

Graph Fourier Neural ODEs: Modeling Spatial-temporal Multi-scales in Molecular Dynamics

Anonymous authors

Paper under double-blind review

Abstract

Accurately predicting long-horizon molecular dynamics (MD) trajectories remains a significant challenge, as existing deep learning methods often struggle to retain fidelity over extended simulations. We hypothesize that one key factor limiting accuracy is the difficulty of capturing interactions that span distinct spatial and temporal scales—ranging from high-frequency local vibrations to low-frequency global conformational changes. To address these limitations, we propose **Graph Fourier Neural ODEs (GF-NODE)**, integrating a graph Fourier transform for spatial frequency decomposition with a Neural ODE framework for continuous-time evolution. Specifically, GF-NODE first decomposes molecular configurations into multiple spatial frequency modes using the graph Laplacian, then evolves the frequency components in time via a learnable Neural ODE module that captures both local and global dynamics, and finally reconstructs the updated molecular geometry through an inverse graph Fourier transform. By explicitly modeling high- and low-frequency phenomena in this unified pipeline, GF-NODE more effectively captures long-range correlations and local fluctuations alike. Experimental results on challenging MD benchmarks, including MD17 and alanine dipeptide, demonstrate that GF-NODE achieves state-of-the-art accuracy while preserving essential geometrical features over extended simulations. These findings highlight the promise of bridging spectral decomposition with continuous-time modeling to improve the robustness and predictive power of MD simulations. Our implementation is publicly available at <https://anonymous.4open.science/r/GF-NODE-code-B289/>

1 Introduction

Molecular dynamics (MD) simulations are indispensable tools for investigating the behavior of molecular systems at the atomic level, offering profound insights into physical (Bear & Blaisten-Barojas, 1998), chemical (Wang et al., 2011), and biological (Salo-Ahen et al., 2020) processes. These simulations must capture interactions occurring across a wide range of spatial and temporal scales—from localized bond vibrations to long-range non-bonded interactions—posing significant computational challenges. Accurately modeling these multiscale interactions is crucial for uncovering the mechanisms underlying complex molecular phenomena but remains prohibitively expensive for large systems and long trajectories (Vakis et al., 2018).

In recent years, Graph Neural ODEs (Zang & Wang, 2020; Huang et al., 2023a) have gained traction for modeling continuous-time dynamics in multi-agent systems, including MD. By learning an Ordinary Differential Equation (ODE) function via Graph Neural Networks and solving it numerically, these methods allow flexible sampling at arbitrary time points. Such a continuous-time formulation is well-suited for capturing multiple temporal scales inherent in molecular simulations. However, significant challenges persist in accounting for the rich spatial multiscale effects that span localized bond vibrations to extended nonbonded interactions. On another front, Fourier Neural Operators (FNOs) (Li et al., 2020) have demonstrated success in learning operators by decomposing signals into different frequency modes, thereby capturing various spatial scales effectively. Yet, they are not tailored for graph-structured molecular data or continuous-time temporal evolution, limiting their direct applicability to general MD simulations.




Figure 1: Overview of the Graph Fourier Neural ODE framework. The molecular graph \mathcal{G}^{t_0} is first transformed into the spectral domain using a Graph Fourier Transform (\mathcal{GF}), decomposing the spatial structure into frequency components. Neural ODEs are then applied to evolve the Fourier coefficients over time. The evolved coefficients are finally transformed back into the physical domain using an inverse Graph Fourier Transform (\mathcal{GF}^{-1}), reconstructing the molecular graph at future time t_k .

Despite the promise of Graph Neural ODEs for handling multiple temporal scales via continuous time modeling, they alone are inadequate for fully capturing the complex spatial frequency components of molecular systems. Conversely, while approaches based on Fourier transforms can model multiple spatial scales, they do not naturally handle the intricacies of molecular graphs or continuous-time dynamics. To address these limitations, we introduce *Graph Fourier Neural ODEs (GF-NODE)*. As demonstrated in Figure 1, our framework explicitly integrates a graph Fourier transform—for decomposing and encoding spatial multiscale interactions—with a Neural ODE framework for continuous-time modeling of each spatial frequency. By leveraging an inverse graph Fourier transform at the end of the pipeline, GF-NODE reconstructs the molecular state in physical space, thereby enabling a unified approach to spatial and temporal multiscale simulation.

We conduct extensive experiments on benchmark molecular dynamics datasets, including MD17 and alanine dipeptide. Empirical results show that GF-NODE achieves state-of-the-art accuracy in predicting molecular trajectories over long-horizon simulations, preserves essential geometric properties such as bond lengths and angles, and demonstrates stable performance




Figure 2: Illustration of spatial and temporal multiscale interactions in molecular dynamics. The x-axis indicates spatial scales ranging from local bond vibrations to global conformational changes. The y-axis represents timescales ranging from femtoseconds (bond oscillations) to nanoseconds (conformational rearrangements). Local interactions, such as bond vibrations and angle bending, occur over short spatial scales (angstroms) and fast timescales (femtoseconds). Non-local interactions, including hydrogen bonding and van der Waals forces, span larger spatial scales (nanometers) and intermediate timescales (picoseconds). Conformational changes involve the longest spatial and temporal scales, potentially up to nanoseconds or longer.

over temporal super-resolution tasks. These findings underscore the importance of explicitly decomposing molecular configurations into spatial frequency modes and evolving them continuously in time. Our analysis suggests that this multiscale perspective is instrumental for capturing both rapid local fluctuations and slow global conformational changes.

Our contribution can be summarized as follows.

(a) *New perspective.* We provide a new perspective on spatial-temporal multiscale modeling for molecular dynamics by jointly capturing spatial and temporal interactions within a single framework.

(b) *Novel architecture.* Building on this perspective, we introduce the *Graph Fourier Neural ODEs (GF-NODE)* architecture, which combines a graph Fourier transform for decomposing molecular interactions into distinct spatial frequencies with a Neural ODE formulation to evolve these frequencies continuously in time.

(c) *Good performance.* Through extensive experiments on benchmark molecular dynamics datasets, we demonstrate that GF-NODE not only outperforms existing approaches in predictive accuracy but also effectively preserves critical molecular geometrical features, thus offering a promising solution for multiscale MD simulations.

2 Related Work

We review relevant works on multi-scale modeling in molecular dynamics, focusing on neural operator models and graph neural ODEs.

2.1 Classical Molecular Simulation Methods

Traditional molecular simulation methods, including force-field based MD (e.g., AMBER (Cornell et al., 1995; Brooks et al., 2009), GAMD (Li et al., 2022b)) and *ab initio* techniques such as Car-Parrinello MD (Hutter, 2012), have been foundational in exploring molecular behavior. However, these methods face significant limitations: force-field simulations require extremely small time steps to accurately resolve high-frequency bond vibrations, which hampers long-term stability and computational efficiency; *ab initio* MD, though offering first-principles accuracy, is computationally prohibitive for large systems and long trajectories; and while coarse-grained models (e.g., MARTINI (Souza et al., 2021)) enable more efficient multiscale simulations, they often compromise on molecular detail and accuracy, particularly in reproducing local interactions and maintaining seamless force consistency at multiscale interfaces.

2.2 Neural Operator Models

Neural operator models (Kovachki et al., 2023) have emerged as powerful tools for learning mappings between infinite-dimensional function spaces, demonstrating success in modeling complex dynamical systems. Among these, *Fourier Neural Operators (FNOs)* (Li et al., 2020; 2022a; Liu & Jafarzadeh, 2023; Kovachki et al., 2021; Koshizuka et al.) are particularly notable for handling spatial multi-scale interactions in partial differential equation (PDE) data by learning representations in the Fourier domain. However, while FNOs efficiently capture spatial hierarchies, they do not inherently model temporal dynamics, making them suboptimal for time-evolving molecular systems.

In contrast, recent operator-based methods such as the *Implicit Transfer Operator (ITO)* (Schreiner et al., 2024), *Timewarp* (Klein et al., 2024), and *Equivariant Graph Neural Operator (EGNO)* (Xu et al., 2024) focus on temporal multi-scale modeling in molecular dynamics. ITO and Timewarp introduce coarse-graining and adaptive time-stepping mechanisms to accelerate long-horizon simulations. EGNO employs neural operators with SE(3) equivariance to capture rotational and translational symmetries, yet it primarily addresses temporal evolution without explicitly handling spatial multi-scale effects. While these models successfully extend the applicability of neural operators to molecular simulations, none jointly addresses both spatial and temporal multi-scales in molecular dynamics.

2.3 Graph Neural ODE Models

Graph Neural ODEs combine Graph Neural Networks with Neural ODE frameworks (Chen et al., 2018; Kidger, 2022; Goyal & Benner, 2023; Holt et al., 2022; Luo et al., 2023; Luo et al.) to model continuous-time dynamics on graph-structured data (Zang & Wang, 2020; Kim et al., 2021; Huang et al., 2020; 2021; 2023a;b). These methods excel in capturing temporal multi-scale behavior by allowing flexible time integration, making them well-suited for systems with varying temporal resolutions. However, they primarily focus on modeling temporal dependencies, with little emphasis on explicitly handling spatial multi-scales in molecular dynamics.

A key limitation of existing Graph Neural ODE models is their reliance on local message passing, which inherently constrains their ability to capture long-range spatial dependencies within molecular systems. As a result, they may fail to adequately represent the interplay between localized, high-frequency interactions (e.g., bond vibrations) and global, low-frequency effects (e.g., large-scale conformational changes). Unlike spectral-based approaches that can decompose spatial hierarchies, standard Graph Neural ODEs lack a mechanism to explicitly encode spatial multi-scale structures, limiting their effectiveness in modeling complex molecular dynamics.




Figure 3: An overview of the proposed GF-NODE architecture. The model first encodes the initial molecular graph \mathcal{G}^{t_0} from the physical domain into the spectral domain via a graph Fourier transform (\mathcal{GF}). Neural ODE solvers then propagate the dynamics of the Fourier coefficients continuously across time points t_0, t_1, \dots, t_K . The transformed coefficients are subsequently decoded back to the physical domain using an inverse graph Fourier transform (\mathcal{GF}^{-1}), reconstructing the molecular graphs $\mathcal{G}^{t_1}, \mathcal{G}^{t_2}, \dots, \mathcal{G}^{t_K}$ at future time points. This design enables efficient modeling of spatial and temporal multiscale dynamics in molecular systems.

3 The Proposed Approach

We propose **GF-NODE**, a framework specifically designed to address the limitations of existing methods in capturing both spatial and temporal multiscale dynamics in molecular systems. As discussed in Sections 1 and 2, current approaches either handle spatial scales using Fourier-based methods or focus on temporal scales using Graph Neural ODEs, but they do not jointly model these scales in a unified framework. GF-NODE directly addresses this gap by integrating the *Graph Fourier Transform (GFT)* and *Neural Ordinary Differential Equations (Neural ODEs)* to simultaneously decompose spatial interactions and model their continuous-time evolution.

Specifically, as illustrated in Figure 3, GF-NODE first applies a Graph Fourier Transform to decompose the molecular graph into different frequency components, effectively separating localized, high-frequency

interactions from global, low-frequency patterns. This spectral representation allows the model to process molecular structures in a frequency-adaptive manner, capturing both fine-grained local interactions and large-scale conformational changes. The decomposed spectral coefficients are then evolved continuously over time using Neural ODEs, ensuring flexible and adaptive modeling of multiscale temporal dynamics. Finally, the inverse Graph Fourier Transform reconstructs the molecular graph in the physical domain, preserving both local and global structural information over long-horizon molecular simulations.

This design enables GF-NODE to overcome key challenges in molecular dynamics modeling: (1) explicitly encoding spatial multi-scale interactions via spectral decomposition, (2) leveraging continuous-time evolution to capture complex temporal dependencies, and (3) integrating these spatial and temporal scales into a single end-to-end framework. Below, we detail the core components and operations of GF-NODE.

3.1 Notation and problem setup

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represent a molecular graph, where N is the number of nodes $i \in \{1, \dots, N\}$ (atoms), and \mathcal{E} is the set of edges representing chemical bonds or interactions. Edges are identified by checking whether the distance between atoms falls below a threshold. Each node i has:

- Invariant (scalar) features $\mathbf{h}_i \in \mathbb{R}^F$, which include the velocity magnitude $\|\mathbf{v}_i\|$ and a normalized version of the atomic number Z_i ;
- Vector features $\mathbf{z}_i \in \mathbb{R}^{m \times 3}$, which contain position \mathbf{x}_i and velocity \mathbf{v}_i (i.e. $\mathbf{z}_i = (\mathbf{x}_i, \mathbf{v}_i)$ when $m = 2$).

At time t_0 , the molecular system’s state is given as $\{\mathbf{h}, \mathbf{z}\}$, and we aim to predict the future configuration $\mathcal{G}^{(t_k)}$ for any $t_k > t_0$.

3.2 Graph Neural Network encoder

We use a Graph Neural Network (GNN) to encode scalar features $\mathbf{h}_i \in \mathbb{R}^F$ and vector features $\mathbf{z}_i \in \mathbb{R}^{m \times 3}$ for each node i . The initial scalar feature $\mathbf{h}_i^{(0)}$ for each atom i is formed by concatenating $\|\mathbf{v}_i\|$ and $\frac{Z_i}{Z_{\max}}$, where Z_i is the atomic number of atom i , and Z_{\max} is a reference maximum. This concatenated vector is then mapped by a linear embedding layer, producing the hidden dimension used by the GNN. The vector feature $\mathbf{z}_i^{(0)}$ contains \mathbf{x}_i and \mathbf{v}_i . Each GNN layer performs message passing, where node i ’s features are updated based on its neighbors $\mathcal{N}(i)$:

$$\mathbf{m}_{ij}^{(l)} = \phi_e \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{r}_{ij}^{(l)} \right), \quad (1)$$

$$\mathbf{h}_i^{(l+1)} = \phi_h \left(\mathbf{h}_i^{(l)}, \sum_{j \in \mathcal{N}(i)} \mathbf{m}_{ij}^{(l)} \right), \quad (2)$$

$$\mathbf{x}_i^{(l+1)} = \mathbf{x}_i^{(l)} + \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \psi \left(\mathbf{m}_{ij}^{(l)} \right), \quad (3)$$

where $\mathbf{r}_{ij}^{(l)} = \mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}$, and ϕ_e, ϕ_h , and ψ are neural networks. Similarly, we also update the velocity $\mathbf{v}_i^{(l)}$ if present. After L layers, the GNN produces the encoded features $\mathbf{h}_i^{(L)}$ and $\mathbf{z}_i^{(L)}$, capturing both local and global molecular information. These features are then passed to the Graph Fourier Transform (GFT) for spectral decomposition.

3.3 Graph Fourier Transform

Molecular dynamics exhibit behavior on multiple spatial scales: large-scale “global” deformations can be viewed as low-frequency modes, whereas fast “local” vibrations correspond to high frequency modes. In classical signal processing, Fourier analysis decomposes signals into sinusoids of different frequencies. Similarly, on graphs, we can decompose node-based signals into eigenmodes of a suitable operator (often the graph Laplacian). By retaining or emphasizing certain frequency bands, one can explicitly model global vs. local patterns.

Graph Signal as Combination of Laplacian Eigenvectors. A scalar function $f : \mathcal{V} \rightarrow \mathbb{R}$ on the nodes can be seen as a vector $\mathbf{f} \in \mathbb{R}^N$. The graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$ (or a symmetrized variant) admits an eigen-decomposition:

$$\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top, \quad (4)$$

where $\mathbf{U} = [\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}]$ is an orthonormal matrix of eigenvectors, i.e. $\mathbf{u}_k \in \mathbb{R}^N$, and $\mathbf{\Lambda} = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{N-1})$ contains ascending eigenvalues $0 \leq \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{N-1}$.

In graph signal processing, the eigenvector \mathbf{u}_k is viewed as the “ k -th frequency basis.” Smaller eigenvalues ($\lambda_0, \lambda_1, \dots$) correspond to low-frequency (more global, smooth) variations on the graph, while larger eigenvalues correspond to high-frequency (local, rapidly changing) modes. Hence, any signal $\mathbf{f} \in \mathbb{R}^N$ can be written as a linear combination:

$$\mathbf{f} = \sum_{k=0}^{N-1} \alpha_k \mathbf{u}_k, \quad \text{where } \alpha_k = \mathbf{u}_k^\top \mathbf{f}. \quad (5)$$

This collection of coefficients $\{\alpha_k\}$ is the Graph Fourier Transform (GFT) of \mathbf{f} . We **now denote** α_k as \tilde{f}_k .

Truncation at M Bases. For large N , we often use only the first M eigenvectors $\{\mathbf{u}_0, \dots, \mathbf{u}_{M-1}\}$ to approximate \mathbf{f} . This yields a band-limited or multiscale representation:

$$\mathbf{f} \approx \sum_{k=0}^{M-1} \alpha_k \mathbf{u}_k, \quad (6)$$

where $M \leq N$. Sorting λ_k in ascending order ensures the lowest-frequency modes (i.e., global scales) appear first, and higher-frequency (local scales) modes appear last. By choosing M appropriately, we focus on the most critical modes for modeling global vs. local behavior.

Applying the GFT to Scalar vs. Vector Features. For *scalar* features $\mathbf{H} \in \mathbb{R}^{N \times F}$, where F is the latent feature size, we substitute \tilde{f}_k for $\tilde{\mathbf{H}}_k$ in the equation $\tilde{f}_k = \mathbf{u}_k^\top \mathbf{f}$ (from Equation 5) to get:

$$\tilde{\mathbf{H}} = [\tilde{\mathbf{H}}_k]_{k=0}^{M-1} = \mathbf{U}_{(:,0:M)}^\top \mathbf{H}, \quad (7)$$

where $\mathbf{U}_{(:,0:M)}$ denotes the first M eigenvectors. For *vector* features $\mathbf{Z} \in \mathbb{R}^{N \times m \times 3}$, where m is the feature size of the vector feature in the 3D space, we first remove the mean to ensure translational invariance (since the 0-th eigenvector \mathbf{u}_0 corresponds to the constant mode):

$$\mathbf{Z}_c = \mathbf{Z} - \bar{\mathbf{Z}}. \quad (8)$$

where $\bar{\mathbf{Z}}$ is the global mean over all nodes. We then apply the same truncated basis $\mathbf{U}_{(:,0:M)}^\top$ to each coordinate dimension, substituting \tilde{f}_k for $\tilde{\mathbf{Z}}_k$ in the equation $\tilde{f}_k = \mathbf{u}_k^\top \mathbf{f}$ (from Equation 5) to get:

$$\tilde{\mathbf{Z}} = \mathbf{U}_{(:,0:M)}^\top \mathbf{Z}_c \in \mathbb{R}^{M \times m \times 3}. \quad (9)$$

Indices closer to $k = 0$ indicate more global motions, while larger k (up to $M - 1$) captures more local, high-frequency fluctuations. In practice, M can be a hyperparameter that determines how many eigenmodes we keep, balancing efficiency (fewer modes to evolve) and accuracy (how many scales are captured).

Theoretical Analysis. The Graph Fourier Transform (GFT) decomposes molecular dynamics into modes corresponding to global and local spatial patterns, as determined by the eigenvalues and eigenvectors of the graph Laplacian. As demonstrated in Proposition 3.1, low-frequency modes capture smooth, global deformations, while high-frequency modes represent localized structural variations. This decomposition enables efficient modeling of multiscale spatial dynamics.

Proposition 3.1 (Global vs. Local Spatial Scales). *Let \mathbf{u}_k be the k -th eigenvector of \mathbf{L} with eigenvalue λ_k . Suppose \mathbf{x} encodes atomic coordinates or their latent features. Then:*

1. **If λ_k is small**, the corresponding mode \mathbf{u}_k represents slowly varying (global) deformations across the molecule.

2. **If λ_k is large**, the corresponding mode \mathbf{u}_k represents rapidly changing (local) structural variations.

Regarding the number of modes M we need to use, Theorem 3.2 guarantees that we can use the first few modes to represent the entire system accurately:

Theorem 3.2 (Spectral Truncation Error). *Truncating the spectral representation to the first M modes, $\mathbf{x}_{(M)} = \mathbf{U}_{(:,0:M)} \mathbf{U}_{(:,0:M)}^\top \mathbf{x}$, yields an ℓ^2 -norm error:*

$$\|\mathbf{x} - \mathbf{x}_{(M)}\|_2^2 = \sum_{k=M}^{N-1} |\mathbf{u}_k^\top \mathbf{x}|^2. \quad (10)$$

For α -bandlimited signals, choosing M such that $\lambda_{M-1} \leq \alpha$ guarantees exact recovery.

We provide detailed proofs and derivations in Appendix C, where we analyze the spectral decomposition’s efficacy and demonstrate its suitability for multiscale molecular modeling.

3.4 Neural ODEs in the Spectral Domain

We propose a novel approach that models the temporal evolution of the **Fourier coefficients** using Neural ODEs. By representing molecular interactions in the frequency domain, this approach enables the decomposition of dynamics across different spatial scales and provides a more compact representation for modeling temporal evolution. Neural ODEs then learn the dynamics of these Fourier coefficients over time, offering a continuous-time framework that captures multiscale spatial and temporal interactions simultaneously.

NeuralODE Preliminaries. Neural Ordinary Differential Equations (Neural ODEs) (Chen et al., 2018) provide a framework for modeling continuous-time dynamics by learning the evolution of a system as a set of differential equations parameterized by a neural network. Specifically, for a system’s state $\mathbf{h}(t)$, its temporal evolution is defined as:

$$\frac{d\mathbf{h}(t)}{dt} = f_\theta(\mathbf{h}(t), t), \quad (11)$$

where f_θ is a neural network parameterized by θ that learns the dynamics of $\mathbf{h}(t)$ over time. Given an initial state $\mathbf{h}(t_0)$, the state at any future time t is computed by solving the ODE:

$$\mathbf{h}(t) = \mathbf{h}(t_0) + \int_{t_0}^t f_\theta(\mathbf{h}(\tau), \tau) d\tau. \quad (12)$$

This integral can be evaluated numerically using adaptive ODE solvers, such as Runge-Kutta methods, allowing Neural ODEs to handle irregularly sampled data and continuously model temporal dynamics.

Joint Scalar-Vector Block-Diagonal Formulation. We update the spectral coefficients $[\tilde{\mathbf{H}}, \tilde{\mathbf{Z}}]$ jointly. We can write the combined spectral features as

$$\begin{pmatrix} \tilde{\mathbf{H}} \\ \tilde{\mathbf{Z}} \end{pmatrix} \in \mathbb{R}^{N \times F} \oplus \mathbb{R}^{N \times m \times 3}, \quad (13)$$

and define Neural ODE dynamics over continuous time t :

$$\frac{d}{dt} \begin{pmatrix} \tilde{\mathbf{H}}(t) \\ \tilde{\mathbf{Z}}(t) \end{pmatrix} = \begin{pmatrix} f_\theta(\tilde{\mathbf{H}}(t), t) \\ g_\theta(\tilde{\mathbf{Z}}(t), t) \end{pmatrix}. \quad (14)$$

Mathematically, in the Fourier space, this equates to a block-diagonal operator \mathbf{M}_θ . Let $\tilde{\mathbf{H}}_\omega$ and $\tilde{\mathbf{Z}}_\omega$ denote the coefficients for each frequency mode ω . A single ODE step can be represented as:

$$\begin{pmatrix} \tilde{\mathbf{H}}_\omega \\ \tilde{\mathbf{Z}}_\omega \end{pmatrix} \mapsto \begin{bmatrix} \mathbf{M}_\theta^{(\mathbf{h})} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_\theta^{(\mathbf{z})} \end{bmatrix} \cdot \begin{pmatrix} \tilde{\mathbf{H}}_\omega \\ \tilde{\mathbf{Z}}_\omega \end{pmatrix}. \quad (15)$$

Here, $\mathbf{M}_\theta^{(\mathbf{h})}$ operates on the scalar channels and $\mathbf{M}_\theta^{(\mathbf{z})}$ on the vector channels, allowing them to evolve in a coordinated but distinct manner. This is beneficial to preserving the 3D interactions while keeping track of the nuanced representations in the latent space – thus preserving the 3D interactions while capturing the nuanced latent representations in a block-diagonal design.

Implementation wise, each ODE function f_θ or g_θ takes the current spectral modes ($\tilde{\mathbf{H}}_\omega$ or $\tilde{\mathbf{Z}}_\omega$) and the time t , and produces their rate of change for the integrator. Specifically, we use multi-head self-attention across these modes so that each frequency mode can attend to others. Formally, if $\mathbf{h}_\omega(t)$ denotes the coefficients corresponding to frequency mode ω at time t , the self-attention step can be written as

$$\mathbf{h}'_\omega(t) = \text{MHAttn}(\mathbf{h}_\omega(t)), \quad (16)$$

where MHAttn is the multi-head attention layer across different frequency modes. Next, the time t is directly concatenated along the feature dimension via a learned embedding $\gamma(t) \in \mathbb{R}^d$, yielding

$$\mathbf{h}''_\omega(t) = [\mathbf{h}'_\omega(t) \parallel \gamma(t)]. \quad (17)$$

Finally, each mode $\mathbf{h}''_\omega(t)$ is transformed by a mode-wise linear weight \mathbf{W}_ω , resulting in the derivative:

$$f_\theta(\mathbf{h}_\omega(t), t) = \mathbf{W}_\omega \mathbf{h}''_\omega(t). \quad (18)$$

A similar procedure applies for g_θ , handling any additional vector channels by suitably reshaping and performing attention across modes.

To predict at time t_k , we numerically integrate $\frac{d}{dt}[\tilde{\mathbf{H}}, \tilde{\mathbf{Z}}]$ from t_0 to t_k (e.g., via `dopri5`). Low-frequency components capture long-range, slower dynamics, while high-frequency components capture faster local fluctuations.

3.5 From Spectral Domain to Physical Domain

Inverse GFT. After evolving the spectral coefficients $\tilde{\mathbf{H}}(t_k)$ and $\tilde{\mathbf{Z}}(t_k)$, we recover the node-level signals via the inverse GFT:

$$\mathbf{H}(t_k) = \mathbf{U}_{(:,0:M)} \tilde{\mathbf{H}}(t_k), \quad \mathbf{Z}_c(t_k) = \mathbf{U}_{(:,0:M)} \tilde{\mathbf{Z}}(t_k), \quad (19)$$

where $\mathbf{U}_{(:,0:M)}$ (the matrix of the first M eigenvectors) is the same truncated basis used during the forward GFT. Finally, we restore the global translation by adding back the mean $\bar{\mathbf{Z}}$ that was subtracted earlier:

$$\mathbf{Z}(t_k) = \mathbf{Z}_c(t_k) + \bar{\mathbf{Z}}. \quad (20)$$

Graph Neural Network Decoder. We refine local interactions at each predicted time t_k with a GNN decoder that again operates on $[\mathbf{H}(t_k), \mathbf{Z}(t_k)]$. This step can help capture short-range correlations that may not be fully resolved in the spectral update. The decoder GNN has a structure similar to the encoder:

$$\mathbf{m}_{ij}^{(\text{dec})} = \phi'_e(\mathbf{h}_i(t_k), \mathbf{h}_j(t_k), \mathbf{r}_{ij}^{(\text{dec})}(t_k)), \quad (21)$$

$$\mathbf{h}_i^{(\text{dec})}(t_k) = \phi'_h(\mathbf{h}_i(t_k), \sum_{j \in \mathcal{N}(i)} \mathbf{m}_{ij}^{(\text{dec})}), \quad (22)$$

$$\mathbf{x}_i^{(\text{dec})}(t_k) = \mathbf{x}_i(t_k) + \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \psi'(\mathbf{m}_{ij}^{(\text{dec})}), \quad (23)$$

and similarly for \mathbf{v}_i . Here, $\mathbf{r}_{ij}^{(\text{dec})}(t_k) = \mathbf{x}_i(t_k) - \mathbf{x}_j(t_k)$, and ϕ'_e, ϕ'_h, ψ' are learnable functions.

4 Experiments

In this section, we provide a comprehensive evaluation of our proposed model on molecular dynamics datasets, comparing its performance against several baselines and conducting detailed ablation studies to assess the contributions of the different components. Our experimental evaluation is designed to address the following key **research questions**:

- **RQ1: Prediction Accuracy.** Does the proposed GF-NODE framework deliver improved molecular dynamics prediction accuracy compared to state-of-the-art methods?
- **RQ2: Continuous Time Modeling.** How effectively does the continuous-time evolution component capture multiscale temporal dynamics—including long-horizon forecasting and super-resolution—compared to variants without continuous-time propagation?
- **RQ3: Spatial Multiscale Modeling.** How crucial is the explicit spectral decomposition for capturing spatial multiscale interactions, and how do different Fourier mode interaction schemes and the number of retained modes affect overall performance?

4.1 Dataset

We evaluate our model using the MD17 dataset (Chmiela et al., 2017), which contains molecular dynamics trajectories for eight small molecules: Aspirin, Benzene, Ethanol, Malonaldehyde, Naphthalene, Salicylic Acid, Toluene, and Uracil. The dataset provides atomic trajectories simulated under quantum mechanical forces, capturing realistic molecular motions.

We also evaluated our model on the alanine dipeptide dataset (Schreiner et al., 2024), a standard benchmark for studying conformational dynamics in proteins. Our task is to predict the future positions of atoms given the initial state of the molecular system.

4.2 Experimental Setup

For each molecule, we partition the trajectory data into training, validation, and test sets, using 500 samples for training, 2000 for validation, and 2000 for testing. The time scope ΔT for each piece of data is set to 3000 simulation steps, providing a challenging prediction horizon.

A key aspect of our experimental setup is the use of **irregular timestep sampling**, in contrast to the equi-timestep sampling used in some baseline models like EGNO, to better mimic the variable time intervals in real-world physical systems. This setting tests the models’ ability to handle irregular temporal data. Although each trajectory spans 3000 simulation steps, we randomly sample only 8 data points per instance. Because these samples are drawn from different points in the 3000-step window across data instances, this strategy enables the model to learn the dynamics over the entire time span without requiring training on every timestep, thereby significantly enhancing **efficiency**. Nevertheless, we also provide evaluations based on equi-timestep sampling in Appendix B for completeness.

4.3 Baseline Models

We compare our model against several state-of-the-art approaches:

- **NDCN** (Zang & Wang, 2020): A Graph Neural ODE model that integrates graph neural networks into the ODE framework to learn continuous-time dynamics of networked systems.
- **LG-ODE** (Huang et al., 2020): A latent graph-based ODE model that integrates latent representations into continuous-time evolution.
- **EGNN** (Satorras et al., 2021): An Equivariant Graph Neural Network that models molecular systems using 3D equivariant message passing but without explicit time propagation.
- **EGNO** (Xu et al., 2024): An Equivariant Graph Neural Operator that captures temporal dynamics using neural operators with regular timesteps.
- **ITO** (Schreiner et al., 2024): An Implicit Time-stepping Operator that integrates differential equations into the learning process for temporal evolution.

These baselines represent a range of approaches for modeling molecular dynamics, including methods that focus primarily on spatial modeling (EGNN) or temporal modeling (NDCN, LG-ODE, EGNO, ITO).

4.4 Results and Analysis (RQ1)

Units and Calculations: Note that the alanine dipeptide dataset operates in nanometers (nm), whereas the MD17 dataset uses angstroms (Å). When calculating bond lengths and bond angles, we consider all heavy atoms (excluding hydrogen) to focus on the core structure.




Figure 4: Representative snapshots of alanine dipeptide, comparing our model’s predicted conformations (top row) against the ground-truth simulation (bottom row) at four different timestamps. The close agreement illustrates the model’s ability to preserve key structural features over an extended trajectory.

To address **RQ1**, we evaluate the performance of our model and the baselines on the MD17 dataset and alanine dipeptide dataset, and Figure 4 provides visualization of our predictions. The test Mean Squared Error (MSE) on the eight molecules in MD17 and alanine dipeptide for our model and the baseline methods are summarized in Tables 1 and 3, under **irregular** timestep sampling. For completeness, we also provide results in the equi-timestep setting in Table 7 in Appendix B.

Table 1: MSE ($\times 10^{-2} \text{ \AA}^2$) on the MD17 dataset with **irregular** timestep sampling. Best results are in **bold**, and second-best are underlined.

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
NDCN	29.75 \pm 0.02	70.13 \pm 0.98	10.05 \pm 0.02	42.28 \pm 0.07	2.30 \pm 0.00	3.43 \pm 0.05	12.33 \pm 0.00	2.39 \pm 0.00
LG-ODE	51.65 \pm 0.01	68.29 \pm 0.21	12.32 \pm 0.05	43.95 \pm 0.07	2.38 \pm 0.02	2.85 \pm 0.08	18.11 \pm 0.09	2.38 \pm 0.07
EGNN	<u>9.09</u> \pm 0.10	<u>49.15</u> \pm 1.68	<u>4.46</u> \pm 0.01	<u>12.52</u> \pm 0.05	<u>0.40</u> \pm 0.02	<u>0.89</u> \pm 0.01	<u>8.98</u> \pm 0.09	<u>0.64</u> \pm 0.00
EGNO	10.60 \pm 0.01	52.53 \pm 2.40	4.52 \pm 0.06	12.89 \pm 0.06	0.46 \pm 0.01	1.07 \pm 0.00	9.31 \pm 0.10	0.67 \pm 0.01
ITO	12.74 \pm 0.10	57.84 \pm 0.86	7.23 \pm 0.00	19.53 \pm 0.01	1.77 \pm 0.01	2.53 \pm 0.03	9.96 \pm 0.04	1.71 \pm 0.15
Ours	6.46 \pm 0.03	1.52 \pm 0.08	2.74 \pm 0.05	10.54 \pm 0.01	0.23 \pm 0.02	0.63 \pm 0.01	1.80 \pm 0.05	0.41 \pm 0.01

Table 2: Mean Absolute Errors (MAEs) and relative errors for bond lengths and bond angles on alanine dipeptide. Best results are in **bold**.

Model	Bond Length MAE (nm)	Rel. Err. (%)	Bond Angle MAE ($^\circ$)	Rel. Err. (%)
EGNN	0.0209 \pm 0.0006	15.32 \pm 0.49	12.44 \pm 0.91	10.48 \pm 0.76
EGNO	0.0229 \pm 0.0018	16.75 \pm 1.23	10.54 \pm 0.11	8.89 \pm 0.11
Ours	0.0188 \pm 0.0022	13.74 \pm 1.66	10.47 \pm 1.03	8.80 \pm 0.89

From Table 1, we observe that our model consistently outperforms the baseline methods across all eight molecules under irregular timestep sampling. This demonstrates the effectiveness of our approach in jointly modeling spatial and temporal multiscale interactions. The performance gains are particularly pronounced for molecules such as Benzene and Aspirin, where our model significantly reduces the MSE compared to the baselines.

Benzene Drifting. Interestingly, we find that Benzene exhibits substantial drift during simulation. In our data, the maximum

Table 3: Mean Squared Error (MSE) ($\times 10^{-3} \text{ nm}^2$) on the alanine dipeptide dataset. Best results are in **bold**.

Model	MSE ($\times 10^{-3} \text{ nm}^2$)
NDCN	12.27 \pm 0.19
ITO	26.95 \pm 0.19
EGNO	6.92 \pm 0.26
EGNN	5.67 \pm 0.08
Ours	4.48 \pm 0.07

x -coordinate is 197.981 Å, while the minimum x -coordinate is -178.112 Å, indicating large translations and rotations of the entire molecule. In contrast, other molecules exhibit minimal net translation, typically remaining within ± 3 Å from the origin. Methods that enforce strict invariance to translations and rotations (e.g., EGNN) may underperform on such drifting systems, since the global drift is part of the actual dynamics. Indeed, we discovered that *replacing EGNN-based layers with standard message passing layers* (e.g., SAGEConv) *can further boost performance on drifting molecules* like Benzene. We provide details of the dataset statistics in Table 6 of Appendix A, and experiment results using different GNN architectures in Table 9 of Appendix B.

Experimental Results on Alanine Dipeptide. In addition, from Table 3, our model achieves the lowest MSE among all compared methods. This indicates a significant improvement on more complex molecular dynamics data and demonstrates the robustness of our approach.

Analysis of Molecular Structure Recovery. To further evaluate predictions at a structural level, we analyzed the bond lengths and bond angles for alanine dipeptide. Table 2 shows the Mean Absolute Errors (MAEs) and relative errors for bond lengths and angles. Our method achieves the lowest errors, indicating superior recovery of internal molecular structures compared to baselines.

TICA Analysis. Time-lagged Independent Component Analysis (TICA) (Molgedey & Schuster, 1994) is used to extract slow collective motions from the trajectories. Figure 5 shows that, across various numbers of TICA components, our model consistently achieves lower TICA scores (i.e., better alignment with the underlying slow modes) compared to EGNN and EGNO, indicating more effective capture of the underlying multiscale dynamics.




Figure 5: TICA scores for varying numbers of components. Lower scores indicate better alignment with slow modes.

4.5 Continuous-Time Dynamics and Temporal Super-Resolution (RQ2)

Ablation on continuous time components. To assess our model’s ability to capture temporal multiscale dynamics, we perform an ablation study on the continuous-time components. Specifically, we investigate:

1. **No ODEs Evolution:** Removing the continuous-time evolution altogether.
2. **No ODEs on Scalar Channels:** Freeze the ODEs for among scalar features \mathbf{h} during time propagation.
3. **No ODEs on Vector Channels:** Freeze the ODEs among vector features \mathbf{x} during time propagation.

Table 4 reports the test errors (MSE) for these variants across several molecules. The inferior performance of the ablated models confirms the importance of the continuous-time dynamics and the effectiveness of the block-diagonal architecture that jointly propagates scalar and vector features.

Table 4: Ablation study on continuous-time evolution components ($\text{MSE} \times 10^{-2}$ Å² values). Lower values indicate better performance. All results are **inferior** to our standard model (w/ ode).

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
no_ode	6.56±0.03	1.87±0.08	3.09±0.05	11.58±0.02	0.42±0.02	0.86±0.01	3.06±0.05	0.61±0.01
no_ode_h	6.77±0.05	1.89±0.06	3.11±0.03	10.70±0.02	0.42±0.01	0.88±0.02	3.48±0.04	0.59±0.02
no_ode_x	6.95±0.03	1.79±0.07	3.74±0.05	10.61±0.03	0.42±0.01	0.87±0.01	2.89±0.04	0.59±0.01
w/ ode	6.46±0.03	1.52±0.08	2.74±0.05	10.54±0.01	0.23±0.02	0.63±0.01	1.80±0.05	0.41±0.01

Super-Resolution Task. To further validate temporal generalization, we perform a super-resolution experiment on the alanine dipeptide dataset by predicting trajectories at a $10\times$ finer temporal resolution than the training samples, under the equi-timestep setting. Figure 6 compares the MSE for the original versus the super-resolved predictions. Our model maintains low error under super-resolution, demonstrating its ability to interpolate continuous dynamics effectively.

Long-Term Prediction Stability. We evaluate the stability of long-term predictions. We test models at 1000, 2000, 3000, 4000, and 5000 simulation steps on two representative molecules (Benzene and Malonaldehyde). Figure 7 illustrates how errors evolve over these extended horizons. Our model maintains superior performance with slower error growth, indicating better capture of global low-frequency dynamics essential for accurate long-term predictions. In contrast, baselines lacking explicit multiscale modeling accumulate errors more rapidly.

Table 4 in Appendix B provides additional ablations on the different types of temporal embeddings used in the ODE functions. A simple concatenation of the timestamp would work well enough.




Figure 6: Comparison of MSE on the alanine dipeptide dataset: Original vs. 10× Super-resolution.




Figure 7: Error growth over long-term forecasts (up to 5000 steps) for Benzene and Malonaldehyde. Models with too high an error not presented here.

4.6 Spatial Frequency Decomposition and Multiscale Analysis (RQ3)

To examine the role of spatial multiscale modeling, we perform ablations on the spectral decomposition and mode interaction components. We consider two sets of modifications:

1. Ablation of the Spatial Decomposition:

- *No Fourier-based Decomposition:* The model is run without any frequency-based transformation. The spectral decomposition is replaced by a multi-layer perceptron.
- *Replacing GFT with FFT:* The Graph Fourier Transform is substituted with a standard Fast Fourier Transform, disregarding the graph structure.

2. Fourier Mode Interaction Schemes:

- We compare different strategies for inter-mode communication:
- *Attention-based Interaction:* Each Fourier mode interacts with others via a multi-head self-attention mechanism (Used in our standard model).
 - *Concatenation-based Interaction:* Modes are concatenated before being processed.
 - *No Interaction:* Each mode is propagated independently.

Table 5 compiles the results for these spatial ablations. We observe that using the GFT for spectral decomposition—rather than an MLP or FFT—is most effective, underscoring the importance of capturing the inherent graph structure. Moreover, interaction schemes (whether via attention or concatenation) improve performance over treating modes independently.

Table 5: MSE ($\times 10^{-2} \text{ \AA}^2$) on the MD17 dataset with **irregular** timestep sampling for variants in spatial decomposition and Fourier modes interactions. Best results are in **bold**, the standard model (GFT & Attn.).

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
FFT	6.53±0.03	1.94±0.08	3.36±0.05	10.83±0.01	0.43±0.02	0.88±0.01	3.36±0.05	0.60±0.01
No Fourier	6.63±0.04	1.99±0.10	3.28±0.06	10.68±0.07	0.43±0.02	0.88±0.01	3.73±0.09	0.60±0.05
No Interaction	6.51±0.03	1.78±0.08	3.28±0.05	10.64±0.01	0.41±0.02	0.87±0.02	3.08±0.05	0.59±0.01
Concat.	6.55±0.03	1.76±0.07	3.13±0.05	10.58±0.01	0.43±0.02	0.88±0.01	2.80±0.05	0.58±0.01
GFT & Attn.	6.46±0.03	1.52±0.08	2.74±0.05	10.54±0.01	0.23±0.02	0.63±0.01	1.80±0.05	0.41±0.01

Impact of the Number of Fourier Modes. Finally, we investigate how the number of retained Fourier modes affects performance. Figure 8 plots the MSE against different numbers of modes used in the spectral decomposition. We observe that the performance improves as more modes are included up to a threshold, beyond which additional modes yield diminishing returns. This behavior is consistent with the theoretical analysis presented earlier in Theorem 3.2. The number of modes used to get the optimal results for each type of molecule can be found in Table 6 of Appendix A.

5 Conclusion

In this work, we presented Graph Fourier Neural ODEs (GF-NODE), a novel framework that unifies spatial spectral decomposition with continuous-time evolution to effectively model the multiscale dynamics inherent in molecular systems. By decomposing molecular graphs via the Graph Fourier Transform, our approach explicitly disentangles global conformational changes from local vibrational modes, and the subsequent Neural ODE-based propagation enables flexible, continuous-time forecasting of these dynamics.

Our extensive evaluations on benchmark datasets, including MD17 and alanine dipeptide, demonstrate that GF-NODE achieves state-of-the-art performance in long-horizon trajectory prediction while accurately preserving essential molecular geometries. The ablation studies further highlight the critical role of both the spectral decomposition and the continuous-time dynamics in capturing complex spatial-temporal interactions.

Overall, our results suggest that integrating spectral methods with neural ODEs holds significant promise for advancing MD simulation. We anticipate that this approach will not only improve the predictive accuracy and robustness of molecular simulations but also inspire future research on applying similar multiscale modeling techniques to other dynamic, graph-structured systems.

6 Broader Impact Statement

Our work contributes to advancing molecular dynamics simulations, which have broad implications in scientific discovery, particularly in drug design, materials science, and biomolecular modeling. By improving the accuracy and efficiency of multiscale molecular predictions, our approach could accelerate the discovery of novel therapeutics and facilitate the design of functional materials with tailored properties.

While our method relies on data-driven modeling, it does not replace physics-based simulations but rather augments them, reducing computational costs while maintaining interpretability. As with any machine learning-driven approach in scientific domains, care must be taken to ensure model reliability, particularly in high-stakes applications such as drug development. Further validation and collaboration with domain experts will be essential to maximize the positive societal impact of our work while mitigating risks related to model uncertainty.




Figure 8: Effect of the number of Fourier modes on prediction performance. Performance plateaus beyond $n = 12$.

References

- Michael Bear and Estela Blaisten-Barojas. Atomic and ion collisions in solids and at surfaces: Theory, simulation and applications. *American Scientist*, 86(2):188, 1998.
- Bernard R Brooks, Charles L Brooks III, Alexander D Mackerell Jr, Lennart Nilsson, Robert J Petrella, Benoît Roux, Youngdo Won, Georgios Archontis, Christian Bartels, Stefan Boresch, et al. Charmm: the biomolecular simulation program. *Journal of computational chemistry*, 30(10):1545–1614, 2009.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Stefan Chmiela, Alexandre Tkatchenko, Huziel E Sauceda, Igor Poltavsky, Kristof T Schütt, and Klaus-Robert Müller. Machine learning of accurate energy-conserving molecular force fields. *Science advances*, 3(5): e1603015, 2017.
- Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, Providence, R.I., 1997.
- Wendy D Cornell, Piotr Cieplak, Christopher I Bayly, Ian R Gould, Kenneth M Merz, David M Ferguson, David C Spellmeyer, Thomas Fox, James W Caldwell, and Peter A Kollman. A second generation force field for the simulation of proteins, nucleic acids, and organic molecules. *Journal of the American Chemical Society*, 117(19):5179–5197, 1995.
- Pawan Goyal and Peter Benner. Neural ordinary differential equations with irregular and noisy data. *Royal Society Open Science*, 10(7):221475, 2023.
- Samuel I Holt, Zhaozhi Qian, and Mihaela van der Schaar. Neural laplace: Learning diverse classes of differential equations in the laplace domain. In *International Conference on Machine Learning*, pp. 8811–8832. PMLR, 2022.
- Zijie Huang, Yizhou Sun, and Wei Wang. Learning continuous system dynamics from irregularly-sampled partial observations. *Advances in Neural Information Processing Systems*, 33:16177–16187, 2020.
- Zijie Huang, Yizhou Sun, and Wei Wang. Coupled graph ode for learning interacting system dynamics. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 705–715, 2021.
- Zijie Huang, Yizhou Sun, and Wei Wang. Generalizing graph ode for learning complex system dynamics across environments. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 798–809, 2023a.
- Zijie Huang, Wanjia Zhao, Jingdong Gao, Ziniu Hu, Xiao Luo, Yadi Cao, Yuanzhou Chen, Yizhou Sun, and Wei Wang. Tango: Time-reversal latent graphode for multi-agent dynamical systems. *arXiv preprint arXiv:2310.06427*, 2023b.
- Jürg Hutter. Car-parrinello molecular dynamics. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(4):604–612, 2012.
- Patrick Kidger. On neural differential equations. *arXiv preprint arXiv:2202.02435*, 2022.
- Suyong Kim, Weiqi Ji, Sili Deng, Yingbo Ma, and Christopher Rackauckas. Stiff neural ordinary differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(9), 2021.
- Leon Klein, Andrew Foong, Tor Fjelde, Bruno Mlodozieniec, Marc Brockschmidt, Sebastian Nowozin, Frank Noé, and Ryota Tomioka. Timewarp: Transferable acceleration of molecular dynamics by learning time-coarsened dynamics. *Advances in Neural Information Processing Systems*, 36, 2024.
- Takeshi Koshizuka, Masahiro Fujisawa, Yusuke Tanaka, and Issei Sato. Understanding the expressivity and trainability of fourier neural operator: A mean-field perspective. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

- Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for fourier neural operators. *Journal of Machine Learning Research*, 22(290):1–76, 2021.
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- Z Li, DZ Huang, B Liu, and A Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. arxiv, 2022a.
- Zijie Li, Kazem Meidani, Prakarsh Yadav, and Amir Barati Farimani. Graph neural networks accelerated molecular dynamics. *The Journal of Chemical Physics*, 156(14), 2022b.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Ning Liu and Siavash Jafarzadeh. Yue yu. domain agnostic fourier neural operators. *arXiv preprint arXiv:2305.00478*, 2023.
- Xiao Luo, Yiyang Gu, Huiyu Jiang, Hang Zhou, Jinsheng Huang, Wei Ju, Zhiping Xiao, Ming Zhang, and Yizhou Sun. Pgode: Towards high-quality system dynamics modeling. In *Forty-first International Conference on Machine Learning*.
- Xiao Luo, Jingyang Yuan, Zijie Huang, Huiyu Jiang, Yifang Qin, Wei Ju, Ming Zhang, and Yizhou Sun. Hope: High-order graph ode for modeling interacting dynamics. In *International Conference on Machine Learning*, pp. 23124–23139. PMLR, 2023.
- Lutz Molgedey and Heinz Georg Schuster. Separation of a mixture of independent signals using time delayed correlations. *Physical review letters*, 72(23):3634, 1994.
- Outi MH Salo-Ahen, Ida Alanko, Rajendra Bhadane, Alexandre MJJ Bonvin, Rodrigo Vargas Honorato, Shakhawath Hossain, André H Juffer, Aleksei Kabedev, Maija Lahtela-Kakkonen, Anders Støttrup Larsen, et al. Molecular dynamics simulations in drug discovery and pharmaceutical development. *Processes*, 9(1): 71, 2020.
- Victor Garcia Satorras, Emiel Hoogetboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.
- Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: Multiple time-resolution models for molecular dynamics. *Advances in Neural Information Processing Systems*, 36, 2024.
- Paulo CT Souza, Riccardo Alessandri, Jonathan Barnoud, Sebastian Thallmair, Ignacio Faustino, Fabian Grünewald, Ilias Patmanidis, Haleh Abdizadeh, Bart MH Bruininks, Tsjerk A Wassenaar, et al. Martini 3: a general purpose force field for coarse-grained molecular dynamics. *Nature methods*, 18(4):382–388, 2021.
- Antonis I Vakis, Vladislav A Yastrebov, Julien Scheibert, Lucia Nicola, Daniele Dini, Clotilde Minfray, Andreas Almqvist, Marco Paggi, Seunghwan Lee, Georges Limbert, et al. Modeling and simulation in tribology across scales: An overview. *Tribology International*, 125:169–199, 2018.
- Quan-De Wang, Jing-Bo Wang, Juan-Qin Li, Ning-Xin Tan, and Xiang-Yuan Li. Reactive molecular dynamics simulation and chemical kinetic modeling of pyrolysis and combustion of n-dodecane. *Combustion and Flame*, 158(2):217–226, 2011.
- Minkai Xu, Jiaqi Han, Aaron Lou, Jean Kossaifi, Arvind Ramanathan, Kamyar Azizzadenesheli, Jure Leskovec, Stefano Ermon, and Anima Anandkumar. Equivariant graph neural operator for modeling 3d dynamics. *arXiv preprint arXiv:2401.11037*, 2024.
- Chengxi Zang and Fei Wang. Neural dynamics on complex networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 892–902, 2020.

A Experiment Setup

In this section, we describe key details about the datasets used (MD17 and alanine dipeptide), including the simulation step sizes, and then outline the hyperparameter choices and training procedures.

A.1 Datasets

We use the MD17 dataset (Chmiela et al., 2017) for small-molecule dynamics, and the alanine dipeptide dataset (Schreiner et al., 2024) for conformational analysis in proteins. Statistics for MD17 are provided in Table 6. From the table, we observe Benzene has a very marked drifting that is much larger than the scale of the molecule itself. The MD17 dataset was generated using ab initio molecular dynamics (MD) simulations with a time step of 0.5 femtoseconds (fs). The simulations were conducted in the NVT ensemble at 500 K for a total duration of 200 picoseconds (ps). The final dataset was created by subsampling the full trajectory, preserving the Maxwell-Boltzmann distribution for the energies. For the alanine dipeptide dataset, we use the protein fragment’s trajectory recorded at a step size of 1.0 ps.

Table 6: Summary statistics for molecular structures, including the number of atoms, position extrema (X_{\min} , X_{\max} , X_{mean}), and velocity extrema (V_{\min} , V_{\max} , V_{mean}). The numbers of modes used to get the optimal results are also listed.

Molecule	#Atoms	#Modes used	X_{\min}	X_{\max}	X_{mean}	V_{\min}	V_{\max}	V_{mean}
benzene	6	4	-178.112	197.981	-27.737	-0.004	0.003	-0.000
aspirin	13	6	-3.720	3.105	0.026	-0.011	0.012	0.000
ethanol	3	3	-1.398	1.417	-0.004	-0.011	0.010	-0.000
malonaldehyde	5	5	-2.397	2.370	0.000	-0.010	0.009	0.000
naphthalene	10	4	-2.597	2.593	-0.000	-0.012	0.011	0.000
salicylic	10	8	-2.734	2.581	-0.051	-0.013	0.012	-0.000
toluene	7	7	-1.990	2.630	-0.015	-0.010	0.012	0.000
uracil	8	6	-2.338	2.558	0.012	-0.012	0.011	0.000

A.2 Training Setup

Hyperparameters. We train all models using the Adam optimizer at a learning rate of 1×10^{-4} and apply a weight decay of 1×10^{-15} for regularization. Each experiment runs for 5000 epochs, processing batches of size 50 at each training step. For molecular trajectory prediction, we set the sequence length (the number of timesteps) to 8, meaning each training sample contains 8 frames from the overall simulation. In our main experiments, we fix the hidden feature dimensionality to 64 and use a maximum future horizon of 3000 simulation steps when constructing training samples. We also specify the dopri5 ODE solver with relative and absolute tolerances of 1×10^{-3} and 1×10^{-4} , respectively, to integrate the continuous-time model components. Although many of these choices (e.g., total layers, solver tolerances) can be altered, we find these particular settings maintain a good balance of accuracy and computational efficiency.

Training Procedure. During training, each mini-batch is formed by sampling short segments of length 8 from the molecule’s dynamics trajectory. The model then predicts future positions of atoms after continuous-time evolution, and the Mean Squared Error (MSE) between predicted coordinates and ground-truth coordinates is minimized. We checkpoint models whenever validation performance improves, and at the end of training, we report results using the best-performing checkpoint according to the validation set. In addition, regular evaluations on the test set help track the model’s generalization to unseen trajectories. The overall process can be summarized as: (1) load training samples, (2) form mini-batches of molecular frames, (3) perform forward pass through the model to generate predictions, (4) compute the MSE loss, (5) update model parameters via backpropagation, repeating for each epoch until convergence.

By following these procedures with our chosen hyperparameters, we have observed stable convergence across both MD17 and alanine dipeptide datasets, as well as strong generalization to different segments of the trajectory during test runs.

B Additional Experiment Results

Below, we provide four tables corresponding to different experimental comparisons. The first table reports performance on the MD17 dataset with regular (equi-) timesteps. The remaining three tables focus on ablation experiments conducted under irregular sampling conditions: (i) ablations on the alanine dipeptide dataset, (ii) ablations comparing different GNN architectures, and (iii) ablations on different temporal embedding approaches.

Table 7: MD17 with Regular (Equi-) Timestep Sampling. MSE ($\times 10^{-2}$ Å²) on the MD17 dataset using regular timesteps. Best results are in **bold**, and second-best are underlined.

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
NDCN	31.73±0.40	56.21±0.30	10.74±0.02	46.55±0.28	2.25±0.01	3.58±0.11	13.92±0.02	2.38±0.00
LG-ODE	19.36±0.12	53.92±1.32	7.08±0.01	24.41±0.03	1.73±0.02	3.82±0.04	11.18±0.01	2.11±0.02
EGNN	<u>9.24</u> ±0.07	57.85±2.70	<u>4.63</u> ±0.00	<u>12.81</u> ±0.01	<u>0.38</u> ±0.01	<u>0.85</u> ±0.00	<u>10.41</u> ±0.04	<u>0.56</u> ±0.02
EGNO	9.41±0.09	<u>55.13</u> ±3.21	<u>4.63</u> ±0.00	<u>12.81</u> ±0.01	0.40±0.01	0.93±0.01	10.43±0.10	0.59±0.01
ITO	20.56±0.03	57.25±0.58	8.60±0.27	28.44±0.73	1.82±0.17	2.48±0.34	12.47±0.30	1.33±0.12
Ours	6.07 ±0.09	1.51 ±0.07	2.74 ±0.01	9.43 ±0.02	0.24 ±0.02	0.63 ±0.05	1.80 ±0.03	0.41 ±0.02

Explanation. Table 7 shows results for MD17 under a regular (evenly spaced) sampling scheme. Although the dataset inherently has fine-grained timesteps, we constrain both training and evaluation to equidistant frames to compare methods fairly. Our approach demonstrates consistent improvements over baselines on nearly all molecules.

Table 8: **Ablation Results on Alanine Dipeptide (Irregular Sampling).** MSE ($\times 10^{-3}$ nm²). Best results in **bold**.

	standard	no_ode	no_ode_h	no_ode_x	no_interaction
MSE	4.48 ±0.07	4.72±0.05	4.60±0.05	4.64±0.05	4.80±0.04
	interaction_concat	time_posenc	time_mlp	FFT	no_fourier
MSE	4.98±0.06	4.62±0.05	4.60±0.05	4.57±0.05	4.51±0.04

Explanation. Table 8 organizes the ablation settings into two rows, each containing five columns. The first row compares our “standard” model to variants that remove specific ODE blocks or modify scalar/vector-only ODE updates (“no_ode”, “no_ode_h”, “no_ode_x”), and the second row compares different interaction modes, time embeddings, and Fourier settings. The “standard” configuration achieves the best overall MSE.

Table 9: **Ablation on GNN Architectures (Irregular Sampling).** MSE ($\times 10^{-2}$ Å²) on MD17 comparing different GNN layers (SAGEConv, GCNConv, EGNNConv). Best results in **bold**.

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
SAGEConv	6.46 ±0.03	1.79±0.08	2.74 ±0.05	10.54 ±0.01	0.23 ±0.02	0.63 ±0.01	3.08±0.05	0.41 ±0.01
GCNConv	6.91±0.02	1.52 ±0.08	3.09±0.06	10.85±0.03	0.42±0.01	0.88±0.00	1.80 ±0.05	0.61±0.02
EGNNConv	8.85±0.02	40.86±0.98	4.41±0.06	12.49±0.00	0.40±0.01	0.87±0.01	8.63±0.04	0.62±0.02

Explanation. Table 9 shows how our model performs with different GNN backbones on MD17 under irregular sampling. Overall, SAGEConv yields robust performance for most molecules, whereas GCNConv provides better results specifically on Benzene and Toluene. EGNNConv performs well on some local metrics but struggles on large translations (i.e., Benzene).

Table 10: **Ablation on Time Embedding Approaches (Irregular Sampling).** MSE ($\times 10^{-2}$ Å²) on MD17 across different time encoding schemes.

	Aspirin	Benzene	Ethanol	Malonaldehyde	Naphthalene	Salicylic	Toluene	Uracil
posenc	6.91±0.08	1.81±0.02	3.11±0.05	10.69±0.03	4.19±0.04	0.87±0.01	3.56±0.07	0.59±0.02
mlp	6.61±0.06	1.61±0.03	3.08±0.02	10.62±0.04	0.41±0.01	0.87±0.02	3.25±0.06	0.56±0.01
concat	6.46 ±0.03	1.52 ±0.08	2.74 ±0.05	10.54 ±0.01	0.23 ±0.02	0.63 ±0.01	1.80 ±0.05	0.41 ±0.01

Explanation. Table 10 compares three different time-embedding methods under irregular timestep sampling: positional encoding (posenc), a small MLP (mlp), and a direct concatenation of time tokens (concat). Concatenation achieves the lowest MSE, suggesting that a straightforward inclusion of time in the feature vector can be beneficial, though the MLP variant also achieves competitive performance on several molecules.

C Theoretical Guarantees on Spectral Decompositions

Below, we present a concise mathematical exposition on the theoretical underpinnings of the Graph Fourier Transform (GFT) decomposition used in our framework. We explain how the eigenvalue–eigenvector structure of the graph Laplacian \mathbf{L} induces a decomposition of graph signals into low-frequency (global) and high-frequency (local) modes, and we justify truncating to the first M modes. We follow standard nomenclature in spectral graph theory (Chung, 1997).

C.1 Preliminaries and Definitions

Definition C.1 (Graph Laplacian). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with $N = |\mathcal{V}|$ vertices. Let $\mathbf{A} \in \mathbb{R}^{N \times N}$ be its adjacency matrix, and let \mathbf{D} be the diagonal degree matrix, where

$$\mathbf{D}(i, i) = \sum_{j=1}^N \mathbf{A}(i, j). \quad (24)$$

The *graph Laplacian* is defined as

$$\mathbf{L} = \mathbf{D} - \mathbf{A}. \quad (25)$$

It is well known that \mathbf{L} is real symmetric and positive semidefinite. In fact, the eigenvalues of the Laplacian matrix are real and non-negative.

Definition C.2 (Graph Fourier Transform (GFT)). Given the eigen-decomposition

$$\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top, \quad (26)$$

where

$$\mathbf{\Lambda} = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{N-1}), \quad 0 = \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{N-1}, \quad (27)$$

and $\mathbf{U} = [\mathbf{u}_0 \mid \mathbf{u}_1 \mid \dots \mid \mathbf{u}_{N-1}]$ stores the corresponding orthonormal eigenvectors in columns. For a graph signal

$$\mathbf{x} = (x_1, x_2, \dots, x_N)^\top \in \mathbb{R}^N, \quad (28)$$

the *Graph Fourier Transform (GFT)* of \mathbf{x} is given by

$$\hat{\mathbf{x}} = \mathbf{U}^\top \mathbf{x}, \quad (29)$$

and the *inverse GFT* is

$$\mathbf{x} = \mathbf{U} \hat{\mathbf{x}}. \quad (30)$$

C.2 Truncation and Mode Selection

Lemma C.3 (Approximation Error for Spectral Truncation). *Let $\mathbf{x} \in \mathbb{R}^N$ be any graph signal, and let $\mathbf{x}_{(M)}$ be its spectral approximation obtained by keeping the first M modes. Then*

$$\|\mathbf{x} - \mathbf{x}_{(M)}\|_2^2 = \sum_{k=M}^{N-1} |\mathbf{u}_k^\top \mathbf{x}|^2. \quad (31)$$

Moreover, if \mathbf{x} is α -bandlimited in the sense that

$$\mathbf{u}_k^\top \mathbf{x} = 0 \quad \text{for all } \lambda_k > \alpha, \quad (32)$$

then choosing M such that $\lambda_{M-1} \leq \alpha$ yields an exact recovery $\mathbf{x} = \mathbf{x}_{(M)}$.

Proof. See the main text for details. We expand the signal in the Laplacian eigenbasis $\{\mathbf{u}_k\}$, and observe that discarding all modes with $k \geq M$ removes the corresponding frequency components. \square

C.3 Low-Frequency vs. High-Frequency Modes

Because \mathbf{L} is positive semidefinite and the eigenvalues $\{\lambda_i\}$ increase with i , smaller eigenvalues correspond to slow, global variations, while larger eigenvalues capture more oscillatory, local phenomena.

Proposition C.4 (Global vs. Local Spatial Scales). *Let \mathbf{u}_k be the k -th eigenvector of \mathbf{L} with eigenvalue λ_k . Suppose \mathbf{x} encodes atomic coordinates or their latent features. Then:*

1. *If λ_k is small, the corresponding mode \mathbf{u}_k represents slowly varying (global) deformations across the molecule.*
2. *If λ_k is large, the corresponding mode \mathbf{u}_k represents rapidly changing (local) structural variations.*

Proof. From standard results in spectral graph theory (Chung, 1997). The low-frequency (small λ) modes vary smoothly across edges, whereas high-frequency (large λ) modes exhibit large differences across edges. \square

C.4 Practical Mode Truncation Criteria

Definition C.5 (Mode Retention Threshold). For a desired tolerance $\epsilon > 0$, select M such that

$$\sum_{k=M}^{N-1} |\mathbf{u}_k^\top \mathbf{x}|^2 \leq \epsilon \|\mathbf{x}\|_2^2. \quad (33)$$

In practice, one may also pick M based on $\lambda_{M-1} \leq \alpha$, ignoring modes where $\lambda_k > \alpha$.

Corollary C.6 (Error Control via Low-Pass Approximation). *Under the same notation as above, if*

$$\sum_{k=M}^{N-1} |\mathbf{u}_k^\top \mathbf{x}|^2 \leq \epsilon \|\mathbf{x}\|_2^2, \quad (34)$$

then

$$\|\mathbf{x} - \mathbf{x}_{(M)}\|_2 \leq \epsilon \|\mathbf{x}\|_2. \quad (35)$$

Hence, discarding high-frequency modes exceeding this threshold leads to a bounded approximation error.