GLINKX: A SCALABLE UNIFIED FRAMEWORK FOR HOMOPHILOUS AND HETEROPHILOUS GRAPHS

Anonymous authors
Paper under double-blind review

ABSTRACT

In graph learning, there have been two predominant inductive biases regarding graph-inspired architectures: On the one hand, higher-order interactions and message passing work well on homophilous graphs and are leveraged by GCNs and GATs. Such architectures, however, cannot easily scale to large real-world graphs. On the other hand, shallow (or node-level) models using ego features and adjacency embeddings work well in heterophilous graphs. In this work, we propose a novel scalable shallow method – GLINKX – that can work both on homophilous and heterophilous graphs. GLINKX leverages (i) novel monophilous label propagations (ii) ego/node features, (iii) knowledge graph embeddings as positional embeddings, (iv) node-level training, and (v) low-dimensional message passing. Formally, we prove novel error bounds and justify the components of GLINKX. Experimentally, we show its effectiveness of it on several homophilous and heterophilous datasets.

1 INTRODUCTION

In recent years, graph learning methods have emerged with strong performance for various ML tasks. Graph ML methods leverage the topology of graphs underlying the data (Battaglia et al., 2018) to improve their performance. Two very important design options for proposing graph ML based architectures in the context of *node classification* are related to whether the data is *homophilous* or *heterophilous*.

For homophilous data – where neighboring nodes share similar labels (McPherson et al., 2001; Altenburger & Ugander, 2018a) – Graph Convolutional Network (GCN)-based methods (GCN, GAT, etc.) (Kipf & Welling, 2016; Velivcković et al., 2017; Zhu et al., 2020) are able to achieve high accuracy. In the GCN paradigm, *message passing* and *higher-order interactions* help node classification tasks in the homophilous setting since such inductive biases tend to bring the learned representations of linked nodes close to each other. However, GCN-based architectures suffer from *scalability issues*. Performing (higher-order) propagations during the training stage is hard to scale in large graphs because the number of nodes grows exponentially with the increase of the filter receptive field. Thus, for practical purposes, GCN-based methods require *node sampling* which substantially increases their training time. For this reason, architectures (Huang et al., 2020; Zhang et al., 2022b; Sun et al., 2021; Maurya et al., 2021; Rossi et al., 2020) that leverage propagations outside of the training loop (as a preprocessing step) have shown promising results in terms of scaling to large graphs.

In *heterophilous* datasets (Rogers et al., 2014), the nodes that are connected tend to have different labels. Currently, many works that address heterophily can be classified into two categories concerning scale. On the one hand, recent successful architectures (in terms of accuracy) (Jin et al., 2022a; Di Giovanni et al., 2022; Zheng et al., 2022b; Luan et al., 2021; Chien et al., 2020) that address heterophily resemble GCNs in terms of design and thus suffer from the same scalability issues. On the other hand, *shallow or node-level models* (see e.g. (Lim et al., 2021; Zhong et al., 2022)), i.e. models that are treating graph data as tabular data and do not involve propagations during training, have shown a lot of promise for large heterophilous graphs. In (Lim et al., 2021), it is shown that combining ego embeddings (node features) and adjacency embeddings works in the heterophilous setting. One element that LINKX exploits via the adjacency embeddings is monophily (Altenburger & Ugander, 2018a;b), namely the similarity of the labels of a node's neighbors. However, their design is still impractical in real-world data since the method (LINKX) is *not* inductive (see Section 2),

Algorithm 1 GLINKX Algorithm

Input: Graph $G(V, E)$ with train set $V_{\text{train}} \subseteq V$, node features $\boldsymbol{X}$, labels $\boldsymbol{Y}$
Output: Label Predictions $\boldsymbol{Y}_{\text{final}}$
1st Stage. Pre-train knowledge graph embeddings $\boldsymbol{P}$ with Pytorch Biggraph.
2nd Stage. Propagate labels and predict neighbor distribution

1. **Forward Propagation:** Calculate $\hat{\boldsymbol{y}}_i = \frac{\sum_{j \in V_{\text{train}}: j \to i} \boldsymbol{y}_j}{|\{j \in V_{\text{train}}: j \to i\}|}$ for all $i \in V_{\text{train}}$
2. **Learn distribution of a node's neighbors:** For each epoch, calculate $\tilde{\boldsymbol{y}}_i = f_1(\boldsymbol{\xi}_i, \boldsymbol{p}_i; \boldsymbol{\theta}_1)$ for $i \in V_{\text{train}}$
   Update the parameters s.t. $\mathcal{L}_{\text{CE},1}(\boldsymbol{\theta}_1) = \sum_{i \in V_{\text{train}}} \text{CE}(\hat{\boldsymbol{y}}_i, \tilde{\boldsymbol{y}}_i; \boldsymbol{\theta}_1)$ is maximized.
   Let $\boldsymbol{\theta}_1^*$ be the parameters at the end of the training that correspond to the epoch with the best validation accuracy.

3rd Stage. Propagate labels backwards and train model to predict a node's own labels

1. **Backward Propagation:** Calculate $\boldsymbol{y}'_i = \frac{\sum_{j \in V: j \to i} \tilde{\boldsymbol{y}}_j}{|\{j \in V: i \to j\}|}$ for all $i \in V_{\text{train}}$, where $\tilde{\boldsymbol{y}}_j = f_1(\boldsymbol{\xi}_j, \boldsymbol{p}_j; \boldsymbol{\theta}_1^*)$
2. **Learn a node's own distribution:** For each epoch, calculate $y_{\text{final},i} = f_2(\boldsymbol{\xi}_i, \boldsymbol{p}_i, \boldsymbol{y}'_i; \boldsymbol{\theta}_2)$.
   Update the parameters s.t. $\mathcal{L}_{\text{CE},2}(\boldsymbol{\theta}_2) = \sum_{i \in V_{\text{train}}} \text{CE}(y_i, y_{\text{final},i}; \boldsymbol{\theta}_2)$ is maximized.

Return $\boldsymbol{Y}_{\text{final}}$

and embedding the adjacency matrix directly requires many parameters in a model. In LINKX, the adjacency embedding of a node can alternatively be thought of as a *positional embedding* (PE) of the node in the graph, and recent developments (Kim et al., 2022; Dwivedi et al., 2021; Lim et al., 2021) have shown the importance of PEs in both homophilous and heterophilous settings. However, most of these works suggest PE parametrizations that are difficult to compute in large-scale settings. Accordingly, more scalable ways of computing PEs via *knowledge graph embeddings* (El-Kishky et al., 2022; Lerer et al., 2019; Bordes et al., 2013; Yang et al., 2014) are useful in practical settings.

Goal & Contribution: In this work, we develop a scalable method for node classification that: (i) works both on homophilous and heterophilous graphs (ii) is *simpler and faster* than conventional message passing networks (by avoiding the neighbor sampling and message passing overhead during training), and (iii) can work in both a *transductive* and an *inductive*[1] setting. For a method to be scalable, we argue that it should: (i) run models on node-scale (thus leveraging i.i.d. minibatching), (ii) avoid doing message passing during training and do it a constant number of times before training, and (iii) transmit small messages along the edges. Our proposed method – GLINKX (see Section 3) – combines all the above desiderata. GLINKX has three components: (i) ego embeddings [2], (ii) PEs inspired by architectures suited for heterophilous settings, and (iii) *scalable* 2nd-hop-neighborhood propagations inspired by architectures suited for monophilous settings. We provide novel theoretical error bounds and justify components of our method (Section 3.4). Finally, we evaluate GLINKX's empirical effectiveness on several homophilous and heterophilous datasets (Section 4).

2 PRELIMINARIES

2.1 NOTATION

We denote scalars with lower-case, vectors with bold lower-case letters, and bold upper-case letters are used to denote matrices. We consider a directed graph $G = G(V, E)$ with vertex set V with $|V| = n$ nodes, and edge set E with $|E| = m$ edges, and adjacency matrix A. Let $\boldsymbol{X} \in \mathbb{R}^{n \times d_X}$ represent the d_X-dimensional node feature matrix and $\boldsymbol{P} \in \mathbb{R}^{n \times d_P}$ represent the d_P-dimensional node positional embedding matrix. A node i has a feature vector $\boldsymbol{x}_i \in \mathbb{R}^{d_X}$ and a positional embedding $\boldsymbol{p}_i \in \mathbb{R}^{d_P}$ and belongs to a class $y_i \in \{1, \ldots, c\}$. The training set is denoted by V_{train}, validation set by V_{valid}, and test set by V_{test}. $\mathbb{I}\{\cdot\}$ denotes the indicator function. Γ_c is the c-dimensional simplex.

[1] For this paper, we operate in the *transductive setting*. See App. B for the inductive setting.
[2] We use ego embeddings and node features interchangeably.

2.2 GRAPH CONVOLUTIONAL NEURAL NETWORKS

In homophilous datasets, GCN-based methods have been used for node classification. GCNs (Kipf & Welling, 2016) utilize feature propagations together with non-linearities to produce node embeddings. More specifically, a GCN consists of multiple layers where each layer i collects i-th hop information from the nodes, through propagations, and forwards this information to the $i+1$-th layer. More specifically, if G has a symmetrically-normalized adjacency matrix $\boldsymbol{A}_{sym}$ (ignoring the directionality of edges), then a GCN layer has the form

$$\boldsymbol{H}^{(0)} = \boldsymbol{\xi},\ \boldsymbol{H}^{(i+1)} = \sigma\left(\boldsymbol{A}_{sym}\boldsymbol{H}^{(i)}\boldsymbol{W}^{(i)}\right) \forall i \in [L],\ \boldsymbol{Y} = \mathrm{softmax}\left(\boldsymbol{H}^{(L)}\right).$$

Here $\boldsymbol{H}^{(i)}$ is the embedding from the previous layer, $\boldsymbol{W}^{(i)}$ is a learnable projection matrix and $\sigma(\cdot)$ is a non-linearity (e.g. ReLU, sigmoid, etc.).

2.3 LINKX

In heterophilous datasets, the simple method of LINKX has been shown to perform well. LINKX combines two components – MLP on the node features $\boldsymbol{X}$ and LINK regression (Altenburger & Ugander, 2018a) on the adjacency matrix – as follows:

$$\boldsymbol{H}_X = \text{MLP}_X(\boldsymbol{X})\,,\ \boldsymbol{H}_A = \text{MLP}_A(\boldsymbol{A})\,,\ \boldsymbol{Y} = \mathrm{ResNet}(\boldsymbol{H}_X, \boldsymbol{H}_A).$$

2.4 NODE CLASSIFICATION

In node classification problems on graphs, we have a model $f(\boldsymbol{X}, \boldsymbol{Y}_{\text{train}}, \boldsymbol{A}; \boldsymbol{\theta})$ that takes as an input the node features $\boldsymbol{X}$, the training labels $\boldsymbol{Y}_{\text{train}}$ and the graph topology $\boldsymbol{A}$ and produces a prediction for each node i of G, which corresponds to the probability that a given node belongs to any of c classes (with the sum of such probabilities being one). The model is trained with back-propagation. Once trained, the model can be used for the prediction of labels of nodes in the test set.

There are two training regimes: *transductive* and *inductive*. In the transductive training regime, we have full knowledge of the graph topology (for the train, test, and validation sets) and the node features and the task is to predict the labels of the validation and test set. In the inductive regime, only the graph induced by V_{train} is known at the time of training and then the full graph is revealed for prediction on the validation and test sets. In real-world scenarios, such as online social networks, the dynamic nature of problems makes the inductive regime particularly useful.

2.5 HOMOPHILY, HETEROPHILY & MONOPHILY

Homophily and Heterophily: There are various measures of homophily in the GNN literature like node homophily, and/or edge homophily (Lim et al., 2021). Intuitively, homophily in a graph implies that nodes with similar labels are connected to each other. GNN based approaches like GCN, GAT, etc. leverage this property to improve the node classification performance. Alternatively, if a graph has a low homophily – namely nodes that connect to each other tend to have different labels – it is said to be *heterophilous*. In other words, a graph is said to be heterophilous if neighboring node do not share similar labels.

Monophily: Generally, a graph is said to to be monophilous if the label of a node is similar to that of its neighbors' neighbors. In the context of a directed graph, monophily can be thought of a structure that resembles Fig. 2(a) where similar nodes (in this case 3 green nodes connected to a yellow node) are connected to a node with a different label. We argue that encoding monophily into a model can be useful for both heterophilous and homophilpous graphs. In homophilous graphs, monophily will fundamentally encode 2nd-hop neighbor's label information and since in such graphs neighboring nodes have similar labels it can provide useful signal for node classification. In heterophily, neighboring nodes have dissimilar labels but the 2nd-hop neighbors may share the same label thus providing useful information for node classification. In fact, monophily has been shown

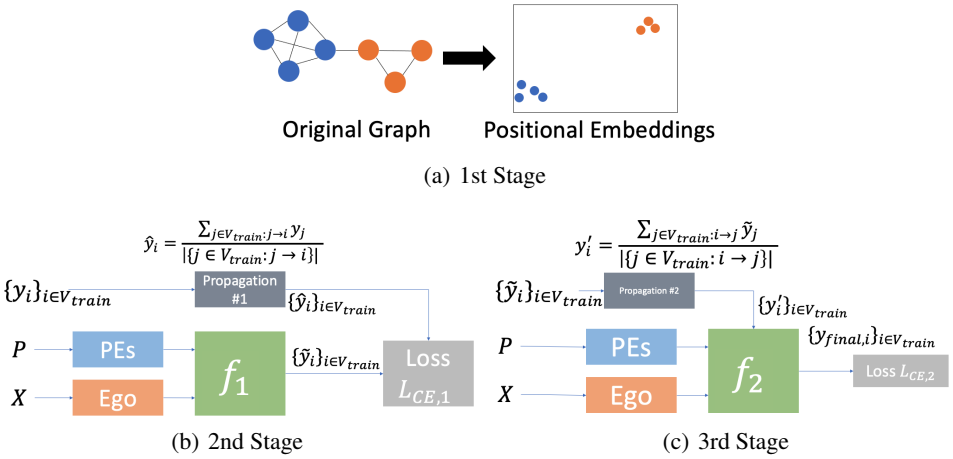

(a) 1st Stage

(b) 2nd Stage

(c) 3rd Stage

Figure 1: Block Diagrams of GLINKX stages.

to be effective for heterophilous graphs (Lim et al., 2021). Thus, an approach encoding monophily has an advantage over methods designed specifically for homophilous and heterophilous graphs especially when varying levels of homophily can exist between different sub-regions in the same graph (see Section 3.3).

2.6 Graph Minibatching & Shallow Methods

In order to train a GCN-based model (or generally, whenever propagations based on the graph topology are involved in the model) on a large-network (that cannot fit in the GPU memory), one has to do *minibatching* through *neighbor sampling*. For large-scale networks, minibatching takes much longer than the full-batch training, which is one of the reasons that graph GCNs are not preferred in real-world settings (see e.g. (Jin et al., 2022b; Zhang et al., 2022a; Zheng et al., 2022a; Lim et al., 2021; Maurya et al., 2021; Rossi et al., 2020)).

For this reason, most methods that can scale on real-world settings are *shallow*. *Shallow* (or node-level) models are based on manipulating the node features $\boldsymbol{X}$ and the graph topology $\boldsymbol{A}$ in a way that propagations do not take place during training. Examples of these methods are LINKX, FSGNN (Maurya et al., 2021), and SIGN (Rossi et al., 2020). Such methods treat the input embeddings as tabular data and pass them through a feed-forward neural network (MLP) to produce the predictions. Thus, they avoid the need for neighborhood sampling, and instead rely on simple tabular minibatching.

3 Method

3.1 Components

GLINKX is described in Alg. 1 and consists of three main components which are detailed as block diagrams in Fig. 1. Fig. 2 shows the GLINKX stages from Alg. 1 on a toy graph.

Positional Embeddings: We use PEs to provide our model information about the position of each node and hypothesize that PEs are an important piece of information in the context of large-scale node classification. PEs have been used to help discriminate isomorphic graph (sub)-structures (Kim et al., 2022; Dwivedi et al., 2021; Srinivasan & Ribeiro, 2019). This is useful for both homophily (Kim et al., 2022; Dwivedi et al., 2021) and heterophily (Lim et al., 2021) because isomorphic (sub)-structures can exist in both the settings. In the homophilous case, adding positional information can help distinguish nodes that have the same neighborhood but distinct position (Dwivedi et al., 2021; Morris et al., 2019; Xu et al., 2019), circumventing the to do higher-order propagations (Dwivedi et al., 2021; Li et al., 2019; Bresson & Laurent, 2017) which are prone to over-squashing (Alon & Yahav, 2021). In heterophily, structural similarity among nodes is important for classification as in the case of LINKX – where adjacency embedding can be thought of as a PE. However, in large

graphs, using adjacency embeddings or Laplacian eigenvectors (as methods such as (Kim et al., 2022) suggest) can be a computational bottleneck and may be infeasible.

In this work, we leverage *knowledge graph embeddings* (KGEs) to encode positional information about the nodes. Using KGEs has the following two benefits: Firstly, KGEs can be trained quickly for large graphs. This is because KGEs essentially compress the adjacency matrix into a fixed sized embedding and adjacency matrices have been shown to be effective in heterophilous cases. Further, KGEs are low-dimensional compared to the adjacency matrix (e.g. $d_P \sim 10^2$) which allows for faster training and inference times. Secondly, KGEs can be pre-trained efficiently on such graphs (Lerer et al., 2019) and can be used off-the-shelf for other downstream tasks including node classification (see e.g. (El-Kishky et al., 2022)). So, in the 1st Stage of our methods in Alg. 1 (**??**) we train KGEs model on the available graph structure. In this work, we fix this positional encoding once they are pre-trained for downstream usage. One can fine-tune these along with learning (Dwivedi et al., 2021) in the downstream task but we leave this for future work.

Ego Embeddings: We obtain ego embeddings from the node features. Such embeddings have been used both in homophilous and heterophilous settings (e.g. (Lim et al., 2021; Zhu et al., 2020)). Node embeddings are useful for tasks where the graph structure provides little/no information for the task.

Monophilous Label Propagations: We now propose a novel monophily (refer Section 2.5) inspired label propagation which we refer to as Monophilous Label Propagation (MLaP). MLaP has the advantage that it can be used both for hetero(homo)philous graphs or in a scenario with varying levels of graph homophily (see Section 3.3) as it encodes monophily (Section 2.5).

In order to understand how MLaP encodes monophily, we consider the example in Fig. 2. In this example, we have 3 green nodes connected to a yellow node and two nodes of different color connected to the yellow node. Then, one way to encode monophily in Fig. 2(a) while predicting label for $j_\ell, \ell \in [5]$, is to get a *distribution* of labels of nodes connected to node i thus encoding its neighbors' distribution. The fact that there are more nodes with green color than other colors can be used by the model to make a prediction. But this information may not always be present or there may be few labelled nodes around node i. Consequently, we propose to use a model that predicts the label distribution of nodes connected to i. We use the node features ($\boldsymbol{x}_i$) and PE ($\boldsymbol{p}_i$) of node i to build this model, since nodes that are connected to node i share similar labels and, thus, the features of node i must be predictive of its neighbors. So, in Fig. 2(a) we train a model to predict a distribution of i's neighbors. Next, we provide j_ℓ the learned distribution of i's neighbors by propagating the learned distribution from i back to j_ℓ. Eqs. (1) to (3) correspond to MLaP. We train a final model that leverages this information together with node features and PEs (Fig. 2(b)).

3.2 GLINKX

Following, we present the individual stages of GLINKX (see Alg. 1):

<u>*1st Stage:*</u> We train KGEs as PEs by using Pytorch-Biggraph and the DistMult method (Yang et al., 2014).

<u>*2nd Stage:*</u> First (2nd Stage in Alg. 1, Fig. 1(b), and Fig. 2(a)), for a node we want to learn the distribution of *its neighbors*. To achieve this, we propagate the labels from a node's neighbors, i.e. calculate

$$\hat{\boldsymbol{y}}_i = \frac{\sum_{j \in V_{\text{train}}: j \to i} \boldsymbol{y}_j}{|\{j \in V_{\text{train}} : j \to i\}|} \qquad \forall i \in V_{\text{train}}. \tag{1}$$

We train a model that predicts the distribution of neighbors, which we denote with $\tilde{\boldsymbol{y}}_i$ using the ego features $\{\boldsymbol{x}_i\}_{i \in V_{\text{train}}}$ and the PEs $\{\boldsymbol{p}_i\}_{i \in V_{\text{train}}}$ and maximize the cross-entropy with treating $\{\hat{\boldsymbol{y}}_i\}_{i \in V_{\text{train}}}$ as ground truth labels, namely we maximize

$$\mathcal{L}_{\text{CE, 1}}(\boldsymbol{\theta}_1) = \sum_{i \in V_{\text{train}}} \sum_{l \in [c]} \hat{\boldsymbol{y}}_{i,l} \log(\tilde{\boldsymbol{y}}_{i,l}), \tag{2}$$

where $\tilde{\boldsymbol{y}}_i = f_1(\boldsymbol{x}_i, \boldsymbol{p}_i; \boldsymbol{\theta}_1)$ and $\boldsymbol{\theta}_1 \in \Theta_1$ is a learnable parameter vector. Although in this paper we assume to be in the *transductive setting*, this step allows us to be inductive (see App. B). In

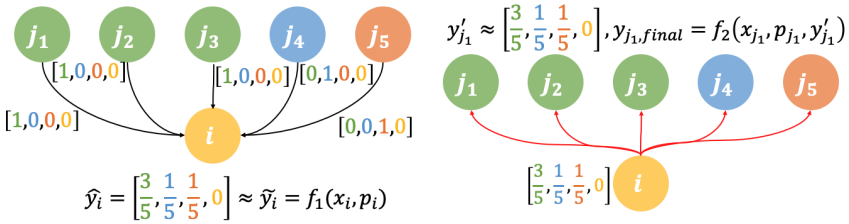

(a) 2nd Stage (b) 3rd Stage

Figure 2: Example. For node i we want to learn a model that takes i's features $\boldsymbol{x}_i \in \mathbb{R}^{d_X}$, and PEs $\boldsymbol{p}_i \in \mathbb{R}^{d_P}$ and predict a value $\widetilde{\boldsymbol{y}}_i \in \mathbb{R}^c$ that matches the label distribution of it's neighbors neighbors $\hat{\boldsymbol{y}}_i$ using a shallow model. Next, we want to propagate (outside the training loop) the (predicted) distribution of a node back to its neighbors and use it together with the ego features and the PEs to make a prediction about a node's own label. We propagate $\tilde{\boldsymbol{y}}_i$ to its neighbors j_1 to j_5. For example, for j_1, we encode the propagated distribution estimate $\tilde{\boldsymbol{y}}_i$ from i to form $\boldsymbol{y}'_{j_1}$. We predict the label by using $\boldsymbol{y}'_{j_1}, \boldsymbol{x}_{j_1}, \boldsymbol{p}_{j_1}$.

Section 3.4 we give a theoretical justification of this step, namely *"why is it good to use a parametric model to predict the distribution of neighbors?"*.

3rd Stage: Then (3rd Stage in Alg. 1, Fig. 1(c), and Fig. 2(b)), we propagate the predicted soft-labels $\tilde{\boldsymbol{y}}_i$ back to the original nodes, i.e. calculate

$$\boldsymbol{y}'_i = \frac{\sum_{j \in V: j \to i} \tilde{\boldsymbol{y}}_j}{|\{j \in V : i \to j\}|} \quad \forall i \in V_{\text{train}}, \tag{3}$$

where the soft labels $\{\tilde{\boldsymbol{y}}_i\}_{i \in V_{\text{train}}}$ have been computed with the parameter $\boldsymbol{\theta}_1^*$ of the epoch with the best validation accuracy of Stage 2. Then, we make the final predictions $\boldsymbol{y}_{\text{final, i}} = f_2(\boldsymbol{x}_i, \boldsymbol{p}_i, \boldsymbol{y}'_i; \boldsymbol{\theta}_2)$ by combining the ego embeddings, PEs, and the (back)-propagated soft labels ($\boldsymbol{\theta}_2$ is a learnable parameter vector). We use the soft-labels $\tilde{\boldsymbol{y}}_i$ instead of the actual labels one-hot (y_i) in order to avoid label leakage, which hurts performance (see also (Shi et al., 2020) for a different way to combat label leakage). Finally, we maximize the cross-entropy with respect to a node's own labels,

$$\mathcal{L}_{\text{CE, 2}}(\boldsymbol{\theta}_2) = \sum_{i \in V_{\text{train}}} \sum_{l \in [c]} \mathbb{I}\{y_i = l\} \log(\boldsymbol{y}_{\text{final, i},l}), \tag{4}$$

Overall, Stage 2 corresponds to learning the neighbor distributions and Stage 3 uses these distributions to train a new model which predicts a node's own labels. In Section 3.4, we prove that such a procedure incurs lower error than using the features to predict a node's own labels.

Relationship to Label Propagation: From a label propagation perspective, one could argue that our method is related to the HITS algorithm of Kleinberg (1999). Besides, on a directed graph a potential solution would be to perform label propagation with the Gram matrix AA^T and use this information for as the propagated labels instead of $\boldsymbol{y}'_i$. This, however, would not be inductive.

Scalability: GLINKX is highly scalable as it performs message passing a constant number of times by paying an $O(mc)$ cost, where the dimensionality of classes c is usually small (compared to d_X that GCNs rely on). In both Stages 2 and 3 of Alg. 1 we train node-level MLPs which allow us to leverage i.i.d. (row-wise) minibatching, like tabular data, and thus our complexity is similar to other shallow methods (LINKX, FSGNN) (Lim et al., 2021; Maurya et al., 2021). This, combined with the propagations that happen outside of the training loops, circumvent the scalability issues of GCNs. For more details, refer App. A.1.

3.3 Varying Homophily

Graphs with monophily experience homophily, heterophily, or both. For instance, in the yelp-chi dataset – where we classify a review as spam/non-spam (see Fig. 3) – we observe a case of monophily together with varying homophily. Specifically in this dataset, spam reviews are linked to non-spam

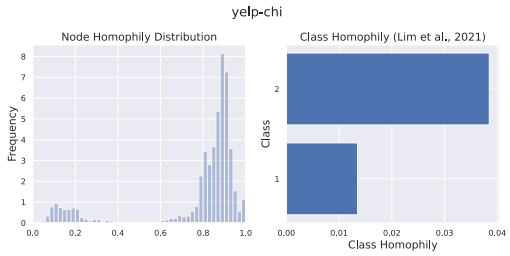

Figure 3: Node Homophily Distribution and Class (or Edge-insensitive) Homophily (Lim et al., 2021).

reviews and non-spam reviews usually connect to other non-spam reviews, which makes the node homophily distribution bimodal. Here the 2nd-order similarity makes both the MLaP mechanism particularly effective and PEs, since the PEs can be used to distinguish nodes that have similar features but are located in different regions of the graph (homophilous/heterophilous region).

3.4 Theoretical Analysis

Justification of Stage 2: In Stage 2, we train a parametric model to learn the distribution of a node's neighbors from the node features $\boldsymbol{\xi}_i$[3]. Arguably, such a distribution can be learned naïvely by counting the neighbors i that belong to each class. This motivates our first theoretical result. In a nutshell, we show that training a parametric model for learning the distribution of a node's neighbors (as we do in Stage 2) yields a lower error than the naïve solution. Below we present the Thm. 1 (proof in App. F), for undirected graphs (the case of directed graphs is the same, but we omit it for simplicity of exposition):

Theorem 1. *Let $G([n], E)$ be an undirected graph of minimum degree $K > c^2$ and let $\boldsymbol{Q}_i \in \Gamma_c$ be the likelihood, from the viewpoint of node i, of any node in its neighborhood $\mathcal{N}(i)$ to be assigned to different classes for every node $i \in [n]$. The following two facts are true (under standard assumptions for SGD and the losses):*

1. *Let $\widehat{\boldsymbol{Q}}_i$ be the sample average of $\boldsymbol{Q}_i$, i.e. $\widehat{Q}_{i,j} = \frac{1}{|\mathcal{N}(i)|}\sum_{k\in\mathcal{N}(i)} \mathbb{I}\{y_k = j\}$. Then, for every $i \in [n]$, we have that $\max_{j\in[c]} \mathbb{E}[|Q_{i,j} - \widehat{Q}_{i,j}|] \leq \mathbb{E}[\|\boldsymbol{Q}_i - \widehat{\boldsymbol{Q}}_i\|_\infty] \leq O\left(\sqrt{\frac{\log(Kc)}{K}}\right)$.*
2. *Let $q(\cdot|\boldsymbol{\xi}_i;\boldsymbol{\theta})$ be a model parametrized by $\boldsymbol{\theta} \in \mathbb{R}^D$ that uses the features $\boldsymbol{\xi}_i$ of each node i to predict $\boldsymbol{Q}_i$. We estimate the parameter $\boldsymbol{\theta}_1$ by running SGD for $t = n$ steps to maximize $\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{c} Q_{i,j} \log q(j|\boldsymbol{\xi}_i;\boldsymbol{\theta})$. Then, for every $i \in [n]$, we have that $\max_{j\in[c]} \mathbb{E}[|q(j;\boldsymbol{\xi}_i;\boldsymbol{\theta}_1) - Q_{i,j}|] \leq O\left(\sqrt{\frac{\log n}{n}}\right)$.*

It is evident here that if the minimum degree K is much smaller than n, then the parametric model has lower error than the naïve approach, namely $\tilde{O}(n^{-1/2})$ compared to $\tilde{O}(K^{-1/2})$.

Justification of Stages 2 and 3: We now provide theoretical foundations for the 2 stage (stages 2 and 3) approach. Specifically, we argue that a two-stage procedure involving learning the distribution of a node's 2nd-hop neighbor distributions (we assume for simplicity, again, that the graph is undirected) first with a parametric model such as in Thm. 1, and then running a two-phase algorithm to learn a parametric model that predicts a node's own label, yields a lower error than naïvely training a shallow parametric model to learn a node's own labels. The first phase of the two-phase algorithm involves training the model first by minimizing the cross-entropy between the predictions and the 2nd-hop neighborhood distribution. Then the model trains a joint objective which uses the learned neighbor distributions and the true labels starting from the model learned in the previous phase. Our result follows (proof in App. F):

Theorem 2. *Let $G([n], E)$ be an undirected graph of minimum degree $K > c^2$ and, let $\boldsymbol{P}_i$ be the likelihood of node i to be assigned to a different class, and let $\boldsymbol{Q}_i, q(\cdot|\boldsymbol{\xi}_i;\boldsymbol{\theta}_1)$ defined as in*

[3] In Section 3.1, $\boldsymbol{\xi}_i$s correspond to the augmented features $\boldsymbol{\xi}_i = [\boldsymbol{x}_i; \boldsymbol{p}_i]$

Thm. 1. Let $p(\cdot|\boldsymbol{\xi}_i;\boldsymbol{w})$ be a model parametrized by $\boldsymbol{w} \in \mathbb{R}^D$ that is used to predict the class assignments $y_i \sim p(\cdot|\boldsymbol{\xi}_i;\boldsymbol{w})$. Let $\boldsymbol{w}_*$ be the optimal parameter. The following are true (under standard assumptions for SGD and the losses):

1. *The naïve optimization scheme that runs SGD to maximize $\mathcal{G}(\boldsymbol{w}) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{c} P_{i,j}\log p(j|\boldsymbol{\xi}_i;\boldsymbol{w})$ for n steps. Then $\mathbb{E}[\mathcal{G}(\boldsymbol{w}_{n+1}) - \mathcal{G}(\boldsymbol{w}_*)] \leq O\left(\frac{\log n}{n}\right)$.*
2. *The two-phase optimization scheme that runs SGD to maximize $\widehat{\mathcal{G}}(\boldsymbol{w}) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{c}\left(\frac{1}{|\mathcal{N}(i)|}\sum_{k\in\mathcal{N}(i)} q(j|\boldsymbol{\xi}_k;\boldsymbol{\theta}_1)\right)\log p(j|\boldsymbol{\xi}_i;\boldsymbol{w})$ for n_1 steps, to estimate a solution $\boldsymbol{w}'$ and then runs SGD on the objective $\lambda\widehat{\mathcal{G}}(\boldsymbol{w}) + (1-\lambda)\mathcal{G}(\boldsymbol{w})$ for n steps starting from $\boldsymbol{w}'$, achieves error $\mathbb{E}[\mathcal{G}(\boldsymbol{w}_{n+1}) - \mathcal{G}(\boldsymbol{w}_*)] \leq O\left(\frac{\sqrt{\log n \log\log n}}{n}\right)$.*

We observe that the two-phase optimization scheme is able to reduce the error by a factor of $\sqrt{\log n/\log\log n}$ highlighting the importance of using the distribution of the 2nd-hop neighbors of a node in order to predict its own label. Also, note that the above two-phase optimization scheme is different from the description of the method we gave in Alg. 1. The difference is that the distribution of a node's neighbors is embedded into the model in the case of Alg. 1, and the distribution of a node's neighbors is embedded into the loss function in Thm. 2 as a regularizer. In Alg. 1, we chose to incorporate this information in the model because using multiple losses harms scalability and makes training harder in practice. In the same spirit, the conception of GCNs (Kipf & Welling, 2016) replaces explicit regularization with the graph Laplacian with the graph topology into the model (see also (Hamilton et al., 2017; Yang et al., 2016)), and, similarly to GCNs, label information as a feature still acts a regularizer in our model.

3.5 Complementarity

Different components of GLINKX provide a *complementary* signal to components proposed in the GNN literature (Maurya et al., 2021; Zhang et al., 2022b; Rossi et al., 2020). One can combine GLINKX with existing architectures (e.g. feature propagations (Maurya et al., 2021; Rossi et al., 2020), label propagations (Zhang et al., 2022b)) for potential metric gains. For example, SIGN computes a series of $r \in \mathbb{N}$ feature propagations $[X, \Phi X, \Phi^2 X, \ldots, \Phi^r X]$ where Φ is a matrix (e.g. normalized adjacency or normalized Laplacian) as a preprocessing step. We can include this complementary signal, namely embed each of the propagated features and combine them in the final layer of the 3rd Stage, to GLINKX. Overall, even though in this paper we want to keep GLINKX simple to highlight its main components, we conjecture that adding more components to GLINKX would improve its performance on datasets with highly variable homophily (see Section 3.3).

4 Experiments

4.1 Comparisons

We experiment with homophilous and heterophilous datasets (see Tab. 1 and App. D.3). We train KGEs with Pytorch-Biggraph (Lerer et al., 2019; Yang et al., 2014). For homophilous datasets we compare with vanilla GCN and GAT, FSGNN and Label Propagation (LP). For a fair comparison, we compare with one-layer GCN/GAT/FSGNN/LP since our method is one-hop. We also compare with higher-order (h.o.) GCN/GAT/FSGNN/LP with 2 and 3 layers. In the heterophilous case, we compare with LINKX[4] because it is scalable and is shown to work better than other baselines (e.g. H2GCN), as well as with FSGNN. Note that we do not compare GLINKX with other more complex methods because GLINKX is complementary to methods (see Section 3), and design principles from these methods can be incorporated into GLINKX. We use a *ResNet* module to combine the components from Stages 2 and 3 of our algorithm. Details about the hyperparameters we use are in App. C.

In the heterophilous datasets, GLINKX outperforms LINKX (except arxiv-year where we are within the confidence interval). Moreover, the performance gap between using KGEs and adjacency

[4]We have run our method with hyperparameter space that is a subset of the sweeps reported in (Lim et al., 2021) due to resource constraints. A bigger hyperparameter search would improve our results.

Table 1: Experimental results. (*) = results from the OGB leaderboard.

	Homophilous Datasets		Heterophilous Datasets		
	PubMed	ogbn-arxiv	squirrel	yelp-chi	arxiv-year
n	19.7K	169.3K	5.2K	169.3K	45.9K
m	44.3K	1.16M	216.9K	7.73M	1.16M
Edge-insensitive homophily (Lim et al., 2021)	0.66	0.41	0.02	0.05	0.27
d_X/c	500 / 27	128 / 40	2089 / 5	32 / 2	128 / 5
GLINKX w/ KGEs	87.95 $\pm$0.30	**69.27 $\pm$0.25**	45.83 $\pm$2.89	87.82 $\pm$0.20	**54.09 $\pm$0.61**
GLINKX w/ Adjacency	**88.03 $\pm$0.30**	69.09 $\pm$0.13	**69.15 $\pm$1.87**	**89.32 $\pm$0.45**	53.07 $\pm$0.29
Label Propagation (1-hop)	83.02 $\pm$0.35	69.59 $\pm$0.00	32.22 $\pm$1.45	85.98 $\pm$0.28	43.71 $\pm$0.22
LINKX (from (Lim et al., 2021))	87.86 $\pm$0.77	67.32 $\pm$0.24	61.81 $\pm$1.80	85.86 $\pm$0.40	**56.00 $\pm$1.34**
LINKX (our runs)	87.55 $\pm$0.37	63.91 $\pm$0.18	61.46 $\pm$1.60	88.25 $\pm$0.24	53.78 $\pm$0.06
GCN w/ 1 Layer	86.43 $\pm$0.74	50.76 $\pm$0.20		N/A	
GAT w/ 1 Layer	86.41 $\pm$0.53	54.42 $\pm$0.10		N/A	
FSGNN w/ 1 Layer	88.93 $\pm$0.31	61.82 $\pm$0.84	64.06 $\pm$2.69	86.36 $\pm$0.36	42.86 $\pm$0.22
Higher-order GCN	86.29 $\pm$0.46	71.18 $\pm$0.27 (*)		N/A	
Higher-order GAT	86.64 $\pm$0.40	**73.66 $\pm$0.11** (*)		N/A	
Higher-order FSGNN	**89.37 $\pm$0.49**	69.26 $\pm$0.36	68.04 $\pm$2.19	86.33 $\pm$0.30	44.89 $\pm$0.29
Label Propagation (2-hop)	83.44 $\pm$0.35	69.78 $\pm$0.00	43.41 $\pm$1.44	85.95 $\pm$0.26	46.30 $\pm$0.27
Label Prop. on $\mathbb{I}[A^2 - A - I \geq 0]$	82.14 $\pm$0.33	9.87 $\pm$0.00	24.43 $\pm$1.18	85.68 $\pm$0.32	23.08 $\pm$0.13

Table 2: Ablation Study. We use the hyperparameters of the best run from Tab. 1 with KGEs.

	Ablation Type	Stages	All	Remove ego embeddings	Remove propagation	Remove PEs
Heterophilous	arxiv-year	All Stages	54.09 $\pm$0.61	53.52 $\pm$0.77	50.83 $\pm$0.24	39.06 $\pm$0.35
	arxiv-year	3rd Stage	54.09 $\pm$0.61	53.69 $\pm$0.65	50.83 $\pm$0.24	49.13 $\pm$1.10
Homophilous	ogbn-arxiv	All Stages	69.27 $\pm$0.25	61.26 $\pm$0.33	62.70 $\pm$0.34	65.64 $\pm$0.18
	ogbn-arxiv	3rd Stage	69.27 $\pm$0.25	67.60 $\pm$0.39	62.70 $\pm$0.34	69.62 $\pm$0.15

embeddings shrinks as the dataset grows. In the homophilous datasets GLINKX outperforms 1-layer GCN/GAT/LP/FSGNN and LINKX. In PubMed, GLINKX outperforms h.o. GCN/GAT and in arxiv-year GLINKX is very close to the performance of GCN/GAT.

Finally, we note that our method produces consistent results across regime shifts. In detail, in the heterophilous regime our method performs on par with LINKX, however when we shift to the homophilous regime, LINKX's performance drops, whereas our method's performance continues to be high. Similarly, while FSGNN performs similar to GLINKX on the homophilous datasets, we observe a big performance drop on the heterophilous datasets (see arxiv-year).

4.2 Ablation Study

We ablate each component of Alg. 1 to see the performance contribution from each of the components. We use the hyperparameters of the best model from Tab. 1. We perform two types of ablations: (i) we remove each of the components from all stages of the training, and (ii) we remove the corresponding components only from the 3rd Stage. With the exception of removing the PEs from the 3rd Stage only on ogbn-arxiv, all components contribute to increasing performance on both datasets. Note that adding PEs on the 1st Stage does improve performance, suggesting the primary use case of PEs.

5 Conclusion

We present GLINKX, a scalable method for node classification in homophilous and heterophilous graphs that combines 3 components: (i) ego embeddings, (ii) PEs, and (iii) monophilous propagations. Our method is complementary to what other methods propose, since we can incorporate extra components such as feature propagations, label propagations, and attention to GLINKX. As future work, (i) GLINKX can be extended in heterogeneous graphs, (ii) use more expressive methods such as attention or Wasserstein barycenters (Cuturi & Doucet, 2014) for averaging the low-dimensional messages, and (iii) add complimentary signals.

REFERENCES

Marjan Albooyeh, Rishab Goel, and Seyed Mehran Kazemi. Out-of-sample representation learning for knowledge graphs. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 2657–2666, 2020.

Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=i80OPhOCVH2.

Kristen M Altenburger and Johan Ugander. Monophily in social networks introduces similarity among friends-of-friends. *Nature human behaviour*, 2(4):284–290, 2018a.

Kristen M Altenburger and Johan Ugander. Node attribute prediction: An evaluation of within-versus across-network tasks. In *NeurIPS Workshop on Relational Representation Learning*, 2018b.

Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

K. Bhatia, K. Dahiya, H. Jain, P. Kar, A. Mittal, Y. Prabhu, and M. Varma. The extreme classification repository: Multi-label datasets and code, 2016. URL http://manikvarma.org/downloads/XC/XMLRepository.html.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26, 2013.

Xavier Bresson and Thomas Laurent. Residual gated graph convnets. *arXiv preprint arXiv:1711.07553*, 2017.

Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. Adaptive universal generalized pagerank graph neural network. *arXiv preprint arXiv:2006.07988*, 2020.

Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 685–693, Bejing, China, 22–24 Jun 2014. PMLR. URL https://proceedings.mlr.press/v32/cuturi14.html.

Francesco Di Giovanni, James Rowbottom, Benjamin P Chamberlain, Thomas Markovich, and Michael M Bronstein. Graph neural networks as gradient flows. *arXiv preprint arXiv:2206.10991*, 2022.

Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pp. 315–324, 2020.

Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph neural networks with learnable structural and positional representations. *arXiv preprint arXiv:2110.07875*, 2021.

Ahmed El-Kishky, Thomas Markovich, Serim Park, Chetan Verma, Baekjin Kim, Ramy Eskander, Yury Malkov, Frank Portman, Sofía Samaniego, Ying Xiao, et al. Twhin: Embedding the twitter heterogeneous information network for personalized recommendation. *arXiv preprint arXiv:2202.05387*, 2022.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.

Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*, 2020.

Di Jin, Rui Wang, Meng Ge, Dongxiao He, Xiang Li, Wei Lin, and Weixiong Zhang. Raw-gnn: Random walk aggregation based graph neural network. *arXiv preprint arXiv:2206.13953*, 2022a.

Wei Jin, Lingxiao Zhao, Shichang Zhang, Yozen Liu, Jiliang Tang, and Neil Shah. Graph condensation for graph neural networks. In *International Conference on Learning Representations*, 2022b. URL https://openreview.net/forum?id=WLEx3Jo4QaB.

Jinwoo Kim, Tien Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. Pure transformers are powerful graph learners. *arXiv preprint arXiv:2207.02505*, 2022.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Jon M Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.

Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. Pytorch-biggraph: A large scale graph embedding system. *Proceedings of Machine Learning and Systems*, 1:120–131, 2019.

Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. Deepgcns: Can gcns go as deep as cnns? In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9267–9276, 2019.

Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and Ser Nam Lim. Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods. *Advances in Neural Information Processing Systems*, 34:20887–20902, 2021.

Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. Is heterophily a real nightmare for graph neural networks to do node classification? *arXiv preprint arXiv:2109.05641*, 2021.

Sunil Kumar Maurya, Xin Liu, and Tsuyoshi Murata. Improving graph neural networks with simple architecture design. *arXiv preprint arXiv:2105.07634*, 2021.

Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pp. 415–444, 2001.

Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 4602–4609, 2019.

Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020.

Everett M Rogers, Arvind Singhal, and Margaret M Quinlan. Diffusion of innovations. In *An integrated approach to communication theory and research*, pp. 432–448. Routledge, 2014.

Emanuele Rossi, Fabrizio Frasca, Ben Chamberlain, Davide Eynard, Michael Bronstein, and Federico Monti. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*, 7:15, 2020.

Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-scale attributed node embedding. *Journal of Complex Networks*, 9(2):cnab014, 2021.

Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.

Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*, 2020.

Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June Hsu, and Kuansan Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, pp. 243–246, 2015.

Balasubramaniam Srinivasan and Bruno Ribeiro. On the equivalence between positional node embeddings and structural graph representations. *arXiv preprint arXiv:1910.00452*, 2019.

Chuxiong Sun, Hongming Gu, and Jie Hu. Scalable and adaptive graph neural networks with self-label-enhanced training. *arXiv preprint arXiv:2104.09376*, 2021.

Petar Velivcković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=ryGs6iA5Km`.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.

Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pp. 40–48. PMLR, 2016.

Shichang Zhang, Yozen Liu, Yizhou Sun, and Neil Shah. Graph-less neural networks: Teaching old MLPs new tricks via distillation. In *International Conference on Learning Representations*, 2022a. URL `https://openreview.net/forum?id=4p6_5HBWPCw`.

Wentao Zhang, Ziqi Yin, Zeang Sheng, Yang Li, Wen Ouyang, Xiaosen Li, Yangyu Tao, Zhi Yang, and Bin Cui. Graph attention multi-layer perceptron. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022b.

Wenqing Zheng, Edward W Huang, Nikhil Rao, Sumeet Katariya, Zhangyang Wang, and Karthik Subbian. Cold brew: Distilling graph node representations with incomplete or missing neighborhoods. In *International Conference on Learning Representations*, 2022a. URL `https://openreview.net/forum?id=1ugNpm7W6E`.

Xin Zheng, Yixin Liu, Shirui Pan, Miao Zhang, Di Jin, and Philip S Yu. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082*, 2022b.

Zhiqiang Zhong, Sergey Ivanov, and Jun Pang. Simplifying node classification on heterophilous graphs with compatible label propagation. *arXiv preprint arXiv:2205.09389*, 2022.

Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33:7793–7804, 2020.

Jiong Zhu, Ryan A Rossi, Anup Rao, Tung Mai, Nedim Lipka, Nesreen K Ahmed, and Danai Koutra. Graph neural networks with heterophily. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11168–11176, 2021.