

# A SIMPLE YET EFFECTIVE BASELINE FOR NON-ATTRIBUTED GRAPH CLASSIFICATION

**Chen Cai & Yusu Wang**

Department of Computer Science

Ohio State University

cai.507@osu.edu, yusu@cse.ohio-state.edu

## ABSTRACT

To better understand the machinery of various methods for the graph classification task, we develop a simple yet meaningful graph representation and explore its effectiveness and limitation. Interestingly, this degree based simple representation achieves similar performance as the state-of-the-art graph kernels and graph neural networks for non-attributed graph classification. Its connection to graph neural networks and Weisfeiler-Lehman kernel is also presented.

## 1 INTRODUCTION

Graph-type data are ubiquitous across many scientific fields. Recently, there has been a surge of approaches that aim to learn representations that encode structure information about the graph. On the high level, the methods can be categorized to be graph-kernel based, or graph neural network based. The problem remains challenging. For example, the design of graph kernel and graph neural networks are often influenced by the (sub-)graph isomorphism problem, where one aims to make sure that features of non-isomorphic graphs are likely different. On the one hand, graph isomorphism is computationally hard<sup>1</sup>. On the other hand, there is no guarantee that graph kernels/neural networks with the strongest expressive power in terms of differentiating non-isomorphic graphs will generalize the best.

In this paper, we are interested in better understanding existing approaches for graph classification tasks, in an attempt to gain a gradual understanding of their powers and limitations. As the first step, we develop a simple graph representation based on local information for non-attributed graphs, which we refer to as Local Degree Profile (LDP). One of our initial goals was to understand in which scenarios the simple summary shall *fail*.

Interestingly, this simple graph representation achieves similar performance as the state-of-the-art graph kernels and graph neural networks for non-attributed graph classification, and in fact, outperforms many existing more sophisticated representations. We report these results as we believe they make the following contributions:

- While our LDP graph representation is simple, it is intuitive and we show its competitive performance in graph classification for a range of graph datasets. This graph representation is computationally efficient (linear-time) to compute. We also provide a simple connection between our representation with the graph neural networks.
- We note that these observations are only for the task of graph classification; while existing methods (especially graph neural network based approaches) are often designed for a broader scope, including for node embedding and for link prediction purposes. The results are also likely biased due to the limited amount of benchmark datasets available, and thus do not form the basis to dismiss any existing graph classification methodology.

Nevertheless, the good performance of our simple graph representation raises concerns about the effectiveness of current benchmark datasets for evaluating different algorithms for non-attribute graph classification. It calls for the development of new, more comprehensive,

---

<sup>1</sup>The graph isomorphism testing itself has only very recently shown to be solvable in quasi-polynomial time Babai (2016). The subgraph isomorphism, on the other hand, is NP-complete.

benchmark datasets for the better evaluation of different graph learning methods as well as for more rigorous analysis of the power and limitation of graph representations.

## 2 A SIMPLE BASELINE FOR GRAPH CLASSIFICATION

### 2.1 OUR MODEL

We denote a graph as  $G(V, E)$  where  $V$  is the set of nodes and  $E$  is the set of edges in  $G$ . For each graph  $G$  we extract features for each node in the following way. For each node  $v \in G(V)$ , let  $DN(v)$  denote the multiset of the degree of all the neighboring nodes of  $v$ , i.e.,  $DN(v) = \{\text{degree}(u) | (u, v) \in E\}$ . We take five node features, which are  $(\text{degree}(v), \min(DN(v)), \max(DN(v)), \text{mean}(DN(v)), \text{std}(DN(v)))$ . In other words, each node feature summarizes the degree information of this node and its 1- neighborhood. We aggregate features over different nodes in the same graph by performing either a histogram or an empirical distribution function (edf) operation, i.e, mapping all node feature into a histogram or an empirical distribution.

We then repeat the same procedure for all five node features and concatenate all the feature vectors as the input feature for SVM classifier. For a fair comparison, we follow the convention in the graph kernel literature. We perform 10-fold cross validation ten times and report the average accuracy. Our model takes linear time to extract features. For simplicity, in the rest of the paper, we denote our baseline as **Local Degree Profile (LDP)**. For the ease of the presentation, we leave the details about hyper-parameters of our model, complexity comparison, and discussion on its variants and limitations in the appendix.

### 2.2 RELATION TO GRAPH NEURAL NETWORKS

Our baseline can be seen as a variant of GNN that is used to learn useful graph representations in an end-to-end manner. GNNs use the graph structure and node features  $X_v$  to learn a representation vector of a node,  $h_v$ , or the entire graph,  $h_G$ . Modern GNNs follow a neighborhood aggregation strategy, which involves iteratively updating the representation of a node by aggregating the representations of its neighbors. Formally, the  $k$ -th layer of a GNN is

$$a_v^{(k)} = \text{AGGREGATE}^{(k)}(\{h_u^{(k-1)} : u \in N(v)\}), h_v^{(k)} = \text{COMBINE}^{(k)}(h_v^{(k-1)}, a_v^{(k)}) \quad (1)$$

where  $h_v^{(k)}$  is the feature vector of node  $v$  at the  $k$ -th iteration/layer.  $h_v^{(0)} = X_v$  initially, and  $N(v)$  is a set of nodes adjacent to  $v$ . The choices of  $\text{AGGREGATE}^{(k)}(\cdot)$  and  $\text{COMBINE}^{(k)}(\cdot)$  in GNNs are crucial. A number of architectures for AGGREGATE have been proposed (Kipf & Welling, 2016; Hamilton et al., 2017).

For graph classification the READOUT function aggregates node features from the final iteration to obtain the entire graph’s representation  $h_G = \text{READOUT}(\{h_v^{(K)} | v \in G\})$ . READOUT can be a simple permutation invariant function such as summation or a more sophisticated graph-level pooling function (Ying et al., 2018; Zhang et al., 2018a).

For our baseline, AGGREGATION is a simple function that summarizes the statistics of neighboring degree distribution by computing min, max, average, and standard deviation. We do not introduce trainable weights. The COMBINE function is a simple concatenation and READOUT is either histogram or empirical distribution operation. The number of iteration is  $K = 2$  in our case. From this point of view, LDP captures the essential elements of GNN, and this may partially explain its effectiveness. However, it is quite surprising to us that without any learning we can still achieve results comparable to many GNNs. We suspect that the use of AGGREGATION or READOUT function may play an important role. Interestingly, in a recent paper, Xu et al. (2018) explore the expressiveness of different pooling strategies and conclude that sum pooling is more powerful than mean and max pooling. To leverage this, we introduce “sum(DN( $v$ ))” as an extra node feature in the hope of achieving better results. We did a preliminary experiment on REDDIT 5K and REDDIT 12K but only achieved marginal improvement (0.5 percent). This might be due to the fact that our baseline does not involve any learning. Thus in the reported results, this “sum(DN( $v$ ))” feature is not deployed.

**Discussion on comparison to Weisfeiler-Lehman Kernel** We can also see the similarity between LDP and the Weisfeiler-Lehman kernel/isomorphism test with two iterations. Both methods start from the local node feature and build new feature from the previous step through graph topology. However, a key difference between WL kernel and LDP is that the hashing step (label compression) of WL kernel **does not** necessarily capture the local similarities: even two nodes with very similar neighborhood could be mapped to totally different labels, and perturbing the edges by a small amount will lead to completely different hashing features. LDP instead uses the statistics of the degree distribution of local neighborhood that is more robust and able to capture the node similarity. We can see in the experiment section that after incorporating the features that able to capture similarity of graph structures, LDP outperforms WL kernel by a large margin on social network dataset.

### 3 EXPERIMENTAL RESULTS

**Datasets.** We collect commonly used graph datasets from existing graph classification literature as our benchmark data. For non-attribute graphs, we use the movie collaboration datasets IMDB BINARY and IMDB MULTI, the scientific collaboration data COLLAB, the social network datasets REDDIT BINARY, REDDIT 5K, and REDDIT 12K. For non-attribute graphs, we use the protein dataset ENZYMES, PROTEINS and D&D, chemical compounds datasets MUTAG, PTC, and NCI1. See appendix for more details.

**Experimental setup** All experiments are performed on a single Intel Xeon CPU E5-2630 v4@2.20GHz  $\times$  40 and 64GB RAM machine. We compare our baseline with 6 state-of-the-art graph kernels (Shervashidze et al., 2011; Yanardag & Vishwanathan, 2015; Zhang et al., 2018b; Verma & Zhang, 2017; Kriege et al., 2016): Weisfeiler-Lehman Kernel (WL), Graphlet kernel (GK), Deep Graph Kernel (DGK), RetGK, FGSD, Weisfeiler-Lehman optimal assignment kernel (WL-OA), and 5 graph neural networks (Niepert et al., 2016; Hamilton et al., 2017; Ying et al., 2018; Xu et al., 2018; Atwood & Towsley, 2016): PATCHYSAN (PSCN), GRAPHSAGE, DIFFPOOL, Graph Isomorphism Network (GIN), Diffusion-convolutional neural networks (DCNN), Deep Graph CNN (DGCNN). In particular, among above graph kernels, only FGSD and GK does not utilize label information for chemical/protein graphs, and for GNNs, the size of chemical/protein graphs are too small and their performance is not reported for attribute graphs in the most paper. All the result is taken from published paper except the baseline. The code is available on github.<sup>2</sup>

**Results.** The results are shown in Table 2-4. For non-attribute graphs, our results show that combining our simple degree-based features and the kernel machine, we can beat WL, GK, DGK kernels by a large margin consistently, and achieve the results on par with more recent graph kernels such as RetGK and WL-OA (no more worse than 2.5 percent) and perform even better on certain datasets. On average, we are slightly better than RetGK for non-attribute graphs. However, our method is much simpler and faster than all of the previous graph kernels. What is more, we find that even using the only linear kernel, which yielding our final model equivalent to local feature + linear SVM, performance on REDDIT BINARY, REDDIT 5K, and REDDIT 12K rarely degrades.

Comparing with GNN, although many accuracy data is not available, we can still clearly see that the state-of-the-art GNNs do not show significant improvement over our baseline. Since our model can be seen as a simple variant of GNN where no end-to-end learning is involved, we interpret current result more as a dataset problem instead of an algorithm problem: current benchmark is no longer suitable for evaluating different algorithms.

For chemical graphs, we can treat them as graphs without the label, and apply the same method. We can also introduce one more feature for each node, which is simply the node label (scalar) in the original data. Adding label information in this way is certainly not principled as our results are no longer invariant to label permutation. However, our goal in this paper is not to handle attribute graphs so we settle on this hack. As we can see, LDP is still quite good on MUTAG, PTC, and PROTEIN datasets that do not have rich label information. For graphs with rich labels such as ENZYME, DD, and NCI1, LDP and its variant is much worse than the other kernels that can better utilize label information. This also suggests that MUTAG, PTC, and PROTEIN are not sufficient

<sup>2</sup><https://github.com/Chen-Cai-OSU/LDP>

Table 1: Comparison with Graph Kernel for non-attribute graph classification. Baseline \* means using only linear SVM.

	WL	GK	DGK	RetGk	FGSD	WL-OA	LDP	LDP*
COLLAB	74.8	72.8	73.1	81.0	80.0	80.7	78.1	73.9
IMDB BINARY	70.8	65.9	67.0	71.9	71.0	-	75.4	67.7
IMDB MULTI	49.8	43.9	44.6	47.7	45.2	-	50.0	45.4
REDDIT BINARY	68.2	77.3	78.0	92.6	86.5	89.3	92.1	89.8
REDDIT 5K	51.2	41.0	41.3	56.1	47.8	-	55.9	54.2
REDDIT 12K	32.6	31.8	32.2	48.7	-	44.4	47.8	46.7
Average	57.90	55.45	56.03	66.33	-	-	66.55	62.95

Table 2: Comparison with graph neural networks for non-attribute graph classification. Baseline \* means using only linear SVM.

	PSCN	GRAPHSAGE	DIFFPOOL	GIN	DCNN	DGCNN	LDP	LDP*
COLLAB	72.6	68.25	75.50	80.2	52.1	73.7	78.1	73.9
IMDB BINARY	71.0	-	-	75.1	49.1	70.0	75.4	67.7
IMDB MULTI	45.2	-	-	52.3	33.5	47.8	50.0	45.4
REDDIT BINARY	86.3	-	-	92.4	-	-	92.1	89.8
REDDIT 5K	49.1	-	-	57.5	-	-	55.9	54.2
REDDIT 12K	41.3	42.24	47.04	-	-	-	47.8	46.7

to evaluate different methods. More discussion on the variants and limitations of the baseline is presented in the appendix.

Table 3: Comparison with other graph kernels for attribute graph classification. LDP + Label means adding the label as an extra node feature.

	WL	GK	DGK	PSCN	RetGk	FGSD	WL-OA	LDP	LDP + Label
MUTAG	84.4	81.6	87.4	89.0	90.3	92.1	84.5	90.1	90.3
PTC	55.4	57.3	60.1	62.3	62.5	62.8	63.6	61.7	64.5
ENZYME	53.4	-	53.4	-	60.4	-	59.9	35.3	40.9
PROTEIN	71.2	71.7	75.7	75.0	75.8	73.4	76.4	72.7	73.7
DD	78.6	78.5	-	76.2	81.6	77.1	79.2	75.5	77.1
NCII	85.4	62.3	80.3	76.3	84.5	79.8	86.1	73.0	74.3

## 4 CONCLUSION AND DISCUSSION

Most graph kernels aim to capture graph topology and graph similarity in the hope of improving classification accuracy. Our experiment suggests that this is not yet well-reflected on current benchmark datasets for non-attribute graphs. This calls for the construction of better and more comprehensive benchmark datasets for understanding and analyzing graph representations. On the other hand, we emphasize that in scenarios where both graph nodes and edges have certain associated features, proper handling labels can significantly improve the classification accuracy, as is shown clearly for the NCII dataset. Graph kernels have been rather effective in incorporating node/edge attributes. It will be an interesting question to see how to incorporate attributes in a more effective yet still simple manner in our graph representation.

Also, although there are large scale chemical graphs datasets available (Hachmann et al., 2011; Ruddigkeit et al., 2012), a benchmark dataset that contains many large graphs is still missing. We plan to create such benchmark dataset for future use. In general, while not addressed in this paper, we note that understanding the power and limitation of various graph representations, as well as

the types of datasets with respect to which they are most effective, are crucial, yet challenging and remain largely open.

#### ACKNOWLEDGEMENT

This work is partially supported by NSF under grants CCF-1740761 and RI-1815697. We thank Siyuan Ma, Han Fu and Shuaichen Chang for helpful comments and discussions.

## A RELATED WORK

### A.1 GRAPH KERNELS

There are various graph kernels, many of which explore the R-convolutional framework (Haussler, 1999). The key idea is to decompose a whole graph into small substructures and build graph kernels based on the similarities defined for these components. For example, graphlet kernels (Shervashidze et al., 2009) are based on small subgraphs up to a fixed size. Weisfeiler-Lehman graph kernels (Shervashidze et al., 2011) is based on subtree patterns. More recently, new methods have been proposed to compare graphs, which is done by quantifying the dissimilarity between the distributions of pairwise distances between nodes. Schieber et al. (2017) uses the shortest path distance, and Verma & Zhang (2017) uses the diffusion distance.

### A.2 GRAPH NEURAL NETWORK

Another way to tackle graph classification involves developing graph neural network (GNN). GNN broadly follows a recursive neighborhood aggregation(or message passing) scheme, where each node aggregates feature vectors of its neighbors to compute its new feature vector. Repeating the above procedure  $k$  times, a node is represented by its transformed feature vector. The representation of an entire graph can be obtained by pooling, for example, by summing the feature vectors over all nodes in the graph.

Many GNN variants with different neighborhood aggregation and graph-level pooling schemes have been proposed (Defferrard et al., 2016; Duvenaud et al., 2015; Hamilton et al., 2017; Li et al., 2015; Kearnes et al., 2016; Kipf & Welling, 2016; Velickovic et al., 2017; Ying et al., 2018). Empirically, these GNNs have achieved state-of-the-art performance in many tasks such as node classification, link prediction, and graph classification. However, the design of new GNNs is often based on empirical intuition, heuristics, and experimental trial-and-error. The theoretical understanding of the properties and limitations of GNNs is somewhat limited, although there exists very recent work starting to address this issue (Xu et al., 2018).

### A.3 BASELINE METHOD

We remark that there has been a few works that draw attention to the potential problems of widely used graph datasets. Specifically, Orlova et al. observe that even by using simple graph features such as the number of nodes, they could achieve similar classification performance on common benchmark datasets compared to early graph kernels. Their results show that we cannot solely rely on these data sets to show the performance of a graph kernel. However, their work focuses on attributed graphs, while our work focuses on non-attributed graphs. Very recently, in an independent work by Wu et al. (2019), the authors hypothesize that the nonlinearity between GCN layers is not critical. They remove the non-linearities and develop a simple network called Simple Graph Convolution. This network works well on text classification, semi-supervised user geolocation, relation extract, zero-shot image classification, and graph classification. While our work shares some similarities with theirs, our paper was developed independently.

## B DETAILS OF LDP

**Computational complexity:** In feature extraction, we only need to count the degree for each node and save the statistics of 1- neighborhood for each node. This can be done in  $O(E)$  time. To map  $V$  numbers into  $B$  bin takes  $O(V)$  time so the total complexity is  $O(E)$  time. This matches the lower bound of reading a graph. In comparison, we attach the table summarizing the complexity of computing various graph kernels, both exactly and approximately.<sup>3</sup>

<sup>3</sup>Deep graph kernel requires corpus generation to build the kernel and we do not find explicit computational complexity in the original paper so it is not listed in the table. The kernels in the table are WL (Weisfeiler-Lehman kernel), WL-OA (Weisfeiler-Lehman optimal assignment kernel), GK (Graphlet Kernel), RetGK (Graph Kernels based on Return Probabilities of Random Walks), MLG (Multiscale Laplacian Kernel), FGSD (family of graph spectral distances).

The more computationally expensive part of our algorithm is the SVM. For linear SVM it takes  $O(nd^2)$  time where  $n$  is the number of features and  $d$  is the dimension of features. For non-linear kernel it is  $O(n^2d)$ . Although there is an efficient algorithm available to approximate the feature map (Rahimi & Recht, 2008), we can still afford running the original algorithm in a relatively short time.

Table 4: Complexity of various graph kernels. With a slight abuse of notation,  $V$  is the number of nodes and  $E$  is the number of edges in the larger graph among two graphs. For WL,  $h$  is the number of iterations. For GK,  $d < V$  and  $k \in \{3, 4, 5\}$ . For RetGK, return probabilities of random walks need to be calculated before computing the kernel, which takes  $O(V^3 + (S + 1)V^2)$  exactly where  $S$  is the number of steps, and takes  $O(VSM)$  approximately where  $M$  is number of Monte Carlo simulations used for simulation of random walks.  $D$  is the number of random Fourier features and  $d$  is the dimension of input feature. For MLG,  $\tilde{V}$  is the number of sampled vertices and  $\tilde{V} < V$ . For FGSD,  $r$  the number of terms of polynomials used to approximate  $f(L)$ .

Complexity	WL/WL-OA	GK	RetGK	MLG	FGSD	LDP
Approximate	-	$O(Vd^{k-1})$	$O(D + d)$	$O(V^3)$	$O(rE)$	-
Worst-Case	$O(hE)$	$O(V^k)$	$O(V^2)$	$O(\tilde{V}^3)$	$O(V^3)$	$O(E)$

**Hyperparameter:** Below we describe the hyper-parameters of our model. They arise naturally when discretizing continuous node features on graphs for the down-stream classifier. In practice, they are robust and easy to tune.

*Bin size.* In our experiment, we map the neighborhood degree distribution into a different number of  $n$  bins of uniform width. We try different sizes from  $\{30, 50, 70, 100\}$  to discretize the distribution.

*Normalization.* There are two natural ways of normalization for our method. The first one is to normalize every graph separately so that the value represents the relative degree. The second one is to normalize the whole dataset by finding the largest degree value across all the graphs. In practice, we do not see the consistent advantage of one normalization over the other so we try both and pick the one gives best training accuracy.

*Empirical Distribution versus Histogram.* We try to represent node features over a graph by both a histogram or an empirical distribution because empirical distribution is more stable with respect with the particular choice of bin size. It indeed yields better results than the histogram on certain dataset so in experiments we treat this choice as a hyperparameter.

*Linear vs logarithmic scales.* The degree distribution of many real-life networks follows the power law, which is usually visualized by log-log scale. We try both log scale and linear scale and notice that log scale yields better results for REDDIT 5K and REDDIT 12K.

*SVM parameter.* The  $C$  parameter is selected from  $\{10^{-3}, 10^{-2}, \dots, 10^2, 10^3\}$  and the Gaussian bandwidth is selected from  $\{10^{-2}, 10^{-1}, 1, 10^1, 10^2\}$ .

*Remark.* In general, all those hyper-parameters except the SVM parameters are not sensitive. Fine tuning of the above hyper-parameters usually yields an improvement of about 2 percentage.

## C VARIANTS AND LIMITATIONS OF BASELINE:

Surprised by the performance of our baseline that is based on purely local node feature, one natural extension is to incorporate more sophisticated 1) node features and 2) edge features in the hope of capturing the local and global graph topology better and therefore improving the accuracy.

To test this, we have experimented adding other node features (also more expansive to compute) such as closeness centrality, Fiedler vector (the second smallest eigenvectors of graph Laplacian), and Ricci curvature (Lin et al., 2011) of graphs. Interestingly, we observe that none of the above features yields any significant improvement consistently across all datasets.

For edge features, we compute all-pair-shortest path distance and add the histogram of distance distribution along with degree-based features for small chemical graphs (For the social networks

dataset, it is too costly to compute all pair shortest path distance and we do not observe any improvement in the preliminary experiments). There is about 2 percentage improvement consistently over different datasets. <sup>4</sup> This indicates that our local method fails to capture more global information which is shown to be useful for chemical/protein classification <sup>5</sup>.

Table 5: Adding distance distribution improves the accuracy. Compared with FGSD and GK, which are the only two models that also does not use label information, we can see after adding all pair shortest path distance, we can match their result on all datasets except the NCI1 for FGSD. This indicates that using only local degree-based features is not enough for chemicals/protein classification

	FGSD	GK	LDP	LDP + distance
ENEZYME	-	-	35.3	37.2
PROTEIN	73.4	71.7	72.7	74.7
DD	77.1	78.5	75.5	77.5
NCI1	79.8	62.3	73.0	75.6

## D DATASETS DESCRIPTION

The statistics of the benchmark graph datasets used in the paper are reported in Table 6. We describe these datasets in detail in the next section.

### D.1 NON-ATTRIBUTED GRAPH DATASETS

**IMDB-BINARY** (Yanardag & Vishwanathan, 2015) is a movie collaboration dataset that consists of the ego-networks of 1,000 actors/actresses who played roles in movies in IMDB. In each graph, nodes represent actors/actress, and there is an edge between them if they appear in the same movie. These graphs are derived from the Action and Romance genres.

**IMDB-MULTI** (Yanardag & Vishwanathan, 2015) is generated in a similar way to IMDB-BINARY. The difference is that it is derived from three genres: Comedy, Romance, and Sci-Fi.

**REDDIT-BINARY** (Yanardag & Vishwanathan, 2015) consists of graphs corresponding to online discussions on Reddit. In each graph, nodes represent users, and there is an edge between them if at least one of them respond to the other’s comment. There are four popular subreddits, namely, IAmA, AskReddit, TrollXChromosomes, and atheism. IAmA and AskReddit are two question/answer based subreddits, and TrollXChromosomes and atheism are two discussion-based subreddits. A graph is labeled according to whether it belongs to a question/answer-based community or a discussion-based community.

**REDDIT-MULTI(5K)** (Yanardag & Vishwanathan, 2015) is generated in a similar way to REDDIT-BINARY. The difference is that there are five subreddits involved, namely, worldnews, videos, AdviceAnimals, aww, and mildlyinteresting. Graphs are labeled with their corresponding subreddits.

**REDDIT-MULTI(12K)** (Yanardag & Vishwanathan, 2015) is generated in a similar way to REDDIT-BINARY and REDDITMULTI(5K). The difference is that there are eleven subreddits involved, namely, AskReddit, AdviceAnimals, atheism, aww, IAmA, mildlyinteresting, Showthoughts, videos, todayilearned, worldnews, and TrollXChromosomes. Still, graphs are labeled with their corresponding subreddits.

<sup>4</sup>We ignore MUTAG and PTC for the reason that the datasets are too small, and that even LDP + Label can achieve the best result.

<sup>5</sup>For graphs, two vertex sets are called non-homometric if the multi-sets of distances determined by them are different. It is unknown whether there exists any distance metric under which two vertex sets of non-isomorphic graphs are always non-homometric; But it is easy to show that the shortest path distance does not satisfy the requirement: a cycle of four vertices and a triangle with a pendant edge are non-isomorphic but have the same multi-set of all pairwise shortest path distances, i.e.,  $\{1, 1, 1, 1, 2, 2\}$



Table 6: Statistics of the benchmark graph datasets

Datasets	graph #	class #	average_nodes #	average edges #	label #
MUTAG	188	2	17.93	19.79	7
PTC	344	2	14.29	14.69	19
ENZYME	600	6	32.63	64.14	3
PROTEIN	1113	2	39.06	72.82	3
DD	1178	2	284.32	715.66	81
NCI1	4110	2	29.87	32.30	37
IMDB BINARY	1000	2	19.77	96.53	-
IMDB MULTI	1500	3	13.00	65.94	-
REDDIT BINARY	2000	2	429.63	497.75	-
REDDIT 5K	4999	5	508.82	594.87	-
REDDIT 12K	12929	11	391.41	456.89	-

## D.2 GRAPHS WITH DISCRETE ATTRIBUTES

**MUTAG** (Debnath et al., 1991) consists of graph representations of 188 mutagenic aromatic and heteroaromatic nitro chemical compounds. These graphs are labeled according to whether or not they have a mutagenic effect on the Gram negative bacterium *Salmonella typhimurium*.

**PTC** (Helma et al., 2001) consists of graph representations of chemical molecules. In each graph, nodes represent atoms, and edges represent chemical bonds. Graphs are labeled according to carcinogenicity on rodents, divided into male mice (MM), male rats (MR), female mice (FM), and female rats (FR).

**ENZYMES** and **PROTEINS** (Borgwardt et al., 2005) consist of graph representations of proteins. Nodes represent secondary structure elements (SSE), and there is an edge if they are neighbors along the amino acid sequence or one of three nearest neighbors in space. The discrete attributes are SSE types. The continuous attributes are the 3D length of the SSE. Graphs are labeled according to which EC top-level class they belong to.

**DD** (Dobson & Doig, 2003) consists of graph representations of 1,178 proteins. In each graph, nodes represent amino acids, and there is an edge if they are less than six Angstroms apart. Graphs are labeled according to whether they are enzymes or not.

**NCI1** (Shervashidze et al., 2011) consists of graph representations of 4,110 chemical compounds screened for activity against non-small cell lung cancer and ovarian cancer cell lines, respectively.

## REFERENCES

- James Atwood and Don Towsley. Diffusion-convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1993–2001, 2016.
- László Babai. Graph isomorphism in quasipolynomial time. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 684–697. ACM, 2016.
- 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.
- 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.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pp. 3844–3852, 2016.
- Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4):771–783, 2003.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pp. 2224–2232, 2015.
- Johannes Hachmann, Roberto Olivares-Amaya, Sule Atahan-Evrenk, Carlos Amador-Bedolla, Roel S Sánchez-Carrera, Aryeh Gold-Parker, Leslie Vogt, Anna M Brockway, and Alán Aspuru-Guzik. The harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *The Journal of Physical Chemistry Letters*, 2(17):2241–2251, 2011.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- David Haussler. Convolution kernels on discrete structures. Technical report, Technical report, Department of Computer Science, University of California at Santa Cruz, 1999.
- Christoph Helma, Ross D. King, Stefan Kramer, and Ashwin Srinivasan. The predictive toxicology challenge 2000–2001. *Bioinformatics*, 17(1):107–108, 2001.
- Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8): 595–608, 2016.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Nils M Kriege, Pierre-Louis Giscard, and Richard Wilson. On valid optimal assignment kernels and applications to graph classification. In *Advances in Neural Information Processing Systems*, pp. 1623–1631, 2016.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- Yong Lin, Linyuan Lu, and Shing-Tung Yau. Ricci curvature of graphs. *Tohoku Mathematical Journal, Second Series*, 63(4):605–627, 2011.
- Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pp. 2014–2023, 2016.
- Yuliia Orlova, Morteza Alamgir, and Ulrike von Luxburg. Graph kernel benchmark data sets are trivial! *FEAST 2015: ICML Workshop on Features and Structures*.

- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pp. 1177–1184, 2008.
- Lars Ruddigkeit, Ruud Van Deursen, Lorenz C Blum, and Jean-Louis Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17. *Journal of chemical information and modeling*, 52(11):2864–2875, 2012.
- Tiago A Schieber, Laura Carpi, Albert Díaz-Guilera, Panos M Pardalos, Cristina Masoller, and Martín G Ravetti. Quantification of network structural dissimilarities. *Nature communications*, 8:13928, 2017.
- Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. Efficient graphlet kernels for large graph comparison. In *Artificial Intelligence and Statistics*, pp. 488–495, 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(Sep):2539–2561, 2011.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Saurabh Verma and Zhi-Li Zhang. Hunt for the unique, stable, sparse and fast feature learning on graphs. In *Advances in Neural Information Processing Systems*, pp. 88–98, 2017.
- Felix Wu, Tianyi Zhang, Amauri Holanda de Souza Jr, Christopher Fifty, Tao Yu, and Kilian Q Weinberger. Simplifying graph convolutional networks. *arXiv preprint arXiv:1902.07153*, 2019.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- Pinar Yanardag and SVN Vishwanathan. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1365–1374. ACM, 2015.
- Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*, 2018.
- Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2018a.
- Zhen Zhang, Mianzhi Wang, Yijian Xiang, Yan Huang, and Arye Nehorai. Retgk: Graph kernels based on return probabilities of random walks. *arXiv preprint arXiv:1809.02670*, 2018b.