Proceedings of the Thirty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-26)
Doctoral Consortium Track

Equivariant Graph Neural Networks for Protein Interaction Modeling and Structure-Aware Molecular Design

Animesh Animesh
Department of Artificial Intelligence, IIT Kharagpur
animesh.sachan24794@kgpian.iitkgp.ac.in

Abstract

Equivariant graph neural networks (EGNNs) have revolutionized molecular property prediction by encoding 3D structures while preserving physical symmetry. However, existing EGNNs struggle to capture fine-grained geometric features needed for tasks like understanding protein interactions, and designing molecules that modulate them are central challenges in computational drug discovery. This thesis develops a progressive framework based on *equivariant graph neural networks* (EqGNNs) that respects the geometric symmetries of molecular structures. Our contributions progress along two axes: from discriminative to generative and from basic representations to physics-informed, language-conditioned generation. In this study, we propose novel methods for both discriminative (E(Q)AGNN-PPIS, GDEGAN) and generative tasks (PhysFlow). Together, these advance symmetry-aware geometric deep learning for drug discovery.

1 Introduction and Motivation

Proteins serve as the primary functional units in biological systems, with their three-dimensional structures—ranging from complex proteins to organic compounds—fundamentally determining their biological functions and being crucial for drug discovery and molecular biology applications. However, computationally capturing these geometric features from non-Euclidean molecular data remains a significant challenge for traditional machine learning approaches. Graph Neural Networks (GNNs) have emerged as a natural framework for molecular representation, with Equivariant GNNs (EGNNs) specifically designed to preserve rotational and translational symmetries critical for 3-D molecular systems [Jing *et al.*, 2020].

Despite progress, key gaps remain. First, existing EqGNNs for protein surfaces underutilize directional geometric information. Second, binding site prediction methods treat proteins as static, ignoring conformational dynamics. Third, generative molecular models often lack physics-based inductive biases, producing chemically implausible candidates. Fourth, specifying molecular design objectives remains inaccessible to domain experts without ML expertise.

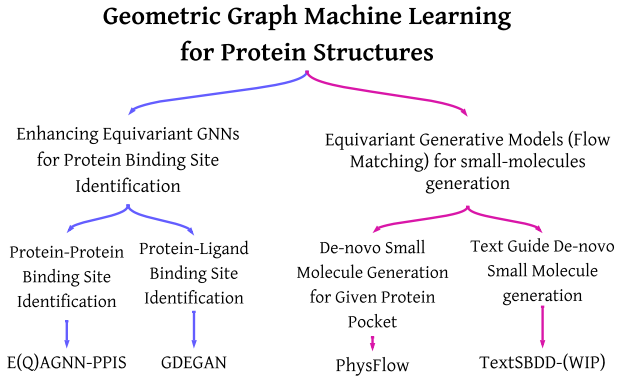

Figure 1: Taxonomy of research contributions.

This thesis addresses these gaps through the progressive research as shown in Figure 1. The overarching question is: **How can geometric symmetry constraints be leveraged across the full drug discovery pipeline, from protein surface understanding to controllable 3D molecule generation?**

2 Research Questions and Contributions

RQ1: *How can attention mechanisms within equivariant GNNs leverage directional vector features to improve protein-protein interaction site prediction?*
Contribution: E(Q)AGNN-PPIS [Animesh *et al.*, 2025]: In this work we hypothesize that incorporating geometric features of the amino acids into model architecture will enhance performance of protein-protein interaction site prediction. To validate this, in this study, we proposed E(Q)AGNN-PPIS[Animesh *et al.*, 2025], a method which takes geometric features into consideration. Our approach leverages equivariant GNNs, specifically augmenting the GVP-GNN architecture [Jing *et al.*, 2020] through the integration of an attention mechanism. Our method preserves rotational and translational equivariance while incorporating both scalar and vector features through attention-enhanced message passing. We formally prove that attention integration maintains model equivariance. Empirical evaluation demonstrates **8.33% improvement in AUPRC** (0.65) and **10% improvement in**

MCC (0.55) over the previous state-of-the-art on the standard benchmark datasets. Additionally, our method demonstrates $4\times$ faster inference and robust generalization across proteins of varying lengths, with **up to 32.15% F1-score improvement** on large proteins.

RQ2: *How can adaptive, variance-aware attention replace static dot-product attention in equivariant GNNs for ligand binding site prediction while offering superior computational efficiency and interpretability?*

While dot-product attention mechanisms have proven effective, as we have shown in our earlier work [Animesh *et al.*, 2025] in modelling protein data, their computational overhead and neighboring molecule understanding challenges motivate further exploration. In this work we built our mechanism based on the key insight that binding sites often appear as statistically distinct, tightly clustered regions compared to the rest of the protein surface.

Contribution: GDEGAN [Animesh *et al.*, 2026]: We propose GDEGAN, a Gaussian Dynamic Equivariant Graph Attention Network built on the GotenNet [Aykent and Xia, 2025] backbone with higher-degree steerable features ($L_{\max} = 2$). Unlike standard dot-product attention that applies uniform similarity metrics, GDEGAN computes attention scores by measuring how statistically probable a neighboring atom's features are, given a Gaussian distribution defined by the target atom's local neighborhood. By dynamically computing local neighborhood statistics (μ_i, σ_i^2) of each atom's local environment at every layer, our attention mechanism becomes inherently adaptive, and we use the variance as an adaptive bandwidth parameter with learnable per-head temperatures. This adaptation to the emergent local geometry of the protein graphs provides a more powerful and physically grounded inductive bias, enabling the model to learn more robust representations of complex protein structures. We prove that binding sites exhibit significantly higher learned variance ($\sigma^2 = 1.41 \pm 0.52$) than non-binding regions (0.75 ± 0.35, $p < 0.001$, Cohen's $d = 1.85$). GDEGAN achieves **37–66% relative improvement in DCC** and **7–19% in DCA** success rates over state-of-the-art methods across COACH420, HOLO4K, and PDBBind2020, while being **19.5× faster** at inference.

RQ3: *How can physics-informed distance supervision and higher-order equivariant representations address the stability-affinity trade-off in structure-based drug design?*

Contribution: PhysFlow (Under Submission): We uncover a critical limitation in existing SBDD methods: the state-of-the-art PAFlow [Zhou *et al.*, 2025] achieves strong binding affinity but only 2.7% molecular stability. We attribute this to position-level MSE losses that penalize all coordinate errors equally, ignoring chemical significance. PhysFlow replaces coordinate-level MSE with a novel *pocket-aware energy loss* that supervises pairwise interatomic distances with exponential decay weighting, combined with protein-ligand anchoring terms and clash-aware guided sampling. Using a higher-order equivariant backbone ($\ell = 2$), PhysFlow achieves **6.3× higher molecular stability** (18% vs. 2.9%), produces **6.6× more usable drug candidates** (1,450 vs. 221), and is **2.8× more parameter-efficient**, without sacrificing binding affinity on the CrossDocked2020 dataset.

RQ4: *How can natural language instructions be integrated with pocket-conditioned equivariant flow matching for controllable 3D molecular generation?*

Proposed Solution: TextSBDD (Work in progress): Existing text-guided molecular methods operate without pocket conditioning (generating SMILES, not 3D coordinates), while pocket-conditioned methods lack natural language control. We identify this as a genuine gap: no method combines text guidance + 3D pocket conditioning + 3D coordinate generation. We develop TextSBDD, which integrates text embeddings with equivariant flow matching to enable both *de novo* generation from text prompts (e.g., "increase solubility while maintaining binding") and *editing* of existing molecules, all conditioned on target pocket geometry.

3 Future Direction and Impact

Future work: Beyond RQ4 includes multi-target molecular optimization under geometric constraints and integration with experimental feedback loops. **Scientific impact:** This thesis advances geometric deep learning by demonstrating how symmetry-aware architectures can systematically address increasingly complex molecular tasks. **Practical impact:** The methods are directly applicable to early-stage drug discovery acceleration. TextSBDD (RQ4) has the potential to make generative drug design accessible to medicinal chemists without ML expertise.

Acknowledgments

I am thankful to my advisors, Prof. Plaban Kumar Bhowmick and Prof. Pralay Mitra, and to the TCS Research Fellowship (TCS Foundation) for supporting my PhD.

References

[Animesh *et al.*, 2025] Animesh, Rishi Suvvada, Plaban Kumar Bhowmick, and Pralay Mitra. E(q)agnn-ppis: Attention enhanced equivariant graph neural network for protein-protein interaction site prediction. *IEEE Transactions on Signal and Information Processing over Networks*, 11:740–751, 2025.

[Animesh *et al.*, 2026] Animesh, Plaban Kumar Bhowmick, and Pralay Mitra. GDEGAN: Gaussian dynamic equivariant graph attention network for ligand binding site prediction. *arXiv preprint arXiv:2603.19817*, 2026.

[Aykent and Xia, 2025] Sarp Aykent and Tian Xia. Gotennet: Rethinking efficient 3d equivariant graph neural networks. In *The Thirteenth International Conference on Learning Representations*, 2025.

[Jing *et al.*, 2020] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror. Learning from protein structure with geometric vector perceptrons. In *International Conference on Learning Representations*, 2020.

[Zhou *et al.*, 2025] Jingyuan Zhou, Hao Qian, Shikui Tu, and Lei Xu. Prior-guided flow matching for target-aware molecule design with learnable atom number. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.