

MMPG: MoE-based Adaptive Multi-Perspective Graph Fusion for Protein Representation Learning

Yusong Wang^{1*}, Jialun Shen^{2*}, Zhihao Wu³, Yicheng Xu², Shiyin Tan², Mingkun Xu^{1†}, Changshuo Wang⁴, Zixing Song⁵, Prayag Tiwari⁶

¹Guangdong Institute of Intelligence Science and Technology, Zhuhai, China

²Institute of Science Tokyo, Tokyo, Japan

³Zhejiang University, Hangzhou, China

⁴University College London, London, U.K.

⁵University of Cambridge, Cambridge, U.K.

⁶Halmstad University, Halmstad, Sweden

Abstract

Graph Neural Networks (GNNs) have been widely adopted for Protein Representation Learning (PRL), as residue interaction networks can be naturally represented as graphs. Current GNN-based PRL methods typically rely on single-perspective graph construction strategies, which capture partial properties of residue interactions, resulting in incomplete protein representations. To address this limitation, we propose MMPG, a framework that constructs protein graphs from multiple perspectives and adaptively fuses them via Mixture of Experts (MoE) for PRL. MMPG constructs graphs from physical, chemical, and geometric perspectives to characterize different properties of residue interactions. To capture both perspective-specific features and their synergies, we develop an MoE module, which dynamically routes perspectives to specialized experts, where experts learn intrinsic features and cross-perspective interactions. We quantitatively verify that MoE automatically specializes experts in modeling distinct levels of interaction—from individual representations, to pairwise inter-perspective synergies, and ultimately to a global consensus across all perspectives. Through integrating this multi-level information, MMPG produces superior protein representations and achieves advanced performance on four different downstream protein tasks.

Introduction

Protein Representation Learning (PRL) has emerged as a fundamental methodology in bioinformatics, aiming to encode the structural, biochemical, and functional properties of proteins into informative representations. These learned representations are crucial across a wide range of downstream applications, including drug discovery, functional annotation, and protein design (Hoang et al. 2024). Since protein properties arise from complex interactions among amino acids (residues), it is essential for PRL methods to capture

*These authors contributed equally.

†Corresponding author: xumingkun@gdiist.cn

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

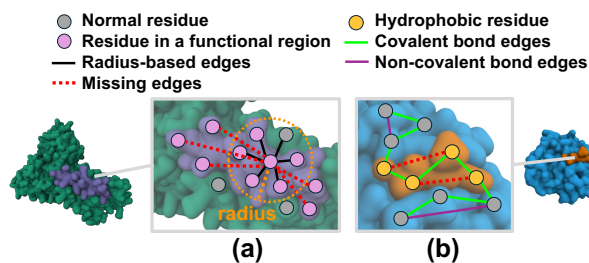


Figure 1: Limitations of protein graph construction from a single perspective. (a) A radius-based graph misses long-range connections. (b) A chemical-bond-based graph fails to capture the association of adjacent hydrophobic residues.

these patterns. A natural way to model such residue-level interactions is through graphs, where residues are represented as nodes and their relationships (e.g., spatial proximity or chemical bonds) as edges. Therefore, Graph Neural Networks (GNNs) have been widely adopted for PRL, as they enable message passing and aggregation across residues to learn expressive protein representations (Scarselli et al. 2009; Zhang et al. 2023; Wang et al. 2025b).

GNN-based PRL methods have achieved remarkable success in recent years (Wang et al. 2023; Zhang et al. 2023; Wang et al. 2025b). They vary in their protein graph construction strategies with specialized GNN architectures to capture protein properties effectively. For instance, ProNet (Wang et al. 2023) leverages radius-based graphs with hierarchical encoders for protein structural modeling, excelling in predicting enzyme reaction type. GearNet (Zhang et al. 2023) stacks sequential, radius, and K-Nearest Neighbors (KNN) edges in a graph to capture geometric proximity relationships of residues at different scales, and then employs a relational Graph Convolutional Network (GCN) to learn protein graph representation, achieving strong performance across multiple protein tasks. However, these methods are limited to single-perspective graph construction, capturing

only one semantic aspect of residue interactions. As illustrated in Figure 1 (a), a radius-based graph may miss functionally coupled distant residues, and a chemical-bond-based graph might overlook adjacent hydrophobic residues driven by the hydrophobic effect (Figure 1(b)) (Tanford 1978). This loss of information limits the expressive power of learned protein representations, constraining their performance at downstream protein tasks.

These limitations motivate us to explore a multi-perspective-based representation learning method. We begin by constructing a set of graphs, each encoding a distinct semantic perspective of residue interactions: 1) a physical-energetic perspective capturing interaction stability via knowledge-based potentials; 2) a chemical-functional perspective encoding residue similarities based on biochemical properties; and 3) a geometric-structural perspective modeling local spatial relationships. Critically, these perspectives are not isolated but exhibit complex interdependencies. Each perspective offers unique information and their synergistic interactions form the basis for a comprehensive understanding of protein properties. For example, a geometric graph identifies approximate neighbor residues, and a physical-energetic graph helps validate the stability of their interactions. To learn both information, we develop a Mixture-of-Experts (MoE) module. Through automatic routing of perspectives to experts, each expert learns different knowledge ranging from intra-perspective features to inter-perspective synergies. Our analysis of expert selection frequencies verifies that MoE specializes its experts to capture interaction levels, ranging from individual representations to pairwise inter-perspective synergies, and ultimately a global consensus across all perspectives. Through this **MoE-based adaptive Multi-Perspective Graph fusion framework**, MMPG produces expressive protein representations.

Main contributions are summarized as:

- We construct protein graphs from three semantic domains of physical, chemical, and geometric perspectives, providing comprehensive coverage of residue interactions beyond single-perspective limitations.
- We develop an MoE module that discovers and leverages multi-level interactions among perspectives to achieve effective multi-perspective information integration.
- We provide quantitative evidence that MoE module can specialize its experts to capture cross-perspective interactions at multiple levels. This multi-perspective integrating mechanism enables MMPG to achieve advanced performance on four different downstream protein tasks, demonstrating its effectiveness on PRL.

Related Work

Protein Representation Learning

PRL is commonly categorized into sequence-based and structure-based approaches. Sequence-based methods use architectures like word embeddings to capture individual residue features (Yang et al. 2018), 1D-CNNs to extract local sequence motifs (Kulmanov and Hoehndorf 2019), and Transformers to model long-range dependencies between residues (Lin et al. 2023). Structure-based approaches,

mainly based on GNNs, explicitly capture structural information, such as spatial relationships between residues, for a more comprehensive understanding of protein structures (Zhang et al. 2023; Wang et al. 2023; Jamasb et al. 2024; Wang et al. 2025b). The performance of GNN-based methods depends heavily on the initial protein graph construction, which remains a key bottleneck.

Protein Graph Construction

Protein graph construction abstracts a protein’s 3D conformation into a graph. In this work, we focus on residue-level graphs where a node corresponds to a residue. The residue interaction network is reflected in the edge connectivity patterns (Fasoulis, Paliouras, and Kavraki 2021). Constructing protein graphs is non-trivial, as multiple chemical, physical, and geometric properties of proteins should be considered. Based on these properties, various methods are developed. Geometric-based methods, such as radius, KNN, or Delaunay triangulation, are used to capture spatial relationships (Quan et al. 2024; Wang et al. 2025b; Jamasb et al. 2024). Chemical-based methods define edges based on chemical relationships of residues, such as chemical bonds (Baldassarre et al. 2020). Most PRL methods adopt only one perspective, thereby partially encoding residue interactions. Few attempts try multi-perspective construction, merely adding another perspective on top of a primary one (Baldassarre et al. 2020; Chiang, Hui, and Chang 2022). They use coarse fusion strategies, such as edge stacking or perspective feature concatenation (Baldassarre et al. 2020), where indiscriminate merging of cross-perspective information introduces noise and dilutes key structural signals. Instead, MMPG constructs protein graphs from physical, chemical, and geometric perspectives, using an MoE model to adaptively fuse their information, yielding richer protein representations.

Mixture of Experts

Recent advances highlight MoE’s effectiveness in capturing both shared and specialized data patterns. In multi-task learning, MoE is introduced to balance common and task-dependent representations (Ma et al. 2018; Chen et al. 2023). In multi-modal learning, Wang et al. (2025a) and Liang et al. (2025) employ MoE to disentangle modality-shared/specific patterns, enabling more effective cross-modal interaction modeling. Motivated by its success, we explore MoE in PRL as a principled module to capture both shared and specific information from multi-perspective protein graphs.

Preliminary

Notations & problem definition. Let $G = (\mathcal{V}, \mathcal{E}, \mathcal{P})$ be a protein graph, where $\mathcal{V} = \{v_i\}_{i=1}^n$ denotes the set of nodes (residues, each represented by its central carbon atom C_α), and $\mathcal{E} = \{e_{ij}\}$ represents edge set. $\mathcal{P} = \{\mathbf{p}_i\}_{i=1, \dots, n}$ denotes coordinate set, where each $\mathbf{p}_i \in \mathbb{R}^{3 \times 1}$ represents the coordinate of node i . Node features are a concatenation of learned embeddings for both residue type and side-chain conformation. Edge features are built from a rotationally invariant relative spatial encoding scheme (Fan et al. 2023), combining sequence information with normalized relative

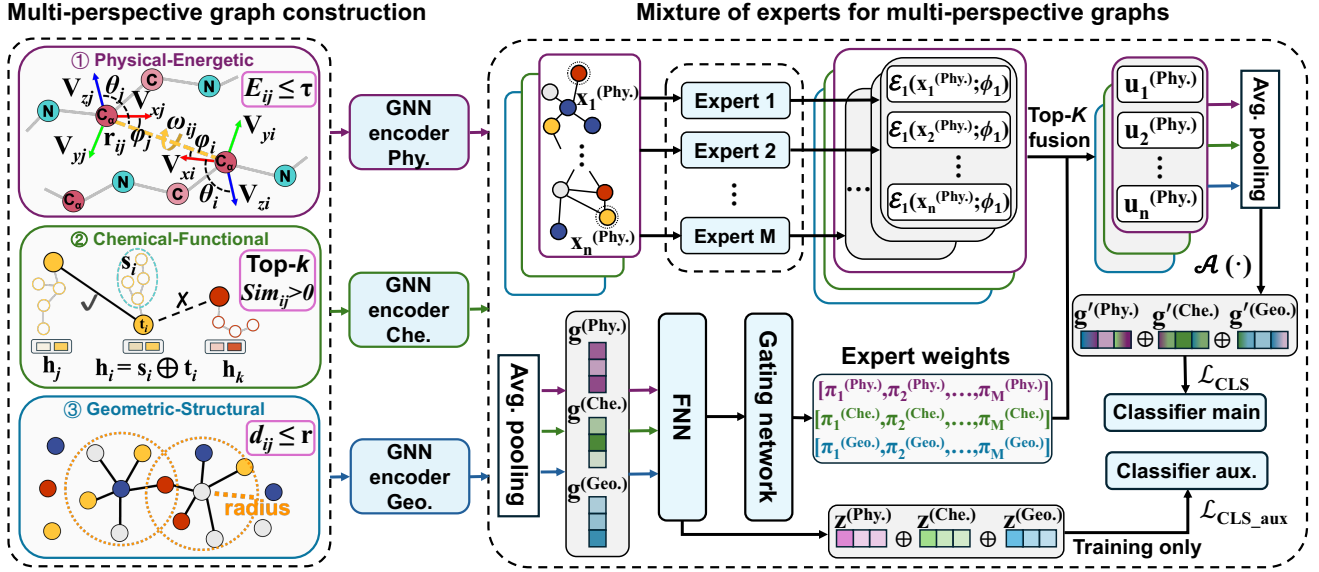


Figure 2: Overview of the proposed MMPG framework, which consists of two stages: (1) Multi-Perspective Graph Construction. Three graphs are constructed to model physical, chemical, and geometric properties of residue interaction, and (2) MoE learning scheme. Perspectives are routed to specialized experts, enabling dynamic representation learning across different perspectives.

position vectors, projections in the local coordinate system, and pairwise distances. Objective of protein prediction is to learn a mapping $f: G \rightarrow \mathcal{Y}$, where \mathcal{Y} is the set of labels.

Graph encoder. All graph encoders in our framework use an edge-aware GCN (Wang et al. 2025b). The hidden feature of node i is updated from layer l to layer $l + 1$ via:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{j \in \mathcal{N}(i)} \frac{1}{\sqrt{|\mathcal{N}(i)| |\mathcal{N}(j)|}} (\mathbf{h}_j^{(l)} \mathbf{W}_n) \odot (\mathbf{e}_{ij} \mathbf{W}_e) \right), \quad (1)$$

where $\mathbf{h}_i^{(l)}$ is the feature vector of node i at layer l , $\mathcal{N}(i)$ denotes the neighbors of i , \mathbf{e}_{ij} is the edge feature between nodes i and j , \mathbf{W}_n and \mathbf{W}_e are learnable weights, \odot denotes element-wise product, and $\sigma(\cdot)$ is the activation function.

Methodology

In this section, we first present our multi-perspective protein graph construction strategies. We then introduce an MoE framework for integrating these multi-perspective protein graphs. The overall architecture is illustrated in Figure 2.

Multi-Perspective Protein Graph Construction

Physical-Energetic Perspective. To endow our protein graphs with rich physical-energetic semantics, we adopt KORP (López-Blanco and Chacón 2019), a knowledge-based pairwise potential which quantitatively captures effective inter-residue interaction energies from experimentally solved structures, thus directly reflecting conformational preferences and stability determinants of proteins (Sippl 1995; Miyazawa and Jernigan 1996).

KORP is grounded on a coarse-grained residue representation and a 6-D joint probability distribution that considers

the relative orientation and position of residue pairs. As illustrated in the left top of Figure 2, each residue pair (i, j) is described by: a distance parameter \mathbf{r}_{ij} ; four angular parameters $\theta_i, \phi_i, \theta_j, \phi_j$, which are the polar coordinates of \mathbf{r}_{ij} in the local 3D reference frame of each amino acid, defined as:

$$\begin{aligned} \mathbf{V}_z &= (\mathbf{r}_{CC_\alpha} + \mathbf{r}_{NC_\alpha}) / |\mathbf{r}_{CC_\alpha} + \mathbf{r}_{NC_\alpha}|, \\ \mathbf{V}_y &= (\mathbf{V}_z \times \mathbf{r}_{NC_\alpha}) / |\mathbf{V}_z \times \mathbf{r}_{NC_\alpha}|, \\ \mathbf{V}_x &= (\mathbf{V}_y \times \mathbf{V}_z), \end{aligned} \quad (2)$$

where $\mathbf{r}_{CC_\alpha} = \mathbf{r}_C - \mathbf{r}_{C_\alpha}$ and $\mathbf{r}_{NC_\alpha} = \mathbf{r}_N - \mathbf{r}_{C_\alpha}$ are vectors from C_α to carbonyl carbon (C) and nitrogen (N) atoms of the same residue, respectively; torsional angle ω_{ij} describes relative rotation of vectors \mathbf{V}_{zi} and \mathbf{V}_{zj} along \mathbf{r}_{ij} axis. Potential for residue pair (i, j) with types (a, b) is given by inverse-Boltzmann equation (Choulli and Stefanov 1996):

$$E_{ij} = -RT \ln \frac{P_{ab}^{obs}(r_{ij}, \theta_i, \phi_i, \theta_j, \phi_j, \omega_{ij}) + z}{P^{ref}(r_{ij}, \theta_i, \phi_i, \theta_j, \phi_j, \omega_{ij}) + z}, \quad (3)$$

where R is the molar gas constant, T is the temperature, P_{ab}^{obs} is the joint probability at the given relative distance and orientation of observing two amino acids i and j of type a and b , P^{ref} is the reference probability using the classical reference state (Samudrala and Moult 1998), and z is a smoothing constant to stabilize low-count statistics. We connect nodes i and j if $E_{ij} \leq \tau$, where τ is a predefined threshold. This yields a protein graph whose edges encode stable structural motifs and protein backbone rigidity, capturing critical long-range and non-covalent interactions (e.g., hydrophobic effects), thereby providing physical insights for PRL.

Chemical-Functional Perspective. This perspective aims to model residue relationships based on chemical-functional similarities. Thus, we construct residue embeddings based

on amino acid type and side chain conformation. The amino acid type determines the residue’s intrinsic biochemical attributes (e.g., polarity, charge) (Kyte and Doolittle 1982), and the side-chain conformation dictates the spatial presentation of chemical properties (Janin et al. 1978). They together define a residue’s chemical semantics, determining its biological function, such as being part of a catalytic center. The similarity among these embeddings identifies residues with analogous roles rather than those merely interacting. We implement these two embeddings as follows:

Amino acid type embedding. Each residue is characterized by its underlying amino acid type, which is mapped to a unique index and embedded into a learnable vector:

$$\mathbf{t}_i = \text{Embedding}(\text{type}_i). \quad (4)$$

These embeddings learn to reflect the intrinsic chemical properties in a data-driven manner. *Side chain conformation embedding* is based on a residue’s first four torsion angles ($\chi^1, \chi^2, \chi^3, \chi^4$), which uniquely determine the side chain conformation given a fixed protein backbone (Wang et al. 2023). Each torsion angle uses a sine–cosine encoding:

$$\mathbf{s}_i = \text{concat}([\sin \chi_i^n, \cos \chi_i^n]_{n=1}^4), \quad (5)$$

where the raw embedding \mathbf{s}_i is subsequently projected by a linear layer to the same dimension as \mathbf{t}_i . To encode their combined effect, \mathbf{t}_i and \mathbf{s}_i are concatenated and passed through a feed-forward neural network (FNN), yielding the chemical-functional embedding:

$$\mathbf{h}_i = \text{FNN}(\mathbf{t}_i \oplus \mathbf{s}_i). \quad (6)$$

The similarity between residues i and j is then quantified by the cosine similarity of their state embeddings:

$$\text{Sim}_{ij} = \frac{\mathbf{h}_i^T \mathbf{h}_j}{\|\mathbf{h}_i\|_2 \|\mathbf{h}_j\|_2}. \quad (7)$$

We construct edges using a hybrid thresholded top- k strategy: For each residue i , we first identify its k most similar neighbors, and then form edges only with those neighbors j with a positive similarity score ($\text{Sim}_{ij} > 0$), ensuring both graph sparsity and semantically meaningful connections.

Geometric-Structural Perspective. We construct the graph based on geometric proximity to represent the local structural context of the protein. Specifically, an edge is added between residues i and j if $d_{ij} \leq r$, where d_{ij} is the Euclidean distance between their C_α atoms and r is a pre-defined radius. These geometric edges complement the sequential connectivity by capturing residues that may be distant in sequence but are brought into close proximity in the folded structure. This construction is critical for encoding the three-dimensional spatial relationships that are essential for protein function (Osadchy and Kolodny 2011).

After constructing graphs from three distinct perspectives, we apply a separate graph encoder to each, yielding node embeddings that capture perspective-specific features.

Mixture of Experts for Multi-Perspective Graphs

To effectively integrate multi-perspective features, we leverage an MoE module to process three protein graphs. This design enables learning both perspective-specific information

and cross-perspective interaction semantics through different experts, yielding comprehensive protein graph representations. MoE consists of two components: 1) a shared pool of experts, where each expert implements a graph encoder to process graphs from any perspective, but through training naturally develops specialization for certain intra- or inter-perspective patterns; and 2) a gating network that produces weights for all experts based on the input perspective graph and routes the graph to the most suitable experts accordingly.

Given the graph-level embedding of a perspective $\mathbf{g}^{(p)} = \frac{1}{n} \sum_{i=1}^n \mathbf{h}_i^{(p)}$ (where p denotes perspective), the gating weight of the m -th expert for perspective p is:

$$\pi_m^{(p)}(\mathbf{g}; \boldsymbol{\theta}) = \frac{\exp(\mathcal{G}(\mathbf{g}^{(p)}; \boldsymbol{\theta})[m])}{\sum_{i=1}^M \exp(\mathcal{G}(\mathbf{g}^{(p)}; \boldsymbol{\theta})[i])}, \quad \forall m \in M, \quad (8)$$

where $\mathcal{G}(\cdot)$ is the gating function, a one-layer FNN parameterized by $\boldsymbol{\theta}$ that outputs a weight vector of dimension M ; $[m]$ denotes its m -th element. Following Shazeer et al. (2017), we adopt a top- K strategy, selecting the K experts with the highest gating weights for each input. The node representation for perspective p is a weighted sum of K experts:

$$\mathbf{u}_i^{(p)} = \sum_{k=1}^K \pi_k^{(p)} \mathcal{E}_k(\mathbf{x}_i^{(p)}; \boldsymbol{\phi}_k), \quad (9)$$

where $\mathcal{E}_k(\cdot)$ denotes the k -th expert parameterized by $\boldsymbol{\phi}_k$, $\mathbf{x}_i^{(p)}$ is the input feature of node i in perspective p . This weighted fusion allows each node to integrate information from both its perspective-specific experts and shared experts across perspectives, thereby capturing both unique and shared semantic features. Furthermore, a load-balancing regularizer is imposed to encourage balanced expert selection (Shazeer et al. 2017), which is excluded from the task loss. The updated graph-level embedding for each perspective is obtained by a mean pooling readout over all nodes:

$$\mathbf{g}'^{(p)} = \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i^{(p)}. \quad (10)$$

Finally, we use concatenation as aggregation function $\mathcal{A}(\cdot)$ to obtain protein representation $\mathbf{g}_{\text{fused}}$ from P views:

$$\mathbf{g}_{\text{fused}} = \mathcal{A}\left(\left\{\mathbf{g}'^{(p)}\right\}_{p=1}^P\right). \quad (11)$$

Optimization

MMPG is trained using a joint task loss that integrates both global and auxiliary supervision¹:

$$\mathcal{L}_{\text{TASK}} = \mathcal{L}_{\text{CLS}} + \lambda \mathcal{L}_{\text{CLS_aux}}, \quad (12)$$

\mathcal{L}_{CLS} provides global supervision to $\mathbf{g}_{\text{fused}}$. $\mathcal{L}_{\text{CLS_aux}}$ provides auxiliary supervision to the concatenated FNN outputs $\text{concat}(\mathbf{z}^{(p)})_{p=1}^3$, which encourages the upstream GNN encoders to produce task-discriminative features that are then fed to the gating network for stable expert routing. Coefficient λ balances the contributions of the two losses.

¹Following (Fan et al. 2023), we use NLL loss for single-label protein classification tasks (RC and FOLD) and BCE loss for multi-label protein classification tasks (GO and EC).

Method	EC	BP	GO MF	CC	Fold	FOLD Super.	Fam.	Reaction
<i>Protein Representation Learning Methods</i>								
3DCNN	0.077	0.240	0.147	0.305	31.6	45.4	92.5	72.2
GraphQA	0.509	0.308	0.329	0.413	23.7	32.5	84.4	60.8
ProtBERT-BFD	0.838	0.279	0.456	0.408	26.6	55.8	97.6	72.2
GVP	0.489	0.326	0.426	0.420	16.0	22.5	82.8	65.5
LM-GVP	0.664	0.417	0.545	0.527	48.3	70.3	99.5	85.3
DeepFRI	0.631	0.399	0.465	0.460	15.3	20.6	73.2	63.3
IEConv	-	0.421	0.624	0.431	47.6	70.2	99.2	87.2
ProNet	-	-	-	-	52.7	70.3	99.3	86.4
GearNet	0.810	0.400	0.581	0.430	48.3	70.3	99.5	85.3
ESM-2	0.861	0.460	0.662	0.427	38.5	81.5	99.2	-
CDCConv	0.820	0.453	0.654	0.479	56.7	77.7	99.6	88.5
EPGGCL	0.885	0.454	0.659	0.477	59.8	80.8	99.5	89.0
<i>Protein Graph Construction Strategies</i>								
Chemical bond	0.620	0.301	0.552	0.411	13.6	14.3	68.0	46.6
KNN	0.821	0.434	0.610	0.415	55.4	75.6	98.9	85.2
Radius	0.840	0.419	0.631	0.419	55.3	75.7	99.3	84.1
Delaunay triangulation	0.774	0.332	0.567	0.424	24.9	47.0	94.2	81.6
MMPG	0.893	0.463	0.663	0.489	60.9	79.5	99.6	89.0

Table 1: Comparison of MMPG with two different types of baselines across multiple protein-related tasks. “-” indicates results not reported in the original paper. **Bold** shows the best performance.

Experiments

Experiment Setting

Tasks. We evaluate on four representative protein tasks: (1) Protein Fold Classification (FOLD) (Hermosilla et al. 2021) at fold/superfamily/family levels, (2) Enzyme Reaction Classification (Reaction) (Hermosilla et al. 2021), (3) Gene Ontology Term Prediction (GO) (Gligorijević et al. 2021) of three sub-ontologies: biological process (BP), molecular function (MF), and cellular component (CC), (4) Enzyme Commission (EC) (Gligorijević et al. 2021) number prediction, using established dataset splits (Fan et al. 2023). **Evaluation Metrics.** Following (Fan et al. 2023), we use Top-1 accuracy for single-label tasks (FOLD, Reaction); and F_{max} for multi-label tasks (GO, EC), computing the optimal F1-score across prediction thresholds per protein.

Baselines. We select two types of baselines: 1) *protein representation learning methods*: 3DCNN (Derevyanko et al. 2018), GraphQA (Baldassarre et al. 2020), ProtBERT-BFD (Elnaggar et al. 2020), LM-GVP (Wang et al. 2021), DeepFRI (Gligorijević et al. 2021), GVP (Jing et al. 2021), IEConv (Hermosilla and Ropinski 2022), ProNet (Wang et al. 2023), GearNet (Zhang et al. 2023), ESM-2 (Lin et al. 2023), CDCConv (Fan et al. 2023), and EPGGCL (Wang et al. 2025b); 2) *protein graph construction strategies*: KNN (Jamashb et al. 2024), Radius (Guo et al. 2025), Delaunay triangulation (Khade et al. 2023), and Chemical bond (Jamashb et al. 2022). For fairness, the graph encoder architecture in MMPG is kept consistent with baselines, differing only in the graph construction strategy.

Details of Training Setup. MMPG is trained on RTX 3090

GPUs. We use SGD (momentum=0.9, weight decay=5e-4) with an initial LR=1e-2 and a multi-step scheduler. All results are averaged over 5 random seeds. See code for details.

Quantitative Results

Performance comparison across four protein tasks is shown in Table 1. We have the following observations: 1) MMPG demonstrates superior performance across most tasks, particularly on complex functional prediction. Notably, MMPG achieves a best score of 0.893 on EC and 0.489 on GO-CC. This suggests that MMPG accurately models enzymatic function and cellular roles by integrating diverse information sources (physical, chemical, and geometric). While ESM-2 excels in FOLD-Superfamily, an evolutionary task, MMPG dominates Fold and Family levels, highlighting its superior capability of precise geometric and structural characterization. 2) MMPG significantly outperforms single-perspective graph construction strategies, demonstrating that combining physical, chemical, and geometric information yields more expressive protein representations.

Ablation Study

Ablation study in Table 2 validates each design of MMPG. We have the following observations: 1) Each perspective provides unique and task-critical information. Removing any perspective leads to a significant performance drop. 2) MoE proves more effective than other fusion strategies. Edge stack performs poorly (e.g., 0.794 on EC task) as dense, conflicting connections obscure structural patterns, while simple concatenation fails to adequately model inter-

Method	EC	BP	GO MF	CC	Fold	FOLD Super.	Fam.	Reaction
w/o Physical-energetic	0.874	0.448	0.652	0.466	59.0	77.4	99.5	84.9
w/o Chemical-functional	0.851	0.440	0.635	0.454	57.0	76.7	99.3	86.5
w/o Geometric-structural	0.821	0.422	0.623	0.436	52.0	72.1	99.0	84.3
w/o MoE (edge stack)	0.794	0.372	0.582	0.446	54.5	75.7	99.3	85.6
w/o MoE (concatenation)	0.865	0.426	0.600	0.452	56.9	76.2	99.3	84.9
w/o \mathcal{L}_{CLS_aux}	0.881	0.456	0.648	0.474	58.8	77.4	99.5	87.8
Complete	0.893	0.463	0.663	0.489	60.9	79.5	99.6	89.0

Table 2: Ablation study evaluates the effectiveness of each component by removing the three semantic perspectives and the MoE fusion module one by one. The general drop in performance across all tasks demonstrates the necessity of each component.

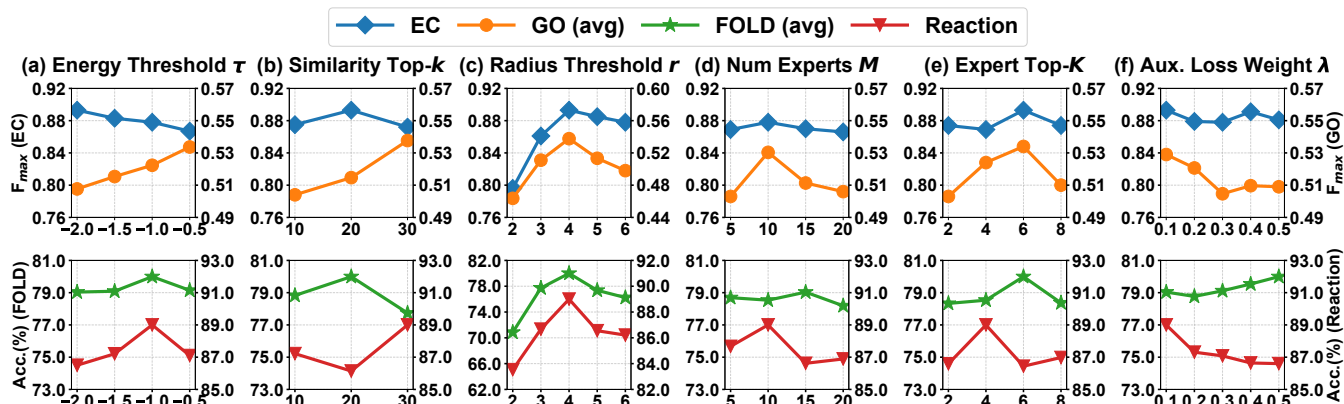


Figure 3: The plots show MMPG’s performance as key hyperparameters for the graph construction and MoE module are varied.

perspective interactions, resulting in poor performance on GO tasks. 3) The auxiliary supervision \mathcal{L}_{CLS_aux} provides beneficial refinement. By direct alignment with downstream tasks, \mathcal{L}_{CLS_aux} ensures that the gating network receives task-informative representations, thereby reducing noise from irrelevant features to facilitate reliable expert routing.

Parameter Analysis

We analyze the sensitivity of MMPG to key hyperparameters, grouped into multi-perspective protein graph construction and MoE, with results shown in Figure 3.

Graph Construction Hyperparameters. Results (Figure 3 (a), (b), and (c)) indicate that the optimal graph construction hyperparameters align with the level of informational granularity required by each task. For the energy threshold τ , optimal values range from -2.0 (for EC) to -0.5 (for GO). EC hinges on identifying the enzyme’s active site. A strict energy threshold filters out weak interactions, retaining structurally stable regions (e.g., enzyme active sites) crucial for catalytic function (Ribeiro et al. 2018). In contrast, GO often describes systemic functions involving multi-domain collaboration. Looser energy threshold preserves the weaker, long-range contacts that form the interaction pathways essential for these synergistic effects (Sanyal et al. 2012). For the similarity, the optimal neighbor count k is 20 for EC and FOLD, and increases to 30 for GO and Reaction. EC and

FOLD focus on specific residue combinations (e.g., the catalytic triad in EC), requiring strict similarity filtering (Martin et al. 1998), whereas GO and Reaction involve broad functional categories or types of reactions, benefiting from capturing more diverse residues (Ashburner et al. 2000). For the radius threshold r , a small radius ($r < 4$) cannot capture sufficient structural information for robust geometric encoding, while a large radius ($r > 4$) decreases the discriminativeness of local spatial patterns by including structurally irrelevant residues, thus leading to inferior performance across tasks.

MoE Hyperparameters. Results are shown in Figure 3 (d), (e), and (f). We first explore the optimal total number of experts M , by keeping the ratio of selected experts K fixed at 40%. Performance peaked at $M = 10$ for most tasks (GO, FOLD, and EC), suggesting this provides a balance between expert diversity and model complexity. The performance degradation beyond $M = 10$ likely stems from expert underutilization: each expert receives proportionally fewer samples, resulting in insufficient specialization. With the expert pool fixed at $M = 10$, we analyzed the optimal number of selected experts K . We find an optimal range where K is 4 or 6. A small K (e.g., 2) is likely insufficient to draw knowledge from these perspectives for a given input; while a large K (e.g., 8) may force the gating network to select less relevant experts, harming the precision of the fusion. This reflects a trade-off between coverage and specialization. Fi-




Figure 4: Expert selection frequency of MoE module for each input perspective on the FOLD task.

nally, we analyze the weight of the auxiliary loss λ . Most tasks achieve optimal performance with a small weight of $\lambda = 0.1$, but FOLD task benefits from a larger weight of $\lambda = 0.5$. Strong auxiliary supervision provides a direct learning signal, suggesting that FOLD’s expert routing requires explicit task guidance rather than implicit learning.

Analysis of Expert Specialization

The expert selection patterns of the MoE module provide insights into interactions among different semantic perspectives. For a detailed analysis, we investigate the FOLD task under the optimal hyperparameter configuration, where $K = 6$ experts are selected from a pool of $M = 10$. Figure 4 shows that the MoE learns a sophisticated, structured division of labor, manifested as three distinct expert-role types: 1) generalist experts (Experts 3, 4, and 10) are consistently utilized across all perspectives; 2) collaborative experts (Experts 2, 6, and 7) are utilized for specific perspective pairs; 3) specialized experts (Experts 1, 8, and 9) are used for a single perspective. These patterns reveal a hierarchical organization of interactions among perspectives with three distinct levels: At the global level, all perspectives are correlated through fundamental protein properties, such as the intrinsic relationship between structure and function. At the intermediate level, certain perspective pairs exhibit intrinsic interactions. For example, the physical and geometric perspectives both characterize protein structural properties. At the granular level, each perspective preserves its own irreplaceable semantic information that cannot be substituted or inferred from others. Overall, MoE can uncover and disentangle deep inter-perspective interactions, yielding a more informative protein representation.

Visualization

We use UMAP projection (McInnes et al. 2018) to visualize learned representations for the Reaction task, comparing MMPG with edge stack fusion strategy to assess representation quality. To demonstrate the results clearly, we randomly select seven classes from the Reaction task. As visualized in Figure 5, MMPG exhibits distinct, compact clusters with clear boundaries, while edge-stacking yields scattered distributions with significant class overlap. This superior clustering (quantified by tighter intra-class distance: 13.25 vs 20.33




Figure 5: UMAP projection of learned protein representations with quantified intra-class and inter-class distances.

and better inter-class separation: 52.30 vs 50.43) demonstrates MoE’s ability to selectively combine perspectives through expert routing, whereas the edge stack strategy indiscriminately superimposes all perspective-specific edges, introducing noise from redundant or conflicting connections that degrade representation quality.




Figure 6: Robustness analysis of MMPG under random residue masking (0-40%) across four protein tasks.

Robustness Analysis

We evaluate the robustness of MMPG under conditions mimicking low-resolution experimental data by randomly masking $u\%$ of input residue embeddings, reflecting scenarios where residue chemical identities are uncertain or unresolved (DiMaio et al. 2009). As shown in Figure 6, MMPG demonstrates strong resilience. Even with 40% masking, performance degradation remains modest, particularly for EC and FOLD tasks. This robustness might stem from the multi-perspective design. MMPG can reconstruct missing information by leveraging the diverse information patterns that remain visible across other perspectives.

Conclusion

This work addresses the limitations of single-perspective protein graph construction in capturing comprehensive protein information. We build physical, chemical, and geometric protein graphs for a comprehensive representational foundation. Building upon this, we design an MoE module that discovers and leverages deep interactions among perspectives, dynamically integrating them into task-specific representations. Our method achieves advanced performance across different protein downstream tasks, demonstrating its effectiveness in PRL.

Acknowledgments

This work was supported by Government Special Support Funds for the Guangdong Institute of Intelligence Science and Technology, and National Natural Science Foundation of China (NSFC) under Grant No.62506084.

References

- Ashburner, M.; Ball, C. A.; Blake, J. A.; Botstein, D.; Butler, H.; Cherry, J. M.; Davis, A. P.; Dolinski, K.; Dwight, S. S.; Eppig, J. T.; et al. 2000. Gene ontology: tool for the unification of biology. *Nature genetics*, 25(1): 25–29.
- Baldassarre, F.; Hurtado, D. M.; Elofsson, A.; and Azizpour, H. 2020. GraphQA: protein model quality assessment using graph convolutional networks. *Bioinformatics*, 37: 360–366.
- Chen, Z.; Shen, Y.; Ding, M.; Chen, Z.; Zhao, H.; Learned-Miller, E. G.; and Gan, C. 2023. Mod-Squad: Designing Mixtures of Experts As Modular Multi-Task Learners. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11828–11837.
- Chiang, Y.; Hui, W.-H.; and Chang, S.-W. 2022. Encoding protein dynamic information in graph representation for functional residue identification. *Cell Reports Physical Science*, 3(7): 100975.
- Choulli, M.; and Stefanov, P. 1996. Inverse scattering and inverse boundary value problems for the linear Boltzmann equation. *Communications in Partial Differential Equations*, 21(5-6): 763–785.
- Derevyanko, G.; Grudin, S.; Bengio, Y.; and Lamoureaux, G. 2018. Deep convolutional networks for quality assessment of protein folds. *Bioinformatics*, 34: 4046–4053.
- DiMaio, F.; Tyka, M. D.; Baker, M. L.; Chiu, W.; and Baker, D. 2009. Refinement of protein structures into low-resolution density maps using rosetta. *Journal of molecular biology*, 392(1): 181–190.
- Elnaggar, A.; Heinzinger, M.; Dallago, C.; Rehawi, G.; Wang, Y.; Jones, L.; Gibbs, T.; Fehér, T. B.; Angerer, C.; Steinegger, M.; Bhowmik, D.; and Rost, B. 2020. ProtTrans: Towards Cracking the Language of Life’s Code Through Self-Supervised Deep Learning and High Performance Computing. *bioRxiv*.
- Fan, H.; Wang, Z.; Yang, Y.; and Kankanhalli, M. 2023. Continuous-Discrete Convolution for Geometry-Sequence Modeling in Proteins. In *The Eleventh International Conference on Learning Representations*.
- Fasoulis, R.; Paliouras, G.; and Kavraki, L. E. 2021. Graph representation learning for structural proteomics. *Emerging Topics in Life Sciences*, 5: 789–802.
- Gligorijević, V.; Renfrew, P. D.; Kosciólek, T.; Leman, J. K.; Berenberg, D.; Vatanen, T.; Chandler, C.; Taylor, B. C.; Fisk, I.; Vlamakis, H.; Xavier, R. J.; Knight, R.; Cho, K.; and Bonneau, R. 2021. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12.
- Guo, P.; Correia, B.; Vandergheynst, P.; and Probst, D. 2025. Boosting Protein Graph Representations through Static-Dynamic Fusion. In *Forty-second International Conference on Machine Learning*.
- Hermosilla, P.; and Ropinski, T. 2022. Contrastive Representation Learning for 3D Protein Structures. *ArXiv*, abs/2205.15675.
- Hermosilla, P.; Schäfer, M.; Lang, M.; Fackelmann, G.; Vázquez, P.-P.; Kozlikova, B.; Krone, M.; Ritschel, T.; and Ropinski, T. 2021. Intrinsic-Extrinsic Convolution and Pooling for Learning on 3D Protein Structures. In *International Conference on Learning Representations*.
- Hoang, T. L.; Sbodio, M. L.; Martinez Galindo, M.; Zayats, M.; Fernandez-Diaz, R.; Valls, V.; Picco, G.; Berrospi, C.; and Lopez, V. 2024. Knowledge Enhanced Representation Learning for Drug Discovery. 38: 10544–10552.
- Jamasb, A.; Viñas Torné, R.; Ma, E.; Du, Y.; Harris, C.; Huang, K.; Hall, D.; Lió, P.; and Blundell, T. 2022. Graphein - a Python Library for Geometric Deep Learning and Network Analysis on Biomolecular Structures and Interaction Networks. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 27153–27167. Curran Associates, Inc.
- Jamasb, A. R.; Morehead, A.; Joshi, C. K.; Zhang, Z.; Didi, K.; Mathis, S. V.; Harris, C.; Tang, J.; Cheng, J.; Lio, P.; and Blundell, T. L. 2024. Evaluating Representation Learning on the Protein Structure Universe. In *The Twelfth International Conference on Learning Representations*.
- Janin, J.; Wodak, S.; Levitt, M.; and Maigret, B. 1978. Conformation of amino acid side-chains in proteins. *Journal of molecular biology*, 125(3): 357–386.
- Jing, B.; Eismann, S.; Suriana, P.; Townshend, R. J. L.; and Dror, R. 2021. Learning from Protein Structure with Geometric Vector Perceptrons. In *International Conference on Learning Representations*.
- Khade, P. M.; Maser, M.; Gligorijević, V.; and Watkins, A. 2023. Mixed structure- and sequence-based approach for protein graph neural networks with application to antibody developability prediction. *bioRxiv*.
- Kulmanov, M.; and Hoehndorf, R. 2019. DeepGOplus: improved protein function prediction from sequence. *Bioinformatics*, 36: 422–429.
- Kyte, J.; and Doolittle, R. F. 1982. A simple method for displaying the hydropathic character of a protein. *Journal of Molecular Biology*, 157(1): 105–132.
- Liang, W.; YU, L.; Luo, L.; Iyer, S.; Dong, N.; Zhou, C.; Ghosh, G.; Lewis, M.; tau Yih, W.; Zettlemoyer, L.; and Lin, X. V. 2025. Mixture-of-Transformers: A Sparse and Scalable Architecture for Multi-Modal Foundation Models. *Transactions on Machine Learning Research*.
- Lin, Z.; Akin, H.; Rao, R.; Hie, B.; Zhu, Z.; Lu, W.; Smetanin, N.; Verkuil, R.; Kabeli, O.; Shmueli, Y.; dos Santos Costa, A.; Fazel-Zarandi, M.; Sercu, T.; Candido, S.; and Rives, A. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637): 1123–1130.

- López-Blanco, J. R.; and Chacón, P. 2019. Korp: knowledge-based 6D potential for fast protein and loop modeling. *Bioinformatics*, 35(17): 3013–3019.
- Ma, J.; Zhao, Z.; Yi, X.; Chen, J.; Hong, L.; and Chi, E. H. 2018. Modeling Task Relationships in Multi-task Learning with Multi-gate Mixture-of-Experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, 1930–1939. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355520.
- Martin, A. C.; Orengo, C. A.; Hutchinson, E. G.; Jones, S.; Karmirantzou, M.; Laskowski, R. A.; Mitchell, J. B.; Taroni, C.; and Thornton, J. M. 1998. Protein folds and functions. *Structure*, 6(7): 875–884.
- McInnes, L.; Healy, J.; Saul, N.; and Großberger, L. 2018. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software*, 3(29).
- Miyazawa, S.; and Jernigan, R. L. 1996. Residue–residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading. *Journal of molecular biology*, 256(3): 623–644.
- Osadchy, M.; and Kolodny, R. 2011. Maps of protein structure space reveal a fundamental relationship between protein structure and function. *Proceedings of the National Academy of Sciences*, 108(30): 12301–12306.
- Quan, R.; Wang, W.; Ma, F.; Fan, H.; and Yang, Y. 2024. Clustering for Protein Representation Learning. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 319–329.
- Ribeiro, A. J. M.; Holliday, G. L.; Furnham, N.; Tyzack, J. D.; Ferris, K.; and Thornton, J. M. 2018. Mechanism and Catalytic Site Atlas (M-CSA): a database of enzyme reaction mechanisms and active sites. *Nucleic acids research*, 46(D1): D618–D623.
- Samudrala, R.; and Moult, J. 1998. An all-atom distance-dependent conditional probability discriminatory function for protein structure prediction. *Journal of molecular biology*, 275(5): 895–916.
- Sanyal, A.; Lajoie, B. R.; Jain, G.; and Dekker, J. 2012. The long-range interaction landscape of gene promoters. *Nature*, 489(7414): 109–113.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 20(1): 61–80.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- Sippl, M. J. 1995. Knowledge-based potentials for proteins. *Current opinion in structural biology*, 5(2): 229–235.
- Tanford, C. 1978. The Hydrophobic Effect and the Organization of Living Matter. *Science*, 200(4345): 1012–1018.
- Wang, L.; Liu, H.; Liu, Y.; Kurtin, J.; and Ji, S. 2023. Learning Hierarchical Protein Representations via Complete 3D Graph Networks. In *The Eleventh International Conference on Learning Representations*.
- Wang, P.; Hu, H.; Tong, B.; Zhang, Z.; Yao, F.; Feng, Y.; Zhu, Z.; Chang, H.; Diao, W.; Ye, Q.; and Sun, X. 2025a. RingMoGPT: A Unified Remote Sensing Foundation Model for Vision, Language, and Grounded Tasks. *IEEE Transactions on Geoscience and Remote Sensing*, 63: 1–20.
- Wang, Y.; Tan, S.; Shen, J.; Xu, Y.; Song, H.; Xu, Q.; Tiwari, P.; and Xu, M. 2025b. Enhancing Graph Contrastive Learning for Protein Graphs from Perspective of Invariance. In *Forty-second International Conference on Machine Learning*.
- Wang, Z.; Combs, S. A.; Brand, R.; Rebollar, M. C.; Xu, P.; Price, G. D.; Golovach, N.; Salawu, E. O.; Wise, C.; Ponnapalli, S. P.; and Clark, P. M. 2021. LM-GVP: A Generalizable Deep Learning Framework for Protein Property Prediction from Sequence and Structure. *bioRxiv*.
- Yang, K. K.; Wu, Z.; Bedbrook, C. N.; Arnold, F. H.; and Wren, J. 2018. Learned protein embeddings for machine learning. *Bioinformatics*, 34: 2642–2648.
- Zhang, Z.; Xu, M.; Jamasb, A. R.; Chenthamarakshan, V.; Lozano, A.; Das, P.; and Tang, J. 2023. Protein Representation Learning by Geometric Structure Pretraining. In *The Eleventh International Conference on Learning Representations*.