BOOSTING PROTEIN GRAPH REPRESENTATIONS THROUGH STATIC-DYNAMIC FUSION

Pengkang Guo[1], **Bruno Correia**[1], **Pierre Vandergheynst**[1] **& Daniel Probst**[1,2*]
[1]École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
[2]Wageningen University & Research
Wageningen, The Netherlands

ABSTRACT

Machine learning for protein modeling faces significant challenges due to proteins' inherently dynamic nature, yet most graph-based machine learning methods rely solely on static structural information. Recently, the growing availability of molecular dynamics trajectories provides new opportunities for understanding the dynamic behavior of proteins; however, computational methods for utilizing this dynamic information remain limited. We propose a novel graph representation that integrates both static structural information and dynamic correlations from molecular dynamics trajectories, enabling more comprehensive modeling of proteins. By applying relational graph neural networks (RGNNs) to process this heterogeneous representation, we demonstrate significant improvements over structure-based approaches across three distinct tasks: atomic adaptability prediction, binding site detection, and binding affinity prediction. Our results validate that combining static and dynamic information provides complementary signals for understanding protein-ligand interactions, offering new possibilities for drug design and structural biology applications.

1 INTRODUCTION

With the recent surge and successes of deep learning methods in protein structure prediction, attention is rapidly turning towards the prediction of the temporal behavior of these highly dynamic macromolecules. Combined with quantitative and qualitative advances in molecular dynamics simulations (Joshi & Deshmukh, 2021; Zeng et al., 2021; Majewski et al., 2023; Nam & Wolf-Watz, 2023), this attention is resulting in the increased availability and accessibility of simulated molecular dynamics trajectories (Vander Meersche et al., 2024; Siebenmorgen et al., 2024a; Liu et al., 2024). Consequently, various approaches are being developed to train predictive and generative models capable of producing molecular dynamics trajectories or sample specific conformations (López-Correa et al., 2023; Jing et al., 2024; Lewis et al., 2024). So far, the potential of these increasingly large trajectory datasets to enhance property predictions in proteins and protein-ligand complexes, such as binding site identification and affinity prediction, remains largely unexplored (Dhakal et al., 2022).

Despite these advances, representing and exploiting molecular dynamics trajectories of proteins for machine learning remains challenging due to the diverse and complex nature of protein structures. One effective alternative is to focus on a higher-order representation of protein dynamics through correlation patterns derived from molecular motion. These dynamic correlations are essential to protein function, and the resulting correlation matrices have long been used to analyze protein dynamics (Agarwal et al., 2002; Lange & Grubmüller, 2008).

In this work, we propose combining molecular structure and simulated molecular trajectories through residue-based correlation matrices and relational graph neural networks (Schlichtkrull et al., 2017). We show that this approach enables the exploitation of the rapidly expanding collection of readily available protein dynamics trajectories for protein and protein–ligand property prediction. In summary:

*Corresponding author: daniel.probst@wur.nl

- We propose a novel heterogeneous graph representation that integrates both static structural information and dynamic correlations from molecular trajectories, enabling more comprehensive modeling of protein properties.
- We introduce the first application of relational graph neural networks to directly process molecular dynamics trajectories, demonstrating clear benefits over graph neural networks (GNNs) based on structure alone.
- We validate our approach across three distinct tasks: atomic adaptability prediction, binding site detection, and binding affinity prediction, showing consistent benefits of combining static and dynamic information.

2 Related Work

2.1 Dynamic Correlations in Protein Analysis

Dynamic correlations derived from molecular dynamics trajectories have been extensively applied in protein analysis, particularly for understanding allosteric mechanisms and signal propagation (McClendon et al., 2009; Long & Brüschweiler, 2011; Wang et al., 2020), investigating tRNA-protein complex interactions (Sethi et al., 2009), and identifying catalytically important regions for enzyme engineering (Bunzel et al., 2021; Gao et al., 2024). However, they have not been used as a representation of trajectories when training predictive models on large data sets but mainly as a means to investigate the propagation of structural changes in a single, or a class of proteins through methods such as dynamical network analysis (Melo et al., 2020; Calvó-Tusell et al., 2022).

2.2 Graph Neural Networks for Protein Structure and Dynamics

Graph neural networks have been widely used to predict properties and functions of proteins as well as properties of protein-ligand or protein-protein interactions based on structure (Gligorijević et al., 2021a; Li et al., 2021b; Réau et al., 2023). More recently, graph neural networks have been applied to enhance and accelerate molecular dynamics simulations (Wang et al., 2022; Yue et al., 2024).

Chiang et al. (2022) explored incorporating dynamic information into protein graphs by using normal mode analysis to generate correlation edges, combining this with 1D and 2D persistence diagrams of α-carbons for molecular function classification using graph convolutional networks (GCN) (Defferrard et al., 2017; Kipf & Welling, 2017). In parallel, relational graph neural networks have shown promise in small-molecule molecular graph generation (Zou et al., 2023), and protein representation learning, integrating sequential and spatial distance in proteins (Zhang et al., 2022).

3 Methodology

3.1 Graph Construction Framework

We represent a protein complex as a tuple (V, E_d, E_c), where V represents the set of nodes, E_d represents distance-based edges, and E_c represents correlation-based edges derived from molecular dynamics trajectories.

We propose a novel approach for incorporating both static structural and dynamic motion information into protein graph representations. Our method consists of two key components: (1) a heterogeneous graph construction framework that combines spatial proximity with dynamic correlations from molecular dynamics simulations, and (2) the application of relational neural networks to effectively process the heterogeneous graphs enriched by both structural and dynamic information.

As illustrated in Figure 1, our approach derives two complementary edge types from protein data: distance edges based on static structure, and correlation edges from molecular dynamics trajectories. These correlation edges provide direct links between dynamically coupled regions of the protein, enabling more efficient information flow than in the original graph structure. The mechanism is akin to graph rewiring, which is known to mitigate over-squashing in GNNs (Attali et al., 2024).

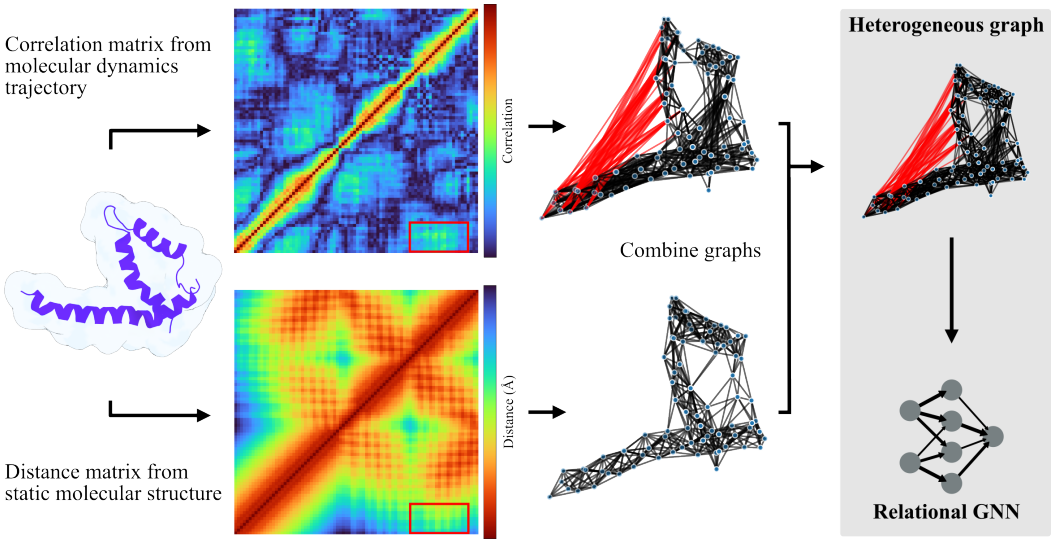

Figure 1: **Fusion of static structure and molecular dynamics information**. The left side shows the transformation of protein structure (PDB ID 5GMU) into a distance-based matrix, while the right side presents the correlation matrix derived from molecular dynamics trajectories, which shows motion correlations between different regions. The correlation edges create direct connections between dynamically coupled regions that may be spatially distant (shown in red), enabling efficient information flow across the protein structure. The fusion of these structural and dynamic features creates a heterogeneous graph representation, which serves as input to relational graph neural networks.

3.1.1 Node Definition and Features

Nodes are defined based on the specific requirements of each task:

For atomic-level predictions (e.g., atomic property prediction), each node represents a non-hydrogen atom, capturing detailed molecular interactions at the atomic scale.

For residue-level tasks (e.g., binding site detection), each node represents a residue, where the coordinates of its C_α atom are used to determine the residue's spatial position.

Each node $v_i \in V$ is associated with a feature vector $\mathbf{h}_i \in \mathbb{R}^d$ consisting of the one-hot encoding of the atom/residue type and the atom charge (for atomic-level graphs).

3.1.2 Distance-based Edge Construction

The distance-based edges E_d capture spatial proximity in the static structure:

$$E_d = \{(v_i, v_j) \mid d(v_i, v_j) < \tau_d\} \tag{1}$$

where $d(v_i, v_j)$ represents the Euclidean distance between nodes, and τ_d is a distance threshold (4.5 Å for atomic-level and 10 Å for residue-level graphs). These thresholds are widely used in protein modeling: the 4.5 Å threshold captures meaningful atomic interactions (Bouysset & Fiorucci, 2021), while the 10 Å threshold is commonly adopted for residue-level contacts (Gligorijević et al., 2021b).

3.1.3 Dynamic Correlation Edge Construction

To capture dynamic behaviors, we analyze classical force-field molecular dynamics trajectories to construct correlation-based edges E_c. Before computing correlations, all trajectory frames are aligned to the initial structure through rigid-body superposition optimized to minimize the root-mean-square deviation (RMSD) between equivalent atomic positions. The alignment eliminates global translations and rotations while preserving internal conformational changes.

Unlike distance-based representations that primarily capture local structural relationships, correlation-based edges can identify dynamically coupled regions regardless of spatial proximity, creating direct pathways between motion-related but spatially distant parts of the protein (as shown in Figure 1). For each pair of nodes, we compute their motion correlation across simulation frames:

$$C_{ij} = \frac{1}{T}\sum_{t=1}^{T}\frac{\Delta\mathbf{r}_i^t \cdot \Delta\mathbf{r}_j^t}{|\Delta\mathbf{r}_i^t||\Delta\mathbf{r}_j^t|} \tag{2}$$

where $\Delta\mathbf{r}_i^t$ represents the displacement vector of node i at frame t, and T is the total number of frames. The correlation edges are then defined as:

$$E_c = \{(v_i, v_j) \mid |C_{ij}| > \tau_c\} \tag{3}$$

where τ_c is the correlation threshold (0.6 for atomic-level and 0.3 for residue-level graphs). These thresholds are chosen to maintain similar graph sparsity, thereby achieving a fairer comparison when either Correlation or Distance Graph is used.

3.1.4 COMBINED GRAPH

The final graph representation integrates both distance-based and correlation-based edges, yielding a heterogeneous graph that captures both static structural information and dynamic behavioral patterns. This combined representation enables the model to utilize local spatial relationships and potential long-range dynamic interactions simultaneously.

3.2 RELATIONAL GRAPH NEURAL NETWORK ARCHITECTURE

The heterogeneous nature of our Combined Graph, containing both distance-based and correlation-based edges, requires a neural network architecture capable of processing different types of relationships. We therefore employ two established relational neural networks: the Relational Graph Convolutional Network (RGCN) (Schlichtkrull et al., 2018) and the Relational Graph Attention Network (RGAT) (Busbridge et al., 2019). These architectures are particularly suited for our approach as they handle heterogeneous edges by learning different weight matrices for different edge types.

The RGCN extends the traditional Graph Convolutional Network by introducing relation-specific transformations. For each layer l, the message passing operation is defined as:

$$\mathbf{h}_i^{(l+1)} = \sigma\Big(\sum_{r\in\mathcal{R}}\sum_{j\in\mathcal{N}_r(i)}\frac{1}{|\mathcal{N}_r(i)|}\mathbf{W}_r^{(l)}\mathbf{h}_j^{(l)} + \mathbf{W}_0^{(l)}\mathbf{h}_i^{(l)}\Big) \tag{4}$$

where $\mathcal{N}_r(i)$ denotes neighbors of node i connected by edges of type r, $\mathbf{W}_r^{(l)}$ is the relation-specific transformation matrix, and $\mathbf{W}_0^{(l)}$ is the self-connection weight matrix. In our case, $\mathcal{R}$ represents the set of edge types (distance and correlation). This formulation allows the network to learn distinct transformations for distance-based and correlation-based relationships, enabling it to capture the unique characteristics of each edge type.

The RGAT extends this formulation by incorporating an attention mechanism. This formulation allows the network to learn distinct transformations for distance-based and correlation-based relationships, enabling it to capture the unique characteristics of each edge type. For each layer l, the attention-based message passing operation is defined as:

$$\mathbf{h}_i^{(l+1)} = \sigma\Big(\sum_{r\in\mathcal{R}}\sum_{j\in\mathcal{N}_r(i)}\alpha_{ij}^{(r)}\mathbf{W}_r^{(l)}\mathbf{h}_j^{(l)} + \mathbf{W}_0^{(l)}\mathbf{h}_i^{(l)}\Big) \tag{5}$$

The attention coefficients $\alpha_{ij}^{(r)}$ are computed using query and key kernels for each relation type r:

$$\mathbf{q}_i^{(r)} = \mathbf{W}_1^{(r)}\mathbf{x}_i \cdot \mathbf{Q}^{(r)} \quad \text{and} \quad \mathbf{k}_i^{(r)} = \mathbf{W}_1^{(r)}\mathbf{x}_i \cdot \mathbf{K}^{(r)} \tag{6}$$

These kernels are used to compute attention logits:

$$\mathbf{a}_{i,j}^{(r)} = \text{LeakyReLU}(\mathbf{q}_i^{(r)} + \mathbf{k}_j^{(r)}) \tag{7}$$

The final attention coefficients are obtained as:

$$\alpha_{i,j}^{(r)} = \frac{\exp(\mathbf{a}_{i,j}^{(r)})}{\sum_{r' \in \mathcal{R}} \sum_{k \in \mathcal{N}_{r'}(i)} \exp(\mathbf{a}_{i,k}^{(r')})} \tag{8}$$

This attention mechanism enables the model to automatically determine the relative importance of different relationships, potentially providing insights into the contributions of structural and dynamic information in protein modeling.

4 EXPERIMENTS

4.1 DATASET

We evaluate our approach using the MISATO dataset (Siebenmorgen et al., 2024b), which contains 19,443 protein-ligand complexes derived from PDBbind (Su et al., 2018; Liu et al., 2017; Wang et al., 2005). Each complex undergoes quantum mechanical relaxation and 10 ns molecular dynamics simulation using the Amber20 software package with gaff2 force field for ligands and ff14SB for proteins. The dataset also provides key physicochemical properties, forming a high-quality benchmark for machine learning tasks.

4.2 RESULTS AND DISCUSSION

We evaluate our approach on three distinct prediction tasks: atomic adaptability, binding site identification, and binding affinity prediction. For each task, we analyze how different graph representations (Distance, Correlation, and Combined) affect model performance using RGCN and RGAT architectures.

4.2.1 ATOMIC ADAPTABILITY PREDICTION

Atomic adaptability quantifies the conformational plasticity of atoms within a protein structure, where higher values indicate greater flexibility and lower values indicate rigidity (see Appendix A.5). This property helps identify key regions of motion, making it crucial for understanding protein dynamics and molecular design. We formulate adaptability prediction as a node-level regression task, where each atom is annotated with an adaptability score from the MISATO dataset (see Figure 2).

Table 1 presents the performance comparison across different graph representations using two baseline architectures (RGCN and RGAT). The Correlation Graph, which captures dynamic motion patterns derived from MD trajectories, consistently outperforms the Distance Graph across all metrics for both architectures. Using RGCN, we observe improvements in both error metrics (MAE reduces from 0.2658 to 0.2311) and correlation coefficients (Pearson R increases from 0.5259 to 0.6426). Similar comprehensive improvements are observed with RGAT, demonstrating the value of dynamic information for atomic adaptability prediction.

When both types of information are integrated in the Combined Graph, we observe further significant improvements across all metrics: the Pearson correlation coefficient reaches 0.7326 (RGCN) and 0.7153 (RGAT), representing improvements of 39.3% and 50.7% respectively over the Distance Graph baseline (0.5259 using RGCN and 0.4746 using RGAT). Similar improvements are observed across other metrics, as evidenced by the reduction in MAE from 0.2658 to 0.1981 (RGCN) and 0.2766 to 0.2074 (RGAT).

These performance improvements align with the physical nature of atomic adaptability. While spatial proximity (captured by the Distance Graph) provides important structural constraints, atomic

Table 1: **Atomic adaptability prediction**. Node-level regression task predicting atomic adaptability values using data from the MISATO dataset. Results show mean ± standard deviation over 5 runs (↓ indicates lower is better, ↑ indicates higher is better). The Combined Graph, which integrates both distance and correlation information, consistently outperforms single-information approaches across all metrics using both RGCN and RGAT. Notably, the Correlation Graph alone achieves better performance than the Distance Graph on every evaluation metric, suggesting the importance of dynamic information for this task.

Model	Graph Type	MAE (↓)	RMSE (↓)	Pearson R (↑)	Spearman R (↑)
RGCN	Distance	0.2658 ± 0.0061	0.4274 ± 0.0008	0.5259 ± 0.0015	0.5543 ± 0.0017
	Correlation	0.2311 ± 0.0014	0.3846 ± 0.0011	0.6426 ± 0.0026	0.6990 ± 0.0019
	Combined	**0.1981** ± 0.0020	**0.3417** ± 0.0008	**0.7326** ± 0.0014	**0.7922** ± 0.0010
RGAT	Distance	0.2766 ± 0.0038	0.4419 ± 0.0018	0.4746 ± 0.0066	0.4762 ± 0.0085
	Correlation	0.2443 ± 0.0013	0.3976 ± 0.0013	0.6106 ± 0.0037	0.6521 ± 0.0068
	Combined	**0.2074** ± 0.0030	**0.3511** ± 0.0018	**0.7153** ± 0.0033	**0.7699** ± 0.0024

Table 2: **Binding site detection**. Node-level binary classification task identifying binding site residues (those within 10 Å from the ligand) using data from the MISATO dataset. Results show mean ± standard deviation over 5 runs (↑ indicates higher is better). The Combined Graph demonstrates superior performance across all metrics for both architectures. The Correlation Graph's improved performance over the Distance Graph baseline on every metric highlights the value of molecular dynamics information in binding site identification.

Model	Graph Type	Acc (↑)	Precision (↑)	Recall (↑)	F1 score (↑)
RGCN	Distance	0.7112 ± 0.0092	0.1678 ± 0.0024	0.4464 ± 0.0164	0.2428 ± 0.0027
	Correlation	0.7282 ± 0.0069	0.1808 ± 0.0022	0.4552 ± 0.0102	0.2578 ± 0.0012
	Combined	**0.7433** ± 0.0067	**0.2005** ± 0.0030	**0.4889** ± 0.0111	**0.2834** ± 0.0023
RGAT	Distance	0.6602 ± 0.0120	0.1475 ± 0.0032	0.4439 ± 0.0234	0.2089 ± 0.0040
	Correlation	0.6938 ± 0.0111	0.1664 ± 0.0031	0.4441 ± 0.0182	0.2294 ± 0.0030
	Combined	**0.7226** ± 0.0067	**0.1861** ± 0.0029	**0.4750** ± 0.0137	**0.2574** ± 0.0032

adaptability is inherently a dynamical property that highly depends on atomic fluctuations and conformational changes, which cannot be fully captured by spatial proximity alone. The Correlation Graph, leveraging dynamical information derived from MD trajectories, better captures elements tied to motion and complements the structural information. When combined, these two edge types enable the model to learn from both spatial constraints and dynamic coupling patterns, resulting in the Combined Graph's superior performance.

The consistent improvement across both RGCN and RGAT architectures suggests that the performance gains primarily stem from the richer graph structure rather than specific architectural choices. This robustness validates our approach of incorporating dynamic information through correlation edges as an effective strategy for enhancing protein graph representations in dynamical property prediction.

4.2.2 Binding Site Detection

Binding site detection aims to identify key residues in proteins that directly interact with ligands, specifically those residues within 10 Å from the ligand, following PDBbind. This task is essential for understanding protein functionality and facilitating early-stage drug design. We formulate this as a binary node classification problem at the residue level, where each node represents a residue and is classified as either a binding site or a non-binding site (see Figure 3). For this task, the ligand coordinates are removed from the complex and not provided to the model.

Table 2 presents the classification performance across different graph representations using RGCN and RGAT architectures. The incorporation of dynamic information through the Correlation Graph proves beneficial, showing better performance than the Distance Graph on all evaluation metrics for both architectures. Using RGCN, we observe improvements in both precision (from 0.1678 to 0.1808) and F1 score (from 0.2428 to 0.2578). Similar improvements are seen with RGAT, where precision increases from 0.1475 to 0.1664 and F1 score from 0.2089 to 0.2294, demonstrating the value of dynamic information for binding site detection.

The fusion of static and dynamic information in the Combined Graph leads to stronger results. The Combined Graph achieves the best performance across all metrics, demonstrating the complementary advantages of integrating these two types of information. For RGCN, the F1 score increases from 0.2428 (Distance Graph) and 0.2578 (Correlation Graph) to 0.2834, representing improvements of 16.7% and 9.9% respectively. Similar patterns emerge with RGAT, where the F1 score improves from 0.2089 (Distance Graph) and 0.2294 (Correlation Graph) to 0.2574.

Compared to atomic adaptability, which relies more directly on dynamic information, binding site identification depends heavily on static structural features such as protein surfaces and binding pockets. While our results still show that the Correlation Graph consistently outperforms the Distance Graph across all metrics, the magnitude of improvement is more moderate compared to the adaptability task, reflecting the balanced importance of both static and dynamic features in this context. When combined, the model can utilize both spatial proximity and motion patterns, leading to more accurate binding site identification.

The consistent improvement across both RGCN and RGAT architectures demonstrates that these improvements result from the complementary nature of static and dynamic features rather than specific architectural choices. These results validate our approach of incorporating both types of information into protein graph representations, offering new possibilities for studying complex protein-ligand interactions.

4.2.3 Binding Affinity Prediction

Binding affinity prediction represents a critical task in drug design and virtual screening, as it quantifies the interaction strength between proteins and ligands. We formulate this as a graph-level regression task, where each graph represents a protein pocket and its ligand, with experimentally measured binding affinities as targets. Following previous work (Li et al., 2021a), we evaluate our approach on the PDBbind 2020 benchmark (details in Appendix A.4).

Table 3: **Binding Affinity Prediction**. Graph-level regression task predicting protein-ligand binding affinity values using MISATO and PDBbind datasets. Results show mean ± standard deviation over 5 runs (↓ indicates lower is better, ↑ indicates higher is better). While the Correlation Graph performs comparably to the Distance Graph baseline, their integration in the Combined Graph yields consistent improvements across all metrics for both architectures, demonstrating the complementary value of static and dynamic information.

Model	Graph Type	MAE (↓)	RMSE (↓)	Pearson R (↑)	Spearman R (↑)
RGCN	Distance	1.3046 ± 0.0267	1.6653 ± 0.0336	0.6596 ± 0.0156	0.6352 ± 0.0234
	Correlation	1.3572 ± 0.0792	1.6974 ± 0.0827	0.6360 ± 0.0428	0.6185 ± 0.0440
	Combined	**1.2439** ± 0.0256	**1.5798** ± 0.0447	**0.6983** ± 0.0193	**0.6773** ± 0.0208
RGAT	Distance	1.3028 ± 0.0261	1.6427 ± 0.0459	0.6694 ± 0.0222	0.6417 ± 0.0225
	Correlation	1.3249 ± 0.0341	1.6623 ± 0.0335	0.6643 ± 0.0212	0.6481 ± 0.0254
	Combined	**1.2596** ± 0.0290	**1.6012** ± 0.0411	**0.6931** ± 0.0157	**0.6752** ± 0.0157

Table 3 presents the regression performance across different graph representations using RGCN and RGAT architectures. For this task, the Correlation Graph shows performance comparable to the Distance Graph, with differences varying across architectures and metrics. This suggests that for binding affinity prediction, dynamic information alone may not provide additional advantages over static structural features.

However, the Combined Graph achieves consistent improvements across all metrics, demonstrating the value of fusing both types of information. Using RGCN, the Combined Graph reaches a Pearson correlation of 0.6983 and reduces MAE to 1.2439, improving upon both the Distance Graph (0.6596, 1.3046) and Correlation Graph (0.6360, 1.3572). Similar patterns emerge with RGAT, where the Combined Graph achieves a Pearson correlation of 0.6931 and MAE of 1.2596, outperforming both single-information approaches (Distance Graph:0.6694, 1.3028; Correlation Graph: 0.6596, 1.3046).

These results reflect the complex nature of protein-ligand binding affinity, which requires both structural and dynamic information for accurate prediction. While static distance information captures essential geometric constraints, it cannot reflect potential conformational adjustments and long-range interactions during binding. Similarly, dynamic correlations alone, though capturing important motion patterns, cannot fully characterize the binding pocket geometry, leaving room for additional improvements. The integration of both information types enables the model to simultaneously consider geometric constraints and dynamic interaction patterns, achieving better performance across both error metrics and correlation coefficients and demonstrating the value of this combined approach.

The consistent improvement pattern across both RGCN and RGAT architectures validates that these performance gains arise from the complementary nature of static and dynamic features rather than specific architectural choices. These results show that the fusion of both static and dynamic information enhances the accuracy of binding affinity prediction, offering valuable insights for drug design and molecular screening applications.

5 Conclusion

This work enhances protein graph representations by overcoming a key limitation in current approaches: their exclusive reliance on static structural information without incorporating crucial information about protein dynamics. We propose a novel heterogeneous graph representation that integrates static structural information and dynamic correlations from molecular simulation trajectories, and apply relational graph neural networks to process these enriched representations. Our systematic evaluation examines three distinct tasks: atomic adaptability prediction, binding site detection, and binding affinity prediction. The experimental results show task-dependent patterns: the Correlation Graph demonstrates clear advantages over the Distance Graph for atomic adaptability prediction and binding site detection, while showing comparable performance for binding affinity prediction. Notably, the Combined Graph consistently achieves superior performance across all tasks and metrics, maintaining this advantage across different graph neural network architectures. These results demonstrate that static and dynamic information provide complementary signals for understanding protein behavior. Our approach opens new possibilities for protein modeling and design by effectively capturing both static structural constraints and dynamic correlations in a unified framework.

Future directions include exploring advanced architectures like graph transformers to enhance heterogeneous information processing, and investigating additional correlation measures such as Mutual Information to enrich dynamic feature representation. As a broader direction, integration with emerging generative models for molecular dynamics could further expand the applicability of our approach by trajectory generation, especially when molecular dynamics trajectories are not readily available. These developments will further strengthen our approach's capability in protein modeling, advancing applications in drug design and structural biology.

Meaningfulness Statement

A meaningful representation of life should capture both structural patterns and the dynamic behaviors crucial to biological function. Our work advances this goal by proposing a heterogeneous graph representation that fuses structural information from static conformations and dynamic correlations from molecular dynamics trajectories. We provide a framework for creating these heterogeneous protein representations and processing them with relational graph neural networks, and validate the effectiveness of our framework through enhanced performance across three distinct tasks.

REFERENCES

Pratul K. Agarwal, Salomon R. Billeter, P. T. Ravi Rajagopalan, Stephen J. Benkovic, and Sharon Hammes-Schiffer. Network of coupled promoting motions in enzyme catalysis. *Proceedings of the National Academy of Sciences*, 99(5):2794–2799, March 2002. doi: 10.1073/pnas.052005999.

Hugo Attali, Davide Buscaldi, and Nathalie Pernelle. Rewiring Techniques to Mitigate Oversquashing and Oversmoothing in GNNs: A Survey, November 2024.

Cédric Bouysset and Sébastien Fiorucci. Prolif: a library to encode molecular interactions as fingerprints. *Journal of cheminformatics*, 13(1):72, 2021.

H. Adrian Bunzel, J. L. Ross Anderson, Donald Hilvert, Vickery L. Arcus, Marc W. van der Kamp, and Adrian J. Mulholland. Evolution of dynamical networks enhances catalysis in a designer enzyme. *Nature Chemistry*, 13(10):1017–1022, October 2021. ISSN 1755-4349. doi: 10.1038/s41557-021-00763-6.

Dan Busbridge, Dane Sherburn, Pietro Cavallo, and Nils Y Hammerla. Relational graph attention networks. *arXiv preprint arXiv:1904.05811*, 2019.

Carla Calvó-Tusell, Miguel A. Maria-Solano, Sílvia Osuna, and Ferran Feixas. Time Evolution of the Millisecond Allosteric Activation of Imidazole Glycerol Phosphate Synthase. *Journal of the American Chemical Society*, 144(16):7146–7159, April 2022. ISSN 0002-7863. doi: 10.1021/jacs.1c12629.

Yuan Chiang, Wei-Han Hui, and Shu-Wei Chang. Encoding protein dynamic information in graph representation for functional residue identification. *Cell Reports Physical Science*, 3(7):100975, July 2022. ISSN 2666-3864. doi: 10.1016/j.xcrp.2022.100975.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering, 2017. URL https://arxiv.org/abs/1606.09375.

Ashwin Dhakal, Cole McKay, John J Tanner, and Jianlin Cheng. Artificial intelligence in the prediction of protein–ligand interactions: Recent advances and future directions. *Briefings in Bioinformatics*, 23(1):bbab476, January 2022. ISSN 1477-4054. doi: 10.1093/bib/bbab476.

Chun-Yu Gao, Gui-Ying Yang, Xu-Wei Ding, Jian-He Xu, Xiaolin Cheng, Gao-Wei Zheng, and Qi Chen. Engineering of Halide Methyltransferase BxHMT through Dynamic Cross-Correlation Network Analysis. *Angewandte Chemie International Edition*, 63(25):e202401235, 2024. ISSN 1521-3773. doi: 10.1002/anie.202401235.

Vladimir Gligorijević, P. Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C. Taylor, Ian M. Fisk, Hera Vlamakis, Ramnik J. Xavier, Rob Knight, Kyunghyun Cho, and Richard Bonneau. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12(1):3168, May 2021a. ISSN 2041-1723. doi: 10.1038/s41467-021-23303-9.

Vladimir Gligorijević, P Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C Taylor, Ian M Fisk, Hera Vlamakis, et al. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 12(1):3168, 2021b.

Bowen Jing, Hannes Stärk, Tommi Jaakkola, and Bonnie Berger. Generative Modeling of Molecular Dynamics Trajectories, September 2024.

Soumil Y. Joshi and Sanket A. Deshmukh. A review of advancements in coarse-grained molecular dynamics simulations. *Molecular Simulation*, 47(10-11):786–803, July 2021. ISSN 0892-7022. doi: 10.1080/08927022.2020.1828583.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. URL https://arxiv.org/abs/1609.02907.

Oliver F. Lange and Helmut Grubmüller. Full correlation analysis of conformational protein dynamics. *Proteins: Structure, Function, and Bioinformatics*, 70(4):1294–1312, 2008. ISSN 1097-0134. doi: 10.1002/prot.21618.

Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew Y. K. Foong, Victor García Satorras, Osama Abdin, Bastiaan S. Veeling, Iryna Zaporozhets, Yaoyi Chen, Soojung Yang, Arne Schneuing, Jigyasa Nigam, Federico Barbero, Vincent Stimper, Andrew Campbell, Jason Yim, Marten Lienen, Yu Shi, Shuxin Zheng, Hannes Schulz, Usman Munir, Cecilia Clementi, and Frank Noé. Scalable emulation of protein equilibrium ensembles with generative deep learning, December 2024.

Shuangli Li, Jingbo Zhou, Tong Xu, Liang Huang, Fan Wang, Haoyi Xiong, Weili Huang, Dejing Dou, and Hui Xiong. Structure-aware interactive graph neural networks for the prediction of protein-ligand binding affinity, 2021a. URL https://arxiv.org/abs/2107.10670.

Shuangli Li, Jingbo Zhou, Tong Xu, Liang Huang, Fan Wang, Haoyi Xiong, Weili Huang, Dejing Dou, and Hui Xiong. Structure-aware Interactive Graph Neural Networks for the Prediction of Protein-Ligand Binding Affinity. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD '21, pp. 975–985, New York, NY, USA, August 2021b. Association for Computing Machinery. doi: 10.1145/3447548.3467311.

Ce Liu, Jun Wang, Zhiqiang Cai, Yingxu Wang, Huizhen Kuang, Kaihui Cheng, Liwei Zhang, Qingkun Su, Yining Tang, Fenglei Cao, Limei Han, Siyu Zhu, and Yuan Qi. Dynamic PDB: A New Dataset and a SE(3) Model Extension by Integrating Dynamic Behaviors and Physical Properties in Protein Structures, September 2024.

Zhihai Liu, Minyi Su, Li Han, Jie Liu, Qifan Yang, Yan Li, and Renxiao Wang. Forging the basis for developing protein–ligand interaction scoring functions. *Accounts of chemical research*, 50 (2):302–309, 2017.

Dong Long and Rafael Brüschweiler. Atomistic Kinetic Model for Population Shift and Allostery in Biomolecules. *Journal of the American Chemical Society*, 133(46):18999–19005, November 2011. ISSN 0002-7863. doi: 10.1021/ja208813t.

Juan Manuel López-Correa, Caroline König, and Alfredo Vellido. GPCR molecular dynamics forecasting using recurrent neural networks. *Scientific Reports*, 13(1):20995, November 2023. ISSN 2045-2322. doi: 10.1038/s41598-023-48346-4.

Maciej Majewski, Adrià Pérez, Philipp Thölke, Stefan Doerr, Nicholas E. Charron, Toni Giorgino, Brooke E. Husic, Cecilia Clementi, Frank Noé, and Gianni De Fabritiis. Machine learning coarse-grained potentials of protein thermodynamics. *Nature Communications*, 14(1):5739, September 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-41343-1.

Christopher L. McClendon, Gregory Friedland, David L. Mobley, Homeira Amirkhani, and Matthew P. Jacobson. Quantifying Correlations Between Allosteric Sites in Thermodynamic Ensembles. *Journal of Chemical Theory and Computation*, 5(9):2486–2502, September 2009. ISSN 1549-9618. doi: 10.1021/ct9001812.

Marcelo C. R. Melo, Rafael C. Bernardi, Cesar de la Fuente-Nunez, and Zaida Luthey-Schulten. Generalized correlation-based dynamical network analysis: A new high-performance approach for identifying allosteric communications in molecular dynamics trajectories. *The Journal of Chemical Physics*, 153(13):134104, October 2020. ISSN 0021-9606. doi: 10.1063/5.0018980.

Kwangho Nam and Magnus Wolf-Watz. Protein dynamics: The future is bright and complicated! *Structural Dynamics*, 10(1):014301, February 2023. ISSN 2329-7778. doi: 10.1063/4.0000179.

Manon Réau, Nicolas Renaud, Li C Xue, and Alexandre M J J Bonvin. DeepRank-GNN: A graph neural network framework to learn patterns in protein–protein interfaces. *Bioinformatics*, 39(1): btac759, January 2023. ISSN 1367-4811. doi: 10.1093/bioinformatics/btac759.

Daniel R Roe and Thomas E Cheatham III. Ptraj and cpptraj: software for processing and analysis of molecular dynamics trajectory data. *Journal of chemical theory and computation*, 9(7):3084–3095, 2013.

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks, October 2017.

Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15*, pp. 593–607. Springer, 2018.

Anurag Sethi, John Eargle, Alexis A. Black, and Zaida Luthey-Schulten. Dynamical networks in tRNA:protein complexes. *Proceedings of the National Academy of Sciences*, 106(16):6620–6625, April 2009. doi: 10.1073/pnas.0810961106.

Till Siebenmorgen, Filipe Menezes, Sabrina Benassou, Erinc Merdivan, Kieran Didi, André Santos Dias Mourão, Radosław Kitel, Pietro Liò, Stefan Kesselheim, Marie Piraud, Fabian J. Theis, Michael Sattler, and Grzegorz M. Popowicz. MISATO: Machine learning dataset of protein–ligand complexes for structure-based drug discovery. *Nature Computational Science*, 4(5):367–378, May 2024a. ISSN 2662-8457. doi: 10.1038/s43588-024-00627-2.

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

Minyi Su, Qifan Yang, Yu Du, Guoqin Feng, Zhihai Liu, Yan Li, and Renxiao Wang. Comparative assessment of scoring functions: the casf-2016 update. *Journal of chemical information and modeling*, 59(2):895–913, 2018.

Yann Vander Meersche, Gabriel Cretin, Aria Gheeraert, Jean-Christophe Gelly, and Tatiana Galochkina. ATLAS: Protein flexibility description from atomistic molecular dynamics simulations. *Nucleic Acids Research*, 52(D1):D384–D392, January 2024. ISSN 0305-1048. doi: 10.1093/nar/gkad1084.

Jian Wang, Abha Jain, Leanna R. McDonald, Craig Gambogi, Andrew L. Lee, and Nikolay V. Dokholyan. Mapping allosteric communications within individual proteins. *Nature Communications*, 11(1):3862, July 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-17618-2.

Renxiao Wang, Xueliang Fang, Yipin Lu, Chao-Yie Yang, and Shaomeng Wang. The pdbbind database: methodologies and updates. *Journal of medicinal chemistry*, 48(12):4111–4119, 2005.

Zun Wang, Chong Wang, Sibo Zhao, Yong Xu, Shaogang Hao, Chang Yu Hsieh, Bing-Lin Gu, and Wenhui Duan. Heterogeneous relational message passing networks for molecular dynamics simulations. *npj Computational Materials*, 8(1):53, 2022.

Yang Yue, Shu Li, Yihua Cheng, Lie Wang, Tingjun Hou, Zexuan Zhu, and Shan He. Integration of molecular coarse-grained model into geometric representation learning framework for protein-protein complex property prediction. *Nature Communications*, 15(1):9629, November 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-53583-w.

Jinzhe Zeng, Timothy J. Giese, Şölen Ekesan, and Darrin M. York. Development of Range-Corrected Deep Learning Potentials for Fast, Accurate Quantum Mechanical/Molecular Mechanical Simulations of Chemical Reactions in Solution. *Journal of Chemical Theory and Computation*, 17(11):6993–7009, November 2021. ISSN 1549-9618. doi: 10.1021/acs.jctc.1c00201.

Zuobai Zhang, Minghao Xu, Arian Rokkum Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein Representation Learning by Geometric Structure Pretraining. In *The Eleventh International Conference on Learning Representations*, September 2022.

Xiandong Zou, Xiangyu Zhao, Pietro Lio, and Yiren Zhao. Will More Expressive Graph Neural Networks do Better on Generative Tasks? In *The Second Learning on Graphs Conference*, November 2023.

A Appendix

A.1 Additional Visualizations

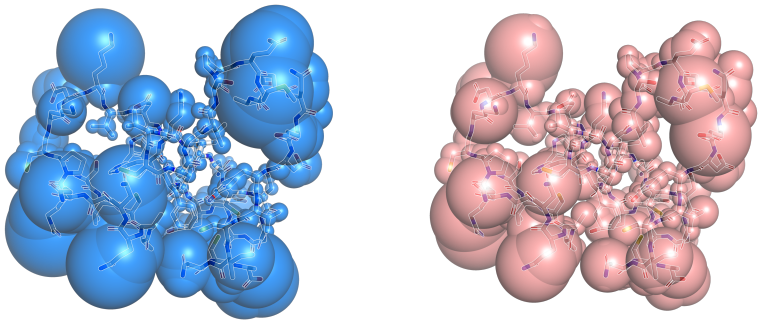

Figure 2: **Atomic Adaptability Prediction.** Visualization of per-atom adaptability in a protein structure (PDB ID `5C11`). Left: ground-truth (target) adaptability values shown as blue spheres. Right: predicted adaptability values shown as pink spheres. Sphere size indicates the magnitude of adaptability, with larger spheres corresponding to more flexible (higher adaptability) regions.

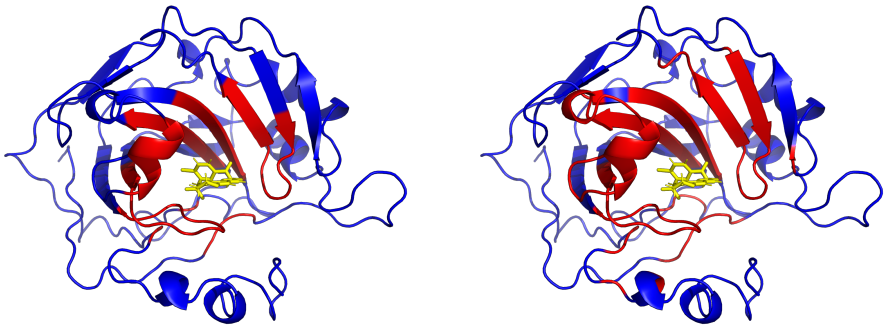

Figure 3: **Binding Site Prediction**. Visualization of binding site detection task using protein structure (PDB ID `3M67`). Left: Ground truth binding site residues (shown in red) and the bound ligand (yellow). Right: Predicted binding site residues using the Combined Graph, showing reasonable agreement with the true binding regions for this example. This case illustrates how the model aims to identify residues within 10 Åfrom the ligand. The ligand is shown only for reference and is not provided to the model.

A.2 Trajectory Alignment

During preprocessing, we aligned all MD trajectories to their initial frames using PyTraj's `align` function (Roe & Cheatham III, 2013). The alignment eliminates global translations and rotations, ensuring that $\Delta \mathbf{r}_i^t$ captures meaningful conformational changes rather than rigid-body motions. By focusing on intrinsic protein dynamics, this preprocessing step improves the quality of our correlation-based edges and leads to more informative graph representations.

A.3 Implementation Details and Hyperparameters

We implemented our models using PyTorch Geometric. Each model consists of 5 GNN layers followed by a two-layer MLP for prediction. We trained models using the Adam optimizer with a learning rate of 1e-4 and batch size of 32. Training epochs were task-specific: 50 for atomic adaptability prediction, 200 for binding site detection, and 500 for binding affinity prediction.

For model architecture optimization, we explored different hidden dimensions for each model-task-graph type combination. The dimension ranges were selected based on architectural differences and memory constraints. For example, in atomic adaptability prediction using RGCN, we explored hidden dimensions $\{26, 32, 53, 64\}$, while for RGAT we tested $\{17, 20, 24\}$ due to its higher memory requirements.

A.4 Dataset Details

For atomic adaptability and binding site detection tasks, we used the data splitting in the MISATO dataset splits, with 13,597 samples for training, 1,582 for validation, and 1,593 for test. These splits were created using sequence-based clustering with BlastP (similarity threshold of 30%) to prevent information leakage through structural similarities. The dataset contains molecular dynamics trajectories generated using the Amber20 software package with a simulation length of 10ns.

For binding affinity prediction, following previous work (Li et al., 2021a), we used the PDBbind 2020 refined set for training and validation, and evaluated on the core set. The refined set consists of 5,316 protein-ligand complexes specifically selected for high-quality binding data and crystal structures through a comprehensive filtering process (Liu et al., 2017). This dataset construction ensures reliable binding affinity values derived from carefully curated experimental measurements.

A.5 Atomic adaptability

Atomic adaptability (γ_x) for each atom x is calculated as the mean distance from its initial position across all simulation frames after alignment:

$$\gamma_x = \frac{1}{N_{\text{frames}}} \sum_{i}^{N_{\text{frames}}} \|\mathbf{r}_{\text{ref},x} - \mathbf{r}_{i,x}\| \tag{9}$$

where $\mathbf{r}_{\text{ref},x}$ is the initial position of atom x and $\mathbf{r}_{i,x}$ is its position in frame i. This measure quantifies each atom's mobility throughout the simulation, providing insight into conformational flexibility at atomic resolution.

While atomic adaptability can be directly computed from molecular dynamics trajectories, predicting it from graph representations provides a valuable benchmark for evaluating how effectively different graph structures capture dynamic information. The Correlation Graph represents a compressed encoding of the full trajectory information, so the ability to accurately predict adaptability demonstrates that this encoding successfully preserves essential dynamic features. This task also maintains continuity with the established MISATO benchmark, facilitating direct comparison with current and future approaches.