

Formulating Node Labelling as Node Classification or Link Prediction in Different Graph Representations

Anonymous authors

Paper under double-blind review

Abstract

Message-passing Graph Neural Networks (GNNs) are increasingly used for predictive tasks on graphs. Much work has been done to improve GNN architectures, but how the actual data graph should be designed is not well studied. In this paper, we investigate how two different graph representations impact the performance of GNN models across datasets with varying characteristics grouped by homophily, heterogeneity, and number of labels per node. A unique phenomenon is that the same abstract predictive task of labelling nodes is formulated as a node classification problem on one representation and as a link prediction problem on the other. Our work is the first to blur the line between these two basic and fundamental tasks in graph learning. Our experiments on 12 real-world datasets suggest that different representations (and tasks) are optimal for different datasets, models, and hyperparameters. We derive empirical heuristics of choosing between the two and pave the way towards a criterion of choosing the optimal graph representations and towards formally understanding the interconnection between node classification and link prediction.

1 Introduction


Message-Passing Graph Neural Networks (GNNs) (Hamilton et al., 2018; Veličković et al., 2018; Kipf & Welling, 2017) have been proven to be effective for analyzing graph-structured data and are widely used in numerous domains (Zhou et al., 2021; Hamilton, 2020).

Two of the common tasks on large graphs are node classification (NC) and link prediction (LP). Identifying fraudulent accounts in a financial transaction network (Wang et al., 2023), assigning academic papers with subject categories in a citation graph (McCallum et al., 2000) and classifying types of businesses or products in a Web-based graph (Zeng et al., 2020) are typically defined as NC problems. Recommending friends and users to follow in a social network (El-Kishky et al., 2022) and predicting missing physical or chemical associations in a biomedical graph (Wishart et al., 2017; Szklarczyk et al., 2019) are typical LP problems. These tasks are formulated not only on homogeneous graphs, but also on real-world heterogeneous graphs that have rich node- and edge-type information as well as feature attributes. Several extensions of classical GNN models explicitly leverage heterogeneity (Wang et al., 2019; Chen & Chen, 2021) in an attempt to improve learning performance.


Although GNNs have shown good performance for NC and LP on a variety of graphs, there are also known issues such as over-smoothing (Rusch et al., 2023) and over-squashing (Topping et al., 2022) which limit the expressiveness of GNNs and consequently their performance. The root cause stems from the fact that the computation graph is exactly the same as the data graph, which allows GNNs to possess strong inductive bias. Rewiring methods (Topping et al., 2022; Arnaiz-Rodriguez et al., 2022) that usually involve adding edges or virtual nodes have been introduced. There are also bags of tricks (Wang et al., 2021; Chen et al., 2022) that incrementally improve GNN performance.

Despite all the progress, very little effort has been put into **how the data graph should be represented** in the first place. Existing work benchmarks on predefined NC and LP tasks on the given representations of datasets. However, there is no canonical way of constructing the data graph and there are numerous ways of representing the same abstract information. Consider the example of modeling "users tagging music

with genre" in a graph, one could define User and Music nodes, and assign Tag edges with properties being one-hot encoding of genre. An alternative might be having User, Music and Genre nodes, and assigning User-Music-Genre hyperedge (an edge connecting three nodes). The optimal representation (or schema) depends on the workload, downstream models, and predictive tasks.



(a) *The node classification representation:* A graph \mathcal{G}_1 where the specific information is stored as node features in every node.



(b) *The link prediction representation:* A graph \mathcal{G}_2 where the specific information is represented by nodes of different types connected with new edges.

Figure 1: Two ways of representing the same abstract graph \mathcal{G} . Colors represent different node and edge types. Letters represent information about nodes.

In this paper, we systematically compare two representations that are general and fundamental, as shown in Figure 1. When we have discrete properties about the nodes, we could represent them as labels on the nodes (Figure 1a), or as distinct nodes with edges connecting them (Figure 1b). This means that for a predictive task on the discrete properties, the task is formulated as NC using the data graph in Figure 1a while it becomes LP using the data graph in Figure 1b. Note that even the formal task differs; a phenomenon of these two representations. This implies that such a decision of picking the data graph early in the machine learning pipeline has profound impact, since entirely separate training and testing process, such as data splits and metrics, as well as model architectures, can be used and compared across NC and LP for the same end problem. Concretely, we investigate the following research questions:

- How do different representations of the same abstract graph affect the performance of a predictive task defined on it?
- What are the characteristics of a graph that determine whether a node classification representation or a link prediction representation is preferred?

Contributions. We give answers to the research questions by evaluating how well different GNNs perform across a wide variety of graphs using NC and LP representations in the transductive setting. We make three key contributions:

- We pose the basic and important question of how the abstract graph should be represented and highlight two general and fundamental representations where even the formal predictive task of NC and LP differs for the same end problem.
- We rigorously map out ways to fairly examine NC and LP representations by defining comparable data splits, evaluation metrics, negative sampling (as a unique step of LP) and model architectures.
- We conduct extensive experiments using datasets of varying characteristics based on homophily index, level of heterogeneity, and number of labels per node.

Our results show that the optimal representation varies for different datasets and depends on the specific GNN used. Although most of the existing benchmarks adopt the NC representation by default, we found that the LP representation indeed sometimes yields better results in fair comparisons. For example, we found that

GraphSAGE(Hamilton et al., 2018) is consistently better on the LP version of widely benchmarked ogbn-arxiv(Hu et al., 2020) and similarly GAT(Veličković et al., 2018) on LP versions of DBLP(multi)(Akujuobi et al., 2019) and CiteSeer(Giles et al., 1998). Our work implies that the question of constructing optimal data graphs shall not be ignored and simply using the default representations in benchmarks is insufficient. Our result provides the groundwork for future research towards a formal and general criterion of designing optimal graph representations and a new perspectives which unifies node classification and link prediction.

2 Preliminary

Definition 2.1 (Graph). A graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{T}, \mathcal{R})$ consists of a set of nodes \mathcal{V} , a set of node types \mathcal{A} , a set of edges $\mathcal{T} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$, and a set of relation types \mathcal{R} . Each node $v \in \mathcal{V}$ has one or more node types $\alpha(v) \subseteq \mathcal{A}$. Each edge $(v_i, r, v_j) \in \mathcal{T}$, has a relation type $r \in \mathcal{R}$ which specifies the type of relation between nodes v_i and v_j . Each node $v \in \mathcal{V}$ has a initial node embedding $z_v \in \mathbb{R}^d$ for some dimension d .

Features for all nodes are abbreviated as $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$. A homogeneous graph has $|\mathcal{A}| = |\mathcal{R}| = 1$. Let \mathcal{L} be a set of labels for the nodes. The partial labelling function $L : \mathcal{V} \rightarrow \mathcal{P}(\mathcal{L})$ assigns some of the nodes their labels. Let $\mathcal{V}_L \subseteq \mathcal{V}$ denote the subset of nodes that are labelled. If the graph is not multilabeled, then $|L(v)| = 1$ for all defined $v \in \mathcal{V}_L$.

Definition 2.2 (Transductive Node Classification). Given a graph \mathcal{G} with node features \mathbf{X} and a partial labelling function \mathcal{L} where $\mathcal{V}_L \subset \mathcal{V}$, we want to assign labels to the nodes in $\mathcal{V} \setminus \mathcal{V}_L$ by some $f|_{\mathcal{G}, \mathbf{X}, L} : \mathcal{V} \setminus \mathcal{V}_L \rightarrow \mathcal{P}(\mathcal{L})$.

Definition 2.3 (Transductive Link Prediction). Given a graph \mathcal{G} with node features \mathbf{X} and an incomplete set of edges $\mathcal{T} \subset \mathcal{V} \times \mathcal{R} \times \mathcal{V}$, we want to predict a new set of edges $\mathcal{T}' \subset (\mathcal{V} \times \mathcal{R} \times \mathcal{V}) \setminus \mathcal{T}$ by some $f|_{\mathcal{G}, \mathbf{X}} : \mathcal{T}' \rightarrow \{0, 1\}$.

Definition 2.4 (K -hop Homophily Index). The K -hop homophily index β_K of a graph is given by

$$\beta_K = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \frac{|u \in \mathcal{N}^K(u) : |L(u) \cap L(v)| > 0|}{|u \in \mathcal{N}^K(u)|}.$$

When $K = 1$ and the graph is single-labelled, our K -hop homophily index reduces to the common homophily index β that measures the proportion of nodes that share labels with their neighbours defined in Pei et al. (2020).

2.1 Graph Neural Network

A message passing GNN iteratively aggregates node embedding from a node’s neighbourhood and upate it’s embedding. These embeddings are initialized as some feature vectors of the nodes. The embeddings and learnable weights are optimised with some objective on the graph. Formally, the message passing step as

$$\mathbf{h}_v^{(k+1)} = \text{UPDATE}^{(k)}\left(\mathbf{h}_v^{(k)}, \text{AGGREGATE}^{(k)}(\{\mathbf{h}_u^{(k)}, \forall u \in \mathcal{N}(v)\})\right) \quad (1)$$

MEAN, MAX, MIN are commonly used for AGGREGATE and learnable non-linear MLPs for UPDATE. Note that the neighbourhood function $\mathcal{N}(v)$ does not have to include the full neighborhood of v . GraphSAGE(Hamilton et al., 2018) samples the neighbourhoods to improve model scalability. GAT(Veličković et al., 2018) learns attention coefficients between pairs of nodes. AGGREGATE being a set-function ensures the embeddings are permutation equivariant. Since the computation structure directly follows the underlying data graph, a strong inductive bias of the underlying graph is baked into the model.

The GNN formula could be applied to heterogeneous graphs \mathcal{G} by ignoring the node and edge type information. A natural extension that takes advantage of heterogeneity is to apply Equation 1 separately for each triplet type $\mathcal{A} \times \mathcal{R} \times \mathcal{A}$ as distinct channels and aggregate the messages at the node level. A pictorial illustration of the computation graph is given in Appendix A.1. HAN(Wang et al., 2019) and MAGNN(Fu et al., 2020) use predefined metapaths (list of node and relation types) to induce metapath-induced neighbourhoods as another way to explicitly leverage heterogeneity.

3 Node classification and link prediction on two representations

Among many different ways of representing \mathcal{G} for various predictive tasks, we will focus on studying the impact of using the NC representation (Figure 1a) and the LP representation (Figure 1b) for predictive tasks on nodes in the transductive setting.

Given an abstract graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{T}, \mathcal{R})$ and the labelling partial function $L : \mathcal{V} \rightarrow \mathcal{P}(\mathcal{L})$, we define our representations formally. For clarity, we use calligraphic font in Section 2 to denote the abstract graph and problems, and we use mostly default font in Section 3 whenever we refer to concrete data and computation graphs and formal tasks.

3.1 Description

Definition 3.1 (Node classification representation). $G_{NC} = (V, A, T, R)$ is our data graph.

Definition 3.2 (Link prediction representation). $G_{LP} = (V \cup \mathcal{L}, A \cup \{*\}, T \cup T', R \cup \{*\})$ is our data graph. \mathcal{L} is a new set of nodes disjoint from the original set, $V \cap \mathcal{L} = \emptyset$. $*$ is a new node type where $a(v) = *$ for all $v \in \mathcal{L}$. $*$ ' is a new edge type. $T' = \{(v, *, l) : v \in V, l \in \mathcal{L}, L(v) \ni l\}$.

Definition 3.3 (P_{NC} : Transductive node classification on G_{NC}). Given G_{NC} with node features \mathbf{X}_V and partial labelling function L , we want to learn $f|_{G_{NC}, X, \mathcal{L}} : V \setminus V_L \rightarrow \mathcal{L}$.

Definition 3.4 (P_{LP} : Transductive link prediction on G_{LP}). Given G_{LP} with node features $\mathbf{X}_{V \cup \mathcal{L}}$. We want to learn a boolean function $f|_{G_{LP}, X} : (V \times *' \times \mathcal{L}) \setminus T' \rightarrow \{0, 1\}$. The set $\{(v, *, l) : v \in V, l \in \mathcal{L}, f(v, *, l) = 1\}$ therefore represents the predicted edges.

The learned function $f|_{G_{NC}, X, L}$ from task P_{NC} using G_{NC} naturally answers the questions on assigning labels to unlabelled nodes in \mathcal{G} . We can interpret the output of $f|_{G_{LP}, X}$ as a function $\tilde{f}|_{G_{LP}, X} : \mathcal{V} \setminus V_L \rightarrow \mathcal{P}(\mathcal{L})$ where $\tilde{f}(v) \ni l \iff f(v, *, l) = 1$. $\tilde{f}|_{G_{LP}, X}$ therefore solves transductive node classification in Definition 2.2.

3.2 Graph Data Splits

Now that we have defined the two tasks P_{NC} and P_{LP} on two graphs G_{NC} and G_{LP} , we need to define how splits can be done so that trained models are generalisable (to assign labels for nodes). In order to measure the impact of different representations, we need to ensure that the splits are comparable.

The split for P_{NC} on G_{NC} is simply partitioning the train, validation, and test set where $V_{train} \cup V_{valid} \cup V_{test} = V_L$ where the random partitioning is determined by some seed. While the computation graph provided by G_{LP} is different from that of G_{NC} , we need to somehow ensure that the information available for training the models remains equivalent. This allows comparisons to be drawn purely on the representations of the information.

The split is less apparent for G_{LP} . The set of edges in each split are further divided into message-passing edges (for the computation graph) and supervision edges (for the objective function), both as information used by the learning method. Given any $V_{train}, V_{valid}, V_{test}$ for G_{NC} , we construct the set of edges for G_{LP} and their message-passing and supervision subsets as:

$$T'_{train} = \{(v, *, l) : L(v) \ni l \wedge v \in V_{train}\}, T'_{train_mp} \cup T'_{train_sup} = T'_{train}, T'_{train_mp} \cap T'_{train_sup} = \emptyset \quad (2)$$

$$T_{train_mp} = T'_{train_mp} \cup T, T_{train_sup} = T'_{train_sup} \quad (3)$$

$$T_{valid_mp} = T_{train_mp} \cup T_{train_sup}, T_{valid_sup} = \{(v, *, l) : L(v) \ni l \wedge v \in V_{valid}\} \quad (4)$$

$$T_{test_mp} = T_{valid_mp} \cup T_{valid_sup}, T_{test_sup} = \{(v, *, l) : L(v) \ni l \wedge v \in V_{test}\} \quad (5)$$

T'_{train_mp} and T'_{train_sup} in Equation 2 are created with some supervision ratio hyperparameter from T'_{train} . Note that as a phenomenon of the equivalence of P_{LP} and P_{NC} for solving the same end task on different graph representations G_{LP} , G_{NC} , a unique new step of message passing and supervision data splitting appears since the models solving P_{LP} use information from the graph topology for both the computation graph and the objective, whereas for P_{NC} , only the computation graph is based on the topology. An illustration of the equivalent splits is shown in Figure 2.

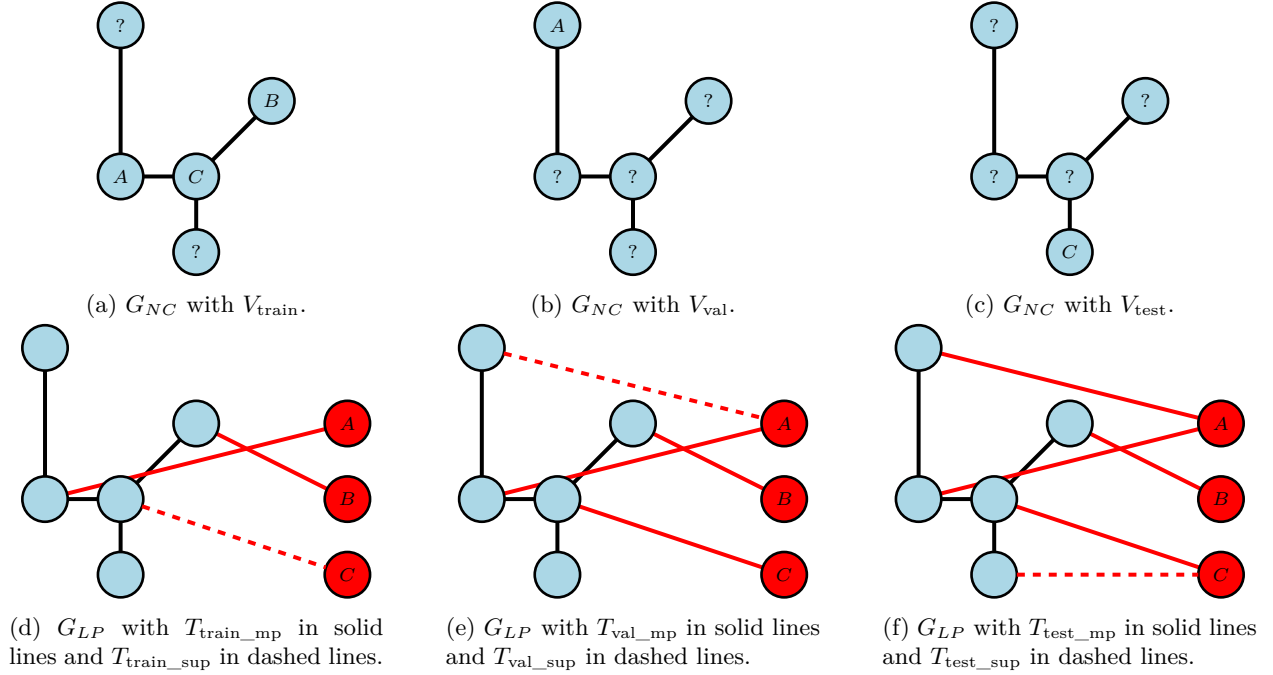


Figure 2: An example of equivalent splits for P_{NC} on G_{NC} and P_{LP} on G_{LP} .

3.3 Negative Sampling

The existence of supervision edges ensures a model generalises beyond seen edges and is able to predict new ones. However, without negative data present in the learning process, the models are likely to converge towards classifying all edges as positive. Since the graph topology is used explicitly as the computation graph for inductive bias in the architecture, as a unique advantage of GNNs, negative samples are normally utilised by the objective function (e.g binary cross-entropy).

There is a rich literature of negative sampling techniques (Nguyen & Fang, 2024; Ying et al., 2018; Yang et al., 2020) for the general link prediction task (of the abstract Definition 2.3) where heuristically "representative and difficult" node pairs are chosen as negative edges. However, for our P_{LP} on G_{LP} , we only want to learn a function $f|_{G_{LP}, X}$ over $(V \times *' \times \mathcal{L}) \setminus T'$ (i.e between the V nodes to the \mathcal{L} nodes) and not the entire $(V \cup \mathcal{L}) \times (R \cup \{*\}) \times (V \cup \mathcal{L})$. As a result, we indeed have a true set of negative samples $\{(v, *, l) : v \in V_{\mathcal{L}} \wedge L(v) \not\supseteq l\}$. We do not require any negative samples (nor supervision edges) over the subgraph $V \times R \times V$.

4 Experiments

We experimentally compare how the two different representations and tasks perform on a variety of 12 real-world datasets and 7 synthetic graphs of different characteristics, using two backbone models (GraphSAGE, GAT) and their heterogeneous extensions as specified in Section 2.1 (Appendix A.1). All experiments use a T4 16G GPU and our code is available on Github¹.

4.1 Setup

For each dataset G , and for every seed, an 80/10/10 split is obtained on G_{NC} and the equivalent split is created on G_{LP} following Equation 2 3 4 5. 30% of the label edges are encoded as message passing edges T'_{train_mp} , $|T'_{train_mp}/T'_{train_sup}| = 3/7$. The mutual hyper-parameters (number of epochs and learning

¹Anonymised Github: https://anonymous.4open.science/r/nc_lp_paper-3FB0/README.md

rate) and the optimizer (Adam) are the same in both cases. Early stopping is applied on the validation score. The test score with the highest validation score is recorded. Batching is used on larger datasets. For G_{LP} , we make sure that all the \mathcal{L} nodes are present in every batch.

Accuracy is calculated for single-labeled datasets, and average precision (AP) is used for multi-labelled ones. We report the average scores for each task/representation and model on each dataset together with the standard deviation. We also report a direct comparison of scores for each fixed split and model, on the two representations, and calculate the ratio of LP being better than NC in column $LP > NC$. Additional experimental details are specified in Appendix A.2.

4.2 Datasets

We compare three sets of real-world datasets that exhibit different levels of homophily, heterogeneity, and number of labels per node, to identify generalisable trends that correlate to specific dataset characteristics. We additionally use a set of generated random graphs to validate the heuristics we observe.

For homophilic and heterophilic data, six datasets from Rossi et al. (2023) with varying degrees of 1-hop homophily index were used. The homophilic datasets are CiteSeer (Giles et al., 1998; Bojchevski & Günnemann, 2018), Cora-ML (McCallum et al., 2000; Bojchevski & Günnemann, 2018), and ogbn-Arxiv (Hu et al., 2020), and the heterophilic datasets are Chameleon (Rozemberczki et al., 2021), Squirrel (Rozemberczki et al., 2021), and Roman-Empire (Platonov et al., 2024). We additionally provide statistics of K-hop homophily index and neighbourhood sizes for them in Appendix A.3 (Table 8, 9).

Heterogeneous datasets are ACM, DBLP, and IMDB from Wang et al. (2019) and multi-labelled datasets DBLP-multi (Akujuobi et al., 2019), BlogCat (Tang & Liu, 2009), and Yelp (Zeng et al., 2020).

All of the datasets by default adopt the NC representation and NC benchmarks in the original work. All datasets in G_{NC} form have node features \mathbf{X}_V , except BlogCat, which we initialise randomly. For BlogCat, we also found that the split from Tang & Liu (2009) is the only split where the score remained consistent over the training, validation, and test sets, which we use instead of random splits. We initialise $X_{\mathcal{L}}$ for the \mathcal{L} nodes in G_{LP} with one-hot encodings.

4.3 Results on homophilic and heterophilic datasets

Tables 1 and 2 present comparisons between NC and LP on homophilic and heterophilic datasets. Both models consistently perform better on the NC representations for heterophilic datasets. This is perhaps counterintuitive since the extra edges T' in G_{LP} directly connect the same-labelled nodes far apart in G_{NC} . For homophilic datasets, we observe that GAT on CiteSeer_{LP} performs better than with CiteSeer_{NC}. This gap can be as significant as +4%. Furthermore, the advantage of CiteSeer_{LP} is observed consistently across every seed (with $LP > NC$ at 100%).

More interestingly, both two-layer models perform better on the LP representations of the widely benchmarked ogbn-arxiv datasets, as opposed to its default NC presentation. The highest average score is obtained by GraphSAGE on ogbn-arxiv_{LP}. The behaviour of three-layer models varies significantly. For example, GAT on ogbn-arxiv_{NC} obtains consistent scores when using 2 or 3 layers, but drops significantly on ogbn-arxiv_{LP} with 3 layers. The behaviour is different for GraphSAGE, where 2- and 3-layer GraphSAGE consistently and on average performs better on ogbn-arxiv_{LP} than on ogbn-arxiv_{NC}.

4.4 Results on heterogeneous datasets

We experiment with fundamentally heterogeneous datasets and metapath-based heterogeneous GNNs to understand whether the NC and LP representations make a difference in this case. Performance results for heterogeneous datasets using HAN are summarized in Table 3. LP and NC achieve comparable results on DBLP (single), with LP slightly outperforming NC in average accuracy. In contrast, NC significantly outperforms LP on IMDB and ACM. The detailed dataset description and the metapaths used are specified in Appendix A.4.

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
CiteSeer	96.36 ± 0.90	94.59 ± 1.02	0%	90.02 ± 0.94	94.54 ± 0.85	100%
Cora-ML	88.63 ± 1.18	86.00 ± 2.80	0%	86.60 ± 1.86	84.57 ± 1.86	10%
OGBN-Arxiv	69.53 ± 0.32	72.17 ± 0.30	100%	69.76 ± 0.52	70.40 ± 0.59	60%
Chameleon	65.13 ± 2.60	57.54 ± 3.30	0%	73.38 ± 2.18	67.76 ± 3.16	0%
Squirrel	43.88 ± 1.66	40.88 ± 1.09	20%	44.73 ± 2.54	38.15 ± 2.26	0%
Roman Empire	78.18 ± 0.69	76.37 ± 0.92	0%	62.07 ± 1.09	56.02 ± 4.01	10%

Table 1: Test scores on the homophilic and heterophilic datasets, using models with 2 layers. LP > NC represent the percentage of runs where G_{LP} is a better representation than G_{NC} .

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
CiteSeer	95.86 ± 0.85	95.44 ± 1.08	50%	90.12 ± 1.44	94.35 ± 0.87	100%
Cora-ML	87.37 ± 1.49	82.33 ± 4.73	0%	85.67 ± 2.12	84.70 ± 2.32	20%
OGBN-Arxiv	70.37 ± 0.32	72.14 ± 0.41	100%	69.84 ± 0.42	65.59 ± 1.12	0%
Chameleon	63.95 ± 4.07	56.93 ± 2.79	0%	73.16 ± 3.28	59.56 ± 3.79	0%
Squirrel	44.12 ± 2.47	37.73 ± 2.59	0%	42.15 ± 2.64	31.96 ± 4.01	0%
Roman Empire	78.99 ± 0.65	74.47 ± 1.22	0%	43.67 ± 1.79	37.05 ± 3.55	10%

Table 2: Test scores on the homophilic and heterophilic datasets, using models with 3 layers. LP > NC represent the percentage of runs where G_{LP} is a better representation than G_{NC} .

	NC	LP	LP > NC
DBLP (single)	93.15 ± 1.13	93.45 ± 1.61	50%
IMDB	67.99 ± 2.26	56.80 ± 3.27	0%
ACM	92.15 ± 1.71	36.39 ± 6.83	0%

Table 3: Test scores on the heterogeneous datasets with HAN as the model.

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
DBLP (multi)	93.41 \pm 0.55	94.24 \pm 0.26	100%	91.46 \pm 0.41	93.36 \pm 0.32	100%
BlogCat	99.83 \pm 0.21	11.94 \pm 1.20	0%	58.08 \pm 0.95	9.05 \pm 2.07	0%
Yelp	62.32 \pm 0.55	39.21 \pm 2.80	0%	42.59 \pm 0.75	26.34 \pm 1.09	0%

Table 4: Test scores on the multi-labelled datasets, using models with 2 layers. LP > NC represent the percentage of runs where G_{LP} is a better representation than G_{NC} .

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
DBLP (multi)	93.68 \pm 0.42	93.78 \pm 0.49	80%	91.63 \pm 0.49	91.57 \pm 0.39	60%
BlogCat	99.58 \pm 0.29	12.59 \pm 1.81	0%	34.00 \pm 1.86	11.36 \pm 1.32	0%
Yelp	61.41 \pm 0.62	37.77 \pm 3.02	0%	40.83 \pm 0.53	27.05 \pm 1.02	0%

Table 5: Test scores on the multi-labelled datasets, using models with 3 layers. LP > NC represent the percentage of runs where G_{LP} is a better representation than G_{NC} .

4.5 Results on multi-labeled datasets

Results presented in Tables 4 and 5 demonstrate competitive LP performance on DBLP(multi). For BlogCat and Yelp, NC outperforms LP. The extremely large gap shown by BlogCat is not surprising as the dataset is particularly sensitive to splits.

4.6 Results on random graphs

We plot distributions of the node degrees of the 12 real-world datasets in Figures 5, 6, 7 in Appendix A.3. We observe a pattern in that most of the real-world datasets on which LP performs better (CiteSeer, OGBN-Arxiv, DBLP (multi)) are sparse with many low-degree nodes and few mega-nodes.

Heuristically, there seems to be a trend of *LP is sometimes better than NC on sparse and weakly homophilic graphs with few high-degree nodes*. We extend our experiments to a set of random graphs generated under different distributions and various label assignments.


A barbell graph consists of two cliques of m_1 number of nodes connected by a path of length m_2 . We assign binary labels with a method *hom*, where each clique with half of the path being one label. We assign binary labels with another method *homsplit*, where both cliques are assigned the same label, and the path the other label. Both methods produce homophilic graphs. A stochastic block model (sbm) graph consists of k blocks with probabilities p_{same} and p_{diff} of creating edges over nodes of the same and different blocks respectively. We create sbm graphs with 4 blocks of varying sizes and edge probabilities. We assign binary labels to each of the two blocks of nodes. We randomly flip node labels for 10% of nodes for all graphs to ensure the graph functions to be learnt are non-trivial. We generate all G_{NC} graphs and transform them into G_{LP} .

Our result is shown in Table 6. Similar trends observed in real-world datasets can be found again. For example, NC is better for dense (homophilic) graphs bb-100-10-*. For sparse and weakly homophilic graphs such as sbm-0.005-0.0005-2000, LP is +30% better for GraphSAGE than NC on average and for every split. As the graphs become more sparse, the performance of GAT improves on LP but not NC. As an example, we plot in Figure 3 the one-hop and two-hop neighbourhoods of a particular node in sbm-0.005-0.0005-2000_{NC} which LP classified correctly but not NC with GraphSAGE. The interpretation is that while the node is homophilic (over both 1 and 2 hops), there is still a portion of nodes of the opposite label. However, under the G_{LP} representation, it's 2-hop neighbourhood of same-labelled nodes will explode disproportionately. This dominant scaling of neighbourhoods over higher k -hops imply more message passing from same-labelled nodes in G_{LP} .

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
bb-100-10-hom	91.43 \pm 5.95	90.0 \pm 6.88	0%	92.30 \pm 5.30	90.0 \pm 6.19	0%
bb-100-10-homsplit	84.29 \pm 5.65	82.85 \pm 5.71	10%	83.81 \pm 6.10	84.29 \pm 5.65	10%
sbm-0.05-0.2-50	78.50 \pm 8.96	74.50 \pm 11.5	50%	84.50 \pm 10.11	75.00 \pm 10.72	10%
sbm-0.05-0.01-50	75.0 \pm 9.21	78.0 \pm 13.8	60%	81.0 \pm 8.31	66.0 \pm 8.31	10%
sbm-0.01-0.002-1000	57.2 \pm 2.49	89.8 \pm 1.14	100%	89.7 \pm 1.23	75.9 \pm 9.77	0%
sbm-0.01-0.002-1000	61.4 \pm 2.31	89.6 \pm 1.77	100%	89.3 \pm 1.39	81.8 \pm 4.29	10%
sbm-0.005-0.0005-2000	60.32 \pm 3.81	90.69 \pm 0.51	100%	90.51 \pm 0.58	85.89 \pm 6.32	10%

Table 6: Test scores on the random graphs. bb- m_1 - m_2 -mode represents a barbell graph of m_1 clique sizes and m_2 path length with mode either *hom* or *homsplit*. sbm- p_{same} - p_{diff} -blocksize represents a sbm graph of 4 blocks of blocksize number of nodes, p_{same} probability of edges in a block and p_{diff} probability of edges across block. When $p_{\text{same}} > p_{\text{diff}}$ the graphs are generally homophilic. When $p_{\text{same}} < p_{\text{diff}}$ the graphs are generally heterophilic. Two sbm-0.01-0.002-1000 graphs are generated with different seeds.

1-Hop Ego Graph (Colored by True Label) — Center Node 720



2-Hop Ego Graph (Colored by True Label) — Center Node 720



Figure 3: Example 1-hop and 2-hop neighbourhoods of a node id 720 in the sbm-0.005-0.0005-2000 graph in G_{NC} representation. Node 720 is highlighted in black. Nodes share the same label are in lightblue. Nodes of the other label are in orange.

	GraphSAGE		GAT	
	LP	$\downarrow \Delta$	LP	$\downarrow \Delta$
CiteSeer	95.34 ± 0.97	0.1%	94.06 ± 1.49	0.29%
OGBN-Arxiv	69.63 ± 1.52	2.51%	59.84 ± 5.23	5.75%
Roman Empire	66.19 ± 1.86	8.28%	34.74 ± 5.65	2.31%

Table 7: Test scores when using 80% message passing edges, $|T'_{train_mp}/T'_{train_sup}| = 8/2$ instead of $3/7$. $\downarrow \Delta$ represents the changes in average scores.

4.7 Hyperparameter analysis

There is a spectrum of model architecture and learning hyperparameter choices that impact the performance of different models to solve P_{NC} on G_{NC} and P_{LP} on G_{LP} . A hyperparameter unique to LP but not NC is the training message passing and supervision edge ratio. Table 7 shows the results when we use 80% edges as message passing instead of 30%. We also examine the impact of embedding dimensions in representative datasets in Appendix A.5 Table 11. Variations for different datasets validate that the common impact of hyperparameter choices carry over to the different graph representations.

5 Related Work

Rewiring. There are several recent works that modify the computation graphs from the underlying data graph in order to remediate over-smoothing or over-squashing effects of GNNs, or to improve empirical performance. Topping et al. (2022) reduces over-squashing by removing bottlenecks in the graph by examining their curvature. Barbero et al. (2024) improves connectivity while preserving locality (distant nodes are kept separate). Qian et al. (2024) introduces virtual-nodes that allows efficient message-passing over originally long-distance nodes. Rossi et al. (2023) show that treating heterophilic graphs as directed can improve the performance of GNNs when they are appropriately extended to directed-GNN. The approach of rewiring often introduces nodes or edges that are not part of the data whereas we take the perspective of considering different data representations on an abstract graph. Unique to our two representations, the appropriate formal task becomes different, whereas none of the rewiring methods affect the end task.

Using label information. Our LP representation encodes the label information of nodes as new nodes and edges in the concrete data graph itself. There are several other works that utilise node-label information for learning. Wang et al. (2021) combines node label and feature information together by propagating them in parallel. This approach dates back to the classical label propagation algorithm (Zhu & Ghahramani, 2002), with Sato (2024) extending them by introducing training-free GNNs that approximate label propagation.

Node classification and link prediction. There is very limited work that connects these two basic and fundamental learning tasks on graphs. Abboud & İsmail İlkan Ceylan (2021) studies jointly the task of transductive node classification over incomplete graphs and link prediction over graphs with node features. Wu et al. (2022) similarly jointly learns node classification and link prediction. Daza et al. (2021) evaluates entity representations that are learnt with LP objective.

6 Discussion and Future Work

Conclusion. In this paper, we investigate how different representations of the same abstract graph affect the performance of a predicted task defined on it. We study one of the most common predictive tasks of labelling nodes. It turns out that among many different possible general representations, one could assign nodes with labels and formulate the task as a node classification problem, or construct new label-nodes, with edge connections representing label assignment to other nodes, hence formulate the task as a link prediction problem.

Our work is the first to make a connection between the two basic and fundamental tasks in graph learning at this level. The line between benchmarking on each task is now blurred. Our experiments show that different representations and their respective tasks are optimal for different graphs, models and hyperparameters. The majority of such predictive tasks have been only benchmarked as a node classification problem, which we now know is insufficient. In line with many recent works that challenges graph learning benchmarks (Luo et al., 2024; Lv et al., 2021; Bechler-Speicher et al., 2025; Coupette et al., 2025) on the use of complex GNNs, graph structures themselves and hyperparameter choices, our work adds a new consideration to all past and future benchmarks.

Our results also empirically suggest heuristics of when a representation is better than the other. It appears that relatively homophilic sparse graphs with few high-degree nodes may be suitable for link prediction representations. We also observe that several of the graphs that link prediction perform better on are subsets of citation networks, suggesting correlation with more nuanced characteristics, such as node features or higher-order graph properties.

We pave the way for formally and fundamentally understanding the interconnection between the key problems of node classification and link prediction. We demonstrate the importance of considering different graph representations and call for more comprehensive future benchmarks. We hope our work is a step towards a general and provable criterion of defining the optimal representation and task.

Limitations and Future work. Our paper is empirical in nature. Given that there is no prior work on the question studied, we attempt to first answer the question of whether different representations matter. The answer turns out to be positive. A valuable future direction will be formally establishing the connections between different representations, in particular for the node classification and link prediction representations, as well as the respective tasks.

Our work focuses only on the transductive setting due to the amount of interdependent components in the entire learning process, such as initialisation, model choices, model and training process hyperparameters, and datasets. Future empirical work could extend our problem definitions to inductive cases. Such an extension will be nontrivial, for example, the optimal way of defining and sampling negative edges in the inductive case is unobvious.

References

- Ralph Abboud and İsmail İlkan Ceylan. Node classification meets link prediction on knowledge graphs, 2021. URL <https://arxiv.org/abs/2106.07297>.
- Uchenna Akujuobi, Han Yufei, Qiannan Zhang, and Xiangliang Zhang. Collaborative graph walk for semi-supervised multi-label node classification, 2019. URL <https://arxiv.org/abs/1910.09706>.
- Adrian Arnaiz-Rodriguez, Ahmed Begga, Francisco Escolano, and Nuria Oliver. Diffwire: Inductive graph rewiring via the lovász bound, 2022. URL <https://arxiv.org/abs/2206.07369>.
- Federico Barbero, Ameya Velingker, Amin Saberi, Michael Bronstein, and Francesco Di Giovanni. Locality-aware graph-rewiring in gnns, 2024. URL <https://arxiv.org/abs/2310.01668>.
- Maya Bechler-Speicher, Ben Finkelshtein, Fabrizio Frasca, Luis Müller, Jan Tönshoff, Antoine Siraudin, Viktor Zaverkin, Michael M. Bronstein, Mathias Niepert, Bryan Perozzi, Mikhail Galkin, and Christopher Morris. Position: Graph learning will lose relevance due to poor benchmarks, 2025. URL <https://arxiv.org/abs/2502.14546>.
- Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking, 2018. URL <https://arxiv.org/abs/1707.03815>.
- Jun Chen and Haopeng Chen. Edge-featured graph attention network, 2021. URL <https://arxiv.org/abs/2101.07671>.
- Tianlong Chen, Kaixiong Zhou, Keyu Duan, Wenqing Zheng, Peihao Wang, Xia Hu, and Zhangyang Wang. Bag of tricks for training deeper graph neural networks: A comprehensive benchmark study, 2022. URL <https://arxiv.org/abs/2108.10521>.

- Corinna Coupette, Jeremy Wayland, Emily Simons, and Bastian Rieck. No metric to rule them all: Toward principled evaluations of graph-learning datasets, 2025. URL <https://arxiv.org/abs/2502.02379>.
- Daniel Daza, Michael Cochez, and Paul Groth. Inductive entity representations from text via link prediction. In *Proceedings of the Web Conference 2021*, WWW '21, pp. 798–808. ACM, April 2021. doi: 10.1145/3442381.3450141. URL <http://dx.doi.org/10.1145/3442381.3450141>.
- Ahmed El-Kishky, Thomas Markovich, Serim Park, Chetan Verma, Baekjin Kim, Ramy Eskander, Yury Malkov, Frank Portman, Sofia Samaniego, Ying Xiao, and Aria Haghighi. Twihin: Embedding the twitter heterogeneous information network for personalized recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, pp. 2842–2850. ACM, August 2022. doi: 10.1145/3534678.3539080. URL <http://dx.doi.org/10.1145/3534678.3539080>.
- Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020*, WWW '20. ACM, April 2020. doi: 10.1145/3366423.3380297. URL <http://dx.doi.org/10.1145/3366423.3380297>.
- C. Lee Giles, Kurt D. Bollacker, and Steve Lawrence. Citeseer: an automatic citation indexing system. In *Proceedings of the Third ACM Conference on Digital Libraries*, DL '98, pp. 89–98, New York, NY, USA, 1998. Association for Computing Machinery. ISBN 0897919653. doi: 10.1145/276675.276685. URL <https://doi.org/10.1145/276675.276685>.
- William L. Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs, 2018. URL <https://arxiv.org/abs/1706.02216>.
- Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 22118–22133. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/fb60d411a5c5b72b2e7d3527cfc84fd0-Paper.pdf.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. URL <https://arxiv.org/abs/1609.02907>.
- Yuankai Luo, Lei Shi, and Xiao-Ming Wu. Classic gnns are strong baselines: Reassessing gnns for node classification, 2024. URL <https://arxiv.org/abs/2406.08993>.
- Qingsong Lv, Ming Ding, Qiang Liu, Yuxiang Chen, Wenzheng Feng, Siming He, Chang Zhou, Jianguo Jiang, Yuxiao Dong, and Jie Tang. Are we really making much progress? revisiting, benchmarking, and refining heterogeneous graph neural networks, 2021. URL <https://arxiv.org/abs/2112.14936>.
- Andrew K. McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3:127–163, 2000. doi: 10.1023/A:1009953814988.
- Trung-Kien Nguyen and Yuan Fang. Diffusion-based negative sampling on graphs for link prediction, 2024. URL <https://arxiv.org/abs/2403.17259>.
- Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *CoRR*, abs/2002.05287, 2020. URL <https://arxiv.org/abs/2002.05287>.
- Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of gnns under heterophily: Are we really making progress?, 2024. URL <https://arxiv.org/abs/2302.11640>.
- Chendi Qian, Andrei Manolache, Christopher Morris, and Mathias Niepert. Probabilistic graph rewiring via virtual nodes, 2024. URL <https://arxiv.org/abs/2405.17311>.

- Emanuele Rossi, Bertrand Charpentier, Francesco Di Giovanni, Fabrizio Frasca, Stephan Günnemann, and Michael Bronstein. Edge directionality improves learning on heterophilic graphs, 2023. URL <https://arxiv.org/abs/2305.10498>.
- Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-scale attributed node embedding, 2021. URL <https://arxiv.org/abs/1909.13021>.
- T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A survey on oversmoothing in graph neural networks, 2023. URL <https://arxiv.org/abs/2303.10993>.
- Ryoma Sato. Training-free graph neural networks and the power of labels as features, 2024. URL <https://arxiv.org/abs/2404.19288>.
- Damian Szklarczyk, Annika L Gable, David Lyon, Alexander Junge, Stefan Wyder, Jaime Huerta-Cepas, Milan Simonovic, Nadezhda T Doncheva, John H Morris, Peer Bork, Lars J Jensen, and Christian von Mering. String v11: protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. *Nucleic Acids Research*, 47(D1):D607–D613, 2019. ISSN 0305-1048. doi: 10.1093/nar/gky1131.
- Lei Tang and Huan Liu. Relational learning via latent social dimensions. pp. 817–826, 06 2009. doi: 10.1145/1557019.1557109.
- Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature, 2022. URL <https://arxiv.org/abs/2111.14522>.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks, 2018. URL <https://arxiv.org/abs/1710.10903>.
- Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Peng Cui, Philip S. Yu, and Yanfang Ye. Heterogeneous graph attention network. *CoRR*, abs/1903.07293, 2019. URL <http://arxiv.org/abs/1903.07293>.
- Yangkun Wang, Jiarui Jin, Weinan Zhang, Yong Yu, Zheng Zhang, and David Wipf. Bag of tricks for node classification with graph neural networks, 2021. URL <https://arxiv.org/abs/2103.13355>.
- Yuchen Wang, Jinghui Zhang, Zhengjie Huang, Weibin Li, Shikun Feng, Ziheng Ma, Yu Sun, Dianhai Yu, Fang Dong, Jiahui Jin, Beilun Wang, and Junzhou Luo. Label information enhanced fraud detection against low homophily in graphs. In *Proceedings of the ACM Web Conference 2023*, WWW ’23, pp. 406–416. ACM, April 2023. doi: 10.1145/3543507.3583373. URL <http://dx.doi.org/10.1145/3543507.3583373>.
- David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, Nazanin Assempour, Ithayavani Iynkkaran, Yifeng Liu, Adam Maciejewski, Nicola Gale, Alex Wilson, Lucy Chin, Ryan Cummings, Diana Le, Allison Pon, Craig Knox, and Michael Wilson. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic Acids Research*, 46(D1):D1074–D1082, 11 2017. ISSN 0305-1048. doi: 10.1093/nar/gkx1037. URL <https://doi.org/10.1093/nar/gkx1037>.
- Zongqian Wu, Mengmeng Zhan, Haiqi Zhang, Qimin Luo, and Kun Tang. Mtgen: A multi-task approach for node classification and link prediction in graph data. *Information Processing & Management*, 59(3): 102902, 2022. ISSN 0306-4573. doi: <https://doi.org/10.1016/j.ipm.2022.102902>. URL <https://www.sciencedirect.com/science/article/pii/S0306457322000292>.
- Zhen Yang, Ming Ding, Chang Zhou, Hongxia Yang, Jingren Zhou, and Jie Tang. Understanding negative sampling in graph representation learning, 2020. URL <https://arxiv.org/abs/2005.09863>.
- Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD ’18, pp. 974–983. ACM, July 2018. doi: 10.1145/3219819.3219890. URL <http://dx.doi.org/10.1145/3219819.3219890>.

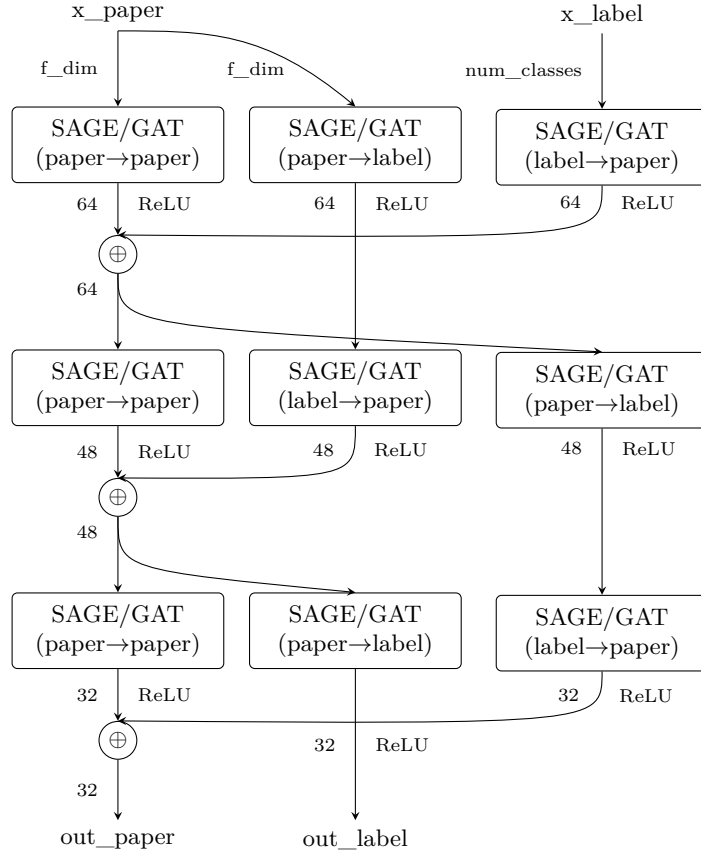


Figure 4: The architecture of a heterogeneous GNN for a graph with two node types $\mathcal{A} = \{\text{paper}, \text{label}\}$ and one relation type $\mathcal{R} = \{*\}$ with no edges between label nodes. This example GNN has three layers. A separate channel is instantiated for each triplet type. \oplus denotes an aggregation of embeddings, in this paper, PyG’s default aggregator, summation, is used.

Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graphsaint: Graph sampling based inductive learning method, 2020. URL <https://arxiv.org/abs/1907.04931>.

Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications, 2021. URL <https://arxiv.org/abs/1812.08434>.

Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. 2002. URL <https://api.semanticscholar.org/CorpusID:15008961>.

A Appendix

A.1 Computation graph of heterogeneous GNN

An example computation graph of a heterogeneous GNN is illustrated in Figure 4. This implementation is directly supported by PyG’s `to_hetero()` function on *HeteroData*. This implementation is a basic and natural extension of homogeneous GNNs without adding metapaths that represent additional heuristics that strengthen the models. The models with three layers have layer sizes (64, 48, 32), while the models with two layers have layer sizes (64, 32). In the GAT models, 4 attention heads are used. For both task representations, the decoder is a single linear layer on either nodes (for NC) or node-pairs (for LP).

		1-hop	2-hop	3-hop	RW 2-hop	RW 3-hop
Homophilic	CITSEER	0.96	0.94	0.92	0.95	0.93
	CORA-ML	0.81	0.73	0.58	0.77	0.61
	OGBN-ARXIV	0.64	0.50	0.37	0.58	0.42
Heterophilic	CHAMELEON	0.25	0.25	0.23	0.30	0.23
	SQUIRREL	0.22	0.21	0.20	0.26	0.21
	ROMAN-EMPIRE	0.05	0.07	0.07	0.29	0.11

Table 8: Homophily numbers for the different datasets in the heterophilic and homophilic comparison, before and after rewiring with an 80/10/10 split. RW denotes homophily numbers for the rewired graph.

	1-hop	2-hop	3-hop	RW 1-hop	RW 2-hop	RW 3-hop
CITSEER	2.5	16.7	42.5	2.76	57.6	186.9
CORA-ML	5.4	71.8	294.7	5.7	99.3	422.3
OGBN-ARXIV	13.7	3484.2	18469.6	13.9	4151.7	22579.6
CHAMELEON	27.6	558.7	1066.3	27.8	580.8	1312.3
SQUIRREL	76.3	1692.1	3538.5	76.5	1736.1	3836.9
ROMAN-EMPIRE	2.9	7.1	12.8	3.1	121.2	750.6

Table 9: Neighbourhood sizes for the different datasets in the heterophilic and homophilic comparison, before and after rewiring with an 80/10/10 split. RW denotes neighborhood sizes for the rewired graph.

A.2 Additional experimental setups

For the smaller real-world datasets and all synthetic graphs we used 10 random seeds. For larger datasets, use 5 seeds for obgn-arxiv and Squirrel and 3 seeds for BlogCat and Yelp if the results are consistent. Default learning rate = 0.01. Default hidden channels = 64 and embedding size = 32.

The homogeneous GraphSAGE and GAT encoders consist of either two or three SAGE / GAT-layers with RELU-activation functions between them. The models with three layers have layer sizes (64, 48, 32), while the models with two layers have layer sizes (64, 32). In the GAT models, 4 attention heads are used. For both task representations, the decoder is a single linear layer on either nodes (for NC) or nodepairs (for LP).

The HAN model only consists of one layer, since the metapaths allow for message-passing in neighbourhoods larger than 1-hop. Following Wang et al., 8 attention heads are used together with a dropout rate of 0.6, and 128 hidden channels. The decoder architectures are the same as for GraphSAGE and GAT.


A.3 Dataset Statistics

Number of nodes, number of edges, number of classes, average degree, and node feature size for every dataset can be found in Table 10. For each dataset, the node degree distribution is visualised in Figures 5, 6, and 7. For the heterogeneous datasets, the node degree distribution is calculated on the computation graph induced by the meta-paths.


A.4 Heterogeneous Dataset Information

ACM is a network of papers, authors, and subjects, where the papers are divided into three classes (*database*, *wireless communication*, *data mining*). Paper features correspond to elements of a bag-of-words representation of keywords. The meta-path set $\{PAP, PSP\}$ is used in Wang et al. (2019).


DBLP is a network of authors, papers, conferences and terms. The labels indicate research areas of the authors, and are one of *database*, *data mining*, *machine learning*, and *information retrieval*. Author features contain the elements of a bag-of-words represented by keywords. The meta-path set employed in Wang et al. (2019) is $\{APA, APCPA, APTPA\}$.




(a) CiteSeer




(b) Cora-ML



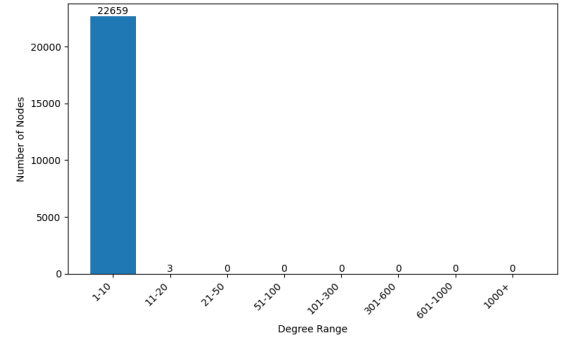
(c) OGBN-Arxiv



(d) Chameleon



(e) Squirrel




(f) Roman-Empire


Figure 5: Node degree distributions for the six datasets with varying label homophily levels.

	$ V $	$ T $	Number of Classes	d	Average Degree
CiteSeer	4,230	10,674	6	602	2.5
Cora-ML	2,995	16,316	7	2,879	5.4
OGBN-Arxiv	169,343	1,166,243	40	128	13.7
Chameleon	2,277	36,101	5	2,325	27.6
Squirrel	5,201	217,073	5	2,089	76.3
Roman Empire	22,662	65,854	18	300	2.9
DBLP (single)	4,057	119,783	4	334	5,942.9
IMDB	4,932	20,172	2	3,489	48.1
ACM	3,025	273,000	3	1,902	4,124.3
DBLP (multi)	28,702	68,335	4	223	4.8
BlogCat	10,312	333,983	39	0	64.8
Yelp	716,847	7,340,000	100	300	19.5


Table 10: Statistics of all real-world datasets used in the paper, where $|V|$ is the number of nodes, $|T|$ is the number of edges, and d is the size of node features. All statistics are based on the G_{NC} representation.



(a) DBLP (single)



(b) IMDB



(c) ACM

Figure 6: Node degree distributions for the computational graph (only involving meta-paths) of heterogeneous datasets.




Figure 7: Node degree distributions for the multi-labeled datasets.

	GraphSAGE			GAT		
	NC	LP	LP > NC	NC	LP	LP > NC
CiteSeer	96.24 ± 0.97	94.3 ± 1.32	10%	90.26 ± 1.26	94.09 ± 0.93	100%
OGBN-Arxiv	67.99 ± 3.76	70.95 ± 1.00	80%	69.57 ± 0.17	65.69 ± 0.66	0%
Roman Empire	80.27 ± 1.36	72.44 ± 1.11	0%	42.41 ± 2.90	38.05 ± 2.52	10%

Table 11: Test scores on representative datasets using models with 3 layers. Hidden channel dimensions are smooth interpolations of the input and output dimensions.

IMDB is a dataset containing movies, actors, and directors. Each movie is divided into one out of three classes (*Action*, *Comedy*, *Drama*). Movie features correspond to elements of a bag-of-words representation of plots. The meta-path set $\{MAM, MDM\}$ is used in Wang et al. (2019).

The new G_{LP} representations of these graphs make it possible to define new meta-paths which include label nodes. Studying this is out of the scope of this paper. The same meta-path sets are used as in the original paper. The added nodes and message-passing edges will still yield added information in the LP representation.

A.5 Hyperparameter analysis on embedding dimensions

Table 11 shows the test scores on selected datasets when the hidden channel embedding dimension are not fixed to 64, but instead smoothly interpolate between input and output embedding sizes. For example, if input feature has dimension 4000, output embedding has dimension 1000, then the three-layer models have embeddings 4000-3000-2000-1000.