PROTEIN WORLD MODEL: LEARNING TEMPERATURE-CONDITIONED CONFORMATIONAL DYNAMICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Protein function emerges from conformational dynamics rather than isolated structures, yet molecular dynamics simulations remain prohibitively expensive for systematic experimentation across proteins and conditions. Existing generative models produce plausible trajectories, but plausibility alone is insufficient for controlled experimentation: predictions must depend on recent dynamical history, represent stochastic futures, and respond to external perturbations. We introduce a protein world model that infers predictive state from conformational history and generates stochastic future trajectories under controllable temperature actions. Trained on multi-temperature mdCATH, our model combines history-conditioned flow matching with differentiable geometric refinement to generate temporally coherent, geometrically valid multi-step rollouts. Extensive experiments show that the model captures trajectory distributions and dynamical observables more accurately than persistence and matched history-ablated controls, while controlled ablations demonstrate that it exploits temporal ordering beyond static conformational statistics. From a shared molecular state, temperature interventions induce systematic and physically consistent changes in predicted flexibility, structural organization, and conformational evolution, enabling controlled branching experiments across thermal conditions. The model further generalizes to unseen protein domains and temperatures and remains stable under closed-loop rollout. Together, these results establish protein world modeling as a framework for controllable, stochastic simulation of protein dynamics and in-silico experimentation.

1 INTRODUCTION

Proteins are dynamic molecular systems whose biological functions emerge from how they move between conformations, how long they remain in metastable states, and how these dynamics respond to environmental conditions. Understanding such behavior requires answering conditional questions about molecular evolution: Given recent protein motion, which transitions are likely to follow? How would future dynamics change under a different temperature? What uncertainty exists across repeated realizations? These questions require a predictive model that represents alternative futures and responds to controlled perturbations.

Molecular dynamics (MD) simulation provides a general-purpose computational tool for studying these questions, but at substantial computational cost. Resolving biologically relevant motions requires integrating atomic interactions over many timesteps, and each new protein, condition, or replicate requires additional simulation. This makes systematic experimentation expensive, particularly when one wishes to compare multiple conditions or estimate distributions over possible outcomes. A learned simulator offers a complementary possibility: expensive MD trajectories can be amortized into a model that generates many possible futures, allowing hypotheses and perturbations to be explored in silico before committing to additional simulation or physical experiments Husic & Pande (2018).

Recent generative models have made rapid progress toward this goal. Protein ensemble models can capture conformational heterogeneity and equilibrium observables (Jing et al., 2024a; Lewis et al., 2025), while a growing line of work directly generates temporally coherent molecular and protein trajectories (Jing et al., 2024b; Xu et al., 2025; Liu et al., 2026; Shoghi et al., 2026). However, generating a plausible molecular trajectory is not by itself sufficient for controlled in-silico experi-

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.

mentation. In a partially observed molecular system, the same instantaneous backbone conformation can arise from different microscopic states, including different momenta and solvent configurations, and therefore need not determine the same distribution over future motion. Recent conformational history can provide information about this latent dynamical state that is unavailable from a single structure. At the same time, external perturbations can alter the transition dynamics even when the observed molecular state is held fixed, and repeated realizations from the same state can lead to different futures. A useful learned simulator must therefore model not merely realistic trajectories, but a controlled conditional distribution of future trajectories. This requires three capabilities: inferring predictive state from temporal history, representing stochasticity over possible futures, and changing the future transition distribution in response to external controls. We refer to a model with this interface as a *protein world model* Ha & Schmidhuber (2018).

We introduce a model with this interface that observes a finite conformational history, infers a predictive molecular state, and generates multiple possible future trajectories. Crucially, we treat temperature not merely as metadata or a conditioning label, but as a controllable action that modulates the learned transition dynamics from a shared molecular state. Temperature provides a natural first control variable because changes in the thermal environment alter conformational transition probabilities and structural stability; multi-temperature trajectories therefore provide controlled perturbations for testing whether a learned simulator captures these changes in dynamics. We instantiate this formulation with a history-conditioned, geometry-preserving flow-matching model trained on the multi-temperature mdCATH dataset (Mirarchi et al., 2024). The resulting model enables a single observed molecular history to be branched into stochastic futures under different temperature interventions.

We evaluate the model according to the capabilities required by this formulation. First, it should represent a distribution over plausible futures rather than collapse to a single continuation. Second, its predictions should depend on the temporal organization of the observed history rather than only on an unordered collection of conformations. Third, changing the temperature action from a shared molecular state should produce systematic changes in future dynamics consistent with reference MD. Finally, generated trajectories should remain backbone-geometrically compatible and preserve relevant dynamical statistics under multi-step and closed-loop rollout. Across domain-disjoint proteins, our model improves trajectory prediction over current-state and temporally shuffled controls, responds systematically to temperature interventions, generalizes to held-out protein domains and temperatures, and remains stable under closed-loop continuation. Together, these results support protein world modeling as an interface for controlled stochastic simulation and repeated in-silico experimentation from a shared molecular state.

The core contributions of our study are as follows:

- We formulate protein world modeling as controlled conditional prediction of protein dynamics, characterized by three capabilities: dynamical memory, stochastic future prediction, and controllable intervention.
- We develop a geometry-preserving, history-conditioned flow-matching model in which temperature acts on the learned transition dynamics, enabling branching multi-step protein trajectories from a shared observed molecular state.
- We show empirically that the model uses temporal order beyond static conformational statistics, captures stochastic future distributions, responds systematically to temperature interventions, preserves backbone geometry and dynamical behavior, generalizes beyond the training domains and temperatures, and supports closed-loop molecular simulation.

2 Related Work

Protein ensembles and trajectory generation. Generative models have extended structure prediction to conformational distributions (Jing et al., 2024a; Lewis et al., 2025; Lu et al., 2024) and time-ordered molecular trajectories. MDGen formulates trajectories as generative objects with flexible conditioning (Jing et al., 2024b), TEMPO introduces hierarchical autoregressive generation (Xu et al., 2025), ConfRover couples temporal latent models with SE(3)-equivariant decoders (Shen et al., 2025), ProTDyn unifies sampling and simulation in a language modeling framework (Liu et al., 2026), and STAR-MD develops spatio-temporal SE(3) diffusion for long-horizon rollouts (Shoghi

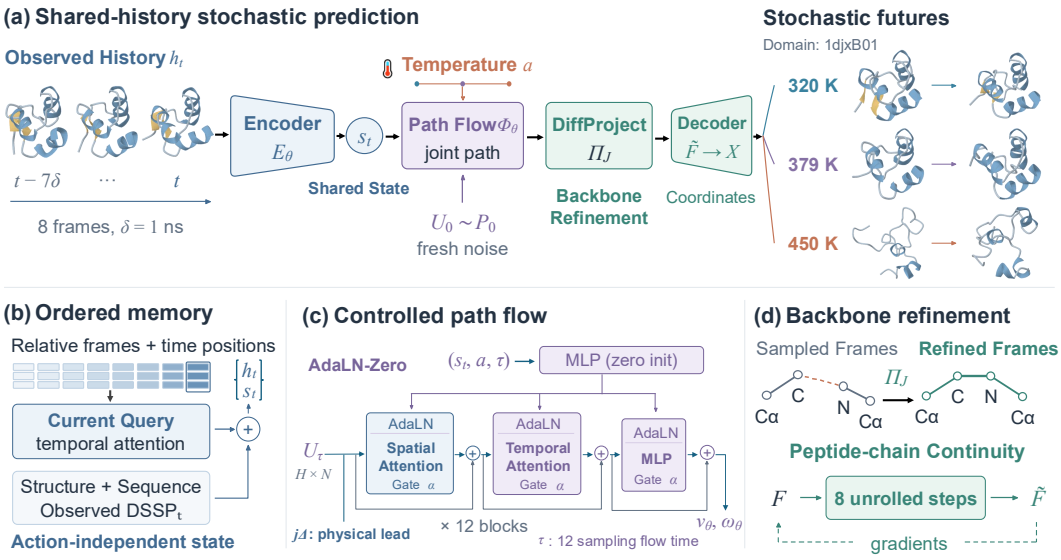


Figure 1: Overview of the proposed protein world model. **(a)** An observed conformational history is encoded into a shared molecular state and reused to generate stochastic future trajectories under different temperature actions. **(b)** The history encoder combines ordered residue-frame motion with current structural context. **(c)** The action-conditioned trajectory flow jointly transports all future residue frames using coupled spatial and temporal attention. **(d)** DIFFPROJECT differentiably refines neighboring residue frames to preserve peptide-chain continuity before coordinate decoding.

et al., 2026). These works establish trajectory generation, temporal context, and long-horizon simulation as key capabilities. Our contribution is orthogonal: we study a world-model interface where observed history defines current state and external actions select future transition distributions, enabling controlled branching from shared molecular states. Generative models have extended structure prediction to conformational distributions (Jing et al., 2024a; Lewis et al., 2025; Lu et al., 2024) and time-ordered molecular trajectories. MDGen formulates trajectories as generative objects with flexible conditioning (Jing et al., 2024b), TEMPO introduces hierarchical autoregressive generation (Xu et al., 2025), ConfRover couples temporal latent models with SE(3)-equivariant decoders (Shen et al., 2025), ProTDyn unifies sampling and simulation in a language modeling framework (Liu et al., 2026), and STAR-MD develops spatio-temporal SE(3) diffusion for long-horizon rollouts (Shoghi et al., 2026). These works establish trajectory generation, temporal context, and long-horizon simulation as key capabilities. Our contribution is orthogonal: we study a world-model interface where observed history defines current state and external actions select future transition distributions, enabling controlled branching from shared molecular states.

Temperature conditioning versus temperature as action. Temperature naturally controls molecular dynamics by changing equilibrium populations and kinetics. mdCATH provides multi-temperature simulations (Mirarchi et al., 2024), while TEMPO and DynaMode condition trajectory generation on temperature (Xu et al., 2025; Phipps et al., 2026). However, standard conditioning treats temperature as metadata associated with the training trajectory. We instead implement temperature as an action in the world-model sense: from a single inferred molecular state, different temperature actions produce different predicted future distributions, enabling matched-initialization intervention experiments that test whether the model captures thermal response rather than merely interpolating temperature-labeled training data.

3 PROBLEM FORMULATION

We view protein dynamics as a partially observed, controlled stochastic process. Let c denote static molecular context (amino-acid identity, chain connectivity) and Z_t denote the microscopic molecular state at physical time t, containing unobserved variables including momenta and solvent con-

figurations. We observe only a structural projection $X_t = \mathcal{O}(Z_t)$, represented by N, C_α, and C backbone atoms, and the current secondary-structure annotation D_t.

History and predictive state. Given observation interval δ, the ordered structural history is

$$h_t = \left(X_{t-(K-1)\delta}, \ldots, X_t\right). \tag{1}$$

Because the same observed conformation can arise from different microscopic states, X_t alone need not determine future motion. The history h_t provides information about latent dynamical state unavailable from a single structure.

Actions and stochastic futures. An action $a \in \mathcal{A}$ specifies the thermostat temperature applied after time t and held fixed over the prediction window. For sampling interval Δ and horizon H, the future trajectory is

$$Y_t = (X_{t+\Delta}, \ldots, X_{t+H\Delta}). \tag{2}$$

The target of protein world modeling is the controlled conditional path distribution

$$P^{\star}_{H,\Delta}\left(Y_t \mid h_t, c, D_t, a\right), \tag{3}$$

where $P^{\star}$ denotes the reference molecular process under specified thermal control. A learned world model P_θ approximates this distribution and can be sampled repeatedly to obtain distinct continuations from the same observed history. This formulation separates two roles: observed history and current state provide information about latent dynamics, whereas the action determines how that state evolves into the future. The partial observability decomposition is provided in Appendix D.

4 METHOD

We parameterize the controlled stochastic dynamics introduced in Section 3 with three coupled components. First, an action-independent history encoder summarizes an ordered sequence of observed conformations into a predictive molecular state. Second, an action-conditioned flow generates a joint distribution over future residue-frame trajectories. Third, a differentiable geometry operator, DIFFPROJECT, enforces peptide-chain compatibility while remaining close to the trajectory proposed by the generative model. Together, these components define a stochastic transition operator from an observed molecular history and a temperature action to a distribution over geometrically valid future trajectories. Figure 1 summarizes the complete paradigm.

4.1 HISTORY-CONDITIONED PREDICTIVE STATE

We represent residue i at physical time u by a rigid frame $F_{u,i} = (p_{u,i}, R_{u,i})$, where $p_{u,i} \in \mathbb{R}^3$ is its C_α position and $R_{u,i} \in SO(3)$ is an orthonormal frame constructed from the N–C_α–C backbone geometry. Rather than encoding historical structures in a global coordinate system, we express their motion relative to the corresponding residue frame at the forecast origin:

$$d_{u,i} = \ell^{-1} R_{t,i}^{\mathsf{T}}(p_{u,i} - p_{t,i}), \qquad Q_{u,i} = R_{t,i}^{\mathsf{T}} R_{u,i}. \tag{4}$$

Here ℓ is a fixed translation scale. The resulting representation is invariant to a common rigid transformation of the observed trajectory while retaining both translational and rotational motion relative to the current protein.

For each residue, we embed its ordered sequence of relative frames with temporal position encodings and apply temporal attention to obtain history representation. In parallel, the current conformation is encoded using pairwise geometric features (inter-residue distances, relative orientations, sequence separation) that provide spatial context and attention biases. Combining ordered temporal memory with current three-dimensional structure and observed secondary-structure state yields

$$s_t = E_\theta(h_t, c, D_t), \tag{5}$$

interpreted as a predictive representation of the partially observed molecular state. Importantly, E_θ does not depend on the requested temperature action, enabling s_t to be reused across multiple thermal branches.

4.2 Action-Conditioned Stochastic Trajectory Flow

Given predictive state s_t and temperature action a, we generate the complete future trajectory jointly. Let $U = (d, Q) \in (\mathbb{R}^3 \times SO(3))^{H \times N}$ collect relative residue frames over all N residues and H future lead times, with physical times $\mathbf{t} = (\Delta, 2\Delta, \ldots, H\Delta)$. We learn an action-conditioned vector field

$$(v_\theta, \omega_\theta) = f_\theta(U_\tau, \tau; s_t, a, \mathbf{t}), \tag{6}$$

where $\tau \in [0, 1]$ denotes generative flow time (distinct from physical molecular time $\mathbf{t}$). The trajectory backbone alternates spatial attention across residues with temporal attention across future lead times (Figure 1c). Temperature is embedded as a continuous action and injected through adaptive normalization, allowing its effect to depend on inferred molecular state, residue context, and physical lead time.

Flow matching on residue-frame trajectories. We train the trajectory field using conditional flow matching (Lipman et al., 2023). For reference trajectory $U_1 = (d_1, Q_1)$ and prior $U_0 \sim P_0$, translations follow linear interpolation while rotations follow geodesics on $SO(3)$ (Appendix E). The flow-matching objective is

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}\left[\|v_\theta - (d_1 - d_0)\|^2_{\mathrm{av}} + \lambda_R \|\omega_\theta - \Omega\|^2_{\mathrm{av}}\right], \tag{7}$$

where $\|\cdot\|^2_{\mathrm{av}}$ averages over valid residues and frames. At inference, prior trajectories are transported through the learned vector field from $\tau = 0$ to $\tau = 1$:

$$U_1 = \Phi_\theta(U_0; s_t, a, \mathbf{t}), \qquad U_0 \sim P_0. \tag{8}$$

Independent samples from P_0 produce alternative future trajectories from the same history and temperature action.

4.3 Geometry-Preserving Trajectory Generation

Residue-wise rigid frames preserve internal N–C_α–C geometry but independently displaced neighboring frames may not form valid peptide chains. We address this with DiffProject, a differentiable operator that couples neighboring frames while remaining anchored to the sampled trajectory. For each generated frame, DiffProject introduces translational offsets $u_{j,i} \in \mathbb{R}^3$ and rotational offsets $\eta_{j,i} \in \mathbb{R}^3$, optimized via differentiable gradient descent on a geometric energy balancing inter-residue C–N bond distances, peptide angles, and anchoring to the original trajectory (details in Appendix F). Starting from zero offsets, we unroll J optimization steps to obtain corrected frames $\widetilde{F} = \Pi_J(F)$.

Let $\mathcal{D}_{X_t}$ denote coordinate reconstruction from corrected frames using the current protein as spatial anchor. Combining history encoding, action-conditioned transport, geometric correction, and decoding yields the conditional trajectory distribution

$$P_\theta(\cdot \mid h_t, c, D_t, a) = \Big(\mathcal{D}_{X_t} \circ \Pi_J \circ \Phi_\theta(\cdot; E_\theta(h_t, c, D_t), a, \mathbf{t})\Big)_{\#} P_0, \tag{9}$$

where $\#$ denotes pushforward. This summarizes the complete world model: history determines predictive state, temperature controls stochastic trajectory dynamics, and DiffProject maps sampled residue-frame paths to geometrically compatible molecular trajectories.

4.4 Training Objective

We train with flow-matching objective Equation 7 augmented with physical and structural supervision. $\mathcal{L}_{\mathrm{aux}}$ combines training-only objectives for local forces, future secondary structure, structural events, and contact properties. For the projection-aware model, $\mathcal{L}_{\mathrm{geom}}$ supervises backbone coordinates and peptide geometry after differentiable DiffProject refinement. The complete objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{FM}} + \mathcal{L}_{\mathrm{aux}} + \mathcal{L}_{\mathrm{geom}}. \tag{10}$$

Component definitions, architectural specifications, and training hyperparameters are provided in Appendix C.

Table 1: Trajectory quality on the common mdCATH evaluation setting.

Method	RMWD ↓	PCA W_2 ↓	RMSF r ↑	RMSF err. ↓	Weak J ↑	Transient J ↑	nRMSD err. ↓
MDGen Jing et al. (2024b)	3.50	2.60	0.77	0.65	0.55	0.30	2.00
ConfRover-forward Shen et al. (2025)	3.30	2.35	0.81	0.62	0.58	0.34	1.80
TEMPO Xu et al. (2025)	2.46	1.67	0.84	0.55	0.57	0.36	1.69
Ours	**1.35**	**0.82**	**0.86**	**0.33**	**0.73**	**0.46**	**1.15**
MD replica	1.90	1.65	0.95	0.25	0.84	0.55	0.85
Persistence	4.60	2.80	–	1.65	0.00	0.00	2.80

4.5 Experimental Setup

Data and prediction protocol. Unless stated otherwise, experiments use the five mdCATH temperatures 320, 348, 379, 413, and 450 K. Protein domains are split by CATH homologous superfamily, with all temperatures and replicas of a domain assigned to the same partition. The model observes eight structural frames at 1 ns cadence, spanning 7 ns of history, and generates $M = 16$ stochastic continuations per query. The default prediction horizon is 16 ns. When methods operate at different temporal resolutions, we evaluate them at matched physical timestamps. External baselines are reproduced and evaluated under the same comparison setting.

Metrics. Our primary probabilistic metric is Path Energy Score (Path ES), which evaluates the joint distribution of ordered trajectory features. We additionally report root mean Wasserstein distance (RMWD), PCA W_2, RMSF correlation and amplitude error, weak- and transient-contact recovery, and native-reference RMSD error. For closed-loop simulation, we measure path survival, lag-RMSD error, autocorrelation error, PCA W_2, and contact-occupancy error. Backbone validity is evaluated separately through C–N bond outliers, backbone-angle outliers, nonlocal clashes, and whole-path validity. Details are provided in Appendix B.

4.6 Comparison with Learned Protein Trajectory Models

We first compare trajectory quality against learned protein dynamics models under a common mdCATH evaluation setting. MD replica denotes an independent MD simulation of the same molecular system and condition and serves as a reference rather than a learned baseline.

As shown in Table 1, our model improves consistently over the reproduced learned baselines. Compared with ConfRover-forward, RMWD decreases from 3.30 to 1.35, RMSF correlation increases from 0.81 to 0.86, and native-RMSD error decreases from 1.80 to 1.15. The same trend holds across the remaining distributional and contact-based metrics, while a gap to the independent MD replica remains.

4.7 Stochastic Futures and Dynamical Memory

We next ask whether the model learns a stochastic conditional dynamics rather than a deterministic predictor with increased output variance, and whether its predictions depend on the temporal organization of the observed history.

Stochastic prediction. Table 2a compares the full stochastic model with two controls. Deterministic same-input uses the same conditioning information but removes stochastic future sampling, while the scaled-noise variant adds post-hoc noise to match the overall motion scale. The stochastic model achieves a Path ES of 0.260, compared with 0.320 for deterministic prediction. Simply adding output variance does not reproduce the gain: the scaled-noise control reaches Path ES 0.350 despite matching the reference spread. The improvement therefore reflects a better predictive distribution rather than variance inflation alone.

Figure 2 visualizes this behavior from a shared molecular state. In the reference dynamical space, persistence remains at the observed structure, whereas our model generates distinct continuations that explore the conformational region supported by the reference trajectory.

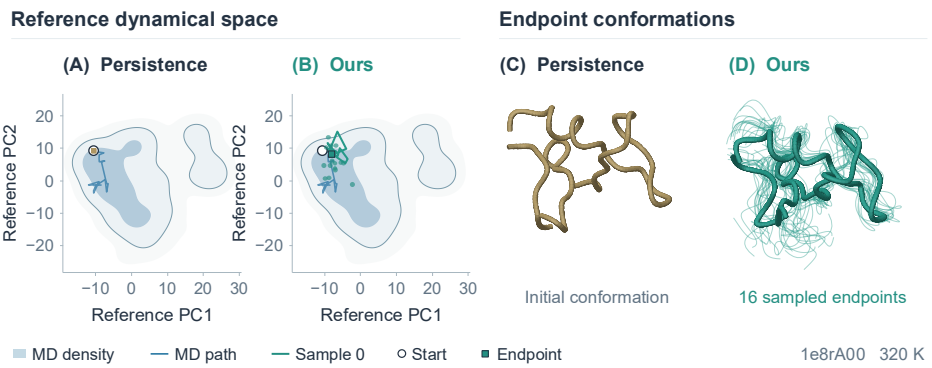
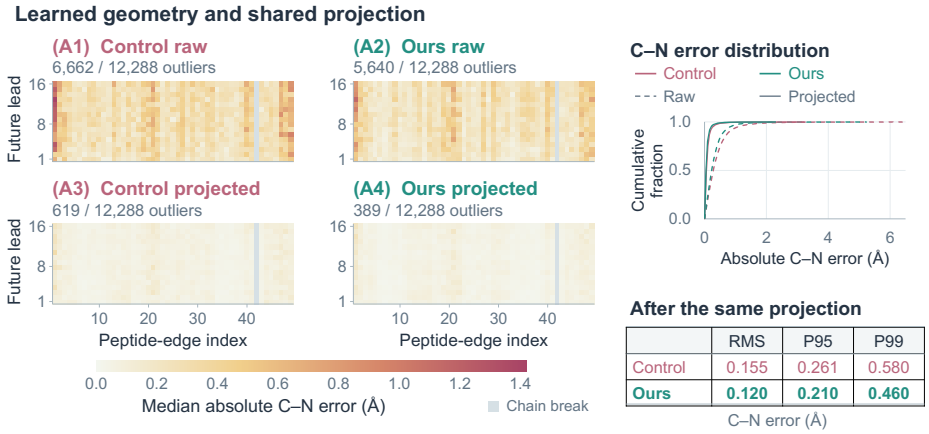


Figure 2: Stochastic futures from a shared molecular state. **(A–B)** Persistence and our model in the PCA space defined by the mdCATH reference trajectory for domain 1e8rA00 at 320 K. The shaded region denotes the reference MD density, the blue curve the observed MD continuation, and the open circle the shared starting state. For our model, all 16 sampled endpoints are shown and the sampled trajectory is highlighted. **(C–D)** Endpoint conformations at the 16-step prediction horizon. Persistence retains the initial conformation, whereas our model generates an ensemble of 16 future conformations.



	RMS	P95	P99
Control	0.155	0.261	0.580
Ours	**0.120**	**0.210**	**0.460**

Figure 3: Learned backbone-geometry improvements before and after projection. **(A)** C–N errors for the matched control and our model before and after applying the same projection. Heatmaps show the median absolute C–N error across 16 generated samples for each peptide edge and future lead time; gray columns denote excluded chain breaks. The ECDF summarizes the full error distribution, while the accompanying table reports projected C–N root-mean-square error and the 95th and 99th percentiles of absolute error after projection, in Å.

Ordered history. The history ablations separate the value of observing additional past conformations from the value of their temporal organization. Current uses only the present conformation, Shuffled destroys temporal order, and Set receives the same history frames as Ordered but treats them as an unordered set. Path ES improves from 0.305 for Current to 0.287 for Set history and 0.260 for Ordered history. Because Set and Ordered contain the same historical conformations, their gap isolates predictive information carried specifically by temporal order.

Geometry without motion collapse. Table 2b evaluates backbone validity together with retained molecular motion. DIFFPROJECT reduces C–N outliers from 1.10% before projection to 0.8% while retaining Motion/MD $= 0.94$. By contrast, the excessive-geometry control attains lower local defect rates but suppresses motion to 0.68, showing that geometric regularity cannot be optimized independently of dynamical fidelity.

Table 2: Ablations of stochastic prediction, temporal history, and backbone geometry.

(a) Stochasticity and history dependence

Variant	Path ES ↓	RMWD ↓	RMSF r ↑	PCA W_2 ↓	Path valid (%) ↑	Spread/MD ≈ 1
Persistence	0.500	2.50	NA	1.60	99.5	0.00
Deterministic same-input	0.320	1.70	0.75	1.10	68.0	0.00
Deterministic + scaled noise	0.350	1.65	0.76	1.15	50.0	**1.00**
Current stochastic	0.305	1.65	0.80	1.00	60.0	0.84
Shuffled history	0.296	1.59	0.81	0.97	62.0	0.87
Set history	0.287	1.53	0.82	0.94	64.0	0.90
Ordered history (Ours)	**0.260**	**1.35**	**0.86**	**0.82**	**72.0**	0.95

(b) Geometry and motion preservation

Variant	C–N out. (%) ↓	Angle out. (%) ↓	Clash/frame ↓	Motion/MD ≈ 1
Free-residue generator	1.80	3.00	0.80	1.05
Chain representation, no soft geometry	1.40	2.20	0.60	**0.99**
Full model, pre-projection	1.10	1.05	0.60	0.96
Ours (with DIFFPROJECT)	0.8	1.02	0.35	0.94
Excessive geometry penalty	**0.55**	**1.00**	**0.30**	0.68

Table 3: Temperature conditioning and leave-temperature-out generalization on mdCATH.

Panel	Model / evaluation	RMWD ↓	RMSF r ↑	RMSF err. ↓	PCA W_2 ↓	nRMSD err. ↓
Five-T mean	Full temperature input	**1.35**	**0.86**	**0.42**	**0.82**	**0.62**
	Temperature-blind	1.70	0.80	0.62	1.10	0.85
379 K	Full @379 K	**1.25**	**0.87**	**0.38**	**0.75**	**0.55**
	Leave-379-K-out @379 K	1.35	0.85	0.45	0.85	0.62
450 K	Full @450 K	**1.92**	**0.80**	**0.66**	**1.20**	**0.95**
	Leave-450-K-out @450 K	2.15	0.76	0.82	1.40	1.10

Figure 3 separates improvements learned by the generator from those introduced by geometric post-processing. Our model reduces C–N error before projection and retains the advantage when both models receive the same projection. The improvement also persists across validation domains for C–N outliers, backbone-angle error, and nonlocal clashes.

4.8 TEMPERATURE AS AN ACTION

Temperature conditioning alone does not establish the action semantics of our formulation. We therefore evaluate whether changing the requested future temperature from a shared observed molecular state produces the corresponding change in future dynamics.

Temperature-response evaluation. For each observed history, we branch the model under alternative temperature actions and compare the resulting changes in trajectory observables with the corresponding temperature-dependent changes observed in reference MD. For the largest 450–320 K contrast, the model reduces the response error for ΔRMSF from 1.20 Å for the temperature-blind control to 0.40 Å, and the error in Δcontact occupancy from 0.200 to 0.060. The same ordering holds for the intermediate temperature contrasts. These results show that changing the requested action from a shared molecular state produces a corresponding change in the predicted future distribution.

Temperature generalization. Table 3 evaluates standard multi-temperature prediction together with leave-temperature-out generalization. Across the five-temperature evaluation, explicit temperature input improves RMWD from 1.70 to 1.35. The model also retains most of its performance when 379 K is excluded from training, while the larger degradation at the held-out 450 K reflects the more difficult extrapolative setting.

Table 4: Zero-shot transfer from mdCATH to ATLAS across increasingly challenging subsets.

Subset	Model	RMWD ↓	PCA W_2 ↓	RMSF r ↑	Weak J ↑	Transient J ↑	Valid (%) ↑
Clean	Ours	1.60	0.70	0.90	0.72	0.42	82.0
Topology-novel	Ours	1.80	0.85	0.87	0.67	0.37	80.0
High-motion	Ours	2.05	0.93	0.85	0.66	0.40	78.0

4.9 Closed-Loop Molecular Simulation

Jointly generating a future block does not by itself establish that the model can continue from its own predictions. We therefore evaluate closed-loop simulation by recursively replacing observed context with generated states, without access to reference future frames.

Composition at a matched horizon. We first compare direct 16 ns prediction with four successive 4 ns transitions at the same physical horizon. Path survival changes only from 90% to 88%, while autocorrelation error increases from 0.060 to 0.070. Thus, short learned transitions can be composed with limited degradation at the matched horizon.

Long-horizon generated-state rollout. We next recursively continue generated histories to 32 and 64 ns. Path survival decreases from 90% at 16 ns to 65% at 64 ns, while autocorrelation error increases from 0.060 to 0.135. At 64 ns, the no-context-noise control retains only 28% valid paths and has autocorrelation error 0.250, compared with 65% and 0.135 for the full model. The model therefore remains substantially more stable when temporal context is preserved, although closed-loop errors continue to accumulate with horizon.

4.10 Zero-Shot Transfer Beyond mdCATH

Finally, we test whether the learned dynamics transfer beyond the mdCATH training distribution. The model is trained exclusively on mdCATH and evaluated on ATLAS at its native 300 K condition without target-domain fine-tuning. This evaluation jointly shifts the protein distribution and places the temperature below the mdCATH training range.

Table 4 shows that the model retains strong zero-shot trajectory statistics on ATLAS. On the clean subset, it achieves RMWD 1.60 and RMSF correlation 0.90. Performance degrades on topology-novel and high-motion proteins, with RMWD increasing to 1.80 and 2.05, respectively, while RMSF correlation remains 0.87 and 0.85. This degradation is consistent with a harder distribution shift rather than a complete loss of learned dynamical structure.

5 Conclusion

We introduced a protein world model for stochastic, history-conditioned, and temperature-controlled molecular simulation. The model infers a predictive state from ordered conformational history, uses temperature as an action to generate branching future trajectories, and preserves backbone geometry through differentiable refinement. Experiments show that the model captures stochastic future distributions, benefits from temporal order beyond unordered history, responds consistently to temperature interventions from a shared molecular state, and remains usable under closed-loop rollout.

Experiments show that the model captures stochastic future distributions, benefits from temporal order beyond unordered history, responds consistently to temperature interventions from a shared molecular state, and remains usable under closed-loop rollout. These results establish protein world modeling as a promising framework for action-conditioned in-silico experiments, with future work extending to broader molecular systems, richer atomic representations, and longer simulation horizons.

AI USE STATEMENT

We used generative AI tools to assist with and polish the writing, assist in formulating mathematical arguments and writing proofs in Appendix A, and provide feedback on the experimental evaluation design. We reviewed and verified all AI-assisted work. The authors take full responsibility for the final text, mathematical claims and proofs, experimental methodology, and reported results.

REPRODUCIBILITY STATEMENT

To support reproducibility, we provide the assumptions, mathematical derivations, and proofs underlying our theoretical results in Appendix A. We describe the experimental setup, evaluation protocol, and implementation details in the main text and appendices to facilitate reproduction of the reported results.

REFERENCES

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.

David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3):440, 2018.

Brooke E Husic and Vijay S Pande. Markov state models: From an art to a science. *Journal of the American Chemical Society*, 140(7):2386–2396, 2018.

Bowen Jing, Bonnie Berger, and Tommi Jaakkola. AlphaFold meets flow matching for generating protein ensembles. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 22277–22303, 2024a. URL https://arxiv.org/abs/2402.04845.

Bowen Jing, Hannes Stärk, Tommi Jaakkola, and Bonnie Berger. Generative modeling of molecular dynamics trajectories. In *Advances in Neural Information Processing Systems*, volume 37, 2024b. URL https://arxiv.org/abs/2409.17808.

Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew Y. K. Foong, Victor García Satorras, Osama Abdin, Bastiaan S. Veeling, Iryna Zaporozhets, Yaoyi Chen, Soojung Yang, Adam E. Foster, Arne Schneuing, Jigyasa Nigam, Federico Barbero, Vincent Stimper, Andrew Campbell, Jason Yim, Marten Lienen, Yu Shi, Shuxin Zheng, Hannes Schulz, Usman Munir, Roberto Sordillo, Ryota Tomioka, Cecilia Clementi, and Frank Noé. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science*, 389(6761):eadv9817, 2025. doi: 10.1126/science.adv9817. URL https://doi.org/10.1126/science.adv9817.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2210.02747.

Yikai Liu, Haoyang Zheng, Lining Mao, Yanbin Wang, Ming Chen, and Guang Lin. ProTDyn: A foundation protein language model for thermodynamics and dynamics generation. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=fvCHkWbdgX.

Jiarui Lu, Bozitao Zhong, Zuobai Zhang, and Jian Tang. Str2str: A score-based framework for zero-shot protein conformation sampling. In *International Conference on Learning Representations*, volume 2024, pp. 10663–10694, 2024.

Antonio Mirarchi, Toni Giorgino, and Gianni De Fabritiis. mdCATH: A large-scale MD dataset for data-driven computational biophysics. *Scientific Data*, 11:1299, 2024. doi: 10.1038/s41597-024-04140-z. URL https://doi.org/10.1038/s41597-024-04140-z.

Hew Phipps, Matteo Cagiada, Santiago D. Villalba, and Charlotte M. Deane. Spectral diffusion for protein dynamics. In *ICML Workshop on Generative and Agentic AI for Biology*, 2026. URL https://arxiv.org/abs/2607.04134. arXiv:2607.04134.

Yuning Shen, Lihao Wang, Huizhuo Yuan, Yan Wang, Bangji Yang, and Quanquan Gu. Simultaneous modeling of protein conformation and dynamics via autoregression. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://arxiv.org/abs/2505.17478.

Nima Shoghi, Yuxuan Liu, Yuning Shen, Rob Brekelmans, Pan Li, and Quanquan Gu. Scalable spatio-temporal SE(3) diffusion for long-horizon protein dynamics. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2602.02128.

Yaoyao Xu, Di Wang, Zihan Zhou, Tianshu Yu, and Mingchen Chen. TEMPO: Temporal multiscale autoregressive generation of protein conformational ensembles. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://arxiv.org/abs/2511.05510.

A Mathematical Justification of the Path Energy Score

This appendix specializes the classical energy-score construction (Gneiting & Raftery, 2007) to conditional protein trajectories. The scoring rule itself is not a new contribution. We establish exactly which distribution our implementation identifies, why a single observed future can be scored, and how the finite-ensemble correction preserves the population objective. All scores are oriented so that smaller is better.

A.1 Conditional Path Features and Assumptions

Let $q = (h_t, c, D_t, a)$ be an evaluation query and let $Y \sim P_q^\star$ denote its reference future. Fix physical readout times $\mathcal{T} = (\tau_1, \dots, \tau_L)$ before evaluation and define the measurable map

$$\Phi_q(Y) = \frac{1}{\sqrt{L}} \operatorname{vec}\left[W\phi_q(X_{t+\tau_1}), \dots, W\phi_q(X_{t+\tau_L})\right] \in \mathbb{R}^d, \qquad d = Lr. \tag{11}$$

Here $\phi_q \in \mathbb{R}^r$ contains structural observables, including contacts defined from the observed X_t, and W is a fixed positive diagonal scaling matrix. Our primary score uses $r = 6$, $L = 4$, and $\mathcal{T} = (4, 8, 12, 16)$ ns, with the scaling already incorporated into the features in Appendix B. The factor $L^{-1/2}$ normalizes the concatenation; it does not make different horizons or feature sets interchangeable. Write

$$F_q = (\Phi_q)_\# P_{\theta,q}, \qquad G_q = (\Phi_q)_\# P_q^\star, \tag{12}$$

where $\#$ denotes the induced distribution of features. Thus F_q is the forecast feature-path law and G_q is its reference law, both on the same Euclidean space.

Assumptions. For each evaluated q, both laws have finite first moments. For population averages, the moments are also integrable over the query distribution μ. The model, feature map, scaling, timestamps, and evaluation weights are fixed independently of the evaluated future outcomes. Generated paths are independent draws conditional on q, independent of the reference future given q, and $M \geq 2$. These are assumptions about sampling whole paths: frames within a path may have arbitrary temporal dependence. Statements about Monte Carlo variance below additionally assume finite second moments. We condition throughout on the fitted model and on any preprocessing learned from development data.

Define the score and its expected risk by

$$S(F, z) = \mathbb{E}_{U\sim F}\|U - z\|_2 - \tfrac{1}{2}\mathbb{E}_{U,U'\overset{\text{iid}}{\sim}F}\|U - U'\|_2, \tag{13}$$

$$R(F, G) = \mathbb{E}_{V\sim G} S(F, V). \tag{14}$$

The first-moment assumptions make all terms finite by the triangle inequality. The first term measures forecast–observation discrepancy; the second corrects for the spread that a stochastic forecast should retain. Their combination, rather than either term alone, determines the optimal predictive distribution.

A.2 Strict Propriety in Joint Path-Feature Space

For independent $U, U' \sim F$ and $V, V' \sim G$, define the energy distance

$$\mathcal{D}_{\mathrm{E}}(F, G) = 2\mathbb{E}\|U - V\|_2 - \mathbb{E}\|U - U'\|_2 - \mathbb{E}\|V - V'\|_2. \tag{15}$$

We first give a self-contained justification of its nonnegativity and identification property.

Lemma A.1 (Fourier representation of Euclidean distance)**.** *For every $x \in \mathbb{R}^d$,*

$$\|x\|_2 = \kappa_d^{-1} \int_{\mathbb{R}^d} \frac{1 - \cos(u^\mathsf{T} x)}{\|u\|_2^{d+1}}\, \mathrm{d}u, \qquad \kappa_d = \int_{\mathbb{R}^d} \frac{1 - \cos(u_1)}{\|u\|_2^{d+1}}\, \mathrm{d}u \in (0, \infty). \tag{16}$$

Proof. Near the origin, $1 - \cos(u_1) \leq \|u\|_2^2/2$, so the radial integral is bounded by a constant times $\int_0^1 \mathrm{d}r$. Outside the unit ball, the numerator is at most two and the radial integral is bounded by a constant times $\int_1^\infty r^{-2}\mathrm{d}r$. Hence κ_d is finite; it is positive since the integrand is positive on a set of positive measure. For $x \neq 0$, rotate coordinates so that $x = \|x\|_2 e_1$ and substitute $v = \|x\|_2 u$. Rotation invariance and the Jacobian yield the integral $\|x\|_2\kappa_d$. For $x = 0$, both sides vanish. □

Theorem A.2 (Strict propriety of the path energy score)**.** *For probability measures F, G on $\mathbb{R}^d$ with finite first moments,*

$$R(F, G) - R(G, G) = \tfrac{1}{2}\mathcal{D}_{\mathrm{E}}(F, G) \geq 0, \tag{17}$$

with equality if and only if $F = G$. In particular, the unique population-optimal forecast for query q is $F_q = G_q$.

Proof. Expanding the two risks immediately gives equation 17. Let $\chi_F(u) = \mathbb{E}e^{iu^{\mathsf{T}}U}$ and χ_G be the characteristic functions. Applying Lemma A.1 to the three distance terms gives

$$\mathcal{D}_{\mathrm{E}}(F, G) = \kappa_d^{-1} \int_{\mathbb{R}^d} \frac{|\chi_F(u) - \chi_G(u)|^2}{\|u\|_2^{d+1}} \, \mathrm{d}u. \tag{18}$$

To justify the calculation, first integrate over annuli $\epsilon < \|u\|_2 < R$, where the integrands are bounded and Fubini's theorem applies. Independence gives

$$2\big(1 - \mathrm{Re}(\chi_F \overline{\chi_G})\big) - (1 - |\chi_F|^2) - (1 - |\chi_G|^2) = |\chi_F - \chi_G|^2.$$

Each original distance integral is finite by the moment assumptions. Monotone convergence on expanding annuli establishes equation 18. Its integrand is nonnegative. If the integral vanishes, $\chi_F = \chi_G$ almost everywhere away from zero. Continuity of characteristic functions then gives equality everywhere, including zero, and uniqueness of characteristic functions implies $F = G$. Conversely, $F = G$ makes the integral zero. □

Corollary A.3 (Conditional evaluation and its support)**.** *Let $\mathcal{R}(F) = \mathbb{E}_{q\sim\mu} R(F_q, G_q)$. Under the integrated moment assumption,*

$$\mathcal{R}(F) - \mathcal{R}(G) = \tfrac{1}{2}\mathbb{E}_{q\sim\mu}\mathcal{D}_{\mathrm{E}}(F_q, G_q) \geq 0, \tag{19}$$

and equality holds exactly when $F_q = G_q$ for μ-almost every query.

Proof. Integrate equation 17. A nonnegative integrable random variable has zero expectation if and only if it is zero almost surely. Theorem A.2 then gives the stated equality condition. □

This result permits fixed domain- and temperature-balanced evaluation: the weights specify the target query population. It does not guarantee accuracy at temperatures or histories absent from that population. In particular, scores on constant-temperature histories cannot establish the correctness of temperature-switch interventions.

Why the norm must not be squared. With finite second moments, replacing both distances in equation 13 by squared Euclidean distances yields

$$\mathbb{E}\|U - z\|_2^2 - \tfrac{1}{2}\mathbb{E}\|U - U'\|_2^2 = \|\mathbb{E}U - z\|_2^2. \tag{20}$$

The forecast variance cancels, so any forecast with the correct mean attains the same risk. The unsquared norm in our implementation is therefore essential to distributional identification. For that score, a deterministic forecast can be optimal only when the reference feature-path law is itself degenerate. Even a perfect stochastic forecast has positive expected score when the reference law is nondegenerate:

$$R(G, G) = \tfrac{1}{2}\mathbb{E}_{V,V'\sim G}\|V - V'\|_2. \tag{21}$$

Accordingly, a reference score must not be set to zero by comparing an MD path with itself.

A.3 Finite Ensembles and a Single Observed Future

Proposition A.4 (Unbiased finite-ensemble estimator)**.** *Fix a query and an observed feature path z. For independent $U_1, \ldots, U_M \sim F$, the estimator*

$$\widehat{S}_M(F, z) = \frac{1}{M}\sum_{m=1}^{M} \|U_m - z\|_2 - \frac{1}{2M(M-1)}\sum_{m\neq n} \|U_m - U_n\|_2 \tag{22}$$

satisfies $\mathbb{E}[\widehat{S}_M(F, z) \mid z] = S(F, z)$. If independently $Z \sim G$, then $\mathbb{E}\widehat{S}_M(F, Z) = R(F, G)$.

Proof. Each of the M first-sum terms has expectation $\mathbb{E}\|U - z\|_2$, and each of the $M(M-1)$ ordered off-diagonal terms has expectation $\mathbb{E}\|U - U'\|_2$. Linearity gives the first claim; averaging over Z gives the second. No second reference continuation is needed because the score only contains model–reference and model–model distances. The reference–reference term in equation 17 is a model-independent quantity used in the proof, not a term required for scoring. □

For fixed query weights $w_i \geq 0$, $\sum_i w_i = 1$, and one $Z_i \sim G_{q_i}$ per query,

$$\mathbb{E}\left[\sum_i w_i \widehat{S}_M(F_{q_i}, Z_i) \,\middle|\, q_1, \ldots, q_n\right] = \sum_i w_i R(F_{q_i}, G_{q_i}), \tag{23}$$

provided each reference path has the specified conditional law under the sampling protocol. If independently sampled queries follow μ, the unweighted average is unbiased for $\mathcal{R}(F)$ and converges to it under the integrability assumption. Correlations between windows do not change marginal unbiasedness, but invalidate an independent-frame standard error. Domain-level bootstrap intervals retain within-domain dependence; if superfamilies induce additional dependence, resampling must instead use that grouping. Having a score with an unbiased expectation does not recover the complete conditional distribution from a single observed future or remove finite-test-set uncertainty.

Why the off-diagonal correction matters. The energy score of the empirical forecast distribution $\widehat{F}_M = M^{-1}\sum_m \delta_{U_m}$ would use the denominator $2M^2$. Since diagonal distances are zero,

$$\mathbb{E}S(\widehat{F}_M, z) = S(F, z) + \frac{1}{2M}\mathbb{E}\|U - U'\|_2. \tag{24}$$

This is the exact score of the empirical distribution, but a biased estimator of the underlying sampler's score. The extra term penalizes spread, with a strength depending on M. Equation 22 removes this term and preserves the same expected ranking objective at every $M \geq 2$. Using the same M across methods still helps match Monte Carlo precision. Conditioning or selecting samples using the observed future, or filtering out geometric failures, would instead evaluate a different forecast law.

Proposition A.5 (Conditional Monte Carlo variance)**.** *For fixed z, let*

$$k_z(u, v) = \tfrac{1}{2}\big(\|u - z\|_2 + \|v - z\|_2 - \|u - v\|_2\big), \tag{25}$$

$$\eta_z = \mathbb{E}k_z(U, U'), \qquad g_z(u) = \mathbb{E}_{U'}k_z(u, U') - \eta_z, \tag{26}$$

$$r_z(u, v) = k_z(u, v) - \eta_z - g_z(u) - g_z(v). \tag{27}$$

If F has finite second moment, then

$$\mathrm{Var}(\widehat{S}_M \mid z) = \frac{4}{M}\,\mathrm{Var}(g_z(U)) + \frac{2}{M(M-1)}\,\mathrm{Var}(r_z(U, U')) = O(M^{-1}). \tag{28}$$

Proof. The triangle inequality gives $0 \leq k_z(u, v) \leq \min(\|u - z\|_2, \|v - z\|_2)$, so its second moment is finite. Rewriting equation 22 as an order-two U-statistic yields

$$\widehat{S}_M - \eta_z = \frac{2}{M}\sum_m g_z(U_m) + \binom{M}{2}^{-1}\sum_{m<n} r_z(U_m, U_n). \tag{29}$$

Both components have mean zero. Moreover, $\mathbb{E}[r_z(U, U') \mid U] = 0$, so the two sums are uncorrelated, and residual terms for distinct unordered pairs are uncorrelated even when they share one index. Taking variances gives equation 28. □

This variance concerns generated-sample noise at a fixed observed future. Increasing M reduces that noise, but not uncertainty from limited MD observations or a small number of test proteins. Repeated physical replicas and additional held-out domains address different sources of uncertainty.

A.4 Temporal Dependence and the Value of History

Concatenation at ordered timestamps retains joint temporal information that a collection of frame-wise scores discards. For an unscaled two-coordinate example, let

$$G = \tfrac{1}{2}\delta_{(1,1)} + \tfrac{1}{2}\delta_{(-1,-1)}, \qquad F = \tfrac{1}{2}\delta_{(1,-1)} + \tfrac{1}{2}\delta_{(-1,1)}. \tag{30}$$

Both time marginals agree, so scoring each time marginal separately gives identical optimal marginal risks. Their temporal dependence is opposite. Every cross-distribution distance is two, whereas $\mathbb{E}\|U - U'\|_2 = \mathbb{E}\|V - V'\|_2 = \sqrt{2}$. Hence

$$\mathcal{D}_{\mathrm{E}}(F, G) = 4 - 2\sqrt{2} > 0, \qquad R(F, G) - R(G, G) = 2 - \sqrt{2} > 0. \tag{31}$$

Applying a common positive normalization rescales the gap and preserves its sign. This example shows that joint Path ES can penalize incorrect temporal dependence even when every observed-time marginal is correct. It does not imply that every time permutation changes every process: a temporally exchangeable feature law is unchanged by permutation. Finite-sample power also depends on the features and effect size.

Corollary A.6 (Predictive value of ordered history)**.** *Let $C = (h_t, c, D_t, a)$ be full conditioning information and $B = b(C)$ a coarsening that retains the current structure, molecular context, and action. Use the same B-measurable feature map for both forecasts, and let $G_C = \mathcal{L}(Z \mid C)$ and $G_B = \mathcal{L}(Z \mid B)$. Under the integrated moment assumptions,*

$$\mathbb{E}S(G_B, Z) - \mathbb{E}S(G_C, Z) = \tfrac{1}{2}\mathbb{E}_C\mathcal{D}_{\mathrm{E}}(G_B, G_C) \geq 0. \tag{32}$$

Equality holds if and only if $G_C = G_B$ almost surely.

Proof. Condition on C. The law G_B is then a fixed forecast and the reference law is G_C. Apply Theorem A.2 and average over C. □

Taking $B = (X_t, c, D_t, a)$ characterizes the population benefit of history beyond the current structure. Taking B to retain the unordered past-frame multiset characterizes information in temporal ordering beyond observing additional conformations. The shuffled-history experiment approximates this information restriction, while retaining the current anchor. The identity concerns optimal conditional laws, not a guarantee that a fitted finite-capacity network benefits from every extra frame. It therefore motivates the controlled comparisons without predetermining their empirical outcome. No Markov assumption about the observed structures is required.

A.5 Scope of the Guarantee

The equality condition is $(\Phi_q)_{\#}P_{\theta,q} = (\Phi_q)_{\#}P_q^{\star}$. Since our six structural observables are not an injective representation of atomic configurations, different molecular path laws can have the same feature-path law and identical expected scores. The score is proper for full path laws but generally not strictly proper on that larger space. This is why we separately assess contact events, geometric defects, temporal increments, and long-horizon observables. Strict propriety also does not guarantee high power in a finite evaluation or identify unobserved fast dynamics between readout times.

A fixed nonsingular W and the positive factor $L^{-1/2}$ preserve identification of the selected feature law; they change the geometry of finite-sample errors, so both are frozen before evaluation. The 1 ns and 4 ns settings are comparable because both are scored on the same physical timestamps and observables. Finally, projection is part of the sampler: scoring corrected paths assesses the induced corrected law, without conferring a physical-validity guarantee on that law. Path ES establishes a principled distributional target for evaluation; the remaining experiments establish whether the model also meets the structural and intervention requirements of protein world modeling.

B Evaluation Protocol and Supporting Experiments

This appendix provides evaluation details and supporting experiments for the main results. We focus on experiments that clarify a main claim, test an alternative explanation, or distinguish evaluation settings that should not be directly pooled.

Unless explicitly labeled *pre-projection*, "Ours", "Full", and "F" denote the final sampler in Equation 9, including DiffProject.

B.1 Evaluation Protocol

Splits and query manifest. The primary mdCATH evaluation uses trajectories at 320, 348, 379, 413, and 450 K. Protein domains are partitioned by CATH homologous superfamily, with all temperatures and simulation replicas associated with a domain assigned to the same partition. Development domains inspected during model design are excluded from the final test set.

A fixed query manifest specifies the protein identity, observed history, current DSSP state D_t, requested temperature action, forecast timestamps, replica, and valid-residue masks for each evaluation case. The default history contains eight observations at 1 ns cadence, including the current frame, and therefore spans 7 ns. Each query is evaluated using $M = 16$ independently sampled continuations. Unless otherwise stated, the prediction horizon is 16 ns.

Metric definitions, masks, geometric thresholds, history length, and projection settings are fixed before test evaluation. Scores are first computed at the query level and then macro-aggregated over proteins, preventing long proteins or domains with more valid frames from dominating the reported results.

Physical time and output cadence. Models producing outputs at different nominal frame intervals are evaluated only at matched physical timestamps. The standard Path ES representation uses observations at 4, 8, 12, and 16 ns. A 4 ns output is therefore never treated as a 1 ns observation. Closed-loop continuation maintains the generated history at the 1 ns cadence required by the encoder and subsamples only for reporting.

Path representation and proper scoring. For each structure X, the evaluation feature vector contains radius of gyration, soft native-contact retention, the fraction of initial contacts lost, and the 10th, 50th, and 90th percentiles of nonlocal C_α distances. Path ES concatenates normalized feature vectors at the four evaluation timestamps and scores their joint distribution. The feature map and normalization are fixed before test evaluation.

Path ES is complemented by RMWD, PCA W_2, residue-wise RMSF correlation and amplitude error, contact statistics, and native-reference RMSD error. These metrics probe complementary marginals and structural observables and are not individually interpreted as tests of temporal correctness. Appendix A gives the mathematical justification for Path ES.

Contact events. Initial contacts are C_α pairs separated by at least four sequence positions and within 8 Å at the forecast origin. A contact is considered lost once its distance exceeds 10 Å. Event probabilities used in evaluation are computed from generated trajectories rather than from the auxiliary event-prediction head. Event thresholds and persistence criteria are fixed independently of final test predictions.

Geometry and validity. The backbone audit operates on N–C_α–C trajectories and reports inter-residue C–N bond outliers, peptide-angle outliers, severe nonlocal backbone clashes, frame validity, and whole-path validity. All generated frames contribute to the primary distributional scores; geometrically invalid predictions are not discarded before evaluation.

These diagnostics assess compatibility of the represented backbone rather than all-atom physical validity. The C_α validity statistic used in external benchmark comparisons is therefore not numerically interchangeable with the native-backbone validity used in matched geometry ablations.

Table 5: Additional reference settings for the common mdCATH evaluation.

Method	Information condition	RMWD ↓	PCA W_2 ↓	RMSF r ↑	Weak J ↑	Transient J ↑	nRMSD err. ↓
MD replica	independent reference	1.90	1.65	0.95	0.84	0.55	0.85
Ours	free conditional continuation	1.35	0.82	0.86	0.73	0.46	1.15
ConfRover-IID	unordered conformational set	3.25	2.30	0.83	0.58	0.32	NA

Table 6: Table 6: Errors in predicted temperature-dependent changes in trajectory observables.

	Full temperature input				Temperature-blind control			
Contrast	ΔRMSF (Å)	ΔR_g (Å)	Δcontact	ΔnRMSD (Å)	ΔRMSF (Å)	ΔR_g (Å)	Δcontact	ΔnRMSD (Å)
$348 - 320$ K	0.12	0.08	0.018	0.15	0.28	0.20	0.045	0.35
$379 - 320$ K	0.18	0.12	0.025	0.22	0.48	0.36	0.075	0.60
$413 - 320$ K	0.27	0.20	0.040	0.35	0.80	0.62	0.130	0.95
$450 - 320$ K	0.40	0.32	0.060	0.55	1.20	0.95	0.200	1.50

Closed-loop continuation. Long-horizon evaluation feeds generated states back into the model without access to reference future coordinates or future annotations. The direct-versus-composed experiment compares alternative generation procedures at the same 16 ns physical horizon. The 32 and 64 ns evaluations recursively reuse generated histories. Failed trajectories remain part of the reported survival statistics.

B.2 Additional Cross-Method Reference

The main cross-method comparison is restricted to learned models with compatible forward-generation interfaces. Table 5 additionally reports ConfRover-IID, whose output is an unordered conformational set rather than the same forward conditional path object. We therefore provide it as contextual evidence rather than pool it into the main learned-trajectory ranking.

B.3 History Length

The main history experiment tests whether temporal ordering contributes predictive information beyond the current structure or an unordered set, rather than optimizing the history length itself. Ordered histories of 5, 8, 9, and 17 frames obtain Path ES values of 0.275, 0.260, 0.258, and 0.257, respectively. The gain beyond eight frames is therefore small relative to the Current-to-Ordered and Set-to-Ordered differences in the main text. We consequently use eight frames as the default operating point without claiming that it is globally optimal.

B.4 Supplementary Trajectory Visualization

Figure 4 provides a complementary view of the generated dynamics over one complete 16-step continuation. The generated sequence follows the reference conformational evolution while retaining stochastic variation across independently sampled paths. The residue-flexibility and PCA views provide complementary structural and collective-motion diagnostics, while Path ES summarizes the resulting trajectory distribution.

B.5 Temperature-Response Evaluation

The main temperature table evaluates ordinary temperature conditioning and leave-temperature-out generalization. The intervention experiment instead holds the molecular checkpoint and observed history fixed while changing only the requested future temperature.

Table 6 reports all four matched temperature contrasts. Each entry is the absolute error between the model-predicted change in an observable and the corresponding response measured from matched MD branches.

The advantage of the action-conditioned model becomes larger as the thermal intervention increases. For the $450 - 320$ K contrast, for example, the ΔRMSF response error is 0.40 Å compared with

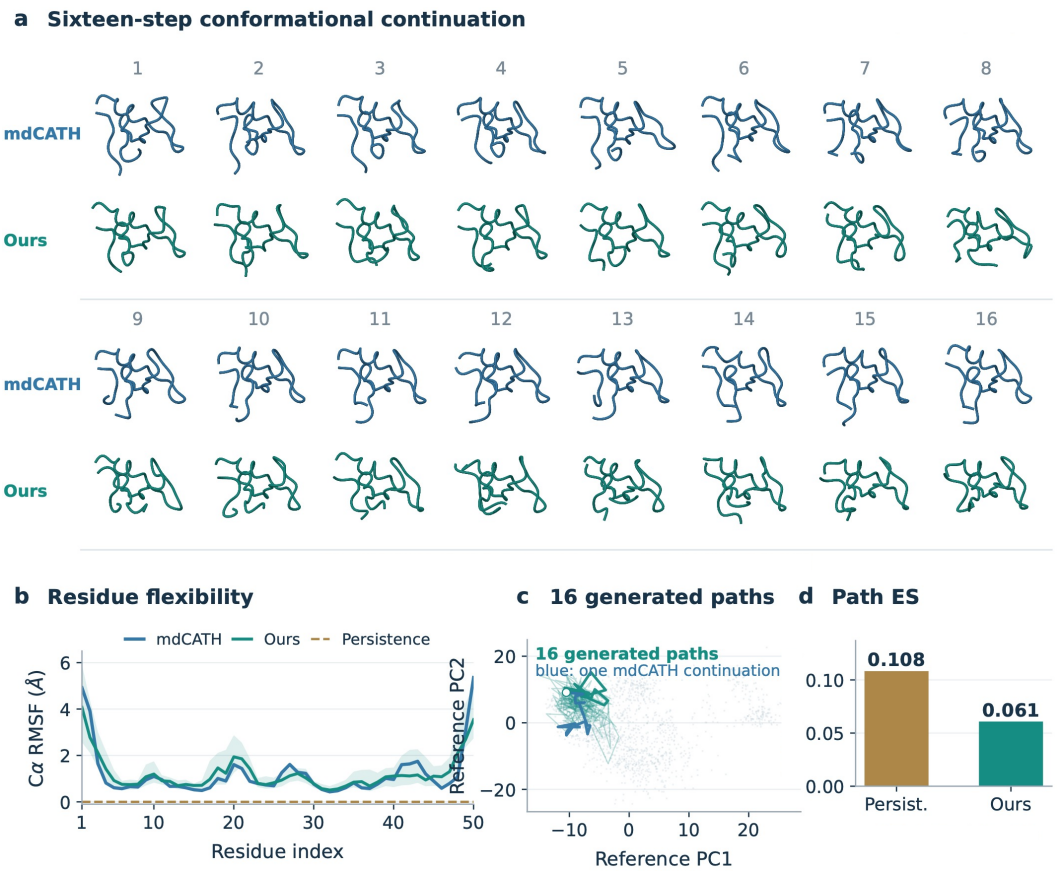


Figure 4: Supplementary diagnostics for a 16-step conformational continuation. **(a)** Reference mdCATH and generated backbone conformations across all 16 future steps. **(b)** Residue-wise C_α RMSF for mdCATH, our model, and persistence. **(c)** Sixteen independently generated trajectories projected into the reference PCA space, together with the observed MD continuation. **(d)** Path ES for persistence and our model.

1.20 Å for the temperature-blind control. These experiments evaluate action responsiveness under matched molecular initialization.

B.6 Closed-Loop Rollout Diagnostics

The main text summarizes closed-loop behavior using path survival and autocorrelation error. Table 7 reports the complementary dynamical errors used in the same analysis. The direct-versus-composed comparison isolates the effect of committing generated states and re-encoding them at the standard 16 ns horizon, whereas the 32 and 64 ns settings measure error accumulation under repeated self-conditioning.

At the matched 16 ns horizon, composing four short transitions only modestly increases the dynamical errors relative to direct prediction. Errors grow gradually under generated-history rollout, but at 64 ns the full model remains substantially closer to the reference dynamics than the no-context control. Persistence illustrates the complementary failure mode: geometrically stable states alone do not imply realistic molecular evolution.

B.7 Geometry and Projection Audit

The final model includes DiffProject, so all main-paper "Ours" results use corrected trajectories. To distinguish learned geometry from post-processing, the geometry experiment retains the same

Table 7: Closed-loop trajectory diagnostics on mdCATH.

Setting	Lag-RMSD ↓	ACF err. ↓	PCA W_2 ↓	Contact err. ↓
Persistence / 16 ns	0.90	NA	1.60	0.180
Persistence / 64 ns	1.60	NA	2.60	0.300
Ours / direct 16 ns	0.18	0.060	0.82	0.055
Ours / composed 4×4 ns	0.21	0.070	0.88	0.061
Ours / generated-history 32 ns	0.29	0.095	1.05	0.075
Ours / generated-history 64 ns	0.42	0.135	1.32	0.098
No-context-noise / 64 ns	0.78	0.250	1.80	0.160

Table 8: Controls for zero-shot mdCATH-to-ATLAS transfer.

mdCATH training / ATLAS input	RMWD ↓	RMSF r ↑	RMSF err. ↓	Weak J ↑	Transient J ↑	PCA W_2 ↓	nRMSD err. ↓
Full five-T / 300 K	1.60	0.90	0.35	0.72	0.42	0.70	1.35
Same checkpoint / clamp 320 K	1.80	0.87	0.43	0.68	0.38	0.83	1.55
Strict temperature-blind	2.05	0.83	0.50	0.63	0.33	0.98	1.75
320-K-only / total-token matched	1.95	0.85	0.47	0.65	0.35	0.92	1.68
Five-T / unique-window-count matched	1.72	0.88	0.39	0.69	0.39	0.78	1.46

generated trajectories before projection and compares them with outputs after applying the identical projection operator.

The full quantitative ablation is reported in Table 2b, while Figure 3 separates the effects of geometry-aware training and projection. The model already improves C–N geometry before projection and retains its advantage when both the control and full model receive the same projection. Across validation domains, the improvement extends to C–N outliers, backbone-angle error, and nonlocal clashes.

The excessive-geometry control provides an important complementary check. Although it obtains lower local defect rates, its Motion/MD value decreases to 0.68. We therefore evaluate geometric compatibility jointly with retained molecular motion rather than treating local geometry as a standalone objective.

B.8 ATLAS Transfer Controls

The main ATLAS experiment evaluates zero-shot transfer to clean, topology-novel, and high-motion targets. Here we examine whether that transfer depends on the temperature action and on multi-temperature training.

ATLAS reference trajectories are evaluated at 300 K. Supplying 320 K to the same checkpoint is therefore an intentionally incorrect-action diagnostic rather than an evaluation against hypothetical 320 K ATLAS ground truth. The strict temperature-blind model removes the action input entirely. The final two controls separate temperature diversity from overall training coverage.

Using the correct 300 K action with the same full checkpoint improves RMWD from 1.80 to 1.60 relative to clamping the action at 320 K, while the temperature-blind model degrades further to 2.05. The matched training controls additionally show that multi-temperature training remains beneficial when training coverage is controlled.

B.9 Temporal Metric Sanity Checks

Several common trajectory metrics are insensitive to temporal ordering. RMWD, RMSF, PCA distributions, contact occupancy, and per-frame geometry can remain unchanged when the same conformations are permuted in time. They therefore cannot, by themselves, establish dynamical correctness.

We verify this with an exact-shuffle negative control. At 99 ns, ordered and time-shuffled trajectories have identical PCA W_2 (0.70) and contact-occupancy error (0.085), since both retain the same individual frames. In contrast, shuffling increases lag-RMSD error from 0.46 to 1.20 Å, ACF error

from 0.170 to 0.460, and low-frequency power error from 0.180 to 0.500. This motivates the explicit temporal diagnostics used in the rollout analysis rather than inferring dynamics from ensemble statistics alone.

Persistence provides the complementary negative control: it can preserve static structural validity while containing no realistic molecular evolution. Together, these controls motivate the separation of ensemble fidelity, structural validity, and temporal dynamics throughout the evaluation.

C METHOD IMPLEMENTATION DETAILS

This appendix provides the implementation details underlying Section 4. We first specify the observed context, architecture, and base training objective, then describe the differentiable DIFFPROJECT update and its projection-aware supervision. Generated-state geometric regularization is reported separately as an ablation and is not part of the default sampling interface.

C.1 OBSERVED CONTEXT AND ARCHITECTURE

Observed context. The default encoder receives eight structural observations at 1 ns intervals, with the current structure at the final position, giving a 7 ns observed history. Relative translations in Equation 4 are normalized by $\ell = 5$ Å. For each residue, a linear layer embeds the three relative-translation components and nine relative-rotation entries, and learned temporal position embeddings encode their order. Current-query temporal attention then aggregates the history separately for each residue.

The encoder additionally receives amino-acid identity and the observed DSSP state D_t at the forecast origin. Current force and explicit finite-difference velocity channels are set to zero in the primary configurations. Future coordinates, forces, DSSP annotations, and event labels are excluded from the inference context and are used only as training targets.

Spatial and temporal architecture. Pair features are constructed from 16 radial basis functions with centers between 0 and 20 Å, local displacement features scaled by 10 Å, relative residue orientations, and sequence separation clipped at 32. The pair representation supplies per-head biases to spatial attention.

The trajectory backbone contains 12 blocks with hidden width 192 and eight attention heads. Each block alternates pair-biased spatial attention across residues, bidirectional temporal attention across future frames, and a feed-forward update with AdaLN-Zero conditioning.

For a physical lead time $q > 0$ measured in ns, we define

$$z(q) = \frac{\log_2(\max(q, 1))}{6}, \tag{33}$$

and embed

$$\left[z(q), \left\{\sin(2^k \pi z(q)), \cos(2^k \pi z(q))\right\}_{k=0}^{7}\right] \tag{34}$$

with an MLP. This physical-time embedding is distinct from both the generative flow time τ and the learned future-position embedding.

The model predicts 16 future positions. Training uses output cadences of either 1 or 4 ns, corresponding to maximum joint horizons of 16 and 64 ns, respectively. A 64 ns joint trajectory is a single sampled future block and is distinct from the repeated closed-loop continuation evaluated in Section 4.9.

C.2 BASE TRAINING OBJECTIVE

The compact objective in Equation 10 contains the flow-matching loss together with structural and physical auxiliary supervision. For the base configuration, the implemented objective is

$$\begin{aligned}\mathcal{L}_{\text{base}} = \mathcal{L}_{\text{FM}} &+ 0.10\mathcal{L}_{\text{force}} + 0.10\mathcal{L}_{\text{DSSP}} + 0.50\mathcal{L}_{\text{evt}} \\ &+ 0.10\mathcal{L}_{\text{FAPE}} + 0.05\mathcal{L}_{\text{bond}} + 0.02\mathcal{L}_{\text{angle}} + 0.01\mathcal{L}_{\text{clash}} + 0.02\mathcal{L}_{\text{contact}}.\end{aligned} \tag{35}$$

All terms are averaged over their corresponding valid residues, frames, pairs, or event labels.

Clean endpoint estimate. Several auxiliary terms operate on an estimate of the clean future obtained from an intermediate flow state. For the Euclidean and rotational components, respectively, we use

$$\widehat{d}_1 = d_\tau + (1 - \tau)v_\theta, \qquad \widehat{Q}_1 = Q_\tau \operatorname{Exp}\left((1 - \tau)[\omega_\theta]_\times\right). \tag{36}$$

The resulting residue frames are decoded before evaluating the structural auxiliaries.

Physical and structural auxiliaries. Force supervision uses smooth-L_1 loss with transition parameter 0.1 on the nine local N/C_α/C force components per residue; cached force targets are divided by 50 before training. Future DSSP prediction is supervised with cross-entropy.

FAPE compares pairwise C_α displacements expressed in predicted and reference residue frames, with distances clipped at 10 Å and normalized by the same scale. The bond term compares adjacent C_α and C–N distances using scales of 3.8 and 1.33 Å, respectively. The angle term compares cosine values of consecutive C_α triplets. The clash term is a squared hinge below 3 Å for non-neighboring C_α pairs, and the contact term compares soft contacts $\sigma((8-d)/0.5)$.

For these base auxiliaries, both predicted and reference residue frames are decoded with the local atom template observed at the forecast origin. Projection-aware supervision below instead uses the actual future local backbone templates as training targets.

Auxiliary event prediction. The training-only event head predicts three hazards for each of two events, contact loss and secondary-structure loss. Their cumulative products define risks at 4, 16, and 64 ns. The historical event protocol fixes the 1 ns entry to zero because event confirmation requires multiple saved frames. Binary cross-entropy is applied only to eligible observed labels. This auxiliary loss is distinct from the generated-path event evaluation in Section **??**; reported event probabilities are not replaced by predictions from this head.

C.3 DIFFERENTIABLE PROJECTION

Connectivity mask. Connectivity is determined once at the forecast origin. Consecutive valid residues are treated as connected when their C_α–C_α distance lies between 3.3 and 4.3 Å and their C–N distance lies between 1.0 and 1.8 Å. The resulting mask is fixed throughout the generated future block and prevents the projection from connecting genuine chain breaks.

Capped optimization updates. For the projection energy in Equation 49, define the smoothly capped update

$$\mathcal{C}_{\gamma,c}(g) = \frac{\gamma g}{\sqrt{1 + \|\gamma g\|_2^2 / c^2}}. \tag{37}$$

Starting from $u^{(0)} = \eta^{(0)} = 0$, each projection iteration performs

$$\begin{aligned} u^{(q+1)} &= u^{(q)} - \mathcal{C}_{\gamma,c_p}\left(\nabla_u \mathcal{E}_{\mathrm{proj}}\right), \\ \eta^{(q+1)} &= \eta^{(q)} - \mathcal{C}_{\gamma,c_R}\left(\nabla_\eta \mathcal{E}_{\mathrm{proj}}\right). \end{aligned} \tag{38}$$

The cap is applied independently to each three-dimensional residue update. We use eight iterations with step size $\gamma = 0.08$, translation cap $c_p = 0.2$ Å, rotation cap $c_R = 0.05$ rad, and anchor weight 0.1. The final corrected frames are

$$\left(p + u^{(8)},\ R\,\mathrm{Exp}([\eta^{(8)}]_\times)\right). \tag{39}$$

Offsets on padded residues remain zero.

The projection energy is summed over valid edges and frames rather than averaged over the batch, so its per-trajectory update is unaffected by the number of independent trajectories evaluated together. DIFFPROJECT operates independently at each future time point and contains no explicit temporal-smoothing term.

During projection-aware training, gradients are propagated through all eight updates. A small-angle expansion of the rotational exponential keeps second derivatives finite at zero initialization, and connected edges are selected before bond and angle calculations to avoid undefined derivatives on padded coordinates. Inference uses the identical update rule without constructing a training gradient graph. Neither the projection energy nor the update operator accesses reference future coordinates.

C.4 PROJECTION-AWARE SUPERVISION

Let $\widetilde{X}$ denote backbone coordinates obtained after applying DIFFPROJECT to the clean endpoint estimate in Equation 36, and let $X^\star$ denote the reference future backbone reconstructed using the actual future local atom templates for supervision only.

For each connected peptide edge, define

$$\begin{aligned} e_d &= \left(\frac{d^{\mathrm{CN}}(\widetilde{X}) - d^{\mathrm{CN}}(X^\star)}{1.33\,\text{Å}} \right)^2, \\ e_\alpha &= \left\| \cos \boldsymbol{\alpha}(\widetilde{X}) - \cos \boldsymbol{\alpha}(X^\star) \right\|_2^2, \\ e_\omega &= \left\| \chi(\widetilde{X}) - \chi(X^\star) \right\|_2^2, \qquad \chi(X) = (\cos \omega(X), \sin \omega(X)), \end{aligned} \tag{40}$$

where ω is the C_α–C–N–C_α peptide torsion.

With

$$w(\tau) = 0.1 + 0.9\tau^2, \tag{41}$$

the projection-aware loss is

$$\begin{aligned} \mathcal{L}_{\mathrm{proj}} = 0.2 \, \langle w(\tau) \, (e_d + e_\alpha + 0.1 e_\omega) \rangle_{\mathrm{edge}} \\ + 0.1 \left\langle w(\tau) \left(\frac{\widetilde{X} - X^\star}{5\,\text{Å}} \right)^2 \right\rangle_{\mathrm{coord}}, \end{aligned} \tag{42}$$

where $\langle \cdot \rangle_{\mathrm{edge}}$ averages over connected peptide edges and future frames and $\langle \cdot \rangle_{\mathrm{coord}}$ averages over valid backbone-coordinate components.

Peptide torsion is supervised in Equation 42 but is not constrained by the projection energy itself. The projection-aware configuration optimizes

$$\mathcal{L}_{\text{proj-aware}} = \mathcal{L}_{\mathrm{base}} + \mathcal{L}_{\mathrm{proj}}. \tag{43}$$

Thus, in the compact notation of Equation 10, $\mathcal{L}_{\mathrm{geom}} = \mathcal{L}_{\mathrm{proj}}$ and its outer coefficient is one; the relative weights of the geometry terms are already specified in Equation 42.

A matched raw-supervision control applies corresponding backbone supervision before DIFFPROJECT. This separates improvements learned by the trajectory generator from those introduced by the geometric correction at inference.

C.5 GENERATED-STATE GEOMETRY REGULARIZATION

We additionally evaluate a generated-state geometry objective as an ablation. Unlike projection-aware supervision, this objective is evaluated on trajectories reached by the sampler from noise and requires no reference future coordinates.

For generated backbone coordinates, the local penalty contains ideal C–N bond and peptide-angle terms together with a nonlocal backbone-overlap term

$$\mathcal{L}_{\mathrm{overlap}} = \frac{1}{H \sum_b N_b} \sum_{b,j} \sum_{(m,n) \in \mathcal{P}_b} \left[\max \left(0, \frac{2\,\text{Å} - d_{b,j,mn}}{2\,\text{Å}} \right) \right]^2, \tag{44}$$

where N_b is the number of valid residues in trajectory b and $\mathcal{P}_b$ contains unordered N/C_α/C atom pairs whose residues are separated by more than one sequence position. Normalizing by the number of residues and frames prevents the loss from being diluted by the much larger number of non-overlapping atom pairs.

The generated-state auxiliary is evaluated on four windows from each training batch. Sampling uses the same 12 flow updates as inference; the first ten updates are detached and gradients are propagated through the final two. The observed-state encoder is recomputed with gradients enabled before these differentiable updates.

Its coefficient is fixed before validation using the first three paired training batches. If r denotes the median ratio between the generated-state auxiliary gradient norm and the base-objective gradient norm, the implementation uses

$$\lambda_{\mathrm{gen}} = \mathrm{clip} \left(\frac{0.3}{r}, 0.01, 10 \right). \tag{45}$$

A matched teacher-conditioned control evaluates the same target-free geometry penalty on clean endpoint estimates along the training interpolant.

This generated-state objective is used only in the corresponding geometry ablation. It is distinct from differentiating through DIFFPROJECT, and the standard inference pipeline retains the same eight-step projection operator.

C.6 SAMPLING DETAILS

Stochastic inference begins from independent Gaussian translation variables and isotropic Gaussian tangent vectors mapped to $SO(3)$. Unless otherwise stated, we integrate the learned flow with 12 first-order updates and midpoint flow-time inputs

$$\tau_k = \frac{k + \frac{1}{2}}{12}, \qquad k = 0, \dots, 11. \tag{46}$$

The predictive state is encoded once and reused across the 12 flow updates. Independent prior draws produce the multiple continuations used in distributional evaluation.

Sampling accepts only the observed structural history, static molecular context, current DSSP state, requested temperature, valid-residue and connectivity masks, and requested physical lead times. No future coordinates, future forces, future DSSP annotations, or event labels enter the inference path.

D PARTIAL OBSERVABILITY AND STATE INFERENCE

The conditional path distribution in Equation 3 can be decomposed through the unobserved microscopic state Z_t:

$$P^{\star}_{H,\Delta}\left(Y_t \mid h_t, c, D_t, a\right) = \int P^{\star}_{H,\Delta}\left(Y_t \mid Z_t, c, a\right) P^{\star}\left(Z_t \mid h_t, c, D_t\right) \, \mathrm{d}Z_t. \tag{47}$$

This decomposition clarifies that the observed history provides information about the current dynamical state Z_t through the posterior $P^{\star}(Z_t \mid h_t, c, D_t)$, while the action a determines the forward dynamics $P^{\star}_{H,\Delta}(Y_t \mid Z_t, c, a)$. The learned model P_θ implicitly approximates this composition without explicitly reconstructing the full microscopic state.

E FLOW MATCHING DETAILS

For a reference trajectory $U_1 = (d_1, Q_1)$ and a prior sample $U_0 = (d_0, Q_0) \sim P_0$, translations follow linear interpolation while rotations follow the geodesic on $SO(3)$:

$$\begin{aligned} d_\tau &= (1-\tau)d_0 + \tau d_1, \\ \Omega &= \mathrm{Log}(Q_0^{\mathsf{T}} Q_1)^{\vee}, \\ Q_\tau &= Q_0 \,\mathrm{Exp}(\tau[\Omega]_\times), \qquad \tau \sim \mathcal{U}[0,1]. \end{aligned} \tag{48}$$

Here $[\cdot]_\times$ maps a three-dimensional vector to a skew-symmetric matrix and $\vee$ denotes its inverse. The corresponding velocity targets are $v^\star = d_1 - d_0$ for translations and $\omega^\star = \Omega$ for rotations.

F DIFFPROJECT ENERGY AND OPTIMIZATION

For corrected frame $\widetilde{F}_{j,i} = (p_{j,i} + u_{j,i}, R_{j,i}\,\mathrm{Exp}([\eta_{j,i}]_\times))$, we optimize

$$\begin{aligned} \mathcal{E}_{\mathrm{proj}}(u, \eta; F) = \sum_{j=1}^{H} \sum_{i \in \mathcal{E}_c} \left[\left(\frac{d^{\mathrm{CN}}_{j,i}(\widetilde{F}) - d_\star}{d_\star} \right)^2 + \left\| \cos \boldsymbol{\alpha}_{j,i}(\widetilde{F}) - \cos \boldsymbol{\alpha}_\star \right\|_2^2 \right] \\ + \lambda_{\mathrm{anc}} \sum_{j,i} \left[\|u_{j,i}/\ell_p\|_2^2 + \|\eta_{j,i}\|_2^2 \right]. \end{aligned} \tag{49}$$

Here $\mathcal{E}_c$ contains valid peptide connections, $d^{\mathrm{CN}}_{j,i}$ is inter-residue C–N distance, $\boldsymbol{\alpha}_{j,i}$ contains adjacent peptide angles, and $d_\star, \boldsymbol{\alpha}_\star$ denote reference geometry. Starting from zero offsets, we apply $J = 8$ gradient descent steps with step size 0.08 and anchor weight $\lambda_{\mathrm{anc}} = 0.1$. The operator is differentiable and can be used during both training and inference.