
A Flexible, Equivariant Framework for Subgraph GNNs via Graph Products and Graph Coarsening

Guy Bar-Shalom*

Computer Science
Technion - Israel Institute of Technology
guy.b@campus.technion.ac.il

Yam Eitan*

Electrical & Computer Engineering
Technion - Israel Institute of Technology
yameitan1997@gmail.com

Fabrizio Frasca

Electrical & Computer Engineering
Technion - Israel Institute of Technology
fabrizio.frasca.efe@gmail.com

Haggai Maron

Electrical & Computer Engineering
Technion - Israel Institute of Technology
NVIDIA Research
haggaimaron@gmail.com

Abstract

Subgraph GNNs enhance message-passing GNNs expressivity by representing graphs as sets of subgraphs, demonstrating impressive performance across various tasks. However, their scalability is hindered by the need to process large numbers of subgraphs. While previous approaches attempted to generate smaller subsets of subgraphs through random or learnable sampling, these methods often yielded suboptimal selections or were limited to small subset sizes, ultimately compromising their effectiveness. This paper introduces a new Subgraph GNN framework to address these issues. Our approach diverges from most previous methods by associating subgraphs with node clusters rather than with individual nodes. We show that the resulting collection of subgraphs can be viewed as the product of coarsened and original graphs, unveiling a new connectivity structure on which we perform generalized message passing.

Crucially, controlling the coarsening function enables meaningful selection of any number of subgraphs. In addition, we reveal novel permutation symmetries in the resulting node feature tensor, characterize associated linear equivariant layers, and integrate them into our Subgraph GNN. We also introduce novel node marking strategies and provide a theoretical analysis of their expressive power and other key aspects of our approach. Extensive experiments on multiple graph learning benchmarks demonstrate that our method is significantly more flexible than previous approaches, as it can seamlessly handle any number of subgraphs, while consistently outperforming baseline approaches. Our code is available at <https://github.com/BarSGuy/Efficient-Subgraph-GNNs>.

1 Introduction

Subgraph GNNs [4, 12, 39, 8, 27, 29, 38, 3] have recently emerged as a promising direction in graph neural network research, addressing the expressiveness limitations of Message Passing Neural Networks (MPNNs) [24, 35, 25]. In essence, a Subgraph GNN operates on a graph by transforming it into a collection of subgraphs, generated based on a specific selection policy. Examples of such policies include removing a single node from the original graph or simply marking a node without changing the graph’s original connectivity [26]. The model then processes these subgraphs using an equivariant architecture, aggregates the derived representations, and makes graph- or node-level predictions. The growing popularity of Subgraph GNNs stems not only from their enhanced

*Equal contribution.

expressive capabilities over MPNNs but also from their impressive empirical results, as notably demonstrated on well-known molecular benchmarks [38, 12, 3].

Unfortunately, Subgraph GNNs are hindered by substantial computational costs as they necessitate message-passing operations across all subgraphs within the bag. Typically, the number of subgraphs is the number of nodes in the graph, n — for bounded degree graphs, this results in a time complexity scaling quadratically ($\mathcal{O}(n^2)$), in contrast to the linear complexity of a standard MPNN. This significant computational burden makes Subgraph GNNs impractical for large graphs, hindering their applicability to important tasks and widely used datasets. To overcome this challenge, various studies have explored methodologies that process only a subset of subgraphs from the bag. These methods range from simple random sampling techniques [8, 4, 40, 3] to more advanced strategies that learn to select the most relevant subset of the bag to process [5, 20, 29]. However, while random sampling of subgraphs yields subpar performance, more sophisticated learnable selection strategies also have significant limitations. Primarily, they rely on *training-time* discrete sampling which complicates the optimization process, as evidenced by the high number of epochs required to train them [20, 5, 29]. As a result, these methods often allow only a very small bag size, yielding only modest performance improvements compared to random sampling and standard MPNNs.

Our approach. The goal of this paper is to devise a Subgraph GNN architecture that can flexibly generate and process variable-sized bags, and deliver strong experimental results while sidestepping intricate and lengthy training protocols. Specifically, our approach aims to overcome the common limitation of restricting usage to a very small set of subgraphs.

Our proposed method builds upon and extends an observation made by Bar-Shalom et al. [3], who draw an analogy between using Subgraph GNNs and performing message-passing operations over a larger “product graph”. Specifically, it was shown that when considering the maximally expressive (node-based) Subgraph GNN suggested by [38]², the bag of subgraphs and its update rules can be obtained by transforming a graph through the *graph cartesian product* of the original graph with itself, i.e., $G \square G$, and then processing the resulting graph using a standard MPNN. In our approach, we propose to modify the first term of the product and replace it with a *coarsened* version of the original graph, denoted $\mathcal{T}(G)$, obtained by mapping nodes to *super-nodes* (e.g., by applying graph clustering, see Figure 1(left)), making the resulting product graph $\mathcal{T}(G) \square G$ significantly smaller. This construction is illustrated in Figure 1(right). This process effectively associates each subgraph – a row in Figure 1(right) – with a set of nodes produced by the coarsening function \mathcal{T} . Different choices of \mathcal{T} allow for both flexible bag sizes and a simple, meaningful selection of the subgraphs.




Figure 1: Product graph construction. **Left:** Transforming of the graph into a coarse graph; **Right:** Cartesian product of the coarsened graph with the original graph. The vertical axis corresponds to the subgraph dimension (super-nodes), while the horizontal axis corresponds to the node dimension (nodes).

While performing message passing on $\mathcal{T}(G) \square G$ serves as the core update rule in our architecture, we augment our message passing operations with another set of operations derived from the symmetry structure of the resulting node feature tensor, which we call *symmetry-based updates*. Specifically, our node feature tensor is indexed by pairs (S, v) where S is a super-node and v is an original node. Accordingly, \mathcal{X} is a $T \times n \times d$ tensor, where d is the feature dimension, and T is the number of super-nodes (a constant hyper-parameter). As super-nodes are sets of nodes, \mathcal{X} can also be viewed as a (very) sparse $2^n \times n \times d$ tensor where 2^n is the number of all subsets of the vertex set. Since the symmetric group S_n acts naturally on this representation, we use it to develop symmetry based updates.

Interestingly, we find that this node feature tensor, \mathcal{X} , adheres to a specific set of symmetries, which, to the best of our knowledge, is yet unstudied in the context of machine learning: applying

²The architecture suggested in [38] was shown to be at least as expressive as all previously studied node-based Subgraph GNNs

a permutation $\sigma \in S_n$ to the nodes in S and to v results in an equivalent representation of our node feature tensor. We formally define the symmetries of this object and characterize all the affine equivariant operations in this space. We incorporate these operations into our message-passing by encoding the parameter-sharing schemes [30] as additional edge features. These additional update rules significantly improve experimental results. We note that our symmetry analysis may be useful for processing bags derived from other high-order generation policies [29, 20] by treating tuples of nodes as sets.

Inspired by these symmetries and traditional binary-based [4] and shortest path-based [38] *node-marking* strategies, we propose four natural marking strategies for our framework. Interestingly, unlike the full-bag scenario, they vary in expressiveness, with the shortest path-based technique being the most expressive.

The flexibility and effectiveness of our full framework are illustrated in Figure 2, depicting detailed experimental results on the popular ZINC-12K dataset [31]. Our method demonstrates a significant performance boost over baseline models in the *small bag* setting (for which they are designed), while achieving results that compare favourably to state-of-the-art Subgraph GNNs in the *full bag* setting. Additionally, we can obtain results in-between these two regimes.

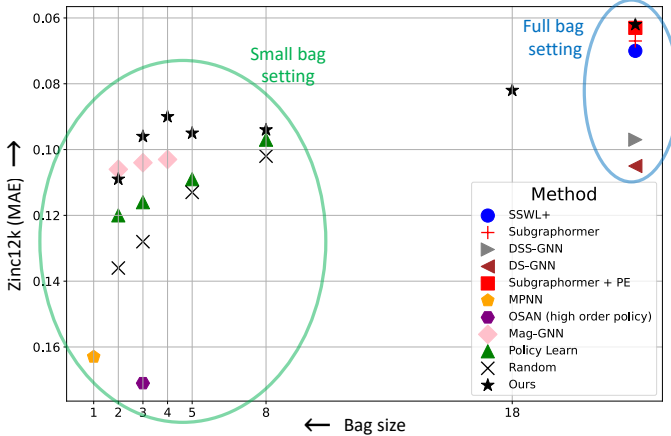



Figure 2: The performance landscape of Subgraph GNNs with varying number of subgraphs: Our method leads in the lower bag-size set, outperforming other approaches in nearly all cases. Additionally, our method matches the performance of state-of-the-art Subgraph GNNs in the full-bag setting. The full mean absolute error (MAE) scores along with standard deviations are available in Table 9 in the appendix.

Contributions. The main contributions of this paper are: (1) the development of a novel, flexible Subgraph GNN framework that enables meaningful construction and processing of bags of subgraphs of any size; (2) a characterization of all affine invariant/equivariant layers defined on our node feature tensors; (3) a theoretical analysis of our framework, including the expressivity benefits of our node-marking strategy; and (4) a comprehensive experimental evaluation demonstrating the advantages of the new approach across both small and large bag sizes, achieving state-of-the-art results, often by a significant margin.

2 Related work

Subgraph GNNs. Subgraph GNNs [39, 8, 27, 4, 40, 26, 12, 29, 17, 38, 3] represent a graph as a collection of subgraphs, obtained by a predefined generation policy. For example, each subgraph can be generated by marking exactly one node in the original graph (see inset ³) – an approach commonly referred to as *node marking* [26]; this marked node is considered the root node in its subgraph. Several recent papers focused on scaling these methods to larger graphs, starting with basic random selection of subgraphs from the bag, and extending beyond with more sophisticated techniques that aim to learn how to select subgraphs. To elaborate, [5] introduced *Policy-Learn* (PL), an approach based on two models, where the first model predicts a distribution over the nodes of the original graph, and the second model processes bags of subgraphs sampled from this distribution. *MAG-GNN* [20] employs a similar approach utilizing Reinforcement Learning. Similarly to our approach, this method permits high-order policies by associating subgraphs with tuples rather than individual nodes, allowing for the marking of several nodes within a subgraph.



³The Figure was taken with permission from [3]

However, as mentioned before, these approaches involve discrete sampling while training, making them very hard to train (1000-4000 epochs vs. ~ 400 epochs of state-of-the-art methods [3, 38] on the ZINC-12K dataset), and limiting their usage to very small bags. Finally, we mention another high-order method, *OSAN*, introduced by [29], which learns a distribution over tuples that represent subgraphs with multiple node markings. In contrast to these previous approaches, we suggest a simpler and more effective way to select subgraphs and also show how to leverage the resulting symmetry structure to augment our message-passing operations.

Symmetries in graph learning. Many previous works have analyzed and utilized the symmetry structure that arises from graph learning setups [22, 23, 18, 2]. Specifically relevant to our paper is the work of [22] that characterized basic equivariant linear layers for graphs, the work of [1] that characterizes equivariant maps for many other types of incidence tensors that arise in graph learning, and the works [4, 12] that leveraged group symmetries for designing Subgraph GNNs in a principled way.

3 Preliminaries

Notation. Let \mathcal{G} be a family of undirected graphs, and consider a graph $G = (V, E)$ within this family. The adjacency matrix $A \in \mathbb{R}^{n \times n}$ defines the connectivity of the graph⁴, while the feature matrix $X \in \mathbb{R}^{n \times d}$ represents the node features. Here, V and E represent the sets of nodes and edges, respectively, with $|V| = n$ indicating the number of nodes. We use the notation $v_1 \sim_A v_2$ to denote that v_1 and v_2 are neighboring nodes according to the adjacency A . Additionally, we define $[n] := \{1, 2, \dots, n\}$, and $\mathcal{P}([n])$ as the power set of $[n]$.

Subgraph GNNs as graph products. In a recent work, [3] demonstrated that various types of update rules used by current Subgraph GNNs can be simulated by employing the *Cartesian graph product* between the original graph and another graph, and running standard message passing over that newly constructed product graph. Formally, the cartesian product of two graphs G_1 (n_1 nodes) and G_2 (n_2 nodes), denoted by $G_1 \square G_2$, forms a graph with vertex set $V(G_1) \times V(G_2)$. Two vertices (u_1, u_2) and (v_1, v_2) are adjacent if either $u_1 = v_1$ and u_2 is adjacent to v_2 in G_2 , or $u_2 = v_2$ and u_1 is adjacent to v_1 in G_1 . We denote by $\mathcal{A} \in \mathbb{R}^{n_1 \cdot n_2 \times n_1 \cdot n_2}$ and $\mathcal{X} \in \mathbb{R}^{n_1 \cdot n_2 \times d}$ the adjacency and node feature matrices of the product graph; in general, we use calligraphic letters to denote the adjacency and feature matrices of product graphs, while capital English letters are used for those of the original graphs. In particular, for the graph cartesian product, $G_1 \square G_2$, the following holds:

$$\mathcal{A}_{G_1 \square G_2} = A_1 \otimes I + I \otimes A_2. \quad (1)$$

For a detailed definition of the cartesian product of graphs, please refer to Definition A.1. As a concrete example for the analogy between Subgraph GNNs and the Cartesian product of graphs, we refer to a result by [3], which states that the maximally expressive node-based Subgraph GNN architecture GNN-SSWL+ [38], can be simulated by an MPNN on the Cartesian product of the original graph with itself, denoted as $G \square G$. As we shall see, our framework utilizes a cartesian product of the original graph and a coarsened version of it, as illustrated in Figure 1 (right).

Equivariance. A function $L : U \rightarrow W$ is called equivariant if it commutes with the group action. More formally, given a group element, $g \in \mathbb{G}$, the function L should satisfy $L(g \cdot v) = g \cdot L(v)$ for all $v \in U$ and $g \in \mathbb{G}$. L is said to be invariant if $L(g \cdot v) = L(v)$.

4 Coarsening-based Subgraph GNN

Overview. This section introduces the *Coarsening-based Subgraph GNN* (CS-GNN) framework. The main idea is to select and process subgraphs in a principled and flexible manner through the following approach: (1) coarsen the original graph via a coarsening function, \mathcal{T} – see Figure 1(left); (2) Obtain the product graph – Figure 1(right) defined by the combination of two adjacencies, $\mathcal{A}_{\mathcal{T}(G)}$ (red edges), \mathcal{A}_G (grey edges), which arise from the graph Cartesian product operation (details follow); (3) leveraging the symmetry of this product graph to develop *symmetry-based* updates, described by $\mathcal{A}_{\text{Equiv}}$ (this part is not visualized in Figure 1). The general update of our suggested layer takes the following form⁴,

⁴Edge features are also allowed but are omitted here for simplicity

$$\mathcal{X}^{t+1}(S, v) = f^t \left(\mathcal{X}(S, v)^t, \right. \tag{2}$$

$$\left. \underbrace{\{\{\mathcal{X}(S', v')^t\}\}_{(S', v') \sim_{\mathcal{A}_G}(S, v)}}_{\text{Original connectivity (horizontal)}}, \underbrace{\{\{\mathcal{X}(S', v')^t\}\}_{(S', v') \sim_{\mathcal{A}_{\mathcal{T}(G)}}(S, v)}}_{\text{Induced connectivity (vertical)}}, \underbrace{\{\{\mathcal{X}(S', v')^t\}\}_{(S', v') \sim_{\mathcal{A}_{\text{Equiv}}}(S, v)}}_{\text{Symmetry-based updates}} \right),$$

where the superscript t indicates the layer index. In what follows, we further elaborate on these three steps (in Sections 4.1 to 4.2).

We note that each connectivity in Equation (2) is processed using a distinct MPNN, and after stacking of those layers, we apply a pooling layer⁵ to obtain a graph representation; that is, $\rho(\mathcal{X}^T) = \text{MLP}^T \left(\sum_S \left(\sum_{v=1}^n \mathcal{X}^T(S, v) \right) \right)$; T denotes the final layer.

For more specific implementation details, we refer to Appendix F.

4.1 Construction of the coarse product graph

As mentioned before, a maximally expressive node-based Subgraph GNN can be realized via the Cartesian product of the original graph with itself $G \square G$. In this work, we extend this concept by allowing the left operand in the product to be the coarsened version of G , denoted as $\mathcal{T}(G)$, as defined next. This idea is illustrated in Figure 1.

Graph coarsening. Consider a graph $G = (V, E)$ with n nodes and an adjacency matrix A . Graph coarsening is defined by the function $\mathcal{T} : \mathcal{G} \rightarrow \mathcal{G}$, which maps G to a new graph $\mathcal{T}(G) = (V^T, E^T)$ with an adjacency matrix $A^T \in \mathbb{R}^{2^n \times 2^n}$ and a feature matrix $X^T \in \mathbb{R}^{2^n \times d}$. Here, V^T , the vertex set of the new graph represents super-nodes – defined as subsets of $[n]$. Additionally, we require that nodes in V^T induce a partition over the nodes of the original graph⁶. The connectivity E^T is extremely sparse and induced from the original graph’s connectivity via the following rule:

$$A^T(S_1, S_2) = \begin{cases} 1 & \text{if } \exists v \in S_1, \exists v \in S_2 \text{ s.t. } A(v, u) = 1, \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

To clarify, in our running example (Figure 1), it holds that $A^T(\{a, b, c, d\}, \{e\}) = 1$, while $A^T(\{e\}, \{f\}) = 0$. For a more formal definition, refer to Definition A.3.

More specifically, our implementation of the graph coarsening function \mathcal{T} employs spectral clustering⁷ [33] to partition the graph into T clusters, which in our framework controls the size of the bag. This results in a coarsened graph with fewer nodes and edges than G . We highlight and stress that the space complexity of this sparse graph, $\mathcal{T}(G)$, is upper bounded by that of the original graph G (we do not store 2^n nodes).

Defining the (coarse) product graph $\mathcal{T}(G) \square G$. We define the connectivity of the product graph, see Figure 1(right), by applying the cartesian product between the coarsened graph, $\mathcal{T}(G)$, and the original graph, G . The product graph is denoted by $\mathcal{T}(G) \square G$, and is represented by the matrices $\mathcal{A}_{\mathcal{T}(G) \square G} \in \mathbb{R}^{(2^n \times n) \times (2^n \times n)}$ and $\mathcal{X} \in \mathbb{R}^{2^n \times n \times d \times 8}$, where by recalling Equation (1), we obtain,

$$\mathcal{A}_{\mathcal{T}(G) \square G} = \overbrace{A^T}^{\triangleq \mathcal{A}_{\mathcal{T}(G)}} \otimes I + I \otimes \overbrace{A}^{\triangleq \mathcal{A}_G}. \tag{4}$$

The connectivity in this product graph induces the horizontal (\mathcal{A}_G) and vertical updates ($\mathcal{A}_{\mathcal{T}(G)}$) in Equation (2), visualized in Figure 1(right) via grey and red edges, respectively.

⁵For some of the theoretical analysis, this pooling operation is expressed as: $\rho(\mathcal{X}^T) = \text{MLP}^T \left(\sum_S \left(\text{MLP}^T \left(\sum_{v=1}^n \mathcal{X}^T(S, v) \right) \right) \right)$

⁶Our method also supports the case of which it is not a partition.

⁷Other graph coarsening or clustering algorithms can be readily used as well.

⁸We note that while the node matrix of the product graph, \mathcal{X} , can be initialized in various ways, e.g., deep sets-based architecture [37], in our implementation we simply use the original node features, i.e., $\mathcal{X}(S, v) = X(v)$, given that X is the node feature matrix of the original graph.

4.2 Symmetry-based updates

In the previous subsection, we used a combination of a coarsening function and the graph Cartesian product to derive the two induced connectivities $\mathcal{A}_G, \mathcal{A}_{\mathcal{T}(G)}$ of our product graph. We use these connectivities to perform message-passing on our product graph (see Equation (2)).

Inspired by recent literature on Subgraph GNNs [12, 3, 38], which incorporates and analyzes additional non-local updates arising from various symmetries (e.g., updating a node’s representation via all nodes in its subgraphs), this section aims to identify potential new updates that can be utilized over our product graph. To that end, we study the symmetry structure of the node feature tensor in our product graph, $\mathcal{X}(S, v)$. The new updates described below will result in the third term in Equation (2), dubbed *Symmetry-based updates* ($\mathcal{A}_{\text{Equiv}}$). For better clarity in this derivation, we change the notation from nodes (v) to indices (i).

4.2.1 Symmetries of our product graph

Since the order of nodes in the original graph G is arbitrary, each layer in our architecture must exhibit equivariance to any induced changes in the product graph. This requires maintaining equivariance to permutations of nodes in both the original graph and its transformation $\mathcal{T}(G)$. As a result, recalling that $\mathcal{A} \in \mathbb{R}^{(2^n \times n) \times (2^n \times n)}$ and $\mathcal{X} \in \mathbb{R}^{2^n \times n \times d}$ represent the adjacency and feature matrices of the product graph, the symmetries of the product graph are defined by an action of the symmetric group S_n . Formally, a permutation $\sigma \in S_n$ acts on the adjacency and feature matrices by:

$$(\sigma \cdot \mathcal{A})(S_1, i_1, S_2, i_2) = \mathcal{A}(\sigma^{-1}(S_1), \sigma^{-1}(i_1), \sigma^{-1}(S_2), \sigma^{-1}(i_2)), \quad (5)$$

$$(\sigma \cdot \mathcal{X})(S, i) = \mathcal{X}(\sigma^{-1}(S), \sigma^{-1}(i)), \quad (6)$$

where we define the action of $\sigma \in S_n$ on a set $S = \{i_1, i_2, \dots, i_k\}$ of size k as: $\sigma \cdot S := \{\sigma^{-1}(i_1), \sigma^{-1}(i_2), \dots, \sigma^{-1}(i_k)\} := \sigma^{-1}(S)$.

4.2.2 Derivation of linear equivariant layers for the node feature tensor

We now characterize the linear equivariant layers with respect to the symmetry defined above, focusing on Equation (6). We adopt a similar notation to [22], and assume for simplicity that the number of feature channels is $d = 1$ (extension to multiple features is straightforward [22]). In addition, our analysis considers the case where $V^{\mathcal{T}}$ encompasses all potential super-nodes formed by subsets of $[n]$ (i.e we use the sparse coarsened adjacency⁹).

Our main tool is the characterization of linear equivariant layers for permutation symmetries as parameter-sharing schemes [34, 30, 22]. In a nutshell, this characterization states that the parameter vectors of the biases, invariant layers, and equivariant layers can be expressed as a learned weighted sum of basis tensors, where the basis tensors are indicators of the orbits induced by the group action on the respective index spaces. We focus here on presenting the final results and summarize them in Proposition 4.1 at the end of this subsection. Detailed discussion and derivations are available in Appendix E.

Equivariant bias and invariant layers. The bias vectors of the linear layers in our space are in $\mathbb{R}^{2^n \times n}$. As shown in Figure 3(right), the set of orbits induced by the action of S_n satisfies:

$$(\mathcal{P}([n]) \times [n])/S_n := \{\gamma^{k*} : k = 1, \dots, n; * \in \{+, -\}\}. \quad (7)$$

Here, γ^{k+} corresponds to all pairs $(S, i) \in \mathcal{P}([n]) \times [n]$ with $|S| = k$ and $i \notin S$, and γ^{k-} to all pairs with $|S| = k$ and $i \in S$.

As stated in [34, 30, 22], the tensor set $\{\mathbf{B}_{S,i}^\gamma\}_{\gamma \in (\mathcal{P}([n]) \times [n])/S_n}$ where:

$$\mathbf{B}_{S,i}^\gamma = \begin{cases} 1, & \text{if } (S, i) \in \gamma; \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

are a basis of the space of bias vectors of the invariant linear layers induced by the action of S_n .

⁹This is because the action of S_n is well defined over the index set $\mathcal{P}(V[[n]]) \times [n]$ but not over $V \times V^{\mathcal{T}}$

colors. To clarify, the nodes (S, v) that satisfy $v \in S$ “send messages” (i.e., broadcast their representation) to all the nodes (S', v') such that $v = v'$. A more formal discussion regarding our implementation of those symmetry based updates is given in Appendix F.4.

Maintaining sparsity. While the updates above are defined over the sparse representation of the coarse product graph, in practice we use its dense representation, treating it as a graph over the set of nodes $V \times V^T$, which requires space complexity $\mathcal{O}(T \cdot |V|)$. The update rules above are adapted to this representation simply by masking all nodes (S, v) in the sparse representation such that $S \notin V^T$. We note the models using the resulting update rule remain invariant to the action of S_n . See discussion in [1].

4.3 Marking Strategies and Theoretical Analysis

One of the key components of subgraph architectures is their marking strategy. Two widely used approaches in node-based subgraph architectures are binary-based node marking [4] and distance-based marking [38], which were proven to be equally expressive in the full-bag setup [38]. Empirically, distance-based marking has been demonstrated to outperform other strategies across several standard benchmarks. In this section, our aim is to develop and theoretically justify an appropriate marking strategy, specifically tailored to the structure of our product graph. We present and discuss here our main results, and refer to Appendix C for a more formal discussion.

Building on existing marking strategies and considering the unique structure of our product graph, we propose two natural extensions to both the binary node marking [4] and distance-based marking strategies [38]. Extending binary node marking, we first suggest **Simple Marking** (π_S), where an element (S, v) is assigned a binary feature that indicates whether node v belongs to subgraph S ($v \in S$). The second extension, **Node + Size Marking** (π_{SS}), builds on the *simple marking* by assigning an additional feature that encodes the size of the super-node S .

For distance-based strategies, we propose **Minimum Distance** (π_{MD}), where each element (S, v) is assigned the smallest (minimal) shortest path distance (SPD) from node v to any node $u \in S$. Finally, **Learned Distance Function** (π_{LD}) extends this further by assigning to each element (S, v) the output of a permutation-invariant learned function, which takes the set of SPDs between node v and the nodes in S as input.

Surprisingly, unlike the node-based full-bag case, we find that these marking strategies are not all equally expressive. We conveniently gather the first three strategies as $\Pi = \{\pi_S, \pi_{SS}, \pi_{MD}\}$ and summarize the relation between all variants as follows:

Proposition 4.2 (Informal – Expressivity of marking strategies.). *(i) Strategies in Π are all equally expressive, independently of the transformation function \mathcal{T} . (ii) The strategy π_{LD} is at least as expressive as strategies in Π . Additionally, there exists transformation functions s.t. it is strictly more expressive than all of them.*

The above is formally stated in Propositions C.1 and C.2, and more thoroughly discussed in Appendix C. In light of the above proposition, we instantiate the learned distance function π_{LD} strategy when implementing our model, as follows,

$$\mathcal{X}_{S,v} \leftarrow \sum_{u \in S} z_{d_G(v,u)} \quad (10)$$

where $d_G(v, u)$ denotes the *shortest path distance* between nodes v and u in the original G ¹⁰.

Coarsening Function and Expressivity. We investigate whether our CS-GNN framework offers more expressiveness compared to directly integrating information between the coarsened graph and the original graph.

The two propositions below illustrate that a simple, straight forward integration of the coarsen graph with the original graph (this integration is referred to as the *sum graph* – formally defined in Definition D.2), and further processing it via standard message-passing, results in a less expressive architecture. Furthermore, when certain coarsening functions are employed within the CS-GNN framework, our resulting architecture becomes strictly more expressive than conventional node-based

¹⁰To facilitate this, we maintain a lookup table where each index corresponds to a shortest path distance, assigning a learnable embedding, $z_{d_G(v,u)} \in \mathbb{R}^d$, to each node (S, v) .

subgraph GNNs. These results suggest that the interplay between the coarsening function and the subgraph layers we have developed enhances the model’s overall performance. We summarize this informally below and provide a more formal discussion in Appendix D.

Proposition 4.3 (Informal – CS-GNN goes beyond coarsening). *For any transformation function \mathcal{T} , CS-GNN can implement message-passing on the sum graph, hence being at least as expressive. Also, there exist transformations \mathcal{T} ’s s.t. CS-GNN is strictly more expressive than that.*

Proposition 4.4 (Informal – CS-GNN vs node based subgraphs). *There exist transformations \mathcal{T} ’s s.t. our CS-GNN model using \mathcal{T} as its coarsening function is strictly more expressive than GNN-SSWL+.*

5 Experiments

We experimented extensively over seven different datasets to answer the following questions: (Q1) Can CS-GNN outperform efficient Subgraph GNNs operating on small bags? (Q2) Does the additional symmetry-based updates boost performance? (Q3) Does CS-GNN offer a good solution in settings where full-bag Subgraph GNNs cannot be applied? (Q4) Does CS-GNN in the full-bag setting validate its theory and match state-of-the-art full-bag Subgraph GNNs?

In the following sections, we present our main results and refer to Appendix F for additional experiments and details.

Baselines. For each task, we include several baselines. The RANDOM baseline corresponds to random subgraph selection. We report the best performing random baseline from all prior work [5, 20, 29, 3]. The other two (non-random) baselines are: (1) LEARNED [5, 20, 29], which represents methods that learn the specific subgraphs to be used; and (2) FULL [38, 3], which corresponds to full-bag Subgraph GNNs.

ZINC. We experimented with both the ZINC-12k and ZINC-FULL datasets [31, 14, 10], adhering to a 500k parameter budget as prescribed. As shown in Table 1, CS-GNN outperforms all efficient baselines by a significant margin, with at least a +0.008 MAE improvement for bag sizes $T \in \{3, 4, 5\}$. Additionally, in the full-bag setting, our method recovers state-of-the-art results. The results for ZINC-FULL are available in Table 8 in the Appendix.

OGB. We tested our framework on several datasets from the OGB benchmark collection [16]. Table 4 shows the performance of our method compared to both efficient and full-bag Subgraph GNNs. Our CS-GNN outperforms all baselines across all datasets for bag sizes $T \in \{2, 5\}$, except for the MOLHIV dataset with $T = 2$, where PL achieves the best results and our method ranks second. In the full-bag setting, CS-GNN is slightly outperformed by the top-performing Subgraph GNNs but still offers comparable results.

Peptides. We experimented on the PEPTIDES-FUNC and PEPTIDES-STRUCT datasets [9] – which full-bag Subgraph GNNs already struggle to process – evaluating CS-GNN’s ability to scale to larger graphs. The results are summarized in Table 2. CS-GNN outperforms all MPNN variants, even when incorporating structural encodings such as GATEDGCN+RWSE. Additionally, our method surpasses the random¹¹ baseline on both datasets.

Table 1: Results on ZINC-12k dataset. Top two results are reported as **First** and **Second**.

Method	Bag size	ZINC (MAE ↓)
GIN [35]	$T = 1$	0.163 ± 0.004
OSAN [29]	$T = 2$	0.177 ± 0.016
Random [20]	$T = 2$	0.131 ± 0.005
PL [5]	$T = 2$	0.120 ± 0.003
Mag-GNN [20]	$T = 2$	0.106 ± 0.014
Ours	$T = 2$	0.109 ± 0.005
Random [20]	$T = 3$	0.124 ± N/A
Mag-GNN [20]	$T = 3$	0.104 ± N/A
Ours	$T = 3$	0.096 ± 0.005
Random [20]	$T = 4$	0.125 ± N/A
Mag-GNN [20]	$T = 4$	0.101 ± N/A
Ours	$T = 4$	0.090 ± 0.003
Random [5]	$T = 5$	0.113 ± 0.006
PL [5]	$T = 5$	0.109 ± 0.005
Ours	$T = 5$	0.095 ± 0.003
GNN-SSWL+ [38]	Full	0.070 ± 0.005
Subgraphormer [3]	Full	0.067 ± 0.007
Subgraphormer+PE [3]	Full	0.063 ± 0.001
Ours	Full	0.062 ± 0.0007

Table 2: Results on PEPTIDES dataset.

Model ↓ / Dataset →	PEPTIDES-FUNC (AP ↑)	PEPTIDES-STRUCT (MAE ↓)
GCN [19]	0.5930±0.0023	0.3496±0.0013
GIN [35]	0.5498±0.0079	0.3547±0.0045
GatedGCN [7]	0.5864±0.0077	0.3420±0.0013
GatedGCN+RWSE [9]	0.6069±0.0035	0.3357±0.0006
Random [3]	0.5924±0.005	0.2594±0.0021
Ours	0.6156 ±0.0080	0.2539 ±0.0015

¹¹For the PEPTIDES datasets, we benchmarked our model against the random variant of Subgraphormer + PE, which similarly incorporates information from the Laplacian eigenvectors. To ensure a fair comparison,

Ablation study – symmetry-based updates. We assessed the impact of the symmetry-based update on the performance of CS-GNN. Specifically, we ask, *do the symmetry-based updates significantly contribute to the performance of CS-GNN?* To evaluate this, we conducted several experiments using the ZINC-12K dataset across various bag sizes, $T \in \{2, 3, 4, 5\}$, comparing CS-GNN with and without the symmetry-based update. The results are summarized in Table 3. It is clear that the symmetry-based updates play a key role in the performance of CS-GNN. For a bag size of $T = 2$, the inclusion of the symmetry-based update improves the MAE by a significant 0.034. For other bag sizes, the improvements range from 0.005 to 0.016, clearly demonstrating the benefits of including the symmetry-based updates.

Table 3: Ablation study.

Bag size	w/	w/o
T=2	0.109 \pm 0.005	0.143 \pm 0.003
T=3	0.096 \pm 0.005	0.101 \pm 0.006
T=4	0.090 \pm 0.003	0.106 \pm 0.001
T=5	0.095 \pm 0.003	0.104 \pm 0.005

Discussion. In what follows, we address research questions **Q1** to **Q4**. (A1) Tables 1, 2 and 4 clearly demonstrate that we outperform efficient Subgraph GNNs (which operate on a small bag) in 10 out of 12 dataset and bag size combinations. (A2) Our ablation study on the ZINC-12K dataset, as shown in Table 3, clearly demonstrates the benefits of the symmetry-based updates across all the considered bag sizes. (A3) Table 2 demonstrates that CS-GNN provides an effective solution when the full-bag setting cannot be applied, outperforming all baselines. (A4) On the ZINC-12K dataset (see Table 1), CS-GNN achieves state-of-the-art results compared to Subgraph GNNs. On the OGBG datasets (see Table 4), our performance is comparable to these top-performing Subgraph GNNs.

Table 4: Results on OGB datasets. The top two results are reported as **First** and **Second**.

Model \downarrow / Dataset \rightarrow	Bag size	MOLHIV (ROC-AUC \uparrow)	MOLBACE (ROC-AUC \uparrow)	MOLESOL (RMSE \downarrow)
GIN [35]	$T = 1$	75.58 \pm 1.40	72.97 \pm 4.00	1.173 \pm 0.057
Random [5]	$T = 2$	77.55 \pm 1.24	75.36 \pm 4.28	0.951 \pm 0.039
PL [5]	$T = 2$	79.13 \pm 0.60	78.40 \pm 2.85	0.877 \pm 0.029
Mag-GNN [20]	$T = 2$	77.12 \pm 1.13	-	-
Ours	$T = 2$	77.72 \pm 0.76	80.58 \pm 1.04	0.850 \pm 0.024
OSAN [29]	$T = 3$	-	-	0.959 \pm 0.184
OSAN [29]	$T = 5$	-	76.30 \pm 3.00	-
PL [5]	$T = 5$	78.49 \pm 1.01	78.39 \pm 2.28	0.883 \pm 0.032
Random [5]	$T = 5$	77.30 \pm 2.56	78.14 \pm 2.36	0.900 \pm 0.032
Ours	$T = 5$	79.09 \pm 0.90	79.64 \pm 1.43	0.863 \pm 0.029
GNN-SSWL+ [38]	Full	79.58 \pm 0.35	82.70 \pm 1.80	0.837 \pm 0.019
Subgraphormer	Full	80.38 \pm 1.92	81.62 \pm 3.55	0.832 \pm 0.043
Subgraphormer + PE	Full	79.48 \pm 1.28	84.35 \pm 0.65	0.826 \pm 0.010
Ours	Full	79.44 \pm 0.87	80.71 \pm 1.76	0.814 \pm 0.021

6 Conclusions

In this work, we employed graph coarsenings and leveraged the insightful connection between Subgraph GNNs and the graph Cartesian product to devise CS-GNN, a novel and flexible Subgraph GNN that can effectively generate and process any desired bag size. Several directions for future research remain open. Firstly, we experimented with spectral clustering based coarsening, but other strategies are possible and are interesting to explore. Secondly, in our symmetry-based updates, we have only considered a portion of the whole equivariant basis we derived: evaluating the impact of other basis elements deserve further attention, both theoretically and in practice. Finally, whether Higher-Order Subgraph GNNs can benefit from our developed parameter-sharing scheme remains an intriguing open question.

Limitations. Our method operates over a product graph. Although we provide control over the size of this product graph, achieving better performance requires a larger bag size. This can become a complexity bottleneck, particularly when the original graph is large.

Acknowledgements

The authors are grateful to Beatrice Bevilacqua, for helpful discussions, and constructive conversations about the experiments. HM is the Robert J. Shillman Fellow and is supported by the Israel Science Foundation through a personal grant (ISF 264/23) and an equipment grant (ISF 532/23). FF is funded by the Andrew and Erna Finci Viterbi Post-Doctoral Fellowship; FF partially performed this work while visiting the Machine Learning Research Unit at TU Wien led by Prof. Thomas Gartner.

we used a single vote and the same exact bag size of 35 subgraphs. Additionally, since Subgraphormer + PE employs GAT [32] as the underlying MPNN, we also utilized GAT for this specific experiment to maintain consistency and fairness.

References

- [1] Marjan Albooyeh, Daniele Bertolini, and Siamak Ravanbakhsh. Incidence networks for geometric deep learning. *arXiv preprint arXiv:1905.11460*, 2019.
- [2] Weiss Azizian and Marc Lelarge. Expressive power of invariant and equivariant graph neural networks. *arXiv preprint arXiv:2006.15646*, 2020.
- [3] Guy Bar-Shalom, Beatrice Bevilacqua, and Haggai Maron. Subgraphormer: Unifying subgraph GNNs and graph transformers via graph products. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=6djdWVTUEq>.
- [4] Beatrice Bevilacqua, Fabrizio Frasca, Derek Lim, Balasubramaniam Srinivasan, Chen Cai, Gopinath Balamurugan, Michael M Bronstein, and Haggai Maron. Equivariant subgraph aggregation networks. *International Conference on Learning Representations*, 2022.
- [5] Beatrice Bevilacqua, Moshe Eliasof, Eli Meir, Bruno Ribeiro, and Haggai Maron. Efficient subgraph gnn by learning effective selection policies. *International Conference on Learning Representations*, 2024.
- [6] Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- [7] Xavier Bresson and Thomas Laurent. Residual gated graph convnets. *arXiv preprint arXiv:1711.07553*, 2017.
- [8] Leonardo Cotta, Christopher Morris, and Bruno Ribeiro. Reconstruction for powerful graph representations. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [9] Vijay Prakash Dwivedi, Ladislav Rampásek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. Long range graph benchmark. *Advances in Neural Information Processing Systems*, 35:22326–22340, 2022.
- [10] Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023.
- [11] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [12] Fabrizio Frasca, Beatrice Bevilacqua, Michael Bronstein, and Haggai Maron. Understanding and extending subgraph gnn by rethinking their symmetries. *Advances in Neural Information Processing Systems*, 35:31376–31390, 2022.
- [13] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.
- [14] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.
- [15] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. *arXiv preprint arXiv:1905.12265*, 2019.
- [16] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.
- [17] Yinan Huang, Xingang Peng, Jianzhu Ma, and Muhan Zhang. Boosting the cycle counting power of graph neural networks with i^2 -gnns. In *The Eleventh International Conference on Learning Representations*, 2022.

- [18] Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [19] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2016.
- [20] Lecheng Kong, Jiarui Feng, Hao Liu, Dacheng Tao, Yixin Chen, and Muhan Zhang. Mag-gnn: Reinforcement learning boosted graph neural network. *Advances in Neural Information Processing Systems*, 36, 2024.
- [21] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, 2021.
- [22] Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. *arXiv preprint arXiv:1812.09902*, 2018.
- [23] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *Advances in neural information processing systems*, 32, 2019.
- [24] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4602–4609, 2019.
- [25] Christopher Morris, Yaron Lipman, Haggai Maron, Bastian Rieck, Nils M Kriege, Martin Grohe, Matthias Fey, and Karsten Borgwardt. Weisfeiler and leman go machine learning: The story so far. *arXiv preprint arXiv:2112.09992*, 2021.
- [26] Pál András Papp and Roger Wattenhofer. A theoretical comparison of graph neural network extensions. In *International Conference on Machine Learning*, pages 17323–17345. PMLR, 2022.
- [27] Pál András Papp, Karolis Martinkus, Lukas Faber, and Roger Wattenhofer. Dropgnn: Random dropouts increase the expressiveness of graph neural networks. *Advances in Neural Information Processing Systems*, 34:21997–22009, 2021.
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [29] Chendi Qian, Gaurav Rattan, Floris Geerts, Mathias Niepert, and Christopher Morris. Ordered subgraph aggregation networks. *Advances in Neural Information Processing Systems*, 35: 21030–21045, 2022.
- [30] Siamak Ravanbakhsh, Jeff Schneider, and Barnabas Poczos. Equivariance through parameter-sharing. In *International conference on machine learning*, pages 2892–2901. PMLR, 2017.
- [31] Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.
- [32] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2017.
- [33] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17:395–416, 2007.
- [34] Jeffrey Wood and John Shawe-Taylor. Representation theory and invariant neural networks. *Discrete applied mathematics*, 69(1-2):33–60, 1996.
- [35] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *International Conference on Learning Representations*, 2018.

- [36] Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small relu networks are powerful memorizers: a tight analysis of memorization capacity. *Advances in Neural Information Processing Systems*, 32, 2019.
- [37] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.
- [38] Bohang Zhang, Guhao Feng, Yiheng Du, Di He, and Liwei Wang. A complete expressiveness hierarchy for subgraph gnns via subgraph weisfeiler-lehman tests. *International Conference on Machine Learning*, 2023.
- [39] Muhan Zhang and Pan Li. Nested graph neural networks. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [40] Lingxiao Zhao, Wei Jin, Leman Akoglu, and Neil Shah. From stars to subgraphs: Uplifting any GNN with local structure awareness. In *International Conference on Learning Representations*, 2022.

Appendix

The appendix is organized as follows:

- In Appendix A, we provide some basic definitions that will be used in later sections of the paper.
- In Appendix B we discuss some theoretical aspects of our model implementation, and its relation to Equation (2).
- In Appendix C we define four natural general node marking policies and analyze their theoretical effects on our model, as well as their relation to some node-based node marking policies. Finally, we provide a principled derivation of one of these policies using the natural symmetry of our base object.
- In Appendix D.1 we compare our model to node-based subgraph GNNs, which are the most widely used variant of subgraph GNNs. Additionally, we demonstrate that different choices of coarsening functions can recover various existing subgraph GNN designs.
- In Appendix D.2 we demonstrate how our model can leverage the information provided by the coarsening function in an effective way, comparing its expressivity to a natural baseline which also leverages the coarsening function. We show that for all coarsening functions, we are at least as expressive as the baseline and that for some coarsening functions, our model is strictly more expressive.
- In Appendix E we delve deeper into the characterization of all linear maps $L : \mathbb{R}^{\mathcal{P}([n]) \times [n]} \rightarrow \mathbb{R}^{\mathcal{P}([n]) \times [n]}$ that are equivariant to the action of the symmetric group.
- In Appendix F we provide experimental details to reproduce the results in Section 5, as well as a comprehensive set of ablation studies.
- In Appendix G we provide detailed proofs to all propositions in this paper.

A Basic Definitions

We devote this section to formally defining the key concepts of this paper, as well as introducing new useful notation. We start by defining the two principle components of our pipeline, the cartesian product graph and the coarsening function:

Definition A.1 (Cartesian Product Graph). *Given two graphs G_1 and G_2 , their Cartesian product $G_1 \square G_2$ is defined as:*

- The vertex set $V(G_1 \square G_2) = V(G_1) \times V(G_2)$.
- Vertices (u_1, u_2) and (v_1, v_2) in $G_1 \square G_2$ are adjacent if:
 - $u_1 = v_1$ and u_2 is adjacent to v_2 in G_2 , or
 - $u_2 = v_2$ and u_1 is adjacent to v_1 in G_1 .

Definition A.2 (Coarsening Function). *A Coarsening function $\mathcal{T}(\cdot)$ is defined as a function that, given a graph $G = (V, E)$ with vertex set $V = [n]$ and adjacency matrix $A \in \mathbb{R}^{n \times n}$, takes A as input and returns a set of "super-nodes" $\mathcal{T}(A) \subseteq \mathcal{P}([n])$. The function $\mathcal{T}(\cdot)$ is considered equivariant if, for any permutation $\sigma \in S_n$, the following condition holds:*

$$\mathcal{T}(\sigma \cdot A) = \sigma \cdot \mathcal{T}(A). \quad (11)$$

Here, $\sigma \cdot A$, and $\sigma \cdot \mathcal{T}(A)$ represent the group action of the symmetric group S_n on $\mathbb{R}^{n \times n}$, and $\mathcal{P}([n])$ respectively.

A coarsening function allows us to naturally define a graph structure on the "super-nodes" obtained from a given graph in the following way:

Definition A.3 (Coarsened Graph). *Given a coarsening function $\mathcal{T}(\cdot)$ and a graph $G = (V, E)$ with vertex set $V = [n]$, adjacency matrix $A \in \mathbb{R}^{n \times n}$, we abuse notation and define the coarsened graph $\mathcal{T}(G) = (V^\mathcal{T}, E^\mathcal{T})$ as follows:*

- $V^\mathcal{T} = \mathcal{T}(A)$

$$\bullet E^{\mathcal{T}} = \{\{S, S'\} \mid S, S' \in \mathcal{T}(A), \exists i \in S, i' \in S' \text{ s.t. } A_{i,i'} = 1\}.$$

The adjacency matrix of the coarsened graph can be expressed in two ways. The dense representation $A_{dense}^{\mathcal{T}} \in \mathbb{R}^{|V^{\mathcal{T}}| \times |V^{\mathcal{T}}|}$ is defined by:

$$A_{dense}^{\mathcal{T}}(S, S') = \begin{cases} 1 & \{S, S'\} \in E^{\mathcal{T}} \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

The sparse representation $A_{sparse}^{\mathcal{T}} \in \mathbb{R}^{\mathcal{P}([n]) \times \mathcal{P}([n])}$ is defined by:

$$A_{sparse}^{\mathcal{T}}(S, S') = \begin{cases} 1 & S, S' \in V^{\mathcal{T}}, \{S, S'\} \in E^{\mathcal{T}} \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

We note that if the coarsened graph $\mathcal{T}(G)$ has a corresponding node feature map $\mathcal{X} : V^{\mathcal{T}} \rightarrow \mathbb{R}^d$, it also has sparse and dense vector representations defined similarly. Though the dense representation seems more natural, the sparse representation is also useful, as the symmetric group S_n acts on it by:

$$\sigma \cdot A_{sparse}^{\mathcal{T}}(S, S') = A_{sparse}^{\mathcal{T}}(\sigma^{-1}(S), \sigma^{-1}(S')). \quad (14)$$

When the type of representation is clear from context, we abuse notation and write $A^{\mathcal{T}}$. Note also that in the above discussion, we have used the term "node feature map". Throughout this paper, in order to denote the node features of a graph $G = (V, E)$ with $|V| = n$, we use both the vector representation $X \in \mathbb{R}^{n \times d}$ and the map representation $\mathcal{X} : V \rightarrow \mathbb{R}^d$ interchangeably. Now, recalling that our pipeline is defined to create and update a node feature map $\mathcal{X}(S, v)$ supported on the nodes of the product graph $G \square \mathcal{T}(G)$, we define a general node marking policy, the following way:

Definition A.4 (General Node Marking Policy). *A general node marking policy $\pi(\cdot, \cdot)$, is a function which takes as input a graph $G = (V, E)$, and a coarsening function $\mathcal{T}(\cdot)$, and returns a node feature map $\mathcal{X} : V^{\mathcal{T}} \times V \rightarrow \mathbb{R}^d$.*

In Appendix C We provide four different node marking policies, and analyze the effect on our pipeline. We now move on to define the general way in which we update a given node feature map on the product graph.

Definition A.5 (General CS-GNNLayer Update). *Given a graph $G = (V, E)$ and a coarsening function $\mathcal{T}(\cdot)$, let $\mathcal{X}^t(S, v) : V \times V^{\mathcal{T}} \rightarrow \mathbb{R}^d$ denote the node feature map at layer t . The general CS-GNNlayer update is defined by:*

$$\begin{aligned} \mathcal{X}^{t+1}(S, v) = f^t \left(\right. & \mathcal{X}^t(S, v), \\ & \text{agg}_1^t \{ \{ \mathcal{X}^t(S, v'), e_{v,v'} \mid v' \sim_G v \} \}, \\ & \text{agg}_2^t \{ \{ \mathcal{X}^t(S', v), \tilde{e}_{S,S'} \mid S' \sim_{G^{\mathcal{T}}} S \} \}, \\ & \text{agg}_3^t \{ \{ \mathcal{X}^t(S', v), z(S, v, S', v) \mid S' \in V^{\mathcal{T}} \text{ s.t. } v \in S' \} \}, \\ & \left. \text{agg}_4^t \{ \{ \mathcal{X}^t(S, v'), z(S, v, S, v') \mid v' \in V \text{ s.t. } v' \in S \} \} \right). \end{aligned} \quad (15)$$

Here, f^t is an arbitrary (parameterized) continuous function, agg_i^t , $i = 1, \dots, 4$ are learnable permutation invariant aggregation functions, $e_{v,v'}$, $\tilde{e}_{S,S'}$ are the (optional) edge features of G and $\mathcal{T}(G)$ respectively and the function $z : \mathcal{P}([n]) \times [n] \times \mathcal{P}([n]) \times [n] \rightarrow \mathbb{R}^d$ maps each tuple of indices $\mathbf{v} = (S, v, S', v')$ to a vector uniquely encoding the orbit of \mathbf{v} under the action of S_n as described in 73.

We note that for brevity, the notation used in the main body of the paper omits the aggregation functions $\text{agg}_1^t, \dots, \text{agg}_4^t$ and the edge features from the formulation of some of the layer updates. However, we explicitly state each component of the update, as we heavily utilize them in later proofs. We also note that this update is different than the general layer update presented in Equation (2), as it doesn't use all global updates characterized in 9. The reason for this is that some of the global updates have an asymptotic runtime of $\tilde{O}(n^2)$ where n is the number of nodes in the input graph. As our goal was to create models that improve on the scalability of standard subgraph GNNs which have

an asymptotic runtime of $\tilde{O}(n^2)$. We decided to discard some of the general global updates and keep only the ones that are induced by the last two entries in equation 15 which all have a linear runtime. After a stacking of the layers in Equation (15), we employ the following pooling procedure on the final node feature map \mathcal{X}^T :

$$\rho(\mathcal{X}^T) = \text{MLP}_2 \left(\sum_{S \in V^\mathcal{T}} \left(\text{MLP}_1 \left(\sum_{v \in V} \mathcal{X}^T(S, v) \right) \right) \right). \quad (16)$$

Finally, we define the set of all functions that can be expressed by our model:

Definition A.6 (Expressivity of Family of Graph Functions). *Let \mathcal{F} be a family of graph functions, we say that \mathcal{F} can express a graph function $g(\cdot)$ if for every finite family of graphs \mathcal{G} there exists a function $f \in \mathcal{F}$ such that:*

$$f(G) = g(G) \quad \forall G \in \mathcal{G}. \quad (17)$$

Here, \mathcal{G} is a finite family of graphs if all possible values of node/edge features of the graphs in \mathcal{G} form a finite set, and the maximal size of the graphs within \mathcal{G} is bounded.

Definition A.7 (Family of Functions Expressed By CS-GNN). *Let π be a general node marking policy and \mathcal{T} be a coarsening function. Define $\mathcal{S}(\mathcal{T}, \pi)$ to be the family of graph functions, which when given input graph $G = (V, E)$, first compute $\mathcal{X}^0(S, v)$ using $\pi(G, \mathcal{T})$, then update this node feature map by stacking T layers of the form 15, and finally pooling $\mathcal{X}^0(S, v)$ using equation 16. We define CS-GNN(\mathcal{T}, π) to be the set of all functions that can be expressed by $\mathcal{S}(\mathcal{T}, \pi)$.*

B Theoretical Validation of Implementation Details

In this section, we provide implementation details of our model and prove that they enable us to recover the conceptual framework of the model discussed thus far. First, we note that in Section 4.2, we characterized all equivariant linear maps $L : \mathbb{R}^{\mathcal{P}([n]) \times [n]} \rightarrow \mathbb{R}^{\mathcal{P}([n]) \times [n]}$ in order to incorporate them into our layer update. Given the high dimensionality of the space of all such linear maps, and in order to save parameters, we demonstrate that it is possible to integrate these layers into our layer update by adding edge features to a standard MPNN model. This is formalized in the following proposition:

Lemma B.1 (Parameter Sharing as MPNN). *Let $B_1, \dots, B_k : \mathbb{R}^{n \times n}$ be orthogonal matrices with entries restricted to 0 or 1, and let $W_1, \dots, W_k \in \mathbb{R}^{d \times d'}$ denote a sequence of weight matrices. Define $B_+ = \sum_{i=1}^k B_i$ and choose $z_1, \dots, z_k \in \mathbb{R}^{d^*}$ to be a set of unique vectors representing an encoding of the index set. The function that represents an update via parameter sharing:*

$$f(X) = \sum_{i=1}^k B_i X W_i, \quad (18)$$

can be implemented on any finite family of graphs \mathcal{G} , by a stack of MPNN layers of the following form [13],

$$m_v^l = \sum_{u \in N_{B_+}(v)} M^l(X_u^l, e_{u,v}), \quad (19)$$

$$X_v^{l+1} = U^l(X_v^l, m_v^l), \quad (20)$$

where U^l, M^l are multilayer perceptrons (MLPs). The inputs to this MPNN are the adjacency matrix B_+ , node feature vector X , and edge features – the feature of edge (u, v) is given by:

$$e_{u,v} = \sum_{i=1}^k z_i \cdot B_i(u, v). \quad (21)$$

Here, $B_i(u, v)$ denotes the (u, v) entry to matrix B_i .

The proof is given in Appendix G. The analysis in Section 4.2 demonstrates that the basis of the space of all equivariant linear maps $L : \mathbb{R}^{\mathcal{P}([n]) \times [n]} \rightarrow \mathbb{R}^{\mathcal{P}([n]) \times [n]}$ satisfies the conditions of Lemma F.1. Additionally, we notice that some of the equivariant linear functions have an asymptotic runtime of

$\tilde{O}(n^2)$ where n is the number of nodes in the input graph. As our main goal is to construct a more scalable alternative to node-based subgraph GNNs, which also have a runtime of $\tilde{O}(n^2)$, we limit ourselves to a subset of the basis for which all maps run in linear time. This is implemented by adding edge features to the adjacency matrices A_{P_1} and A_{P_2} , defined later in this section.

We now move on to discussing our specific implementation of the general layer update from Definition A.5.

Given a graph $G = (V, E)$ and a coarsening function \mathcal{T} , we aim to implement this general layer update by combining several standard message passing updates on the product graph $G \square \mathcal{T}(G)$. In the next two definitions, we define the adjacency matrices supported on the node set $V \times V^{\mathcal{T}}$, which serve as the foundation for these message passing procedures, and formalize the procedures themselves.

Definition B.1 (Adjacency Matrices on Product Graph). *Let $G = (V, E)$ be a graph with adjacency matrix A and node feature vector X , and let $\mathcal{T}(\cdot)$ be a coarsening function. We define the following four adjacency matrices on the vertex set $V^{\mathcal{T}} \times V$:*

$$A_G(S, v, S', v') = \begin{cases} 1 & v \sim_G v', S = S' \\ 0 & \text{otherwise.} \end{cases} \quad (22)$$

$$A_{\mathcal{T}(G)}(S, v, S', v') = \begin{cases} 1 & S \sim_{\mathcal{T}(G)} S', v = v' \\ 0 & \text{otherwise.} \end{cases} \quad (23)$$

$$A_{P_1}(S, v, S', v') = \begin{cases} 1 & v \in S', v = v' \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

$$A_{P_2}(S, v, S', v') = \begin{cases} 1 & v' \in S, S' = S \\ 0 & \text{otherwise.} \end{cases} \quad (25)$$

Given edge features $\{e_{v,v'} \mid v \sim_G v'\}$ and $\{\tilde{e}_{S,S'} \mid S \sim_{\mathcal{T}(G)} S'\}$ corresponding to the graphs G and $\mathcal{T}(G)$, respectively, we can trivially define the edge features corresponding to A_G and $A_{G\mathcal{T}}$ as follows:

$$e_G(S, v, S', v') = e_{v,v'}, \quad (26)$$

$$e_{\mathcal{T}(G)}(S, v, S', v') = \tilde{e}_{S,S'}. \quad (27)$$

In addition, for $i = 1, 2$, we define the edge features corresponding to adjacency matrices A_{P_i} as follows:

$$e_{P_i}(S, v, S', v') = z(S, v, S', v'). \quad (28)$$

Here, the function $z : \mathcal{P}([n]) \times [n] \times \mathcal{P}([n]) \times [n] \rightarrow \mathbb{R}^d$ maps each tuple $\mathbf{v} = (S, v, S', v')$ to a vector uniquely encoding the orbit of \mathbf{v} under the action of S_n as described in Equation 73.

Definition B.2 (CS-GNN Update Implementation). *Given a graph $G = (V, E)$, and a coarsening function $\mathcal{T}(\cdot)$, let $A_1 \dots A_4$ enumerate the set of adjacency matrices $\{A_G, A_{\mathcal{T}(G)}, A_{P_1}, A_{P_2}\}$. We define a CS-GNN layer update in the following way:*

$$\mathcal{X}_i^t(S, v) = U_i^t \left((1 + \epsilon_i^t) \cdot \mathcal{X}^t(S, v) + \sum_{(S', v') \sim_{A_i}(S, v)} M^t(\mathcal{X}^t(S', v') + e_i(S, v, S', v')) \right). \quad (29)$$

$$\mathcal{X}^{t+1}(S, v) = U_{\text{fin}}^t \left(\sum_{i=1}^4 \mathcal{X}_i^t(S, v) \right). \quad (30)$$

Here $\mathcal{X}^t(S, v)$ and $\mathcal{X}^{t+1}(S, v)$ denote the node feature maps of the product graph at layers t and $t + 1$, respectively. $e^1(S, v, S', v'), \dots, e^4(S, v, S', v')$ denote the edge features associated with adjacency matrices A_1, \dots, A_4 . $\epsilon_1^t, \dots, \epsilon_4^t$ represent learnable parameters in \mathbb{R} , and $U_1^t, \dots, U_4^t, U_{\text{fin}}^t, M^t$ all refer to multilayer perceptrons.

The next proposition states that using the layer update defined in equations 29 and 30 is enough to efficiently recover the general layer update defined in equation 15.

Proposition B.1 (Equivalence of General Layer and Implemented Layer). *Let $\mathcal{T}(\cdot)$ be a coarsening function, π be a generalized node marking policy, and \mathcal{G} be a finite family of graphs. Applying a stack of t general layer updates as defined in Equation 15 to the node feature map $\mathcal{X}(S, v)$ induced by $\pi(G, \mathcal{T})$, can be effectively implemented by applying a stack of t layer updates specified in Equations 29 and 30 to $\mathcal{X}(S, v)$. Additionally, the depths of all MLPs that appear in 29 and 30 can be bounded by 4.*

C Node Marking Policies – Theoretical Analysis

In this section, we define and analyze various general node marking policies, starting with four natural choices.

Definition C.1 (Four General Node Marking policies). *Let $G = (V, E)$ be a graph with adjacency matrix $A \in \mathbb{R}^{n \times n}$ and node feature vector $X \in \mathbb{R}^{n \times d}$, and let $\mathcal{T}(\cdot)$ be a coarsening function. All of the following node marking policies take the form:*

$$\pi(G, \mathcal{T}) = \mathcal{X}(S, v) = [X_u, b_\pi(S, v)], \quad (31)$$

where $[\cdot, \cdot]$ denotes the concatenation operator. We focus on four choices for $b_\pi(S, v)$:

1. Simple Node Marking:

$$b_\pi(S, v) = \begin{cases} 1 & \text{if } v \in S, \\ 0 & \text{if } v \notin S. \end{cases} \quad (32)$$

We denote this node marking policy by π_S .

2. Node + Size Marking:

$$b_\pi(S, v) = \begin{cases} (1, |S|) & \text{if } v \in S, \\ (0, |S|) & \text{if } v \notin S. \end{cases} \quad (33)$$

We denote this node marking policy by π_{SS} .

3. Minimum Distance:

$$b_\pi(S, v) = \min_{v' \in S} d_G(v, v') \quad (34)$$

where $d_G(v, v')$ is the shortest path distance between nodes v and v' in the original graph. We denote this node marking policy by π_{MD} .

4. Learned Distance Function:

$$b_\pi(S, v) = \phi(\{d_G(v, v') \mid v' \in S\}) \quad (35)$$

where $\phi(\cdot)$ is a learned permutation-invariant function. We denote this node marking policy by π_{LD} .

We note that when using the identity coarsening function $\mathcal{T}(G) = G$, our general node marking policies output node feature maps supported on the product $V \times V$. Thus, they can be compared to node marking policies used in node-based subgraph GNNs. In fact, in this case, both π_S and π_{SS} reduce to classical node-based node marking, while π_{MD} and π_{LD} reduce to distance encoding. The definitions of these can be found in [38]. Interestingly, even though in the case of node-based subgraph GNNs, both distance encoding and node marking were proven to be maximally expressive [38], in our case for some choices of \mathcal{T} , π_{LD} is strictly more expressive than the other three choices. The exact effect of each generalized node marking policy on the expressivity of our model is explored in the following two propositions.

Proposition C.1 (Equal Expressivity of Node Marking Policies). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$CS\text{-GNN}(\mathcal{T}, \pi_S) = CS\text{-GNN}(\mathcal{T}, \pi_{SS}) = CS\text{-GNN}(\mathcal{T}, \pi_{MD}). \quad (36)$$

Proposition C.2 (Expressivity of Learned Distance Policy). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$CS\text{-}GNN(\mathcal{T}, \pi_S) \subseteq CS\text{-}GNN(\mathcal{T}, \pi_{LD}). \quad (37)$$

In addition, for some choices of $\mathcal{T}(\cdot)$ the containment is strict.

The proofs of both propositions can be found in Appendix G. Finally, we provide a principled approach to deriving a generalized node marking policy based on symmetry invariance, and prove its equivalence to π_{SS} . Given a graph $G = (V, E)$ with $V = [n]$, adjacency matrix A , and node feature vector $X \in \mathbb{R}^{n \times d}$, along with a coarsening function $\mathcal{T}(\cdot)$, We define an action of the symmetric group S_n on the space $\mathbb{R}^{\mathcal{P}([n]) \times [n]}$ as follows:

$$\sigma \cdot \mathcal{X}(S, v) = \mathcal{X}(\sigma^{-1}(S), \sigma^{-1}(v)) \quad \text{for } \sigma \in S_n, \mathcal{X} \in \mathbb{R}^{\mathcal{P}([n]) \times [n]}. \quad (38)$$

Now, for each orbit $\gamma \in (\mathcal{P}([n]) \times [n])/S_n$, we define $\mathbf{1}_\gamma \in \mathbb{R}^{\mathcal{P}([n]) \times [n]}$ as follows:

$$\mathbf{1}_\gamma(S, v) = \begin{cases} 1 & (S, v) \in \gamma, \\ 0 & \text{otherwise.} \end{cases} \quad (39)$$

Choosing some enumeration of the orbit set $(\mathcal{P}([n]) \times [n])/S_n = \{\gamma_1, \dots, \gamma_k\}$, We now define the invariant generalized node marking policy π_{inv} by first setting:

$$b_{\pi_{inv}}^{\text{sparse}}(S, v) : \mathcal{P}([n]) \times [n] \rightarrow \mathbb{R}^k$$

and

$$b_{\pi_{inv}} : V^{\mathcal{T}} \times V \rightarrow \mathbb{R}^k$$

as follows:

$$b_{\pi_{inv}}^{\text{sparse}}(S, v) = [\mathbf{1}_{\gamma_1}(S, v), \dots, \mathbf{1}_{\gamma_k}(S, v)] \quad S \in \mathcal{P}(V), v \in V, \quad (40)$$

$$b_{\pi_{inv}}(S, v) = b_{\pi_{inv}}^{\text{sparse}}(S, v) \quad S \in V^{\mathcal{T}}, v \in V. \quad (41)$$

Then, we define the node feature map induced by π_{inv} as:

$$\mathcal{X}^{\pi_{inv}}(S, v) = [X_v, b_{\pi_{inv}}(S, v)]. \quad (42)$$

Interestingly, π_{inv} , derived solely from the group action of S_n on $\mathcal{P}([n]) \times [n]$, is equivalent to the generalized node marking policy π_{SS} . This is stated more rigorously in the following proposition:

Proposition C.3 (Node + Size Marking as Invariant Marking). *Given a graph $G = (V, E)$ with node feature vector $X \in \mathbb{R}^{n \times d}$, and a coarsening function $\mathcal{T}(\cdot)$, let $\mathcal{X}^{\pi_{SS}}, \mathcal{X}^{\pi_{inv}}$ be the node feature maps induced by π_{SS} and π_{inv} respectively. Recall that:*

$$\mathcal{X}^{\pi_{SS}}(S, v) = [X_v, b_{\pi_{SS}}(S, v)], \quad (43)$$

$$\mathcal{X}^{\pi_{inv}}(S, v) = [X_v, b_{\pi_{inv}}(S, v)]. \quad (44)$$

The following now holds:

$$b_{\pi_{inv}}(S, v) = OHE(b_{\pi_{SS}}(S, v)) \quad \forall S \in V^{\mathcal{T}}, \forall v \in V. \quad (45)$$

Here, OHE denotes a one-hot encoder, independent of the choice of both G and \mathcal{T} .

The proof of proposition C.3 can be found in Appendix G.

D Expressive Power of CS-GNN

D.1 Recovering Subgraph GNNs

In this section, we demonstrate that by choosing suitable coarsening functions, our architecture can replicate various previous subgraph GNN designs. We begin by focusing on node-based models, which are the most widely used type. We define a variant of these models which was proven in [38] to be maximally expressive, and show that our approach can recover it.

Definition D.1 (Maximally Expressive Subgraph GNN). *We define $MSGNN(\pi_{NM})$ as the set of all functions expressible by the following procedure:*

1. **Node Marking:** The representation of tuple $(u, v) \in V \times V$ is initially given by:

$$\mathcal{X}^0(u, v) = \begin{cases} 1 & \text{if } u = v, \\ 0 & \text{if } u \neq v. \end{cases} \quad (46)$$

2. **Update:** The representation of tuple (u, v) is updated according to:

$$\begin{aligned} \mathcal{X}^{t+1}(u, v) = f^t & \left(\mathcal{X}^t(u, v), \mathcal{X}^t(u, u), \mathcal{X}^t(v, v), \right. \\ & \text{agg}_1^t \{ \{ \mathcal{X}^t(u, v'), e_{v,v'} \} \mid v' \sim v \}, \\ & \left. \text{agg}_2^t \{ \{ \mathcal{X}^t(v, u'), e_{u,u'} \} \mid u' \sim u \} \right). \end{aligned} \quad (47)$$

3. **Pooling:** The final node feature vector $\mathcal{X}^T(u, v)$ is pooled according to:

$$\text{MLP}_2 \left(\sum_{u \in V} \text{MLP}_1 \left(\sum_{v \in V} \mathcal{X}^T(u, v) \right) \right). \quad (48)$$

Here, for any $t \in [T]$, f^t is any continuous (parameterized) functions, $\text{agg}_1^t, \text{agg}_2^t$ are any continuous (parameterized) permutation-invariant functions and $\text{MLP}_1, \text{MLP}_2$ are multilayer perceptrons.

Proposition D.1 (CS-GNN Can Implement MSGNN). *Let $\mathcal{T}(\cdot)$ be the identity coarsening function defined by:*

$$\mathcal{T}(G) = \{ \{v\} \mid v \in V \} \quad \forall G = (V, E). \quad (49)$$

The following holds:

$$\text{CS-GNN}(\mathcal{T}, \pi_S) = \text{MSGNN}(\pi_{NM}). \quad (50)$$

The proof of proposition D.1 can be found in Appendix G. We observe that, similarly, by selecting the coarsening function:

$$\mathcal{T}(G) = E \quad \forall G = (V, E), \quad (51)$$

one can recover edge-based subgraph GNNs. An example of such a model is presented in [4] (DS-GNN), where it was proven capable of distinguishing between two 3-WL indistinguishable graphs, despite having an asymptotic runtime of $\tilde{O}(m^2)$, where m is the number of edges in the input graph. This demonstrates our model’s ability to achieve expressivity improvements while maintaining a (relatively) low asymptotic runtime by exploiting the graph’s sparsity through the coarsening function. Finally, we note that by selecting the coarsening function:

$$\mathcal{T}(G) = \{ S \in \mathcal{P}(V) \mid |S| = k \} \quad G = (V, E), \quad (52)$$

We can recover an unordered variant of the k -OSAN model presented in [29].

D.2 Comparison to Natural Baselines

In this section, we demonstrate how our model can leverage the information provided by the coarsening function $\mathcal{T}(\cdot)$ in an effective way. First, we define a baseline model that incorporates \mathcal{T} in a straightforward manner. We then prove that, for any $\mathcal{T}(\cdot)$, our model is at least as expressive as this baseline. Additionally, we show that for certain choices of $\mathcal{T}(\cdot)$, our model exhibits strictly greater expressivity. To construct the baseline model, we first provide the following definition:

Definition D.2 (Coarsened Sum Graph). *Given a graph $G = (V, E)$ and a coarsening function $\mathcal{T}(\cdot)$, we define the coarsened sum graph $G_+^T = (V_+^T, E_+^T)$ by:*

- $V_+^T = V \cup V^T$.
- $E_+^T = E \cup E^T \cup \{ \{S, v\} \mid S \in V^T, v \in V, v \in S \}$.

If graph G had a node feature vector $X \in \mathbb{R}^{n \times d}$, we define the node feature vector of G_+^T as:


$$X_v = \begin{cases} [X_v, 1] & v \in V \\ 0_{d+1} & v \in V^T. \end{cases} \quad (53)$$

Here we concatenated a 1 to the end of node features of V to distinguish them from the nodes of V^T .

The connectivity of the sum graph (for our running example Figure 1) is visualized inset.

We now define our baseline model:

Definition D.3 (Coarse MPNN). *Let $\mathcal{T}(\cdot)$ be a coarsening function. Define $MPNN_+(\mathcal{T})$ as the set of all functions which can be expressed by the following procedure:*



1. **Preprocessing:** We first construct the sum graph G_+^T of the input graph G , along with a node feature map $\mathcal{X}^0 : V_+^T \rightarrow \mathbb{R}^d$ defined according to equation 53.
2. **Update:** The representation of node $v \in V_+^T$ is updated according to:

$$\begin{aligned}
 \text{For } v \in V : \quad & \mathcal{X}^{t+1}(v) = f_V^t(\mathcal{X}^t(v), \text{agg}_1^t\{\{\mathcal{X}^t(u), e_{u,v} \mid u \sim_G v\}\}, \\
 & \text{agg}_2^t\{\{\mathcal{X}^t(S) \mid S \in V^T, v \in S\}\}), \\
 \text{For } S \in V^T : \quad & \mathcal{X}^{t+1}(S) = f_{V^T}^t(\mathcal{X}^t(S), \text{agg}_1^t\{\{\mathcal{X}^t(S'), e_{S,S'} \mid S' \sim_{\mathcal{T}(G)} S\}\}, \\
 & \text{agg}_2^t\{\{\mathcal{X}^t(v) \mid v \in V, v \in S\}\}).
 \end{aligned} \tag{54}$$

3. **Pooling:** The final node feature vector $\mathcal{X}^T(\cdot)$ is pooled according to:

$$MLP \left(\sum_{v \in V_+^T} \mathcal{X}^T(v) \right). \tag{55}$$

Here, for $t \in [T]$, $f_V^t, f_{V^T}^t$ are continuous (parameterized) functions and $\text{agg}_1^t, \text{agg}_2^t$ are continuous (parameterized) permutation invariant functions. Finally, we notice that for the trivial coarsening function defined by

$$\mathcal{T}_\emptyset(G) = \emptyset, \tag{56}$$

the update in Equation (54) devolves into a standard MPNN update, as defined in [13] and so we define:

$$MPNN = MPNN_+(\mathcal{T}_\emptyset). \tag{57}$$

In essence, given an input graph $G = (V, E)$, the $MPNN_+(\mathcal{T})$ pipeline first constructs the coarsened graph $\mathcal{T}(G)$. It then adds edges between each super-node $S \in V^T$ and the nodes it is comprised of (i.e., any $v \in S$). This is followed by a standard message passing procedure on the graph. The following two propositions suggest that this simple approach to incorporating \mathcal{T} into a GNN pipeline is less powerful than our model.

Proposition D.2 (CS-GNN Is at Least as Expressive as Coarse MPNN). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$MPNN \subseteq MPNN_+(\mathcal{T}) \subseteq CS\text{-}GNN(\mathcal{T}, \pi_S) \tag{58}$$

Proposition D.3 (CS-GNN Can Be More Expressive Than $MPNN_+$). *Let $\mathcal{T}(\cdot)$ be the identity coarsening function defined by:*

$$\mathcal{T}(G) = \{\{v\} \mid v \in V\} \quad G = (V, E). \tag{59}$$

The following holds:

$$MPNN = MPNN_+(\mathcal{T}). \tag{60}$$

Thus:

$$MPNN_+(\mathcal{T}) \subset CS\text{-}GNN(\mathcal{T}, \pi_S), \tag{61}$$

where this containment is strict.

The proofs to the last two propositions can be found in Appendix G. Proposition D.3 demonstrates that CS-GNNs are strictly more expressive than MPNN_+ when using the identity coarsening function. However, this result extends to more complex coarsening functions as well. We briefly discuss one such example. Let $\mathcal{T}_\Delta(\cdot)$ be the coarsening function defined by:

$$\mathcal{T}_\Delta(G) = \{v_1, v_2, v_3 \mid G[v_1, v_2, v_3] \cong \Delta\}, \quad (62)$$

i.e. for an input graph G , the set of super-nodes is composed of all triplets of nodes whose induced subgraph is isomorphic to a triangle. To see that CS-GNN is strictly more expressive than MPNN_+ when using $\mathcal{T}_\Delta(\cdot)$, we look at the two graphs G and H depicted in Figure 4. In the figure, we see the two original graphs, G and H , their corresponding sum graphs $G_+^{\mathcal{T}_\Delta}$ and $H_+^{\mathcal{T}_\Delta}$, and a subgraph of their corresponding product graphs $G \square \mathcal{T}_\Delta(G)$ and $H \square \mathcal{T}_\Delta(H)$ induced by the sets $\{(S_0, v) \mid v \in V_G\}$ and $\{(S_0, v) \mid v \in V_H\}$ respectively (this can be thought of as looking at a single subgraph from the bag of subgraphs induced by CS-GNN). One can clearly see that both the original graphs and their respective sum graphs are 1-WL indistinguishable. On the other hand, the subgraphs induced by our method are 1-WL distinguishable. Since for both G and H the "bag of subgraphs" induced by CS-GNN is composed of 6 isomorphic copies of the same graph, this would imply that our method can distinguish between G and H , making it strictly more expressive than MPNN_+ .




Figure 4: Rows 1 and 3 depict two 1-WL indistinguishable graphs. Rows 2 and 4 depict the sum graph of each of these graphs, as well as one subgraph of their product graphs induced by all node, super-node tuples whose super-node is fixed.

We conclude this section with the following proposition, showing there exists coarsening functions which, when combined with CS-GNN, results in an architecture that is strictly more expressive than node-based subgraph GNNs.

Proposition D.4 (CS-GNN can be strictly more expressive than node-based subgraph GNNs). *Let \mathcal{T} be the coarsening function defined by:*

$$\mathcal{T}(G) = \{\{v\} \mid v \in V\} \cup E \quad G = (V, E). \quad (63)$$

The following holds:

1. Let G_1, G_2 be a pair of graphs such that there exists a node-based subgraph GNN model M where $M(G_1) \neq M(G_2)$. There exists a CS-GNN model M' which uses \mathcal{T} such that $M'(G_1) \neq M'(G_2)$.
2. There exists a pair of graphs G_1, G_2 such that for any subgraph GNN model M it holds that $M(G_1) = M(G_2)$, but there exists a CS-GNN model M' which uses \mathcal{T} such that $M'(G_1) \neq M'(G_2)$.

This proposition is proved in Appendix G.

E Linear Invariant (Equivariant) Layer – Extended Section

We introduce some key notation. In the matrix \mathcal{X} , the i -th row corresponds to the i -th subset S arranged in the lexicographic order of all subsets of $[n]$, namely, $\{\{0\}, \{0, 1\}, \{0, 2\}, \dots, \{0, 1, 2, \dots, n\}\}$. Each i -th position in this sequence aligns with the i -th row index in \mathcal{X} . It follows, that the standard basis for such matrices in $\mathbb{R}^{2^n \times n}$ is expressed as $\mathbf{e}^{(S)} \cdot \mathbf{e}^{(i)^T}$, where $\mathbf{e}^{(S)}$ is a 1-hot vector, with the value 1 positioned according to S in the lexicographic order. For a matrix $X \in \mathbb{R}^{a \times b}$, the operation of vectorization, denoted by $\text{vec}(X)$, transforms X into a single column vector in $\mathbb{R}^{ab \times 1}$ by sequentially stacking its columns; in the context of \mathcal{X} , the basis vectors of those vectors are $\mathbf{e}^{(i)} \otimes \mathbf{e}^{(S)}$. The inverse process, reshaping a vectorized matrix back to its original format, is denoted as $[\text{vec}(X)] = X$. We also denote an arbitrary permutation by $\sigma \in S_n$. The actions of permutations on vectors, whether indexed by sets or individual indices, are represented by $\mathbf{P}_S \in \text{GL}(2^n)$ and $\mathbf{P}_{\mathcal{I}} \in \text{GL}(n)$, respectively. This framework acknowledges S_n as a subgroup of the larger permutation group S_{2^n} , which permutes all 2^n positions in a given vector $\mathbf{v}_S \in \mathbb{R}^{2^n}$.

Let $\mathbf{L} \in \mathbb{R}^{1 \times 2^n \cdot n}$ be the matrix representation of a general linear operator $\mathcal{L} : \mathbb{R}^{2^n \times n} \rightarrow \mathbb{R}$ in the standard basis. The operator \mathcal{L} is order-invariant iff

$$\mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_{\mathcal{I}}) = \mathbf{L} \text{vec}(\mathcal{X}). \quad (64)$$

Similarly, let $\mathbf{L} \in \mathbb{R}^{2^n \cdot n \times 2^n \cdot n}$ denote the matrix for $\mathcal{L} : \mathbb{R}^{2^n \times n} \rightarrow \mathbb{R}^{2^n \times n}$. The operator \mathcal{L} is order-equivariant if and only if

$$[\mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_{\mathcal{I}})] = \mathbf{P}_S^T [\mathbf{L} \text{vec}(\mathcal{X})] \mathbf{P}_{\mathcal{I}}. \quad (65)$$

Using properties of the Kronecker product (see Appendices E.1 and E.2 for details), we derive the following conditions for invariant and equivariant linear layers:

$$\text{Invariant } \mathbf{L} : \quad \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \text{vec}(\mathbf{L}) = \text{vec}(\mathbf{L}), \quad (66)$$

$$\text{Equivariant } \mathbf{L} : \quad \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \otimes \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \text{vec}(\mathbf{L}) = \text{vec}(\mathbf{L}). \quad (67)$$

Solving Equations (66) and (67). Let $\sigma \in S_n$ denote a permutation corresponding to the permutation matrix \mathbf{P} . Let $\mathbf{P} \star \mathbf{L}$ denote the tensor that results from expressing \mathbf{L} after renumbering the nodes in $V^{\mathcal{T}}, V$ according to the permutation σ . Explicitly, for $\mathbf{L} \in \mathbb{R}^{2^n \times n}$, the $(\sigma(S), \sigma(i))$ -entry of $\mathbf{P} \star \mathbf{L}$ equals to the (S, i) -entry of \mathbf{L} . The matrix that corresponds to the operator $\mathbf{P} \star$ in the standard basis, $\mathbf{e}^{(i)} \otimes \mathbf{e}^{(S)}$ is the kronecker product $\mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S$. Since $\text{vec}(\mathbf{L})$ is exactly the coordinate vector of the tensor \mathbf{L} in the standard basis we have,

$$\text{vec}(\mathbf{P} \star \mathbf{L}) = \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \text{vec}(\mathbf{L}), \quad (68)$$

following the same logic, the following holds for the equivariant case, where $\mathbf{L} \in \mathbb{R}^{2^n \cdot n \times 2^n \cdot n}$,

$$\text{vec}(\mathbf{P} \star \mathbf{L}) = \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \otimes \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_S \text{vec}(\mathbf{L}). \quad (69)$$

Given Equations (66) and (68) and Equations (67) and (69), it holds that we should focus on solving,

$$\mathbf{P} \star \mathbf{L} = \mathbf{L}, \quad \forall \mathbf{P} \text{ permutation matrices}, \quad (70)$$

for both cases where $\mathbf{L} \in \mathbb{R}^{2^n \times n}$ and $\mathbf{L} \in \mathbb{R}^{2^n \times n \times 2^n \times n}$, corresponding to the bias term, and linear term.

Bias. To this end, let us define an equivalence relation in the index space of a tensor in $\mathbb{R}^{2^n \times n}$. Given a pair $(S, i) \in \mathcal{P}([n]) \times [n]$, we define γ^{k^+} to correspond to all pairs (S, i) such that $|S| = k$ and $i \notin S$. Similarly, γ^{k^-} corresponds to all pairs (S, i) such that $|S| = k$ and $i \in S$. We denote this equivalence relation as follows:

$$(\mathcal{P}([n]) \times [n]) / \sim \triangleq \{\gamma^{k^*} : k = 1, \dots, n; * \in \{+, -\}\}. \quad (71)$$

For each set-equivalence class $\gamma \in (\mathcal{P}([n]) \times [n]) / \sim$, we define a basis tensor, $\mathbf{B}^\gamma \in \mathbb{R}^{2^n \times n}$ by setting:

$$\mathbf{B}_{S,i}^\gamma = \begin{cases} 1, & \text{if } (S, i) \in \gamma; \\ 0, & \text{otherwise.} \end{cases} \quad (72)$$

Following similar reasoning, consider elements $(S_1, i_1, S_2, i_2) \in (\mathcal{P}([n]) \times [n] \times \mathcal{P}([n]) \times [n])$. We define a partition according to six conditions: the relationship between i_1 and i_2 , denoted as $i_1 \leftrightarrow i_2$, which is determined by the condition: $i_1 = i_2$ or $i_1 \neq i_2$; the cardinalities of S_1 and S_2 , denoted as k_1 and k_2 , respectively; the size of the intersection $S_1 \cap S_2$, denoted as k^\cap ; the membership of i_l in S_l for $l \in \{1, 2\}$, denoted as $\delta_{\text{same}} \in \{1, 2, 3, 4\}$; and the membership of i_{l_1} in S_{l_2} for distinct $l_1, l_2 \in \{1, 2\}$, denoted as $\delta_{\text{diff}} \in \{1, 2, 3, 4\}$. The equivalence relation thus defined can be represented as:

$$(\mathcal{P}([n]) \times [n] \times \mathcal{P}([n]) \times [n]) / \sim \triangleq \{\Gamma^{\leftrightarrow; k_1; k_2; k^\cap; \delta_{\text{same}}; \delta_{\text{diff}}}\}. \quad (73)$$

For each set-equivalence class $\Gamma \in (\mathcal{P}([n]) \times [n] \times \mathcal{P}([n]) \times [n]) / \sim$, we define a basis tensor, $\mathbf{B}^\Gamma \in \mathbb{R}^{2^n \times n \times 2^n \times n}$ by setting:

$$\mathbf{B}_{S_1, i_1; S_2, i_2}^\Gamma = \begin{cases} 1, & \text{if } (S_1, i_1, S_2, i_2) \in \Gamma; \\ 0, & \text{otherwise.} \end{cases} \quad (74)$$

The following two proposition summarizes the results in this section,

Lemma E.1 (γ (Γ) are orbits). *The sets $\{\gamma^{k^*} : k = 1, \dots, n; * \in \{+, -\}\}$ and $\{\Gamma^{\leftrightarrow; k_1; k_2; k^\cap; \delta_{\text{same}}; \delta_{\text{diff}}}\}$ are the orbits of S_n on the index space $(\mathcal{P}([n]) \times [n])$ and $(\mathcal{P}([n]) \times [n] \times (\mathcal{P}([n]) \times [n]))$, respectively.*

Proposition E.1 (Basis of Invariant (Equivariant) Layer). *The tensors \mathbf{B}^γ (\mathbf{B}^Γ) in Equation (72) (Equation (74)) form an orthogonal basis (in the standard inner product) to the solution of Equation (66) (Equation (67)).*

The proofs are given in Appendix G.

E.1 Full Derivation of Equation (66).

Our goal is to transition from the equation,

$$\mathbf{L} \text{vec}(\mathbf{P}_{\mathcal{I}}^T \mathcal{X} \mathbf{P}_{\mathcal{I}}) = \mathbf{L} \text{vec}(\mathcal{X}) \quad (64)$$

to the form,

$$\mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_{\mathcal{S}} \text{vec}(\mathbf{L}) = \text{vec}(\mathbf{L}) \quad (66)$$

We introduce the following property of the Kronecker product,

$$\text{vec}(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A}) \text{vec}(\mathbf{B}). \quad (75)$$

Using Equation (75) on the left side of Equation (64), we obtain

$$\mathbf{LP}_{\mathcal{I}}^T \otimes \mathbf{P}_{\mathcal{S}}^T \text{vec}(\mathcal{X}) = \mathbf{L} \text{vec}(\mathcal{X}), \quad (76)$$

since this should be true for any $\mathcal{X} \in \mathbb{R}^{2^n \times n}$, we derive

$$\mathbf{LP}_{\mathcal{I}}^T \otimes \mathbf{P}_{\mathcal{S}}^T = \mathbf{L}. \quad (77)$$

Applying the transpose operation on both sides, and noting that $(\mathbf{P}_{\mathcal{I}}^T \otimes \mathbf{P}_{\mathcal{S}}^T)^T = \mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_{\mathcal{S}}$, we obtain

$$\mathbf{P}_{\mathcal{I}} \otimes \mathbf{P}_{\mathcal{S}} \mathbf{L}^T = \mathbf{L}^T. \quad (78)$$

Recalling that $\mathbf{L} \in \mathbb{R}^{1 \times 2^n \cdot n}$, and thus $\mathbf{L}^T \in \mathbb{R}^{2^n \cdot n \times 1}$, we find that $\mathbf{L}^T = \text{vec}(\mathbf{L})$. Substituting this back into the previous equation we achieve Equation (66).

E.2 Full Derivation of Equation (67).

Our goal is to transition from the equation,

$$[\mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_I)] = \mathbf{P}_S^T [\mathbf{L} \text{vec}(\mathcal{X})] \mathbf{P}_I \quad (65)$$

to the form,

$$\mathbf{P}_I \otimes \mathbf{P}_S \otimes \mathbf{P}_I \otimes \mathbf{P}_S \text{vec}(\mathbf{L}) = \text{vec}(\mathbf{L}). \quad (67)$$

Applying the property in Equation (75), after the reverse operation of the vectorization, namely,

$$[\text{vec}(\mathbf{ABC})] = [(\mathbf{C}^T \otimes \mathbf{A})\text{vec}(\mathbf{B})] \quad (79)$$

on the right hand side of Equation (65), for

$$\mathbf{A} \triangleq \mathbf{P}_S^T; \quad (80)$$

$$\mathbf{B} \triangleq [\mathbf{L} \text{vec}(\mathcal{X})]; \quad (81)$$

$$\mathbf{C} \triangleq \mathbf{P}_I, \quad (82)$$

we obtain,

$$[\mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_I)] = [\mathbf{P}_I^T \otimes \mathbf{P}_S^T \mathbf{L} \text{vec}(\mathcal{X})]. \quad (83)$$

Thus, by omitting the reverse-vectorization operation,

$$\mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_I) = \mathbf{P}_I^T \otimes \mathbf{P}_S^T \mathbf{L} \text{vec}(\mathcal{X}). \quad (84)$$

Noting that $(\mathbf{P}_I^T \otimes \mathbf{P}_S^T)^{-1} = \mathbf{P}_I \otimes \mathbf{P}_S$, and multiplying by this inverse both sides (from the left), we obtain,

$$\mathbf{P}_I \otimes \mathbf{P}_S \mathbf{L} \text{vec}(\mathbf{P}_S^T \mathcal{X} \mathbf{P}_I) = \mathbf{L} \text{vec}(\mathcal{X}). \quad (85)$$

Applying, again, the property in Equation (75), we obtain,

$$\mathbf{P}_I \otimes \mathbf{P}_S \mathbf{L} \mathbf{P}_I^T \otimes \mathbf{P}_S^T \text{vec}(\mathcal{X}) = \mathbf{L} \text{vec}(\mathcal{X}). \quad (86)$$

Since this should be true for any $\mathcal{X} \in \mathbb{R}^{2^n \times n}$, we derive,

$$\mathbf{P}_I \otimes \mathbf{P}_S \mathbf{L} \mathbf{P}_I^T \otimes \mathbf{P}_S^T = \text{vec}(\mathbf{L}). \quad (87)$$

Again, applying Equation (75) on the left side, where,

$$\mathbf{A} \triangleq \mathbf{P}_I \otimes \mathbf{P}_S; \quad (88)$$

$$\mathbf{B} \triangleq \mathbf{L}; \quad (89)$$

$$\mathbf{C} \triangleq \mathbf{P}_I^T \otimes \mathbf{P}_S^T, \quad (90)$$

we get the following equality,

$$\mathbf{P}_I \otimes \mathbf{P}_S \mathbf{L} \mathbf{P}_I^T \otimes \mathbf{P}_S^T = \mathbf{P}_I \otimes \mathbf{P}_S \otimes \mathbf{P}_I \otimes \mathbf{P}_S \text{vec}(\mathbf{L}). \quad (91)$$

By substituting this to the left side of Equation (87) we obtain Equation (67).

E.3 Comparative Parameter Reduction in Linear Equivariant Layers

To demonstrate the effectiveness of our parameter-sharing scheme, which results from considering unordered tuples rather than ordered tuples, we present the following comparison. 3-IGNs [22] are structurally similar to our approach, with the main difference being that they consider indices as ordered tuples, while we consider them as sets. Both approaches use a total of six indices, as shown in the visualized block in Figure 3, making 3-IGNs a natural comparator. By leveraging our scheme, we reduce the number of parameters from 203 (the number of parameters in 3-IGNs) to just 35!

Table 5: Overview of the graph learning datasets.

Dataset	# Graphs	Avg. # nodes	Avg. # edges	Directed	Prediction task	Metric
ZINC-12K [31]	12,000	23.2	24.9	No	Regression	Mean Abs. Error
ZINC-FULL [31]	249,456	23.2	49.8	No	Regression	Mean Abs. Error
OGBG-MOLHIV [16]	41,127	25.5	27.5	No	Binary Classification	AUROC
OGBG-MOLBACE [16]	1513	34.1	36.9	No	Binary Classification	AUROC
OGBG-MOLESOL [16]	1,128	13.3	13.7	No	Regression	Root Mean Squ. Error
PEPTIDES-FUNC [9]	15,535	150.9	307.3	No	10-task Classification	Avg. Precision
PEPTIDES-STRUCT [9]	15,535	150.9	307.3	No	11-task Regression	Mean Abs. Error

F Extended Experimental Section

F.1 Dataset Description

In this section we overview the eight different datasets considered; this is summarized in Table 5.

ZINC-12K and ZINC-FULL Datasets [31, 14, 10]. The ZINC-12K dataset includes 12,000 molecular graphs sourced from the ZINC database, a compilation of commercially available chemical compounds. These molecular graphs vary in size, ranging from 9 to 37 nodes, where each node represents a heavy atom, covering 28 different atom types. Edges represent chemical bonds and there are three types of bonds. The main goal when using this dataset is to perform regression analysis on the constrained solubility (logP) of the molecules. The dataset is divided into training, validation, and test sets with 10,000, 1,000, and 1,000 molecular graphs respectively. The full version, ZINC-FULL, comprises approximately 250,000 molecular graphs, ranging from 9 to 37 nodes and 16 to 84 edges per graph. These graphs also represent heavy atoms, with 28 distinct atom types, and the edges indicate bonds between these atoms, with four types of bonds present.

OGBG-MOLHIV, OGBG-MOLBACE, OGBG-MOLESOL Datasets [16]. These datasets are used for molecular property prediction and have been adopted by the Open Graph Benchmark (OGB, MIT License) from MoleculeNet. They use a standardized featurization for nodes (atoms) and edges (bonds), capturing various chemophysical properties.

PEPTIDES-FUNC and PEPTIDES-STRUCT Datasets [9]. The PEPTIDES-FUNC and PEPTIDES-STRUCT datasets consist of atomic graphs representing peptides released with the Long Range Graph Benchmark (LRGB, MIT License). In PEPTIDES-FUNC, the task is to perform multi-label graph classification into ten non-exclusive peptide functional classes. Conversely, PEPTIDES-STRUCT is focused on graph regression to predict eleven three-dimensional structural properties of the peptides.

We note that for all datasets, we used the random splits provided by the public benchmarks.

F.2 Experimental Details

Implementation Details. Our implementation of Equation (2) is given by:

$$\mathcal{X}^{(l+1)} = \text{MLP} \left(\sum_{i=1}^3 \text{MPNN}^{(l+1,i)} (\mathcal{X}, \mathcal{A}_i) \right), \quad (92)$$

where $\mathcal{A}_1 = \mathcal{A}_G$, $\mathcal{A}_2 = \mathcal{A}_{\mathcal{T}(G)}$, and $\mathcal{A}_3 = \mathcal{A}_{\text{Equiv}}$.

For all considered datasets, namely, ZINC-12K, ZINC-FULL, OGBG-MOLHIV, OGBG-MOLBACE, and OGBG-MOLESOL, except for the PEPTIDES-FUNC and PEPTIDES-STRUCT datasets, we use a GINE [15] base encoder. Given an adjacency matrix \mathcal{A} , and defining $e_{(S',v'),(S,v)}$ to denote the edge features from node (S', v') to node (S, v) , it takes the following form:

$$\mathcal{X}(S, v) = \text{MLP} \left((1 + \epsilon) \cdot \mathcal{X}(S, v) + \sum_{(S',v') \sim_{\mathcal{A}} (S,v)} \text{ReLU}(\mathcal{X}(S', v') + e_{(S',v'),(S,v)}) \right). \quad (93)$$

We note that for the symmetry-based updates, we switch the ReLU to an MLP¹² to align with the theoretical analyses¹³ (Appendix B), stating that we can implement the equivariant update developed

¹²With the exception of the OGB datasets, to avoid overfitting.

¹³The theoretical analysis assumes the usage of an MLP for all three considered updates.

in Section 4.2. A more thorough discussion regarding the implementation of the symmetry-based updates is given in Appendix F.4.

When experimenting with the PEPTIDES-FUNC and PEPTIDES-STRUC datasets, we employ GAT [32] as our underlying MPNN to ensure a fair comparison with the random baseline—the random variant of Subgraphormer + PE [3]. To clarify, we consider the random variant of Subgraphormer + PE as a natural random baseline since it incorporates the information in the eigenvectors of the Laplacian (which we also do via the coarsening function). To maintain a fair comparison, we use a single vote for this random baseline, and maintained the same hyperparameters.

Our experiments were conducted using the PyTorch [28] and PyTorch Geometric [11] frameworks (resp. BSD and MIT Licenses), using a single NVIDIA L40 GPU, and for every considered experiment, we show the mean \pm std. of 3 runs with different random seeds. Hyperparameter tuning was performed utilizing the Weight and Biases framework [6] – see Appendix F.3. All our MLPs feature a single hidden layer equipped with a ReLU non-linearity function. For the encoding of atom numbers and bonds, we utilized learnable embeddings indexed by their respective numbers.

In the case of the OGBG-MOLHIV, OGBG-MOLE SOL, OGBG-MOLBACE datasets, we follow Frasca et al. [12], therefore adding a residual connection between different layers. Additionally, for those datasets (except OGBG-MOLHIV), we used linear layers instead of MLPs inside the GIN layers. Moreover, for these four datasets, and for the PEPTIDES datasets, the following pooling mechanism was employed

$$\rho(\mathcal{X}) = \text{MLP} \left(\sum_S \left(\frac{1}{n} \sum_{v=1}^n \mathcal{X}(s, v) \right) \right). \quad (94)$$

For the PEPTIDES datasets, we also used a residual connection between layers.

F.3 HyperParameters

In this section, we detail the hyperparameter search conducted for our experiments. Besides standard hyperparameters such as learning rate and dropout, our specific hyperparameters are:

1. **Laplacian Dimension:** This refers to the number of columns used in the matrix U , where $L = U^T \lambda U$, for the spectral clustering in the coarsening function.
2. **SPD Dimension:** This represents the number of indices used in the node marking equation. To clarify, since $|S|$ might be large, we opt for using the first k indices that satisfy $i \in S$, sorted according to the SPD distance.

SPD Dimension. For the Laplacian dimension, we chose a fixed value of 10 for all bag sizes for both ZINC-12K and ZINC-FULL datasets. For OGBG-MOLHIV, we used a fixed value of 1, since the value 10 did not perform well. For the PEPTIDES datasets, we also used the value 1. For the OGBG-MOLE SOL and OGBG-MOLBACE datasets, we searched over the two values $\{1, 2\}$.

Laplacian Dimension. For the Laplacian dimension, we searched over the values $\{1, 2\}$ for all datasets.

Standard Hyperparameters. For ZINC-12K, we used a weight decay of 0.0003 for all bag sizes, except for the full bag size, for which we used 0.0001.

All of the hyperparameter search configurations are presented in Table 6, and the selected hyperparameters are presented in Table 7.

Table 6: Hyperparameters search for CS-GNN.

Dataset	Bag size	Num. layers	Learning rate	Embedding size	Epochs	Batch size	Dropout	Laplacian dimension	SPD dimension
ZINC-12K	$T = 2$	6	0.0005	96	400	128	0	$\{1, 2\}$	10
ZINC-12K	$T \in \{3, 4, 5, 8, 18\}$	6	0.0007	96	400	128	0	$\{1, 2\}$	10
ZINC-12K	$T = \text{"full"}$	6	0.0007	96	500	128	0	$\{1, 2\}$	10
ZINC-FULL	$T = 4$	6	0.0007	96	400	128	0	$\{1, 2\}$	10
ZINC-FULL	$T = \text{"full"}$	6	$\{0.001, 0.0005\}$	96	500	128	0	$\{1, 2\}$	10
OGBG-MOLHIV	$T \in \{2, 5, \text{"full"}\}$	2	0.01	60	100	32	0.5	$\{1, 2\}$	1
OGBG-MOLE SOL	$T \in \{2, 5, \text{"full"}\}$	3	0.001	60	100	32	0.3	$\{1, 2\}$	$\{1, 2\}$
OGBG-MOLBACE	$T \in \{2, 5, \text{"full"}\}$	$\{2, 3\}$	0.01	60	100	32	0.3	$\{1, 2\}$	$\{1, 2\}$
PEPTIDES-FUNC	$T = 30$	5	$\{0.01, 0.005\}$	96	200	128	0	$\{1, 2\}$	1
PEPTIDES-STRUC	$T = 30$	4	$\{0.01, 0.005\}$	96	200	128	0	$\{1, 2\}$	1

Table 7: Chosen Hyperparameters for CS-GNN.

Dataset	Bag size	Num. layers	Learning rate	Embedding size	Epochs	Batch size	Dropout	Laplacian dimension	SPD dimension
ZINC-12K	$T = 2$	6	0.0005	96	400	128	0	1	10
ZINC-12K	$T = 3$	6	0.0007	96	400	128	0	2	10
ZINC-12K	$T = 4$	6	0.0007	96	400	128	0	1	10
ZINC-12K	$T = 5$	6	0.0007	96	400	128	0	1	10
ZINC-12K	$T = 8$	6	0.0007	96	400	128	0	1	10
ZINC-12K	$T = 18$	6	0.0007	96	400	128	0	1	10
ZINC-12K	$T = \text{"full"}$	6	0.0007	96	500	128	0	N/A	10
ZINC-FULL	$T = 4$	6	0.0007	96	400	128	0	1	10
ZINC-FULL	$T = \text{"full"}$	6	0.0005	96	500	128	0	N/A	N/A
OGBG-MOLHIV	$T = 2$	2	0.01	60	100	32	0.5	1	1
OGBG-MOLHIV	$T = 5$	2	0.01	60	100	32	0.5	1	1
OGBG-MOLHIV	$T = \text{"full"}$	2	0.01	60	100	32	0.5	N/A	N/A
OGBG-MOLESOL	$T = 2$	3	0.001	60	100	32	0.3	1	2
OGBG-MOLESOL	$T = 5$	3	0.001	60	100	32	0.3	1	2
OGBG-MOLESOL	$T = \text{"full"}$	3	0.001	60	100	32	0.3	N/A	N/A
OGBG-MOLBACE	$T = 2$	3	0.01	60	100	32	0.3	1	1
OGBG-MOLBACE	$T = 5$	3	0.01	60	100	32	0.3	1	2
OGBG-MOLBACE	$T = \text{"full"}$	3	0.01	60	100	32	0.3	N/A	N/A
PEPTIDES-FUNC	$T = 30$	5	0.005	96	200	128	0	1	1
PEPTIDES-STRUC	$T = 30$	4	0.01	96	200	128	0	1	1

Optimizers and Schedulers. For the ZINC-12K and ZINC-FULL datasets, we employ the Adam optimizer paired with a ReduceLRonPlateau scheduler, factor set to 0.5, patience at 40^{14} , and a minimum learning rate of 0. For the OGBG-MOLHIV dataset, we utilized the ASAM optimizer [21] without a scheduler. For both OGBG-MOLESOL and OGBG-MOLBACE, we employed a constant learning rate without any scheduler. Lastly, for the PEPTIDES-FUNC and PEPTIDES-STRUCT datasets, the AdamW optimizer was chosen in conjunction with a cosine annealing scheduler, incorporating 10 warmup epochs.

F.4 Implementation of Linear Equivariant and Invariant layers – Extended Section

In this section, in a more formal discussion, we specify how to integrate those invariant and equivariant layers to our proposed architecture. We start by drawing an analogy between parameter sharing in linear layers and the operation of an MPNN on a fully connected graph with edge features in the following lemma,

Lemma F.1 (Parameter Sharing as MPNN). *Let $B_1, \dots, B_k : \mathbb{R}^{n \times n}$ be orthogonal matrices with entries restricted to 0 or 1, and let $W_1, \dots, W_k \in \mathbb{R}^{d \times d}$ denote a sequence of weight matrices. Define $B_+ = \sum_{i=1}^k B_i$ and choose $z_1, \dots, z_k \in \mathbb{R}^{d^*}$ to be a set of unique vectors representing an encoding of the index set. The function, which represents an update via parameter sharing:*

$$f(X) = \sum_{i=1}^k B_i X W_i, \quad (95)$$

can be implemented by a stack of MPNN layers of the following form [13],

$$m_u^l = \sum_{v \in N_{B_+}(u)} M^l(X_v^l, e_{u,v}), \quad (96)$$

$$X_u^{l+1} = U^l(X_u^l, m_u^l), \quad (97)$$

where U^l, M^l are multilayer perceptrons (MLPs). The inputs to this MPNN are the adjacency matrix B_+ , node feature vector X , and edge features – the feature of edge (u, v) is given by:

$$e_{u,v} = \sum_{i=1}^k z_i \cdot B_i(u, v). \quad (98)$$

Here, $B_i(u, v)$ denotes the (u, v) entry to matrix B_i .

The proof is given in Appendix G.

¹⁴For ZINC-12K, $T \in \{2, \text{"full"}\}$, we used a patience of 50.

Table 8: Comparison over the ZINC-FULL molecular dataset under 500k parameter budget. The best performing method is highlighted in **blue**, while the second best is highlighted in **red**.

Model ↓ / Dataset →	ZINC-FULL (MAE ↓)
MAG-GNN [20] ($T = 4$)	0.030 ±0.002
Ours ($T = 4$)	0.027 ±0.002
GNN-SSWL [38] ($T = \text{“full”}$)	0.026±0.001
GNN-SSWL+ [38] ($T = \text{“full”}$)	0.022±0.001
Subgraphormer [3] ($T = \text{“full”}$)	0.020 ±0.002
Subgraphormer + PE [3] ($T = \text{“full”}$)	0.023±0.001
Ours ($T = \text{“full”}$)	0.021 ±0.001

Thus, our implementation for the global update is as follows,

$$\mathcal{X}(S, i) = \text{MLP} \left((1 + \epsilon) \cdot \mathcal{X}(S, i) + \sum_{(S', i') \sim_{\mathcal{A}_{\text{Equiv}}}(S, i)} \text{MLP} \left(\mathcal{X}(S', i') + e_{(S', i'), (S, i)} \right) \right), \quad (99)$$

where $e_{(S', i'), (S, i)} = \sum_{\Gamma} z_{\Gamma} \cdot \mathbf{B}_{S, i; S', i'}^{\Gamma}$ and z_{Γ} are orthogonal 1-hot vectors for different Γ 's. The connectivity $\mathcal{A}_{\text{Equiv}}$ is such that $\mathcal{A}_{\text{Equiv}}(S, v, S', v')$ contains the value one iff $v \in S, v = v'$. This corresponds to choosing only several Γ 's in the partition, and since each Γ is invariant to the permutation, this choice still maintains equivariance.

F.5 Additional Results

ZINC-FULL. Below, we present our results on the ZINC-FULL dataset for a bag size of $T = 4$ and the full-bag. For the bag size $T = 4$, we benchmark against MAG-GNN [20], which in their experiments used the best out of the bag sizes $T \in \{2, 3, 4\}$; however, they did not specify which one performed the best. The results are summarized in Table 8.

ZINC-12K – additional results. We present all the results from Figure 2, along with some additional ones, in Table 9.

Runtime comparison. We compare the training time and prediction performance on the ZINC-12K dataset. For all methods, we report the training and inference times on the entire training and test sets, respectively, using a batch size of 128. Our experiments were conducted using an NVIDIA L40 GPU, while for the baselines, we used the timing reported in [5], which utilized an RTX A6000 GPU. The runtime comparison is presented in Table 10.

F.6 ZINC12K Product Graph Visualization

In this subsection, we visualize the product graph derived from the first graph in the ZINC12K dataset. Specifically, we present the right part of Figure 1, for the case of the real-world graphs in the ZINC12K dataset. We perform this visualization for different cluster sizes, $T \in \{2, 3, 4, 5, 8, 12\}$, which also define the bag size, hence the notation T . The nodes in the product graph, $\mathcal{T}(G) \square G$, are (S, v) , where S is the coarsened graph node (again a tuple), and v is the node index (of a node from the original graph). For better clarity, we color the nodes (S, v) with $v \in S$ using different colors, while reserving the gray color exclusively for nodes (S, v) where $v \notin S$. The product graphs are visualized in Figures 5 to 10 below.

Table 9: Test results on the ZINC-12K molecular dataset under 500k parameter budget. The top two results are reported as **First** and **Second**.

Method	Bag size	ZINC (MAE ↓)
GCN [19]	$T = 1$	0.321 ± 0.009
GIN [35]	$T = 1$	0.163 ± 0.004
OSAN [29]	$T = 2$	0.177 ± 0.016
Random [20]	$T = 2$	0.131 ± 0.005
PL [5]	$T = 2$	0.120 ± 0.003
Mag-GNN [20]	$T = 2$	0.106 ± 0.014
Ours	$T = 2$	0.109 ± 0.005
Random [20]	$T = 3$	$0.124 \pm \text{N/A}$
Mag-GNN [20]	$T = 3$	$0.104 \pm \text{N/A}$
Ours	$T = 3$	0.096 ± 0.005
Random [20]	$T = 4$	$0.125 \pm \text{N/A}$
Mag-GNN [20]	$T = 4$	$0.101 \pm \text{N/A}$
Ours	$T = 4$	0.090 ± 0.003
Random [5]	$T = 5$	0.113 ± 0.006
PL [5]	$T = 5$	0.109 ± 0.005
Ours	$T = 5$	0.095 ± 0.003
Random [5]	$T = 8$	0.102 ± 0.003
PL [5]	$T = 8$	0.097 ± 0.005
Ours	$T = 8$	0.094 ± 0.006
Ours	$T = 18$	0.082 ± 0.003
NGNN [39]	Full	0.111 ± 0.003
DS-GNN [4]	Full	0.116 ± 0.009
DSS-GNN [4]	Full	0.102 ± 0.003
GNN-AK [40]	Full	0.105 ± 0.010
GNN-AK+ [40]	Full	0.091 ± 0.002
SUN [12]	Full	0.083 ± 0.003
OSAN [29]	Full	0.154 ± 0.008
GNN-SSWL+ [38]	Full	0.070 ± 0.005
Subgraphormer [3]	Full	0.067 ± 0.007
Subgraphormer+PE [3]	Full	0.063 ± 0.001
Ours	Full	0.062 ± 0.0007

Table 10: Run time comparison over the ZINC-12K dataset. Time taken at train for one epoch and at inference on the test set. All values are in milliseconds.

Method	Train time (for a single epoch; ms)	Test time (ms)	MAE ↓
GIN [35]	1370.10 ± 10.97	84.81 ± 0.26	0.163 ± 0.004
OSAN [29] ($T = 2$)	2964.46 ± 30.36	227.93 ± 0.21	0.177 ± 0.016
PL [5] ($T = 2$)	2489.25 ± 9.42	150.38 ± 0.33	0.120 ± 0.003
Ours ($T = 2$)	2764.60 ± 234	383.14 ± 15.74	0.109 ± 0.005




Figure 5: $T = 2$.




Figure 6: $T = 3$.




Figure 7: $T = 4$.




Figure 8: $T = 5$.




Figure 9: $T = 8$.




Figure 10: $T = 12$.

G Proofs

G.1 Proofs of Appendix B

We first state the memorization theorem, proven in [36], which will be heavily used in a lot of the proofs in this section.

Theorem G.1 (Memorization Theorem). *Consider a dataset $\{x_j, y_j\}_{j=1}^N \in \mathbb{R}^d \times \mathbb{R}^{d_y}$, with each x_j being distinct and every $y_j \in \{0, 1\}^{d_y}$. There exists a 4-layer fully connected ReLU neural network $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^{d_y}$ that perfectly maps each x_j to its corresponding y_j , i.e., $f_\theta(x_j) = y_j$ for all j .*

We now restate and prove the propositions and lemmas of Appendix B.

Lemma B.1 (Parameter Sharing as MPNN). *Let $B_1, \dots, B_k : \mathbb{R}^{n \times n}$ be orthogonal matrices with entries restricted to 0 or 1, and let $W_1, \dots, W_k \in \mathbb{R}^{d \times d'}$ denote a sequence of weight matrices. Define $B_+ = \sum_{i=1}^k B_i$ and choose $z_1, \dots, z_k \in \mathbb{R}^{d^*}$ to be a set of unique vectors representing an encoding of the index set. The function that represents an update via parameter sharing:*

$$f(X) = \sum_{i=1}^k B_i X W_i, \quad (18)$$

can be implemented on any finite family of graphs \mathcal{G} , by a stack of MPNN layers of the following form [13],

$$m_v^l = \sum_{u \in N_{B_+}(v)} M^l(X_u^l, e_{u,v}), \quad (19)$$

$$X_v^{l+1} = U^l(X_v^l, m_v^l), \quad (20)$$

where U^l, M^l are multilayer perceptrons (MLPs). The inputs to this MPNN are the adjacency matrix B_+ , node feature vector X , and edge features – the feature of edge (u, v) is given by:

$$e_{u,v} = \sum_{i=1}^k z_i \cdot B_i(u, v). \quad (21)$$

Here, $B_i(u, v)$ denotes the (u, v) entry to matrix B_i .

Proof. Since we are concerned only with input graphs G from a finite family of graphs (where "finite" means that the maximal graph size is bounded and all possible node and edge feature values come from a finite set), we assume that for any $v \in [n]$, $i \in [k]$, both the input feature vectors $X_v \in \mathbb{R}^d$ and the encoding vectors $z_i \in \mathbb{R}^{d^*}$ are one-hot encoded. We aim to show that under these assumptions, any function $f(\cdot)$ of the form 95 can be realized through a single-layer update detailed in Equations 97, 96, where M is a 4 layer MLP, and U is a single linear layer. The proof involves the following steps:

1. Compute $[B_1 X, \dots, B_k X]$ using the message function M .
2. Compute $f(X)$ using the update function U .

Step 1: We notice that for every $i \in [k]$, $v \in [n]$ we have:

$$(B_i X)_v = \sum_{B_i(v,u)=1} X_u = \sum_{u \in N_{B_+}(v)} X_u \cdot \mathbf{1}_{z_i}(e_{u,v}). \quad (100)$$

Here $\mathbf{1}_{z_i}$ is the indicator function of the set $\{z_i\}$. We notice that since X_u and z_i are one-hot encoded, there is a finite set of possible values for the pair $(X_u, e_{u,v})$. In addition, the function:

$$\text{enc}(X_u, e_{u,v}) = [X_u \cdot \mathbf{1}_{z_1}(e_{u,v}), \dots, X_u \cdot \mathbf{1}_{z_k}(e_{u,v})] \quad (101)$$

outputs vectors in the set $\{0, 1\}^{d \times k}$. Thus, employing the memorization theorem G.1, we define a dataset $\{x_j, y_j\}_{j=1}^N$ by taking the x_j s to be all possible (distinct) values of $(X_u, e_{u,v})$ with each

corresponding y_i being the output $\text{enc}(x_i)$. We note that there are finitely many such values as both X_u and $e_{u,v}$ are one-hot encoded. The theorem now tells us that there exists a 4-layer fully connected ReLU neural network M such that:

$$M(X_u, e_{u,v}) = \text{enc}(X_u, e_{u,v}). \quad (102)$$

and so, equation 100 implies:

$$m_v = \sum_{u \in N_{B_+}(v)} M(X_u, e_{u,v}) = [(B_1 X)_v, \dots, (B_k X)_v]. \quad (103)$$

Step 2: Define $P_i : \mathbb{R}^{k \times d} \rightarrow \mathbb{R}^d$ as the projection operator, extracting coordinates $d \cdot i + 1$ through $d \cdot (i + 1)$ from its input vector:

$$P_i(V) = V|_{d \cdot i + 1 : d \cdot (i + 1)}. \quad (104)$$

We define the update function to be the following linear map:

$$U(X_v, m_v) = \sum_{i=1}^k P_i(m_v) W_i. \quad (105)$$

Combining equations 103 and 105 we get:

$$\tilde{X}_v = U(X_v, m_v) = \sum_{i=1}^k (B_i X)_v \cdot W_i = f(X)_v. \quad (106)$$

□

Proposition B.1 (Equivalence of General Layer and Implemented Layer). *Let $\mathcal{T}(\cdot)$ be a coarsening function, π be a generalized node marking policy, and \mathcal{G} be a finite family of graphs. Applying a stack of t general layer updates as defined in Equation 15 to the node feature map $\mathcal{X}(S, v)$ induced by $\pi(G, \mathcal{T})$, can be effectively implemented by applying a stack of t layer updates specified in Equations 29 and 30 to $\mathcal{X}(S, v)$. Additionally, the depths of all MLPs that appear in 29 and 30 can be bounded by 4.*

Proof. For convenience, let us first restate the general layer update:

$$\begin{aligned} \mathcal{X}^{t+1}(S, v) = f^t \left(\mathcal{X}^t(S, v), \right. \\ \left. \text{agg}_1^t \{ \{ \mathcal{X}^t(S, v'), e_{v,v'} \mid v' \sim_G v \} \}, \right. \\ \left. \text{agg}_2^t \{ \{ \mathcal{X}^t(S', v), \tilde{e}_{S,S'} \mid S' \sim_{G^T} S \} \}, \right. \\ \left. \text{agg}_3^t \{ \{ \mathcal{X}^t(S', v), z(S, v, S', v) \mid S' \in V^T \text{ s.t. } v \in S' \} \}, \right. \\ \left. \text{agg}_4^t \{ \{ \mathcal{X}^t(S, v'), z(S, v, S, v') \mid v' \in V \text{ s.t. } v' \in S \} \} \right), \end{aligned} \quad (15)$$

as well as the two step implemented layer update:

$$\mathcal{X}_i^t(S, v) = U_i^t \left((1 + \epsilon_i^t) \cdot \mathcal{X}^t(S, v) + \sum_{(S', v') \sim_{A_i}(S, v)} M^t(\mathcal{X}^t(S', v') + e_i(S, v, S', v')) \right). \quad (29)$$

$$\mathcal{X}^{t+1}(S, v) = U_{\text{fin}}^t \left(\sum_{i=1}^4 \mathcal{X}_i^t(S, v) \right). \quad (30)$$

We aim to demonstrate that any general layer, which updates the node feature map $\mathcal{X}^t(S, v)$ at layer t to node feature map $\mathcal{X}^{t+1}(S, v)$ at layer $t + 1$ as described in equation 15, can be effectively implemented using the layer update processes outlined in equations 29 and 30.

As we are concerned only with input graphs belonging to the finite graph family \mathcal{G} (where "finite" indicates that the maximal graph size is bounded and all node and edge features have a finite set of possible values), we assume that the values of the node feature map $\mathcal{X}^t(S, v)$ and the edge feature vectors $e_i(S, v, S', v')$ are represented as one-hot vectors in \mathbb{R}^k . We also assume that the parameterized functions f^t and $\text{agg}_1^t, \dots, \text{agg}_4^t$, which are applied in Equation 15 outputs one-hot vectors. Finally, we assume that there exists integers d, d^* , such that the node feature map values are supported on coordinates $1, \dots, d$, the edge feature vectors are supported on coordinates $d + 1, \dots, d + d^*$, and coordinates $d + d^* + 1, \dots, k$ are used as extra memory space, with:

$$k > d \times d^* + d + d^*. \quad (107)$$

We note that the last assumption can be easily achieved using padding. The proof involves the following steps:

1. For $i = 1, \dots, 4$, Use the term:

$$m_i^t = \sum_{(S', v') \sim_{A_i}(S, v)} M^t(\mathcal{X}^t(S', v') + e_i(S, v, S', v')) \quad (108)$$

to uniquely encode:

$$\{(\mathcal{X}^t(S', v'), e_i(S, v, S', v')) \mid (S', v') \sim_{A_i}(S, v)\}. \quad (109)$$

2. Use the term:

$$\mathcal{X}_*^t = \sum_{i=1}^4 \mathcal{X}_i^t(S, v) \quad (110)$$

to uniquely encode the input of f^t as a whole.

3. Implement the parameterized function f^t .

Step 1: Since we assume that node feature map values and edge feature vectors are supported on orthogonal sub-spaces of \mathbb{R}^k , the term:

$$\mathcal{X}^t(S, v) + e_i(S, v, S', v') \quad (111)$$

uniquely encodes the value of the tuple:

$$(\mathcal{X}^t(S, v), e_i(S, v, S', v')). \quad (112)$$

Since $\mathcal{X}^t(S, v)$ is a one-hot encoded vector with d possible values, while $e_i(S, v, S', v')$ is a one-hot encoded vector with d^* possible values, their sum has $d \cdot d^*$ possible values. Thus there exists a function:

$$\text{enc} : \mathbb{R}^k \rightarrow \mathbb{R}^k$$

which encodes each such possible value as a one-hot vector in \mathbb{R}^k supported on the last $k - d - d^*$ coordinates (this is possible because of equation 107). Now, employing theorem G.1, we define the x_j s as all possible (distinct) values of 111, with each corresponding y_j being the output $\text{enc}(x_j)$. The theorem now tells us that there exists a 4-layer fully connected ReLU neural network capable of implementing the function $\text{enc}(\cdot)$. We choose M^t to be this network. Now since m_i^t , defined in equation 108 is a sum of one-hot encoded vectors, it effectively counts the number of each possible value in the set 109. This proves step 1.

Step 2: First, we note that:

$$\begin{aligned} & \{(\mathcal{X}^t(S, v'), e_{v, v'}) \mid v' \sim_G v\} \\ &= \{(\mathcal{X}^t(S', v'), e_G(S, v, S', v')) \mid (S, v) \sim_{A_G}(S', v')\} \end{aligned} \quad (113)$$

$$\begin{aligned} & \{(\mathcal{X}^t(S', v), e_{s', s}) \mid S' \sim_{\mathcal{T}(G)} S\} \\ &= \{(\mathcal{X}^t(S', v'), e_{\mathcal{T}(G)}(S, v, S', v')) \mid (S, v) \sim_{A_{\mathcal{T}(G)}}(S', v')\} \end{aligned} \quad (114)$$

$$\begin{aligned} & \{(\mathcal{X}^t(S', v), z(S, v, S', v')) \mid v \in S'\} \\ &= \{(\mathcal{X}^t(S', v'), e_{P_1}(S, v, S', v')) \mid (S, v) \sim_{A_{P_1}}(S', v')\} \end{aligned} \quad (115)$$

$$\begin{aligned} & \{ \{ \mathcal{X}^t(S, v'), z(S, v, S', v') \} \mid v' \in S \} \\ & = \{ \{ \mathcal{X}^t(S', v'), e_{P_2}(S, v, S', v') \} \mid (S, v) \sim_{A_{P_2}} (S', v') \} \end{aligned} \quad (116)$$

Now, since m_i^t and $\mathcal{X}^t(S, v)$ are supported on orthogonal sub-spaces of \mathbb{R}^k , the sum $\mathcal{X}^t(S, v) + m_i^t$ uniquely encodes the value of:

$$(\mathcal{X}^t(S, v), \{ \{ \mathcal{X}^t(s, v), e_i(S, v, S', v') \} \mid (S, v) \sim_{A_i} (S', v') \}). \quad (117)$$

Thus, we choose $\epsilon_1^t, \dots, \epsilon_4^t$ to be all zeroes. To compute the aggregation functions $\text{agg}_1^t, \dots, \text{agg}_4^t$ using these unique encodings, and to avoid repetition of the value $\mathcal{X}^t(S, v)$, we define auxiliary functions $\text{agg}_i^t : \mathbb{R}^k \rightarrow \mathbb{R}^{k_i}$ for $i = 1, \dots, 4$ as follows:

$$\text{agg}_1^t(\mathcal{X}^t(S, v) + m_1^t) = (\mathcal{X}^t(S, v), \text{agg}_1^t(\{ \{ \mathcal{X}^t(S, v'), e_1(S, v, S', v') \} \mid (S, v) \sim_{A_1} (S', v') \})) \quad (118)$$

and for $i > 1$:

$$\text{agg}_i^t(\mathcal{X}^t(S, v) + m_i^t) = \text{agg}_i^t(\{ \{ \mathcal{X}^t(s, v), e_i(S, v, S', v') \} \mid (S, v) \sim_{A_i} (S', v') \}). \quad (119)$$

Here, since we avoided repeating the value of $\mathcal{X}^t(S, v)$ by only adding it to the output of $\text{agg}_1^t(\cdot)$, the expression:

$$(\text{agg}_1^t(\mathcal{X}^t(S, v) + m_1^t), \dots, \text{agg}_4^t(\mathcal{X}^t(S, v) + m_4^t)) \quad (120)$$

is exactly equal to the input of f^t . In addition, since the function agg_i^t outputs one-hot encoded vectors, and the vector $\mathcal{X}^t(S, v)$ is one-hot encoded, the output of agg_i^t is always within the set $\{0, 1\}^{k_i}$. Now for any input vector $X \in \mathbb{R}^k$ define:

$$V_1^t(X) = (\text{agg}_1^t(X), 0_{k_2}, 0_{k_3}, 0_{k_4}). \quad (121)$$

$$V_2^t(X) = (0_{k_1}, \text{agg}_2^t(X), 0_{k_3}, 0_{k_4}). \quad (122)$$

$$V_3^t(X) = (0_{k_1}, 0_{k_2}, \text{agg}_3^t(X), 0_{k_4}). \quad (123)$$

$$V_4^t(X) = (0_{k_1}, 0_{k_2}, 0_{k_3}, \text{agg}_4^t(X)). \quad (124)$$

We note that since the output of agg_i^t is always within the set $\{0, 1\}^{k_i}$, the outputs of V_i^t is always within $\{0, 1\}^{k_1 + \dots + k_4}$. Now for $i = 1, \dots, 4$, employing theorem G.1 we define a dataset $\{x_j, y_j\}_{j=1}^N$ by taking the x_j s as all possible (distinct) values of $\mathcal{X}^t(S, v) + m_i^t$, with each corresponding y_j being the output $V_i^t(x_j)$. We note that there are finitely many such values as both $\mathcal{X}^t(S, v)$ and m_i^t are one-hot encoded vectors. The theorem now tells us that there exists a 4-layer fully connected ReLU neural network capable of implementing the function $V_i^t(\cdot)$. We choose U_i^t to be this network. Equations 121 - 124 now give us:

$$\sum_{i=1}^4 \mathcal{X}_i^t(S, v) = (\text{agg}_1^t(\mathcal{X}^t(S, v) + m_1^t), \dots, \text{agg}_4^t(\mathcal{X}^t(S, v) + m_4^t)). \quad (125)$$

which as stated before, is exactly the input to f^t . This proves step 2.

Step 3: We employ theorem G.1 for one final time, defining a dataset $\{x_j, y_j\}_{j=1}^N$ by taking the x_j s as all possible (distinct) values of:

$$\sum_{i=1}^4 \mathcal{X}_i^t(S, v)$$

(which we showed is a unique encoding to the input of $f^t(\cdot)$), with each corresponding y_j being the output $f^t(x_j)$. We note that Given the finite nature of our graph set, there are finitely many such values. Recalling that $f^t(\cdot)$ outputs one-hot encoded vectors, The theorem now tells us that there exists a 4-layer fully connected ReLU neural network capable of implementing the function $f^t(\cdot)$. We choose U_{fin}^t to be this network. This completes the proof. \square

G.2 Proofs of Appendix C

Proposition C.1 (Equal Expressivity of Node Marking Policies). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$\text{CS-GNN}(\mathcal{T}, \pi_S) = \text{CS-GNN}(\mathcal{T}, \pi_{SS}) = \text{CS-GNN}(\mathcal{T}, \pi_{MD}). \quad (36)$$

Proof. Let $\Pi = \{\pi_S, \pi_{SS}, \pi_{MD}\}$ be the set of all relevant node initialization policies, and assume for simplicity that our input graphs have no node features (the proof can be easily adjusted to account for the general case). For each $\pi \in \Pi$, let $\mathcal{X}^\pi(S, v)$ denote the node feature map induced by general node marking policy π , as per Definition C.1. We notice it is enough to prove for each $\pi_1, \pi_2 \in \Pi$ that $\mathcal{X}^{\pi_1}(S, v)$ can be implemented by updating $\mathcal{X}^{\pi_2}(S, v)$ using a stack of T layers of type 54. Thus, we prove the following four cases:

- Node + Size Marking \Rightarrow Simple Node Marking.
- Minimum Distance \Rightarrow Simple Node Marking.
- Simple Node Marking \Rightarrow Node + Size Marking.
- Simple Node Marking \Rightarrow Minimum Distance.

Node + Size Marking \Rightarrow Simple Node Marking:

In this case, we aim to update the node feature map:

$$\mathcal{X}^0(S, v) = \mathcal{X}^{\pi_{SS}}(S, v) = \begin{cases} (1, |S|) & v \in S \\ (0, |S|) & v \notin S. \end{cases} \quad (126)$$

We notice that:

$$\mathcal{X}^0(S, v) = \langle (1, 0), \mathcal{X}^{\pi_S}(S, v) \rangle, \quad (127)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard inner product in \mathbb{R}^2 . Using a CS-GNN update as per equation 15, with the update function:

$$f^1(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \cdot) = \langle (1, 0), \mathcal{X}^0(S, v) \rangle, \quad (128)$$

where $f(a, \cdot, \cdot, \cdot, \cdot)$ indicates that the function f depends solely on the parameter a , we obtain:

$$\mathcal{X}^1(S, v) = f^1(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \cdot) = \mathcal{X}^{\pi_S}(S, v). \quad (129)$$

This implies that for any coarsening function $\mathcal{T}(\cdot)$, the following holds:

$$\text{CS-GNN}(\mathcal{T}, \pi_S) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_{SS}). \quad (130)$$

Minimum Distance \Rightarrow Simple Node Marking:

In this case, we aim to update the node feature map:

$$\mathcal{X}^0(S, v) = \mathcal{X}^{\pi_{MD}}(S, v) = \min_{u \in S} d_G(u, v) \quad (131)$$

We notice that:

$$\mathcal{X}^S(S, v) = g(\mathcal{X}^0(S, v)) \quad (132)$$

where $g : \mathbb{R} \rightarrow \mathbb{R}$ is any continuous function such that:

1. $g(x) = 1 \forall x > \frac{1}{2}$,
2. $g(x) = 0 \forall x < \frac{1}{4}$.

Using a CS-GNN update as per equation 15, with the update function:

$$f^1(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \cdot) = g(\mathcal{X}^0(S, v)), \quad (133)$$

we obtain:

$$\mathcal{X}^1(S, v) = f^1(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \cdot) = \mathcal{X}^{\pi_S}(S, v). \quad (134)$$

This implies that for any coarsening function $\mathcal{T}(\cdot)$ the following holds:

$$\text{CS-GNN}(\mathcal{T}, \pi_S) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_{\text{MD}}). \quad (135)$$

Simple Node Marking \Rightarrow Node + Size Marking: In this case, we aim to update the node feature map:

$$\mathcal{X}^0(S, v) = \mathcal{X}^{\pi_S}(S, v) = \begin{cases} 1 & v \in S \\ 0 & v \notin S. \end{cases} \quad (136)$$

We notice that:

$$\sum_{v' \in S} \mathcal{X}^0(S, v') = |S|. \quad (137)$$

Using a CS-GNN update as per Equation (15), with aggregation function:

$$\text{agg}_4^l \{ \{ (\mathcal{X}^0(S, v'), z(S, v, S, v')) \mid v' \in S \} \} = \sum_{v' \in S} \mathcal{X}^0(S, v'), \quad (138)$$

and update function:

$$f^1 \left(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \sum_{v' \in S} \mathcal{X}^0(S, v') \right) = \left(\mathcal{X}^0(S, v), \sum_{v' \in S} \mathcal{X}^0(S, v') \right), \quad (139)$$

we obtain:

$$\mathcal{X}^1(S, v) = f^1 \left(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \sum_{v' \in S} \mathcal{X}^0(S, v') \right) = \mathcal{X}^{\pi_{SS}}(S, v). \quad (140)$$

This implies that for any coarsening function $\mathcal{T}(\cdot)$ the following holds:

$$\text{CS-GNN}(\mathcal{T}, \pi_{SS}) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_S). \quad (141)$$

Simple Node Marking \Rightarrow Minimum Distance:

In this case, we aim to update the node feature map:

$$\mathcal{X}^0(S, v) = \mathcal{X}^{\pi_S}(S, v) = \begin{cases} 1 & v \in S \\ 0 & v \notin S. \end{cases} \quad (142)$$

We shall prove that $\mathcal{X}^{\pi_{\text{MD}}}$ can be expressed by updating $\mathcal{X}^0(S, v)$ with a stack of CS-GNN layers. We do this by inductively showing that this procedure can express the following auxiliary node feature maps:

$$\mathcal{X}_*^t(S, v) = \begin{cases} \min_{v' \in S} d_G(v, v') + 1 & \min_{v' \in S} d_G(v, v') \leq t \\ 0 & \text{otherwise.} \end{cases} \quad (143)$$

We notice first that:

$$\mathcal{X}^0(S, v) = \mathcal{X}_*^0(S, v). \quad (144)$$

Now for the induction step, assume that there exists a stack of t CS-GNN layers such that:

$$\mathcal{X}^t(S, v) = \mathcal{X}_*^t(S, v). \quad (145)$$

We observe that equation:

$$\min_{v' \in S} d_G(v, v') = t + 1 \quad (146)$$

holds if and only if the following two conditions are met:

$$\min_{v' \in S} d_G(v, v') > t \quad (147)$$

$$\exists u \in N_G(v) \text{ s.t. } \min_{u' \in S} d_G(u, u') = t. \quad (148)$$

Equations 143 imply:

$$\min_{v' \in S} d_G(v, v') > t \Leftrightarrow \mathcal{X}^t(S, v) = 0. \quad (149)$$

In addition, since the node feature map $\mathcal{X}^t = \mathcal{X}_*^t$ is bounded by $t + 1$, Equation (143) implies:

$$\exists u \in N_G(v) \text{ s.t. } \min_{u' \in S} d_G(u, u') = t \Leftrightarrow \max\{\mathcal{X}^t(s, u) \mid v \sim_G u\} = t + 1. \quad (150)$$

Now, let $g_t : \mathbb{R}^2 \rightarrow \mathbb{R}$ be any continuous function such that for every pair of natural numbers $a, b \in \mathbb{N}$:

1. $g_t(a, b) = t + 2$ if $a = 0, b = t + 1$,
2. $g_t(a, b) = a$ otherwise.

Equations 146 - 150 imply:

$$\mathcal{X}_*^{t+1}(S, v) = g_t(\mathcal{X}^t(S, v), \max\{\mathcal{X}^t(s, u) \mid v \sim_G u\}). \quad (151)$$

Using a CS-GNN update as per Equation (15), with aggregation function:

$$\text{agg}_1^t \{ \{ \mathcal{X}^t(S, v'), e_{v, v'} \mid v' \sim_G v \} \} = \max_{v' \sim_G v} \mathcal{X}^t(S, v'). \quad (152)$$

and update function:

$$f^t(\mathcal{X}^t(S, v), \max_{v' \sim_G v} \mathcal{X}^t(S, v'), \cdot, \cdot, \cdot) = g_t(\mathcal{X}^t(S, v), \max_{v' \sim_G v} \mathcal{X}^t(S, v')) \quad (153)$$

we obtain:

$$\mathcal{X}^{t+1}(S, v) = f^t(\mathcal{X}^t(S, v), \max_{v' \sim_G v} \mathcal{X}^t(S, v'), \cdot, \cdot, \cdot) = \mathcal{X}_*^{t+1}(S, v). \quad (154)$$

This completes the induction step. Now, let \mathcal{G} be a finite family of graphs, whose maximal vertex size is n . We notice that:

$$\mathcal{X}^{\pi_{\text{MD}}}(S, v) = \mathcal{X}_*^n(S, v) - 1, \quad (155)$$

Which implies that there exists a stack of n CS-GNN layers such that:

$$\mathcal{X}^0(S, v) = \mathcal{X}^{\pi_{\text{S}}}(S, v) \quad \text{and} \quad \mathcal{X}^n(S, v) = \mathcal{X}^{\pi_{\text{MD}}}(S, v). \quad (156)$$

This implies:

$$\text{CS-GNN}(\mathcal{T}, \pi_{\text{MD}}) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_{\text{S}}). \quad (157)$$

This concludes the proof. \square




Figure 11: Graphs G and H defined in the proof of Proposition C.2. In each graph, the circle marks the single super-node induced by \mathcal{T} , while the number next to each node u is the maximal SPD between u and the nodes that compose the super-node.

Proposition C.2 (Expressivity of Learned Distance Policy). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$\text{CS-GNN}(\mathcal{T}, \pi_S) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_{LD}). \quad (37)$$

In addition, for some choices of $\mathcal{T}(\cdot)$ the containment is strict.

Proof. First, since we are concerned with input graphs belonging to a finite graph family \mathcal{G} , the learned function $\phi(\cdot)$ implemented by an MLP can express any continuous function on \mathcal{G} . This follows from Theorem G.1 (see the proof of Proposition B.1 for details). By choosing $\phi = \min(\cdot)$ in equation 35, it is clear that for any coarsening function $\mathcal{T}(\cdot)$ we have:

$$\text{CS-GNN}(\mathcal{T}, \pi_S) = \text{CS-GNN}(\mathcal{T}, \pi_{MD}) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_{LD}). \quad (158)$$

We now construct a coarsening function $\mathcal{T}(\cdot)$ along with two graphs, G and H , and demonstrate that there exists a function in $\text{CS-GNN}(\mathcal{T}, \pi_{LD})$ that can separate G and H . However, every function in $\text{CS-GNN}(\mathcal{T}, \pi_S)$ cannot separate the two.

For an input graph $G = (V, E)$ define:

$$\mathcal{T}(G) = \{\{u \in V \mid \deg_G(u) = 3\}\}. \quad (159)$$

i.e., $\mathcal{T}(\cdot)$ returns a single super-node composed of all nodes with degree 3. Now, define $G = (V_G, E_G)$ as the graph obtained by connecting two cycles of size four by adding an edge between a single node from each cycle. Additionally, define $H = (V_H, E_H)$ as the graph formed by joining two cycles of size five along one of their edges. See Figure 11 for an illustration of the two graphs. By choosing $\phi = \max(\cdot)$ in equation 35 a quick calculation shows that:

$$\sum_{S \in \mathcal{V}_{\mathcal{T}(G)}} \sum_{v \in V_G} \mathcal{X}^{\pi_{LD}}(S, v) = 16, \quad (160)$$

while:

$$\sum_{S \in \mathcal{V}_{\mathcal{T}(H)}} \sum_{v \in V_H} \mathcal{X}^{\pi_{LD}}(S, v) = 14. \quad (161)$$

Refer to Figure 11 for more details. Observe that:

$$f(G) = \sum_{s \in \mathcal{V}_{\mathcal{T}(H)}} \sum_{u \in V_H} \mathcal{X}^{\pi_{LD}}(S, v) \in \text{CS-GNN}(\mathcal{T}, \pi_{LD}) \quad (162)$$

Thus it is enough to show that:

$$f(G) = f(H), \quad \forall f \in \text{CS-GNN}(\mathcal{T}, \pi_S). \quad (163)$$

To achieve this, we use the layer update as per Definition B.2, which was demonstrated in Proposition B.1 to be equivalent to the general equivariant message passing update in Definition A.5. First, we observe that the graphs G and H are WL-indistinguishable. We then observe that since $|V^{\mathcal{T}}| = 1$, the graphs induced by the adjacency matrices A_G and A_H in Definition B.1 are isomorphic to the original graphs G and H , respectively, and therefore they are also WL-indistinguishable. Additionally, we notice that the graphs induced by the adjacency matrices $A_{\mathcal{T}(G)}$ and $A_{\mathcal{T}(H)}$ in Definition B.1 are both isomorphic to the fully disconnected graph with 8 nodes, making them WL-indistinguishable as well. We also observe that there exists a bijection $\sigma : V_G \rightarrow V_H$ that maps all nodes of degree 3 in G to all nodes of degree 3 in H . The definition of $\mathcal{T}(\cdot)$ implies that σ is an isomorphism between the adjacency matrices A_{P_i} corresponding to G and H , where $i = 1, 2$. Finally, we notice that for both G , and H , the node feature map induced by π_S satisfies:

$$\mathcal{X}^{\pi_S}(S, v) = \deg(v) - 2. \quad (164)$$

This node feature map can be easily implemented by the layer update in definition B.2 and so it can be ignored. Since all four graphs corresponding to G that are induced by the adjacency matrices in Definition B.1, are WL-indistinguishable from their counterpart corresponding to H , and equation 29 in definition B.2 is an MPNN update, which is incapable of distinguishing graphs that are WL-indistinguishable, we see that equation 163 holds, concluding the proof. \square

Proposition C.3 (Node + Size Marking as Invariant Marking). *Given a graph $G = (V, E)$ with node feature vector $X \in \mathbb{R}^{n \times d}$, and a coarsening function $\mathcal{T}(\cdot)$, let $\mathcal{X}^{\pi_{SS}}, \mathcal{X}^{\pi_{inv}}$ be the node feature maps induced by π_{SS} and π_{inv} respectively. Recall that:*

$$\mathcal{X}^{\pi_{SS}}(S, v) = [X_v, b_{\pi_{SS}}(S, v)], \quad (43)$$

$$\mathcal{X}^{\pi_{inv}}(S, v) = [X_v, b_{\pi_{inv}}(S, v)]. \quad (44)$$

The following now holds:

$$b_{\pi_{inv}}(S, v) = OHE(b_{\pi_{SS}}(S, v)) \quad \forall S \in V^{\mathcal{T}}, \forall v \in V. \quad (45)$$

Here, *OHE* denotes a one-hot encoder, independent of the choice of both G and \mathcal{T} .

Proof. Let $G = (V, E)$ be a graph with $V = [n]$, and let $\mathcal{T}(\cdot)$ be a coarsening function. Recall that the maps $b_{\pi_{SS}}(\cdot, \cdot)$ and $b_{\pi_{inv}}(\cdot, \cdot)$ are both independent of the connectivity of G and are defined as follows:

$$b_{\pi_{SS}}(S, v) = \begin{cases} (1, |S|) & v \in S, \\ (0, |S|) & v \notin S. \end{cases} \quad (165)$$

$$b_{\pi_{inv}}(S, v) = [\mathbf{1}_{\gamma_1}(S, v), \dots, \mathbf{1}_{\gamma_k}(S, v)]. \quad (166)$$

Here, $v \in [n]$, $S \in \mathcal{T}([n]) \subseteq \mathcal{P}([n])$, $\gamma_1, \dots, \gamma_k$ is any enumeration of the set of all orbits $(\mathcal{P}([n]) \times [n])/S_n$, and $\mathbf{1}_{\gamma_i}$ denotes the indicator function of orbit γ_i . Since any tuple $(S, v) \in \mathcal{P}([n]) \times [n]$ belongs to exactly one orbit γ_i , we note that the right hand side of Equation (166) is a one-hot encoded vector. Thus, it suffices to show that for every $v, v' \in [n]$ and $S, S' \in \mathcal{P}([n])$, we have:

$$b_{\pi_{SS}}(S, v) = b_{\pi_{SS}}(S', v') \Leftrightarrow b_{\pi_{inv}}(S, v) = b_{\pi_{inv}}(S', v'). \quad (167)$$

This is equivalent to:

$$(\mathcal{P}([n]) \times [n])/S_n = \{(S, v) \mid |S| = i, \mathbf{1}_S(v) = j \mid i \in [n], j \in \{0, 1\}\}. \quad (168)$$

Essentially, this means that each orbit corresponds to a choice of the size of s and whether $v \in S$ or not. To conclude the proof, it remains to show that for any two pairs $(S, v), (S', v') \in \mathcal{P}([n]) \times [n]$, there exists a permutation $\sigma \in S_n$ such that:

$$\sigma \cdot (S, v) = (S', v') \quad (169)$$

if and only if

$$|S| = |S'| \text{ and } \mathbf{1}_S(v) = \mathbf{1}_{S'}(v'). \quad (170)$$

Assume first that $\sigma \cdot (S, v) = (S', v')$, then $\sigma^{-1}(S) = S'$ and since σ is a bijection, $|S| = |S'|$. In addition $\sigma^{-1}(v) = v'$ thus:

$$v \in S \Leftrightarrow v' = \sigma^{-1}(v) \in \sigma^{-1}(S) = S'. \quad (171)$$

Assume now that:

$$|S| = |S'| \quad (172)$$

$$\mathbf{1}_S(v) = \mathbf{1}_{S'}(v') \quad (173)$$

It follows that for some $r, m \in [n]$:

$$|S \setminus \{v\}| = |S' \setminus \{v'\}| = r \quad \text{and} \quad |[n] \setminus (S \cup \{v\})| = |[n] \setminus (S' \cup \{v'\})| = m \quad (174)$$

Write:

$$\begin{aligned} S \setminus \{v\} &= \{i_1, \dots, i_r\}, & S' \setminus \{v'\} &= \{i'_1, \dots, i'_r\}, \\ [n] \setminus (S \cup \{v\}) &= \{j_1, \dots, j_m\}, & [n] \setminus (S' \cup \{v'\}) &= \{j'_1, \dots, j'_m\} \end{aligned}$$

and define:

$$\sigma(x) = \begin{cases} v' & x = v \\ i'_l & x = i_l, l \in [r] \\ j'_l & x = j_l, l \in [m] \end{cases} \quad (175)$$

We now have:

$$\sigma \cdot (S, v) = (S', v'). \quad (176)$$

This concludes the proof. \square

G.3 Proofs of Appendix D.1

Proposition D.1 (CS-GNN Can Implement MSGNN). *Let $\mathcal{T}(\cdot)$ be the identity coarsening function defined by:*

$$\mathcal{T}(G) = \{\{v\} \mid v \in V\} \quad \forall G = (V, E). \quad (49)$$

The following holds:

$$\text{CS-GNN}(\mathcal{T}, \pi_S) = \text{MSGNN}(\pi_{NM}). \quad (50)$$

Proof. Abusing notation, for a given graph $G = (V, E)$ we write $\mathcal{T}(G) = G$, $V^{\mathcal{T}} = V$. First, we observe that:

$$v \in \{u\} \Leftrightarrow u = v, \quad (177)$$

This implies that the initial node feature map $\mathcal{X}^0(u, v)$ induced by π_S is equivalent to the standard node marking described in equation 46. Additionally, we note that the pooling procedures for both models, as described in equations 16 and 55, are identical. Therefore, it is sufficient to show that the CS-GNN and MSGNN layer updates described in equations 15 and 47 respectively are also identical. For this purpose, let $\mathcal{X}^t(v, u)$ be a node feature map supported on the set $V \times V$. The inputs to the MSGNN layer are the following:

1. $\mathcal{X}^t(u, v)$.
2. $\mathcal{X}^t(u, u)$.
3. $\mathcal{X}^t(v, v)$.
4. $\text{agg}_1^t \{\{\mathcal{X}^t(u, v'), e_{v, v'} \mid v' \sim v\}\}$.
5. $\text{agg}_2^t \{\{\mathcal{X}^t(u', v), e_{u, u'} \mid u' \sim u\}\}$.

The inputs to the CS-GNN layer are the following:

1. $\mathcal{X}^t(S, v) \Rightarrow \mathcal{X}^t(u, v)$.
2. $\text{agg}_1^t \{\{\mathcal{X}^t(S, v'), e_{v, v'} \mid v' \sim_G v\}\} \Rightarrow \text{agg}_1^t \{\{\mathcal{X}^t(u, v'), e_{v, v'} \mid v' \sim v\}\}$.
3. $\text{agg}_2^t \{\{\mathcal{X}^t(S', v), \tilde{e}_{S, S'} \mid S' \sim_{G^{\mathcal{T}}} S\}\} \Rightarrow \text{agg}_2^t \{\{\mathcal{X}^t(u, u'), e_{u, u'} \mid v' \sim v\}\}$.
4. $\text{agg}_3^t \{\{\mathcal{X}^t(S', v), z(S, v, S', v) \mid \forall s' \in V^{\mathcal{T}} \text{ s.t. } v \in S'\}\} \Rightarrow \{\{X^t(v, v), z(u, v, v, v)\}\}$.
5. $\text{agg}_4^t \{\{\mathcal{X}^t(S, v'), z(S, v, S, v') \mid \forall u' \in V \text{ s.t. } v' \in S\}\} \Rightarrow \{\{X^t(u, u), z(u, v, u, u)\}\}$.

The terms $z(u, v, v, v)$ and $z(u, v, u, u)$ appearing in the last two input terms of the CS-GNN layer uniquely encode the orbit tuples (u, v, v, v) and (u, v, u, u) belong to respectively. Since these orbits depend solely on whether $u = v$, these values are equivalent to the node marking feature map $\mathcal{X}^0(u, v)$. Therefore, these terms can be ignored. Observing the two lists above, we see that the inputs to both update layers are identical (ignoring the $z(\cdot)$ terms). Thus, as both updates act on these inputs in the same way, the updates themselves are identical. and so

$$\text{MSGNN}(\pi_{NM}) = \text{CS-GNN}(\mathcal{T}, \pi_S). \quad (178)$$

□

G.4 Proofs of Appendix D.2

Proposition D.2 (CS-GNN Is at Least as Expressive as Coarse MPNN). *For any coarsening function $\mathcal{T}(\cdot)$ the following holds:*

$$\text{MPNN} \subseteq \text{MPNN}_+(\mathcal{T}) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_S) \quad (58)$$

Proof. For convenience, let us first restate the CS-GNN layer update:

$$\begin{aligned}
\mathcal{X}^{t+1}(S, v) = f^t \Big(& \mathcal{X}^t(S, v), \\
& \text{agg}_1^t \{ \{ \mathcal{X}^t(S, v'), e_{v, v'} \} \mid v' \sim_G v \}, \\
& \text{agg}_2^t \{ \{ \mathcal{X}^t(S', v), \tilde{e}_{S, S'} \} \mid S' \sim_{G^\mathcal{T}} S \}, \\
& \text{agg}_3^t \{ \{ \mathcal{X}^t(S', v), z(S, v, S', v) \} \mid S' \in V^\mathcal{T} \text{ s.t. } v \in S' \}, \\
& \text{agg}_4^t \{ \{ \mathcal{X}^t(S, v'), z(S, v, S, v') \} \mid u' \in V \text{ s.t. } v' \in S \} \Big),
\end{aligned} \tag{15}$$

as well as the MPNN₊ layer update:

$$\begin{aligned}
\text{For } v \in V : \quad & \mathcal{X}^{t+1}(v) = f_V^t \left(\mathcal{X}^t(v), \text{agg}_1^t \{ \{ \mathcal{X}^t(v'), e_{v, v'} \} \mid v \sim_G v' \}, \right. \\
& \left. \text{agg}_2^t \{ \{ \mathcal{X}^t(S) \mid S \in V^\mathcal{T}, v \in S \} \}, \right) \\
\text{For } S \in V^\mathcal{T} : \quad & \mathcal{X}^{t+1}(S) = f_{V^\mathcal{T}}^t \left(\mathcal{X}^t(S), \text{agg}_1^t \{ \{ \mathcal{X}^t(S'), e_{S, S'} \} \mid S \sim_{G^\mathcal{T}} S' \}, \right. \\
& \left. \text{agg}_2^t \{ \{ \mathcal{X}^t(v) \mid v \in V, v \in S \} \} \right).
\end{aligned} \tag{54}$$

We note that by setting $f_{V^\mathcal{T}}^t$ to be a constant zero and choosing f_V^t to be any continuous function that depends only on its first two arguments, the update in equation 54 becomes a standard MPNN layer. This proves:

$$\text{MPNN} \subseteq \text{MPNN}_+(T). \tag{179}$$

Next, we prove the following 2 Lemmas:

Lemma G.1. *Given a graph $G = (V, E)$ such that $V = [n]$ with node feature vector $X \in \mathbb{R}^{n \times d}$, and a coarsening function $\mathcal{T}(\cdot)$, there exists a CS-GNN(\mathcal{T}, π_S) layer such that:*

$$\mathcal{X}^1(S, v) = [0_{d+1}, X_v, 1] = [\tilde{\mathcal{X}}^0(S), \tilde{\mathcal{X}}^0(v)]. \tag{180}$$

Here $[\cdot, \cdot]$ denotes concatenation and $\tilde{\mathcal{X}}^0(\cdot)$ denotes the initial node feature map of the coarsened sum graph $G_+^\mathcal{T}$.

Lemma G.2. *Let $\tilde{\mathcal{X}}^t(\cdot)$ denote the node feature maps of $G_+^\mathcal{T}$ at layers t of a stack of MPNN₊(\mathcal{T}) layers. There exists a stack of $t + 1$ CS-GNN(\mathcal{T}, π_S) layers such that:*

$$\mathcal{X}^{t+1}(S, v) = [\tilde{\mathcal{X}}^t(S), \tilde{\mathcal{X}}^t(v)]. \tag{181}$$

proof of Lemma G.1. Recall that the initial node feature map of CS-GNN(\mathcal{T}, π_S) is given by:

$$\mathcal{X}^0(S, v) = \begin{cases} [X_v, 1] & v \in S \\ [X_v, 0] & v \notin S. \end{cases} \tag{182}$$

In addition, the initial node feature map of MPNN₊(\mathcal{T}) is given by:

$$\tilde{\mathcal{X}}^0(v) = \begin{cases} [X_v, 1] & v \in V \\ 0_{d+1} & v \in V^\mathcal{T}. \end{cases} \tag{183}$$

Thus, we choose a layer update as described in equation 15 with:

$$\mathcal{X}^1(S, v) = f^0(\mathcal{X}^0(S, v), \cdot, \cdot, \cdot, \cdot) = [0_{d+1}, \mathcal{X}^0(S, v)_{1:d}, 1] \tag{184}$$

Here, $f(a, \cdot, \cdot, \cdot, \cdot)$ denotes that the function depends only on the parameter a , and $X_{a:b}$ indicates that only the coordinates a through b of the vector X are taken. This gives us:

$$\mathcal{X}^1(S, v) = [\tilde{\mathcal{X}}^0(S), \tilde{\mathcal{X}}^0(v)]. \tag{185}$$

□

proof of Lemma G.2. We prove this Lemma by induction on t . We note that Lemma G.1 provides the base case $t = 0$. Assume now that for a given stack of $t + 1$ $\text{MPNN}_+(\mathcal{T})$ layer updates, with corresponding node feature maps:

$$\tilde{\mathcal{X}}^i : V_+^{\mathcal{T}} \rightarrow \mathbb{R}^{d_i} \quad i = 1 \dots, t + 1, \quad (186)$$

there exists a stack of $t + 1$ $\text{CS-GNN}(\mathcal{T}, \pi_S)$ layers with node feature maps:

$$\mathcal{X}^i : V^{\mathcal{T}} \times V \rightarrow \mathbb{R}^{2d_i} \quad i = 1, \dots, t + 1, \quad (187)$$

such that:

$$\mathcal{X}^{t+1}(S, v) = [\tilde{\mathcal{X}}^t(S), \tilde{\mathcal{X}}^t(v)]. \quad (188)$$

We shall show that there exists a single additional $\text{CS-GNN}(\mathcal{T}, \pi_S)$ layer update such that:

$$\mathcal{X}^{t+2}(S, v) = [\tilde{\mathcal{X}}^{t+1}(S), \tilde{\mathcal{X}}^{t+1}(v)]. \quad (189)$$

For that purpose we define the following $\text{CS-GNN}(\mathcal{T}, \pi_S)$ update (abusing notation, the left hand side refers to components of the $\text{CS-GNN}(\mathcal{T}, \pi_S)$ update at layer $t + 1$, while the right hand side refers to components of the $\text{MPNN}_+(\mathcal{T})$ update at layer t):

$$\begin{aligned} \text{agg}_1^{t+1} &= \text{agg}_1^t|_{1:d_t}, \\ \text{agg}_2^{t+1} &= \text{agg}_1^t|_{d_t+1:2d_t}, \\ \text{agg}_3^{t+1} &= \text{agg}_2^t|_{1:d_t}, \\ \text{agg}_4^{t+1} &= \text{agg}_2^t|_{d_t+1:2d_t}, \end{aligned} \quad (190)$$

$$f^{t+1}(a, b, c, d, e) = [f_V^t(a_{1:d_t}, b, d), f_{V^{\mathcal{T}}}^t(a_{d_t+1:2d_t}, c, e)]. \quad (191)$$

Here the operation $\text{agg}_{a:b}^t$ initially projects all vectors in the input multi-set onto coordinates a through b , and subsequently passes them to the function agg . equations 190, 191 guarantee that:

$$\begin{aligned} \mathcal{X}^{t+2}(S, v)_{1:d_{t+1}} &= f_V^t(\mathcal{X}^t(S, v)_{1:d_t}, \\ &\quad \text{agg}_1^t\{(S, v')_{1:d_t} \mid v \sim_G v'\}, \\ &\quad \text{agg}_2^t\{(S', v)_{1:d_t} \mid v \in S'\}) \\ &= \tilde{\mathcal{X}}^{t+1}(v), \\ \mathcal{X}^{t+2}(S, v)_{d_{t+1}+1:2d_{t+1}} &= f_{V^{\mathcal{T}}}^t(\mathcal{X}^t(S, v)_{d_t+1:2d_t}, \\ &\quad \text{agg}_1^t\{(S', v)_{d_t+1:2d_t} \mid S' \sim_{\mathcal{T}(G)} S\}, \\ &\quad \text{agg}_2^t\{(S, v')_{d_t+1:2d_t} \mid v' \in S\}) \\ &= \tilde{\mathcal{X}}^{t+1}(S). \end{aligned} \quad (192)$$

This proves the Lemma. \square

Now, for a given finite family of graphs \mathcal{G} and a function $f \in \text{MPNN}_+(\mathcal{T})$, there exists a stack of T $\text{MPNN}_+(\mathcal{T})$ layers such that:

$$f(G) = U \left(\sum_{v \in V_+^{\mathcal{T}}} \tilde{\mathcal{X}}^T(v) \right) \quad \forall G \in \mathcal{G}. \quad (193)$$

Here, $\tilde{\mathcal{X}}^T : V_+^{\mathcal{T}} \rightarrow \mathbb{R}^{d_T}$ denotes the final node feature map, and U is an MLP. Lemma G.2 now tells us that there exists a stack of $T + 1$ $\text{CS-GNN}(\mathcal{T}, \pi_S)$ layers such that:

$$\mathcal{X}^{T+1}(S, v) = [\tilde{\mathcal{X}}^T(S), \tilde{\mathcal{X}}^T(v)]. \quad (194)$$

Similarly to Lemma G.1, we use one additional layer to pad $\mathcal{X}^{T+1}(S, v)$ as follows:

$$\mathcal{X}^{T+2}(S, v) = [\tilde{\mathcal{X}}^T(S), \tilde{\mathcal{X}}^T(v), 1]. \quad (195)$$

We notice that:

$$\begin{aligned} \sum_{s \in V^{\mathcal{T}}} \mathcal{X}^{T+2}(S, v) &= \left[\sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(S), \sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(v), \sum_{S \in V^{\mathcal{T}}} 1 \right] \\ &= \left[\sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(S), |V^{\mathcal{T}}| \cdot \tilde{X}^T(v), |V^{\mathcal{T}}| \right]. \end{aligned} \quad (196)$$

Thus, in order to get rid of the $|V^{\mathcal{T}}|$ term, We define:

$$\text{MLP}_1(a, b, c) = \left[a, \frac{1}{c} \cdot b, 1 \right], \quad a, b \in \mathbb{R}^{d_L}, c > 0. \quad (197)$$

We note that since we are restricted to a finite family of input graphs, the use of an MLP in equation 200 can be justified using Theorem G.1 (see the proof of Proposition B.1 for a detailed explanation).

Equations 196 and 200 imply:

$$\text{MLP}_1 \left(\sum_{s \in V^{\mathcal{T}}} \mathcal{X}^{T+2}(S, v) \right) = \left[\sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(S), \tilde{X}^T(v), 1 \right] \quad (198)$$

Thus, similarly to equation 196:

$$\sum_{v \in V} \text{MLP}_1 \left(\sum_{S \in V^{\mathcal{T}}} \mathcal{X}^{T+2}(S, v) \right) = \left[|V| \cdot \sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(S), \sum_{v \in V} \tilde{X}^T(v), |V| \right] \quad (199)$$

And so, in order to get rid of the $|V|$ term, We define:

$$\text{MLP}_2(a, b, c) = U \left(a \cdot \frac{1}{c} + b, 1 \right), \quad a, b \in \mathbb{R}^{d_T}, c > 0. \quad (200)$$

Thus for all $G \in \mathcal{G}$:

$$\begin{aligned} &\text{MLP}_2 \left(\sum_{v \in V} \text{MLP}_1 \left(\sum_{S \in V^{\mathcal{T}}} \mathcal{X}^{T+2}(S, v) \right) \right) \\ &= \text{MLP}_2 \left(\left[|V| \cdot \sum_{S \in V^{\mathcal{T}}} \tilde{X}^T(S), \sum_{v \in V} \tilde{X}^T(v), |V| \right] \right) \\ &= U \left(\sum_{v \in V^{\mathcal{T}}} \tilde{X}^T(v) \right) \\ &= f(G). \end{aligned} \quad (201)$$

and so $f \in \text{CS-GNN}(\mathcal{T}, \pi_S)$. This proves:

$$\text{MPNN}_+(T) \subseteq \text{CS-GNN}(\mathcal{T}, \pi_S). \quad (202)$$

□

Proposition D.3 (CS-GNN Can Be More Expressive Than MPNN+). *Let $\mathcal{T}(\cdot)$ be the identity coarsening function defined by:*

$$\mathcal{T}(G) = \{\{v\} \mid v \in V\} \quad G = (V, E). \quad (59)$$

The following holds:

$$\text{MPNN} = \text{MPNN}_+(\mathcal{T}). \quad (60)$$

Thus:

$$\text{MPNN}_+(\mathcal{T}) \subset \text{CS-GNN}(\mathcal{T}, \pi_S), \quad (61)$$

where this containment is strict.

Proof. First, using the notation \tilde{v} to mark the single element set $\{v\} \in V^T$, We notice that the $\text{MPNN}_+(\mathcal{T})$ layer update described in equation 54, becomes:

$$\begin{aligned} \text{For } v \in V : \quad \mathcal{X}^{t+1}(v) &= f_V^t \left(\mathcal{X}^t(v), \mathcal{X}^t(\tilde{v}), \text{agg}^t \{ \{ \mathcal{X}^t(v'), e_{v,v'} \} \mid v' \sim_G v \} \right), \\ \text{For } \tilde{v} \in V^T : \quad \mathcal{X}^{t+1}(\tilde{v}) &= f_{V^T}^t \left(\mathcal{X}^t(\tilde{v}), \mathcal{X}^t(v), \text{agg}^t \{ \{ \mathcal{X}^t(\tilde{v}'), e_{\tilde{v},\tilde{v}'} \} \mid v \sim_G v' \} \right). \end{aligned} \quad (203)$$

Now, for a given finite family of graphs \mathcal{G} and a function $f \in \text{MPNN}_+(\mathcal{T})$, there exists a stack of T $\text{MPNN}_+(\mathcal{T})$ layers such that:

$$f(G) = U \left(\sum_{v \in V_+^T} \mathcal{X}^T(v) \right) \quad \forall G \in \mathcal{G}. \quad (204)$$

Here, $\mathcal{X}^T : V_+^T \rightarrow \mathbb{R}^d$ denotes the final node feature map, and U is an MPL. We now prove by induction on t that there exists a stack of t standard MPNN layers, with corresponding node feature map $X^t : V \rightarrow \mathbb{R}^{2d_t}$ such that :

$$X^t(v) = [\mathcal{X}^t(v), \mathcal{X}^t(\tilde{v})]. \quad (205)$$

Here, $[\cdot, \cdot]$ stands for concatenation. We assume for simplicity that the input graph G does not have node features, though the proof can be easily adapted for the more general case. We notice that for the base case $t = 0$, equation 53 in definition D.2 implies:

$$\mathcal{X}^0(v) = \begin{cases} 1 & v \in V, \\ 0 & v \in V^T. \end{cases} \quad (206)$$

Thus, we define:

$$X^0(v) = (1, 0). \quad (207)$$

This satisfies Equation (205), establishing the base case of the induction. Assume now that Equation (205) holds for some $t \in [T]$. Let $\text{agg}^t, f_V^t, f_{V^T}^t$ be the components of layer t , as in equation 203. We define:

$$\text{agg}^t = [\text{agg}^t|_{1:d_t}, \text{agg}^t|_{d_t+1:2d_t}]. \quad (208)$$

Here the operation $\text{agg}|_{a:b}$ initially projects all vectors in the input multi-set onto coordinates a through b , and subsequently passes them to the function agg .

Additionally, let d^* denote the dimension of the output of the function agg^t . We define:

$$\tilde{f}^t(a, b) = [f_V^t(a|_{1:d_t}, a|_{d_t+1:2d_t}, b|_{1:d^*}), f_{V^T}^t(a|_{d_t+1:2d_t}, a|_{1:d_t}, b|_{d^*+1:2d^*})]. \quad (209)$$

Finally, we update our node feature map X^t using a standard MPNN update according to:

$$X^{t+1}(v) = \tilde{f}^t(X^t(v), \{ \{ X^t(v'), e_{v,v'} \} \mid v' \sim_G v \}). \quad (210)$$

equations 203, 205 and 210 now guarantee that:

$$X^{t+1}(v) = [\mathcal{X}^t(v), \mathcal{X}^{t+1}(\tilde{v})]. \quad (211)$$

This concludes the inductive proof. We now define:

$$\text{MLP}(x) = U(x|_{1:d_T}) + U(x|_{d_T+1:2d_T}). \quad (212)$$

This gives us:

$$U \left(\sum_{v \in V_+^T} \mathcal{X}^T(v) \right) = \text{MLP} \left(\sum_{v \in V} X^T(v) \right) = f(G). \quad (213)$$

We have thus proven that $f \in \text{MPNN}$ and so:

$$\text{MPNN}_+(\mathcal{N}) \subseteq \text{MPNN}. \quad (214)$$

Combining this result with Proposition D.2, we obtain:

$$\text{MPNN} = \text{MPNN}_+(\mathcal{T}). \quad (215)$$

Finally, since Proposition D.1 tells us that $\text{CS-GNN}(\mathcal{T}, \pi_S)$ has the same implementation power as the maximally expressive node policy subgraph architecture MSGNN, which is proven to be strictly more expressive than the standard MPNN, we have:

$$\text{MPNN}_+(\mathcal{T}) \subset \text{CS-GNN}(\mathcal{T}, \pi_S). \quad (216)$$

□

Proposition D.4 (CS-GNN can be strictly more expressive than node-based subgraph GNNs). *Let \mathcal{T} be the coarsening function defined by:*

$$\mathcal{T}(G) = \{\{v\} \mid v \in V\} \cup E \quad G = (V, E). \quad (63)$$

The following holds:

1. Let G_1, G_2 be a pair of graphs such that there exists a node-based subgraph GNN model M where $M(G_1) \neq M(G_2)$. There exists a CS-GNN model M' which uses \mathcal{T} such that $M'(G_1) \neq M'(G_2)$.
2. There exists a pair of graphs G_1, G_2 such that for any subgraph GNN model M it holds that $M(G_1) = M(G_2)$, but there exists a CS-GNN model M' which uses \mathcal{T} such that $M'(G_1) \neq M'(G_2)$.

Proof. First, notice that the super-nodes produced by \mathcal{T} are either of size 1, in which case they correspond to nodes, or they are of size two, in which case they correspond to edges. Since an CS-GNN model processes feature maps $\mathcal{X}^t(S, v)$ where in the initial layer the set size of S is encoded in $\mathcal{X}^t(S, v)$, we can easily use the CS-GNNupdate in Definition A.5 to ignore all values of $\mathcal{X}^t(S, v)$ where $|S| = 2$ (This can be done by using $f^t, \text{agg}_1^t, \dots, \text{agg}_1^t$ in Definition A.5 to zero out these values at each update). This means CS-GNN using \mathcal{T} is able to simulate an CS-GNNupdate with the identity coarsening function, which was shown in Proposition D.1 to be as expressive as GNN-SSWL+ (Definition D.1) which is a maximally expressive node-based subgraph GNN, thus proving part (1) of the proposition. To prove part (2), notice that using the same reasoning as before, an CS-GNN model using \mathcal{T} as a coarsening function can implement an CS-GNN model using the edge coarsening function:

$$\mathcal{T}'(G) = E \quad G = (V, E). \quad (217)$$

An CS-GNN model with the identity coarsening function can be interpreted as a GNN-SSWL+ model. Similarly, an CS-GNN model using the edge coarsening function \mathcal{T}' generalizes the GNN-SSWL+ framework by extending it from node-based subgraph GNNs to edge-based subgraph GNNs. In fact, the same proof in [38, 12] showing that GNN-SSWL+ is at least as expressive as a DSS subgraph GNN using the node deletion policy (see [4] for a definition of the DSS subgraph GNN), can be used to show that CS-GNN using the edge coarsening function \mathcal{T}' is at least as expressive as a DSS subgraph GNN with an edge deletion policy. The latter model was shown in [4] to be able to separate a pair of 3-WL indistinguishable graphs. In contrast, node-based subgraph GNNs were shown in [12] to not be able to separate any pair of 3-WL indistinguishable graphs. Thus, there exists a pair of graphs which CS-GNN using \mathcal{T} can separate while node-based subgraph GNNs can't, proving part (2) of the proposition.

□

G.5 Proofs of Appendix E

Lemma E.1 ($\gamma(\Gamma)$ are orbits). *The sets $\{\gamma^{k^*} : k = 1, \dots, n; * \in \{+, -\}\}$ and $\{\Gamma^{\leftrightarrow; k_1; k_2; k^\cap; \delta_{\text{same}}; \delta_{\text{diff}}}\}$ are the orbits of S_n on the index space $(\mathcal{P}([n]) \times [n])$ and $(\mathcal{P}([n]) \times [n] \times (\mathcal{P}([n]) \times [n]))$, respectively.*

Proof. We will prove this lemma for γ . The proof for Γ follows similar reasoning; we also refer the reader to [22] for a general proof.

We will prove this lemma through the following three steps.

- (1). Given indices $(S, i) \in \mathcal{P}([n]) \times [n]$, there exists $\gamma \in (\mathcal{P}([n]) \times [n])_\sim$ such that $(S, i) \in \gamma$.
- (2). Given indices $(S, i) \in \gamma$, for any $\sigma \in S_n$, it holds that $(\sigma^{-1}(S), \sigma^{-1}(i)) \in \gamma$.

(3). Given $(S, i) \in \gamma$ and $(S', i') \in \gamma$ (the same γ), it holds that there exists a $\sigma \in S_n$ such that $\sigma \cdot (S, i) = (S', i')$.

We prove in what follows.

(1). Given indices $(S, i) \in \mathcal{P}([n]) \times [n]$, w.l.o.g. we assume that $|S| = k$, thus if $i \in S$ ($i \notin S$) it holds that $(S, i) \in \gamma^{k^-}$ ($(S, i) \in \gamma^{k^+}$), recall Equation (71).

(2). Given indices $(S, i) \in \gamma$, note that any permutation $\sigma \in S_n$ does not change the cardinality of S nor the inclusion (or exclusion) of i in S . Recalling Equation (71), we complete this step.

(3). Given that $(S, i) \in \gamma$ and $(S', i') \in \gamma$, and recalling Equation (71), we note that $|S| = |S'|$ and that either both $i \in S$ and $i' \in S'$, or both $i \notin S$ and $i' \notin S'$.

(3.1). In (3.1) we focus on the case where $i \notin S$ and $i' \notin S'$. Let $S = \{i_1, \dots, i_k\}$ and $S' = \{i'_1, \dots, i'_k\}$. Then, we have $(\{i_1, \dots, i_k\}, j)$ and $(\{i'_1, \dots, i'_k\}, j')$. Define $\sigma \in S_n$ such that $\sigma(i_l) = i'_l$ for $l \in [k]$, and $\sigma(j) = j'$. Since $(\{i_1, \dots, i_k\}, j)$ consists of $k + 1$ distinct indices and $(\{i'_1, \dots, i'_k\}, j')$ also consists of $k + 1$ distinct indices, this is a valid $\sigma \in S_n$.

(3.2). Here, we focus on the case where $i \in S$ and $i' \in S'$. This proof is similar to (3.1), but without considering the indices j and j' , as they are included in S and S' , respectively.

□

Proposition E.1 (Basis of Invariant (Equivariant) Layer). *The tensors \mathbf{B}^γ (\mathbf{B}^Γ) in Equation (72) (Equation (74)) form an orthogonal basis (in the standard inner product) to the solution of Equation (66) (Equation (67)).*

Proof. We prove this proposition for the invariant case. The equivariant case is proved similarly – we also refer the reader for [22] for a general proof. We will prove this in three steps,

(1). For any $\gamma \in (\mathcal{P}([n]) \times [n])_\sim$ it holds that $\mathbf{B}_{S,i}^\gamma$ solves Equation (66).

(2). Given a solution \mathbf{L} to Equation (66), it is a linear combination of the basis elements.

(3). We show that the basis vectors are orthogonal and thus linearly independent.

We prove in what follows.

(1). Given $\gamma \in (\mathcal{P}([n]) \times [n])_\sim$, we need to show that $\mathbf{B}_{S,i}^\gamma = \mathbf{B}_{\sigma^{-1}(S), \sigma^{-1}(i)}^\gamma$. Since any $\gamma \in (\mathcal{P}([n]) \times [n])_\sim$ is an orbit in the index space (recall Lemma E.1), and $\mathbf{B}_{S,i}^\gamma$ are indicator vectors of the orbits this always holds.

(2). Given a solution \mathbf{L} to Equation (66), it must hold that $\mathbf{L}_{S,i} = \mathbf{L}_{\sigma^{-1}(S), \sigma^{-1}(i)}$. Since the set $\{\gamma^{k^*} : k = 1, \dots, n; * \in \{+, -\}\}$ corresponds to the orbits in the index space with respect to S_n , \mathbf{L} should have the same values over the index space of these orbits. Let's define these values as α^γ for each $\gamma \in \{\gamma^{k^*} : k = 1, \dots, n; * \in \{+, -\}\}$. Thus, we obtain that $\mathbf{L}' = \sum_{\gamma \in (\mathcal{P}([n]) \times [n])_\sim} \alpha^\gamma \cdot \mathbf{B}^\gamma$, since \mathbf{B}^γ are simply indicator vectors of the orbits. This completes this step.

(3). Once again, since the basis elements are indicator vectors of disjoint orbits we obtain their orthogonality, and thus linearly independent. □

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and precede the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist".**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract spells out all the main contributions in the present paper, both theoretical and empirical ones. These are extensively discussed and recapitulated in the Introduction Section 1 (see paragraphs "Our approach" and "Contributions"). The scope of the paper is well defined in the first periods of the abstract and comprehensively articulated in the first two paragraphs of the Introduction Section 1.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please refer to paragraph “Limitations” in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Please refer to Appendices B to E and F.4, which include precise and contextualized statements of all theoretical results and derivations, and to Appendix G for proofs thereof.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Please refer to Appendix F.1 for a description of the employed datasets and splitting procedure, Appendices F.2 and F.3 for a list of experimental details and hyperparameter settings, and Appendix F.5 for a series of complementary results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code to reproduce our results can be found in the following GitHub repository: <https://github.com/BarSGuy/Efficient-Subgraph-GNNs>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please refer to Appendix F.1 for a description of the employed datasets and splitting procedure, Appendices F.2 and F.3 for a list of experimental details and hyperparameter settings, including the utilized training procedures. The results for baselines approaches are reported according to what stated in Section 5 and Appendix F.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Our results in Section 5 are reported in terms of mean and standard deviation calculated over different model initializations (i.e., by setting different random seeds prior to code execution). Table 9 reports error bars for the results illustrated in Figure 2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The hardware employed to obtain all experimental results, as well as a runtime comparison, are described in Appendix F.2 (see “Implementation Details”).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn’t make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have made sure to comply to the Code of Ethics and to preserve our anonymity.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The models we developed are not generative, hence not posing risks of malicious use as for what concerns fabricating misleading or otherwise fake information, online profiles and media. Additionally, although our approach improves the efficiency of certain Graph Neural Networks, the models we developed are not scalable enough to apply and impact (online) social networks: represented as graphs, their scale is way beyond that considered in our experiments.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our model(s) do not have a high risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: As explicitly mentioned in the main corpus of the paper, the inset figure in Section 2 is taken with permission by the original authors. Creators of datasets employed in this study, as well as the benchmark frameworks used are properly referenced and cited. For these last we report license information in Appendix F.1, also reported for code assets Appendix F.2.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.