QPD: ORDER-AWARE LEARNING OF ALL-ATOM PROTEIN DYNAMICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative models of molecular dynamics learn which conformations a protein visits, yet a trajectory is also defined by the order in which it visits them. A residue can reach the same orientation through different sequences of local rotations, and losses on endpoints or on the distribution of frames cannot tell these sequences apart. We introduce Quaternion Path Dynamics (QPD), a training objective that supervises the ordered rotations of residue backbone frames across time. QPD compares rotation increments, their products, and their commutators, and contrasts each reference path with a Hurwitz negative that keeps the endpoint and the rotation angles but changes the intermediate orientation. Added to an all-atom trajectory generator, QPD improves protein ensembles on ATLAS and first-passage kinetics on held-out CATH2 domains over the same model with QPD switched off, raising the correlation of mean first-passage times from 0.763 to 0.805. It also improves on MDGen for tetrapeptide dynamics and generates ligand dissociation paths with few clashes. Ablations and frame-shuffling tests tie these gains to the order of frames. In short, **supervising how residues rotate teaches a generator how proteins move, not only where they go**.

1 INTRODUCTION

Generative models now produce protein structural ensembles and molecular dynamics (MD) trajectories at a fraction of the cost of simulation (Jing et al., 2024a; Lewis et al., 2025; Jing et al., 2024b; Feng et al., 2026b). Their training signals, however, mostly compare individual frames or the distribution of frames. Two trajectories can visit the same set of conformations and still describe different dynamics, because function often depends on the route: which domain moves first, which side of a pocket opens, which contacts break before others form. In reference MD of the tetrapeptide FLRH, the same frames in true and shuffled order give different lagged torsion distributions, and some backbone torsions stay correlated for about 1 ns (Appendix B.3). A model that matches the ensemble but not the order of intermediate states captures the thermodynamics while missing the mechanism.

Local backbone orientations make this order measurable. The orientation of each residue lies on $\mathrm{SO}(3)$, where rotations do not commute: applying the same two rotations in opposite order generally leads to different intermediate and final orientations. Endpoint or per-frame coordinate errors therefore summarize only part of a rotation sequence, and a denoiser trained on them receives no direct signal about composition.

In this work, we introduce Quaternion Path Dynamics (QPD), an order-aware training objective that supervises ordered local rotations in all-atom trajectory generators. QPD builds N–Cα–C frames for every residue, compares predicted and reference rotation increments, their short-window products, and their commutators, and contrasts the reference with matched-endpoint Hurwitz negatives. A Hurwitz move reorders two increments while preserving the net rotation and both rotation angles, so the reference and the negative differ only in the intermediate orientation (Figure 1). The objective has no learnable parameters and here trains a Protenix-based spatiotemporal denoiser (Protenix Team, 2024). On ATLAS (Vander Meersche et al., 2024), adding QPD during adaptation improves all seven flexibility and distributional metrics, and on CATH2-OOD20 kinetics it raises the Pearson correlation of mean first-passage times from 0.763 to 0.805 at the same training budget.

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.

We find that (1) supervising ordered rotation increments makes the accuracy of a generator depend on the chronology of its frames, (2) this dependence relies most on the Hurwitz negatives, which a plain swap of increments cannot replace, (3) order learned on short windows carries over to transitions between metastable states and to protein flexibility, leading to better kinetic and ensemble statistics, and (4) the gain in order costs a small amount of accuracy in static pair distances. Our findings highlight a gap between matching the conformations a protein visits and learning the order in which it visits them. Trajectory generators should be trained on this order and evaluated with controls that remove it, such as frame permutation, alongside ensemble statistics.

2 Related work

Conformational ensemble generation. Alternative conformations can be drawn from structure predictors by subsampling or clustering the multiple sequence alignment, or learned directly from MD data with flow- and diffusion-based generators, physical feedback, and large-scale emulation (del Alamo et al., 2022; Wayment-Steele et al., 2024; Noé et al., 2019; Jing et al., 2024a; Zheng et al., 2024; Lu et al., 2024; Wang et al., 2024; Lewis et al., 2025; Lu et al., 2025). These models describe which conformations a protein can adopt. QPD targets the order in which a trajectory visits them.

Trajectory learning for MD simulation. Time-ordered generators learn transition densities over a fixed lag, whole trajectories as molecular videos, frames autoregressively with causal attention, structure tokens, or probabilistic autoregression, dynamics at two time scales, through Markov-state transitions, or in state-space and learned latent models, and all-atom protein–ligand trajectories with forecasting and interpolation (Klein et al., 2023; Schreiner et al., 2023; Costa et al., 2024; Jing et al., 2024b; Shen et al., 2025; Shoghi et al., 2026; Liu et al., 2026; Cheng et al., 2026; Xu et al., 2025; Kapuśniak et al., 2026; Shi et al., 2026; Sengar et al., 2026; Liu et al., 2025; Feng et al., 2026b;a). They are trained with losses on coordinates, internal coordinates, tokens, or frame statistics. QPD adds a loss on the ordered rotation path of each residue and leaves the generator architecture unchanged. Appendix A compares QPD with the closest of these models one by one.

Rotations, ordered motion, and kinetics. Structure models represent each residue as a rigid backbone frame (Jumper et al., 2021; Yim et al., 2023; Watson et al., 2023; Bose et al., 2024), and the order in which rotations compose has been studied in domain motion, quaternionic protein deformation, Lie-group integration, and path signatures (Hayward & Berendsen, 1998; Chen et al., 2026; Blanes et al., 2009; Lyons, 1998). Markov state models, TICA, path samplers such as DIMS, FRODA, and metadynamics, and flow-based transition-path generators summarize slow dynamics and rare transitions (Prinz et al., 2011; Pérez-Hernández et al., 2013; Husic & Pande, 2018; Perilla et al., 2011; Wells et al., 2005; Laio & Parrinello, 2002; Seyler et al., 2015; Xu et al., 2026). QPD follows backbone frames along physical time and uses the Hurwitz move (Hurwitz, 1891) to build negative paths that share the net rotation of the reference. We evaluate it against these kinetic and path references.

3 Preliminaries

Let S be a protein sequence and $X_{0:F-1} \in \mathbb{R}^{F\times A\times 3}$ a reference trajectory with F frames and A atoms. Frames are ordered by physical time, while τ indexes denoising of the entire window. The observed-frame indices are $\mathcal{C}$, and (s, z) are single and pair features from a frozen Protenix trunk (Protenix Team, 2024). Given noisy coordinates $\widetilde{X}^{(\tau)}$ and per-frame noise levels $\boldsymbol{\sigma}(\tau)$, a spatiotemporal denoiser with parameters θ predicts

$$\widehat{X}^{(\tau)} = D_\theta\Big(\widetilde{X}^{(\tau)}, \boldsymbol{\sigma}(\tau), \mathcal{C}; s, z\Big)\,. \tag{1}$$

Observed frames enter as clean conditioning, so the same network forecasts from an initial frame and interpolates between two endpoints (Chen et al., 2024; Feng et al., 2026b). Coordinate denoising compares $\widehat{X}$ with the reference frame by frame. This constrains where the atoms are in each frame and gives no direct signal about how local rotations compose from frame to frame, which is the signal QPD adds.

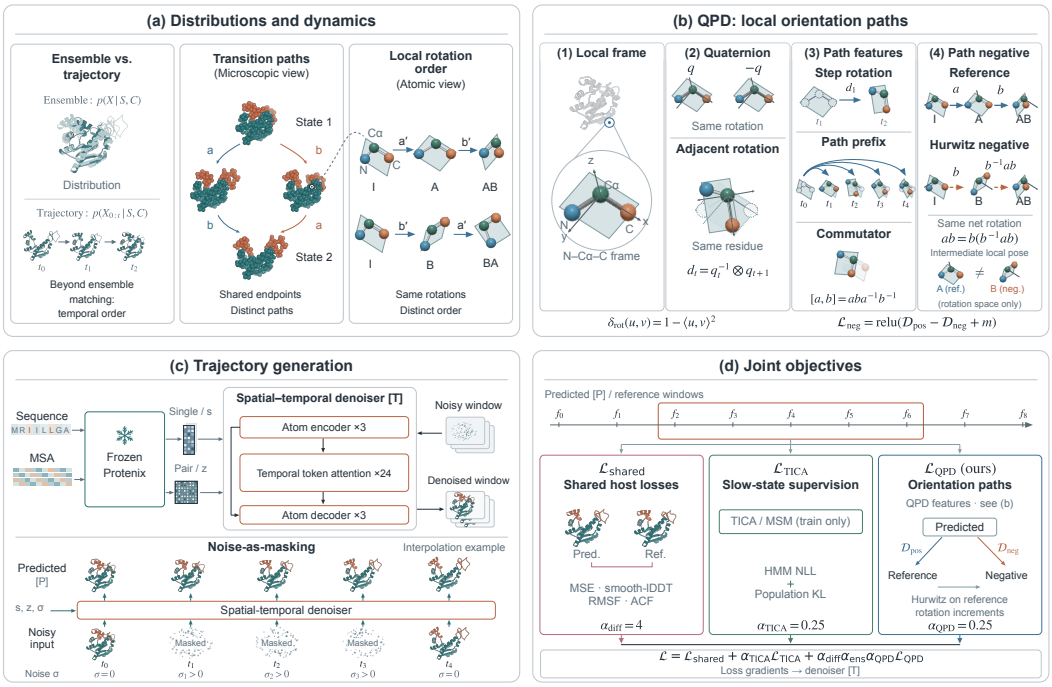

Figure 1: **Quaternion Path Dynamics.** (a) Ensembles and endpoints leave the order of intermediate states undetermined. (b) Backbone frames define quaternion rotation paths, and a Hurwitz move changes an intermediate orientation while keeping the net rotation and increment angles. (c) A frozen Protenix trunk conditions the spatiotemporal denoiser. (d) QPD augments the denoising and slow-state objectives.

4 Quaternion Path Dynamics

QPD makes the training objective order-aware by comparing each residue across the physical frames of the prediction $\widehat{X}$ in Equation (1) and of the paired reference window Y at a fixed denoising iteration. We omit τ below and write t for the frame index. Figure 1 summarizes the computation, and Algorithm 1 lists one training step.

4.1 Local frames and rotations

For protein residue r at frame t, let $x_{t,r}^N$, $x_{t,r}^{CA}$, and $x_{t,r}^C$ be its backbone atom positions, $u = x^C - x^{CA}$ and $v = x^N - x^{CA}$. For complete, non-collinear triplets we define

$$e_1 = \frac{u}{\|u\|}, \qquad e_2 = \frac{v - (v^\top e_1)e_1}{\|v - (v^\top e_1)e_1\|}, \qquad e_3 = e_1 \times e_2, \qquad R_{t,r} = [e_1\ e_2\ e_3], \tag{2}$$

and represent each frame by a sign-aligned unit quaternion $q_{t,r}$. With the Hamilton product $\otimes$, the rotation increment and the discrepancy between unit quaternions q, p are

$$d_{t,r} = q_{t,r}^{-1} \otimes q_{t+1,r}, \qquad \delta_{\mathrm{rot}}(q, p) = 1 - \langle q, p \rangle^2. \tag{3}$$

The squared inner product removes the quaternion sign ambiguity, and $\delta_{\mathrm{rot}} = \sin^2(\vartheta/2)$ for relative angle ϑ. Increments are invariant to any common rigid transformation of the trajectory.

4.2 Composing rotations over a window

Increments describe adjacent motion, while products and commutators describe how motions compose. For a window starting at frame j with ℓ increments,

$$\begin{aligned} P_{j,\ell} &= d_j \otimes \cdots \otimes d_{j+\ell-1}, \qquad A_{j,\ell} = d_j \otimes \cdots \otimes d_{j+\ell-2}, \qquad B_{j,\ell} = d_{j+\ell-1}, \\ C_{j,\ell} &= A_{j,\ell} \otimes B_{j,\ell} \otimes A_{j,\ell}^{-1} \otimes B_{j,\ell}^{-1}. \end{aligned} \tag{4}$$

Algorithm 1 QPD training step (per trajectory window).

Require: Reference window Y (F frames), observed indices $\mathcal{C}$, frozen features (s, z), denoiser D_θ
Ensure: Loss $\mathcal{L}$ and parameter update of θ

/ Stage I: Denoising prediction */*
$\widetilde{X} \leftarrow Y + \boldsymbol{\sigma} \odot \epsilon$ on frames $\notin \mathcal{C}$ ▷ noise-as-masking
$\widehat{X} \leftarrow D_\theta(\widetilde{X}, \boldsymbol{\sigma}, \mathcal{C}; s, z)$ ▷ Eq. 1
/ Stage II: Ordered rotation features */*
for $Z \in \{\widehat{X}, Y\}$, residue r, frame t **do**
  $R_{t,r} \leftarrow \mathrm{Frame}(x^N_{t,r}, x^{CA}_{t,r}, x^C_{t,r})$; $q_{t,r} \leftarrow \mathrm{Quat}(R_{t,r})$ ▷ Eq. 2
  $d_{t,r} \leftarrow q^{-1}_{t,r} \otimes q_{t+1,r}$ ▷ Eq. 3
  $P_{j,\ell}$, $C_{j,\ell} \leftarrow$ products and commutators, $2 \le \ell \le 4$ ▷ Eq. 4
/ Stage III: Hurwitz negatives */*
$HY \leftarrow$ apply $H(d_j, d_{j+1})$ to the first pair of each window with $\ell \ge 3$ ▷ Eq. 5
$\Delta_Y \leftarrow \mathcal{D}(\Phi(Y), \Phi(HY))$; keep windows with $\Delta_Y \ge \epsilon_{\mathrm{neg}}$
$\mathcal{L}_{\mathrm{neg}} \leftarrow \mathrm{mean\ ReLU}(\mathcal{D}_{\mathrm{pos}} - \mathcal{D}_{\mathrm{neg}} + m(Y))$ ▷ Eq. 6
/ Stage IV: Objective and update */*
$\mathcal{L}_{\mathrm{QPD}} \leftarrow \mathcal{L}_{\mathrm{inc}} + \mathcal{L}_{\mathrm{prod}} + 0.25\mathcal{L}_{\mathrm{comm}} + 0.25\mathcal{L}_{\mathrm{neg}} + 0.10\mathcal{L}_{\mathrm{cond}}$ ▷ Eq. 7
Sum losses and valid counts across workers, then normalize
$\mathcal{L} \leftarrow \mathcal{L}_{\mathrm{shared}} + \alpha_{\mathrm{TICA}}\mathcal{L}_{\mathrm{TICA}} + \alpha_{\mathrm{diff}}\alpha_{\mathrm{ens}}\alpha_{\mathrm{QPD}}\mathcal{L}_{\mathrm{QPD}}$; $\theta \leftarrow \theta - \eta\nabla_\theta\mathcal{L}$ ▷ Eq. 8

We use all valid starts with $2 \le \ell \le 4$. Products accumulate orientation over a window, and commutators measure how strongly the order of composition matters.

4.3 Negative paths with the same endpoint

To make the model distinguish intermediate order from final orientation, we contrast each reference window with a path that shares its endpoint. For two successive increments a, b, the Hurwitz move

$$H(a, b) = \big(b,\ b^{-1} \otimes a \otimes b\big) \tag{5}$$

preserves the product, $b(b^{-1}ab) = ab$, and both rotation angles, because conjugation preserves the angle (Appendix B.1). The intermediate orientation changes. A plain swap (b, a) would instead change the endpoint whenever $ab \neq ba$. We apply H of Equation (5) to the first pair of every window with three or four increments. Windows start at every valid frame, so each adjacent pair except the last pair of a segment is perturbed in some window. For one residue and window, Φ collects the products and commutators of all contiguous sub-windows with $\ell \ge 2$, except the products that H leaves unchanged, and $\mathcal{D}(\Phi, \Phi') = |\mathcal{V}_\Phi|^{-1} \sum_{k \in \mathcal{V}_\Phi} \delta_{\mathrm{rot}}(\Phi_k, \Phi'_k)$ averages δ_{rot} over the valid entries $\mathcal{V}_\Phi$, with products and commutators weighted equally. With $\mathcal{D}_{\mathrm{pos}} = \mathcal{D}(\Phi(\widehat{X}), \Phi(Y))$, $\mathcal{D}_{\mathrm{neg}} = \mathcal{D}(\Phi(\widehat{X}), \Phi(HY))$, and separation $\Delta_Y = \mathcal{D}(\Phi(Y), \Phi(HY))$,

$$m(Y) = \min\{m_0, \tfrac{1}{2}\,\mathrm{stopgrad}(\Delta_Y)\}, \qquad \mathcal{L}_{\mathrm{neg}} = \mathrm{mean\ ReLU}(\mathcal{D}_{\mathrm{pos}} - \mathcal{D}_{\mathrm{neg}} + m(Y))\,, \tag{6}$$

where the mean runs over residue windows with $\Delta_Y \ge \epsilon_{\mathrm{neg}}$. We use $m_0 = 0.1$ and $\epsilon_{\mathrm{neg}} = 10^{-3}$. The capped margin gives zero loss at an exact prediction.

4.4 Training objective

The QPD loss is

$$\mathcal{L}_{\mathrm{QPD}} = \mathcal{L}_{\mathrm{inc}} + \mathcal{L}_{\mathrm{prod}} + 0.25\,\mathcal{L}_{\mathrm{comm}} + 0.25\,\mathcal{L}_{\mathrm{neg}} + 0.10\,\mathcal{L}_{\mathrm{cond}}, \tag{7}$$

where $\mathcal{L}_{\mathrm{inc}} = \mathrm{mean}_{\mathcal{V}_d}\,\delta_{\mathrm{rot}}(\widehat{d}, d)$ averages Equation (3) over all valid increments, $\mathcal{L}_{\mathrm{prod}}$ and $\mathcal{L}_{\mathrm{comm}}$ apply the same mean to the products and commutators of Equation (4), and $\mathcal{L}_{\mathrm{cond}}$ penalizes short or near-collinear predicted N–Cα and C–Cα vectors at valid reference frames. Valid comparisons exclude missing atoms, disconnected segments, and ill-conditioned reference frames. Loss sums and counts are aggregated across workers before normalization. $\mathcal{L}_{\mathrm{neg}}$ is the contrast of Equation (6). QPD augments the host objective,

$$\mathcal{L} = \mathcal{L}_{\mathrm{shared}} + \alpha_{\mathrm{TICA}}\mathcal{L}_{\mathrm{TICA}} + \alpha_{\mathrm{diff}}\alpha_{\mathrm{ens}}\alpha_{\mathrm{QPD}}\mathcal{L}_{\mathrm{QPD}}, \tag{8}$$

Table 1: **Protein ensembles on the ATLAS test set** (82 proteins). Bold: best per metric. QPD rows average three seeds.

	Flexibility			Distributional accuracy				Ensemble observables			
Method	RMSD $r\uparrow$	gRMSF $r\uparrow$	tRMSF $r\uparrow$	RMWD $\downarrow$	MD-$W_2\downarrow$	Jnt-$W_2\downarrow$	PC% $\uparrow$	Weak $J\uparrow$	Trans $J\uparrow$	Exp $J\uparrow$	MI $\rho\uparrow$
AlphaFlow-MD	0.48	0.60	0.85	2.61	1.52	2.25	44	0.62	0.41	0.50	0.25
ESMFlow-MD	0.19	0.31	0.76	3.60	1.51	3.19	26	0.55	0.34	0.49	0.20
ConfDiff	0.59	0.67	0.85	2.76	1.44	2.25	35	0.59	0.36	0.50	0.24
BioEmu	0.46	0.57	0.71	4.32	1.97	3.98	**51**	0.33	–	–	0.07
Str2Str	0.23	0.38	0.57	4.05	2.04	3.55	12	0.43	–	–	0.21
MDGen	0.48	0.50	0.71	2.69	1.89	–	10	0.51	–	0.29	–
EBA	0.62	0.71	0.90	2.43	1.19	2.04	44	**0.65**	0.41	**0.70**	0.36
BioMD	0.70	0.76	0.91	2.18	1.24	1.82	46	0.61	**0.46**	**0.70**	0.34
STAR-MD	0.52	0.56	0.86	2.85	1.50	2.33	35	0.52	0.37	0.57	0.24
BioKinema	**0.80**	0.82	0.90	2.25	1.33	1.82	46	0.59	0.41	0.62	0.32
QPD (ours)	0.76	**0.84**	**0.93**	**1.93**	**1.10**	**1.56**	44	0.62	0.43	0.65	**0.39**
w/o QPD in adaptation ($\alpha_{\text{QPD}}=0$)	0.75	0.83	0.92	2.01	1.14	1.65	43	0.61	0.43	0.65	**0.39**

with $\alpha_{\text{diff}}=4$, $\alpha_{\text{ens}}=1$, and $\alpha_{\text{QPD}}=0.25$. $\mathcal{L}_{\text{shared}}$ contains coordinate denoising and trajectory fluctuation terms, and $\mathcal{L}_{\text{TICA}}$ is a TICA hidden-Markov slow-state term fitted to training data, weighted by α_{TICA} (Appendix C.1). QPD acts only through predicted coordinates and adds no learnable parameters, and all data domains use the same QPD weights without per-benchmark tuning. Setting $\alpha_{\text{QPD}}=0$ keeps every other term and gives the controls of Sections 5.1 and 5.3.

5 Experiments

We evaluate protein ensembles on ATLAS (Vander Meersche et al., 2024), torsional and Markov-state dynamics on the MDGen tetrapeptides (Jing et al., 2024b), kinetics on held-out domains of CATH2-OOD20 (Feng et al., 2026a), domain-opening paths of adenylate kinase (ADK), and ligand dissociation on DD-13M (Li et al., 2025). Baselines fall into three groups: ensemble generators such as AlphaFlow, ConfDiff, BioEmu, and EBA (Jing et al., 2024a; Wang et al., 2024; Lewis et al., 2025; Lu et al., 2025), trajectory generators such as MDGen, STAR-MD, ConfRover, ProTDyn, and MarS-FM (Jing et al., 2024b; Shoghi et al., 2026; Shen et al., 2025; Liu et al., 2026; Kapuśniak et al., 2026), and all-atom protein–ligand models such as NeuralMD, BioMD, and BioKinema (Liu et al., 2025; Feng et al., 2026b;a). Baseline values are quoted from the original or benchmark papers, and baselines without published results on a benchmark are run with their official code and weights. Replicate MD, metadynamics, and the static input structure serve as references and are shaded gray. Inference settings are selected on validation data only (Appendix C.2). The physical validity of generated structures on ATLAS, MISATO (Siebenmorgen et al., 2024), and DD-13M is reported in Appendix D.6.

5.1 Protein ensembles on ATLAS

To evaluate whether order supervision helps a generator predict protein flexibility, we use ATLAS (Vander Meersche et al., 2024), which provides three 100 ns simulations for each of 1,390 protein chains. We train on 1,265 proteins, select inference settings on 39 validation proteins, and follow the AlphaFlow protocol (Jing et al., 2024a) on the 82 test proteins: we condition on the initial frame of each of three replicas, generate 100 frames at 1 ns spacing, and score 250 of the 300 generated structures per protein. Both arms start from the same QPD-trained CATH2 checkpoint and differ only in the QPD weight during ATLAS adaptation ($\alpha_{\text{QPD}}=0.25$ or 0), with three training seeds each. In Table 1, r, ρ, and J denote Pearson and Spearman correlation and Jaccard similarity, g/tRMSF global and per-target RMSF, RMWD and MD-/Jnt-W_2 the root mean and PCA W_2 distances (Å), and PC% the share of proteins with PC similarity above 0.5. Weak, Trans, and Exp J compare weak contacts, transient contacts, and exposed residues, and MI ρ correlates exposure mutual-information matrices.

As shown in Table 1, QPD gives the best value on six of eleven metrics: global and per-target RMSF correlation (0.84 and 0.93), RMWD (1.93 Å), MD-PCA and joint-PCA W_2 (1.10 and 1.56 Å), and exposure MI correlation (0.39). The $\alpha_{\text{QPD}}=0$ arm already leads the baselines on these six metrics, and QPD during adaptation widens these leads further. BioKinema, a concurrent preprint evaluated

Table 2: **Forward simulation on the MDGen tetrapeptide benchmark** (100 test peptides). Bold: best among generative models.

	JSD↓									
Method	Torsion bb	Torsion sc	Torsion all	TICA 0	TICA 0,1	MSM states	ΔG MAE↓	Flux $\rho\uparrow$	Speed-up↑	Runtime (s)↓
Replicate MD										
100 ns	0.103	0.055	0.076	0.201	0.268	0.208	0.88	0.88	1	10,800
10 ns	0.145	0.111	0.125	0.323	0.424	0.363	1.70	0.71	1	1,067
1 ns	0.212	0.261	0.240	0.432	0.568	0.493	3.00	0.42	1	107
100 ps	0.311	0.403	0.364	0.477	0.643	0.527	4.10	0.21	1	11
MDGen	0.130	0.093	0.109	0.230	0.316	0.235	1.13	0.67	**100**	**60**
MarS-FM	**0.11**	**0.09**	**0.10**	**0.21**	**0.29**	**0.22**	**0.63**	–	–	–
QPD (ours)	0.116	**0.091**	**0.102**	**0.214**	0.296	**0.224**	0.74	**0.69**	80	90

on 81 proteins, leads pairwise-RMSD correlation (0.80), BioEmu leads PC similarity, EBA leads weak-contact Jaccard, BioMD leads transient-contact Jaccard, and EBA and BioMD share the best exposed-residue Jaccard. Among trajectory generators that take the initial frame, QPD improves every reported metric over MDGen and STAR-MD, 8 of 11 over BioMD, and 9 of 11 over BioKinema.

Switching the QPD loss off during adaptation makes all seven flexibility and distributional metrics worse, for example RMWD rises to 2.01 Å and joint-PCA W_2 to 1.65 Å. Paired over proteins, the gains in per-target RMSF correlation, RMWD, and MD- and joint-PCA W_2 have 95% bootstrap intervals that exclude zero, and QPD wins on 52 to 60 of the 82 proteins for each of them (Appendix D.1). The four contact and exposure metrics differ by at most 0.002 over proteins (Table A3). Template-conditioned models, results under the MDGen protocol, and concurrent preprints are listed in Appendix D.1.

5.2 Tetrapeptide dynamics

To test dynamics on small systems with long reference simulations, we use the MDGen tetrapeptide benchmark (Jing et al., 2024b), which scores 100 test peptides against 100 ns explicit-solvent simulations. Its force field, temperature, and frame spacing differ from our other data, so we train QPD on the MDGen tetrapeptide training split. We report Jensen–Shannon divergences (JSD) of backbone (bb), side-chain (sc), and all torsions and of TICA and Markov-state distributions, the macrostate free-energy error (ΔG MAE) as defined by Kapuśniak et al. (2026), the Spearman correlation of inter-state flux, the wall-clock decorrelation speedup along TICA-0 relative to MD, and runtime. Baseline values are from Jing et al. (2024b) and Kapuśniak et al. (2026), who report the ΔG MAE of a retrained MDGen. We compute the ΔG MAE and flux of replicate MD and all speedups.

As shown in Table 2, QPD improves on MDGen in every torsion, TICA, and Markov-state JSD and lowers the macrostate free-energy error from 1.13 to 0.74. MarS-FM leads backbone-torsion and TICA-0,1 JSD and the free-energy error, QPD matches it on the other JSDs at the two decimals MarS-FM reports, and replicate 100 ns MD remains closest to the reference on every JSD. QPD is closer to the reference than 10 ns replicate MD on every JSD. In terms of dynamics, QPD reaches a flux correlation of 0.69 against 0.67 for MDGen, while MDGen samples faster, with a 100-fold decorrelation speedup and a 60 s runtime against 80-fold and 90 s for QPD. For transition paths between the two most weakly connected states of each peptide, QPD paths have a lower visited-state JSD and a higher valid-path rate than Markov models built on replicate MD of up to 100 ns, and MDGen interpolation is the best generative model (Appendix D.2).

5.3 Protein kinetics on CATH2-OOD20 and component ablation

Kinetics describe how a protein travels between metastable states, which is where the order of intermediate states matters most. CATH2-OOD20 (Feng et al., 2026a) contains 20 CATH domains held out from training and 761 initial structures, from each of which a model generates five 1 μs trajectories at 10 ns spacing, scored against a 10-state Markov model of the reference MD estimated at 100 ns lag. We report the Pearson (r) and Spearman (ρ) correlations of mean first-passage times (MFPT), the KL divergence and mean absolute error of stationary populations (π MAE), and the W_2 distance of slow coordinates. The MD reference row scores five 1 μs trajectories per domain

Table 3: **Protein kinetics on CATH2-OOD20 and component ablation.** Bold: best per metric within each block, excluding † (training set not filtered against the test domains); (b) gives the mean over three seeds.

(a) Comparison with baselines

	MFPT		Stationary		
				π	
Method	r ↑	ρ ↑	KL↓	MAE↓	W_2 ↓
MD reference	0.910	0.920	0.105	0.0230	0.140
BioKinema	0.768	0.789	**0.234**	**0.0414**	**0.225**
MDGen	0.460	0.490	0.580	0.0750	0.520
ConfRover-W	0.720	0.750	0.305	0.0500	0.280
ProTDyn†	0.825	0.835	0.215	0.0385	0.210
BioEmu†	–	–	0.275	0.0460	0.255
QPD (ours)	**0.805**	**0.811**	0.248	0.0428	0.232

(b) Component ablation

	Tetrapeptides					CATH2-OOD20	
	10 ns	20 ns	Order	Pair	TICA-0,1	MFPT	
Variant	err.↓	err.↓	gap↑	W_1 ↓	JSD↓	r ↑	KL↓
QPD (full)	**0.158**	**0.143**	**0.012**	0.470	**0.296**	**0.805**	0.248
w/o increments	0.163	0.149	0.006	0.465	0.304	0.779	0.259
w/o products	0.161	0.151	0.007	0.478	0.308	0.771	0.266
w/o commutators	0.159	0.145	0.009	0.471	0.299	0.797	0.251
w/o Hurwitz negatives	0.162	0.148	0.004	0.468	0.303	0.784	0.257
w/ plain swap	0.161	0.147	0.003	0.466	0.301	0.792	**0.246**
w/o QPD loss	0.166	0.154	0.002	**0.456**	0.312	0.763	0.271

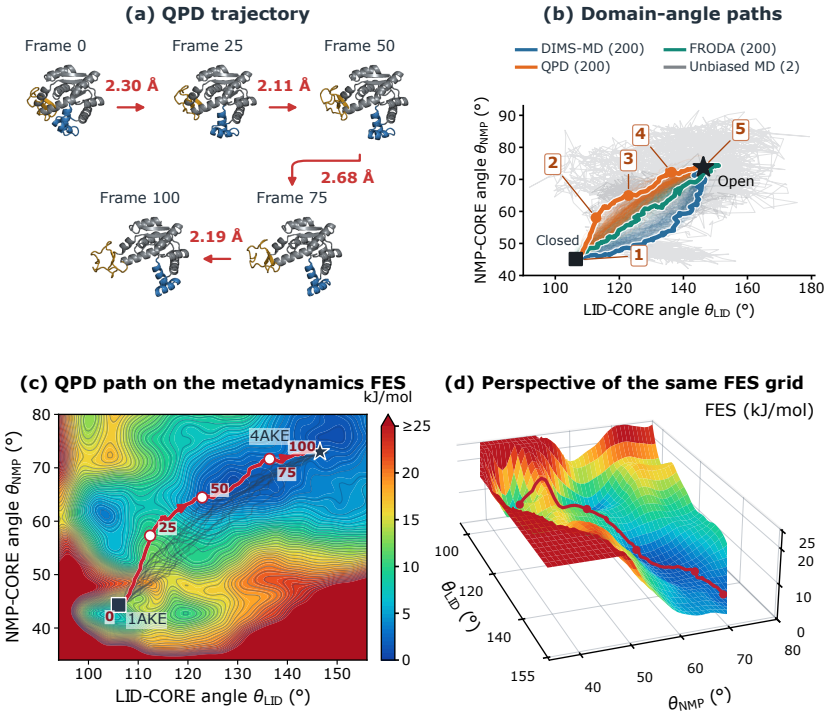

Figure 2: **ADK closed-to-open paths.** (a) Five saved all-atom frames of one QPD path, with Cα RMSD between consecutive frames. (b) Domain-angle paths, 200 per method, with two unbiased MD trajectories. (c,d) The path from (a) on the free-energy surface from 1,230 ns of well-tempered metadynamics, in 2D and 3D.

drawn from its 32 to 41 reference simulations, and ConfRover-W is the ConfRover variant defined by Shoghi et al. (2026). The CATH2 QPD model is trained from Protenix on CATH2 training domains with the same QPD weights as on the other benchmarks.

As shown in Table 3a, QPD reaches MFPT Pearson and Spearman correlations of 0.805 and 0.811, the highest among models trained with the test domains held out. ProTDyn, whose CATH2 training set is not filtered against these domains, reaches 0.825. BioEmu draws independent samples without a time axis, so its MFPTs are not defined. The same training without the QPD loss ($\alpha_{\mathrm{QPD}} = 0$) gives 0.763 (Table 3b), and the QPD loss raises it by 0.042 at the same training budget. MFPTs depend on the order of frames, and shuffling frames within each generated trajectory lowers the correlation of QPD to 0.614. Shuffling also raises the 10 and 20 ns backbone-angle correlation errors of QPD

by 0.011 on average, against 0.004 for the $\alpha_{\mathrm{QPD}} = 0$ control, a difference with a 95% interval of [0.001, 0.014] over domains (Appendix D.3). In terms of stationary statistics, BioKinema leads KL divergence, population error, and W_2 among held-out models, and QPD follows within 0.014 in KL. The number of occupied Markov states is listed in Appendix D.3.

Which parts of QPD matter. To isolate the parts of the objective, we retrain the tetrapeptide and CATH2 models with one change at a time, with the same training budget and three seeds per variant. Table 3b reports the mean over seeds, and Appendix D.3 gives the standard deviations. We remove the increment, product, or commutator term or the Hurwitz negatives of Equation (7), build negatives with a plain swap (b, a) in place of the Hurwitz move, or switch the QPD loss off ($\alpha_{\mathrm{QPD}} = 0$). On tetrapeptides we report the error of backbone-angle circular correlation at 10 and 20 ns lags against the reference MD, the order gap, Cα pair-distance W_1, and TICA-0,1 JSD. The order gap is the angular correlation error with frame order shuffled within each segment minus the error in true order, averaged over both lags. Shuffling keeps every static angle and removes chronology (Appendix B.2), so a larger gap means that accuracy depends more on the order of frames, the property an order-aware objective should raise.

As shown in Table 3b, every variant has a higher angular error at both lags, a higher TICA-0,1 JSD, and a lower MFPT correlation than full QPD, with MFPT correlations between 0.763 and 0.797. The order gap is largest for full QPD (0.012) and falls to 0.002 without the QPD loss. Among single-term removals, dropping the Hurwitz negatives lowers the gap most, to 0.004, and a plain swap, which changes the endpoint, gives 0.003. The plain swap gives the lowest KL divergence on CATH2-OOD20 (0.246 against 0.248). The $\alpha_{\mathrm{QPD}} = 0$ control has the lowest pair-distance W_1 (0.456 Å against 0.470 Å), so order supervision trades a small amount of static distance accuracy for temporal order.

5.4 Transition paths of adenylate kinase

For ADK closed-to-open, QPD generates 200 paths with endpoint-conditioned domain guidance (checkpoint in Appendix C.2), and each path is projected to valid geometry and relaxed with ff14SB/OBC2. We compare with 200 DIMS-MD and 200 FRODA paths (Seyler et al., 2015) and score every interior all-atom frame (Appendix D.4).

Figure 2 follows one QPD path through five saved all-atom frames. The 200 QPD paths advance jointly in the NMP–CORE and LID–CORE angles within the region sampled by unbiased MD, and the example path connects the closed and open basins of a metadynamics surface built from 1,230 ns of well-tempered sampling. On an independent umbrella-sampling PMF of apo ADK (Beckstein et al., 2009), the median over paths of the 95th-percentile PMF along interior frames is 8.9 kcal/mol for the guided QPD paths, against 10.4 for DIMS-MD and 10.9 for FRODA (Appendix D.4).

5.5 Ligand dissociation on DD-13M

To test whether QPD generates complete ligand dissociation events, we fine-tune on DD-13M (Li et al., 2025), which contains 26,612 metadynamics dissociation trajectories across 565 complexes, and evaluate on the 36 test complexes of BioMD (Feng et al., 2026b). Following BioMD, a trajectory succeeds when the ligand centroid leaves the convex hull of the initial protein heavy atoms, and Success@k is the fraction of complexes with at least one success among k trajectories. Path RMSD (Å) compares generated ligand-centroid paths with the closest metadynamics path. Baseline values are those of Feng et al. (2026b). As shown in Table 4, QPD reaches Success@1, @5, and @10 of 0.65, 0.85, and 0.89, above BioMD-abs at every k and below BioMD-rel, and its path RMSD of 0.63 Å lies between those of the two BioMD variants. Both BioMD variants end closer to the reference exit sites (Appendix D.5).

Table 4: **Unbinding on DD-13M** (36 complexes). Bold: best generative model.

	Success↑			Path
Method	@1	@5	@10	RMSD↓
Static	0.00	0.00	0.00	0.65
Metadynamics	–	–	–	0.42
BioMD-rel (AR-5)	**0.71**	**0.93**	**0.97**	0.71
BioMD-abs (AR-5)	0.57	0.74	0.79	**0.56**
QPD (ours)	0.65	0.85	0.89	0.63

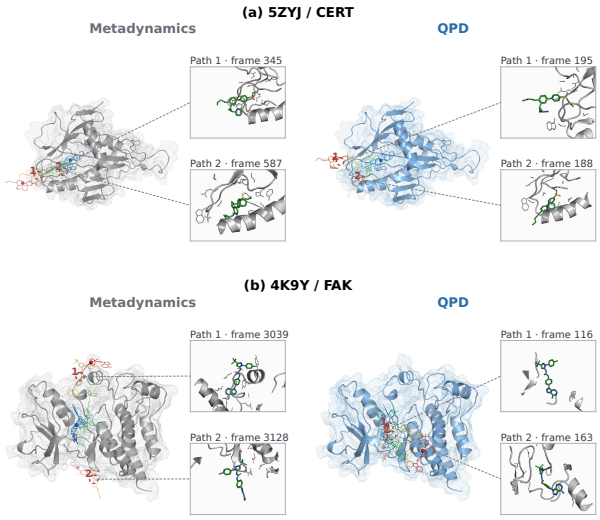

Figure 3: **Ligand dissociation paths.** (a) 5ZYJ/CERT and (b) 4K9Y/FAK, two recorded paths per method in a common protein-aligned view. Six saved ligand states per path progress from blue to red; callouts enlarge exact frames. The mesh is the initial protein surface.

The two BioMD variants mark the ends of a trade-off between success and physical quality. BioMD-rel has the highest success, 0.97 at $k = 10$, with 0.6375 protein–ligand clashes per conformation, while BioMD-abs has the fewest clashes and the lowest path RMSD at lower success (Tables 4 and 5).

QPD sits between them. It keeps most of the success of BioMD-rel with 0.0240 clashes, and its ligand bond error of 0.0350 Å is lower than that of both BioMD variants and closer to the 0.0246 Å of metadynamics.

For 5ZYJ/CERT and 4K9Y/FAK (Figure 3, checkpoint in Appendix C.2), QPD produces distinct exit routes, and four of eight paths per complex carry the ligand out of the initial protein Cα hull (Appendix D.5).

Table 5: **Ligand geometry on DD-13M.** Bold: best generative model.

Method	Bond (Å)↓	Clashes/conf.↓
Static	–	0.0000
Metadynamics	0.0246	0.0000
BioMD-rel (AR-5)	0.0580	0.6375
BioMD-abs (AR-5)	0.0728	**0.0009**
QPD (ours)	**0.0350**	0.0240

6 Discussion

Coordinate and distribution losses reward a model for visiting the right conformations. QPD additionally rewards the right sequence of local rotations, and the Hurwitz negatives remove the shortcut of matching endpoints alone. Order supervision improves most distributional metrics, on ATLAS and in TICA space on tetrapeptides, and its cost appears in static pair distances, where the $\alpha_{\mathrm{QPD}} = 0$ control is best. The QPD loss shapes how frames are sequenced. It cannot supply conformations the base model never samples, which is consistent with the lower recovery of alternative states on OOD60 (Appendix D.7) and with MarS-FM leading or matching QPD on the tetrapeptide distributions. QPD supervises rotations within generated windows at the saved-frame resolution of the training data and does not model motion between saved frames or microscopic momenta.

7 Conclusion

QPD makes the order of intermediate states a training target for all-atom trajectory generators. Rotation increments, their compositions, and Hurwitz negatives give a training signal without learnable parameters that improves protein ensembles on ATLAS, first-passage kinetics on CATH2-OOD20, tetrapeptide dynamics, and ligand dissociation with few clashes.

Reproducibility statement

The supplementary material contains our implementation of the QPD loss. Appendix C.1 lists the training settings and loss weights, Appendix C.2 maps checkpoints and data splits to benchmarks, Appendix B.2 defines the permutation protocol, and Appendix D gives additional results for each benchmark.

AI use statement

In this work, we used generative AI tools only for tasks for which the ICLR 2027 AI Policy recommends disclosure. We used them to organize our ideas, to search for and identify relevant literature, and to edit the language of the paper for readability. We did not use generative AI tools for any task that requires disclosure, including developing the conceptual framework, formulating or proving mathematical claims, proposing hypotheses, designing experiments, implementing methods, generating or processing data, analyzing data and interpreting results. The authors read every suggested reference in its original source and checked every edited passage against the intended meaning. We take responsibility for the final content of this work, including all text, claims and artifacts.

References

Oliver Beckstein, Elizabeth J. Denning, Juan R. Perilla, and Thomas B. Woolf. Zipping and unzipping of adenylate kinase: Atomistic insights into the ensemble of open ↔ closed transitions. *Journal of Molecular Biology*, 394(1):160–176, 2009. doi: 10.1016/j.jmb.2009.09.009.

Oliver Beckstein, Elizabeth J. Denning, Juan R. Perilla, and Thomas B. Woolf. Adenylate kinase potential of mean force. Dataset, 2014. URL https://zenodo.org/records/29617.

Sergio Blanes, Fernando Casas, José Angel Oteo, and José Ros. The Magnus expansion and some of its applications. *Physics Reports*, 470(5–6):151–238, 2009. doi: 10.1016/j.physrep.2008.11.001.

Avishek Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong. SE(3)-stochastic flow matching for protein backbone generation. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.02391.

Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. In *Advances in Neural Information Processing Systems*, 2024. URL https://arxiv.org/abs/2407.01392.

Xiaoting Chen, Chon-Fai Kam, Yu Li, David Medina-Ortiz, Cedric Damour, Jean Pierre Chabriat, Alain Miranville, Miloud Bessafi, and Frederic Cadet. Quaternionic response geometry for proteins: Toward a noncommutative theory of ordered deformations. *arXiv preprint arXiv:2607.29101*, 2026. URL https://arxiv.org/abs/2607.29101v1.

Kaiwen Cheng, Yutian Liu, Zhiwei Nie, Mujie Lin, Yanzhen Hou, Yiheng Tao, Chang Liu, Jie Chen, Youdong Mao, and Yonghong Tian. ProAR: Probabilistic autoregressive modeling for molecular dynamics. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 147–155, 2026. doi: 10.1609/aaai.v40i1.36974.

Allan dos Santos Costa, Ilan Mitnikov, Franco Pellegrini, Ameya Daigavane, Mario Geiger, Zhonglin Cao, Karsten Kreis, Tess Smidt, Emine Kucukbenli, and Joseph Jacobson. EquiJump: Protein dynamics simulation via SO(3)-equivariant stochastic interpolants, 2024. URL https://arxiv.org/abs/2410.09667.

Diego del Alamo, Davide Sala, Hassane S. Mchaourab, and Jens Meiler. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *eLife*, 11:e75751, 2022. doi: 10.7554/eLife.75751.

Bin Feng, Jiying Zhang, Xinni Zhang, He Cao, Ming Zhang, Patrick Barth, Zijing Liu, and Yu Li. Physically grounded generative modeling of all-atom biomolecular dynamics. *bioRxiv*, 2026a. doi:

10.64898/2026.02.15.705956. URL https://www.biorxiv.org/content/10.64898/2026.02.15.705956v2.

Bin Feng, Jiying Zhang, Xinni Zhang, Zijing Liu, and Yu Li. BioMD: All-atom generative model for biomolecular dynamics simulation. In *International Conference on Learning Representations*, 2026b. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/776f8da4b9e7007f92797ddfdf063131-Paper-Conference.pdf.

Steven Hayward and Herman J. C. Berendsen. Systematic analysis of domain motions in proteins from conformational change: New results on citrate synthase and T4 lysozyme. *Proteins: Structure, Function, and Bioinformatics*, 30(2):144–154, 1998. doi: 10.1002/(SICI)1097-0134(19980201)30:2⟨144::AID-PROT4⟩3.0.CO;2-N.

Adolf Hurwitz. Ueber Riemann'sche Flächen mit gegebenen Verzweigungspunkten. *Mathematische Annalen*, 39(1):1–60, 1891. doi: 10.1007/BF01199469.

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. doi: 10.1021/jacs.7b12191.

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*, 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. doi: 10.52202/079017-1282.

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. doi: 10.1038/s41586-021-03819-2.

Kacper Kapuśniak, Cristian Gabellini, Michael Bronstein, Prudencio Tossou, and Francesco Di Giovanni. MarS-FM: Generative modeling of molecular dynamics via Markov state models. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=jP3HnYXoIp.

Leon Klein, Andrew Y. K. Foong, Tor Erlend Fjelde, Bruno Mlodozeniec, Marc Brockschmidt, Sebastian Nowozin, Frank Noé, and Ryota Tomioka. Timewarp: Transferable acceleration of molecular dynamics by learning time-coarsened dynamics. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/a598c367280f9054434fdcc227ce4d38-Abstract-Conference.html.

Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99(20):12562–12566, 2002. doi: 10.1073/pnas.202427399.

Minji Lee, Colin Kalicki, Minkyu Jeon, Aymen Qabel, Alisia Fadini, and Mohammed AlQuraishi. ConforNets: Latents-based conformational control in OpenFold3, 2026. URL https://arxiv.org/abs/2604.18559.

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.

Maodong Li, Jiying Zhang, Zhe Wang, Bin Feng, Wenqi Zeng, Dechin Chen, Zhijun Pan, Yu Li, Zijing Liu, and Yi Isaac Yang. A novel 4-d dataset paradigm for studying complete ligand-protein dissociation dynamics, 2025. URL https://arxiv.org/abs/2504.18367.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574.

Shengchao Liu, Weitao Du, Hannan Xu, Yanjing Li, Zhuoxinran Li, Vignesh Bhethanabotla, Divin Yan, Christian Borgs, Anima Anandkumar, Hongyu Guo, and Jennifer Chayes. A multi-grained symmetric differential equation model for learning protein-ligand binding dynamics. *Nature Communications*, 17:1049, 2025. doi: 10.1038/s41467-025-67808-z.

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*, 2024.

Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Aurelie Lozano, Vijil Chenthamarakshan, Payel Das, and Jian Tang. Aligning protein conformation ensemble generation with physical feedback. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 40436–40451. PMLR, 2025. URL https://proceedings.mlr.press/v267/lu25b.html.

Terry J. Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310, 1998. doi: 10.4171/RMI/240.

Microsoft Research. BioEmu Benchmarks. Software, version 0.0.2, 2025. URL https://github.com/microsoft/bioemu-benchmarks.

Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019. doi: 10.1126/science.aaw1147.

Guillermo Pérez-Hernández, Fabian Paul, Toni Giorgino, Gianni De Fabritiis, and Frank Noé. Identification of slow molecular order parameters for Markov model construction. *The Journal of Chemical Physics*, 139(1):015102, 2013. doi: 10.1063/1.4811489.

Juan R. Perilla, Oliver Beckstein, Elizabeth J. Denning, and Thomas B. Woolf. Computing ensembles of transitions from stable states: Dynamic importance sampling. *Journal of Computational Chemistry*, 32(2):196–209, 2011. doi: 10.1002/jcc.21564.

Jan-Hendrik Prinz, Hao Wu, Marco Sarich, Bettina G. Keller, Martin Senne, Martin Held, John D. Chodera, Christof Schütte, and Frank Noé. Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics*, 134(17):174105, 2011. doi: 10.1063/1.3565032.

Protenix Team. Protenix. Software repository, 2024. URL https://github.com/bytedance/Protenix.

Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: Multiple time-resolution models for molecular dynamics. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/7274ed909a312d4d869cc328ad1c5f04-Abstract-Conference.html.

Aditya Sengar, Jiying Zhang, Pierre Vandergheynst, and Patrick Barth. Beyond ensembles: Simulating all-atom protein dynamics in a learned latent space. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://openreview.net/forum?id=AwowReRWXI. arXiv:2509.02196.

Sean L. Seyler, Avishek Kumar, M. F. Thorpe, and Oliver Beckstein. Path similarity analysis: A method for quantifying macromolecular pathways. *PLOS Computational Biology*, 11(10): e1004568, 2015. doi: 10.1371/journal.pcbi.1004568. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004568.

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*, 2025. URL https://openreview.net/forum?id=jj0nJQYFlW.

Liang Shi, Jiarui Lu, Junqi Liu, Chence Shi, Zhi Yang, and Jian Tang. Atomic trajectory modeling with state space models for biomolecular dynamics, 2026. URL https://arxiv.org/abs/2603.17633.

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://openreview.net/forum?id=Q1JpRZkR3S.

Till Siebenmorgen, Filipe Menezes, Sabrina Benassou, Erinc Merdivan, Kieran Didi, André Santos Dias Mourão, Radosław Kitel, Pietro Liò, Stefan Kesselheim, Marie Piraud, Fabian J. Theis, Michael Sattler, and Grzegorz M. Popowicz. MISATO: machine learning dataset of protein–ligand complexes for structure-based drug discovery. *Nature Computational Science*, 4(5):367–378, 2024. doi: 10.1038/s43588-024-00627-2.

Yann Vander Meersche, Gabriel Cretin, Aria Gheeraert, Jean-Christophe Gelly, and Tatiana Galochkina. ATLAS: protein flexibility description from atomistic molecular dynamics simulations. *Nucleic Acids Research*, 52(D1):D384–D392, 2024. doi: 10.1093/nar/gkad1084.

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*, 2024.

Yusong Wang, Youjun Xu, Wentao Li, Haoyu Yu, Wenjuan Tan, Shaoning Li, Qiaojing Huang, Nanjun Chen, Xuan Wu, Qilong Wu, and Kai Liu. Learning the all-atom equilibrium distribution of biomolecular interactions at scale. *bioRxiv*, 2026. doi: 10.64898/2026.03.10.710952. URL https://www.biorxiv.org/content/10.64898/2026.03.10.710952v2.

Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, et al. De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100, 2023. doi: 10.1038/s41586-023-06415-8.

Hannah K. Wayment-Steele, Adedolapo M. Ojoawo, Renee Otten, Julia M. Apitz, Warintra Pitsawong, Marc Hömberger, Sergey Ovchinnikov, Lucy J. Colwell, and Dorothee Kern. Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature*, 625(7996):832–839, 2024. doi: 10.1038/s41586-023-06832-9.

Stephen A. Wells, Scott Menor, Brandon M. Hespenheide, and M. F. Thorpe. Constrained geometric simulation of diffusive motion in proteins. *Physical Biology*, 2(4):S127–S136, 2005. doi: 10.1088/1478-3975/2/4/S07.

Christopher J. Williams, Jeffrey J. Headd, Nigel W. Moriarty, Michael G. Prisant, Lizbeth L. Videau, Lindsay N. Deis, Vishal Verma, Daniel A. Keedy, Bradley J. Hintze, Vincent B. Chen, Swati Jain, Steven M. Lewis, W. Bryan Arendall, Jack Snoeyink, Paul D. Adams, Simon C. Lovell, Jane S. Richardson, and David C. Richardson. MolProbity: More and better reference data for improved all-atom structure validation. *Protein Science*, 27(1):293–315, 2018. doi: 10.1002/pro.3330.

Kai Xu, Likun Zhao, Yanan Tian, Kewei Zhou, Jieying Zang, Jianmin Wang, Yuquan Li, Shihang Wang, Shaolong Lin, Huanxiang Liu, and Xiaojun Yao. TPS-Flow: Physics-guided flow-based generative modeling of protein transition paths. *Journal of Chemical Information and Modeling*, 66(14):7952–7971, 2026. doi: 10.1021/acs.jcim.6c00807.

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

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. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 40001–40039. PMLR, 2023. URL https://proceedings.mlr.press/v202/yim23a.html.

Shuxin Zheng, Jiyan He, Chang Liu, Yu Shi, Ziheng Lu, Weitao Feng, Fusong Ju, Jiaxi Wang, Jianwei Zhu, Yaosen Min, He Zhang, Shidi Tang, Hongxia Hao, Peiran Jin, Chi Chen, Frank Noé, Haiguang Liu, and Tie-Yan Liu. Predicting equilibrium distributions for molecular systems with deep learning. *Nature Machine Intelligence*, 6(5):558–567, 2024. doi: 10.1038/s42256-024-00837-3.

A Extended related work

This section extends Section 2 and states how QPD differs from the trajectory generators closest to it. QPD is a training objective, so each comparison separates the host model from the signal used to train it. We start with the two all-atom models that share our setting, BioKinema and BioMD, and then turn to MDGen, STAR-MD, TEMPO, and ProTDyn.

BioKinema. BioKinema (Feng et al., 2026a) generates continuous-time, all-atom trajectories with Protenix features, spatiotemporal attention whose temporal bias is derived from Langevin dynamics, and hierarchical forecasting and interpolation, trained with an EDM denoising loss and auxiliary structural and distributional terms. QPD uses a Protenix-based spatiotemporal denoiser of the same design, initialized from Protenix (Appendix C.1), and adds the QPD loss of Equation (7) to the training objective.

BioMD. BioMD (Feng et al., 2026b) introduced all-atom forecasting and interpolation within a single conditional flow-matching model with per-frame noise, trained with a velocity loss and auxiliary losses on predicted coordinates, and evaluated it on MISATO, DD-13M, and ATLAS. QPD operates in the same all-atom, per-frame-noise setting and uses DD-13M for ligand dissociation. Its added signal compares products and commutators of residue rotation increments across frames and contrasts each reference window with a Hurwitz-reordered window that has the same net rotation.

MDGen. MDGen (Jing et al., 2024b) represents a trajectory as SE(3)-invariant tokens, namely residue roto-translation offsets and torsion angles relative to conditioning key frames, and generates the token array with stochastic interpolants for forward simulation, interpolation, upsampling, and inpainting. QPD also works with residue frames. It uses the increments $q_t^{-1} \otimes q_{t+1}$ between consecutive frames, their window products, and their commutators as a training target for a model that outputs all-atom coordinates.

STAR-MD. STAR-MD (Shoghi et al., 2026) is an autoregressive SE(3) diffusion model with joint spatio-temporal attention and a key–value cache over all previously generated frames, which extends backbone trajectories to microsecond horizons. STAR-MD diffuses residue frames as the state of each new frame. QPD supervises how residue frames compose across frames inside a window denoiser that forecasts and interpolates with the same network.

TEMPO and ProTDyn. TEMPO (Xu et al., 2025) generates backbone trajectories autoregressively at two temporal scales, with a low-resolution model for slow collective motion and a high-resolution model for fast fluctuations, and ProTDyn (Liu et al., 2026) generates ensembles and trajectories autoregressively over ESM3 structure tokens with a likelihood objective. Both methods change how the generator factorizes time. QPD keeps the host factorization and changes the training signal, and it requires only predicted backbone coordinates across frames.

B Method details

This section supports Section 4. Appendix B.1 derives the properties of the QPD features and the Hurwitz negatives used there, Appendix B.2 defines the shuffled-order reference behind the order gap of Table 3b, and Appendix B.3 shows on a reference trajectory why this reference is needed.

B.1 Properties of the rotation objective

Rigid transformations and quaternion signs. For a common rigid transformation $x \mapsto Gx + c$, Equation (2) gives $R'_{t,r} = GR_{t,r}$. If g represents G, then $(gq_t)^{-1}(gq_{t+1}) = q_t^{-1}q_{t+1}$, and translation cancels in the backbone vectors. Sign changes may flip product signs, but $1 - \langle q, p\rangle^2$ is unchanged.

Endpoint and angle preservation. Associativity gives $b(b^{-1}ab) = ab$, so the Hurwitz move preserves every later accumulated orientation, while the first intermediate orientation changes from q_0a to q_0b. Since $b^{-1}ab$ has the same trace as a, conjugation preserves its rotation angle. A plain

swap (b, a) changes the endpoint when $ab \neq ba$. The separation threshold removes degenerate negatives.

Zero contrastive penalty at an exact prediction. At $\Phi(\widehat{X}) = \Phi(Y)$ we have $\mathcal{D}_{\text{pos}} = 0$ and $\mathcal{D}_{\text{neg}} = \Delta_Y$. Since $m(Y) \leq \Delta_Y/2$, the hinge argument is at most $-\Delta_Y/2$ and the loss is zero.

Reconstructing sampled orientations. Given q_0 and exact increments, $q_{t+1} = q_t d_t$ reconstructs the sampled orientation sequence up to sign. The reversed and inverted increments satisfy $d_{n-1}^{-1} \cdots d_0^{-1} = (d_0 \cdots d_{n-1})^{-1}$, so the same features describe forward and backward paths.

B.2 Temporal-permutation control

The order gap in Table 3b measures how much the angular accuracy of a model depends on the chronology of its frames. It compares each angular error with the error expected after the frames are shuffled. For angular observations $\theta_1, \ldots, \theta_N$ within one continuous segment, the expected circular correlation between two distinct positions after a uniform random permutation is

$$\mathbb{E}_\pi\left[\cos(\theta_{\pi(i)} - \theta_{\pi(j)})\right] = \frac{|\sum_{t=1}^{N} e^{i\theta_t}|^2 - N}{N(N-1)}, \quad i \neq j. \tag{9}$$

We compute the expectation in Equation (9) exactly for the reference and generated trajectories and report the absolute difference, keeping segment boundaries, lag-eligible pair counts, and equal seed weights. The order gap in Table 3b subtracts the error in true order from this shuffled-order error.

B.3 Order in a reference trajectory

A reference trajectory of the tetrapeptide FLRH shows how much of the dynamics the set of frames leaves open. All panels of Figure A1 come from one 100 ns reference simulation. Panels (a), (b), and (d) place two conformers, A and B, in the backbone TICA space and in the Leu2 torsion distribution, and panel (c) gives the autocorrelation of the first two TICA components. Panel (e) compares the joint distribution of the Leu2 ψ angle at times t and $t + \tau$, with $\tau = 0.1$ ns, for the frames in true order and after a random permutation of the same frames. The frames and every static distribution are identical, yet the lagged distributions differ. Panel (f) shows that backbone torsions lose their memory at different rates, and the ψ angles of Leu2 and Arg3 remain correlated near 1 ns while other angles approach the permuted baseline. A loss or a metric that sees only the set of frames cannot separate the two orders in panel (e), and QPD and the order gap of Appendix B.2 are built to see this difference.

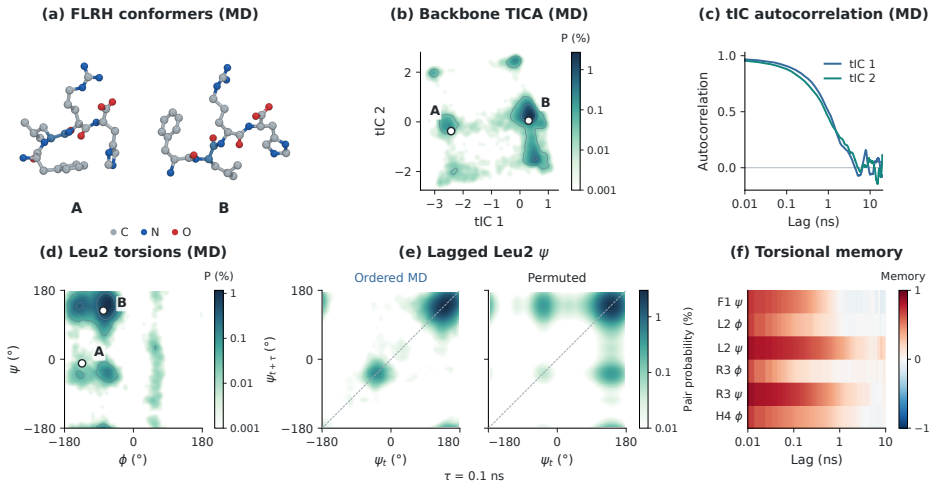

Figure A1: **Temporal structure of a reference FLRH trajectory.** All panels use the reference MD, and (e) compares true and permuted frame order at $\tau = 0.1$ ns.

C EXPERIMENTAL SETUP

This section first describes how the models are trained and then which checkpoint and data split each benchmark uses.

C.1 TRAINING SETTINGS

All models in the benchmark tables except Table A11 start from the released Protenix weights (Protenix Team, 2024). The trunk that produces (s, z) stays frozen as in Section 3, and only the denoiser D_θ is updated. The denoiser is initialized from Protenix with parameters loaded by name and shape, and the temporal modules that Protenix lacks are newly initialized. Training follows the configuration of BioKinema (Feng et al., 2026a). Because the data domains differ in force field, temperature, and frame spacing, we train them separately along three lines, Protenix→MISATO→DD-13M for protein–ligand complexes, Protenix→CATH2→ATLAS for proteins, and Protenix→tetrapeptide forward simulation→transition-path interpolation for tetrapeptides. MISATO, CATH2, and tetrapeptide forward simulation are main stages, and DD-13M, ATLAS, and tetrapeptide transition paths are adaptation stages. Each stage starts from the model of the same seed in the preceding stage, and the Initialization column of Table A1 names that stage.

Table A1 lists the data, updates, window, and temporal attention exponent β of each stage, and Table A2 lists the optimization settings shared by all stages. Each update uses 32 windows sampled with replacement, so a 100,000-update stage draws 3,200,000 windows. The training, validation, and test splits of each stage are described with its checkpoint in Appendix C.2.

Table A1: **Training stages** (updates per seed). Window is frames × frame spacing, and β is the temporal attention exponent.

Stage	Training data	Val. / test	Initialization	Updates	Window	β
MISATO	Official training split	Official val. / BioMD panel	Protenix	100,000	16 × 80 ps	0.5
CATH2	Filtered training domains	Fixed val. / OOD20	Protenix	100,000	16 × 10 ns	0.25
Tetrapeptide forward	3,109 MDGen peptides	100 / 100	Protenix	100,000	1,000 × 10 ps	0.5
Tetrapeptide TPS	Same 3,109 peptides	100 / 100	Forward	10,000	100 × 10 ps	0.5
DD-13M	476 complexes	53 / 36	MISATO	20,000	16 × 1 ps	0.5
ATLAS	1,265 proteins	39 / 82	CATH2	20,000 per arm	32 × 1 ns	0.5

Table A2: **Shared optimization settings.**

Setting	Value
Seeds	42, 43, 44
Optimizer	Adam, $\beta_1 = 0.9$, $\beta_2 = 0.95$, weight decay 10^{-8}
Learning rate, main stages	Peak 10^{-4}, 200 warmup updates, cosine decay to $0.1\times$ peak
Learning rate, adaptation stages	Peak 2×10^{-5}, 200 warmup updates, cosine decay to 2×10^{-6}
Precision	BF16
EMA decay	0.999
Gradient clipping	Global norm 10
Global batch	32 windows (8 GPUs × 1 window × 4 accumulation steps)
Diffusion noise	One noise sample per window
Cropping	Up to 256 protein tokens; complexes keep the full ligand and crop the protein around the pocket
Checkpoints	Every 1,000 updates, validated with EMA weights, no early stopping
Evaluated model	EMA weights at the end of the stage
Validation data	Selection of inference settings only

Each main stage switches the objective on gradually. The first 10,000 updates use per-frame coordinate denoising only, and the remaining 90,000 updates use the full objective of Equation (8). Over the

first 5,000 updates of the full objective, α_{QPD} rises linearly from 0 to 0.25 and then stays fixed. Within $\mathcal{L}_{\mathrm{shared}}$, the trajectory terms have weight 1 for RMSF, 8 for relative and local RMSF, and 8 for autocorrelation. The base bond term is kept, and the velocity, center, and additional ligand-bond terms are switched off. The TICA weight α_{TICA} is 0.25 on CATH2 and tetrapeptides and 0 on MISATO, DD-13M, and both ATLAS arms. In Equation (8), $\alpha_{\mathrm{diff}} = 4$ weights the host diffusion loss group and $\alpha_{\mathrm{ens}} = 1$ its ensemble subgroup, so the QPD loss enters with an effective weight of $4 \times 1 \times 0.25 = 1$.

The QPD loss uses the component weights 1, 1, 0.25, 0.25, and 0.10 of Equation (7), windows of up to four increments, margin $m_0 = 0.1$, and minimum separation $\epsilon_{\mathrm{neg}} = 10^{-3}$. Valid QPD frames require complete N–CA–C atoms and connected temporal support, and backbone vector norms and the sine of their angle must exceed 10^{-4} on the reference. An invalid predicted feature contributes 1 to its positive distance and 0 to its negative distance, so the prediction cannot evade the contrast. The implementation evaluates every valid contiguous start and builds longer commutators from a prefix and the last increment. QPD builds paths only on protein N–Cα–C frames, and ligand atoms are supervised by the host coordinate and geometry terms. The separation threshold keeps 30.2% of candidate negative windows on tetrapeptides, 52.4% on CATH2, 39.1% in ATLAS adaptation, 18.6% on MISATO, and 2.7% in DD-13M adaptation.

The frame spacing of each stage sets the physical time that a QPD window covers. DD-13M trajectories saved every 0.1 ps are subsampled to every tenth frame, giving the 1 ps spacing of Table A1. The four increments of a tetrapeptide QPD window span 40 ps, so the angular errors at 10 and 20 ns lags in Table 3b test transfer to longer times. Transition-path interpolation fixes both endpoints, taken from one continuous training trajectory, and applies the generative loss to interior frames only. Positive and negative QPD paths share these endpoints. Test endpoints and the reference Markov models follow the public protocol of Jing et al. (2024b). At inference, sampling on ATLAS uses a geometric guard that keeps each bond within its training-derived range and preserves Cα chirality relative to the initial frame.

C.2 Model checkpoints

Each benchmark table except Table A11 is evaluated with a checkpoint from one of the three training lines of Appendix C.1, and inference settings are selected on validation data only. We describe the checkpoints line by line, then the ablation variants, and last the separate checkpoint used for the ADK paths, the ligand case studies, and OOD60.

On the protein line, the CATH2 model is trained on all eligible training domains except the 20 OOD20 test domains and any training domain above 40% sequence identity to them. The ATLAS validation and test proteins are also excluded from this CATH2 training, and CATH2 inference settings are selected on a fixed validation split. ATLAS training uses all eligible proteins of the AlphaFlow training split (Jing et al., 2024a). Its training list contains 1,266 entries, and the data pipeline excludes `16pk_A`, leaving 1,265 training systems. The ATLAS models with $\alpha_{\mathrm{QPD}} = 0.25$ and $\alpha_{\mathrm{QPD}} = 0$ start from the same QPD-trained CATH2 checkpoint of each seed, so their comparison measures QPD during ATLAS adaptation, and train for 20,000 updates each with the same data, window order, and noise seeds. The $\alpha_{\mathrm{QPD}} = 0$ model keeps every other term of the objective. ATLAS inference settings are selected on the 39 validation proteins only.

On the tetrapeptide line, forward simulation and transition paths both use the MDGen split of 3,109 training, 100 validation, and 100 test peptides, each simulated for 100 ns with amber14/TIP3P-FB at 350 K and saved every 100 fs, which we subsample to 10 ps (Jing et al., 2024b). The transition-path model continues from the forward model of the same seed.

On the protein–ligand line, MISATO is trained on the official training split, validated on the official validation split (Siebenmorgen et al., 2024), and tested on the MISATO test panel of BioMD. For DD-13M, the 36 test complexes of BioMD (Feng et al., 2026b) form the test set, and the remaining 529 of the 565 complexes are split by complex into 476 training and 53 validation complexes. The MISATO training set excludes the DD-13M validation and test complexes.

The component ablations of Table 3b cover the CATH2 and tetrapeptide domains. For each domain and seed, the seven variants start from one shared state after the first 10,000 updates, copying its model, optimizer, scheduler, and EMA state, and each variant then trains for 90,000 updates with

paired window order and noise seeds. Removing a term keeps the other weights unchanged, and removing the product or commutator term also removes that feature from the distances of Equation (6). The variant without Hurwitz negatives sets the negative term to zero. The plain-swap variant keeps the number and positions of the negatives and replaces the Hurwitz move by a swap that does not preserve the net rotation. The full QPD row is the main CATH2 or tetrapeptide model. The variant without the QPD loss sets $\alpha_{\mathrm{QPD}} = 0$ and keeps every other term.

The ADK paths, the two ligand case studies, and OOD60 use a separate QPD checkpoint, adapted on ATLAS from the released BioKinema weights (Feng et al., 2026a). For the 5ZYJ/CERT case study, this checkpoint is further fine-tuned on replicas 4–18 of that complex, selected on replicas 19 and 20, and evaluated with replicas 1–3 held out.

D Additional results

Each subsection below extends one result of Section 5, in the order of the main text. The last two subsections report the physical validity of generated structures and the coverage of alternative conformations on OOD60.

D.1 Additional results on ATLAS

This subsection extends Table 1 with paired statistics of the two QPD arms and with results obtained outside its setting. Table A3 compares the arms, which share the QPD-trained CATH2 starting checkpoint and differ only in the QPD weight during adaptation. For each metric we give the paired difference between the arms with a 95% bootstrap interval over the 82 proteins, the number of proteins on which QPD wins, ties, or loses after averaging the three seeds, and the standard deviation of each arm over seeds. The gains in per-target RMSF correlation, RMWD, and MD- and joint-PCA W_2 have intervals that exclude zero, and QPD wins on 52 to 60 proteins for each of them. The two global correlations are computed over all proteins and have no per-protein outcome.

Table A3: **QPD against the $\alpha_{\mathrm{QPD}} = 0$ arm on ATLAS** (82 proteins, three seeds per arm). Δ: QPD minus $\alpha_{\mathrm{QPD}} = 0$.

Metric	Δ	95% interval	Win / tie / loss	SD, QPD	SD, $\alpha_{\mathrm{QPD}} = 0$
Pairwise RMSD r ↑	$+0.012$	$[-0.011, 0.034]$	–	0.008	0.010
Global RMSF r ↑	$+0.011$	$[-0.002, 0.023]$	–	0.005	0.006
Per-target RMSF r ↑	$+0.008$	$[0.002, 0.015]$	52 / 1 / 29	0.003	0.004
RMWD (Å)↓	-0.078	$[-0.139, -0.024]$	57 / 0 / 25	0.027	0.035
MD-PCA W_2 (Å)↓	-0.039	$[-0.074, -0.005]$	54 / 0 / 28	0.016	0.018
Joint-PCA W_2 (Å)↓	-0.095	$[-0.158, -0.035]$	60 / 0 / 22	0.031	0.040
PC-sim > 0.5 (%)↑	$+0.81$	$[-2.85, 4.47]$	7 / 70 / 5	1.22	1.86
Weak contacts J ↑	$+0.002$	$[-0.003, 0.008]$	46 / 2 / 34	0.002	0.003
Transient contacts J ↑	-0.001	$[-0.006, 0.005]$	40 / 2 / 40	0.002	0.002
Exposed residues J ↑	$+0.001$	$[-0.005, 0.007]$	44 / 3 / 35	0.002	0.002
Exposed MI ρ ↑	$+0.002$	$[-0.010, 0.014]$	43 / 0 / 39	0.004	0.005

Table A4 lists ATLAS results obtained outside the setting of Table 1, as reported by each source (Jing et al., 2024a; Wang et al., 2026; Xu et al., 2025; Feng et al., 2026a), with BioKinema and QPD repeated from Table 1. AlphaFlow-MD+Templates and AnewSampling use template structures at inference, and AnewSampling scores 1,000 frames per protein. TEMPO reports all methods trained and tested under the MDGen protocol, and its reported baseline values differ from those published by the baselines themselves (for example, 0.62 instead of 0.51 weak-contact Jaccard for MDGen). BioKinema, a concurrent preprint, is listed with the values released with its public checkpoint (81 proteins, 300 frames). Among these entries, AlphaFlow-MD+Templates has the highest pairwise-RMSD and global RMSF correlations (0.94 and 0.91), and TEMPO reports the lowest RMWD (1.49 Å) and MD-PCA W_2 (0.60 Å) under the MDGen protocol.

Table A4: **Additional ATLAS comparisons.** Metrics as in Table 1.

Method	RMSDr	gRMSFr	tRMSFr	RMWD	MD-W_2	Jnt-W_2	PC%	Weak.J	Trans.J	Exp.J	MIρ
AlphaFlow-MD+Templates	0.94	0.91	0.90	1.95	1.25	1.58	44	0.62	0.47	0.50	0.25
AnewSampling	0.85	0.89	0.93	1.88	1.16	1.44	51	0.70	0.52	0.75	0.41
TEMPO	0.91	0.89	–	1.49	0.60	–	76	0.74	0.38	–	–
BioKinema	0.80	0.82	0.90	2.25	1.33	1.82	46	0.59	0.41	0.62	0.32
QPD (ours)	0.76	0.84	0.93	1.93	1.10	1.56	44	0.62	0.43	0.65	0.39

D.2 Tetrapeptide dynamics and transition paths

This subsection gives bootstrap intervals for Table 2 and the transition-path results summarized at the end of Section 5.2. Resampling the 100 test peptides, with the training seeds kept within each peptide, gives the 95% intervals of Table A5 for the QPD row.

Table A5: **Bootstrap intervals for QPD on the MDGen tetrapeptides** (100 test peptides).

Metric	QPD	95% interval	Metric	QPD	95% interval
Torsion bb JSD↓	0.116	$[0.109, 0.124]$	MSM-state JSD↓	0.224	$[0.208, 0.242]$
Torsion sc JSD↓	0.091	$[0.080, 0.103]$	ΔG MAE↓	0.74	$[0.64, 0.86]$
Torsion all JSD↓	0.102	$[0.094, 0.110]$	Flux ρ ↑	0.69	$[0.63, 0.74]$
TICA-0 JSD↓	0.214	$[0.197, 0.233]$	Speedup↑	80	$[64, 101]$
TICA-0,1 JSD↓	0.296	$[0.277, 0.316]$	Runtime (s)↓	90	$[85, 96]$

For transition paths, we follow the protocol of MDGen (Jing et al., 2024b). It samples paths between the two most weakly connected states of each of the 100 test peptides and scores them under the reference Markov model. MDGen reports these path metrics only as a figure, so we ran its official interpolation checkpoint ourselves. The Markov-model rows are from Jing et al. (2024b, Fig. 3). As shown in Table A6, QPD paths have a lower visited-state JSD (0.23) and a higher valid-path rate (0.84) than paths from Markov models built on replicate MD of any length up to 100 ns, and a higher path probability than those built on up to 30 ns. MDGen interpolation is the best generative model on all three metrics.

Table A6: **Tetrapeptide transition paths** under the reference Markov model (100 test peptides). Bold: best among generative models.

	MSM from replicate MD of length								
Metric	2 ns	5 ns	10 ns	20 ns	30 ns	50 ns	100 ns	MDGen	QPD
Visited-state JSD↓	0.92	0.78	0.64	0.51	0.44	0.38	0.31	**0.18**	0.23
Valid-path rate↑	0.13	0.32	0.49	0.64	0.69	0.75	0.82	**0.88**	0.84
Mean path probability↑	0.0007	0.0057	0.0057	0.0066	0.0065	0.0081	0.0086	**0.0082**	0.0078

D.3 CATH2-OOD20 and component ablation

This subsection extends Table 3 with occupied Markov states, a frame-shuffling test on the generated CATH2 trajectories, and the seed variation of the component ablation. The number of occupied Markov states complements the stationary statistics of Table 3a. The reference MD subsample occupies 9.6 of the 10 Markov states on average. QPD occupies 9.3, ProTDyn 9.4, BioKinema 9.2, BioEmu 9.0, ConfRover-W 8.7, and MDGen 6.9. The BioEmu row of Table 3a uses BioEmu v1.1.

To test whether QPD uses chronology on proteins, we shuffle frames within each generated CATH2 trajectory and recompute the angular correlation errors of Appendix B.2 and the MFPTs, with the reference MD kept in true order (Table A7). Shuffling raises the 10 and 20 ns errors of QPD from 0.083 and 0.118 to 0.092 and 0.131, an average order gap of 0.011 with a 95% interval of [0.004, 0.019] over domains, against 0.004 ($[-0.001, 0.009]$) for the $\alpha_{\mathrm{QPD}} = 0$ control. The difference between the two gaps is 0.007 ($[0.001, 0.014]$). Shuffling lowers the MFPT correlation of QPD from 0.805 to 0.614 and that of the control from 0.763 to 0.602.

Table A7: **Frame shuffling on CATH2-OOD20.** Angular correlation errors and MFPT correlation of generated trajectories in true and shuffled order.

Model	Frame order	10 ns error↓	20 ns error↓	MFPT r ↑
QPD	True	0.083	0.118	0.805
	Shuffled	0.092	0.131	0.614
$\alpha_{\mathrm{QPD}} = 0$	True	0.090	0.130	0.763
	Shuffled	0.093	0.134	0.602

Table A8 repeats the component ablation of Table 3b with the standard deviation over the three training seeds. Full QPD reaches an MFPT correlation of 0.805 with a standard deviation of 0.010, against 0.763 and 0.016 without the QPD loss.

Table A8: **Component ablation with seed variation** (mean$_{\pm\mathrm{SD}}$ over three seeds).

	Tetrapeptides					CATH2-OOD20	
Variant	10 ns err.↓	20 ns err.↓	Order gap↑	Pair W_1 ↓	TICA-0,1 JSD↓	MFPT r ↑	KL↓
QPD (full)	$0.158_{\pm 0.002}$	$0.143_{\pm 0.002}$	$0.012_{\pm 0.002}$	$0.470_{\pm 0.007}$	$0.296_{\pm 0.004}$	$0.805_{\pm 0.010}$	$0.248_{\pm 0.005}$
w/o increments	$0.163_{\pm 0.003}$	$0.149_{\pm 0.003}$	$0.006_{\pm 0.002}$	$0.465_{\pm 0.007}$	$0.304_{\pm 0.004}$	$0.779_{\pm 0.013}$	$0.259_{\pm 0.006}$
w/o products	$0.161_{\pm 0.002}$	$0.151_{\pm 0.003}$	$0.007_{\pm 0.002}$	$0.478_{\pm 0.008}$	$0.308_{\pm 0.005}$	$0.771_{\pm 0.015}$	$0.266_{\pm 0.007}$
w/o commutators	$0.159_{\pm 0.002}$	$0.145_{\pm 0.003}$	$0.009_{\pm 0.002}$	$0.471_{\pm 0.007}$	$0.299_{\pm 0.004}$	$0.797_{\pm 0.011}$	$0.251_{\pm 0.005}$
w/o Hurwitz negatives	$0.162_{\pm 0.003}$	$0.148_{\pm 0.003}$	$0.004_{\pm 0.001}$	$0.468_{\pm 0.007}$	$0.303_{\pm 0.005}$	$0.784_{\pm 0.015}$	$0.257_{\pm 0.006}$
w/ plain swap	$0.161_{\pm 0.002}$	$0.147_{\pm 0.003}$	$0.003_{\pm 0.001}$	$0.466_{\pm 0.008}$	$0.301_{\pm 0.004}$	$0.792_{\pm 0.014}$	$0.246_{\pm 0.006}$
w/o QPD loss	$0.166_{\pm 0.003}$	$0.154_{\pm 0.004}$	$0.002_{\pm 0.001}$	$0.456_{\pm 0.006}$	$0.312_{\pm 0.005}$	$0.763_{\pm 0.016}$	$0.271_{\pm 0.007}$

D.4 ADK TRANSITION PATHS

This subsection extends the ADK comparison of Section 5.4 with the all-atom quality of the paths, their placement on an independent PMF, their frame-to-frame continuity, and the metadynamics sampling behind the free-energy surface of Figure 2. Figure A3 scores the interior frames of the 200 paths per method by path-level clashscore, bond and angle outliers, Ramachandran outliers, and rotamer outliers. After projection and relaxation, the median path-level clashscore is 0.33 for QPD, 10.77 for DIMS-MD, and 66.23 for FRODA. The median Ramachandran outlier rates are 0.62%, 2.97%, and 5.63%.

Figure A2 places the paths of all three methods on an independent PMF of apo ADK, computed by two-dimensional umbrella sampling and WHAM (Beckstein et al., 2009; 2014). For each path we take the 95th percentile of the PMF over its interior frames. The median over paths is 8.9 kcal/mol for QPD, 10.4 kcal/mol for DIMS-MD, and 10.9 kcal/mol for FRODA. The QPD distribution is lower as a whole and still contains individual paths with high values. The paths are projected with the Cα-based domain angles of Figure 2, whereas the PMF uses backbone and Cβ centers.

Figure A4 shows frame-to-frame continuity along one path per method as the pairwise Cα RMSD between its saved frames. The smooth band near the diagonal shows gradual progression between neighboring frames. Figure A5 shows how the eight metadynamics walkers covered the ADK angle space during the first 609 ns of the 1,230 ns run. The top panel gives the fraction of the 713 target 2° angle bins visited by all walkers, by the walkers with closed starts, and by the walkers with open starts. The bottom panel gives the RMS change between successive cumulative surfaces for each walker.

D.5 EXIT SITES AND LIGAND PATHS ON DD-13M

This subsection extends the ligand-dissociation results of Section 5.5 with the accuracy of exit sites on the 36 test complexes and with contact diagnostics for the two case studies. Of the two BioMD baselines (Feng et al., 2026b), BioMD-rel predicts coordinate changes relative to an anchor frame, BioMD-abs predicts absolute coordinates, and AR-5 marks their autoregressive variants. Following BioMD, the endpoint distance is the minimum distance between the generated ligand centroid path and the final ligand centroid of the reference simulation, and the correct rate is the fraction of trajectories whose endpoint distance falls below a threshold. Table A9 compares QPD with the values

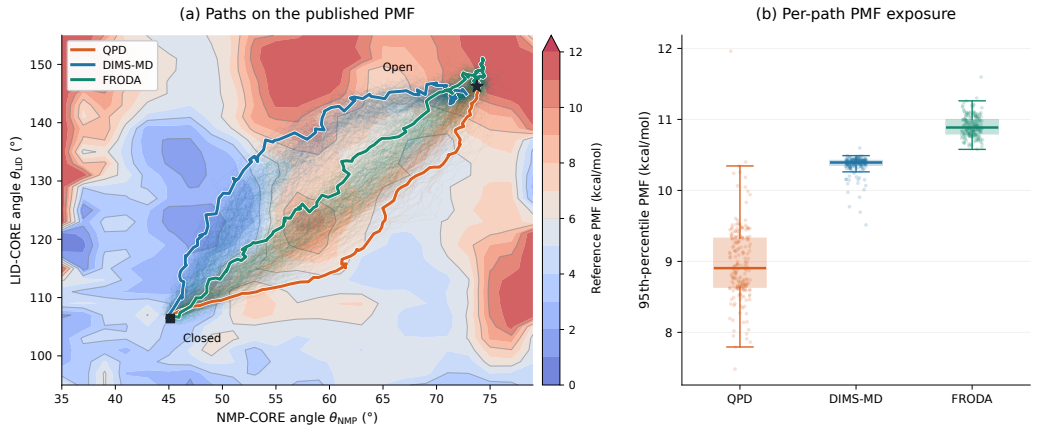

Figure A2: **ADK paths on an independent PMF.** (a) QPD, DIMS-MD, and FRODA paths on the umbrella-sampling PMF of apo ADK. (b) 95th-percentile PMF along the interior frames of each path; boxes span the quartiles.

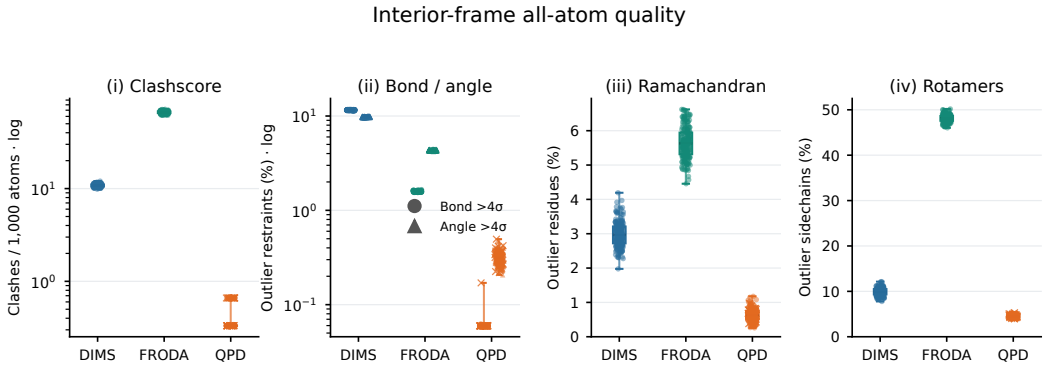

Figure A3: **All-atom quality of interior ADK frames** (200 paths per method). Each point is one path.

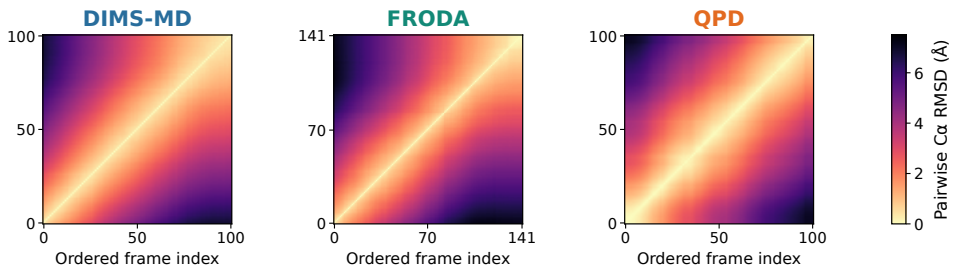

Figure A4: **Frame continuity along ADK paths.** Pairwise Cα RMSD (Å) between saved frames of one path per method.

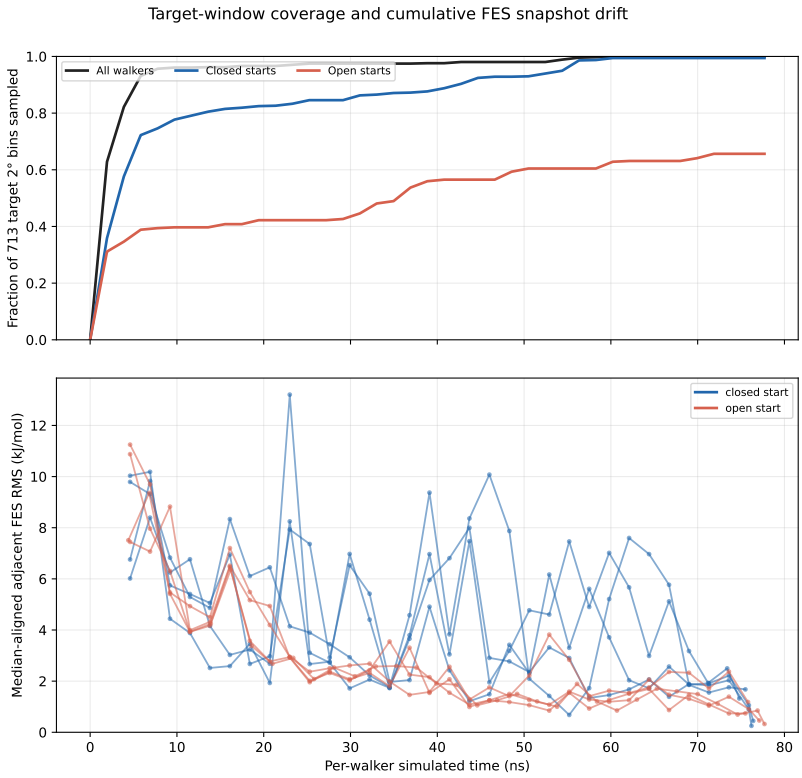

Figure A5: **Metadynamics sampling of the ADK landscape during the first 609 ns.** Top, visited fraction of the target angle bins. Bottom, RMS change (kJ/mol) between successive cumulative surfaces.

of Feng et al. (2026b) on the same 36 complexes. Both BioMD variants pass closer to the reference exit than QPD, in mean distance and at every threshold. Resampling the 36 complexes gives 95% bootstrap intervals for QPD of [0.54, 0.76], [0.75, 0.94], and [0.80, 0.97] for Success@1, @5, and @10, [0.54, 0.74] Å for path RMSD, [0.0290, 0.0420] Å for ligand bond error, and [0.0140, 0.0390] for protein–ligand clashes per conformation.

We also report the precision and recall at 2 Å defined by Feng et al. (2026a). Precision is the fraction of generated trajectories that end within 2 Å of a metadynamics exit point, and recall is the fraction of metadynamics trajectories matched by at least one generated trajectory. QPD reaches a precision of 0.82 and a recall of 0.71. BioKinema reports a precision of 0.874 and a recall of 0.777 on its own 31-system DD-13M test set, which differs from the 36 complexes used here.

Table A9: **Exit-site accuracy on DD-13M** (36 complexes). Bold: best among generative models, gray: references.

Method	Distance (Å)↓	< 0.5 Å↑	< 1 Å↑	< 2 Å↑	< 4 Å↑
Static	3.59	0.18	0.32	0.41	0.74
Metadynamics	0.97	0.51	0.75	0.90	0.96
BioMD-rel (AR-5)	1.80	0.35	0.59	0.75	**0.87**
BioMD-abs (AR-5)	**1.73**	**0.49**	**0.68**	**0.79**	0.86
QPD (ours)	2.05	0.32	0.53	0.69	0.82

For the case studies of Figure 3, QPD generates eight paths per complex for 5ZYJ/CERT, whose replicas 1–3 are held out from its fine-tuning (Appendix C.2), and for 4K9Y/FAK, without complex-specific tuning, shown with metadynamics trajectories from the same dataset. Four of the eight QPD

paths per complex leave the convex hull of the initial protein Cα atoms, and Figure A6 shows protein contacts and van der Waals overlap along these paths.

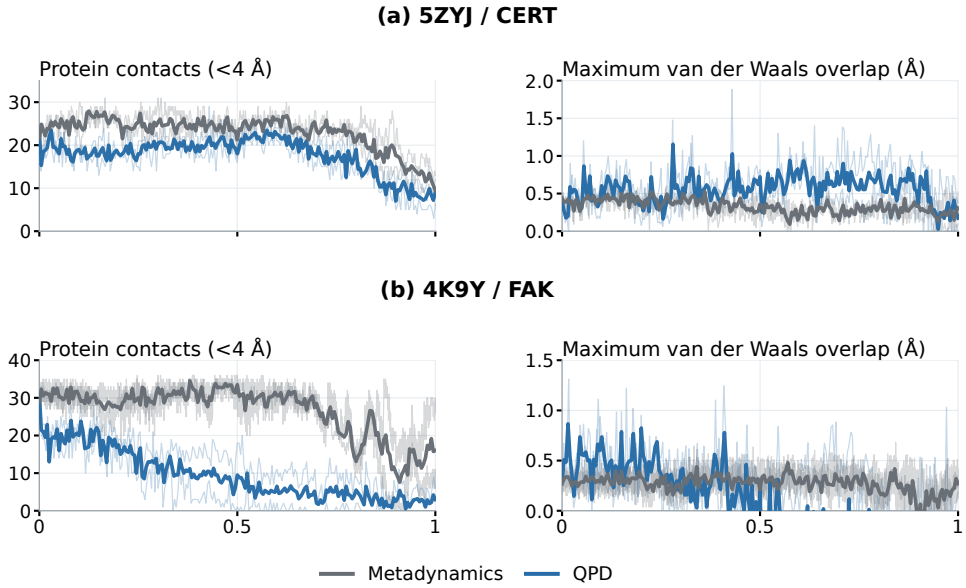

Figure A6: **Contacts and overlap along the ligand paths of Figure 3.** Protein heavy-atom contacts within 4 Å and maximum van der Waals overlap against normalized saved-frame progress.

D.6 Physical validity

This subsection gives the physical-validity comparison referred to in Section 5 and extends the ligand geometry of Table 5 with more baselines and metrics. Table A10 scores the geometry of generated structures on three test sets. On ATLAS, bond and angle errors are measured relative to the initial frame, and near contacts count frames with a close nonlocal heavy-atom pair. MolProbity, Ramachandran, and rotamer statistics follow Williams et al. (2018), with Rama and Rotamer giving the share of favored residues, and the last column gives the RMSD between a structure and its energy-minimized version. On MISATO (Siebenmorgen et al., 2024) and DD-13M, bond and angle errors and clashes are computed on ligand heavy atoms as in Feng et al. (2026b). All ATLAS rows, the BioKinema rows, and NeuralMD on DD-13M are evaluated by us with a common protocol, and the remaining MISATO and DD-13M baselines are from Feng et al. (2026b).

On ATLAS, QPD frames, sampled with the geometric guard of Appendix C.1, have a bond-length error of 0.030 Å relative to the initial frame, a MolProbity score of 1.85 against 1.45 for MD, and 98.0% Ramachandran-favored residues, level with BioKinema. AlphaFlow leads the bond, angle, contact, and minimization metrics, and BioKinema leads MolProbity and rotamers. QPD has the most near-contact frames, with ConfRover. On the MISATO test panel of BioMD, BioMD-rel leads four of five metrics, BioKinema leads angle error, and QPD has the highest bond error. On its own 35-system MISATO-OOD subset, BioKinema reports ligand bond and angle errors of 0.046 Å and 0.054 rad against 0.042 Å and 0.048 rad for MD. On DD-13M, BioKinema has the lowest ligand bond and angle errors and BioMD-abs the fewest clashes, and QPD is second in bond error.

Table A10: **Physical validity of generated structures.** Bold: best among generative models, gray: references.

(a) ATLAS

Method	Bond (Å)↓	Angle (°)↓	Contacts (%)↓	MolProbity score↓	Rama (%)↑	Rotamer (%)↑	Min. RMSD (Å)↓
MD reference	0.018	1.7	1.5	1.45	98.5	97.2	0.32
MDGen	0.030	3.4	8.0	2.15	96.2	93.8	0.82
BioKinema	0.032	2.8	6.5	**1.75**	**98.0**	**96.5**	0.70
ConfRover	0.034	3.1	8.5	1.95	97.0	95.0	0.78
AlphaFlow	**0.027**	**2.5**	**5.0**	1.80	97.5	95.8	**0.66**
QPD (ours)	0.030	2.7	8.5	1.85	**98.0**	96.0	0.71

(b) MISATO test panel of BioMD

Method	Bond MAE (Å)↓	Angle MAE (rad)↓	Intra-lig clashes↓	Prot–lig clashes↓	Min. RMSD (Å)↓
MD	0.0377	0.0575	0.0000	0.0000	0.69
BioMD-rel	**0.0395**	0.0655	**0.0003**	**0.0006**	**0.72**
BioMD-abs	0.0495	0.0709	0.0019	0.0023	–
NeuralMD-ODE	0.0483	0.0605	0.0114	0.0578	–
BioKinema	0.0480	**0.0580**	0.0010	0.0030	0.76
QPD (ours)	0.0500	0.0610	0.0014	0.0040	0.80

(c) DD-13M test set

Method	Bond MAE (Å)↓	Angle MAE (rad)↓	Intra-lig clashes↓	Prot–lig clashes↓
Metadynamics	0.0246	0.0452	0.2777	0.0000
BioMD-rel (AR-5)	0.0580	0.0918	0.4021	0.6375
BioMD-abs (AR-5)	0.0728	0.0802	**0.2943**	**0.0009**
NeuralMD	0.0500	0.0750	0.3600	0.0800
BioKinema	**0.0330**	**0.0580**	0.3000	0.0150
QPD (ours)	0.0350	0.0630	0.3200	0.0240

D.7 Conformational coverage on OOD60

This subsection gives the OOD60 result cited in Section 6. We run QPD (checkpoint in Appendix C.2) on the BioEmu-benchmarks `MULTICONF_OOD60` panel (Microsoft Research, 2025; Lewis et al., 2025), which asks whether a sequence-only start recovers alternative conformational states. Each of 22 sequence variants in 19 cases starts from an ESMFold prediction (Lin et al., 2023). We generate 40 trajectories of 100 frames per variant and score success at a 100-structure budget over 38 reference states (100 resamples). Table A11 lists the share of reference states recovered with 100 structures, with the published values of structure predictors specialized for multi-state sampling (Lee et al., 2026). QPD recovers 32.2% of the reference states, below every listed method, whose values range from 43.1% for BioEmu v1.2 to 60.7% for ConforNets-distance.

Table A11: **Conformational recovery on OOD60** (22 sequence variants). Mean±std over resamples.

Method	Recovered states@100 (%) ↑
BioEmu v1.2	43.1 ± 2.5
OF3p	45.3 ± 1.9
AFsample3	54.0 ± 3.0
ConforMix	57.7 ± 3.3
ConforNets-distance	60.7 ± 2.8
QPD (ours)	32.2 ± 2.7