Full-Atom Protein-Protein Interaction Prediction via Atomic Equivariant Attention Network

Chunchen Wang
wangchunchen@bupt.edu.cn
Beijing University of Posts and Telecommunications
Beijing, China
Mohamed Bin Zayed Artificial Intelligence University
Abu Dhabi, UAE

Cheng Yang
yangcheng@bupt.edu.cn
Beijing University of Posts and Telecommunications
Beijing, China

Wenchuan Yang
wenchuanyang97@163.com
National University of Defense Technology
Changsha, China
Beijing University of Posts and Telecommunications
Beijing, China

Le Song
le.song@mbzuai.ac.ae
Mohamed Bin Zayed Artificial Intelligence University
Abu Dhabi, UAE
GenBio AI
CA, USA

Chuan Shi*
shichuan@bupt.edu.cn
Beijing University of Posts and Telecommunications
Beijing, China

Abstract

Protein-protein Interaction (PPI) prediction, which aims to identify the interactions between proteins within a biological system, is an important problem in understanding disease mechanisms and drug discovery. Recently, Equivariant Graph Neural Networks (E3-GNNs) are advanced computational models that provide a powerful solution for accurately predicting PPIs by preserving the geometric integrity of protein interactions. However, most E3-GNNs model protein interactions at the residue level, potentially neglecting critical atomic details and side-chain conformations. In this paper, we propose a novel model, MEANT, designed to adaptively extract atom-level geometric information from varying numbers of atoms within different residues for PPI prediction. Specifically, we define a full-atom graph that contains atomic geometry and guides the message passing under the structure of residues. We also design a geometric relation extractor to integrate geometric information from different residues and adaptively handle variations in the number of atoms within each residue. Finally, we adopt the attention mechanism to update the residue representation and the atomic coordinates within a residue. Experimental results show that our proposed model, MEANT, significantly outperforms state-of-the-art methods on three typical PPI prediction tasks. Our code and data are available on GitHub at https://github.com/BUPT-GAMMA/MEANT.

*Corresponding Author

CIKM '25, Seoul, Republic of Korea

ACM ISBN 979-8-4007-2040-6/2025/11
https://doi.org/10.1145/3746252.3761352

CCS Concepts

• **Computing methodologies** → **Machine learning**.

Keywords

Graph Neural Network, Protein-Protein Interaction, Geometric Learning

ACM Reference Format:
Chunchen Wang, Cheng Yang, Wenchuan Yang, Le Song, and Chuan Shi. 2025. Full-Atom Protein-Protein Interaction Prediction via Atomic Equivariant Attention Network. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25), November 10–14, 2025, Seoul, Republic of Korea.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3746252.3761352

1 Introduction

Protein-protein interactions (PPIs), which are physical contacts with specific binding affinities between protein molecules, play a pivotal role in nearly all biological processes [18]. PPI prediction refers to the computational process of identifying potential interactions between protein pairs, including binary interactions and changes in binding free energy [16, 20]. Accurate prediction of PPIs is important for improving our understanding of disease mechanisms and supporting the development of new drugs [10, 40]. However, traditional experimental methods, such as those based on mass spectrometry, are often time-consuming, expensive, and limited in their ability to process large numbers of samples, which underscores the importance of reliable computational solutions.

The development of deep learning has greatly advanced PPI research by offering more effective modeling and prediction capabilities than traditional approaches [11, 33]. Sequence-based models rely on amino acid sequences to predict interactions, using the sequential order of residues to capture meaningful patterns [36, 41]. Graph Neural Networks (GNNs) represent proteins as nodes within

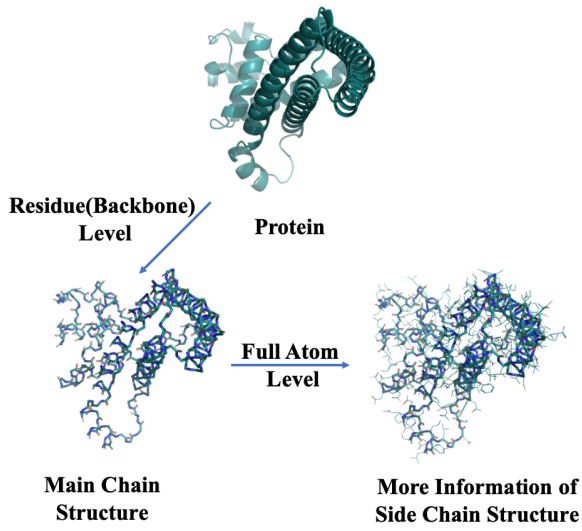

Figure 1: The illustration of the difference between residue level and full-atom level.

a graph, focusing on the relationships between them to infer potential interactions [3, 4, 20, 39]. A recent breakthrough in this area is the introduction of Equivariant Graph Neural Networks (E3-GNNs), which improve upon standard GNNs by maintaining invariance to node order and sensitivity to geometric symmetries present in biological data [17, 30]. These properties are particularly important for accurately modeling the three-dimensional structures of proteins and predicting how they interact.

Despite the promising capabilities of E3-GNNs, existing methods that operate at the residue level [17, 30] often fail to capture detailed atomic-level information. As illustrated in Fig. 1, residue-level modeling focuses primarily on the protein backbone structure, effectively representing the overall fold but neglecting the fine-grained structural features of side chains [22, 37]. These side chain details play a crucial role in protein-protein interactions, influencing binding specificity and interaction strength. In contrast, atom-level modeling provides a more complete view of the protein's three-dimensional structure, including both main chain and side chain configurations. To the best of our knowledge, there have been no prior studies that explore PPI prediction from a full-atom perspective, highlighting a significant gap in current research.

In practice, designing an atomic equivariant graph neural network is non-trivial, we need to carefully consider how to apply EGNNs to full-atom protein graphs. This requires us to address the following two fundamental problems:

(1) *Efficiency Challenges in Protein Graph Modeling: From Residue to Atomic Level.* This shift dramatically increases the complexity and size of protein graphs, as each amino acid residue is decomposed into its constituent atoms, resulting in a much denser and more detailed graph representation. HSRN [15] makes an initial attempt to incorporate all side-chain atoms into a hierarchical graph; however, it suffers from efficiency limitations and inevitably omits certain structural details of the protein.

(2) *How to Extract Full-atom Geometric Information between Residues.* Due to the varying number of atoms in different residues, representing the geometric information between residues using inter-atomic distances results in distance matrices of different dimensions. Consequently, integrating atomic-level binding information between different residues presents a significant challenge that necessitates careful consideration.

In this paper, we propose a novel model named Ato**M**ic **E**quivariant **A**ttention **N**e**T**work (MEANT), an E(3)-equivariant graph neural network designed to adaptively extract atom-level geometric information from varying numbers of atoms within different residues for PPI prediction. To address the first challenge, we define a full-atom protein graph that allows us to extract atomic-level geometric information between different residues while maintaining residue-level network structure for message passing. To address the second challenge, we propose a geometric relationship extractor to capture the atomic geometry between residues. This extractor is capable of integrating geometric information across different residues and adaptively manages variations in the number of atoms within each residue. Finally, we aggregate neighborhood information using an attention mechanism to update the residue-level representations and the coordinates of atoms within the residues. We evaluate our model, (MEANT), on three representative PPI tasks. The experimental outcomes indicate that our proposed MEANT model significantly surpasses baseline methods, demonstrating its effectiveness in navigating the intricacies of PPI prediction with enhanced accuracy and efficiency.

To conclude, our contributions are as follows:

• For the first time, we first try to design a solution for PPI prediction from a full-atom perspective and tested it on five datasets across three main PPI prediction tasks.

• Our proposed model, MEANT, adaptively accommodates variations in the number of atoms across different residues while considering the full atomic geometric structure, thereby enabling effective message passing along the 3D structure.

• Experimental results on three PPI prediction tasks show that MEANT achieves a relative improvement of up to 3.10% in terms of ACC in binary classification and significantly outperforms baseline methods in the other two tasks.

2 Related Work

2.1 GNN-based PPI Method

Graph neural networks (GNNs) have played a crucial role in advancing protein-protein interaction (PPI) prediction by effectively capturing both sequence-derived and structural features within protein graphs. S-VGAE [39] combines sequence information with graph structure using a variational graph autoencoder, achieving superior performance compared to purely sequence-based methods. Struct2Graph [3] introduces a graph attention network that incorporates three-dimensional structural features of proteins to improve the accuracy of PPI prediction. GRABP [4] adopts a graph signal processing (GSP) approach, modeling PPI networks through a Markovian framework to analyze node connectivity and signal flow. Kabir et al. [20] propose a deep learning framework for semi-supervised learning on multi-relational graphs, which adaptively learns the contribution of different types of relations to the prediction outcomes. Furthermore, Jha et al. [14] integrate molecular

protein GNNs with language models by using per-residue embeddings from sequence data as node features, enabling a more comprehensive representation of proteins for interaction prediction. Collectively, these studies highlight the versatility and effectiveness of GNN-based methods in capturing complex relationships underlying PPIs. SemiGNN-PPI [42] introduces a self-ensembling multi-graph neural network that leverages both feature and label graph structures with consistency-based semi-supervised learning, achieving strong generalization under label-scarce and domain-shifted PPI scenarios.

2.2 Equivariant Graph Neural Networks

Equivariant Graph Neural Networks (E3-GNNs) are a type of geometric graph neural network designed with the necessary inductive bias that their outcomes should not depend on the observer's viewpoint, achieving E(3)-equivariance. EGNNs are specifically designed to leverage information from geometric data, such as the three-dimensional structures of molecules or spatial relationships within physical systems. The SE(3)-Transformer [34] utilizes the advantages of self-attention to operate on large point clouds and graphs with varying numbers of points, ensuring SE(3)-equivariance for robustness. DimeNet [8] employs spherical Bessel functions and spherical harmonics to construct theoretically grounded orthogonal representations. EGNN [6] introduces a novel model to learn graphs equivariant to rotations, translations, reflections, and permutations, capable of modeling protein graphs at the residue level and predicting molecular properties. Similar efforts include GVP-GNN [17], which introduces a geometric vector perceptron for geometric and relational reasoning on large molecules' effective representations. MEAN [21] proposes a multi-channel equivariant graph network where each residue node references multiple coordinates for different atoms, primarily modeling atoms on the main chain. An advanced version, dyMEAN [22], introduces a more powerful multi-channel equivariant message-passing model capable of modeling protein structures at the full atomic level, applied in antibody design tasks.

3 Preliminary

In this section, we will formalize the protein-protein interaction prediction problem, *i.e.*, binding binary classification, binding energy prediction ($\triangle G$), and prediction of mutation effect ($\triangle\triangle G$). Firstly, a full-atom protein graph can be defined as follows [22]:

Definition 1. Full-atom Protein Graph. We describe the single protein as a graph $\mathcal{G}_{\mathcal{P}} = (\mathcal{V}_{\mathcal{P}}, \mathcal{E}_{\mathcal{P}})$, where $\mathcal{V}_{\mathcal{P}}$ refer to the vertices (*i.e., the residues*) and $\mathcal{E}_{\mathcal{P}}$ is the set of edges. Each residue $v_i = (s_i, \boldsymbol{X_i}), v_i \in \mathcal{V}_{\mathcal{P}}$, where s_i is the amino acid type, $\boldsymbol{X_i} \in \mathbb{R}^{3\times c_i}$ is the 3D coordinate matrix and c_i is the number of atoms in v_i. The edges in $\mathcal{E}_{\mathcal{P}}$ are constructed by finding the k-nearest Neighbors (kNN) of each residue, using the minimum pair-wise distance between v_i and v_j:

$$d\left(v_i, v_j\right) = \left\| \frac{1}{c_i} \sum_{p=1}^{c_i} \boldsymbol{X}_i^{(l)}(:, p) - \frac{1}{c_j} \sum_{q=1}^{c_j} \boldsymbol{X}_j^{(l)}(:, q) \right\|_2, \tag{1}$$

where $\frac{1}{c_i} \sum_{p=1}^{c_i} \boldsymbol{X}_i^{(l)}(:, p)$ represents the average of all atomic coordinates for the residue v_i , which we consider it as the coordinates of residue v_i.

After we define the single full-atom protein graph, we can further define the tasks in the protein-protein interaction problem.

Definition 2. Protein-protein Interaction(PPI) Netwok. A protein-protein interaction network can be indicated as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where V and E are the sets of nodes and edges. Each node $v_i \in V$ represents a type of protein, and each edge $e_i \in E$ represents the interaction between proteins.

Definition 3. Binding Binary Classification in PPI Networks. Given a PPI network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where each node $v_i \in V$ can represent a protein graph, our goal is to predict whether there will be an interaction between two proteins.

Definition 4. Free energy of binding($\triangle G$) in PPI networks. Given a protein-protein(*i.e., protein-ligand*) complex as a protein graph, our goal is to predict the binding energy E by the neural network.

Definition 5. Mutation effect ($\triangle\triangle G$) prediction in PPI Networks. The goal of the prediction of the mutational effect of PPI is to understand how a specific mutation (or mutations) alters the affinity of the protein-protein interaction. Formally, it can be expressed as:

$$\triangle\triangle G = \triangle G_{mut} - \triangle G_{wt}, \tag{2}$$

where $\triangle\triangle G$ is the change in the free energy of binding due to mutation, $\triangle G_{mut}$ is the free energy of binding of the mutated protein complex, $\triangle G_{wt}$ is the free energy of binding of the wild-type protein complex.

4 Methodology

4.1 Overview

Our model consists of two components. The first component is the Atomic Equivariant Attention Network(MEANT), an E3-equivariant graph neural network designed to adaptively capture geometric information at the atomic level between different residues. The second component describes how our model can be applied to solve PPI prediction problems, including binding binary classification, free energy ($\triangle G$) prediction, and mutation effect ($\triangle\triangle G$) prediction. Lastly, we provide proof that our model demonstrates E3 equivariance at the atomic structural level of proteins.

In this work, our primary objective is to model the atomic-level representation of proteins $\boldsymbol{h}^{\mathcal{G}}$ based on the given protein sequence s, initial residue features $\boldsymbol{h}^{(0)}$, and the 3D structure of the protein $\boldsymbol{X}$. For the initial embedding of the $\boldsymbol{h}^{(0)}$, we can adopt random initialization or use a large protein language model to initialize the embedding of each node of the protein. Next, we will describe how to use our proposed method to model proteins at the atomic level.

4.2 Atomic Equivariant Attention Network

The Atomic Equivariant Attention Network(MEANT) is designed to acquire the representation of residues and handle the problem of message passing between residues due to the different number of atoms in different residues. In this section, we first introduce the

MEANT architecture, which generates hidden representations and updates 3D coordinates for each residue. Then, we will explain how to extract the geometric relationships between residues and update their representations and coordinates.

Architecture. The illustration of the framework is presented in Figure 2. Given the initial embedding $\boldsymbol{h}_i^{(0)}$ and the coordinate matrix $\boldsymbol{X}_i$, the l_{th} layer updates the hidden representation $\boldsymbol{h}_i$ and the coordinate matrix $\boldsymbol{X}_i$ as follows:

$$\boldsymbol{m}_{ij} = \phi_m\left(\boldsymbol{h}_i^{(l)}, \boldsymbol{h}_j^{(l)}, \boldsymbol{R}_{ij}\right), \tag{3}$$

$$\boldsymbol{X}_{ij} = \boldsymbol{X}_i^{(l)} - \phi_x\left(\boldsymbol{m}_{ij}\right)\frac{1}{c_j}\sum_{k=1}^{c_j}\boldsymbol{X}_j^{(l)}(:,k), \tag{4}$$

$$\boldsymbol{h}_i^{(l+1)} = \phi_h(\boldsymbol{h}_i^{(l)}, \sum_{j\in\mathcal{N}(i)}\alpha_{ij}\boldsymbol{m}_{ij}), \tag{5}$$

$$\boldsymbol{X}_i^{(l+1)} = \boldsymbol{X}_i^{(l)} + \sum_{j\in\mathcal{N}(i)}\alpha_{ij}\boldsymbol{X}_{ij}, \tag{6}$$

where, ϕ_m, ϕ_x, ϕ_h are MLPs, $\mathcal{N}(i)$ is the neighbor set of residue i, $\boldsymbol{m}_{ij}$ and $\boldsymbol{X}_{ij}$ are scalar information(*a.k.a.* non-geometric information) and vector information(*a.k.a.* geometric information). $\boldsymbol{R}_{ij}$ is the processed message between two distinct-shape matrices $\boldsymbol{X}_i \in \mathbb{R}^{3\times c_i}$ and $\boldsymbol{X}_j \in \mathbb{R}^{3\times c_j}$, The attention score α_{ij} measures the impact of scalar information on the current residue representation. Next, we will introduce how to calculate $\boldsymbol{R}_{ij}$ and attention weight α_{ij}.

Geometric Relation $\boldsymbol{R}_{ij}$. The main goal of geometric relation $\boldsymbol{R}_{ij}$ is to extract atomic geometry between residues. We first calculate the distance map $\boldsymbol{D}_{ij}(p,q) = \left\|\boldsymbol{X}_i(:,p) - \boldsymbol{X}_j(:,q)\right\|_2$ between two residues. Then we use two learnable attribute matrices $\boldsymbol{A}_i \in \mathbb{R}^{c_i\times d}$ and $\boldsymbol{c_j \times d}$ and a residue-specific parameter $\boldsymbol{\omega}_i \in \mathbb{R}^{c_i\times d}$ to extract useful information from residue coordinates. Formally, the $\boldsymbol{R}_{ij} \in \mathbb{R}^{d\times d}$ can be written as:

$$\boldsymbol{R}_{ij} = \frac{\boldsymbol{\omega}_i^\top(\boldsymbol{A}_i + \boldsymbol{D}_{ij}\boldsymbol{A}_j)}{\|\boldsymbol{\omega}_i^\top(\boldsymbol{A}_i + \boldsymbol{D}_{ij}\boldsymbol{A}_j)\|_F}, \tag{7}$$

where $\|\dot\|_F$ is Frobenius norm [12]. From the above equation, it can be seen that the dimension of $\boldsymbol{R}_{ij}$ does not change with the variation of residue types, thereby achieving the alignment of geometric information between different residues.

Attention weight $\boldsymbol{\alpha_{ij}}$. The primary purpose of the attention weight α_{ij} is to measure the impact of neighboring residue representations on the current residue representation. We need a shared attention weight $\boldsymbol{a} \in \mathbb{R}^{2d}$ to calculate the attention score. The calculation method for α_{ij} is as follows:

$$\alpha_{ij} = \frac{\exp\left(\boldsymbol{a}^\top\left(\boldsymbol{h}_i^{(l)} \bigoplus \boldsymbol{h}_j^{(l)}\right)\right)}{\sum_{j\in\mathcal{N}(i)}\exp\left(\boldsymbol{a}^\top\left(\boldsymbol{h}_i^{(l)} \bigoplus \boldsymbol{h}_j^{(l)}\right)\right)}, \tag{8}$$

where $\bigoplus$ means vector concatenation, $\boldsymbol{h}_i^{(l)}$ and $\boldsymbol{h}_j^{(l)}$ are the representation of residues i and j in $l-th$ layer.

4.3 Prediction and Training Losses

Prediction. We calculate the protein representation through the graph pooling [38], after the representations $\boldsymbol{h}_i$ are obtained for every residue by MEANT encoder.

$$\begin{aligned}\boldsymbol{h}_k^{\mathcal{G}} &= \textbf{READOUT}(\mathcal{G}_k)\\ &= \text{CONCAT}\left(\sum_{i\in\mathcal{V}_k} h_i^{(l)} \mid l = 0, 1, \ldots, L\right),\end{aligned} \tag{9}$$

where CONCAT is vector concatenation operation.

To predict binary classification between protein k and protein k' in the PPI network, we need to obtain the representations of both proteins $\mathbf{h}_k^{\mathcal{G}}$ and $\mathbf{h}_{k'}^{\mathcal{G}}$. Then we can calculate the affinity score by $p_{kk'} = \phi_p(\mathbf{h}_k^{\mathcal{G}}, \mathbf{h}_{k'}^{\mathcal{G}})$, where ϕ_p is MLPs.

When we prepare to predict the binding free energy ($\triangle G$) and the mutation effect ($\triangle\triangle G$), our research subject is the protein complex, which can be seen as a whole protein with multiple chains. In protein complexes, the distance between interacting residues from different peptide chains is smaller than the boundary distance used for sampling neighboring residues during the construction of the protein graph. As a result, our message-passing mechanism can propagate information across different peptide chains without restriction.

To predict the binding free energy of protein complex k, we first calculate the complex representation $\boldsymbol{h}_k^{\mathcal{G}}$ through the MEANT encoder. And then predict the free energy through function ϕ_e: $\triangle G = \phi_e(\boldsymbol{h}_k^{\mathcal{G}})$, where ϕ_e is MLPs. What's more, we follow the settings in work [24] to predict mutation effect through the function ϕ_{mut}: $\triangle\triangle G = \phi_{mut}(\boldsymbol{h}_{mut}^{\mathcal{G}} - \boldsymbol{h}_{wl}^{\mathcal{G}})$, where $\boldsymbol{h}_{mut}^{\mathcal{G}}$ is the protein representation after the mutation and $\boldsymbol{h}_{wl}^{\mathcal{G}}$ is the protein representation of wild type.

Loss function. For three different tasks, we have different loss functions: binary classification prediction loss $\mathcal{L}_b$, binding free energy prediction loss $\mathcal{L}_e$, and mutation effect loss $\mathcal{L}_m$. We formally define the binary classification loss function as:

$$\mathcal{L}_b = -y_{kk'}\log p_{kk'} - (1 - y_{kk'})\log(1 - p_{kk'}), \tag{10}$$

where $y_{kk'}$ is the ground-truth of the interaction between protein k and protein k'.

For binding free energy prediction and mutation effect prediction, which are regression problems, we calculate the error between predicted values and actual values using MSE (Mean Squared Error) loss.

5 Analysis of E(3)-equivariance

We analyze the equivariance of MEANT with respect to E(3) symmetries, which consist of arbitrary 3D rotations and translations. Preserving such symmetries is essential for molecular modeling, since the physical properties of a structure should not depend on its absolute position or orientation in space. An E(3)-equivariant encoder guarantees that the learned representation is geometry-aware yet independent of arbitrary coordinate choices.

In MEANT, the geometric relations $\boldsymbol{R}_{ij}$ and the associated messages $\boldsymbol{m}_{ij}$ are constructed solely from inter-atomic distances. Distances remain unchanged under any rigid motion $(\boldsymbol{Q}, \boldsymbol{t})$, with $\boldsymbol{Q} \in$

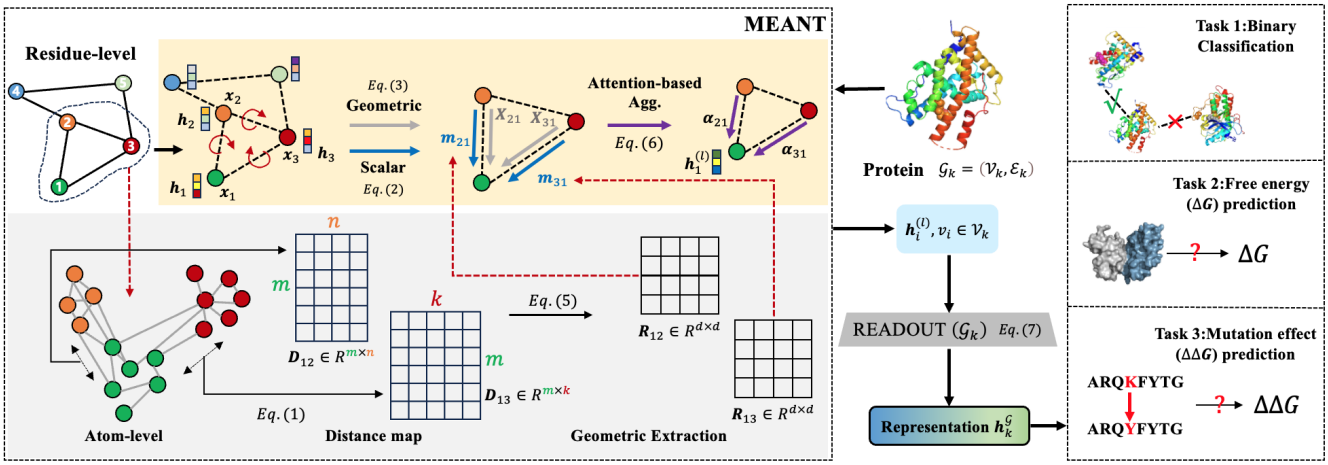

Figure 2: The overall framework of our proposed MEANT.

O(3) and $\boldsymbol{t} \in \mathbb{R}^3$. Consequently, the invariant design ensures that the attention weights and message-passing steps are unaffected by transformations of the inputs.

Furthermore, the coordinate update rule is designed so that the outputs transform consistently with the inputs. Specifically, if the input coordinates follow $\boldsymbol{X}_i^{(l)} \mapsto \boldsymbol{Q}\boldsymbol{X}_i^{(l)} + \boldsymbol{t}$, the updated coordinates satisfy $\boldsymbol{X}_i^{(l+1)} \mapsto \boldsymbol{Q}\boldsymbol{X}_i^{(l+1)} + \boldsymbol{t}$, while the hidden features $\boldsymbol{h}_i^{(l+1)}$ remain invariant.

Taken together, these properties show that MEANT is E(3)-equivariant: the encoder consistently respects rotational and translational symmetries, ensuring stable and physically meaningful representations across different molecular conformations.

6 Experiments

6.1 Experimental Setup

6.1.1 Datasets. For the prediction of binary classification in the PPI network, we extracted three datasets to evaluate the model.

Dataset	H-PPI	M2H-PPI	AM-PPI
Unique Nodes	11403	624	545
Number of PPIs	156066	39041	14664

H-PPI is extracted from the complete human PPI network, specifically focussing on interactions classified as physical binding. In this dataset, each node represents a distinct protein, and each edge indicates that there exists a physical binding between two proteins.

M2H-PPI is a subset of the H-PPI dataset, consisting exclusively of human cell membrane proteins. In this dataset, each protein represents a human cell membrane protein, and each edge also means that there exists a physical binding between two proteins.

AM-PPI is a PPI dataset specific to Actinosynnema mirum, a novel species. Similarly, we have filtered this dataset to include only interactions categorised as "physical binding" between proteins.

Then we choose two public datasets for the binding free energy prediction and mutation effect prediction.

SAbDab [31] is the Structured Antibody Database, which has 566 antibody-antigen complexes with binding affinity labels.

SKEMPI [13] has 348 protein complexes and approximately 6000 $\triangle\triangle G$ data points, which are used to train the models for $\triangle\triangle G$ prediction.

6.1.2 Baseline Approaches. To evaluate the effectiveness of our proposed method, we compare it with a comprehensive set of baselines across three tasks: binary PPI classification, binding free energy prediction, and mutation effect prediction.

For binary classification in PPI networks, we categorise the baselines into three groups:

Traditional GNNs: (1) GraphSAGE [9] is an inductive framework that generates node embeddings by sampling and aggregating information from local neighbourhoods. (2) GAT [35] introduces attention mechanisms into message passing, enabling the model to assign dynamic importance to neighbouring nodes.

GNN-based PPI methods: (3) GNN-PPI [25] models protein sequences as graphs and applies GNNs to predict binary interactions. (4) High-PPI [7] enhances this approach by incorporating hierarchical structural information from protein domains. (5) SemiGNN-PPI [42] combines multigraph modelling with a Mean Teacher framework to improve generalisation in semi-supervised PPI prediction.

Geometric GNNs: (6) EGNN [30] is an E(3)-equivariant GNN that maintains geometric symmetries such as rotation and translation. (7) GVP-GNN [17] uses geometric vector perceptrons to model both scalar and vector features in protein structures. (8) GBPNet [2] integrates geometry-aware attention and physical priors for structure-based PPI prediction.

For binding free energy prediction, we compare against the following **supervised learning baselines:** (1) Rosetta [1], a physics-based modelling suite that estimates binding affinity via energy minimisation and scoring functions. (2) FoldX [5], an empirical force-field-based tool for assessing mutation-induced stability changes. (3) ESM-1v [26], a protein language model trained on sequence

Table 1: Experimental results of binary classification in PPI networks.

Settings	Methods	M2H-PPI			H-PPI			AM-PPI		
		ACC(%)	AUC(%)	F1(%)	ACC(%)	AUC(%)	F1(%)	ACC(%)	AUC(%)	F1(%)
Random Initialization	GraphSAGE	80.32 ±0.18	87.07 ±0.21	79.89 ±0.16	82.65 ±0.15	89.12 ±0.17	81.82 ±0.12	80.81 ±0.09	86.32 ±0.06	69.74 ±0.08
	GAT	80.97 ±0.17	87.84 ±0.19	79.49 ±0.16	82.12 ±0.21	89.27 ±0.25	81.16 ±0.17	81.79 ±0.13	87.04 ±0.11	70.69 ±0.08
	GNN-PPI	83.32 ±0.21	90.41 ±0.18	81.78 ±0.15	84.5 ±0.15	91.02 ±0.19	84.28 ±0.13	83.59 ±0.09	89.47 ±0.13	72.61 ±0.08
	High-PPI	83.97 ±0.09	91.02 ±0.11	82.68 ±0.13	85.03 ±0.11	91.62 ±0.13	85.11 ±0.12	84.32 ±0.14	90.45 ±0.18	73.46 ±0.15
	SemiGNN-PPI	84.12 ±0.14	91.56 ±0.13	82.84 ±0.11	85.47 ±0.16	92.21 ±0.18	85.53 ±0.13	84.51 ±0.12	90.83 ±0.10	73.88 ±0.10
	EGNN	84.23 ±0.06	91.32 ±0.08	83.04 ±0.05	85.34 ±0.08	92.08 ±0.11	85.46 ±0.06	84.69 ±0.05	90.72 ±0.07	73.74 ±0.06
	GVP-GNN	84.54 ±0.15	91.78 ±0.18	83.32 ±0.12	85.43 ±0.13	92.14 ±0.11	85.72 ±0.13	84.53 ±0.09	90.61 ±0.13	73.59 ±0.11
	GBPNet	84.62 ±0.15	91.91 ±0.19	83.54 ±0.13	85.86 ±0.16	92.59 ±0.23	85.98 ±0.15	84.82 ±0.07	90.83 ±0.08	73.92 ±0.09
	MEANT	**86.96** ±0.21	**93.47** ±0.26	**85.28** ±0.19	**87.15** ±0.25	**94.15** ±0.31	**86.92** ±0.22	**86.24** ±0.15	**92.31** ±0.18	**75.43** ±0.13
Improvements(%)		2.77	1.70	2.03	1.36	1.68	1.09	1.67	1.65	2.04
ESM-2 Initialization	GraphSAGE	83.81 ±0.18	90.95 ±0.21	82.33 ±0.14	84.78 ±0.15	91.28 ±0.19	84.85 ±0.13	84.07 ±0.13	89.92 ±0.16	73.07 ±0.14
	GAT	83.97 ±0.21	91.23 ±0.24	82.56 ±0.17	85.11 ±0.23	91.79 ±0.27	84.98 ±0.24	84.66 ±0.17	90.87 ±0.15	73.84 ±0.14
	GNN-PPI	84.37 ±0.13	91.46 ±0.15	82.81 ±0.11	85.52 ±0.16	92.03 ±0.17	85.24 ±0.12	84.63 ±0.09	90.52 ±0.08	73.67 ±0.09
	High-PPI	84.85 ±0.15	91.93 ±0.13	83.37 ±0.11	85.83 ±0.16	92.33 ±0.18	85.92 ±0.13	85.19 ±0.11	91.05 ±0.12	74.18 ±0.08
	SemiGNN-PPI	85.27 ±0.12	92.34 ±0.15	84.27 ±0.12	86.34 ±0.11	93.27 ±0.14	86.44 ±0.13	85.46 ±0.16	91.64 ±0.13	74.63 ±0.09
	EGNN	85.27 ±0.14	92.16 ±0.15	84.07 ±0.14	86.29 ±0.13	92.98 ±0.16	86.37 ±0.12	85.83 ±0.09	91.78 ±0.11	74.93 ±0.09
	GVP-GNN	85.54 ±0.06	92.48 ±0.08	84.32 ±0.04	86.43 ±0.08	93.14 ±0.11	86.72 ±0.07	85.67 ±0.05	91.83 ±0.04	74.88 ±0.05
	GBPNet	85.86 ±0.09	92.71 ±0.11	84.62 ±0.08	86.82 ±0.13	93.56 ±0.15	86.95 ±0.11	85.88 ±0.04	92.03 ±0.05	75.07 ±0.05
	MEANT	**88.52** ±0.16	**94.89** ±0.19	**86.46**±0.13	**88.79** ±0.21	**96.26** ±0.25	**89.32** ±0.22	**87.79** ±0.15	**94.56** ±0.18	**76.63** ±0.12
Improvements(%)		3.10	2.35	1.84	1.71	2.26	2.89	2.22	2.75	2.08

Table 2: Experimental results of mutation effect prediction on SKEMPI datasets.

Category	Method	Per-Structure		Overall				
		Pearson	Spearman	Pearson	Spearman	RMSE	MAE	AUROC
Sequence Based	ESM-1v	0.0073	-0.0118	0.1921	0.1572	1.9609	1.3683	0.5414
	PSSM	0.0826	0.0822	0.0159	0.0666	1.9978	1.3895	0.5261
	MSA Transf.	0.1031	0.0868	0.1173	0.1313	1.9835	1.3816	0.5768
	Tranception	0.1348	0.1236	0.1141	0.1402	2.0382	1.3883	0.5885
Energy Function	Rosetta	0.3284	0.2988	0.3113	0.3468	1.6173	1.1311	0.6562
	FoldX	0.3789	0.3693	0.3120	0.4071	1.9080	1.3089	0.6582
Supervised	DDGPred	0.3752	0.3407	0.6580	0.4687	1.4998	1.0821	0.6992
	EGNN	0.3735	0.3423	0.6137	0.4665	1.9265	1.3321	0.6876
	RDE(End-to-End)	0.3873	0.3587	0.6373	0.4882	1.6198	1.1761	0.7172
	MEANT	**0.3892**	**0.3764**	**0.6682**	**0.4891**	**1.4217**	**1.0386**	**0.7189**

data without requiring structural input. (4) AF2 [19], a structure prediction model that provides transferable representations for downstream tasks. (5) FANN [28], a geometric model using frame-aligned representations to estimate binding free energy. (6) EGNN [6], adapted for regression-based prediction using equivariant representations.

For mutation effect prediction, we follow prior work and group baselines into three categories:

Traditional empirical energy functions: (1) Rosetta [23]; (2) FoldX [5].

Sequence/evolution-based methods: (3) ESM-1v [26]; (4) PSSM, which uses position-specific scoring matrices derived from evolutionary alignments; (5) MSA Transformer [29], which incorporates evolutionary information through attention to MSAs; (6) Tranception [27], which combines autoregressive modelling with alignment-based priors for predicting zero-shot fitness.

End-to-end learning models: (7) DDGPred [32], a deep model for predicting mutation effects directly; (8) EGNN [6]; (9) RDE [24], originally unsupervised, here adapted for supervised learning to predict PPI labels based on side-chain conformational changes.

6.1.3 Evaluation Task. We selected three representative PPI tasks to assess the effectiveness of our method in capturing protein expression capabilities.

Binary classification in the PPI network. The evaluation task is to predict whether there is an interaction between two proteins. All of the methods will output an affinity score for each link and the prediction can be treated as a binary classification problem. Therefore, we employ three metrics to fully evaluate our model and baseline methods: **accuracy (ACC)**, **AUC**, and **Macro-F1**.

Binding free energy prediction. We tested our model on SAbDab, which has 566 antibody-antigen complexes with binding affinity labels. We follow the settings in the work [16] and train the model on all binding affinity labels in SKEMPI except those appearing in the test set. Finally, we test the model with five random seeds and report the average **Spearman** correlation on the SAbDab.

Mutation effect prediction. We follow the settings in the work [24] and partition the dataset into three distinct folds, each comprising exclusive protein complexes absent in the other folds. We employ five metrics to assess the accuracy of $\triangle\triangle G$ prediction: **Pearson** correlation coefficient, **Spearman** correlation coefficient, **RMSE** (root mean squared error), minimised **MAE** (mean absolute error), and **AUROC** (area under the receiver operating characteristic curve).

6.1.4 Hyperparameter Settings. We implement MEANT using PyTorch 2.1.2 and train it with the Adam optimiser, setting the initial learning rate at 1×10^{-3} and applying a decay rate of 0.96. The model uses a hidden state size of 128, an embedding dimension of 64, and an atom attribute dimension of 16. Each residue connects to its 10 nearest spatial neighbours, and MEANT is composed of two message-passing layers. A batch size of 32 is used for all experiments and all hyperparameters are kept consistent between tasks to ensure fair comparisons. Experiments are conducted on a Linux server equipped with a single NVIDIA A800 GPU and an Intel(R) Xeon(R) Platinum 8358 CPU, enabling efficient training on full-atom protein graphs.

6.2 Main Experiments

This section evaluates our proposed MEANT method against several baseline approaches in three types of PPI prediction tasks. The results are shown in Table 1, Fig. 3, and Table 2 respectively. Testing across three different PPI prediction tasks showcases our method's ability to effectively represent proteins in PPI tasks.

6.2.1 Binary Prediction in PPI Networks. We evaluated the binary classification performance of our proposed model, MEANT, in comparison with three categories of baseline models in three benchmark PPI datasets (M2H-PPI, H-PPI, and AM-PPI) under two initialisation settings: random initialization and ESM-2 initialization. The experimental results, summarised in Table 1, lead to several key observations:

Our model significantly outperforms the baseline models in all three datasets under both settings. In terms of AUC, MEANT shows 2. 35%, 2. 26% and 2. 75% relative improvements over the best-performing baseline method, demonstrating the effectiveness of our design based on geometric graph neural networks at the atomic level for the prediction of PPI. Additionally, the performance drop of our model with two types of initial node representation is not as pronounced as with traditional graph neural networks, indicating the stability of our model and its ability to capture more useful information from the spatial structure of proteins.

What's more, MEANT shows strong robustness across different input initializations. Although traditional GNNs and even some geometric GNNs exhibit notable performance fluctuations between random and pretrained initializations, MEANT maintains stable and high performance. This stability indicates that our model can effectively extract meaningful structural features from raw inputs and is less dependent on external sequence-based embeddings, making it more adaptable in real-world scenarios with limited annotation or pretraining resources.

Compared to traditional GNNs, our approach offers a better representation of a protein at both the feature and structural levels through the use of geometric graph neural networks at the atomic level. While traditional GNNs focus on the connectivity of nodes in the PPI network and treat proteins as abstract nodes without internal structure, MEANT leverages detailed 3D geometric configurations within each protein. This allows it to capture fine-grained interaction cues that are critical for PPI prediction but are often overlooked in conventional approaches.

Furthermore, unlike GNN-based PPI methods, our approach integrates spatial structural information from proteins into our representations. Protein interactions depend on their spatial structures, and our model incorporates more spatial structural information than GNN-based methods, resulting in improved performance.

In contrast to other geometric graph neural network methods, we construct the entire protein graph at the atomic level. This enables the model to capture interaction-specific atomic configurations that are critical in determining the binding affinity and specificity. By modeling proteins at such fine granularity, MEANT is better equipped to distinguish subtle structural variations that influence PPIs, resulting in its consistently superior performance across all tasks.

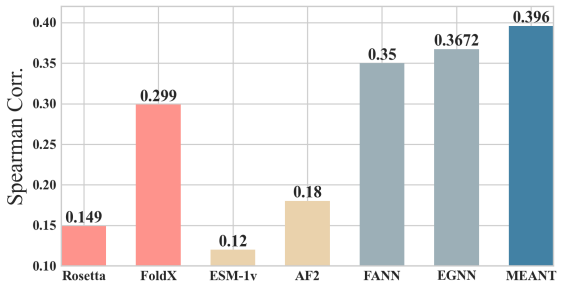

Figure 3: The results of binding free energy prediction. Larger scores indicate better performances.

6.2.2 Binding free energy ($\triangle G$) prediction. Our model, MEANT, was evaluated on the SabDab dataset for the prediction of binding free energy against six baseline methods, as illustrated in Figure 3. The results indicate that MEANT significantly outperforms energy function-based models, such as Rosetta and FoldX, as well as protein language models like AF2. Additionally, sequence-based protein language models exhibit poor performance in predicting binding free energy, as merely modeling sequence information is insufficient for accurately predicting the binding of two proteins. This experiment demonstrates that incorporating more atomic-level structural information in MEANT is advantageous for predicting binding free energy.

6.2.3 Mutation effect ($\triangle\triangle G$) prediction. We conducted comparisons on the SKEMPI dataset against three categories of baseline methods, evaluating the Pearson correlation coefficient, Spearman correlation coefficient, minimized RMSE, minimized MAE, and AUROC. The results are presented in Table 2. Our method outperforms all baselines, particularly in demonstrating improvements in each structural correlation metric, indicating higher reliability in practical applications.

Compared to sequence-based methods, our model captures more effective information from the structural level of proteins. Sequence-based models are based on the mining of evolutionary history and statistical analysis of multiple sequence alignments from large sequence databases. These approaches are not ideal for predicting the impact of mutations on general protein-protein interactions, especially when the two proteins involved may not belong to the same species. This leads to the underperformance of methods like MSA and ESM. In contrast to energy-based methods, our approach is specifically designed for the prediction of mutation effects, whereas baseline energy-based methods are primarily aimed at predicting protein structure and stability, making them less suitable for predicting binding free energy. Compared to supervised methods, our approach incorporates information at the original atomic level, integrating the structure of side chains to more effectively model the contact positions between proteins, thereby achieving better results.

6.3 Analysis of Model Components

Effect of atomic information. We conducted ablation experiments to substantiate the importance of incorporating atomic-level geometric structural information for PPI prediction tasks. Table 3 shows the results for the binary classification and the predictions of the mutation effect. In these experiments, the variant that excludes atomic-level structural information is denoted as "-Atomic." The results indicate that the absence of atomic structural information leads to the most pronounced decline in performance. This underscores the crucial role of atomic-level structural information in enhancing the accuracy of PPI predictions.

Effect of geometric relation extractor. To demonstrate the effectiveness of our designed geometric relation extractor, we replaced it with the average distance of all atoms between two residues. Table 3 presents the results for binary classification and mutation effect predictions, where the variant with the geometric relation extractor replaced is labeled as "$-\boldsymbol{R}_{ij}$". The results show that replacing the geometric relation extractor still has a significant impact on the experimental outcomes. This indicates that merely changing the graph structure to an all-atom level is insufficient; it is crucial to model the atomic-level geometric information between residues accurately.

Effect of residue-specific learnable matrix. In addition to the geometric relation extractor, we also investigated the impact of the residue-specific parameter matrix on the experimental results. We fixed the parameters of the matrix, denoted as "$-\boldsymbol{\omega_i}$", to create a variant. Table 3 shows that fixing the parameter matrix has a noticeable impact on the model performance. This suggests that there are differences in the geometric information extracted when considering different residues as central nodes, and it is essential to account for these differences.

Effect of attention aggregation. As shown in Equation 6, we employ an attention mechanism to update the representations and coordinate information of the residues. To validate the effectiveness of this attention mechanism, we created a variant removing it, denoted as '-atten'. The results in Table 3 indicate that the attention mechanism significantly contributes to the performance in our selected ablation experiments. Therefore, we incorporated this mechanism into our final model.

6.4 Analysis of Hyperparameter

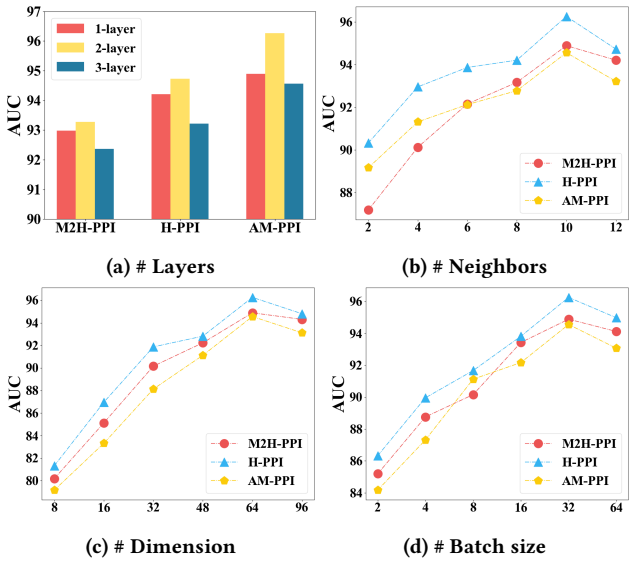

Figure 4: Performance under different hyperparameters.

In this subsection, we discuss the impact of primary hyperparameters on the performance of MEANT. Specifically, we analyse the number of layers in MEANT, the number of residue neighbours in the atomic protein graph, the dimension of the final representations, and the batch size. The results, presented in Fig. 4, reveal several key observations. First, utilising additional layers in the model enhances performance as a result of the central residue node's increased ability to aggregate information from more distant residues. However, as the number of layers continues to increase, the performance gains diminish and the time cost for model training and inference

Table 3: Ablation Study on binary PPI classification and mutation effect prediction. Best results are highlighted in bold.

Binary Classification						
Model	M2H-PPI		H-PPI		AM-PPI	
	AUC	F1	AUC	F1	AUC	F1
MEANT	**94.89** ±0.19	**86.46** ±0.13	**96.26** ±0.25	**89.32** ±0.22	**94.56** ±0.18	**76.63** ±0.12
-Atomic	92.87 ±0.13	84.11 ±0.15	93.27 ±0.21	86.42 ±0.20	92.39 ±0.14	73.91 ±0.11
-$\boldsymbol{R_{ij}}$	93.17 ±0.22	84.62 ±0.16	94.42 ±0.27	87.83 ±0.18	92.96 ±0.21	74.10 ±0.10
-$\boldsymbol{\omega_i}$	93.64 ±0.16	85.33 ±0.14	95.92 ±0.22	88.67 ±0.19	93.47 ±0.16	75.02 ±0.09
-Atten	93.68 ±0.18	85.27 ±0.12	95.26 ±0.21	88.19 ±0.17	93.82 ±0.18	75.34 ±0.10
Mutation Effect Prediction						
Model	Per-Structure		Overall			
	Pearson	Spearman	Pearson	Spearman	RMSE	AUROC
MEANT	**0.3892**	**0.3764**	**0.6682**	**0.4891**	**1.4217**	**0.7189**
-Atomic	0.3784	0.3660	0.6513	0.4726	1.4615	0.7032
-$\boldsymbol{R_{ij}}$	0.3816	0.3705	0.6549	0.4754	1.4478	0.7067
-$\boldsymbol{\omega_i}$	0.3852	0.3738	0.6601	0.4813	1.4332	0.7110
-Atten	0.3846	0.3722	0.6592	0.4807	1.4359	0.7103

grows non-linearly. Consequently, we have set the number of model layers to 2. Additionally, the model performance stabilises when the number of neighbours sampled for residues increases to 10, the dimension of final representations is set to 64, and the batch size is increased to 32. It is important to note that while increasing the dimension of the final representations and the batch size improves performance, it also significantly increases the computational resource requirements of our model.

6.5 Computation Efficiency Analysis

The computational efficiency of graph-based protein modeling is closely related to the granularity of the graph representation. Traditional residue-level GNNs construct coarse-grained graphs, where each node corresponds to a residue. These graphs typically contain N_r nodes and a limited number of edges reflecting sequential or structural proximity, leading to a per-layer complexity of $O(N_r \cdot d^2)$, where d is the hidden feature dimension. This design is computationally efficient but lacks fine-grained structural resolution.

In contrast, full-atom GNNs model proteins at the atomic level, where the number of nodes increases to $N_a \approx 10 \cdot N_r$, and pairwise interactions among atoms often result in dense connectivity. This leads to a message passing complexity of $O(N_a^2 \cdot d)$, which is significantly more expensive, especially for large proteins. Our method, MEANT, strikes a balance between expressiveness and efficiency by retaining atomic-level features, but restricting message passing to a residue-level graph. Specifically, we used atomic nodes for geometric encoding within each residue, followed by a residue-level message passing across the protein graph. This hybrid design reduces the overall time complexity to $O(N_r \cdot d^2 + N_a \cdot d)$, avoiding quadratic growth while preserving structural detail. Empirically, our model introduces only a modest overhead compared to residue-level GNNs, yet delivers significantly improved predictive performance.

7 Conclusion

In this paper, we introduce MEANT, a novel model designed for full-atom PPI prediction. Our method utilises learnable parameterized matrices to construct atom-level geometric knowledge between residues. This approach enables us to capture comprehensive structural information at the atomic level while avoiding the need to learn individual representations for each atom. Consequently, we ensure the computational efficiency of the model without compromising its expressive capability. The results of three tasks on five datasets demonstrate that MEANT consistently outperforms baseline models.

While the results are promising, MEANT currently considers only relative positional relationships between residues, ignoring more complex geometric details like side-chain rotational angles. Future work will address these limitations and explore using structures of unlabeled proteins and protein complexes to enhance the model's representation. Advancing protein affinity research can accelerate drug discovery, and this study may inspire the AI4Science community to tackle practical challenges in this field, promoting further development and having significant real-world implications.

8 Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (No. 62192784, U22B2038, 62472329). Cheng Yang is also supported by the Young Elite Scientists Sponsorship Program (No.2023QNRC001) by CAST.

GenAI Usage Disclosure

This paper includes text that was revised with the assistance of generative AI tools, specifically ChatGPT-4o, to improve grammar, clarity, and expression. The tools were used exclusively for language refinement purposes and did not contribute to the generation of technical content, experimental design, model architecture, or results. All scientific ideas, methods, experiments, and analyses presented in this paper are original contributions by the authors.

References

[1] Rebecca F Alford, Andrew Leaver-Fay, Jeliazko R Jeliazkov, Matthew J O'Meara, Frank P DiMaio, Hahnbeom Park, Maxim V Shapovalov, P Douglas Renfrew, Vikram K Mulligan, Kalli Kappel, et al. 2017. The Rosetta all-atom energy function for macromolecular modeling and design. Journal of chemical theory and computation 13, 6 (2017), 3031–3048.

[2] Sarp Aykent and Tian Xia. 2022. Gbpnet: Universal geometric representation learning on protein structures. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 4–14.

[3] Mayank Baranwal, Abram Magner, Jacob Saldinger, Emine S Turali-Emre, Paolo Elvati, Shivani Kozarekar, J Scott VanEpps, Nicholas A Kotov, Angela Violi, and Alfred O Hero. 2022. Struct2Graph: a graph attention network for structure based predictions of protein–protein interactions. BMC bioinformatics 23, 1 (2022), 1–28.

[4] Stefania Colonnese, Manuela Petti, Lorenzo Farina, Gaetano Scarano, and Francesca Cuomo. 2021. Protein-protein interaction prediction via graph signal processing. IEEE Access 9 (2021), 142681–142692.

[5] Javier Delgado, Leandro G Radusky, Damiano Cianferoni, and Luis Serrano. 2019. FoldX 5.0: working with RNA, small molecules and a new graphical interface. Bioinformatics 35, 20 (2019), 4168–4169.

[6] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. 2020. Se (3)-transformers: 3d roto-translation equivariant attention networks. Advances in neural information processing systems 33 (2020), 1970–1981.

[7] Ziqi Gao, Chenran Jiang, Jiawen Zhang, Xiaosen Jiang, Lanqing Li, Peilin Zhao, Huanming Yang, Yong Huang, and Jia Li. 2023. Hierarchical graph learning for protein–protein interaction. Nature Communications 14, 1 (2023), 1093.

[8] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. 2020. Directional message passing for molecular graphs. arXiv preprint arXiv:2003.03123 (2020).

[9] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. Advances in neural information processing systems 30 (2017).

[10] Kristin J Hope, Liqing Jin, and John E Dick. 2004. Acute myeloid leukemia originates from a hierarchy of leukemic stem cell classes that differ in self-renewal capacity. Nature immunology 5, 7 (2004), 738–743.

[11] Xiaotian Hu, Cong Feng, Tianyi Ling, and Ming Chen. 2022. Deep learning frameworks for protein–protein interaction prediction. Computational and Structural Biotechnology Journal 20 (2022), 3223–3233.

[12] Wenbing Huang, Jiaqi Han, Yu Rong, Tingyang Xu, Fuchun Sun, and Junzhou Huang. 2022. Equivariant graph mechanics networks with constraints. arXiv preprint arXiv:2203.06442 (2022).

[13] Justina Jankauskaitė, Brian Jiménez-García, Justas Dapkūnas, Juan Fernández-Recio, and Iain H Moal. 2019. SKEMPI 2.0: an updated benchmark of changes in protein–protein binding energy, kinetics and thermodynamics upon mutation. Bioinformatics 35, 3 (2019), 462–469.

[14] Kanchan Jha, Sriparna Saha, and Hiteshi Singh. 2022. Prediction of protein–protein interaction using graph neural networks. Scientific Reports 12, 1 (2022), 1–12.

[15] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2022. Antibody-antigen docking and design via hierarchical structure refinement. In International Conference on Machine Learning. PMLR, 10217–10227.

[16] Wengong Jin, Caroline Uhler, and Nir Hacohen. 2023. SE (3) denoising score matching for unsupervised binding energy prediction and nanobody design. In NeurIPS 2023 Generative AI and Biology (GenBio) Workshop.

[17] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael JL Townshend, and Ron Dror. 2020. Learning from protein structure with geometric vector perceptrons. arXiv preprint arXiv:2009.01411 (2020).

[18] Susan Jones and Janet M Thornton. 1996. Principles of protein-protein interactions. Proceedings of the National Academy of Sciences 93, 1 (1996), 13–20.

[19] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. 2021. Highly accurate protein structure prediction with AlphaFold. Nature 596, 7873 (2021), 583–589.

[20] Anowarul Kabir and Amarda Shehu. 2022. Graph neural networks in predicting protein function and interactions. In Graph Neural Networks: Foundations, Frontiers, and Applications. Springer, 541–556.

[21] Xiangzhe Kong, Wenbing Huang, and Yang Liu. 2022. Conditional antibody design as 3d equivariant graph translation. arXiv preprint arXiv:2208.06073 (2022).

[22] Xiangzhe Kong, Wenbing Huang, and Yang Liu. 2023. End-to-End Full-Atom Antibody Design. arXiv preprint arXiv:2302.00203 (2023).

[23] Julia Koehler Leman, Brian D Weitzner, Steven M Lewis, Jared Adolf-Bryfogle, Nawsad Alam, Rebecca F Alford, Melanie Aprahamian, David Baker, Kyle A Barlow, Patrick Barth, et al. 2020. Macromolecular modeling and design in Rosetta: recent methods and frameworks. Nature methods 17, 7 (2020), 665–680.

[24] Shitong Luo, Yufeng Su, Zuofan Wu, Chenpeng Su, Jian Peng, and Jianzhu Ma. 2023. Rotamer Density Estimator is an Unsupervised Learner of the Effect of Mutations on Protein-Protein Interaction. bioRxiv (2023), 2023–02.

[25] Guofeng Lv, Zhiqiang Hu, Yanguang Bi, and Shaoting Zhang. 2021. Learning unknown from correlations: graph neural network for inter-novel-protein interaction prediction. arXiv preprint arXiv:2105.06709 (2021).

[26] Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. 2021. Language models enable zero-shot prediction of the effects of mutations on protein function. Advances in Neural Information Processing Systems 34 (2021), 29287–29303.

[27] Pascal Notin, Mafalda Dias, Jonathan Frazer, Javier Marchena Hurtado, Aidan N Gomez, Debora Marks, and Yarin Gal. 2022. Tranception: protein fitness prediction with autoregressive transformers and inference-time retrieval. In International Conference on Machine Learning. PMLR, 16990–17017.

[28] Omri Puny, Matan Atzmon, Heli Ben-Hamu, Ishan Misra, Aditya Grover, Edward J Smith, and Yaron Lipman. 2021. Frame averaging for invariant and equivariant network design. arXiv preprint arXiv:2110.03336 (2021).

[29] Roshan M Rao, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. 2021. Msa transformer. In International Conference on Machine Learning. PMLR, 8844–8856.

[30] Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. 2021. E (n) equivariant graph neural networks. In International conference on machine learning. PMLR, 9323–9332.

[31] Constantin Schneider, Matthew IJ Raybould, and Charlotte M Deane. 2022. SAbDab in the age of biotherapeutics: updates including SAbDab-nano, the nanobody structure tracker. Nucleic acids research 50, D1 (2022), D1368–D1372.

[32] Sisi Shan, Shitong Luo, Ziqing Yang, Junxian Hong, Yufeng Su, Fan Ding, Lili Fu, Chenyu Li, Peng Chen, Jianzhu Ma, et al. 2022. Deep learning guided optimization of human antibody against SARS-CoV-2 variants with broad neutralization. Proceedings of the National Academy of Sciences 119, 11 (2022), e2122954119.

[33] Farzan Soleymani, Eric Paquet, Herna Viktor, Wojtek Michalowski, and Davide Spinello. 2022. Protein–protein interaction prediction with deep learning: A comprehensive review. Computational and Structural Biotechnology Journal (2022).

[34] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. 2018. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. arXiv preprint arXiv:1802.08219 (2018).

[35] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. Proceedings of ICLR 2018 (2017).

[36] Yan-Bin Wang, Zhu-Hong You, Xiao Li, Tong-Hai Jiang, Xing Chen, Xi Zhou, and Lei Wang. 2017. Predicting protein–protein interactions from protein sequences by a stacked sparse autoencoder deep neural network. Molecular BioSystems 13, 7 (2017), 1336–1344.

[37] Tianqi Wu and Jianlin Cheng. 2022. Atomic protein structure refinement using all-atom graph representations and SE (3)-equivariant graph neural networks. bioRxiv (2022), 2022–05.

[38] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? arXiv preprint arXiv:1810.00826 (2018).

[39] Fang Yang, Kunjie Fan, Dandan Song, and Huakang Lin. 2020. Graph-based prediction of Protein-protein interactions with attributed signed graph embedding. BMC bioinformatics 21, 1 (2020), 1–16.

[40] Andy GX Zeng, Suraj Bansal, Liqing Jin, Amanda Mitchell, Weihsu Claire Chen, Hussein A Abbas, Michelle Chan-Seng-Yue, Veronique Voisin, Peter van Galen, Anne Tierens, et al. 2022. A cellular hierarchy framework for understanding heterogeneity and predicting drug response in acute myeloid leukemia. Nature medicine 28, 6 (2022), 1212–1223.

[41] Lingling Zhao, Junjie Wang, Yang Hu, and Liang Cheng. 2020. Conjoint feature representation of GO and protein sequence for PPI prediction based on an inception RNN attention network. Molecular Therapy-Nucleic Acids 22 (2020), 198–208.

[42] Ziyuan Zhao, Peisheng Qian, Xulei Yang, Zeng Zeng, Cuntai Guan, Wai Leong Tam, and Xiaoli Li. 2023. SemiGNN-PPI: Self-Ensembling Multi-Graph Neural Network for Efficient and Generalizable Protein-Protein Interaction Prediction. arXiv preprint arXiv:2305.08316 (2023).