

SE3Set: HARNESSING EQUIVARIANT HYPERGRAPH NEURAL NETWORKS FOR MOLECULAR REPRESENTATION LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

In this paper, we develop SE3Set, an SE(3) equivariant hypergraph neural network architecture tailored for advanced molecular representation learning. Hypergraphs are not merely an extension of traditional graphs; they are pivotal for modeling high-order relationships, a capability that conventional equivariant graph-based methods lack due to their inherent limitations in representing intricate many-body interactions. To achieve this, we first construct hypergraphs via proposing a new fragmentation method that considers both chemical and three-dimensional spatial information of molecular system. We then design SE3Set, which incorporates equivariance into the hypergraph neural network. This ensures that the learned molecular representations are invariant to spatial transformations, thereby providing robustness essential for accurate prediction of molecular properties. SE3Set has shown performance on par with state-of-the-art (SOTA) models for small molecule datasets like QM9 and MD17. It excels on the MD22 dataset, achieving a notable improvement of approximately 20% in accuracy across all molecules, which highlights the prevalence of complex many-body interactions in larger molecules. This exceptional performance of SE3Set across diverse molecular structures underscores its transformative potential in computational chemistry, offering a route to more accurate and physically nuanced modeling.

1 INTRODUCTION

Molecular representation (Mathews & Chaffee, 2012; David et al., 2020; Wigh et al., 2022) is pivotal for cheminformatics (Fourches et al., 2010), impacting the prediction of molecular properties in drug discovery and material science. Traditional descriptors like fingerprints capture basic structural and energetic aspects of molecules by considering mainly one- and two-body interactions. However, they often miss complex electronic correlations and collective behaviors important for understanding phenomena such as chemical reactivity and protein folding. To address this, advanced methods that include many-body interactions are crucial for a more comprehensive molecular characterization. These methods enhance the predictive capabilities of computational models by more accurately reflecting the intricate dynamics and properties of molecules, which are essential for a deeper understanding of their functionality and reactivity in cheminformatics.

Graph neural networks (GNNs) (Zhou et al., 2020; Wu et al., 2020) a foundational tool for representing structured data in molecular sciences with atoms as nodes and chemical bonds as edges, respectively. GNN models excel in tasks ranging from property prediction to reaction simulation (Do et al., 2019; Xiong et al., 2021; Reiser et al., 2022). GNNs can capture higher-order molecular interactions through message passing (Gilmer et al., 2017) but face overfitting and inefficiency challenges (Godwin et al., 2021; Rusch et al., 2023). Architectural improvements in GNNs facilitate the modeling of complex interactions, overcoming some limitations of deep networks (Gasteiger et al., 2019; Schütt et al., 2021; Batzner et al., 2022). Advances demonstrate the potential of architectural enhancements in GNNs to represent complex interactions (Gasteiger et al., 2020; 2021; Thölke & De Fabritiis, 2021; Batatia et al., 2022; Musaelian et al., 2023; Wang et al., 2024), but efficiently integrating many-body interactions into these networks is an ongoing challenge (Wang et al., 2023).

To address the complexities of many-body interactions in molecular systems, hypergraphs offer a compelling alternative to complex GNN architectures. Hypergraphs, with hyperedges connecting multiple vertices, can naturally represent many-body phenomena like electronic delocalization and collective vibrations. This allows for a more accurate modeling of molecular intricacies beyond the limitations of traditional graphs. Integrating hypergraphs with machine learning, particularly through Hypergraph Neural Networks (HGNNs), is an emerging research area. HGNNs manage the flow of information across hyperedges, capturing complex multi-atom interactions and enriching molecular representations. This technique promises to balance model expressiveness with computational efficiency. By innately encoding many-body interactions, HGNNs stand to significantly advance cheminformatics, offering a new approach to molecular property prediction and simulation that resonates with the actual behavior of chemical systems.

In this work, we introduce SE3Set, an innovative approach that enhances traditional GNNs by exploiting hypergraphs for modeling many-body interactions, while ensuring SE(3) equivariant representations that remain consistent regardless of molecular orientation. Our key contributions are:

- A new fragmentation method for hypergraph construction that seamlessly integrates 2D chemical and 3D spatial information, enriching the molecular structure representation.
- The deployment of hypergraph neural networks to capture many-body interactions, providing a deeper insight into molecular behavior that surpasses conventional pairwise modeling.
- The incorporation of SE(3) equivariance within our hypergraph framework, guaranteeing orientation-independent molecular representations.
- SE3Set underwent a comprehensive benchmarking process, exhibiting comparable outcomes to state-of-the-art (SOTA) models on small molecule datasets QM9 and MD17. It demonstrated exceptional performance on the larger molecule dataset MD22, where higher-order interactions are more evident, surpassing SOTA models with a significant reduction in mean absolute errors (MAEs) by an average of roughly 20%. This confirms SE3Set’s efficacy in capturing the complexity of molecular representations.

These advances establish SE3Set as a formidable tool for molecular representation learning, with implications for computational chemistry and beyond.

2 RELATED WORKS

2.1 GRAPH NEURAL NETWORKS

Message passing neural networks (MPNNs), a class of graph neural networks, are essential for learning node features by transmitting information along graph edges, a process crucial for interpreting structured data like molecules (Gilmer et al., 2017). Equivariant GNNs are especially important for molecular modeling. They adopt either group representation methods, aligning architectures to symmetry groups for improved interaction modeling (Thomas et al., 2018; Anderson et al., 2019; Fuchs et al., 2020; Batzner et al., 2022; Liao & Smidt, 2022; Liao et al., 2023; Musaelian et al., 2023), or direction-based methods that incorporate spatial information for accurate molecular representations (Schütt et al., 2017; Kindermans & Müller, 2018; Coors et al., 2018; Gasteiger et al., 2019; 2020; Schütt et al., 2021; Thölke & De Fabritiis, 2021; Gasteiger et al., 2021; Wang et al., 2022; Du et al., 2024; Aykent & Xia, 2024; Wang et al., 2024) and have been engineered to handle intricate up to five-body interactions (Wang et al., 2023).

2.2 HYPERGRAPH NEURAL NETWORKS

Hypergraph Neural Networks (HGNNs) enhance GNNs by incorporating multi-node hyperedges, better capturing complexity in data from various domains. They advance GNNs’ binary interactions with methods like clique expansion for compatibility with existing algorithms (Agarwal et al., 2005; Zhou et al., 2006) and employ tensor techniques for improved hypergraph-based feature learning (Li et al., 2013; Pearson & Zhang, 2014; Benson et al., 2017; Chien et al., 2021a; Tudisco et al., 2021). While equivariant HGNNs adeptly handle node permutations, preserving data symmetries (Kim et al., 2021; 2022a), they often miss 3D spatial transformations, crucial for physical system modeling. In computational chemistry, hypergraph algorithms simulate complex behaviors and optimize molecules

through hypergraph grammar (Cui et al., 2023; Tavakoli et al., 2022; Kajino, 2019), providing multidimensional insights into molecular structures (Nachmani & Wolf, 2020; Chen et al., 2021; Chen & Schwaller, 2023). Despite their promise, these methods still face hurdles in integrating spatial information effectively.

2.3 FRAGMENTATION METHODS

Fragmentation methods break down complex molecules for simpler *ab initio* QM computations of properties, later combining these for a holistic view (Gordon et al., 2012; Collins & Bettens, 2015). Leveraging the localized nature of chemical reactions, these techniques aim for scalable algorithms suitable for large molecule analysis. While instrumental in computational pretraining (Du et al., 2021; Kim et al., 2022b; Luong & Singh, 2023), they typically neglect the fusion of 2D structural with 3D spatial data. Hence, we advocate for a refined fragmentation approach that merges chemical properties with spatial context, potentially advancing hypergraph-based chemical modeling.

3 PRELIMINARIES

3.1 EQUIVARIANCE

Consider a function \mathcal{L} that maps inputs from space \mathcal{X} to outputs in space \mathcal{Y} . \mathcal{L} is called G -equivariant if it preserves the symmetry of a group G across mappings, meaning for each $g \in G$, we have:

$$\mathcal{L} \circ D^{\mathcal{X}}(g) = D^{\mathcal{Y}}(g) \circ \mathcal{L}, \quad (1)$$

where $D^{\mathcal{X}}$ represents the group G 's action on \mathcal{X} . This ensures that the function \mathcal{L} reflects changes made to inputs by G in its outputs.

3.2 HYPERGRAPH

Hypergraphs elegantly capture the essence of higher-order interactions among multiple entities, making them an invaluable tool for representing complex relational data. Let $\mathcal{G} = (V, E)$ be a hypergraph with N vertices and M hyperedges, where V represents a set of nodes and E is a set of hyperedges. Distinguishing itself from a traditional graph, a hyperedge can encompass multiple nodes, not limited to two, i.e. each hyperedge $e \in E$ is a non-empty subset of V .

3.3 ALLSET

The AllSet framework (Chien et al., 2021b), an advanced HGNN model, addresses heuristic propagation rule limitations in HGNNs by integrating Deep Sets (Zaheer et al., 2017) and Set Transformers (Lee et al., 2019) principles. It uses task-optimized dual multiset functions that maintain permutation invariance, crucial for hypergraph learning. The update rules in AllSet are:

$$Z_{e,:}^{(t+1),v} = f_{\mathcal{V} \rightarrow \mathcal{E}}(V_{e \setminus v, X^{(t)}}; Z_{e,:}^{(t),v}, X_{v,:}^{(t)}), \quad (2)$$

$$X_{v,:}^{(t+1)} = f_{\mathcal{E} \rightarrow \mathcal{V}}(E_{v, Z^{(t+1),v}}; X_{v,:}^{(t)}). \quad (3)$$

Here, $f_{\mathcal{V} \rightarrow \mathcal{E}}$ and $f_{\mathcal{E} \rightarrow \mathcal{V}}$ are the key multiset functions mapping node and hyperedge features. For example, $f_{\mathcal{V} \rightarrow \mathcal{E}}(S) = \text{MLP}(\sum_{s \in S} \text{MLP}(s))$ is used in AllDeepSets. The notation $V_{e,X}$ and $E_{v,Z}$ represent multisets of node and hyperedge features, respectively. The AllSet approach updates nodes and hyperedges in the hypergraph by leveraging their features in conjunction with those of adjacent hyperedges or nodes, enabling a rich representation of the hypergraph structure. The method could differentiate node v from its multiset, allowing for sophisticated feature aggregation.

4 METHODS

We introduce the SE3Set model to leverage hypergraph neural networks for capturing complex molecular interactions, integrating both 2D chemical and 3D spatial structures (Sec. 4.1). It builds upon the AllSet framework (Chien et al., 2021b) and the Equiformer (Liao & Smidt, 2022). Upcoming sections will delve into the specifics of molecular fragmentation and the SE3Set architecture.




Figure 1: Folic acid fragmentation illustrated with CID 135398658 from PubChem. (a) Preprocessing to identify cleavable bonds for fragmentation. (b) Initial fragments formed using BFS, color-coded by functional groups (blue), rings (orange), and single atoms (green). (c) Fragments merged to satisfy atom count criteria, detailed in C. (d) Expansion of fragments shown with directional arrows.

4.1 FRAGMENTATION ALGORITHM

To harness the power of hypergraph neural networks for molecular representations, we need to map molecules onto hypergraph structures through a refined fragmentation algorithm. Our strategy intertwines molecular topology and spatial geometry to create hyperedges that capture groups of atoms, reflecting their functional and spatial characteristics. In crafting this fragmentation approach for hypergraph-based molecular representation, the methodology must adhere to a set of fundamental principles:

1. The design should merge topological chemistry with 3D structural data into a unified hypergraph representation, ensuring hyperedges accurately embody the molecule’s chemical and spatial properties.
2. Controlling fragment size is vital for the SE3Set model to balance capturing meaningful interactions and computational efficiency. Optimal fragment sizes are key for model performance and learning capabilities.
3. The fragmentation could only selectively break single bonds and must maintain functional groups and ring integrity to preserving key chemical information critical for the molecule’s properties and behavior.
4. Fragment overlap is essential to maintain functional group effects on local charge distribution and to ensure hyperedge interaction within the SE3Set model for improved molecular learning.

Before delving deeper into the specifics of our fragmentation method, it’s important to establish a foundational understanding through key definitions and concepts,

Definition 4.1. The bond order represents the multiplicity or the number of shared electron pairs that constitute a covalent bond between two atoms.

The bond order matrix \mathbf{B} is an $N \times N$ representation of bond strength between atoms in a molecule, with higher bond order values indicating stronger bonds. This symmetric matrix ($B_{ij} = B_{ji}$) is crucial for studying molecular structure and reactivity, capturing bond nuances including delocalized and resonance bonds in computational chemistry.

Our fragmentation algorithm improves molecular representations by combining bond order, functional groups, and substructures, including SMARTS-identified smaller rings and merged adjacent groups. Overcoming the drawbacks of non-overlapping fragmentation, it uses 3D spatial data and allows overlaps, preserving local effects for precise charge distribution and enhancing hypergraph neural network learning of molecular interactions.

Definition 4.2. A molecular fragment, denoted as \mathfrak{F} , is defined as a specific subset of atoms within a molecule, characterized by being a cohesive assembly of predefined substructures linked in a sequential concatenation.

Our fragmentation method meticulously dissects a given molecule into meaningful subsets of atoms, and this process unfolds through four steps (corresponding to the pipeline in Fig. 1):

1. Pre-processing by analyzing the molecule’s bond order matrix to mask high-order bonds and those within functional groups or rings, and merging adjacent functional groups for a streamlined structural representation.
2. Core substructures are delineated from the remaining bonds using a Breadth-First Search algorithm, establishing the basic units of the molecular framework.
3. These substructures are then aggregated into larger molecular fragments according to predefined rules that maintain a minimum atom count within each fragment. This step could be optional.
4. To enhance fragment connectivity, we expand each by incorporating adjacent groups, using interaction strength metrics based on interatomic distances to guide this process. Here we set the cutoff value denoted as c_w of an interaction strength metrics to intercept the extended fragment.

Furthermore, step 4 leads to a substantial computational overhead for hypergraph neural networks when processing larger molecular systems. To enhance the efficiency of our model for such expansive molecular systems, we introduce a revised strategy for step 4:

- 4* For each atom i , identify the neighboring atoms \mathcal{N}_i that fall within a specified radial cutoff r_c . A fragment \mathfrak{F} is considered to be adjacent to atom i if there is an overlap of at least one atom between \mathfrak{F} and \mathcal{N}_i . For ease of reference, the set of fragments adjacent to atom i is represented as $\mathcal{N}_i^{\mathfrak{F}}$, which implies that $\mathcal{N}_i^{\mathfrak{F}} = \{\mathfrak{F} | \mathfrak{F} \cap \mathcal{N}_i \neq \emptyset\}$.

We designate the application of step 4 as an explicit overlap and the application of step 4* as an implicit overlap. These approaches introduce nuanced variations in the mathematical expressions of our model, as reflected in the Eq. 8, Eq. 9, and Eq. 16. For the detailed step-by-step methodology, please refer to Appendix A.

4.2 SE3SET

Building upon our aforementioned fragmentation algorithm, we now turn to outline the architecture of SE3Set. SE3Set model, influenced by AllSet (Chien et al., 2021b) and built on the Equiformer (Liao & Smidt, 2022), incorporates 3D spatial equivariance (proof refers to Appendix E) in our hypergraph neural network, improving capture of many-body interactions for precise molecular structure representation. SE3Set consists of an embedding layer, attention blocks, and an output head, as shown in Fig. 2 (a).

4.2.1 EMBEDDING

As depicted in Fig. 2 (b), the embedding block generates detailed node and hyperedge features reflecting molecular structures. Node features blend intrinsic properties with degree embeddings from connected hyperedges, while hyperedge features aggregate node embeddings and the corresponding relative position vectors depending on the node on which the hyperedge feature located. SH (Spherical Harmonic) functions are used to project normalized relative position vectors \vec{r}_{ij} between node i and j into the irreducible representations (irreps) feature space with different order l , i.e. $\text{SH}(\vec{r}_{ij}) = Y^l \left(\frac{\vec{r}_{ij}}{\|\vec{r}_{ij}\|} \right)$. The features are also mapped onto same l -order irreps space for SE3 equivariance and

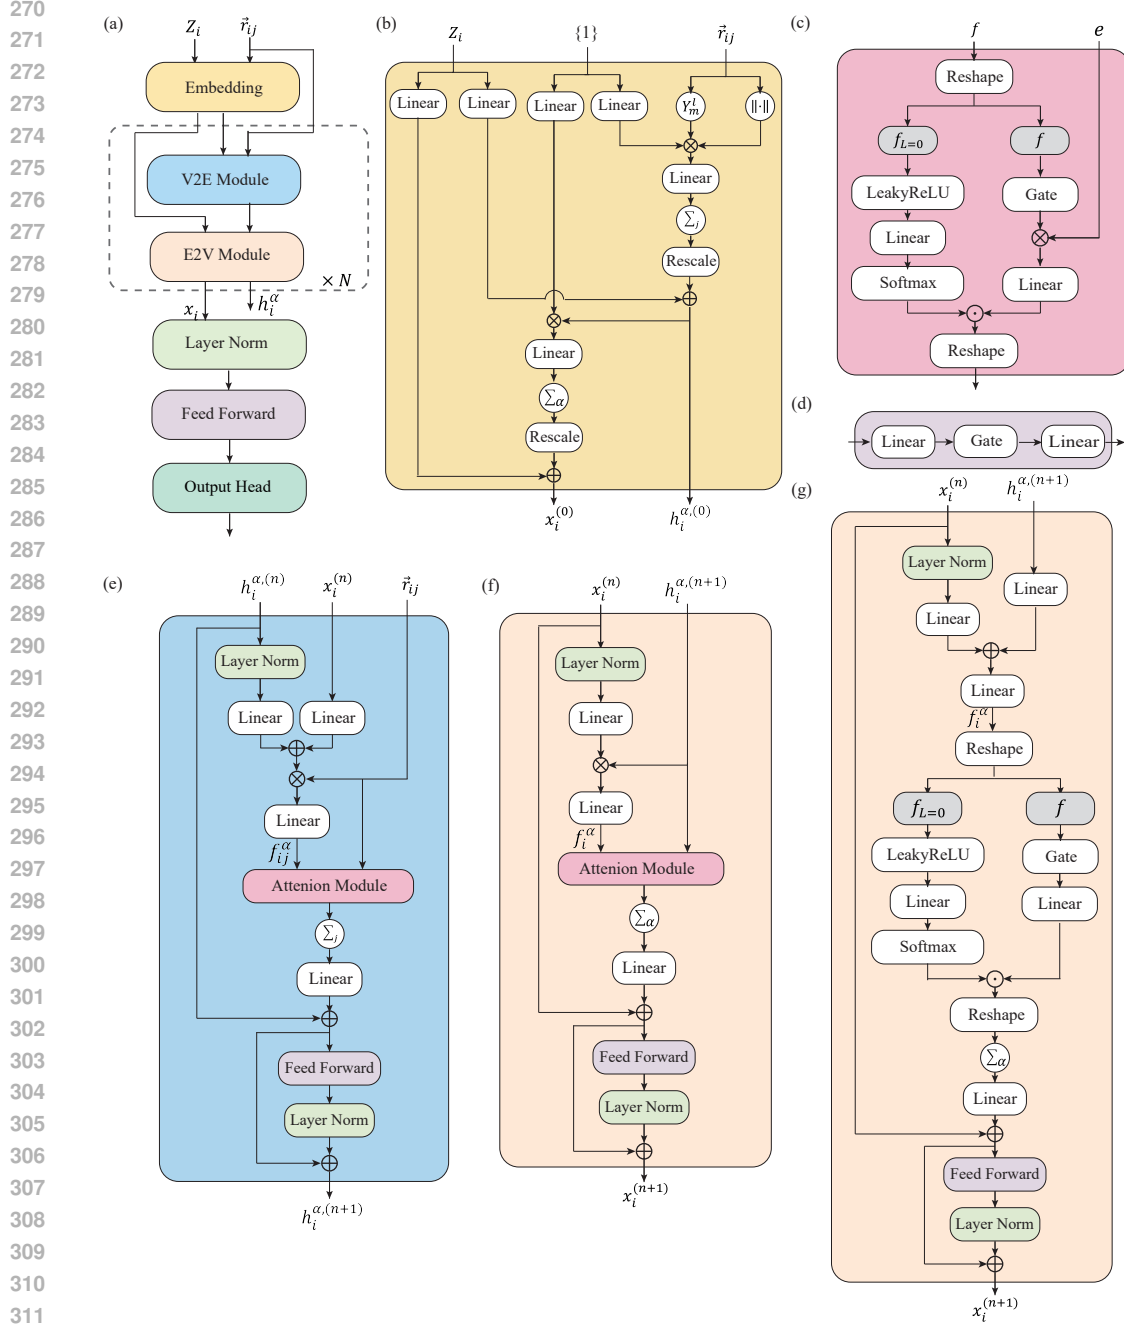


Figure 2: Overall architecture of SE3Set. (a) SE3Set begins with node and hyperedge embeddings, cycles through V2E and E2V attention modules for iterative updates, and concludes with normalization and a feed-forward block for output. (b) Embedding. Atomic numbers and position vectors are transformed into initial embeddings for nodes and hyperedges. (c) Attention Block. Merges feature sets with positional or hyperedge data for feature processing. (d) Feed-Forward Block. Enhances feature sets through a streamlined network. (e) V2E Module. Utilizes node features and their relative positions to update hyperedge features. (f) E2V Module. Employs hyperedge features to refresh node features, using tensor products (left) or summation (right) for updates. Symbols \otimes , \oplus , and \odot in figures denote depth-wise tensor product, summation, and Hadamard multiplication, respectively. h_i^α represents hyperedge features, x_i is for node features, superscript n indicates the number of updates, and \vec{r}_{ij} is the relative position vector between nodes i and j .

updated separately. Hyperedges capture nodes’ positional relationships, assigning a distinct feature h_i^α to each node i in hyperedge \mathfrak{F}_α , reinforcing structural fidelity. Nodes x_i integrate hyperedge information, harmonizing uniqueness with interconnections. Attention mechanisms then refine node and hyperedge interactions for accurate molecular and structural representation.

4.2.2 EQUIVARIANT HYPERGRAPH ATTENTION BLOCKS

As presented in Fig. 2 (c)-(f), the attention mechanism comprises two essential components: the Vertex-to-Edge (V2E) and Edge-to-Vertex (E2V) attention blocks, based on the AllSet framework (Chien et al., 2021b). The V2E block refines hyperedge features, while the E2V block updates node features, both operating with an equivariant hypergraph attention mechanism. To improve training and enable deeper network structures, we incorporate normalization layers and residual connections to prevent gradient issues. The attention module’s output passes through a feed forward block (Fig. 2 (d)), enhancing representation complexity. Node and hyperedge features maintain equivariance to molecular geometry, preserving data symmetries and the integrity of representations, thus boosting the model’s expressiveness in capturing complex structural interactions. (The concepts of irreducible representations and tensor products can be referenced in the Appendix F.)

V2E attention The SE3Set model uses geometrically invariant attention weights a_{ij} , derived from $l = 0$ irreps acting as scalars under geometric transformations. These weights are computed from scalar features $f_{ij,l=0}$ using an MLP with LeakyReLU activation and softmax normalization, reflecting node relationships within the hypergraph. Node and hyperedge features undergo non-linear transformations represented by tensor products of irreps with quantum number l . The features combine through direct tensor products (DTP), yielding non-linear values v_{ij} (Fig. 2 (c)). Hyperedge features are updated by aggregating features from connected nodes, utilizing SH and radial basis functions on hyperedge features. The model calculates initial features f_{ij}^α and V2E attention weights a_{ij}^α via MLPs, with non-linear values v_{ij}^α emerging from similar transformations.

$$t_{ij}^\alpha = (\text{Linear}(x_i) + \text{Linear}(x_j)) \quad (4)$$

$$f_{ij}^\alpha = \text{Linear}(t_{ij}^\alpha \otimes_w^{\text{DTP}} \text{SH}(\vec{r}_{ij})) \quad (5)$$

$$a_{ij}^\alpha = \text{Softmax}_j(a^\top \text{LeakyReLU}(f_{ij,l=0}^\alpha)) \quad (6)$$

$$v_{ij}^\alpha = \text{Linear}(\text{Gate}(f_{ij}^\alpha) \otimes_w^{\text{DTP}} \text{SH}(\vec{r}_{ij})) \quad (7)$$

Ultimately, the SE3Set model updates hyperedge features h_i^k by accumulating the weighted features of nodes within the same hyperedge and applying a linear transformation to the aggregated information. For the explicit overlap fragmentation method,

$$\Delta h_i^\alpha = \text{Linear} \left(\sum_{j:n_i \in \mathfrak{F}_\alpha \wedge n_j \in \mathfrak{F}_\alpha} a_{ij}^\alpha v_{ij}^\alpha \right) \quad (8)$$

where n_i denotes the node with index i and \mathfrak{F}_α denotes the fragment with index α as each fragment could be considered as a hyperedge in the hypergraph. Due to the frequent occurrence of a high number of explicitly overlapping atoms, this scenario commonly results in increased computational complexity. Consequently, when adopting the implicit overlap approach, we may opt for an equation of the form:

$$\Delta h_i^\alpha = \text{Linear} \left(\sum_{j:j \in \mathfrak{F}_\alpha \wedge \mathfrak{F}_\alpha \in \mathcal{N}_i^\mathcal{F}} a_{ij}^\alpha v_{ij}^\alpha \right) \quad (9)$$

where $\mathcal{N}_i^\mathcal{F}$ is delineated in step 4* of the implicit overlap method. This characteristic renders it a more computationally efficient scheme for Vertex-to-Edge (V2E) attention mechanisms. The detailed architecture of V2E attention block is shown in Fig. 2 (e).

E2V attention Following the V2E attention module, the E2V attention module (Fig. 2 (f)) updates node features by transforming them with a tensor product of the updated hyperedge feature, followed by a linear layer. Attention weights are then calculated using softmax-applied, LeakyReLU-activated

378 features, ensuring node features are refined after hyperedge updates.

$$379 \quad f_i^\alpha = \text{Linear}(\text{Linear}(x_i) \otimes^{\text{DTP}} h_i^\alpha) \quad (10)$$

$$380 \quad a_i^\alpha = \text{Softmax}_\alpha(a^\top \text{LeakyReLU}(f_{i,l=0}^\alpha)) \quad (11)$$

381 These attention weights direct the synthesis of information, culminating in the calculated value:

$$382 \quad v_i^\alpha = \text{Linear}(\text{Gate}(f_i^\alpha) \otimes^{\text{DTP}} h_i^\alpha). \quad (12)$$

383 Furthermore, we propose an alternative method for constructing the E2V attention block as shown in
384 Fig. 2 (g).

$$385 \quad f_i^\alpha = \text{Linear}(\text{Linear}(x_i) + \text{Linear}(h_i^\alpha)) \quad (13)$$

$$386 \quad a_i^\alpha = \text{Softmax}_\alpha(a^\top \text{LeakyReLU}(f_{i,l=0}^\alpha)) \quad (14)$$

$$387 \quad v_i^\alpha = \text{Linear}(\text{Gate}(f_i^\alpha)) \quad (15)$$

388 However, practical experiments reveal that the previous method yields superior results, with detailed
389 findings presented in Sec. 5.3.

390 Then the node aggregates all the hyperedge features corresponding to itself to update the node feature,

$$391 \quad \text{Explicit overlap: } \Delta x_i = \text{Linear} \left(\sum_{\alpha: i \in \mathfrak{F}_\alpha} a_i^\alpha v_i^\alpha \right), \quad (16)$$

$$392 \quad \text{Implicit overlap: } \Delta x_i = \text{Linear} \left(\sum_{\alpha: \mathfrak{F}_\alpha \in \mathcal{N}_i^{\mathcal{F}}} a_i^\alpha v_i^\alpha \right) \quad (17)$$

403 4.2.3 OUTPUT HEAD

404 The SE3Set model employs node features to generate predictions, using a feed-forward network to
405 transform these features into the target label’s irreps dimension. A summation strategy aggregates
406 node features into a single hypergraph-level representation, which is then processed by a linear layer
407 to output the model’s final predictions.

408 5 RESULTS

409 We tested our equivariant hypergraph neural network on QM9 (Ruddigkeit et al., 2012; Ramakrishnan
410 et al., 2014), MD17 (Chmiela et al., 2017) (see Appendix G), and MD22 (Chmiela et al., 2023)
411 to assess its molecular representation learning. QM9 and MD17 gauge small molecule property
412 prediction, while MD22 evaluates larger systems with complex many-body interactions (Wang et al.,
413 2023). An ablation study was also conducted to pinpoint the contributions of fragmentation and
414 architecture to our method’s performance, offering insights into the network’s efficacy and areas
415 for enhancement. In the Appendix H, we have also provided a detailed analysis of the model’s
416 complexity.

417 5.1 QM9

418 The QM9 dataset (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) consists of 134k small organic
419 molecules calculated at the B3LYP/6-31G(2df, p) level. SE3Set, after training on 110k QM9
420 molecules and validation on 10k, achieves low mean absolute errors (MAEs) in 12 tasks, performing
421 on par with leading models, as detailed in Table 1. In small molecular systems, higher-order many-
422 body interactions are less pronounced, and as a result, SE3Set does not significantly outperform other
423 state-of-the-art (SOTA) models.

424 5.2 MD22

425 Recognizing the prominence of higher-order many-body interactions in larger molecules (Wang
426 et al., 2023), SE3Set was tested on the comprehensive MD22 dataset (Chmiela et al., 2023). This
427

Table 1: A comparative analysis was performed to assess the Mean Absolute Errors (MAEs) on the QM9 dataset when training SE3Set on a configuration comprising 110,000 training samples and 1,000 validation samples. Bolding shows the best model and underlining shows the second best model and the underlining tilde shows third best model.

	UNIT	SCHNET	DIMENET++	PAINN	SPHERENET	COMENET	ET	ALLEGRO	VISNET	QUINNET	EQUIFORMER	SE3SET
μ	D	0.033	0.030	0.012	0.026	0.0245	<u>0.011</u>	-	0.010	0.771	<u>0.011</u>	<u>0.011</u>
α	a_0^3	0.235	<u>0.044</u>	<u>0.045</u>	0.046	0.0452	0.059	-	0.041	0.047	0.046	<u>0.045</u>
HOMO	meV	41	25	20	23	23	20	-	<u>17.3</u>	20.4	15	15
LUMO	meV	34	20	28	18	20	18	-	<u>14.8</u>	17.6	<u>14</u>	13
GAP	meV	63	33	46	32	32	36	-	31.7	28.2	<u>30</u>	<u>29</u>
R^2	a_0^2	0.073	0.331	<u>0.066</u>	0.292	0.259	<u>0.033</u>	-	0.030	0.194	0.251	0.197
ZPVE	meV	1.70	<u>1.21</u>	1.28	1.12	<u>1.20</u>	1.84	-	1.56	1.26	1.26	1.40
U_0	meV	14	6	5.85	6	6.59	6.15	<u>4.7</u>	4.23	7.6	6.59	<u>5.74</u>
U	meV	19	6	5.83	7	6.82	6.38	<u>4.4</u>	4.25	8.4	6.74	<u>5.69</u>
H	meV	14	7	5.98	6	6.86	6.16	4.4	<u>4.52</u>	7.8	6.63	<u>5.70</u>
G	meV	14	8	7.35	8	7.98	7.62	5.7	<u>5.86</u>	8.5	7.63	<u>6.63</u>
C_v	$\frac{\text{kcal}}{\text{mol}\cdot\text{K}}$	0.033	<u>0.023</u>	0.024	0.021	0.024	0.026	-	<u>0.023</u>	0.024	<u>0.023</u>	0.025

Table 2: A comparison of Mean Absolute Errors (MAEs) across various benchmarked models. SE3Set is trained on the five molecules of MD22 dataset with specific number of training/validation. Bolding shows the best model and underlining shows the second best model. The improvements column shows the improvement of our model over the previous SOTA model in percentage terms. The MAEs reflect the precision of energy predictions in units of kcal/mol and forces in units of kcal/(mol·Å). The results of TorchMD-Net, Allegro, and Equiformer are extracted from Ref. Li et al. (2024).

MOLECULE	# TRAIN/VAL		SGDML	TORCHMD-NET	ALLEGRO	MAE	EQUIFORMER	VISNET	QUINNET	EQUIFORMER-LSRM	VISNET-LSRM	SE3SET	IMPROVEMENTS
AC-ALA3-NHME	5500/500	ENERGY	0.3902	0.1121	0.1019	<u>0.0620</u>	0.0828	0.0796	0.084	0.0780	0.0654	0.0499	19.5%
		FORCE	0.7968	0.1879	0.1068	0.0876	0.0804	0.0972	<u>0.0681</u>	0.0877	0.0902	0.0545	20.0%
DHA	7500/500	ENERGY	1.3117	0.1205	0.1153	0.1317	0.1788	0.1526	0.12	0.0878	<u>0.0873</u>	0.0826	5.4%
		FORCE	0.7474	0.1209	0.0732	0.0646	<u>0.0506</u>	0.0668	0.0515	0.0534	0.0598	0.0360	28.9%
STACHYOSE	7500/500	ENERGY	4.0497	0.1393	0.2485	0.1244	0.1404	0.1283	0.23	0.1252	<u>0.1055</u>	0.0762	27.8%
		FORCE	0.6744	0.1921	0.0971	0.0876	0.0635	0.0869	<u>0.0543</u>	0.0632	0.0767	0.0424	21.9%
AT-AT	2500/500	ENERGY	0.7235	0.1120	0.1428	0.1093	0.1309	0.1688	0.14	0.1007	<u>0.0772</u>	0.0585	24.2%
		FORCE	0.6911	0.2036	0.0952	0.0992	0.0960	0.1070	<u>0.0687</u>	0.0811	0.0781	0.0556	19.1%
AT-AT-CG-CG	1500/500	ENERGY	1.3885	0.2072	0.3933	0.1578	0.1510	0.1995	0.38	0.1335	<u>0.1135</u>	0.1002	11.7%
		FORCE	0.7028	0.3259	0.1280	0.1153	0.1252	0.1563	0.1273	0.1065	<u>0.1063</u>	0.0825	22.4%

dataset spans four classes of biomolecules and supramolecules, from a 42-atom peptide to a 370-atom nanotube, with high-resolution sampling at 400-500 K using the PBE+MBD (Perdew et al., 1996; Tkatchenko et al., 2012) framework for energy and force computations. Our fragmentation method, which maintains functional groups and rings, selectively excludes structures like the Buckyball catcher and Double-walled nanotube from MD22, thus concentrating on the other five molecular types. We partition the training/test set following QuinNet (Wang et al., 2023). As Table 2 shows, SE3Set outperforms other SOTA models in these cases, reducing MAEs by an average of 20%, underscoring its exceptional ability to capture molecular intricacies. Moreover, our results indicate that incorporating higher-order many-body interactions is crucial for representing the non-local features of larger molecules within the MD22 dataset.

5.3 ABLATION STUDIES

To better understand SE3Set, we conduct ablation studies focusing on fragmentation and model architecture. For explicit overlap fragmentation method, we explore how different fragmentation techniques affects SE3Set’s training and compare with the non-overlapping BRICS (Degen et al., 2008; Landrum et al., 2020) strategy on QM9’s homo energy task. As Fig. 3 indicates, tests on QM9’s homo energy task showed SE3Set’s robustness to c_w variations in fragmentation method. The results Our method surpasses BRICS demonstrates the importance of hyperedge interaction. Furthermore, we per-




Figure 3: Ablation studies on the QM9 dataset’s HOMO task (units: meV). The variable c_w represents the threshold for expansion in the fourth step of fragmentation, guided by the fragment bond order defined in Eq. 18. The term BRICS denotes another fragmentation method implemented in RDKit. Additionally, the E2V summation refers to the architectural framework specified from Eq. 13 to Eq. 15.

Table 3: The MAE for energy (unit: kcal/mol) and forces (unit: kcal/mol \cdot Å) on the AT-AT-CG-CG dataset using 3 layers SE3Set with different cutoff radii (r_c) in the implicit overlap method.

IMPLICIT r_c (Å)	4.0	5.0	6.0
ENERGY	0.2123	0.1153	0.1103
FORCE	0.1559	0.1019	0.0937

formed ablation studies on the model architecture. Among two design variants in the E2V attention section, the one using tensor product interactions between nodes and hyperedges proved superior, emphasizing the value of our tensor product-based mechanism and architecture design in enhancing molecular property predictions. Additionally, a 6-layer SE3Set model outperformed its 3-layer counterpart.

Unlike the explicit overlap method, the fragment size of implicit overlap method depends on the choice of r_c . We tests the effect of r_c on the model training results on the AT-AT-CG-CG molecule of the MD22 dataset (Table 3). The SE3Set performs better when using higher r_c as it will include more fragments to generate implicit overlaps. The effect of r_c on fragment size is more pronounced for the implicit method than for the explicit method, but it gives a good performance for different parameters comparing with the baseline models.

6 CONCLUSION

In conclusion, this study demonstrate the efficacy of SE3Set, a cutting-edge hypergraph neural network architecture, in the realm of molecular representation learning. By meticulously crafting a fragmentation method that coalesces two-dimensional chemical knowledge with three-dimensional spatial information, we establish a robust foundation for constructing hypergraphs that faithfully capture the complex nature of molecular structures. The SE3Set architecture, drawing inspiration from the AllSet framework and the Equiformer, adeptly processing these hypergraphs and preserving the essential invariances and symmetries. SE3Set demonstrates performance comparable to SOTA models in small molecular systems and significantly outperforms SOTA models in large molecular systems where higher-order many-body interactions are pronounced. The results of our research affirm the potential of SE3Set to model high-order many-body interactions, providing a powerful tool for molecular representation.

REFERENCES

- Sameer Agarwal, Jongwoo Lim, Lihi Zelnik-Manor, Pietro Perona, David Kriegman, and Serge Belongie. Beyond pairwise clustering. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pp. 838–845. IEEE, 2005.
- Brandon Anderson, Truong Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *Advances in neural information processing systems*, 32, 2019.
- Sarp Aykent and Tian Xia. Savenet: a scalable vector network for enhanced molecular representation learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, 2022.
- 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. *Nat. Commun.*, 13(1):2453, 2022.
- Austin R Benson, David F Gleich, and Lek-Heng Lim. The spacey random walk: A stochastic process for higher-order data. *SIAM Rev.*, 59(2):321–345, 2017.
- Adam J. Bridgeman and Christopher J Empson. Bond orders between molecular fragments. *Chemistry*, 12 8:2252–62, 2006.

- 540 Fangying Chen, Junyoung Park, and Jinkyoo Park. A hypergraph convolutional neu-
541 ral network for molecular properties prediction using functional group. Preprint at
542 <http://arxiv.org/abs/2106.01028>, 2021.
- 543
- 544 Junwu Chen and Philippe Schwaller. Molecular hypergraph neural networks. Preprint at
545 <http://arxiv.org/abs/2312.13136>, 2023.
- 546 Eli Chien, Pan Li, and Olgica Milenkovic. Landing probabilities of random walks for seed-set
547 expansion in hypergraphs. In 2021 IEEE Information Theory Workshop (ITW), pp. 1–6. IEEE,
548 2021a.
- 549 Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenkovic. You are AllSet: A multiset function frame-
550 work for hypergraph neural networks. In International Conference on Learning Representations,
551 2021b.
- 552
- 553 Stefan Chmiela, Alexandre Tkatchenko, Huziel E Sauceda, Igor Poltavsky, Kristof T Schütt, and
554 Klaus-Robert Müller. Machine learning of accurate energy-conserving molecular force fields. Sci.
555 Adv., 3(5):e1603015, 2017.
- 556 Stefan Chmiela, Valentin Vassilev-Galindo, Oliver T Unke, Adil Kabylda, Huziel E Sauceda, Alexan-
557 dre Tkatchenko, and Klaus-Robert Müller. Accurate global machine learning force fields for
558 molecules with hundreds of atoms. Science Advances, 9(2):eadf0873, 2023.
- 559
- 560 Michael A Collins and Ryan PA Bettens. Energy-based molecular fragmentation methods. Chem.
561 Rev., 115(12):5607–5642, 2015.
- 562 Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. SphereNet: Learning spherical
563 representations for detection and classification in omnidirectional images. In Proceedings of the
564 European conference on computer vision (ECCV), pp. 518–533, 2018.
- 565
- 566 Shicheng Cui, Qianmu Li, Deqiang Li, Zhichao Lian, Jun Hou, et al. Hyper-mol: Molecular
567 representation learning via fingerprint-based hypergraph. Comput. Intell. Neurosci., 2023, 2023.
- 568 Laurianne David, Amol Thakkar, Rocío Mercado, and Ola Engkvist. Molecular representations in
569 AI-driven drug discovery: a review and practical guide. J. Cheminform., 12(1):1–22, 2020.
- 570
- 571 Jörg Degen, Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. On the art of
572 compiling and using ‘drug-like’ chemical fragment spaces. ChemMedChem: Chemistry Enabling
573 Drug Discovery, 3(10):1503–1507, 2008.
- 574 Kien Do, Truyen Tran, and Svetha Venkatesh. Graph transformation policy network for chemi-
575 cal reaction prediction. In Proceedings of the 25th ACM SIGKDD international conference on
576 knowledge discovery & data mining, pp. 750–760, 2019.
- 577
- 578 Boxin Du, Changhe Yuan, Robert Barton, Tal Neiman, and Hanghang Tong. Hypergraph pre-training
579 with graph neural networks. Preprint at <http://arxiv.org/abs/2105.10862>, 2021.
- 580 Yuanqi Du, Limei Wang, Dieqiao Feng, Guifeng Wang, Shuiwang Ji, Carla P Gomes, Zhi-Ming Ma,
581 et al. A new perspective on building efficient and expressive 3d equivariant graph neural networks.
582 Advances in Neural Information Processing Systems, 36, 2024.
- 583
- 584 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric.
585 Preprint at <http://arxiv.org/abs/1903.02428>, 2019.
- 586 Denis Fourches, Eugene Muratov, and Alexander Tropsha. Trust, but verify: on the importance of
587 chemical structure curation in cheminformatics and qsar modeling research. J. Chem. Inf. Model.,
588 50(7):1189, 2010.
- 589 Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. SE (3)-transformers: 3D Roto-
590 translation equivariant attention networks. Advances in Neural Information Processing Systems,
591 33:1970–1981, 2020.
- 592
- 593 Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular
graphs. In International Conference on Learning Representations, 2019.

- 594 Johannes Gasteiger, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and
595 uncertainty-aware directional message passing for non-equilibrium molecules. Preprint at
596 <http://arxiv.org/abs/2011.14115>, 2020.
- 597
- 598 Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional
599 graph neural networks for molecules. Advances in Neural Information Processing Systems, 34:
600 6790–6802, 2021.
- 601
- 602 Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural
603 message passing for quantum chemistry. In International conference on machine learning, pp.
604 1263–1272. PMLR, 2017.
- 605
- 606 Jonathan Godwin, Michael Schaarschmidt, Alexander L Gaunt, Alvaro Sanchez-Gonzalez, Yulia
607 Rubanova, Petar Veličković, James Kirkpatrick, and Peter Battaglia. Simple GNN regularisation
608 for 3D molecular property prediction and beyond. In International Conference on Learning
609 Representations, 2021.
- 610
- 611 Mark S Gordon, Dmitri G Fedorov, Spencer R Pruitt, and Lyudmila V Slipchenko. Fragmentation
612 methods: A route to accurate calculations on large systems. Chem. Rev., 112(1):632–672, 2012.
- 613
- 614 Hiroshi Kajino. Molecular hypergraph grammar with its application to molecular optimization. In
615 International Conference on Machine Learning, pp. 3183–3191. PMLR, 2019.
- 616
- 617 Jinwoo Kim, Saeyoon Oh, and Seunghoon Hong. Transformers generalize deepsets and can be
618 extended to graphs & hypergraphs. Advances in Neural Information Processing Systems, 34:
619 28016–28028, 2021.
- 620
- 621 Jinwoo Kim, Saeyoon Oh, Sungjun Cho, and Seunghoon Hong. Equivariant hypergraph neural
622 networks. In European Conference on Computer Vision, pp. 86–103. Springer, 2022a.
- 623
- 624 Seojin Kim, Jaehyun Nam, Junsu Kim, Hankook Lee, Sungsoo Ahn, and Jinwoo Shin. Contrastive
625 learning of molecular representation with fragmented views. 2022b.
- 626
- 627 P-J Kindermans and K-R Müller. Schnet—a deep learning architecture for molecules and materials. J.
628 Chem. Phys., 148(24), 2018.
- 629
- 630 Greg Landrum, Paolo Tosco, Brian Kelley, sriniker, gedec, NadineSchneider, Riccardo Vianello, Ric,
631 Andrew Dalke, Brian Cole, AlexanderSavelyev, Matt Swain, Samo Turk, Dan N, Alain Vaucher,
632 Eisuke Kawashima, Maciej Wójcikowski, Daniel Probst, guillaume godin, David Cosgrove, Axel
633 Pahl, JP, Francois Berenger, strets123, JLVarjo, Noel O’Boyle, Patrick Fuller, Jan Holst Jensen,
634 Gianluca Sforna, and DoliathGavid. rdkit/rdkit: 2020_03_1 (q1 2020) release, March 2020. URL
635 <https://doi.org/10.5281/zenodo.3732262>.
- 636
- 637 Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set trans-
638 former: A framework for attention-based permutation-invariant neural networks. In International
639 conference on machine learning, pp. 3744–3753. PMLR, 2019.
- 640
- 641 György Lendvay. On the correlation of bond order and bond length. J. Mol. Struct., 501:389–393,
642 2000.
- 643
- 644 Guoyin Li, Liqun Qi, and Gaohang Yu. The Z-eigenvalues of a symmetric tensor and its application
645 to spectral hypergraph theory. Numer. Linear Algebra Appl., 20(6):1001–1029, 2013.
- 646
- 647 Yunyang Li, Yusong Wang, Lin Huang, Han Yang, Xinran Wei, Jia Zhang, Tong Wang, Zun Wang,
Bin Shao, and Tie-Yan Liu. Long-short-range message-passing: A physics-informed framework
to capture non-local interaction for scalable molecular dynamics simulation. In The Twelfth
International Conference on Learning Representations, 2024. URL <https://openreview.net/forum?id=rvDQtMnOl>.
- Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3D atomistic
graphs. In The Eleventh International Conference on Learning Representations, 2022.

- 648 Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. EquiformerV2: Im-
649 proved equivariant transformer for scaling to higher-degree representations. Preprint at
650 <http://arxiv.org/abs/2306.12059>, 2023.
- 651 Kha-Dinh Luong and Ambuj Singh. Fragment-based pretraining and finetuning on molecular graphs.
652 Preprint at <http://arxiv.org/abs/2310.03274>, 2023.
- 653 Jonathan P Mathews and Alan L Chaffee. The molecular representations of coal—a review. Fuel, 96:
654 1–14, 2012.
- 655 David L Mobley, Caitlin C Bannan, Andrea Rizzi, Christopher I Bayly, John D Chodera, Victoria T
656 Lim, Nathan M Lim, Kyle A Beauchamp, David R Slochower, Michael R Shirts, et al. Escaping
657 atom types in force fields using direct chemical perception. J. Chem. Theory Comput., 14(11):
658 6076–6092, 2018.
- 659 Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai
660 Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic
661 dynamics. Nat. Commun., 14(1):579, 2023.
- 662 Eliya Nachmani and Lior Wolf. Molecule property prediction and classification with graph hypernet-
663 works. Preprint at <http://arxiv.org/abs/2002.00240>, 2020.
- 664 Kelly J Pearson and Tan Zhang. On spectral hypergraph theory of the adjacency tensor. Graphs
665 Combin., 30:1233–1248, 2014.
- 666 John P Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made
667 simple. Phys. Rev. Lett., 77(18):3865, 1996.
- 668 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum
669 chemistry structures and properties of 134 kilo molecules. Sci. Data, 1(1):1–7, 2014.
- 670 Patrick Reiser, Marlen Neubert, André Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam
671 Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. Graph neural networks for
672 materials science and chemistry. Commun. Mater., 3(1):93, 2022.
- 673 Lars Ruddigkeit, Ruud Van Deursen, Lorenz C Blum, and Jean-Louis Reymond. Enumeration of 166
674 billion organic small molecules in the chemical universe database GDB-17. J. Chem. Inf. Model.,
675 52(11):2864–2875, 2012.
- 676 T Konstantin Rusch, Michael M Bronstein, and Siddhartha Mishra. A survey on oversmoothing in
677 graph neural networks. Preprint at <http://arxiv.org/abs/2303.10993>, 2023.
- 678 Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre
679 Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network
680 for modeling quantum interactions. Advances in neural information processing systems, 30, 2017.
- 681 Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction
682 of tensorial properties and molecular spectra. In International Conference on Machine Learning,
683 pp. 9377–9388. PMLR, 2021.
- 684 Mohammadamin Tavakoli, Alexander Shmakov, Francesco Ceccarelli, and Pierre Baldi. Rxn
685 hypergraph: a hypergraph attention model for chemical reaction representation. Preprint at
686 <http://arxiv.org/abs/2201.01196>, 2022.
- 687 Philipp Thölke and Gianni De Fabritiis. Equivariant transformers for neural network based molecular
688 potentials. In International Conference on Learning Representations, 2021.
- 689 Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley.
690 Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds.
691 Preprint at <http://arxiv.org/abs/1802.08219>, 2018.
- 692 Alexandre Tkatchenko, Robert A DiStasio Jr, Roberto Car, and Matthias Scheffler. Accurate and
693 efficient method for many-body van der Waals interactions. Phys. Rev. Lett., 108(23):236402,
694 2012.

- 702 Francesco Tudisco, Austin R Benson, and Konstantin Prokopchik. Nonlinear higher-order label
703 spreading. In *Proceedings of the Web Conference 2021*, pp. 2402–2413, 2021.
704
- 705 Jeff Wagner, Matt Thompson, David L. Mobley, John Chodera, Caitlin Bannan, Andrea Rizzi,
706 trevorgokey, David L. Dotson, Josh A. Mitchell, jaimergp, Camila, Pavan Behara, Christopher
707 Bayly, JoshHorton, Iván Pulido, Lily Wang, Victoria Lim, Sukanya Sasmal, SimonBoothroyd,
708 Andrew Dalke, Daniel Smith, Brent Westbrook, Josh Horton, Lee-Ping Wang, Richard Gowers,
709 Ziyuan Zhao, Connor Davel, and Yutong Zhao. openforcefield/openff-toolkit: 0.15.1. Testing
710 updates, January 2024. URL <https://doi.org/10.5281/zenodo.10593535>.
- 711 Limei Wang, Yi Liu, Yuchao Lin, Haoran Liu, and Shuiwang Ji. ComENet: Towards complete and
712 efficient message passing for 3D molecular graphs. In Alice H. Oh, Alekh Agarwal, Danielle
713 Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems, 2022*.
714
- 715 Yusong Wang, Tong Wang, Shaoning Li, Xinheng He, Mingyu Li, Zun Wang, Nanning Zheng,
716 Bin Shao, and Tie-Yan Liu. Enhancing geometric representations for molecules with equivariant
717 vector-scalar interactive message passing. *Nat. Commun.*, 15(1):313, 2024.
- 718 Zun Wang, Guoqing Liu, Yichi Zhou, Tong Wang, and Bin Shao. Efficiently incorporating quintuple
719 interactions into geometric deep learning force fields. In *Thirty-seventh Conference on Neural
720 Information Processing Systems, 2023*.
- 721 Daniel S Wigh, Jonathan M Goodman, and Alexei A Lapkin. A review of molecular representation
722 in the age of machine learning. *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, 12(5):e1603, 2022.
723
- 724 Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
725 comprehensive survey on graph neural networks. *IEEE Trans. Neural. Netw. Learn. Syst.*, 32(1):
726 4–24, 2020.
- 727 Jiacheng Xiong, Zhaoping Xiong, Kaixian Chen, Hualiang Jiang, and Mingyue Zheng. Graph neural
728 networks for automated de novo drug design. *Drug Discov. Today*, 26(6):1382–1393, 2021.
729
- 730 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and
731 Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.
- 732 Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. Learning with hypergraphs: Clustering,
733 classification, and embedding. *Advances in neural information processing systems*, 19, 2006.
734
- 735 Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang,
736 Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications.
737 *AI open*, 1:57–81, 2020.
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

APPENDIX

A DETAILS OF FRAGMENTATION STEPS

Based on the design principles in Sec. 4.1, the detailed step-by-step methodology of explicit overlap fragmentation method is shown as follows,

1. The pre-processing step begins by analyzing the given molecule through its bond order matrix, denoted as B . Identify and mask bonds that are part of functional groups or rings, as well as those with a bond order of $B_{ij} \geq 2$. Functional groups are then identified using predefined SMARTS patterns for accurate matching. To achieve a more generalized representation of functional groups, topologically adjacent functional groups are merged into a single entity. This aggregation allows to focus on specific subfunctional groups that are of particular interest, simplifying the complexity of the molecular structure for subsequent analysis.
2. Following the masking of selected bonds, the Breadth-First Search (BFS) algorithm is employed to reconstruct the substructures, denoted as $\{\mathfrak{S}\}$, from the remaining unmasked bonds. These groups represent the core structural units of the molecule as discussed at the outset of this section.
3. Consolidate the previously identified groups $\{\mathfrak{S}\}$ into larger molecular fragments, applying a set of predefined rules to guide the merging process. These rules are meticulously designed to ensure that each resulting fragment, now denoted as $\{\mathfrak{F}\}$, contains at least a minimum specified number of atoms. For a comprehensive understanding of the merging criteria, one can refer to the detailed rules outlined in C.
4. Extend each fragment $\{\mathfrak{F}\}$ by incorporating adjacent groups from $\{\mathfrak{S}\}$ to enrich the connectivity between molecular fragments, thus intentionally creating regions of overlap among the fragments. This expansion is controlled by a cutoff threshold, denoted as c_w , which is typically a function based on interatomic distances. the fragment bond order (Lendvay, 2000; Bridgeman & Empson, 2006), symbolized by W_{fs} , is used to quantitatively assess the interaction strength between a fragment \mathfrak{F}_i and an adjacent substructure \mathfrak{S}_j . This method reflects the interaction strength based on the proximity of atoms in different fragments, expressed by the following equation:

$$W_{fs} = \sum_{i \in \mathfrak{F}_f, j \in \mathfrak{S}_s} \exp\left(-\frac{(d_{ij} - d_{ij}^e) \cdot d_{ij}^e}{(0.25 \text{ \AA})^2}\right), \quad (18)$$

where d_{ij} represents the interatomic distance between atoms i and j , and d_{ij}^e stands for the equilibrium distance typically expected for such a bond. This equation is utilized to determine which substructures should be included in the expansion of a fragment, based on the strength of their interactions as governed by the distance function. Additionally, in alignment with Pauling’s concept of "chemist’s bond order" (Lendvay, 2000), an alternative method is introduced to calculate the bond order using a single exponential function,

$$W_{fs} = \sum_{i \in \mathfrak{F}_f, j \in \mathfrak{S}_s} \exp(-(d_{ij} - d_{ij}^e)), \quad (19)$$

where W_{fs} encapsulates the bond order between atoms belonging to a fragment \mathfrak{F}_f and a substructure \mathfrak{S}_s . In this context, d_{ij} signifies the actual measured distance between atom i and atom j . The term d_{ij}^e refers to the theoretical equilibrium covalent bond length, which is estimated by summing the empirical covalent radii of the two atoms involved, given by:

$$d_{ij}^e = r_{z_i} + r_{z_j}, \quad (20)$$

where r_{z_i} is the empirical covalent radius of an atom with atomic number z_i . This function provides a simplified yet effective representation of bond order, allowing us to gauge the bonding interactions within the molecular structure with respect to the proximity of the atoms.

The implicit overlap fragmentation method only change the step 4, the details of the changed fourth step has been spelled out in 4.1.

B FUNCTIONAL GROUPS SMARTS

In the initial phase of our fragmentation approach, we identify functional groups using the SMARTS pattern matching language. In Table 4, we present the complete list of SMARTS patterns utilized, which have been expanded upon from the default set found within the Open Force Field toolkit (Mobley et al., 2018; Wagner et al., 2024) (accessible at: <https://github.com/openforcefield/openff-fragmenter/blob/main/openff/fragmenter/data/default-functional-groups.json>).

Table 4: SMARTS patterns for functional groups employed in the preprocessing stage of fragmentation.

FUNCTIONAL GROUPS NAME	SMARTS
HYDRAZINE	[NX3:1][NX3:2]
HYDRAZONE	[NX3:1][NX2:2]
NITRIC OXIDE	[N:1]-[O:2]
AMIDE	[#7:1][#6:2](=[#8:3]), [NX3:1][CX3:2](=[OX1:3])[NX3:4]
AMIDE NEGATIVE ION	[#7:1][#6:2](-[O:-:3])
ALDEHYDE	[CX3H1:1](=[O:2])[#6:3]
SULFOXIDE	[#16X3:1]=[OX1:2], [#16X3+:1][OX1-:2]
SULFONYL	[#16X4:1](=[OX1:2])=[OX1:3]
SULFINIC ACID	[#16X3:1](=[OX1:2])[OX2H,OX1H0-:3]
SULFONIC ACID	[#16X4:1](=[OX1:2])(=[OX1:3])[OX2H,OX1H0-:4]
SULFINAMIDE	[#16X4:1](=[OX1:2])(=[OX1:3])([NX3R0:4])
PHOSPHINE OXIDE	[PX4:1](=[OX1:2])([#6:3])([#6:4])([#6:5])
PHOSPHONATE	[P:1](=[OX1:2])([OX2H,OX1-:3])([OX2H,OX1-:4])
PHOSPHATE	[PX4:1](=[OX1:2])([#8:3])([#8:4])([#8:5])
CARBOXYLIC ACID	[CX3:1](=[O:2])[OX1H0-,OX2H1:3]
NITRO	[NX3+:1](=[O:2])[O:-:3], [NX3:1](=[O:2])=[O:3]
ESTER	[CX3:1](=[O:2])[OX2H0:3]
TRI-HALIDE	[#6:1]([F,CL,I,Br:2])([F,CL,I,Br:3])([F,CL,I,Br:4])
HYDROXYL	[#8:1]-[#1:2]

C MERGE PROCESS OF FRAGMENTATION

During the third step of our fragmentation method, we introduce a strategy to enlarge substructures, ensuring that each initial fragment contains at least n_{\min} atoms, with n_{\min} being a predefined integer. To maintain permutation invariance for a molecule, we incorporate weights, W_{fs} , to guide the sequence of merging. The process is outlined in the pseudocode (Algorithm 1). The calculation of W is based on either Eq. 18 or Eq. 19. By considering the sum of bond orders to other groups, we assess each group’s centrality. The groups are then ordered first by the number of atoms they contain, followed by the summation of their bond orders, ensuring that the fragmentation merge process is permutation invariant when following this specified sequence. The algorithm then assists smaller groups in merging with others to achieve a size of at least n_{\min} atoms. Initially, we consider topologically adjacent groups with the fewest atoms. If a target group lacks topological neighbors, we proceed to merge based on the bond order from W . We introduce a threshold c_{is} that allows a group to remain isolated if it is significantly distant from others. It should be noted that isolated groups may not meet the minimum atom number requirement; however, they could be further expanded in the subsequent fragmentation step, depending on the chosen thresholds for c_{is} and c_w (refer to Sec. 4.1). Overall, this algorithm ensures a permutation invariant merging process.

D DISTRIBUTION OF FRAGMENTATION DATASET

Different parameters used in the fragmentation process can lead to a variety of hyperedges, which in turn result in distinct hypergraphs utilized for training our model. To illustrate the variances attributed to different fragmentation parameters or methods (such as BRICS implemented in RDKit (Degen

Algorithm 1 Pseudo code of fragmentation merge step.

```

864 Input: groups  $\{\mathcal{G}\}$ , minimum atoms number  $n_{\min}$ , maximum atoms number  $n_{\max}$ , Topological
865 bond order matrix  $B$ , isolated threshold  $c_{is}$ 
866  $m = |\{\mathcal{G}\}|$ 
867 Isolate groups  $\{\mathcal{G}^j\} = \{\}$ 
868 Calculate fragmentation bond order matrix  $W_{\mathcal{G}_i\mathcal{G}_j}$ .
869 Sort  $\{\mathcal{G}\}$  in descending order based on the following attributes: number of atoms,
870  $\sum_{\mathcal{G}', \mathcal{G}' \neq \mathcal{G}} W_{\mathcal{G}\mathcal{G}'}$ .
871 repeat
872   Pop last fragment as  $\mathcal{G}_k$  from  $\{\mathcal{G}\}$ 
873   for  $i = m - 1$  to 1 do
874      $a = \text{MAX\_INT}, \text{merge\_idx} = -1$ 
875     if any  $B_{ij} \geq 1, i \in \mathcal{G}_i, j \in \mathcal{G}_k$  and  $|\mathcal{G}_i| < a$  and  $a + |\mathcal{G}_i| \leq n_{\max}$  then
876        $a = |\mathcal{G}_i|, \text{merge\_idx} = i$ 
877     end if
878   end for
879   if  $\text{merge\_idx} == -1$  then
880     for  $i = m - 1$  to 1 do
881       if any  $W_{\mathcal{G}_i\mathcal{G}_k} \geq c_{is}$  and  $|\mathcal{G}_i| < a$  then
882          $a = |\mathcal{G}_i|, \text{merge\_idx} = -1$ 
883       end if
884     end for
885   end if
886   if  $\text{merge\_idx} \neq -1$  then
887     Merge  $\mathcal{G}_k$  to  $\mathcal{G}_{\text{merge\_idx}}$ 
888     Resort  $\{\mathcal{G}\}$  by the same priority and update  $W$ .
889   else
890     Add  $\{\mathcal{G}_k\}$  to  $\{\mathcal{G}^j\}$ 
891   end if
892 until  $|\{\mathcal{G}_k\}| \geq n_{\min}$ 
893  $\{\mathcal{F}\} = \{\mathcal{G}\} \cup \{\mathcal{G}^j\}$ 

```

et al., 2008; Landrum et al., 2020)), we use the QM9 dataset (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) to demonstrate how the data distributions attached to hypergraphs may change.

The impact of adjusting fragmentation parameters on the composition of hyperedges can be observed in Fig. 4. Altering the expansion threshold c_w within a certain range has a minimal effect on fragment expansion. However, when utilizing the Lendvay bond order (Eq. 18), fragments tend to comprise fewer atoms compared to when using the Exponential bond order (Eq. 19). This difference is likely due to the more gradual decline in the exponential function, which results in a greater cumulative contribution to the weights W_{fs} .

Our ablation study (Sec. 5.3) also includes a comparison with the BRICS fragmentation method. Fragments generated by the BRICS method are observed to contain significantly fewer atoms since this approach does not create overlapping regions between different fragments.

E PROOF OF SE(3) EQUIVARIANCE

The SE3Set consists of the basic modules including linear, depth-wise tensor product, gate activation and layer normalization. Here we will prove the SE3 equivariance for these modules. As the inputs of SE3Set are all invariant to translation in 3D Euclidean space, we only need to prove the $SO(3)$ equivariance.

Let $g \in SO(3)$ and the \mathcal{D} denotes the group representation of $SO(3)$. Besides, we denote the irreps feature as \mathbf{f} , and \mathbf{f}^l denotes the k^{th} feature vector in irreducible representations space of $SO(3)$ with l order.

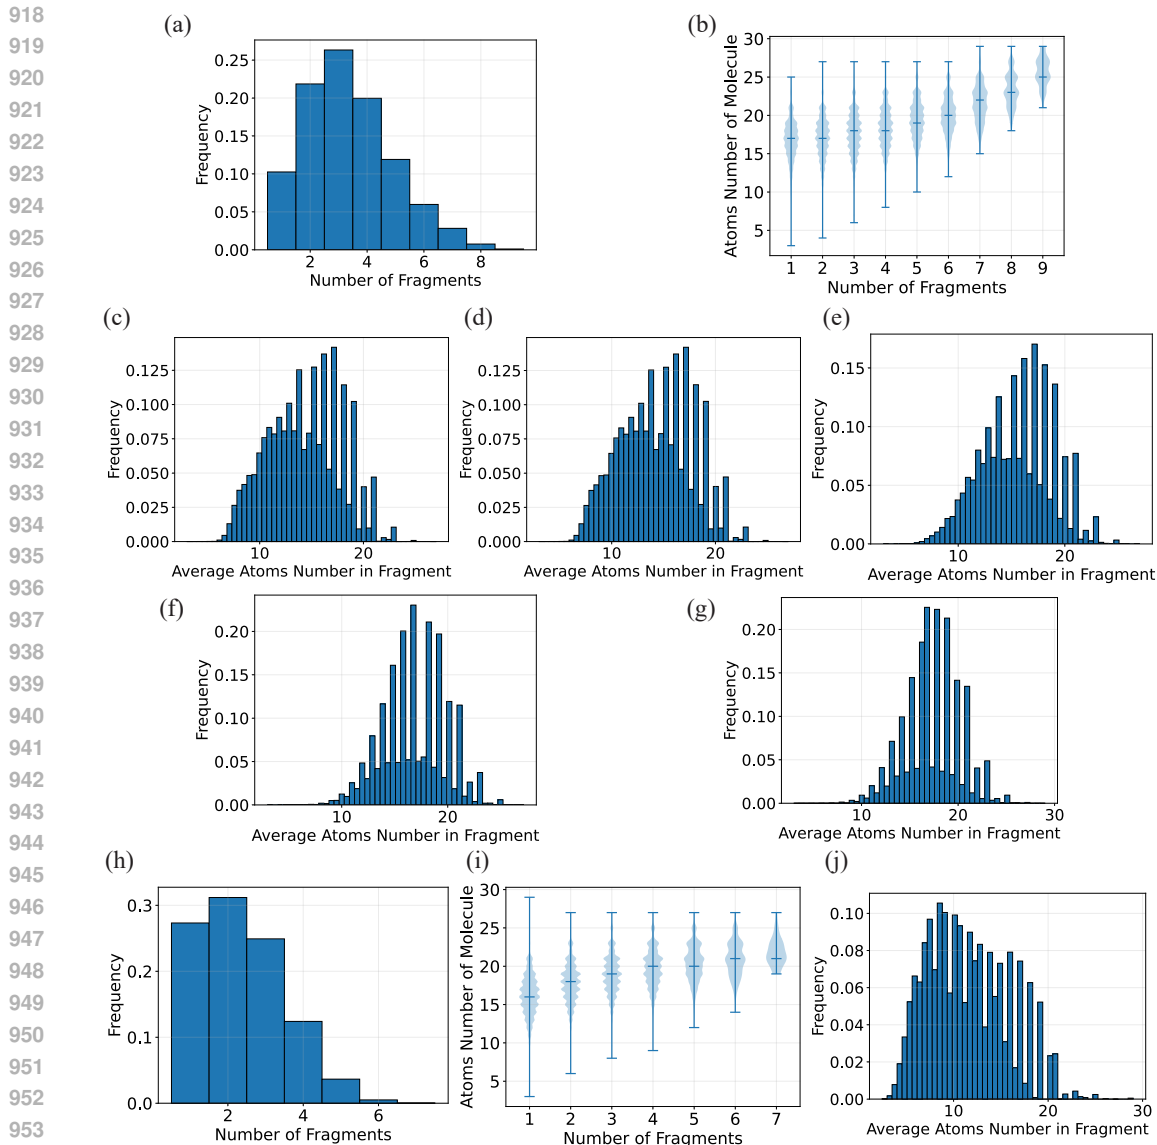


Figure 4: Distribution of fragments in QM9 dataset. (a) Fragment Count Distribution. The distribution remains consistent regardless of the value of c_w or the bond order calculation method employed. (b) Molecule Size vs. Fragment Count Distribution. Generally, the more atoms molecule has, the more fragments will generate. It is also invariant for c_w or bond order calculation scheme. Average Atom Count per Fragment Distribution (c) $c_w = 0.1$, (d) $c_w = 0.05$, (e) $c_w = 0.01$ for Lendvay bond order and (f) $c_w = 0.4$ and (g) $c_w = 0.2$ for exponential bond order, respectively. (h) BRICS Fragment Count Distribution. (i) BRICS Molecule Size vs. Fragment Count Distribution (j) BRICS Average Atom Count per Fragment.

Linear Linear module deploys separated linear operations for each l in the irreps feature. For each l , we consider the output channel \mathbf{f}_k^l . Then we have

$$\mathbf{f}_j^l = \text{Linear}(\mathbf{f}^l) = \sum_k w_{kj}^l \mathbf{f}_k^l \quad (21)$$

where w_{kj}^l denotes the linear combination weight. When acting $\mathcal{D}(g)$ at the input f^l , we can find that

$$(\mathbf{f}_j^l)' = \text{Linear}(\mathcal{D}(g)\mathbf{f}^l) \quad (22)$$

$$= \sum_k w_{kj}^l \mathcal{D}(g)\mathbf{f}_k^l \quad (23)$$

$$= \mathcal{D}(g) \sum_k w_{kj}^l \mathbf{f}_k^l \quad (24)$$

$$= \mathcal{D}(g)\mathbf{f}_j^l \quad (25)$$

Therefore, the linear module is $SO(3)$ equivariant.

Depth-wise Tensor-product Tensor product is an equivariant operation for $\mathcal{D}(g)$

$$\mathcal{D}(g)(\mathbf{f}_{k_1}^l \otimes \mathbf{f}_{k_2}^l) = (\mathcal{D}(g)\mathbf{f}_{k_1}^l) \otimes (\mathcal{D}(g)\mathbf{f}_{k_2}^l) \quad (26)$$

The depth-wise tensor product differs from the tensor product only in that one order l vector in our irreps feature depends only on one order l' feature, where l' is equal to or different from l . Hence the $SO(3)$ equivariance still holds for depth-wise tensor product.

Gate As the $l = 0$ vector is invariant for $\mathcal{D}(g)$, we have

$$\text{Gate}(\mathcal{D}(g)\mathbf{f}_k^0) = \text{Activation}(\mathcal{D}(g)\mathbf{f}_k^0) \quad (27)$$

$$= \text{Activation}(\mathbf{f}_k^0) \quad (28)$$

$$= \mathcal{D}(g)\text{Activation}(\mathbf{f}_k^0) \quad (29)$$

$$= \mathcal{D}(g)\text{Gate}(\mathbf{f}_k^0) \quad (30)$$

We use the non-linear output from the $\text{Activation}(\mathcal{D}(g)\mathbf{f}_{k_1}^0)$ (i.e. Sigmoid in SE3Set) as the weight which multiplies $l > 0$ vector to implement the gate function. Thus for $l > 0$, we also have

$$\text{Gate}(\mathcal{D}(g)\mathbf{f}_{k_1}^l) = \text{Activation}(\mathcal{D}(g)\mathbf{f}_{k_1}^0)\mathcal{D}(g)\mathbf{f}_{k_1}^l \quad (31)$$

$$= \text{Activation}(\mathbf{f}_{k_1}^0)\mathcal{D}(g)\mathbf{f}_{k_1}^l \quad (32)$$

$$= \mathcal{D}(g)\text{Gate}(\mathbf{f}_{k_1}^l) \quad (33)$$

So the gate function is $SO(3)$ equivariant for all l vector.

Layer Normalization For $l = 0$, as it is $SO(3)$ invariant, the general layer normalization is adapted

$$\text{LN}(\mathbf{f}_k^0) = \left(\frac{\mathbf{f}_k^0 - \mu}{\text{RMS}_C(\|\mathbf{f}_k^0\|)} \right) \gamma + \beta \quad (34)$$

where C denotes the channel number corresponding and $\|\cdot\|$ denotes the 2-norm for each channel vector. μ is the mean value of \mathbf{f}_k^0 . Learnable weight γ and learnable bias β are also deployed. The module is $SO(3)$ invariant as \mathbf{f}_k^0 is $SO(3)$ invariant. The layer normalization for $l > 0$ vector has the following form

$$\text{LN}(\mathbf{f}_k^l) = \left(\frac{\mathbf{f}_k^l}{\text{RMS}_C(\|\mathbf{f}_k^l\|)} \right) \gamma \quad (35)$$

We can prove the module is $SO(3)$ equivariant as

$$\text{LN}(\mathcal{D}(g)\mathbf{f}_k^l) = \left(\frac{\mathcal{D}(g)\mathbf{f}_k^l}{\text{RMS}_C(\|\mathcal{D}(g)\mathbf{f}_k^l\|)} \right) \gamma \quad (36)$$

$$= \left(\frac{\mathcal{D}(g)\mathbf{f}_k^l}{\text{RMS}_C(\|\mathbf{f}_k^l\|)} \right) \gamma \quad (37)$$

$$= \mathcal{D}(g)\text{LN}(\mathbf{f}_k^l) \quad (38)$$

F CONCEPTS OF IRREPS FEATURES AND TENSOR PRODUCT

Irreps features The SE3Set model utilizes the special orthogonal group $SO(3)$ to capture three-dimensional rotational symmetries in molecular structures. This approach is similar to

Equiformer (Liao & Smidt, 2022; Liao et al., 2023). It employs irreducible representations (irreps) of $SO(3)$, parameterized by an integer l , which correspond to spherical harmonics (SH) functions Y_l^m . These functions imbue feature vectors with rotational information, ensuring the model’s equivariance to rotations and enabling consistent geometric property analysis. This approach is key to the model’s ability to accurately represent and predict molecular and other rotationally invariant systems.

Tensor product To boost the model’s expressive power, we consider interactions between irrep features of different angular momenta l through the tensor product, which merges two irreps l_1 and l_2 into a new irrep with angular momentum l_3 . This is achieved using Clebsch-Gordan coefficients in an expansion weighted by w_{m_1, m_2} .

$$\begin{aligned} f_{m_3}^{l_3} &= (f_{m_1}^{l_1} \otimes f_{m_2}^{l_2})_{m_3} \\ &= \sum_{m_1, m_2} w_{m_1, m_2} C_{l_1, m_1, l_2, m_2}^{l_3, m_3} f_{m_1}^{l_1} f_{m_2}^{l_2}. \end{aligned} \quad (39)$$

To reduce complexity, a depth-wise tensor product \otimes^{DTP} is adopted from the Equiformer (Liao & Smidt, 2022; Liao et al., 2023), utilizing internal weights to streamline computations. Input-dependent tensor product weights are denoted as \otimes_w^{DTP} , ensuring computational efficiency while preserving equivariance for feature interactions.

G RESULTS OF MD17

The MD17 dataset (Chmiela et al., 2017) features a wide variety of molecular configurations simulated at 500 K, with high-resolution trajectories and labeled with energies and forces from the PBE+vdW-TS method (Perdew et al., 1996; Tkatchenko et al., 2012). SE3Set’s performance on this dataset is shown in Table 5. SE3Set outperforms Equiformer in accuracy, highlighting its refined force calculation capabilities. In small molecular systems, higher-order many-body interactions are less pronounced, and as a result, SE3Set does not significantly outperform other state-of-the-art (SOTA) models.

Table 5: A comparison of Mean Absolute Errors (MAEs) across various benchmarked models. SE3Set is trained on the MD17 dataset with a configuration of 950 training samples and 50 validation samples. Bolding shows the best model and underlining shows the second best model and the underlining tilde shows third best model. The MAEs reflect the precision of energy predictions in units of kcal/mol and forces in units of kcal/(mol·Å).

		SCHNET	DIMENET	PAIINN	ET	GEMNET	NEQUIP ($l=3$)	VISNET	QUINNET	EQUIFORMER	SE3SET
ASPIRIN	ENERGY	0.37	0.204	0.167	0.123	-	0.131	0.116	<u>0.119</u>	<u>0.122</u>	0.130
	FORCE	1.35	0.499	0.338	0.253	0.217	0.184	0.155	0.145	<u>0.152</u>	<u>0.153</u>
ETHANOL	ENERGY	0.08	0.064	0.064	0.052	-	<u>0.051</u>	<u>0.051</u>	0.050	<u>0.051</u>	0.054
	FORCE	0.39	0.230	0.224	0.109	0.085	0.071	0.060	0.060	0.067	<u>0.062</u>
MALONALDEHYDE	ENERGY	0.13	0.104	0.091	0.077	-	0.076	<u>0.075</u>	0.078	0.074	0.074
	FORCE	0.66	0.383	0.319	0.169	0.155	0.129	<u>0.100</u>	0.097	0.125	<u>0.103</u>
NAPHTHALENE	ENERGY	0.16	0.122	0.116	0.085	-	0.113	0.085	<u>0.101</u>	0.085	0.113
	FORCE	0.58	0.215	0.077	0.061	0.051	0.039	0.039	0.039	0.046	0.039
SALICYLIC ACID	ENERGY	0.20	0.134	0.116	0.093	-	0.106	0.092	0.101	<u>0.099</u>	0.108
	FORCE	0.85	0.374	0.195	0.129	0.125	<u>0.090</u>	<u>0.084</u>	0.080	<u>0.090</u>	<u>0.090</u>
TOLUENE	ENERGY	0.12	0.102	0.095	0.074	-	0.092	0.074	<u>0.080</u>	<u>0.085</u>	0.093
	FORCE	0.57	0.216	0.094	0.067	0.060	<u>0.046</u>	0.039	0.039	0.048	<u>0.046</u>
URACIL	ENERGY	0.14	0.115	0.106	<u>0.095</u>	-	0.104	0.095	<u>0.096</u>	0.099	0.103
	FORCE	0.56	0.301	0.139	0.095	0.097	0.076	0.062	0.062	0.076	<u>0.067</u>

H COMPLEXITY ANALYSIS

The computational complexity mainly depends on V2E module and E2V module.

The V2E module aggregates the information from each atom in one fragment to generate atom-wise hyperedge feature. Considering the system are split into m fragments and each fragment has n_i atoms, this module includes $\sum_i^m n_i(n_i - 1)$ pair-wise message for attention architecture. Actually, this

Table 6: The inference speed (unit: iterations/s) of SE3Set with different c_w or r_c . The tests on MD22 dataset only use the AT-AT-CG-CG subset.

	DATASET	EXPLICIT c_w			IMPLICIT r_c		
		0.1	0.05	0.01	4.0	5.0	6.0
INFERENCE SPEED	QM9	11.60	11.54	11.37	-	-	-
	MD22 (AT-AT-CG-CG)	2.62	2.29	1.68	2.92	2.90	2.87

number depends on the fragment hyperparameter n_{min} and n_{max} , in particular for explicit overlap method, also on c_w . For explicit overlap method, this module would have higher complexity than implicit overlap method because the explicit overlap has more average atoms number in one fragment.

For each atom, the E2V module aggregates all the hyperedge features that an atom possesses. Thus the calculated pair-wise message number depends on the number of fragments per atom shared, corresponding to the introduced overlap degrees. Then the computational complexity of explicit overlap method in this modules depends on the c_w for explicit overlap method and r_c for implicit overlap method.

To give a general approximate complexity, we consider a system with N atoms. When using explicit overlap method, if the system has m fragments with average n_{exp} atoms in one fragment, the V2E module’s complexity will be $O(mn_{exp}^2)$ according to the previous analysis. For E2V module, the number of fragments to which each atom belongs on average can be represented as mn_{exp}/N . Thus the complexity will be $O(mn_{exp})$ when considering N atoms.

Similarly, for the implicit overlap method, the complexity of V2E module will be $O(n_{\mathcal{F}}n_{imp}N)$, where $n_{\mathcal{F}}$ represents the average number of neighbor fragments within cutoff per atom and n_{imp} denotes the average atoms in one fragment. And the complexity of E2V module is $O(n_{\mathcal{F}}N)$.

Compare to the explicit overlap method, the implicit overlap method has lower computational complexity. As we can approximate that $n_{\mathcal{F}}n_{imp} \approx n_{exp}$ (overlap degree is similar) and determine $mn_{imp} = N$ (implicit overlap method has same fragment number m but without extra count of atoms in n_{imp}). On the other hand we can determine $mn_{exp} \geq N$ as the explicit overlap will count atoms more times. Hence the complexity of explicit overlap method has higher complexity than implicit overlap on V2E Module according to the deduction below

$$mn_{exp}^2 \geq Nn_{exp} \approx n_{\mathcal{F}}n_{imp}N \quad (40)$$

$$mn_{exp} \approx mn_{\mathcal{F}}n_{imp} \approx n_{\mathcal{F}}N \quad (41)$$

To better support the analysis, we have tested the inference speed of SE3Set on part of QM9 and MD22 datasets as shown in Table 6. Larger fragments, with smaller c_w , reduce speed, more so in larger molecules. Speed drops with increased fragment size (smaller c_w or larger r_c), and the explicit overlap method is slower than the implicit, leading to our preference for the latter in MD22. This highlights the need to optimize fragment size for a trade-off between detailed interaction capture and speed, particularly in big molecular systems. Additionally, the explicit overlap method is slower than implicit overlap method, which confirms our previous analysis.

Moreover, we have evaluates the inference time with different baseline models. The results in Figure 5 shows SE3Set runs slightly slower than the other models. This is reasonable as we have included more pairwise node information in our hypergraph. The improvement in our model’s performance indicates that our model requires further optimization to achieve a better balance between efficiency and performance.

I TRAINING DETAILS

This section outlines the training specifics, encompassing the fragmentation parameters, SE3Set hyperparameters, and certain implementation nuances utilized in our experimental setup.

Our dataset construction is founded on PyTorch Geometric (Fey & Lenssen, 2019) augmented with our fragmentation process (Sec. 4.1). Due to inconsistencies in molecular topology identified through

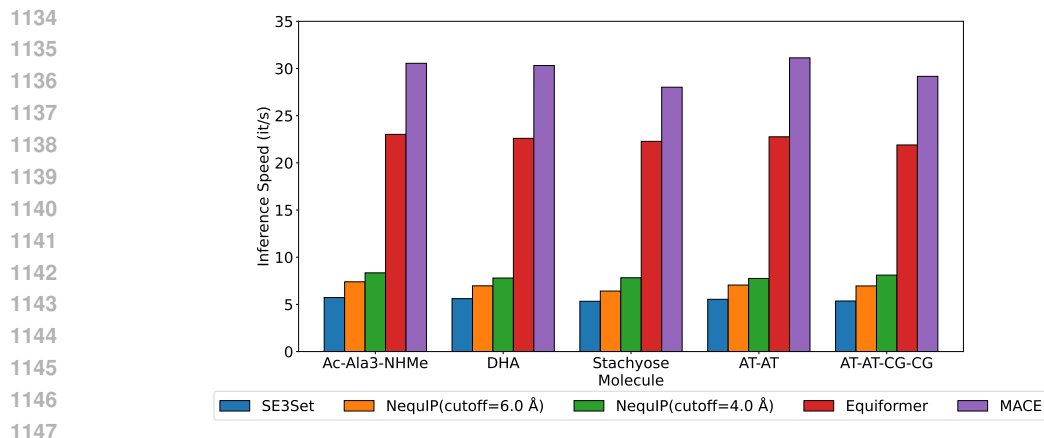


Figure 5: Inference speeds of different models on the MD22 dataset in units iteration/s. The average inference speeds were calculated for each molecule dataset. All tests were performed on a single Tesla A100 80G. The SE3Set includes 3 layers and different cutoff is set (4 Å and 6 Å) for NeuqIP. The other models use parameter configurations that corresponded to the results reported in their respective studies.

RDKit’s sanitization routine (Landrum et al., 2020), 1,403 data points were excised from the original dataset. We designated 110,000 data points for the training set and 10,000 for the validation set, selected at random.

We used an explicit overlap scheme on the QM9 and MD17 datasets because of their relatively small molecular systems. We implemented two distinct schemes for calculating fragment bond orders, following either Eq. 18 or Eq. 19. The parameters for fragmentation are detailed below. Given that the MD17 molecules are relatively small, the merge step in the fragmentation process was not actually utilized. However, we present the fragmentation parameters here for the sake of completeness.

Besides, we adopt an implicit overlap scheme on the MD22 dataset to reduce computational resource consumption. The details of cutoff r_c introduced in 4* can be found in Table 8.

On QM9 and MD17 dataset, our model was trained using a single Tesla V100 GPU with 32GB of memory, except for the 6-layer model employing the exponential bond order on QM9 dataset, which was trained on two Tesla V100 GPUs with 32GB each. For MD22 dataset, our model was trained on a single Tesla A100 GPU with 80GB of memory.

We selected $l = 2$ for our irreducible representations (irreps) feature, which includes both node and hypergraph features. For the radial basis function (RBF), we utilized Gaussian basis functions or Bessel basis functions for the QM9 dataset (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) and exponential basis functions for the MD17 dataset and MD22 dataset (Chmiela et al., 2017). Details could be found in Table 7 and Table 9.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Table 7: Hyper-parameters for fragmentation. The expand threshold does not work for models training on MD22 dataset because they adapt implicit overlap scheme.

BOND ORDER METHODS	BOND ORDER BY LENDVAY (18)	FRAGMENTATION BY EXPONENTIAL (19)
MINIMUM ATOMS NUMBER n_{\min}	2	2
MAXIMUM ATOMS NUMBER n_{\max}	6	6
ISOLATED THRESHOLD (c_{is})	0.1	0.4
EXPAND THRESHOLD (c_w)	0.1	0.2, 0.4

Table 8: Hyper-parameters for step 4* of implicit overlap scheme in MD22 experiments.

MOLECULES	AC-ALA3-NHME	DHA	STACHYOSE	AT-AT	AT-AT-CG-CG
DISTANCE CUTOFF r_c (Å)	5.0	4.0	4.0	6.0	6.0

Table 9: Hyper-parameters for training SE3Set model. In the context of hyperparameter settings for dimensions, the symbols e and o are used to denote even and odd parity, respectively.

HYPER-PARAMETERS	QM9	MD17	MD22
	VALUE OR DISCRIPTIONS	VALUE OR DISCRIPTIONS	VALUE OR DISCRIPTIONS
OPTIMIZER n_{\min}	ADAMW	ADAMW	ADAMW
LEARNING RATE SCHEDULER	COSINE	COSINE	COSINE
WARM UP EPOCHS n_{\max}	5	10	10
MINIMUM LEARNING RATE	1.0×10^{-6}	1.0×10^{-6}	1.0×10^{-6}
BATCH SIZE	32, 128	8	8
NUMBER OF EPOCHS	400	1500	1500
WEIGHT DECAY	5.0×10^{-3}	1.0×10^{-6}	1.0×10^{-6}
DROPOUT RATE	0.1	0.0	0.0
RBF CUTOFF (Å)	42.0	MAX DISTANCE OF USED ATOM PAIRS	
NUMBER OF BASIS	128(GAUSSIAN), 8(BESSEL)	32(EXPONENTIAL)	32(EXPONENTIAL)
NUMBER OF BLOCKS	3, 6	3, 6	6
NODE EMBEDDING DIMENSION		[(128, 0e), (64, 1o), (32, 2e)]	
HYPEREDGE EMBEDDING DIMENSION		[(128, 0e), (64, 1o), (32, 2e)]	
ATTENTION HEAD DIMENSION		[(32, 0e), (16, 1o), (8, 2e)]	
FEED FORWARD DIMENSION		[(384, 0e), (192, 1o), (96, 2e)]	
OUTPUT FEATURE DIMENSION		[(512, 0e)]	