
P-tensors: a General Framework for Higher Order Message Passing in Subgraph Neural Networks

Andrew Hands

Department of Computer Science
University of Chicago
hands@uchicago.edu

Tianyi Sun

Computational and Applied Math
Department of Statistics
University of Chicago
tianyisun@uchicago.edu

Risi Kondor

Department of Computer Science
Department of Statistics
University of Chicago
risi@uchicago.edu

Abstract

Several recent papers have proposed increasing the expressive power of graph neural networks by exploiting subgraphs or other topological structures. In parallel, researchers have investigated higher order permutation equivariant networks. In this paper we tie these two threads together by providing a general framework for higher order permutation equivariant message passing in subgraph neural networks. In this paper we introduce a new type of mathematical object called P -tensors, which provide a simple way to define the most general form of permutation equivariant message passing in both the above two categories of networks. We show that the P -tensors paradigm can achieve state-of-the-art performance on benchmark molecular datasets.

1 INTRODUCTION

Graph Neural Networks (GNNs) have proved to be remarkably successful in a wide range of domains from filling in edges in social networks to predicting the chemical properties of molecules. Amongst graph neural networks, so-called Message Passing Neural Nets (MPNNs), which loosely imitate convolution on the graph, are particularly popular and have found use in almost every branch of science (Gilmer et al., 2017). However, it has also been shown that there are severe theoretical limitations on the expressive power of MPNNs, and that in some cases they fail to capture

even relatively simple topological features such as cycles (Morris et al., 2019; Xu et al., 2019).

To develop more expressive GNNs, researchers have developed models that consider higher order representations of graphs, such as Invariant Graph Neural Networks (IGNs) (Maron et al., 2019b,a,c), and models that explicitly account for specific types of subgraphs (Frasca et al., 2022; Bodnar et al., 2021b). In the first category of models, at the cost of increased time and space complexity, in principle, one can use arbitrarily high order permutation equivariant representations (Maron et al., 2019a,c). In the second category one can encode a rich set of substructures allowing the network to learn different sets of weights for each type of structure (Bodnar et al., 2021a). However, the question of how to bring these two lines of research together and design a general architecture where neurons corresponding to varied substructures can pass messages to each other in the most general possible way has so far been elusive.

Main contributions. In this paper we present a general framework for extending the message passing paradigm to a setting where (a) the titular “neurons” can correspond to not just individual vertices, but specific types of subgraphs or other topological structures related to the underlying graph, as long as these structures are selected by a permutation invariant selection policy; (b) the neurons communicate with each other via *higher order* message passing in the sense that with respect to local permutations the messages behave as permutation covariant vectors, matrices or tensors.

Our first technical contribution is to give a precise definition of equivariance in this setting, which, to the best of our knowledge has not been previously done (Definitions 3–4). Next, we introduce P -tensors, a new mathematical device that makes it easy to describe and implement higher order message passing between subgraphs. Here we have two technical results. First we derive the form of all possible equivariant linear maps

between P -tensors, which turns out to be a generalization of the results of (Maron et al., 2019b), but affording a larger set of possible interactions (Theorem 3). Second, we show that any subgraph neural network utilizing these maps is equivariant in the sense of our earlier definition (Theorem 1). In our experiments we find that the flexibility and generality of the P -tensors framework does indeed pay off, in particular, on the ZINC 12K benchmark we reach state-of-the-art results.

2 BACKGROUND: MESSAGE PASSING NEURAL NETWORKS

Let \mathcal{G} be an undirected graph with n vertices and $A \in \mathbb{R}^{n \times n}$ be its adjacency matrix. Graph neural networks (GNNs) learn a function $\Phi: A_{\mathcal{G}} \mapsto \Phi(A_{\mathcal{G}})$ embedding such graphs in some Euclidean space \mathbb{R}^D . The fundamental constraint on Φ is that it must be invariant to relabeling the vertices. If we change the order in which the vertices are numbered by a permutation σ , the adjacency matrix transforms as

$$A \mapsto A', \quad \text{with} \quad [A']_{i,j} = A_{\sigma^{-1}(i), \sigma^{-1}(j)}. \quad (1)$$

However, A' still represents the same graph \mathcal{G} , so overall the network must satisfy $\Phi(A) = \Phi(A')$.

Requiring permutation invariance at the level of individual neurons would be too restrictive. Instead, modern GNNs are designed in such a way that their internal layers are *equivariant* (rather than invariant) to permutations, meaning that under (1) the outputs of the layers do change, but do so in a specific, controlled way. The last layer of the network is designed to cancel out these transformations, typically by pooling over the vertices, and thus guarantees that the final output is invariant.

Presently the most popular approach to building equivariant GNNs is the *Message Passing Neural Network* (MPNN) paradigm (Gilmer et al., 2017), where the titular “neurons” are attached to the vertices of the graph and communicate with each other by sending messages to their neighbors. In the simplest case, the output of the neuron at vertex i in layer ℓ is a vector $f_i^\ell \in \mathbb{R}^{D_\ell}$, and the update rule is

$$f_i^\ell = \eta \left(W_\ell \sum_{j \in \mathcal{N}(i)} f_j^{\ell-1} + b_i^\ell \right). \quad (2)$$

Here $\mathcal{N}(i)$ denotes the neighbors of node i , W_ℓ is a learnable weight matrix, b_i^ℓ is a learnable bias term and η is a pointwise non-linearity.

It is easy to see that under the transformation (1), the output of each neuron in such an MPNN becomes

$f'_i{}^\ell = f_{\sigma^{-1}(i)}^\ell$, and in this sense the network as a whole is permutation equivariant. The message passing process also bears some similarity to classical convolution in e.g., image processing, and captures the intuitive idea that information in complex networks should propagate via local connections.

Despite these attractive properties, classical MPNN do have significant limitations, the most obvious of which is that each vertex i just *sums* all the incoming messages from its neighbors. The associative and commutative nature of summation is critical for ensuring equivariance, but it also makes MPNNs myopic in the sense that once the activations have been summed, downstream neurons have no way to distinguish between which part of the incoming message came from which neighbor (Hy et al., 2018; Morris et al., 2019; Xu et al., 2019; Chen et al., 2020; You et al., 2021). From a theoretical point of view, the consequence is that classical MPNNs are only as powerful as the first order Weisfeiler-Leman test (Xu et al., 2019; Weisfeiler and Leman, 1968). In response to these criticisms, the community has been exploring various ways to generalize the message passing idiom.

2.1 Higher order equivariance

One way to make MPNNs more expressive is to design architectures that are equivariant to higher order actions of the group of permutations, technically called the symmetric group, \mathbb{S}_n . Studying this problem at the most general level involves considering equivariance to each irreducible representation of \mathbb{S}_n , a rich but mathematically involved subject (Sagan, 2001; Sannai et al., 2019; Keriven and Peyré, 2019; Thiede et al., 2020). For practical GNNs however it is usually sufficient to consider k 'th order equivariance in the sense of how \mathbb{S}_n acts on k 'th order tensors:

$$T \xrightarrow{\sigma} T' \quad [T']_{i_1, \dots, i_k} = [T]_{\sigma^{-1}(i_1), \dots, \sigma^{-1}(i_k)}. \quad (3)$$

Maron et al. (2019b) derived the general form of linear neural network layers that are invariant to this action and showed that for higher values of k it captures much richer interactions than simple first-order message passing, which is essentially what classical MPNNs do. The fundamental limitation of this approach however is that the size of the tensor T grows exponentially with k . Therefore, even for moderate sized graphs, considering more than second or third order permutation equivariance becomes infeasible.

2.2 Subgraph neural networks

The other natural way to increase the expressiveness of graph neural networks is to extend the MPNN paradigm to passing messages to/from edges and other

subgraphs (Alsentzer et al., 2020; Thiede et al., 2021; Bevilacqua et al., 2022; Frasca et al., 2022). This class of approaches is also attractive because in many applications subgraphs have explicit semantic meaning. In organic chemistry, for example, they can correspond to functional groups.

At the extreme, subgraph neural networks allow combining two (or more) separate GNN algorithms in a recursive fashion, with one GNN running at the level of the subgraphs, and the other one combining their results at the global level. Subgraph neural networks are also related to hypergraph neural nets and simplicial complex networks (Feng et al., 2019; Dong et al., 2020; Ebli et al., 2020; Bodnar et al., 2021b; Zhao et al., 2022; Frasca et al., 2022).

There is a natural connection between subgraph networks and higher order permutation equivariance, which is that to capture the underlying combinatorial structures, subgraph neurons cannot just communicate via scalar (invariant) messages. Rather, the messages must be indexed by the vertices of the sending and receiving subgraph, i.e., they must be permutation covariant objects. This complicates deriving the rules of equivariant message passing, and most existing architectures employ somewhat ad hoc solutions. The goal of the present paper is to combine the higher order message passing and subgraph neural network frameworks and derive the general laws of permutation equivariant message passing between subgraphs in higher order subgraph networks (HOSNNs).

3 EQUIVARIANCE IN HIGHER ORDER SUBGRAPH NEURAL NETWORKS

The symmetric group acts on subgraph neural networks on two distinct levels: the vertices inside individual subgraphs get permuted, and the subgraphs themselves are permuted with each other. Our first challenge is to define what equivariance even means in this setting. We begin by formalizing how the subgraphs are selected.

Definition 1. (Selection policy) Let \mathcal{G} be a graph with vertex set $V = \{1, \dots, n\}$, adjacency matrix $A \in \mathbb{R}^{n \times n}$ and optionally an input feature matrix $L \in \mathbb{R}^{n \times d}$. A subgraph selection policy is a function $\psi: A \mapsto \mathbb{D}$ or $\psi: (A, L) \mapsto \mathbb{D}$ where \mathbb{D} is a set of subsets of V .

The purpose of \mathbb{D} is to demark the subgraphs to which we assign higher order neurons. For example, we might define ψ to return the set of all edges in \mathcal{G} , all paths of a given length, or all cycles. Each of these selection policies is *invariant* in the following sense.

Definition 2. (Invariant selection policy) Let \mathcal{G} ,




Figure 1: A subgraph neural network must be equivariant to two different ways that permutations act on it: changing the set of vertices assigned to a given subgraph, and reordering the vertices of the subgraph internally. This is especially important when the subgraph neurons produce matrix/tensor valued outputs indexed by the vertices of the subgraph itself. The P -tensors formalism allows us to handle this situation in a simple way, defining the most general form of equivariant linear messages between such tensors, without making reference to the global ordering.

A and L be as above. Assume that under permuting the vertices $A \mapsto A^{(\sigma)}$ and $L \mapsto L^{(\sigma)}$ with

$$A_{i,j}^{(\sigma)} = A_{\sigma^{-1}(i),\sigma^{-1}(j)} \quad \text{and} \quad L_{i,c}^{(\sigma)} = L_{\sigma^{-1}(i),c}.$$

Then a selection policy $\psi: (A, L) \mapsto \mathbb{D}$ is said to be invariant if for any permutation $\sigma \in \mathbb{S}_n$ and any selected set $\{i_1, \dots, i_p\} \in \psi(A, L)$, we have $\{\sigma(i_1), \dots, \sigma(i_p)\} \in \psi(A^{(\sigma)}, L^{(\sigma)})$.

An invariant selection policy is a policy that depends only on the graph topology and the input vertex features, but not the (arbitrary) ordering of the vertices. All the subgraph selection policies considered in this paper are invariant policies.

3.1 Two-level equivariance

The output of each layer of a higher order subgraph neural network is a collection of vectors, matrices or tensors $\mathcal{F} = \{T_1, \dots, T_m\}$ corresponding to the subgraphs $\mathcal{S}_{\{i_1^1, \dots, i_{p_1}^1\}}, \dots, \mathcal{S}_{\{i_1^m, \dots, i_{p_m}^m\}}$ induced by the vertex sets picked out by the selection policy $\psi(A, L)$. The difficulty with this setup is two-fold:

1. Since ψ picks out subgraphs in an arbitrary order, the order in which the T_1, \dots, T_m tensors are listed in \mathcal{F} is also arbitrary. In particular, we cannot guarantee that the order will not change if we permute the vertices of the underlying graph by σ .
2. Each tensor T_a is sensitive to the order in which the $\{i_1^a, \dots, i_{p_a}^a\}$ vertices are listed, and this can also change with σ . In particular, if T is a vector type quantity (w.r.t. permutations of the subgraph), then under a local permutation $(i_1^a, \dots, i_{p_a}^a) \xrightarrow{\tau} (i_{\tau(1)}^a, \dots, i_{\tau(p)}^a)$ it will change as $T_j^{(\tau)} = T_{\tau^{-1}(j)}$. If it is a matrix type quantity

then it will change as $T_{j,j'}^{(\tau)} = T_{\tau^{-1}(j),\tau^{-1}(j')}$. More generally, if T is a k 'th order tensor, then $T_{i_1,i_2,\dots,i_k}^{(\tau)} = T_{\tau^{-1}(i_1),\dots,\tau^{-1}(i_k)}$.

Remarkably, despite all these degrees of freedom, it is still possible to define covariance and equivariance in higher order subgraph neural networks in a meaningful way.

Definition 3. (Permutation covariant layer) Let $\mathcal{F} = \{T_1, \dots, T_m\}$ be the output of a layer in a higher order subgraph neural network induced by an invariant subgraph selection policy ψ and $\mathcal{F}' = \{T'_1, \dots, T'_m\}$ be the output of the same layer after permuting the vertices of the underlying graph \mathcal{G} by a permutation σ . By the invariance of ψ , for each T_a in \mathcal{F} there is a corresponding T'_a in \mathcal{F}' such that $\{\sigma(i_1^a), \dots, \sigma(i_{p_a}^a)\}$ and $\{i'_1{}^a, \dots, i'_{p_a}{}^a\}$ are equal as sets (intuitively, T_a and T'_a correspond to the same underlying subgraph despite the relabeling by σ). In particular, there is a permutation $\tau \in \mathbb{S}_{p_a}$ that elementwise aligns these two sets of vertex indices in the sense that $i'_j{}^a = \sigma(i_{\tau^{-1}(j)}^a)$. \mathcal{F} is said to be a permutation covariant layer if

$$[T'_a]_{j_1,j_2,\dots,j_k} = [T_a]_{\tau^{-1}(j_1),\dots,\tau^{-1}(j_k)} \quad (4)$$

for any permutation $\sigma \in \mathbb{S}_n$.

As usual, a permutation equivariant map is defined as one which preserves this covariance property.

Definition 4. (Permutation equivariant higher order subgraph neural network) A subgraph layer $\phi: \mathcal{F}^{in} \mapsto \mathcal{F}^{out}$, or more generally $\phi: (\mathcal{F}_1^{in}, \dots, \mathcal{F}_r^{in}) \mapsto \mathcal{F}^{out}$ is said to be **permutation equivariant** if whenever $\mathcal{F}_1^{in}, \dots, \mathcal{F}_r^{in}$ are covariant, \mathcal{F}^{out} is covariant as well. The entire network is equivariant if all the maps connecting its different layers are equivariant.

While these definitions appear quite technical, they are necessary for capturing the interaction between the effect of permutations at the subgraph selection level and the level of individual subgraphs. In the special case of ψ just picking out the individual vertices of \mathcal{G} (as trivial subgraphs consisting of a single vertex) and the activations being scalars, the definitions reduce to the classical case where \mathcal{F} can be represented as an $n \times c$ matrix, where n is the number of vertices and c the number of channels. This base case is what one typically uses in the input layer and readout layer of an equivariant GNN. The purpose of the next section is to introduce a mathematical formalism that makes implementing the intermediate, higher order layers as straightforward as possible.

4 P-TENSORS

The mathematical device that we introduce to derive the rules of higher order message passing are a type of object that we call *P-tensors*. To define *P-tensors* first we need to define a finite or countably infinite set \mathcal{U} of base objects called *atoms* that permutations act on. In the case of graph neural networks the atoms are just the vertices. However, the *P-tensor* formalism is also applicable to other permutation equivariant learning scenarios, such as relational learning, in which case \mathcal{U} might for example be a set of individuals or the words in a given language.

A given *P-tensor* T is defined relative to an ordered subset $\mathcal{D} = (x_1, \dots, x_d)$ of atoms called its **reference domain**, which, in the case of subgraph neurons, is just the vertex set of the given subgraph. Reordering the reference domain by a permutation $\tau \in \mathbb{S}_d$ changes it to

$$\mathcal{D}' = \tau \circ \mathcal{D} = (x_{\tau^{-1}(1)}, \dots, x_{\tau^{-1}(d)}). \quad (5)$$

The defining property of *P-tensors* is how they transform under this action.

Definition 5 (P-tensors). Let \mathcal{U} be a finite or countably infinite set of atoms and $\mathcal{D} = (x_1, \dots, x_d)$ an ordered subset of \mathcal{U} . We say that a k 'th order tensor $T \in \mathbb{R}^{d \times d \times \dots \times d}$ is a k 'th order permutation covariant tensor (or *P-tensor* for short) with reference domain \mathcal{D} if under reordering \mathcal{D} as in (5) it transforms to

$$[\tau \circ T]_{i_1,i_2,\dots,i_k} = T_{\tau^{-1}(i_1),\dots,\tau^{-1}(i_k)}. \quad (6)$$

The neurons in modern neural networks typically have many channels, so we also allow *P-tensors* to have a channel dimension. Thus, a k 'th order *P-tensor* can actually be a $k+1$ 'th order tensor $T \in \mathbb{R}^{d \times d \times \dots \times d \times C}$. The channel dimension is not affected by permutations. We will sometimes also write (T, \mathcal{D}) to denote a *P-tensor* with reference domain \mathcal{D} .

To derive the rules of equivariant message passing from one *P-tensor* T_1 to another *P-tensor* T_2 , we need to consider the three cases when their respective reference domains \mathcal{D}_1 and \mathcal{D}_2 in the unordered sense are (a) the same (b) partially overlap (c) are disjoint. The advantage of the *P-tensors* formalism compared to the previous section is that in each of these cases we need only consider the action of permutations that are *internal* to \mathcal{D}_1 and \mathcal{D}_2 . In particular, we have the following definition.

Definition 6 (Permutation equivariant maps between P-tensors). Let \mathcal{D}_1 and \mathcal{D}_2 be two fixed reference domains, $\overline{\mathcal{D}}_1$ be any reordering of \mathcal{D}_1 and $\overline{\mathcal{D}}_2$ any reordering of \mathcal{D}_2 . Consider a family of linear maps

$$\phi_{\overline{\mathcal{D}}_1, \overline{\mathcal{D}}_2}: (T_1, \overline{\mathcal{D}}_1) \mapsto (T_2, \overline{\mathcal{D}}_2).$$

We say that this is a permutation equivariant family of linear maps between P -tensors if

$$\phi_{\tau_1 \circ \mathcal{D}_1, \tau_2 \circ \mathcal{D}_2}(\tau_1 \circ T_1) = \tau_2 \circ (\phi_{\mathcal{D}_1, \mathcal{D}_2}(T_1))$$

for any P -tensor T_1 with reference domain \mathcal{D}_1 and any pair of permutations $\tau_1 \in \mathbb{S}_{|\mathcal{D}_1|}$ and $\tau_2 \in \mathbb{S}_{|\mathcal{D}_2|}$.

In the next section we will see that such families of equivariant maps can be defined in a relatively straightforward manner by transforming both the source and destination P -tensors to a canonical position, where the shared atoms between their reference domains occupy the first $d^\cap = |\mathcal{D}_1 \cap \mathcal{D}_2|$ positions.

In addition to equivariance to local permutations, we also need to consider global relabelings $\sigma: \mathcal{U} \rightarrow \mathcal{U}$ of the entire universe of atoms. The effect of such a relabeling is to map $\mathcal{D} = (x_1, \dots, x_d)$ to $\sigma \bullet \mathcal{D} = (\sigma(x_1), \dots, \sigma(x_d))$.

Definition 7 (Relabeling invariant maps between P -tensors). A family of linear maps between P -tensors

$$\phi_{\mathcal{D}_1, \mathcal{D}_2}: (T_1, \mathcal{D}_1) \mapsto (T_2, \mathcal{D}_2)$$

is said to be invariant to global relabelings if

$$\phi_{\sigma \bullet \mathcal{D}_1, \sigma \bullet \mathcal{D}_2} = \phi_{\mathcal{D}_1, \mathcal{D}_2}$$

for any permutation σ of the universe \mathcal{U} of atoms.

A key result of our paper, proved in the Appendix, is the following theorem, showing that the above two constraints are exactly what is needed to build equivariant higher order MPNNs out of P -tensors.

Theorem 1. Any higher order MPNN in which

- the subgraphs in each layer are selected using an invariant subgraph selection rule;
- the output of each subgraph neuron is a P -tensor;
- the messages sent from the P -tensors in each layer \mathcal{F}^{in} to the P -tensors in following layer \mathcal{F}^{out} are linear and satisfy the conditions of Definitions 6 and 7

is a permutation equivariant MPNN in the sense of Definition 4.

4.1 Message passing between P -tensors with the same reference domain

If $\mathcal{D}_1 = (x_1, \dots, x_d)$ and $\mathcal{D}_2 = (x'_1, \dots, x'_d)$ are equal as sets, they can only differ by a permutation μ mapping $x'_1 = x_{\mu(1)}$, $x'_2 = x_{\mu(2)}$, and so on. In any kind of permutation equivariant learning scenario such as graph neural networks, it is always possible to determine which labels refer to the same object, so, without loss of generality, we can assume that the elements of

\mathcal{P}	ϕ
$\{\{1\}, \{2\}, \{3\}, \{4\}\}$	$T_{a,b}^{\text{out}} = \sum_{c,d} T_{c,d}^{\text{in}}$
$\{\{1\}, \{2\}, \{3, 4\}\}$	$T_{a,b}^{\text{out}} = \sum_c T_{c,c}^{\text{in}}$
$\{\{1\}, \{2, 4\}, \{3\}\}$	$T_{a,b}^{\text{out}} = \sum_c T_{c,b}^{\text{in}}$
$\{\{1\}, \{2, 3\}, \{4\}\}$	$T_{a,b}^{\text{out}} = \sum_c T_{b,c}^{\text{in}}$
$\{\{2\}, \{1, 4\}, \{3\}\}$	$T_{b,a}^{\text{out}} = \sum_c T_{c,b}^{\text{in}}$
$\{\{2\}, \{1, 3\}, \{4\}\}$	$T_{b,a}^{\text{out}} = \sum_c T_{b,c}^{\text{in}}$
$\{\{1, 2\}, \{3\}, \{4\}\}$	$T_{a,a}^{\text{out}} = \sum_{b,c} T_{b,c}^{\text{in}}$
$\{\{1\}, \{2, 3, 4\}\}$	$T_{a,b}^{\text{out}} = T_{b,b}^{\text{in}}$
$\{\{2\}, \{1, 3, 4\}\}$	$T_{b,a}^{\text{out}} = T_{b,b}^{\text{in}}$
$\{\{1, 2, 3\}, \{4\}\}$	$T_{a,a}^{\text{out}} = \sum_b T_{a,b}^{\text{in}}$
$\{\{1, 2, 4\}, \{3\}\}$	$T_{a,a}^{\text{out}} = \sum_b T_{b,a}^{\text{in}}$
$\{\{1, 2\}, \{3, 4\}\}$	$T_{a,a}^{\text{out}} = \sum_c T_{c,c}^{\text{in}}$
$\{\{1, 3\}, \{2, 4\}\}$	$T_{a,b}^{\text{out}} = T_{a,b}^{\text{in}}$
$\{\{1, 4\}, \{2, 3\}\}$	$T_{a,b}^{\text{out}} = T_{b,a}^{\text{in}}$
$\{\{1, 2, 3, 4\}\}$	$T_{a,a}^{\text{out}} = T_{a,a}^{\text{in}}$

Table 1: The $B(4) = 15$ possible partitions of the set $\{1, 2, 3, 4\}$ and the corresponding permutation equivariant linear maps $\phi: \mathbb{R}^{k \times k} \rightarrow \mathbb{R}^{k \times k}$ as derived by Maron et al. Maron et al. (2019b).

\mathcal{D}_2 have been rearranged so that $x'_1 = x_1, \dots, x'_d = x_d$. This reduces the problem of satisfying Definition 6 to that of “ordinary” permutation equivariant vector, matrix, etc., valued neural network layers, which, by now, is a well studied subject.

In the first order case, P -tensors are simply vectors, so deriving the rules of permutation equivariant message passing reduces to finding the space of linear maps $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfying

$$\phi([\mathbf{v}]_{\tau^{-1}(i)}) = [\phi(\mathbf{v})]_{\tau^{-1}(i)}$$

for any $\mathbf{v} \in \mathbb{R}^d$ and any $\tau \in \mathbb{S}_d$. The seminal Deep Sets paper Zaheer et al. (2017) proved that in this case ϕ can have at most two (learnable) parameters λ_1 and λ_2 , and must be of the form

$$\phi(\mathbf{v}) = \lambda_1 \mathbf{v} + \lambda_2 \mathbf{1} \mathbf{1}^\top \mathbf{v}.$$

In the more general case, ϕ maps a k_1 ’th order tensor T^{in} to a k_2 ’th order tensor T^{out} , both transforming under permutations as in (6), leading to the equivariance condition

$$[\phi(T^{\text{in}})]_{\tau^{-1}(i_1), \dots, \tau^{-1}(i_{k_2})} = \phi(T_{\tau^{-1}(i_1), \dots, \tau^{-1}(i_{k_1})}^{\text{in}}). \quad (7)$$

The characterization of the space of equivariant maps for this case was given in a similarly influential paper (Maron et al., 2019b).

Proposition 2 (Maron et al.). The space of linear maps $\phi: \mathbb{R}^{d^{k_1}} \rightarrow \mathbb{R}^{d^{k_2}}$ that is equivariant to permutations $\tau \in \mathbb{S}_d$ in the sense of (7) is spanned by a basis indexed by the partitions of the set $\{1, 2, \dots, k_1 + k_2\}$.

Since the number of partitions of $\{1, 2, \dots, m\}$ is given by the so-called *Bell number* $B(m)$, according to this

result, the number of learnable parameters in such a map is $B(k_1 + k_2)$. To be specific, the equivariant linear map corresponding to a given partition \mathcal{P} can be written as a composition of three operations: (1) *summing* over specific dimensions or diagonals of T^{in} ; (2) *transferring* this result to the output tensor by identifying some combinations of input indices with output indices; (3) *broadcasting* the result along certain dimensions or diagonals of the output tensor T^{out} . These three operations correspond to the three different types of parts that can appear in \mathcal{P} : those parts that only involve the second k_2 numbers, those that involve a mixture of the first k_1 and second k_2 , and those that only involve the first k_1 . We shall say that \mathcal{P} is of type (p_1, p_2, p_3) if it has p_1 parts of the first category, p_2 of the second and p_3 of the third. In all three categories, a part $\{j_1, \dots, j_\ell\}$ appearing in \mathcal{P} implies that the corresponding indices are tied together. The way to distinguish between input and output indices is that if $1 \leq j_q \leq k_1$ then it refers to the j_q 'th index of T^{out} , whereas if $k_1 + 1 \leq j_q \leq k_1 + k_2$, then it refers to the $j_q - k_1$ index of T^{in} .

As an example, in the $k_1 = k_2 = 3$ case, the partition $\mathcal{P} = \{\{1, 3\}, \{2, 5, 6\}, \{4\}\}$ corresponds to (a) summing T^{in} along its first dimension (corresponding to $\{4\}$) (b) transferring the diagonal along the second and third dimensions of T^{in} to the second dimension of T^{out} (corresponding to $\{2, 5, 6\}$), (c) broadcasting the result along the diagonal of the first and third dimensions (corresponding to $\{1, 3\}$). Explicitly, this gives the equivariant map

$$T_{a,b,a}^{\text{out}} = \sum_c T_{c,b,b}^{\text{in}}. \quad (8)$$

Since $B(6) = 203$, listing all other possible maps for $k_1 = k_2 = 3$ would be very laborious. In Table 1 we list the possible equivariant maps for the $k_1 = k_2 = 2$ case.

The presence of multiple channels enriches this picture only to the extent that each input channel can be linearly mixed with each output channel. For example, in the case of C_{in} channels in T^{in} and C_{out} channels in T^{out} , (8) becomes

$$T_{a,b,a,\alpha}^{\text{out}} = \sum_\beta W_{\alpha,\beta} \sum_c T_{c,b,b,\beta}^{\text{in}}$$

for some (learnable) weight matrix $W \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}}}$. This type of linear mixing across channels can be separated from the equivariant message passing operation itself, which consists of applying (8) to each channel separately.




Figure 2: Given two P -tensors T^{in} and T^{out} whose reference domains have d^\cap atoms in common, without loss of generality we can rearrange the two tensors so that the indices corresponding to the common atoms appear first. Mapping the corresponding $d^\cap \times d^\cap \times \dots \times d^\cap$ subtensor of T^{in} to the analogous $d^\cap \times d^\cap \times \dots \times d^\cap$ subtensor of T^{out} with any of the $B(k_1 + k_2)$ linear maps described in Section 4.1 is an equivariant operation. The additional equivariant operations correspond to similar maps except with the summations and broadcast operations extending over not just the overlapping part of the tensors but the entirety of T^{in} or T^{out} .

4.2 Message passing between P-tensors with different reference domains

Our second main result is the following theorem, which generalizes Proposition 2 to when the reference domains of the input and output tensors only partially overlap.

Theorem 3. *Let T_1 and T_2 be two P -tensors with reference domains \mathcal{D}_1 and \mathcal{D}_2 such that $|\mathcal{D}_1 \cap \mathcal{D}_2| \geq 2$ and $\mathcal{D}_1 \not\subseteq \mathcal{D}_2$ and $\mathcal{D}_2 \not\subseteq \mathcal{D}_1$. Then for each partition \mathcal{P} of $\{1, \dots, k_1 + k_2\}$ of type (p_1, p_2, p_3) there are $2^{p_1 + p_3}$ linearly independent permutation equivariant linear maps $\phi: T_1 \mapsto T_2$.*

To derive the form of the actual maps, without loss of generality, we assume that \mathcal{D}_1 and \mathcal{D}_2 have been reordered in such a way that the $d^\cap = |\mathcal{D}_1 \cap \mathcal{D}_2|$ atoms that they have in common occupy the first d^\cap positions in both, and are listed in the same order. An easy generalization of our previous results is to associate to each partition the same map as before, except we now only transfer information from the subtensor of T^{in} cut out by the common atoms to the subtensor of T^{out} cut out by the same (Figure 2 left). For example, the counterpart of (8) would be

$$T_{a,b,a}^{\text{out}} = \begin{cases} \sum_{c=1}^{d^\cap} T_{c,b,b}^{\text{in}} & a, b \leq d^\cap \\ 0 & \text{otherwise.} \end{cases}$$

This, however, would only give us $B(k_1 + k_2)$ equivariant maps, like in the previous section.

The additional factor of $2^{p_1 + p_3}$ comes from the fact that for any partition that has parts purely involving

(k_1, k_2)	# of maps in $\mathcal{D}_1 = \mathcal{D}_2$ case	# of maps in $\mathcal{D}_1 \neq \mathcal{D}_2$ case
(0, 0)	1	1
(1, 1)	2	5
(1, 2)	5	17
(2, 2)	15	61
(2, 3)	52	321
(3, 3)	203	769

Figure 3: The number of independent equivariant linear maps from a k_1 'th order P -tensor to a k_2 'th order P -tensor when the reference domains are the same vs. when they overlap only partially.

indices of T^{in} or T^{out} , each of the corresponding operations can extend across either just the common atoms, or all atoms of the given tensor. For our running example $\mathcal{P} = \{\{1, 3\}, \{2, 5, 6\}, \{4\}\}$, we have the three additional equivariant maps

$$\begin{aligned}
 T_{a,b,a}^{\text{out}} &= \begin{cases} \sum_{c=1}^{d^\cap} T_{c,b,b}^{\text{in}} & b \leq d^\cap \\ 0 & \text{otherwise,} \end{cases} & a \in \{1, \dots, d_2\} \\
 T_{a,b,a}^{\text{out}} &= \begin{cases} \sum_{c=1}^{d_1} T_{c,b,b}^{\text{in}} & a, b \leq d^\cap \\ 0 & \text{otherwise,} \end{cases} & a \in \{1, \dots, d^\cap\} \\
 T_{a,b,a}^{\text{out}} &= \begin{cases} \sum_{c=1}^{d_1} T_{c,b,b}^{\text{in}} & b \leq d^\cap \\ 0 & \text{otherwise.} \end{cases} & a \in \{1, \dots, d_2\}.
 \end{aligned}$$

Table 3 contrasts the number of possible equivariant maps in this case to the number of maps described in the previous section. As before, in the presence of multiple channels, each of the above maps can also involve mixing the different channels by a learnable weight matrix.

5 P -TENSORS IN GRAPH NEURAL NETWORKS

In this section we describe how the P -tensor formalism can be used to build expressive graph neural networks, and how various existing GNNs reduce to special cases.

5.1 Zeroth order message passing networks

In classical message passing networks the reference domain of each neuron consists of just a single vertex. Talking about the transformation properties of an object with respect to permutations of a single atom is vacuous, so in this case all the activations are zeroth order P -tensors. The only type of equivariant message that a zeroth order vertex neuron $T_c^{v_1}$ can send to another vertex neuron $T_c^{v_2}$ is $T_c^{v_2} \leftarrow T_c^{v_1}$ (here and in the following c stand for the channel index). Applying this operation to all the neighbors of a given vertex and adding a bias term plus mixing with a learnable

matrix W leads exactly to the update rule (2). Thus, classical message passing networks just correspond to message passing between zeroth order P -tensors.

5.2 Edge networks

One of the first extensions of the MPNN formalism was the introduction of networks that can pass messages not just from vertices to vertices but also from vertices to edges and edges to vertices Gilmer et al. (2017). In the P -tensor formalism a neuron corresponding to the edge $\overline{v_1 v_2}$ has receptive domain $\mathcal{D} = (v_1, v_2)$. If the edge P -tensor $T^{\overline{v_1 v_2}}$ is zeroth order, then the rule for sending messages from the vertex P -tensors T^{v_1} and T^{v_2} will essentially be the same as above.

However, when $T^{\overline{v_1 v_2}}$ is a first order P -tensor we have two different possible equivariant maps: the ‘‘concatenating map’’

$$T_{i,c}^{\overline{v_1 v_2}} \leftarrow T_c^{v_i} \quad i \in \{1, 2\},$$

and the ‘‘averaging map’’

$$T_{i,c}^{\overline{v_1 v_2}} \leftarrow T_c^{v_1} + T_c^{v_2} \quad i \in \{1, 2\}.$$

To maximize expressivity, the edge neuron would concatenate these two messages, effectively doubling the number of channels as we pass from the ‘‘vertex layer’’ neurons to the ‘‘edge layer’’.

Edge-to-vertex message passing similarly affords two distinct types of linear maps:

$$T_c^{v_i} \leftarrow T_{i,c}^{\overline{v_1 v_2}} \quad \text{and} \quad T_c^{v_i} \leftarrow T_{1,c}^{\overline{v_1 v_2}} + T_{2,c}^{\overline{v_1 v_2}}.$$

Finally, for passing a message from a first order edge $T^{(v_1, v_2)}$ to another first order edge $T^{(v_2, v_3)}$ (with $v_1 \neq v_3$) we have one linear map corresponding to the $\{\{1, 2\}\}$ partition,

$$T_{1,c}^{\overline{v_2 v_3}} \leftarrow T_{2,c}^{\overline{v_1 v_2}}$$

and three different maps corresponding to the $\{\{1\}, \{2\}\}$ partition:

$$\begin{aligned}
 T_{i,c}^{\overline{v_2 v_3}} &\leftarrow T_{2,c}^{\overline{v_1 v_2}} & i \in \{1, 2\} \\
 T_{1,c}^{\overline{v_2 v_3}} &\leftarrow T_{1,c}^{\overline{v_1 v_2}} + T_{2,c}^{\overline{v_1 v_2}} \\
 T_{i,c}^{\overline{v_2 v_3}} &\leftarrow T_{1,c}^{\overline{v_1 v_2}} + T_{2,c}^{\overline{v_1 v_2}} & i \in \{1, 2\}.
 \end{aligned}$$

While individually these operations are simple and could be derived by hand on a case-by-case basis, as the size of the reference domains (as well as the order of the tensors) increases, enumerating and separately implementing all possibilities in code becomes unwieldy.

5.3 Message passing between subgraphs

The true power of the P -tensors model manifests in message passing between larger subgraphs. Take for example the paradigmatic case from chemistry of two adjacent benzene rings (six-atom rings of carbon atoms) made up of vertices $\{v_1, \dots, v_6\}$ and $\{v_1, v_2, v_7, v_8, v_9, v_{10}\}$ represented by first order P -tensors T^{src} and T^{dest} .

The linear map corresponding to the $\{\{1, 2\}\}$ partition,

$$T_{i,c}^{\text{dest}} \leftarrow \begin{cases} T_{i,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise} \end{cases}$$

transfers the rows of T^{src} corresponding to the two shared carbon atoms to the corresponding rows of T^{dest} . We have four maps corresponding to $\{\{1\}, \{2\}\}$:

1. Transferring the sum of the rows of the shared carbons to the corresponding rows:

$$T_{i,c}^{\text{dest}} \leftarrow \begin{cases} \sum_{j=1}^2 T_{j,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise,} \end{cases}$$

2. Transferring the sum of the rows of the shared carbons to all rows:

$$T_{i,c}^{\text{dest}} \leftarrow \sum_{j=1}^2 T_{j,c}^{\text{src}},$$

3. Transferring the sum of all rows to the rows of the shared carbons:

$$T_{i,c}^{\text{dest}} \leftarrow \begin{cases} \sum_{j=1}^6 T_{j,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise,} \end{cases}$$

4. Transferring the sum of all rows to all rows:

$$T_{i,c}^{\text{dest}} \leftarrow \sum_{j=1}^6 T_{j,c}^{\text{src}}.$$

If we concatenate the results of all of these maps, T^{dest} will have five times as many channels as T^{src} .

The second-order to second-order case is even more interesting, since in this case the i, j slice of $T_{i,j,c}^{\text{src}}$ and $T_{i,j,c}^{\text{dest}}$ can effectively represent interactions between the i 'th and j 'th atoms in the two rings. Space limitations prevent us from listing all 61 possible maps, but some examples are the following:

1. Transferring the sum of the interactions of shared carbon i with the other carbons in the ring to the corresponding slice:

$$T_{i,i,c}^{\text{dest}} \leftarrow \begin{cases} \sum_{k=1}^6 T_{i,k,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise,} \end{cases}$$

2. Transferring the sum of all interactions to the shared carbons:

$$T_{i,i,c}^{\text{dest}} \leftarrow \begin{cases} \sum_{k=1}^6 \sum_{\ell=1}^6 T_{k,\ell,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise,} \end{cases}$$

3. Transferring the sum of all self-interactions to the shared carbons:

$$T_{i,i,c}^{\text{dest}} \leftarrow \begin{cases} \sum_{k=1}^6 T_{k,k,c}^{\text{src}} & i \in \{1, 2\} \\ 0 & \text{otherwise,} \end{cases}$$

Note that it is often the linear combinations of these maps, e.g., the second map above minus the third map, that have the most intuitive interpretations. Also note that many of these maps occur naturally in other subgraph, hypergraph and simplicial complex networks. The advantage of our formalism is in being to enumerate (and efficiently implement) all possible equivariant maps in a systematic way.

6 EXPERIMENTAL RESULTS

We validated our model on several molecular datasets with one of the simplest possible realizations of the P -tensors framework. The subgraphs are limited to vertices, edges and cycles (of any length) and the order of the corresponding neurons is either zero or one.

The interactions between vertices and edges can be conceptualized as a classical MPNN: vertices send messages to edges, then messages are passed back to vertices, and vertices update from the incoming messages and their previous state. In this interaction, we are also able to get a partial update for the edges using the information from the vertices. We similarly perform message passing between edges and cycles, with two notable distinctions: (1) since cycles are first order representations, we consider both internal linear maps when updating them, as apposed to the single linear map for vertex and edge updates; (2) We limit the edge-cycle interactions to edges that are fully encapsulated by cycles that they are interacting with. In some of our models we also used cycle-cycle interactions. In other respects, such as the placement of MLPs, etc., our model is similar to Cellular Isomorphism Networks (CIN) (Bodnar et al., 2021a). For details, please see the supplementary materials.

We report results on four molecular datasets that are standard benchmarks in the literature: (1) The full ZINC dataset of almost 250K organic molecules (Sterling and Irwin, 2015; Gómez-Bombarelli et al., 2018); (2) its more commonly used subset of just 12K molecules; (3) the OGBG-MolHIV classification dataset of 41K molecules Hu et al. (2020a); (4) the TOX21 property prediction benchmark on 78311 molecules (Hu et al., 2020b). As baselines we use the best performing alternative algorithms from the literature of the same class, in particular, we do not compare to transformers or algorithms that explicitly take into account the 3D positions of atoms.

Remarkably, our P -tensors based equivariant message

	ZINC-12K MAE(↓)	ZINC-Full MAE(↓)	OGBG-MOLHIV ROC-AUC(% ↑)	TOX21 ROC-AUC(% ↑)
RP-NGF (Murphy et al., 2019)	–	–	–	0.79.4 ± 1.00
GCN (Kipf and Welling, 2017)	0.321 ± 0.009	–	76.07 ± 0.97	–
GIN (Xu et al., 2019)	0.408 ± 0.008	0.088 ± 0.002	75.58 ± 1.40	–
GINE (Hu et al., 2020b)	0.252 ± 0.014	0.088 ± 0.002	75.58 ± 1.40	86.68 ± 0.77
PNA (Corso et al., 2020)	0.133 ± 0.011	0.320 ± 0.032	79.05 ± 1.32	–
HIMP (Fey et al., 2020)	0.151 ± 0.002	0.036 ± 0.002	78.80 ± 0.82	87.36 ± 0.50
CIN (Bodnar et al., 2021a)	0.079 ± 0.006	0.022 ± 0.002	80.94 ± 0.57	–
DS-GNN (EGO+) (Bevilacqua et al., 2022)	0.105 ± 0.003	–	77.40 ± 2.19	76.39 ± 1.18
DSS-GNN (EGO+) (Bevilacqua et al., 2022)	0.097 ± 0.006	–	76.78 ± 1.66	77.95 ± 0.40
GNN-AK+ (Zhao et al., 2022)	0.091 ± 0.011	–	79.61 ± 1.19	–
SUN (EGO+) (Frasca et al., 2022)	0.084 ± 0.002	–	80.03 ± 0.55	–
First order P-tensors (our model)	0.071 ± 0.004	0.024 ± 0.001	80.76 ± 0.82	84.95 ± 0.58

Table 2: Experimental results on molecular datasets with baselines taken from (Frasca et al., 2022; Bodnar et al., 2021a; Bevilacqua et al., 2022).

passing algorithm is competitive with the best performing algorithms on each dataset, and beats all the other algorithms on ZINC 12K. We hypothesize that this is a direct result of the algorithm’s greater expressivity. Additional experimental data can be found in the Supplementary Materials.

7 CONCLUSIONS

The P -tensors framework unifies and generalizes equivariant message passing across a range of subgraph neural network models, as well as some other models, such as certain simplicial complex neural networks, that are not based on subgraphs *per se*, but ultimately still depend on the same equivariance constraints. The experimental results suggest that even in the first order case, the increased expressive power of our model helps improve upon the performance of other graph neural networks on standard molecular benchmarks.

One of the advantages of our framework is that instead of implementing each type of subgraph interaction separately, in a piecemeal fashion, it makes it possible to formulate higher order message passing as a generic computational paradigm that is reusable across many models. In ongoing work we are developing a software library for higher order message passing that follows this philosophy.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). Ten-

sorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.

Alsentzer, E., Finlayson, S., Li, M., and Zitnik, M. (2020). Subgraph neural networks. *Advances in Neural Information Processing Systems*, 33:8017–8029.

Atwood, J. and Towsley, D. (2016). Diffusion-convolutional neural networks. *Advances in Neural Information Processing Systems*, 29.

Bevilacqua, B., Frasca, F., Lim, D., Srinivasan, B., Cai, C., Balamurugan, G., Bronstein, M. M., and Maron, H. (2022). Equivariant subgraph aggregation networks. In *International Conference on Learning Representations*.

Bodnar, C., Frasca, F., Otter, N., Wang, Y., Lio, P., Montufar, G. F., and Bronstein, M. (2021a). Weisfeiler and Lehman go cellular: CW networks. *Advances in Neural Information Processing Systems*, 34:2625–2640.

Bodnar, C., Frasca, F., Wang, Y., Otter, N., Montufar, G. F., Lio, P., and Bronstein, M. (2021b). Weisfeiler and Lehman go topological: Message passing simplicial networks. In *International Conference on Machine Learning*.

Chen, Z., Chen, L., Villar, S., and Bruna, J. (2020). Can graph neural networks count substructures? *Advances in Neural Information Processing Systems*, 33.

Corso, G., Cavalleri, L., Beaini, D., Liò, P., and Velivcković, P. (2020). Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33.

de Haan, P., Cohen, T. S., and Welling, M. (2020). Natural graph networks. *Advances in Neural Information Processing Systems*, 33:3636–3646.

Dong, Y., Sawin, W., and Bengio, Y. (2020).

- HNHN: hypergraph networks with hyperedge neurons. *CoRR*, abs/2006.12278.
- Ebli, S., Defferrard, M., and Spreemann, G. (2020). Simplicial neural networks. In *Topological Data Analysis and Beyond workshop at NeurIPS*.
- Falcon, W. and The PyTorch Lightning team (2019). PyTorch Lightning.
- Feng, Y., You, H., Zhang, Z., Ji, R., and Gao, Y. (2019). Hypergraph neural networks. *33rd AAAI Conference on Artificial Intelligence*.
- Ferreira, R., Grossi, R., Rizzi, R., Sacomoto, G., and Sagot, M.-F. (2014). Amortized-delay algorithm for listing chordless cycles in undirected graphs. In *European Symposium on Algorithms*, pages 418–429. Springer.
- Fey, M. and Lenssen, J. E. (2019). Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Fey, M., Yuen, J. G., and Weichert, F. (2020). Hierarchical inter-message passing for learning on molecular graphs. In *ICML Graph Representation Learning and Beyond (GRL+) Workshop*.
- Frasca, F., Bevilacqua, B., Bronstein, M., and Maron, H. (2022). Understanding and extending subgraph GNNs by rethinking their symmetries. *Advances in Neural Information Processing Systems*, 35.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International conference on machine learning*.
- Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. (2018). Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276.
- Hu, W., Chan, Z., Liu, B., Zhao, D., Ma, J., and Yan, R. (2019). GSN: A graph-structured network for multi-party dialogues. *CoRR*, abs/1905.13637.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. (2020a). Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. (2020b). Strategies for pre-training graph neural networks. In *International Conference on Learning Representations*.
- Hy, T. S., Trivedi, S., Pan, H., Anderson, B. M., and Kondor, R. (2018). Predicting molecular properties with covariant compositional networks. *The Journal of Chemical Physics*, 148(24):241745.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*.
- Keriven, N. and Peyré, G. (2019). Universal invariant and equivariant graph neural networks. In *Advances in Neural Information Processing Systems 32*.
- Kipf, T. N. and Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Kwon, J., Kim, J., Park, H., and Choi, I. K. (2021). Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*.
- Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y. (2019a). Provably powerful graph networks. *Advances in Neural Information Processing Systems*, 32.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. (2019b). Invariant and equivariant graph networks. In *International Conference on Learning Representations*.
- Maron, H., Fetaya, E., Segol, N., and Lipman, Y. (2019c). On the universality of invariant networks. In *International conference on machine learning*.
- Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. (2019). Weisfeiler and Leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*.
- Murphy, R., Srinivasan, B., Rao, V., and Ribeiro, B. (2019). Relational pooling for graph representations. In *International Conference on Machine Learning*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*.
- Rong, Y., Huang, W., Xu, T., and Huang, J. (2020). Dropedge: Towards deep graph convolutional networks on node classification. In *International Conference on Learning Representations*.
- Sagan, B. E. (2001). *The Symmetric Group*. Springer New York.

- Sannai, A., Takai, Y., and Cordonnier, M. (2019). Universal approximations of permutation invariant/equivariant functions by deep neural networks. *CoRR*, abs/1903.01939.
- Sardellitti, S., Barbarossa, S., and Testa, L. (2021). Topological signal processing over cell complexes. In *2021 55th Asilomar Conference on Signals, Systems, and Computers*.
- Sterling, T. and Irwin, J. J. (2015). ZINC 15 Ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337. PMID: 26479676.
- Thiede, E., Zhou, W., and Kondor, R. (2021). Autobahn: Automorphism-based graph neural nets. *Advances in Neural Information Processing Systems*, 34.
- Thiede, E. H., Hy, T., and Kondor, R. (2020). The general theory of permutation equivariant neural networks and higher order graph variational encoders. *CoRR*, abs/2004.03990.
- Weisfeiler, B. and Leman, A. (1968). The reduction of a graph to canonical form and the algebra which appears therein. *NTI, Series*, pages 12–16.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2019). How powerful are graph neural networks? In *International Conference on Learning Representations*.
- You, J., Gomes-Selman, J. M., Ying, R., and Leskovec, J. (2021). Identity-aware graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., and Smola, A. J. (2017). Deep sets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*.
- Zhang, M., Cui, Z., Neumann, M., and Chen, Y. (2018). An end-to-end deep learning architecture for graph classification. In *Proceedings of the AAAI conference on artificial intelligence*.
- Zhao, L., Jin, W., Akoglu, L., and Shah, N. (2022). From stars to subgraphs: Uplifting any GNN with local structure awareness. In *International Conference on Learning Representations*.
- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [No]
 - (c) Clear explanations of any assumptions. [Yes]
 3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Not Applicable]

Appendix

1 Proofs

Theorem 1. *Any higher order MPNN in which*

(a) *the subgraphs in each layer are selected using an invariant subgraph selection rule;*

(b) *the output of each subgraph neuron is a P-tensor;*

(c) *the messages sent from the P-tensors in each layer \mathcal{F}^{in} to the P-tensors in following layer \mathcal{F}^{out} are linear and satisfy the conditions of Definitions 6 and 7*

is a permutation equivariant MPNN in the sense of Definition 4.

Proof of Theorem 1. Let $\mathcal{F}^{\text{in}} = \{T_1^{\text{in}}, \dots, T_m^{\text{in}}\}$ be a layer of P-tensors that sends messages to another layer $\mathcal{F}^{\text{out}} = \{T_1^{\text{out}}, \dots, T_m^{\text{out}}\}$. Let $\mathcal{F}^{\text{in}'}$ and $\mathcal{F}^{\text{out}'}$ be the corresponding layers after permuting the vertices of the underlying graph by some permutation σ .

By the invariance of the subset selection rule, for each T_a^{in} with reference domain $\mathcal{D} = (x_1, \dots, x_d)$, there is a corresponding $T_a^{\text{in}'}$ with reference domain $\mathcal{D}' = (\sigma(x_{\tau^{-1}(1)}), \dots, \sigma(x_{\tau^{-1}(d)})) = \sigma \bullet \tau \circ \mathcal{D}$ for some permutation $\tau \in \mathbb{S}_d$. Similarly, for any T_a^{out} with reference domain $\mathcal{D} = (\bar{x}_1, \dots, \bar{x}_{\bar{d}})$, there is a corresponding $T_a^{\text{out}'}$ with reference domain $\mathcal{D}' = (\sigma(\bar{x}_{\bar{\tau}^{-1}(1)}), \dots, \sigma(\bar{x}_{\bar{\tau}^{-1}(\bar{d})})) = \sigma \bullet \bar{\tau} \circ \mathcal{D}$ for some $\bar{\tau} \in \mathbb{S}_{\bar{d}}$.

Assuming that \mathcal{F}^{in} is a covariant layer, its P-tensors before and after permutation are related by $T_a^{\text{in}'} = \tau \circ T_a^{\text{in}}$. If the message passing process mapping $\mathcal{F}^{\text{in}} \mapsto \mathcal{F}^{\text{out}}$ is relabeling invariant and permutation equivariant in the sense of Definitions 6 and 7,

$$T_a^{\text{out}'} = \phi_{\mathcal{D}', \mathcal{D}'}(T_a^{\text{in}'}) = \phi_{\sigma \bullet \tau \circ \mathcal{D}, \sigma \bullet \bar{\tau} \circ \mathcal{D}}(T_a^{\text{in}'}) = \bar{\tau} \circ \phi_{\mathcal{D}, \mathcal{D}}(T_a^{\text{in}}) = \bar{\tau} \circ T_a^{\text{out}}$$

showing that \mathcal{F}^{out} is also covariant. Since this relationship holds for any $(\mathcal{F}^{\text{in}}, \mathcal{F}^{\text{out}})$ pair of layers, the network as a whole is equivariant. ■

2 Additional experiments

Our experimental approach is designed to show how introducing higher order features can improve performance amongst a wide range of datasets. To this end, we pick a model closely related to various ones found in the literature, but modified to included these extra features and make use of the given P-Tensor operations. In addition to the results presented in the main body of the paper we also ran experiments on the classic TU datasets (Table 3).

Tested Variations

Cycle-Cycle interactions. In addition to having first order and zeroth order interactions, we also considered cycle-cycle interactions, with their corresponding five linear maps. We mainly considered this on ZINC, MolTox21, and OGBG-MolHIV, but noticed a spike in validation volatility towards the end of training, damaging the reliability for the model in many cases. A similar issue came up in (Frasca et al., 2022) for training on OGBG-MolHIV, and the issue was alleviated by using the ASAM optimizer (Kwon et al., 2021), which is designed to reduce sharpness during training. We found that this had limited success on the same dataset for our model. Our hypothesis is that for highly irregular interactions, batch normalization struggles to converge because of the natural volatility that comes with such irregular structures.

Dataset	MUTAG	PTC	PROTEINS	NCI1	IMDB-B	IMDB-M
DCNN(Atwood and Towsley, 2016)	–	–	61.3 ± 1.6	56.6 ± 1.0	49.1 ± 1.4	33.5 ± 1.4
DGCNN(Zhang et al., 2018)	85.8 ± 1.8	58.6 ± 2.5	75.5 ± 0.9	74.4 ± 0.5	70.0 ± 0.9	47.8 ± 0.9
IGN(Keriven and Peyré, 2019)	83.9 ± 13.0	58.5 ± 6.9	76.6 ± 5.5	74.3 ± 2.7	72.0 ± 5.5	48.7 ± 3.4
PPGNs(Maron et al., 2019a)	90.6 ± 8.7	66.2 ± 6.6	77.2 ± 4.7	83.2 ± 1.1	73.0 ± 5.8	50.5 ± 3.6
Natural GN(de Haan et al., 2020)	89.4 ± 1.6	66.8 ± 1.7	71.7 ± 1.0	82.4 ± 1.3	73.5 ± 2.0	51.3 ± 1.5
GSN(Hu et al., 2019)	92.2 ± 7.5	68.2 ± 7.2	76.6 ± 5.0	83.5 ± 2.0	77.8 ± 3.3	54.3 ± 3.3
SIN(Bodnar et al., 2021b)	–	–	76.4 ± 3.3	82.7 ± 2.1	75.6 ± 3.2	52.7 ± 3.1
CIN(Bodnar et al., 2021a)	92.7 ± 6.1	68.2 ± 5.6	77.0 ± 4.3	83.6 ± 1.4	75.6 ± 3.7	52.7 ± 3.1
GIN(Xu et al., 2019)	89.4 ± 5.6	64.6 ± 7.0	76.2 ± 2.8	82.7 ± 1.7	75.1 ± 5.1	52.3 ± 2.8
GIN + ID-GNN(You et al., 2021)	90.4 ± 5.6	67.2 ± 4.3	75.4 ± 2.7	82.6 ± 1.6	76.0 ± 2.7	52.7 ± 4.2
DropEdge(Rong et al., 2020)	91.0 ± 5.7	64.5 ± 2.6	73.5 ± 4.5	82.0 ± 2.6	76.5 ± 3.3	52.8 ± 2.8
DS-GNN (GIN) (ND)(Bevilacqua et al., 2022)	89.4 ± 4.8	66.3 ± 7.0	77.1 ± 4.6	83.8 ± 2.4	75.4 ± 2.9	52.7 ± 2.0
DS-GNN (GIN) (EGO)(Bevilacqua et al., 2022)	89.9 ± 6.5	68.6 ± 5.8	76.7 ± 5.8	81.4 ± 0.7	76.1 ± 2.8	52.6 ± 2.8
DS-GNN (GIN) (EGO+)(Bevilacqua et al., 2022)	91.1 ± 7.0	69.2 ± 6.5	75.9 ± 4.3	83.7 ± 1.8	77.1 ± 3.0	53.2 ± 2.8
DSS-GNN (GIN) (ND) (Bevilacqua et al., 2022)	91.0 ± 3.5	66.3 ± 5.9	76.1 ± 3.4	83.6 ± 1.5	76.1 ± 2.9	53.3 ± 1.9
DSS-GNN (GIN) (EGO) (Bevilacqua et al., 2022)	91.0 ± 4.7	68.2 ± 5.8	76.7 ± 4.1	83.6 ± 1.8	76.5 ± 2.8	53.3 ± 3.1
DSS-GNN (GIN) (EGO+)(Bevilacqua et al., 2022)	91.1 ± 7.0	69.2 ± 6.5	75.9 ± 4.3	83.7 ± 1.8	77.1 ± 3.0	53.2 ± 2.4
GIN-AK+ (Zhao et al., 2022)	91.3 ± 7.0	67.8 ± 8.8	77.1 ± 5.7	85.0 ± 2.0	75.0 ± 4.2	–
SUN (GIN) (NULL)(Frasca et al., 2022)	91.6 ± 4.8	67.5 ± 6.8	76.8 ± 4.4	84.1 ± 2.0	76.2 ± 1.9	52.6 ± 3.2
SUN (GIN) (NM) (Frasca et al., 2022)	91.0 ± 4.7	67.0 ± 4.8	75.7 ± 3.4	84.2 ± 1.5	76.1 ± 2.9	53.1 ± 2.5
SUN (GIN) (EGO) (Frasca et al., 2022)	92.7 ± 5.8	67.2 ± 5.9	76.8 ± 5.0	83.7 ± 1.3	76.6 ± 3.4	52.7 ± 2.3
SUN (GIN) (EGO+)(Frasca et al., 2022)	92.1 ± 5.8	67.6 ± 5.5	76.1 ± 5.1	84.2 ± 1.5	76.3 ± 1.9	52.9 ± 2.8
Ours	92.9 ± 1.7	71.7 ± 5.2	75.9 ± 2.5	84.2 ± 1.7	77.9 ± 3.2	54.3 ± 2.0

Table 3: Summary of results on TUDatasets with baselines taken from (Frasca et al., 2022). Top three scores are given in red, purple, and bold.

Implementation

We implemented our algorithm in Pytorch (Paszke et al., 2019) using PyTorch Geometric (Fey and Lenssen, 2019). To add some encapsulation and improve runtime, we utilized pytorch lightning (Falcon and The PyTorch Lightning team, 2019) along with TensorBoard for visualization (Abadi et al., 2015). For cycle finding, we implemented (Ferreira et al., 2014), and ran it along with all other structure map finding during preprocessing. Thus, during runtime the only computations involving the structure of the graphs were the scatter operations themselves. We ran our experiments on an NVIDIA GeForce GTX 1080. As for running time, on MolHIV a training runs averaged 69.3 ± 1.2 minutes. The source code of our implementation can be found at <https://github.com/arhands/ptensors>.

Hyperparameter Selection

Aside from some initial variations for designing our model, we limited our hyperparameter search space to primarily consider learning rates, the number of layers used, and the reductions used between cycles and edges. Initial testing revealed high validation volatility, similar to that described in (Frasca et al., 2022), but we found that this was partially mitigated by reducing the momentum used in batch normalization, so we also considered that as part of our hyperparameters.

Our experimental setup for ZINC/ZINC-10K, and MolHIV is based on (Bodnar et al., 2021a; Sardellitti et al., 2021), while our setup for Tox21 is based on (Fey et al., 2020). We base our experimental setup for the TUDatasets on (Frasca et al., 2022), utilizing the same hyperparameter grid and base node/edge embeddings.

3 Elementary discussion of our models

For completeness in the following we give an elementary description of one of our models, showing how most P-tensor operations, at least in the first order case, can be implemented with customary GNN operations.

We will use μ_i to denote a generic multilayer perceptron (MLP) with the input parameters concatenated together channel-wise. In our experiments, each linear layer is followed by a batch normalization layer (Ioffe and Szegedy, 2015) and then a ReLU activation, unless specified otherwise. We also set multilayer MLPs to have double the hidden channels within their respective internal layers. Let $\mathcal{G} = (V, E)$ be a simple undirected graph. Let X^v and X^e be a zeroth order representation of the vertices and edges in \mathcal{G} , respectively. To help simplify notation, we shall consider a given edge $e_{i,j} = \{i, j\} = e_{j,i}$.

Our model can be then be described as a sequence of layers, each allowing for interactions between subgraphs of \mathcal{G} . To begin with, we shall define how cycles and vertices interact. Then, in each layer of our model, we capture

the expressive power of MPNNs by transferring from vertices to edges, and then back to vertices.

$$Y_i^v = \mu_1 \left(X_i^v, \sum_{e_{ij} \in E} \mu_2 \left(X_{e_{i,j}}^e, X_i^v + X_j^v \right) \right) \quad (9)$$

Where $i \in V$, μ_1 contains two layers, and μ_2 has a single layer. As can be seen in the above equation, the inner sum corresponds to the single linear map from vertices to edges, and the outer sum corresponds to the other direction. Similarly, we get the incoming messages from vertices to edges using a single transfer operation.

$$Y_{e_{i,j}}^{v \rightarrow e} = \mu_3 \left(X_{e_{i,j}}^e, X_i^v + X_j^v \right) \quad (10)$$

So far, this can largely be seen as classical message passing. However, where things get more interesting when we consider the interactions between edges. Let X^c denote a representation of selected cycles on \mathcal{G} (see hyperparameter subsection for per-dataset details) and, in similar fashion to (Bodnar et al., 2021a), let $\beta_\downarrow(e)$ denote the cycles that contain a given edge e and $\beta_\uparrow(c)$ be the edges covered entirely by a given cycle c . Edges interact with cycles in a way analogous to vertices to edges, with the key difference being cycles use a higher order representation. First, consider the following equation depicting how edges send messages to cycles, denoting $H_{c,i}^{e \rightarrow c}$ as the value received at index $i \in V \cap c$ for the first order P-tensor corresponding to a given cycle c .

$$H_{c,i}^{e \rightarrow c} = \sum_{\substack{e \in \beta_\uparrow(c) \\ i \in e}} X_e^e \sum_{e \in \beta_\uparrow(c)} X_e^e \quad (11)$$

This can be equated to the first summation in equation 9, but where we gain an additional linear map by considering the overlapping vertices separately. It is worth noting for completeness that in practice we considered both mean and sum reductions for the summations in 11. We can then get the update to a given cycle by combining the incoming representation with the representation obtained by computing the internal linear maps from X_c^c to itself.

$$Y_{c,i}^c = \mu_4 \left((1 + \varepsilon_1) \left(X_{c,i}^c \sum_{j \in c} X_{c,j}^c \right) + H_{c,i}^{e \rightarrow c} \right) \quad (12)$$

Where μ_4 has two layers and ε_1 is a learnable scalar. We also use $H^{e \rightarrow c}$ to compute messages sent back to edges in a similar way to equation 9.

$$Y_e^{c \rightarrow e} = \mu_5 \left((1 + \varepsilon_2) Y_e^{e \rightarrow v} + (1 + \varepsilon_3) \sum_{c \in \beta_\downarrow(e)} \sum_{i \in e} \mu_6 (H_{c,i}^c, X_{c,i}^c) + \sum_{c \in \beta_\downarrow(e)} \sum_{i \in c} \mu_6 (H_{c,i}^c, X_{c,i}^c) \right) \quad (13)$$

Where μ_5 has two layers, μ_6 has one, and ε_2 and ε_3 are learnable weights. Here we can see a direct analogy between this and equation 9, where μ_5 corresponds to μ_1 and μ_6 corresponds to μ_2 . It is again worth noting that there are many options for the given reductions, but we simply least summations for ease of notation. Since edges receive messages from both vertices and cycles, we compute their new state via a single layer perceptron, μ_7 :

$$Y_e^e = \mu_7 (Y_e^{v \rightarrow e}, Y_e^{c \rightarrow e}) \quad (14)$$