

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BIoMD: ALL-ATOM GENERATIVE MODEL FOR BIOMOLECULAR DYNAMICS SIMULATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Molecular dynamics (MD) simulations are essential tools in computational chemistry and drug discovery, offering crucial insights into dynamic molecular behavior. However, their utility is significantly limited by substantial computational costs, which severely restrict accessible timescales for many biologically relevant processes. Despite the encouraging performance of existing machine learning (ML) methods, they struggle to generate extended biomolecular system trajectories, primarily due to the lack of MD datasets and the large computational demands of modeling long historical trajectories. Here, we introduce BioMD, the first all-atom generative model to simulate long-timescale protein-ligand dynamics using a hierarchical framework of forecasting and interpolation. We demonstrate the effectiveness and versatility of BioMD on the DD-13M (ligand unbinding) and MISATO datasets. For both datasets, BioMD generates highly realistic conformations, showing high physical plausibility and low reconstruction errors. Besides, BioMD successfully generates ligand unbinding paths for 97.1% of the protein-ligand systems within ten attempts, demonstrating its ability to explore critical unbinding pathways. Collectively, these results establish BioMD as a tool for simulating complex biomolecular processes, offering broad applicability for computational chemistry and drug discovery.

1 INTRODUCTION

Molecular dynamics (MD) simulations have emerged as an indispensable tool in computational chemistry and drug discovery, offering insights into the dynamic behavior of biomolecular systems. Through numerical integration of Newton’s equations of motion, MD simulations directly produce atomic trajectories that reveal the time evolution of molecular structures (Hollingsworth & Dror, 2018). These trajectories enable the exploration of conformational ensembles, optimization of small molecule structures, and identification of potential binding sites, significantly accelerating the design and development of novel therapeutics (Karplus & McCammon, 2002).

Despite their utility, traditional MD simulations face substantial computational limitations. **Although algorithms like Particle Mesh Ewald (PME) have improved scaling to approximately $O(N \log N)$** (Darden et al., 1993), these calculations remain the most computationally demanding component (Dror et al., 2012; Adcock & McCammon, 2006). Furthermore, accurately resolving high-frequency atomic vibrations necessitates extremely small time steps (on the order of femtoseconds), severely limiting the accessible simulation timescales (Shaw et al., 2010; 2009). Exploring biologically relevant processes, which often span microseconds to milliseconds, remains computationally intensive, restricting the practical application of atomistic MD to obtain trajectories.

Recently, machine learning (ML) methods have emerged as computational alternatives to molecular dynamics (MD) simulations. Key advances include models for generating protein conformation ensembles (Lewis et al., 2025) and neural network potentials trained on quantum mechanical data (Wang et al., 2024a). For biomolecular systems, AlphaFold 3 (Abramson et al., 2024) has demonstrated promising accuracy in predicting protein–ligand interactions. Despite these achievements, generating full MD trajectories for complex protein–ligand systems using ML remains a major challenge. Existing approaches tend to fall into two categories: (i) methods that can generate protein conformation ensembles but cannot produce time-resolved trajectories Jing et al. (2024b); Wang et al. (2024b), or (ii) methods that attempt trajectory modeling but struggle to capture protein–ligand interactions. For

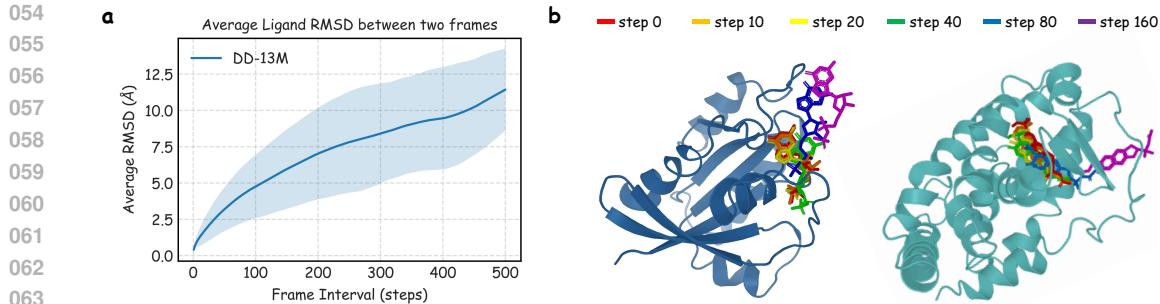


Figure 1. Average Ligand RMSD between two frames. (a) Line plot showing that the average ligand RMSD between two frames in the same trajectory increases with the frame interval. (b) Examples of ligand unbinding trajectories at time steps 0, 10, 20, 40, 80, and 160.

example, NeuralMD (Liu et al., 2024) treats protein atoms as static and only models ligand dynamics, while MDGen (Jing et al., 2024c) is specifically designed for peptides and proteins and does not handle small-molecule ligands. This limitation arises from both the complexity of protein-ligand energy landscapes and the scarcity of high-quality trajectory data for training generative models.

To address these limitations, we propose BioMD, a hierarchical framework for generating all-atom biomolecular trajectories. Building upon the insight that short-timescale conformational changes exhibit little conformational change (**Figure 1**), BioMD decomposes long trajectory generation into two synergistic stages: forecasting of large-step conformations, followed by interpolation to refine intermediate steps. This strategy reduces sequence length by decoupling long-term evolution from local dynamics and helps manage the error accumulation problem for generating long trajectories. Crucially, BioMD unifies forecasting and interpolation within a conditional flow matching model, where we use the “noising-as-masking” methods following Diffusion Forcing (Chen et al., 2024) to our time-scale transformer. We apply independent noise to each frame, which enables flexible conditioning on partial trajectory segments, and we implement different tasks simply by using different masking schedules. Inspired by the success of AlphaFold 3, BioMD generates all-atom trajectories using a velocity network that adapts its core transformer architecture, while employing an SE(3)-equivariant graph transformer to encode the initial conformation as conditional embeddings.

To evaluate the effectiveness of BioMD, we conducted experiments on two datasets: MISATO (Siebenmorgen et al., 2024) and DD-13M (Li et al., 2025). Our results show that BioMD generates highly realistic conformations with promising physical stability, evidenced by low energy and reconstruction errors across both benchmarks. On the MISATO dataset, which focuses on ligand dynamics within the binding pocket, our model accurately captures the system’s conformational flexibility, outperforming existing methods. For the more challenging task of ligand unbinding on the DD-13M dataset, BioMD successfully generates complete unbinding paths for up to 97.1% of the protein-ligand systems, demonstrating a robust ability to explore critical and long-timescale biomolecular pathways. Collectively, these results establish BioMD as a powerful and efficient tool for simulating complex biomolecular processes, offering broad applicability for computational chemistry and drug discovery.

2 RELATED WORKS

Conformational Ensemble and Binding Pose Generation. One major line of research uses ML to generate a biomolecule’s conformational ensemble by modeling the equilibrium distribution of its dynamic structures. Early efforts like AlphaFold2 (Jumper et al., 2021) produce a set of diverse conformations primarily through multiple sequence alignment (MSA) subsampling and masking techniques (Stein & Mchaourab, 2022; del Alamo et al., 2022; Wayment-Steele et al., 2024). More advanced approaches now directly learn the conformational distribution from large-scale MD datasets using flow-based (Noé et al., 2019; Jing et al., 2024b) or diffusion-based (Wang et al., 2024b; Jing et al., 2023; Lu et al., 2024a; Zheng et al., 2024; Lu et al., 2025a) generative models. Models such as BioEmu (Lewis et al., 2025) can effectively generate diverse and physically plausible conformations, providing a powerful alternative to extensive MD sampling to understand a conformational space. **Beyond proteins, recent works have extended these generative frameworks to protein-ligand systems.** DynamicBind (Lu et al., 2024b) and DynamicFlow (Zhou et al., 2025) leverage diffusion and flow

108 matching models, respectively, to recover key protein-ligand binding poses while accounting for
 109 receptor flexibility. However, these methods are fundamentally time-agnostic; they can sample what
 110 conformations are possible (or the final binding states) but lack the temporal information to show the
 111 kinetic pathways between them.
 112

113 **Trajectory Learning for MD Simulation.** To capture these kinetic pathways, a complementary
 114 research direction aims to generate full, time-ordered trajectories. Approaches like EquiJump (Costa
 115 et al., 2024) learn to sample future states based solely on the current conformation. To capture
 116 higher-order dependencies between the frames, MDGen (Jing et al., 2024c) models the joint dis-
 117 tribution of entire trajectories via masked frame modeling. ConfRover (Shen et al., 2025) models
 118 these dependencies auto-regressively by conditioning each frame on its entire history through a
 119 causal transformer. While powerful, these methods are often specialized for protein-only dynamics.
 120 Conversely, methods that model protein-ligand interactions often introduce other simplifications. For
 121 instance, NeuralMD (Liu et al., 2024) treats the protein receptor as static. These simplifications limit
 122 their scope of accessible dynamics.
 123

124 3 PRELIMINARIES

125 **Notations.** A complex \mathcal{C} is composed of a protein \mathcal{P} and a ligand ℓ . The trajectory of a complex
 126 contains $T + 1$ frames of coordinates, denoted as $\mathbf{X}_T = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T\} \in \mathbb{R}^{(T+1) \times N \times 3}$, where
 127 $\mathbf{x}_t = [\mathbf{x}_t^{\mathcal{P}}, \mathbf{x}_t^{\ell}] \in \mathbb{R}^{N \times 3}$ represents the concatenation of protein coordinates $\mathbf{x}_t^{\mathcal{P}}$ and ligand coordinates
 128 \mathbf{x}_t^{ℓ} at time-step t , and N is the number of atoms in the complex. The complex trajectory prediction
 129 task is defined as generating subsequent conformations (coordinates) of a complex trajectory given
 130 its initial conformation (i.e., the first frame).
 131

132 **Molecular dynamics.** Molecular dynamics (MD) simulates the time evolution of a particle system
 133 under classical mechanics. It leverages numerical schemes such as Verlet integration (Verlet, 1967) or
 134 Langevin dynamics to generate trajectories approximating the Boltzmann distribution. In the simplest
 135 deterministic case with no friction or noise, each particle i evolves according to $\mathrm{d}\mathbf{x}_i = \frac{p_i}{m_i} \mathrm{d}t$, $\mathrm{d}p_i =$
 136 $-\nabla_{\mathbf{x}_i} E(\mathbf{x}) \mathrm{d}t$, where p_i and m_i are the momentum and mass, and $E(\mathbf{x})$ is the potential energy
 137 function. Metadynamics (Laio & Parrinello, 2002; Barducci et al., 2011; Li et al., 2025) extends
 138 MD by introducing a history-dependent bias potential $V(s, t)$, constructed over collective variables
 139 $s(\mathbf{x})$ as $V(s, t) = \sum_{t' < t} w \exp\left(-\frac{\|s(\mathbf{x}(t)) - s(\mathbf{x}(t'))\|^2}{2\sigma^2}\right)$, where Gaussians of height w and width σ
 140 are periodically added to discourage revisiting explored states. This bias fills free-energy wells and
 141 enhances sampling of rare events and transition pathways beyond the reach of standard MD.
 142

143 **Flow matching based models.** Flow matching (FM) (Lipman et al., 2023) is an efficient and
 144 simulation-free method for training continuous normalizing flows (CNFs), a class of generative mod-
 145 els based on ordinary differential equations (ODEs). In Euclidean space, CNFs define a transformation
 146 $\phi_{\tau}(\cdot) : \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}^{N \times 3}$ via an ODE governed by a time-dependent velocity field v_{τ} :
 147

$$\frac{d}{d\tau} \phi_{\tau}(\mathbf{x}^0) = v_{\tau}(\phi_{\tau}(\mathbf{x}^0)), \quad \phi_0(\mathbf{x}^0) = \mathbf{x}^0, \quad \tau \in [0, 1], \quad (1)$$

148 where \mathbf{x}^0 is sampled from a simple distribution p_0 , and ϕ_{τ} evolves it over time $\tau \in [0, 1]$ to match
 149 the target distribution p_1 at $\tau = 1$. Since v_{τ} is unknown, FM learns v_{τ} by regressing the conditional
 150 flow $u(\phi_{\tau}(\mathbf{x}^0 | \mathbf{x}^1)) = \frac{d}{d\tau} \phi_{\tau}(\mathbf{x}^0 | \mathbf{x}^1)$, where $\phi_{\tau}(\mathbf{x}^0 | \mathbf{x}^1)$ interpolates between $\mathbf{x}^0 \sim p_0$ and $\mathbf{x}^1 \sim p_1$.
 151 In our setting, each conformation $\mathbf{x}_t \in \mathbb{R}^{N \times 3}$ represents a frame in a complex trajectory, and FM is
 152 used to generate future frames from an initial structure.
 153

154 4 BIOMD METHOD

155 4.1 A UNIFIED GENERATIVE FRAMEWORK VIA FLOW MATCHING

156 Our model capitalizes on a fundamental insight into molecular dynamics: conformational changes
 157 are typically subtle over short timescales but can involve significant global movements over longer
 158

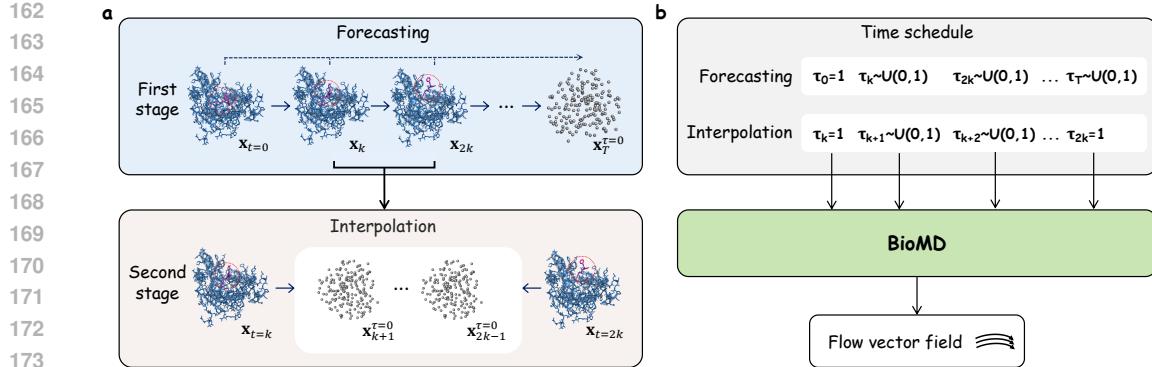


Figure 2. Model framework. (a) The hierarchical framework, showing the two-stage process of coarse-grained forecasting followed by fine-grained interpolation. (b) The time scheduling mechanism for forecasting and interpolation tasks, where known frames are noise-free ($\tau = 1$) and generated frames are noised ($\tau \in [0, 1]$).

timescales (Figure 1). This principle underpins our hierarchical prediction framework, which decomposes the generation of long trajectories into two principal stages: coarse-grained forecasting and fine-grained interpolation (Figure 2).

Notably, this entire framework is implemented within a single model architecture that processes the sequence of the whole trajectory at once. We adopt a “noise as mask” strategy, where the distinction between the two stages is made simply by varying the input masking patterns (Figure 2b). In this unified framework, each frame in an input sequence is independently perturbed by noise according to a time variable τ . Known or conditioning frames are kept clean (equivalent to setting their corresponding $\tau = 1$, i.e., “unmasked”), while frames to be generated are initialized from pure noise (equivalent to $\tau = 0$, i.e., “masked”) and then iteratively denoised.

Let a trajectory sequence be denoted by $\mathbf{X} = \{\mathbf{x}_{t_1}, \mathbf{x}_{t_2}, \dots, \mathbf{x}_{t_L}\}$. During training, we sample a vector of independent time steps $\mathbf{T} = \{\tau_{t_1}, \tau_{t_2}, \dots, \tau_{t_L}\}$, where each $\tau_{t_i} \sim U(0, 1)$. The sequence is then noised to $\mathbf{X}^T = \{\mathbf{x}_{t_1}^T, \dots, \mathbf{x}_{t_L}^T\}$, where each frame is an interpolation between the real coordinates and Gaussian noise $\mathbf{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$: $\mathbf{x}_{t_i}^T = \tau_{t_i} \mathbf{x}_{t_i} + (1 - \tau_{t_i}) \mathbf{\epsilon}_i$. The corresponding ground-truth velocity field for the sequence is $\mathbf{U}^T = \{\mathbf{u}_{t_1}^T, \dots, \mathbf{u}_{t_L}^T\}$, with $\mathbf{u}_{t_i}^T = (\mathbf{x}_{t_i} - \mathbf{x}_{t_i}^T) / (1 - \tau_{t_i})$.

Our velocity model u_θ takes the entire noisy sequence and conditioning information to predict the velocities for all frames simultaneously. The training objective is a Mean Squared Error loss over the entire sequence:

$$\mathcal{L}_{\text{flow}} = \text{MSE}(u_\theta(\mathbf{X}^T, \mathbf{Z}, \mathbf{T}), \mathbf{U}^T). \quad (2)$$

where \mathbf{Z} contains static information including the first frame coordinate \mathbf{x}_0 , amino acid sequence \mathbf{s} , and ligand atom types \mathbf{a} . We explore two modeling approaches: BioMD-rel, which predicts coordinate changes relative to an anchor frame, and BioMD-abs, which predicts absolute atomic coordinates. For clarity, we focus on the absolute coordinate prediction task below.

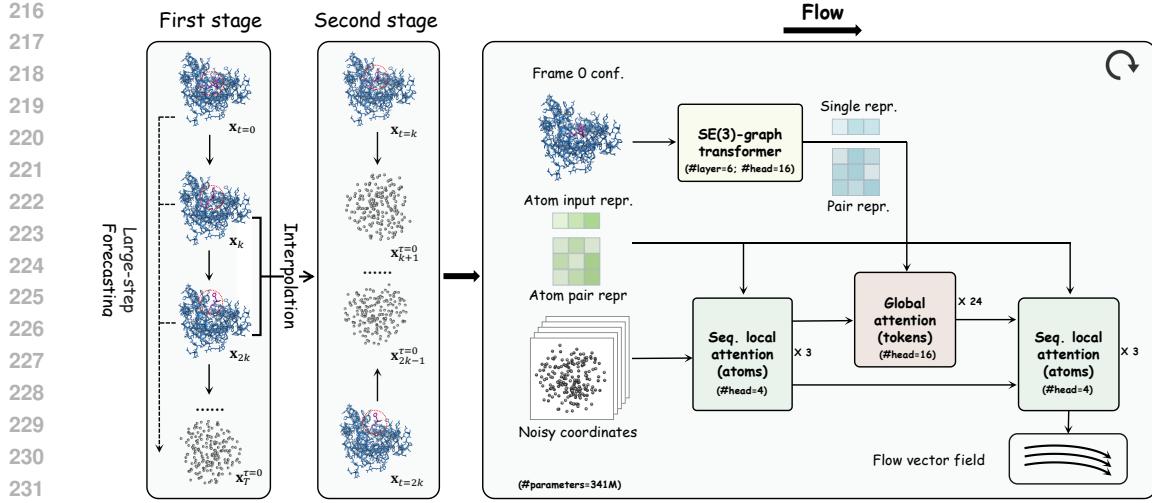
4.2 HIERARCHICAL GENERATION WITH FORECASTING AND INTERPOLATION

The two stages of our hierarchical framework are realized simply by applying different masking schedules to our unified model during training and inference.

4.2.1 COARSE-GRAINED FORECASTING

The first stage generates a coarse-grained trajectory, constructed by sampling every $k = 10$ steps (empirically chosen, see ablation study in Appendix A.5) from the full trajectory, resulting in a sequence $\mathbf{X}_C = \{\mathbf{x}_0, \mathbf{x}_k, \mathbf{x}_{2k}, \dots\}$. This task is framed as a forecasting problem where, given the initial frame \mathbf{x}_0 , the model must generate all subsequent frames.

This is achieved by applying a specific masking schedule to our unified framework. During training, the time step for the initial frame is always fixed at $\tau_0 = 1$ (making it a known, “unmasked” condition), while the time steps for all other frames $\{\tau_k, \tau_{2k}, \dots\}$ are sampled independently from $U(0, 1)$. The



Forecasting: first frame is known, other frames begin with noise;
Interpolation: first frame and last frame are known, other frames begin with noise

Figure 3. Detailed architecture of BioMD. The model operates in two modes, **Forecasting** and **Interpolation**, set up by the hierarchical framework (left). The core velocity network (right) processes noisy coordinates, conditioned on features from an SE(3)-Graph Transformer. A local-global-local attention pathway generates the final flow vector field used for trajectory generation.

model u_θ is trained to predict the velocities for all frames in the sequence, conditioned on the clean initial frame.

During inference, this setup supports multiple generation strategies:

- **All-at-once:** All future frames $\{x_k, x_{2k}, \dots\}$ are generated concurrently. We set $\tau_0 = 1$, initialize all other frames from noise (i.e., their τ values start at 0), and use an ODE solver like the Euler method to integrate all frames simultaneously to $\tau = 1$.
- **Auto-regressive (AR):** Frames are generated in sequential blocks of size j . To generate one such block, the model conditions on the previously generated history. This is controlled by the time variable τ : the τ values for all frames in the history are set to be constant 1, making them clean, “unmasked” inputs. The τ values for all j frames within the current target block are then jointly evolved from 0 to 1 by the ODE solver. This process simultaneously denoises all frames in the block, using the generated history as context. Once generated, this block is added to the history, and the process is repeated for the next block until the full trajectory is complete.

4.2.2 FINE-GRAINED INTERPOLATION

After obtaining the coarse-grained trajectory $\{x_0, x_k, x_{2k}, \dots\}$, the second stage replenishes the intermediate frames. This is an interpolation task, where for each coarse interval, we generate the frames $\{x_{ik+1}, \dots, x_{(i+1)k-1}\}$ conditioned on the two “anchor” frames, x_{ik} and $x_{(i+1)k}$.

This task uses the exact same velocity model u_θ and training framework, differing only in the data and masking schedule. The input sequence is now a fine-grained segment $\mathbf{X}_I = \{x_{ik}, x_{ik+1}, \dots, x_{(i+1)k}\}$. During training, the anchor frames are designated as known by fixing their time steps $\tau_{ik} = 1$ and $\tau_{(i+1)k} = 1$. The time steps for all intermediate frames are sampled independently from $U(0, 1)$. The model learns to generate the intermediate trajectory conditioned on the start and end conformations.

During inference, this task is always performed in an “all-at-once” manner. The anchor frames x_{ik} and $x_{(i+1)k}$ are provided as clean inputs (their $\tau = 1$), while all intermediate frames are initialized from noise (their $\tau = 0$). The model then simultaneously generates all $k - 1$ intermediate frames by integrating them to $\tau = 1$. This process is described by:

$$\hat{\mathbf{Y}}_{ik}^{\tau+\Delta\tau} = \hat{\mathbf{Y}}_{ik}^\tau + u_\theta(\hat{\mathbf{X}}_I^T, \mathbf{Z}_{\text{seq}}, \mathbf{T}) \cdot \Delta\tau, \quad (3)$$

270 where \mathbf{Y}_{ik} represents the block of intermediate frames, and the velocity predictions are extracted for
 271 only those frames. This hierarchical approach allows BioMD to efficiently generate long, physically
 272 plausible trajectories.
 273

274 4.3 VELOCITY MODEL ARCHITECTURE 275

276 BioMD is a generative model that operates directly on all-atom Cartesian coordinates. In contrast to
 277 approaches that rely on internal coordinates such as coarse-grained backbones and torsion angles, our
 278 method directly models all atoms, enabling it to capture subtle structural variations that are critical
 279 for realistic biomolecular dynamics. The effectiveness of this all-atom modeling strategy has been
 280 demonstrated by state-of-the-art biomolecular structure models like AlphaFold3 (Abramson et al.,
 281 2024). Notably, our unified model architecture is capable of performing both the forecasting and
 282 interpolation tasks (subsec. 4.2.1 and 4.2.2) within the same framework.
 283

284 Our velocity model architecture is specifically tailored for generating trajectories from a single initial
 285 structure (**Figure 3**). The model first employs an SE(3) Graph Transformer to encode the initial
 286 conformation, creating rich single and pair representations. Subsequently, our core generative module,
 287 the `FlowTrajectoryTransformer` (**Algorithm 6**), operates on the entire trajectory sequence. To
 288 effectively capture complex biomolecular dynamics, each block of this transformer incorporates
 289 two primary attention mechanisms: `AttentionPairBias` is responsible for modeling intra-
 290 frame spatial interactions, while `TemporalAttention` specifically addresses inter-frame temporal
 291 dependencies by focusing on the same atom or token across different time steps. By stacking these
 292 two attention mechanisms, the model can simultaneously process spatial and temporal information,
 293 which is crucial for accurate trajectory prediction.
 294

295 4.4 AUXILIARY LOSSES 296

297 In addition to the primary flow-matching objective, we incorporate several auxiliary losses to improve
 298 the physical plausibility of the generated structures. These losses are applied to the final predicted
 299 coordinates, which are obtained using the model’s output velocity field.
 300

- 301 • **Ligand Bond Loss:** To preserve the ligand’s local structure, we introduce a bond loss following
 302 AlphaFold 3 (Abramson et al., 2024). For each bonded atom pair in the ligand, we compute the
 303 mean squared error between the predicted inter-atomic distance and its ground-truth value, ensuring
 304 that the generated ligand structure maintains correct bond lengths.
 305
- 306 • **Collision Loss:** To ensure physical plausibility and prevent steric clashes, we implement a collision
 307 loss that applies a squared penalty to non-bonded atom pairs that are unrealistically close. This loss
 308 operates on both protein-ligand and intra-ligand interactions, and penalizes inter-atomic distance
 309 that falls below a predefined threshold.
 310
- 311 • **Ligand Geometric Center Loss:** To penalize unrealistic rigid-body movements of ligands, we
 312 define a geometric center loss. This loss calculates the mean squared error between the geometric
 313 center of the predicted ligand atoms and that of the ground-truth ligand atoms, penalizing large and
 314 unrealistic movements of the entire molecule.
 315

316 5 EXPERIMENTS 317

318 We evaluate BioMD on three MD trajectory datasets: the MISATO Dataset (Siebenmorgen et al.,
 319 2024), which comprises protein-ligand interaction trajectories focusing on ligand movement within
 320 the protein binding pocket; the DD-13M Dataset (Li et al., 2025), which contains trajectories of ligand
 321 unbinding from protein binding pockets and ultimately reaching the protein surface; **the ATLAS**
 322 **dataset (Vander Meersche et al., 2024), which contains 100 ns simulations for 1390 protein chains.**
 323 Examples of predicted trajectories can be obtained from Zenodo.¹

324 To comprehensively evaluate our model’s performance in generating all-atom biomolecular trajec-
 325 tories, we first evaluate the physical stability of the generated structures for both MISATO and
 326 DD-13M. **To enable direct comparison with methods that don’t consider protein flexibility, physical**
 327

¹<https://doi.org/10.5281/zenodo.16979768>

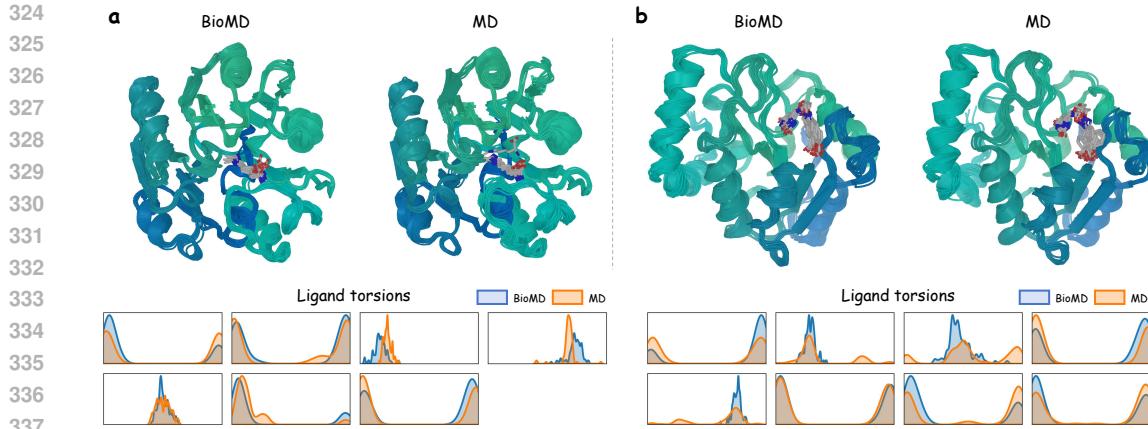


Figure 4. Conformation ensemble on the MISATO test set. A comparison of the distributions of conformations and ligand torsion angles generated by BioMD and MD simulation for 6DGE (a) and 3FCF (b).

Table 1. Results on the MISATO test set. Comparison of all methods on physical stability (first six metrics) and conformational flexibility (last four metrics). Mean values on the test samples are reported.

Method	Bond Geometry ^a		Angle Geometry ^a		Steric Clashes		RMSF Correlation ^b		RMSF Value ^{a,c}	
	MAE	MSE	MAE	MSE	Intra-Lig	Prot-Lig	Ligand	Protein	Ligand (1.211)	Protein (1.002)
Molecular Dynamics	.0377	.0023	.0575	.0053	0	0	-	-	-	-
DenoisingLD	$> 10^{10}$	$> 10^{27}$.1018	.0431	.0160	.0295	-0.0290	-	$> 10^{12}$	-
GNNMD	.2123	.1032	.2115	.1072	.3626	.0028	-0.0103	-	.2165	-
NeuralMD-ODE	.0483	.0076	.0605	.0086	.0114	.0578	.3405	-	.3220	-
NeuralMD-SDE	.0483	.0076	.0604	.0086	.0114	.0578	.3405	-	.3220	-
VerletMD	19.73	1050	.5847	.5482	.1983	3.111	.3356	-	.3226	-
BioMD-rel	.0395	.0026	.0655	.0075	.0003	.0006	.4861	.5945	.5369	.5177
BioMD-abs	.0495	.0155	.0709	.0097	.0019	.0023	.4789	.6854	.7023	.6242

^a Bond geometry (bond length) and RMSF values are in angstroms (Å). Angle geometry (bond angle) is in radians.

^b RMSF Correlation is reported using the Pearson correlation coefficient.

^c RMSF values for reference trajectories are given in parentheses. Values closer to those of the reference indicate better results.

stability metrics are computed for all heavy atoms of ligands. For the MISATO dataset, given that this dataset provides conformational ensembles, we further evaluate our model’s ability to predict the conformational flexibility of both proteins and ligands. For the DD-13M ligand unbinding dataset, we also included several metrics to assess the accuracy of the predicted unbinding pathways. For experiments on the Atlas dataset, we evaluate all methods using metrics provided by Jing et al. (2024a). In this paper, we compare BioMD with several established ML methods, including methods developed for protein-ligand MD (e.g. NeuralMD (Liu et al., 2024)) and methods developed for protein conformation sampling (e.g. ESMFLow Jing et al. (2024a)). We also include a Static model as a baseline, where the initial conformation of the system is held constant throughout the entire trajectory. For the ATLAS dataset, results of Str2Str (Lu et al., 2024a) are obtained using the released code from the official repositories with T_δ set to 0.10 and 0.15 following their instructions, and evaluation results on other methods are provided by Lu et al. (2025b).

5.1 RESULTS ON MISATO

To evaluate BioMD’s ability to generate realistic protein-ligand interaction trajectories, we first conduct experiments on the MISATO dataset, which focuses on ligand dynamics within the protein binding pocket. MISATO comprises nearly 20,000 protein-ligand interaction trajectories, each containing 100 frames sampled from an 8 ns MD simulation. We evaluate all methods on the MISATO dataset using 1,031 targets, each with a protein sequence length ≤ 800 and a ligand size of ≤ 100 heavy atoms. As shown in Table 1 and Figure 7, both variants of our model, BioMD-rel and BioMD-abs, produce trajectories with promising physical stability. The bond and angle geometry errors closely approach the values of the static input structure, and the steric clash scores are orders of magnitude lower than all competing models. Further validation on relaxation consistency (Figure 8)

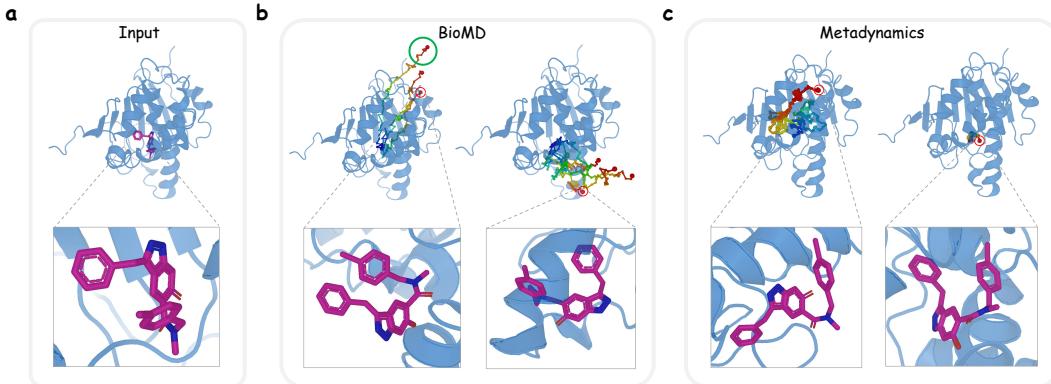


Figure 5. Ligand unbinding path on 6EY8. (a) The input conformation. (b) The unbinding pathways generated by BioMD, with an unbinding path RMSD of 0.24 Å (Metadynamics=0.12 Å). The novel pathway discovered by BioMD is highlighted in green. (c) The reference unbinding pathways obtained using metadynamics simulations.

Table 2. Results on the DD-13M test set. Comparison of methods on physical stability (first six metrics), ligand unbinding path reconstruction metric (Unbinding Path RMSD), and ligand unbinding success rates. Mean values on the test samples are reported.

Method	Bond Geometry ^a		Angle Geometry ^a		Steric Clashes		Unbinding Path ^a	Unbinding Success		
	MAE	MSE	MAE	MSE	Intra-Lig	Prot-Lig		@1	@5	@10
Static	-	-	-	-	.2778	0	.6504	0	0	0
Metadynamics ^b	.0246	.0012	.0452	.0030	.2777	0	.4217	-	-	-
BioMD-rel	.0308	.0018	.0606	.0077	.2943	.0004	.6845	.0029	.0147	.0294
BioMD-abs	.0369	.0026	.0545	.0061	.2941	.0003	.6802	.0176	.0440	.0588
BioMD-rel (AR-5)	.0580	.0100	.0918	.0184	.4021	.6375	.7055	.7088	.9295	.9706
BioMD-abs (AR-5)	.0728	.0111	.0802	.0132	.2943	.0009	.5645	.5676	.7419	.7941

^a Bond geometry (bond length) and unbinding path RMSD values are in angstroms (Å), and angle geometry (bond angle) is in radians.

^b The metadynamics trajectory serves as the lower-bound. The metrics are calculated among trajectories of multiple repeating simulations.

and pocket-ligand interactions (**Figures 9, 10**) confirms the effectiveness of BioMD to generate physically plausible structures.

In terms of conformational flexibility, BioMD demonstrates a superior ability to capture the system’s dynamic behavior. We measure Pearson’s correlation between the Root Mean Square Fluctuation (RMSF) of our generated trajectories and the reference MD trajectories. BioMD achieves the highest correlation score for ligand atoms, outperforming NeuralMD by 42.8%. Besides, BioMD achieves the correlation score of 0.685 for protein atoms, while other comparing methods fail to simulate protein conformation changes. Visual analysis in **Figures 4, 11** further corroborates these findings, showing that BioMD’s predicted atomic fluctuations closely trace the ground truth profiles and that the generated conformational ensemble is qualitatively similar to that of a traditional MD simulation. Collectively, these results indicate that BioMD can accurately simulate the flexibility of the entire protein-ligand complex.

5.2 RESULTS ON DD-13M

We further evaluate BioMD on the more challenging task of ligand unbinding using the DD-13M dataset, which comprises 26,612 dissociation trajectories across 565 complexes, each with an average of 480 frames. 36 complexes were held out as a test set for evaluation, while the remaining were used for training. **As DD-13M is generated via metadynamics, the task focuses on reproducing sampling pathways, which do not necessarily represent true thermodynamic or kinetic behavior.** A key advantage of our architecture is its flexibility in supporting multiple generation strategies. A concurrent denoising of all future frames, as used on MISATO, results in minimal ligand movement because the model lacks historical guidance and averages over many potential paths. To overcome this, we generate the trajectory auto-regressively, which breaks the long-range prediction into steps and uses previously generated frames to help predict subsequent ones.

432 **Table 3. Results on the ATLAS test set.** Performance comparison of methods for modeling protein dynamics
 433 across different metrics. Bold values indicate the best performance for each metric.

		ESMFLow-MD								
		Metrics	Full	Distilled	ConfDiff	BioEmu	Str2Str	MDGen	EBA	BioMD
Predicting flexibility	Pairwise RMSD $r \uparrow$	0.19	0.19	0.59	0.46	0.23	0.48	0.62	0.70	
	Global RMSF $r \uparrow$	0.31	0.33	0.67	0.57	0.38	0.50	0.71	0.76	
	Per-target RMSF $r \uparrow$	0.76	0.74	0.85	0.71	0.57	0.71	0.90	0.91	
Distributional accuracy	Root mean W_2 -dist. \downarrow	3.60	4.23	2.76	4.32	4.05	2.69	2.43	2.18	
	\hookrightarrow Trans. contrib. \downarrow	3.13	3.75	2.23	4.04	3.73	-	2.03	1.89	
	\hookrightarrow Var. contrib. \downarrow	1.74	1.90	1.40	1.77	1.43	-	1.20	1.10	
	MD PCA W_2 -dist. \downarrow	1.51	1.87	1.44	1.97	2.04	1.89	1.19	1.24	
	Joint PCA W_2 -dist. \downarrow	3.19	3.79	2.25	3.98	3.55	-	2.04	1.82	
Ensemble observables	% PC-sim $> 0.5 \uparrow$	26	33	35	51	12	-	44	46	
	Weak contacts $J \uparrow$	0.55	0.48	0.59	0.33	0.43	0.51	0.65	0.61	
	Transient contacts $J \uparrow$	0.34	0.30	0.36	-	-	-	0.41	0.46	
	Exposed residue $J \uparrow$	0.49	0.43	0.50	-	-	0.29	0.70	0.70	
	Exposed MI matrix $\rho \uparrow$	0.20	0.16	0.24	0.07	0.21	-	0.36	0.34	

447 The results, summarized in **Table 2**, highlight the effectiveness of this auto-regressive strategy.
 448 **While maintaining high physical stability, the BioMD-abs (AR-5) model significantly improved path**
 449 **accuracy, reducing the Unbinding Path RMSD to 0.5645 Å and achieving a high unbinding correct rate**
 450 **(Table 9).** Most importantly, the AR strategy enabled the successful generation of complete unbinding
 451 events. The BioMD-rel (AR-5) model achieved a remarkable unbinding success rate, identifying a
 452 valid path in 70.9% of cases with a single attempt (@1), increasing to 97.1% with ten attempts (@10).
 453 This demonstrates BioMD’s reliability in exploring critical biomolecular pathways.

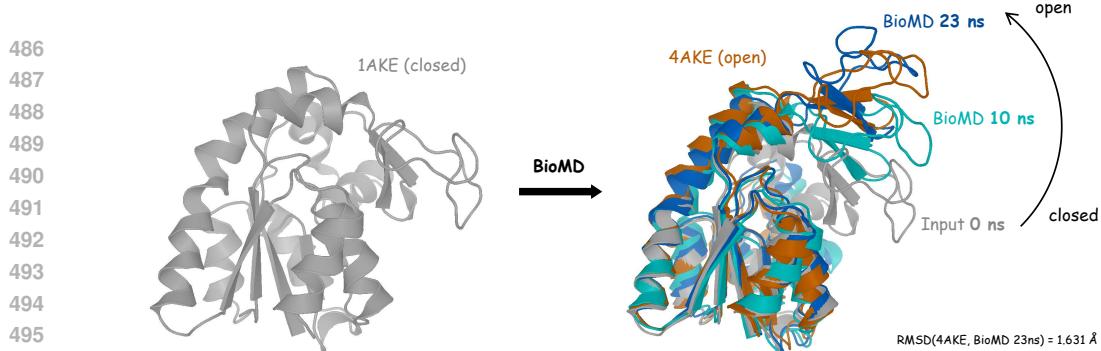
454 On the qualitative analysis for the 6EY8 system (**Figure 5**), our model not only reproduced the two
 455 distinct unbinding pathways found by metadynamics simulations with high fidelity but also discovered
 456 a novel third pathway, highlighting the exploratory power of our generative approach. Furthermore,
 457 BioMD achieves this with remarkable computational efficiency. While metadynamics required 2654
 458 steps (approx. 1 hour on a single GPU) to find the first path, our model generated a complete path in
 459 under 10 seconds using just 50 coarse-grained steps. **Additionally, in an out-of-distribution (OOD)**
 460 **case with low sequence similarity (Figure 12), BioMD successfully covers multiple escape clusters**
 461 **with an unbinding path RMSD of 0.77 Å (Metadynamics = 0.65 Å).**

462 5.3 RESULT ON ATLAS

464 We finally used the Atlas dataset (Vander Meersche et al., 2024) to evaluate the model’s ability to
 465 capture long-timescale protein dynamics with BioMD-rel. The Atlas dataset comprises 1,390 single-
 466 chain targets, each associated with three independent 100 ns Molecular Dynamics (MD) trajectories.
 467 We evaluated its performance on the Atlas test set, which comprises 82 targets. As shown in **Table 3**,
 468 BioMD achieves state-of-the-art performance on 9 out of 13 metrics, demonstrating its superiority in
 469 capturing both structural flexibility and distributional accuracy. Compared to MDGen, which shares
 470 the same input setting (sequence + initial frame), BioMD exhibits substantial improvements across
 471 all metrics. Notably, we observe a 52% increase in the Global RMSF correlation coefficient (r),
 472 highlighting the effectiveness of our all-atom architecture in modeling residue-level fluctuations.
 473 Furthermore, BioMD surpasses EBA (Lu et al., 2025a), the leading sequence-based method, on the
 474 majority of distributional metrics. This confirms that conditioning on the initial structure provides
 475 critical geometric guidance for producing high-fidelity and dynamically relevant ensembles.

476 Beyond statistical metrics, BioMD successfully reproduces functional conformational transitions. A
 477 prime example is the domain motion of Adenylate Kinase (**Figure 6**). Starting from the closed state
 478 (1AKE, 0 ns), BioMD simulates the opening process, reaching a conformation at 23 ns that closely
 479 aligns with the experimental open state (4AKE) with an RMSD of 1.631 Å. This indicates that our
 480 model can capture major protein conformational transitions that are biologically significant.

481 To further analyze the exploration capability, we projected the generated conformations onto a 2D
 482 space using Time-lagged Independent Component Analysis (TICA). As shown in **Figure 13**, BioMD
 483 explores the conformational space more sufficiently than short-timescale MD simulations (1 ns and
 484 10 ns), capturing a diversity closer to long MD trajectories (100 ns). Remarkably, BioMD generates a
 485 full 100 ns trajectory in approximately 56 seconds, offering orders-of-magnitude acceleration over
 traditional MD simulations.



496 **Figure 6. Case study on the adenylate kinase.** BioMD simulates the domain motion of the adenylate kinase
497 from close (0 ns) to open (23 ns). The test protein is not included in the training set.

500 5.4 ANALYSIS

501 The success of the auto-regressive (AR) strategy in modeling long-range dynamics simultaneously
502 exposes a fundamental challenge in generative trajectory modeling: the error accumulation problem.
503 As shown in **Table 1, 2 and Figure 7**, while the non-AR models produce local geometries with
504 errors comparable to the metadynamics reference, the AR models exhibit a notable increase in error.
505 However, thanks to our hierarchical framework, these errors remain manageable. The bond and angle
506 MAEs for our AR models remain below 0.1 Å and 0.1 radians, respectively—a threshold well within
507 the range of thermal fluctuations for molecular systems. These geometrical errors can be readily
508 corrected via a simple local refinement step with minor structural deviations (< 0.1 Å), similar to
509 the relaxation procedure used in AlphaFold. In contrast, non-hierarchical approaches are trapped
510 between two failure modes: large AR steps yield nearly static trajectories, while small AR steps cause
511 significant error accumulation that results in physically unrealistic structures.

512 Our results also reveal a distinct trade-off between predicting relative coordinate changes (BioMD-rel)
513 and absolute coordinates (BioMD-abs). The absolute coordinate prediction method (BioMD-abs)
514 demonstrates a superior grasp of the global conformational landscape, evidenced by its higher protein
515 RMSF correlation on MISATO and a more accurate centroid path RMSD on DD-13M, making it the
516 preferred choice for tasks requiring the precise reproduction of specific dynamic pathways. In contrast,
517 the relative coordinate prediction method (BioMD-rel) excels at encouraging more exploratory
518 behavior while preserving local chemical fidelity. Its strength is highlighted by the significantly
519 higher unbinding success rate on DD-13M, which makes it more effective for applications focused on
520 sampling large-scale conformational changes and discovering novel dynamic events. This functional
521 duality means BioMD can be flexibly adapted to the specific goals of a simulation, whether the
522 priority is accuracy in reproducing known dynamics or exploration to discover new ones.

523 6 CONCLUSION

524 In this work, we introduce BioMD, a novel all-atom generative model that overcomes the com-
525 putational limitations of traditional molecular dynamics to simulate long-timescale biomolecular
526 events. Our hierarchical framework, which synergistically combines coarse-grained forecasting with
527 fine-grained interpolation, effectively mitigates error accumulation and enables the generation of
528 physically realistic trajectories. We demonstrated BioMD’s capabilities on two challenging datasets,
529 showing it can produce stable conformations that accurately capture protein-ligand flexibility on the
530 MISATO dataset and successfully generate complete ligand unbinding pathways for up to 97.1% of
531 systems on the DD-13M dataset. Notably, BioMD achieves this with remarkable computational effi-
532 ciency, identifying unbinding paths in seconds compared to the hours required by traditional methods
533 like metadynamics. By offering distinct modes optimized for either accurate pathway reproduction
534 or broad exploratory sampling, BioMD provides a powerful, flexible, and efficient tool poised to
535 accelerate research in computational chemistry and drug discovery. **Nevertheless, BioMD’s ability to**
536 **generalize to substantially longer trajectories (μs or ms) or rare events beyond the training distribution**
537 **remains limited, and we highlight this as an important direction for future work.**²

538
539 ²This paper is written with assistance from large language models (LLM) for proofreading and polishing.

540 REFERENCES
541

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

545 Stewart A Adcock and J Andrew McCammon. Molecular dynamics: survey of methods for simulating
546 the activity of proteins. *Chemical reviews*, 106(5):1589–1615, 2006.

547 Alessandro Barducci, Massimiliano Bonomi, and Michele Parrinello. Metadynamics. *Wiley Interdis-*
548 *ciplinary Reviews: Computational Molecular Science*, 1(5):826–843, 2011.

549 Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann.
550 Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural*
551 *Information Processing Systems*, 37:24081–24125, 2024.

552 Allan dos Santos Costa, Ilan Mitnikov, Franco Pellegrini, Ameya Daigavane, Mario Geiger, Zhonglin
553 Cao, Karsten Kreis, Tess Smidt, Emine Kucukbenli, and Joseph Jacobson. Equijump: Protein dy-
554 namics simulation via so (3)-equivariant stochastic interpolants. *arXiv preprint arXiv:2410.09667*,
555 2024.

556 Thomas A. Darden, Darrin M. York, and Lee G. Pedersen. Particle mesh ewald: An $n \log(n)$ method
557 for ewald sums in large systems. *Journal of Chemical Physics*, 98:10089–10092, 1993. URL
558 <https://api.semanticscholar.org/CorpusID:55356282>.

559 Diego del Alamo, Davide Sala, Hassane S Mchaourab, and Jens Meiler. Sampling alternative
560 conformational states of transporters and receptors with alphafold2. *eLife*, 11:e75751, mar 2022.
561 doi: 10.7554/eLife.75751. URL <https://doi.org/10.7554/eLife.75751>.

562 Ron O Dror, Robert M Dirks, JP Grossman, Huafeng Xu, and David E Shaw. Biomolecular simulation:
563 a computational microscope for molecular biology. *Annual review of biophysics*, 41:429–452,
564 2012.

565 Scott A Hollingsworth and Ron O Dror. Molecular dynamics simulation for all. *Neuron*, 99(6):
566 1129–1143, 2018.

567 Bowen Jing, Ezra Erives, Peter Pao-Huang, Gabriele Corso, Bonnie Berger, and Tommi Jaakkola.
568 Eigenfold: Generative protein structure prediction with diffusion models. *arXiv preprint*
569 *arXiv:2304.02198*, 2023.

570 Bowen Jing, Bonnie Berger, and T. Jaakkola. Alphafold meets flow matching for generating pro-
571 tein ensembles. *ArXiv*, abs/2402.04845, 2024a. URL <https://api.semanticscholar.org/CorpusID:267522949>.

572 Bowen Jing, Bonnie Berger, and Tommi Jaakkola. Alphafold meets flow matching for generating pro-
573 tein ensembles. In *Proceedings of the 41st International Conference on Machine Learning*, pp.
574 22277–22303, 2024b.

575 Bowen Jing, Hannes Stärk, Tommi Jaakkola, and Bonnie Berger. Generative modeling of molecular
576 dynamics trajectories. *Advances in Neural Information Processing Systems*, 37:40534–40564,
577 2024c.

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

581 Martin Karplus and J Andrew McCammon. Molecular dynamics simulations of biomolecules. *Nature*
582 *structural biology*, 9(9):646–652, 2002.

583 Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the national*
584 *academy of sciences*, 99(20):12562–12566, 2002.

594 Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew YK Foong,
 595 Victor García Satorras, Osama Abdin, Bastiaan S Veeling, Iryna Zaporozhets, et al. Scalable
 596 emulation of protein equilibrium ensembles with generative deep learning. *Science*, pp. eadv9817,
 597 2025.

598 Maodong Li, Jiying Zhang, Bin Feng, Wenqi Zeng, Dechin Chen, Zhijun Pan, Yu Li, Zijing Liu,
 599 and Yi Isaac Yang. Enhanced sampling, public dataset and generative model for drug-protein
 600 dissociation dynamics. *arXiv preprint arXiv:2504.18367*, 2025.

602 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
 603 matching for generative modeling. In *ICLR*, 2023.

605 Shengchao Liu, Weitao Du, Yanjing Li, Zhuoxinran Li, Vignesh Bhethanabotla, Nakul Rampal,
 606 Omar Yaghi, Christian Borgs, Anima Anandkumar, Hongyu Guo, et al. A multi-grained sym-
 607 metric differential equation model for learning protein-ligand binding dynamics. *arXiv preprint*
 608 *arXiv:2401.15122*, 2024.

609 Jiarui Lu, Bozitao Zhong, Zuobai Zhang, and Jian Tang. Str2str: A score-based framework for
 610 zero-shot protein conformation sampling. In *The Twelfth International Conference on Learning
 611 Representations*, 2024a.

612 Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Aurelie Lozano, Vijil Chenthamarakshan, Payel Das,
 613 and Jian Tang. Aligning protein conformation ensemble generation with physical feedback. In
 614 *Forty-second International Conference on Machine Learning*, 2025a. URL <https://openreview.net/forum?id=Asr955jcuZ>.

617 Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Aurélie C. Lozano, Vijil Chenthamarakshan, Payel
 618 Das, and Jian Tang. Aligning protein conformation ensemble generation with physical feedback.
 619 *ArXiv*, abs/2505.24203, 2025b. URL <https://api.semanticscholar.org/CorpusID:279070408>.

620 Wei Lu, Jixian Zhang, Weifeng Huang, Ziqiao Zhang, Xiangyu Jia, Zhenyu Wang, Leilei Shi,
 621 Chengtao Li, Peter G. Wolynes, and Shuangjia Zheng. Dynamicbind: predicting ligand-specific
 622 protein-ligand complex structure with a deep equivariant generative model. *Nature Communica-
 623 tions*, 15, 2024b. URL <https://api.semanticscholar.org/CorpusID:267411416>.

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

627 David E Shaw, Ron O Dror, John K Salmon, JP Grossman, Kenneth M Mackenzie, Joseph A Bank,
 628 Cliff Young, Martin M Deneroff, Brannon Batson, Kevin J Bowers, et al. Millisecond-scale
 629 molecular dynamics simulations on anton. In *Proceedings of the conference on high performance
 630 computing networking, storage and analysis*, pp. 1–11, 2009.

632 David E Shaw, Paul Maragakis, Kresten Lindorff-Larsen, Stefano Piana, Ron O Dror, Michael P
 633 Eastwood, Joseph A Bank, John M Jumper, John K Salmon, Yibing Shan, et al. Atomic-level
 634 characterization of the structural dynamics of proteins. *Science*, 330(6002):341–346, 2010.

635 Yuning Shen, Lihao Wang, Huizhuo Yuan, Yan Wang, Bangji Yang, and Quanquan Gu. Simul-
 636 taneous modeling of protein conformation and dynamics via autoregression. *arXiv preprint*
 637 *arXiv:2505.17478*, 2025.

639 Till Siebenmorgen, Filipe Menezes, Sabrina Benassou, Erinc Merdivan, Kieran Didi, André San-
 640 tos Dias Mourão, Radosław Kitel, Pietro Liò, Stefan Kesselheim, Marie Piraud, et al. Misato:
 641 machine learning dataset of protein–ligand complexes for structure-based drug discovery. *Nature
 642 Computational Science*, pp. 1–12, 2024.

643 Richard A Stein and Hassane S Mchaourab. Speach_af: Sampling protein ensembles and confor-
 644 mational heterogeneity with alphafold2. *PLoS computational biology*, 18(8):e1010483, 2022.

646 Yann Vander Meersche, Gabriel Cretin, Aria Gheeraert, Jean-Christophe Gelly, and Tatiana Galochk-
 647 ina. Atlas: protein flexibility description from atomistic molecular dynamics simulations. *Nucleic
 acids research*, 52(D1):D384–D392, 2024.

648 Loup Verlet. Computer "experiments" on classical fluids. i. thermodynamical properties of lennard-
649 jones molecules. *Physical Review*, 159(1):98–103, 1967.
650

651 Tong Wang, Xinheng He, Mingyu Li, Yatao Li, Ran Bi, Yusong Wang, Chaoran Cheng, Xiangzhen
652 Shen, Jiawei Meng, He Zhang, et al. Ab initio characterization of protein molecular dynamics with
653 ai2bmd. *Nature*, 635(8040):1019–1027, 2024a.

654 Yan Wang, Lihao Wang, Yuning Shen, Yiqun Wang, Huizhuo Yuan, Yue Wu, and Quanquan Gu. Pro-
655 tein conformation generation via force-guided se (3) diffusion models. In *Forty-first International*
656 *Conference on Machine Learning*, 2024b.

657

658 Hannah K Wayment-Steele, Adedolapo Ojoawo, Renee Otten, Julia M Apitz, Warintra Pitsawong,
659 Marc Hömberger, Sergey Ovchinnikov, Lucy Colwell, and Dorothee Kern. Predicting multiple
660 conformations via sequence clustering and alphafold2. *Nature*, 625(7996):832–839, 2024.

661 Shuxin Zheng, Jiyan He, Chang Liu, Yu Shi, Ziheng Lu, Weitao Feng, Fusong Ju, Jiaxi Wang, Jianwei
662 Zhu, Yaosen Min, et al. Predicting equilibrium distributions for molecular systems with deep
663 learning. *Nature Machine Intelligence*, 6(5):558–567, 2024.

664 Xiangxin Zhou, Yi Xiao, Haowei Lin, Xinheng He, Jiaqi Guan, Yang Wang, Qiang Liu, Feng Zhou,
665 Liang Wang, and Jianzhu Ma. Integrating protein dynamics into structure-based drug design via
666 full-atom stochastic flows. *ArXiv*, abs/2503.03989, 2025. URL <https://api.semanticscholar.org/CorpusID:276813705>.

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A TECHNICAL APPENDICES AND SUPPLEMENTARY MATERIAL
703704 A.1 DETAILED MODEL ARCHITECTURE
705706 **Hierarchical Generation Framework.** As illustrated in **Figure 3**, BioMD employs a hierarchical
707 framework to perform both coarse-grained forecasting and fine-grained interpolation within a unified
708 model. The specific task is controlled by applying noise selectively. For **Forecasting**, the initial
709 frame \mathbf{x}_0 is provided without noise, while all subsequent frames are initialized from a standard
710 Gaussian distribution. For **Interpolation**, two anchor frames (e.g., \mathbf{x}_k and \mathbf{x}_{2k}) are kept clean, while
711 the intermediate frames are initialized from noise. The model’s objective is to denoise the masked
712 frames conditioned on the known ones.
713714 **Input Representation and Conditioning.** The core of the model is the **FlowModule**
715 (**Algorithm 4**), which processes three primary inputs. The main dynamic input is the set of Noisy
716 Coordinates ($\{\bar{\mathbf{x}}_l^{\text{noisy}}\}$), representing the current state of the trajectory. To provide structural con-
717 text, the initial conformation (Frame 0 conf.) is processed by an SE(3)-Graph Transformer, as
718 detailed in the main inference loop (**Algorithm 1**). This produces static Single ($\{\mathbf{s}_i^{\text{trunk}}\}$) and Pair
719 ($\{\mathbf{z}_{ij}^{\text{trunk}}\}$) representations. These representations, along with other atom features, are processed by
720 the **FlowConditioning** module (**Algorithm 5**) to generate the final conditioning signals.
721722 **Spatial-Temporal Attention Pathway.** The **FlowModule** uses a local-global-local attention
723 pathway to predict the velocity field. First, the noisy coordinates and conditioning features are passed
724 to an **AtomAttentionEncoder**, which models local atomic environments. The resulting repre-
725 sentations are aggregated into tokens and fed into the central **FlowTrajectoryTransformer**
726 (**Algorithm 6**). This module integrates spatial and temporal information using two key mechanisms:
727 **AttentionPairBias** resolves intra-frame spatial relationships, while **TemporalAttention**
728 captures inter-frame dynamics. The globally-aware token representations are then broadcast back to
729 the atomic level, where an **AtomAttentionDecoder** computes the final per-atom updates.
730731 **TemporalAttention Module.** The **TemporalAttention** module is a key component of our
732 architecture, designed to capture dynamic dependencies across the time dimension of the trajectory.
733 As shown in **Algorithm 7**, it operates on the single representation, c_s , by treating the time axis as the
734 sequence length for the attention mechanism. This allows the model to integrate information from all
735 frames of the trajectory for each residue or atom.
736737 **AttentionPairBias Module.** The **AttentionPairBias** module (**Algorithm 8**) originally
738 comes from AlphaFold. It projects the pair representation c_z into a bias term. This term is added di-
739 rectly to the attention logits before the softmax operation, effectively steering the attention mechanism
740 to focus on spatially relevant residue pairs.
741742 **Velocity Field Prediction and Trajectory Generation.** The output of the **FlowModule** is the
743 Flow vector field ($\{\bar{\mathbf{u}}_l\}$), which represents the predicted velocity for each atom. During training
744 (**Algorithm 2**), the model is optimized via a mean squared error loss between the predicted velocity
745 and the true velocity. During inference (**Algorithm 3**), this vector field is used in an Euler integration
746 step, $\bar{\mathbf{x}}_l^{\tau+1} \leftarrow \bar{\mathbf{x}}_l^{\tau} + dt \cdot \bar{\mathbf{u}}_l^{\tau}$, to iteratively update the coordinates from a noisy state to a final trajectory.
747748 A.2 AUXILIARY LOSSES
749750 After we get the estimated vector field \mathbf{u}_θ , we can get the predicted structure coordinates via
751

752
$$\hat{\mathbf{x}}_t^1 = \hat{\mathbf{x}}_t^\tau + \mathbf{u}_\theta(1 - \tau), \quad (4)$$

753 and then we get the predicted protein and ligand structure $[\hat{\mathbf{x}}_t^P, \hat{\mathbf{x}}_t^\ell] = \hat{\mathbf{x}}_t^1$.
754

755 **Ligand geometric center loss.** To stabilize the global placement of the ligand and prevent spurious
756 rigid translations, we align the predicted and reference geometric centers of ligand atoms. Let
757 $\mathbf{x}_t^\ell = \{\mathbf{x}_t^{\ell,i}\}_{i=1}^{N_\ell}$ and $\hat{\mathbf{x}}_t^\ell = \{\hat{\mathbf{x}}_t^{\ell,i}\}_{i=1}^{N_\ell}$ denote ground-truth and predicted ligand coordinates at step t .
758

756 The geometric center is
 757

$$758 C(\mathbf{x}_t^\ell) = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \mathbf{x}_t^{\ell,i}, \quad C(\hat{\mathbf{x}}_t^{\ell,i}) = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \hat{\mathbf{x}}_t^{\ell,i},$$

760 and the loss is the mean-squared discrepancy
 761

$$762 \mathcal{L}_{\text{center}} = \|C(\hat{\mathbf{x}}_t^\ell) - C(\mathbf{x}_t^\ell)\|_2^2.$$

763 This term softly anchors the ligand’s global position while remaining agnostic to its internal geometry.
 764

765 **Collision loss.** To penalize steric clashes, we define a collision loss between protein–ligand atoms
 766 and within ligand atoms. Let \mathbf{x}_t^ℓ and \mathbf{x}_t^P denote ligand and protein atom coordinates at step t , and $\hat{\mathbf{x}}_t^\ell$,
 767 $\hat{\mathbf{x}}_t^P$ their predictions. We compute predicted distances
 768

$$769 d_{ij}^{PL} = \|\hat{\mathbf{x}}_t^{\mathcal{P},i} - \hat{\mathbf{x}}_t^{\ell,j}\|_2, \quad d_{ij}^L = \|\hat{\mathbf{x}}_t^{\ell,i} - \hat{\mathbf{x}}_t^{\ell,j}\|_2,$$

770 and corresponding ground-truth minimal distances
 771

$$772 d_{ij}^{PL,gt} = \min_t \|\mathbf{x}_t^{\mathcal{P},i} - \mathbf{x}_t^{\ell,j}\|_2, \quad d_{ij}^{LL,gt} = \min_t \|\mathbf{x}_t^{\ell,i} - \mathbf{x}_t^{\ell,j}\|_2.$$

773 Protein–ligand and ligand–ligand thresholds are set as
 774

$$775 \zeta_{ij}^{PL} = \min(0.9 d_{ij}^{PL,gt}, \zeta_{pl}) \quad \zeta_{ij}^{LL} = \min(0.9 d_{ij}^{LL,gt}, \zeta_{ll}),$$

776 where $\zeta_{pl} = 3.0 \text{ \AA}$ and $\zeta_{ll} = 2.0 \text{ \AA}$.
 777

778 The collision loss is then defined as
 779

$$780 \mathcal{L}_{\text{collision}} = \sum_{i,j} \mathbf{1}(d_{ij}^{PL} < \zeta_{ij}^{PL}) (\zeta_{ij}^{PL} - d_{ij}^{PL})^2 + \sum_{i \neq j} \mathbf{1}(d_{ij}^{LL} < \zeta_{ij}^{LL}) (1 - b_{ij}) (\zeta_{ij}^{LL} - d_{ij}^{LL})^2,$$

782 where $\mathbf{1}(\cdot)$ represents the indicator function and b_{ij} is the ligand bond mask to exclude bonded pairs.
 783

784 **Ligand bond loss.** To preserve ligand bond lengths, we penalize deviations between predicted and
 785 ground-truth bonded atom distances. Let \mathcal{B} denote the set of bonded atom pairs according to the
 786 ligand bond mask. For each bond $(i,j) \in \mathcal{B}$, we compute the predicted and ground-truth distances
 787

$$788 d_{ij}^\ell = \|\hat{\mathbf{x}}_t^{\ell,i} - \hat{\mathbf{x}}_t^{\ell,j}\|_2, \quad d_{ij}^{\ell,gt} = \|\mathbf{x}_t^{\ell,i} - \mathbf{x}_t^{\ell,j}\|_2.$$

789 The bond loss is then defined as the mean squared deviation:
 790

$$791 \mathcal{L}_{\text{bond}} = \frac{1}{|\mathcal{B}|} \sum_{(i,j) \in \mathcal{B}} (d_{ij}^\ell - d_{ij}^{\ell,gt})^2.$$

793 **Geometric constraint loss.** We combine the above terms into a single geometric regularizer
 794

$$795 \mathcal{L}_{\text{geom}} = \lambda_{\text{col}} \mathcal{L}_{\text{collision}} + \lambda_{\text{bond}} \mathcal{L}_{\text{bond}} + \lambda_{\text{ctr}} \mathcal{L}_{\text{center}},$$

796 where $\lambda_{\text{col}}, \lambda_{\text{bond}}, \lambda_{\text{ctr}} > 0$ balance steric clash avoidance, bond-length preservation, and global
 797 ligand anchoring, respectively.
 798

799 A.3 EVALUATION METRICS

800 A.3.1 PHYSICAL STABILITY

802 This metric assesses whether the generated trajectories preserve physically stable conformations,
 803 which is essential to ensure chemical validity and avoid unrealistic molecular structures. We evaluate
 804 stability from two complementary perspectives:
 805

- 806 **1. Local Structure Stability.** To assess whether the generated trajectories maintain chemically
 807 reasonable local geometries, we calculate the deviations of bond lengths and bond angles with
 808 respect to the initial frame of the reference trajectories. Both the Mean Absolute Error (MAE) and
 809 Mean Squared Error (MSE) are reported. Lower values indicate that the generated conformations
 810 remain close to the idealized covalent structure and are thus more chemically stable.

810
 811 2. **Steric Clashes.** We further quantify the presence of steric conflicts, which occur when non-bonded
 812 atoms are unrealistically close to each other. Specifically, a clash is counted if the interatomic
 813 distance (excluding bonded pairs and angle-related atoms) is less than a threshold of 1.5 Å.
 814 We compute clash scores for both intra-ligand and protein–ligand interactions, where the score
 815 corresponds to the average number of clashes per generated conformation. Lower clash scores
 816 indicate physically more plausible conformations.

817 **A.3.2 CONFORMATIONAL FLEXIBILITY**

818 In addition to stability, it is important that generated trajectories capture the dynamic flexibility of
 819 molecular systems. For the MISATO protein–ligand interaction dataset, we adopt the Root Mean
 820 Square Fluctuation (RMSF) to quantify the extent of atomic motion over time after trajectory
 821 alignment:

$$823 \quad 824 \quad 825 \quad 826 \quad \text{RMSF}_i = \sqrt{\frac{1}{T} \sum_{t=1}^T \|\mathbf{r}_i(t) - \bar{\mathbf{r}}_i\|^2},$$

827 where $\mathbf{r}_i(t)$ is the position of atom i at time t , and $\bar{\mathbf{r}}_i$ is its time-averaged position.

828 We evaluate flexibility from two perspectives: 1. **Global Consistency.** We compute the Pearson
 829 correlation coefficient between the RMSFs of generated and reference trajectories, where higher
 830 correlation indicates better agreement in the fluctuation profiles. 2. **Magnitude Accuracy.** We also
 831 report the average RMSF of the generated trajectories. Values closer to the reference average RMSF
 832 imply that the model produces realistic levels of conformational motion rather than being overly rigid
 833 or excessively flexible.

834 **A.3.3 UNBINDING PATH DISTANCE**

835 For the DD-13M ligand unbinding dataset, we evaluate whether generated unbinding trajectories
 836 follow realistic spatial pathways compared to reference simulations. We compute the Root Mean
 837 Square Deviation (RMSD) between generated and reference ligand centroid trajectories with the
 838 following procedure:

839 1. **Trajectory Standardization.** All ligand centroid trajectories are resampled to a uniform length
 840 ($L = 100$ frames) using linear interpolation, ensuring comparability between different sequences.

841 2. **Best-Match Search.** For each generated trajectory, we identify the reference trajectory that yields
 842 the minimum RMSD. This accounts for the possibility of multiple plausible unbinding pathways.
 843 Let $\mathbf{C}_{\text{gen}} = (\mathbf{c}_1, \dots, \mathbf{c}_L)$ denote the sequence of ligand centroids in a generated trajectory, and let
 844 $\mathcal{T}_{\text{ref}} = \{\mathbf{C}_{\text{ref}}^{(k)}\}_{k=1}^K$ be the set of K reference trajectories. The distance metric is calculated as:

$$845 \quad 846 \quad 847 \quad \text{RMSD}_{\text{best}} = \min_k \sqrt{\frac{1}{L} \sum_{t=1}^L \|\mathbf{c}_t - \mathbf{c}_t^{(k)}\|_2^2} \quad (5)$$

848 To ensure a fair comparison for the static baseline, we construct a pseudo-trajectory by replicating
 849 the initial bound pose L times (i.e., $\mathbf{c}_t = \mathbf{c}_1$ for all t).

850 3. **Final Score.** The reported metric is the average of these best-match RMSDs across all generated
 851 trajectories. Lower RMSD values indicate that the model generates ligand motions more consistent
 852 with physically realistic unbinding paths.

853 **A.3.4 UNBINDING SUCCESS**

854 This metric evaluates whether the generated ligand trajectories successfully capture the unbinding
 855 event. Specifically, we construct the convex hull of the protein heavy atoms in the initial bound
 856 state. If at least one predicted ligand centroid position lies outside this convex hull, the trajectory is
 857 considered as a successful unbinding case.

858 We report the **Success@ k** , which measures the probability that at least one out of k independently
 859 generated trajectories for the same protein–ligand complex achieves successful unbinding. A higher

success rate indicates a better capability of the model to reproduce realistic ligand unbinding processes. Formally, for each complex with k attempts, Success@ k is defined as

$$\text{Success}@k = \frac{1}{N} \sum_{n=1}^N \mathbb{I} \left[\max_{1 \leq j \leq k} s_n^{(j)} = 1 \right],$$

where $s_n^{(j)}$ is the binary success indicator (1 if the j -th trajectory of complex n achieves unbinding, 0 otherwise), and N is the total number of complexes. We report Success@1, Success@5, and Success@10, which reflect performance under varying generation attempts, respectively.

A.3.5 UNBINDING CORRECT RATES

To provide a more rigorous evaluation that directly addresses whether our model identifies the correct exit pathway, we perform an analysis on the successful unbinding trajectories. We define the ground-truth exit position, $\mathbf{c}_{\text{exit}}^{\text{GT}}$, as the ligand's centroid coordinates in the final frame of the reference simulation. We introduce two metrics to quantify the spatial accuracy of the generated exit paths:

1. **Endpoint Distance.** This metric measures how closely the generated pathway approaches the known exit location. For a generated trajectory consisting of L frames, denoted as $\mathbf{C}_{\text{gen}} = (\mathbf{c}_1, \dots, \mathbf{c}_L)$, we calculate the minimum Euclidean distance between any centroid position along the predicted path and the ground-truth exit position:

$$D_{\text{end}} = \min_{t \in \{1, \dots, L\}} \|\mathbf{c}_t - \mathbf{c}_{\text{exit}}^{\text{GT}}\|_2 \quad (6)$$

We report the mean of this minimum distance across all test samples. Lower values indicate that the generated trajectory passes closer to the true exit point.

2. **Correct Rate.** To provide an intuitive measure of success, we calculate the percentage of trajectories that successfully navigate within a specific proximity of the correct exit. Given a distance threshold τ (e.g., 0.5 Å or 1.0 Å), the correct rate is defined as:

$$\text{Rate}(\tau) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(D_{\text{end}}^{(i)} < \tau) \quad (7)$$

where N is the total number of samples, $D_{\text{end}}^{(i)}$ is the endpoint distance for the i -th sample, and $\mathbb{I}(\cdot)$ is the indicator function.

A.4 EXPERIMENTAL DETAILS

A.4.1 SAMPLING SETTINGS

For the first coarse-grained forecasting stage, we set the sampling step $k = 10$ for all three datasets, with different underlying simulation time intervals:

- (i) For the MISATO dataset, the all-atom MD simulations were performed with a 80 ps timestep. Therefore, our sampling of every $k = 10$ steps corresponds to a physical time interval of 0.8 ns (10 steps \times 0.08 ns /step).
- (ii) For the DD-13M dataset, the simulations utilized a 0.1 ps timestep. Consequently, sampling every $k=10$ steps resulted in a time interval of 1 ps between saved frames (10 steps \times 0.1 ps/step). For two case studies (6EY8 and 3PCU), we sampled 100 trajectories to compute the unbinding correct rate metric.
- (iii) For the Atlas dataset, each trajectory has a total simulation length of 100 ns and is provided as a sequence of 10,000 frames. Consequently, sampling every $k=10$ steps resulted in a time interval of 1 ns between saved frames (10 steps \times 100 ps/step).

A.4.2 HYPERPARAMETERS

We use 8 NVIDIA RTX A6000 GPUs for training. For each dataset, we use the Adam optimizer and train BioMD for 20,000 steps with a fixed learning rate of 0.0001 and a batch size of 32. The hyperparameters for FlowTrajectoryTransformer, AtomAttentionEncoder, AtomAttentionDecoder, and GraphTransformer is shown in **Table 4,5,6**.

918 **Table 4. Key hyperparameters for the FlowTrajectoryTransformer model.**

Hyperparameter	Symbol	Value	Description
Number of Transformer Blocks	N_{block}	24	Number of sequential transformer layers.
Number of Attention Heads	N_{head}	16	Number of heads in the multi-head attention.
Single Representation Dim.	c_s	384	Dimension of the per-residue representation.
Pair Representation Dim.	c_z	128	Dimension of the residue-pair representation.

926 **Table 5. Key hyperparameters for the AtomAttentionEncoder and AtomAttentionDecoder model.**

Hyperparameter	Symbol	Value	Description
Number of Transformer Blocks	N_{block}	3	Number of sequential transformer layers.
Number of Attention Heads	N_{head}	4	Number of heads in the multi-head attention.
Single Atom Representation Dim.	c_{atom}	128	Dimension of the per-atom representation.
Pair Atom Representation Dim.	c_{atompair}	16	Dimension of the atom-pair representation.

934 **Table 6. Key hyperparameters for the GraphTransformer model.**

Hyperparameter	Symbol	Value	Description
Number of Encoder Layers	L	6	Number of encoder layers.
Encoder Embedding Dimension	H	384	Encoder embedding dimension.
FFN Embedding Dimension	F	1536	Encoder embedding dimension for FFN.
Number of Attention Heads	A	16	Number of encoder attention heads.

941 **Table 7. Ablation study on the MISATO dataset.** Comparison of all methods on physical stability (first six 942 metrics) and conformational flexibility (last four metrics). Mean values on the test samples are reported.

Method	Bond Geometry ^a		Angle Geometry ^a		Steric Clashes		RMSF Correlation ^b		RMSF Value ^{a,c}	
	MAE	MSE	MAE	MSE	Intra-Lig	Prot-Lig	Ligand	Protein	Ligand (1.211)	Protein (1.002)
Molecular Dynamics	.0377	.0023	.0575	.0053	0	0	-	-	-	-
BioMD (k=1)	.0356	.0021	.0568	.0053	0	0	.4526	.6291	.3717	.3825
BioMD (k=5)	.0387	.0024	.0627	.0066	0	.0001	.4982	.6023	.4552	.4287
BioMD (k=10)	.0495	.0155	.0709	.0097	.0019	.0023	.4789	.6854	.7023	.6242
BioMD (k=20)	.0421	.0030	.0698	.0087	.0002	.0008	.4267	.5309	.5403	.4959

^a Bond geometry (bond length) and RMSF values are in angstroms (Å). Angle geometry (bond angle) is in radians.

^b RMSF Correlation is reported using the Pearson correlation coefficient.

^c RMSF values for reference trajectories are given in parentheses. Values closer to those of the reference indicate better results.

954 **Table 8. Ablation study on the DD-13M dataset.** Comparison of methods on physical stability (first six 955 metrics), ligand unbinding path reconstruction metric (Unbinding Path RMSD), and ligand unbinding success 956 rates. Mean values on the test samples are reported.

Method	Bond Geometry ^a		Angle Geometry ^a		Steric Clashes		Unbinding Path ^a		Unbinding Success		
	MAE	MSE	MAE	MSE	Intra-Lig	Prot-Lig	RMSD	@1	@5	@10	
Static	-	-	-	-	.2778	0	.6504	0	0	0	
Metadynamics ^b	.0246	.0012	.0452	.0030	.2777	0	.4217	-	-	-	
BioMD (AR-1, k=10)	.1880	.0612	.2011	.0761	1.727	4.079	1.131	.2285	.4728	.5822	
BioMD (AR-5, k=10)	.0728	.0111	.0802	.0132	.2943	.0009	.5645	.5676	.7419	.7941	
BioMD (AR-10, k=10)	.1067	.0221	.1478	.0450	.7098	.5356	.7465	.4008	.5719	.6239	
BioMD (w.o. AR, k=10)	.0308	.0018	.0606	.0077	.2944	.0004	.6845	.0029	.0147	.0294	
BioMD (AR-5, k=10, L*2) ^c	.1457	.0382	.1770	.0595	1.407	2.029	.7766	.5735	.7512	.8010	
BioMD (AR-5, k=1) ^d	.0682	.0179	.1231	.0321	.3686	7.380	.7313	.3428	.5128	.6293	
BioMD (AR-5, k=5)	.0575	.0164	.1062	.0246	.4700	1.346	.5703	.5214	.6582	.7023	
BioMD (AR-5, k=10)	.0728	.0111	.0802	.0132	.2943	.0009	.5645	.5676	.7419	.7941	
BioMD (AR-5, k=20)	.0497	.0053	.0914	.0184	.3736	.0967	.5909	.4714	.7065	.8043	

^a Bond geometry (bond length) and unbinding path RMSD values are in angstroms (Å), and angle geometry (bond angle) is in radians.

^b The metadynamics trajectory serves as the lower-bound. The metrics are calculated among trajectories of multiple repeating simulations.

^c "L*2" means that the generated trajectory length is doubled (i.e., 1000 frames).

^d $k = 1$ represents a "forecasting-only" version of BioMD.

972 A.5 ABLATION STUDY
973

974 We conducted ablation studies using BioMD-rel, generating 8 ns trajectories for MISATO and 500
975 frames for DD-13M. During forecasting, we employed a moving window of the 50 most recent
976 frames as historical context to mitigate GPU memory constraints. The quantitative results, presented
977 in **Tables 7 and 8**, demonstrate that the model’s performance is highly sensitive to both the auto-
978 regressive block size (AR- x), which is the number of concurrently denoised frames during sampling,
979 and the hierarchical step size (k), which is the frame interval in the coarse-grained forecasting stage.

980 We first observe that the auto-regressive block size is a critical determinant of physical stability. A
981 minimal block size (e.g., AR-1) necessitates a high number of sequential inference steps, leading
982 to severe error accumulation. This accumulation manifests as catastrophic physical violations, with
983 protein-ligand clashes increasing by over 4000 times compared to the baseline. Conversely, an
984 excessively large block size (e.g., w.o. AR) appears to increase the task complexity, confusing the
985 model and degrading predictive accuracy. Consequently, AR-5 emerges as the optimal trade-off,
986 providing sufficient context for stable predictions without introducing unnecessary computational
987 difficulty.

988 Furthermore, a larger step size (k) proves essential for mitigating error propagation. Our experiments
989 show that a non-hierarchical approach ($k = 1$) suffers from extreme instability due to the vast number
990 of required inference steps, resulting in protein-ligand clashes exceeding 7000 times the baseline
991 levels. While the model is generally less sensitive to k than to the AR setting, the results indicate that
992 $k = 10$ yields superior efficacy for unbinding tasks. Therefore, the specific combination of AR-5 and
993 $k = 10$ was selected as the final configuration, as it effectively balances unbinding metrics with the
994 maintenance of physical realism in complex biomolecular dynamics.

995 A.6 ADDITIONAL EXPERIMENTAL RESULTS
996997 A.6.1 BOND AND ANGLE ERROR
998

999 To evaluate the geometric accuracy and temporal stability of our generative model, we analyzed
1000 the distributions of bond length and angle errors, as well as the error accumulation problem. All
1001 experiments are conducted with BioMD-rel. The results in **Figure 7** show that on the MISATO test
1002 set, 97.9% of bond length errors are below 0.2 Å, and 95.3% of bond angle errors are below 0.2
1003 radians (11.5°). The performance on the DD-13M test set also demonstrates 99.7% of bond errors
1004 and 98.2% of angle errors falling below these thresholds.

1005 Furthermore, we examined the trend of errors over time. On the MISATO dataset, due to the short
1006 generated trajectories (within 100 frames), the errors across different time spans are similar. In
1007 contrast, on the DD-13M dataset involving long trajectory generation, as the number of generated
1008 frames increases (up to 500 frames), a significant increase in error is observed, indicating the
1009 occurrence of the error accumulation problem.

1010 A.6.2 RELAXATION CONSISTENCY ANALYSIS
1011

1012 To assess the energetic quality of the generated conformations, we conducted a relaxation consistency
1013 analysis on them. For each trajectory generated by BioMD-rel on the MISATO test set, we took each
1014 frame and performed energy minimization using the Amber99SB force field. We then calculated the
1015 all-atom RMSD between the pre- and post-relaxation structures. A low RMSD indicates that our
1016 generated conformations are near a local energy minimum, confirming their physical plausibility. Our
1017 analysis reveals a mean all-atom RMSD of 0.72 Å (**Figure 8**) after relaxation across all generated
1018 frames (0.69 Å for MD conformations, **Figure 8**), demonstrating the ability of BioMD to generate
1019 energetically favorable structures and providing strong evidence for the physical plausibility of our
1020 generated trajectories beyond purely geometric metrics.

1021 A.6.3 POCKET-LIGAND INTERACTION ANALYSIS
1022

1023 To evaluate the accuracy of the predicted binding interactions, we perform an analysis on the
1024 MISATO test set (1,031 complexes), focusing on the conformation of key binding residues and their
1025 interactions with the ligand. We identified residues located within 5 Å of the ligand and compared

their conformational properties against ground truth MD simulations. All experiments are conducted with BioMD-rel. As shown in **Figure 9**, the chi1 dihedral angle distributions predicted by BioMD (red dashed line) closely align with those observed in the MD trajectories (blue solid line). Furthermore, we visualized the interaction patterns by calculating the contact probability between protein residues and ligand atoms using a 5 Å cutoff (**Figure 10**). The resulting contact probability map from BioMD exhibits a high degree of consistency with the MD reference, confirming that the model accurately captures both the side-chain geometries and the dynamic contact profiles within the binding pocket.

A.6.4 TICA ANALYSIS

For each protein in the test set, we projected the generated conformations onto a 2D space using Time-lagged Independent Component Analysis (TICA). We train the BioMD-rel model on the Atlas dataset, which contains 100 ns MD trajectories for 1,390 protein chains. The first two independent components (tIC1 and tIC2), which capture the slowest and most significant motions of the system, were used to define a 2D conformational space. We first fit TICA on a 100 ns MD trajectory and then projected the conformations generated by BioMD onto this 2D space. Specifically, we used BioMD to sample three 100 ns trajectories and projected all generated conformations onto the TICA space. For comparison, we also projected conformations from other MD simulation replicas (100 ns, 10 ns, and 1 ns) onto the same space. As shown in **Figure 13**, BioMD explores the sample space more sufficiently than 1 ns and 10 ns MD simulations, capturing a broader range of conformational diversity. Notably, the average sampling time for one trajectory is only 56 seconds, making BioMD much faster than MD simulation.

Table 9. Comparison of methods on endpoint distance and unbinding correct rates. Mean values on the test samples are reported.

Method	Endpoint Distance	Correct Rate (< X Å)			
		RMSD (Å)	< 0.5 Å	< 1 Å	< 2 Å
Static	3.5852	0.1765	0.3235	0.4118	0.7353
Metadynamics ^b	0.9739	0.5136	0.7532	0.9022	0.9641
BioMD-rel	2.3433	0.4265	0.5265	0.6971	0.8000
BioMD-abs	2.3897	0.3765	0.5324	0.6882	0.8059
BioMD-rel (AR-5)	1.8001	0.3471	0.5882	0.7529	0.8706
BioMD-abs (AR-5)	1.7253	0.4912	0.6824	0.7853	0.8588

^b The metadynamics trajectory serves as the lower-bound. The metrics are calculated among trajectories of multiple repeating simulations.

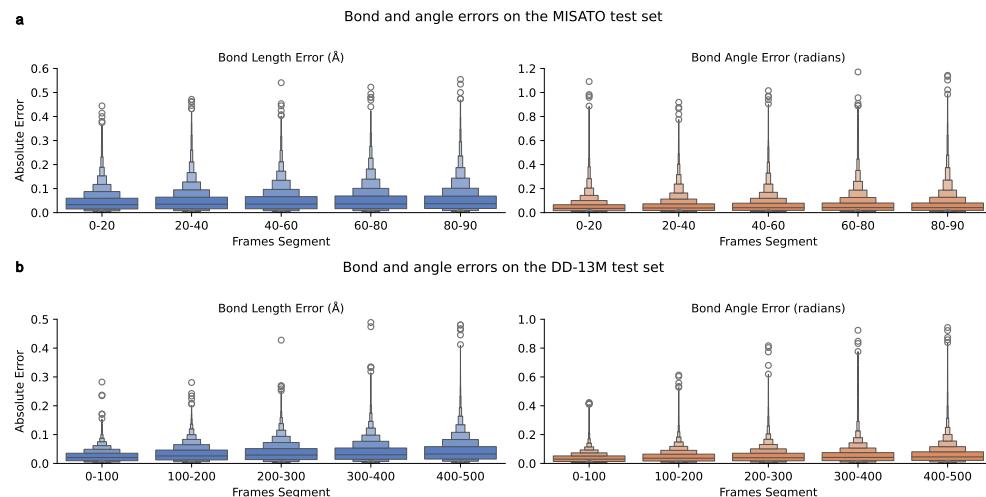


Figure 7. Bond and angle error distribution. a. Boxen plot showing error distribution on the MISATO test set. b. Boxen plot showing error distribution on the DD-13M test set.

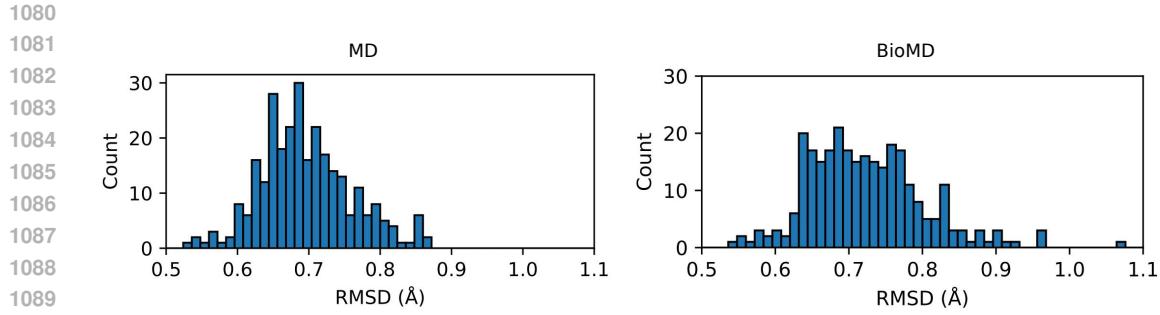


Figure 8. All-atom RMSD after relaxation. All-atom RMSD when using the Amber99SB force field for relaxation on MD conformations (mean=0.69 Å) and BioMD-generated conformations (mean=0.72 Å).

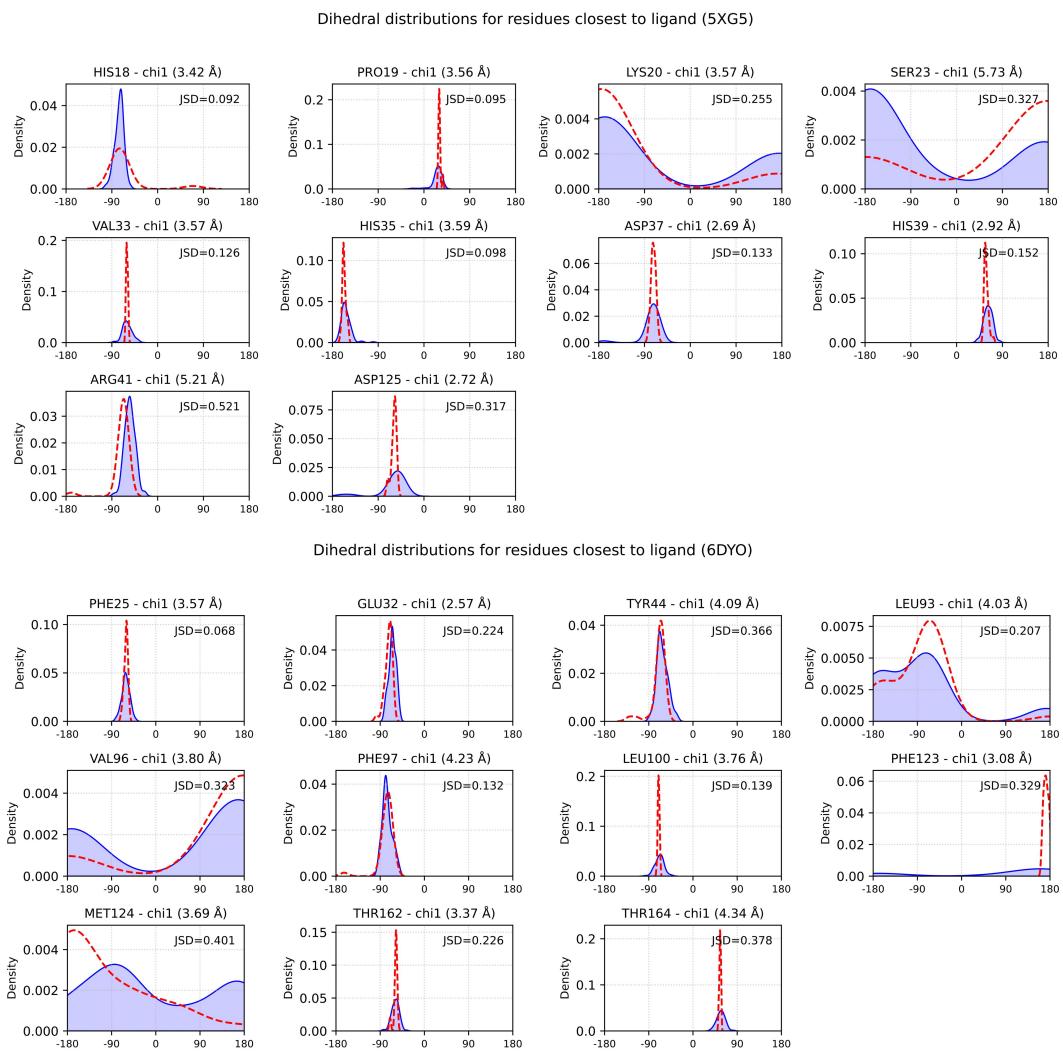
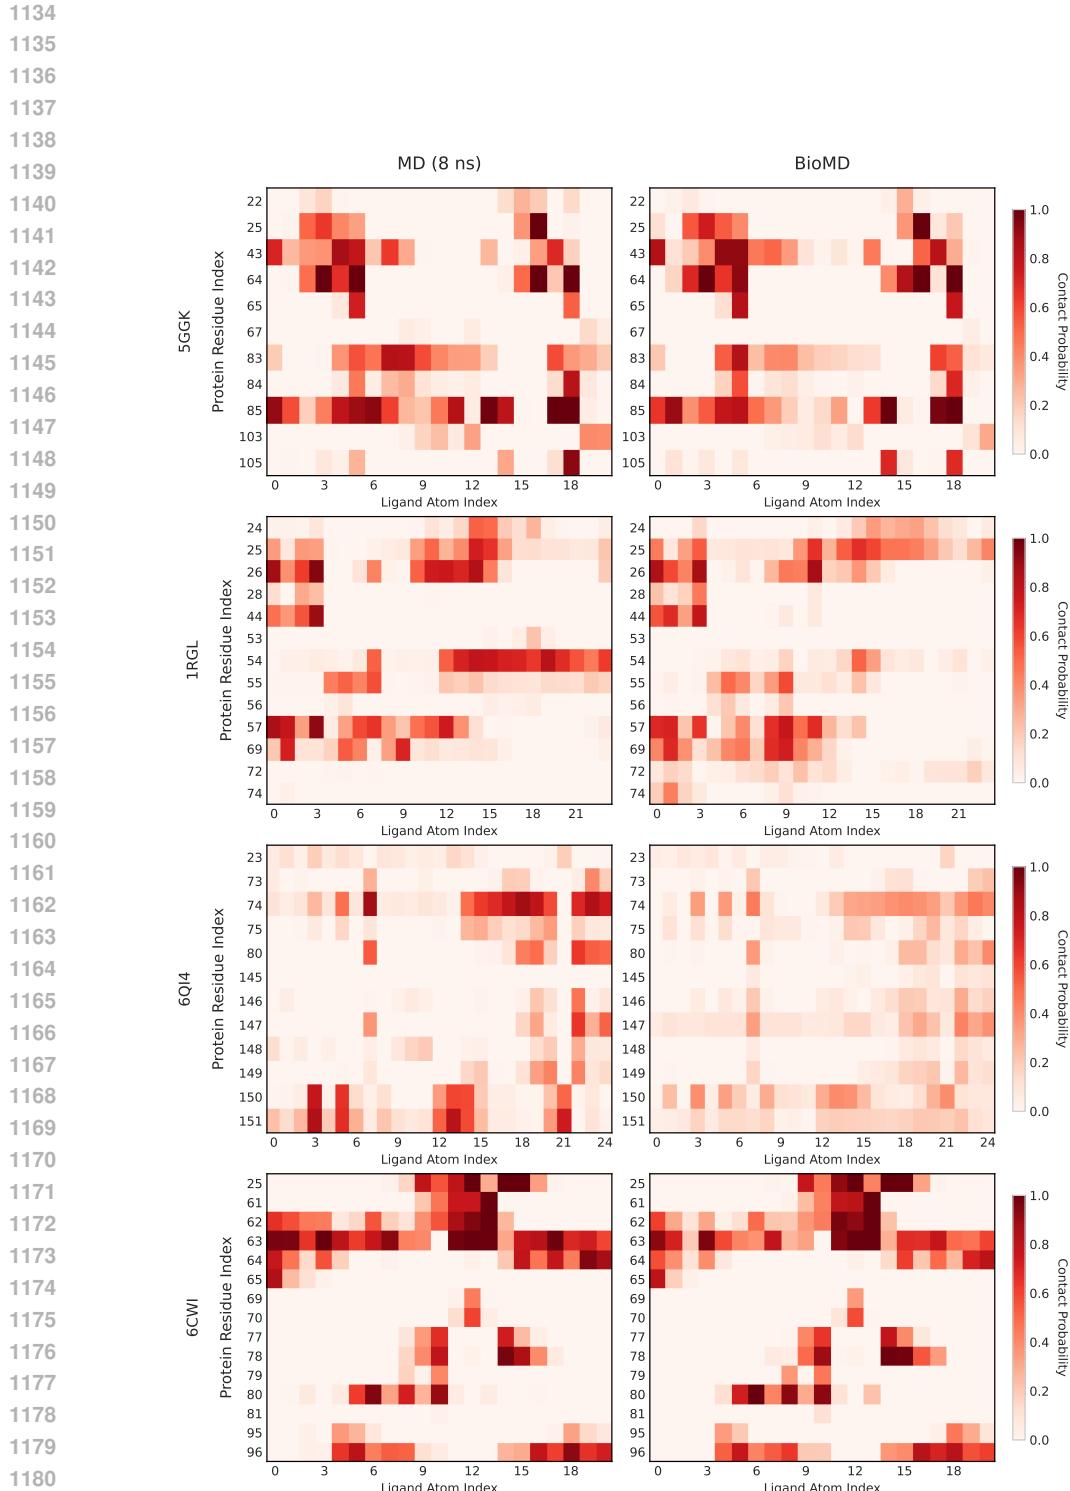
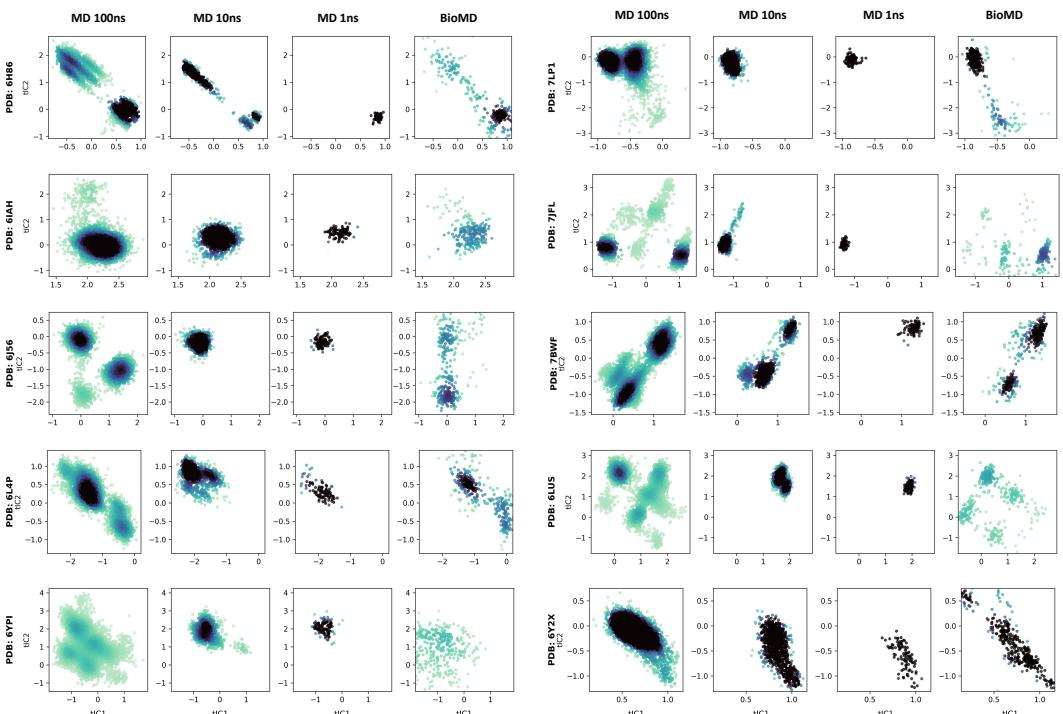
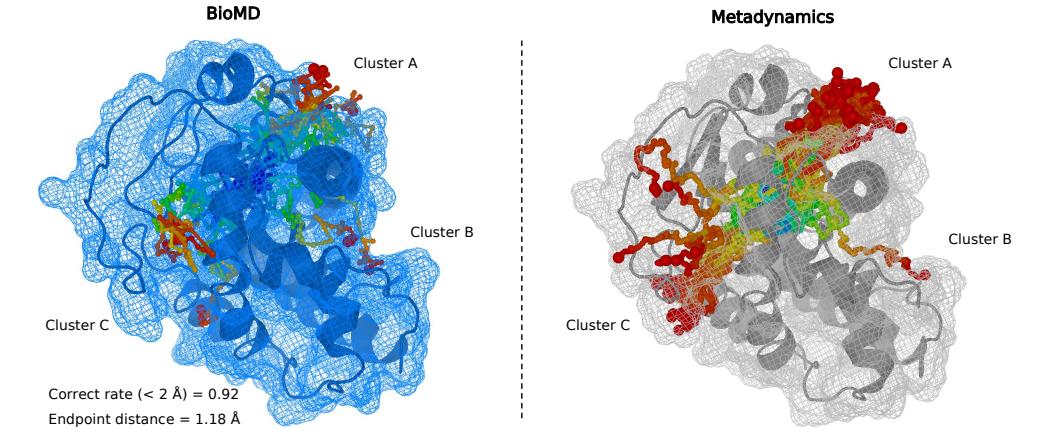
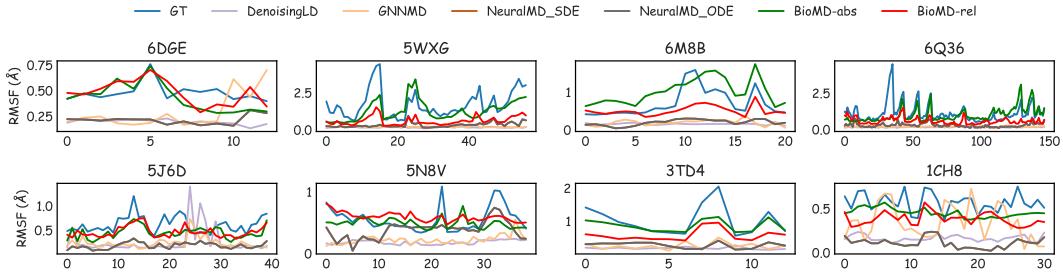


Figure 9. Dihedral distributions for residues closest to the ligand. The dashed red line represents BioMD, while the solid blue line represents the ground truth MD simulations. The value in parentheses indicates the minimum distance between the considered residue atoms and all ligand atoms. Test cases (6DYO and 5XG5) are from the MISATO test set.





1242 **Algorithm 1: Main Inference Loop**

1243 **Input:** $\{\mathbf{f}^*\}$, $\{\vec{\mathbf{x}}_{0,l}\}$, $N_{cycle} = 4$, $c_s = 384$, $c_z = 128$

1244 1 $\{\mathbf{s}_i^{\text{inputs}}\} \leftarrow \text{InputFeatureEmbedder}(\{\mathbf{f}^*\})$;

1245 2 $\mathbf{s}_i^{\text{init}} \leftarrow \text{LinearNoBias}(\mathbf{s}_i^{\text{inputs}})$;

1246 3 $\mathbf{z}_{ij}^{\text{init}} \leftarrow \text{LinearNoBias}(\mathbf{s}_i^{\text{inputs}}) + \text{LinearNoBias}(\mathbf{s}_j^{\text{inputs}})$;

1247 4 $\{\mathbf{z}_{ij}\}, \{\mathbf{s}_i\} \leftarrow 0, 0$;

1248 5 **foreach** $c \in \{1, \dots, N_{cycle}\}$ **do**

1249 6 $\mathbf{z}_{ij} \leftarrow \mathbf{z}_{ij}^{\text{init}} + \text{LinearNoBias}(\text{LayerNorm}(\mathbf{z}_{ij}))$;

1250 7 $\{\mathbf{z}_{ij}\}, \{\mathbf{s}_i\} \leftarrow \text{GraphTransformer}(\{\vec{\mathbf{x}}_{0,l}\}, \{\mathbf{s}_i\}, \{\mathbf{z}_{ij}\}, \{\mathbf{s}_i^{\text{inputs}}\})$;

1251 8 $\mathbf{s}_i \leftarrow \mathbf{s}_i^{\text{init}} + \text{LinearNoBias}(\text{LayerNorm}(\mathbf{s}_i))$;

1252 9 traj_list = $\{\vec{\mathbf{x}}_{0,l}\}$;

1253 10 **foreach** $t \in \{1, \dots, T\}$ **do**

1254 11 $\{\vec{\mathbf{x}}_{t,l}^{\text{pred}}\} \leftarrow \text{SampleFlow}(\{\vec{\mathbf{x}}_{his}\}, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i\}, \{\mathbf{z}_{ij}\})$;

1255 12 traj_list.add($\vec{\mathbf{x}}_{t,l}^{\text{pred}}$);

1256 13 $\{\vec{\mathbf{x}}_{his}\} = \text{traj_list}$;

1257 14 **return** traj_list

1261 **Algorithm 2: TrainFlow**

1262 **Input:** $\{\vec{\mathbf{x}}_l\}, \{\vec{\mathbf{x}}_{his}\}, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\}$

1263 1 **# Independent noise levels**;

1264 2 $\tau \sim (\mathcal{U}(0, 1), \mathcal{U}(0, 1), \dots, \mathcal{U}(0, 1))$;

1265 3 $\{\vec{\mathbf{x}}_l^0\} \sim \mathcal{N}(\vec{0}, \mathbf{I}_3)$;

1266 4 $\{\vec{\mathbf{x}}_l\} \leftarrow \text{CentreRandomAugmentation}(\{\vec{\mathbf{x}}_l\})$;

1267 5 $\{\vec{\mathbf{x}}_l^\tau\} = \tau \{\vec{\mathbf{x}}_l\} + (1 - \tau) \{\vec{\mathbf{x}}_l^0\}$;

1268 6 $\{\vec{\mathbf{u}}_l^\tau\} \leftarrow \text{FlowModule}(\{\vec{\mathbf{x}}_l^\tau\}, \{\vec{\mathbf{x}}_{his}\}, \tau, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\})$;

1269 7 $\mathcal{L}_{flow} = \text{MSE}(\{\vec{\mathbf{u}}_l^\tau\}, \{\frac{\vec{\mathbf{x}}_l - \vec{\mathbf{x}}_l^\tau}{1 - \tau}\})$;

1270 8 **return** \mathcal{L}_{flow}

1273 **Algorithm 3: SampleFlow**

1274 **Input:** $\{\vec{\mathbf{x}}_{his}\}, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\}$

1275 1 $\vec{\mathbf{x}}_l^0 \sim \mathcal{N}(\vec{0}, \mathbf{I}_3)$;

1276 2 **foreach** τ in $\{0, 0.1, 0.2, \dots, 0.9\}$ **do**

1277 3 $\{\vec{\mathbf{u}}_l^\tau\} \leftarrow \text{FlowModule}(\{\vec{\mathbf{x}}_l^\tau\}, \{\vec{\mathbf{x}}_{his}\}, \tau, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\})$;

1278 4 $\vec{\mathbf{x}}_l^{\tau+1} \leftarrow \vec{\mathbf{x}}_l^\tau + dt \cdot \vec{\mathbf{u}}_l^\tau$;

1279 5 **return** $\{\vec{\mathbf{x}}_l^1\}$

1283 1284 **Usage Note:** The returned bias is added to the attention logits inside the main transformer block before the softmax activation: $\text{logits} \leftarrow (Q \cdot K^T) / \sqrt{d_k} + \text{bias}$.

1285

1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296
1297
1298

Algorithm 4: FlowModule

1300 **Input:** $\{\vec{x}_l^{\text{noisy}}\}, \{\vec{x}_{\text{his}}\}, t, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\}$
 1301 $\sigma_{\text{data}} = 16, c_{\text{atom}} = 128, c_{\text{atompair}} = 16, c_{\text{token}} = 768$
 1302 1 $\{\mathbf{s}_i\}, \{\mathbf{z}_{ij}\} \leftarrow \text{FlowConditioning}(t, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\}, \sigma_{\text{data}});$
 1303 2 **# Sequence-local Atom Attention with history info and aggregation to coarse-grained tokens ;**
 1304 3 $\{a_i\}, \{q_k^{\text{skip}}\}, \{p_k^{\text{skip}}\}, \{t_k^{\text{skip}}\} \leftarrow$
 1305 $\text{AtomAttentionEncoder}(\{\vec{x}_{\text{his}}\}, \{\mathbf{f}^*\}, \{\vec{x}_l^{\text{noisy}}\}, \{\mathbf{s}_i\}, \{\mathbf{z}_{ij}\}, c_{\text{atom}}, c_{\text{atompair}}, c_{\text{token}});$
 1306 4 **# Full self-attention on token level;**
 1307 5 $a_i \leftarrow \text{LinearNoBias}(\text{LayerNorm}(a_i));$
 1308 6 $\{a_k\} \leftarrow \text{FlowTrajectoryTransformer}(\{a_i\}, \{\mathbf{s}_i\}, \{\mathbf{z}_{ij}\}, \beta_{ij} = 0, N_{\text{block}} = 24, N_{\text{head}} = 16);$
 1309 7 $a_i \leftarrow \text{LayerNorm}(a_i);$
 1310 8 **# Broadcast token activations to atoms and run Atom Attention.;**
 1311 9 $\{\vec{u}_l\} \leftarrow \text{AtomAttentionDecoder}(\{a_i\}, \{q_k^{\text{skip}}, p_k^{\text{skip}}, t_k^{\text{skip}}\});$
 1312 10 **return** $\{\vec{u}_l\}$

1314
1315
1316
1317
1318
1319

Algorithm 5: FlowConditioning

1320 **Input:** $\hat{t}, \{\mathbf{f}^*\}, \{s_i^{\text{inputs}}\}, \{s_i^{\text{trunk}}\}, \{z_{ij}^{\text{trunk}}\}, \sigma_{\text{data}}, c_z = 128, c_s = 384$
 1321 1 **# Pair conditioning;**
 1322 2 $z_{ij} \leftarrow \text{LinearNoBias}(\text{LayerNorm}(z_{ij}));$
 1323 3 **foreach** $b \in \{1, 2\}$ **do**
 1324 4 $z_{ij} += \text{Transition}(z_{ij}, n = 2);$
 1325 5 **# Single conditioning;**
 1326 6 $s_i \leftarrow \text{concat}([\mathbf{s}_i^{\text{trunk}}, \mathbf{s}_i^{\text{inputs}}]);$
 1327 7 $s_i \leftarrow \text{LinearNoBias}(\text{LayerNorm}(s_i));$
 1328 8 **foreach** $b \in \{1, 2\}$ **do**
 1329 9 $s_i += \text{Transition}(s_i, n = 2);$
 1330 10 **return** $\{s_i\}, \{z_{ij}\}$

1334
1335
1336
1337
1338
1339

Algorithm 6: FlowTrajectoryTransformer

1340 **Input:** $\{a_i\}, \{s_i\}, \{z_{ij}\}, \{\beta_{ij}\}, N_{\text{block}} = 24, N_{\text{head}} = 16$
 1341 1 **for** $n \in [1, \dots, N_{\text{block}}]$ **do**
 1342 2 $\{b_i\} \leftarrow \text{AttentionPairBias}(\{a_i\}, \{s_i\}, \{z_{ij}\}, \{\beta_{ij}\}, N_{\text{head}});$
 1343 3 $\{b_i\} \leftarrow \text{TemporalAttention}(\{a_i + b_i\});$
 1344 4 $a_i \leftarrow b_i + \text{ConditionedTransitionBlock}(a_i, s_i);$
 1345 5 **return** $\{a_i\}$

1347
1348
1349

1350

1351

1352

1353

1354

1355

1356

1357

Algorithm 7: TemporalAttention Module

1358

Input: Single representation c_s of shape (T, N, c_s) , where T is time, N is residues.

1359

Output: Updated single representation output of shape (T, N, c_s) .

1360

1 # Permute dimensions to make time the sequence axis for attention

1361

 $c'_s \leftarrow \text{Permute}(c_s, \text{dims} = (1, 0, 2))$;

1362

2 # Project to Query, Key, Value for each residue independently;

1363

 3 $Q \leftarrow \text{Linear}_Q(c'_s)$;

1364

 4 $K \leftarrow \text{Linear}_K(c'_s)$;

1365

 5 $V \leftarrow \text{Linear}_V(c'_s)$;

1366

6 # Calculate scaled dot-product attention scores across time;

1367

7 $d_k \leftarrow \text{dimension of } K$;

1368

8 logits $\leftarrow (Q \cdot K^T) / \sqrt{d_k}$;

1369

9 weights $\leftarrow \text{Softmax}(\text{logits}, \text{dim} = -1)$;

1370

10 # Apply attention weights to values;

1371

11 output' $\leftarrow \text{weights} \cdot V$;

1372

12 # Permute back to the original dimension order;

1373

13 output $\leftarrow \text{Permute}(\text{output}', \text{dims} = (1, 0, 2))$;

1374

14 **return** output;

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

Algorithm 8: AttentionPairBias Module

1390

Input: Pair representation c_z of shape (N, N, c_z) , where N is residues.

1391

Output: Attention bias bias of shape (N_{head}, N, N) .

1392

1 # Project c_z to match the number of attention heads;

1393

2 bias' $\leftarrow \text{Linear}_{\text{bias}}(c_z)$;

1394

3 # Permute dimensions to align with attention logits shape;

1395

4 bias $\leftarrow \text{Permute}(\text{bias}', \text{dims} = (2, 0, 1))$;

1396

5 **return** bias;

1397

1398

1399

1400

1401

1402

1403