

---

# EquiHGNN: Scalable Rotationally Equivariant Hypergraph Neural Networks

---

**Tien Dang**

Department of Computer Science  
The University of Alabama at Birmingham  
Birmingham, Alabama 35294, USA  
tien.danganh02@gmail.com

**Truong-Son Hy\***

Department of Computer Science  
The University of Alabama at Birmingham  
Birmingham, Alabama 35294, USA  
thy@uab.edu

## Abstract

Molecular interactions often involve high-order relationships that cannot be fully captured by traditional graph-based models limited to pairwise connections. Hypergraphs naturally extend graphs by enabling multi-way interactions, making them well-suited for modeling complex molecular systems. In this work, we introduce **EquiHGNN**, an **Equivariant HyperGraph Neural Network** framework that integrates symmetry-aware representations to improve molecular modeling. By enforcing the equivariance under relevant transformation groups, our approach preserves geometric and topological properties, leading to more robust and physically meaningful representations. We examine a range of equivariant architectures and demonstrate that integrating symmetry constraints leads to notable performance gains on large-scale molecular datasets. Experiments on both small and large molecules show that high-order interactions offer limited benefits for small molecules but consistently outperform 2D graphs on larger ones. Adding geometric features to these high-order structures further improves the performance, emphasizing the value of spatial information in molecular learning. Our source code is available at <https://github.com/HySonLab/EquiHGNN/>.

## 1 Introduction

Molecular systems exhibit complex, high-order interactions, including conjugated  $\pi$ -systems, hydrogen bonding networks, and ring strain effects [1, 2]. Taking advantage of the inherent benefits of graph-based representations, Graph Neural Networks (GNNs) have been widely used to model molecular interactions [3] due to their ability to efficiently learn relational structures. GNNs leverage message passing mechanisms to aggregate information from neighboring atoms, making them well suited to encode local connectivity and bond-based interactions [4, 5]. However, standard GNNs primarily model pairwise relationships between nodes, limiting their ability to represent the multi-body dependencies inherent in molecular interactions. Furthermore, they often lack explicit geometric information, such as spatial coordinates, bond angles, and torsional relationships, that is crucial for accurately capturing the three-dimensional structure and properties of molecules.

Topological Deep Learning (TDL) [6, 7, 8] offers a robust framework to overcome the limitations of GNN by integrating higher-order structures that extend beyond simple node-to-node connections. Techniques such as simplicial complexes [9], cell complexes [10], combinatorial complexes [6, 11] and hypergraphs [12, 13, 14] enrich traditional graph representations by capturing complex multi-body interactions intrinsic to molecular systems. By embedding these topological structures within deep learning architectures, we can achieve a more holistic and principled understanding of molecular interactions. This approach has already led to state-of-the-art performance in a variety of machine

---

\*Corresponding Author

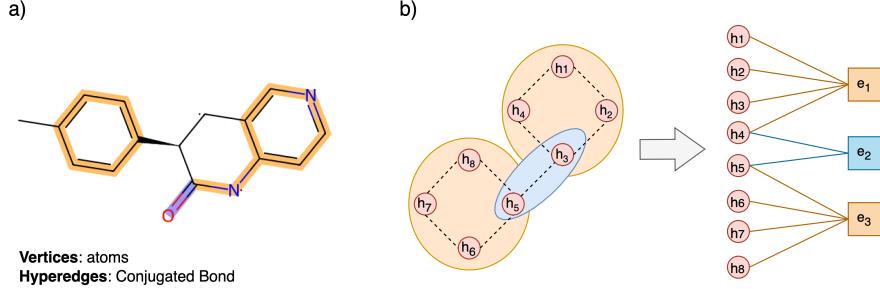
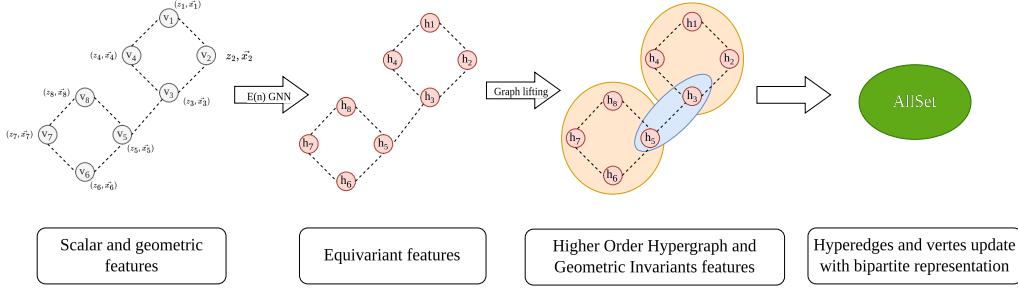


Figure 1: a) Illustration of a hypergraph constructed from a molecule, where vertices represent atoms and hyperedges represent conjugated bonds, highlighted in blue and orange. b) Hypergraph to Bipartite representations



learning tasks [12, 15, 11, 16], and holds significant promise to advance research and applications in the applied sciences and beyond.

In addition to topological considerations, molecular interactions are inherently governed by geometric constraints. Geometric Deep Learning (GDL) [17] incorporates geometric priors, such as structural and symmetry information about the input space, which are essential to maintain spatial symmetries and to ensure that the learned representations accurately reflect the physical properties of molecules. A key requirement in molecular modeling is the equivariance to transformations such as rotations, translations, and reflections. Recent advancements in equivariant architectures enable the integration of geometric constraints into deep learning models, ensuring that molecular representations align with the principles of three-dimensional spatial organization [18, 19, 20, 21, 22, 23, 24].

Although previous work has explored integrating symmetry into simplicial [25] and combinatorial complexes[26], incorporating equivariant features into hypergraphs offers a more expressive and adaptable framework for modeling complex multi-body interactions in molecular systems. Unlike other complex-based representations, hypergraphs provide a more direct and scalable approach to capturing these interactions, enabling more intuitive molecular modeling. Given these advantages, in this study we focus on hypergraphs to model high-level interactions in molecular systems.

We introduce EquiHGNN, a unified framework that seamlessly integrates topological and geometric learning to model molecular interactions while preserving geometric consistency. Instead of designing a complex message-passing framework, we initialize the hypergraph features with symmetry-aware geometric representations and invariant scalar information, enabling the model to effectively capture structural and functional properties without sacrificing equivariance. Compared to baseline models that do not incorporate symmetry into hypergraph representations, our approach encodes both scalar and geometric features within a unified framework.

We evaluated our model on QM9 [27], OPV [28], PCQM4Mv2 [29], and Molecule3D [30] to assess both accuracy and scalability. While QM9 and OPV contain small molecules, PCQM4Mv2 and Molecule3D test performance on large-scale graphs. Our model achieved competitive results across all datasets, demonstrating strong generalization and scalability to complex molecular structures.

Our contributions are as follows.

- We propose Equivariant Hypergraph Neural Network (EquiHGNN) for molecular property prediction that effectively captures both scalar and geometric features through hypergraph

representations. Its modular and intuitive design enables seamless integration with existing frameworks and allows reuse of publicly available models.

- We conducted an empirical study to analyze the impact of various architectural choices and configurations, exploring different equivariant approaches. Specifically, we examine the use of EGNN [31] in the Euclidean domain, Equiformer [32] in the Fourier domain, and FAFormer [24] in the frame domain.
- We observe that high-order interactions in hypergraphs, although slightly less effective on small molecules, consistently enhance performance on large-scale graphs. Additionally, integrating 3D geometric information into the hypergraph significantly improves the accuracy of the model.

## 2 Method

This section presents our intuitive approach to modeling high-order interactions with symmetry-aware features, emphasizing key components and detailing their seamless integration for optimal performance.

### 2.1 AllSet

The AllSet framework [13] formulates HGNNs using multiset functions, ensuring permutation invariance and expression of the message passing. It models hypergraphs as bipartite graphs, enabling flexible message propagation through two learnable set functions.

Formally, let  $\mathcal{H} = (\mathcal{V}, \mathcal{E})$  be a hypergraph, where  $\mathcal{V}$  is the set of nodes and  $\mathcal{E}$  is the set of hyperedges, each of which connects a subset of nodes. The AllSet framework updates representations through a two-step message-passing mechanism. First, hyperedge embeddings are computed by aggregating features from incident nodes using a set function  $f_{V \rightarrow E}$ , defined as follows:

$$Z_e^{(t+1)} = f_{V \rightarrow E} \left( V_{e \setminus v, X^{(t)}}; Z_{e,:}^{(t),v} \right),$$

where  $X^{(t)}$  represents the features of nodes in iteration  $t$ ,  $Z_e^{(t)}$  denotes hyperedge embeddings, and  $V_e$  is the set of nodes belonging to hyperedge  $e$ . This function aggregates node information into a hyperedge representation while preserving permutation invariance.

The node features are then updated based on the embeddings of the hyperedges through a second set function  $f_{E \rightarrow V}$ , which propagates information back from the hyperedges to the nodes:

$$X_v^{(t+1)} = f_{E \rightarrow V} \left( E_{v, Z_e^{(t+1),v}}; X_{v,:}^{(t)} \right),$$

where  $E_v$  is the set of hyperedges containing the node  $v$ . This formulation allows the message-passing process to flexibly capture complex dependencies between nodes and hyperedges.

After  $T$  steps of the message passing, the hypergraph-level prediction is calculated in the readout part on the final hidden states of hyperedges and nodes:

$$y = \text{MLP} \left( \sum_{v \in G} X_v^{(T)} \sum_{e \in G} Z_e^{(T)} \right).$$

This architecture ensures permutation invariance while allowing expressive transformations of hypergraph features.

### 2.2 Equivariant Hypergraph Neural Network

Our approach enhances AllSet by initializing node features with both scalar attributes and 3D geometric properties, naturally embedding equivariant information into the model. This simple yet effective design leads to a more expressive and symmetry-aware framework for hypergraph learning.

To explore the best strategy for symmetry-aware representation, we evaluated models across three geometric domains. Scalar-based models such as EGNN [31] preserve pairwise distances and relational structure, ensuring invariance. Frame-based methods such as FAFormer [24] apply frame

averaging to enforce equivariance while capturing higher-order geometric patterns. Fourier-based models such as Equiformer [32] transform geometric features into the spectral domain to model long-range dependencies. This comparative analysis helps identify the most effective backbone for incorporating geometric symmetry into hypergraph learning.

As illustrated in Figure 2, we first extract symmetry-aware embeddings using an equivariant model (e.g., EGNN or Equiformer). These embeddings are then used as input features in a hypergraph module that captures complex higher-order interactions. The hyperedges embedding is concatenated with the vertexes features and passed through several MLP layers to generate the final predictions. Although the pipeline remains conceptually simple, it significantly improves performance over other baselines.

### 3 Experiments

Density Functional Theory (DFT) is a widely-used quantum mechanical method for predicting molecular properties, such as structure, reactivity, and responses to electromagnetic fields. Although DFT provides high accuracy, its computational cost increases rapidly with the size of the system, making it impractical for large-scale molecular screening. As a result, DFT calculations can be time-consuming, often requiring several hours for even small molecules. This limitation makes it challenging to explore large chemical spaces or conduct extensive property predictions. In this study, we utilize the QM9 [27] and OPV [28] datasets for molecular property prediction. QM9 consists of small organic molecules and is widely used to benchmark fundamental molecular properties. The OPV dataset focuses on organic photovoltaic molecules, which typically contain larger conjugated systems relevant for electronic and optoelectronic applications. To assess the scalability and expressiveness of our model, we evaluate large-scale graph datasets such as PCQM4Mv2 [29] and Molecule3D [30], both derived from PubChemQC [33]. These datasets are designed to predict critical molecular properties for fields such as drug discovery and materials science, featuring complex molecular structures and substantial graph data. While PCQM4Mv2 is focused primarily on predicting the HOMO-LUMO gap from SMILE strings, Molecule3D emphasizes 3D molecular geometry prediction, integrating molecular properties prediction as a secondary task.

#### 3.1 Results

We perform experiments comparing our approach to the MHNN baseline [14], which represents the hypergraph using a bipartite graph structure, similar to AllSet [13]. To incorporate symmetry-aware features, we evaluated three different setups: EGNN [31], FAFormer [24], and Equiformer [32]. Furthermore, we compare performance with 2D graph models to assess whether incorporating higher-order interactions with symmetry awareness can improve the model performance. All results show significantly better performance compared to 2D graph models, including GIN [34] and GAT [35], highlighting the importance of capturing geometric and topological information. The tables below report the MAE scores across all datasets, with bold values indicating the best-performing models and underlined values denoting the second-best. For the PCQM4Mv2 and Molecule3D datasets, only EGNN-MHNN is included among the equivariant integration models, as training FaFormer and Equiformer on these large-scale datasets is prohibitively time-consuming. In particular, in PCQM4Mv2, 3D geometric information is available only for the training set; therefore, all experiments are conducted on this subset to ensure a fair comparison between 2D and 3D graph representations.

##### 3.1.1 QM9 dataset

Table 1 presents the Mean Absolute Error (MAE) for six molecular properties from the QM9 dataset. The baseline MHNN, which captures higher-order interactions through hypergraph representations, underperforms compared to standard 2D graph-based models. For example, on  $\epsilon_{\text{HOMO}}$ , MHNN yields an MAE of 55.38 meV, while GIN and GAT achieve 47.67 meV and 51.56 meV, respectively. GIN also records the lowest error on  $\mu$  (0.2 meV). These results suggest that modeling higher-order relations alone does not improve the performance of small molecules.

In contrast, incorporating geometric inductive biases leads to significant gains. Both EGNN-MHNN and FAFormer-MHNN consistently reduce MAE across tasks. FAFormer-MHNN achieves the best results in three of the six properties, including  $\epsilon_{\text{LUMO}}$  (51.9 meV),  $\Delta\epsilon$  (73.3 meV), and  $\langle R^2 \rangle$  (2602.89 meV).

Table 1: MAE on the QM9 test set.

Task Units ( $\downarrow$ )	$\mu$ D	$\alpha$ $a_0^3$	$\epsilon_{\text{HOMO}}$ meV	$\epsilon_{\text{LUMO}}$ meV	$\Delta\epsilon$ meV	$\langle R^2 \rangle$ $a_0^2$
GIN	<b>0.2 ± 0.003</b>	4.09 ± 0.04	47.67 ± 0.4	99.62 ± 0.9	147.87 ± 1.3	6279.8 ± 86.91
GAT	0.65 ± 0.006	6.17 ± 0.07	51.56 ± 0.5	111.24 ± 0.9	158.26 ± 1.1	8772.5 ± 97.43
MHNN	0.67 ± 0.005	9.29 ± 0.1	55.38 ± 0.5	124.23 ± 1.07	166.6 ± 1.3	9301.44 ± 138.369
EGNN-MHNN	0.59 ± 0.005	<b>2.02 ± 0.02</b>	44.82 ± 0.3	92.54 ± 0.9	<u>140.06 ± 1.5</u>	<u>4293.09 ± 66.44</u>
FAFormer-MHNN	<u>0.3 ± 0.003</u>	4.85 ± 0.04	<u>26.47 ± 0.2</u>	<b>51.9 ± 0.4</b>	<b>73.3 ± 0.7</b>	<b>2602.89 ± 25.04</b>
Equiformer-MHNN	0.34 ± 0.003	<u>2.48 ± 0.02</u>	<b>25.57 ± 0.2</b>	<u>67.19 ± 0.5</u>	230.77 ± 1.8	102815 ± 445.433

$a_0^2$ ), halving the error compared to the baselines of the MHNN and 2D graph. EGNN-MHNN obtains the lowest MAE in  $\alpha$  ( $2.02 a_0^3$ ) and competitive results on  $\langle R^2 \rangle$  ( $4293.09 a_0^2$ ). In particular, Equiformer-MHNN achieves the best performance in  $\epsilon_{\text{HOMO}}$  (25.57 meV). These findings highlight that while higher-order modeling alone is insufficient, combining topological and geometric priors yields a more powerful and accurate framework for molecular property prediction in the QM9 setting.

### 3.1.2 OPV dataset

Table 2: MAE on the OPV test set.

Methods Units ( $\downarrow$ )	Molecular				Polymer			
	$\Delta\epsilon$ meV	$\epsilon_{\text{HOMO}}$ meV	$\epsilon_{\text{LUMO}}$ meV	$I_{\text{overlap}}$ W/mol	$\Delta\epsilon$ meV	$\epsilon_{\text{HOMO}}$ meV	$\epsilon_{\text{LUMO}}$ meV	$O_{\text{LUMO}}$ meV
GIN	50.45 ± 0.9	39.16 ± 0.5	53.29 ± 0.8	206.53 ± 3.6	53.69 ± 1.0	61.65 ± 0.9	78.48 ± 1.5	64.64 ± 0.6
GAT	55.8 ± 0.9	32.2 ± 0.5	46.68 ± 0.7	204.03 ± 4.2	47.91 ± 0.83	58.47 ± 0.92	71.84 ± 1.3	56.61 ± 0.7
MHNN	34.02 ± 0.4	26.21 ± 0.4	24.46 ± 0.3	139.58 ± 2.3	48.95 ± 1.1	<b>49.93 ± 0.8</b>	<u>60.71 ± 1.1</u>	<u>48.41 ± 0.7</u>
EGNN-MHNN	<b>28.27 ± 0.3</b>	20.97 ± 0.2	20.03 ± 0.3	<b>99.7 ± 1.5</b>	<b>45.63 ± 0.9</b>	66.67 ± 1.1	69.32 ± 1.2	67.28 ± 0.9
FAFormer-MHNN	36.4 ± 0.6	<u>20.5 ± 0.2</u>	<b>18.84 ± 0.3</b>	<u>100.52 ± 1.5</u>	46.12 ± 0.9	54.85 ± 1.0	72.05 ± 1.4	52.74 ± 0.8
Equiformer-MHNN	<b>28.12 ± 0.4</b>	<b>20.24 ± 0.2</b>	20.59 ± 0.3	107.346 ± 1.7	<b>45.42 ± 0.9</b>	<u>50.08 ± 0.8</u>	<b>58.17 ± 1.0</b>	<b>43.6 ± 0.6</b>

The OPV dataset includes both small (molecular) and large (polymer) compounds, offering a robust benchmark for evaluating the scalability and generalization of various graph-based representations. Table 2 highlights the performance (in MAE) of several models in four tasks for each category.

For small molecules, the proposed Equiformer-MHNN outperforms all baselines in two out of four tasks. It achieves the lowest MAE for  $\Delta\epsilon$  (28.12 meV) and  $\epsilon_{\text{HOMO}}$  (20.24 meV). FaFormer-MHNN performs best in  $\epsilon_{\text{LUMO}}$  with 18.84 meV, while EGNN-MHNN leads in  $I_{\text{overlap}}$ . When applied to molecular tasks, incorporating symmetry awareness into the hypergraph consistently outperforms 2D graphs and models that rely solely on higher-order interactions.

In contrast, for larger polymer molecules, Equiformer-MHNN continues to demonstrate strong performance, outperforming all other models in three of four tasks. It achieves the best MAE for  $\Delta\epsilon$  (45.42 meV),  $\epsilon_{\text{LUMO}}$  (58.17 meV) and  $O_{\text{LUMO}}$  (43.6 meV), with a close second for  $\epsilon_{\text{HOMO}}$  (50.08 meV). Although EGNN-MHNN and FAFormer-MHNN do not surpass Equiformer-MHNN, they still significantly outperform traditional 2D GNNs across all polymer-related tasks, highlighting the importance of incorporating geometric and equivariant representations when modeling complex macromolecules.

The baseline MHNN, which models high-order interactions without geometric inductive bias, achieves a moderate performance boost over GIN and GAT for small molecules (e.g., 24.46 meV in  $\epsilon_{\text{LUMO}}$  vs 53.29 and 46.68 meV for GIN and GAT, respectively). However, its improvements diminish in the polymer regime, where long-range dependencies and complex geometry require more expressive representations.

### 3.1.3 PCQM4Mv2

Table 3 reports the MAE for the prediction of the HOMO-LUMO gap in the PCQM4Mv2 test set. The baseline models, GIN and GAT, achieve errors of 117.65 meV and 116.93 meV, respectively. MHNN reduces the error to 108.11 meV by incorporating high-order molecular interactions through hypergraph representations. EGNN-MHNN further improves the performance to 98.45 meV by

Table 3: MAE on the PCQM4Mv2 test set in meV.

Model	gap ( $\downarrow$ )
GIN	$117.65 \pm 0.23$
GAT	$116.93 \pm 0.21$
MHNN	$108.11 \pm 0.25$
EGNN-MHNN	<b><math>98.45 \pm 0.2</math></b>

Table 4: MAE on the Molecule3D test set in meV.

Model	gap ( $\downarrow$ )
GIN	$129.61 \pm 0.25$
GAT	$137.22 \pm 0.36$
MHNN	<b><math>117.55 \pm 0.28</math></b>
EGNN-MHNN	<u><math>122.25 \pm 0.24</math></u>

integrating 3D geometric features with equivariant message passing. These results indicate that geometric inductive biases and spatial information are beneficial for learning quantum chemical properties in large-scale molecular graphs.

### 3.1.4 Molecule3D

Table 4 shows the MAE results in the Molecule3D test set. GIN and GAT yield errors of 129.61 meV and 137.22 meV, respectively. MHNN achieves the lowest error of 117.55 meV, demonstrating the effectiveness of high-order interaction modeling to capture structural dependencies. EGNN-MHNN obtains a slightly higher error of 122.25 meV, suggesting that the added geometric modeling does not consistently improve performance in datasets with high conformational flexibility.

## 4 Conclusion

In this work, we have explored the integration of symmetry-aware features into hypergraph representations for molecular modeling. Our approach focuses on a simple yet effective strategy of preparing node features through an embedding that combines both symmetry-aware geometric representations and invariant scalar information. We experimented with several equivariant techniques, including spatial domain modeling, frame averaging, and Fourier domain methods. Although we also attempted to directly modify the hypergraph message passing framework, this approach proved to be not only challenging but also yielded poor results, highlighting the elegance and practicality of our embedding-based strategy. Our results demonstrate that hypergraphs consistently outperform pairwise graph approaches on larger molecules, showcasing their scalability. Additionally, incorporating symmetry-aware features significantly improves model performance, emphasizing the importance of capturing both high-order interactions and geometric consistency. Overall, our model demonstrates strong generalization capabilities to large graphs while maintaining a simple and robust framework.

## 5 Limitations and Future work

In molecular systems, various higher-order interactions beyond conjugated bonds are important, such as those explored by ETNNs [26]. Furthermore, SE3Set [36] employs a fragmentation method using a BFS-like algorithm to identify connected subgraphs, which could be valuable for capturing complex interactions. Our experiments with using rings as higher-order interactions resulted in significantly worse performance compared to conjugated bonds, highlighting the need for further exploration of effective high-order interactions. Moreover, while several state-of-the-art equivariant models outperform Equiformer on certain tasks with better computational and parameter efficiency, our work demonstrates a simple and flexible framework where node features for hypergraph neural networks are initialized with embeddings from equivariant models. This design offers a straightforward plug-and-play capability, allowing integration with other more advanced equivariant architectures to further

enhance performance and scalability. Future work will explore the incorporation of these models to unlock better efficiency and generalization.

## References

- [1] Laurianne David, Amol Thakkar, Rocío Mercado, and Ola Engkvist. Molecular representations in ai-driven drug discovery: a review and practical guide. *Journal of Cheminformatics*, 12(1):56, 2020.
- [2] Daniel S Wigh, Jonathan M Goodman, and Alexei A Lapkin. A review of molecular representation in the age of machine learning. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 12(5):e1603, 2022.
- [3] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.
- [4] Marinka Zitnik, Monica Agrawal, and Jure Leskovec. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13):i457–i466, 2018.
- [5] Vladimir Gligorijević, P Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C Taylor, Ian M Fisk, Hera Vlamakis, et al. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 12(1):3168, 2021.
- [6] Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, Karthikeyan Natesan Ramamurthy, Tolga Birdal, Tamal K Dey, Soham Mukherjee, Shreyas N Samaga, et al. Topological deep learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2022.
- [7] Theodore Papamarkou, Tolga Birdal, Michael M. Bronstein, Gunnar E. Carlsson, Justin Curry, Yue Gao, Mustafa Hajij, Roland Kwitt, Pietro Lio, Paolo Di Lorenzo, Vasileios Maroulas, Nina Miolane, Farzana Nasrin, Karthikeyan Natesan Ramamurthy, Bastian Rieck, Simone Scardapane, Michael T Schaub, Petar Veličković, Bei Wang, Yusu Wang, Guowei Wei, and Ghada Zamzmi. Position: Topological deep learning is the new frontier for relational learning. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 39529–39555. PMLR, 21–27 Jul 2024.
- [8] Mathilde Papillon, Sophia Sanborn, Mustafa Hajij, and Nina Miolane. Architectures of topological deep learning: A survey of message-passing topological neural networks. *arXiv preprint arXiv:2304.10031*, 2023.
- [9] Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lio, and Michael Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In *International Conference on Machine Learning*, pages 1026–1037. PMLR, 2021.
- [10] Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Lio, Guido F Montufar, and Michael Bronstein. Weisfeiler and lehman go cellular: Cw networks. *Advances in neural information processing systems*, 34:2625–2640, 2021.
- [11] Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, and Karthikeyan Natesan Ramamurthy. Higher-order attention networks. *arXiv preprint arXiv:2206.00606*, 2(3):4, 2022.
- [12] Yihe Dong, Will Sawin, and Yoshua Bengio. Hnnh: Hypergraph networks with hyperedge neurons. *ICML Graph Representation Learning and Beyond Workshop*, 2020.
- [13] Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenkovic. You are allset: A multiset function framework for hypergraph neural networks. *arXiv preprint arXiv:2106.13264*, 2021.

- [14] Junwu Chen and Philippe Schwaller. Molecular hypergraph neural networks. *The Journal of Chemical Physics*, 160(14):144307, 2024.
- [15] Sergio Barbarossa and Stefania Sardellitti. Topological signal processing over simplicial complexes. *IEEE Transactions on Signal Processing*, 68:2992–3007, 2020.
- [16] Yuzhou Chen, Yulia R Gel, and H Vincent Poor. Bscnets: Block simplicial complex neural networks. In *Proceedings of the aaai conference on artificial intelligence*, volume 36, pages 6333–6341, 2022.
- [17] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- [18] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *Advances in neural information processing systems*, 30, 2017.
- [19] Brandon Anderson, Truong Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [20] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):2453, 2022.
- [21] Shuo Zhang, Yang Liu, and Lei Xie. A universal framework for accurate and efficient geometric deep learning of molecular systems. *Scientific Reports*, 13(1):19171, 2023.
- [22] Yusong Wang, Tong Wang, Shaoning Li, Xinheng He, Mingyu Li, Zun Wang, Nanning Zheng, Bin Shao, and Tie-Yan Liu. Enhancing geometric representations for molecules with equivariant vector-scalar interactive message passing. *Nature Communications*, 15(1):313, 2024.
- [23] Yi-Lun Liao, Brandon Wood, Abhishek Das\*, and Tess Smidt\*. EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations. In *International Conference on Learning Representations (ICLR)*, 2024.
- [24] Tinglin Huang, Zhenqiao Song, Rex Ying, and Wengong Jin. Protein-nucleic acid complex modeling with frame averaging transformer. In *Neural Information Processing Systems*, 2024.
- [25] Floor Eijkelboom, Rob Hesselink, and Erik J Bekkers. E ( $n$ ) equivariant message passing simplicial networks. In *International Conference on Machine Learning*, pages 9071–9081. PMLR, 2023.
- [26] Claudio Battiloro, Mauricio Tec, George Dasoulas, Michelle Audirac, Francesca Dominici, et al. E ( $n$ ) equivariant topological neural networks. *arXiv preprint arXiv:2405.15429*, 2024.
- [27] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- [28] Peter C St John, Caleb Phillips, Travis W Kemper, A Nolan Wilson, Yanfei Guan, Michael F Crowley, Mark R Nimlos, and Ross E Larsen. Message-passing neural networks for high-throughput polymer screening. *The Journal of chemical physics*, 150(23), 2019.
- [29] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021.
- [30] Zhao Xu, Youzhi Luo, Xuan Zhang, Xinyi Xu, Yaochen Xie, Meng Liu, Kaleb Dickerson, Cheng Deng, Maho Nakata, and Shuiwang Ji. Molecule3d: A benchmark for predicting 3d geometries from molecular graphs. *arXiv preprint arXiv:2110.01717*, 2021.

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

[32] Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. In *International Conference on Learning Representations*, 2023.

[33] Maho Nakata and Tomomi Shimazaki. Pubchemqc project: a large-scale first-principles electronic structure database for data-driven chemistry. *Journal of chemical information and modeling*, 57(6):1300–1308, 2017.

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

[35] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.

[36] Hongfei Wu, Lijun Wu, Guoqing Liu, Zhirong Liu, Bin Shao, and Zun Wang. Se3set: Harnessing equivariant hypergraph neural networks for molecular representation learning. *arXiv preprint arXiv:2405.16511*, 2024.

[37] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI open*, 1:57–81, 2020.

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

[39] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[40] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations (ICLR)*, 2020.

[41] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. *Advances in Neural Information Processing Systems*, 34:6790–6802, 2021.

[42] Truong Son Hy. Graph representation learning, deep generative models on graphs, group equivariant molecular neural networks and multiresolution machine learning. page 366.

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

[44] Alexandre Agm Duval, Victor Schmidt, Alex Hernández-García, Santiago Miret, Fragkiskos D Malliaros, Yoshua Bengio, and David Rolnick. Faenet: Frame averaging equivariant gnn for materials modeling. In *International Conference on Machine Learning*, pages 9013–9033. PMLR, 2023.

[45] Taco S. Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical CNNs. In *International Conference on Learning Representations*, 2018.

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

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

[48] Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J Bekkers, and Max Welling. Geometric and physical quantities improve e(3) equivariant message passing. In *International Conference on Learning Representations*, 2022.

- [49] Chaitanya K Joshi, Cristian Bodnar, Simon V Mathis, Taco Cohen, and Pietro Lio. On the expressive power of geometric graph neural networks. In *International conference on machine learning*, pages 15330–15355. PMLR, 2023.
- [50] Truong Son Hy, Shubhendu Trivedi, Horace Pan, Brandon M. Anderson, and Risi Kondor. Predicting molecular properties with covariant compositional networks. *The Journal of Chemical Physics*, 148(24):241745, 06 2018.
- [51] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3558–3565, 2019.
- [52] Cong Liu, David Ruhe, Floor Eijkelboom, and Patrick Forré. Clifford group equivariant simplicial message passing networks. *arXiv preprint arXiv:2402.10011*, 2024.
- [53] Kevin A Murgas, Emil Saucan, and Romeil Sandhu. Hypergraph geometry reflects higher-order dynamics in protein interaction networks. *Scientific Reports*, 12(1):20879, 2022.
- [54] Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical science*, 11(12):3316–3325, 2020.
- [55] Justin Airas and Bin Zhang. Scaling graph neural networks to large proteins. *Journal of Chemical Theory and Computation*, 2024.
- [56] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

## A Related Work

### A.1 Graph Neural Network

Graph Neural Networks (GNNs) have been widely developed to improve representation learning in graph-structured data, enabling effective modeling of relational and structural information [37]. GraphSAGE [38] introduced an inductive framework that aggregates information from the local neighborhood of a node, allowing generalization to unseen graphs. More general, Message Passing Neural Networks (MPNNs) [3] are a foundational class of GNNs that iteratively update node representations by aggregating and transforming information from their neighbors, enabling effective learning on graph-structured data. Graph Convolutional Networks (GCN) [39] leveraged spectral graph theory to perform efficient message passing through neighborhood-based feature propagation. Graph Isomorphism Networks (GIN) [34] maximized expressive power by using sum aggregation, making them as discriminative as the Weisfeiler-Lehman test for graph isomorphism. Graph Attention Networks (GAT) [35] integrated attention mechanisms to dynamically weight neighbor contributions, improving the model’s ability to capture important structural dependencies. However, these models operate on discrete graph structures, lack inherent equivariance or invariance to geometric transformations such as rotations, translations, and reflections, and can only model pairwise interactions, making them impractical for capturing high-order molecular interactions such as conjugate bonds. This work addresses these limitations by proposing EquiHGNN, a framework that incorporates both geometric awareness and higher-order interaction modeling.

### A.2 Geometric Graph Neural Networks

Geometric graph [17] is a special kind of graph with geometric information, e.g. the positions of the atoms in 3D coordinates, encapsulating rich directional information that depicts the geometry of the system, making the system ineffectively processed by GNNs. Researchers propped a variety of Geometric Graph Neural Networks quipped with invariant/equivariant properties to better characterize the geometry of geometric graph.

Many tasks require models to be invariant under Euclidean transformations, which is often achieved by converting equivariant coordinates into invariant scalars. Early works like Cormorant [19] introduced the idea of using covariant tensorial representations for molecular graphs, ensuring that the learned features transform predictably under rotations and translations. Using spherical harmonics and tensor contractions, Cormorant demonstrated how symmetry-preserving architectures can substantially improve molecular property predictions. SchNet [18] uses continuous filter convolutions with filter weights conditioned on relative distances but lacks directional encoding. DimeNet [40] addresses this by introducing directional message passing, incorporating both distances and angles between adjacent edges. GemNet [41] extends this further by incorporating dihedral angles, enabling more expressive two-hop directional message passing based on quadruplets of nodes.

Equivariant graph neural networks, on the contrary, simultaneously update invariant and equivariant features, as many tasks require equivariant output [42]. EGNN [31], a well-known scalarization-based model, constrains messages to invariant distances and multiplies them by relative coordinates to ensure equivariant updates. Frame Averaging (FA) [43, 44] ensures equivariance by encoding coordinates in multiple reference frames and averaging their representations. Since summing over all group elements is computationally difficult, FA selects a representative subset using a frame function [43]. This method has been further explored in material design, offering a scalable alternative to traditional equivariant architectures [44]. FAFormer [24] incorporates the Transformer with frame averaging within each layer, offers superior performance in the prediction of contact maps and the detection of aptamers.

Spherical harmonics-based models use functions derived from spherical harmonics and irreducible representations, leveraging tensor product operations to ensure equivariant data transformations [45, 19]. Tensor Field Network (TFN) [46] and NequIP[20] utilize equivariant graph convolutions with linear messages derived from tensor products, with NequIP further enhancing this approach using equivariant gate activations. The SE(3)-Transformer [47] extends SEGNN [48] by replacing equivariant gate activations with equivariant dot product attention for dynamic interaction weighting, while Equiformer [32] further enhances it with MLP-based attention, equivariant layer normalization, and regularizations such as dropout and stochastic depth.

As shown in [49], rotationally equivariant GNNs are more expressive than invariant GNNs, especially for sparse geometric graphs. In this work, we focus on equivariant methods, specifically EGNN [31] for the scalarization-based approach, FAFormer [24] for the frame-averaging-based approach, and Equiformer [32] for the spherical harmonic-based approach.

### A.3 Topological Deep Learning

Topological Deep Learning (TDL) [6, 7] extends beyond traditional graphs by leveraging higher-order structures, enabling a more expressive framework for modeling complex interactions among multiple entities simultaneously. Beyond rotational symmetry, molecular graphs also exhibit rich permutation symmetries in their relational structure. Predicting molecular properties with Covariant Compositional Networks (CCNs) [50] proposed a framework that preserves higher-order permutation equivariance during message passing by modeling interactions as higher-order tensors. This approach enables the network to learn more expressive and physically meaningful representations compared to first-order (pairwise) GNNs, especially when modeling complex molecular systems with many-body interactions. The Weisfeiler-Lehman graph isomorphism test has been extended to simplicial and regular cell complexes [9, 10], providing a theoretical foundation for higher-order graph structures. HGNN [51] introduced a spectral-based framework that utilizes the Laplacian hypergraph to pass messages across hyperedges. To improve flexibility, AllSet [13] proposed a more general approach that models hypergraphs as multi-sets, employing learnable permutation-invariant set functions for adaptive message aggregation. Compared to baseline pretrained GNNs, MHNN [14] takes advantage of the hypergraph to achieve better performance under limited training data. CCNN [6] further advances this direction by introducing Combinatorial Complexes, which capture hierarchical order and enable structured dependencies across multiple levels. A comprehensive review of these advances can be found in [8].

Despite these developments, there has been limited work that incorporates symmetry with topological structures. Recent efforts have explored simplicial complexes with equivariant message passing [25, 52], integrating symmetry-aware mechanisms into higher-order networks [52]. Additionally, ETNN [26] extends equivariant message passing to combinatorial complexes, providing a more structured approach to learning equivariant representations in topological deep learning.

Hypergraphs offer a powerful framework for modeling higher-order interactions, particularly in domains such as molecular modeling and complex systems. However, equivariant hypergraph neural networks remain largely unexplored. This work introduces a novel hypergraph-equivariant framework that inherits the scalability of graph-based methods, making it suitable for large molecular systems, while also incorporating equivariant geometric features to enhance expressive power and robustness.

## B Background

### B.1 Message Passing

Message Passing Neural Networks (MPNNs) play a fundamental role in learning node representations by propagating information along graph edges. As inherently permutation-invariant architectures, MPNNs effectively capture relational structures in graph-structured data, making them particularly well-suited for applications such as molecular modeling [3].

Given a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with nodes  $v_i \in \mathcal{V}$  and edges  $e_{ij} \in \mathcal{E}$ , each node  $v_i$  is associated with a feature vector  $\mathbf{h}_i \in \mathbb{R}^{c_n}$  and each edge  $e_{ij}$  with a feature vector  $\mathbf{a}_{ij} \in \mathbb{R}^{c_e}$ , where  $c_n$  and  $c_e$  represent the dimensionalities of the node and edge features, respectively. The nodes representation are iteratively updated by:

$$\begin{aligned} \mathbf{m}_{ij} &= \phi_e(\mathbf{h}_i^l, \mathbf{h}_j^l, \mathbf{a}_{ij}) \\ \mathbf{m}_i &= \text{AGGREGATE}\left(\{\mathbf{m}_{ij}\}_{j \in \mathcal{N}(i)}\right) \\ \mathbf{h}_i^{l+1} &= \phi_u(\mathbf{h}_i^l, \mathbf{m}_i) \end{aligned} \tag{1}$$

where  $\mathcal{N}(i)$  denotes the set of neighbors of node  $v_i$ , and the AGGREGATE function is a permutation-invariant operation over the neighbors (e.g., summation). The functions  $\phi_m$  and  $\phi_e$  are the message computation and the feature update function, respectively, typically parameterized by multilayer

Table 5: Model Architecture Overview

Model	No. parameters	No. layers	No. attention heads	Hidden dimension
GIN	965K	2	-	256
GAT	3.7M	2	4	256
MHNN	2.5M	2	-	256
EGNN	844K	2	8	256
FAFormer	3.2M	2	2	256
Equiformer	13M	1	1	256

perceptrons (MLP). To obtain the final graph representation, a permutation-invariant aggregator is applied to the final hidden states of all nodes. However, this operation itself does not inherently preserve the  $E(3)$  equivariance.

## B.2 Equivariance

Equivariance is a property of functions that ensures that information is preserved under transformations from a group  $G$ . A function  $\phi : \mathcal{X} \rightarrow \mathcal{Y}$  is equivariant with respect to  $G$  if for all  $g \in G$  and  $x \in \mathcal{X}$ , the following holds:

$$\phi(g \cdot x) = g \cdot \phi(x).$$

This property ensures that the output transforms in the same way as the input under transformations like rotation, translation, or scaling. In molecular modeling, the equivariance with these symmetries (e.g., the rotation group  $SO(3)$ ) is crucial because molecules retain their physical properties regardless of their orientation or position.

## B.3 Hypergraph

A *hypergraph* is a generalization of a graph where an edge, called a *hyperedge*, can connect any number of nodes. Formally, a hypergraph is defined as:

$$\mathcal{H} = (\mathcal{V}, \mathcal{E}),$$

where  $\mathcal{V}$  is the set of nodes and  $\mathcal{E} \subseteq 2^{\mathcal{V}} \setminus \{\emptyset\}$  is the set of hyperedges, each representing a subset of nodes. Unlike traditional graphs that capture pairwise relationships, hypergraphs model higher-order interactions, making them ideal for complex systems such as protein interactions [53] and chemical reactions [54].

Hypergraph-based models extend GNNs to process hyperedges, allowing better representation of structured data. In molecular chemistry, hypergraphs effectively model *conjugated structures*, where delocalized electrons form multi-atomic interactions crucial for optoelectronic properties. Recent studies demonstrate that HGNNs outperform traditional GNNs and 3D-based models in predicting molecular properties, offering a powerful approach for data-scarce applications like organic semiconductor design.

## C Model details

Table 5 outlines the architectural configurations of the models used in our experiments, including the number of parameters, layers, attention heads, and hidden dimensions.

## D Datasets

**QM9** [27] dataset is a widely used reference for the prediction of chemical properties. It comprises approximately 134,000 small organic molecules, each containing up to 29 atoms. The data set includes five atomic species including hydrogen, carbon, oxygen, nitrogen, and fluorine, structured as molecular graphs where the atoms are connected by four types of chemical bonds: single, double,

Table 6: Overview of the datasets

Dataset	Graphs	Task type	Task number	Metric
QM9	134K	regression	12	MAE
OPV	91K	regression	8	MAE
PCQM4Mv2	3.7M	regression	1	MAE
Molecule3D	3.9M	regression	6	MAE

triple, and aromatic. In addition, the 3D coordinates of each atom, measured in angstroms, are provided.

**Organic Photovoltaic (OPV)** [28] comprises 90,823 unique molecules, providing their SMILES representations, 3D geometries, and optoelectronic properties computed by DFT calculations. It includes four molecular-level tasks for monomers: the highest occupied molecular orbital ( $\varepsilon_{\text{HOMO}}$ ), the lowest unoccupied molecular orbital ( $\varepsilon_{\text{LUMO}}$ ), the HOMO-LUMO gap ( $\Delta\varepsilon$ ), and the spectral overlap ( $I_{\text{overlap}}$ ). Furthermore, OPV features four polymer-level tasks: polymer  $\varepsilon_{\text{HOMO}}$ , polymer  $\varepsilon_{\text{LUMO}}$ , the polymer energy gap ( $\Delta\varepsilon$ ), and the optical LUMO ( $O_{\text{LUMO}}$ ).

**PCQM4Mv2** [29] is a large-scale quantum chemistry dataset consisting of approximately 3.7 million molecular graphs, derived from the PubChemQC project [33]. It is designed for predicting the DFT-calculated HOMO-LUMO energy gap from SMILES representations, and additionally provides 3D structures for the training molecules.

**Molecule3D** [30] is a large-scale benchmark designed to predict 3D molecular geometries from 2D molecular graphs and to assess their effectiveness in downstream prediction of quantum properties. It supports two main tasks: (1) prediction of DFT-optimized 3D atomic coordinates from SMILES strings or molecular graphs, and (2) prediction of quantum properties such as total energy, HOMO/LUMO energies, and the HOMO-LUMO gap using either ground truth or predicted 3D structures. Each sample includes a SMILES string, molecular graph, 3D coordinates, and quantum properties sourced from PubChemQC [33]. In this work, we focus specifically on the prediction of the HOMO-LUMO gap.

Table 6 provides an overview of the experimental dataset. We use RDKit to identify conjugated bonds, which serve as hyperedges, with atoms as vertices, as illustrated in Section 1a. For all experiments, the data are split into training, validation and test sets using an 80-10-10 ratio. The model is trained in the training set, the best model is selected based on the performance in the validation set, and the final evaluation is performed in the test set.

## E Training details

Equivariant models utilize radial distances, where a larger radius enables the capture of high-level features crucial for complex molecules such as polymers and proteins. In such molecules, long-range interactions, such as electrostatic and hydrophobic effects, are key to determining their electronic and structural properties. A study on scaling GNNs [55] shows that increasing the number of message-passing layers and the cutoff radius helps GNNs incorporate distant atomic interactions, thereby enhancing expressiveness for large proteins.

However, in transformer-based architectures, a larger radius significantly increases computational costs due to the quadratic scaling of the attention mechanism with the number of nodes. Based on these empirical insights, we adopt a consistent configuration with a 5 Angstrom radius cutoff and 16 neighboring nodes for EGN, FAFormer, and Equiformer, achieving an optimal balance between expressiveness and computational efficiency. Appendix C provides a summary of the model architecture configurations.

We train the models for 400 epochs with a batch size of 16, using the Adam optimizer with a fixed learning rate of  $1 \times 10^{-4}$ . Training is carried out on 2xRTX 3060 GPUs, enabling parallel processing for efficiency. The models are optimized to minimize the loss of MSE, and the checkpoint with the lowest MAE validation is selected for the final evaluation on the test set. Our implementation is built using PyTorch Geometric [56].