Learning to Rank Features to Enhance Graph Neural Networks for Graph Classification

Anonymous authors
Paper under double-blind review

Abstract

A common strategy to enhance the predictive performance of graph neural networks (GNNs) for graph classification is to extend input graphs with node- and graph-level features. However, identifying the optimal feature set for a specific learning task remains a significant challenge, often requiring domain-specific expertise. To address this, we propose a general two-step method that automatically selects a compact, informative subset from a large pool of candidate features to improve classification accuracy. In the first step, a GNN is trained to estimate the importance of each feature for a given graph. In the second step, the model generates feature rankings for the training graphs, which are then aggregated into a global ranking. A top-ranked subset is selected from this global ranking and used to train a downstream graph classification GNN. Experiments on real-world and synthetic datasets show that our method outperforms various baselines, including models using all candidate features, and achieves state-of-the-art results on several benchmarks.

1 Introduction

Graph neural networks (GNNs) (see, e.g., Hamilton, 2020) have become a widely adopted approach for learning from graph-structured data, particularly for node and graph classification tasks. While GNNs have shown remarkable success in scenarios where rich feature information is available, their effectiveness often suffers when such features are absent or difficult to define. This is especially true for graph classification tasks, where learning effective graph representations depends on both node- and graph-level features.

Feature selection is therefore a critical challenge in learning with GNNs. Traditional approaches often rely on predefined static feature sets (Barceló et al., 2021; Bouritsas et al., 2023; Cui et al., 2022; Duong et al., 2019), which require substantial domain expertise and fail to generalize across different datasets. More recent work on GNN node classification has explored dynamic feature selection methods (Naik et al., 2024), offering flexibility but focusing on local graph patterns and neglecting global properties that could be crucial for certain tasks. Accordingly, there is a clear demand for scalable, data-driven methods capable of identifying the most informative local and global features, regardless of the underlying graph structure.

Addressing these limitations, the recent work by Alkhoury et al. (2025) proposed a feature-ranking GNN framework that learns to select relevant features for node classification on unseen graphs, thereby improving predictive accuracy and reducing computational cost. However, this framework is limited to node classification and cannot be directly adapted to graph classification tasks.

In this paper, we build upon this framework by introducing a novel method for automatic, scalable feature selection for *graph classification*. Our approach proceeds in two major steps. First, given a set of training graphs, we sample a small, random subset of the data and compute all features from a universal pool for each graph in this subset, resulting in feature vectors of uniform dimensionality. These vectors, along with their corresponding class labels, are then used for feature ranking. To generate these rankings, we considered three different approaches: random forests, support vector machines, and local explanation methods from explainable AI. Among them, random forests consistently delivered the most reliable feature rankings. Following the framework of Alkhoury et al. (2025), we then train a feature-ranking GNN (FR-GNN) using the original graphs and their feature rankings. This enables the FR-GNN to predict feature importance for

all remaining graphs in the training set without the need to compute any features. The local rankings are aggregated into a single global ranking, which is then applied to all training graphs in the second step and to unseen graphs later during class prediction.

In the second step, the top K features from the global ranking are selected and computed for all training graphs. A graph classification GNN (GC-GNN) is then trained using these graphs, augmented with only the K selected features. Since both feature selection and ranking rely on a small subset of data and a limited number of features, the overall process remains computationally feasible.

To assess the effectiveness of our approach, we conducted experiments on both real-world and synthetic datasets, utilizing 124 local and global graph features as well as three GNN architectures (GCN (Kipf & Welling, 2016), GAT (Velickovic et al., 2018), and GraphSAGE (Hamilton et al., 2017)), and compared our method to three baseline approaches.

On the real-world datasets, using only the top six features selected by our method ($K = 6$), the trained GC-GNN consistently outperformed the baseline that used no features, demonstrating the importance of incorporating features. Furthermore, our method surpassed the baseline that used all 124 features by an average margin of 7–11%, suggesting that including irrelevant or redundant features can lead to overfitting and degrade performance. Additionally, it consistently outperformed GC-GNNs trained on six randomly selected features, indicating that the features chosen by our method are genuinely informative rather than selected by chance. Notably, for the real-world datasets where state-of-the-art (SOTA) algorithm results are publicly available, our method either outperforms or achieves comparable predictive performance.

To further evaluate robustness, we generated synthetic graphs using various random graph models with increasing classification difficulty. This was achieved by attaching graph motifs (graphlets) and defining target classes through logical formulas of varying complexity, where literals indicate the presence or absence of specific motifs. Our method consistently outperformed the baselines that used no features or random features, even on the most challenging tasks. While the baseline using all features achieved high accuracy ($\geq 95\%$) on synthetic data, our method using just six selected features often performed slightly better. In one particularly complex case involving four classes defined by nested logical conditions, the all-features baseline outperformed our approach; however, increasing our selection to the top 20 features ($K = 20$) narrowed the performance gap. This result confirms that our method remains competitive even as task complexity increases.

The rest of the paper is organized as follows. Section 2 overviews the related work. Section 3 describes the proposed method. The experimental results are reported in Section 4. Finally, in Section 5 we conclude and mention some problems for future work.

2 Related Work

Feature initialization strategies have been extensively studied to enhance GNN performance in graph classification. A common approach is the graphlet kernel (Shervashidze et al., 2009), which measures graph similarity based on small subgraphs (graphlets) counts. While polynomial-time graphlet enumeration for small graphs is possible, it becomes computationally infeasible for large graphs. Efficient subgraph estimation techniques (Rossi et al., 2018) approximate graphlet counts without requiring full enumeration. Unlike these methods, our approach learns a compact, dataset-specific set of informative features such as graphlet counts, structural properties, and statistical summaries without relying on exhaustive enumeration or estimation.

Real-world graphs often exhibit noise and sparsity (Dai et al., 2022) which may contain redundant or irrelevant connections, making it difficult to extract meaningful patterns. By selecting only the most informative features, our method mitigates the impact of such noise, enabling the downstream graph classification model to train on a compact and effective feature set.

Several works employ dual-network architectures to improve feature selection and classification. Akyol et al. (2021) use two GNNs within a variational graph autoencoder for manipulation action recognition and prediction. Similarly, our method leverages two GNNs: one performs feature ranking, thereby enhancing the discriminative power of the second GNN used for graph classification. Maurya et al. (2023) propose a

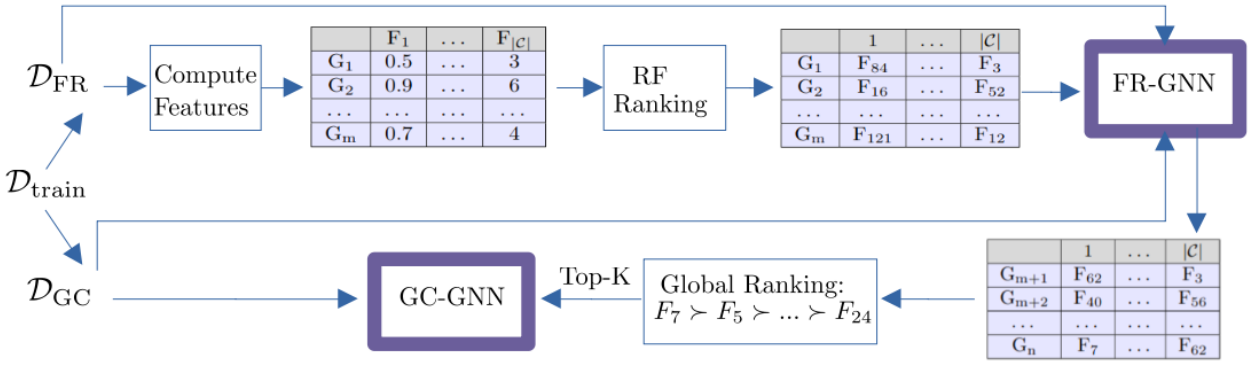

Figure 1: High-level overview of the FR-GNN training step (top) and the GC-GNN training step (bottom).

Dual-Net GNN for node classification, where the selector identifies optimal node features for the classifier. While both networks are trained jointly, our method separates the feature ranking and classification stages, and focus on graph-level (not node-level) features. DualNetGO (Chen & Luo, 2024) also uses a dual-network design for protein function prediction, training a selector and classifier alternately to optimize feature selection. Unlike our method, DualNetGO is tailored to biological datasets and starts with randomly combined features. Other approaches aim to learn graph-level representations for classification tasks. Fei & Huan (2008) select frequent subgraphs based on spatial distribution consistency, while we use subgraph counts as explicit features. HGP-SL (Zhang et al., 2019) combines hierarchical graph pooling with structure learning, implicitly refining graph topology via learned node importance scores. In contrast, our approach follows a feature-driven strategy: we rank and select informative graph-level features prior to training using an interpretable model, namely random forests. Cui et al. (2022) evaluate the impact of different hand-crafted features on GNN performance without ranking or selecting features, unlike our approach.

A closely related recent work by Alkhoury et al. (2025) proposes a framework to improve predictive performance in node classification by identifying important features using a random forest-based selection mechanism. In contrast, our approach focuses on graph classification and explores both node-level and graph-level features, requiring entirely different techniques. In particular, while the feature-ranking GNN classifier in (Alkhoury et al., 2025) is trained on synthetic datasets designed to exhibit structural diversity, our method operates directly on the training graphs of the target dataset, ensuring that the feature ranking reflects the specific patterns relevant to the classification task at hand. Another key difference lies in the aggregation strategy: whereas their method produces a separate feature ranking for each input graph (as needed for node classification), our approach aggregates feature importance scores across all training graphs, resulting in a dataset-level feature ranking suitable for graph classification.

3 The Method

This section presents our method[1] for automatically selecting a compact subset of features from a large pool of candidates, with the goal of improving the predictive performance of a GC-GNN using the selected additional features. Our method consists of two steps (see Fig. 1 for a schematic overview of the process).

(Step 1) Training a Feature Ranking GNN (FR-GNN): Given a set of training graphs $\mathcal{D}_{\text{train}}$, we first select a small random set $\mathcal{D}_{\text{FR}} \subseteq \mathcal{D}_{\text{train}}$, compute a feature ranking for each graph in $\mathcal{D}_{\text{FR}}$ using random forests, and train an FR-GNN on these graphs and their corresponding feature rankings. This enables the FR-GNN to learn to predict feature importance rankings for unseen graphs.

(Step 2) Training a Graph Classification GNN (GC-GNN): We then take the remaining training data $\mathcal{D}_{\text{GC}} = \mathcal{D}_{\text{train}} \setminus \mathcal{D}_{\text{FR}}$, use the trained FR-GNN to predict a feature ranking for each graph in

[1]Code is available at github.com/FouadAlkhoury/GraphClassificationRankingFeatures.

Table 1: Overview of the 124 features used in the experiments, grouped into three categories: graphlets, aggregated node features, and global graph properties.

Category	**Description**
Graphlet Features (62)	31 count features and 31 binary indicator features for all connected, non-isomorphic subgraphs (graphlets) of 2-5 nodes, plus the hexagon
Aggregated Node Features (56)	Statistics (mean, std, skewness, kurtosis) for degree, centrality measures (eigenvector, closeness, harmonic, and betweenness), largest first coloring, number of edges within egonet, node clique number, number of cliques, clustering coefficient, square clustering coefficient, pagerank, hubs, core number
Global Graph Properties (6)	number of nodes, number of edges, diameter, avg. and std. of path length, random feature

$\mathcal{D}_{\text{GC}}$, and aggregate these individual rankings into a global ranking. The top K features from this global ranking are then selected, computed for all graphs in $\mathcal{D}_{\text{GC}}$, and used to augment them with corresponding feature vectors of length K. A GC-GNN is subsequently trained on these augmented graphs to classify unseen graphs using only the most informative features.

We discuss these two steps in detail in Sections 3.2 and 3.3 below. Before that, in Section 3.1, we first discuss the types of candidate graph-level features used in our experimental evaluation.

3.1 Candidate Features

The success of our approach depends, among other factors, on the appropriate selection of a *universal* pool of candidate features. This pool is not tailored to a specific task. Instead, it should capture a broad range of global graph statistics and structural properties, which can either be intrinsic to the graph or derived by aggregating node-level features. Capturing such structural properties is crucial for the quality of the trained model, especially when graphs from the same target class are expected to share common patterns that distinguish them from graphs of other classes. Aggregating node-level features into graph-level representations allows us to embed graphs into a common feature space, thereby facilitating effective graph classification.

Graphlet Features These features are computed directly over the entire graph and are used in both binary and counting forms. They encode the presence (binary form) or frequency (counting form) of small subgraphs, known as *graphlets*, such as, for example, triangles or 4-cliques. The counts of these graphlets capture how often characteristic structural patterns appear as subgraphs, providing highly indicative information about the graph's properties and potentially capturing structural differences between graphs from different classes.

Aggregated Node Features In contrast to graphlets, aggregated node features (e.g., node degree, betweenness centrality) are computed at the node level before being summarized to form graph-level features. To derive graph-level features from a particular node feature for a graph G, we compute statistical summaries characterizing the feature's distribution across all nodes of G. Specifically, we use the following four descriptive statistics for aggregation: (1) *mean*, (2) *standard deviation*, (3) *skewness*, defined by $\frac{n}{(n-1)(n-2)} \sum_{i=1}^{n} \left(\frac{x_i - \bar{x}}{s}\right)^3$, measuring the asymmetry of the distribution, and (4) *kurtosis*, defined as $\frac{1}{n} \sum_{i=1}^{n} \left(\frac{x_i - \bar{x}}{s}\right)^4 - 3$, quantifying the peakedness or flatness of the distribution relative to the normal distribution. Here, n denotes the number of nodes in G, x_i the node feature value for node i, $\bar{x}$ the mean of the node feature values, and s their standard deviation.

Global Graph Properties These graph-level features provide high-level summaries of a graph's connectivity and topology. Examples of such global properties include the number of nodes and the average shortest path length.

Algorithm 1 Learning Feature Ranking GNN

Input: set $\mathcal{D}_{\text{FR}}$ of labeled graphs, candidate feature set $\mathcal{C}$
Output: FR-GNN model Φ_{FR}

1: **for all** $G \in \mathcal{D}_{\text{FR}}$ **do**
2: compute the feature vector $\vec{x}_G$ defined by $\mathcal{C}$ for G
3: $\mathcal{D} \leftarrow \emptyset$
4: **for all** $G \in \mathcal{D}_{\text{FR}}$ **do**
5: $\mathcal{G}_G = \{G' : G' \in \mathcal{D}_{\text{FR}} \text{ and } \textsc{Class}(G') \neq \textsc{Class}(G)\}$
6: **for** $i = 1, \ldots, t$ **do**
7: select a number k uniformly at random from $\{3, 4, 5\}$
8: select a random subset $\mathcal{G}'$ of $\mathcal{G}_G$ with $|\mathcal{G}'| = k$
9: train a random forest F_i for the training set $\{(\vec{x}_G, +)\} \cup \{(\vec{x}_{G'}, -) : G' \in \mathcal{G}'\}$
10: compute a ranking $\pi_{G,i}$ of $\mathcal{C}$ from their "importance" in F_i
11: compute an aggregated ranking π_G of $\mathcal{C}$ from $\{\pi_{G,1}, \ldots, \pi_{G,t}\}$
12: add (G, π_G) to $\mathcal{D}$
13: $\Phi_{\text{FR}} \leftarrow$ learn a a feature ranking GNN model from $\mathcal{D}$
14: **return** Φ_{FR}

Altogether, these three feature types yield 124 candidate features (see Table 1). Given the diversity of graph datasets, the importance of these features can vary significantly across graphs. The key challenge is to *automatically* identify a small subset for the task at hand that is most informative for the majority of training graphs.

3.2 Step 1: Training the FR-GNN Model

The first step of our method (see Algorithm 1 and the top part of Fig. 1) begins by selecting a random subset $\mathcal{D}_{\text{FR}} \subseteq \mathcal{D}_{\text{train}}$. In the experiments reported in Section 4, we consistently used 10% of the graphs in $\mathcal{D}_{\text{train}}$.

For each graph in $\mathcal{D}_{\text{FR}}$, the method computes all features from the pool $\mathcal{C}$ of candidate features. Node-level features are then aggregated into graph-level representations, as described in Section 3.1. In this way, each graph is mapped to a feature vector of dimension $|\mathcal{C}|$ (see the top-left matrix in Fig. 1 and lines 1–2 of Alg. 1).

To compute a feature ranking specific to each graph in $\mathcal{D}_{\text{FR}}$, we adopt a binary classification strategy based on random forests (see the top-right feature ranking matrix in Fig. 1 and lines 5–11 of Algorithm 1). More precisely, for all $G \in \mathcal{D}_{\text{FR}}$, we sample a set of graphs from $\mathcal{D}_{\text{FR}}$ that have a different target label than G. In our experiments, the number of these graphs randomly varies between 3 and 5. This sampling strategy ensures that the random forest learns to distinguish between graphs with different target labels, facilitating the identification of discriminative features. The training data for the random forest classifier consists of the graph-level feature vectors extracted from the sampled graphs, where the feature vector of G is assigned the label '+', and those of the other graphs are assigned the label '–'. A random forest classifier with T trees ($T = 128$ in the experiments) is then trained on this data, using $\sqrt{|\mathcal{C}|}$ candidate features at each split. Feature importance scores are computed using the Gini gain, which measures the total reduction in node impurity due to splits on that feature, normalized over all trees in the forest.

To ensure robustness, this procedure is repeated t times ($t = 10$ in the experiments), each time with a newly sampled set of graphs from $\mathcal{D}_{\text{FR}}$. The final ranking π_G for graph G is obtained by aggregating the importance scores across the t independent runs, resulting in a feature ranking for G (see the top-right matrix in Fig. 1, which contains the individual rankings for each graph).

This process is repeated for all graphs in $\mathcal{D}_{\text{FR}}$, producing a dataset, denoted $\mathcal{D}$ in the algorithm, that consists of training examples of the form (G, π_G) for all $G \in \mathcal{D}_{\text{FR}}$, where π_G is the feature ranking computed for G as described above. The algorithm then trains an FR-GNN on $\mathcal{D}$ and returns the resulting model Φ_{FR} (see lines 13–14 of Algorithm 1 and the top-right box in Fig. 1).

Remark During preliminary experiments, we observed that the feature importance vectors produced by random forests often contain many values close to zero. These low scores can introduce unnecessary complexity into the Φ_{FR} model. To address this, we introduce a dynamic thresholding strategy that adaptively identifies and removes low-importance features, retaining only those that make a clear contribution to the classification task. Specifically, for each graph, the feature importance values are first sorted in ascending order. We then compute the differences between consecutive values, defined as $\Delta_i = f_{i+1} - f_i$. If Δ_i exceeds the importance value f_i itself, we set f_i and all smaller values to zero. This adaptive mechanism effectively filters out features whose contribution is negligible, eliminating the need for a manually defined threshold. The remaining non-zero values correspond to features that exhibit a sufficiently sharp increase in importance relative to the preceding values, signaling their discriminative potential. These final rankings are then used as input to the Φ_{FR} model.

3.3 Step 2: Training the GC-GNN Model

Given the feature learning model Φ_{FR} obtained in the previous step, we apply it to each graph G in the subset $\mathcal{D}_{\mathrm{GC}} = \mathcal{D}_{\mathrm{train}} \setminus \mathcal{D}_{\mathrm{FR}}$ to obtain a predicted feature ranking for G. These individual rankings are then aggregated to derive a global ranking of the candidate features. From this aggregated ranking, the top-K features are selected and subsequently computed for the graphs in $\mathcal{D}_{\mathrm{GC}}$. We again stress that this is the step where the benefits of our method become apparent: it computes only the top-K features from the pool $\mathcal{C}$ for each training graph in $\mathcal{D}_{\mathrm{GC}}$, rather than calculating all features in $\mathcal{C}$. Using the training set formed by the labeled graphs in $\mathcal{D}_{\mathrm{GC}}$, augmented with these top-K features, the method trains a GC-GNN model to classify unseen graphs (see the bottom pipeline in Fig. 1). If the training graphs in $\mathcal{D}_{\mathrm{train}}$ are originally associated with node-level features, these features are used throughout the process and aggregated in the same way as the node-level features in $\mathcal{C}$ (see Section 3.1).

4 Experimental Evaluation

In this section, we experimentally evaluate the performance of the proposed method using real-world benchmark and synthetic datasets. The experiments are designed to address the following research questions:

(Q1) To what extent do features improve the predictive performance of GNNs for graph classification? If they do, how much improvement can be expected?

(Q2) How does the predictive performance of GC-GNN using the feature set selected by our method compare to using all features in the pool? How robust is this performance as the classification task becomes increasingly difficult?

(Q3) Are the K features selected by our method genuinely relevant to the learning task, or does a randomly selected subset of K features perform similarly?

4.1 Datasets

To address the above questions, we conducted experiments on a diverse collection of real-world and synthetic datasets. This section provides a detailed description of the datasets used.

Real-World Datasets To evaluate our method, we selected a diverse set of nine real-world graph classification benchmarks spanning the domains of molecular chemistry, bioinformatics, and computer vision. For the molecular datasets, we considered MUTAG (Debnath et al., 1991; Kriege & Mutzel, 2012), BZR (Sutherland et al., 2003), DHFR (Sutherland et al., 2003), COX2 (Sutherland et al., 2003), NCI1 (Wale et al., 2008; Kim et al., 2025; Shervashidze et al., 2011), PTCMM and PTCFR (Helma et al., 2001; Kriege & Mutzel, 2012). These datasets consist of graphs where nodes represent atoms and edges denote chemical bonds, with target labels indicating biological activity. Specifically, MUTAG, BZR, DHFR, and COX2 focus on classifying molecules based on mutagenicity, activity against the benzodiazepine receptor, enzyme inhibition, and cyclooxygenase-2 activity, respectively. NCI1 targets activity against human non-small cell lung cancer, while PTCMM and PTCFR predict rodent carcinogenicity in male mice and female rats, respectively. The

PROTEINS dataset (Borgwardt et al., 2005; Dobson & Doig, 2003) from bioinformatics consists of protein structures, where nodes represent secondary structure elements and edges denote spatial proximity between them. The classification task in PROTEINS is to distinguish enzymes from non-enzymes. Finally, the MSRC9 dataset (Neumann et al., 2016) from computer vision, derived from image segmentation, represents segmented images as graphs, with nodes and edges corresponding to superpixel relationships, and labels indicating object classes. Dataset statistics are summarized in Table 2 in Appendix B.

Synthetic Datasets To evaluate our method's ability to detect structural features defining target classes, we conducted experiments on synthetic graph datasets with known ground truth. For each learning problem defined below, we generated four datasets, each with 600 connected base graphs from one of four standard random graph models: Barabási–Albert (BA) (Barabási & Albert, 1999), Erdős–Rényi (ER) (Erdős & Rényi, 1959), Watts–Strogatz (WS) (Watts & Strogatz, 1998), and Power-Law (PL) cluster (Holme & Kim, 2002).

We defined the target classes using four distinct graphlet motifs: the 4-clique (G_8), the path $v_1 - v_2 - v_3 - v_4 - v_5$ on five vertices with chords v_1v_4 and v_2v_5 (G_{20}), the house graph (G_{21}), and the hexagon (G_{30}) (see Fig. 6 in Appendix A). The target classes were constructed by applying logical expressions of varying complexity to the presence or absence of these motifs in the graphs. The base graphs were generated to exclude all four motifs, and graphlets were attached so that each graph satisfies exactly one class condition. Note that the motifs are mutually exclusive (none is a subgraph of another).

For each base graph type, we considered five learning tasks, with target classes defined by the following logical formulas ($\overline{G}_i$ below denotes the absence of motif G_i):

$$\begin{aligned}\Phi_1 : &(\textsc{Class} = c_1 \iff G_8 \wedge \overline{G}_{20} \wedge \overline{G}_{21}) \wedge \\ &(\textsc{Class} = c_2 \iff \overline{G}_8 \wedge G_{20} \wedge \overline{G}_{21}) \wedge \\ &(\textsc{Class} = c_3 \iff \overline{G}_8 \wedge \overline{G}_{20} \wedge G_{21}) \wedge \\ &(\textsc{Class} = c_4 \iff \overline{G}_8 \wedge \overline{G}_{20} \wedge \overline{G}_{21}) \end{aligned} \tag{1}$$

$$\begin{aligned}\Phi_2 : &(\textsc{Class} = c_1 \iff G_8 \wedge G_{20} \wedge \overline{G}_{21}) \wedge \\ &(\textsc{Class} = c_2 \iff G_8 \wedge \overline{G}_{20} \wedge G_{21}) \wedge \\ &(\textsc{Class} = c_3 \iff \overline{G}_8 \wedge G_{20} \wedge G_{21}) \end{aligned} \tag{2}$$

$$\begin{aligned}\Phi_3 : &(\textsc{Class} = c \iff G_8 \wedge G_{20} \wedge G_{21}) \wedge \\ &(\textsc{Class} = \overline{c} \iff \overline{G}_8 \vee \overline{G}_{20} \vee \overline{G}_{21}) \end{aligned} \tag{3}$$

$$\begin{aligned}\Phi_4 : &(\textsc{Class} = c_0 \iff \overline{G}_8 \wedge \overline{G}_{20} \wedge \overline{G}_{21}) \wedge \\ &(\textsc{Class} = c_1 \iff (G_8 \wedge \overline{G}_{20} \wedge \overline{G}_{21}) \vee (\overline{G}_8 \wedge G_{20} \wedge \overline{G}_{21}) \vee (\overline{G}_8 \wedge \overline{G}_{20} \wedge G_{21})) \wedge \\ &(\textsc{Class} = c_2 \iff (G_8 \wedge G_{20} \wedge \overline{G}_{21}) \vee (G_8 \wedge \overline{G}_{20} \wedge G_{21}) \vee (\overline{G}_8 \wedge G_{20} \wedge G_{21})) \wedge \\ &(\textsc{Class} = c_3 \iff G_8 \wedge G_{20} \wedge G_{21}) \end{aligned} \tag{4}$$

$$\begin{aligned}\Phi_5 : &(\textsc{Class} = c_1 \iff (G_8 \vee G_{20}) \wedge G_{21} \wedge G_{30}) \wedge \\ &(\textsc{Class} = c_2 \iff (G_8 \vee G_{20} \vee G_{21}) \wedge \overline{G}_{30}) \wedge \\ &(\textsc{Class} = c_3 \iff \overline{G}_{21} \wedge G_{30}) \wedge \\ &(\textsc{Class} = c_4 \iff \overline{G}_8 \wedge \overline{G}_{20} \wedge ((G_{21} \wedge G_{30}) \vee (\overline{G}_{21} \wedge \overline{G}_{30}))) \end{aligned} \tag{5}$$

In Φ_1, the four classes are pairwise disjoint and determined by the presence of exactly one motif or none, corresponding to a one-hot encoding over the three motifs. Φ_2 defines three classes, each by the presence of exactly two motifs, making the problem more challenging due to feature overlaps. Φ_3 is a concept learning task where the target class requires the presence of all three motifs, highlighting the difficulty for random forests with the Gini index, which evaluates features individually and is therefore less sensitive to feature interactions. In Φ_4, four classes are defined by the presence of exactly i motifs ($i = 0, 1, 2, 3$). Φ_5 introduces a fourth motif and higher logical formula depth, testing robustness as class definitions become more complex.

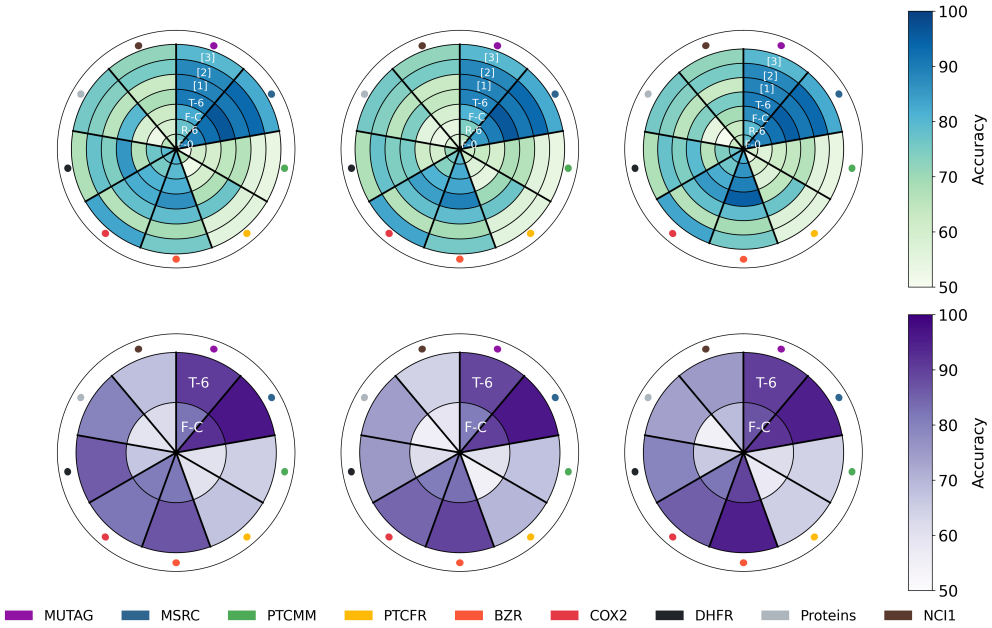

Figure 2: Circular heatmap diagrams summarizing the results from Table 3 in Appendix B the real-world datasets across three GNN architectures (left: GCN, middle: GAT, right: GraphSAGE). In each top diagram, performance is shown as concentric rings, ordered from innermost to outermost: F-0, R-6, F-$\mathcal{C}$, T-6, with three methods achieving SOTA results (Akyol et al., 2021; Vincent-Cuaz et al., 2022; Zhang et al., 2019), visually denoted in the diagrams by markers [1], [2], and [3], respectively. Each diagram is divided into nine major circular sectors, corresponding to the datasets (color-coded). The bottom diagrams make the performance difference between F-$\mathcal{C}$ and our method T-6 more pronounced.

4.2 Experimental Setting

For the experiments, we use three widely adopted GNN architectures: Graph Convolution Networks (GCN) (Kipf & Welling, 2016), Graph Attention Networks (GAT) (Velickovic et al., 2018), and GraphSAGE (Hamilton et al., 2017). The dataset is split such that 10% of the graphs are used to train the FR-GNN, while the remaining 90% are further divided 70/30 for training and testing the GC-GNN. We report the average accuracy over 5 independent runs. For feature ranking, for each graph $G \in \mathcal{D}_{\text{FR}}$, $k \in \{3, 4, 5\}$ graphs with different labels are sampled. Feature importance is computed via a random forest with 128 trees and $\sqrt{|\mathcal{C}|}$ features per split.

All GNN models are optimized using Adam, with hyperparameters selected via grid search (learning rate of 0.005, dropout rate of 0.5). Experiments are run on an AMD Ryzen 9 5950X 16-core CPU @ 3.40GHz with 125 GB of memory. The PyTorch Geometric library (Fey & Lenssen, 2019) is used for data handling.

For all experiments, except those involving synthetic data with target classes defined by Eq. (5), we consistently used $K = 6$, i.e., selected the top six features. (We discuss the rationale for this choice in Sect. 4.3.1.) To address questions Q1–Q3, we compare the proposed method, denoted as T-6 (short for "top 6"), against three baselines F-0, F-$\mathcal{C}$, and R-6. These baselines correspond to using no features (cf. question Q1), all features (cf. question Q2), and six features selected randomly from $\mathcal{C}$ (cf. question Q3), respectively.

4.3 Experimental Results

The results are presented as circular heatmaps in Figs. 2 and 3, for the real-world and synthetic datasets, respectively. For the real-world datasets, results are shown for all three GNN architectures. For the synthetic datasets, we report only the GCN results in the main figure, as all three architectures exhibit similar trends; full results for all architectures are provided in Figs. 7–9 in Appendix A.

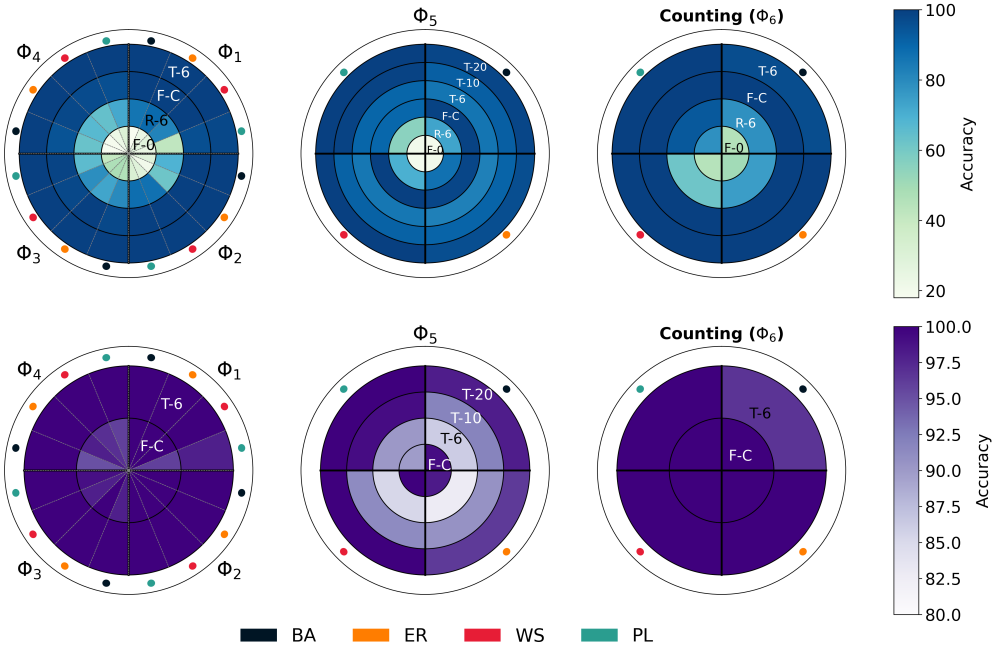

Figure 3: Circular heatmaps visualizing the results from Tables 4–6 in Appendix B for the synthetic datasets using GCN. The left heatmap displays performance for problems Φ_1–Φ_4, the middle for Φ_5, and the right for the count-based problem (Φ_6). In each top heatmap, concentric rings (from innermost to outermost) represent: F-0, R-6, F-$\mathcal{C}$, and T-6 (left and right); and F-0, R-6, F-$\mathcal{C}$, T-6, T-10, and T-20 (middle). The left heatmap is divided into four major sectors corresponding to Φ_1–Φ_4, each subdivided into four color-coded segments for the different random graph models. The middle and right heatmaps are divided into four color-coded segments, each representing one of the four graph models. The bottom heatmaps highlight differences between F-$\mathcal{C}$ and T-6 (and also T-10, T-20 in the middle) using a narrowed color scale for improved clarity.

Answer to Q1: The results on the synthetic data provide a clear answer to this question (see the left and middle heatmaps in Fig. 3 and Tables 4 and 5 in Appendix B): baselines that utilize all features (F-$\mathcal{C}$), or even just six randomly selected features (R-6), significantly outperform the baseline that uses no features (F-0). Across the synthetic experiments, we observe improvements of up to 80% for F-$\mathcal{C}$ and up to 70% for R-6 relative to F-0.

The results for real-world datasets are more mixed (see Fig. 2 and Table 3 in Appendix B). For all three GNN architectures, F-0 is outperformed by both F-$\mathcal{C}$ and R-6 on some datasets (e.g., MUTAG), while on others (e.g., DHFR), it actually outperforms the other two baselines. These results suggest that using either too many features or a small number of randomly selected features can, in some cases, diminish predictive performance.

Furthermore, our method T-6 consistently outperforms F-0 across all real-world and synthetic datasets and all three GNN architectures, with average improvements of about at least 50% on the synthetic and about 8.4% for GCN, 10.3% for GAT, and 8.4% for GraphSAGE on the real-world datasets. On DHFR, where all three GNNs using no features perform better than when using either all or six random features, T-6 achieves a 6% improvement over F-0 with GCN. The improvement is even more dramatic on MUTAG, where T-6 outperforms F-0 by around 23% with GAT.

In summary, our answer to question Q1 is conditionally affirmative: features can substantially enhance predictive performance, as long as they are not overly numerous or randomly selected.

Answer to Q2: The accuracy results in Fig. 2 (see also Table 3 in Appendix B) clearly show that T-6 consistently outperforms F-$\mathcal{C}$ on real-world datasets across all three GNN architectures. Specifically, our method achieves average accuracies of 80.2% (GCN), 78.9% (GAT), and 80.4% (GraphSage)

across nine real-world benchmarks, compared to F-$\mathcal{C}$'s accuracies of 71.9% (GCN), 69.6% (GAT), and 73.3% (GraphSage). This corresponds to average improvements of 8.3%, 9.3%, and 7.1%, respectively. The most significant improvement, around 20% across all GNN architectures, was observed on the Proteins dataset. These results support our conjecture that too many features can lead to *overfitting*.

In contrast, on synthetic datasets (see the left heatmap for Φ_1–Φ_4 and the middle for Φ_5 in Fig. 3 and Tables 4 and 5 in Appendix B), F-$\mathcal{C}$ achieves at least 90% accuracy on all five problems (Φ_1–Φ_5), across all four types of random graphs. Remarkably, for Φ_1–Φ_4, the six features selected by our method perform at least as well as F-$\mathcal{C}$, which attains nearly 96% accuracy, and sometimes even improve upon these results by up to 2%.

However, for Φ_5 (see the middle diagrams in Fig. 3, as well as Table 5 and Fig. 8 in the Appendix), which involves four classes and more complex logical formulas, our method with six features achieves lower predictive performance than F-$\mathcal{C}$. Nevertheless, as the number of selected features increases, the performance difference becomes negligible (see the comparison between F-$\mathcal{C}$ and T_{20}).

In summary, GC-GNNs using the six features selected by our method consistently outperform those using all features on real-world datasets, demonstrating improved accuracy and reduced overfitting. On synthetic datasets, our method matches or exceeds the performance of using all features for simpler tasks and remains robust, with only a negligible performance gap as task complexity and the number of selected features increase.

Answer to Q3: Regarding the question of whether the six features selected by our method (T-6) are genuinely relevant to the learning task, we compared the accuracy results obtained by T-6 with those obtained using six features selected uniformly at random from the candidate pool. A closer look at Figs. 2 and 3 (see, also, Tables 3–5) reveals that T-6 consistently outperforms R-6 across all datasets and all three GNNs. In particular, on the real-world datasets, T-6 outperforms R-6 on average by 10.9% (GCN), 11.7% (GAT), and 10.0% (GraphSAGE). In some cases, the difference is dramatic (e.g., R-6 achieves 52.5% (GCN) on Proteins, while T-6 achieves 80%).

In summary, based on these results, we can provide a clear affirmative answer to Q3: The feature set returned by our method is genuinely relevant to the learning task and not selected by chance.

4.3.1 Further Results on the Real-World Data

We now present and discuss additional results obtained on the real-world benchmark datasets.

Comparison to SOTA Results In Fig. 2 (see also the bottom three rows of Table 3 in Appendix B), we present results for three algorithms from (Akyol et al., 2021; Vincent-Cuaz et al., 2022; Zhang et al., 2019) that achieved state-of-the-art (SOTA) performance on at least one of the nine datasets and for which code is publicly available.[2] In Table 3 in Appendix B, these SOTA results are marked with a box. We evaluated each of the three algorithms not only on the dataset where it originally achieved SOTA, but also across all nine benchmarks. Our method, T-6, outperforms the SOTA result of (Akyol et al., 2021) on MSRC by approximately 5%, and performs comparably to (Vincent-Cuaz et al., 2022) on NCI1 and to (Zhang et al., 2019) on Proteins. Moreover, T-6 clearly outperforms all three algorithms in terms of average accuracy: across all nine datasets and architectures, T-6 achieves an average accuracy of 79.8%, compared to 75.9% for (Akyol et al., 2021), 73.1% for (Vincent-Cuaz et al., 2022), and 71.8% for (Zhang et al., 2019).

Statistical Significance We evaluated the statistical significance of the observed differences in predictive performance between T-6 and the three baselines (F-0, F-$\mathcal{C}$, and R-6) on the real-world datasets. For this purpose, we first applied the Kruskal-Wallis test. The null hypothesis (i.e., no performance differences) was rejected at $\alpha = 0.01$ (99% confidence level) with $p = 0.0031$, indicating strong evidence of performance variation across the methods. Given this result, we proceeded with Dunn's post-hoc test (Dunn, 1961) with Holm adjustment to identify pairwise differences. We found that T-6 significantly outperforms R-6 at the 99% confidence level, and achieves statistically significant improvements over both F-0 and F-$\mathcal{C}$ at the 95% confidence level. Thus, the improvements by T-6 are not only consistent but also statistically significant.

[2] github.com/gamzeakyol/GNet for (Akyol et al., 2021), github.com/cedricvincentcuaz/TFGW for (Vincent-Cuaz et al., 2022), and github.com/cszhangzhen/HGP-SL for (Zhang et al., 2019).

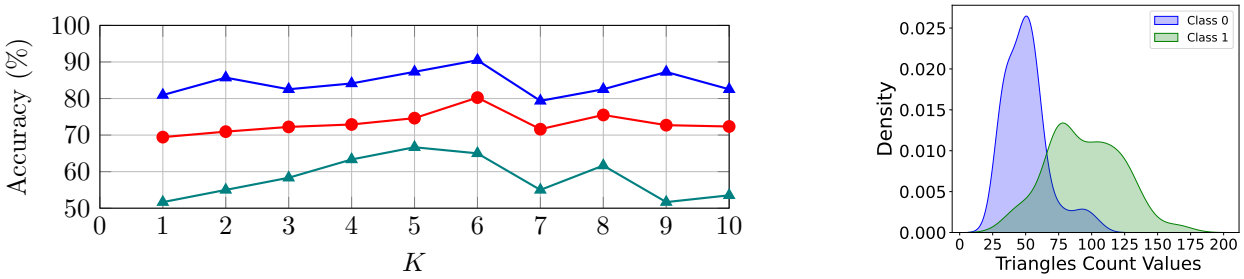

Figure 4: Left: Model accuracy of the top K features selected by our method for different values of K on two datasets: MUTAG (blue) and PTCMM (green), along with the average accuracy across all datasets using GCN (red). Right: Kernel density estimation (KDE) plots of the top-ranked feature for MUTAG (triangle count), showing the distribution of feature values across the two graph classes.

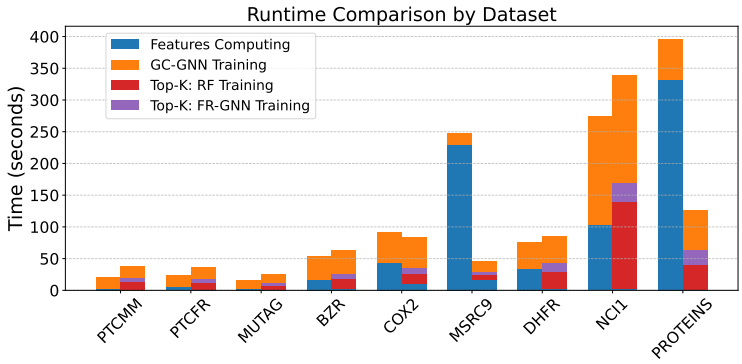

Figure 5: Runtime breakdown per dataset (sorted by size). Each dataset has two bars: the left for full feature setting (stacked: feature computation and GC-GNN training), and the right for the top-K setting (stacked: top-K computation, RF training, FR-GNN training, GC-GNN training).

Choice of K All experimental results presented in this paper, except those for the problems defined by Φ_5, were obtained using $K = 6$. This value was selected by computing the average accuracy for all values of $K = 1, \ldots, 10$ and choosing the value (i.e., 6) that yielded the highest *average* performance. In Fig. 4 (left), we present not only the overall average accuracies but also the accuracy results for different K values on the MUTAG and PTCMM datasets. While the optimal value of K is 6 for MUTAG, it is 5 for PTCMM. We note that K can alternatively be selected *automatically* for each dataset individually, by using a small validation subset of the training data and choosing the smallest value of K that yields the highest predictive performance on this subset.

Feature Analysis To further validate our feature selection approach, we conducted a qualitative analysis of several top-ranked features identified by our method for the real-world datasets. Specifically, we examined the distributions of these features across graph classes using kernel density estimation (KDE), which allowed us to assess their discriminative power: The selected features exhibit meaningful and class-discriminative patterns. As an illustrative example, in the MUTAG dataset, the triangle count emerged as one of the top-ranked features (see Fig. 4, right). The KDE plot shows that class 1 graphs consistently exhibit higher triangle counts, suggesting that the triangle graphlet structure is a strong indicator of mutagenic activity and is effectively captured by our selection method.

Runtime We also evaluated the time required to identify the top-K features on real-world datasets, including both random forest and FR-GNN training (see Fig. 5). For small graphs, the overhead of training our model exceeds the time needed to compute the full feature set. However, as graph size increases, computing all

features becomes significantly more expensive. In these cases, our approach achieves substantial speedups by avoiding full feature computation, while maintaining—or even improving—predictive accuracy.

4.3.2 Further Results on the Synthetic Data

Other Feature Ranking Approaches We evaluated the performance of three methods: random forests (RF), support vector machines (SVM) (Cortes & Vapnik, 1995), and neural networks coupled with the LIME local explainer (Ribeiro et al., 2016) (NN+LIME), in identifying key graphlet features across the five learning problems Φ_1–Φ_5 defined in Section 4.1. RF consistently identified the three key graphlets (G_8, G_{20}, and G_{21}) among the top-ranked features for the problems defined by Φ_1–Φ_4. Notably, other features, such as the total number of cliques and the node clique number, were also frequently selected as important. Upon investigation, we found these features to be semantically aligned with the identified graphlets, particularly since G_8 corresponds to a 4-clique. SVM exhibited similar, though slightly less consistent, behavior, with the key graphlets appearing lower in the ranking compared to RF. In contrast, NN+LIME could not consistently identify the relevant graphlets. For problem Φ_5, all three methods struggled to consistently rank the four relevant graphlets among the top features. However, RF demonstrated comparatively stronger performance, as the four features associated with Φ_5 were ranked higher than by SVM and NN+LIME.

We also compared the runtime of these three algorithms. NN+LIME required the longest training time, followed by RF, while SVM was the fastest but less accurate in identifying the most informative features. Since accuracy was our primary concern, we chose RF for all subsequent experiments.

The Count-Based Problem We also explored a learning problem (denoted Φ_6) where the class label was determined by the *number* of graphlets from G_8, G_{20}, and G_{21} present in the graph. Specifically, graphs containing at most one motif from this set were assigned to class 0, those containing two to four motifs to class 1, and graphs with five or more motifs to class 2. The results exhibited behavior similar to that observed in learning problems Φ_1–Φ_4 (see the right of Fig. 3, and Table 6 and Fig. 9 in the Appendix).

5 Concluding Remarks

We have proposed a method that automatically selects a small subset of relevant features from a large pool of candidate features to improve the predictive performance of GNNs for graph classification. Our experimental results on nine benchmark datasets and three GNN architectures clearly demonstrate that: (1) incorporating node features can significantly enhance predictive performance compared to vanilla GNNs, (2) to prevent overfitting, it is crucial to select a small, informative subset of features from a large candidate pool rather than using all available features, and (3) the features selected by our method are genuinely relevant to the learning task, rather than being chosen by chance.

We also evaluated the relevance of the selected features in a more controlled and precise setting. Specifically, we generated synthetic base graphs and defined target classes using logical formulas of increasing complexity, based on the presence or absence of a small set of motifs. Our results show that the approach can accurately identify relevant features and successfully learn these classification tasks. However, as the complexity of the formulas increases, more features must be selected to maintain high accuracy.

While our method demonstrates effective feature selection for improving GNN-based graph classification, several promising avenues remain for future exploration. First, scaling the approach to handle very large graphs is a natural next step. In our experiments, the FR-GNN was trained separately for each dataset. We also explored training a single, global FR-GNN across all datasets. Although this global model led to GC-GNNs with higher average accuracy than the three baselines, its performance gains were less pronounced than those achieved by individually trained FR-GNNs. Thus, another question is how to adapt and extend our approach to train a global FR-GNN that can produce task-specific GC-GNNs with predictive performance at least comparable to those reported here. Additionally, investigating the robustness of our method to noisy or adversarial features and developing theoretical guarantees for feature subset optimality would further deepen the understanding and reliability of our approach. Finally, it would be useful to modify the method to account for feature computation time, selecting the most informative features with the smallest cumulative runtime.

References

Gamze Akyol, Sanem Sariel, and Eren Erdal Aksoy. A variational graph autoencoder for manipulation action recognition and prediction. In *2021 20th International Conference on Advanced Robotics (ICAR)*, pp. 968–973. IEEE, 2021.

Fouad Alkhoury, Tamás Horváth, Christian Bauckhage, and Stefan Wrobel. Improving graph neural networks through feature importance learning. *Machine Learning*, 114(8):178, 2025. ISSN 1573-0565. doi: 10.1007/s10994-025-06815-z. URL https://doi.org/10.1007/s10994-025-06815-z.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439): 509–512, 1999.

Pablo Barceló, Floris Geerts, Juan L. Reutter, and Maksimilian Ryschkov. Graph neural networks with local graph parameters. In *35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pp. 25280–25293, 2021.

Karsten M Borgwardt, Cheng Soon Ong, Stefan Schönauer, SVN Vishwanathan, Alex J Smola, and Hans-Peter Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1):i47–i56, 2005.

Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45 (1):657–668, 2023.

Zhuoyang Chen and Qiong Luo. Dualnetgo: a dual network model for protein function prediction via effective feature selection. *Bioinformatics*, 40(7):btae437, 2024.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Mach. Learn.*, 20(3):273–297, 1995. doi: 10.1007/BF00994018. URL https://doi.org/10.1007/BF00994018.

Hejie Cui, Zijie Lu, Pan Li, and Carl Yang. On positional and structural node features for graph neural networks on non-attributed graphs. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 3898–3902, 2022.

Enyan Dai, Wei Jin, Hui Liu, and Suhang Wang. Towards robust graph neural networks for noisy graphs with sparse labels. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pp. 181–191, 2022.

Asim Kumar Debnath, Rosa L Lopez de Compadre, Gargi Debnath, Alan J Shusterman, and Corwin Hansch. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of medicinal chemistry*, 34(2):786–797, 1991.

Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4):771–783, 2003.

Olive Jean Dunn. Multiple comparisons among means. *Journal of the American statistical association*, 56 (293):52–64, 1961.

Chi Thang Duong, Thanh Dat Hoang, Ha The Hien Dang, Quoc Viet Hung Nguyen, and Karl Aberer. On node features for graph neural networks. *arXiv preprint arXiv:1911.08795*, 2019.

P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

Hongliang Fei and Jun Huan. Structure feature selection for graph classification. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pp. 991–1000, 2008.

Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *CoRR*, abs/1903.02428, 2019.

William L. Hamilton. *Graph Representation Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, San Rafael, 2020.

William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pp. 1024–1034, 2017.

Christoph Helma, Ross D. King, Stefan Kramer, and Ashwin Srinivasan. The predictive toxicology challenge 2000–2001. *Bioinformatics*, 17(1):107–108, 2001.

Petter Holme and Beom Jun Kim. Growing scale-free networks with tunable clustering. *Physical review E*, 65(2):026107, 2002.

S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, and E. E. Bolton. Pubchem 2025 update. nucleic acids res., 2025. PubChem.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016.

Nils Kriege and Petra Mutzel. Subgraph matching kernels for attributed graphs. *arXiv preprint arXiv:1206.6483*, 2012.

Sunil Kumar Maurya, Xin Liu, and Tsuyoshi Murata. Feature selection: Key to enhance node classification with graph neural networks. *CAAI Transactions on Intelligence Technology*, 8(1):14–28, 2023.

Harish Naik, Jan Polster, Raj Shekhar, Tamás Horváth, and György Turán. Iterative graph neural network enhancement via frequent subgraph mining of explanations. *CoRR*, abs/2403.07849, 2024.

Marion Neumann, Roman Garnett, Christian Bauckhage, and Kristian Kersting. Propagation kernels: efficient graph kernels from propagated information. *Machine learning*, 102:209–245, 2016.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you? explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144. ACM, 2016.

Ryan A Rossi, Rong Zhou, and Nesreen K Ahmed. Estimation of graphlet counts in massive networks. *IEEE transactions on neural networks and learning systems*, 30(1):44–57, 2018.

Nino Shervashidze, SVN Vishwanathan, Thomas Petri, Kurt Mehlhorn, and Karsten M Borgwardt. Efficient graphlet kernels for large graph comparison. In *Artificial intelligence and statistics*, pp. 488–495. PMLR, 2009.

Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.

Jeffrey J Sutherland, Lee A O'brien, and Donald F Weaver. Spline-fitting with a genetic algorithm: A method for developing classification structure- activity relationships. *Journal of chemical information and computer sciences*, 43(6):1906–1915, 2003.

Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018. URL `https://openreview.net/forum?id=rJXMpikCZ`.

Cédric Vincent-Cuaz, Rémi Flamary, Marco Corneli, Titouan Vayer, and Nicolas Courty. Template based graph neural network with optimal transport distances. *Advances in Neural Information Processing Systems*, 35:11800–11814, 2022.

Nikil Wale, Ian A Watson, and George Karypis. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14:347–375, 2008.

Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684): 440–442, 1998.

Zhen Zhang, Jiajun Bu, Martin Ester, Jianfeng Zhang, Chengwei Yao, Zhi Yu, and Can Wang. Hierarchical graph pooling with structure learning. *arXiv preprint arXiv:1911.05954*, 2019.

A Appendix: Figures

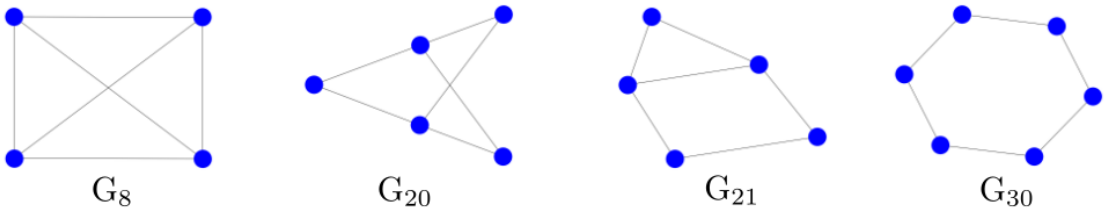

Figure 6: The four graphlets used in our experiments. They are pairwise incomparable (i.e., none is subgraph isomorphic to another).

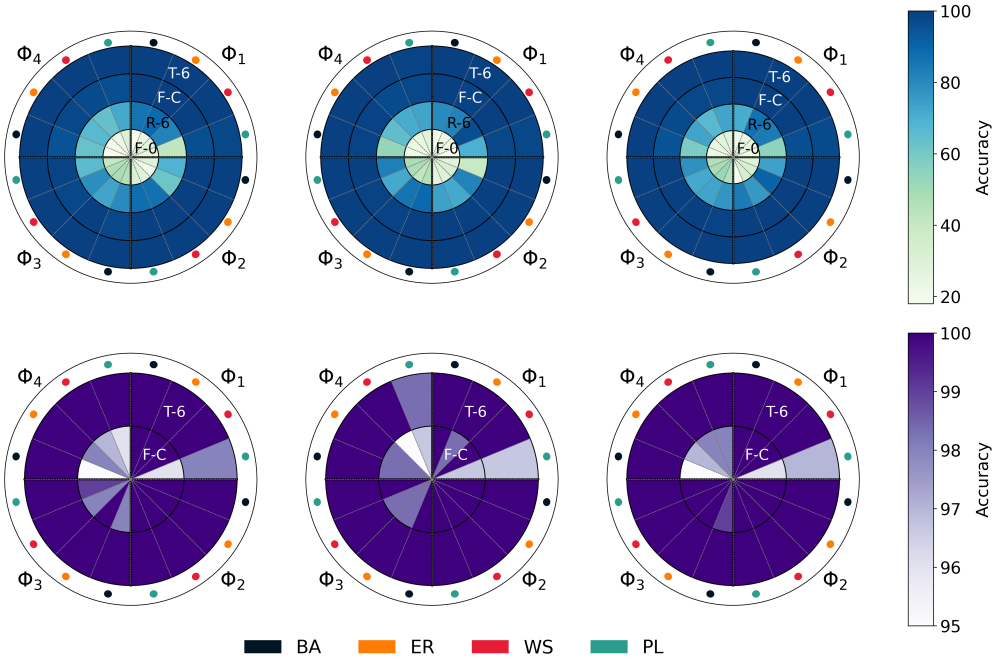

Figure 7: Circular heatmaps summarizing model performance on synthetic datasets for the four classification tasks Φ_1–Φ_4. The left, middle, and right panels show results with GCN, GAT, and GraphSAGE, respectively. In each top diagram, concentric annuli (from innermost to outermost) represent F-0, R-6, F-$\mathcal{C}$, and T-6. Each diagram is divided into four major circular sectors corresponding to the class definitions in Eq. (1)–(4), with each sector further subdivided into four color-coded subsectors representing the four random graph models. The bottom diagrams display the performances of F-$\mathcal{C}$ (innermost cycles) and T-6 (outermost annuli), using a color scale restricted to values between 95% and 100%.

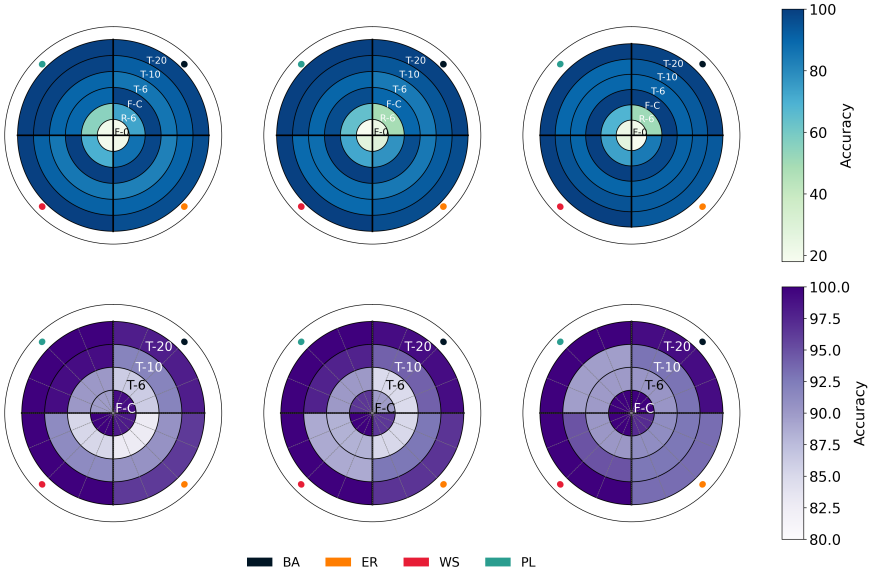

Figure 8: Circular heatmaps illustrating model performance on the synthetic dataset for the classification task Φ_5 defined by Eq. (5). Left: GCN results; middle: GAT results; right: GraphSAGE results. In each top diagram, results are shown in concentric annuli (from innermost to outermost): F-0, R-6, F-$\mathcal{C}$, T-6, T-10, and T-20. Each diagram is divided into four major circular sectors, corresponding to the four random graph models (indicated by color). The bottom diagrams show the performances of F-$\mathcal{C}$ (innermost cycles), T-6, T-10, and T-20 (outermost annuli).

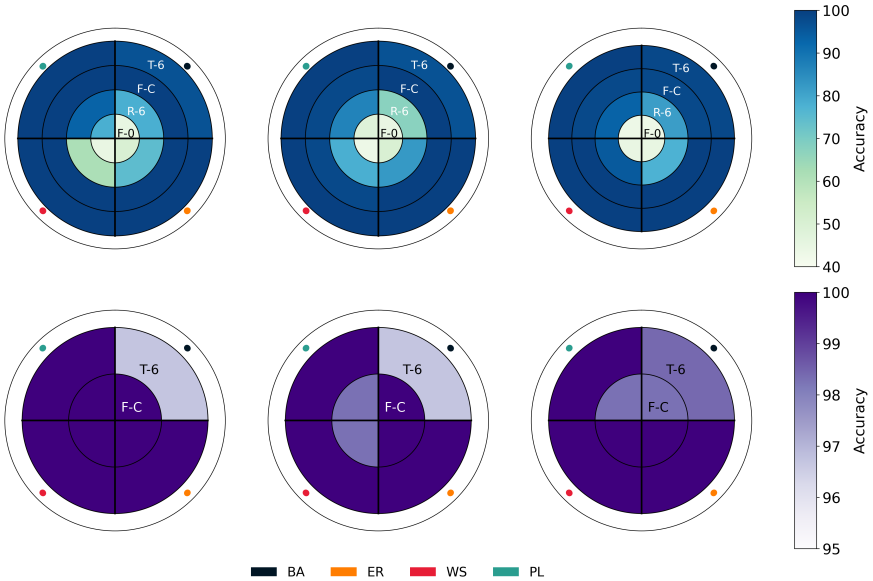

Figure 9: Circular heatmap diagrams for the classification task defined by the count-based learning problem (Φ_6). Left: GCN results; middle: GAT results; right: GraphSAGE results. In each top diagram, results are shown in concentric annuli (from innermost to outermost): F-0, R-6, F-$\mathcal{C}$, and T-6. Each diagram is divided into four major circular sectors, corresponding to the four random graph models (indicated by color). The bottom diagrams show the performances of F-$\mathcal{C}$ (innermost cycles) and T-6 (outermost ring).

B Appendix: Tables

Table 2: Summary of datasets used in our experiments. For each dataset, we provide the number of graphs, average nodes and edges per graph, number of node attributes, and number of graph classes.

dataset	#graphs	avg. #nodes	avg. #edges	#attributes	#classes
MUTAG	188	17.90	39.60	7	2
BZR	405	35.75	38.36	53	2
DHFR	756	42.43	44.54	53	2
COX2	467	41.22	43.45	35	2
NCI1	4110	29.87	32.30	37	2
PTCMM	336	13.97	14.32	20	2
PTCFR	351	14.56	15.00	19	2
Proteins	1113	39.10	145.60	3	2
MSRC9	221	40.58	97.94	10	8

Table 3: Classification accuracy (%) on real-world datasets using GCN, GAT, and GraphSAGE. Results are shown for our method (T-6) and the baselines (F-0, F-$\mathcal{C}$, and R-6). Bold indicates the best result per dataset within each GNN architecture, while underlining denotes the overall best result for each dataset. Results for the algorithms in (Akyol et al., 2021; Vincent-Cuaz et al., 2022; Zhang et al., 2019), which achieve SOTA performance on *some* of the nine benchmark datasets, are also presented for *all* nine datasets. SOTA results are marked with a box.

	MUTAG	MSRC	PTCMM	PTCFR	BZR	COX2	DHFR	Proteins	NCI1
					GCN				
F-0	77.0±1.1	88.7±5.3	50.0±4.7	66.7±3.7	80.7±1.1	77.7±2.2	79.5±0.3	65.5±4.9	60.2±0.6
F-$\mathcal{C}$	82.5±2.2	92.5±3.5	60.1±4.7	60.0±3.8	81.9±0.9	81.1±1.5	66.9±2.6	59.5±9.3	62.3±1.1
R-6	77.6±1.3	89.2±4.6	60.8±5.9	52.5±3.2	79.4±0.9	75.5±1.4	76.9±0.9	52.5±8.8	59.4±0.3
T-6	**90.5**±1.1	**96.3**±1.8	**65.0**±2.4	**67.5**±2.6	**86.9**±0.9	**82.2**±3.2	**85.9**±0.5	**80.0**±5.7	**68.0**±1.1
					GAT				
F-0	65.9±5.6	87.5±3.5	61.7±2.3	62.5±2.4	79.4±2.7	75.6±1.8	70.9±9.1	57.5±4.9	56.9±1.9
F-$\mathcal{C}$	81.1±3.6	90.0±7.1	60.0±9.4	55.0±7.6	83.1±4.4	81.1±3.3	61.6±7.5	55.7±9.9	59.1±1.7
R-6	74.7±1.3	85.0±0.2	63.1±0.3	57.5±3.3	82.5±2.4	73.3±2.5	60.3±4.0	53.7±9.8	54.9±5.8
T-6	**88.9**±1.8	**96.2**±5.3	**67.6**±3.5	**70.0**±3.2	**89.4**±0.9	**84.5**±1.2	**75.3**±1.3	**74.5**±2.1	**64.2**±3.4
					GraphSAGE				
F-0	78.5±3.3	91.2±1.7	60.3±5.1	52.5±4.5	81.9±6.2	80.1±4.4	77.1±5.0	62.1±0.1	64.2±1.1
F-$\mathcal{C}$	87.2±0.7	91.7±1.0	61.6±2.3	57.5±1.9	89.3±4.4	82.2±2.2	66.2±5.3	54.7±9.3	69.0±0.4
R-6	79.7±1.6	88.7±5.3	56.6±0.8	60.0±0.9	86.2±1.7	76.6±1.4	74.3±0.8	48.5±6.3	63.0±0.7
T-6	**90.4**±1.4	**95.0**±1.1	**64.1**±1.2	**65.0**±1.4	**95.0**±1.2	**85.5**±2.4	**79.6**±6.6	**74.0**±2.8	**75.1**±0.5
					Results of Algorithms Achieving SOTA Performance on Some Datasets				
(Akyol et al., 2021)	88.8±2.3	90.9±3.4	66.6±2.2	77.1±1.9	78.7±3.4	78.2±2.8	74.3±3.6	65.7±5.4	63.0±1.4
(Vincent-Cuaz et al., 2022)	88.2±1.5	93.2±3.1	56.2±1.9	57.5±2.3	69.2±3.7	66.6±1.1	77.9±2.2	73.8±4.7	75.2±0.8
(Zhang et al., 2019)	80.0±1.9	82.6±2.2	54.2±3.5	55.5±2.2	75.6±4.0	83.3±1.1	67.5±3.7	75.8±3.9	71.8±0.7

Table 4: Classification accuracy (%) of the four classification tasks defined by Eq. (1)–(4) on synthetic datasets generated using four graph models: Barabási–Albert (BA), Erdős–Rényi (ER), Watts–Strogatz (WS), and the Power-Law (PL) Cluster. Bold indicates the best result per graph model within each GNN architecture, while underlining highlights the overall best performance for each task.

	GCN				GAT				GraphSAGE			
	F-0	F-$\mathcal{C}$	R-6	T-6	F-0	F-$\mathcal{C}$	R-6	T-6	F-0	F-$\mathcal{C}$	R-6	T-6
						Φ_1 defined by Eq. (1)						
BA	18.2±7.1	**100**±0.0	86.1±1.2	**100**±0.0	19.9±6.4	**100**±0.0	80.3±1.1	**100**±0.0	22.4±5.3	**100**±0.0	70.6±2.1	**100**±0.0
ER	24.3±2.4	**100**±0.0	88.5±2.8	**100**±0.0	24.3±3.2	98.9±0.3	82.1±2.5	**100**±0.0	24.1±1.6	**100**±0.0	87.0±2.5	**100**±0.0
WS	27.6±7.8	**100**±0.0	82.2±6.4	**100**±0.0	25.4±6.9	**100**±0.0	96.2±1.1	**100**±0.0	26.3±4.5	**100**±0.0	84.1±4.1	**100**±0.0
PL	24.3±3.5	96.2±0.8	42.1±9.9	**98.3**±1.2	17.3±2.3	**98.2**±1.0	69.8±5.6	98.1±0.3	21.8±3.2	95.9±0.9	55.3±3.2	**97.3**±1.1
						Φ_2 defined by Eq. (2)						
BA	29.1±2.1	**100**±0.0	69.0±7.7	**100**±0.0	27.8±3.6	**100**±0.0	40.1±4.8	**100**±0.0	29.1±2.4	**100**±0.0	73.6±3.3	**100**±0.0
ER	20.8±11.3	**100**±0.0	61.8±3.6	**100**±0.0	33.3±7.4	**100**±0.0	82.3±2.7	**100**±0.0	32.8±5.6	**100**±0.0	90.5±3.9	**100**±0.0
WS	28.9±4.2	**100**±0.0	87.4±4.8	**100**±0.0	29.7±2.1	**100**±0.0	72.2±2.3	**100**±0.0	34.2±3.9	**100**±0.0	76.3±3.2	**100**±0.0
PL	31.7±1.7	**100**±0.0	88.1±1.3	**100**±0.0	27.3±1.8	**100**±0.0	73.6±1.3	**100**±0.0	17.3±2.1	**100**±0.0	84.2±1.7	**100**±0.0
						Φ_3 defined by Eq. (3)						
BA	46.2±2.1	98.2±0.6	80.1±5.7	**100**±0.0	45.9±1.8	**100**±0.0	89.4±4.5	**100**±0.0	46.3±2.7	99.0±0.1	77.0±2.5	**100**±0.0
ER	48.3±3.5	**100**±0.0	73±0.7	**100**±0.0	42.2±3.2	99.1±0.2	86.1±3.6	**100**±0.0	51.5±2.4	**100**±0.0	72.8±2.9	**100**±0.0
WS	46.7±1.7	98.3±0.7	79.1±2.1	**100**±0.0	48.2±2.2	99.0±0.2	71.9±3.2	**100**±0.0	48.6±3.9	**100**±0.0	71.9±3.0	**100**±0.0
PL	40.9±1.9	99.1±0.2	66.2±1.6	**100**±0.0	47.8±3.7	**100**±0.0	82.8±2.8	**100**±0.0	46.4±1.6	**100**±0.0	77.8±2.4	**100**±0.0
						Φ_4 defined by Eq. (4)						
BA	17.8±1.4	94.8±1.1	62.8±3.6	**100**±0.0	22.0±1.3	97.4±0.2	52.7±2.9	**100**±0.0	23.3±3.1	95.0±0.8	58.7±6.3	**100**±0.0
ER	18.8±8.4	97.7±0.4	66.6±3.5	**100**±0.0	22.8±1.5	98.0±0.1	65.0±3.1	**100**±0.0	25.2±2.0	97.1±1.1	72.9±4.1	**100**±0.0
WS	22.9±5.3	97.1±0.6	63.9±2.7	**100**±0.0	24.9±2.1	96.6±0.3	71.3±2.6	**100**±0.0	26.8±2.1	97.2±0.5	67.3±2.9	**100**±0.0
PL	30.9±2.1	96.2±1.3	72.1±1.9	**100**±0.0	29.4±2.3	98.2±0.2	73.1±2.1	**99.1**±0.2	29.1±3.1	98.2±0.2	73.2±2.3	**100**±0.0

Table 5: Classification accuracy (%) of the classification task defined by Eq. (5) on synthetic datasets generated using four graph models: Barabási–Albert (BA), Erdős–Rényi (ER), Watts–Strogatz (WS), and Power-Law (PL) Cluster. Bold indicates the best result per graph model within each GNN architecture, while underlining highlights the overall best performance.

	Φ_5 defined by Eq. (5)			
	BA	ER	WS	PL
		GCN		
F-0	20.2±3.4	21.1±4.1	19.8±2.7	20.8±2.3
F-$\mathcal{C}$	**99.2**±0.1	**98.3**±0.2	**100**±0.0	89.9±0.2
R-6	73.4±4.7	86.9±2.9	70.1±3.9	55.3±2.4
T-6	86.2±2.3	82.8±3.4	85.3±2.1	90.2±1.1
T-10	91.9±1.1	90.7±1.3	91.3±1.4	99.2±0.3
T-20	98.1±0.2	96.4±0.7	**100**±0.0	**100** ±0.0
		GAT		
F-0	23.1±3.2	29.2±3.1	27.4±1.6	17.7±1.9
F-$\mathcal{C}$	89.9±0.4	96.5±0.2	**100**±0.0	96.6±0.4
R-6	48.8±3.1	77.8±1.4	72.2±1.6	64.3±2.7
T-6	84.9±2.6	85.2±1.8	88.3±2.0	90.1±1.7
T-10	94.1±1.5	92.8±1.3	88.9±1.6	97.8±0.3
T-20	**99.2**±0.2	**96.7**±0.2	**100**±0.0	**100**±0.0
		GraphSAGE		
F-0	20.7±2.0	18.3±2.6	27.2±1.4	23.1±1.5
F-$\mathcal{C}$	**100**±0.0	**97.2**±0.6	**100**±0.0	**100**±0.0
R-6	50.8±2.6	84.6±2.9	73.8±3.1	68.8±2.0
T-6	90.4±2.5	91.1±3.1	90.3±2.8	90.0±1.1
T-10	93.0±1.2	92.8±2.7	94.9±0.7	89.7±0.2
T-20	99.2±0.2	93.4±1.4	**100**±0.0	**100**±0.0

Table 6: Classification accuracy (%) of the count-based learning problem on synthetic datasets generated using four graph models: Barabási–Albert (BA), Erdős–Rényi (ER), Watts–Strogatz (WS), and Power-Law (PL) Cluster. Bold indicates the best result per graph model within each GNN architecture, while underlining highlights the overall best performance.

	Φ_6 (the count-based problem)			
	BA	ER	WS	PL
		GCN		
F-0	45.3±1.7	50.1±3.3	44.7±1.9	78.3±3.5
F-$\mathcal{C}$	**100**±0.0	**100**±0.0	**100**±0.0	**100**±0.0
R-6	78.3±2.9	74.8±1.8	61.7±2.0	93.3±2.1
T-6	96.7±1.1	**100**±0.0	**100**±0.0	**100**±0.0
		GAT		
F-0	44.7±2.6	50.3±1.4	43.3±1.9	48.3±1.5
F-$\mathcal{C}$	**100**±0.0	**100**±0.0	98.3±0.5	98.3±0.4
R-6	67.5±4.4	82.5±2.1	78.3±0.9	86.7±1.5
T-6	96.7±1.2	**100**±0.0	**100**±0.0	**100**±0.0
		GraphSAGE		
F-0	43.3±1.7	45.0±1.3	44.8±1.4	43.3±2.8
F-$\mathcal{C}$	98.3±0.3	**100**±0.0	**100**±0.0	98.3±0.2
R-6	81.7±3.5	77.5±1.7	95±1.2	93.3±1.5
T-6	**98.4**±1.3	**100**±0.0	**100**±0.0	**100**±0.0