
Edge Representation Learning with Hypergraphs

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Graph neural networks have recently achieved remarkable success in representing
2 graph-structured data, with rapid progress in both the node embedding and graph
3 pooling methods. Yet, they mostly focus on capturing information from the nodes
4 considering their connectivity, and not much work has been done in representing
5 the *edges*, which are essential components of a graph. However, for tasks such
6 as graph reconstruction and generation, as well as graph classification tasks for
7 which the edges are important for discrimination, accurately representing edges
8 of a given graph is crucial to the success of the graph representation learning.
9 To this end, we propose a novel edge representation learning framework based
10 on *Dual Hypergraph Transformation* (DHT), which transforms the edges of a
11 graph into the nodes of a *hypergraph*. This dual hypergraph construction allows
12 us to apply message passing techniques for node representations to edges. After
13 obtaining edge representations from the hypergraphs, we then cluster or drop
14 edges to obtain holistic graph-level edge representations. We validate our edge
15 representation learning method with hypergraphs on diverse graph datasets for
16 graph representation and generation performance, on which our method largely
17 outperforms existing graph representation learning methods. Moreover, our edge
18 representation learning and pooling method also largely outperforms state-of-the-
19 art graph pooling methods on graph classification, not only because of its accurate
20 edge representation learning, but also due to its lossless compression of the nodes
21 and removal of irrelevant edges for effective message passing.

22 1 Introduction

23 The recent demand in representing graph-structured data, such as molecular, social, and knowledge
24 graphs, has brought remarkable progress in the *Graph Neural Networks* (GNNs) [44, 35]. Early
25 works on GNNs [22, 16, 37] aim to accurately represent each node to reflect the graph topology, by
26 transforming, propagating, and aggregating information from their neighborhoods based on message
27 passing schemes [13]. More recent works focus on learning holistic graph-level representations,
28 by proposing graph pooling techniques that condense the node-level representations into a smaller
29 graph or a single vector. While such state-of-the-art node embedding or graph pooling methods
30 have achieved impressive performances on graph-related tasks (e.g., node classification and graph
31 classification), they have largely overlooked the *edges*, which are essential components of a graph.

32 Most existing GNNs, including ones that consider categorical edge features [29, 13], only implicitly
33 capture the edge information in the learned node/graph representations when updating them. While a
34 few of them aim to obtain explicit representations for edges [20, 14, 39], they mostly use them only to
35 augment the node-level representations, and thus suboptimally capture the edge information. This is
36 partly because many benchmark tasks for GNN performance evaluation, such as graph classification,
37 do not require the edge information to be accurately preserved. Thus, on these benchmarks, simple
38 MLPs without any connectivity information can sometimes outperform GNNs [9, 18]. However, for




Figure 1: **(Left)**: The two molecular graphs have the identical set of nodes, but exhibit completely different properties, due to their difference in the edges. **(Right)**: An illustration of the proposed edge representation learning framework, with two novel edge pooling schemes. The grey box in the center describes the proposed **Dual Hypergraph Transformation**, where the numbers (letters) denote the corresponding edges (nodes) in the graph and nodes (hyperedges) in the hypergraph. The two boxes in the right illustrate the proposed edge pooling methods, **HyperCluster** which clusters similar edges, and **HyperDrop** which drops unnecessary edges.

39 tasks such as graph reconstruction and generation, accurately representing the edges of a graph is
 40 crucial to the success, as incorrectly reconstructing/generating edges may result in complete failure
 41 of the tasks. For example, the two molecules (a) and (b) in Figure 1 have exactly the same set of
 42 nodes and are only different in their edges (bond types), but exhibit extremely different properties.

43 To overcome such limitations of existing GNN methods in edge representation learning, we propose
 44 a simple yet effective scheme to represent the edges. The main challenge of handling edges is the
 45 absence or suboptimality of the message passing scheme for edges. We tackle this challenge by
 46 representing the edges as nodes in a *hypergraph*, which is a generalization of a graph where that can
 47 model higher-order interactions among nodes (a single hyperedge can connect an arbitrary number of
 48 nodes). Specifically, we propose *Dual Hypergraph Transformation* (DHT) to transform edges of the
 49 original graph to nodes of a *hypergraph* (Figure 1), and nodes to hyperedges. This hypergraph-based
 50 approach is effective since it allows us to apply any off-the-shelf message-passing schemes designed
 51 for node-level representation learning, for learning the representation of the edges of a graph.

52 However, representing each edge well alone is insufficient in obtaining an accurate representation of
 53 the entire graph. Thus we propose two novel graph pooling methods for the hypergraph to obtain
 54 compact graph-level edge representations, namely *HyperCluster* and *HyperDrop*. Specifically, for
 55 obtaining global edge representations for an entire graph, *HyperCluster* coarsens similar edges into
 56 a single edge under the global graph pooling scheme (see *HyperCluster* in Figure 1). On the other
 57 hand, *HyperDrop* drops unnecessary edges from the original graph by calculating pruning scores on
 58 the hypergraph (see *HyperDrop* in Figure 1). *HyperCluster* is more useful for graph reconstruction
 59 and generation as it does not result in removal of any edges, while *HyperDrop* is more useful for
 60 classification as it learns to remove edges that are less useful for graph discrimination.

61 We first experimentally validate the effectiveness of the DHT with *HyperCluster*, on the reconstruction
 62 of synthetic and molecular graphs. Our method obtains extremely high performance on these tasks,
 63 largely outperforming baselines, which shows its effectiveness in accurately representing the edges.
 64 Then, we validate DHT with *HyperCluster* on molecular graph generation tasks, and show that it
 65 largely outperforms base generation methods, as it allows to generate molecules with more correct
 66 bonds (edges). Further, we validate *HyperDrop* on 10 benchmark datasets for graph classification, on
 67 which *HyperDrop* outperforms all hierarchical pooling baselines, with larger gains on social graphs,
 68 for which the edge features are important. Our main contributions are as follows:

- 69 • We introduce a novel edge representation learning scheme using *Dual Hypergraph Transformation*,
 70 which exploits the dual hypergraph whose nodes are edges of the original graph, on which we can
 71 apply off-the-shelf message-passing schemes designed for node-level representation learning.
- 72 • We propose novel edge pooling methods for graph-level representation learning, namely *Hyper-*
 73 *Cluster* and *HyperDrop*, to overcome the limitations of existing node-based pooling methods.
- 74 • We validate our methods on graph reconstruction, generation, and classification tasks, on which
 75 they largely outperform existing graph representation learning methods.

76 2 Related Work

77 **Graph neural networks** Graph neural networks (GNNs) mostly use the message passing
 78 scheme [13] to aggregate features from their neighbors. Particularly, Graph Convolutional Net-
 79 work (GCN) [22] generalizes the convolution operation in the spectral domain of graphs, and updates
 80 the representation of each node by applying the shared weights on it and its neighbors’ representations.

81 Similarly, GraphSAGE [16] propagates the features of each node’s neighbors to itself, based on
 82 simple aggregation operations (e.g., mean). Graph Attention Network (GAT) [32] considers the
 83 relative importance on neighboring nodes with attention, to update each node’s representation as
 84 the weighted combination of its neighbors’. Xu et al. [37] show that a simple sum on neighborhood
 85 aggregation makes GNNs as powerful as the Weisfeiler-Lehman (WL) test [34], which is effective
 86 for distinguishing different graphs. While GNNs have achieved impressive success on graph-related
 87 tasks, most of them only focus on learning node-level representations, with less focus on the edges.

88 **Edge-aware graph neural networks** Some existing works on GNNs consider edge features while
 89 updating the node features [29, 31], however, they only use the edges as auxiliary information and
 90 restrict the representation of edges as the discrete features with categorical values. While a few
 91 methods [20, 13, 14, 39] explicitly represent edges by introducing edge-level GNN layers, they use the
 92 obtained edge features solely for enhancing node features. Also, existing message passing schemes
 93 for nodes are not directly applicable to edge-level layers, as they are differently designed from the
 94 node-level layers, which makes it challenging to combine them with graph pooling methods [40] for
 95 graph-level representation learning. We overcome these limitations by proposing a dual hypergraph
 96 transformation scheme, to obtain a hypergraph whose nodes are edges of the original graph.

97 **Hypergraph** Hypergraphs can model higher-order interactions among nodes, by grouping multi-
 98 node relationships into a single hyperedge [3]. Zhou et al. [43] first represent hypergraphs based on
 99 the spectral clustering technique, normalized cut [30]. Then, Feng et al. [11], Yadati et al. [38] propose
 100 methods that are similar to GCN [22] for representing hypergraphs, by generalizing a convolution
 101 operation on the spectral domain of the hypergraphs. There are a few studies that consider hypergraph
 102 duality [3, 28], which transforms the hyperedges into the nodes, to more conveniently deal with them.
 103 Following this line of research, Lugo-Martinez and Radivojac [25] use the duality to cast a hyperlink
 104 prediction task, as an instance of node classification from the dual form of the original hypergraph.
 105 On the other hand, Kajino [21] uses the duality to extract useful rules from the hypergraph structures
 106 by transforming molecular graphs, for their generation. However, none of the existing works exploit
 107 the relation between the original graph and the dual hypergraph for edge representation learning.

108 **Graph pooling** Graph pooling methods aim to learn accurate graph-level representation, by com-
 109 pressing a graph into a smaller graph or a vector with pooling operations. The simplest pooling
 110 approaches are using mean, sum or max over all node representations [1, 37]. However, they treat all
 111 nodes equally, and cannot adaptively adjust the size of graphs for downstream tasks. More advanced
 112 methods, such as node clustering methods, coarsen the graph by clustering similar nodes based on
 113 their embeddings [40, 4], whereas the node pruning methods reduce the number of nodes from the
 114 graph by dropping unimportant nodes based on their scores [12, 23]. Ranjan et al. [27] combine both
 115 node pruning and clustering approaches, by dropping meaningless clusters after grouping nodes. Baek
 116 et al. [2] propose to use attention-based operations for considering relationships between clusters.
 117 Note that all of those pooling schemes not only ignore edge representations, but also alter the node
 118 set by dropping, clustering, or merging nodes, which result in inevitable loss of node information.

119 3 Edge Representation Learning with Hypergraphs

120 In this section, we first introduce our novel edge representation learning framework with dual
 121 hypergraphs, which we refer to as *Edge HyperGraph Neural Network* (EHGNN), and then propose
 122 two novel edge pooling schemes for holistic graph-level representation learning: *HyperCluster* and
 123 *HyperDrop*. We begin with the descriptions of graph neural networks for node representation learning.

124 **Graph Neural Networks.** A graph G with n nodes and m edges, is defined by its node features
 125 $\mathbf{X} \in \mathbb{R}^{n \times d}$, edge features $\mathbf{E} \in \mathbb{R}^{m \times d'}$, and the connectivity among the nodes represented by
 126 an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$. Here, d and d' are the dimension of node and edge features,
 127 respectively. Then, given a graph, the goal of a *Graph Neural Network* (GNN) is to learn the
 128 node-level representation with message passing between neighboring nodes [13] as follows:

$$\mathbf{X}_v^{(l+1)} = \text{UPDATE} \left(\mathbf{X}_v^{(l)}, \text{AGGREGATE} \left(\left\{ \mathbf{X}_u^{(l)} : \forall u \in \mathcal{N}(v; \mathbf{A}) \right\} \right) \right), \quad (1)$$

129 where $\mathbf{X}^{(l)}$ is the node features at l -th layer, AGGREGATE is the function that aggregates messages
 130 from a set of neighboring nodes of the node v , UPDATE is the function that updates the representation
 131 of the node v from the aggregated messages, and $\mathcal{N}(v; \mathbf{A})$ is the set of neighboring nodes for node v ,
 132 obtained from the adjacency matrix \mathbf{A} . Such message passing schemes can incorporate the graph
 133 topology into each node by updating its representation with the representation of its neighbors.




Figure 2: **Dual Hypergraph Transformation.** Illustration of the proposed graph-to-hypergraph transformation.

3.1 Edge-wise representation learning with dual hypergraph transformation

Edge representation learning. To reflect the edge information on message passing, some works on GNNs first obtain the categorical edge features between nodes, and then use them on the AGGREGATE function in equation 1, by adding or multiplying the edge features to the neighboring node’s features [13, 29], for example $\text{AGGREGATE} \left(\left\{ \mathbf{X}_u^{(l)} + \mathbf{E}_{u,v}^{(l)} : \forall u \in \mathcal{N}(v; \mathbf{A}) \right\} \right)$ (see Section A.1 of the supplementary file for more details). Similarly, few recent works aim to obtain explicit edge representations, to use them as auxiliary information to augment the node features, by adding or multiplying edge features to them [14, 39]. Thus, existing works only implicitly capture the edge information in the learned node representations. Although this could be sufficient for most benchmark graph classification tasks, many real-world tasks with graphs (e.g., graph reconstruction and generation) require the edges to be accurately represented as they largely affect the task performance.

Even worse, to define a message passing function for edge representation learning, existing works propose to additionally create the adjacency matrix for edges (see Table 1). However, this is highly suboptimal since the transformation of the node adjacency to the edge adjacency requires $\mathcal{O}(m^2)$ time complexity (see Section A.2 of the supplementary file for detailed descriptions), as shown in Table 1. This is the main obstacle for directly applying the existing message passing schemes for nodes to edges. To this end, we propose a simple yet effective method to represent the edges of a graph, using a hypergraph.

Hypergraph. A *hypergraph* is a generalization of a graph which can model graph-structured data with higher-order interactions among nodes, wherein a single hyperedge connects an arbitrary number of nodes, unlike in conventional graphs where an edge can only connect two nodes. For example, in Figure 2, the hyperedge B defines the relation among three different nodes. To denote such higher-order relations among arbitrary number of nodes defined by a hyperedge, we use an *incidence matrix* $M \in \{0, 1\}^{n \times m}$, which represents the interaction between n nodes and m hyperedges, instead of using an adjacency matrix $A \in \{0, 1\}^{n \times n}$ that only considers interactions among n nodes. Each entry in the incidence matrix indicates whether the node is incident to the hyperedge. We can formally define a hypergraph G^* with n nodes and m hyperedges, as a triplet of three components $G^* = (\mathbf{X}^*, M^*, \mathbf{E}^*)$, where $\mathbf{X}^* \in \mathbb{R}^{n \times d}$ is the node features, $\mathbf{E}^* \in \mathbb{R}^{m \times d'}$ is the hyperedge features, and $M^* \in \{0, 1\}^{n \times m}$ is the incidence matrix of the hypergraph. We can also represent conventional graphs in the form of a hypergraph, $G = (\mathbf{X}, M, \mathbf{E})$, in which a hyperedge in the incidence matrix M is associated with only two nodes. In the following paragraph, we will describe how to transform the edges of a graph into nodes of a hypergraph, for edge representation learning.

Dual Hypergraph Transformation. If we can change the role of the nodes and edges of the graph with a shared connectivity pattern across the nodes and edges, while accurately preserving their information, then we can use any node-based message passing schemes for learning edges. To this end, inspired by the hypergraph duality [3, 28], we propose to transform an edge of the original graph into a node of a hypergraph, and a node of the original graph into a hyperedge of the same hypergraph. We refer to this graph-to-hypergraph transformation as *Dual Hypergraph Transformation* (DHT) (see Figure 2). To be more precise, during the transformation, we interchange the structural role of nodes and edges from the given graph, obtaining the incidence matrix for the new dual hypergraph simply by *transposing* the incidence matrix of the original graph (see the incidence matrix in Figure 2 and Table 1). Along with the structural transformation through the incidence matrix, the DHT naturally interchanges node and edge features across G and G^* (see the feature matrices in Figure 2). Formally, given a triplet representation of a graph, DHT is defined as the following transformation:

$$\text{DHT} : G = (\mathbf{X}, M, \mathbf{E}) \mapsto G^* = (\mathbf{E}, M^T, \mathbf{X}), \quad (2)$$

Table 1: Comparison of the transformation of edge-aware GNNs [20, 14, 39] and the DHT, where the N->E denotes the transformation cost from nodes to edges for obtaining connectives for edges.

Model	Connectivity Patterns Node (N)	Edge (E)	Cost N->E
Edge-aware GNNs	$A \in \mathbb{R}^{n \times n}$	$A \in \mathbb{R}^{m \times m}$	$\mathcal{O}(m^2)$
EHGNN (Ours)	$M \in \mathbb{R}^{n \times m}$	$M^T \in \mathbb{R}^{m \times n}$	$\mathcal{O}(1)$

181 where we refer to the transformed G^* as the *dual hypergraph* of the input graph G . Since the dual
 182 hypergraph $G^* = (\mathbf{E}, \mathbf{M}^T, \mathbf{X})$ retains all the information of the original graph, we can recover the
 183 original graph from the dual hypergraph with the same DHT operation as follows:

$$DHT : G^* = (\mathbf{E}, \mathbf{M}^T, \mathbf{X}) \mapsto G = (\mathbf{X}, \mathbf{M}, \mathbf{E}). \quad (3)$$

184 This implies that DHT is a *bijective transformation*. DHT is simple to implement, does not incur the
 185 loss of any features or topological information of the input graph, and does not require additional
 186 memory for feature representations. Moreover, DHT can be sparsely implemented using the edge list,
 187 which is the sparse form of the adjacency matrix, by only reshaping the edge list of the graph into the
 188 hyperedge list of the dual hypergraph (see Section A.3 of the supplementary file for details), which
 189 is highly efficient in terms of time and memory. Thanks to DHT, we define the message passing
 190 between edges of the graph as the message passing between nodes of its dual hypergraph.

191 **Message passing on the dual hypergraph for edge representation learning.** After transforming
 192 the original graph into its corresponding dual hypergraph using DHT, we can perform the message
 193 passing between the edges of the input graph, by performing the message passing between the nodes
 194 of its dual hypergraph $G^* = (\mathbf{E}, \mathbf{M}^T, \mathbf{X})$, which is formally denoted as follows:

$$\mathbf{E}_e^{(l+1)} = \text{UPDATE} \left(\mathbf{E}_e^{(l)}, \text{AGGREGATE} \left(\left\{ \mathbf{E}_f^{(l)} : \forall f \in \mathcal{N}(e; \mathbf{M}^T) \right\} \right) \right), \quad (4)$$

195 where $\mathbf{E}^{(l)}$ is the node features of G^* at l -th layer, the AGGREGATE function summarizes the
 196 neighboring messages of the node e of the dual hypergraph G^* , and the UPDATE function updates
 197 the representation of the node e from the aggregated messages. Here $\mathcal{N}(e; \mathbf{M}^T)$ is the neighboring
 198 node set of the node e in G^* , which we obtain using the incidence matrix \mathbf{M}^T of G^* . Although
 199 acquiring the neighboring node set from the incidence matrix costs $\mathcal{O}(m)$, compared to using the
 200 adjacency matrix which costs $\mathcal{O}(n)$, this cost can be ignored by using the sparse implementation of
 201 DHT explained above. Note that, since the form of equation 4 is the same as the form of equation 1,
 202 we can use any graph neural networks which realize the message passing operation in equation 1,
 203 such as GCN [22], GAT [32], GraphSAGE [16], and GIN [37] for equation 4. In other words, to learn
 204 the edge representations \mathbf{E} of the original graph, we do not require any specially designed layers, but
 205 simply need to perform DHT to directly apply existing off-the-shelf message passing schemes.

206 To simplify, we summarize the equation 1 as follows: $\mathbf{X}^{(l+1)} = \text{GNN}(\mathbf{X}^{(l)}, \mathbf{M}, \mathbf{E}^{(l)})$, and the
 207 equation 4 as follows: $\mathbf{E}^{(l+1)} = \text{GNN}(\mathbf{E}^{(l)}, \mathbf{M}^T, \mathbf{X}^{(l)}) = \text{EHGNN}(\mathbf{X}^{(l)}, \mathbf{M}, \mathbf{E}^{(l)})$, where
 208 EHGNN indicates our edge representation learning framework using DHT. After updating the edge
 209 features $\mathbf{E}^{(L)}$ with EHGNN, $\mathbf{E}^{(L)}$ is returned to the original graph by applying DHT on the dual
 210 hypergraph G^* . Then, the remaining step is how to make use of these edge-wise representations to
 211 accurately represent the edges of the entire graph, which we describe in the next subsection.

212 3.2 Graph-level edge representation learning with edge pooling

213 Existing graph pooling methods do not explicitly represent edges. To overcome this limitation, we
 214 propose two novel edge pooling schemes: *HyperCluster* and *HyperDrop*.

215 **Graph pooling.** The goal of graph pooling is to learn a holistic representation of the entire graph.
 216 The most straightforward approach for this, is to aggregate all the node features with mean or sum
 217 operations [1, 37], but they treat all nodes equally without consideration of which nodes are important
 218 for the given task. To tackle this limitation, recent graph pooling methods propose to either cluster and
 219 coarsen nodes [40, 4] or drop unnecessary nodes [12, 23]. While they yield improved performances
 220 on graph classification tasks, they suffer from an obvious drawback: inevitable loss of both node and
 221 edge information. The node information is lost as the nodes are dropped and coarsened, and the edge
 222 information is lost as the edges for the dropped nodes or the internal edges for the coarsened nodes
 223 are removed. To overcome this limitation, we propose a graph-level edge representation learning.

224 **HyperCluster.** We first introduce *HyperCluster*, which is a novel edge clustering method to coarsen
 225 similar edges into a single edge, for obtaining the global edge representation. Generally, a clustering
 226 scheme for nodes of the graph [40, 4] is defined as follows:

$$\mathbf{X}^{pool} = \mathbf{C}^T \mathbf{X}', \quad \mathbf{M}^{pool} = \mathbf{C}^T \mathbf{M}, \quad (5)$$

227 where $\mathbf{X}^{pool} \in \mathbb{R}^{n_{pool} \times d}$ and $\mathbf{M}^{pool} \in \mathbb{R}^{n_{pool} \times m}$ denote the pooled representations, $\mathbf{X}' =$
 228 $\text{GNN}(\mathbf{X}, \mathbf{M}, \mathbf{E}) \in \mathbb{R}^{n \times d}$ is the updated node features, and $\mathbf{C} \in \mathbb{R}^{n \times n_{pool}}$ is the cluster assignment

229 matrix that is generated from the \mathbf{X}' . Following this approach, the proposed HyperCluster clusters
 230 similar edges into a single edge, by clustering nodes of the dual hypergraph obtained from the
 231 original graph via DHT. In other words, we first obtain the node representation of the dual hypergraph
 232 $\mathbf{E}' = \text{EHGNN}(\mathbf{X}, \mathbf{M}, \mathbf{E}) \in \mathbb{R}^{m \times d'}$, and then cluster the nodes of the dual hypergraph as follows:

$$\mathbf{E}^{pool} = \mathbf{C}^T \mathbf{E}', \quad (\mathbf{M}^{pool})^T = \mathbf{C}^T \mathbf{M}^T \quad (6)$$

233 where $\mathbf{E}^{pool} \in \mathbb{R}^{m_{pool} \times d'}$ and $\mathbf{M}^{pool} \in \mathbb{R}^{n \times m_{pool}}$ denote the pooled edge representation and the
 234 incidence matrix of the input graph respectively, and $\mathbf{C} \in \mathbb{R}^{m \times m_{pool}}$ is the cluster assignment
 235 matrix generated from the input edge features \mathbf{E}' . Since HyperCluster coarsens the edges rather than
 236 dropping them, this edge pooling method is more appropriate for tasks such as graph reconstruction.

237 **HyperDrop.** We propose another edge pooling scheme, *HyperDrop*, which drops unnecessary edges
 238 to identify task-relevant edges, while performing lossless compression of nodes. Conventional node
 239 drop methods [12, 23] for graph pooling remove less relevant nodes based on their scores, as follows:

$$\mathbf{X}^{pool} = \mathbf{X}_{idx}, \quad \mathbf{M}^{pool} = \mathbf{M}_{idx}; \quad idx = \text{top}_k(\text{score}(\mathbf{X})), \quad (7)$$

241 where idx is the row-wise (i.e., node-wise) indexing vector, $\text{score}(\cdot)$ computes the score of each node
 242 with learnable parameters, and $\text{top}_k(\cdot)$ selects the top k elements in terms of the score. However,
 243 this approach results in the inevitable loss of node information, as it drops nodes. Thus, we propose
 244 to coarsen the graph by dropping edges instead of nodes, exploiting edge representations obtained
 245 from our EHGNN. *HyperDrop* selects the top ranked edges of the original graph, by selecting the top
 246 ranked nodes of the dual hypergraph. The pooling procedure for HyperDrop is formalized as follows:

$$\mathbf{E}^{pool} = \mathbf{E}_{idx}, \quad (\mathbf{M}^{pool})^T = (\mathbf{M}^T)_{idx}; \quad idx = \text{top}_k(\text{score}(\mathbf{E})). \quad (8)$$

247 Then, we can obtain the pooled graph $G^{pool} = (\mathbf{X}, \mathbf{M}^{pool}, \mathbf{E}^{pool})$ by applying DHT to the pooled
 248 dual hypergraph. HyperDrop is most suitable for graph classification tasks, as it identifies discrimi-
 249 native edges for the given task. Since HyperDrop preserves the nodes intact, it can also be used for
 250 node-level classification tasks, which is impossible with exiting graph pooling methods that modify
 251 nodes. A notable advantage of such HyperDrop is that it alleviates the over-smoothing problem in
 252 deep GNNs [24] (i.e., the features of all nodes converge to the same values when stacking a large
 253 number of GNN layers). As HyperDrop *learns* to remove unnecessary edges, the message passing
 254 only happens across relevant nodes, which alleviates over-smoothing.

256 4 Experiments

257 We experimentally validate our EHGNN that is coupled with either HyperCluster or HyperDrop on
 258 four different tasks: graph reconstruction, generation, classification, and node classification.

259 4.1 Graph reconstruction

260 Accurately reconstructing the edges is crucial for graph reconstruction tasks, and thus we validate the
 261 efficacy of our method on the graph reconstruction tasks first.

262 **Experimental setup.** We first validate our EHGNN with HyperCluster on the *edge reconstruction*
 263 tasks. Then, we evaluate it on the graph reconstruction tasks to validate our method’s effectiveness in
 264 holistic graph-level learning. We start with edge reconstruction of a synthetic two-moon graph, where
 265 node features (coordinates) are fixed and edge features are colors. For edge and graph reconstruction
 266 of real-world graphs, we use a ZINC dataset [19] that consists of 12K molecular graphs [7], where
 267 node features are atom types and edge features are bond types. We use the accuracy, validity, and
 268 exact match as evaluation metrics. For more details, please see Section C.1 of the supplementary file.

269 **Implementation details and baselines.** We compare the proposed EHGNN framework against
 270 edge-aware GNNs, namely EGCN [17], MPNN [13], R-GCN [29], and EGNN [14], which use the
 271 edge features as auxiliary information for updating node features. We further combine them with
 272 an existing graph pooling method, namely GMPool [2], to obtain a graph-level edge representation
 273 for a given graph. In contrast, for our method, we first obtain edge representations with EHGNN,
 274 using GCN [22] as the message passing function, and then coarsen the edge-wise representations
 275 using HyperCluster, whose cluster assignment matrices are obtained using GMPool [2]. For node
 276 reconstruction, we set message passing to GCN and graph pooling to GMPool [2] for all models. We
 277 provide further details of the baselines and our model in Section C.1 of the supplementary file.

Figure 3: **Edge reconstruction results** on the ZINC molecule dataset by varying the compression ratio. Solid lines denote the mean of 5 different runs. High scores indicate the better.

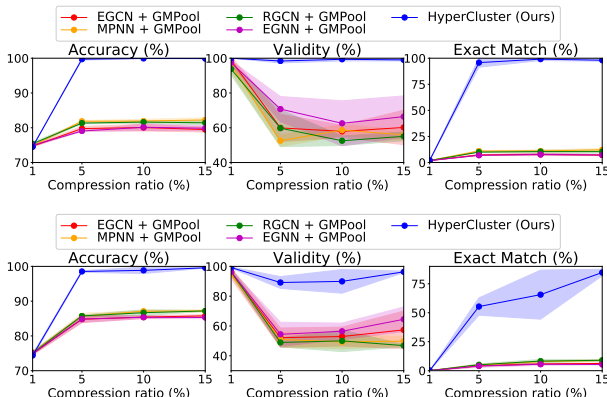


Figure 5: **Graph reconstruction results** on the ZINC molecule dataset by varying the compression ratio. Solid lines denote the mean, and shaded areas denote the standard deviation of 5 runs.

Figure 4: **Edge reconstruction results** of the synthetic two-moon graph. The edge features are represented as colors.





Figure 6: **Graph reconstruction examples.** Red dashed circles and squares indicate the incorrectly predicted edges and nodes, respectively. (c.2) shows the assigned clusters of edges as colors using our method.

278 **Edge reconstruction results.** Figure 4 shows the original two-moon graph and edge-reconstructed
 279 graphs, where edge features are colors, exhibiting clustered pattern. The baselines fail to reconstruct
 280 the edge colors, since they implicitly learn edge representations by using edge features as auxiliary
 281 information to update nodes, hence mixing the colors of the neighboring edges. On the other hand,
 282 our method distinguishes each edge cluster, which shows that our method can capture meaningful
 283 edge information by clustering similar edges. Moreover, as shown in Figure 3, our model obtains
 284 significantly higher performance over all baselines on the edge reconstruction task of molecular
 285 graphs, in all evaluation metrics. The performance gain of our method over baselines is notably
 286 large in exact match, which demonstrates that explicit learning of edge representation is essential for
 287 accurate encoding of the edge information.

288 **Graph reconstruction results.** We now validate our method on graph reconstruction for reconstructing
 289 both the nodes and edges of molecular graphs in Figure 5. Combining our edge representation
 290 learning method (EHGNN + HyperCluster) with existing node representation learning method (GCN
 291 + GMPool) yields incomparably high reconstruction performance compared to the baselines in exact
 292 match, which demonstrates that learning accurate edge representation, as well as node representation,
 293 is crucial to the success of the graph representation learning methods on graph reconstruction.

294 **Qualitative analysis.** We visualize the original and reconstructed molecular graphs in Figure 6. As
 295 shown in Figure 6 (b), the baseline cannot reconstruct the ring structures of the molecule, whereas
 296 our method perfectly reconstructs the rings as well as the atom types. The generated edge clusters in
 297 Figure 6 (c.2) further show that our method captures the detailed substructures of the molecule, as we
 298 can see in the cluster patterns of hexagonal and pentagonal rings.

299 **Graph compression.** To validate the effectiveness of our
 300 method in dense graph compression, we further apply EHGNN
 301 + HyperCluster to the Erdos-Renyi random graph [8] having
 302 six discrete edge features, with the number of nodes fixed to
 303 10^3 while the number of edges increases from 10^3 to 10^4 . In
 304 Figure 7, we report the relative size of compressed graphs from
 305 our method against the node pooling method, GMT [2]. As the
 306 number of edges increases, we observe that compressing only
 307 node features is insufficient for obtaining compact representa-
 308 tions, whereas our method is able to obtain highly compact but
 309 accurate representation which can be assured from the sufficiently high reconstruction accuracy.




Figure 7: **Graph compression results.** For ours, we report the relative size to the full graph where the edge reconstruction accuracy is higher than 95% and 75%.

310 4.2 Graph generation

311 As shown in Figure 1 (left), graph generation depends heavily on the edge representations, as the
 312 model may generate incorrect graphs (e.g., toxic chemicals rather than drugs) if the edge information
 313 is inaccurate. Thus, we further validate our EHGNN on the graph generation tasks.




Figure 8: **Graph generation results on MolGAN.** Table 2: **Graph generation results on MARS.** The solid lines denote the mean scores of 3 different runs. The results are the mean and standard deviation of 3 runs.

Datasets	Metrics	MARS [36]	MARS + EHGNN (Ours)
ZINC15	Success Rate	59.53 ± 2.11	64.30 ± 1.54
	QED (≥ 0.67)	95.71 ± 0.09	96.36 ± 0.49
	GSK3 β (≥ 0.6)	86.52 ± 1.67	90.63 ± 2.57
	JNK3 (≥ 0.6)	71.52 ± 4.15	73.60 ± 1.29
ChEMBL	Success Rate	56.64 ± 5.79	58.25 ± 6.07
	QED (≥ 0.67)	91.01 ± 2.79	91.13 ± 4.84
	GSK3 β (≥ 0.6)	87.45 ± 1.73	90.34 ± 2.65
	JNK3 (≥ 0.6)	70.57 ± 4.75	70.01 ± 4.83

314 **Experimental setup.** We directly forward the edge representation from the EHGNN to molecule generation networks, namely MolGAN [5] and Markov molecular Sampling (MARS) [36]. MolGAN uses the Generative Adversarial Network (GAN) [15], to generate the molecular graph by balancing weights between its generator and discriminator. MolGAN uses R-GCN [29] for node-level message passing, whereas, for ours, we first obtain the edge representations using EHGNN, and use them with mean pooling in the graph encoder. For evaluation metrics, we use the Synthetic Accessibility (SA) and Druglikeness (QED) scores. We further apply EHGNN to MARS [36] that generates the molecule by sequentially adding or deleting its fragment, with MCMC sampling. While the original model uses MPNN [13] to implicitly obtain edge representations for the actions, we use EHGNN to explicitly learn edge representations. We train models to maximize the four molecule properties: the inhibition scores against two proteins, namely GSK3 β and JNK3 (Biological); QED and SA scores (Non-biological). Then we report the success rate of which the molecule satisfies all the properties. For more details, please see Section C.2 of the supplementary file.

327 **MolGAN result.** Figure 8 shows the SA and QED scores of the generated molecules, of the MolGAN architecture with different encoders. Our EHGNN framework obtains significantly improved generation performance, over the original MolGAN which uses the R-GCN encoder and the MolGAN with GMPool, a state-of-the-art global node pooling encoder. This is because EHGNN learns explicit edge representation which enhances the ability of the discriminator to distinguish between real and generated graphs. The improvement in the discriminator also leads to notably more stable results compared to the baseline, which shows large variance in the quality of the generated molecules.

334 **MARS result.** To perform correct editing actions to generate graphs with MARS, we need accurate edge representations, as edges determine the structure of the generated molecule. Table 2 shows that using our EHGNN achieves significantly higher generation performance over original MARS, that uses edges as auxiliary information only to enhance node representations. Notably, performance gain on the GSK3 β metric for which structural binding is important, suggests that accurate learning of edges is beneficial in generating more effective molecules that interact with the target protein.

340 4.3 Graph and node classification

341 Now, we validate the performance of our EHGNN with HyperDrop on classification tasks. Our approach is effective for classification of graphs with or without edge features, since it allows lossless compression of nodes and drops edges to allow message passing only across relevant nodes.

344 **Experimental setup.** Following the experimental setting of Baek et al. [2], we use the GCN as the node-level message passing layers for all models, and compare our edge pooling method against existing graph pooling methods. For this experiment, our HyperDrop uses SAGPool [23] on the hypergraph, which is a node drop pooling method based on self-attention. We use 6 datasets from TU datasets [26] including three from the biochemical domains (i.e., DD, PROTEINS, MUTAG) and the remaining half from the social domains (i.e., IMDB-BINARY, IMDB-MULTI, COLLAB). Also, we further use the 4 molecule datasets (i.e., HIV, Tox21, ToxCast, BBBP) from the recently released OGB datasets [17]. We evaluate the accuracy of each model with 10-fold cross validation [42] on the TU datasets, and use ROC-AUC as the evaluation metric for OGB datasets. For both datasets, we follow the standard experimental settings, from the feature extraction to the dataset splitting. We provide additional details of the experiments in Section C.3 of the supplementary file.

355 **Baselines.** We compare our EHGNN with HyperDrop, against the set encoding (DeepSet [41]), GNNs with naive pooling baselines (GCN and GIN [22, 37]), and state-of-the-art hierarchical pooling methods (DiffPool [40], SAGPool [23], TopKPool [12], MinCutPool [4], ASAP [27], EdgePool [6], and HaarPool [33]) that drop or coarsen node representations. We also additionally compare or combine the state-of-the-art global node pooling methods (SortPool [42], GMT [2]) with our model, for example HyperDrop + GMT. For more details, see Section C.3 of the supplementary file.

Table 3: **Graph classification results.** The results are the mean and standard deviation over 10 different runs. Best performance and its comparable results ($p > 0.05$) from the t-test are highlighted in bold. Hyphen (-) denotes out-of-resources that take more than 10 days. The results for the baselines are taken from Baek et al. [2].

		TU : Biochemical			TU : Social			OGB : Molecule			Average	
		D&D	PROTEINS	MUTAG	IMDB-B	IMDB-M	COLLAB	HIV	Tox21	ToxCast		BBBP
# graphs		1178	1113	188	1000	1500	5000	41127	7831	8576	2039	
# classes		2	2	2	2	3	3	2	12	617	2	
Avg # nodes		284.32	39.06	17.93	19.77	13.00	74.49	25.51	18.57	18.78	24.06	
Avg # edges		715.66	72.82	19.79	96.53	65.94	2457.78	27.47	19.27	19.26	25.95	
Set	DeepSet	77.39 ± 0.67	68.95 ± 0.92	72.56 ± 1.09	72.42 ± 0.36	50.24 ± 0.32	75.27 ± 0.21	71.20 ± 1.26	72.25 ± 0.23	59.44 ± 0.39	63.64 ± 0.62	68.34
	GCN	72.05 ± 0.55	73.24 ± 0.73	69.50 ± 1.78	73.26 ± 0.46	50.39 ± 0.41	80.59 ± 0.27	76.81 ± 1.01	75.04 ± 0.80	60.63 ± 0.51	65.47 ± 1.73	69.70
Naive GNN	GIN	70.79 ± 1.17	71.46 ± 1.66	81.39 ± 1.53	72.78 ± 0.86	48.13 ± 1.36	78.19 ± 0.63	75.95 ± 1.35	73.27 ± 0.84	60.83 ± 0.46	67.65 ± 3.00	70.04
	SortPool	75.58 ± 0.72	73.17 ± 0.88	71.94 ± 3.55	72.12 ± 1.12	48.18 ± 0.83	77.87 ± 0.47	71.82 ± 1.63	69.54 ± 0.75	58.69 ± 1.71	65.98 ± 1.70	68.49
Global	GMT	78.72 ± 0.59	75.09 ± 0.59	83.44 ± 1.33	73.48 ± 0.76	50.66 ± 0.82	80.74 ± 0.54	77.56 ± 1.25	77.30 ± 0.59	65.44 ± 0.58	68.31 ± 1.62	73.07
	DiffPool	77.56 ± 0.41	73.03 ± 1.00	79.22 ± 1.02	73.14 ± 0.70	51.31 ± 0.72	78.68 ± 0.43	75.64 ± 1.86	74.88 ± 0.81	62.28 ± 0.56	68.25 ± 0.96	71.40
	SAGPool	74.72 ± 0.82	71.56 ± 1.49	73.67 ± 4.28	72.55 ± 1.28	50.23 ± 0.44	78.03 ± 0.31	71.44 ± 1.67	69.81 ± 1.75	58.91 ± 0.80	63.94 ± 2.59	68.49
	TopKPool	73.63 ± 0.55	70.48 ± 1.01	67.61 ± 3.36	71.58 ± 0.95	48.59 ± 0.72	77.58 ± 0.85	72.27 ± 0.91	69.39 ± 2.02	58.42 ± 0.91	65.19 ± 2.30	67.47
Hierarchical	MinCutPool	78.22 ± 0.54	74.72 ± 0.48	79.17 ± 1.64	72.65 ± 0.75	51.04 ± 0.70	80.87 ± 0.34	75.37 ± 2.05	75.11 ± 0.69	62.48 ± 1.33	65.97 ± 1.13	71.56
	ASAP	76.58 ± 1.04	73.92 ± 0.63	77.83 ± 1.49	72.81 ± 0.50	50.78 ± 0.75	78.64 ± 0.50	72.86 ± 1.40	72.24 ± 1.66	58.09 ± 1.62	63.50 ± 2.47	69.73
	EdgePool	75.85 ± 0.58	75.12 ± 0.76	74.17 ± 1.82	72.46 ± 0.74	50.79 ± 0.59	-	72.66 ± 1.70	73.77 ± 0.68	60.70 ± 0.92	67.18 ± 1.97	-
	HaarPool	-	-	66.11 ± 1.50	73.29 ± 0.34	49.98 ± 0.57	-	-	-	-	66.11 ± 0.82	-
Ours	HyperDrop	78.74 ± 0.68	75.30 ± 0.45	84.00 ± 0.69	73.96 ± 0.41	51.68 ± 0.41	81.29 ± 0.25	76.79 ± 0.86	76.95 ± 0.32	64.21 ± 0.70	69.04 ± 0.86	73.20
	HyperDrop + GMT	78.39 ± 0.33	75.39 ± 0.26	85.72 ± 0.61	74.45 ± 0.61	51.45 ± 0.28	80.59 ± 0.33	77.84 ± 0.37	77.58 ± 0.43	65.15 ± 0.65	69.16 ± 1.04	73.57

Model	MUTAG	PROTEINS	Tox21
HyperDrop	84.00 ± 0.69	75.30 ± 0.45	76.95 ± 0.32
HyperCluster	84.50 ± 1.50	72.76 ± 1.12	76.68 ± 0.56
w/ RandDrop	83.06 ± 1.15	74.92 ± 0.51	76.39 ± 0.47
w/o HyperDrop	83.06 ± 1.20	75.08 ± 0.37	76.60 ± 0.45
w/o EHGNN	69.50 ± 1.78	73.24 ± 0.73	75.04 ± 0.80





Table 4: **Ablation study of HyperDrop** on the MUTAG, PROTEINS, the COLLAB dataset. Colors denote connected components. Figure 9: **Edge pooling results** on the MUTAG, PROTEINS, the COLLAB dataset. Colors denote connected components. Figure 10: **Node classification results.** Lines denote means over 10 runs and shades denote variances.

361 **Classification results.** Table 3 shows that the proposed EHGNN + HyperDrop significantly out-
362 performs all hierarchical pooling baselines. This is because HyperDrop not only retains nodes by
363 removing edges that are less useful for graph discrimination, but also explicitly uses the edge features
364 for graph classification tasks. Since HyperDrop does not remove any nodes on the graph, it can be
365 jointly used with any node pooling methods, and thus, we pair HyperDrop with GMT. This model
366 largely outperforms GMT, obtaining best performance on most of the datasets, which demonstrates
367 that accurate learning of both the nodes and edges is important for classifying graphs. We further
368 visualize the edge pooling process of HyperDrop in Figure 9, which shows that our method accurately
369 captures the substructures of the entire graph, which leads to dividing the large graph into the several
370 connected components, thus adjusting the graph topology for more effective message passing.

371 **Ablation study.** To see how much each component contributes to the performance gain, we conduct
372 an ablation study on EHGNN with HyperDrop. Table 4 shows that, compared with a model that only
373 uses node features obtained by GCN, learning explicit edge representations significantly improves
374 the performance. Our model with HyperCluster, or without HyperDrop, or the model with random
375 edge drop obtains decent performance, but substantially underperforms HyperDrop.

376 **Over-Smoothing with Deep GNNs.** Lastly, we demonstrate that our EHGNN with HyperDrop
377 alleviates the over-smoothing problem of deep GNNs on semi-supervised node classification tasks,
378 following the setting of existing works [22, 32, 10]. We provide the experimental details in Section
379 C.4 of the supplementary file. As shown in Figure 10, HyperDrop retains the accuracy as the number
380 of layers increases, whereas the naive GCN or random drop results in largely degraded performance.
381 Further, our method outperforms BatchNorm which alleviates over-smoothing by yielding differently
382 normalized feature distribution at each node. This is because HyperDrop splits the given graph into
383 smaller subgraphs that capture meaningful message passing substructures as shown in Figure 9.

384 5 Conclusion

385 We tackled the problem of accurately representing the edges of a graph, which has been relatively
386 overlooked over node representation learning. To this end, we proposed a novel edge representation
387 learning framework using *Dual Hypergraph Transformation* (DHT), which transforms the edges of
388 the original graph into nodes on a hypergraph. This allows us to apply a message passing scheme
389 for node representation learning for edge representation learning. Further, we proposed two edge
390 pooling methods to obtain a holistic edge representation for a given graph, where one clusters similar
391 edges into a single edge for graph reconstruction and the other drops unnecessary edges for graph
392 classification. We validated our edge representation learning framework on graph reconstruction,
393 generation, and classification tasks, showing its effectiveness over relevant baselines. We describe the
394 **limitations and societal impacts** of our work in section E of the supplementary file.

395 References

- 396 [1] James Atwood and Don Towsley. Diffusion-convolutional neural networks. In *Advances*
397 *in Neural Information Processing Systems 29: Annual Conference on Neural Information*
398 *Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1993–2001, 2016.
- 399 [2] Jinheon Baek, Minki Kang, and Sung Ju Hwang. Accurate learning of graph representations
400 with graph multiset pooling. In *International Conference on Learning Representations*, 2021.
- 401 [3] Claude Berge. Graphs and hypergraphs. 1973.
- 402 [4] Filippo Maria Bianchi, Daniele Grattarola, and Cesare Alippi. Spectral clustering with graph
403 neural networks for graph pooling. *arXiv preprint*, arXiv:1907.00481, 2019.
- 404 [5] Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular
405 graphs. *arXiv preprint*, arXiv:1805.11973, 2018.
- 406 [6] Frederik Diehl. Edge contraction pooling for graph neural networks. *arXiv preprint*
407 *arXiv:1905.10990*, 2019.
- 408 [7] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Thomas Laurent, Yoshua Bengio, and Xavier
409 Bresson. Benchmarking graph neural networks. *arXiv preprint*, arXiv:2003.00982, 2020.
- 410 [8] P. Erdős and A Rényi. On the evolution of random graphs. In *PUBLICATION OF THE*
411 *MATHEMATICAL INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES*, pages
412 17–61, 1960.
- 413 [9] Federico Errica, Marco Podda, Davide Bacciu, and Alessio Micheli. A fair comparison of
414 graph neural networks for graph classification. In *8th International Conference on Learning*
415 *Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- 416 [10] Wenzheng Feng, Jie Zhang, Yuxiao Dong, Yu Han, Huanbo Luan, Qian Xu, Qiang Yang,
417 Evgeny Kharlamov, and Jie Tang. Graph random neural networks for semi-supervised learning
418 on graphs. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan,
419 and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual*
420 *Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12,*
421 *2020, virtual*, 2020.
- 422 [11] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural
423 networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The*
424 *Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth*
425 *AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu,*
426 *Hawaii, USA, January 27 - February 1, 2019*, pages 3558–3565. AAAI Press, 2019.
- 427 [12] Hongyang Gao and Shuiwang Ji. Graph u-nets. In *Proceedings of the 36th International*
428 *Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*,
429 volume 97 of *Proceedings of Machine Learning Research*, pages 2083–2092. PMLR, 2019.
- 430 [13] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl.
431 Neural message passing for quantum chemistry. In *Proceedings of the 34th International*
432 *Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*,
433 volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272. PMLR, 2017.
- 434 [14] Liyu Gong and Qiang Cheng. Exploiting edge features for graph neural networks. In *IEEE*
435 *Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA,*
436 *June 16-20, 2019*, pages 9211–9219. Computer Vision Foundation / IEEE, 2019.
- 437 [15] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil
438 Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint*,
439 arXiv:1406.2661, 2014.
- 440 [16] William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on
441 large graphs. In *Advances in Neural Information Processing Systems 30: Annual Conference*
442 *on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*,
443 pages 1024–1034, 2017.

- 444 [17] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
445 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
446 *arXiv preprint*, arXiv:2005.00687, 2020.
- 447 [18] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin Benson. Combining label
448 propagation and simple models out-performs graph neural networks. In *International Conference*
449 *on Learning Representations*, 2021.
- 450 [19] John J. Irwin, Teague Sterling, Michael M. Mysinger, Erin S. Bolstad, and Ryan G. Coleman.
451 ZINC: A free tool to discover chemistry for biology. *J. Chem. Inf. Model.*, 52(7):1757–1768,
452 2012.
- 453 [20] Xiaodong Jiang, Pengsheng Ji, and Sheng Li. Censnet: Convolution with edge-node switching
454 in graph neural networks. In *Proceedings of the Twenty-Eighth International Joint Conference*
455 *on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 2656–2662.
456 ijcai.org, 2019.
- 457 [21] Hiroshi Kajino. Molecular hypergraph grammar with its application to molecular optimization.
458 In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International*
459 *Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*,
460 volume 97 of *Proceedings of Machine Learning Research*, pages 3183–3191. PMLR, 2019.
- 461 [22] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional
462 networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon,*
463 *France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- 464 [23] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. In *Proceedings of the*
465 *36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach,*
466 *California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3734–3743.
467 PMLR, 2019.
- 468 [24] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks
469 for semi-supervised learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial*
470 *Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18),*
471 *and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18),*
472 *New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3538–3545. AAAI Press, 2018.
- 473 [25] Jose Lugo-Martinez and Predrag Radivojac. Classification in biological networks with hyper-
474 graphlet kernels. *arXiv preprint*, arXiv:1703.04823, 2017.
- 475 [26] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural
476 networks for graphs. In *Proceedings of the 33rd International Conference on Machine Learning,*
477 *ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and*
478 *Conference Proceedings*, pages 2014–2023. JMLR.org, 2016.
- 479 [27] Ekagra Ranjan, Soumya Sanyal, and Partha P. Talukdar. ASAP: adaptive structure aware
480 pooling for learning hierarchical graph representations. In *The Thirty-Fourth AAAI Conference*
481 *on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial*
482 *Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in*
483 *Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 5470–5477.
484 AAAI Press, 2020.
- 485 [28] Edward Scheinerman and Daniel Ullman. *Fractional graph theory: a rational approach to the*
486 *theory of graphs*. Courier Coporation, 2011.
- 487 [29] Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and
488 Max Welling. Modeling relational data with graph convolutional networks. In *The Semantic*
489 *Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018,*
490 *Proceedings*, volume 10843 of *Lecture Notes in Computer Science*, pages 593–607. Springer,
491 2018.
- 492 [30] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern*
493 *Anal. Mach. Intell.*, 22(8):888–905, 2000.

- 494 [31] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. Composition-based
495 multi-relational graph convolutional networks. In *8th International Conference on Learning*
496 *Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- 497 [32] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
498 Bengio. Graph attention networks. In *6th International Conference on Learning Representations,*
499 *ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings,*
500 2018.
- 501 [33] Yu Guang Wang, Ming Li, Zheng Ma, Guido Montúfar, Xiaosheng Zhuang, and Yanan Fan.
502 Haarpooling: Graph pooling with compressive haar basis. *arXiv preprint arXiv:1909.11580*,
503 2019.
- 504 [34] B. Yu. Weisfeiler and A. A. Leman. Reduction of a graph to a canonical form and an algebra
505 arising during this reduction. 1968.
- 506 [35] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A
507 comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- 508 [36] Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li.
509 Mars: Markov molecular sampling for multi-objective drug discovery. *arXiv preprint*,
510 arXiv:2103.10432, 2021.
- 511 [37] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
512 networks? In *7th International Conference on Learning Representations, ICLR 2019, New*
513 *Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- 514 [38] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and
515 Partha P. Talukdar. Hypergcn: A new method for training graph convolutional networks on
516 hypergraphs. In *Advances in Neural Information Processing Systems 32: Annual Conference on*
517 *Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver,*
518 *BC, Canada*, pages 1509–1520, 2019.
- 519 [39] Yulei Yang and Dongsheng Li. NENN: incorporate node and edge features in graph neural
520 networks. In *Proceedings of The 12th Asian Conference on Machine Learning, ACML 2020,*
521 *18-20 November 2020, Bangkok, Thailand*, volume 129 of *Proceedings of Machine Learning*
522 *Research*, pages 593–608. PMLR, 2020.
- 523 [40] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure
524 Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances*
525 *in Neural Information Processing Systems 31: Annual Conference on Neural Information*
526 *Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 4805–
527 4815, 2018.
- 528 [41] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov,
529 and Alexander J. Smola. Deep sets. In *Advances in Neural Information Processing Systems 30:*
530 *Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long*
531 *Beach, CA, USA*, pages 3391–3401, 2017.
- 532 [42] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning
533 architecture for graph classification. In *Proceedings of the Thirty-Second AAAI Conference*
534 *on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence*
535 *(IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence*
536 *(EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4438–4445. AAAI Press,
537 2018.
- 538 [43] Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. Learning with hypergraphs: Cluster-
539 ing, classification, and embedding. In *Advances in Neural Information Processing Systems 19,*
540 *Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems,*
541 *Vancouver, British Columbia, Canada, December 4-7, 2006*, pages 1601–1608. MIT Press,
542 2006.
- 543 [44] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, and Maosong Sun. Graph
544 neural networks: A review of methods and applications. *arXiv preprint*, arXiv:1812.08434,
545 2018.

546 **Checklist**

- 547 1. For all authors...
- 548 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
549 contributions and scope? [Yes]
- 550 (b) Did you describe the limitations of your work? [Yes] We discuss them in section E of
551 the supplementary file.
- 552 (c) Did you discuss any potential negative societal impacts of your work? [Yes] We
553 describe the potential societal impacts in section E of the supplementary file.
- 554 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
555 them? [Yes]
- 556 2. If you are including theoretical results...
- 557 (a) Did you state the full set of assumptions of all theoretical results? [N/A]
- 558 (b) Did you include complete proofs of all theoretical results? [N/A]
- 559 3. If you ran experiments...
- 560 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
561 mental results (either in the supplemental material or as a URL)? [Yes] We provide the
562 code, data, and instructions in the supplementary materials, with detailed experimental
563 setups in section C of the supplementary file.
- 564 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
565 were chosen)? [Yes] We specify them in section C of the supplementary file.
- 566 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
567 ments multiple times)? [Yes] All of our main results contain variances or standard
568 deviations of multiple runs.
- 569 (d) Did you include the total amount of compute and the type of resources used (e.g.,
570 type of GPUs, internal cluster, or cloud provider)? [Yes] We provide the details of
571 computing resources in section C of the supplementary file.
- 572 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 573 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 574 (b) Did you mention the license of the assets? [N/A] If the assets that we used are brought
575 from the public repository on the web, then we do our best to specify them with links
576 in the section C of the supplementary file.
- 577 (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
578 We do not additionally curate the new assets, but follow the standard settings of existing
579 works on each task.
- 580 (d) Did you discuss whether and how consent was obtained from people whose data you're
581 using/curating? [N/A] Since we do not curate new datasets, we cite the dataset paper.
582 Please see the original dataset paper for details.
- 583 (e) Did you discuss whether the data you are using/curating contains personally identifiable
584 information or offensive content? [N/A] Since we do not curate new datasets, we cite
585 the dataset paper. Please see the original dataset paper for details.
- 586 5. If you used crowdsourcing or conducted research with human subjects...
- 587 (a) Did you include the full text of instructions given to participants and screenshots, if
588 applicable? [N/A]
- 589 (b) Did you describe any potential participant risks, with links to Institutional Review
590 Board (IRB) approvals, if applicable? [N/A]
- 591 (c) Did you include the estimated hourly wage paid to participants and the total amount
592 spent on participant compensation? [N/A]