing strategy combining neural packing (Lee & Kim, 2025) with template attachment.

- For canonical residues in $\mathcal{I}_c$, we train a flow-based packer to sample side-chain torsions. The model corresponds to a velocity field conditioning on backbones and perturbation-aware representations, denoted as $v_\tau(\boldsymbol{\chi}_t, \mid t, \mathcal{B}, \widetilde{\boldsymbol{H}})$, and we train it by torsional flow matching (Lee & Kim, 2025):

$$\tau^* = \arg\min_\tau \mathbb{E}_{t,\{\boldsymbol{\chi}_{i,0}\}_{i\in\mathcal{I}_c},\{\boldsymbol{\chi}_{i,1}\}_{i\in\mathcal{I}_c}}\left[\sum\nolimits_{i\in\mathcal{I}_c} \|v_\tau(\boldsymbol{\chi}_{i,t} \mid t, \mathcal{B}, \widetilde{\boldsymbol{H}}) - u(\boldsymbol{\chi}_{i,0}, \boldsymbol{\chi}_{i,1}, t)\|_2^2\right]. \tag{6}$$

  Here, $\boldsymbol{\chi}_{i,1}$ denotes the torsions of the i-th canonical residue sampled from the training data and $\boldsymbol{\chi}_{i,0}$ denotes the random torsions sampled from noise. $u(\boldsymbol{\chi}_{i,0}, \boldsymbol{\chi}_{i,1}, t)$ denotes interpolation between them at time t, along the path on $\mathbb{T}^4$. This model naturally internalizes the altered ω-torsional distributions and local steric crowding without requiring specialized auxiliary loss functions.
- For each noncanonical residue in $\mathcal{I}_n$, whose complex N-modifications and side-chain cannot be parameterized by torsions, we rigidly superpose its template conformation $\boldsymbol{Q}_i$ onto the generated backbone frame $\boldsymbol{T}_i$ via optimal backbone alignment, leads to an aligned conformation $\boldsymbol{C}_i$.

Following MuCO (Wang et al., 2026), given the backbone $\mathcal{B}$ and side-chains $\mathcal{C} = \{\boldsymbol{\chi}_i\}_{i\in\mathcal{I}_c} \cup \{\boldsymbol{C}_i\}_{i\in\mathcal{I}_n}$, we obtain a conformation via a forward-kinematics operator $\widetilde{\boldsymbol{X}} = \mathrm{FK}(\mathcal{B}, \mathcal{C})$ and refine its force-field by CHARMM36 (Huang & MacKerell Jr, 2013), i.e., $\boldsymbol{X}^\star = \arg\min_{\boldsymbol{X}} E_R(\boldsymbol{X}) + \frac{w}{2}\|\boldsymbol{X} - \widetilde{\boldsymbol{X}}\|_F^2$, where E_R denotes the force-field energy under perturbation specification R, w is the positional restraint weight, and $X^\star$ is the final refined output.

4 Improvements in Learning and Inference

4.1 Curriculum Learning Based on Physics-Aware Data Augmentation

For each cyclic peptide, its perturbation set $\mathcal{R} = \{(m_i, c_i, \eta_i, \mathcal{G}_i, \boldsymbol{Q}_i)\}$ indicates the complexity of its conformation: $k_{\mathrm{Me}} = \sum_i m_i$ denotes the number of N-methylated residues, $k_{\mathrm{D}} = \sum_i \eta_i \mathbb{I}[c_i = -1]$ denotes the number of D-residues, $k_\chi = \sum_i \eta_i$ denote the number of chiral centers, and $b_{\mathrm{D}} = \min(k_{\mathrm{D}}, k_\chi - k_{\mathrm{D}})$ denotes the discrepancy from pure enantiomers. Accordingly, we denote canonical cyclic peptide dataset, e.g., filtered CPSea (Yang et al., 2026; Wang et al., 2026)), as $\mathcal{D}_0$ ($k_{\mathrm{Me}} = b_{\mathrm{D}} = 0$), and we can impose single-site modifications on $\mathcal{D}_0$ and simulate noncanonical conformations by CHARMM36 (Huang & MacKerell Jr, 2013), leading to single-side modification dataset $\mathcal{D}_1$ ($k_{\mathrm{Me}} = 1$ and/or $b_{\mathrm{D}} = 1$). However, conformations with large $\{k_{\mathrm{Me}}, b_{\mathrm{D}}\}$ are very scarce, and simulating such conformations directly by CHARMM36 suffers a high failure rate. To avoid mode collapse, we train MuCO2 in a curriculum learning (Bengio et al., 2009) framework, outlined in Algorithm 1.

Initialization: Canonical Prior Learning. We first pretrain the backbone generator p_θ and side-chain packer p_τ on $\mathcal{D}_0$ and achieve a canonical peptide generation method like MuCO.

Round 1: Generalization for Single-Site Modifications. Given the pretrained model, we first enhance its generalizability to single-site modifications, e.g., isolated N-methylation ($k_{\mathrm{Me}} = 1$) and single D-residue inversion ($k_{\mathrm{D}} = 1$). For each conformation $\boldsymbol{X} \in \mathcal{D}_0 \cup \mathcal{D}_1$, whose center is denoted as $\bar{\boldsymbol{x}}$, we leverage its physical symmetry and augment it via random reflection: $M(\boldsymbol{X}) = (\boldsymbol{X} - \mathbf{1}_N \bar{\boldsymbol{x}}^\top)\boldsymbol{R} + \mathbf{1}_N \bar{\boldsymbol{x}}^\top$, where $\boldsymbol{R}$ is a random reflection matrix, and $M(\boldsymbol{X})$ is the mirror enantiomer of $\boldsymbol{X}$ with inverted chirality labels ($c_i \leftarrow -c_i$). This augmentation method doubles data size without time-consuming physical simulation. Fine-tuning the pretrained model on the augmented data achieves a warmup for more complex multi-side modifications.

Algorithm 1 Curriculum Learning Framework

Require: An initial dataset $\mathcal{D}_0 \cup \mathcal{D}_1$

1: **Initialization:**
$\theta_0, \tau_0 \leftarrow$ Pretrain($\mathcal{D}_0$) by equations 5 & 6.
2: **Round-1 Generalization:**
Augmentation: $\mathcal{D}^{(1)} = M(\mathcal{D}_0 \cup \mathcal{D}_1) \cup \mathcal{D}_0 \cup \mathcal{D}_1$.
$\theta_1, \tau_1 \leftarrow$ Finetune($\mathcal{D}^{(1)}$) by equations 5 & 6.
3: **Round 2-3 Generalization:**
For $r = 2, 3$ **do**
Generate data: $\mathcal{D} \sim p_{\theta_{r-1}} \cdot p_{\tau_{r-1}}$
$\mathcal{D} =$ CHARMM36($\mathcal{D}$) with $k_{\mathrm{Me}}, b_{\mathrm{D}} \leq 2r-1$.
Augmentation: $\mathcal{D}^{(r)} = M(\mathcal{D}) \cup \mathcal{D} \cup \mathcal{D}^{(r-1)}$.
$\theta_r, \tau_r \leftarrow$ Finetune($\mathcal{D}^{(r)}$) by equations 5 & 6
4: **Alignment:**
$\phi^* \leftarrow$ Train($\mathcal{D}_0$) by equation 2.
5: **return** Model θ^*, τ^*, $\phi^* \leftarrow \theta_3$, τ_3, ϕ^*

Rounds 2-3: Generalization for Multi-Site Modifications. Given the noncanonical conformations generated by the fine-tuned model, we add more modifications to them. Compared to the samples in $\mathcal{D}_0 \cup \mathcal{D}_1$, these modified conformations provide good initializations for CHARMM36, and we can refine them with a high success rate. Training on the refined conformations leads to a better model. Repeating the above steps in two rounds, we obtain conformations with $k_{\mathrm{Me}}, b_{\mathrm{D}} \leq 5$ and improve the model generalizability.

4.2 Improving Diversity via Feynman-Kac Steering Exploration

Given a backbone generator and a side-chain packer, MuCO samples K backbones and expand them to $K \times M$ conformations via sampling M rotamers for each backbone. Beyond the greedy tree expansion used in MuCO, we introduce a *Structure-Guided Feynman-Kac Steering* (FKS) strategy. In Stage-1, the model initially proposes an overcomplete pool of K candidate backbones $\boldsymbol{\mathcal{B}} = \{\mathcal{B}_k\}_{k=1}^{K} \sim p_{\theta^*}(\cdot \mid \mathcal{S}, \mathcal{R})$. To select a non-redundant subset $\boldsymbol{\mathcal{B}}'$ ($|\boldsymbol{\mathcal{B}}'| = K_0 \ll K$) before side-chain packing, **FKS models sequential backbone selection as an interacting particle system governed by a temperature-modulated crowding potential.** Each backbone's conditional transition probability is dictated by its kernelized structural crowding relative to already chosen backbones under $SO(3)$-aligned Kabsch RMSD (Kabsch, 1976). This provides gentle repulsive steering that drives backbones toward distinct conformational basins while retaining physical plausibility. Implementation details are provided in Algorithm 3 and Appendix C.

During training, FKS breaks the sample degeneracy inherent in maximum-likelihood generative flows by repelling trajectories from already colonized modes, providing high-quality, diverse backbones for training the generator across successive stages and mitigating model collapse. During inference, FKS improves the hit rate for broad potential wells and avoids severe steric clashes caused by noncanonical side-chain attachment, providing good initial conformations for force-field relaxation.

5 Experiments

5.1 Experimental Setup

Datasets. We construct our evaluation pipeline across canonical, backbone-mutated, and open-set noncanonical cyclic peptides. Full data processing details are given in Appendix D.

1. **CPSea**: To verify compatibility with canonical cases, we adopt the CPSea dataset (Yang et al., 2026) with four splits (Wang et al., 2026): CPTrans, CPBind, CPCore, and CPSea-PDB.
2. **CPSea-MeD**: To benchmark backbone generalization, we randomly select 12% of peptides from CPSea and introduce chemically feasible N-methylation, D-stereochemistry, or their combinations, whose conformations are simulated by CHARMM36 (Huang & MacKerell Jr, 2013).
3. **HighFold2-Compatible**: To benchmark noncanonical side-chain generation, we collect cyclic peptides from the HighFold2 benchmarks (Zhu et al., 2025) and CycPeptMPDB (Li et al., 2023), respectively, whose residues fall within HighFold2's vocabulary and exclude N-methylation.

4. **CycPeptMPDB-MeD**: To isolate the effect of α-backbone modifications, we filter the cyclic peptides without any side-chain modifications from CycPeptMPDB (Li et al., 2023).
5. **CycPeptMPDB-Mixed**: To assess model performance under challenging cases (beyond HighFold2's scope), we filter cyclic peptides from CycPeptMPDB featuring compounding modifications: D-amino acids, noncanonical side-chains, N-methylation, and peptoid-like N-alkylations.

Baselines. We benchmark MUCO2 against state-of-the-art methods. Based on the chemical capabilities shown in Table 1, we compare MUCO2 with MUCO (Wang et al., 2026), HighFold2 (Zhu et al., 2025), HighFold-MeD (Cao et al., 2025b), and HighFold-MeD2 (Cao et al., 2026) on CPSea and CPSea-MeD. For noncanonical side-chain transfer, we evaluate MUCO2 against HighFold2 and HighFold3 (Cao et al., 2025a) on HighFold2-Compatible, and against HighFold3 on CycPeptMPDB-Mixed. For a fair comparison, all the methods are followed by CHARMM36 for force-field relaxation.

For discriminative predictors (HighFold2 and HighFold-MeD), we run inference across the five standard model checkpoints and select the candidate with the highest predicted confidence (pLDDT) as the deterministic result. For generative methods (HighFold3, HighFold-MeD2, MUCO and MUCO2), we consider generating a single conformation with a fixed seed and sampling 10 conformations independently as an ensemble. Comprehensive implementation details are given in Appendix B.

Evaluation Metrics. We evaluate different methods across five dimensions: (i) **Success Rate** (%) measures the proportion of candidates that achieve valid covalent cyclization and yield finite potential energy after relaxation. (ii) **Vacuum Energy** (E_{vac}, kJ/mol) evaluates post-relaxation intramolecular strain, torsional preferences, and non-bonded contacts under the CHARMM36 force field. (iii) **Water Score** (E_{water}, kJ/mol) provides an implicit-solvent energetic score derived from the OBC2/ACE Generalized Born model (Best et al., 2012). (iv) **C_α-RMSD** (Å) assesses backbone folding accuracy against reference conformations.[3] (v) **Chirality** (%) quantifies the preservation of target α-stereochemistry across all non-Glycine residues, computed as the percentage of chirality preservation during force-field relaxation.[4] Metrics are reported as deterministic conformations (**@1**) and best-of-10 ensembles (**@10**). More details are shown in Appendix E.

5.2 Main Results

Stereochemical preservation on canonical cyclic peptides. As shown in Table 2, MUCO2 retains strong conformational sampling performance while substantially improving stereochemical preservation over MUCO. Chirality increases from 61.90% to 99.86%, with 100.00% success@10 and lower vacuum energies and water scores. The improvement also extends to reference fidelity: C_α-RMSD@10 decreases from 3.23Å to 3.19Å. Because our reference conformations come from MD simulations, a lower RMSD suggests that MUCO2 can sample conformations close to the true equilibrium ensemble, underscoring its strong multi-basin exploration capacity.

Generalization on MeD cyclic peptides. On the two MeD benchmarks featuring N-methylations and D-amino acids, MUCO2 consistently achieves the highest cyclization success rates and lowest energies. Especially, on the out-of-distribution (OOD) CycPeptMPDB-MeD benchmark, the leading baseline HighFold-MeD2 suffers a catastrophic stereochemical collapse, while MUCO2 shows superior OOD generalizability, achieving the highest chirality (89.21%) alongside 100.00% success@10 and the lowest C_α-RMSD. In our opinion, this is because the curriculum learning algorithm of MUCO2 reduces the risk of producing unphysical, strain-relieved chiral inversions.

Broad chemical scope and stereochemical integrity on open-set modifications. Compared to HighFold2 and HighFold3, MUCO2 generalizes to a substantially broader noncanonical chemical space. On the CycPeptMPDB-Mixed benchmark featuring challenging combinatorial perturbations, MUCO2 achieves physical energies and structural accuracy comparable to, and slightly outperforming, the AlphaFold3-based HighFold3. Notably, HighFold3 suffers a severe stereochemical breakdown, with its chirality dropping to 49.27%, while MUCO2 preserves 98.41% chirality, underscoring the need to embed stereochemical information explicitly.

[3] For datasets based on CycPeptMPDB, the reference conformations are sampled from explicit-solvent MD ensembles provided by CycPeptMPDB-4D v3 (Liu et al., 2026).

[4] During force-field relaxation, the force field flips unstable chiral centers when the generated conformations suffer from severe local steric clashes or strained backbone geometries. Therefore, post-relaxation chirality preservation reflects whether generated conformations reside in physically stable potential energy wells.

Table 2: Conformation generation performance on canonical and noncanonical cyclic peptides. For each dataset, we bold the best result and underline the second-best one.

Dataset	Method	Success (%)↑		E_{vac}(kJ/mol)↓		E_{water}(kJ/mol)↓		C_α-RMSD↓		Chirality
		@1	@10	@1	@10	@1	@10	@1	@10	(%)↑
CPSea	GT (MD Ref.)	100.00	–	-15.78	–	-835.12	–	0.00	–	100.00
	HighFold2	**98.87**	–	111.35	–	-611.80	–	3.82	–	96.33
	HighFold-MeD	94.93	–	103.83	–	-570.91	–	3.68	–	97.73
	HighFold-MeD2	92.16	97.10	66.87	0.17	-588.25	-704.13	3.86	3.41	85.27
	MuCO	94.03	99.98	30.30	-20.04	-684.79	-881.56	3.71	3.23	61.90
	MuCO2 (ours)	96.25	**100.00**	**20.81**	**-32.13**	**-700.86**	**-902.06**	**3.67**	**3.19**	**99.86**
CPSea-MeD	GT (MD Ref.)	100.00	–	441.03	–	-595.70	–	0.00	–	100.00
	HighFold-MeD	22.63	–	770.47	–	-177.23	–	3.85	–	96.77
	HighFold-MeD2	23.28	69.67	770.69	571.08	-174.56	-355.11	3.72	3.23	80.67
	MuCO2 (ours)	**31.39**	**84.10**	**761.37**	**503.94**	**-185.53**	**-528.44**	**3.25**	**2.62**	**98.90**
CycPeptMPDB -MeD	GT (MD Ref.)	100.00	–	1195.62	–	1068.03	–	0.00	–	100.00
	HighFold-MeD	65.58	–	1459.21	–	1383.26	–	1.39	–	87.89
	HighFold-MeD2	17.63	56.81	1429.96	1390.18	1315.19	1228.22	1.59	1.12	24.18
	MuCO2 (ours)	**67.90**	**100.00**	**1421.32**	**1341.44**	**1307.47**	**1220.40**	**1.35**	**1.07**	**89.21**
HighFold2 -Compatible	GT (NMR/MD Ref.)	100.00	–	-583.52	–	-2930.65	–	0.00	–	100.00
	HighFold2	**100.00**	–	**-1264.12**	–	-3292.24	–	**2.81**	–	98.77
	HighFold3	**100.00**	**100.00**	-1196.81	-1233.43	-3272.89	-3342.69	2.85	**2.36**	87.79
	MuCO2 (ours)	90.00	**100.00**	-1205.08	**-1235.02**	**-3359.07**	**-3424.84**	3.76	3.14	**98.81**
CycPeptMPDB -Mixed	GT (MD Ref.)	100.00	–	760.68	–	623.49	–	0.00	–	100.00
	HighFold3	**100.00**	**100.00**	648.67	521.11	561.51	**443.09**	1.02	0.64	49.27
	MuCO2 (ours)	**100.00**	**100.00**	**643.24**	**518.49**	**548.33**	447.92	**0.96**	**0.56**	**98.41**

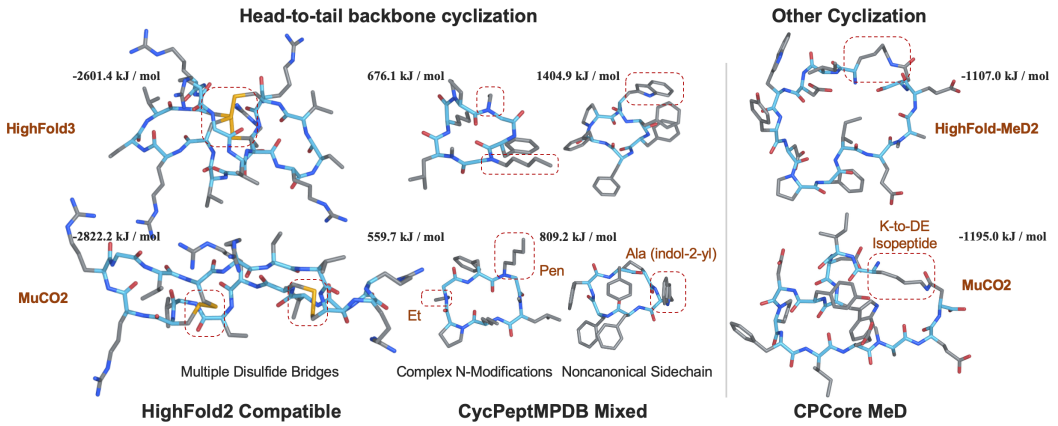


Figure 2: Visual comparisons for various methods in challenging generation cases.

As illustrated in Figure 2, MuCO2 is capable of capturing diverse cyclization modes and multi-type perturbations. Remarkably, it can form intricate internal linkages, such as multiple disulfide bridges, while achieving lowest energies among baselines. More numerical and visual results are provided in Appendix F and Appendix H.

5.3 Ablation Studies

We conduct ablation studies to analyze the contributions of different modules in MuCO2. Some representative results are shown in Figure 3, and full numerical results are provided in Appendix G.

Quaternion Backbone Flow. Compared to SE(3)-Flow (Bose et al., 2024), QFlow improves chirality preservation across all five benchmarks, most notably boosting CycPeptMPDB-MeD from 78.37% to 89.21% alongside a sharp surge in success@1 (42.13% to 67.90%). On CPSea-MeD, QFlow reduces C_α-RMSD@10 from 3.18Å to 2.62Å.

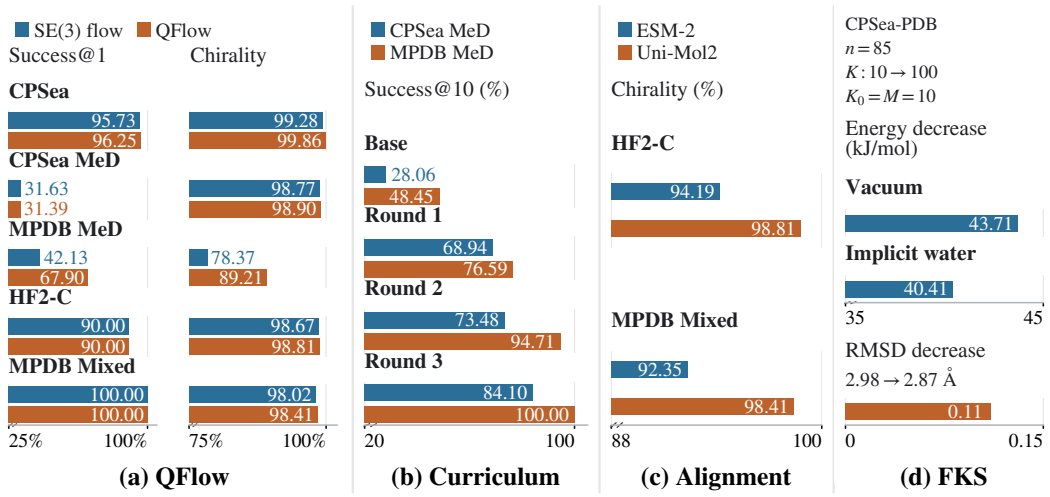


Figure 3: Ablation studies of core modules in MuCO2.

Curriculum Learning. The progressive curriculum is crucial for overcoming the steep kinetic and steric barriers introduced by N-methylation and D-amino acids. Over three training rounds, success@10 increases monotonically from 28.06% to 84.10% on CPSea-MeD and from 48.45% to 100.00% on CycPeptMPDB-MeD, respectively. In parallel, chirality preservation reaches 98.90% and 89.21%, respectively, on the two datasets, accompanied by massive drops in effective energy (e.g., water score@10 decreases by over 1200kJ/mol on CPSea-MeD).

Molecular Representation Alignment. Uni-Mol2 alignment substantially outperforms sequence-only ESM-2 representations across all noncanonical side-chain and complex backbone N-modifications evaluations. It enhances chirality preservation by 4.62% on HighFold2-Compatible and 6.06% on CycPeptMPDB-Mixed, while yielding lower physical relaxation energies (water score@10 drops by 105.94kJ/mol and 56.47kJ/mol, respectively). Backbone accuracy also consistently improves, bringing C_α-RMSD@10 down to 0.56Å on the mixed benchmark.

Structure-guided Feynman-Kac Steering. Comparing FKS against naive tree expansion on CPSea-PDB demonstrates the clear benefit of steered exploration under an identical budget of 100 full-atom candidates. In particular, FKS suppresses sample degeneracy and discovers deeper energy funnels: it drives down vacuum energy by 43.71kJ/mol and implicit water score by 40.41kJ/mol, while simultaneously reducing C_α-RMSD and maintaining near-perfect chirality (>99.8%).

6 Conclusion and Future Work

In this work, we introduced MuCO2, a multi-stage conformation generation framework tailored for noncanonical cyclic peptides across multiple chemical perturbations. Across canonical benchmarks, backbone mutants, and external noncanonical datasets, MuCO2 achieves superior or at least competitive performance compared to existing methods, establishing an efficient foundation for library-scale macrocycle exploration.

Limitations and Future Work. Despite these advances, several limitations remain. Firstly, current MuCO2 focuses on standard α-backbones across the three dominant perturbation types, leaving backbone extensions (β/γ-amino acids) and non-amide linkages (e.g., depsipeptides) for future exploration; moreover, while backbone N-methylation is explicitly parameterized, more complex N-substitutions (e.g., bulky peptoids) currently rely on rigid template attachment rather than specialized generative flow priors. Secondly, because cyclic closure is not enforced as a hard kinematic constraint during generative flow integration, highly strained rings or intricate polycyclic architectures may occasionally suffer from closure failures during downstream force-field relaxation. Finally, due to the severe scarcity of high-resolution crystal structures for noncanonical macrocycles, evaluations primarily rely on molecular dynamics simulations. To overcome limitations describe above, we plan to further improve MuCO2 and collaborate with wet-lab teams in the future.

AI USE STATEMENT

In this work, we used generative AI tools to assist with and polish the writing of the manuscript (specifically, grammar checking, phrasing refinement, and improving clarity). Additionally, generative AI was utilized to assist with code organization and formatting; all code changes were manually inspected, executed, and verified by the authors. We have not used generative AI tools for research ideation or execution, mathematical derivations or proofs, drafting primary scientific sections from scratch, literature retrieval or discovery, or generating synthetic datasets. All AI-assisted text revisions were thoroughly reviewed, edited, and verified for scientific accuracy by the authors. The authors take full and sole responsibility for the final content of this work, including all text, code, claims, methodologies, and results presented herein.

ETHICS STATEMENT

This work introduces a computational, multi-stage generative framework for macrocyclic peptide conformation sampling to facilitate peptide drug discovery and structural biology research. The study is conducted entirely *in silico* and utilizes exclusively publicly available, anonymized macromolecular and structural datasets (CPSea, CycPeptMPDB, and the Protein Data Bank). The research does not involve human participants, private personal data, or vertebrate animals, and thus requires no Institutional Review Board (IRB) approval. Furthermore, the methodology is designed as a foundational tool for biophysical conformation modeling and does not target biological toxins, dual-use hazardous agents, or pathogen engineering. The authors foresee no direct negative societal, discriminatory, or biosecurity implications resulting from this research.

REPRODUCIBILITY STATEMENT

To ensure the full reproducibility of our models, algorithms, and experimental results, we have made the following efforts:

- **Source Code and Configuration:** A complete, self-contained, and anonymized implementation of MUCO2, including the Quaternion Flow (QFlow) backbone generator, Structure-Guided Feynman-Kac Steering (FKS), hybrid side-chain assembly, and custom CHARMM36 topology parameter stream files, is provided in the supplementary material.
- **Training and Hyperparameter Details:** Detailed network architectures, layer dimensions, optimization schedules, loss functions, and curriculum training protocols are comprehensively documented in Appendix B and Table 4.
- **Data Processing and Simulation Protocols:** Complete data filtering pipelines are provided in Appendix D, noncanonical residue mappings, and explicit-solvent molecular dynamics protocols (GROMACS parameters, solvent box setups, force-field specifications, and restraint configurations) are described in Appendix B.1.
- **Experimental Evaluation and Stochasticity:** We provide explicit random seed configurations for backbone generation and side-chain rotamer sampling. While minor numerical variations can occur during downstream force-field relaxation due to hardware-dependent floating-point operations and numerical minimization tolerances, our provided setup and deterministic seed controls ensure that the conformational ensembles, statistical energy distributions, and comparative benchmark metrics reported in this paper can be reliably reproduced.

REFERENCES

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

Nomindari Bayaraa, Maxim Secor, Marc L Descoteaux, and Yu-Shan Lin. Fast generation of simulation-quality structural ensembles of mixed-chirality cyclic peptides via diffusion models. *Journal of Chemical Theory and Computation*, 22(6):3103, 2026.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

Robert B. Best, Xiao Zhu, Jihyun Shim, Pedro E. M. Lopes, Jeetain Mittal, Michael Feig, and Alexander D. MacKerell, Jr. Optimization of the additive CHARMM all-atom protein force field targeting improved sampling of the backbone ϕ, ψ and side-chain χ_1 and χ_2 dihedral angles. *Journal of Chemical Theory and Computation*, 8(9):3257–3273, 2012. doi: 10.1021/ct300400x.

Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks, Chenghao Liu, Andrei Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong. Se (3)-stochastic flow matching for protein backbone generation. In *International Conference on Learning Representations*, volume 2024, pp. 22590–22621, 2024.

Sen Cao, Cheng Zhu, Qingyi Mao, Jingjing Guo, Ning Zhu, and Hongliang Duan. Accurate structure prediction of cyclic peptides containing unnatural amino acids using HighFold3. *Briefings in Bioinformatics*, 26(5):bbaf488, 2025a.

Zhigang Cao, Sen Cao, Linghong Wang, Zhiguo Wang, Qingyi Mao, Jingjing Guo, and Hongliang Duan. HighFold-MeD: a rosetta distillation model to accelerate structure prediction of cyclic peptides with backbone N-methylation and D-amino acids. *Journal of Cheminformatics*, 17(1): 167, 2025b.

Zhigang Cao, Wentong Wang, Sen Cao, Chengyun Zhang, Jingjing Guo, and Hongliang Duan. HighFold-MeD2: An enhanced Boltz-2 model for accurate structure prediction of N-methylated and D-amino acid cyclic peptides. *Journal of Chemical Information and Modeling*, 66(10): 6057–6066, 2026.

Jovan Damjanovic, Jiayuan Miao, He Huang, and Yu-Shan Lin. Elucidating solution structures of cyclic peptides using molecular dynamics simulations. *Chemical reviews*, 121(4):2292, 2021.

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

David Gfeller, Olivier Michielin, and Vincent Zoete. Swisssidechain: a molecular and structural database of non-natural sidechains. *Nucleic acids research*, 41(D1):D327–D332, 2012.

Colin A Grambow, Hayley Weir, Nathaniel L Diamant, Gabriele Scalia, Tommaso Biancalani, and Kangway V Chuang. Accurate and efficient structural ensemble generation of macrocyclic peptides using internal coordinate diffusion. *arXiv preprint arXiv:2305.19800*, 2023.

Colin A Grambow, Hayley Weir, Christian N Cunningham, Tommaso Biancalani, and Kangway V Chuang. CREMP: Conformer-rotamer ensembles of macrocyclic peptides for machine learning. *Scientific Data*, 11(1):859, 2024.

Jennifer L Hickey, Dan Sindhikara, Susan L Zultanski, and Danielle M Schultz. Beyond 20 in the 21st century: prospects and challenges of non-canonical amino acids in peptide drug discovery. *ACS Medicinal Chemistry Letters*, 14(5):557, 2023.

Scott A Hollingsworth and Ron O Dror. Molecular dynamics simulation for all. *Neuron*, 99(6): 1129–1143, 2018.

Jing Huang and Alexander D MacKerell Jr. CHARMM36 all-atom additive protein force field: Validation based on comparison to NMR data. *Journal of computational chemistry*, 34(25): 2135–2145, 2013.

Tiffani Hui, Marc L Descoteaux, Jiayuan Miao, and Yu-Shan Lin. Training neural network models using molecular dynamics simulation results to efficiently predict cyclic hexapeptide structural ensembles. *Journal of Chemical Theory and Computation*, 19(14):4757, 2023.

Xiaohong Ji, Zhen Wang, Zhifeng Gao, Hang Zheng, Linfeng Zhang, Guolin Ke, et al. Uni-Mol2: Exploring molecular pretraining model at scale. *arXiv preprint arXiv:2406.14969*, 2024a.

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

Dapeng Jiang, Xiangzhe Kong, Jiaqi Han, Mingyu Li, Rui Jiao, Wenbing Huang, Stefano Ermon, Jianzhu Ma, and Yang Liu. Zero-shot cyclic peptide design via composable geometric constraints. In *Forty-second International Conference on Machine Learning*, 2025.

Xiaoshu Jing and Kang Jin. A gold mine for drug discovery: Strategies to develop cyclic peptides into therapies. *Medicinal research reviews*, 40(2):753–810, 2020.

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

Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of Crystallography*, 32(5):922–923, 1976.

Takayuki Katoh and Hiroaki Suga. In vitro genetic code reprogramming for the expansion of usable noncanonical amino acids. *Annual Review of Biochemistry*, 91(1):221–243, 2022.

Andy Chi-Lung Lee, Janelle Louise Harris, Kum Kum Khanna, and Ji-Hong Hong. A comprehensive review on current advances in peptide drug development and design. *International journal of molecular sciences*, 20(10):2383, 2019.

Jin Sub Lee and Philip M Kim. Flowpacker: protein side-chain packing with torsional flow matching. *Bioinformatics*, 41(3):btaf010, 2025.

Jianan Li, Keisuke Yanagisawa, Masatake Sugita, Takuya Fujie, Masahito Ohue, and Yutaka Akiyama. Cycpeptmpdb: a comprehensive database of membrane permeability of cyclic peptides. *Journal of chemical information and modeling*, 63(7):2240, 2023.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.

Wei Liu, Nguyen Hung Pham, Chandra S Verma, Hwee Kuan Lee, and Jianguo Li. Multi-solvent conformational ensembles for predicting cyclic peptide permeability. *Scientific Data*, 2026.

Morgane Mannes, Charlotte Martin, Christel Menet, and Steven Ballet. Wandering beyond small molecules: peptides as allosteric protein modulators. *Trends in Pharmacological Sciences*, 43(5): 406–423, 2022.

Qingyi Mao, Tianfeng Shang, Wen Xu, Silong Zhai, Chengyun Zhang, Jingjing Guo, An Su, Chengxi Li, and Hongliang Duan. NCPepFold: Accurate prediction of noncanonical cyclic peptide structures via cyclization optimization with multigranular representation. *Journal of Chemical Theory and Computation*, 21(9):4979–4991, 2025.

Jiayuan Miao, Marc L Descoteaux, and Yu-Shan Lin. Structure prediction of cyclic peptides by molecular dynamics+ machine learning. *Chemical science*, 12(44):14927–14936, 2021.

Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, et al. Boltz-2: Towards accurate and efficient binding affinity prediction. *BioRxiv*, 2025.

Philipp Pracht, Fabian Bohle, and Stefan Grimme. Automated exploration of the low-energy chemical space with fast quantum chemical methods. *Physical Chemistry Chemical Physics*, 22(14):7169–7192, 2020.

Tamar Schlick, Stephanie Portillo-Ledesma, Christopher G Myers, Lauren Beljak, Justin Chen, Sami Dakhel, Daniel Darling, Sayak Ghosh, Joseph Hall, Mikaeel Jan, et al. Biomolecular modeling and simulation: a prospering multidisciplinary field. *Annual review of biophysics*, 50(1):267–301, 2021.

Diana P Slough, Hongtao Yu, Sean M McHugh, and Yu-Shan Lin. Toward accurately modeling N-methylated cyclic peptides. *Physical Chemistry Chemical Physics*, 19(7):5377–5388, 2017.

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

Yitian Wang, Fanmeng Wang, Angxiao Yue, Wentao Guo, Yaning Cui, and Hongteng Xu. MuCO: Generative peptide cyclization empowered by multi-stage conformation optimization. In *International Conference on Machine Learning*, 2026.

John D Westbrook, Chenghua Shao, Zukang Feng, Marina Zhuravleva, Sameer Velankar, and Jasmine Young. The chemical component dictionary: complete descriptions of constituent molecules in experimentally determined 3d macromolecules in the protein data bank. *Bioinformatics*, 31(8): 1274–1278, 2015.

N Kithmini Wijesiri and Benjamin P Brown. Parameterizing non-canonical amino acids for cyclic peptide simulations. *Methods in enzymology*, 723:167, 2025.

Dizhou Wu and Yike Zou. Accurate 3D structure prediction of small cyclic peptides containing non-canonical amino acid residues using an all-atom diffusion model with stereogenic implementation. *Journal of Chemical Information and Modeling*, 66(8):4398–4408, 2026.

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

Jason Yim, Andrew Campbell, Andrew YK Foong, Michael Gastegger, José Jiménez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S Veeling, Regina Barzilay, Tommi Jaakkola, et al. Fast protein backbone generation with se (3) flow matching. *arXiv preprint arXiv:2310.05297*, 2023a.

Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023b.

Angxiao Yue, Zichong Wang, and Hongteng Xu. ReQFlow: Rectified quaternion flow for efficient and high-quality protein backbone generation. In *Forty-second International Conference on Machine Learning*, 2025.

Cheng Zhu, Sen Cao, Tianfeng Shang, Jingjing Guo, An Su, Chengxi Li, and Hongliang Duan. Predicting the structures of cyclic peptides containing unnatural amino acids by HighFold2. *Briefings in Bioinformatics*, 26(3):bbaf202, 2025.

APPENDIX CONTENTS

A Taxonomy **15**

A.1 Five-Category Structural Taxonomy 15

A.2 Empirical Coverage Across Cyclic Peptide Datasets 15

B Model Architecture Training and Inference **17**

B.1 MuCO2 Architecture, Training, and Inference 17

B.2 Baseline Implementation: HighFold-MeD and HighFold-MeD2 20

B.3 MuCO2 Algorithms 22

C Structure-Guided Feynman-Kac Steering: Mathematical Formulation and Statistical Analysis **24**

C.1 Hierarchical Flow Architecture and the Inter-Stage Branching Checkpoint 24

C.2 Sequential Interacting Feynman-Kac Dynamics and Temperature Modulation . . . 25

C.3 Statistical Analysis: Why Feynman-Kac Branching Unlocks Multi-Basin Landscapes 25

D Dataset Curation and Data Processing **27**

D.1 Canonical CPSea Repository and Benchmark Subsets 27

D.2 Backbone-modified CPSea MeD Dataset and Simulation Protocol 27

D.3 HighFold2-Compatible Noncanonical Benchmark 28

D.4 CycPeptMPDB MeD and Mixed External Benchmarks 29

E Metrics **30**

E.1 Vacuum Potential Energy 30

E.2 Implicit Solvent Water Score 30

E.3 Sequence-Ordered Central Carbon RMSD 30

E.4 Residue-Level Chirality Preservation 31

F Additional Experimental Results **32**

G Additional Ablation Results **33**

H MuCO2 Samples Visualization **36**

A TAXONOMY

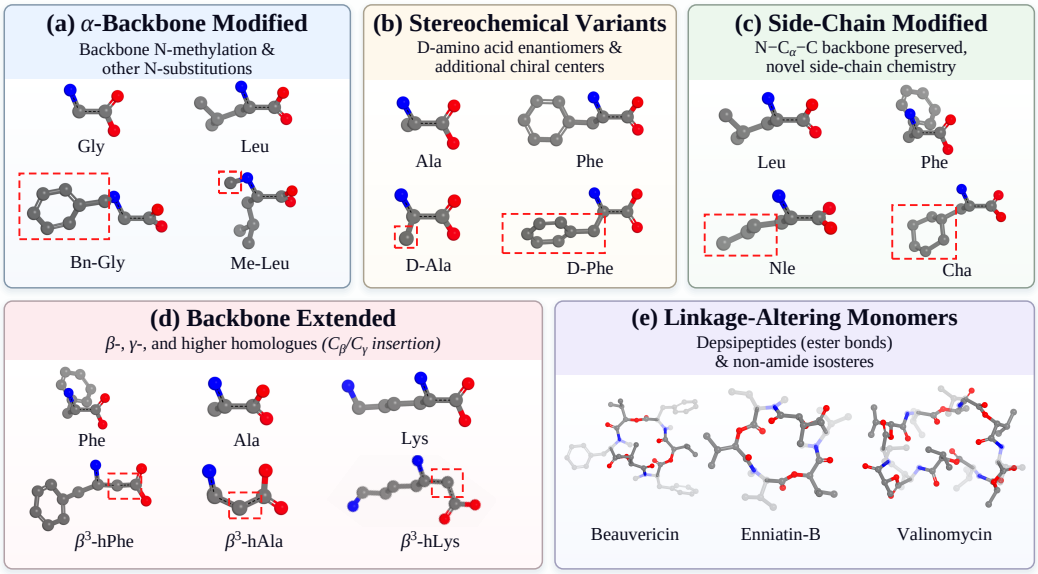


Figure 4: **Taxonomy of noncanonical building blocks in cyclic peptides. (a)** α-Backbone modifications (N-methylation and other complex N-substitutions). **(b)** Stereochemical variants (D-amino acids and chiral inversions). **(c)** Side-chain modifications with standard $\mathrm{N} - \mathrm{C}_\alpha - \mathrm{C}$ backbones. **(d)** Backbone extensions (β- and γ-amino acids). **(e)** Linkage-altering monomers (ester bonds in depsipeptides and non-amide links). MUCO2 focuses on categories (a), (b), and (c), which represent the vast majority of noncanonical modifications in drug discovery.

A.1 FIVE-CATEGORY STRUCTURAL TAXONOMY

To systematically classify noncanonical amino acids (ncAAs) in cyclic peptides, we define five structural perturbation categories based on how they alter the molecular scaffold (Figure 4):

(a) α-Backbone Modifications: Modifications on the standard backbone atoms while maintaining the standard connectivity. The most common form is backbone N-methylation (*e.g.*, Me-Leu), which removes an amide hydrogen-bond donor and alters *cis/trans* isomerization barriers. This class also includes C_α-substitutions like Aib as well as other bulky N-modifications (e.g., Bn-Gly in peptoids).

(b) Stereochemical Transformations: Inversions of C_α chirality from L- to D-amino acids (*e.g.*, D-Ala, D-Phe), as well as additional stereocenters in side chains. These inversions reshape local Ramachandran preferences and macrocyclic ring conformations.

(c) Side-Chain Modifications: Monomers that keep the standard $\mathrm{N} - \mathrm{C}_\alpha - \mathrm{C}(=\mathrm{O})$ backbone and standard chirality, but feature novel side chains (*e.g.*, Nle, Cha, halogenated rings, and cross-linking handles).

(d) Backbone Extensions: Residues with extra carbon atoms inserted into the backbone, such as β- and γ-amino acids (*e.g.*, β^3-hPhe). These expand the standard 3-bond repeating unit to 4 or 5 bonds.

(e) Linkage-Altering Monomers: Building blocks that replace standard amide bonds along the macrocyclic ring, such as ester bonds in depsipeptides (*e.g.*, valinomycin, beauvericin) and non-peptidic linkers.

A.2 EMPIRICAL COVERAGE ACROSS CYCLIC PEPTIDE DATASETS

To verify whether this taxonomy covers real-world macrocyclic peptides, we analyzed noncanonical residues across four major cyclic peptide datasets: CREMP, MacrocycleDB, CyclicPepedia, and CycPeptMPDB. In total, this corpus contains over $500{,}000$ residue units across more than $64{,}000$

entries, as summarized in Table 3. Input records refer to entries containing peptide IDs, while retained entries comprise deduplicated sequences with valid molecular representations (e.g., SMILES).

Comprehensive Coverage ($> 99.6\%$): Across all four datasets, our five categories capture almost all noncanonical residues. Specifically, unmapped or unclassified building blocks account for 0.00% in CREMP, 0.08% in MacrocycleDB, 0.35% in CyclicPepedia, and 0.06% in CycPeptMPDB. This confirms that the proposed taxonomy achieves practically complete coverage across standard macrocycle repositories.

Dominance of Categories (a), (b), and (c): The statistics in Table 3 also show that the vast majority of noncanonical modifications belong to the first three categories. In CycPeptMPDB, N/C_α modification (41.64%), stereochemical variation (27.05%), and side-chain modification (7.80%) together represent over 96% of all noncanonical instances. In CREMP, categories (a) and (b) account for 100% of the noncanonical modifications. Categories (d) and (e) are relatively rare ($< 3.5\%$) and typically occur in specialized natural products like depsipeptides.

These empirical observations justify the design of MuCO2. By decoupling and effectively modeling α-backbone modifications (a), stereochemical inversions (b), and open-ended side-chain chemistry (c), MuCO2 covers the overwhelming majority of noncanonical chemistry needed for macrocyclic drug discovery.

Table 3: **Noncanonical residue profiles and coverage across four cyclic peptide datasets.** Systematic profiling of noncanonical monomers across CREMP, MacrocycleDB, CyclicPepedia, and CycPeptMPDB ($> 500{,}000$ residue units). The proposed five-category taxonomy achieves near-lossless coverage ($> 99.6\%$ across all databases with unmapped residues $\leq 0.35\%$), while demonstrating that categories (a), (b), and (c) represent over 95% of all practical noncanonical modifications.

Metric / Dataset	CREMP	MacrocycleDB	CyclicPepedia	CycPeptMPDB
Dataset composition				
Input records, n	36,198	11,077	8,751	8,466
Retained entries, n	36,198	11,074	8,335	7,988
Detected residue/monomer units, N	167,860	113,684	159,585	63,472
Sequence-only entries, n	0	0	1,386	0
Entries without usable structure or sequence, n	0	0	325	0
Class frequencies				
Any a–e label, n (%)	40,608 (24.19)	36,472 (32.08)	11,107 (6.96)	43,884 (69.14)
a: N/Cα modification, n (%)	26,923 (16.04)	7,287 (6.41)	3,034 (1.90)	26,431 (41.64)
b: Specified stereochemical variant, n (%)	26,048 (15.52)	14,842 (13.06)	4,288 (2.69)	17,167 (27.05)
c: Side-chain modification, n (%)	0 (0.00)	19,009 (16.72)	4,116 (2.58)	4,950 (7.80)
d: Backbone extension, n (%)	0 (0.00)	1,188 (1.05)	534 (0.33)	1,467 (2.31)
e: Linkage / nonpeptide / other, n (%)	0 (0.00)	3,543 (3.12)	2,388 (1.50)	1,709 (2.69)
Label multiplicity and common patterns				
Exactly 1 label, n (%)	28,245 (16.83)	28,291 (24.89)	8,345 (5.23)	36,140 (56.94)
Exactly 2 labels, n (%)	12,363 (7.37)	7,041 (6.19)	2,328 (1.46)	7,649 (12.05)
Exactly 3 labels, n (%)	0 (0.00)	1,064 (0.94)	383 (0.24)	94 (0.15)
Four or more labels, n (%)	0 (0.00)	76 (0.07)	51 (0.03)	1 (0.00)
Recorded pattern rank 1, pattern (%)	*a* (8.67)	*c* (10.59)	*b* (1.41)	*a* (30.27)
Recorded pattern rank 2, pattern (%)	*b* (8.15)	*b* (7.59)	*c* (1.31)	*b* (16.61)
Recorded pattern rank 3, pattern (%)	*ab* (7.37)	*a* (3.66)	*a* (1.25)	*ab* (9.52)
Annotation quality				
Unspecified relevant stereochemistry, n (%)	0 (0.00)	1,411 (1.24)	3,979 (2.49)	15 (0.02)
Other/unmapped subset of e, n (%)	0 (0.00)	91 (0.08)	555 (0.35)	37 (0.06)

B Model Architecture Training and Inference

This section provides comprehensive technical details regarding the architectural design, feature injection mechanisms, curriculum learning workflows, and inference protocols of MuCO2, alongside the implementation, retraining, and evaluation standards for the comparative baselines HighFold-MeD and HighFold-MeD2.

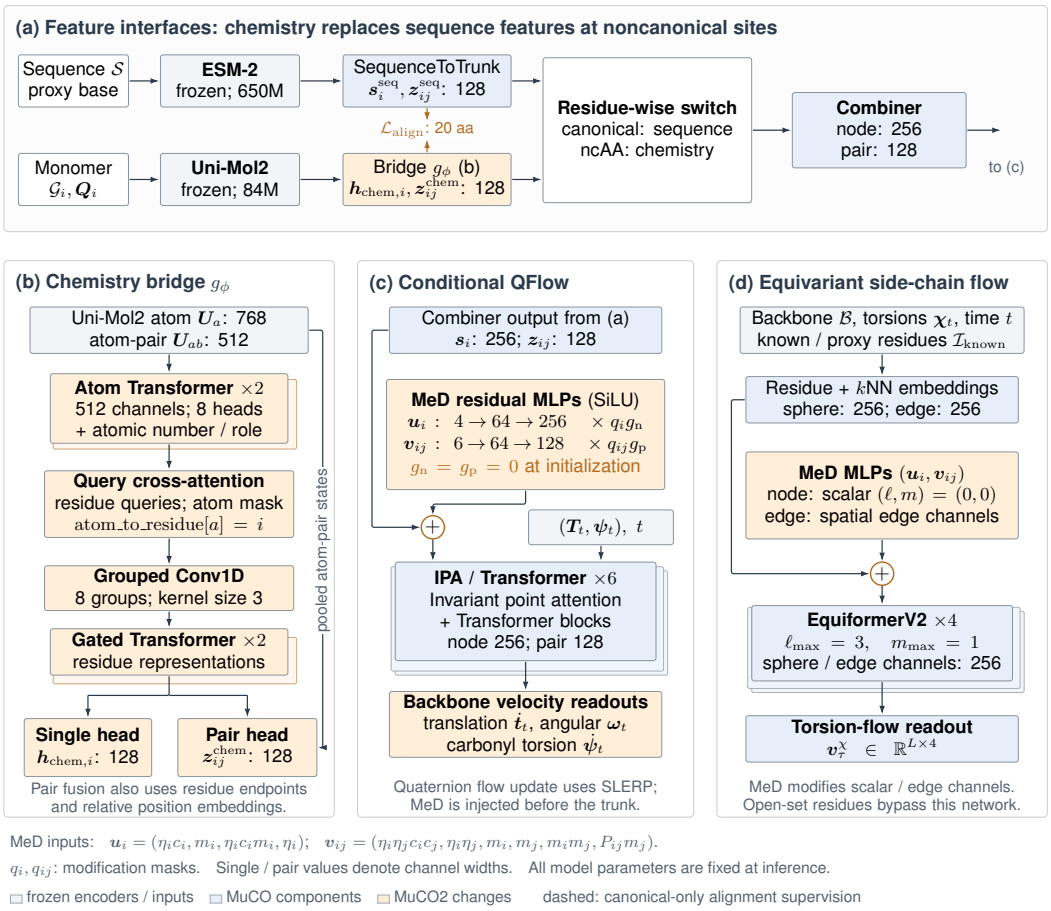


Figure 5: **Overview of MuCO2 architecture and multi-axis conditioning.** (a) Sequence and chemical feature interfaces. (b) Chemistry bridge from atom features to residue representations. (c) QFlow backbone network with gated MeD conditioning. (d) Equivariant side-chain network with scalar and edge conditioning.

B.1 MuCO2 Architecture, Training, and Inference

Network Architecture and Conditioning Mechanisms. MuCO2 establishes a decoupled, multi-stage generative framework comprising an accelerated cyclic backbone generator, a zero-shot cross-modal adapter, a stereochemistry-aware side-chain packer, and downstream force-field relaxation (Figure 5).

(i) Quaternion Backbone Flow (QFlow). Rather than utilizing iterative Riemannian SE(3) geodesic projections, MuCO2 parameterizes cyclic backbone rotations on the quaternion 3-sphere $\mathbb{S}^3$ using Rectified Quaternion Flow (Yue et al., 2025). The state at time $t \in [0, 1]$ comprises residue frames $T_t = (R_t, t_t)_{i=1}^L$ and carbonyl oxygen dihedrals $\psi_t \in \mathbb{T}^L$. Translations $t_t \in \mathbb{R}^3$ are linearly interpolated under a coordinate scaling factor of 0.1. Rotations $R_t \in \mathrm{SO}(3)$ are mapped to unit quaternions $q_t \in \mathbb{S}^3$, with double-cover ambiguities eliminated by enforcing $\langle q_t, q_1 \rangle \geq 0$. Velocity

fields are integrated via spherical linear interpolation (SLERP):

$$\hat{q}_0 = q_t \otimes \exp\left(-\frac{1}{2} t\omega_t\right), \quad q_{t-\Delta t} = \mathrm{SLERP}\left(q_t,\, \hat{q}_0,\, \min\left(\frac{\Delta t}{t}, 1\right)\right), \tag{7}$$

where $\omega_t \in \mathbb{R}^3$ denotes the predicted angular velocity vector. The backbone trunk operates on joint representations combining sequence embeddings and current geometric frames through Invariant Point Attention (IPA) and Transformer blocks.

(ii) Explicit MeD Condition Injection. Backbone N-methylation and stereochemical chirality are parameterized via four primary operational fields: signed chirality $c_i \in \{+1, -1\}$ ($+1$ for L, -1 for D, fixed to $+1$ for achiral Glycine), binary methylation state $m_i \in \{0, 1\}$, chiral mask $\eta_i \in \{0, 1\}$ ($\eta_i = 0$ for Glycine, 1 otherwise), and directed peptide-bond adjacency $P_{ij} \in \{0, 1\}$ ($P_{ij} = 1$ denoting a verified covalent bond from C_i to N_j). These raw indicators are transformed into a 4-dimensional node feature u_i and a 6-dimensional pair feature v_{ij}:

$$u_i = \left(\eta_i c_i,\ m_i,\ \eta_i c_i m_i,\ \eta_i\right) \in \mathbb{R}^4, \tag{8}$$

$$v_{ij} = \left(\eta_i\eta_j c_i c_j,\ \eta_i\eta_j,\ m_i,\ m_j,\ m_i m_j,\ P_{ij} m_j\right) \in \mathbb{R}^6. \tag{9}$$

The term $P_{ij}m_j$ explicitly encodes the directional steric impact of peptide bonds pointing into an N-methylated nitrogen. To prevent optimization instability and avoid disrupting canonical generation priors, modification activation masks are defined as $q_i = \mathbb{I}[(\eta_i > 0 \wedge c_i < 0) \vee m_i > 0]$ and $q_{ij} = q_i \vee q_j \vee (P_{ij} m_j > 0)$. Residual perturbations are injected into the backbone representation immediately downstream of the combiner module and prior to the trunk Transformer:

$$\Delta s_i = g_{\text{node}} \cdot \mathrm{MLP}_{\text{node}}(u_i) \cdot q_i, \tag{10}$$

$$\Delta z_{ij} = g_{\text{pair}} \cdot \mathrm{MLP}_{\text{pair}}(v_{ij}) \cdot q_{ij}, \tag{11}$$

where $\mathrm{MLP}_{\text{node}} : \mathbb{R}^4 \to \mathbb{R}^{64} \to \mathbb{R}^{256}$ and $\mathrm{MLP}_{\text{pair}} : \mathbb{R}^6 \to \mathbb{R}^{64} \to \mathbb{R}^{128}$ employ SiLU activations. Crucially, g_{node} and g_{pair} are initialized strictly to zero, ensuring that the model initially recovers the unperturbed canonical backbone flow prior to parameter updates.

(iii) Uni-Mol2 Cross-Modal Chemical Alignment. To represent open-ended noncanonical side-chain chemistries without retraining the primary generative flows, MuCO2 utilizes the frozen 3D molecular foundation model Uni-Mol2 (84M parameter configuration: 12 layers, atom hidden dimension 768, pair width 512). For any noncanonical monomer, 3D conformers $Q_i^{(p)}$ generated from the Chemical Component Dictionary (CCD) or RDKit ETKDG are encoded into atom representations $U_a \in \mathbb{R}^{768}$ and atom-pair representations $U_{ab} \in \mathbb{R}^{512}$. A specialized residual bridge network (UniMolToTrunkBridge) maps these states directly into the 128-dimensional single and pair latent interfaces of the backbone trunk. The bridge applies: (1) an atom Transformer encoder (2 layers, hidden dimension 512, 8 heads) augmented with atomic number and role embeddings; (2) residue query cross-attention, where attends exclusively to constituent atoms defined by an assignment tensor $\mathrm{atom_to_residue}[a] = i$; (3) residue-level 1D grouped convolutions (8 groups, kernel size 3) followed by 2 gated Transformer blocks; and (4) a dual-head projection yielding contextual single embeddings $h_{\text{chem},i} \in \mathbb{R}^{128}$ and semantic pair embeddings $z_{ij}^{\text{chem}} \in \mathbb{R}^{128}$ (combining mean-pooled atom pairs $\bar{U}_{ij}$, projected residue endpoints, and relative position embeddings copied from the sequence adapter). During inference, noncanonical residues are masked in the sequence prior, and $h_{\text{chem},i}$ and z_{ij}^{chem} directly replace the ESM-2 representations, bypassing sequence language model limitations in a zero-shot manner.

(iv) Hybrid Side-Chain Assembly and Stereochemical Reconstruction. For canonical residues and known natural proxies (I_{known}), side-chain torsions $\chi \in \mathbb{T}^4$ are sampled via the conditional EquiformerV2 architecture (4 equivariant layers, 256 sphere channels, 256 edge channels, $\ell_{\max} = 3, m_{\max} = 1$). The MeD node perturbation $\mathrm{MLP}_{\text{sc,node}}(u_i)$ is injected exclusively into the $\ell = 0, m = 0$ scalar spherical harmonic channels, while edge perturbations $\mathrm{MLP}_{\text{sc,edge}}(v_{ij})$ are added to the spatial edge embeddings constructed over the k-nearest-neighbor residue graph. For open-ended noncanonical residues (I_{open}, which also accommodates complex category-(a) N-substituted residues), neural rotamer prediction is bypassed; rigid monomer templates Q_i^{tpl} (bearing the noncanonical side chain and/or bulky N-substituent) are aligned onto the generated backbone frame $T_i = (R_i, t_i)$ via optimal Kabsch superposition on the backbone atoms (N, Cα, C). To enforce correct D-stereochemistry on reconstructed Cartesian coordinates, the Idealizer module reflects side-chain coordinates across the local N–C_α–C plane whenever $c_i = -1$, followed by analytical dihedral adjustment of χ_1 to accommodate local valence angles.

Pretraining and Progressive Curriculum Learning. MuCO2 is trained via a progressive multi-stage curriculum designed to navigate high-energy steric and torsional barriers without encountering mode collapse:

(i) Canonical Foundation Pretraining. Base generative flows for backbone generation (QFlow and SE(3) Flow) and side-chain packing (EquiformerV2) were trained from scratch on the canonical CPSea dataset (75,337 cyclic peptides, $L \in [7, 16]$) following the foundational MuCO protocol. Stage-1 backbone models were trained for 100 epochs using the AdamW optimizer with a base learning rate of $\text{lr}_{\text{base}} = 2 \times 10^{-4}$, total batch size 256, and exponential moving average (EMA) decay of 0.999. Stage-2 side-chain models were trained for 100 epochs with $\text{lr}_{\text{base}} = 1 \times 10^{-4}$ and batch size 128.

(ii) Progressive Three-Round MeD Curriculum. Fine-tuning across backbone modifications and chiral inversions was conducted across three successive rounds, with backbone and side-chain networks fine-tuned independently. To prevent catastrophic forgetting and stabilize fine-tuning dynamics, all learning rates were scaled down by a factor of 10 relative to pretraining ($\text{lr}_{\text{bb}} = 2 \times 10^{-5}$ for backbone; $\text{lr}_{\text{sc}} = 1 \times 10^{-5}$ for side chains):

Round 1 (Warmup): The model was trained for 100 epochs on the CPSea MeD training subset (9,022 systems featuring single N-methylation and single D-amino acids) augmented with their full enantiomeric mirror reflections ($k_D = L_\chi$, generated at zero simulation cost via the centered reflection operator $M(X)$). To anchor canonical folding priors, this dataset was co-trained with an equivalent volume of canonical cyclic peptides and their corresponding mirror reflections, strictly disjoint from the parent structures of CPSea MeD.

Round 2 (Self-Improving Exploration): The Round 1 checkpoint was deployed to generate conformational candidates for macrocycles with moderate modification densities ($k_{\text{Me}} \leq 3$, chirality boundary distance $b_D = \min(k_D, L_\chi - k_D) < 3$) using FKS sampling ($K = 10, K_0 = 5, M = 2$). Generated candidates underwent regularized CHARMM36 physical relaxation (Appendix D.2). Uncyclized structures and conformations exhibiting steric clashes or energetic blow-ups were purged. For each target, the single conformer attaining the lowest valid potential energy was retained as a pseudo-label. This harvested pool was combined with its enantiomeric reflections and mixed with an equivalent, mutually non-overlapping batch of canonical peptides and their mirror counterparts. The model was fine-tuned on this expanded distribution for 50 epochs.

Round 3 (Multi-Site Expansion): The exploration boundary was expanded to dense, compounding regimes ($k_{\text{Me}} \leq 5, b_D \leq 5$). The Round 2 model sampled diverse backbones and rotamers, followed by identical CHARMM36 relaxation and minimum-energy filtering. Training on this final dataset, balanced with an equivalent volume of unseen canonical and mirror structures, was conducted for 30 epochs, establishing the final production parameters $(\theta^\star, \tau^\star)$.

(iii) Cross-Modal Chemistry Alignment Pretraining. The UniMolToTrunkBridge adapter g_ϕ was trained independently on the 20 canonical amino acid classes embedded in full cyclic scaffolds. Contextual targets were derived from the frozen ESM-2 (650M) and its trained SequenceToTrunk projection. The adapter was optimized for 20 epochs using AdamW ($\text{lr} = 2 \times 10^{-4}$, weight decay 1×10^{-4}) under the alignment objective specified in equation 2 with cosine penalty factor $\lambda_{\cos} = 0.2$. Random sequence masking was applied with probability $p = 0.35$ up to 2 sites per macrocycle.

Inference Pipeline and Hierarchical Sampling. During inference, MuCO2 executes a three-stage hierarchical exploration pipeline:

Stage 1 (Steered Backbone Generation): The model proposes an overcomplete candidate set of K backbone scaffolds $\{B_1, \ldots, B_K\}$ from QFlow in 100 reverse integration steps. Structure-Guided Feynman-Kac Steering (FKS) evaluates the pairwise SO(3)-constrained Kabsch RMSD distance matrix d_{ij} and mean Gaussian kernel crowding $\rho_i(A) = \frac{1}{|A|} \sum_{j \in A} \exp(-d_{ij}^2/(2\sigma^2))$ with structural bandwidth $\sigma = 1.75$ Å. Sequential particle selection under exploratory temperature $\tau = 2.5$ extracts a structurally dispersed subset A containing K_0 diverse backbones ($K_0 \leq K$).

Stage 2 (Rotamer Expansion): For each surviving scaffold $B_k \in A$, the model samples M independent rotamer configurations across 10 reverse steps using EquiformerV2 for known

Table 4: **Comprehensive training and inference hyperparameters of MUCO2.**

Hyperparameter / Setting	Stage-1 (Backbone)	Stage-2 (Side-Chain)	Adapter (g_ϕ)
Model Architecture & Dimensions			
Backbone / Network Family	QFlow / SE(3) Flow	EquiformerV2	Res-Transformer
Layers / Depth	6 IPA / Transformer	4 Equivariant blocks	2 Atom + 2 Residue
Hidden Dimensions	Node: 256, Pair: 128	Sphere: 256, Edge: 256	Hidden: 512, Head: 128
Spherical Harmonics	–	$\ell_{\max} = 3,\ m_{\max} = 1$	–
Optimization & Curriculum Training			
Base LR (lr_{base})	2×10^{-4}	1×10^{-4}	2×10^{-4}
Curriculum LR (lr_{curr})	2×10^{-5}	1×10^{-5}	–
Optimizer / Weight Decay	AdamW / 1×10^{-4}		
Pretrain / Align Epochs	100	100	20
Curriculum Epochs (R1 / R2 / R3)	100 / 50 / 30	100 / 50 / 30	–
Batch Size / Precision	256 / FP32	128 / FP32	32 / FP32
EMA Decay Rate	0.999	0.999	–
Training Hardware	4× NVIDIA A100 80GB (PyTorch DDP)		
Inference & Sampling Protocol			
Integration Steps	100 time points	10 time points	–
FKS Steering (τ, σ)	$\tau = 2.5,\ \sigma = 1.75$ Å	–	–
Sampling Budget @1	$K = 1,\ K_0 = 1$	$M = 1$ (Total: 1)	–
Sampling Budget @10	$K = 5,\ K_0 = 5$	$M = 2$ (Total: 10)	–
Inference Hardware	1× NVIDIA A100 80GB		

residues or template attachment for open-set noncanonical residues, producing a total pool of $K_0 \times M$ all-atom structures $\{\widetilde{X}_m^k\}$.

Stage 3 (Physics Refinement): All $K_0 \times M$ candidates undergo automated covalent topology detection and fast regularized CHARMM36 force-field relaxation (500 steps of steepest descent with harmonic restraint weight $w = 10.0\,\text{kcal}/(\text{mol}\cdot\text{Å}^2)$) to eliminate local van der Waals clashes and resolve covalent closure. Conformations satisfying the closure tolerance $\delta = 0.1$ Å are ranked by total effective potential energy.

For standardized benchmarking, deterministic single-structure evaluation (@1) fixes $K = K_0 = 1$ and $M = 1$ (1 candidate). Ensemble evaluation (@10) fixes $K = K_0 = 5$ and $M = 2$ (10 candidates assembled from 5 distinct scaffolds and 2 rotamer packings). A comprehensive summary of training and sampling hyperparameters is provided in Table 4.

B.2 Baseline Implementation: HighFold-MeD and HighFold-MeD2

Because model weights and pre-processed training sets were not released for HighFold-MeD and HighFold-MeD2, both baselines were retrained from scratch on the CPSea MeD dataset following their respective architectural specifications and published training protocols to establish a rigorous and equitable benchmark.

HighFold-MeD Architecture and Retraining. HighFold-MeD (Cao et al., 2025b) adapts the AlphaFold2 architecture (defaulting to the 48-block Evoformer configuration of five parameters) to support backbone N-methylation and D-stereochemistry.

Chemical Vocabulary Expansion: The primary sequence vocabulary is expanded to 76 discrete residue types (20 standard canonical amino acids, 19 N-methylated variants, 19 D-amino acids, and 18 dual N-methylated D-residues), each indexed via unique single-character symbols. The internal residue coordinate templates extend dense atom representations to 15 slots (accommodating the additional methyl carbon atom CN), with modified rigid group definitions and ideal geometry parameters. However, in standard AlphaFold2 feature preprocessing, input sequence one-hot channels remain constrained to 22/23 dimensions; consequently, extended residue tokens outside

this range project to zero vectors along the one-hot input channel, relying on template frames and recycling mechanisms to convey modification context.

Cyclic Relative Positional Encoding: Pair representations are augmented with a cyclic relative sequence distance defined over sequence length L:

$$d_{ij} = \text{sign}(i-j)\min(|i-j|,\, L-|i-j|), \tag{12}$$

which is binned, one-hot encoded, and projected into pair activations.

Training Objective and Optimization: Training was initialized from the official AlphaFold2 five parameters. The optimization objective is governed by the final-layer backbone Frame Aligned Point Error (FAPE):

$$\mathcal{L}_{\text{MeD1}} = w_{\text{native}} \cdot \frac{1}{|\mathcal{M}|} \sum_{i,j \in \mathcal{M}} \frac{\min\left(d_{\text{clamp}},\, \|\hat{T}_i^{-1}\hat{x}_j - T_i^{-1}x_j\|\right)}{d_{\text{unit}}}, \tag{13}$$

where $\hat{T}_i$ and T_i represent predicted and ground-truth backbone frames, and w_{native} denotes the reference weighting. Auxiliary violation and distogram losses computed during forward passes do not contribute gradients to this target. Following the native JAX/Haiku framework design, HighFold-MeD was fine-tuned on a single NVIDIA A100 (80GB) GPU using the Optax Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$, lr $= 7.5 \times 10^{-5}$, gradient clipping norm 10.0) with exact-length batch bucketing. When subjected to the multi-round curriculum, the model rapidly saturated and reached convergence at Round 2, showing no performance gains in subsequent rounds.

HighFold-MeD2 Architecture and Retraining. HighFold-MeD2 (Cao et al., 2026) is constructed on the all-atom Boltz-2 diffusion architecture, predicting full Cartesian coordinates directly from molecular component graphs.

Chemical Component Representation: Noncanonical residues are ingested via Chemical Component Dictionary (CCD) definitions. The Boltz2 tokenizer assigns a single polymer residue token to modified amino acids while marking an active `modified` flag, with full heavy-atom graphs, reference 3D conformers, and partial charges supplied to local atom attention modules.

Conditioning Embeddings: Token representations incorporate four mandatory condition flags: cyclic boundary flags (`add_cyclic_flag`), experimental method types (`add_method_conditioning`), noncanonical residue flags (`add_modified_flag`), and polymer molecule type embeddings (`add_mol_type_feat`).

Diffusion Loss and Optimization: Training was initialized from the pre-trained Boltz-2 checkpoint. Diffusion score networks were trained using EDM preconditioned denoising over all-atom coordinates perturbed by log-normal noise $\sigma \sim \mathcal{N}(-1.2, 1.2^2)$:

$$\mathcal{L}_{\text{MeD2}} = 4.0 \cdot \mathcal{L}_{\text{diffusion}} + 0.03 \cdot \mathcal{L}_{\text{distogram}}, \tag{14}$$

where $\mathcal{L}_{\text{diffusion}}$ evaluates coordinate MSE under optimal weighted rigid alignment and $\mathcal{L}_{\text{distogram}}$ supervises token pair distance bins. Fine-tuning was executed on 4 $\times$ NVIDIA A100 GPUs using PyTorch Distributed Data Parallel (DDP) with mixed precision (BF16), AdamW optimizer (lr $= 1.8 \times 10^{-3}$, weight decay 0.01, linear warmup of 1,000 steps, gradient clipping 10.0), and a noise multiplicity of 16. Similar to HighFold-MeD, curriculum progression stabilized by Round 2, beyond which validation loss plateaued.

Inference and Evaluation Protocols for Baselines. To maintain strict parity across benchmarks, baseline inferences were standardized without structural templates or MSA database searches, all generated structures from both baselines were subjected to identical downstream CHARMM36 fast energy minimization:

HighFold-MeD: Evaluated using its five retrained AlphaFold2 checkpoint configurations. For deterministic @1 metrics, predictions were generated across all 5 model parameters, and the single candidate exhibiting the highest predicted confidence (mean pLDDT) was selected. Because model sampling with fixed random seeds is deterministic, @10 multi-conformation entries are not applicable and left blank.

HighFold-MeD2: Evaluated using its trained diffusion network over 200 denoising steps per trajectory. For @1 metrics, a single stochastic trajectory was drawn. For @10 ensemble metrics, 10 independent trajectories were drawn across distinct stochastic seeds.

B.3 MuCO2 Algorithms

Algorithm 2 Progressive Physics-Aware Curriculum Training

Input: An initial dataset $\mathcal{D}_0 \cup \mathcal{D}_1$; target perturbation sets $\mathcal{T}_r$ ($r = 2, 3$); frozen ESM-2 and Uni-Mol2 encoders.
Output: Model $(\theta^\star, \tau^\star, \phi^\star) \leftarrow (\theta_3, \tau_3, \phi^\star)$.

Initialization
$(\theta_0, \tau_0) \leftarrow \mathrm{Pretrain}(\mathcal{D}_0; 100 \text{ epochs})$ by equations 5 & 6
$g_{\mathrm{node}}, g_{\mathrm{pair}} \leftarrow 0$; $\quad (E_1, E_2, E_3) \leftarrow (100, 50, 30)$
Round-1 Generalization: supervised warmup
$\mathcal{D}^{(1)} \leftarrow M(\mathcal{D}_0 \cup \mathcal{D}_1) \cup \mathcal{D}_0 \cup \mathcal{D}_1$
$(\theta_1, \tau_1) \leftarrow \mathrm{FineTune}(\theta_0, \tau_0; \mathcal{D}^{(1)}, E_1)$ by equations 5 & 6
Round 2–3 Generalization: physics-guided self-training
for $r = 2, 3$ **do**
  $\mathcal{D} \leftarrow \varnothing$
  for $(S, R) \in \mathcal{T}_r$ with $k_{\mathrm{Me}}, b_{\mathrm{D}} \le 2r - 1$ **do**
    $\widetilde{H} \leftarrow \mathrm{Encode}_{\mathrm{ESM-2}}(S, R)$
    $\{B_k\}_{k=1}^{K} \sim p_{\theta_{r-1}}(\cdot \mid \widetilde{H})$
    $A \leftarrow \mathrm{FKS}(\{B_k\}_{k=1}^{K}; K_0, \sigma, \tau_{\mathrm{fks}})$
    $P \leftarrow \mathrm{HybridPack}_{\tau_{r-1}}(A, \widetilde{H}, R; M)$
    $P^\star \leftarrow \{\mathrm{Relax}_{500}(\widetilde{X}; R, w) : \widetilde{X} \in P\}$
    $P_{\mathrm{valid}} \leftarrow \{X \in P^\star : \mathrm{Valid}(X; R, \delta)\}$
    if $P_{\mathrm{valid}} \neq \varnothing$ **then**
      $X_{\mathrm{pseudo}} \leftarrow \arg\min_{X \in P_{\mathrm{valid}}} E_R(X)$
      $\mathcal{D} \leftarrow \mathcal{D} \cup \{(S, R, X_{\mathrm{pseudo}})\}$
    end if
  end for
  $\mathcal{D}^{(r)} \leftarrow M(\mathcal{D}) \cup \mathcal{D} \cup \mathcal{D}^{(r-1)}$
  $(\theta_r, \tau_r) \leftarrow \mathrm{FineTune}(\theta_{r-1}, \tau_{r-1}; \mathcal{D}^{(r)}, E_r)$ by equations 5 & 6
end for
Alignment
$\phi^\star \leftarrow \mathrm{Align}(\mathcal{D}_0, \text{ESM-2}, \theta_3, \text{Uni-Mol2}; 20 \text{ epochs})$ by equation 2
return Model $(\theta^\star, \tau^\star, \phi^\star) \leftarrow (\theta_3, \tau_3, \phi^\star)$

Operators and settings. $\mathrm{Align}(\mathcal{D}_0, \text{ESM-2}, \theta_3, \text{Uni-Mol2}; 20 \text{ epochs})$ optimizes the cross-modal bridge g_ϕ via the node MSE/cosine and pair MSE loss defined in equation 2. Encode, FKS, and HybridPack (supporting open-set side chains and N-modifications) follow Algorithm 3. $M(\mathcal{D}) = \mathcal{D} \cup \mathcal{M}(\mathcal{D})$ reflects 3D Cartesian coordinates and inverts chiral configurations ($c_i \mapsto -c_i$ for chiral residues with $\eta_i = 1$). Batches during fine-tuning are volume-balanced against mutually disjoint canonical structures. Target complexity constraints: $k_{\mathrm{Me}}, b_{\mathrm{D}} \le 2r - 1$, where $b_{\mathrm{D}} = \min(k_{\mathrm{D}}, L_\chi - k_{\mathrm{D}})$. Sampling and optimization hyperparameters: $(K, K_0, M) = (10, 5, 2)$, $(\sigma, \tau_{\mathrm{fks}}) = (1.75\,\text{Å}, 2.5)$, and $(lr_{\mathrm{bb}}, lr_{\mathrm{sc}}) = (2 \times 10^{-5}, 10^{-5})$. $\mathrm{Relax}_{500}(\widetilde{X}; R, w)$ takes 500 CHARMM36 steepest-descent steps on $E_R(X) + \frac{w}{2}\|X - \widetilde{X}\|_F^2$ with $w = 10\,\mathrm{kcal}/(\mathrm{mol} \cdot \text{Å}^2)$. Valid requires covalent ring closure error $\le \delta = 0.1$ Å and the absence of steric clashes or energy divergences.

Algorithm 3 Hierarchical Multi-Stage Steered Inference

Input: Sequence S; perturbations R; trained $(\theta^\star, \tau^\star, \phi^\star)$;
templates $\{Q_i^{\mathrm{tpl}}\}$; budgets $1 \le K_0 \le K$, $M \ge 1$; $\epsilon > 0$.
Output: Low-energy ensemble $X_{\mathrm{LE}}(R)$ (possibly empty).

Representation encoding
if all residues are canonical **then**
  $\widehat{H} \leftarrow \mathrm{Seq2Trunk}(\text{ESM-2}(S))$
else
  $\widehat{H} \leftarrow g_{\phi^\star}\big(\text{Uni-Mol2}(\{G_i, Q_i\}_{i=1}^{L})\big)$
end if
$(\boldsymbol{u}_{\mathrm{bb}}, \boldsymbol{v}_{\mathrm{bb}}, \boldsymbol{u}_{\mathrm{sc}}, \boldsymbol{v}_{\mathrm{sc}}) \leftarrow \mathrm{MeDFeatures}(R)$
$\widetilde{H} \leftarrow \{\hat{\boldsymbol{h}}_i + \Delta \boldsymbol{s}_i(\boldsymbol{u}_{\mathrm{bb},i}),\ \hat{\boldsymbol{h}}_{ij} + \Delta \boldsymbol{z}_{ij}(\boldsymbol{v}_{\mathrm{bb},ij})\}_{i,j}$
Stage 1: backbone generation and FKS selection
$\{B_k\}_{k=1}^{K} \sim p_{\theta^\star}(\,\cdot \mid \widetilde{H})$ (100 SLERP steps)
$d_{ij} \leftarrow \mathrm{RMSD}_{\mathrm{SO}(3)}(B_i, B_j)$; $\kappa_{ij} \leftarrow e^{-d_{ij}^2/(2\sigma^2)}$
$i_1 \sim \mathrm{Uniform}(\{1, \dots, K\})$; $A \leftarrow \{i_1\}$
for $r = 2, \dots, K_0$ **do**
  $\rho_i \leftarrow |A|^{-1} \sum_{j \in A} \kappa_{ij}$, $\omega_i \leftarrow (\epsilon + \rho_i)^{-1/\tau_{\mathrm{fks}}}$ $(i \notin A)$
  $i_r \sim \mathrm{Categorical}\big(\{\omega_i / \sum_{j \notin A} \omega_j\}_{i \notin A}\big)$; $A \leftarrow A \cup \{i_r\}$
end for
Stage 2: hybrid side-chain assembly
$P \leftarrow \varnothing$
for $k \in A$ and $m = 1, \dots, M$ **do**
  $\boldsymbol{\chi}_{\mathrm{known}} \sim p_{\tau^\star}(\,\cdot \mid B_k, \widetilde{H}, \boldsymbol{u}_{\mathrm{sc}}, \boldsymbol{v}_{\mathrm{sc}})$
  $\mathcal{C}_{\mathrm{open}} \leftarrow \{\mathrm{Attach}_R(Q_i^{\mathrm{tpl}}, T_{i,k}) : i \in \mathcal{I}_{\mathrm{open}}\}$
  $\widetilde{X} \leftarrow \mathrm{FK}_R(B_k, \boldsymbol{\chi}_{\mathrm{known}}, \mathcal{C}_{\mathrm{open}})$
  $P \leftarrow P \cup \{\widetilde{X}\}$
end for
Stage 3: relaxation and multi-basin selection
$P_{\mathrm{valid}} \leftarrow \varnothing$
for $X \in P$ **do**
  $X^\star \leftarrow \mathrm{Relax}_{500}(\widetilde{X}; R, w)$
  if $\mathrm{Valid}(X^\star; R, \delta)$ **then**
    $E_{\mathrm{eff}}(X^\star) \leftarrow E_{\mathrm{vac}}(X^\star) + \Delta G_{\mathrm{solv}}(X^\star)$
    $P_{\mathrm{valid}} \leftarrow P_{\mathrm{valid}} \cup \{X^\star\}$
  end if
end for
return $\mathrm{Select}_{E_{\mathrm{eff}},\mathrm{RMSD}}(P_{\mathrm{valid}})$

Operators and settings. MeDFeatures yields backbone descriptors $(\boldsymbol{u}_{\mathrm{bb}}, \boldsymbol{v}_{\mathrm{bb}})$ and side-chain descriptors $(\boldsymbol{u}_{\mathrm{sc}}, \boldsymbol{v}_{\mathrm{sc}})$, where $\boldsymbol{u}_{\mathrm{bb},i} = (\eta_i c_i, m_i, \eta_i c_i m_i, \eta_i)$ and $\boldsymbol{v}_{\mathrm{bb},ij} = (\eta_i \eta_j c_i c_j, \eta_i \eta_j, m_i, m_j, m_i m_j, P_{ij} m_j)$, with P_{ij} being the directed peptide-bond indicator used by the MeD encoder.

Attach_R aligns a template to the local backbone frame, applies N-methylation, and matches target chirality by a single local reflection when required. Select ranks valid candidates by E_{eff} and removes RMSD-redundant structures; $\mathrm{Select}(\varnothing) = \varnothing$. Use $\sigma = 1.75$ Å, $\tau_{\mathrm{fks}} = 2.5$, and the same Relax_{500}, w, Valid, and δ as Algorithm 2.

C Structure-Guided Feynman-Kac Steering: Mathematical Formulation and Statistical Analysis

This section provides the rigorous mathematical formulation, algorithmic implementation, and theoretical analysis of the Structure-Guided Feynman-Kac Steering (FKS) framework proposed in Section 4.2.

C.1 Hierarchical Flow Architecture and the Inter-Stage Branching Checkpoint

Rather than executing an unstructured monolithic generative trajectory, MuCO2 decomposes macrocycle conformation sampling into a two-level hierarchical Markov jump process. Stage-1 operates on the continuous macrocyclic backbone space $\mathcal{B} \subset SE(3)^L$ parameterized by Quaternion Flow (QFlow) over time $t \in [0, 1]$. Stage-2 operates on the downstream multi-dimensional side-chain rotamer manifold $\mathcal{C} \subset \mathbb{T}^{4L}$ conditioned on the synthesized backbone.

Because side-chain rotamer degrees of freedom exhibit an exponential combinatorial explosion ($4L$ torsional angles per macrocycle), advancing raw, uncurated backbone trajectories directly into downstream side-chain sampling results in severe computational redundancy and mode collapse. Consequently, the transitional interface between Stage-1 and Stage-2 ($t_{\text{backbone}} = 1 \to t_{\text{sidechain}} = 0$) serves as a natural *Feynman-Kac branching checkpoint*. At this temporal boundary, the continuous backbone flow proposes an overcomplete candidate pool of K backbone particles:

$$\mathcal{B} = \{B_1, B_2, \ldots, B_K\} \sim p_{\theta^\star}(\cdot \mid \mathcal{S}, \boldsymbol{\mathcal{G}}, \boldsymbol{c}, \boldsymbol{m}) \tag{15}$$

Each scaffold B_i defines an ordered sequence of local frames and peptide torsions. For structural comparison, we extract its C_α Cartesian trace $\mathbf{X}_i = (\mathbf{x}_{i,1}, \ldots, \mathbf{x}_{i,L})^\top \in \mathbb{R}^{L\times 3}$. Because sequential residue ordering is fixed by the primary cyclic sequence, comparisons must maintain invariance under Euclidean rigid-body transformations while strictly prohibiting permutations, cyclic shifts, and improper reflections. We therefore evaluate metric distance via optimal $SO(3)$-constrained Kabsch alignment:

$$d_{ij} = \min_{\substack{\boldsymbol{R}\in SO(3)\\ \mathbf{t}\in\mathbb{R}^3}} \left(\frac{1}{L} \sum_{\ell=1}^{L} \|\boldsymbol{R}\mathbf{x}_{i,\ell} + \mathbf{t} - \mathbf{x}_{j,\ell}\|^2 \right)^{1/2}. \tag{16}$$

Enforcing $\det(R) = +1$ is mathematically vital: it forbids spatial inversions from collapsing enantiomers (such as L- and D-chiral stereoisomers) into identical representations, preserving the stereochemical fidelity of mixed-chirality rings.

Pairwise metric distances are mapped onto a continuous similarity manifold via a Gaussian radial basis kernel parameterized by structural resolution $\sigma > 0$:

$$k_\sigma(i, j) = \exp\left(-\frac{d_{ij}^2}{2\sigma^2}\right). \tag{17}$$

Let $\mathcal{A} \subset \{1, \ldots, K\}$ denote the active subset of selected surviving scaffolds. For any unselected candidate $i \notin \mathcal{A}$, its structural redundancy relative to the colonized subspace is quantified by its *mean kernel crowding*:

$$\rho_i(\mathcal{A}) = \frac{1}{|\mathcal{A}|} \sum_{j\in\mathcal{A}} k_\sigma(i, j). \tag{18}$$

Unlike heuristic single-linkage clustering that only accounts for the nearest neighbor, mean crowding measures the collective local density, exerting a robust repulsive force against over-sampled conformational clusters. We define the selection potential governing inter-stage particle branching as:

$$S_i(\mathcal{A}) = -\gamma \log\left(\epsilon + \rho_i(\mathcal{A})\right), \tag{19}$$

where $\gamma = 1$ is a scaling exponent and $\epsilon = 10^{-8}$ prevents numerical singularity. This potential is purely geometric, deliberately excluding force-field approximations to preserve uncorrupted structural exploration.

C.2 Sequential Interacting Feynman-Kac Dynamics and Temperature Modulation

We formulate the reduction of the overcomplete backbone pool into an optimal branched seed ensemble $\mathcal{A} \subset \mathcal{B}$ ($|\mathcal{A}| = K_0 \ll K$) as an interacting particle system governed by discrete-time Feynman-Kac path measures.

Let $r \in \{1, \dots, K_0\}$ denote the sequential branching step, and let $\mathcal{A}_{r-1}$ represent the historical particle ensemble selected up to step $r-1$. At step r, each candidate particle $i \notin \mathcal{A}_{r-1}$ is assigned a path-dependent fitness weight governed by an exploratory temperature $\tau > 0$:

$$G_r(\mathcal{A}_{r-1}, i) = \exp\left(\frac{S_i(\mathcal{A}_{r-1})}{\tau}\right) = (\epsilon + \rho_i(\mathcal{A}_{r-1}))^{-\gamma/\tau}. \tag{20}$$

The normalized conditional transition probability for selecting the r-th surviving particle is:

$$P_r(i \mid \mathcal{A}_{r-1}) = \frac{(\epsilon + \rho_i(\mathcal{A}_{r-1}))^{-\gamma/\tau}}{\sum_{j \notin \mathcal{A}_{r-1}} (\epsilon + \rho_j(\mathcal{A}_{r-1}))^{-\gamma/\tau}}. \tag{21}$$

Drawing $i_r \sim P_r(\cdot \mid \mathcal{A}_{r-1})$ augments the active ensemble to $\mathcal{A}_r = \mathcal{A}_{r-1} \cup \{i_r\}$, after which crowding scores and weights are dynamically refreshed until $|\mathcal{A}_{K_0}| = K_0$. At the initialization step ($r = 1, \mathcal{A}_0 = \varnothing$), the ensemble possesses no reference anchors; thus, i_1 is drawn uniformly at random from $\mathcal{B}$. The total probability of realizing the ordered sequence of surviving particles $(i_1, \dots, i_{K_0})$ is:

$$\mathbb{P}(i_1, \dots, i_{K_0}) = \frac{1}{K} \prod_{r=2}^{K_0} \frac{G_r(\mathcal{A}_{r-1}, i_r)}{\sum_{j \notin \mathcal{A}_{r-1}} G_r(\mathcal{A}_{r-1}, j)}. \tag{22}$$

Because the normalization at step r depends nonlinearly on the historical realization path $\mathcal{A}_{r-1}$, this process constitutes an adaptive self-interacting particle system (McKean-Vlasov type) rather than a static Gibbs sampler over independent configurations.

Temperature Limits and Entropic Regularization. The temperature parameter τ smoothly governs the sampling behavior between two classical boundaries. In the high-temperature limit ($\tau \to \infty$), relative selection odds approach unity ($P_r(a \mid \mathcal{A})/P_r(b \mid \mathcal{A}) \to 1$), degenerating into uniform random selection devoid of diversity awareness. Conversely, in the zero-temperature limit ($\tau \to 0^+$), probability mass concentrates exclusively on the maximally isolated outlier $i_r = \arg\min_{i \notin \mathcal{A}_{r-1}} \rho_i(\mathcal{A}_{r-1})$, corresponding to deterministic greedy dispersion. In continuous generative modeling, backbones at the extreme geometric periphery often represent unphysical artifacts (e.g., strained bond angles or un-cyclizable open geometries). Deterministic greedy dispersion acts as an extreme-value filter that is acutely vulnerable to trapping by these artifacts. Setting a finite temperature ($\tau = 2.5$) provides entropic regularization: it optimizes structural dispersion under a Kullback–Leibler divergence penalty relative to the generative proposal, strongly favoring diverse modes while retaining viable physical plausibility.

Resolution Sensitivity of Kernel Bandwidth σ. The bandwidth σ determines the structural granularity at which distinct conformational basins are recognized. In the small bandwidth regime ($\sigma \to 0$), pairwise distances quickly exceed $\sigma\sqrt{2\log(10^8)} \approx 6.07\sigma$, causing crowding factors to drop below ϵ and potentials to saturate at $-\gamma \log \epsilon$, which collapses selection into arbitrary tie-breaking. In the large bandwidth regime ($\sigma \to \infty$), second-order Taylor expansion yields $k_\sigma(i,j) \approx 1 - d_{ij}^2/(2\sigma^2)$; crowding differences scale as $\mathcal{O}(\sigma^{-2})$, which flattens selection probabilities and neutralizes repulsive steering. Setting $\sigma = 1.75\,\text{Å}$ provides optimal resolution matching the characteristic RMSD separation between distinct cyclic peptide backbone clusters.

C.3 Statistical Analysis: Why Feynman-Kac Branching Unlocks Multi-Basin Landscapes

Mitigating Sample Degeneracy Across Multi-Funneled Landscapes. Continuous flow models trained via maximum likelihood or flow matching naturally concentrate probability density around primary modes supported by training data. For cyclic peptides with N-methylations and D-amino

acids, the conformational landscape is partitioned into multiple distinct free-energy basins separated by towering torsional isomerization barriers (e.g., 18–22 kcal/mol for *cis*/*trans* peptide bond isomerization). Standard ancestral sampling from an unsteered generator $p_\theta(B)$ suffers from severe sample degeneracy, repeatedly drawing configurations from the dominant low-barrier funnel. Because subsequent gradient-based force-field minimization cannot cross high torsional barriers, unsteered sampling yields redundant structures. Applying the repulsive potential $G_r(\mathcal{A}_{r-1}, i)$ at the inter-stage interface reshapes the proposal measure into a self-avoiding distribution that actively repels particles from colonized basins, ensuring each surviving backbone seeds a distinct topological family.

Connection to Repulsive Gibbs Ensembles and Determinantal Point Processes. Statistically, FKS functions as an efficient sequential approximation to repulsive Determinantal Point Processes (DPPs), where the inclusion probability of subset $\mathcal{A}$ scales with the determinant of its marginal kernel matrix $\mathcal{P}_{\mathrm{DPP}}(\mathcal{A}) \propto \det(\mathbf{K}_{\mathcal{A}})$. While exact DPP sampling requires spectral decomposition with prohibitive $\mathcal{O}(K^3)$ complexity, FKS interprets crowding $\rho_i(\mathcal{A})$ as the mean-field potential of a continuous repulsive Coulomb gas:

$$U(\mathbf{X}_i; \mathcal{A}) = \frac{1}{|\mathcal{A}|} \sum_{j \in \mathcal{A}} \exp\left(-\frac{\|\mathbf{X}_i - \mathbf{X}_j\|^2}{2\sigma^2} \right). \tag{23}$$

The transition probability in Equation 21 sequentially places particles into regions of minimal interaction energy, achieving near-optimal manifold dispersion with linear per-step complexity $\mathcal{O}(K \cdot K_0)$.

Thermodynamic Basin Coverage and Extreme Value Statistics. At the Stage-1/Stage-2 interface, each of the K_0 selected backbone scaffolds $B_k \in \mathcal{A}$ undergoes rotamer expansion in Stage-2, branching into M independent side-chain trajectories via $Y_{k,m} \sim p_{\mathrm{SC}}(\cdot \mid B_k)$. The resulting full-atom candidates are subsequently relaxed to local potential energy minima:

$$\mathbf{X}^\star_{k,m} = \mathrm{Relax}\left(\mathrm{FK}(B_k, Y_{k,m})\right). \tag{24}$$

Let $\mathcal{H}(\mathbf{X})$ denote the physical force-field energy surface, where each local minimum b is surrounded by an attraction basin Ω_b. The global objective is discovering the thermodynamic ground state:

$$E_{\min} = \min_{k \in \mathcal{A},\, m \in \{1,\dots,M\}} \mathcal{H}(\mathbf{X}^\star_{k,m}). \tag{25}$$

Because geometric selection potential $S_i(\mathcal{A})$ operates without direct access to $\mathcal{H}$, geometric dispersion does not inherently guarantee monotonic energy reduction for any individual scaffold. However, FKS provides a decisive advantage under extreme value statistics. Let $\Omega^\star$ denote the attraction basin of the thermodynamic global minimum. Under unsteered independent sampling, the failure probability of missing $\Omega^\star$ decays with the effective number of independent basins sampled:

$$\mathbb{P}\left(\bigcup_{k=1}^{K_0} \{B_k\} \cap \Omega^\star = \varnothing \right) \approx \prod_{k=1}^{K_{\mathrm{eff}}} \left(1 - \mathbb{P}(B \in \Omega^\star)\right). \tag{26}$$

In unsteered sampling, clustering collapses the effective sample budget ($K_{\mathrm{eff}} \ll K_0$). By repelling backbone particles from colonized funnels prior to rotamer branching, FKS maximizes the union volume of covered attraction basins:

$$\mathrm{Vol}\left(\bigcup_{i \in \mathcal{A}} \Omega(B_i) \right) \approx \sum_{i \in \mathcal{A}} \mathrm{Vol}\left(\Omega(B_i)\right), \tag{27}$$

driving $K_{\mathrm{eff}} \to K_0$. This maximized basin coverage drastically elevates the hit rate for reaching deep, narrow energy wells, reliably lowering the ensemble minimum energy $E_{\min}$ across complex noncanonical cyclic peptides.

D Dataset Curation and Data Processing

This section details the curation pipelines, filtering protocols, simulation settings, and dataset partition strategies across all canonical, backbone-modified, and external noncanonical benchmarks. A comprehensive overview of dataset roles, references, and composition is summarized in Table 5.

Table 5: **Comprehensive overview of benchmark datasets across training, validation, and evaluation in MuCO2.** The overarching CPSea repository comprises five mutually disjoint subsets. The primary CPSea subset (and its modified counterpart CPSea MeD) is employed for model training and validation, whereas the remaining four subsets serve as strictly held-out evaluation benchmarks.

Dataset / Subset	Perturbation Axis	Records (n)	Length (L)	Reference Modality / Role
Canonical CPSea Benchmarks (Inherited from MuCO)				
CPSea subset	Canonical L-peptides	75,337	7–16	Inherited from MuCO; Backbone base training / validation
CPTrans	Canonical L-peptides	6,590	7–16	Independent held-out test set (Membrane permeability)
CPBind	Canonical L-peptides	3,194	7–16	Independent held-out test set (Target binding)
CPCore	Canonical L-peptides	263	7–16	Independent held-out test set (Structurally diverse core)
CPSea-PDB	Canonical L-peptides	85	7–16	Independent held-out test set (Experimental PDB native)
Backbone-modified CPSea MeD Benchmarks				
CPSea MeD subset	N-Me / D-stereochemistry	9,022	7–16	10 ns explicit-solvent MD; Curriculum training / validation
CPTrans MeD	N-Me / D-stereochemistry	799	7–16	10 ns explicit-solvent MD; Held-out test set
CPBind MeD	N-Me / D-stereochemistry	392	7–16	10 ns explicit-solvent MD; Held-out test set
CPCore MeD	N-Me / D-stereochemistry	32	7–16	10 ns explicit-solvent MD; Held-out test set
CPSea-PDB MeD	N-Me / D-stereochemistry	10	7–16	10 ns explicit-solvent MD; Held-out test set
External Noncanonical Benchmarks				
HighFold2-Compatible	Open sidechain	10	6–34	Solution NMR (9 entries) + CycPeptMPDB (1 entry)
CycPeptMPDB MeD	Backbone N-Me + D-aa	947	6–7	50 ns explicit-solvent MD (CycPeptMPDB-4D, water/hexane)
CycPeptMPDB Mixed	N-Me / D-aa / Open sidechain	206	6	50 ns explicit-solvent MD (CycPeptMPDB-4D, water/hexane)

D.1 Canonical CPSea Repository and Benchmark Subsets

Following the setting of MuCO (Wang et al., 2026), we inherit the curated cyclic peptide repository CPSea (Yang et al., 2026). In our nomenclature, the term **CPSea** denotes either the overarching five-subset collection or specifically its primary training subset, depending on context. Crucially, the entire CPSea repository is partitioned into five mutually disjoint subsets categorized by functional domains and structural sources: the large-scale base subset (**CPSea subset**), membrane-permeable macrocycles (**CPTrans**), target-binding binders (**CPBind**), structurally clustered representatives (**CPCore**), and high-resolution crystal structures from the Protein Data Bank (**CPSea-PDB**).

To guarantee zero data leakage across functional domains, these subsets are assigned strictly non-overlapping roles. The **CPSea subset** (75,337 entries) is exclusively utilized for base model training and validation, and its derived variants likewise form our curriculum training pool (**CPSea MeD subset**). Conversely, the remaining four subsets, including **CPTrans**, **CPBind**, **CPCore**, and **CPSea-PDB**, along with their respective MeD counterparts, are strictly quarantined as held-out test sets. This ensures an unbiased evaluation of generalizability across distinct thermodynamic distributions, permeability-driven designs, and experimental native folds.

D.2 Backbone-modified CPSea MeD Dataset and Simulation Protocol

To benchmark backbone conformational adaptation under dense steric and stereochemical transitions, we constructed an in-house augmented dataset (**CPSea MeD**) by performing explicit-solvent molecular dynamics (MD) on systematically modified parent structures.

Stratified Sampling and Modification Rules. Across the five canonical CPSea subsets, we analyzed 85,469 cyclic parent records and identified three valid covalent cyclization topologies: head-to-tail peptide bonds (`LIG1`), disulfide bridges (`DISU`), and Lys–Asp isopeptide linkages (`DKAM`). After filtering out 3,326 records lacking feasible modification sites, a deterministic stratified sampling scheme targeting approximately 12% of the parent collection was executed across all *subset* $\times$ *cyclization topology* strata, selecting 10,256 targets. Each parent scaffold received up to two site-specific modifications: **(1)** backbone N-methylation on an Alanine residue (excluding terminal positions, ring-closing anchors, and Proline-adjacent residues), and **(2)** stereochemical inversion

to a D-enantiomer on an Ala, Phe, Ile, or Val residue (mapped to DALA, DPHE, DILE, or DVAL). Modifications on the same residue as well as configurations where a D-residue directly precedes an N-methylated position were disallowed. A total of 10,255 systems were successfully constructed and simulated (99.990% execution rate).

Force-Field Parameterization and Local NME Patch. All-atom topologies were built via psfgen utilizing CHARMM36 protein parameters (Huang & MacKerell Jr, 2013) and D-amino acid stream topologies. Because standard CHARMM releases lack a general-purpose backbone N-methylation patch for internal amide linkages, we implemented an NME patch derived by analogy from the quantum-mechanically parameterized CHARMM MAL model. To preserve strict group electrical neutrality ($-0.16\,\mathrm{e}$, matching the native peptide NH group) without disrupting adjacent C_α and side-chain charges across arbitrary sequences, the $+0.01\,\mathrm{e}$ partial charge assigned to C_α in MAL was transferred to the methyl carbon (CM, set to $-0.08\,\mathrm{e}$), paired with the backbone nitrogen (N, $-0.35\,\mathrm{e}$) and three methyl hydrogens (HM, $+0.09\,\mathrm{e}$ each). Bonded parameters were transferred from chemically equivalent CHARMM terms. All parameter and topology stream files are provided in the Supplementary Information.

Explicit-Solvent Molecular Dynamics Protocol. Simulations were performed using GROMACS 2026.0. Each modified cyclic peptide was immersed in a rhombic dodecahedron box with a minimum solute-to-box clearance of $1.0\,\mathrm{nm}$ using the CHARMM-modified TIP3P water model and neutralized with $0.15\,\mathrm{mol/L}$ NaCl. The multi-stage protocol proceeded sequentially: **Energy Minimization** was performed via steepest descent for up to 50,000 steps with a force threshold of $1{,}000\,\mathrm{kJ\,mol^{-1}\,nm^{-1}}$. **Equilibration** comprised a $50\,\mathrm{ps}$ NVT run at $300\,\mathrm{K}$ using the velocity-rescaling thermostat ($\tau_t = 1.0\,\mathrm{ps}$, followed by a $50\,\mathrm{ps}$ NPT run at $300\,\mathrm{K}$ and $1.0\,\mathrm{bar}$ with the isotropic C-rescale barostat ($\tau_p = 5.0\,\mathrm{ps}$). **Production MD** was executed for $10\,\mathrm{ns}$ (5,000,000 steps with a $2\,\mathrm{fs}$ time step) under identical NPT controls with LINCS constraints applied to hydrogen bonds. Long-range electrostatics were handled by PME (real-space cutoff $1.2\,\mathrm{nm}$), and van der Waals interactions were force-switched between 1.0 and $1.2\,\mathrm{nm}$. Trajectories were saved every $10\,\mathrm{ps}$ (1,001 frames per system, yielding 10,357,550 solute frames over $100.255\,\mu\mathrm{s}$ of aggregate simulation). Across all 10,255 completed runs, cyclic bonds remained intact without distance violations (mean lengths: $1.356\,\text{Å}$ for head-to-tail, $2.030\,\text{Å}$ for disulfide, and $1.348\,\text{Å}$ for isopeptide linkages). The final frame serves as the physical reference.

Dataset Splits for Curriculum Learning and Testing. The 10,255 completed systems strictly follow the split of their parent sources. Specifically, the 9,022 systems derived from the **CPSea** parent pool serve as warmup and progressive training references within the curriculum learning loop. In contrast, the remaining 1,233 systems originating from the functional parent test sets—**CPTrans MeD** (799), **CPBind MeD** (392), **CPCore MeD** (32), and **CPSea-PDB MeD** (10)—are reserved strictly as held-out test benchmarks for evaluating model generalization on backbone-modified cyclic peptides (Table 6).

D.3 HighFold2-Compatible Noncanonical Benchmark

To evaluate noncanonical side-chain modeling against specialized predictors, we curated the **HighFold2-Compatible** benchmark. This set comprises macrocycles whose noncanonical building blocks fall within HighFold2's vocabulary (Zhu et al., 2025), cyclize via head-to-tail amide or disulfide linkages, and exclude N-methylation. The final benchmark consists of **10 macrocyclic targets** ($L = 6$–34 residues):

Experimental NMR Entries (9 targets): Curated from the official HighFold2 noncanonical cyclic test set (cyclic_unAA_test). Each entry possesses confirmed head-to-tail cyclization (7 also contain internal disulfide bonds) and supported noncanonical α-amino acids, specifically α-aminobutyric acid (ABA) and hydroxyproline (HYP): PDB IDs 1T9E, 2J15, 2M1P, 2M2G, 2M2H, 2M2S, 2M2X, 2MSQ, and 2N8E. All entries are solution NMR structures; the first coordinate model is taken as the reference.

CycPeptMPDB Entry (1 target): Identified from CycPeptMPDB by selecting candidates with D-amino acids and noncanonical side chains without N-methylation whose sequence maps fully into HighFold2's dictionary. A single hexapeptide, CP2337 (*cyclo*-[L-dL-Abu-L-dP-F]), satisfied all criteria. Explicit-solvent MD from CycPeptMPDB-4D serves as its reference.

Table 6: **Compositional and chemical perturbation breakdown of CPSea MeD and CycPeptMPDB benchmarks.** The top panel details the modification distributions across the CPSea MeD subsets. The bottom panel breaks down the co-occurrence of structural modifications across the external CycPeptMPDB MeD and Mixed evaluation sets.

Benchmark / Subset	Single N-Me	Single D-aa	Mixed (N-Me + D)	Total Records	Primary Dataset Role
CPSea MeD Benchmark					
CPSea MeD	395	3,699	4,928	9,022	Curriculum Training / Validation
CPTrans MeD	18	434	347	799	Held-out Evaluation
CPBind MeD	7	124	261	392	Held-out Evaluation
CPCore MeD	0	13	19	32	Held-out Evaluation
CPSea-PDB MeD	0	3	7	10	Held-out Evaluation
Subtotal	*420* (4.10%)	*4,273* (41.67%)	*5,562* (54.24%)	**10,255**	–
CycPeptMPDB Benchmarks					
CycPeptMPDB MeD (All Circular 6–7 mers)					
Standard sidechain + D-aa + N-Me	–	–	947	947	Held-out Evaluation (946 unique mol.)
CycPeptMPDB Mixed (All Circular 6-mers)					
Sidechain + D-aa + N-Me + Peptoid	–	–	128	128	Held-out Evaluation
Sidechain + D-aa + Peptoid (no N-Me)	–	–	32	32	Held-out Evaluation
Sidechain + D-aa + N-Me (no Peptoid)	–	–	20	20	Held-out Evaluation
Sidechain + Peptoid + N-Me (no D-aa)	–	–	10	10	Held-out Evaluation
Sidechain + D-aa only	–	–	15	15	Held-out Evaluation
Sidechain only	–	–	1	1	Held-out Evaluation
Subtotal	–	–	–	**206**	Held-out Evaluation (202 unique mol.)

D.4 CycPeptMPDB MeD and Mixed External Benchmarks

To benchmark generalizability on natural macrocyclic chemistry, we curated two evaluation sets from the 5,160 records in CycPeptMPDB-4D v3 (Liu et al., 2026):

Perturbation Filtering. We categorized macrocyclic chemical space into five axes: backbone N-substitution (N), D-amino acids (D), noncanonical α-sidechains (S), backbone extensions (B, *e.g.*, β-homophenylalanine), and non-peptidic/mono units (L). Entries containing backbone extensions (B) or non-peptidic/mono units (L) were excluded.

CycPeptMPDB MeD (N (Methyl) $\wedge$ $D \wedge \neg S \wedge \neg B \wedge \neg L$): Starting from 3,375 candidates, we retained circular head-to-tail macrocycles possessing standard α-sidechains with both backbone N-methylation and D-amino acids. Candidates bearing non-methyl N-modifications (1,053 peptoid entries) or ester-lariat topologies (1,375 depsipeptides) were removed. The final **CycPeptMPDB MeD** benchmark contains **947 records** (946 unique canonical isomeric molecules), comprising 386 hexapeptides and 561 heptapeptides.

CycPeptMPDB Mixed ($S \wedge \neg B \wedge \neg L \setminus \{$CP2337$\}$): To test compounding modifications across multiple perturbation axes, we curated all macrocycles featuring noncanonical α-sidechains. Excluding `CP2337` (allocated to the Compatible set), this yields **206 records** (202 unique isomeric molecules), all of which are hexapeptides ($L = 6$). As detailed in Table 6, this benchmark features combinations of noncanonical side chains (*e.g.*, `Abu`, `Cha`, `Nle`), D-amino acids (195/206), and backbone N-modifications (180/206, including 148 N-methylations and 160 peptoid-like N-alkylations).

Both CycPeptMPDB benchmarks are evaluated against 50 ns explicit-solvent MD trajectories in water and hexane provided by CycPeptMPDB-4D.

E METRICS

This section provides the rigorous mathematical formulations, force-field evaluation protocols, and stereochemical criteria summarized in Section 5.1.

E.1 VACUUM POTENTIAL ENERGY

We compute the all-atom intramolecular potential energy in vacuum using the CHARMM36 additive force field:

$$E_{\text{vac}}(\boldsymbol{X}) = E_{\text{bond}} + E_{\text{angle}} + E_{\text{torsion}} + E_{\text{improper}} + E_{\text{vdW}} + E_{\text{electrostatic}}, \tag{28}$$

expressed in kJ/mol. Lower potential energy indicates reduced steric crowding, favorable hydrogen-bonding networks, and relaxed valence geometries.

Pairwise Context and Size-Dependence. Because total potential energy scales extensively with peptide sequence length, chemical composition, and net charge, absolute energies cannot be compared across different peptides (e.g., a 30-mer with $-2000\,\text{kJ}/\text{mol}$ is not intrinsically superior to a 6-mer with $+500\,\text{kJ}/\text{mol}$). Meaningful physical comparisons are established strictly in a pairwise setting over identical target sequences:

$$\Delta E_i = E_i^{(\text{Method } A)} - E_i^{(\text{Method } B)}. \tag{29}$$

E.2 IMPLICIT SOLVENT WATER SCORE

To account for dielectric shielding and polar desolvation in physiological environments, we compute an implicit-solvent energy score based on the Generalized Born model with simple surface area approximations (OBC2/ACE):

$$E_{\text{water}}(\boldsymbol{X}) = E_{\text{vac}}(\boldsymbol{X}) + \Delta G_{\text{polar}}(\boldsymbol{X}) + \Delta G_{\text{nonpolar}}(\boldsymbol{X}). \tag{30}$$

Here, ΔG_{polar} is approximated by the Onufriev-Bashford-Case (OBC2) GB model, and $\Delta G_{\text{nonpolar}} = \gamma \cdot \text{SASA}$ captures non-polar hydrophobic burial via solvent-accessible surface area.

Distinction from Solvation Free Energy. We emphasize that E_{water} is a *total effective score* encompassing both intramolecular vacuum potential and solvation free energy:

$$E_{\text{water}} = E_{\text{vac}} + \Delta G_{\text{solvation}}. \tag{31}$$

Hence, negative values reflect favorable combined vacuum and polar interactions rather than solitary solvation stabilization. We refer to this metric as a "score" rather than an experimental free energy, as it omits explicit-water hydrogen bond rearrangements and conformational entropy.

E.3 SEQUENCE-ORDERED CENTRAL CARBON RMSD

For peptides with available reference coordinates $\boldsymbol{Y} = \{\boldsymbol{y}_\ell\}_{\ell=1}^{L}$, we evaluate structural recovery on the L ordered Cα atoms:

$$\text{RMSD}(\boldsymbol{X}, \boldsymbol{Y}) = \min_{\substack{\boldsymbol{R}\in SO(3) \\ \boldsymbol{t}\in\mathbb{R}^3}} \sqrt{\frac{1}{L}\sum_{\ell=1}^{L} \|\boldsymbol{R}\boldsymbol{x}_\ell + \boldsymbol{t} - \boldsymbol{y}_\ell\|^2}, \tag{32}$$

where $\boldsymbol{x}_\ell$ and $\boldsymbol{y}_\ell$ denote the Cα coordinates of the ℓ-th residue of the generated and reference structures, respectively.

- **Rigid Invariance:** Optimal superposition is solved via the Kabsch algorithm restricted strictly to proper rotations ($\det(R) = +1$). Spatial reflections and residue permutations are forbidden to prevent collapsing enantiomers.
- **Ordered vs. Cyclic Superposition:** Unlike permissive cyclic-oracle alignment (which minimizes over cyclic sequence shifts), our evaluation enforces 1-to-1 sequential residue correspondence, providing a stringent benchmark of fold recapitulation.

E.4 Residue-Level Chirality Preservation

Because force-field relaxation algorithms lack explicit non-inversion barriers, unphysical local clashes in generated backbones are frequently alleviated by spontaneously inverting $\mathrm{C}\alpha$ chiral centers. We therefore independently evaluate the survival of intended stereochemistry directly from relaxed 3D coordinates.

Micro-Averaged Metric. For the j-th conformation of target i with L_i^* non-Glycine chiral centers, the overall Chirality metric is micro-averaged across all evaluated residues:

$$\text{Chirality} = \frac{\sum_i \sum_{j \in \mathcal{V}_i} \sum_{\ell=1}^{L_i^*} \mathbb{I}(\chi_{ij\ell} = \chi_{i\ell}^{\text{target}})}{\sum_i |\mathcal{V}_i| L_i^*} \times 100\%, \tag{33}$$

where $\chi_{ij\ell} \in \{\mathrm{L}, \mathrm{D}\}$ is the evaluated stereochemical configuration. We report this micro-average across all chiral residues, rather than the proportion of entirely error-free conformations, to provide a granular measure of stereochemical stability.

Coordinate-Based Stereochemical Perception. Rather than relying on metadata tags in exported files (which often stay static despite atomic movement), stereocenters are determined from 3D coordinates using the signed tetrahedral volume:

$$V = [(\boldsymbol{x}_{\mathrm{N}} - \boldsymbol{x}_{\mathrm{CA}}) \times (\boldsymbol{x}_{\mathrm{C}} - \boldsymbol{x}_{\mathrm{CA}})] \cdot (\boldsymbol{x}_{\mathrm{CB}} - \boldsymbol{x}_{\mathrm{CA}}). \tag{34}$$

Under standard atomic indexing conventions, $V > 0$ defines L-amino acids, while $V < 0$ corresponds to D-enantiomers (with achiral Glycine omitted and degenerate near-planar tetrahedra $|V| < 10^{-4}$ flagged as unphysical). For noncanonical building blocks featuring non-standard substituents, configurations are verified via 3D Cahn-Ingold-Prelog (CIP) perception.

F Additional Experimental Results

Table 7: **Detailed performance breakdown on CPSea and CPSea MeD functional subsets.** Evaluation across CPTrans (membrane permeability), CPBind (target binding), CPCore, and CPSea-PDB sub-benchmarks alongside their N-methylated and D-stereochemical variants (MeD). MuCO2 consistently attains lower physical relaxation energies, enhanced structural recovery, and near-perfect chirality preservation ($\geq 98.6\%$), demonstrating uniform generalizability across distinct functional macrocycle classes.

Dataset	Method	Success (%) ↑		E_{vac}(kJ/mol) ↓		E_{water}(kJ/mol) ↓		Cα RMSD ↓		Chirality ↑
		@1	@10	@1	@10	@1	@10	@1	@10	(%)
CPTrans	GT (Ref.)	100.00	–	6.82	–	-769.05	–	0.00	–	100.00
	HighFold2	**98.91**	–	138.36	–	-529.31	–	3.78	–	96.41
	HighFold-MeD	96.30	–	125.25	–	-482.92	–	**3.62**	–	97.94
	HighFold-MeD2	93.41	97.89	98.37	15.59	-514.35	-643.89	3.92	3.40	84.27
	MuCO	94.23	99.97	53.91	-1.58	-583.59	-804.12	3.69	**3.26**	63.35
	MuCO2 (ours)	96.22	**100.00**	**44.29**	**-15.74**	**-589.90**	**-828.32**	3.64	**3.26**	**99.97**
CPBind	GT (Ref.)	100.00	–	-63.21	–	-966.94	–	0.00	–	100.00
	HighFold2	**98.72**	–	54.46	–	-795.12	–	3.89	–	96.21
	HighFold-MeD	92.20	–	55.94	–	-764.95	–	3.77	–	97.49
	HighFold-MeD2	88.95	95.21	-2.49	-34.94	-740.19	-824.31	3.73	3.46	87.42
	MuCO	93.52	**100.00**	-19.29	-58.12	-893.39	-1034.82	3.74	3.17	59.13
	MuCO2 (ours)	96.06	**100.00**	**-27.34**	**-64.53**	**-924.23**	**-1047.85**	**3.72**	**3.05**	**99.64**
CPCore	GT (Ref.)	100.00	–	48.41	–	-741.38	–	0.00	–	100.00
	HighFold2	**99.62**	–	174.32	–	-384.49	–	3.79	–	95.32
	HighFold-MeD	94.30	–	172.84	–	-441.95	–	3.85	–	95.97
	HighFold-MeD2	98.86	99.62	128.85	73.48	-579.03	-664.70	3.87	3.19	84.95
	MuCO	95.06	**100.00**	97.32	42.19	-561.49	-793.45	**3.69**	3.21	61.34
	MuCO2 (ours)	99.24	**100.00**	**83.58**	**28.83**	**-627.55**	**-813.45**	3.71	**3.16**	**99.65**
CPSea-PDB	GT (Ref.)	100.00	–	-183.96	–	-1294.19	–	0.00	–	100.00
	HighFold2	**100.00**	–	-43.21	–	-837.33	–	3.87	–	97.42
	HighFold-MeD	92.94	–	-47.85	–	-810.56	–	3.88	–	95.54
	HighFold-MeD2	95.29	98.82	-93.12	-141.50	-904.87	-1102.41	3.54	3.10	83.13
	MuCO	94.12	**100.00**	-160.68	-212.32	-1137.20	-1397.12	3.90	3.03	55.56
	MuCO2 (ours)	96.47	**100.00**	**-193.34**	**-273.57**	**-1157.92**	**-1415.06**	**3.51**	**2.87**	**99.86**
CPTrans MeD	GT (MD Ref.)	100.00	–	457.19	–	-569.78	–	0.00	–	100.00
	HighFold-MeD	25.41	–	**774.08**	–	**-170.28**	–	3.76	–	96.59
	HighFold-MeD2	28.91	75.47	794.96	576.97	-149.56	-395.92	3.66	3.17	81.76
	MuCO2 (ours)	**30.91**	**90.11**	784.46	**528.81**	-168.96	**-559.69**	**3.18**	**2.53**	**98.64**
CPBind MeD	GT (MD Ref.)	100.00	–	365.42	–	-627.71	–	0.00	–	100.00
	HighFold-MeD	16.58	–	736.45	–	-202.99	–	4.10	–	96.81
	HighFold-MeD2	11.48	57.40	677.89	562.56	-245.69	-284.47	3.86	3.36	77.73
	MuCO2 (ours)	**32.14**	**71.43**	**646.79**	**449.88**	**-248.42**	**-445.48**	**3.42**	**2.80**	**99.50**
CPCore MeD	GT (MD Ref.)	100.00	–	690.02	–	-546.54	–	0.00	–	100.00
	HighFold-MeD	21.88	–	**937.52**	–	**-113.93**	–	3.90	–	**100.00**
	HighFold-MeD2	25.00	78.13	1009.80	787.27	-87.18	-138.00	3.81	3.31	81.33
	MuCO2 (ours)	**31.25**	**93.75**	1159.78	**704.23**	-25.21	**-298.96**	**3.28**	**2.88**	99.33
CPSea-PDB MeD	GT (MD Ref.)	100.00	–	-468.35	–	-1646.38	–	0.00	–	100.00
	HighFold-MeD	**40.00**	–	-242.58	–	-1385.29	–	3.62	–	**100.00**
	HighFold-MeD2	30.00	60.00	-346.83	-517.99	-1412.48	-1560.10	3.41	3.23	82.86
	MuCO2 (ours)	**40.00**	**70.00**	**-391.25**	**-582.60**	**-1583.29**	**-1639.82**	**3.22**	**3.00**	**100.00**

G Additional Ablation Results

Table 8: **Ablation study comparing SE(3) flow and Quaternion Flow (QFlow) architectures.** Evaluation of MuCO2-SE(3) and MuCO2-Q across canonical, backbone-mutated, and side-chain-modified datasets. MuCO2-Q matches or outperforms the Riemannian SE(3) counterpart across physical energy, cyclization success rate, and structural fidelity while enabling substantially accelerated backbone sampling.

Dataset	Generator	Success (%)↑		E_{vac}(kJ/mol)↓		E_{water}(kJ/mol)↓		C_α-RMSD↓		Chirality
	of MuCO2	@1	@10	@1	@10	@1	@10	@1	@10	(%)↑
CPSea	SE(3) Flow	95.73	99.97	25.60	-27.45	-665.72	-832.45	3.68	3.23	99.28
	QFlow	**96.25**	**100.00**	**20.81**	**-32.13**	**-700.86**	**-902.06**	**3.67**	**3.19**	**99.86**
CPSea MeD	SE(3) Flow	**31.63**	82.24	904.67	522.93	-96.89	-521.16	3.41	3.18	98.77
	QFlow	31.39	**84.10**	**761.37**	**503.94**	**-185.53**	**-528.44**	**3.25**	**2.62**	**98.90**
CycPeptMPDB MeD	SE(3) Flow	42.13	**100.00**	1442.93	**1330.76**	**1306.83**	**1214.62**	1.43	**1.07**	78.37
	QFlow	**67.90**	**100.00**	**1421.32**	1341.44	1307.47	1220.40	**1.35**	**1.07**	**89.21**
HighFold2 Compatible	SE(3) Flow	**90.00**	**100.00**	-1186.12	-1221.67	-3305.74	-3385.23	**3.73**	3.21	98.67
	QFlow	**90.00**	**100.00**	**-1205.08**	**-1235.02**	**-3359.07**	**-3424.84**	3.76	**3.14**	**98.81**
CycPeptMPDB Mixed	SE(3) Flow	**100.00**	**100.00**	686.85	544.57	552.33	**431.59**	0.98	0.63	98.02
	QFlow	**100.00**	**100.00**	**643.24**	**518.49**	**548.33**	447.92	**0.96**	**0.56**	**98.41**

Table 9: **Fine-grained comparison between SE(3) Flow and QFlow across all CPSea subsets.** Performance comparison across canonical and MeD-mutated CPSea subsets. QFlow consistently matches or slightly outperforms the Riemannian SE(3) baseline in structural accuracy and energetic relaxation, demonstrating that quaternion-based flow acceleration incurs no loss in physical fidelity.

Dataset	Generator	Success (%) ↑		E_{vac}(kJ/mol) ↓		E_{water}(kJ/mol) ↓		Cα RMSD ↓		Chirality ↑
	of MuCO2	@1	@10	@1	@10	@1	@10	@1	@10	(%)
CPTrans	SE(3) Flow	95.81	**100.00**	48.73	-7.49	-547.23	-729.04	**3.63**	3.29	99.34
	QFlow	**96.22**	**100.00**	**44.29**	**-15.74**	**-589.90**	**-828.32**	3.64	**3.26**	**99.97**
CPBind	SE(3) Flow	95.34	99.91	-21.47	**-67.24**	-903.48	-1031.34	3.78	3.09	99.12
	QFlow	**96.06**	**100.00**	**-27.34**	-64.53	**-924.23**	**-1047.85**	**3.72**	**3.05**	**99.64**
CPCore	SE(3) Flow	98.48	**100.00**	**82.13**	31.49	-603.74	**-832.84**	3.79	3.32	99.43
	QFlow	**99.24**	**100.00**	83.58	**28.83**	**-627.55**	-813.45	**3.71**	**3.16**	**99.65**
CPSea_PDB	SE(3) Flow	95.29	**100.00**	-189.17	-263.48	**-1162.15**	-1381.66	3.70	3.04	**100.00**
	QFlow	**96.47**	**100.00**	**-193.34**	**-273.57**	-1157.92	**-1415.06**	**3.51**	**2.87**	99.86
CPTrans MeD	SE(3) Flow	**32.92**	87.61	**769.48**	540.79	**-219.46**	-556.24	3.28	3.11	98.61
	QFlow	30.91	**90.11**	784.46	**528.81**	-168.96	**-559.69**	**3.18**	**2.53**	**98.64**
CPBind MeD	SE(3) Flow	28.83	70.66	1183.69	494.44	180.10	-422.65	3.72	3.36	99.04
	QFlow	**32.14**	**71.43**	**646.79**	**449.88**	**-248.42**	**-445.48**	**3.42**	**2.80**	**99.50**
CPCore MeD	SE(3) Flow	**31.25**	**93.75**	**1118.86**	**660.73**	**-58.68**	**-326.08**	3.66	3.30	**100.00**
	QFlow	**31.25**	**93.75**	1159.78	704.23	-25.21	-298.96	**3.28**	**2.88**	99.33
CPSea-PDB MeD	SE(3) Flow	**40.00**	**70.00**	-389.94	**-612.00**	-1529.13	**-1715.41**	3.48	**2.97**	**100.00**
	QFlow	**40.00**	**70.00**	**-391.25**	-582.60	**-1583.29**	-1639.82	**3.22**	3.00	**100.00**

Table 10: **Ablation study of curriculum learning across training rounds.** Progressive evaluation of MuCO2 on CPSea MeD and CycPeptMPDB MeD over three consecutive curriculum rounds. Stepwise progression from unadapted baseline (MuCO-Q) to round 3 results in monotonic gains in cyclization success, substantial energy reduction, and robust stereochemical recovery, confirming the necessity of the curriculum sampling loop.

Dataset	Method	Success (%)↑		E_{vac}(kJ/mol)↓		E_{water}(kJ/mol)↓		C_α-RMSD↓		Chirality
		@1	@10	@1	@10	@1	@10	@1	@10	(%)↑
CPSea MeD	MuCO	8.19	28.06	1442.05	1280.92	806.54	677.41	3.66	2.77	64.65
	$\text{MuCO2}_{\text{round 1}}$	25.71	68.94	867.14	757.89	262.87	144.62	3.36	2.67	89.89
	$\text{MuCO2}_{\text{round 2}}$	29.03	73.48	785.92	622.46	70.79	-105.40	**3.23**	2.65	96.45
	$\text{MuCO2}_{\text{round 3}}$	**31.39**	**84.10**	**761.37**	**503.94**	**-185.53**	**-528.44**	3.25	**2.62**	**98.90**
CycPeptMPDB MeD	MuCO	20.31	48.45	1743.12	1637.49	1651.20	1564.04	1.74	1.52	51.59
	$\text{MuCO2}_{\text{round 1}}$	37.45	76.59	1479.31	1398.90	1396.39	1294.01	1.43	1.19	78.23
	$\text{MuCO2}_{\text{round 2}}$	57.52	94.71	**1405.34**	1350.52	1329.76	1245.19	**1.34**	1.12	88.49
	$\text{MuCO2}_{\text{round 3}}$	**67.90**	**100.00**	1421.32	**1341.44**	**1307.47**	**1220.40**	1.35	**1.07**	**89.21**

Table 11: **Sub-benchmark breakdown of the curriculum learning ablation on CPSea MeD.** Progression of sampling metrics across rounds 1 to 3 across each CPSea MeD functional subset. The curriculum strategy consistently enhances physical validity, lowers intramolecular strain, and rectifies stereochemical instability across all sub-distributions.

Dataset	Method	Success (%) ↑		E_{vac}(kJ/mol) ↓		E_{water}(kJ/mol) ↓		Cα RMSD ↓		Chirality ↑
		@1	@10	@1	@10	@1	@10	@1	@10	(%)
CPTrans MeD	MuCO-Q	8.39	27.28	1521.49	1412.38	835.43	694.54	3.59	2.59	68.74
	$\text{MuCO2}_{\text{round 1}}$	26.91	69.59	893.24	774.58	263.61	183.56	3.23	2.61	93.45
	$\text{MuCO2}_{\text{round 2}}$	28.16	73.47	823.95	668.32	83.82	-79.03	**3.13**	2.55	96.32
	$\text{MuCO2}_{\text{round 3}}$	**30.91**	**90.11**	**784.46**	**528.81**	**-168.96**	**-559.69**	3.18	**2.53**	**98.64**
CPBind MeD	MuCO-Q	7.40	29.59	1340.53	1045.69	852.23	701.56	3.81	3.07	57.12
	$\text{MuCO2}_{\text{round 1}}$	23.47	66.07	834.21	753.67	341.73	145.26	3.62	**2.76**	82.31
	$\text{MuCO2}_{\text{round 2}}$	30.61	**72.45**	723.41	535.18	104.36	-116.21	3.43	2.83	96.43
	$\text{MuCO2}_{\text{round 3}}$	**32.14**	71.43	**646.79**	**449.88**	**-248.42**	**-445.48**	**3.42**	2.80	**99.50**
CPCore MeD	MuCO-Q	9.38	31.25	1712.89	1442.55	778.37	341.70	3.71	3.41	54.58
	$\text{MuCO2}_{\text{round 1}}$	18.75	90.63	1232.12	732.13	171.53	-256.73	**3.25**	2.94	92.62
	$\text{MuCO2}_{\text{round 2}}$	28.13	**93.75**	1178.36	745.77	34.04	-264.44	3.31	2.99	98.91
	$\text{MuCO2}_{\text{round 3}}$	**31.25**	**93.75**	**1159.78**	**704.23**	**-25.21**	**-298.96**	3.28	**2.88**	**99.33**
CPSea-PDB MeD	MuCO-Q	20.00	20.00	-153.56	-213.45	-781.63	-912.46	3.31	3.11	65.62
	$\text{MuCO2}_{\text{round 1}}$	**40.00**	60.00	-325.67	-482.47	-1453.31	-1551.87	3.36	3.09	93.23
	$\text{MuCO2}_{\text{round 2}}$	**40.00**	50.00	-361.37	-543.98	**-1586.69**	-1632.41	3.27	**2.97**	**100.00**
	$\text{MuCO2}_{\text{round 3}}$	**40.00**	**70.00**	**-391.25**	**-582.60**	-1583.29	**-1639.82**	**3.22**	3.00	**100.00**

Table 12: **Ablation on cross-modal molecular representation transfer.** Comparison between sequence-only ESM-2 representation and zero-shot Uni-Mol2 embedding transfer on noncanonical side-chain benchmarks. Incorporating 3D conformation-aware Uni-Mol2 embeddings consistently yields lower relaxation energies, improved backbone recovery (C_α RMSD), and enhanced stereochemical accuracy.

Dataset	Embeddings of MuCO2	Success (%) ↑ @1	Success (%) ↑ @10	E_{vac}(kJ/mol) ↓ @1	E_{vac}(kJ/mol) ↓ @10	E_{water}(kJ/mol)↓ @1	E_{water}(kJ/mol)↓ @10	C_α-RMSD↓ @1	C_α-RMSD↓ @10	Chirality (%)↑
HighFold2	ESM-2	80.00	**100.00**	-1112.85	-1205.39	-3212.56	-3318.90	**3.72**	3.17	94.19
Compatible	Uni-mol2	**90.00**	**100.00**	**-1205.08**	**-1235.02**	**-3359.07**	**-3424.84**	3.76	**3.14**	**98.81**
CycPeptMPDB	ESM-2	**100.00**	**100.00**	701.35	561.98	583.25	504.39	1.18	0.64	92.35
Mixed	Uni-mol2	**100.00**	**100.00**	**643.24**	**518.49**	**548.33**	**447.92**	**0.96**	**0.56**	**98.41**

H MuCO2 Samples Visualization

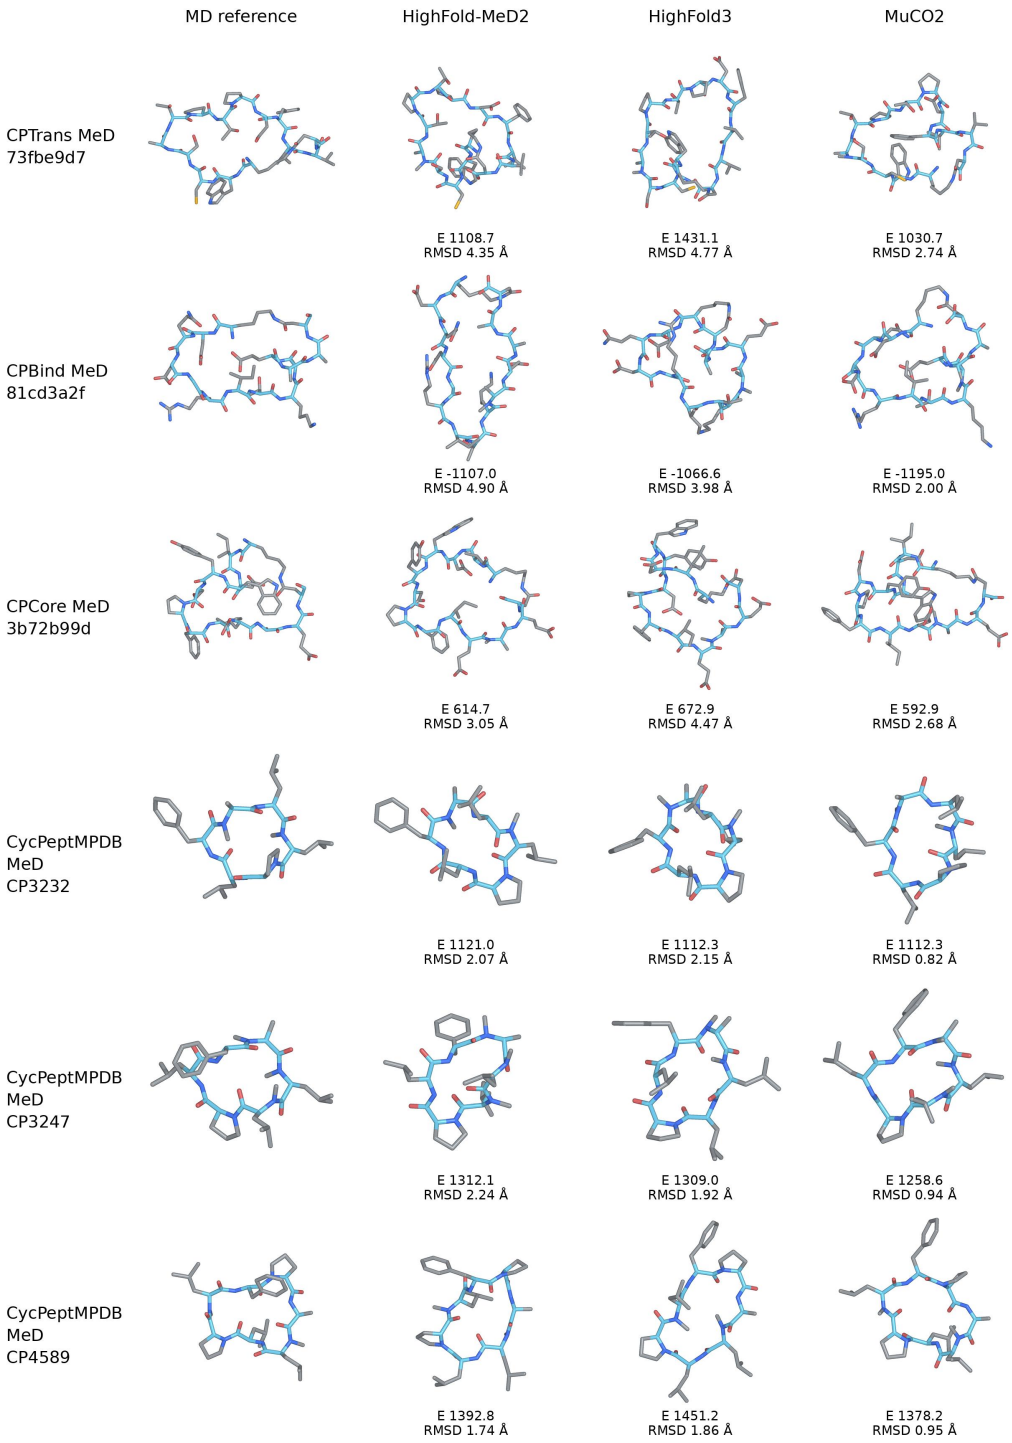


Figure 6: **Structural comparison on backbone-modified cyclic peptides.** Visual comparison of predicted conformations against explicit-solvent MD references on CPSea MeD and CycPeptMPDB MeD (N-methylation and D-amino acids). Values below structures indicate post-relaxation potential energy E (kJ/mol) and C_α-RMSD (Å). MuCO2 achieves superior backbone fidelity and lower physical strain compared to HighFold-MeD2 and HighFold3.

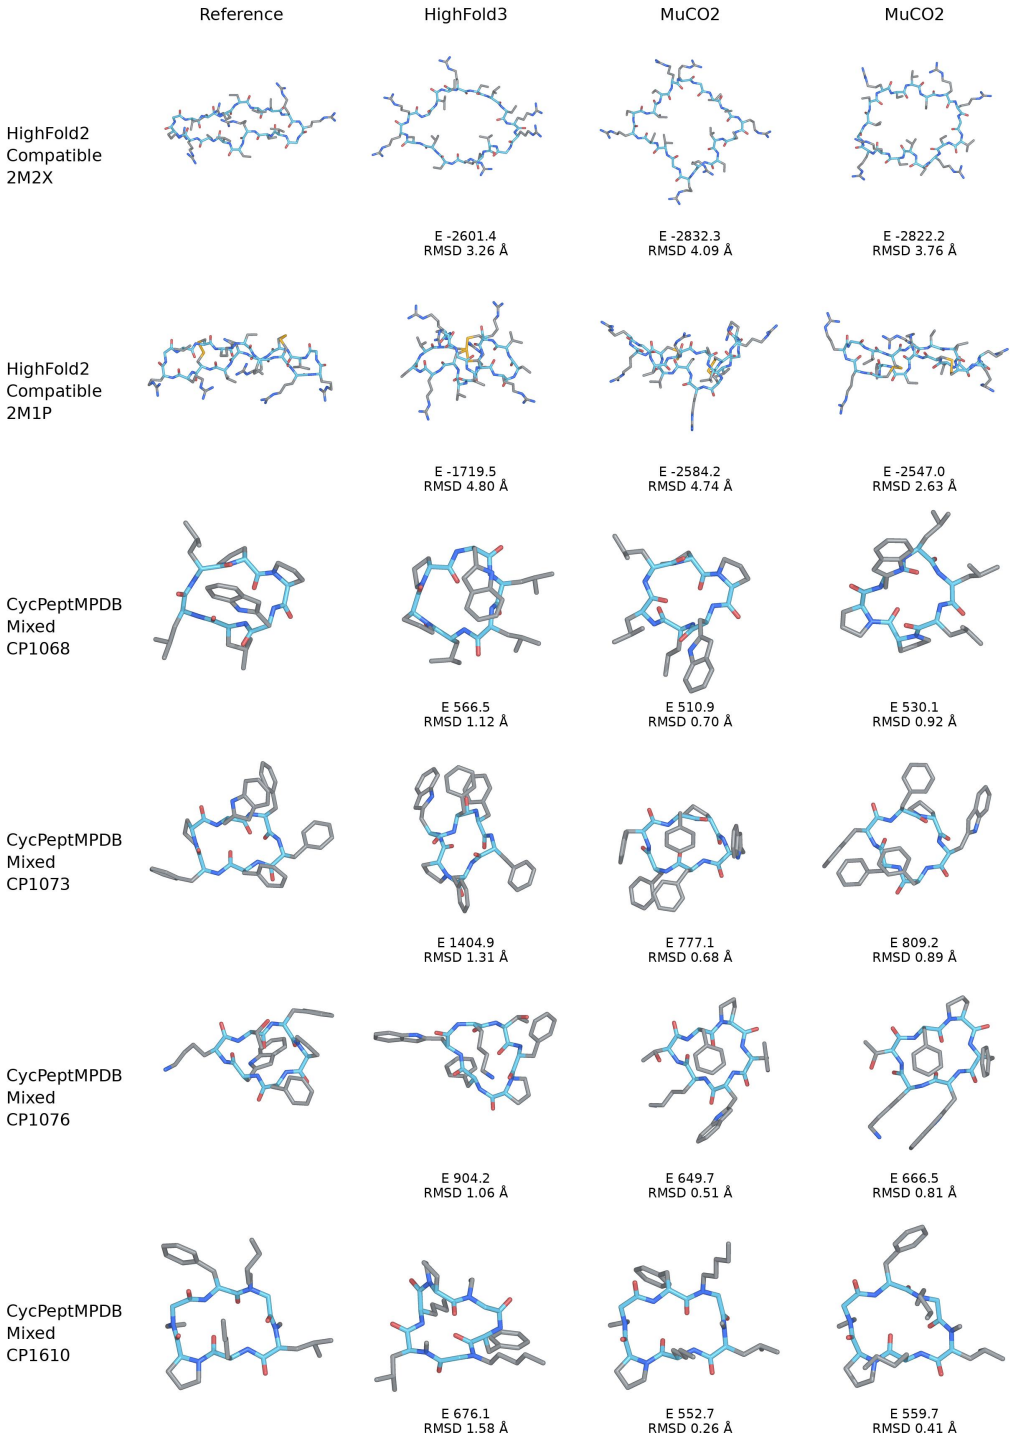


Figure 7: **Structural comparison on side-chain modified and mixed cyclic peptides.** Visual comparison against experimental NMR and MD references across noncanonical side-chain and compound-modification benchmarks. Values report post-relaxation energy E (kJ/mol) and C_α-RMSD (Å). The two rightmost columns show distinct low-energy conformers from MuCO2, demonstrating its multi-basin sampling and strict stereochemical preservation over HighFold3.