

A MULTI-GRAINED GROUP SYMMETRIC FRAMEWORK FOR LEARNING PROTEIN-LIGAND BINDING DYNAMICS

Anonymous authors

Paper under double-blind review

ABSTRACT

In drug discovery, molecular dynamics (MD) simulation for protein-ligand binding provides a powerful tool for predicting binding affinities, estimating transport properties, and exploring pocket sites. There has been a long history of improving the efficiency of MD simulations through better numerical methods and, more recently, by augmenting them with machine learning (ML) methods. Yet, challenges remain, such as accurate modeling of extended-timescale simulations. To address this issue, we propose NeuralMD, the first ML surrogate that can facilitate numerical MD and provide accurate simulations of protein-ligand binding dynamics. We propose a principled approach that incorporates a novel physics-informed multi-grained group symmetric framework. Specifically, we propose (1) a BindingNet model that satisfies group symmetry using vector frames and captures the multi-level protein-ligand interactions, and (2) an augmented neural ordinary differential equation solver that learns the trajectory under Newtonian mechanics. For the experiment, we design ten single-trajectory and three multi-trajectory binding simulation tasks. We show the efficiency and effectiveness of NeuralMD, with a $2000\times$ speedup over standard numerical MD simulation and outperforming all other ML approaches by up to $\sim 80\%$ under the stability metric. We further qualitatively show that NeuralMD reaches more stable binding predictions.

1 INTRODUCTION

The simulation of protein-ligand binding dynamics is one of the fundamental tasks in drug discovery (Kairys et al., 2019; Yang et al., 2020; Volkov et al., 2022). Such simulations of binding dynamics are a key component of the drug discovery pipeline to select, refine, and tailor the chemical structures of potential drugs to enhance their efficacy and specificity. To simulate the protein-ligand binding dynamics, *numerical molecular dynamics (MD)* methods have been extensively developed. However, the numerical MD methods are computationally expensive due to the expensive force calculations on individual atoms in a large protein-ligand system.

To alleviate this issue, *machine learning (ML)* surrogates have been proposed to either augment or replace numerical MD methods to estimate the MD trajectories. However, all prior ML approaches for MD are limited to single-system dynamics (e.g., small molecules or proteins) and not protein-ligand binding dynamics. A primary reason is the lack of large-scale datasets. The first large-scale dataset with binding dynamics was released in May 2023 (Siebenmorgen et al., 2023), and to our knowledge, we are now the first to explore it in this paper. Further, prior ML-based MD approaches limit to studying the MD dynamics on a small time interval (1e-15 seconds), while simulation on a longer time interval (e.g., 1e-9 seconds) is needed for specific tasks, such as detecting the transient and cryptic states (Vajda et al., 2018) in binding dynamics. However, such longer-time MD simulations are challenging due to the catastrophic buildup of errors over longer rollouts (Fu et al., 2022a).

Another critical aspect that needs to be integrated into ML-based modeling is the group symmetry present in protein-ligand geometry. Specifically, the geometric function should be equivariant to rotation and translation (i.e., SE(3)-equivariance). The principled approach to satisfy equivariance is to use vector frames, which have been previously explored for single molecules (Jumper et al.,

*Equal contribution

†Equal advising. Correspondence to hongyu.guo@uottawa.ca, jchayes@berkeley.edu.




Figure 1: Visualization of last-snapshot binding predictions on three PDB complexes. NeuralMD stays more stable than DenoisingLD, exhibiting a lower degree of torsion with the natural conformations. Other methods collapse heavily, including GNN-MD and VerletMD, where atoms extend beyond the frame for the latter.

2021), but not yet for the protein-ligand binding complexes. The vector frame basis achieves SE(3)-equivariance by projecting vectors (*e.g.*, positions and accelerations) into the vector frame basis, and such a projection can maintain the equivariant property with efficient calculations (Liu et al., 2023a).

Our Approach: NeuralMD. We propose NeuralMD, a multi-grained physics-informed approach designed to handle extended-timestep MD simulations for protein-ligand binding dynamics. Our multi-grained method explicitly decomposes the complexes into three granularities: the atoms in ligands, the backbone structures in proteins, and the residue-atom pairs in the complex, to obtain a scalable approach for modeling a large system. We achieve **group symmetry** in BindingNet through the incorporation of vector frames, and include three levels of vector frame bases for multi-grained modeling, from the atom and backbone level to the residue level for binding interactions.

Further, our ML approach NeuralMD preserves the **Newtonian mechanics**. In MD, the movement of atoms is determined by Newton’s second law, $F = m \cdot a$, where F is the force, m is the mass, and a is the acceleration of each atom. By integrating acceleration and velocity w.r.t. time, we can obtain the velocities and positions, respectively. Thus in NeuralMD, we formulate the trajectory simulation as a second-order ordinary differential equation (ODE) problem, and it is solved using neural ODE.

Experiments. To verify the effectiveness and efficiency of NeuralMD, we design ten single-trajectory and three multi-trajectory binding simulation tasks. For evaluation, we adopt the recovery and stability metrics (Fu et al., 2022a). NeuralMD achieves $2000\times$ speedup compared to the numerical methods. We observe that NeuralMD outperforms all other ML methods (Zhang et al., 2018; Musaelian et al., 2023b; Fu et al., 2022b; Wu & Li, 2023; Arts et al., 2023) on 12 tasks using recovery metric, and NeuralMD is consistently better by a large gap using the stability metric (up to $\sim 80\%$). Qualitatively, we illustrate that NeuralMD realizes more stable binding dynamics predictions in three case studies. They are three protein-ligand binding complexes from Protein Data Bank (PDB), as shown in Figure 1.

Related Work. We briefly review the most related work here and include a more comprehensive discussion in Appendix A. **Geometric modelings** can be split into three categories (Liu et al., 2023a; Zhang et al., 2023): invariant models that utilize invariant features (distances and angles) (Schütt et al., 2018; Klicpera et al., 2020), equivariant models with spherical frames (Smidt et al., 2018; Musaelian et al., 2022), and equivariant models with vector frames (Jumper et al., 2021). This also includes the works studying the protein-ligand binding (Stepniewska-Dziubinska et al., 2018; Jiménez et al., 2018; Jones et al., 2021; Yang et al., 2023; Corso et al., 2023), but they are focusing on the equilibrium state, not the dynamics. On the other hand, existing ML works have studied **molecular simulation** (Zhang et al., 2018; Doerr et al., 2020; Musaelian et al., 2023b), and they can be categorized into two groups, each with its respective limitations. (1) The first direction focuses on predicting the energy (or force) (Zhang et al., 2018; Doerr et al., 2020; Musaelian et al., 2023b), which is then fed into the numerical integration algorithm for trajectory simulations. (2) Another research line directly predicts the atomic positions along the trajectories in an auto-regressive




Figure 2: Illustrations of NeuralMD for protein-ligand binding dynamics. The landscape depicts the energy level, and the binding dynamic leads to an equilibrium state with lower energy. For NeuralMD, the force prediction follows the SE(3)-equivariance, and the dynamics prediction follows the Newtonian mechanics.

(AR) manner (Wu & Li, 2023; Arts et al., 2023; Fu et al., 2022b). Our work proposes a geometric ML framework as an MD simulation tool to learn the protein-ligand binding dynamics. It is also important to acknowledge the efforts of the whole community in constructing the datasets and benchmarks, including MISATO (Siebenmorgen et al., 2023) and Geom3D (Liu et al., 2023a).

2 PRELIMINARIES

Ligands. In this work, we consider binding complexes involving small molecules as ligands. Small molecules can be treated as sets of atoms in the 3D Euclidean space, $\{f^{(l)}, \mathbf{x}^{(l)}\}$, where $f^{(l)}$ and $\mathbf{x}^{(l)}$ represent the atomic numbers and 3D Euclidean coordinates for atoms in each ligand, respectively.

Proteins. Proteins are macromolecules, which are essentially chains of amino acids or residues. There are 20 natural amino acids, and each amino acid is a small molecule. Noticeably, amino acids are made up of three components: a basic amino group ($-\text{NH}_2$), an acidic carboxyl group ($-\text{COOH}$), and an organic R group (or side chain) that is unique to each amino acid. Additionally, the carbon that connects all three groups is called C_α . Due to the large number of atoms in proteins, this work proposes a multi-grained method for modeling the protein-ligand complexes. In this regard, the **backbone-level** data structure for each protein is $\{f^{(p)}, \{\mathbf{x}_N^{(p)}, \mathbf{x}_{C_\alpha}^{(p)}, \mathbf{x}_C^{(p)}\}\}$, for the residue type and the coordinates of $N - C_\alpha - C$ in each residue, respectively. (We may omit the superscript in the coordinates of backbone atoms, as these backbone structures are unique to protein residues.) In addition to the backbone level, for a coarser-grained data structure on protein-ligand complex, we further consider **residue-level** modeling for binding interactions, $\{f^{(p)}, \mathbf{x}^{(p)}\}$, where the coordinate of C_α is taken as the residue-level coordinate, *i.e.*, $\mathbf{x}^{(p)} \triangleq \mathbf{x}_{C_\alpha}^{(p)}$.

Molecular Dynamics Simulations. Generally, molecular dynamics (MD) describes how each atom in a molecular system moves over time, following Newton’s second law of motion:

$$F = m \cdot a = m \cdot \frac{d^2 \mathbf{x}}{dt^2}, \quad (1)$$

where F is the force, m is the mass, a is the acceleration, \mathbf{x} is the position, and t is the time. Then, an MD simulation will take Newtonian dynamics, an ordinary differential equation (ODE), to get the trajectories, where such a molecular system can be a small molecule, a protein, a polymer, or a protein-ligand binding complex. The **numerical methods for MD** can be classified into classical MD and *ab-initio* MD, where the difference lies in how the force on each atom is calculated: *ab-initio* MD calculates the forces using a quantum-mechanics-based method, such as density functional theory (DFT), while classical MD uses force-field-based approaches to predict the atomic forces. More recently, **machine learning (ML) methods for MD** have opened a new perspective by utilizing the group symmetric tools for geometric representation and the automatic differential tools for dynamics learning (please check Appendix A for a more detailed discussion).

In MD simulations, the systems are considered to be either in a vacuum or with explicit modeling of solvent or air molecules. The former is impractical in a real-world system, especially when the jostling




Figure 3: Three granularities of vector frame basis in BindingNet: (a) atom-level basis for ligands, (b) backbone-level basis for proteins, and (c) residue-level basis for the protein-ligand complex.

of solvent or air molecules causes friction, and the occasional high-velocity collision can perturb the system, leading to prohibitively high computational costs. Thus, as a surrogate solution, researchers have considered adopting Brownian dynamics or Langevin dynamics for molecular simulation (Wu & Li, 2023): they can reduce the computational cost while keeping track of the stochasticity of solvent or air molecules. In the ML-based MD simulations, adopting Newtonian dynamics or Langevin dynamics can be treated as a selection of different inductive biases when building the framework. In this work, we follow Newtonian dynamics since the experiment dataset (MISATO (Siebenmorgen et al., 2023)) uses classical MD, explicitly including all the solvent molecules.

Problem Setting: Protein-Ligand Binding Dynamics Simulation. In this work, we are interested in the protein-ligand binding dynamics in the semi-flexible setting (Salmaso & Moro, 2018), *i.e.*, proteins with rigid structures and ligands with flexible movements (Siebenmorgen et al., 2023; Corso et al., 2023). Thus, the problem is formulated as follows: suppose we have a fixed protein structure $\{f^{(p)}, \{\mathbf{x}_N^{(p)}, \mathbf{x}_C^{(p)}, \mathbf{x}_C^{(p)}\}\}$ and a ligand with its initial structure and velocity, $\{f^{(l)}, \mathbf{x}_0^{(l)}, \mathbf{v}_0^{(l)}\}$. We want to predict the trajectories of ligands following the Newtonian dynamics, *i.e.*, the movement of $\{\mathbf{x}_t^{(l)}, \dots\}$ over time. We also want to clarify two critical points about this problem setting: (1) Our task is trajectory prediction, *i.e.*, positions as labels, and no explicit energy and force labels are considered. ML methods for energy prediction followed with numerical ODE solver may require smaller time intervals (*e.g.*, 1e-15 seconds), while trajectory prediction, which directly predicts the positions, is agnostic to the magnitude of time intervals. This is appealing for datasets like MISATO with larger time intervals (*e.g.*, 1e-9 seconds), as will be discussed in Section 4. (2) Each trajectory is composed of a series of geometries of molecules, and such geometries are called **snapshots**. We avoid using *frames* since we will introduce the *vector frame* in modeling the binding complex in Section 3.

3 METHOD: BINDINGNET AND NEURALMD

In this section, we will introduce our framework, NeuralMD, for simulating the protein-ligand binding dynamics. It has two main phases: (1) A multi-grained SE(3)-equivariant geometric model, BindingNet. It models the protein-ligand complex from three granularities, employing three levels of vector frame basis: atom level for ligands, backbone level for proteins, and residue level for protein-ligand complexes. (2) A second-order ordinary differential equation (ODE) solver to learn the Newtonian mechanics. We propose a variant of neural ODE with two attributes (Chen et al., 2018): the ODE augmentation for the efficient second-order solver and the joint ODE for lower memory costs. This section outlines the structure: the three levels of vector frames in Section 3.1, the architecture of BindingNet in Section 3.2, the design of NeuralMD in Section 3.3, and implementation details in Section 3.4.

3.1 MULTI-GRAINED VECTOR FRAMES

Proteins are essentially macromolecules composed of thousands of residues (amino acids), where each residue is a small molecule. Thus, it is infeasible to model all the atoms in proteins due to the large volume of the system, and such an issue also holds for the protein-ligand complex. To address this issue, we propose BindingNet, a multi-grained SE(3)-equivariant model, to capture the interactions between a ligand and a protein. The vector frame basis ensures SE(3)-equivariance, and the multi-granularity is achieved by considering frames at three levels.

Vector Frame Basis for SE(3)-Equivariance. Recall that the geometric representation of the whole molecular system needs to follow the physical properties of the equivariance w.r.t. rotation and translation. Such a group symmetric property is called SE(3)-equivariance. We also want to point out that the reflection or chirality property is equivariant for properties like energy, yet it is not for

the ligand modeling with rigid protein structures (*i.e.*, antisymmetric to the reflection). The vector frame basis can handle this naturally, and we leave a more detailed discussion in Appendix C, along with the proof on group symmetry of vector frame basis. In the following, we introduce three levels of vector frames for multi-grained modeling.

Atom-Level Vector Frame for Ligands. For small molecule ligands, we first extract atom pairs (i, j) within the distance cutoff c , and the vector frame basis is constructed using the Gram-Schmidt as:

$$\mathcal{F}_{\text{ligand}} = \left(\frac{\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}}{\|\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}\|}, \frac{\mathbf{x}_i^{(l)} \times \mathbf{x}_j^{(l)}}{\|\mathbf{x}_i^{(l)} \times \mathbf{x}_j^{(l)}\|}, \frac{\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}}{\|\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}\|} \times \frac{\mathbf{x}_i^{(l)} \times \mathbf{x}_j^{(l)}}{\|\mathbf{x}_i^{(l)} \times \mathbf{x}_j^{(l)}\|} \right), \quad (2)$$

where \times is the cross product. Note that both $\mathbf{x}_i^{(l)}$ and $\mathbf{x}_j^{(l)}$ are for geometries at time t - henceforth, we omit the subscript t for brevity. Such an atom-level vector frame allows us to do SE(3)-equivariant message passing to get the atom-level representation.

Backbone-Level Vector Frame for Proteins. Proteins can be treated as chains of residues, where each residue possesses a backbone structure. The backbone structure comprises an amino group, a carboxyl group, and an alpha carbon, delegated as $N - C_\alpha - C$. Such a structure serves as a natural way to build the vector frame. For each residue in the protein, the coordinates are \mathbf{x}_N , \mathbf{x}_{C_α} , and \mathbf{x}_C , then the backbone-level vector frame for this residue is:

$$\mathcal{F}_{\text{protein}} = \left(\frac{\mathbf{x}_N - \mathbf{x}_{C_\alpha}}{\|\mathbf{x}_N - \mathbf{x}_{C_\alpha}\|}, \frac{\mathbf{x}_{C_\alpha} - \mathbf{x}_C}{\|\mathbf{x}_{C_\alpha} - \mathbf{x}_C\|}, \frac{\mathbf{x}_N - \mathbf{x}_{C_\alpha}}{\|\mathbf{x}_N - \mathbf{x}_{C_\alpha}\|} \times \frac{\mathbf{x}_{C_\alpha} - \mathbf{x}_C}{\|\mathbf{x}_{C_\alpha} - \mathbf{x}_C\|} \right). \quad (3)$$

This is built for each residue, providing a residue-level representation.

Residue-Level Vector Frame for Protein-Ligand Complexes. It is essential to model the protein-ligand interaction to better capture the binding dynamics. We achieve this by introducing the residue-level vector frame. More concretely, proteins are sequences of residues, marked as $\{(f_0^{(p)}, \mathbf{x}_0^{(p)}), \dots, (f_i^{(p)}, \mathbf{x}_i^{(p)}), (f_{i+1}^{(p)}, \mathbf{x}_{i+1}^{(p)}), \dots\}$. Here, we use a cutoff threshold c to determine the interactions between ligands and proteins, and the interactive regions on proteins are called pockets. We construct the following vector frame for residues in the pockets sequentially:

$$\mathcal{F}_{\text{complex}} = \left(\frac{\mathbf{x}_i^{(p)} - \mathbf{x}_{i+1}^{(p)}}{\|\mathbf{x}_i^{(p)} - \mathbf{x}_{i+1}^{(p)}\|}, \frac{\mathbf{x}_i^{(p)} \times \mathbf{x}_{i+1}^{(p)}}{\|\mathbf{x}_i^{(p)} \times \mathbf{x}_{i+1}^{(p)}\|}, \frac{\mathbf{x}_i^{(p)} - \mathbf{x}_{i+1}^{(p)}}{\|\mathbf{x}_i^{(p)} - \mathbf{x}_{i+1}^{(p)}\|} \times \frac{\mathbf{x}_i^{(p)} \times \mathbf{x}_{i+1}^{(p)}}{\|\mathbf{x}_i^{(p)} \times \mathbf{x}_{i+1}^{(p)}\|} \right). \quad (4)$$

Through this complex-level vector frame, the message passing enables the exchange of information between atoms from ligands and residues from the pockets. The illustration of the above three levels of vector frames can be found in Figure 3. Once we build up such three vector frames, we then conduct a *scalarization* operation (Hsu, 2002), which transforms the equivariant variables (*e.g.*, coordinates) to invariant variables by projecting them to the three vector bases in the vector frame.

3.2 MULTI-GRAINED SE(3)-EQUIVARIANT BINDING FORCE MODELING: BINDINGNET

In this section, we introduce BindingNet, a multi-grained SE(3)-equivariant geometric model for protein-ligand binding. The input of BindingNet is the geometry of the rigid protein and the ligand at time t , while the output is the force on each atom in the ligand.

Atom-Level Ligand Modeling. We first generate the atom embedding using one-hot encoding and then aggregate each atom’s embedding, $\mathbf{z}^{(l)}$, by aggregating all its neighbor’s embedding within the cutoff distance c . Then, we obtain the atom’s equivariant representation by aggregating its neighborhood’s messages as $(\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}) \cdot \mathbf{z}_i^{(l)}$. A subsequent scalarization is carried out based on the atom-level vector frame as $\mathbf{h}_{ij}^{(l)} = (\mathbf{h}_i^{(l)} \oplus \mathbf{h}_j^{(l)}) \cdot \mathcal{F}_{\text{ligand}}$, where \oplus is the concatenation. Finally, it is passed through several equivariant message-passing layers (MPNN) defined as:

$$\text{vec}_i^{(l)} = \text{vec}_i^{(l)} + \text{Agg}_j(\text{vec}_i^{(l)} \cdot \text{MLP}(\mathbf{h}_{ij}) + (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(p)}) \cdot \text{MLP}(\mathbf{h}_{ij})), \quad (5)$$

where $\text{MLP}(\cdot)$ and $\text{Agg}(\cdot)$ are the multi-layer perceptron and mean aggregation functions, respectively. $\text{vec} \in \mathbb{R}^3$ is a vector assigned to each atom and is initialized as 0. The outputs are atom representation and vector $(\mathbf{h}^{(l)}$ and $\text{vec}^{(l)})$, and they are passed to the complex module.

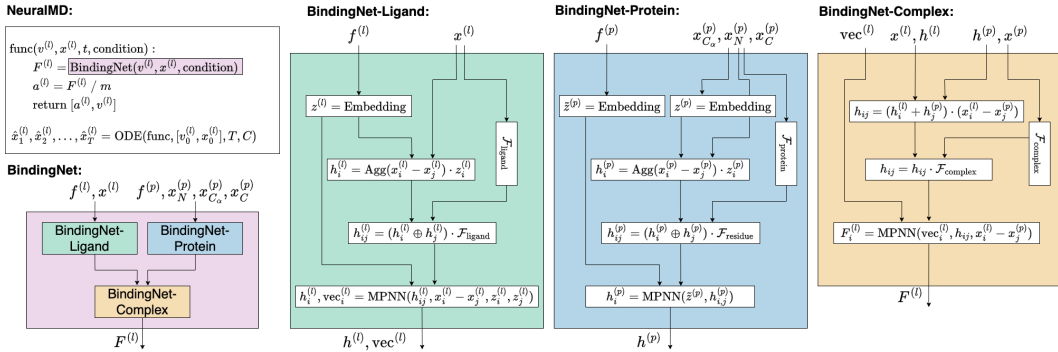


Figure 4: Brief pipeline of NeuralMD. In the three key modules of BindingNet, there are three vertical boxes, corresponding to three granularities of vector frames, as in Equations (2) to (4). More details are in Appendix E.

Backbone-Level Protein Modeling. For the coarse-grained modeling of proteins, we consider three backbone atoms in each residue. We first obtain the atom embedding on three atom types, and then we obtain each atom’s representation $\mathbf{z}^{(p)}$ by aggregating its neighbor’s representation. Then, we obtain an equivariant atom representation by aggregating the edge information, $(\mathbf{x}_i^{(p)} - \mathbf{x}_j^{(p)}) \cdot \mathbf{z}_i^{(p)}$, within cutoff distance c . Following which is the scalarization on the residue frame $\mathbf{h}_{ij}^{(p)} = (\mathbf{h}_i^{(p)} \oplus \mathbf{h}_j^{(p)}) \cdot \mathcal{F}_{\text{protein}}$. Recall that we also have the residue type, and with a type embedding $\tilde{\mathbf{z}}^{(p)}$, we can obtain the final residue-level representation using an MPNN layer as $\mathbf{h}^{(p)} = \tilde{\mathbf{z}}^{(p)} + (\mathbf{h}_{N, C_\alpha}^{(p)} + \mathbf{h}_{C_\alpha, C}^{(p)}) / 2$.

Residue-Level Complex Modeling. Once we obtain the atom-level representation and vector $(\mathbf{h}^{(l)}, \text{vec}^{(l)})$ from ligands, and backbone-level representation $(\mathbf{h}^{(p)})$ from proteins, the next step is to learn the protein-ligand interaction. We first extract the residue-atom pair (i, j) with a cutoff c , based on which we obtain an equivariant interaction edge representation $\mathbf{h}_{ij} = (\mathbf{h}_i^{(l)} + \mathbf{h}_j^{(p)}) \cdot (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(p)})$. After scalarization, we can obtain invariant interaction edge representation $\mathbf{h}_{ij} = \mathbf{h}_{ij} \cdot \mathcal{F}_{\text{complex}}$. Finally, we adopt an equivariant MPNN layer to get the atom-level force as:

$$\text{vec}_{ij}^{(pl)} = \text{vec}_i^{(l)} \cdot \text{MLP}(\mathbf{h}_{ij}) + (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(p)}) \cdot \text{MLP}(\mathbf{h}_{ij}), \quad F_i^{(l)} = \text{vec}_i^{(l)} + \text{Agg}_{j \in \mathcal{N}(i)} \text{vec}_{ij}^{(pl)}. \quad (6)$$

In the last equation, the ultimate force predictions can be viewed as two parts: the internal force from the molecule $\text{vec}_i^{(l)}$ and the external force from the protein-ligand interaction $\text{vec}_{ij}^{(pl)}$. More complete descriptions of the three modules can be found in Appendix E.

3.3 BINDING MOLECULAR DYNAMICS MODELING: NEURALMD

As clarified in Section 2, molecular dynamics follows Newtonian dynamics, and we solve it as an ordinary differential equation (ODE) problem. The BindingNet introduced in Sections 3.1 and 3.2 takes in the molecular system geometry $(\mathbf{x}_t^{(l)}, \mathbf{x}^{(p)})$ at arbitrary time t , and outputs the forces. Then in this section, we describe how we use the neural ODE solver to predict the coordinates at future snapshots.

We want to highlight that one research line is predicting the force field directly (Zhang et al., 2018; Doerr et al., 2020; Musaelian et al., 2023b), which will be fed into the numerical integration algorithms for trajectory simulation. For accuracy, such a simulation must be at the femtosecond level (1e-15). However, as shown in recent works (Unke et al., 2021; Stocker et al., 2022; Fu et al., 2022a), minor errors in the force field can lead to catastrophic failure for long-time simulations. For instance, there can be pathological behavior, such as extreme force predictions or bond breaking, within the low end of the distribution. Our experiments have yielded similar observations, as will be shown in Section 4. In this paper, however, we overcome this issue by directly learning the extended-timescale MD trajectories (nanosecond level, 1e-9).

On the other hand, recent explorations on automatic differentiation have opened new perspectives for solving the ordinary/partial differential equations (ODE/PDEs) (Chen et al., 2018; Raissi et al., 2019; Lu et al., 2019; 2021; Li et al., 2020; Kovachki et al., 2021). Newtonian mechanics in MD is essentially an ODE, and for protein-ligand complexes with a large volume of atoms, we consider using adjoint neural ODE (NODE) to save memory costs (Chen et al., 2018), allowing small step sizes

in the ODE integration. Roughly speaking, the adjoint method calculates the variational derivative of the trajectory w.r.t. our objective function directly. Thus, it avoids calling the auto-differentiation for all intermediate time intervals. We further want to highlight that we are working on a second-order ODE in Newtonian dynamics. The key module of NODE is the ODE function, which returns the first-order derivative. Then, it will be integrated using algorithms like Euler. To learn the MD trajectory following second-order ODE, we propose the following formulation of the second-order ODE within one integration call:

$$\begin{bmatrix} dx/dt \\ dv/dt \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ F/m \end{bmatrix}, \quad (7)$$

where F is the output forces from BindingNet. This means we augment ODE derivative space by concurrently calculating the accelerations and velocities, allowing simultaneous integration of velocities and positions. Ultimately, following Newtonian mechanics, the coordinates at time t are integrated as:

$$\begin{aligned} F_\tau^{(l)} &= \text{BindingNet}(f^{(l)}, \mathbf{x}_\tau^{(l)}, f^{(p)}, \mathbf{x}_N^{(p)}, \mathbf{x}_{C_\alpha}^{(p)}, \mathbf{x}_C^{(p)}), & \mathbf{a}_\tau^{(l)} &= \frac{F_\tau^{(l)}}{m}, \\ \hat{\mathbf{v}}_t^{(l)} &= \mathbf{v}_0^{(l)} + \int_0^t \mathbf{a}_\tau^{(l)} d\tau, & \hat{\mathbf{x}}_t^{(l)} &= \mathbf{x}_0^{(l)} + \int_0^t \hat{\mathbf{v}}_\tau^{(l)} d\tau. \end{aligned} \quad (8)$$

The objective is the mean absolute error between the predicted coordinates and ground-truth coordinates: $\mathcal{L} = \mathbb{E}_t [|\hat{\mathbf{x}}_t^{(l)} - \mathbf{x}_t^{(l)}|]$. An illustration of NeuralMD pipeline is in Figure 4.

From Newtonian dynamics to Langevin dynamics. It is worth noting that the protein-ligand binding dynamics in the MISATO dataset considers the solvent molecules. However, the released dataset does not include the trajectory information of solvent molecules. To compensate such missing information, we leverage Langevin dynamics. Specifically, Langevin dynamics is an extension of the standard Newtonian dynamics with the addition of damping and random noise terms: $\hat{F}_\tau^{(l)} = \text{BindingNet}(f^{(l)}, \mathbf{x}_\tau^{(l)}, f^{(p)}, \mathbf{x}_N^{(p)}, \mathbf{x}_{C_\alpha}^{(p)}, \mathbf{x}_C^{(p)}) - \gamma m \mathbf{v} + \sqrt{2m\gamma k_B T} R(t)$, where γ is the damping constant or collision frequency, T is the temperature, k_B is the Boltzmann’s constant, and $R(t)$ is a delta-correlated stationary Gaussian process with zero-mean. This modified $\hat{F}_\tau^{(l)}$ is then plugged into Equation (8) to train and sample the trajectories, thus compensating the missing solvent effect.

3.4 IMPLEMENTATION DETAILS AND SUMMARY

In this section, we would like to include extra details of NeuralMD. As in Equations (7) and (8), both the coordinates and velocities are required inputs for NeuralMD. Unfortunately, certain datasets may not cover such information. To handle this issue, we propose a surrogate velocity, which is a summation of a predicted velocity by an extra equivariant model and a coordinate momentum, *i.e.*, $\mathbf{v}_t^{(l)} = \text{BindingNet-Ligand}(f^{(l)}, \mathbf{x}_t^{(l)}) + (\mathbf{x}_{t+1}^{(l)} - \mathbf{x}_t^{(l)})$, where the BindingNet-Ligand is the ligand model described in Section 3.2. This is not displayed in Figure 4 for brevity, yet we provide all the critical details for reproducibility. The algorithm of our method is presented in Algorithm 1.

Algorithm 1 Training of NeuralMD

- 1: **Input:** Initial position $\mathbf{x}_0^{(l)}$ and initial velocity $\mathbf{v}_0^{(l)}$ for ligands, atomic features $f^{(l)}$ for ligands, residue types and coordinates $f^{(p)}, \mathbf{x}_N^{(p)}, \mathbf{x}_{C_\alpha}^{(p)}, \mathbf{x}_C^{(p)}$ for proteins, and time T .
 - 2: **for** discretized time $t \in \{1, 2, \dots, T - 1\}$ **do**
 - 3: Centralize the coordinates of the ligand-protein complex for $\mathbf{x}_0^{(l)}$ and $\mathbf{x}^{(p)}$ by removing the mass center.
 - 4: Set T_list = $[t, t + 1, \dots]$ and condition $C = [f^{(l)}, f^{(p)}, \mathbf{x}_N^{(p)}, \mathbf{x}_{C_\alpha}^{(p)}, \mathbf{x}_C^{(p)}]$.
 - 5: Get predicted position $[\hat{\mathbf{x}}_{t+1}^{(l)}, \hat{\mathbf{x}}_{t+2}^{(l)}, \dots] = \text{ODE}(\text{NeuralMD}, [\mathbf{x}_0^{(l)}, \mathbf{v}_0^{(l)}, \text{T_list}, C])$. // Equation (7)
 - 6: Calculate the position prediction loss $\mathcal{L} = \mathbb{E}_\tau [\text{MAE}(\hat{\mathbf{x}}_\tau^{(l)}, \mathbf{x}_\tau^{(l)})]$.
 - 7: Update model BindingNet using gradient descent.
 - 8: **end for**
-

4 EXPERIMENTS

4.1 EXPERIMENT SETTING

Datasets. As described in Appendix A, one of the main bottlenecks of ML for binding dynamics simulation is the available datasets for learning molecular dynamics (MD). Recently, the community

has put more effort into gathering the datasets, and we consider MISATO in our work (Siebenmorgen et al., 2023). It is built on 16,972 experimental protein-ligand complexes extracted from the protein data bank (PDB) (Berman et al., 2000). Such data is obtained using X-ray crystallography, Nuclear Magnetic Resonance (NMR), or Cryo-Electron Microscopy (Cryo-EM), where systematic errors are unavoidable. This motivates the MISATO project, which utilizes semi-empirical quantum mechanics for structural curation and refinement, including regularization of the ligand geometry. For each protein-ligand complex, the trajectory comprises 100 snapshots in 8 nanoseconds under the fixed temperature and pressure. We want to highlight that MD trajectories allow the analysis of small-range structural fluctuations of the protein-ligand complex. Still, in some cases, large-scale rare events can be observed. In Appendix D, we list several basic statistics of MISATO, *e.g.*, the number of atoms in small molecule ligands and the number of residues in proteins.

Baselines. In this work, we mainly focus on machine learning methods for trajectory prediction, *i.e.*, no energy or force labels are considered. GNN-MD is to apply geometric graph neural networks (GNNs) to predict the trajectories in an auto-regressive manner (Siebenmorgen et al., 2023; Fu et al., 2022b). More concretely, GNN-MD takes as inputs the geometries at time t and predicts the geometries at time $t + 1$. DenoisingLD (denoising diffusion for Langevin dynamics) (Arts et al., 2023; Wu & Li, 2023; Fu et al., 2022b) is a baseline method that models the trajectory prediction as denoising diffusion task (Song et al., 2020), and the inference for trajectory generation essentially becomes the Langevin dynamics. CG-MD learns a dynamic GNN and a score GNN (Fu et al., 2022b), which are essentially the hybrid of GNN-MD and DenoisingLD. Here, to make the comparison more explicit, we compare these two methods (GNN-MD and DenoisingLD) separately. Additionally, we consider VerletMD, an energy prediction research line (including DeePMD (Zhang et al., 2018), TorchMD (Doerr et al., 2020), and Allegro-LAMMPS (Musaelian et al., 2023b)), where the role of ML models is to predict the energy, and the MD trajectory is obtained by the velocity Verlet algorithm, a numerical integration method for Newtonian mechanics. We keep the same backbone model (BindingNet) for energy or force prediction for all the baselines.

Experiments Settings. We consider two experiment settings. The first type of experiment is the single-trajectory prediction, where both the training and test data are snapshots from the same trajectory, and they are divided temporally. The second type of experiment is the multi-trajectory prediction, where each data point is the sequence of all the snapshots from one trajectory, and the training and test data correspond to different sets of trajectories.

Evaluation Metrics. For MD simulation, the evaluation is a critical factor for evaluating trajectory prediction (Fu et al., 2022a). For both experiment settings, the trajectory recovery is the most straightforward evaluation metric. To evaluate this, we take both the mean absolute error (MAE) and mean squared error (MSE) between the predicted coordinates and ground-truth coordinates **over all snapshots**. Stability, as highlighted in (Fu et al., 2022a), is an important metric for evaluating the predicted MD trajectory. The intuition is that the prediction on long-time MD trajectory can enter a pathological state (*e.g.*, bond breaking), and stability is the measure to quantify such observation. It is defined as $\mathbb{P}_{i,j} \|\mathbf{x}_i - \mathbf{x}_j\| - \mathbf{b}_{i,j}\| > \Delta$, where $\mathbf{b}_{i,j}$ is the pair distance at the last snapshot (the most equilibrium state), and we take $\Delta = 0.5 \text{ \AA}$. Another metric considered is frames per second (FPS) (Fu et al., 2022a) on a single Nvidia-V100 GPU card, and it measures the MD efficiency.

Ablation Study on Flexible Setting. As introduced in Section 2, we have been focusing on the semi-flexible setting for binding dynamics so far. Yet, we also conduct ablation studies on the flexible setting for small-scale experiments, as a proof-of-concept. For more details, please check Appendix F.

4.2 MD PREDICTION: GENERALIZATION ON ONE SINGLE TRAJECTORY

This type of task has been widely studied in the existing literature for other molecular systems. Specifically, both the training and test data are from the same trajectory of one single protein-ligand complex, and here we take the first 80 snapshots for training and the remaining 20 snapshots for test.

Results are in Table 1. The first observation is that the baseline VerletMD has a clear performance gap compared to the other methods. This verifies that using ML models to predict the energy (or force) at each snapshot, and then using a numerical integration algorithm can fail in the long-time simulations (Fu et al., 2022a). Additionally, we can observe that the recovery error of trajectory (MAE and MSE) occasionally fails to offer a discernible distinction among methods (*e.g.*, for protein-ligand complex 3EOV, 1KT1, and 4G3E), though NeuralMD is slightly better. However, the stability (%) can

Table 1: Results on ten single-trajectory binding dynamics prediction. Results with optimal training loss are reported. Four evaluation metrics are considered: MAE (\AA , \downarrow), MSE (\downarrow), and Stability (% , \downarrow).

PDB ID	Metric	VerletMD	GNN-MD	DenosingLD	NeuralMD ODE (Ours)	NeuralMD SDE (Ours)	PDB ID	Metric	VerletMD	GNN-MD	DenosingLD	NeuralMD ODE (Ours)	NeuralMD SDE (Ours)
5WIJ	MAE	9.618	2.319	2.254	2.118	2.109	1XP6	MAE	13.444	2.303	1.915	1.778	1.822
	MSE	6.401	1.553	1.502	1.410	1.408		MSE	9.559	1.505	1.282	1.182	1.216
	Stability	79.334	45.369	18.054	12.654	13.340		Stability	86.393	43.019	28.417	19.256	22.734
4ZX0	MAE	21.033	2.255	1.998	1.862	1.874	4YUR	MAE	15.674	7.030	6.872	6.807	6.826
	MSE	14.109	1.520	1.347	1.260	1.271		MSE	10.451	4.662	4.520	4.508	4.526
	Stability	76.878	41.332	23.267	18.189	18.845		Stability	81.309	50.238	32.423	23.250	25.008
3EOV	MAE	25.403	3.383	3.505	3.287	3.282	4G3E	MAE	5.181	2.672	2.577	2.548	2.478
	MSE	17.628	2.332	2.436	2.297	2.294		MSE	3.475	1.743	1.677	1.655	1.615
	Stability	91.129	57.363	51.590	44.775	44.800		Stability	65.377	16.365	7.188	2.113	2.318
4K6W	MAE	14.682	3.674	3.555	3.503	3.429	6B7F	MAE	31.375	4.129	3.952	3.717	3.657
	MSE	9.887	2.394	2.324	2.289	2.234		MSE	21.920	2.759	2.676	2.503	2.469
	Stability	87.147	57.852	39.580	38.562	38.476		Stability	87.550	54.900	16.050	3.625	22.750
1KTI	MAE	18.067	6.534	6.657	6.548	6.537	3B9S	MAE	19.347	2.701	2.464	2.351	2.374
	MSE	12.582	4.093	4.159	4.087	4.085		MSE	11.672	1.802	1.588	1.527	1.542
	Stability	77.315	4.691	7.377	0.525	0.463		Stability	41.667	43.889	8.819	0.000	0.000

Table 2: Efficiency comparison of FPS between VerletMD and NeuralMD on single-trajectory prediction.

PDB ID	5WIJ	4ZX0	3EOV	4K6W	1KTI	1XP6	4YUR	4G3E	6B7F	3B9S	Average
VerletMD	12.564	30.320	29.890	26.011	19.812	28.023	31.513	29.557	19.442	31.182	25.831
NeuralMD (Ours)	33.164	39.415	31.720	31.909	24.566	37.135	39.365	39.172	20.320	37.202	33.397

Table 3: Results on three multi-trajectory binding dynamics predictions. Results with optimal validation loss are reported. Four evaluation metrics are considered: MAE (\AA , \downarrow), MSE (\downarrow), and Stability (% , \downarrow).

Dataset	MISATO-100			MISATO-1000			MISATO-All		
	MAE	MSE	Stability	MAE	MSE	Stability	MAE	MSE	Stability
VerletMD	90.326	56.913	86.642	80.187	53.110	86.702	105.979	69.987	90.665
GNN-MD	7.176	4.726	35.431	7.787	5.118	33.926	8.260	5.456	32.638
DenosingLD	7.112	4.684	29.956	7.746	5.090	18.898	15.878	10.544	89.586
NeuralMD-ODE (Ours)	6.852	4.503	19.173	7.653	5.028	15.572	8.147	5.386	17.468
NeuralMD-SDE (Ours)	6.869	4.514	19.561	7.665	5.037	16.501	8.165	5.398	19.012

be a distinctive factor in method comparisons, where we observe NeuralMD outperform on all 10 tasks up to $\sim 80\%$. We further qualitatively show the predicted trajectories of three case studies using three methods and ground truth in Figure 1. It is shown that the GNN-MD collapses occasionally, while DenosingLD stays comparatively structured. Meanwhile, NeuralMD is the most stable in all cases.

One main benefit of using NeuralMD for binding simulation is its efficiency. To show this, we list the computational time in Table 2. We further approximate the wall time of the numerical method for MD simulation (PDB 5WIJ). Concretely, we can get an estimated speed of 1 nanosecond of dynamics every 0.28 hours. This is running the simulation with GROMACS (Van Der Spoel et al., 2005) on 1 GPU with 16 CPU cores and a moderately sized water box at the all-atom level (with 2 femtosecond timesteps). This equivalently shows that NeuralMD is $\sim 2000\times$ faster than numerical methods.

4.3 MD PREDICTION: GENERALIZATION AMONG MULTIPLE TRAJECTORIES

A more challenging task is to test the generalization ability of NeuralMD among different trajectories. The MISATO dataset includes 13,765 protein-ligand complexes, and we first create two small datasets by randomly sampling 100 and 1k complexes, respectively. Then, we take 80%-10%-10% for training, validation, and testing. We also consider the whole MISATO dataset, where the data split has already been provided. After removing the peptide ligands, we have 13,066, 1,357, and 1,357 complexes for training, validation, and testing, respectively.

The quantitative results are in Table 3. First, we can observe that VerletMD has worse performance on all three datasets, and the performance gap with other methods is even larger compared to the single-trajectory prediction. The other two baselines, GNN-MD and DenosingLD, show similar performance, while NeuralMD outperforms in all datasets. Notice that stability (%) remains more distinguishable than the two trajectory recovery metrics (MAE and MSE).

5 CONCLUSION AND OUTLOOK

To sum up, we devise NeuralMD, an ML framework that incorporates a novel multi-grained group symmetric network architecture and second-order ODE Newtonian dynamics, enabling accurate predictions of protein-ligand binding dynamics in a larger time interval. Not only is such a time

interval critical for understanding the dynamic nature of the ligand-protein complex, but our work marks the first approach to learning to predict binding dynamics. We quantitatively and qualitatively verify that NeuralMD achieves superior performance on 13 binding prediction tasks.

One potential limitation of our work is the dataset. Currently, we are using the MISATO dataset, a binding simulation dataset with a large timescale. However, NeuralMD is agnostic to the time interval, and it can also be applied to binding dynamics datasets with time interval as a femtosecond. This is beyond the scope of this work and may require the effort of the whole community for the dataset construction.

One interesting future direction is as follows. Recall that BindingNet only explicitly includes the atom-level information, and no quantum features are considered. Thus, one promising direction is to incorporate the quantum mechanisms into MD. This challenge necessitates deeper integration of quantum physics knowledge into the model and algorithm design, a task we defer for future exploration.

Authors: Please allow a few lines extension during rebuttal. We will fix this in the final version.

REFERENCES

- Marloes Arts, Victor Garcia Satorras, Chin-Wei Huang, Daniel Zugner, Marco Federici, Cecilia Clementi, Frank Noe, Robert Pinsler, and Rianne van den Berg. Two for one: Diffusion models and force fields for coarse-grained molecular dynamics. *Journal of Chemical Theory and Computation*, 2023. 2, 3, 8, 16
- Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1): 235–242, 2000. 8
- J Ceriotti. More, and de manolopoulos. *Comput. Phys. Commun*, 185:1019, 2014. 15
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. 4, 6
- Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *International Conference on Learning Representations (ICLR)*, 2023. 2, 4, 15
- Stefan Doerr, Maciej Majewsk, Adrià Pérez, Andreas Krämer, Cecilia Clementi, Frank Noe, Toni Giorgino, and Gianni De Fabritiis. Torchmd: A deep learning framework for molecular simulations, 2020. 2, 6, 8, 15, 16
- Weitao Du, He Zhang, Yuanqi Du, Qi Meng, Wei Chen, Nanning Zheng, Bin Shao, and Tie-Yan Liu. Se (3) equivariant graph neural networks with complete local frames. In *International Conference on Machine Learning*, pp. 5583–5608. PMLR, 2022. 15
- Hehe Fan, Zhangyang Wang, Yi Yang, and Mohan Kankanhalli. Continuous-discrete convolution for geometry-sequence modeling in proteins. In *The Eleventh International Conference on Learning Representations*, 2022. 15
- Hans Frauenfelder, Guo Chen, Joel Berendzen, Paul W Fenimore, Helén Jansson, Benjamin H McMahon, Izabela R Stroe, Jan Swenson, and Robert D Young. A unified model of protein dynamics. *Proceedings of the National Academy of Sciences*, 106(13):5129–5134, 2009. 16
- Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Academic Press San Diego, 2002. 16
- Xiang Fu, Zhenghao Wu, Wujie Wang, Tian Xie, Sinan Ketten, Rafael Gomez-Bombarelli, and Tommi Jaakkola. Forces are not enough: Benchmark and critical evaluation for machine learning force fields with molecular simulations. *arXiv preprint arXiv:2210.07237*, 2022a. 1, 2, 6, 8
- Xiang Fu, Tian Xie, Nathan J Rebello, Bradley D Olsen, and Tommi Jaakkola. Simulate time-integrated coarse-grained molecular dynamics with geometric machine learning. *arXiv preprint arXiv:2204.10348*, 2022b. 2, 3, 8, 16
- Joe G Greener and David T Jones. Differentiable molecular simulation can learn all the parameters in a coarse-grained force field for proteins. *PloS one*, 16(9):e0256990, 2021. 16
- Jiaqi Guan, Wesley Wei Qian, Xingang Peng, Yufeng Su, Jian Peng, and Jianzhu Ma. 3d equivariant diffusion for target-aware molecule generation and affinity prediction. *arXiv preprint arXiv:2303.03543*, 2023. 15
- Elton P Hsu. *Stochastic analysis on manifolds*. Number 38. American Mathematical Soc., 2002. 5
- John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-based protein design. *Advances in neural information processing systems*, 32, 2019. 15
- José Jiménez, Miha Skalic, Gerard Martinez-Rosell, and Gianni De Fabritiis. K deep: protein–ligand absolute binding affinity prediction via 3d-convolutional neural networks. *Journal of chemical information and modeling*, 58(2):287–296, 2018. 2, 16

- Derek Jones, Hyojin Kim, Xiaohua Zhang, Adam Zemla, Garrett Stevenson, WF Drew Bennett, Daniel Kirshner, Sergio E Wong, Felice C Lightstone, and Jonathan E Allen. Improved protein–ligand binding affinity prediction with structure-based deep fusion inference. *Journal of chemical information and modeling*, 61(4):1583–1592, 2021. 2, 16
- 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. 1, 2, 15
- Visvaldas Kairys, Lina Baranauskiene, Migle Kazlauskiene, Daumantas Matulis, and Egidijus Kazlauskas. Binding affinity in drug design: experimental and computational techniques. *Expert opinion on drug discovery*, 14(8):755–768, 2019. 1
- Martin Karplus and John Kuriyan. Molecular dynamics and protein function. *Proceedings of the National Academy of Sciences*, 102(19):6679–6685, 2005. 16
- Johannes Klicpera, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint arXiv:2011.14115*, 2020. 2, 15
- Divya B Korlepara, CS Vasavi, Shruti Jeurkar, Pradeep Kumar Pal, Subhajit Roy, Sarvesh Mehta, Shubham Sharma, Vishal Kumar, Charuvaka Muvva, Bhuvanesh Sridharan, et al. Plas-5k: Dataset of protein-ligand affinities from molecular dynamics for machine learning applications. *Scientific Data*, 9(1):548, 2022. 16
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces. *arXiv preprint arXiv:2108.08481*, 2021. 6
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*, 2020. 6
- Shengchao Liu, Weitao Du, Yanjing Li, Zhuoxinran Li, Zhiling Zheng, Chenru Duan, Zhiming Ma, Omar Yaghi, Anima Anandkumar, Christian Borgs, Jennifer Chayes, Hongyu Guo, and Jian Tang. Symmetry-informed geometric representation for molecules, proteins, and crystalline materials. *arXiv preprint arXiv:2306.09375*, 2023a. 2, 3, 15
- Shengchao Liu, Weitao Du, Zhi-Ming Ma, Hongyu Guo, and Jian Tang. A group symmetric stochastic differential equation model for molecule multi-modal pretraining. In *International Conference on Machine Learning*, pp. 21497–21526. PMLR, 2023b. 15
- Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019. 6
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021. 6
- Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *arXiv preprint arXiv:2204.05249*, 2022. 2, 15
- Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023a. 15
- Albert Musaelian, Anders Johansson, Simon Batzner, and Boris Kozinsky. Scaling the leading accuracy of deep equivariant models to biomolecular simulations of realistic size. *arXiv preprint arXiv:2304.10061*, 2023b. 2, 6, 8, 15, 16

- Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In *International Conference on Machine Learning*, pp. 17644–17655. PMLR, 2022. 15
- U Deva Priyakumar, Divya B Korlepara, CS Vasavi, Rakesh Srivastava, Pradeep Kumar Pal, Saalim H Raza, Vishal Kumar, Shivam Pandit, Aathira G Nair, Sanjana Pandey, et al. Plas-20k: Extended dataset of protein-ligand affinities from md simulations for machine learning applications. 2023. 16
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. 6
- Dennis C Rapaport. *The art of molecular dynamics simulation*. Cambridge university press, 2004. 16
- Veronica Salmaso and Stefano Moro. Bridging molecular docking to molecular dynamics in exploring ligand-protein recognition process: An overview. *Frontiers in pharmacology*, 9:923, 2018. 4
- Victor Garcia Satorras, Emiel Hooeboom, and Max Welling. E(n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021. 15
- Kristof T Schütt, Huziel E Saucedo, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. SchNet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24):241722, 2018. 2, 15
- Kristof T Schütt, Oliver T Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. *arXiv preprint arXiv:2102.03150*, 2021. 15
- Till Siebenmorgen, Filipe Menezes, Sabrina Benassou, Erinc Merdivan, 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. *bioRxiv*, pp. 2023–05, 2023. 1, 3, 4, 8, 16, 21
- Tess Smidt, Nathaniel Thomas, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018. 2, 15
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 8, 16
- Marta M Stepniewska-Dziubinska, Piotr Zielenkiewicz, and Pawel Siedlecki. Development and evaluation of a deep learning model for protein–ligand binding affinity prediction. *Bioinformatics*, 34(21):3666–3674, 2018. 2, 16
- Sina Stocker, Johannes Gasteiger, Florian Becker, Stephan Günnemann, and Johannes T Margraf. How robust are modern graph neural network potentials in long and hot molecular dynamics simulations? *Machine Learning: Science and Technology*, 3(4):045010, 2022. 6
- Philipp Tholke and Gianni De Fabritiis. Equivariant transformers for neural network based molecular potentials. In *International Conference on Learning Representations*, 2022. 15
- Aidan P Thompson, H Metin Aktulga, Richard Berger, Dan S Bolinteanu, W Michael Brown, Paul S Crozier, Pieter J in’t Veld, Axel Kohlmeyer, Stan G Moore, Trung Dac Nguyen, et al. LAMMPS—a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications*, 271:108171, 2022. 15
- Oliver T Unke, Stefan Chmiela, Huziel E Saucedo, Michael Gastegger, Igor Poltavsky, Kristof T Schütt, Alexandre Tkatchenko, and Klaus-Robert Müller. Machine learning force fields. *Chemical Reviews*, 121(16):10142–10186, 2021. 6

- Sandor Vajda, Dmitri Beglov, Amanda E Wakefield, Megan Egbert, and Adrian Whitty. Cryptic binding sites on proteins: definition, detection, and druggability. *Current opinion in chemical biology*, 44:1–8, 2018. [1](#)
- David Van Der Spoel, Erik Lindahl, Berk Hess, Gerrit Groenhof, Alan E Mark, and Herman JC Berendsen. Gromacs: fast, flexible, and free. *Journal of computational chemistry*, 26(16):1701–1718, 2005. [9](#)
- Mikhail Volkov, Joseph-André Turk, Nicolas Drizard, Nicolas Martin, Brice Hoffmann, Yann Gaston-Mathé, and Didier Rognan. On the frustration to predict binding affinities from protein–ligand structures with deep neural networks. *Journal of medicinal chemistry*, 65(11):7946–7958, 2022. [1](#)
- Fang Wu and Stan Z Li. Diffmd: A geometric diffusion model for molecular dynamics simulations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5321–5329, 2023. [2](#), [3](#), [4](#), [8](#), [16](#)
- Jincai Yang, Cheng Shen, and Niu Huang. Predicting or pretending: artificial intelligence for protein–ligand interactions lack of sufficiently large and unbiased datasets. *Frontiers in pharmacology*, 11:69, 2020. [1](#)
- Ziduo Yang, Weihe Zhong, Qiuji Lv, Tiejun Dong, and Calvin Yu-Chian Chen. Geometric interaction graph neural network for predicting protein–ligand binding affinities from 3d structures (gign). *The Journal of Physical Chemistry Letters*, 14(8):2020–2033, 2023. [2](#), [16](#)
- Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and EJPRL Weinan. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Physical review letters*, 120(14):143001, 2018. [2](#), [6](#), [8](#), [15](#), [16](#)
- Xuan Zhang, Limei Wang, Jacob Helwig, Youzhi Luo, Cong Fu, Yaochen Xie, Meng Liu, Yuchao Lin, Zhao Xu, Keqiang Yan, et al. Artificial intelligence for science in quantum, atomistic, and continuum systems. *arXiv preprint arXiv:2307.08423*, 2023. [2](#), [15](#)

REPRODUCIBILITY STATEMENT

In Sections 3 and 4, we have included the key modules of NeuralMD, and we list more details on the model architecture and hyperparameters in Appendix E. We will release the codes, checkpoints, and log files after the paper acceptance.

A RELATED WORKS AND PRELIMINARIES

SE(3)-Equivariant Representation for Small Molecules and Proteins. The molecular systems are indeed a set of atoms located in the 3D Euclidean space. From a machine learning point of view, the representation function or encoding function of such molecular systems needs to be group-symmetric, *i.e.*, the representation needs to be equivariant when we rotate or translate the whole system. Such symmetry is called the SE(3)-equivariance. Recently works (Liu et al., 2023a; Zhang et al., 2023) on molecules has provided a unified way of equivariant geometric modeling. They categorize the mainstream representation methods into three big venues: SE(3)-invariant models, SE(3)-equivariant models with spherical frame basis, and SE(3)-equivariant models with vector frame basis. (1) Invariant models that utilize invariant features (distances and angles) to predict the energies (Schütt et al., 2018; Klicpera et al., 2020), but the derived forces are challenging for ML optimization after integration. (2) Equivariant models with spherical frames that include a computationally expensive tensor-product operation (Smidt et al., 2018; Musaelian et al., 2022), which is unsuitable for large molecular systems. (3) Equivariant models with vector frames that have been explored for single stable molecules, including molecule representation and pretraining (Satorras et al., 2021; Schütt et al., 2021; Liu et al., 2023b), molecule conformation generation (Du et al., 2022), protein representation (Fan et al., 2022), and protein folding and design (Ingraham et al., 2019; Jumper et al., 2021). However, no one has tried it for binding complexes.

Structure-Based Ligand Pose Generation in Equilibrium State. One seemingly related but fundamentally different research line is the structure-based ligand pose generation. The core idea is to generate the ligand pose based on the protein structure. Statistical generative methods like autoregressive Peng et al. (2022) and denoising diffusion Corso et al. (2023); Guan et al. (2023) have been applied to this task. We want to highlight that these methods generate a trajectory, but they are essentially different for two reasons: (1) These works consider only the conformations at the equilibrium state, and no dynamics information is considered. (2) The inference of denoising diffusion generation is in the form of Monte Carlo (MC), which formulates a trajectory for each sampled sample. However, such trajectories include no physical information (*e.g.*, energy, and force), while the movements of molecules follow the Newtonian mechanics (*e.g.*, force fields) in MD simulation.

Potential Energy and Force Learning for Molecular Dynamics Simulation. One straightforward way of molecular dynamics (MD) simulation is through potential energy modeling. *Numerical methods* for MD simulation can be classified into classical MD and *ab-initio* MD, depending on using the classical mechanism or quantum mechanism to calculate the forces. Alternatively, a machine learning (ML) research line is to adopt geometric representation methods to learn the energies or the forces, *e.g.*, by the geometric methods listed above. The first work is DeePMD (Zhang et al., 2018), which targets learning the potential energy function at each conformation. For inference, the learned energy can be applied to update the atom placement using i-PI software Ceriotti (2014), composing the MD trajectories. TorchMD (Doerr et al., 2020) utilizes TorchMD-Net (Tholke & Fabritiis, 2022) for energy prediction, which will be fed into the velocity Verlet algorithm for MD simulation. Similarly, Musaelian et al. adopts Allegro (Musaelian et al., 2023a) model to learn the force at each conformation. The learned model will be used for MD trajectory simulation using LAMMPS Thompson et al. (2022). In theory, all the geometric models on small molecules Satorras et al. (2021); Schütt et al. (2021); Liu et al. (2023b) and proteins Fan et al. (2022) can be applied to the MD simulation task. However, there are two main challenges: (1) They require the interval between snapshots to be at the femtoseconds (1e-15 seconds) level for integration to effectively capture the motion of the molecules. (2) They take the position-energy pairs independently, and thus, they ignore their temporal correlations during learning.

Trajectory Learning for Molecular Dynamics Estimation. More recent works have explored MD simulation by directly learning the placement along the trajectories. There are two key differences

Table 4: Comparison of different numerical and machine learning (ML) methods for molecular dynamics (MD). AR for autoregressive and denoising for denoising diffusion method.

Category	Method	Energy / Force Calculation	Dynamics	Objective Function	Publications
Numerical Methods	Classical MD	Classical Mechanics: Force Field	Newtonian Dynamics	-	-
	<i>Ab-initio</i> MD	Quantum Mechanics: DFT for Schrodinger Equation	Newtonian Dynamics	-	-
	Langevin MD	Classical Mechanics: Force Field	Langevin Dynamics	-	-
ML Methods	DeepMD (Zhang et al., 2018)	Atom-level Modeling	Newtonian Dynamics (i-PI)	Energy Prediction	PRL'18
	TorchMD (Doerr et al., 2020)	Atom-level Modeling	Newtonian Dynamics (velocity Verlet)	Energy Prediction	ACS'20
	Allegro-LAMMPS (Musaelian et al., 2023b)	Atom-level Modeling	Newtonian Dynamics (LAMMPS)	Force Prediction	ArXiv'23
	VerletMD (Ours, baseline)	Atom-level Modeling	Newtonian Dynamics (velocity Verlet)	Energy Prediction	-
	CGDMS (Greener & Jones, 2021)	Atom-level Modeling	Newtonian Dynamics (velocity Verlet)	Position Prediction	PLOS'21
	DiffMD (Wu & Li, 2023)	Atom-level Modeling	AR + Denoising	Position Prediction	AAAI'23
	DFF (Arts et al., 2023)	Atom-level Modeling	AR + Denoising	Position Prediction	ACS'23
	CG-MD (Fu et al., 2022b)	Atom-level Modeling	AR + Denoising	Position Prediction	TMLR'23
	LigandMD (Ours, baseline)	Atom-level Modeling	AR + Denoising	Position Prediction	-
NeuralMD (Ours)	Atom-level Modeling	Newtonian Dynamics (NODE)	Position Prediction	-	

between energy and trajectory prediction for MD: (1) Energy prediction takes each conformation and energy as IID, while trajectory learning optimizes the conformations along the whole trajectory, enforcing the temporal relation. (2) The time interval of trajectory learning is agnostic to the time interval, and energy prediction can be sensitive to longer MD simulations. More concretely, along such trajectory learning research line, CGDMS (Greener & Jones, 2021) builds an SE(3)-invariant model, followed by the velocity Verlet algorithm for MD simulation. DiffMD (Wu & Li, 2023) is a Markovian method and treats the dynamics between two consecutive snapshots as a coordinate denoising process. It then applies the SDE solver Song et al. (2020) to solve the molecular dynamics. A parallel work, DFF (Arts et al., 2023), applies a similar idea for MD simulation. CG-MD Fu et al. (2022b) encodes a hierarchical graph neural network model for an auto-regressive position generation and then adopts the denoising method for fine-tuning. However, these works disregard the prior knowledge of the Newtonian mechanics governing the motion of atoms.

Protein-Ligand Binding Dynamics. The MD simulation papers discussed so far are mainly for small molecules or proteins, not the binding dynamics. On the other hand, many works have studied the protein-ligand binding problem in the equilibrium state (Stepniewska-Dziubinska et al., 2018; Jiménez et al., 2018; Jones et al., 2021; Yang et al., 2023), but not the dynamics. In this work, we consider a more challenging task, which is the protein-ligand binding dynamics. The viability of this work is also attributed to the efforts of the scientific community, where more binding dynamics dataset has been revealed, including PLAS-5k Korlepara et al. (2022), MISATO Siebenmorgen et al. (2023), and PLAS-20k Priyakumar et al. (2023).

Preliminary on Molecular Dynamics. Molecular dynamics (MD) simulations predict how every atom in a molecular system moves over time, which is determined by the interatomic interactions following Newton’s second law. Such a molecular system includes Small Molecules (Rapaport, 2004; Zhang et al., 2018), proteins (Karplus & Kuriyan, 2005; Frauenfelder et al., 2009; Arts et al., 2023), polymers (Fu et al., 2022b), and protein-ligand complexes (Korlepara et al., 2022; Siebenmorgen et al., 2023). Typically, an MD simulation is composed of two main steps, *i.e.*, (1) the energy and force calculation and (2) integration of the equations of motion governed by Newton’s second law of motion, using the initial conditions and forces calculated in step (1). As the initial condition, the initial positions and velocities are given for all the particles (*e.g.*, atoms) in the molecular system; the MD simulation repeats the two steps to get a trajectory. Such MD simulations can be used to calculate the equilibrium and transport properties of molecules, materials, and biomolecular systems (Frenkel & Smit, 2002).

In such an MD simulation, one key factor is estimating the forces on each atom. The function that gives the energy of a molecular system as a function of its structure (and forces via the gradient of the energy with respect to those atomic coordinates) is referred to as a potential energy surface (PES). In general, MD simulations integrate the equations of motion using a PES from one of two sources: (1) Classical MD using the force fields, which are parameterized equations that approximate the true PES, and are less costly to evaluate, allowing for the treatment of larger systems and longer timesteps. (2) *ab-initio* MD (which calculates the energy of a molecular system via electronic structure methods, *e.g.*, DFT) provide more accurate PES, but are limited in the system size and timesteps that are practically accessible due to the cost of evaluating the PES at a given point.

B GROUP SYMMETRY AND EQUIVARIANCE

In this article, a 3D molecular graph is represented by a collection of 3D point clouds. The corresponding symmetry group is $SE(3)$, which consists of translations and rotations. Recall that we define the notion of equivariance functions in \mathbf{R}^3 in the main text through group actions. Formally, the group $SE(3)$ is said to act on \mathbf{R}^3 if there is a mapping $\phi : SE(3) \times \mathbf{R}^3 \rightarrow \mathbf{R}^3$ satisfying the following two conditions:

1. if $e \in SE(3)$ is the identity element, then

$$\phi(e, \mathbf{r}) = \mathbf{r} \quad \text{for } \forall \mathbf{r} \in \mathbf{R}^3.$$

2. if $g_1, g_2 \in SE(3)$, then

$$\phi(g_1, \phi(g_2, \mathbf{r})) = \phi(g_1 g_2, \mathbf{r}) \quad \text{for } \forall \mathbf{r} \in \mathbf{R}^3.$$

Then, there is a natural $SE(3)$ action on vectors \mathbf{r} in \mathbf{R}^3 by translating \mathbf{r} and rotating \mathbf{r} for multiple times. For $g \in SE(3)$ and $\mathbf{r} \in \mathbf{R}^3$, we denote this action by $g\mathbf{r}$. Once the notion of group action is defined, we say a function $f : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ that transforms $\mathbf{r} \in \mathbf{R}^3$ is equivariant if:

$$f(g\mathbf{r}) = gf(\mathbf{r}), \quad \text{for } \forall \mathbf{r} \in \mathbf{R}^3.$$

On the other hand, $f : \mathbf{R}^3 \rightarrow \mathbf{R}^1$ is invariant, if f is independent of the group actions:

$$f(g\mathbf{r}) = f(\mathbf{r}), \quad \text{for } \forall \mathbf{r} \in \mathbf{R}^3.$$

For some scenarios, our problem is chiral sensitive. That is, after mirror reflecting a 3D molecule, the properties of the molecule may change dramatically. In these cases, it's crucial to include reflection transformations into consideration. More precisely, we say an $SE(3)$ equivariant function f is **reflection anti-symmetric**, if:

$$f(\rho\mathbf{r}) \neq f(\mathbf{r}), \tag{9}$$

for reflection $\rho \in E(3)$.

C EQUIVARIANT MODELING WITH VECTOR FRAMES

Frame is a popular terminology in science areas. In physics, the frame is equivalent to a coordinate system. For example, we may assign a frame to all observers, although different observers may collect different data under different frames, the underlying physics law should be the same. In other words, denote the physics law by f , then f should be an equivariant function.

There are certain ways to choose the frame basis, and below we introduce two main types: the orthogonal basis and the protein backbone basis. The orthogonal basis can be built for flexible 3D point clouds such as atoms, while the protein backbone basis is specifically proposed to capture the protein backbone.

C.1 BASIS

Since there are three orthogonal directions in \mathbf{R}^3 , one natural frame in \mathbf{R}^3 can be a frame consisting of three orthogonal vectors:

$$F = (e_1, e_2, e_3).$$

Once equipped with a frame (coordinate system), we can project all geometric quantities to this frame. For example, an abstract vector $\mathbf{x} \in \mathbf{R}^3$ can be written as $\mathbf{x} = (r_1, r_2, r_3)$ under the frame F , if: $\mathbf{x} = r_1e_1 + r_2e_2 + r_3e_3$. A vector frame further requires the three orthonormal vectors in (e_1, e_2, e_3) to be equivariant. Intuitively, a vector frame will transform according to the global rotation or translation of the whole system. Once equipped with a vector frame, we can project vectors into this frame in an equivariant way:

$$\mathbf{x} = \tilde{r}_1e_1 + \tilde{r}_2e_2 + \tilde{r}_3e_3. \quad (10)$$

We call the process of $\mathbf{x} \rightarrow \tilde{r} := (\tilde{r}_1, \tilde{r}_2, \tilde{r}_3)$ the **scalarization** or **projection** operation. Since $\tilde{r}_i = e_i \cdot \mathbf{x}$ is expressed as an inner product between vector vectors, we know that \tilde{r} consists of scalars.

In this article, we assign a vector frame to each node/edge, therefore we call them the local frames. We note that, in this section, we prove the equivariance property of the vector frame basis using the Gram-Schmidt methods. However, similar equivariance property can be easily guaranteed for the vector frame bases in Section 3 after we remove the mass center of the molecular system.

In the main body, we constructed three vector frames based on three granularities. Here we provide the proof on the protein backbone frame. Say the three backbone atoms in on proteins are $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$ respectively. Then the vector frame is defined by:

$$\text{Vector-Frame}(\mathbf{x}_i, \mathbf{x}_j) := \text{Gram-Schmidt}\{\mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_k, (\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k)\}. \quad (11)$$

The Gram-Schmidt orthogonalization makes sure that the $\text{Vector-Frame}(\mathbf{x}_i, \mathbf{x}_j)$ is orthonormal.

Reflection Antisymmetric Since we implement the cross product \times for building the local frames, the third vector in the frame is a pseudo-vector. Then, the **projection** operation is not invariant under reflections (the inner product between a vector and a pseudo-vector change signs under reflection). Therefore, our model can discriminate two 3D geometries with different chiralities.

Our local frames also enable us to output vectors equivariantly by multiplying scalars (v_1, v_2, v_3) with the frame: $\mathbf{v} = v_1 \cdot e_1 + v_2 \cdot e_2 + v_3 \cdot e_3$.

Equivariance w.r.t. cross-product The goal is to prove that the cross-product is equivariant to the SE(3)-group, *i.e.*:

$$gx \times gy = g(x \times y), \quad g \in \text{SE}(3)\text{-Group} \quad (12)$$

Geometric proof. From intuition, with rotation matrix g , we are transforming the whole basis, thus the direction of $gx \times gy$ changes equivalently with g . And for the value/length of $gx \times gy$, because $|gx \times gy| = \|gx\| \cdot \|gy\| \cdot \sin \theta = \|x\| \cdot \|y\| \cdot \sin \theta = |x \times y|$. So the length stays the same, and the direction changes equivalently. Intuitively, this interpretation is quite straightforward.

Analytical proof. A more rigorous proof can be found below:

Proof. First, we have that for the rotation matrix g :

$$gx \times gy = \begin{bmatrix} \mathbf{g}_1^T \mathbf{x} \\ \mathbf{g}_2^T \mathbf{x} \\ \mathbf{g}_3^T \mathbf{x} \end{bmatrix} \times \begin{bmatrix} \mathbf{g}_1^T \mathbf{y} \\ \mathbf{g}_2^T \mathbf{y} \\ \mathbf{g}_3^T \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{g}_2^T \mathbf{x} \cdot \mathbf{g}_3^T \mathbf{y} - \mathbf{g}_3^T \mathbf{x} \cdot \mathbf{g}_2^T \mathbf{y} \\ -\mathbf{g}_1^T \mathbf{x} \cdot \mathbf{g}_3^T \mathbf{y} + \mathbf{g}_3^T \mathbf{x} \cdot \mathbf{g}_1^T \mathbf{y} \\ \mathbf{g}_1^T \mathbf{x} \cdot \mathbf{g}_2^T \mathbf{y} - \mathbf{g}_2^T \mathbf{x} \cdot \mathbf{g}_1^T \mathbf{y} \end{bmatrix}, \quad (13)$$

where $\mathbf{g}_i, \mathbf{x}, \mathbf{y} \in \mathbb{R}^{3 \times 1}$.

Because $A^T C \cdot B^T D - A^T D \cdot B^T C = (A \times B)^T (C \times D)$, so we can have:

$$gx \times gy = \begin{bmatrix} \mathbf{g}_2^T \mathbf{x} \cdot \mathbf{g}_3^T \mathbf{y} - \mathbf{g}_3^T \mathbf{x} \cdot \mathbf{g}_2^T \mathbf{y} \\ -\mathbf{g}_1^T \mathbf{x} \cdot \mathbf{g}_3^T \mathbf{y} + \mathbf{g}_3^T \mathbf{x} \cdot \mathbf{g}_1^T \mathbf{y} \\ \mathbf{g}_1^T \mathbf{x} \cdot \mathbf{g}_2^T \mathbf{y} - \mathbf{g}_2^T \mathbf{x} \cdot \mathbf{g}_1^T \mathbf{y} \end{bmatrix} = \begin{bmatrix} (\mathbf{g}_2 \times \mathbf{g}_3)^T (\mathbf{x} \times \mathbf{y}) \\ (\mathbf{g}_3 \times \mathbf{g}_1)^T (\mathbf{x} \times \mathbf{y}) \\ (\mathbf{g}_1 \times \mathbf{g}_2)^T (\mathbf{x} \times \mathbf{y}) \end{bmatrix} \quad (14)$$

Then because:

$$\begin{aligned} \det(g) &= (\mathbf{g}_2 \times \mathbf{g}_3)^T \mathbf{g}_1 = \mathbf{g}_1^T \mathbf{g}_1 = 1 \\ \implies (\mathbf{g}_2 \times \mathbf{g}_3)^T \mathbf{g}_1 \mathbf{g}_1^{-1} &= \mathbf{g}_1^T \mathbf{g}_1 \mathbf{g}_1^{-1} \\ \implies (\mathbf{g}_2 \times \mathbf{g}_3)^T &= \mathbf{g}_1^T. \end{aligned} \quad (15)$$

Thus, we can have

$$gx \times gy = \begin{bmatrix} (\mathbf{g}_2 \times \mathbf{g}_3)^T (\mathbf{x} \times \mathbf{y}) \\ (\mathbf{g}_3 \times \mathbf{g}_1)^T (\mathbf{x} \times \mathbf{y}) \\ (\mathbf{g}_1 \times \mathbf{g}_2)^T (\mathbf{x} \times \mathbf{y}) \end{bmatrix} = \begin{bmatrix} \mathbf{g}_1^T (\mathbf{x} \times \mathbf{y}) \\ \mathbf{g}_2^T (\mathbf{x} \times \mathbf{y}) \\ \mathbf{g}_3^T (\mathbf{x} \times \mathbf{y}) \end{bmatrix} = g(\mathbf{x} \times \mathbf{y}). \quad (16)$$

□

Rotation symmetric The goal is to prove

$$\text{Vector-Frame}(g\mathbf{x}_i, g\mathbf{x}_j) = g\text{Gram-Schmidt}\{\mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_k, (\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k)\}. \quad (17)$$

Proof. Thus we can have:

$$\begin{aligned} \text{Vector-Frame}(g\mathbf{x}_i, g\mathbf{x}_j) &= \text{Gram-Schmidt}\{g\mathbf{x}_i - g\mathbf{x}_j, g\mathbf{x}_i - g\mathbf{x}_k, (g\mathbf{x}_i - g\mathbf{x}_j) \times (g\mathbf{x}_i - g\mathbf{x}_k)\} \\ &= \text{Gram-Schmidt}\{g(\mathbf{x}_i - \mathbf{x}_j), g(\mathbf{x}_i - \mathbf{x}_k), g((\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k))\}. \end{aligned} \quad (18)$$

Recall that Gram-Schmidt projection (**Gram-Schmidt** $\{v_1, v_2, v_3\}$) is:

$$\begin{aligned} \mathbf{u}_1 &= \mathbf{v}_1, & \mathbf{e}_1 &= \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}, \\ \mathbf{u}_2 &= \mathbf{v}_2 - \frac{\mathbf{u}_1^T \mathbf{v}_2}{\|\mathbf{u}_1\|} \mathbf{u}_1, & \mathbf{e}_2 &= \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|}, \\ \mathbf{u}_3 &= \mathbf{v}_3 - \frac{\mathbf{u}_1^T \mathbf{v}_3}{\|\mathbf{u}_1\|} \mathbf{u}_1 - \frac{\mathbf{u}_2^T \mathbf{v}_3}{\|\mathbf{u}_2\|} \mathbf{u}_2, & \mathbf{e}_3 &= \frac{\mathbf{v}_3}{\|\mathbf{v}_3\|}. \end{aligned} \quad (19)$$

Thus, the Gram-Schmidt projection on the rotated vector (**Gram-Schmidt** $\{g\mathbf{v}_1, g\mathbf{v}_2, g\mathbf{v}_3\}$) is:

$$\begin{aligned} \mathbf{u}'_1 &= g\mathbf{v}_1, \\ \mathbf{u}'_2 &= g\mathbf{v}_2 - g \frac{\mathbf{u}_1^T \mathbf{v}_2}{\|\mathbf{u}_1\|} \mathbf{u}_1, \\ \mathbf{u}'_3 &= g\mathbf{v}_3 - g \frac{\mathbf{u}_1^T \mathbf{v}_3}{\|\mathbf{u}_1\|} \mathbf{u}_1 - g \frac{\mathbf{u}_2^T \mathbf{v}_3}{\|\mathbf{u}_2\|} \mathbf{u}_2, \end{aligned} \quad (20)$$

Thus, **Gram-Schmidt** $\{g\mathbf{v}_1, g\mathbf{v}_2, g\mathbf{v}_3\} = g\text{Gram-Schmidt}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$.

□

Transition symmetric

$$\text{Vector-Frame}(\mathbf{x}_i + \delta\mathbf{x}, \mathbf{x}_j + \delta\mathbf{x}) = \text{Gram-Schmidt}\{\mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_k, (\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k)\}. \quad (21)$$

Proof. Because the basis is based on the difference of coordinates, it is straightforward to observe that $\text{Gram-Schmidt}\{\mathbf{v}_1 + \mathbf{t}, \mathbf{v}_2 + \mathbf{t}, \mathbf{v}_3 + \mathbf{t}\} = \text{Gram-Schmidt}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$. So the frame operation is transition equivariant. We also want to highlight that for all the other vector frame bases introduced in Section 3, we remove the mass center for each molecular system, thus, we can guarantee the transition equivariance property. \square

Reflection antisymmetric

$$\text{Vector-Frame}(\mathbf{x}_i, \mathbf{x}_j) \neq \text{Vector-Frame}(-\mathbf{x}_i, -\mathbf{x}_j). \quad (22)$$

Proof. From intuition, this makes sense because the cross-product is anti-symmetric.

A simple counter-example is the original geometry R and the reflected geometry by the original point $-R$. Thus the two bases before and after the reflection group is the following:

$$\text{Gram-Schmidt}\{\mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_k, (\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k)\} \quad (23)$$

$$\text{Gram-Schmidt}\{-\mathbf{x}_i + \mathbf{x}_j, -\mathbf{x}_i + \mathbf{x}_k, (\mathbf{x}_i - \mathbf{x}_j) \times (\mathbf{x}_i - \mathbf{x}_k)\}. \quad (24)$$

The bases between $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ and $\{-\mathbf{v}_1, -\mathbf{v}_2, \mathbf{v}_3\}$ are different, thus such frame construction is reflection anti-symmetric. \square

If you are able to get the above derivations carefully with a good understanding, then you can tell that this can be trivially generalized to arbitrary vector frames.

C.2 SCALARIZATION

Once we have the three vectors as the vector frame basis, the next thing is to do modeling. Suppose the frame is $\mathcal{F} = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$, then for a vector (tensor) \mathbf{h} , the scalarization is:

$$\mathbf{h} \odot \mathcal{F} = (\mathbf{h} \odot \mathbf{e}_1, \mathbf{h} \odot \mathbf{e}_2, \mathbf{h} \odot \mathbf{e}_3) = (\mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3). \quad (25)$$

D SPECIFICATIONS ON MISATO

In this section, we provide more details on the MISATO dataset (Siebenmorgen et al., 2023).

For small molecule ligands, we ignore the Hydrogen atoms.




Figure 5: Distribution on # atoms in small molecule ligands for all protein-ligand complex.




Figure 6: Distribution on # residues in proteins for all protein-ligand complex.

We also plot the distribution of the energy gap between each time step and the initial snapshot, *i.e.*, $E_t - E_0$. The distribution is in Figure 7. We can observe that as the time processes, the mean of the energy stays almost the same, yet the variance gets higher.




Figure 7: Distribution on energy $E_t - E_0$.

E DETAILS OF NEURALMD

In this section, we provide more details on the model architecture in Figure 8, and hyperparameter details in Table 5.

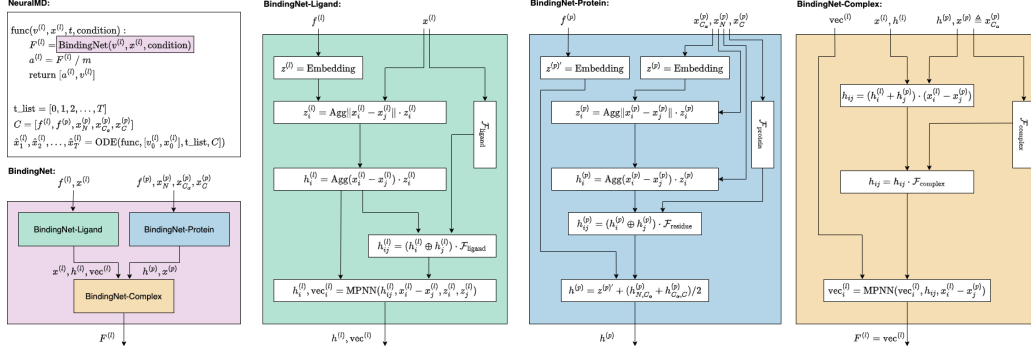


Figure 8: Detailed pipeline of NeuralMD. In the three key modules of BindingNet, there are three vertical boxes, corresponding to three granularities of vector frames, as in Equations (2) to (4).

Here, we explain each of the three modules in detail and list the dimensions of each variable to make it easier for readers to understand. Suppose the representation dimension is d .

BindingNet-Ligand:

- $z^{(l)} = \text{Embedding}(f^{(l)}) \in \mathbb{R}^{N_{\text{atom}} \cdot d}$ is atom type embedding.
- Then for each atom type embedding $z^{(l)}$, we add a normalization by multiplying it with the RBF of distance among the neighborhoods, and the resulting atom type embedding stays the same dimension $z^{(l)} \in \mathbb{R}^{N_{\text{atom}} \times d}$.
- $\{h_i^{(l)} = \text{Agg}_j(\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}) \cdot z_i^{(l)}\} \in \mathbb{R}^{N_{\text{atom}} \cdot d \cdot 3}$ is the equivariant representation of each atom.
- $\{h_{ij}^{(l)}\} \in \mathbb{R}^{N_{\text{edge}} \cdot 2d \cdot 3}$ is the invariant representation after scalarization. Then we will take a simple sum-pooling, followed by an MLP to get the invariant representation $h_{ij}^{(l)} \in \mathbb{R}^{N_{\text{edge}} \cdot d}$.
- Finally, we will repeat L layers of MPNN:

$$\begin{aligned} \text{vec}_i^{(l)} &= \text{vec}_i^{(l)} + \text{Agg}_j(\text{vec}_i^{(l)} \cdot \text{MLP}(h_{ij}^{(l)}) + (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(p)}) \cdot \text{MLP}(h_{ij}^{(l)})), \quad // \{\text{vec}_i^{(l)}\} \in \mathbb{R}^{N_{\text{atom}} \cdot 3} \\ h_i^{(l)} &= h_i^{(l)} + \text{Agg}_j(\text{MLP}(h_{ij}^{(l)})). \quad // \{h_i^{(l)}\} \in \mathbb{R}^{N_{\text{atom}} \cdot d} \end{aligned} \quad (26)$$

BindingNet-Protein:

- $z^{(p)} \in \mathbb{R}^{N_{\text{backbone-atom}} \cdot d}$ is the backbone-atom type representation by aggregating the neighbors without the cutoff c .
- $\tilde{z}^{(p)} \in \mathbb{R}^{N_{\text{backbone-atom}} \cdot d}$ is the backbone-atom type representation.
- $\{h_i^{(p)}\} \in \mathbb{R}^{N_{\text{backbone-atom}} \cdot d \cdot 3}$ is the backbone-atom equivariant representation.
- $\{h_{ij}^{(p)}\} \in \mathbb{R}^{N_{\text{edge}} \cdot 2d \cdot 3}$ is the invariant representation after scalarization. Then we take a simple sum-pooling, followed by an MLP to get the invariant representation $\{h_{ij}^{(p)}\} \in \mathbb{R}^{N_{\text{edge}} \cdot d}$.
- Finally, we get the residue-level representation as $h^{(p)} = \tilde{z}^{(p)} + (h_{N, C_\alpha}^{(p)} + h_{C_\alpha, C}^{(p)})/2 \in \mathbb{R}^{N_{\text{residue}} \cdot d}$.

BindingNet-Complex:

- $\{h_{ij}\} \in \mathbb{R}^{N_{\text{edge}} \cdot d \cdot 3}$ is the equivariant interaction/edge representation.
- $\{h_{ij} = h_{ij} \cdot \mathcal{F}_{\text{Complex}}\} \in \mathbb{R}^{N_{\text{edge}} \cdot d \cdot 3}$ is the scalarization. Then we take a simple sum-pooling, followed by an MLP to get the invariant representation $\{h_{ij}^{(l)}\} \in \mathbb{R}^{N_{\text{edge}} \cdot d}$.
- The final output is obtained by L MPNN layers as:

$$\begin{aligned} \text{vec}_{ij}^{(pl)} &= \text{vec}_i^{(l)} \cdot \text{MLP}(h_{ij}) + (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(p)}) \cdot \text{MLP}(h_{ij}), \quad // \{\text{vec}_{ij}^{(pl)}\} \in \mathbb{R}^{N_{\text{edge}} \cdot 3} \\ F_i^{(l)} &= \text{vec}_i^{(l)} + \text{Agg}_{j \in \mathcal{N}(i)} \text{vec}_{ij}^{(pl)}. \quad // \{F_i^{(l)}\} \in \mathbb{R}^{N_{\text{atom}} \cdot 3} \end{aligned} \quad (27)$$

Table 5: Hyperparameter specifications for NeuralMD.

Hyperparameter	Value
# layers	{5}
cutoff c	{5}
learning rate	{1e-4, 1e-3}
optimizer	{SGD, Adam }

F ABLATION STUDIES: FLEXIBLE BINDING

Recall that, in the main paper, we have discussed using the *semi-flexible* binding setting, *i.e.*, proteins with rigid structures while small molecule ligands with flexible structures, and the goal is to predict the trajectories of the ligands. As discussed in Section 5, if we want to take both proteins and ligands with flexible structures, one limitation is the GPU memory cost. However, we would like to mention that it is possible to do NeuralMD on protein-ligand with small volume, and we take an ablation study to test them as below.

Problem Setup. Both the proteins and ligands are flexible, and we want to predict their trajectories simultaneously. In the main paper, we consider three levels of vector frames. Here in the flexible setting, due to the large volume of atoms in the protein-ligand complex, we are only able to consider two levels, *i.e.*, the atom-level and residue-level. Thus, the backbone model (BindingNet) also changes accordingly. The performance is shown in Table 6, and we can see that NeuralMD is consistently better than the GNN-MD on all three metrics and all 10 single trajectories. We omit the multi-trajectory experiments due to the memory limitation.

Table 6: Results on ten single-trajectory binding dynamics prediction. Results with optimal training loss are reported. Four evaluation metrics are considered: MAE (\AA , \downarrow), MSE (\downarrow), and Stability ($\%$, \downarrow).

		GNN-MD	NeuralMD-ODE (Ours)
5WIJ	MAE	7.126	3.101
	MSE	4.992	2.070
	Stability	68.317	30.655
4ZX0	MAE	9.419	2.580
	MSE	6.269	1.724
	Stability	67.492	29.013
3EOV	MAE	10.695	3.664
	MSE	7.447	2.521
	Stability	67.782	39.714
4K6W	MAE	8.347	3.056
	MSE	5.605	2.007
	Stability	63.839	36.972
1KTI	MAE	8.900	6.815
	MSE	5.820	4.268
	Stability	65.010	26.805
1XP6	MAE	8.496	1.910
	MSE	5.673	1.276
	Stability	70.019	33.907
4YUR	MAE	11.710	7.568
	MSE	7.759	4.966
	Stability	69.163	34.636
4G3E	MAE	1314.425	3.282
	MSE	814.641	2.152
	Stability	65.703	21.095
6B7F	MAE	182.278	3.166
	MSE	115.688	2.121
	Stability	72.027	26.931
3B9S	MAE	3.590	2.477
	MSE	2.431	1.615
	Stability	54.890	18.817

G LANGEVIN DYNAMICS

The **Langevin dynamics** is defined as

$$m\mathbf{a} = -\nabla U(\mathbf{x}) - \gamma m\mathbf{v} + \sqrt{2m\gamma k_B T}R(t), \quad (28)$$

where γ is the damping constant or collision frequency, T is the temperature, k_B is the Boltzmann's constant, and $R(t)$ is a delta-correlated stationary Gaussian process with zero-mean.

The term $m\mathbf{a}$ is called inertial force. In **overdamped Langevin dynamics** (a.k.a. Brownian dynamics), the inertial force is much smaller than the other three terms, thus it is considered negligible. Then the equation for overdamped Langevin dynamics is:

$$-\nabla U(\mathbf{x}) - \gamma m\mathbf{v} + \sqrt{2m\gamma k_B T}R(t) = 0. \quad (29)$$

Thus, the trajectories are given by:

$$\begin{aligned} \mathbf{x}_{t+1} - \mathbf{x}_t &= -\frac{1}{\gamma m} \nabla U(x) + \frac{\sqrt{2m\gamma k_B T}}{\gamma m} R(t) \\ &= -\frac{D}{k_B T} \nabla U(X) + \sqrt{2D}R(t), \end{aligned} \quad (30)$$

where $D = k_B T / \gamma$. However, such an overdamped Langevin dynamics does not hold for small particles like molecules. Because the assumption that inertial force can be ignored is only valid for large particles.