MECHPATH: GEOMETRY-PRESERVING FLOW MATCHING FOR PROTEIN MECHANICAL UNFOLDING

Anonymous authors
Paper under double-blind review

ABSTRACT

Understanding how proteins unfold under force is important for elucidating their mechanical functions and associated mechanopathologies, such as cardiomyopathy and muscular dystrophy. Steered molecular dynamics (SMD) simulations resolve unfolding structural dynamics under physiologically relevant pulling geometries, but their high computational cost limits extensive pathway sampling. Existing generative models for protein dynamics offer a scalable alternative, yet typically construct probability paths in ambient structural spaces that do not preserve the constraints imposed by mechanical pulling. Here, we formulate unfolding trajectory generation as probability transport directly on the path geometry induced by mechanical unfolding. We parameterize each unfolding pathway by separating monotonic unfolding progress along the pulling reaction coordinate from the remaining conformational degrees of freedom. Building on this formulation, we introduce MECHPATH, a flow-matching framework that learns the transport in unconstrained latent tokens and maps it back to structural trajectories through a geometry-preserving realization map. This construction makes the satisfaction of the mechanical constraints an *invariant* of the entire generative probability path. In the absence of a public benchmark, we further establish a large-scale dataset comprising $6{,}750$ SMD trajectories from $1{,}350$ proteins. On this benchmark, MECHPATH consistently outperforms general-purpose protein dynamics generators in structural accuracy, conformational-distribution recovery, whole-pathway similarity, and underlying contact-release dynamics. The results show that incorporating the geometry of externally driven molecular processes into generative probability paths substantially improves structural pathway generation, which provides a scalable route to study molecular mechanisms underlying mechanopathologies.

1 INTRODUCTION

Protein mechanical unfolding is central to many biological processes, particularly in load-bearing and force-transmitting systems such as muscle and the cytoskeleton (Heisenberg & Bellaïche, 2013; Iskratsch et al., 2014). In these systems, proteins are pulled between specific attachment sites and perform their mechanical functions through directional, force-induced conformational changes. For example, dystrophin dissipates mechanical stress between the actin cytoskeleton and the sarcolemma (Le et al., 2018), and talin transduces force between integrins and F-actin (Del Rio et al., 2009). During unfolding, a protein may traverse heterogeneous intermediates and follow multiple pathways (Marszalek et al., 1999; Schönfelder et al., 2016). Resolving these pathways is key to understanding how protein topology determines mechanical stability and how mutations compromise mechanical function in mechanopathologies (Brockwell et al., 2003; Sharma et al., 2007).

Steered molecular dynamics (SMD) resolves the structural details of unfolding pathways by fixing one terminus of the protein and pulling the other along a prescribed direction through a moving harmonic potential, thereby emulating physiologically relevant pulling geometries (Li & Makarov, 2003; Sikora et al., 2010). This geometry naturally defines the protein's end-to-end distance along the pulling direction as a one-dimensional reaction coordinate for unfolding progress. However, because unfolding is intrinsically stochastic (Oberhauser et al., 2001), characterizing the pathway distribution of even a single protein requires many independent replicas, and the computational cost of SMD limits such repeated sampling across large protein sets (Stirnemann, 2022; Zhang et al., 2026). To model the large-scale conformational changes induced by unfolding, a scalable alter-

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.

native is to learn a generative model for the trajectory of representative conformations. Such a model should capture a distribution over high-dimensional structural trajectories while respecting the pulling geometry and its imposed constraints. We call an unfolding trajectory *mechanically admissible* if the fixed residue remains at its prescribed position and the projected end-to-end distance along the pulling direction advances monotonically across the representative conformational states.

Deep generative models have extended protein modeling from static structure prediction (Abramson et al., 2024) to conformational ensembles (Lewis et al., 2025) and molecular dynamics (MD) trajectories. For instance, MDGen (Jing et al., 2024) models $\mathrm{SE}(3)$-invariant trajectory representations with stochastic interpolants; 4D-Diffusion (Cheng et al., 2025) incorporates an initial-structure reference and temporal motion alignment to improve multi-frame consistency; and TEMPO (Xu et al., 2026) models temporal evolution autoregressively across multiple timescales. While these approaches improve temporal coherence and conformational evolution, they target general protein dynamics and do not explicitly preserve the mechanical admissibility required for unfolding. To improve structural plausibility, recent methods have incorporated geometric or physics-based constraints directly into the generation process, e.g., diffusion on the $\mathrm{SE}(3)$ manifold (Yim et al., 2023), force-guided conformational sampling (Wang et al., 2024), and explicit structural constraints (Christopher et al., 2026). However, these constraints act on individual conformations and do not address how conformations admissibly evolve along an entire path.

More fundamentally, we show that standard interpolation applied directly to mechanically admissible unfolding trajectories almost surely leaves the admissible path space at intermediate generative time (Remark 1), requiring the model to correct artificial violations while learning the conformational variability of unfolding from limited data. To address this mismatch, we formulate unfolding trajectory generation directly in the geometry induced by mechanical pulling. We represent each trajectory using a proposed progress–conformation coordinate system, which separates the ordered unfolding progress along the pulling reaction coordinate from the conformational variation compatible with that progress. A realization map is constructed to convert unconstrained latent flows into mechanically admissible structural trajectories. This turns mechanical admissibility from a behavior that needs to be learned from data into an *invariant* of the generative process, allowing the model to focus on the conformational evolution of unfolding pathways. Our contributions are as follows:

- **A geometric formulation of mechanical unfolding trajectories.** We formulate protein mechanical unfolding as conditional generation over a mechanically admissible path space and derive a bijective progress–conformation coordinate system for this space. The resulting representation explicitly encodes the pulling geometry induced by mechanical unfolding while retaining the conformational diversity of the unfolding pathway.
- **Flow matching with mechanical admissibility by construction.** We develop MECHPATH, a flow-matching framework that transports probability in the progress–conformation latent space and realizes every generative state as a mechanically admissible structural trajectory through a geometry-preserving map. We prove that mechanical admissibility is preserved throughout the entire generative probability path (Proposition 2).
- **A large-scale benchmark for generative mechanical unfolding.** We construct a large-scale benchmark for unfolding structural trajectory generation by performing SMD simulations on $1{,}350$ proteins and evaluate models from individual structural accuracy to whole-pathway similarity, contact-release dynamics reconstruction, and ensemble-level conformation coverage. On this benchmark, MECHPATH shows improved performance compared with general-purpose protein dynamics generators and substantially improves the recovery of both the structural pathways and conformational distributions.

2 PRELIMINARIES

Protein and unfolding trajectory representation. We represent a protein of N residues by its coarse-grained C_α coordinates $X = (x_1, \ldots, x_N)^\top \in \mathbb{R}^{N\times 3}$, where $x_i \in \mathbb{R}^3$ is the position of residue i. As illustrated in Figure 1(a), constant-velocity SMD fixes one terminus as anchor and pulls the other through a harmonic spring with stiffness k_s moving at velocity v_pull (Sikora et al., 2010), until the end-to-end distance of the protein approaches its contour length (Ainavarapu et al., 2007). From each simulated SMD trajectory, we retain the native structure X_0 and K

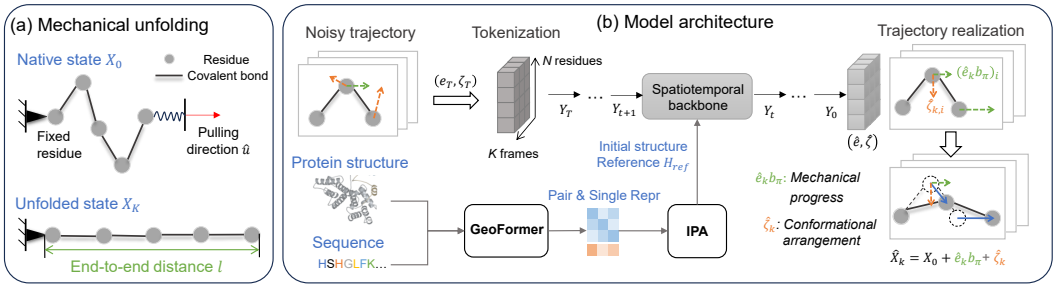

Figure 1: **Overview of the mechanical unfolding setup and MECHPATH architecture. (a)** Mechanical unfolding starts from the native state X_0, with one residue fixed and the protein pulled along $\widehat{u}$; the change of projected end-to-end distance defines mechanical progress toward X_K. **(b)** A noisy trajectory is tokenized into progress–conformation representations, while the initial structure and sequence are encoded into the reference H_{ref}. Conditioned on this reference, the spatiotemporal backbone evolves the trajectory through the generative flow, after which trajectory realization maps the predicted progress and conformational variation back to structural coordinates.

subsequent representative states uniformly along the unfolding progress (Appendix B) to capture large-scale conformational changes. This gives a progress-parameterized unfolding trajectory $X_{0:K} = (X_0, \ldots, X_K) \in \mathbb{R}^{(K+1)\times N\times 3}$. We represent each state by its deformation relative to X_0,

$$\Delta x_k = \mathrm{vec}(X_k - X_0) \in \mathbb{R}^{3N}, \qquad \Delta x_0 = 0, \qquad X_k = X_0 + \mathrm{unvec}(\Delta x_k),$$

where vec stacks the residue coordinates in residue order, so that $(\Delta x_k)_i \in \mathbb{R}^3$ is the displacement of residue i. Referencing every state to the same X_0 turns the fixed-residue condition into a linear constraint on Δx_k, as formalized below.

Mechanical pulling geometry. Let a and b denote the fixed and pulled residues, respectively ($a = 1$ and $b = N$ in this study), and let $\widehat{u} \in \mathbb{R}^3$ be the unit pulling direction. The end-to-end distance projected onto the pulling axis defines the *mechanical reaction coordinate*,

$$\ell_\pi(X) = \widehat{u}^\top (x_b - x_a), \qquad \pi := (a, b, \widehat{u}, v_{\mathrm{pull}}, k_{\mathrm{s}}),$$

where π parametrizes the conditions of the SMD trial. Thus, the projected extension of state k relative to the initial structure is $e_k := \ell_\pi(X_k) - \ell_\pi(X_0) = c_\pi^\top \Delta x_k$, which we refer to as *mechanical progress* of unfolding. Here, $c_\pi \in \mathbb{R}^{3N}$ is the pulling covector that maps a structural deformation to its induced change in projected end-to-end extension.

Mechanically admissible path space. Fixing residue a requires its displacement to vanish, $(\Delta x_k)_a = \mathbf{0}$ for $k = 0, \ldots, K$, which defines the fixed-residue constraint. Every admissible state therefore lies in the fixed-residue subspace $\mathcal{V}_a = \left\{\Delta x \in \mathbb{R}^{3N} \mid (\Delta x)_a = \mathbf{0}\right\}$. Within $\mathcal{V}_a$, each value e of mechanical progress defines an *equal-extension slice* $\mathcal{F}_\pi(e; X_0) = \left\{\Delta x \in \mathcal{V}_a \mid c_\pi^\top \Delta x = e\right\}$. Each slice is an affine subspace, and all slices share the same translation space $\mathcal{T}_\pi = \mathcal{V}_a \cap \ker(c_\pi^\top) = \left\{\zeta \in \mathcal{V}_a \mid c_\pi^\top \zeta = 0\right\}$. Hence, for any $\Delta x \in \mathcal{F}_\pi(e; X_0)$ and $\zeta \in \mathcal{T}_\pi$, $\Delta x + \zeta \in \mathcal{F}_\pi(e; X_0)$.

A progress-parameterized unfolding trajectory traverses these slices in order of increasing mechanical progress. We therefore define the *mechanically admissible path space* as:

$$\mathcal{P}_\pi(X_0) = \left\{ \Delta x_{0:K} \;\middle|\; \begin{array}{l} 0 = e_0 < e_1 < \cdots < e_K, \\ \Delta x_k \in \mathcal{F}_\pi(e_k; X_0), \quad k = 0, \ldots, K \end{array} \right\}. \tag{1}$$

Trajectories in $\mathcal{P}_\pi(X_0)$ are admissible with respect to the prescribed pulling geometry: the anchor residue remains fixed, each state lies on the equal-extension slice defined by the prescribed pulling direction, and these slices are traversed in strictly increasing mechanical progress, while conformational variability is retained within each slice.

Conditional stochastic interpolants. We formulate unfolding trajectory generation with stochastic interpolants (Albergo et al., 2025). Let $\mathcal{C}$ denote the conditioning variables, let $Y_1 \in \mathbb{R}^d$ be a data sample drawn from $p_{\mathrm{data}}(\cdot \mid \mathcal{C})$, and let $Y_0 \sim p_0 = \mathcal{N}(0, I_d)$ be a base sample, where d is the data dimension. The interpolant

$$Y_t = \alpha(t) Y_1 + \sigma(t) Y_0, \qquad t \in [0, 1], \tag{2}$$

with $\alpha(0) = \sigma(1) = 0$ and $\alpha(1) = \sigma(0) = 1$, transports a realization of p_0 to that of $p_{\text{data}}(\cdot \mid \mathcal{C})$. A velocity field v_θ is learned by regressing onto the velocity $\dot{\alpha}(t)Y_1 + \dot{\sigma}(t)Y_0$ with the objective

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t,\, Y_0,\, (Y_1, \mathcal{C})} \left[\left\| v_\theta(Y_t, t \mid \mathcal{C}) - \big(\dot{\alpha}(t)Y_1 + \dot{\sigma}(t)Y_0\big) \right\|_2^2 \right], \qquad t \sim \mathcal{U}[0,1]. \tag{3}$$

At inference, samples are generated by integrating $\mathrm{d}Y_t/\mathrm{d}t = v_\theta(Y_t, t \mid \mathcal{C})$ from $t = 0$ to $t = 1$.

Mismatch of ambient-space probability paths. Our goal is to learn a conditional distribution over mechanically admissible unfolding trajectories,

$$p_\theta(\Delta x_{0:K} \mid \mathcal{C}), \qquad \mathcal{C} = (S, X_0), \qquad \text{such that } \Delta x_{0:K} \in \mathcal{P}_\pi(X_0), \tag{4}$$

where S is the amino-acid sequence. A typical modeling approach applies Eq. 2 to $\mathrm{vec}(\Delta x_{1:K}) \in \mathbb{R}^{3NK}$, which constructs the probability path in the unconstrained ambient deformation space.

Remark 1 (Ambient interpolation leaves the target path space) *Let $\Delta x^\star_{0:K} \in \mathcal{P}_\pi(X_0)$ be an arbitrary admissible trajectory. For $k = 1, \ldots, K$, consider the interpolant $\Delta x_{k,t} := \alpha(t)\Delta x^\star_k + \sigma(t)\varepsilon_k$, where $\varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{3N})$ and for $k = 0$, $\Delta x_{0,t} \equiv \mathbf{0}$. For any $t \in (0,1)$ with $\sigma(t) > 0$:*

(i) *the fixed-residue constraint is violated almost surely. Since $(\Delta x^\star_k)_a := \mathbf{0}$, $(\Delta x_{k,t})_a = \sigma(t)(\varepsilon_k)_a \sim \mathcal{N}(\mathbf{0}, \sigma(t)^2 I_3)$, and hence $\Pr((\Delta x_{k,t})_a = \mathbf{0}) = 0$ for every $k \geq 1$.*

(ii) *the monotonicity of extension is violated with nonzero probability. The extension increment $\delta e_{k,t} = c_\pi^\top(\Delta x_{k,t} - \Delta x_{k-1,t})$ satisfies*

$$\delta e_{k,t} \sim \mathcal{N}\big(\alpha(t)\, \delta e^\star_k,\ \kappa_k \sigma(t)^2 \|c_\pi\|_2^2\big), \qquad \kappa_k = \begin{cases} 1, & k = 1, \\ 2, & k > 1, \end{cases}$$

where $\delta e^\star_k := c_\pi^\top(\Delta x^\star_k - \Delta x^\star_{k-1}) > 0$. Hence $\Pr(\delta e_{k,t} < 0) > 0$ for every $k \geq 1$.

Remark 1 shows that stochastic interpolation on unconstrained structural space leaves the target path geometry: Gaussian perturbations displace the fixed residue and assign nonzero probability to non-monotonicity of extension (see Appendix A.2 for the proof). Although these violations become progressively weaker as t approaches the data endpoint, the model needs to learn from limited data how to transport such intermediate states back to the mechanically admissible distribution while also capturing the conformational variability of unfolding. Empirically, our results show that unconstrained baselines do not fully recover these constraints at generation and face difficulty in capturing conformational distributions (see Table 1). Section 3 instead constructs the probability path in coordinates where both constraints hold by construction.

3 Flow Matching on Mechanically Admissible Path Space

3.1 Mechanical Path Representation

Progress–conformation coordinates. Remark 1 motivates moving the generative process from the unconstrained ambient space to the mechanically admissible path space. At state k, the mechanical progress e_k determines which equal-extension slice the structure occupies, and the remaining freedom lies within that slice. We therefore parameterize each Δx_k by two components: its mechanical progress and the corresponding conformational variation within the equal-extension slice.

To realize the scalar extension e_k as a deformation in $\mathbb{R}^{3N}$, we introduce a *mechanical carrier* $b_\pi \in \mathcal{V}_a$ with $c_\pi^\top b_\pi = 1$, so that one unit of structural deformation along b_π produces exactly one unit of extension along the unit pulling direction. In practice, we choose a carrier b_π that distributes unit extension smoothly along the protein backbone,

$$b_\pi = \arg\min_{\beta \in \mathcal{V}_a} \ \tfrac{1}{2}\, \beta^\top M_G\, \beta \quad \text{s.t.} \quad c_\pi^\top \beta = 1, \qquad M_G = (L_{\text{chain}} + \lambda I_N) \otimes I_3, \tag{5}$$

where L_{chain} is the weighted graph Laplacian of the chain graph with edges $(i, i+1)$ connecting adjacent residues and corresponding edge weights $\{w_i\}_{i=1}^{N-1}$, and $\lambda > 0$ is a small regularizer. In particular, $\beta^\top M_G \beta = \sum_{i=1}^{N-1} w_i \|\beta_{i+1} - \beta_i\|_2^2 + \lambda \sum_{i=1}^{N} \|\beta_i\|_2^2$, so the minimizer favors smoothly varying displacements along the backbone while realizing unit extension. For uniform edge weights, terminal pulling $(a, b) = (1, N)$, and $\lambda \to 0$, this reduces to $b_{\pi,i} = (i-1)\widehat{u}/(N-1)$. The carrier has a closed form and is computed once per protein during generation (Appendix A.3).

Proposition 1 (Progress–conformation direct-sum decomposition) *Let $b_\pi \in \mathcal{V}_a$ satisfy $c_\pi^\top b_\pi = 1$, and define $P_\pi := I_{3N} - b_\pi c_\pi^\top$. Recall $\mathcal{T}_\pi = \mathcal{V}_a \cap \ker(c_\pi^\top)$. Then $\mathcal{V}_a = \mathrm{span}\{b_\pi\} \oplus \mathcal{T}_\pi$, and every $\Delta x_k \in \mathcal{V}_a$ admits the unique decomposition*

$$\boxed{\Delta x_k = e_k b_\pi + \zeta_k, \qquad e_k = c_\pi^\top \Delta x_k, \qquad \zeta_k = P_\pi \Delta x_k \in \mathcal{T}_\pi.} \tag{6}$$

The term $e_k b_\pi$ represents structural deformation along the reaction coordinate that selects the equal-extension slice, and ζ_k captures conformational variability within that slice. Please see Appendix A.4 for the proof.

Latent parameterization. The pair (e_k, ζ_k) provides the unique progress–conformation coordinates of a structural deformation $\Delta x_k \in \mathcal{V}_a$ that preserves the fix-residue property. For generative modeling, we perform flow matching in the latent space $Y = (\gamma_{1:K}, \widetilde{\zeta}_{1:K}) \in \mathcal{Y}_\pi := \mathbb{R}^K \times \mathcal{V}_a^K$, whose components parameterize mechanical progress $e_{0:K}$ and conformational variation $\zeta_{0:K}$, respectively. To guarantee monotonicity, γ_k parameterizes a strictly positive extension increment,

$$\delta e_k = \varepsilon_e + s_e \,\mathrm{softplus}(\gamma_k), \qquad e_k = \sum_{j=1}^{k} \delta e_j, \tag{7}$$

where $\varepsilon_e > 0$ is a small floor and s_e is a length scale, so that $0 = e_0 < e_1 < \cdots < e_K$. The latent conformation variation is projected onto the conformational subspace by $\zeta_k = P_\pi \widetilde{\zeta}_k \in \mathcal{T}_\pi$.

We define the geometry-preserving realization map $\mathcal{R}_\pi : \mathcal{Y}_\pi \to \mathcal{P}_\pi(X_0)$, which decodes a latent state Y for $k = 1, \ldots, K$ as:

$$\widehat{\Delta x}_k = e_k(\gamma)\, b_\pi + P_\pi \widetilde{\zeta}_k, \qquad \widehat{X}_k = X_0 + \mathrm{unvec}(\widehat{\Delta x}_k), \tag{8}$$

Proposition 2 (Admissibility of the realized probability path) *Let $Y^\star = (\gamma^\star_{1:K}, \widetilde{\zeta}^\star_{1:K}) \in \mathcal{Y}_\pi$ be an arbitrary data latent state, and let Y_0 be standard Gaussian noise on the linear latent space $\mathcal{Y}_\pi$, where the fixed-residue coordinates of the conformational component are masked to zero $(\widetilde{\zeta}_{1:K,0})_a = \mathbf{0}$. Consider the stochastic interpolant $Y_t = \alpha(t)Y^\star + \sigma(t)Y_0$ for $t \in [0,1]$:*

(i) Since $\mathcal{Y}_\pi$ is a vector space, $Y_t \in \mathcal{Y}_\pi$. (ii) Then, for every unfolding state $k = 1, \ldots, K$, $\widehat{\Delta x}_{k,t} = \mathcal{R}_\pi(Y_t)$, s.t., $(\widehat{\Delta x}_{k,t})_a = \mathbf{0}$ and $\delta e_{k,t} = \varepsilon_e + s_e \,\mathrm{softplus}(\gamma_{k,t}) > 0$. Thus, since extension $e_{k,t} := \sum_{j=1}^{k} \delta e_{j,t}$ (Eq. 7) and $e_{0,t} := 0$, it follows that $0 = e_{0,t} < e_{1,t} < \cdots < e_{K,t}$.

Recall $\Delta x_{0,t} \equiv \mathbf{0}$. Hence, for every $t \in [0,1]$, $\widehat{\Delta x}_{0:K,t} \in \mathcal{P}_\pi(X_0)$.

Proposition 2 shows that although flow matching remains unconstrained in latent space, every intermediate state is decoded into a trajectory that remains admissible with respect to the prescribed pulling geometry: the anchor residue is fixed, progress is defined along the prescribed pulling direction, and the resulting extension advances monotonically (see Figure 2).

3.2 Mechanical Unfolding Flow-Matching Model

Trajectory tokenization. Our target is to model the conditional distribution of unfolding trajectories parameterized by $Y = (\gamma_{1:K}, \widetilde{\zeta}_{1:K})$, given conditions $\mathcal{C}$. We arrange the latent variables into residue-wise tokens. At state k and residue i, the token is $y_{k,i} = [\mathbf{1}_{i=b}\,\gamma_k,\ \widetilde{\zeta}_{k,i}] \in \mathbb{R}^4$, where $\widetilde{\zeta}_{k,i} \in \mathbb{R}^3$ is the block of residue i. The scalar γ_k is placed only on the token of the pulled residue b, where the pulling is applied. Stacking tokens gives $Y^{\mathrm{tok}} \in \mathbb{R}^{K \times N \times 4}$.

Architecture. As shown in Figure 1(b), we encode the initial structure and sequence into a reference

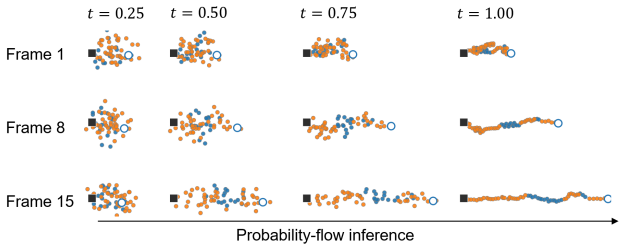

Figure 2: **Probability-flow inference of mechanical unfolding pathways.** Throughout inference, the fixed residue remains stationary, and the extension advances monotonically along the prescribed pulling direction, as guaranteed by Proposition 2.

representation $H_{\text{ref}} \in \mathbb{R}^{N\times d}$ with invariant point attention (IPA) (Jumper et al., 2021). IPA takes the residue frames of X_0 as geometric input, amino-acid embeddings as residue features, and pair representations from the pretrained GeoFormer (Wu et al., 2022) as pair features. The tokens Y_t^{tok} are linearly projected to dimension d and added to H_{ref} at every progress state, giving $H_t^{(0)} \in \mathbb{R}^{K\times N\times d}$. The backbone stacks M diffusion transformer-style blocks (Peebles & Xie, 2023) that alternate two operations: residue attention models interactions among residues within each state, and path attention couples each residue across states. Each block is conditioned on the flow time t through adaptive layer normalization, and a final projection maps $H_t^{(M)}$ back to the token space to predict the velocity $v_\theta(Y_t, t \mid \mathcal{C})$. Details of the model architecture and hyperparameters can be found in Appendix C.2.

Training objective. We use the geometric variance-preserving path (Song et al., 2023), $\alpha(t) = \sin(\pi t/2)$ and $\sigma(t) = \cos(\pi t/2)$, and train v_θ with $\mathcal{L}_{\text{FM}}$ in Eq. 3. To improve the structural validity, we further supervise the clean pathway implied by the prediction. The clean-endpoint estimate $\widehat{Y}_1$ follows analytically from Y_t and v_θ, and decoding it with $\mathcal{R}_\pi$ gives $\widehat{X}_{1:K}$. The training objective is

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{FM}} + g(t)(w_{\text{bond}}\,\mathcal{L}_{\text{bond}} + w_{\text{contact}}\,\mathcal{L}_{\text{contact}}), \qquad g(t) = \mathbf{1}[t > 0.75], \tag{9}$$

where $\mathcal{L}_{\text{bond}}$ penalizes deviations of consecutive C_α distances from their native values, $\mathcal{L}_{\text{contact}}$ supervises the survival and release of native contacts along the pathway, and w_{bond} and w_{contact} are weights of losses. The auxiliary losses are applied only near the data endpoint, where the clean-path estimate is reliable enough to constrain local geometry (Yim et al., 2023). The estimator and loss definitions are given in Appendix D.2.

4 Experiments

4.1 Experimental Setup

Dataset. To our knowledge, no public dataset provides structural trajectories of mechanical unfolding. We therefore construct one using constant-velocity SMD with a coarse-grained C_α structure-based model and a pulling setup similar to BSDB (Sikora et al., 2010), in which the N-terminus is fixed and the C-terminus is pulled. We simulate $1{,}350$ proteins with five independent replicas each ($6{,}750$ trajectories). Since our objective is to model large-scale unfolding pathways instead of short-timescale thermal fluctuations, we subsample each trajectory into X_0 and $K = 15$ subsequent states. Proteins are split $80/10/10$ ($1{,}080/135/135$), with all replicas of a protein assigned to the same split. Data collection and preprocessing details are provided in Appendix B.

Baselines. We compare with three recent generative models for protein dynamics: MDGen (Jing et al., 2024), a flow-based model that jointly generates spatial and temporal motion; 4D-Diffusion (Cheng et al., 2025), a diffusion model over residue $\mathrm{SE}(3)$ frames that generates multiple time steps simultaneously; and TEMPO (Xu et al., 2026), which autoregressively models conformational evolution across multiple timescales. All methods are trained on the same dataset and splits, conditioned on the same initial structure X_0, and evaluated with the same pipeline (Appendix F).

Metrics. We evaluate five aspects: **(1) Structural accuracy:** progress-aligned C_α path RMSD (PA-RMSD, Å), which measures the average positional deviation between corresponding C_α atoms after matching generated and reference states by mechanical progress. **(2) Mechanical admissibility:** monotonicity satisfaction rate (MSR), the fraction of non-decreasing extension steps, and anchor drift (AD, Å), the mean displacement of the fixed residue. **(3) Contact-release dynamics:** contact-resolved Wasserstein-1 distance (CR-W_1) between generated and reference rupture-progress distributions of individual native contacts, and rupture-order Kendall correlation (RO-τ). **(4) Conformational distribution:** Jensen–Shannon divergence between generated and reference conformations in the first two principal components of the SMD ensemble (PCA-JSD) (Shen et al., 2026). **(5) Pathway similarity:** discrete Fréchet path distance (FPD, Å) between complete pathways (Seyler et al., 2015). Full definitions of all metrics are provided in Appendix E.

Evaluation. For each test protein, we independently generate five pathways and compare them with the five SMD replicas. Trajectory-level metrics (PA-RMSD, MSR, AD, RO-τ, FPD) compare each generated pathway with its paired replica and then average over the five pairs to obtain one score per protein; CR-W_1 and PCA-JSD compare the pooled generated and SMD ensembles, yielding one

Table 1: Quantitative comparison of mechanical unfolding trajectory generation methods, with results of mean ± standard deviation over three independent runs. Best results are shown in **bold**, and second-best results are underlined. ↑ indicates higher is better, while ↓ indicates lower is better.

	Structural	*Mechanical Admissibility*		*Contact-Release Dynamics*		*Distribution*	*Pathway Similarity*
Model	**PA-RMSD (Å) ↓**	**MSR ↑**	**AD (Å) ↓**	**CR-W$_1$ ↓**	**RO-τ ↑**	**PCA-JSD ↓**	**FPD (Å) ↓**
MDGen	47.53 ± 5.95	0.87 ± 0.03	0.52 ± 0.24	0.32 ± 0.02	0.29 ± 0.02	0.28 ± 0.02	92.13 ± 12.38
4D-Diffusion	28.77 ± 12.23	0.79 ± 0.05	16.63 ± 9.68	0.22 ± 0.02	$\mathbf{0.37 \pm 0.06}$	0.27 ± 0.09	43.28 ± 18.51
TEMPO	17.95 ± 1.04	$\mathbf{1.00 \pm 0.00}$	8.33 ± 0.67	0.22 ± 0.00	0.29 ± 0.02	0.30 ± 0.03	28.40 ± 0.91
MechPath	$\mathbf{13.06 \pm 0.66}$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{0.00 \pm 0.00}$	$\mathbf{0.21 \pm 0.01}$	0.35 ± 0.02	$\mathbf{0.13 \pm 0.00}$	$\mathbf{21.31 \pm 0.34}$

Table 2: Ablation study of geometric representations for unfolding trajectory generation.

	Structural	*Mechanical Admissibility*		*Contact-Release Dynamics*		*Distribution*	*Pathway Similarity*
Representation	**PA-RMSD (Å) ↓**	**MSR ↑**	**AD (Å) ↓**	**CR-W$_1$ ↓**	**RO-τ ↑**	**PCA-JSD ↓**	**FPD (Å) ↓**
Cartesian	15.27 ± 0.53	0.96 ± 0.02	4.29 ± 1.07	0.35 ± 0.02	0.18 ± 0.01	0.16 ± 0.00	22.77 ± 0.66
Residue Frame	30.14 ± 0.96	0.96 ± 0.01	0.58 ± 0.08	0.27 ± 0.01	0.35 ± 0.01	0.19 ± 0.00	55.50 ± 2.06
Prog.–Conf. (MechPath)	13.06 ± 0.66	$\mathbf{1.00 \pm 0.00}$	$\mathbf{0.00 \pm 0.00}$	0.21 ± 0.01	0.35 ± 0.02	$\mathbf{0.13 \pm 0.00}$	21.31 ± 0.34
Prog.–Conf. with Angles	$\mathbf{12.42 \pm 0.21}$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{0.00 \pm 0.00}$	$\mathbf{0.20 \pm 0.01}$	$\mathbf{0.38 \pm 0.01}$	0.15 ± 0.01	$\mathbf{19.94 \pm 0.13}$

score per protein. For each training run, we average these protein-level scores over the test set and report the values of mean ± standard deviation across three independent runs.

4.2 MechPath Enables Accurate and Robust Unfolding Trajectory Generation

Main benchmark. Table 1 shows that MechPath achieves the best or second-best result on all seven metrics. As guaranteed by Proposition 2, MechPath satisfies both mechanical constraints by construction, yielding MSR 1.00 and AD 0.00 Å. However, no baseline satisfies both constraints: MDGen nearly fixes the anchor (AD 0.52 Å) but shows 13% backward steps (MSR 0.87); TEMPO is fully monotonic but has substantial anchor drift (AD 8.33 Å); and 4D-Diffusion violates both (MSR 0.79, AD 16.63 Å). This exact admissibility of MechPath is accompanied by more accurate pathways and distributions. Relative to the strongest baseline, MechPath reduces PA-RMSD by 27% (17.95 to 13.06 Å) and FPD by 25% (28.40 to 21.31 Å) versus TEMPO, and PCA-JSD by 52% (0.27 to 0.13) versus 4D-Diffusion. In contrast, MDGen performs poorly on both structural metrics (PA-RMSD 47.53 Å, FPD 92.13 Å), indicating difficulty in capturing the large conformational changes of unfolding, while TEMPO shows the weakest PCA-JSD (0.30) and RO-τ (0.29), suggesting that temporal coherence alone is insufficient to recover the conformations and contact-release order. These results show that enforcing mechanical admissibility by construction is compatible with, and empirically accompanied by, more accurate structural pathways and conformational reconstruction.

Why does unfolding path geometry help? We next ask whether the gains of MechPath come from richer geometric features or from incorporating the pulling geometry into the generative state space. Table 2 compares four representations under the same model architecture and training procedure: Cartesian coordinates, residue $\mathrm{SE}(3)$ frames, progress–conformation coordinates (MechPath), and MechPath augmented with local frame geometry (Appendix C.1). Residue frames reduce anchor drift from 4.29 Å to 0.58 Å and CR-W$_1$ by 24% relative to Cartesian coordinates, yet PA-RMSD and FPD remain high (30.14 Å and 55.50 Å), respectively, showing that richer angular geometry alone does not necessarily recover the correct unfolding pathway. The proposed progress–conformation coordinates instead guarantee mechanical admissibility and improve pathway accuracy simultaneously by separating mechanical progress from conformational variation. Relative to Cartesian coordinates, PA-RMSD decreases by 14%, PCA-JSD by 19%, while significantly improving reconstruction of contact-release dynamics (CR-W$_1$ from 0.35 to 0.21). Adding local frame geometry reduces PA-RMSD by 5% and FPD by 6%, and improvs RO-τ by 9%, showing that progress–conformation representation can further benefit from complementary local geometric information. These results support our central argument: preserving mechanical admissibility throughout generation allows the model to devote its modeling capacity to internal conformational changes, improving trajectory fidelity and recovery of conformational and contact-release dynamics.

MechPath scales more robustly with protein size. We further test whether the gains of MechPath persist as proteins grow larger. Test proteins are divided into five approximately balanced

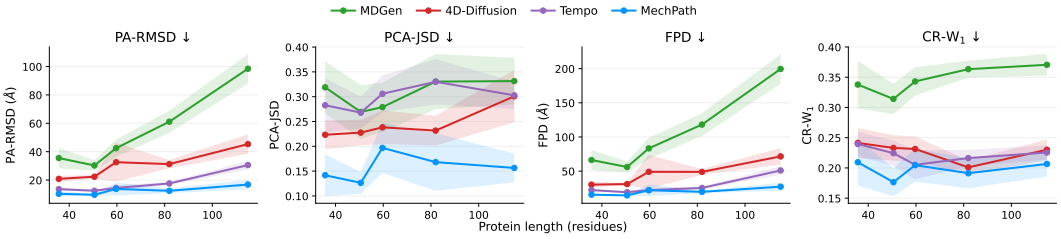

Figure 3: Generalization across protein length bins.

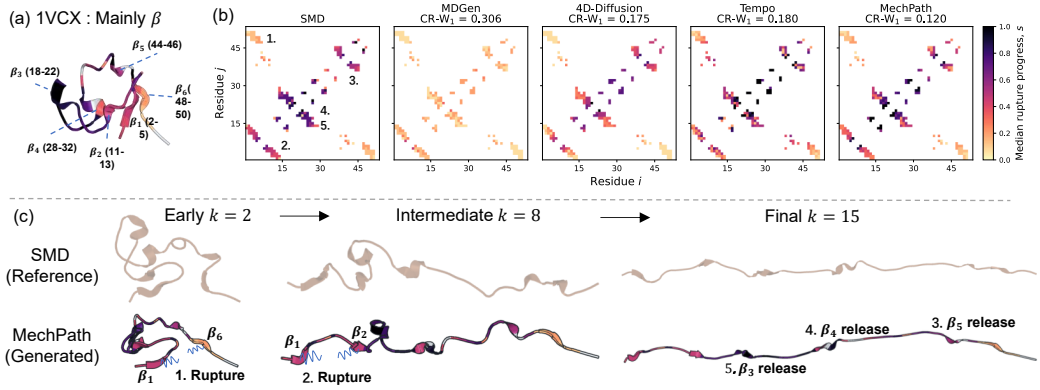

Figure 4: **Representative contact-release mechanism and structural unfolding pathway for 1VCX. (a)** Native structure with labeled β-strand elements. **(b)** Contact rupture progress for SMD and generated pathways, with CR-W_1 measuring agreement with the SMD reference. **(c)** SMD and MECHPATH conformations along unfolding, illustrating the structural realization of the corresponding contact-release sequence.

length bins (16–46, 46–57, 57–64, 64–104, and 104–128 residues; 19–28 proteins per bin), and we report group means with 95% confidence intervals. Figure 3 shows that PA-RMSD increases from the shortest to longest bin by 177% for MDGen (35.51 to 98.45 Å) and 125% for TEMPO (13.60 to 30.55 Å), but only 63% for MECHPATH (10.31 to 16.84 Å). In the longest bin, MECHPATH also achieves 46% lower FPD than TEMPO (27.29 vs. 51.00 Å) and 48–53% lower PCA-JSD than the baselines (0.156 vs. 0.301–0.331). This weaker error growth highlights the robustness of MECHPATH as conformational complexity increases with protein size. Moreover, MECHPATH remains robust across distinct native topologies (see Appendix G.2).

Computational cost. Mechanical admissibility adds only lightweight operations: b_π is computed once per protein, and projection and progress decoding scale linearly with trajectory size. Thus, the dominant computational cost remains the spatiotemporal backbone and ODE integration. In our experiments, training MECHPATH on a single NVIDIA RTX 5090 GPU takes approximately 3 h for 50 epochs, while inference generates one sampled pathway per second. By comparison, one SMD trajectory requires approximately $800 \sim 1,600$ seconds for 50–100-residue proteins, yielding a per-sample speed-up of approximately $\sim 10^3\times$.

4.3 RECOVERING CONTACT-RELEASE DYNAMICS AND CONFORMATIONAL COVERAGE

Aggregate metrics do not reveal whether generated pathways unfold through the correct mechanism. Faithful modeling of mechanical unfolding requires recovering both the contact-release dynamics underlying structural transitions and the conformational variability across repeated unfolding events (Israelewitz et al., 2001; Beedle & Garcia-Manyes, 2023). Figure 4 examines the first aspect using 1VCX, a mainly-β protein characterized by sequential release of native β-strand contacts. MECHPATH most closely reproduces the SMD contact-release map (CR-W_1 0.120 vs. 0.175, 0.180, and 0.306 for 4D-Diffusion, TEMPO, and MDGen). Its structural trajectory follows the same progression as SMD: β_1–β_6 contacts rupture first, followed by β_1–β_2 and further β-strand contacts as

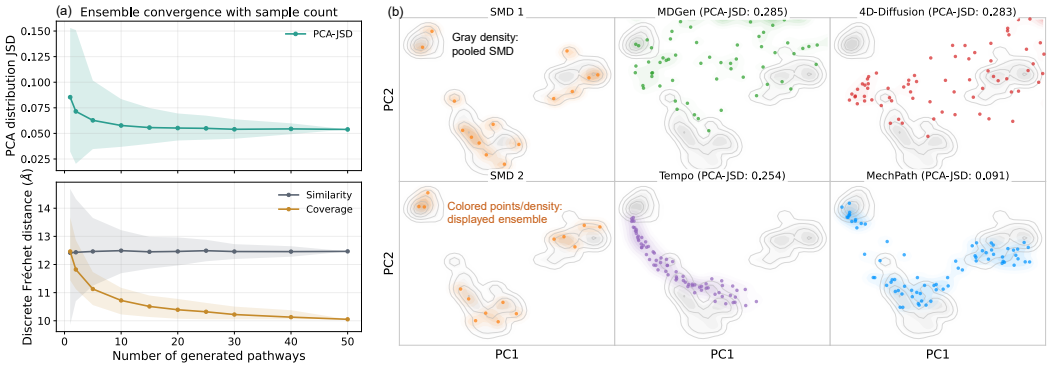

Figure 5: **Sample-count convergence and ensemble coverage of generated mechanical unfolding pathways. (a)** PCA-JSD and FPD-based coverage as the number of generated pathways increases; shaded regions denote uncertainty across repeated subsampling. **(b)** Shared PCA projection of SMD and generated conformations. Gray contours show the pooled SMD conformational density, and colored points show the displayed SMD replica or generated ensemble.

the protein approaches the extended state. MECHPATH thus recovers not only the structures along unfolding but also the topology-dependent contact-release sequence connecting them. Additional examples of proteins with diverse topologies are shown in Appendix G.3.

Because unfolding is stochastic, recovering this mechanism also requires sampling the broader conformational support of independent SMD replicas. As more MECHPATH samples are generated (Figure 5a), PCA-JSD and the real-to-generated coverage distance both decrease, whereas the generated-to-real fidelity distance remains approximately constant. Additional samples therefore populate underrepresented regions of the reference ensemble rather than degrading individual pathways. The shared PCA projection in Figure 5(b) illustrates this behavior. MDGen and 4D-Diffusion scatter broadly and miss several SMD regions. TEMPO captures the dominant elongated manifold but remains confined to a narrow region, possibly because autoregressive conditioning on previously generated states favors locally consistent evolution. In contrast, MECHPATH covers the three major SMD-supported regions, spanning compact, intermediate-extension, and highly extended conformations, while achieving the best fidelity distance (12.25 Å). These results show that MECHPATH captures the ordered contact-release evolution and its conformational variability of unfolding.

5 CONCLUSION

Force-driven conformational change is not unconstrained molecular motion: it evolves under a prescribed pulling geometry that fixes the anchor, orders extension, and restricts admissible pathways. MECHPATH incorporates this geometry directly into generation by decomposing each trajectory into mechanical progress and conformational variation on equal-extension slices, and performing flow matching in this mechanically aligned representation. The resulting realization map preserves mechanical admissibility throughout generation, allowing the model to focus on pathway-dependent internal motion with no need to relearn known constraints. Across a large SMD benchmark, this design yields more accurate pathways, better conformational ensemble recovery, and competitive contact-release dynamics compared to general-purpose protein dynamics generators, with advantages that persist across protein size and structural topology. Ablations further show that these gains arise primarily from aligning the generative coordinates with the pulling geometry instead of increased representational complexity. More broadly, the results suggest that when externally driven molecular processes impose known structure on system evolution, embedding that structure into the generative path can improve both physical consistency and modeling efficiency. The present study is limited to coarse-grained, simulation-derived trajectories under constant-velocity pulling and proteins with fewer than 128 residues. Future work should therefore extend the framework to longer proteins, broader pulling setups, higher-resolution simulations, and wet-lab validation.

AI USE STATEMENT

In this work, we used generative AI tools for polishing the language of the paper and writing code. We have reviewed all AI-assisted work. LLM-generated code was verified and tested for correctness. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide the information required to reproduce the theoretical, training, and evaluation components of this work throughout the main paper and appendices. Appendix A gives the mathematical definitions, assumptions, and complete proofs underlying the mechanically admissible path representation and realization map. Appendix B describes the SMD data generation, trajectory preprocessing, progress-state construction, and protein-level train/validation/test split. Appendix C specifies the structural representations, conditioning features, and model architecture, while Appendix D provides the flow-matching objective, auxiliary losses, optimization settings, and ODE sampling procedure. Appendix E defines all evaluation metrics and aggregation procedures, and Appendix F documents the implementation and adaptation of each baseline together with the fair-comparison protocol. Additional robustness and qualitative results are provided in Appendix G. Code, configuration files, and scripts for data processing, training, inference, and evaluation are provided in the anonymous supplementary material.

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.

Sri Rama Koti Ainavarapu, Jasna Brujić, Hector H Huang, Arun P Wiita, Hui Lu, Lewyn Li, Kirstin A Walther, Mariano Carrion-Vazquez, Hongbin Li, and Julio M Fernandez. Contour length and refolding rate of a small protein controlled by engineered disulfide bonds. *Biophysical journal*, 92(1):225–233, 2007.

Michael Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025.

Amy EM Beedle and Sergi Garcia-Manyes. The role of single-protein elasticity in mechanobiology. *Nature Reviews Materials*, 8(1):10–24, 2023.

David J Brockwell, Emanuele Paci, Rebecca C Zinober, Godfrey S Beddard, Peter D Olmsted, D Alastair Smith, Richard N Perham, and Sheena E Radford. Pulling geometry defines the mechanical resistance of a β-sheet protein. *Nature Structural & Molecular Biology*, 10(9):731–737, 2003.

Kaihui Cheng, Ce Liu, Qingkun Su, Jun Wang, Liwei Zhang, Yining Tang, Yao Yao, Siyu Zhu, and Yuan Qi. 4d diffusion for dynamic protein structure prediction with reference and motion guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 93–101, 2025.

Jacob K Christopher, Austin Seamann, Jingyi Cui, Sagar Khare, and Ferdinando Fioretto. Constrained diffusion for protein design with hard structural constraints. In *International Conference on Learning Representations*, volume 2026, pp. 124089–124108, 2026.

Armando Del Rio, Raul Perez-Jimenez, Ruchuan Liu, Pere Roca-Cusachs, Julio M Fernandez, and Michael P Sheetz. Stretching single talin rod molecules activates vinculin binding. *Science*, 323 (5914):638–641, 2009.

Carl-Philipp Heisenberg and Yohanns Bellaïche. Forces in tissue morphogenesis and patterning. *Cell*, 153(5):948–962, 2013.

Thomas Iskratsch, Haguy Wolfenson, and Michael P Sheetz. Appreciating force and shape—the rise of mechanotransduction in cell biology. *Nature reviews Molecular cell biology*, 15(12):825–833, 2014.

Barry Isralewitz, Mu Gao, and Klaus Schulten. Steered molecular dynamics and mechanical functions of proteins. *Current opinion in structural biology*, 11(2):224–230, 2001.

Bowen Jing, Hannes Stärk, Tommi Jaakkola, and Bonnie Berger. Generative modeling of molecular dynamics trajectories. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. URL https://openreview.net/forum?id=LbwM4VCDUU.

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.

Shimin Le, Miao Yu, Ladislav Hovan, Zhihai Zhao, James Ervasti, and Jie Yan. Dystrophin as a molecular shock absorber. *ACS nano*, 12(12):12140, 2018.

Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew YK Foong, Victor García Satorras, Osama Abdin, Bastiaan S Veeling, Iryna Zaporozhets, et al. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science*, 389(6761): eadv9817, 2025.

Pai-Chi Li and Dmitrii E Makarov. Theoretical studies of the mechanical unfolding of the muscle protein titin: bridging the time-scale gap between simulation and experiment. *The Journal of chemical physics*, 119(17):9260–9268, 2003.

Piotr E Marszalek, Hui Lu, Hongbin Li, Mariano Carrion-Vazquez, Andres F Oberhauser, Klaus Schulten, and Julio M Fernandez. Mechanical unfolding intermediates in titin modules. *Nature*, 402(6757):100–103, 1999.

Andres F Oberhauser, Paul K Hansma, Mariano Carrion-Vazquez, and Julio M Fernandez. Stepwise unfolding of titin under force-clamp atomic force microscopy. *Proceedings of the National Academy of Sciences*, 98(2):468–472, 2001.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

Jörg Schönfelder, Raul Perez-Jimenez, and Victor Munoz. A simple two-state protein unfolds mechanically via multiple heterogeneous pathways at single-molecule resolution. *Nature communications*, 7(1):11777, 2016.

Sean L Seyler, Avishek Kumar, Michael F Thorpe, and Oliver Beckstein. Path similarity analysis: a method for quantifying macromolecular pathways. *PLoS computational biology*, 11(10): e1004568, 2015.

Deepak Sharma, Ognjen Perisic, Qing Peng, Yi Cao, Canaan Lam, Hui Lu, and Hongbin Li. Single-molecule force spectroscopy reveals a mechanically stable protein fold and the rational tuning of its mechanical stability. *Proceedings of the National Academy of Sciences*, 104(22):9278–9283, 2007.

Yuning Shen, Lihao Wang, Huizhuo Yuan, Yan Wang, Bangji Yang, and Quanquan Gu. Simultaneous modeling of protein conformation and dynamics via autoregression. *Advances in Neural Information Processing Systems*, 38:109782–109822, 2026.

Mateusz Sikora, Joanna I Sułkowska, Bartłomiej S Witkowski, and Marek Cieplak. Bsdb: the biomolecule stretching database. *Nucleic acids research*, 39(suppl_1):D443–D450, 2010.

Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36:549–568, 2023.

Guillaume Stirnemann. Recent advances and emerging challenges in the molecular modeling of mechanobiological processes. *The Journal of Physical Chemistry B*, 126(7):1365–1374, 2022.

Vaishali P Waman, Nicola Bordin, Andy Lau, Shaun Kandathil, Jude Wells, David Miller, Sameer Velankar, David T Jones, Ian Sillitoe, and Christine Orengo. Cath v4. 4: major expansion of cath by experimental and predicted structural data. *Nucleic Acids Research*, 53(D1):D348–D355, 2025.

Yan Wang, Lihao Wang, Yuning Shen, Yiqun Wang, Huizhuo Yuan, Yue Wu, and Quanquan Gu. Protein conformation generation via force-guided se(3) diffusion models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Ruidong Wu, Fan Ding, Rui Wang, Rui Shen, Xiwen Zhang, Shitong Luo, Chenpeng Su, Zuofan Wu, Qi Xie, Bonnie Berger, et al. High-resolution de novo structure prediction from primary sequence. *BioRxiv*, pp. 2022–07, 2022.

Yaoyao Xu, Di Wang, Zihan Zhou, Tianshu Yu, and Mingchen Chen. Tempo: Temporal multi-scale autoregressive generation of protein conformational ensembles. *Advances in Neural Information Processing Systems*, 38:104292–104324, 2026.

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*, 2023.

Yiyuan Zhang, Cailong Hua, Vinitendra Singh, Joseph M Muretta, James M Ervasti, and Murti Salapaka. Genunfold: Rapidly predict protein mechanical unfolding trajectory via a physics-guided diffusion model. In *Forty-third International Conference on Machine Learning*, 2026.

A Mathematical Details and Proofs

This section provides the mathematical details underlying the mechanical path representation introduced in the main text. We first formalize the fixed-anchor path space and the geometry induced by the pulling condition, then establish the existence and uniqueness of the canonical mechanical carrier. We subsequently prove the failure of ambient Gaussian interpolation, the progress–conformation direct-sum decomposition, and the admissibility of the realized generative path.

A.1 Notation and Mechanical Path Space

We use a coarse-grained C_α representation of a protein with N residues,

$$X = (x_1, \dots, x_N)^\top \in \mathbb{R}^{N\times 3}, \qquad x_i \in \mathbb{R}^3.$$

Let X_0 denote the initial structure. The pulling condition specifies a fixed residue a, a pulled residue $b \neq a$, and a unit pulling direction $\widehat{u} \in \mathbb{R}^3$. The corresponding projected end-to-end reaction coordinate is

$$\ell_\pi(X) = \widehat{u}^\top (x_b - x_a).$$

Let $\mathbf{r}_i \in \mathbb{R}^N$ denote the i-th standard basis vector, i.e., $(\mathbf{r}_i)_j = \delta_{ij}$. The pulling covector is

$$c_\pi = (\mathbf{r}_b - \mathbf{r}_a) \otimes \widehat{u} \in \mathbb{R}^{3N},$$

such that, for any X, X',

$$\ell_\pi(X) - \ell_\pi(X') = c_\pi^\top \operatorname{vec}(X - X').$$

The fixed-anchor configurations form the affine space

$$\mathcal{A}_a(X_0) := \left\{ X \in \mathbb{R}^{N\times 3} \ : \ x_a = x_{0,a} \right\}.$$

Its associated displacement space is the linear subspace

$$\mathcal{V}_a := \left\{ \delta x \in \mathbb{R}^{3N} \ : \ (\delta x)_a = 0 \right\},$$

where $(\delta x)_a \in \mathbb{R}^3$ denotes the three coordinates corresponding to residue a. Hence

$$X \in \mathcal{A}_a(X_0) \quad \Longleftrightarrow \quad \Delta x = \operatorname{vec}(X - X_0) \in \mathcal{V}_a.$$

For a prescribed extension $e \in \mathbb{R}$, we define the equal-extension slice

$$\mathcal{F}_\pi(e; X_0) := \left\{ \Delta x = \operatorname{vec}(X - X_0) \in \mathcal{V}_a \ : \ \ell_\pi(X) - \ell_\pi(X_0) = e \right\}.$$

Because both the anchor condition and reaction coordinate are affine-linear, $\mathcal{F}_\pi(e; X_0)$ is an affine subspace whose translation space is

$$\mathcal{T}_\pi := \mathcal{V}_a \cap \ker(c_\pi^\top).$$

Thus, $\mathcal{T}_\pi$ contains precisely the anchor-preserving deformations that do not change the projected extension.

For a sequence of $K+1$ progress states $X_{0:K}$, define

$$e_k := \ell_\pi(X_k) - \ell_\pi(X_0), \qquad e_0 = 0.$$

The mechanically admissible path space considered in this work is

$$\mathcal{P}_\pi(X_0) = \left\{ \Delta x_{0:K} \mid \exists\, e_{0:K} \text{ s.t. } 0 = e_0 < e_1 < \cdots < e_K,\ \Delta x_k \in \mathcal{F}_\pi(e_k; X_0),\ k = 0, \dots, K \right\}. \tag{10}$$

The two defining properties are therefore a fixed anchor and strictly ordered progress along the pulling reaction coordinate. Structural variation orthogonal to this scalar progress remains unrestricted within the equal-extension slices.

A.2 Proof of Remark 1

Proof 1 *Let $\Delta x^\star_{0:K} \in \mathcal{P}_\pi(X_0)$ be an arbitrary admissible trajectory. For $k = 1, \dots, K$, consider the ambient Gaussian interpolant*

$$\Delta x_{k,t} = \alpha(t)\Delta x^\star_k + \sigma(t)\varepsilon_k, \qquad \varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{3N}),$$

with $\Delta x_{0,t} \equiv \mathbf{0}$. Fix $t \in (0,1)$ such that $\sigma(t) > 0$.

Because $\Delta x^\star_{0:K} \in \mathcal{P}_\pi(X_0)$, the fixed-residue constraint gives

$$(\Delta x^\star_k)_a = \mathbf{0}, \qquad k = 1, \dots, K.$$

Hence

$$(\Delta x_{k,t})_a = \sigma(t)(\varepsilon_k)_a.$$

Since $(\varepsilon_k)_a \sim \mathcal{N}(\mathbf{0}, I_3)$, we have

$$(\Delta x_{k,t})_a \sim \mathcal{N}\left(\mathbf{0}, \sigma(t)^2 I_3\right),$$

and therefore

$$\Pr((\Delta x_{k,t})_a = \mathbf{0}) = 0, \qquad k = 1, \dots, K.$$

Membership in $\mathcal{P}_\pi(X_0)$ requires the fixed-residue constraint at every progress state. Thus, for any $k \geq 1$,

$$\{\Delta x_{0:K,t} \in \mathcal{P}_\pi(X_0)\} \subseteq \{(\Delta x_{k,t})_a = \mathbf{0}\},$$

which implies

$$\Pr(\Delta x_{0:K,t} \in \mathcal{P}_\pi(X_0)) = 0.$$

We next consider monotonicity of the mechanical progress. For $k = 1, \dots, K$, define the adjacent-state extension increment

$$\delta e_{k,t} := c_\pi^\top (\Delta x_{k,t} - \Delta x_{k-1,t}),$$

and, along the admissible trajectory,

$$\delta e^\star_k := c_\pi^\top \left(\Delta x^\star_k - \Delta x^\star_{k-1}\right) > 0.$$

For $k > 1$, both adjacent states contain independent Gaussian noise, so

$$\begin{aligned}\delta e_{k,t} &= c_\pi^\top \left[\alpha(t)\left(\Delta x^\star_k - \Delta x^\star_{k-1}\right) + \sigma(t)\left(\varepsilon_k - \varepsilon_{k-1}\right)\right] \\ &= \alpha(t)\delta e^\star_k + \sigma(t) c_\pi^\top (\varepsilon_k - \varepsilon_{k-1}).\end{aligned}$$

Independence gives

$$\varepsilon_k - \varepsilon_{k-1} \sim \mathcal{N}(0, 2I_{3N}),$$

and therefore

$$c_\pi^\top (\varepsilon_k - \varepsilon_{k-1}) \sim \mathcal{N}\left(0, 2\|c_\pi\|_2^2\right).$$

For $k = 1$, the initial structure is held fixed, $\Delta x_{0,t} \equiv \mathbf{0}$, so

$$\delta e_{1,t} = \alpha(t)\delta e^\star_1 + \sigma(t) c_\pi^\top \varepsilon_1,$$

with

$$c_\pi^\top \varepsilon_1 \sim \mathcal{N}\left(0, \|c_\pi\|_2^2\right).$$

Thus, for every $k = 1, \dots, K$,

$$\delta e_{k,t} \sim \mathcal{N}\left(\alpha(t)\delta e^\star_k, \kappa_k \sigma(t)^2 \|c_\pi\|_2^2\right), \qquad \kappa_k = \begin{cases} 1, & k = 1, \\ 2, & k > 1. \end{cases} \tag{11}$$

Because $c_\pi \neq 0$ and $\sigma(t) > 0$, the variance in Eq. 11 is strictly positive. Hence the Gaussian distribution has positive density on all of $\mathbb{R}$, and in particular

$$\Pr(\delta e_{k,t} < 0) = \Phi\left(-\frac{\alpha(t)\delta e^\star_k}{\sqrt{\kappa_k}\,\sigma(t)\|c_\pi\|_2}\right) > 0, \qquad k = 1, \dots, K,$$

where Φ denotes the standard Gaussian cumulative distribution function.

Therefore, ambient Gaussian interpolation violates the fixed-residue constraint almost surely and assigns strictly positive probability to negative mechanical-progress increments.

A.3 Canonical Mechanical Carrier

The reaction coordinate determines which deformations change extension but does not uniquely specify how a unit extension should be distributed over all residues. To define a canonical direction, let $M_G \in \mathbb{R}^{3N\times 3N}$ denote the graph-derived mechanical metric. We assume that M_G is symmetric and positive definite on $\mathcal{V}_a$, i.e.,

$$v^\top M_G v > 0 \qquad \text{for all } v \in \mathcal{V}_a \setminus \{0\}.$$

The canonical mechanical carrier is defined by

$$b_\pi := \arg\min_{b\in\mathcal{V}_a} \frac{1}{2} b^\top M_G b \qquad \text{subject to} \qquad c_\pi^\top b = 1. \tag{12}$$

The normalization $c_\pi^\top b_\pi = 1$ means that one unit of displacement along b_π produces exactly one unit of projected extension. In this work, we construct the mechanical metric as

$$M_G = (L_{\text{chain}} + \lambda I_N) \otimes I_3,$$

where L_{chain} is the graph Laplacian of the protein backbone, with edges connecting consecutive residues. All backbone edges are assigned a uniform weight $w_i = 1$, yielding an unweighted chain graph, and we use a small diagonal regularization $\lambda = 10^{-4}$ to remove the Laplacian null mode and ensure numerical stability. The resulting metric penalizes spatially non-smooth displacement fields along the backbone, so that the mechanical carrier b_π distributes the unit extension smoothly across the chain rather than localizing it at the pulled residue.

Existence and uniqueness. Let $J_a \in \mathbb{R}^{3N\times 3(N-1)}$ be the canonical inclusion matrix that inserts a zero three-vector at the anchor coordinates. Every $b \in \mathcal{V}_a$ can be written uniquely as $b = J_a \bar{b}$. Define

$$\bar{M}_G := J_a^\top M_G J_a, \qquad \bar{c}_\pi := J_a^\top c_\pi.$$

Since M_G is positive definite on $\mathcal{V}_a$, $\bar{M}_G \succ 0$. Moreover, because $a \neq b$ and $\|\widehat{u}\|_2 = 1$, we have $\bar{c}_\pi \neq 0$.

Problem 12 is therefore equivalent to

$$\min_{\bar{b}} \frac{1}{2} \bar{b}^\top \bar{M}_G \bar{b} \qquad \text{subject to} \qquad \bar{c}_\pi^\top \bar{b} = 1.$$

Its Lagrangian is

$$\mathcal{L}(\bar{b}, \lambda) = \frac{1}{2} \bar{b}^\top \bar{M}_G \bar{b} - \lambda(\bar{c}_\pi^\top \bar{b} - 1).$$

Stationarity gives $\bar{M}_G \bar{b} = \lambda \bar{c}_\pi$, and hence

$$\bar{b} = \lambda \bar{M}_G^{-1} \bar{c}_\pi.$$

Enforcing the unit-extension constraint yields

$$\lambda = \frac{1}{\bar{c}_\pi^\top \bar{M}_G^{-1} \bar{c}_\pi},$$

so the carrier has the explicit form

$$b_\pi = J_a \frac{\bar{M}_G^{-1} \bar{c}_\pi}{\bar{c}_\pi^\top \bar{M}_G^{-1} \bar{c}_\pi}. \tag{13}$$

Because the objective is strictly convex on $\mathcal{V}_a$ and the constraint is affine and feasible, this solution is unique.

Thus b_π is the unique minimum-M_G-cost anchor-preserving deformation that produces unit projected extension. The precise physical meaning of the cost is determined by the choice of M_G, while the geometric results below require only

$$b_\pi \in \mathcal{V}_a, \qquad c_\pi^\top b_\pi = 1.$$

A.4 Proof of Proposition 1

Recall the linear operator

$$P_\pi := I - b_\pi c_\pi^\top.$$

Proof 2 *We first establish that*

$$\mathrm{span}\{b_\pi\} \cap \mathcal{T}_\pi = \{0\}.$$

Let $v \in \mathrm{span}\{b_\pi\} \cap \mathcal{T}_\pi$. Then $v = \lambda b_\pi$ for some $\lambda \in \mathbb{R}$. Since $v \in \mathcal{T}_\pi$, $c_\pi^\top v = 0$, and therefore

$$0 = c_\pi^\top v = \lambda c_\pi^\top b_\pi = \lambda.$$

Hence $v = 0$.

Now let $\Delta x \in \mathcal{V}_a$ be arbitrary and define

$$e := c_\pi^\top \Delta x, \qquad \zeta := P_\pi \Delta x = \Delta x - e b_\pi.$$

Because both Δx and b_π belong to $\mathcal{V}_a$, we have $\zeta \in \mathcal{V}_a$. Furthermore,

$$c_\pi^\top \zeta = c_\pi^\top \Delta x - e\, c_\pi^\top b_\pi = e - e = 0.$$

Therefore $\zeta \in \mathcal{V}_a \cap \ker(c_\pi^\top) = \mathcal{T}_\pi$, and

$$\Delta x = e b_\pi + \zeta.$$

Since Δx was arbitrary,

$$\mathcal{V}_a = \mathrm{span}\{b_\pi\} + \mathcal{T}_\pi.$$

Together with the trivial intersection established above, this yields the direct sum

$$\mathcal{V}_a = \mathrm{span}\{b_\pi\} \oplus \mathcal{T}_\pi. \tag{14}$$

The decomposition is unique. Suppose

$$\Delta x = e_1 b_\pi + \zeta_1 = e_2 b_\pi + \zeta_2, \qquad \zeta_1, \zeta_2 \in \mathcal{T}_\pi.$$

Then

$$(e_1 - e_2) b_\pi = \zeta_2 - \zeta_1.$$

The left-hand side belongs to $\mathrm{span}\{b_\pi\}$, whereas the right-hand side belongs to $\mathcal{T}_\pi$. By Eq. 14, their intersection contains only the zero vector. Hence

$$e_1 = e_2, \qquad \zeta_1 = \zeta_2.$$

Moreover, applying $c_\pi^\top$ to the decomposition gives

$$c_\pi^\top \Delta x = e\, c_\pi^\top b_\pi + c_\pi^\top \zeta = e,$$

so the two components are recovered explicitly as

$$e = c_\pi^\top \Delta x, \qquad \zeta = P_\pi \Delta x.$$

Thus

$$\Phi_\pi : \mathcal{V}_a \longrightarrow \mathbb{R} \times \mathcal{T}_\pi, \qquad \Phi_\pi(\Delta x) = \big(c_\pi^\top \Delta x, P_\pi \Delta x\big)$$

is a linear bijection with inverse

$$\Phi_\pi^{-1}(e, \zeta) = e b_\pi + \zeta.$$

Therefore the progress–conformation coordinates (e, ζ) are identifiable for every anchor-preserving displacement.

Projection property. The operator P_π is a projection because

$$\begin{aligned} P_\pi^2 &= (I - b_\pi c_\pi^\top)^2 \\ &= I - 2 b_\pi c_\pi^\top + b_\pi (c_\pi^\top b_\pi) c_\pi^\top \\ &= I - b_\pi c_\pi^\top = P_\pi. \end{aligned}$$

Restricted to $\mathcal{V}_a$, it satisfies

$$\mathrm{Im}(P_\pi) = \mathcal{T}_\pi, \qquad \ker(P_\pi) = \mathrm{span}\{b_\pi\}.$$

Indeed, $c_\pi^\top P_\pi = 0$, so $P_\pi(\mathcal{V}_a) \subseteq \mathcal{T}_\pi$; conversely, $P_\pi \zeta = \zeta$ for every $\zeta \in \mathcal{T}_\pi$. Similarly, $P_\pi b_\pi = 0$, and the direct-sum decomposition shows that no other component lies in its kernel. Hence P_π is the projection onto the equal-extension tangent space along the mechanical carrier direction. It is generally an oblique, rather than Euclidean-orthogonal, projection.

A.5 Latent-Space Closure under Stochastic Interpolation

The generative variable is

$$Y = \left(\gamma_{1:K}, \widetilde{\zeta}_{1:K}\right) \in \mathcal{Y}_\pi,$$

with

$$\mathcal{Y}_\pi := \mathbb{R}^K \times \mathcal{V}_a^K$$

Although $\mathcal{Y}_\pi$ incorporates the linear anchor condition in the conformational branch, it is itself a finite-dimensional real vector space. Equivalently, after removing the three fixed-anchor coordinates, it can be treated as an ordinary Euclidean latent space.

For the Gaussian base distribution, the conformational component must be sampled intrinsically in $\mathcal{V}_a$. For example, using the inclusion matrix J_a above, one may draw

$$\bar{\varepsilon}_\zeta \sim \mathcal{N}(0, I_{3(N-1)})$$

and set

$$\varepsilon_\zeta = J_a \bar{\varepsilon}_\zeta \in \mathcal{V}_a.$$

This is equivalent to masking the anchor coordinates of the ambient Gaussian noise.

Lemma 1 (Closure of the latent space under interpolation) *Let $Y_0, Y_1 \in \mathcal{Y}_\pi$, and let*

$$Y_t = \alpha(t) Y_1 + \sigma(t) Y_0$$

for arbitrary scalar functions α, σ. Then $Y_t \in \mathcal{Y}_\pi$ for every t for which $\alpha(t)$ and $\sigma(t)$ are defined.

Proof 3 *The spaces $\mathbb{R}^K$, $\mathcal{V}_a^K$, and $\mathbb{R}^{K\times N\times 3}$ are linear subspaces of their respective ambient Euclidean spaces. Their Cartesian product $\mathcal{Y}_\pi$ is therefore a vector space and is closed under scalar multiplication and addition. Hence, for $Y_0, Y_1 \in \mathcal{Y}_\pi$,*

$$\alpha(t) Y_1 + \sigma(t) Y_0 \in \mathcal{Y}_\pi.$$

The lemma guarantees that every latent state remains within the domain on which the mechanical realization map is well-defined. Physical path admissibility follows after realization, as proved next.

A.6 Proof of Proposition 2

Recall that path ordering is parameterized by

$$\delta e_k = \varepsilon + \mathrm{softplus}(\gamma_k), \qquad e_0 = 0, \qquad e_k = \sum_{j=1}^{k} \delta e_j,$$

where $\varepsilon > 0$. The mechanically admissible realization map is defined by

$$\mathrm{vec}(\widehat{X}_k - X_0) = e_k b_\pi + P_\pi \widetilde{\zeta}_k,$$

Proof 4 *Fix an arbitrary generative time $t \in [0, 1]$. By Lemma 1,*

$$Y_t = \left(\gamma_{1:K,t}, \widetilde{\zeta}_{1:K,t}\right) \in \mathcal{Y}_\pi.$$

We suppress the index t temporarily for clarity.

First, since

$$\mathrm{softplus}(z) = \log(1 + e^z) > 0 \qquad \textit{for every } z \in \mathbb{R},$$

the extension increments satisfy

$$\delta e_k = \varepsilon + \mathrm{softplus}(\gamma_k) > \varepsilon > 0.$$

Consequently,

$$e_k - e_{k-1} = \delta e_k > 0, \qquad k = 1, \ldots, K.$$

Next consider the realized displacement

$$\widehat{\Delta x}_k := \mathrm{vec}(\widehat{X}_k - X_0) = e_k b_\pi + P_\pi \widetilde{\zeta}_k.$$

Because $b_\pi \in \mathcal{V}_a$,

$$(b_\pi)_a = 0.$$

Furthermore, $\widetilde{\zeta}_k \in \mathcal{V}_a$, *and*

$$P_\pi \widetilde{\zeta}_k = \widetilde{\zeta}_k - b_\pi c_\pi^\top \widetilde{\zeta}_k \in \mathcal{V}_a.$$

Hence

$$(\widehat{\Delta x}_k)_a = 0,$$

which is equivalent to

$$\widehat{x}_{k,a} = x_{0,a}.$$

Thus every realized state satisfies the fixed-anchor condition.

For the reaction coordinate, using $c_\pi^\top b_\pi = 1$ *and*

$$c_\pi^\top P_\pi = c_\pi^\top - (c_\pi^\top b_\pi) c_\pi^\top = 0,$$

we obtain

$$\begin{aligned} \ell_\pi(\widehat{X}_k) - \ell_\pi(X_0) &= c_\pi^\top \widehat{\Delta x}_k \\ &= e_k c_\pi^\top b_\pi + c_\pi^\top P_\pi \widetilde{\zeta}_k \\ &= e_k. \end{aligned}$$

Therefore

$$\begin{aligned} \ell_\pi(\widehat{X}_k) - \ell_\pi(\widehat{X}_{k-1}) &= e_k - e_{k-1} \\ &= \delta e_k \\ &> 0. \end{aligned}$$

The realized coordinate path thus satisfies both defining conditions of Eq. 10, and hence

$$\widehat{X}_{0:K} \in \mathcal{P}_\pi(X_0).$$

The argument holds for an arbitrary $t \in [0, 1]$. *Therefore*

$$\mathcal{R}_\pi(Y_t) \in \mathcal{P}_\pi(X_0) \qquad \forall t \in [0, 1]. \tag{15}$$

Thus the stochastic interpolation may be carried out in the intrinsic linear latent space $\mathcal{Y}_\pi$, *while its image under* $\mathcal{R}_\pi$ *remains mechanically admissible at every generative time.*

B Dataset Construction and Preprocessing

B.1 Steered Molecular Dynamics Data

We construct the dataset from constant-velocity steered molecular dynamics (SMD) simulations. Candidate proteins are indexed from a BSDB-derived protein collection Sikora et al. (2010), with the corresponding PDB/mmCIF structures used to initialize the simulations. In total, we simulate 1350 proteins with five independent SMD replicas per protein, yielding 6750 unfolding trajectories. Replicas share the same protein and pulling condition but use independent Langevin seeds, thereby providing multiple stochastic realizations of the same mechanical unfolding process.

Each protein is represented by a one-bead-per-residue C_α structure-based model derived from its native structure. Local backbone geometry is stabilized by native bond, angle, and dihedral interactions, while nonlocal native contacts encode tertiary interactions present in the reference conformation. These contacts are identified from residue-level heavy-atom van der Waals overlap, excluding near-neighbor pairs along the sequence, and are represented by native-centered attractive interactions between the corresponding C_α beads.

Before pulling, each structure is energy minimized and simulated using stochastic dynamics at $T = 300\,\mathrm{K}$ with a $1\,\mathrm{fs}$ integration step. The first residue is fixed at the origin, while the last residue

is initially positioned along the positive x-axis and pulled in the $+x$ direction by a harmonic potential with spring constant $k_{\mathrm{s}} = 100\,\mathrm{kJ\,mol^{-1}\,nm^{-2}} \approx 166\,\mathrm{pN\,nm^{-1}}$ whose anchor is displaced a constant rate $v_{\mathrm{pull}} = 0.005\,\mathrm{nm\,ps^{-1}}$.

The pulling coordinate and force are recorded together with the structural trajectory. The simulation length is chosen such that the moving restraint spans approximately the contour-length scale of the protein, using $0.36\,\mathrm{nm}$ per residue.

B.2 Progress-State Construction

Dense SMD trajectories contain high-frequency thermal fluctuations and need not exhibit strictly monotonic instantaneous molecular extension. Our objective, however, is to model the large-scale structural progression of unfolding rather than these fast fluctuations. We therefore reparameterize each SMD trajectory by its projected extension along the pulling direction and retain K representative states at successively increasing progress levels. The resulting trajectory is thus an ordered structural pathway.

We first construct a pulling progress p^{pull} from the external pulling coordinate recorded by GROMACS. When this signal is unavailable, we use the prescribed pulling rate multiplied by simulation time, and finally the projected molecular extension as a fallback. The resulting coordinate is made nondecreasing by a cumulative maximum and normalized to $[0, 1]$.

To retain structural transitions that are not uniformly distributed along extension, we additionally define a structural progress coordinate. For adjacent raw frames, the structural increment combines four quantities: $\delta_\tau^{\mathrm{str}} = \frac{1}{4}\Big[\mathcal{S}(|\Delta\ell_\tau|) + \mathcal{S}(|\Delta Q_\tau|) + \mathcal{S}(\Delta C_\tau) + \mathcal{S}(\Delta R_\tau)\Big]$, where ℓ is the end-to-end molecular extension, Q is the native-contact fraction, ΔC is the fraction of native contacts whose binary survival state changes between adjacent frames, and $\Delta R_\tau = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \|x_i(\tau+1) - x_i(\tau)\|_2^2}$ is the coordinate RMS displacement. Here $\mathcal{S}(\cdot)$ denotes the robust normalization used to place the four quantities on comparable scales. A native contact is considered present when its instantaneous distance is below 1.5 times its native reference distance. The cumulative structural increments are subsequently normalized to obtain $p^{\mathrm{str}} \in [0, 1]$.

The progress coordinate used for trajectory construction is the hybrid $p = \mathcal{N}_\uparrow\left(0.65\,p^{\mathrm{pull}} + 0.35\,p^{\mathrm{str}}\right)$, where $\mathcal{N}_\uparrow$ denotes monotonic normalization to $[0, 1]$. This construction keeps the global ordering tied primarily to the externally imposed pulling coordinate while allocating additional resolution to regions with substantial conformational or contact rearrangement.

We identify the unfolding endpoint when $Q \le 0.10$, $\frac{\ell}{3.6(N-1)\,\text{Å}} \ge 0.80$ for three consecutive raw frames. Trajectories that do not reach this criterion are retained and marked as incomplete in the processed dataset.

Finally, we place $K = 15$ uniformly spaced target values on the normalized progress interval and select the nearest corresponding *real SMD frames*, subject to strictly increasing source-frame indices. The first and final valid states are retained explicitly. We neither interpolate Cartesian coordinates nor average neighboring structures; consequently, every state used for training corresponds to an actual SMD conformation while the resulting sequence is approximately uniform in mechanical progress.

B.3 Train–Validation–Test Split

We split the dataset at the *protein level*. The reported $80/10/10$ train–validation–test partition contains $1{,}080$, 135, and 135 proteins, respectively. Since each protein has five independent replicas, this corresponds to $5{,}400$, 675, and 675 SMD trajectories before any preprocessing failures or filtering. All replicas belonging to the same protein are assigned to the same partition, preventing structural and trajectory-level leakage between training and evaluation.

The dataset builder makes the assignment deterministically from the protein identifier and a fixed split seed, so the same raw dataset always reproduces the same partition.

C Model and Representation Details

C.1 Residue Frame Construction

Residue orientations are constructed from the CA/CG trace using an initial-reference gauge. For residue i at state k, we first define the local tangent

$$t_{k,i} = \begin{cases} \dfrac{x_{k,2} - x_{k,1}}{\|x_{k,2} - x_{k,1}\|}, & i = 1, \\ \dfrac{x_{k,i+1} - x_{k,i-1}}{\|x_{k,i+1} - x_{k,i-1}\|}, & 1 < i < N, \\ \dfrac{x_{k,N} - x_{k,N-1}}{\|x_{k,N} - x_{k,N-1}\|}, & i = N. \end{cases}$$

Thus the frame x-axis follows the local chain direction while the one-sided definitions at the two termini avoid introducing virtual residues.

The complete frame at $k = 0$,

$$R_{0,i} = [t_{0,i}, y_{0,i}, z_{0,i}] \in SO(3),$$

is initialized using a spatial Bishop construction along the native CA trace. This fixes the otherwise arbitrary transverse y/z gauge once for the reference structure.

For every subsequent state, the frame is constructed independently from the initial frame:

$$R^{\text{ref}}_{k,i} = \text{MinRot}\left(t_{0,i} \to t_{k,i}\right) R_{0,i}.$$

Here $\text{MinRot}(a \to b)$ is the minimum-angle proper rotation mapping unit vector a to b. The transported transverse axis is re-orthogonalized against $t_{k,i}$, followed by

$$z_{k,i} = t_{k,i} \times y_{k,i}, \qquad y_{k,i} = z_{k,i} \times t_{k,i},$$

to maintain a right-handed orthonormal frame. Parallel tangents produce the identity rotation. For the antiparallel case, where the rotation axis is not unique, a deterministic Cartesian fallback axis is used.

The generative orientation variable is then the frame-0-relative rotation

$$\omega_{k,i} = \log_{SO(3)}\left(R_{0,i}^{\top} R_{k,i}\right),$$

and reconstruction uses

$$\widehat{R}_{k,i} = R_{0,i}\,\text{Exp}\left([\omega_{k,i}]_{\times}\right).$$

Unlike adjacent-frame transport, $R_{k,i}$ therefore depends only on the initial frame and the current state and does not accumulate rotational errors along the pathway.

The implementation also supports an optional local-angle block. In the CA-only setting, one virtual dihedral is computed from $(i-1, i, i+1, i+2)$ and encoded relative to its initial value as

$$\left[\sin(\Delta\phi_{k,i}), \cos(\Delta\phi_{k,i}) - 1\right].$$

Boundary residues for which the four-residue dihedral is undefined are masked out.

C.2 Network Architecture

The denoising network operates on the complete noisy pathway

$$Y_t \in \mathbb{R}^{K \times N \times d}$$

and generates all K pathway states jointly. For the main experiments, $K = 15$. Each latent token is first mapped to a 384-dimensional hidden representation by a learned linear projection.

Reference encoding. The initial structure is encoded once by a prepended invariant point attention (IPA) module. Amino-acid identities are embedded as residue features, while the initial residue frames $R_{0,i}, x_{0,i}$ provide the geometric reference to IPA. In the pair-aware configuration, static GeoFormer edge representations are projected to the IPA pair channel and provide residue-pair context for the folded state. The resulting reference representation

$$H_{\mathrm{ref}} \in \mathbb{R}^{N\times 384}$$

is added to the latent tokens at every progress state,

$$H^{(0)}_{k,i} = W_{\mathrm{in}} Y_{t,k,i} + H_{\mathrm{ref},i} + H_{\mathrm{cond},k,i}.$$

This supplies every pathway state with the same folded-structure reference without exposing any future SMD state.

The reference IPA uses four heads, a hidden point-channel dimension of 32, eight query/key points, and eight value points per head. The current main configuration uses the pair-aware reference encoder but does not repeatedly cross-attend to X_0 after every trajectory block.

Factorized spatiotemporal blocks. The backbone consists of four Transformer blocks with hidden dimension 384, 16 attention heads, and an FFN dimension of $4 \times 384 = 1536$. Within each block, the $K \times N$ pathway grid is factorized into two attention operations.

First, residue attention is applied independently at every pathway state:

$$H_{k,:} \leftarrow \mathrm{Attn}_{\mathrm{res}}(H_{k,:}),$$

allowing long-range communication among residues within one conformation. Second, the tensor is transposed and path attention is applied independently for each residue:

$$H_{:,i} \leftarrow \mathrm{Attn}_{\mathrm{path}}(H_{:,i}),$$

which couples the same residue across all unfolding states. This gives the factorized complexity

$$\mathcal{O}(KN^2 + NK^2)$$

rather than full attention over all KN tokens.

Both attention operators use rotary positional embeddings and masked multi-head self-attention. Residue masks remove padded residues, and the same masks are propagated through temporal attention so padded trajectory tokens do not participate in message passing. Dropout is set to zero in the reported configuration.

Generative-time modulation. The flow time t is mapped by a sinusoidal timestep embedder followed by an MLP. Each spatiotemporal block converts this embedding into nine 384-dimensional vectors,

$$(\beta_r, \alpha_r, g_r,\ \beta_p, \alpha_p, g_p,\ \beta_f, \alpha_f, g_f),$$

which provide shift, scale, and residual gating for residue attention, path attention, and the FFN, respectively. Schematically,

$$\begin{aligned} H' &= H + g_r \odot \mathrm{Attn}_{\mathrm{res}}\big(\mathrm{adaLN}(H; t)\big), \\ H'' &= H' + g_p \odot \mathrm{Attn}_{\mathrm{path}}\big(\mathrm{adaLN}(H'; t)\big), \\ H^{+} &= H'' + g_f \odot \mathrm{FFN}\big(\mathrm{adaLN}(H''; t)\big). \end{aligned}$$

The final modulation layers are initialized to zero, following the adaLN-Zero design, such that the residual branches initially make negligible updates. A time-conditioned final projection maps the 384-dimensional features back to the latent representation dimension and predicts the flow velocity over the entire pathway.

C.3 Conditioning Features

The conditioning information enters MechPath through both the neural network and the mechanical representation.

Sequence. The amino-acid sequence $S = (a_1, \ldots, a_N)$ is converted to residue-type indices and embedded inside the reference IPA encoder. Sequence identity therefore conditions the folded reference representation before it is broadcast over the unfolding pathway. We use static GeoFormer pair representations derived from the same sequence as pairwise features in IPA.

Initial structure. The initial structure X_0 enters at several levels. Its coordinates define the zero point of every generated displacement, its residue frames provide the geometric input to the IPA reference encoder, and its initial anchor geometry defines the pulling direction, covector c_π, and mechanical carrier b_π. Thus all generated conformations and orientations remain explicitly referenced to the same folded structure.

Mechanical topology. The mechanical topology G enters the realization geometry through the metric M_G used to construct b_π. The default experiments use the regularized chain Laplacian. The pair-aware IPA additionally receives learned static residue-pair features from GeoFormer. These pair features condition geometric reasoning in the reference encoder but are distinct from the metric M_G defining the PullGauge carrier.

Pulling geometry. The pulling setup enters structurally rather than through a free global embedding. The fixed and pulled residue indices and the initial pulling direction determine

$$c_\pi, \qquad b_\pi, \qquad P_\pi,$$

and therefore determine which latent direction represents mechanical progress and which directions represent equal-extension conformational variation. The progress channel itself is placed only on the pulled-anchor token, while the fixed-anchor constraint is imposed exactly during realization. Because the processed trajectories are expressed in the canonical pulling frame, this representation also removes arbitrary global orientation relative to the pulling axis.

D Training and Sampling Details

D.1 Flow-Matching Objective

We train the model with conditional flow matching in the latent progress–conformation representation . Let Y_1 denote the encoded mechanical unfolding pathway and let

$$Y_0 \sim \mathcal{N}(0, I)$$

denote the Gaussian reference state. We use the geometric variance-preserving interpolant

$$Y_t = \alpha(t)Y_1 + \sigma(t)Y_0, \qquad t \sim \mathcal{U}(0,1),$$

with

$$\alpha(t) = \sin\left(\frac{\pi t}{2}\right), \qquad \sigma(t) = \cos\left(\frac{\pi t}{2}\right).$$

The corresponding conditional velocity is

$$u_t = \dot\alpha(t)Y_1 + \dot\sigma(t)Y_0 = \frac{\pi}{2}\cos\left(\frac{\pi t}{2}\right)Y_1 - \frac{\pi}{2}\sin\left(\frac{\pi t}{2}\right)Y_0.$$

The network $v_\theta(Y_t, t; \mathcal{C})$ predicts this velocity, where $\mathcal{C}$ denotes the sequence, initial structure, and pulling-dependent conditioning information.

D.2 Structural Auxiliary Objectives

Flow matching is supplemented with differentiable structural objectives computed from the clean pathway implied by the predicted velocity. Given Y_t and $v_\theta(Y_t, t)$, the corresponding estimate of the clean endpoint Y_1 is obtained analytically from the interpolant:

$$\widehat{Y}_1 = \frac{\dot\sigma(t)Y_t - \sigma(t)v_\theta(Y_t,t)}{\alpha(t)\dot\sigma(t) - \dot\alpha(t)\sigma(t)}.$$

We then decode $\widehat{Y}_1$ through the same realization map used at inference to obtain $\widehat{X}_{1:K}$. Structural losses therefore act on predicted clean structures rather than directly on the noisy intermediate Y_t.

Backbone bond geometry. Let

$$d_{k,i} = \|x_{k,i+1} - x_{k,i}\|_2$$

denote the consecutive CA/CG distance and let d_i^0 denote its value in the initial structure. We use the relative squared bond error

$$\mathcal{L}_{\text{bond}} = \frac{1}{|\mathcal{E}_{\text{chain}}|K} \sum_{k=1}^{K} \sum_{i \in \mathcal{E}_{\text{chain}}} \left(\frac{\widehat{d}_{k,i} - d_i^0}{d_i^0} \right)^2,$$

where invalid or padded residue pairs are excluded.

Native-contact pathway supervision. We further supervise the evolution of native contacts along the complete decoded pathway. For each native contact $c = (i, j)$, let d_{ij}^0 be its distance in X_0. The generated soft contact-survival probability is

$$\widehat{q}_{k,ij} = \text{sigmoid}\left(\frac{\lambda_c d_{ij}^0 - \widehat{d}_{k,ij}}{\tau_c} \right),$$

with

$$\lambda_c = 1.5, \qquad \tau_c = 0.5\,\text{Å}.$$

Pairs separated by at most two residues along the sequence are excluded. Using the native-contact list and survival targets stored in the processed dataset, the contact objective is

$$\mathcal{L}_{\text{contact}} = \frac{1}{K|\mathcal{C}_0|} \sum_{k=1}^{K} \sum_{(i,j) \in \mathcal{C}_0} (\widehat{q}_{k,ij} - q_{k,ij})^2.$$

For the main experiment, the complete training loss is

$$\boxed{\mathcal{L} = \mathcal{L}_{\text{FM}} + g(t)(0.1\,\mathcal{L}_{\text{bond}} + 0.5\,\mathcal{L}_{\text{contact}})}$$

with

$$w_{\text{bond}} = 0.1,$$
$$w_{\text{contact}} = 0.5,$$

D.3 Optimization

We optimize all trainable parameters using AdamW with initial learning rate

$$10^{-4}$$

and weight decay

$$10^{-5}.$$

The learning rate follows cosine annealing over the full training run, with

$$T_{\max} = 50, \qquad \eta_{\min} = 10^{-5},$$

and the scheduler is updated once per epoch.

Models are trained for 50 epochs with a mini-batch size of 4 trajectories and gradient accumulation factor 1. Gradients are clipped to a maximum norm of 1.0. An exponential moving average of the model parameters is maintained during training with decay

$$\beta_{\text{EMA}} = 0.999.$$

Validation is performed once per epoch, the best validation checkpoint is retained, and additional checkpoints are written every five epochs. No fixed optimization-step budget is specified; the number of parameter updates is determined by the number of training batches in each epoch.

D.4 ODE SAMPLING

At inference, all active latent variables are initialized independently from the standard Gaussian reference distribution,

$$Y(0) \sim \mathcal{N}(0, I),$$

followed by multiplication with the representation active mask. Samples are generated by integrating the learned probability-flow ODE

$$\frac{dY(t)}{dt} = v_\theta(Y(t), t; \mathcal{C}), \qquad t \in [0, 1].$$

We use the adaptive Dormand–Prince solver (`dopri5`) with absolute and relative tolerances

$$\text{atol} = 10^{-6}, \qquad \text{rtol} = 10^{-3}.$$

E EVALUATION DETAILS

We evaluate generated trajectories at two complementary levels. Trajectory-level metrics quantify the structural and mechanical fidelity of individual generated unfolding pathways, whereas ensemble-level metrics compare the distributions formed by repeated unfolding events.

E.1 PATHWAY FIDELITY METRICS

Progress-aligned Cα path RMSD. Direct frame-by-frame RMSD can confound structural error with differences in the rate at which two trajectories advance through unfolding. We therefore align generated and reference trajectories by mechanical progress before comparing their structures.

For a trajectory X, we first compute the projected anchor extension

$$\ell_\pi(X_k) = \hat{u}^\top (x_{k,b} - x_{k,a}),$$

using the pulling axis $\hat{u}$ defined by the initial reference structure. After subtracting the initial value, the resulting progress is made non-decreasing by a cumulative maximum and normalized to $[0, 1]$. Generated and reference coordinates are then linearly interpolated onto a common progress grid $\{s_m\}_{m=1}^M$.

The Progress-Aligned Cα Path RMSD (PA-RMSD) is

$$\text{PA-RMSD} = \frac{1}{M} \sum_{m=1}^{M} \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left\| \widetilde{x}_{m,i}^{\text{gen}} - \widetilde{x}_{m,i}^{\text{ref}} \right\|_2^2}.$$

Fréchet path distance. The main benchmark additionally evaluates whole-path similarity using the discrete Fréchet distance from Path Similarity Analysis (PSA) (Seyler et al., 2015). For two ordered pathways $P = (P_1, \ldots, P_K)$ and $Q = (Q_1, \ldots, Q_L)$,

$$d_F(P, Q) = \min_{\mathcal{C}} \max_{(k,l) \in \mathcal{C}} d_{\text{RMSD}}(P_k, Q_l),$$

where $\mathcal{C}$ ranges over monotone couplings of the two frame sequences. This metric preserves pathway ordering while allowing different rates of progression.

The Fréchet Path Distance (FPD) reported in the main benchmark corresponds to the mean source-matched PSA Fréchet distance: a generated rollout is compared with the SMD replica from which its conditioning example was formed, and these distances are subsequently aggregated at the protein level. As with PA-RMSD, frame RMSDs are computed directly in the pulling-gauge coordinate system without an additional Kabsch fit.

E.2 MECHANICAL ADMISSIBILITY METRICS

Anchor drift. For fixed anchor a, anchor drift is the mean displacement from its generated initial position,

$$\text{AD} = \frac{1}{K} \sum_{k=1}^{K} \|x_{k,a} - x_{0,a}\|_2.$$

We also record the maximum drift

$$\mathrm{AD}_{\max} = \max_k \|x_{k,a} - x_{0,a}\|_2.$$

A mechanically admissible realization with an exactly fixed anchor therefore has $\mathrm{AD} = 0$ up to numerical precision.

Monotonicity satisfaction rate. Using the same projected extension

$$e_k = \hat{u}^\top (x_{k,b} - x_{k,a}),$$

we define the Monotonicity Satisfaction Rate (MSR) as

$$\mathrm{MSR} = \frac{1}{K-1} \sum_{k=1}^{K-1} \mathbf{1}\left[e_{k+1} - e_k \geq -\epsilon_{\mathrm{mono}}\right],$$

where

$$\epsilon_{\mathrm{mono}} = 10^{-5}\ \text{Å}$$

allows only floating-point-scale backtracking. MSR equals one when every generated step respects non-decreasing mechanical extension.

E.3 Contact-Release Metrics

Native contacts provide a topology-resolved description of the structural events underlying mechanical unfolding. A native contact $c = (i, j)$ with initial distance d_c^0 is considered intact at state k when

$$d_{k,c} \leq \lambda_c d_c^0, \qquad \lambda_c = 1.5.$$

Let

$$q_{k,c} \in \{0, 1\}$$

denote the corresponding contact-survival indicator.

Persistent rupture progress. For contact c, its rupture progress $T_c \in [0, 1]$ is the first progress at which the contact becomes broken and remains broken for the configured persistence window. The benchmark uses one stored frame, so T_c reduces to the first detected rupture. Contacts that never rupture are assigned

$$T_c = 1.$$

Rupture-order correlation. Rupture-Order Kendall Correlation (RO-τ) compares the vectors of contact rupture progress between a generated pathway and its reference:

$$\text{RO-}\tau = \tau_b\left(\{T_c^{\mathrm{gen}}\}_{c\in\mathcal{C}}, \{T_c^{\mathrm{ref}}\}_{c\in\mathcal{C}}\right).$$

We use Kendall's τ_b, which accounts for tied rupture progresses. The metric therefore evaluates relative ordering of contact-release events without requiring their absolute rupture progresses to coincide.

Contact-resolved Wasserstein distance. The CR-W_1 reported in the main benchmark evaluates contact release at the *ensemble* level rather than between a single generated/reference pair. For each protein, we first construct a consensus native-contact set from its five independent SMD replicas. A contact is retained when it occurs in at least 60% of the replicas, corresponding to at least three of five trajectories, and its native distance is taken as the median initial distance across the real replicas.

For each consensus contact c, the five real pathways and all generated pathways define empirical rupture-progress distributions,

$$F_c^{\mathrm{ref}}(s) = \Pr(T_c^{\mathrm{ref}} \leq s), \qquad F_c^{\mathrm{gen}}(s) = \Pr(T_c^{\mathrm{gen}} \leq s).$$

Their one-dimensional Wasserstein distance is

$$W_{1,c} = \int_0^1 \left|F_c^{\mathrm{gen}}(s) - F_c^{\mathrm{ref}}(s)\right| ds,$$

which is evaluated directly from the empirical rupture times without histogram binning. We then average over contact identity:

$$\boxed{\text{CR-W}_1 = \frac{1}{|\mathcal{C}_{\text{cons}}|} \sum_{c \in \mathcal{C}_{\text{cons}}} W_{1,c}}.$$

CR-W_1 therefore preserves which contact ruptures at which stage of unfolding. This distinguishes it from $Q(s)$-W_1, where errors from different contacts can cancel after averaging over contact identity, and from RO-τ, which measures ordering but discards absolute rupture progress.

E.4 Ensemble Coverage Metrics

PCA distribution JSD. To evaluate whether repeated generations reproduce the conformational support of repeated SMD trajectories, we compare their distributions in a common PCA space. For each protein, all noninitial SMD conformations are first aligned by proper Cα Kabsch alignment to the first real conformation. Generated conformations are aligned in the same manner. PCA is fitted using *only* the pooled real SMD conformations, and generated structures are projected into this fixed reference basis.

We retain PC1 and PC2 and standardize each component using the standard deviation of its real-SMD projection. Real and generated populations are converted into normalized two-dimensional histograms over

$$[-5, 5] \times [-5, 5].$$

The configured maximum resolution is 32×32 bins; for the finite five-replica reference ensemble, the implementation automatically reduces the number of bins when necessary to avoid a highly sparse reference histogram.

For real and generated histograms P and Q, respectively,

$$\text{PCA-JSD} = \frac{1}{2} D_{\text{KL}}(P \| M) + \frac{1}{2} D_{\text{KL}}(Q \| M), \qquad M = \frac{1}{2}(P + Q).$$

Natural logarithms are used. Lower PCA-JSD therefore indicates closer agreement between the generated and SMD-supported conformational distributions.

PSA pathway coverage and fidelity. The discrete Fréchet distance also provides directional ensemble measures. For the five real pathways $\mathcal{R} = \{R_i\}$ and generated ensemble $\mathcal{G} = \{G_j\}$, we define

$$d_{\text{R2G}} = \frac{1}{|\mathcal{R}|} \sum_i \min_j d_F(R_i, G_j),$$

and

$$d_{\text{G2R}} = \frac{1}{|\mathcal{G}|} \sum_j \min_i d_F(G_j, R_i).$$

R2G is the PSA *coverage distance*: it is small only when every observed SMD pathway has a nearby generated representative. G2R is the corresponding *fidelity distance*: it penalizes generated pathways that lie far from the observed SMD ensemble. These are the coverage quantities used in the sample-count analysis in Figure 5.

E.5 Aggregation across Replicas and Rollouts

The unit of evaluation is a *protein under one pulling condition*, not an individual SMD trajectory. Each test protein contains five independently simulated SMD replicas,

$$\mathcal{R}_p = \{R_{p,1}, \ldots, R_{p,5}\},$$

which together define the reference pathway ensemble. In the main benchmark, we independently generate one stochastic pathway from each replica-conditioned input, giving five generated trajectories

$$\mathcal{G}_p = \{G_{p,1}, \ldots, G_{p,5}\}.$$

More generally, if R stochastic rollouts are generated from each of the five source samples, the generated ensemble contains $5R$ pathways.

For trajectory-level metrics—including PA-RMSD, anchor drift, MSR, and RO-τ—each generated rollout is evaluated against its corresponding source SMD replica. All finite rollout scores belonging to the same protein are then pooled:

$$\bar{m}_p = \frac{1}{|\mathcal{G}_p|} \sum_{G \in \mathcal{G}_p} m(G),$$

with the within-protein standard deviation retained as a diagnostic.

Ensemble quantities—including CR-W_1, PCA-JSD, and PSA —are computed directly after pooling all five real replicas and all generated rollouts belonging to the same protein. Thus trajectory frames are never treated as independent ensemble samples.

Finally, each protein contributes equally to the benchmark. For each independent run r, we first compute the benchmark score for metric m as

$$M_m^{(r)} = \frac{1}{|\mathcal{P}_{\text{test}}|} \sum_{p \in \mathcal{P}_{\text{test}}} \bar{m}_p^{(r)},$$

where $\bar{m}_p^{(r)}$ denotes the metric aggregated over the replicas or generated rollouts associated with protein p, or the corresponding per-protein ensemble metric for distributional quantities. The values reported in the main text are the mean $\pm$ standard deviation of $M_m^{(r)}$ over three independent runs. This per-protein aggregation prevents proteins with additional replicas or generated rollouts from receiving greater statistical weight.

F Baseline Implementations

We adapt all baselines to the same conditional generation task used by MechPath: given the protein sequence and folded structure X_0, generate the remaining states of a $K = 15$ mechanical unfolding pathway.

F.1 MDGen

We implement MDGen (Jing et al., 2024) as an in-house ambient-space flow-matching baseline using the MDGen spatiotemporal backbone on the same initial-reference trajectory representation used throughout our codebase. Unlike MechPath, no pulling-specific geometric decomposition or mechanical constraint is imposed.

Each residue frame is represented relative to its folded reference frame by the full $SE(3)$ logarithm

$$\xi_{k,i} = \log\big(T_{0,i}^{-1} T_{k,i}\big) = [\rho_{k,i}, \omega_{k,i}] \in \mathbb{R}^6,$$

together with relative sine–cosine local-angle variables. Thus translations and rotations are modeled directly in the ambient residue-frame space. The model therefore does not explicitly enforce a fixed anchor, prescribed pulling direction, or monotonic extension.

The network uses four spatiotemporal blocks with hidden dimension 384, 16 residue-attention heads, and four IPA heads. Residue attention and bidirectional temporal attention are factorized as in the MDGen backbone, with rotary positional embeddings. The folded structure X_0 is provided through the initial residue-frame conditioning module. Pair-aware IPA, dynamic pair states, causal temporal residuals, and repeated reference conditioning are disabled.

We use the same geometric variance-preserving flow path as in our flow-matching models and predict the velocity field jointly over the complete trajectory. Models are trained for 50 epochs with AdamW, learning rate 10^{-4}, weight decay 10^{-5}, cosine learning-rate decay to 10^{-5}, batch size 4, and exponential moving average with decay 0.999.

F.2 4D DIFFUSION

We adapt 4D-Diffusion (Cheng et al., 2025) using the released architecture and diffusion implementation. 4D-Diffusion operates on residue $SE(3)$ frames and jointly denoises multiple structural states through spatial and temporal interactions.

For a 16-state trajectory

$$X_{0:15} = (X_0, X_1, \ldots, X_{15}),$$

we use X_0 as the sole structural condition and train the model to generate $X_{1:15}$.

The adapted model uses node and pair embedding dimensions of 128 and 64, respectively, four IPA heads, hidden IPA dimension 128, and two spatial–temporal blocks. GeoFormer residue and pair representations are supplied through the released feature interface. The model is trained for $100{,}000$ optimization steps with learning rate 10^{-4}, and the checkpoint with the lowest validation denoising loss is used for test inference.

During sampling, the model generates the 15 future states and the exact conditioned X_0 is prepended to recover the 16-state trajectory.

F.3 TEMPO

TEMPO (Xu et al., 2026) is a hierarchical autoregressive model originally designed to generate protein dynamics at coupled slow and fast physical-time scales. Because the states in our work are selected according to mechanical progress rather than uniform physical time, assigning TEMPO's original $20\,\mathrm{ns}$ and $1\,\mathrm{ns}$ scales to these states would be artificial. We therefore retain its core `ProteinSDE` and `LatentMDGenModel` architecture but use it as a *single-scale causal baseline* over the 16 ordered mechanical-progress states.

A key modification is the temporal conditioning length. The original TEMPO setup can condition on multiple observed frames, whereas our benchmark permits only the folded structure X_0 at inference. We therefore set

$$\texttt{threshold} = 1,$$

so generation follows

$$X_0 \to \widehat{X}_1 \to \widehat{X}_2 \to \cdots \to \widehat{X}_{15}.$$

Only X_0 is externally supplied. Each generated state is subsequently used as the structural condition for predicting the next state.

During training, we use the same causal one-step objective with scheduled sampling. For target state X_k, the model receives either the preceding ground-truth state X_{k-1} or its own prediction $\widehat{X}_{k-1}$. With scheduled-sampling probability $p_{\mathrm{TF}} = 0.5$, the training context is

$$\widetilde{X}_{k-1} = \begin{cases} X_{k-1}, & \text{with probability } p_{\mathrm{TF}}, \\ \widehat{X}_{k-1}, & \text{otherwise.} \end{cases}$$

The next-state prediction is optimized using the masked latent-space mean squared error employed by the adapted TEMPO `ProteinSDE` objective.

The model retains the released TEMPO architectural defaults: five layers, hidden dimension 384, 16 attention heads, four IPA heads, and the RNN-based temporal model. We train on the same training split as MECHPATH for $100{,}000$ optimization steps with learning rate 10^{-4}, gradient clipping at 1.0, and select the checkpoint with the lowest validation loss. The default adaptation is trained directly on the MechPath trajectories.

At test time, the TEMPO adapter is given only sequence information and X_0. The trajectory is generated autoregressively for 15 additional steps using the upstream inference noise scale of 5.0.

TEMPO internally subtracts a single translation determined from the initial structure. We apply this same X_0-derived translation to all states and invert it after generation:

$$X'_k = X_k - c_0, \qquad c_0 = \frac{1}{N}\sum_i x_{0,i}.$$

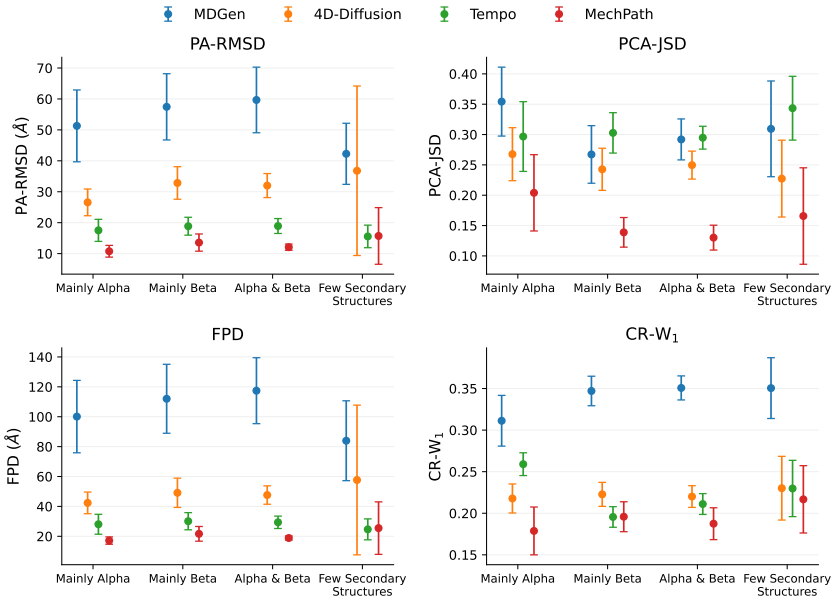

Figure 6: Generalization across protein structural classes. Class-wise performance of MDGen, 4D-Diffusion, TEMPO, and MechPath on the test set, grouped into Mainly Alpha, Mainly Beta, Alpha & Beta, and Few Secondary Structures. Performance is evaluated using progress-aligned RMSD (PA-RMSD), PCA Jensen–Shannon divergence (PCA-JSD), pathway-space Fréchet distance (FPD), and contact-rupture Wasserstein distance (CR-W_1); lower values indicate better agreement with SMD trajectories for all metrics. Points denote class-wise means and error bars indicate standard deviations.

F.4 Fair-Comparison Protocol

All methods use the same model-neutral export of the MechPath dataset, preserving the same protein-level train/validation/test split, the same independent SMD replicas, and exactly the same 16 selected mechanical-progress states. Manifest and protein-split SHA-256 fingerprints are stored to prevent baseline implementations from silently using different dataset partitions.

The common inference task is

$$(S, X_0) \longrightarrow \widehat{X}_{1:15},$$

The methods differ only in how they parameterize and generate the unknown trajectory. MDGen generates all noninitial states jointly through ambient-space flow matching; 4D-Diffusion jointly denoises the 15 future $SE(3)$ states; TEMPO generates them sequentially using causal autoregression; and MechPath transports the complete pathway in its mechanically admissible progress–conformation representation.

All outputs are converted back to the same stored pulling-frame coordinate system before evaluation. We apply no baseline-specific Kabsch alignment, anchor correction, pulling-axis correction, or monotonic-extension projection. Consequently, violations of the prescribed mechanical geometry remain visible to the evaluation metrics.

For the reported benchmark, the same number of stochastic trajectories is generated for each method under each test condition. Trajectory-level quantities and protein-level ensemble metrics are aggregated identically across methods.

Table 3: Robustness of trajectory-level structural comparison to different alignment strategies. Results are reported as mean ± standard deviation over three independent runs. Direct RMSD compares conformations at the same trajectory index, anchor-drift-corrected RMSD removes displacement of the fixed residue before comparison, and progress-aligned RMSD compares conformations at matched mechanical progress. The latter reduces sensitivity to model-dependent pathway parameterization while retaining differences in the generated conformations. Best results are shown in **bold**, and second-best results are underlined. All RMSD and anchor-drift values are in Å.

Model	Direct Cα path RMSD ↓	Anchor-drift-corrected Cα path RMSD ↓	Progress-aligned Cα path RMSD ↓	Anchor drift ↓
MDGen	48.42 ± 5.76	48.42 ± 5.78	47.53 ± 5.94	$\underline{0.52 \pm 0.24}$
4D-Diffusion	27.38 ± 9.86	25.58 ± 8.25	28.77 ± 12.23	16.63 ± 9.68
TEMPO	$\underline{18.92 \pm 0.03}$	$\underline{20.58 \pm 0.30}$	$\underline{17.95 \pm 0.04}$	8.33 ± 0.67
MechPath	$\mathbf{14.07 \pm 0.62}$	$\mathbf{14.07 \pm 0.62}$	$\mathbf{13.06 \pm 0.66}$	$\mathbf{0.00 \pm 0.00}$

G Additional Experiments

G.1 Robustness to trajectory alignment.

To assess whether the improvement measured by progress-aligned Cα path RMSD (PA-RMSD) depends on the alignment procedure, we additionally report direct path RMSD and anchor-drift-corrected path RMSD in Table 3. Direct RMSD compares structures at the same discrete pathway index and therefore assumes that different methods traverse the unfolding process at comparable rates. Anchor-drift correction removes errors arising from displacement of the prescribed fixed residue, but does not account for differences in mechanical progression along the pathway. In contrast, PA-RMSD compares conformations at matched mechanical progress, avoiding penalization caused solely by different parameterizations of the same force-driven structural evolution.

The relative performance is consistent across all three measures. MechPath achieves the lowest direct, anchor-drift-corrected, and progress-aligned RMSDs, with values of 14.07, 14.07, and 13.06 Å, respectively. TEMPO remains second best under each comparison, while MDGen and 4D-Diffusion exhibit substantially larger errors. Notably, the ranking is unchanged when progress alignment is removed, indicating that the improvement of MechPath is not an artifact of the proposed alignment procedure. Instead, progress alignment provides a common mechanical reference for comparing methods whose generated trajectories may traverse unfolding at different rates or use different temporal parameterizations.

G.2 MechPath remains robust across distinct native topologies.

Proteins with CATH annotations are grouped into Mainly Alpha-helix ($n = 24$), Mainly Beta-sheet ($n = 35$), Alpha & Beta ($n = 43$), and Few Secondary Structures ($n = 12$) (Waman et al., 2025). Across diverse protein topologies, MechPath maintains consistently low errors across structural, distributional, and contact-release metrics, as shown in Figure 6. Its PA-RMSD remains within 10.76–13.57 Å across these classes, compared with 17.51–18.92 Å for TEMPO, while similarly stable trends are observed for FPD and PCA-JSD. Importantly, this advantage persists despite the substantially different native contact organizations associated with α-rich, β-rich, and mixed folds, suggesting that the model is not specialized to a particular structural topology. The same consistency is visible in CR-W_1, indicating that the improved generalization extends beyond structural similarity to the ordering of contact release during unfolding. Performance differences are less conclusive for proteins with few secondary structures, where the small sample size and greater structural heterogeneity produce substantially larger uncertainty. Overall, these results show that MechPath preserves both pathway fidelity and contact-release behavior across diverse native topologies.

G.3 Contact-release dynamics.

Figure 7 illustrates the contact-release behavior across representative proteins with distinct native topologies. Despite substantial differences in their contact organization, MechPath consistently recovers both where native contacts rupture and in what order they are released along the unfolding coordinate. For example, in the top-row protein, MechPath reproduces the early rupture of peripheral contact groups such as A–B and A–F, while retaining the later-breaking central β-sheet contacts

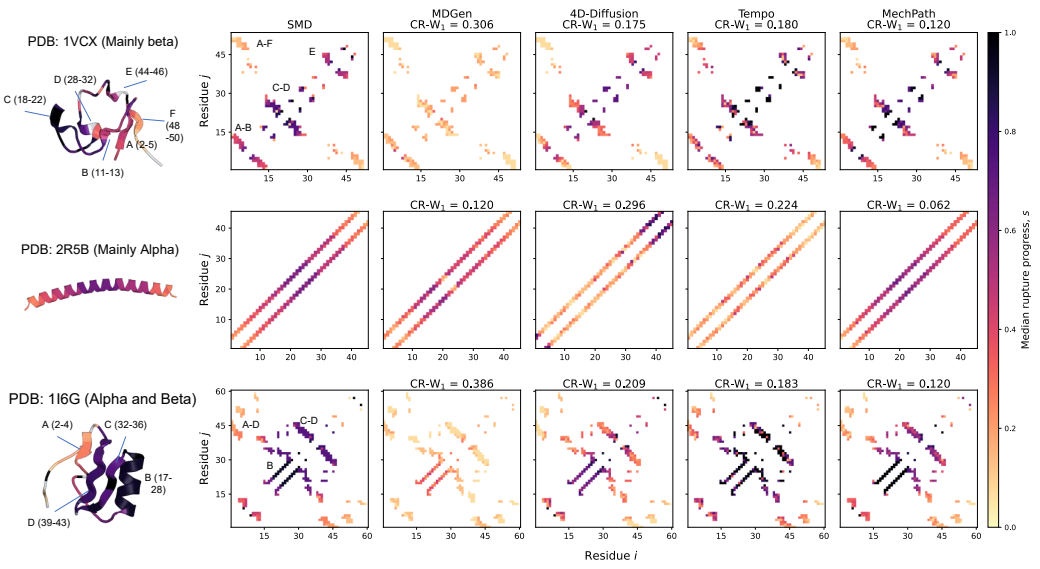

Figure 7: **Topology-dependent contact-release dynamics across representative proteins.** Structures and contact matrices on the right are colored by median normalized rupture progress s, from early (light) to late (dark) release. Labels denote representative secondary-structure elements and their associated contact regions.

around C–D, thereby recovering the mechanically resistant core observed in SMD. In the middle-row α-helical protein, it also captures the gradual, ordered loss of local helical contacts rather than an unstructured or premature collapse. Most notably, in the bottom-row protein (1i6g), MECHPATH recovers that the α-helix remains intact until after the major β-sheet contact network has ruptured. This ordering suggests that the β-sheet acts as a mechanical clamp that protects the helix during stretching. Such behavior is physically plausible and consistent with the well-established greater mechanical resistance of sheared β-sheet motifs relative to α-helices. These examples show that MECHPATH does not merely match a coarse rupture pattern, but reproduces topology-dependent contact-release mechanisms across different protein folds.