

E(n) Equivariant Message Passing Cellular Networks

Veljko Kovac¹ Erik J. Bekkers^{1,2} Pietro Liò³ Floor Eijkelboom^{1,2,4}

Abstract

This paper introduces E(n) Equivariant Message Passing Cellular Networks (EMPCNs), an extension of E(n) Equivariant Graph Neural Networks to CW-complexes. Our approach addresses two aspects of geometric message passing networks: 1) enhancing their expressiveness by incorporating arbitrary cells, and 2) achieving this in a computationally efficient way with a *decoupled EMPCNs* technique. We demonstrate that EMPCNs achieve close to state-of-the-art performance on multiple tasks *without* the need for steerability, including many-body predictions and motion capture. Moreover, ablation studies confirm that *decoupled EMPCNs* exhibit stronger generalization capabilities than their non-topologically informed counterparts. These findings show that EMPCNs can be used as a scalable and expressive framework for higher-order message passing in geometric and topological graphs.

1. Introduction

Graph Neural Networks (GNNs) are a type of neural network designed to operate on graph-structured data, capturing relationships and interconnections between data points. They are particularly useful in domains where data is naturally represented as relations, such as social networks, molecular structures, and recommendation systems. The most common type of GNNs is a Message Passing Neural Network (MPNN), in which the representation of nodes are updated by aggregating information in their 1-hop neighborhoods. While GNNs have shown notable success in applications such as graph classification (Kipf & Welling, 2017), link prediction (Zhang & Chen, 2018), and algorithmic reasoning (Veličković & Blundell, 2021), their expressiveness

is inherently limited by their design to capture only pairwise interactions. Some examples where modeling multi-body relationships is crucial include protein-protein interaction networks (Koyutürk et al., 2004; Girvan & Newman, 2002), which require the identification of protein complexes through group interactions, and social networks (Jiang et al., 2010), where the dynamics of group behavior provide essential insights.

This limitation of GNNs is formally discussed by Xu et al. (2019) and Morris et al. (2019), where they showed that the local neighborhood aggregation used by GNNs is at most as powerful as the Weisfeiler-Lehman (WL) test in distinguishing non-isomorphic graphs (Weisfeiler & Leman, 1968). Several studies have underscored the importance of distinguishing graphs that are indistinguishable by the 1-WL test, pointing out that, for example in molecular structures, distinct local structures around atoms should correspond to distinct representations (Maron et al., 2019; Chen et al., 2019; Dwivedi et al., 2023). One approach to overcome this problem of recognizing higher-dimensional graph features like cliques or cycles, is by generalizing MPNNs to more elaborate topological space, i.e. to take into account these features explicitly. In Bodnar et al. (2021b), MPNNs are extended to operate on the clique complex of the graph – i.e. assigning simplices to each clique in the graph – resulting in an architecture that is provably more powerful than the WL test by creating more complex neighborhood relationships. However, as simplicial complexes have a rigid combinatorial structure that restricts the range of lifting transformations, this work was later extended to CW-complexes – assigning features to arbitrary cells in the graph – defining a strictly more expressive MPNN again (Bodnar et al., 2021a), provably not less powerful than the 3-WL test.

A generalized version of the Weisfeiler-Lehman framework has been proposed for geometric graphs, where the nodes correspond to points in Euclidean space. This framework accounts for a stronger notion of geometric isomorphism, considering spatial symmetries (Joshi et al., 2023). This so-called Geometric Weisfeiler-Leman (GWL) is used to qualify the expressivity of MPNNs with augmented discriminative capabilities of GNNs beyond the 1-WL test by integrating geometric information of the data, e.g. as done by Satorras et al. (2021) and Brandstetter et al. (2021). For instance, Satorras et al. (2021) proposed E(n) Equivariant

¹University of Amsterdam ²AMLab ³Department of Computer Science and Technology, University of Cambridge ⁴UvA-Bosch Delta Lab. Correspondence to: Veljko Kovac <veljko.kovac@student.uva.nl>.

Proceedings of the Geometry-grounded Representation Learning and Generative Modeling at the 41st International Conference on Machine Learning, Vienna, Austria. PMLR Vol Number 251, 2024. Copyright 2024 by the author(s).

Graph Neural Networks (EGNN), in which the messages passed in the MPNN are conditioned on the distance between the nodes, making the model equivariant to rotations, reflections, and translations. This kind of approaches have shown to be highly effective in many domains (Weiler et al., 2018; Gasteiger et al., 2021; Bekkers et al., 2024).

Recently, multiple approaches to combining geometric *and* topological information have been proposed (Eijkelboom et al., 2023; Liu et al., 2024; Anh Trang et al., 2024). In Eijkelboom et al. (2023), E(n) Equivariant Message Passing Simplicial Networks (EMPSNs) are proposed, a method that combines simplicial message passing with geometric attributes invariant information such as angles, volume, and area, which cannot be captured when focusing only on pairwise interactions. Although this method achieved state-of-the-art results, it presents two main drawbacks. Firstly, even though the use of simplicial complexes adds higher-order topological information to the model, these structures are not the most elaborate forms available. More sophisticated structures such as CW complexes can offer a more nuanced representation of topological features, which could be especially beneficial in fields like molecular chemistry where specific topological configurations are crucial. This is specifically the case in \mathbb{R}^3 , because due to the construction of EMPSNs, the highest-dimensional object will be a tetrahedron, i.e., a 4-body interaction. Secondly, the approach involves increased computational complexity due to the higher-dimensional message passing required in each layer. The authors use Vietoris-Rips complexes – a higher-dimensional analogue to a radius graph – and as such introduce a typically infeasible amount of extra features due to the combinatorial nature of such complexes, making this approach hard to use for processing large graphs or datasets where real-time performance is crucial.

In this work, we propose E(n) Equivariant Message Passing Cellular Networks (EMPCNs), a generalization of E(n) Equivariant Message Passing Simplicial Networks to cellular complexes. Our contributions are the following:

- Targetting the limited expressivity of EMPSNs, EMPCNs operate on CW-complexes, hence enabling the identification of a more general set of geometric invariants across a broader range of topological objects, making them strictly more expressive than their simplicial counterparts.
- Addressing the increased computational costs of higher-order message passing, we introduce an efficient method to integrate higher-order topological information in geometric datasets inspired by EMPCNs. These *decoupled* EMPCNs provide a simple, cellular add-on to existing MPNNs without increasing computational complexity.

- We show that EMPCNs achieve close to state-of-the-art performance on two tasks, outperforming many elaborate, steerable methods. Moreover, we demonstrate that the *decoupled* EMPCN model provides improved performance without additional computational costs and leads to stronger generalization, providing a step towards scalable higher-order message passing.

2. Background

Symmetry-Informed Deep Learning Considering the intrinsic structure of data – specifically in terms of symmetries – is a key design principle in Geometric Deep Learning (Bronstein et al., 2021), which has led to many successes in recent years (Zitnik et al., 2018; Gligorijević et al., 2021). Symmetries are mathematically described by groups, and a function respects such a symmetry – formally: is equivariant with respect to the relevant group – if it commutes with its group action, i.e. if

$$f(g \cdot x) = g \cdot f(x) \text{ for all } g \in G. \quad (1)$$

Intuitively, this means that first applying a transformation $g \in G$ and then evaluating the function is equivalent to first applying the function and then performing the transformation. A function that produces the same result when applied to both transformed and untransformed elements is referred to as group invariant. In the setting of graph prediction, many different forms of symmetry come into play, e.g. permutations of the graphs, rotations of the graph in space (if it is embedded in some geometric space), or topological features such as rings.

Message Passing Neural Networks Message Passing Neural Networks (MPNNs) have emerged as a popular class of models for graph representation learning (Gilmer et al., 2017). MPNNs update node representations through their respective 1-hop neighborhoods, i.e.

$$h_i^{\ell+1} = \text{MPNN}(h_i^\ell, \{\{h_j^\ell\}\}_{j \sim i}), \quad (2)$$

where $\{\{h_j^\ell\}\}_{j \sim i}$ denotes the multi-set of neighbours of node i , and ℓ denotes the layer. An MPNN-layer can be decomposed into two main steps:

1. First, a message function $\text{Mes}(h_i, h_j) =: m_{ij}$ computes the messages from nodes j to nodes i , which are then aggregated in a permutation invariant manner to find the aggregated message to i , i.e. $m_i := \text{Agg}_{j \sim i} m_{ij}$.
2. An update function $\text{Update}(h_i, m_i) =: h'_i$ which updates the representation h_i based on the aggregated message.

As mentioned in Section 1, there are two major ways to augment MPNNs: (1) by using the underlying geometry of

the space in which the graph is embedded (e.g., the coordinates of atoms in a molecule), and (2) by using topological or structural information of the graph. We refer to these approaches as Geometric MPNNs and Topological MPNNs, respectively.

Geometric MPNNs If the relevant graph is embedded in some geometric space, such as when each node is associated with coordinates $x_i \in \mathbb{R}^n$, one can leverage the underlying geometry in the predictions, as discussed by [Duval et al. \(2024\)](#). Importantly, it is now possible to utilize more symmetries since rotating, translating, or flipping the graph leads to predictable changes. Arguably, the simplest approach to do this is E(n) Equivariant Graph Neural Networks (EGNNs), where the geometric information is added in the message function by conditioning it on the distance between two nodes, i.e. E(n) *invariant* information:

$$m_{ij} = \text{Mes}(h_i, h_j, \|x_i - x_j\|_2^2). \quad (3)$$

Note that indeed distance is preserved by applying any symmetry of the Euclidean group.

Topological MPNNs Topological MPNNs mitigate the issue of MPNNs not being able to learn higher-order structures directly by learning features on a more elaborate topological space. Two such topological spaces are simplicial and cellular complexes.

A simplex is the generalization of a triangle to arbitrary dimension, i.e. a k -dimensional simplex is defined by a set of $k + 1$ fully connected points. As such, a graph consists of 0-dimensional simplices (nodes) and 1-dimensional simplices (edges), but can be trivially lifted by assigning a k -dimensional simplex of each clique of $k + 1$ nodes. Though useful, in many domains graphs are not best understood in terms of simplices, e.g. in molecular prediction tasks where triangles are much less common than rings. A generalization of simplicial complexes is given by CW complexes, which considers arbitrary cells in the graphs, hence strictly generalizing simplicial complexes.

More precisely, a CW complex is built from *cells* of various dimensions. A cell is a space homeomorphic to an open n -dimensional disk, which informally one can understand as any object obtainable from continuously transforming an n -sphere. The construction of the complex involves a hierarchical glueing process:

1. Start with a set of vertices (0-cells).
2. Attach edges (1-cells) by gluing endpoints of line segments to these vertices.
3. Attach higher-dimensional cells by mapping the boundary of each n -cell, S^{n-1} , to the $(n - 1)$ -dimensional

skeleton $X^{(n-1)}$ via a continuous map $\phi : S^{n-1} \rightarrow X^{(n-1)}$.

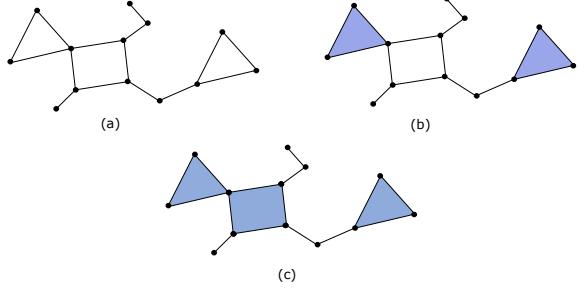


Figure 1. Visualization of a graph (a), a simplicial complex (b), and a cellular complex (c). As observed, a simplicial complex cannot represent arbitrary polygons.

This allows for the inclusion of complex structures, such as rings, in molecular graphs as it can be seen in Figure 1.

To learn features on these topological spaces, a more elaborate set of messages are sent as introduced by [Bodnar et al. \(2021a;b\)](#). Let $\sigma \prec \tau$ denote that σ is a boundary of τ , such as a node of an edge or an edge of a ring. With this notation, one can differentiate four types of relationships:

1. $\mathcal{B}(\sigma) = \{\tau \mid \tau \prec \sigma\}$, i.e. all boundaries of σ , called the *boundaries*.
2. $\mathcal{C}(\sigma) = \{\tau \mid \sigma \prec \tau\}$, i.e. all cells σ is a boundary of, called the *co-boundaries*.
3. $\mathcal{N}_\downarrow(\sigma) = \{\tau \mid \exists \delta \text{ s.t. } \delta \prec \sigma \text{ and } \delta \prec \tau\}$, i.e. the set of cells that share a boundary σ , called the *lower adjacencies*.
4. $\mathcal{N}_\uparrow(\sigma) = \{\tau \mid \exists \delta \text{ s.t. } \sigma \prec \delta \text{ and } \tau \prec \delta\}$, i.e. the set of cells that share a co-boundary σ , called the *upper adjacencies*.

The complex adjacencies defined above equip CW networks with a message passing scheme encompassing such higher-order topological features, such as vertices (0-cells), edges (1-cells), and cycles (2-cells). Consequently, updates incorporate not merely a single message, but four distinct types. However, [Bodnar et al. \(2021a\)](#) demonstrated that retaining only boundary and upper adjacencies achieves equivalent expressivity, and hence the typical messages considered are:

$$m_{\mathcal{B}}(\sigma) = \text{Agg}_{\tau \in \mathcal{B}(\sigma)} \text{Mes}_{\mathcal{B}}(h_\sigma, h_\tau) \quad (4)$$

$$m_\uparrow(\sigma) = \text{Agg}_{\substack{\tau \in \mathcal{N}_\uparrow(\sigma), \\ \delta \in \mathcal{C}(\sigma, \tau)}} \text{Mes}_\uparrow(h_\sigma, h_\tau, h_\delta) \quad (5)$$

Taking molecules as an example, the first type of communication outlines how atoms communicate with bonds, and

how bonds communicate with rings, whereas the second type describes interactions between atoms connected by a bond, and bonds that form a ring.

Similar to the standard message passing framework, the update step considers the incoming messages and updates the features through a more elaborate function:

$$h'_\sigma = \text{Update}(h_\sigma, m_{\mathcal{B}}(\sigma), m_\uparrow(\sigma)). \quad (6)$$

Geometric & Topological MPNNs Recently, Eijkelboom et al. (2023) combined geometric and topological GNNs by realizing that there is a direct relationship between the two: a geometric space allows one to define higher-order topological objects geometrically – e.g. through a Vietoris-Rips complex –, while higher-order topological objects enable multi-body geometry in MPNNs. They propose E(n) Equivariant Message Passing Simplicial Networks, essentially combining EGNNs with simplicial message passing by defining geometric features on simplicial complexes, specifically volumes, angles, and distances, i.e. E(n) invariant information. This work was later extended to the steerable case by Liu et al. (2024), which defines steerable simplicial message passing using clifford algebras.

3. Methodology

Building upon the E(n) Equivariant Message Passing Simplicial Networks, we propose E(n) Equivariant Message Passing Cellular Networks (EMPCNs), the generalized counterpart to EMPSNs of cellular complexes. With EMPCNs, we achieve the following two things:

- 1. Enhanced Expressivity:** Similar to how normal and simplicial message passing were made geometrically equivariant through conditioning on E(n) invariant information, we extend these invariants to the cellular case. Since CW complexes strictly generalize simplicial complexes, this defines a more expressive MPNN.
- 2. Improved Scalability:** Higher-order message passing networks often struggle with scalability due to the complexity of higher-order lifting, such as EMPSNs using Vietoris-Rips complexes to construct simplices. We address this by introducing a simplified version of EMPCNs, called *decoupled EMPCNs*, that explicitly learn this topological information without added computational complexity.

3.1. E(n) Equivariant Cellular Message Passing

As our method serves as a generalization of E(n) Equivariant Simplicial Message Passing, our primary focus is on making the topological objects and invariants used more

expressive, rather than introducing a new method of conditioning. That is, e.g. for the boundary update

$$m_{\mathcal{B}}(\sigma) = \underset{\tau \in \mathcal{B}(\sigma)}{\text{Agg}} \text{Mes}_{\mathcal{B}}(h_\sigma, h_\tau, \text{Inv}(\sigma, \tau)),$$

our aim is to generalize $\text{Inv}(\sigma, \tau)$, denoting the geometric invariants.

3.1.1. INVARIANTS

Barycentric Subdivision Crucially, we make the observation that we can understand the cells in a CW complex as a set of simplices glued together. Formally, we can define all their geometric invariants using them. To explain this, we refer to Barycentric subdivision (Davis & Okun, 2004), a technique typically employed to decompose a simplex into smaller simplices, which can also be applied to CW complexes. Specifically, a regular CW complex can be subdivided into a simplicial complex, resulting in a structure where the simplices are efficiently combined into closed cells, as shown in Figure 2.

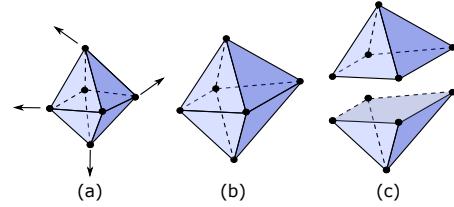


Figure 2. Change of invariants after position updates: a) Initial graph with arrows indicating the future displacement of the respective nodes. b) Displaced graph showing the updated positions of the nodes. c) Cell decomposition into two simplices.

This is key to our approach, as it allows us to directly generalize our existing invariants in EMPSN to the cellular case.

Generalized Volume and Area In the context of graphs embedded within an n -dimensional Euclidean space, the convex hull of some set S is defined as the smallest convex set containing all the points in S . This set – and its volume – are typically obtained through the Quickhull algorithm (Barber et al., 1996), an efficient method in computational geometry that relies on decomposing the hull into simplices, as their volumes are easy to compute, and then summing these subvolumes. For a n -simplex ξ defined by vertices $\xi = (x_0, x_1, \dots, x_n)$, the volume is given by:

$$\text{Vol}(\xi) = \frac{1}{n!} |\det(x_1 - x_0 \quad \dots \quad x_n - x_0)|$$

As the total volume of the hull is the sum of the volumes of all the simplices, note that we trivially obtain the volume of a single simplex through this method, hence strictly generalizing EMPSNs.

Following a similar argument, the surface area of the hull is computed by dividing its area into $(n - 1)$ -simplices, and again the total surface area is obtained by summing the areas of all $(n - 1)$ simplices forming the hull’s surface. Note that again we trivially obtain the simplicial case from here.

Other Invariants For all other types of message passing, one could use the invariants introduced by Eijkelboom et al. (2023), which range from distances between two nodes to calculating the dihedral angles for edge-to-edge communication, as these are not affected by a change in topological space. A detailed list of invariants for each experiment can be found in Appendix B.

Scalable Convex Hull Estimation Computing the convex hull and its invariants during training can be computationally intensive, especially when nodes are moving and their positions are constantly updated, requiring re-calculation of invariants in every iteration. To address this, we can approximate these invariants by assuming the radius of the ring to be half of the maximum pairwise distance between nodes in the ring. By doing so, we are essentially taking the radius of a sphere in \mathbb{R}^3 that includes all the points in the ring. This simplification allows for faster invariant calculations, making it more feasible to track changes in invariants due to node movement.

3.1.2. GEOMETRIC WEISFEILER-LEHMAN

As introduced in Section 1, the Geometric Weisfeiler-Leman framework can be used to compare the expressivity of different geometric MPNNs. The three axes of the plot are the *body order*, *tensor order*, and *depth*, which corresponds to whether equivariant layers are being used. Using this framework as a reference, we position our proposed method, EMPCNs, at least as high as EGNNs (Satorras et al., 2021) and PaiNN (Schütt et al., 2021), given that our features transform equivariantly within the Cartesian system. However, our increased expressiveness comes from the axis corresponding to body order, which denotes the number of nodes needed to compute the geometric invariant information. Since our method calculates the n -dimensional volume, it requires $n + 1$ points and thus can be placed arbitrarily high in terms of body order. Even when restricting ourselves to \mathbb{R}^3 , our method remains more expressive than EGNN, demonstrating its ability to capture complex geometric information.

3.2. Decoupled EMPCNs

As mentioned in Section 1, the most straightforward approach when dealing with geometric graphs is to interconnect all nodes. This allows the network to identify relevant edges on its own, similar to a Graph Transformer architecture (Rampášek et al., 2022; Kreuzer et al., 2021; Dwivedi & Bresson, 2021). Clearly, the number of cells in a graph

grows combinatorially with more edges, meaning that EMPCNs cannot be applied in a feasible way.

To this end, we propose a scalable cellular message passing framework called *decoupled EMPCNs*, which essentially serves as an add-on to existing fully connected geometric MPNNs. *Decoupled EMPCNs* operate on 1) a node graph facilitating direct node communication and 2) a cellular lifted graph derived from the input graph, enabling higher-order message passing based on the input’s graph topology. The decoupling procedure can be seen in Figure 3.

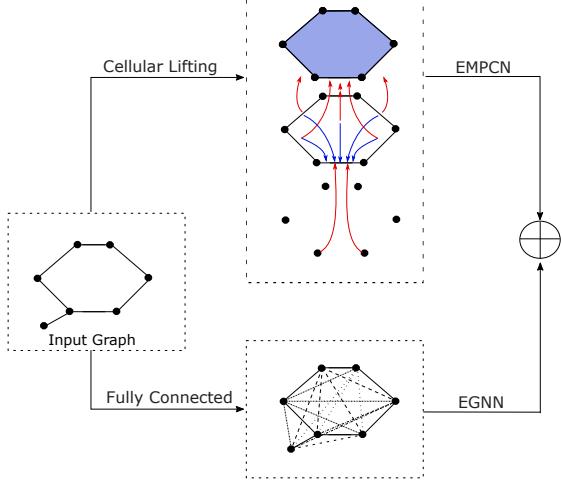


Figure 3. Pipeline of *Decoupled EMPCNs*. The input graph is split into two: the cellular lifted graph for higher-order message passing and the fully connected graph for direct node communication.

In a fully connected setting – which is the typical setting for such tasks – the number of message sent is $\mathcal{O}(|\mathcal{V}|^2)$. To obtain an equal computation complexity, we introduce a new type of cellular adjacency:

$$\mathcal{P}(\sigma) := \{\tau \mid \tau \subset \sigma \text{ and } |\tau| = 1\}, \quad (7)$$

i.e. all vertices that make up σ . Clearly, in the typical case of e.g. molecular prediction tasks, the number of rings $|\mathcal{R}|$ is much smaller than the number of nodes in the graph. If we therefore only perform message passing over the fully connected adjacency for the nodes and over $\mathcal{P}(\sigma)$ for the highest-order features we want to learn, the amount of message sent is:

$$\mathcal{O}(|\mathcal{V}|^2 + |\mathcal{R}| \cdot |\mathcal{V}|) = \mathcal{O}(|\mathcal{V}|^2) \text{ if } |\mathcal{R}| < |\mathcal{V}|. \quad (8)$$

For example, in molecular prediction tasks, this adjacency allows the model to directly propagate information between rings and their adjacent nodes. By incorporating this additional information, which is further enhanced by geometric invariants such as volumes, the model can better understand how adjacent nodes influence ring properties and vice versa.

Since functional groups attached to rings – e.g. hydroxyl, methyl, nitro groups for molecules – can significantly affect ring geometry, calculating these geometric properties could provide a better insights into the complex interactions and structural variations around the rings.

Note that our proposed method is not limited to cellular lifting methods and ring structures. Instead, it allows for the definition of any type of higher-order structure as a template, followed by the calculation of its geometric invariants. This method serves as a general framework that combines geometric GNNs with additional explicit topological information, offering scalability and flexibility in including various higher-order structures.

4. Related Works

Higher-Order Networks: Recent works have expanded Graph Neural Networks (GNNs) to include simplicial and cellular complexes (Papillon et al., 2024), where especially the work by Hajij et al. (2021) and Bodnar et al. (2021a) are closely related to this work. To our knowledge, Eijkelboom et al. (2023) is the first work explicitly integrating topological and geometric MPNNs and as such form the basis of our approach. A similar approach was later proposed to integrate mesh face geometric information into E(3) Equivariant GNNs by altering the update function (Anh Trang et al., 2024). Furthermore, more geometrically complex architectures utilizing Clifford algebras for simplicial complexes have been introduced (Liu et al., 2024). However, here, the improvement in performance comes at a significantly high computational cost.

Efficient Higher-Order Networks: Addressing the computational complexity inherent in higher-order message passing, recent studies have sought ways to enhance efficiency. Notably, SaNN (Gurugubelli & Chepuri, 2023) introduced a method to integrate simplicial inductive bias into networks efficiently, maintaining constant training time and memory requirements. This was achieved by enhancing neural models with pre-aggregated features of simplices across different orders. A different direction was taken by Topo-MLP (Ramamurthy et al., 2023), where they learn representations for simplices across multiple dimensions through the adoption of a higher-order neighborhood contrastive loss, thereby moving away from the traditional message passing framework.

5. Experiments & Results

For all experiments, the implementation and experimental details are provided in Appendix A and B respectively.

5.1. N-Body System

As introduced by Kipf et al. (2018), the N-body system experiment involves tracking the trajectories of five charged particles in three-dimensional space over time. The objective is to predict the positions of all particles after 1,000 time steps, starting from their initial positions and velocities. To ensure a consistent comparison, we adopted the experimental setup as described by Satorras et al. (2021). While the system operates according to physics, with particles being attracted or repelled based on their charges, the task remains E(n) equivariant, maintaining these symmetries throughout the entire trajectory. This means that the predictions are consistent under Euclidean transformations, such as rotations and translations, ensuring the model respects these inherent physical symmetries.

Table 1. Mean Squared Error for the N-body system experiment.

Model	MSE (↓)
TFN (Thomas et al., 2018)	0.0155
PONITA (Bekkers et al., 2024)	0.0043
SEGNN (Brandstetter et al., 2021)	0.0043
CGENN (Ruhe et al., 2023)	0.0039
EMPSN (Eijkelboom et al., 2023)	0.0063
EGNN (Satorras et al., 2021)	0.0071
EGNN + CIN	0.0057 ± 0.0002
EMPCN	0.0046 ± 0.00005

While our method can be applied to E(n) equivariant scenarios, for this experimental setup, we work with E(3) equivariance. As seen in Table 1, our performance is close to on par with state-of-the-art approaches.

To evaluate the significance of geometric invariants in higher-order message passing, we conducted an additional experiment with a modified version of Cellular Message Passing (CIN) (Bodnar et al., 2021a). In this setup, the node-to-node communication (0-cell to 0-cell) included the distance norm, transforming it into EGNN. However, the higher-order message passing layers did not include any geometric invariants, distinguishing it from our proposed EMPCNs. As shown in Table 1, while this approach outperforms EGNN alone, it falls short by 20% in performance compared to EMPCNs, highlighting the important role of geometric invariants in the higher order message passing layers. It is worth noting that the number of parameters has been kept similar to EGNN in all of our experiments.

5.2. QM9

The QM9 dataset, first introduced by Ramakrishnan et al. (2014) and subsequently studied by Gilmer et al. (2017) and Wu et al. (2018), comprises approximately 130,000 graphs, each consisting of around 18 nodes. In this dataset, the graphs represent molecules, with the nodes representing

Table 2. QM9 MAE (\downarrow) Results. Relative improvement with respect to EGNNs (gain) is provided for comparison.

Task Units	α (bohr 3)	$\Delta\epsilon$ (meV)	ϵ_{HOMO} (meV)	ϵ_{LUMO} (meV)	μ (D)	C_v (cal/mol K)	G (meV)	H (meV)	R^2 (bohr 3)	U (meV)	U_0 (meV)	ZPVE (meV)
NMP	.092	69	43	38	.030	.040	19	17	.180	20	20	1.50
SchNet	.235	63	41	34	.033	.033	14	14	.073	19	14	1.70
Cormorant	.085	61	34	38	.038	.026	20	21	.961	21	22	2.02
TFN	.223	58	40	38	.064	.101	-	-	-	-	-	-
SE(3)-Tr.	.142	53	35	33	.051	.054	-	-	-	-	-	-
DimeNet++	.043	32	24	19	.029	.023	7	6	.331	6	6	1.21
SphereNet	.046	32	23	18	.026	.021	8	6	.292	7	6	1.21
PaiNN	.045	45	27	20	.012	.024	7	6	.066	5	5	1.12
SEGNN	.060	42	24	21	.023	.031	15	16	.660	12	15	1.62
MPSN	.266	153	89	77	.101	.122	31	32	.887	33	33	3.02
EMPSN	.066	37	25	20	.023	.024	6	9	.101	7	10	1.37
EGNN	.071	48	29	25	.029	.031	12	12	.106	12	12	1.55
Decoupled EMPCN	.063	40	27	22	.026	.027	10	10	.104	10	11	1.51
Gain	11%	17%	7%	12%	10%	13%	17%	17%	2%	17%	9%	3%

atoms, and edges representing various types of bonds between these atoms. The objective of analyzing this dataset is to predict quantum chemical properties.

For this experiment, we followed the same setting as EGNN (Satorras et al., 2021). We applied our proposed method as introduced in Section 3.2, called *decoupled EMPCNs*, using EGNN for the fully connected node-to-node communication, while employing EMPCNs for the higher-order message passing. As discussed in Section 3.2, different adjacencies and geometric invariants can be used. For this task, we specifically chose only the ring-to-node communication and its respective invariants, including the distance from each node to the ring’s midpoint, the ring’s length, volume and area.

Note that, given the data consists of small molecules – where cells are not present much – it is crucial to optimize parameter usage to avoid unnecessary computations. We hence opt to allocate 75% of the parameters to node-to-node and 25% to ring-to-node message passing, such that the majority of parameters are allocated to the main architecture, which, in our case, is EGNN.

As indicated in Table 2, incorporating this additional ring information yielded an average improvement of approximately 10% over the baseline EGNN model. While one could potentially generate additional rings via the Radius-Graph or Vietoris-Rips complex, our experiment demonstrates that even with the initial graph’s topology, we can achieve significant performance gain *without* extra computational complexity. The reason for not applying the traditional EMPCNs in this experiment is because, in the fully connected setting, the number of rings becomes extremely large, and calculating the respective invariants is computationally infeasible.

5.3. CMU Motion Capture

In this experiment, we evaluated our *decoupled EMPCN* model using the CMU Motion Capture Database (Gross & Shi, 2001), which tracks the trajectories of human motion in various scenarios. The dataset’s graphs are composed of 31 connected nodes, each corresponding to a specific position on the human body during walking. Our goal is to forecast the node positions 30 timesteps into the future based on the positions from a randomly chosen frame. We adopted a similar framework to Kipf & Welling (2017); Huang et al. (2022); Liu et al. (2024), concentrating on the walking motion of a single subject (subject #35) over 23 trials.

Table 3. Comparison of different methods based on MSE (10^{-2}) for the CMU Motion Capture Dataset.

Method	MSE (\downarrow)
TFN (Thomas et al., 2018)	66.9
SE(3)-Tr (Fuchs et al., 2020)	60.9
GNN (Gilmer et al., 2017)	67.3
GMN (200K) (Huang et al., 2022)	17.7
EMPSN (200K) (Eijkelboom et al., 2023)	15.1
CGENN (200K) (Ruhe et al., 2023)	9.41
CSMPN (200K) (Liu et al., 2024)	7.55
EGNN (200K) (Satorras et al., 2021)	31.7
Decoupled-EMPCN (200K)	9.26 ± 0.11

To construct our dataset, we followed a similar approach as outlined in Section 5.2. For node-to-node communication, we employed a traditional EGNN architecture, where all nodes are interconnected. However, for cellular lifting, due to the absence of obvious rings, we adopted a methodology similar to that introduced by Liu et al. (2024). We manually connected nodes in our input graph that we believe will provide additional information. Specifically, in addition to

connecting elbow joint nodes with shoulder and palm nodes to create edges and triangles, and linking hip, knee, and heel nodes to form another set of edges and triangles as proposed, we introduced rings of length four. These rings connect the two elbows with the two knees, two arms with two wrists, two knees with the two feet, and two knees with the feet and the hip. While rings of varying lengths and structures could be utilized, the ultimate goal is to connect nodes that may be significant for the downstream task. Once these ring structures are constructed, cellular lifting is performed, which is then efficiently utilized by our *decoupled EMPCNs* as discussed in Section 3.2. Similar to the experiment described in Section 5.2, we used node-to-node and ring-to-node.

As shown in Table 3, our proposed model achieves scores that are nearly state-of-the-art without relying on highly computationally intensive techniques such as steerable methods or Clifford algebra methods. This demonstrates that our approach, which simply adds an extra message passing layer to learn topological features of the graph combined with the corresponding geometric invariants, can significantly improve the performance of GNNs on geometric datasets.

5.4. Ablation Study

In our ablation study, we compared the performance of EGNN and *decoupled EMPCN* under varying conditions of data availability and model complexity. We evaluated both models in a small-data regime to test generalizability and in a small-model regime to assess parameter efficiency. The results are summarized in Figure 4.

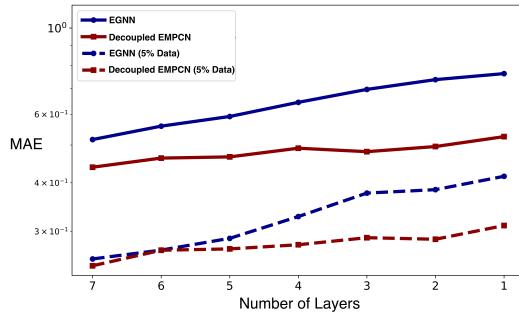


Figure 4. Ablation study on the α property of QM9, comparing decoupled EMPCN (red) and EGNN (blue). Both models were trained with 100% data (solid lines) and 5% data (dashed lines), and evaluated on MAE (\downarrow) across different numbers of layers.

First, we notice that as the number of layers decreases – i.e. a simpler model is considered – the EGNN baselines perform worse than their cellular counterparts, and this performance gap increases as the number of layers decreases. However, when a large enough model *and* enough data are used, this gap can be almost closed. That being said, when

given access to only 5% of the data, EGNN is consistently outperformed, even in large-model settings. This indicates that the cellular counterpart is not only more data-efficient but also has significantly stronger generalization capabilities.

These findings highlight the robustness and efficiency of decoupled EMPCNs, illustrating that they are a more reliable choice for scenarios where data and computational resources are limited as is typical in real-life settings. Robustness and efficiency here refer to maintaining higher performance than EGNN across different data regimes and model complexities.

6. Conclusion

In this study, we introduced E(n) Equivariant Message Passing Cellular Networks (EMPCNs), extending the capabilities of EMPSNs by transitioning from simplicial to cellular complexes. This transition allows for a more flexible representation of topological structures, enabling the inclusion of complex features such as rings in molecular graphs. Consequently, EMPCNs can define geometric invariants over a broader range of objects, enhancing the expressiveness of the model.

We also proposed a method to efficiently integrate topological information into geometric graphs that allow for a fully connected backbone in the standard message-passing phase. By decoupling the input graph into a fully connected graph for direct node-to-node communication and a cellular lifted graph for higher-order message passing, we combined the benefits of both approaches. This technique maintains the original graph’s topology while utilizing any MPNN architecture for direct communication, thereby improving both performance and expressivity.

In summary, decoupled EMPCNs provide a scalable and expressive framework for higher-order message passing in geometric and topological graphs. Although this method is proposed for cellular complexes, the same techniques could be adapted to other higher-order path message passing methods, potentially enhancing the expressiveness by incorporating various structures beneficial for specific tasks. We believe that exploring the intersection between geometric and topological MPNNs is a crucial aspect of developing efficient and strongly generalizing models, especially in low-data settings.

Limitations and Future Work The main limitation of EMPCNs is the computational cost associated with the calculation of the introduced geometric invariants. Although it is possible to retain only those invariants relevant to a specific downstream task, an interesting direction for future research would be to pre-aggregate cell representations and

compute the geometric invariants solely for key structures essential to the task. However, it is worth mentioning that these invariants are significantly cheaper to compute than the ones in methods using e.g. Clifford algebras.

For the decoupled EMPCNs, the primary limitation is that in geometric graphs, initial connectivity may not be available, and only a point cloud may be provided. While pre-domain knowledge, as discussed in Section 5.3, can be used to identify important higher-order structures, a promising research direction would be to develop methods to autonomously identify these structures based on the downstream task. In addition, as we just showed with the decoupled EMPCNs a generic template that can combine geometric and topological approaches, our goal was to beat the performance of the main architecture being used, which was EGNN in our case. Further experiments with more advanced architectures like SEGNN (Brandstetter et al., 2021) should be tested. Finally, although our method is proposed for cellular complexes, the same technique could be applied to higher-order path message passing, potentially resulting in a more expressive framework by incorporating different types of structures that are known to be useful for a particular downstream task.

References

Anh Trang, T., Ngo, N. K., Levy, D. T., Ngoc Vo, T., Ravanbakhsh, S., and Son Hy, T. E(3)-equivariant mesh neural networks. In Dasgupta, S., Mandt, S., and Li, Y. (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 748–756. PMLR, 02–04 May 2024.

Barber, C. B., Dobkin, D. P., and Huhdanpaa, H. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22(4):469–483, dec 1996. ISSN 0098-3500. doi: 10.1145/235815.235821. URL <https://doi.org/10.1145/235815.235821>.

Bekkers, E. J., Vadgama, S., Hesselink, R., der Linden, P. A. V., and Romero, D. W. Fast, expressive SE(n) equivariant networks through weight-sharing in position-orientation space. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=dPHLbUqGbr>.

Bodnar, C., Frasca, F., Otter, N., Wang, Y., Liò, P., Montufar, G. F., and Bronstein, M. Weisfeiler and lehman go cellular: Cw networks. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 2625–2640. Curran Associates, Inc., 2021a. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/157792e4abb490f99dbd738483e0d2d4-Paper.pdf.

Bodnar, C., Frasca, F., Wang, Y., Otter, N., Montufar, G. F., Liò, P., and Bronstein, M. Weisfeiler and lehman go topological: Message passing simplicial networks. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1026–1037. PMLR, 18–24 Jul 2021b.

Brandstetter, J., Hesselink, R., van der Pol, E., Bekkers, E., and Welling, M. Geometric and physical quantities improve e (3) equivariant message passing. *arXiv preprint arXiv:2110.02905*, 2021.

Bronstein, M. M., Bruna, J., Cohen, T., and Velickovic, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *CoRR*, abs/2104.13478, 2021. URL <https://arxiv.org/abs/2104.13478>.

Chen, Z., Villar, S., Chen, L., and Bruna, J. On the equivalence between graph isomorphism testing and function approximation with gnns. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information*

Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/71ee911dd06428a96c143a0b135041a4-Paper.pdf.

Davis, M. and Okun, B. l_2 -homology of right-angled Coxeter groups based on barycentric subdivisions. *Topology and its Applications*, 140:197–202, 05 2004. doi: 10.1016/j.topol.2003.07.008.

Duval, A., Mathis, S. V., Joshi, C. K., Schmidt, V., Miret, S., Malliaros, F. D., Cohen, T., Liò, P., Bengio, Y., and Bronstein, M. A hitchhiker’s guide to geometric GNNs for 3d atomic systems, 2024.

Dwivedi, V. P. and Bresson, X. A generalization of transformer networks to graphs. *AAAI Workshop on Deep Learning on Graphs: Methods and Applications*, 2021.

Dwivedi, V. P., Joshi, C. K., Luu, A. T., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023. URL <http://jmlr.org/papers/v24/22-0567.html>.

Eijkelboom, F., Hesselink, R., and Bekkers, E. J. $E(n)$ equivariant message passing simplicial networks. In *International Conference on Machine Learning*, pp. 9071–9081. PMLR, 2023.

Fuchs, F., Worrall, D., Fischer, V., and Welling, M. $Se(3)$ -transformers: 3d roto-translation equivariant attention networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1970–1981. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/15231a7ce4ba789d13b722cc5c955834-Paper.pdf.

Gasteiger, J., Becker, F., and Günemann, S. Gemnet: Universal directional graph neural networks for molecules. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 6790–6802. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/35cf8659cfcb13224cbd47863a34fc58-Paper.pdf.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. PMLR, 2017.

Girvan, M. and Newman, M. E. J. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002. doi: 10.1073/pnas.122653799. URL <https://www.pnas.org/doi/abs/10.1073/pnas.122653799>.

Gligorijević, V., Renfrew, P. D., Kosciolak, T., Koehler Lehman, J., Berenberg, D., Vatanen, T., Chandler, C., Taylor, B. C., Fisk, I. M., Vlamakis, H., Xavier, R. J., Knight, R., Cho, K., and Bonneau, R. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12(1):3168, 2021. doi: 10.1038/s41467-021-23303-9. URL <https://doi.org/10.1038/s41467-021-23303-9>.

Gross, R. and Shi, J. The cmu motion of body (mobo) database. Technical Report CMU-RI-TR-01-18, Carnegie Mellon University, Pittsburgh, PA, June 2001.

Gurugubelli, S. and Chepuri, S. P. Sann: Simple yet powerful simplicial-aware neural networks. In *The Twelfth International Conference on Learning Representations*, 2023.

Hajij, M., Istvan, K., and Zamzmi, G. Cell complex neural networks, 2021.

Huang, W., Han, J., Rong, Y., Xu, T., Sun, F., and Huang, J. Equivariant graph mechanics networks with constraints. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=SHbhHHfePhP>.

Jiang, C., Coenen, F., and Zito, M. Finding frequent subgraphs in longitudinal social network data using a weighted graph mining approach. In Cao, L., Feng, Y., and Zhong, J. (eds.), *Advanced Data Mining and Applications*, pp. 405–416, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-17316-5.

Joshi, C. K., Bodnar, C., Mathis, S. V., Cohen, T., and Liò, P. On the expressive power of geometric graph neural networks. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.

Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2688–2697. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/kipf18a.html>.

Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks, 2017.

Koyutürk, M., Grama, A., and Szpankowski, W. An efficient algorithm for detecting frequent subgraphs in biological networks. *Bioinformatics*, 20:i200–i207, 08 2004. ISSN 1367-4803. doi: 10.1093/bioinformatics/bth919. URL <https://doi.org/10.1093/bioinformatics/bth919>.

Kreuzer, D., Beaini, D., Hamilton, W., Létourneau, V., and Tossou, P. Rethinking graph transformers with spectral attention. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 21618–21629. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/b4fd1d2cb085390fbbadae65e07876a7-Paper.pdf.

Liu, C., Ruhe, D., Eijkelboom, F., and Forré, P. Clifford group equivariant simplicial message passing networks. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Zz594UBNOH>.

Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y. Provably powerful graph networks. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bb04af0f7ecaee4aae62035497da1387-Paper.pdf.

Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33(01), pp. 4602–4609, 2019.

Papillon, M., Sanborn, S., Hajij, M., and Miolane, N. Architectures of topological deep learning: A survey of message-passing topological neural networks, 2024.

Ramakrishnan, R., Dral, P., Rupp, M., and von Lilienfeld, A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 08 2014. doi: 10.1038/sdata.2014.22.

Ramamurthy, K. N., Guzmán-Sáenz, A., and Hajij, M. Topo-mlp: A simplicial network without message passing. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2023.

Rampášek, L., Galkin, M., Dwivedi, V. P., Luu, A. T., Wolf, G., and Beaini, D. Recipe for a general, powerful, scalable graph transformer. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 14501–14515. Curran Associates, Inc., 2022.

Ruhe, D., Brandstetter, J., and Forré, P. Clifford group equivariant neural networks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=n84bzMrGUD>.

Satorras, V. G., Hoogeboom, E., and Welling, M. E(n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

Schütt, K., Unke, O., and Gastegger, M. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International Conference on Machine Learning*, pp. 9377–9388. PMLR, 2021.

Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7537–7547, 2020.

Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds, 2018.

Veličković, P. and Blundell, C. Neural algorithmic reasoning. *Patterns*, 2(7), 2021.

Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. S. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/488e4104520c6aab692863cc1dba45af-Paper.pdf.

Weisfeiler, B. and Leman, A. The reduction of a graph to canonical form and the algebra which appears therein. *nti, Series*, 2(9):12–16, 1968.

Wu, Z., Ramsundar, B., Feinberg, E., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. Moleculenet: a benchmark for molecular machine learning. *Chem. Sci.*, 9:513–530, 2018. doi: 10.1039/C7SC02664A. URL <http://dx.doi.org/10.1039/C7SC02664A>.

Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.

Zhang, M. and Chen, Y. Link prediction based on graph neural networks. *Advances in neural information processing systems*, 31, 2018.

Zitnik, M., Agrawal, M., and Leskovec, J. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13):i457–i466, 06 2018. ISSN 1367-4803. doi: 10.1093/bioinformatics/bty294. URL <https://doi.org/10.1093/bioinformatics/bty294>.

A. Implementation Details

This section describes the implementation of EMPCNs. After a cellular lifting is performed, we have : 1) a set of features for each cell 2) adjacency relationships between these cells 3) positional information of the nodes $\{x_i\}$, and 4) optionally, initial velocities of the nodes $\{v_i\}$. The learnable functions are as follows:

- Feature Initialization: Initial features are embedded using a linear embedding process. For each dimension n , we apply a separate embedder:

$$\{\text{Initial Feature}_n \rightarrow \{\text{LinearLayer}_n\} \rightarrow \text{Embedded Feature}_n\}$$

- Message Passing: For an adjacency A from τ to σ , we learn a message function. Using the input $[h_\sigma \oplus h_\tau \oplus \text{Inv}(\sigma, \tau)]$, where \oplus denotes concatenation, the message function is defined as follows:

$$\{[h_\sigma \oplus h_\tau \oplus \text{Inv}(\sigma, \tau)] \rightarrow \text{LinearLayer} \rightarrow \text{Swish} \rightarrow \text{LinearLayer} \rightarrow \text{Swish} \rightarrow m_{\sigma, \tau}^A\}$$

- Edge Inference: For each message, we compute the edge importance similar to [Satorras et al. \(2021\)](#):

$$\{m_{\sigma, \tau}^A \rightarrow \text{LinearLayer} \rightarrow \text{Sigmoid}\} \rightarrow e_{\sigma, \tau}^A$$

- Cell Update: For each cell, an update is learned through the different adjacency communications:

$$[h_\sigma, \{m_{\sigma, \tau}^A\}, \{e_{\sigma, \tau}^A\}] \rightarrow \{h_\sigma \oplus \left(\bigoplus_{A_{\sigma, \tau}} e_{\sigma, \tau}^A \cdot m_{\sigma, \tau}^A \right) \rightarrow \text{LinearLayer} \rightarrow \text{Swish} \rightarrow \text{LinearLayer}\} \rightarrow h'_\sigma$$

- Readout: To get the final prediction, a pre-readout and readout phase take place:

$$\begin{aligned} \{h'_\sigma\} &\rightarrow \text{LinearLayer} \rightarrow \text{Swish} \rightarrow \text{LinearLayer} \rightarrow \bigoplus_n \left(\sum_i h_i'^n \right) \rightarrow \text{LinearLayer} \rightarrow \text{Swish} \\ &\rightarrow \text{LinearLayer} \rightarrow \text{Prediction} \end{aligned}$$

This architecture design is related to the $E(n)$ Equivariant Cellular Message Passing Cellular Networks introduced in Section 3.1. In the *decoupled EGNN* discussed in Section 3.2, only 0-cells (nodes) and 2-cells (rings) are used in the experiment described in 5.2. The message passing and update rules are adapted to handle the interactions directly between these two types of cells.

In the experiments of N-Body and CMU Motion from Sections 5.1 and 5.3 respectively, where the final prediction is the displaced position in the next timestep, and the velocity of each node is used within the prediction, the equations for the position update are similar to the ones introduced by [Satorras et al. \(2021\)](#):

- Velocity Update:

$$v_i = \phi_v(h_i)v_i^{\text{init}} + C \sum_{j \neq i} (x_i - x_j) \phi_x(m_{ij})$$

- Position Update:

$$x'_i = x_i + v_i$$

B. Experimental Details

B.1. N-Body System

We used the same setup as in [Satorras et al. \(2021\)](#), i.e. 3,000 training trajectories, 2,000 validation trajectories, and 2,000 test trajectories. Each trajectory contains 1,000 time steps. We used a 4-layer EMPCN of dimension 2 for our experiments, setting the initial number of hidden features to 66, resulting in a total of 200k parameters in total. Additionally, a dropout layer with a probability of 0.2 was used for every single update and message passing. The invariant features are embedded

using Gaussian Fourier features as introduced in [Tancik et al. \(2020\)](#). The optimization is done using the Adam optimizer, with a constant learning rate of $\eta = 5 \times 10^{-4}$, a batch size of 100, and weight decay of 10^{-12} . The loss minimized is the Mean Squared Error (MSE) in the predicted position.

Both the upper and lower adjacencies, as well as the boundaries, were considered for the message passing. The different cell communications that are considered are the following: 0-cell (node) to 0-cell (node), 0-cell (node) to 1-cell (edge), 1-cell (edge) to 0-cell (node), 1-cell (edge) to 1-cell (edge), 1-cell (edge) to 2-cell (ring), and 2-cell (ring) to 1-cell (edge).

The invariants used for message passing are the following: 0-cell to 0-cell: node distance; 0-cell to 1-cell (and vice versa): length of edge; 1-cell to 1-cell: length of sending edge, length of receiving edge, distance from the midpoints of sending and receiving edge; 1-cell to 2-cell (and vice versa): ring approximated radius, ring length, edge length.

B.2. QM9

For the QM9 dataset, we utilized the common split of 100,000 molecules for training, 10,000 molecules for testing, and the remaining molecules for validation. All predicted properties were normalized by subtracting the mean of the target values in the training set and then dividing by the mean absolute deviation (MAD) in the training set to stabilize training.

The models were trained for 1000 epochs each, with the final number of parameters set to one million. For the *decoupled EMPCN* model, the parameters were divided 75:25 between the EGNN and the ring-to-node EMPCN components, respectively. To match the one million parameter target, the number of hidden features in the higher-order message passing was reduced to 44, while the number of layers remained equal to 7. For the node-to-node communication, which is handled by the EGNN, only the norm of the distance was used as an invariant. For the ring-to-node communication, the invariants used were the convex hull volume and area, the norm of the distance between the ring midpoint and each node within the ring, and the ring length. Although one could also use the representations of rings in the final readout, we chose to use only the representations of nodes to ensure a fair comparison with the EGNN model.

The models were optimized using the Adam optimizer with an initial learning rate of 5×10^{-2} and a Cosine Annealing learning rate scheduler. The Mean Absolute Error (MAE) was used as the loss function for optimization. We used a batch size of 96 molecules and applied a weight decay of 10^{-12} . Additionally, dropout layers were included in each message and update layer.

B.3. CMU Motion Capture

As discussed in Section 5.3, for this task the goal is to forecast the node positions 30 timesteps into the future based on the positions from a randomly chosen frame. We adopted a similar framework to [Kipf & Welling \(2017\)](#); [Huang et al. \(2022\)](#); [Liu et al. \(2024\)](#), concentrating on the walking motion of a single subject (subject #35) over 23 trials.

As there are no rings present we use the extra topological information from connecting elbow joint nodes with shoulder and palm nodes to create edges and triangles. Additionally, we link hip, knee, and heel nodes to form another set of edges and triangles as firstly introduced by [Liu et al. \(2024\)](#). As mentioned in Section 5.3, we also introduce rings of length four and five. These rings connect the two elbows with the two knees, the two arms with the two wrists, the two knees with the two feet, and the two knees with the feet and the hip. The specific connections are as follows: [0, 3, 8], [6, 7, 8], [1, 2, 3], [24, 25, 26], [21, 22, 23], [7, 8, 2, 3], [24, 26, 17, 19], [25, 7, 18, 2], and [6, 7, 8, 1, 2, 3].

The models were trained for 5000 epochs each, with the final number of parameters set to 200K. For the *decoupled EMPCN* model, the parameters were divided 80:20 between the EGNN and the ring-to-node EMPCN components, respectively. To achieve the 200K parameter target, the number of hidden features in the higher-order message passing was reduced to 33 in comparison to 66 that was for the node to node communication, while the number of layers was set to 4. Similar to the experimental details for the QM9 dataset mentioned in Section 5.2 and Appendix B.2, for the final readout we used only the representations of nodes.

The models were optimized using the Adam optimizer with an initial learning rate of 5×10^{-2} and a Cosine Annealing learning rate scheduler. The loss minimized was the Mean Squared Error (MSE) in the predicted position. A batch size of 100 was used, along with a weight decay of 10^{-8} . Additionally, dropout layers with a probability of 15% were included in each message and update layer. Finally, the invariants used for the ring-to-node communication were the norm of the distance for node-to-node communication, the ring length, and the norm of the distance from each point to the midpoint of the respective ring.