

Semi-Supervised Node Classification on Graphs: Markov Random Fields vs. Graph Neural Networks

Binghui Wang, Jinyuan Jia, Neil Zhenqiang Gong

ECE Department, Duke University, Durham, NC 27708, USA
 {binghui.wang, jinyuan.jia, neil.gong}@duke.edu

Abstract

Semi-supervised node classification on graph-structured data has many applications such as fraud detection, fake account and review detection, user’s private attribute inference in social networks, and community detection. Various methods such as pairwise Markov Random Fields (pMRF) and graph neural networks were developed for semi-supervised node classification. pMRF is more efficient than graph neural networks. However, existing pMRF-based methods are less accurate than graph neural networks, due to a key limitation that they assume a heuristics-based constant edge potential for all edges. In this work, we aim to address the key limitation of existing pMRF-based methods. In particular, we propose to learn edge potentials for pMRF. Our evaluation results on various types of graph datasets show that our optimized pMRF-based method consistently outperforms existing graph neural networks in terms of both accuracy and efficiency. Our results highlight that previous work may have underestimated the power of pMRF for semi-supervised node classification.

Introduction

Given an undirected graph, a feature vector for each node in the graph, and a small set of labeled nodes, semi-supervised node classification aims to classify the remaining unlabeled nodes simultaneously. Semi-supervised node classification has many applications such as fraud detection (Pandit et al. 2007; Chau et al. 2011; Ye et al. 2011; Tamersoy, Roundy, and Chau 2014; Hooi et al. 2016), fake account and review detection (Wang et al. 2011; Cao et al. 2012; Akoglu, Chandy, and Faloutsos 2013; Rayana and Akoglu 2015; Wang, Zhang, and Gong 2017; Wang, Gong, and Fu 2017; Gao et al. 2018; Wang et al. 2018), user’s private attribute inference in social networks (Gong et al. 2014; Jia et al. 2017; Gong and Liu 2018; Wang, Jia, and Gong 2019), and community detection (He et al. 2018; Jin et al. 2019a,b). Therefore, many methods have been proposed for semi-supervised node classification. These include conventional methods such as label propagation (LP) (Zhu, Ghahramani, and Lafferty 2003), manifold regularization (ManiReg) (Belkin, Niyogi, and Sindhwani 2006), deep semi-supervised embedding (SemiEmb) (Weston et al. 2012), iterative classification algorithm (ICA) (Lu and Getoor 2003), and pairwise

Markov Random Fields (pMRF) (Gatterbauer 2017), as well as recent graph neural networks (Duvenaud et al. 2015; Kipf and Welling 2017; Hamilton, Ying, and Leskovec 2017; Battaglia et al. 2018; Xu et al. 2018; Gao, Wang, and Ji 2018; Ying et al. 2018; Veličković et al. 2018; Wu et al. 2019; Qu, Bengio, and Tang 2019; Ma et al. 2019a; Wang et al. 2020).

In particular, pMRF associates a discrete random variable with each node to model its label. pMRF defines a *node potential* for each node, which is a function of the node’s random variable and models the node’s label using the node’s feature vector. pMRF also defines an *edge potential* for each edge (u, v) , which is a function of the two random variables associated with nodes u and v . Finally, pMRF defines a joint probability distribution of all random variables as the product of the node potentials and edge potentials. The edge potential for an edge (u, v) is essentially a matrix, whose (i, j) th entry indicates the likelihood that u and v have labels i and j , respectively. Therefore, an edge potential is also called *coupling matrix*. Moreover, the standard Belief Propagation (Pearl 1988) or linearized Belief Propagation (Gatterbauer et al. 2015; Gatterbauer 2017; Jia et al. 2017) were often used to perform inference on the pMRF. Compared to graph neural networks, pMRF inferred via linearized Belief Propagation is much more efficient in terms of running time. However, existing pMRF-based methods (Pandit et al. 2007; Tamersoy, Roundy, and Chau 2014; Akoglu, Chandy, and Faloutsos 2013; Rayana and Akoglu 2015; Jia et al. 2017; Wang, Gong, and Fu 2017) are often less accurate than graph neural networks, due to a key limitation: they assume a heuristics-based constant edge potential or coupling matrix for all edges.

In this work, we aim to address the limitation of existing pMRF-based methods. In particular, we propose a novel method to learn the coupling matrix. Specifically, we formulate learning the coupling matrix as an optimization problem, where the objective function is to learn a coupling matrix such that 1) the training loss is small (modeled as cross-entropy loss) and 2) the coupling matrix of an edge (u, v) is consistent with the predicted labels of nodes u and v (modeled as a regularization term). Moreover, we propose a novel algorithm to solve the optimization problem. Our results on various types of graph datasets show that our method is more accurate and efficient than existing graph neural networks.

Problem Definition and Background

Problem Definition

We aim to solve the following *semi-supervised node classification problem* on graph-structured data.

Definition 1. Suppose we are given an undirected graph $G = (V, E)$ with $n = |V|$ nodes and $m = |E|$ edges. Each node v has an F -dimension feature vector \mathbf{x}_v , and each node has a label from a label set $\{1, 2, \dots, C\}$. Moreover, we are given a small set of labeled nodes L , which we call training dataset. Our goal is to predict the labels of the remaining unlabeled nodes using the nodes' feature vectors and graph structure.

Pairwise Markov Random Field (pMRF)

pMRF was traditionally widely used for semi-supervised node classification in various applications such as auction fraud detection (Pandit et al. 2007), fake accounts and reviews detection (Wang et al. 2011; Cao et al. 2012; Akoglu, Chandy, and Faloutsos 2013; Rayana and Akoglu 2015; Wang, Zhang, and Gong 2017; Wang, Gong, and Fu 2017; Wang, Jia, and Gong 2019), malware detection (Tamersoy, Roundy, and Chau 2014), and user's private attribute inference in social networks (Gong et al. 2014; Jia et al. 2017; Gong and Liu 2018; Wang, Jia, and Gong 2019). Specifically, we associate a random variable r_v with each node v to model its label, where $r_v = i$ means that v 's label is i . pMRF essentially models the joint probability distribution of the random variables r_1, r_2, \dots, r_n , where the statistical correlations/independence between the random variables are captured by the graph structure. Formally, pMRF models the joint probability distribution as follows:

$$\Pr(r_1, r_2, \dots, r_n) \propto \prod_{v \in V} \phi_v(r_v) \prod_{(u,v) \in E} \psi_{uv}(r_u, r_v), \quad (1)$$

where $\phi_v(r_v)$ is called node potential and $\psi_{uv}(r_u, r_v)$ is called edge potential. For convenience, we represent the node potential $\phi_v(r_v)$ as a probability distribution $\tilde{\mathbf{q}}_v$ over the C labels $\{1, 2, \dots, C\}$. Specifically, $\tilde{\mathbf{q}}_v$ is a row vector and $\tilde{q}_v(i) = \phi_v(r_v = i)$ for $i \in \{1, 2, \dots, C\}$. If nodes' feature vectors are available, $\tilde{\mathbf{q}}_v$ can be learnt. In particular, given the labeled nodes L and nodes' feature vectors, we learn a multi-class logistic regression classifier and use the classifier to predict the probability distribution $\tilde{\mathbf{q}}_v$ for each node. If nodes's feature vectors are unavailable, we can assign $\tilde{\mathbf{q}}_v$ based on the labeled nodes L . Specifically, if node v is a labeled node, $\tilde{\mathbf{q}}_v$ is set as the node v 's one-hot label vector; otherwise, all entries in $\tilde{\mathbf{q}}_v$ have the equal value $1/C$. Moreover, we represent the edge potential $\psi_{uv}(r_u, r_v)$ as a symmetric matrix $\tilde{\mathbf{H}}_{uv} \in \mathbb{R}^{C \times C}$, where $\tilde{H}_{uv}(i, j) = \psi_{uv}(r_u = i, r_v = j)$ for $i, j \in \{1, 2, \dots, C\}$. $\tilde{\mathbf{H}}_{uv}$ models the coupling strength between the labels of nodes u and v , and thus $\tilde{\mathbf{H}}_{uv}$ is also called *coupling matrix*. In particular, a larger $\tilde{H}_{uv}(i, j)$ means that u and v are more likely to have labels i and j , respectively.

Inference via linearized Belief Propagation (LinBP): Given the joint probability distribution in Equation 1, the probability distribution $\tilde{\mathbf{q}}_v$ for each node,

and the coupling matrix $\tilde{\mathbf{H}}_{uv}$ for each edge, we infer the marginal probability distribution $\tilde{\mathbf{p}}_v$ (a row vector) for each node v and use it to predict v 's label. Specifically, Belief Propagation (BP) (Pearl 1988) is a standard method to infer the marginal probability distributions. However, BP is not guaranteed to converge on loopy graphs and is not scalable due to message maintenance on each edge. Recent work (Gatterbauer et al. 2015; Gatterbauer 2017) proposed to linearize BP (LinBP) to address these limitations. Given LinBP, the *approximate* marginal probability distributions \mathbf{p}_v are solutions to the following linear system (Gatterbauer 2017):

$$\mathbf{p}_v = \mathbf{q}_v + \sum_{u \in \Gamma_v} \mathbf{p}_u \mathbf{H}_{uv}, \quad \forall v \in V, \quad (2)$$

where Γ_v is the set of neighbors of v , $\mathbf{q}_v = \tilde{\mathbf{q}}_v - \frac{1}{C}$, and $\mathbf{H}_{uv} = \tilde{\mathbf{H}}_{uv} - \frac{1}{C}$. Roughly speaking, \mathbf{q}_v and \mathbf{H}_{uv} are centered versions of $\tilde{\mathbf{q}}_v$ and $\tilde{\mathbf{H}}_{uv}$, respectively. Note that \mathbf{p}_v solved by LinBP from Equation 2 is not necessarily a probability distribution any more, e.g., it may not sum to 1. The label of v is predicted as the one that has the largest value in \mathbf{p}_v .

Limitation: The key limitation of existing pMRF-based semi-supervised node classification methods is that they assume a heuristics-based constant coupling matrix for all edges. For example, $\tilde{H}_{uv}(i, i) = 0.9$ and $\tilde{H}_{uv}(i, j) = \frac{0.1}{C-1}$, where $i, j \in \{1, 2, \dots, C\}$ and $i \neq j$, which means that two linked nodes are more likely to have the same label and does not distinguish between labels and between edges. In this work, we address this limitation via proposing a method to learn the coupling matrix. Moreover, we empirically show that, after learning the coupling matrix, pMRF outperforms state-of-the-art graph neural networks on multiple benchmark datasets.

Learning the Coupling Matrix (LCM)

Formulating an Optimization Problem

Constraining the coupling matrices: In pMRF, each edge has a coupling matrix $\tilde{\mathbf{H}}_{uv}$. Therefore, it is challenging to learn a different coupling matrix for each edge as the number of parameters is much larger than the training dataset size. To address the challenge, we constrain the coupling matrices of different edges as different scalings of the same coupling matrix. In particular, we assume $\mathbf{H}_{uv} = W_{uv} \mathbf{H}$ for each edge $(u, v) \in E$. We model the scaling parameter W_{uv} as the weight of the edge (u, v) and we denote all the scaling parameters as a symmetric weight matrix \mathbf{W} of the graph. Given such assumption on the coupling matrices, we can transform Equation 2 as follows:

$$\mathbf{P} = \mathbf{Q} + \mathbf{WPH}, \quad (3)$$

where the matrix $\mathbf{P} = [\mathbf{p}_1; \mathbf{p}_2; \dots; \mathbf{p}_n]$ and the matrix $\mathbf{Q} = [\mathbf{q}_1; \mathbf{q}_2; \dots; \mathbf{q}_n]$. According to (Gatterbauer et al. 2015), Equation 3 has a solution when $\mathbf{I} - \mathbf{H} \otimes \mathbf{W}$ is invertible, where \mathbf{I} is an identity matrix and \otimes is the Kronecker product. LinBP uses the power iteration method to solve Equation 3 and it can find the solution when $\rho(\mathbf{H}) < \frac{1}{\rho(\mathbf{W})}$, where $\rho(\cdot)$ is the spectral radius of a matrix.¹

¹The spectral radius of a square matrix is the maximum of the absolute values of its eigenvalues.

Formulating an optimization problem: Intuitively, we aim to achieve two goals when learning the coupling matrix and weight matrix. The first goal is that the labels predicted for the training nodes based on the learnt coupling matrix and weight matrix should match the ground truth labels of the training nodes. The second goal is that the predicted label distributions $\sigma(\mathbf{p}_u)$ and $\sigma(\mathbf{p}_v)$ of nodes u and v should be *consistent* with the coupling matrix $W_{uv}\mathbf{H}$ for each edge (u, v) . Specifically, if u and v are more likely to be predicted to have labels i and j (i.e., the probabilities $\sigma(\mathbf{p}_u)_i$ and $\sigma(\mathbf{p}_v)_j$ are larger), respectively, then the entry $W_{uv}H(i, j)$ should be larger. This is because $W_{uv}H(i, j)$ encodes the likelihood that u and v have labels i and j , respectively.

To quantify the first goal, we adopt a loss function over the training dataset. Specifically, we consider a standard cross-entropy loss function $L(\mathbf{W}, \mathbf{H})$ as follows:

$$L(\mathbf{W}, \mathbf{H}) = - \sum_{l \in L} \mathbf{y}_l^T \log \sigma(\mathbf{p}_l), \quad (4)$$

where σ is the softmax function, which is defined as $\sigma(\mathbf{p}_l)_i = \frac{\exp(\mathbf{p}_l(i))}{\sum_{j=1}^C \exp(\mathbf{p}_l(j))}$ and converts the vector \mathbf{p}_l to be a probability distribution; the vector \mathbf{y}_l is the one-hot encoding of the node l 's label, i.e., if node l belongs to the i th class, then $y_l(i) = 1$ and $y_l(j) = 0$, for all $j \neq i$. To quantify the second goal, we define the following term $R(\mathbf{W}, \mathbf{H})$:

$$R(\mathbf{W}, \mathbf{H}) = - \sum_{(u,v) \in E} \sigma(\mathbf{p}_u) \cdot W_{uv} \mathbf{H} \cdot \sigma(\mathbf{p}_v)^T. \quad (5)$$

$R(\mathbf{W}, \mathbf{H})$ is smaller if the second goal is better satisfied, i.e., the predicted label distributions of two nodes are more consistent with the coupling matrix of the corresponding edge. Combining the two goals, we propose to learn the coupling matrix and weight matrix via solving the following optimization problem:

$$\min_{\mathbf{W}, \mathbf{H}} \mathcal{L}(\mathbf{W}, \mathbf{H}) = L(\mathbf{W}, \mathbf{H}) + \lambda R(\mathbf{W}, \mathbf{H}), \quad (6)$$

where $\lambda > 0$ is a hyperparameter that balances between the two goals. Our term $R(\mathbf{W}, \mathbf{H})$ can also be viewed as a regularization term, and we call it *consistency regularization*. We note that a similar consistency regularization term was also proposed by (Wang, Jia, and Gong 2019), but only for binary classification.

Challenge for solving the optimization problem: It is computationally challenging to solve the optimization problem in Equation 6. The reason is each \mathbf{p}_v depends on every edge weight and each entry of the coupling matrix. For example, suppose we use gradient descent to solve the optimization problem. In each iteration, we need to compute the gradient $\frac{\partial \mathcal{L}(\mathbf{W}, \mathbf{H})}{\partial W_{uv}}$, which involves the gradient $\frac{\partial \mathbf{P}}{\partial W_{uv}}$. However, according to Equation 3, $\frac{\partial \mathbf{P}}{\partial W_{uv}}$ is a solution to the following system:

$$\frac{\partial \mathbf{P}}{\partial W_{uv}} = \mathbf{W} \frac{\partial \mathbf{P}}{\partial W_{uv}} \mathbf{H} + \frac{\partial \mathbf{W}}{\partial W_{uv}} \mathbf{P} \mathbf{H}. \quad (7)$$

Therefore, in each iteration of gradient descent, we need to solve the above system of equations for *each* edge, which is computationally infeasible for large graphs.

Solving the Optimization Problem

We propose an approximate algorithm to efficiently solve the optimization problem. Our key observation is that \mathbf{P} is often iteratively solved from Equation 3 as follows:

$$\mathbf{P}^{(t)} = \mathbf{Q} + \mathbf{W} \mathbf{P}^{(t-1)} \mathbf{H}, \quad (8)$$

where $\mathbf{P}^{(t)}$ is the matrix \mathbf{P} in the t th iteration. Based on the observation, we propose to *alternately* update \mathbf{P} and learn the weight matrix \mathbf{W} and coupling matrix \mathbf{H} . Specifically, given the current \mathbf{W} , \mathbf{H} , and \mathbf{P} , we update \mathbf{P} in the next iteration. Then, given the current \mathbf{P} , we learn \mathbf{W} and \mathbf{H} such that 1) the matrix \mathbf{P} in the next iteration has a small cross-entropy loss on the training dataset and 2) the consistency regularization term is large for the current \mathbf{P} .

Updating $\mathbf{P}^{(t)}$: Given the weight matrix $\mathbf{W}^{(t-1)}$, the coupling matrix $\mathbf{H}^{(t-1)}$, and the matrix $\mathbf{P}^{(t-1)}$ in the $(t-1)$ th iteration, we compute $\mathbf{P}^{(t)}$ in the t th iteration as follows:

$$\mathbf{P}^{(t)} = \mathbf{Q} + \mathbf{W}^{(t-1)} \mathbf{P}^{(t-1)} \mathbf{H}^{(t-1)}. \quad (9)$$

Learning $\mathbf{W}^{(t)}$ and $\mathbf{H}^{(t)}$: Given the weight matrix $\mathbf{W}^{(t-1)}$ and the coupling matrix $\mathbf{H}^{(t-1)}$ in the $(t-1)$ th iteration as well as the matrix $\mathbf{P}^{(t)}$ in the t th iteration, we learn the weight matrix $\mathbf{W}^{(t)}$ and the coupling matrix $\mathbf{H}^{(t)}$ in the t th iteration as a solution to the following optimization problem:

$$\min_{\mathbf{W}^{(t)}, \mathbf{H}^{(t)}} \mathcal{L}(\mathbf{W}^{(t)}, \mathbf{H}^{(t)}) = - \sum_{l \in L} \mathbf{y}_l^T \log \sigma(\mathbf{p}_l^{(t+1)}) - \lambda \sum_{(u,v) \in E} \sigma(\mathbf{p}_u^{(t)}) \cdot W_{uv}^{(t)} \mathbf{H}^{(t)} \cdot (\sigma(\mathbf{p}_v^{(t)}))^T, \quad (10)$$

where $\mathbf{P}^{(t+1)} = \mathbf{Q} + \mathbf{W}^{(t)} \mathbf{P}^{(t)} \mathbf{H}^{(t)}$. We leverage gradient descent to learn $\mathbf{W}^{(t)}$ and $\mathbf{H}^{(t)}$. Specifically, we have the following gradients:

$$\frac{\partial \mathcal{L}(\mathbf{W}^{(t)}, \mathbf{H}^{(t)})}{\partial W_{uv}^{(t)}} = - \sum_{l \in L} \sum_{j=1}^C (y_l(j) - \sigma(\mathbf{p}_l^{(t+1)})_j) \frac{\partial p_l^{(t+1)}(j)}{\partial W_{uv}^{(t)}} - \lambda \sigma(\mathbf{p}_u^{(t)}) \cdot \mathbf{H}^{(t)} \cdot (\sigma(\mathbf{p}_v^{(t)}))^T, \quad (11)$$

$$\begin{aligned} \frac{\partial \mathcal{L}(\mathbf{W}^{(t)}, \mathbf{H}^{(t)})}{\partial H_{ij}^{(t)}} &= - \sum_{l \in L} \left((y_l(j) - \sigma(\mathbf{p}_l^{(t+1)})_j) \frac{\partial p_l^{(t+1)}(j)}{\partial H_{ij}^{(t)}} \right. \\ &\quad \left. + (y_l(i) - \sigma(\mathbf{p}_l^{(t+1)})_i) \frac{\partial p_l^{(t+1)}(i)}{\partial H_{ij}^{(t)}} \right) \\ &\quad - \lambda \sum_{(u,v) \in E} W_{uv}^{(t)} \cdot \sigma(\mathbf{p}_u^{(t)})_i \cdot \sigma(\mathbf{p}_v^{(t)})_j, \end{aligned} \quad (12)$$

where we have:

$$\frac{\partial p_l^{(t+1)}(j)}{\partial W_{uv}^{(t)}} = \begin{cases} (\mathbf{P}^{(t)} \mathbf{H}^{(t)})_{vj}, & \text{if } u = l \\ (\mathbf{P}^{(t)} \mathbf{H}^{(t)})_{uj}, & \text{if } v = l \\ 0, & \text{otherwise,} \end{cases} \quad (13)$$

$$\frac{\partial p_l^{(t+1)}(j)}{\partial H_{ij}^{(t)}} = (\mathbf{W}^{(t)} \mathbf{P}^{(t)})_{li}, \quad (14)$$

$$\frac{\partial p_l^{(t+1)}(i)}{\partial H_{ij}^{(t)}} = (\mathbf{W}^{(t)} \mathbf{P}^{(t)})_{lj}. \quad (15)$$

We have both $\partial p_i^{(t+1)}(j)/\partial H_{ij}^{(t)}$ and $\partial p_i^{(t+1)}(i)/\partial H_{ij}^{(t)}$ when computing $\partial \mathcal{L}(\mathbf{W}^{(t)}, \mathbf{H}^{(t)})/\partial H_{ij}^{(t)}$ because \mathbf{H} is symmetric. We initialize $\mathbf{W}^{(t)}$ and $\mathbf{H}^{(t)}$ as $\mathbf{W}^{(t-1)}$ and $\mathbf{H}^{(t-1)}$, respectively. Then, we use gradient descent to update $\mathbf{W}^{(t)}$ and $\mathbf{H}^{(t)}$ with learning rates γ_1 and γ_2 , respectively.

We note that the learnt weight matrix \mathbf{W} and coupling matrix \mathbf{H} are bounded during optimization. This is because the edge potential matrix $\tilde{\mathbf{H}}_{uv}$ is non-negative and each row sums to 1. Moreover, $W_{uv}\mathbf{H} = \mathbf{H}_{uv} = \tilde{\mathbf{H}}_{uv} - 1/C$. Therefore, $W_{uv}\mathbf{H}$ is bounded. During optimization, we can clip $W_{uv}\mathbf{H}$ to satisfy such constraints if they are too large. However, we found that $W_{uv}\mathbf{H}$ is small even if we do not clip in our experiments. This is because of our initializations of W_{uv} and $\tilde{\mathbf{H}}_{uv}$ and that our LCM method requires a small number of iterations.

Evaluation

Benchmark Datasets

Following previous work (Yang, Cohen, and Salakhutdinov 2016; Kipf and Welling 2017), we use three benchmark citation graphs (i.e., Cora, Citeseer, and Pubmed) (Sen et al. 2008) and one benchmark knowledge graph NELL (Carlson et al. 2010; Yang, Cohen, and Salakhutdinov 2016). We also adopt two real-world large-scale social graphs (i.e., Google+ (Gong et al. 2012; Jia et al. 2017) and Twitter (Wang, Zhang, and Gong 2017)) to evaluate our method. Descriptions of these datasets and their basic statistics are shown in (Wang, Jia, and Gong 2020).

Training, Validation, and Testing

We split each dataset into a training dataset, a validation set, and a testing dataset. We use the validation dataset to tune hyperparameters and use the testing dataset for evaluation.

For the citation graphs, previous work (Kipf and Welling 2017; Veličković et al. 2018) has set 1,000 nodes in each graph as the testing dataset. For consistent comparisons, we adopt these testing datasets and fix them. Moreover, we sample 20 nodes from each class uniformly at random as the training dataset, and we sample 500 nodes in total uniformly at random as the validation dataset. For the NELL graph, previous work (Yang, Cohen, and Salakhutdinov 2016) has set 969 nodes as the testing dataset, which we adopt. Moreover, like previous work (Kipf and Welling 2017), we consider the extreme case where just one node from each class is sampled for training. We also randomly sample 500 nodes in total as the validation dataset. We repeat sampling the training and validation datasets 5 trials and report the average accuracies and standard deviations on the testing dataset.²

For the Google+ dataset, for each city, we randomly select 1,000 positive users, i.e., who live/lived in the city, and 1,000 negative users, i.e., who do/did not live in the city, as the training dataset; we randomly select 1,000 positive users and

1,000 negative users as the validation dataset; and we treat the remaining users who disclosed at least one city as the testing dataset. Note that we have 50 cities in the Google+ dataset, where each city is treated as a binary classification problem. The accuracy on the Google+ dataset is averaged over the 50 cities. For the Twitter dataset, we randomly sample 1,000 genuine users and 1,000 fake users as the training dataset; we randomly sample 1,000 genuine users and 1,000 fake users as the validation dataset; and we randomly select another 5,000 genuine users and 5,000 fake users as the testing dataset. We sample a relatively larger number of training users in Google+ and Twitter, due to their large graph sizes.

Compared Methods

We compare our method with graph embedding methods, graph neural networks, and belief propagation methods.

- **Graph embedding methods:** These methods first learn an embedding vector for each node; then they train a standard classifier (we consider multi-class logistic regression) using the training dataset and the embedding vectors; finally, they use the classifier to predict labels for the testing nodes. We consider two representative graph embedding methods, i.e., DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) and Planetoid (Yang, Cohen, and Salakhutdinov 2016). We use 128-dimension embedding vectors.
- **Graph neural network methods:** These methods learn the node embeddings and classify nodes simultaneously. In particular, we compare with Chebyshev (Kipf and Welling 2017), graph convolutional network (GCN) (Kipf and Welling 2017), GraphSAGE (Hamilton, Ying, and Leskovec 2017), and graph attention network (GAT) (Veličković et al. 2018).
- **Belief propagation methods:** We use the standard BP (Pearl 1988) or the linearized BP (LinBP) (Gatterbauer et al. 2015) to perform inference in the pMRF. In particular, BP and LinBP infer the approximate marginal probability distributions \mathbf{P} and use them to predict labels of testing nodes. We obtained the implementation of BP and LinBP from (Gatterbauer et al. 2015).
- **Our method LCM and its variants:** In order to show the effectiveness of our consistency regularization term, we consider several variants of LCM. Specifically, LCM-wo is the variant of our method that does not use the consistency regularization term. Moreover, we compare our consistency regularization with the standard L_1 and L_2 regularizations, which are denoted as $R(\mathbf{W}, \mathbf{H}) = \sum_{(u,v) \in E} |W_{uv}|_1 + \sum_{i,j} |H_{ij}|_1$ and $R(\mathbf{W}, \mathbf{H}) = \sum_{(u,v) \in E} W_{uv}^2 + \sum_{i,j} H_{ij}^2$, respectively. Accordingly, we denote the two variants as LCM-L1 and LCM-L2, respectively.

We observe that BP and LinBP outperform other traditional methods such as label propagation (LP) (Zhu, Ghahramani, and Lafferty 2003), manifold regularization (ManiReg) (Belkin, Niyogi, and Sindhwani 2006), deep semi-supervised embedding (SemiEmb) (Weston et al. 2012), and iterative classification algorithm (ICA) (Lu and Getoor 2003). For instance, using the same experimental setting as Veličković et al. (Veličković et al. 2018),

²Note that the results of the baseline methods reported in previous work only use one training and validation dataset without reporting the averages among multiple trials. Therefore, their results are different from what we show in Table 1.

Methods		Cora	Citeseer	Pubmed	NELL	Google+	Twitter
Graph Embedding	DeepWalk	0.669±0.017	0.428±0.019	0.651±0.015	0.352±0.013	–	–
	Planetoid	0.746±0.011	0.645±0.029	0.699±0.035	0.435±0.015	–	–
Graph Neural Network	Chebyshev	0.786± 0.018	0.694± 0.023	0.701± 0.024	–	–	–
	GCN	0.823± 0.009	0.714± 0.016	0.756± 0.032	0.596±0.026	–	–
	GraphSAGE	0.811± 0.009	0.664± 0.010	0.757± 0.014	0.574±0.015	–	–
	GAT	0.824± 0.021	0.708± 0.016	0.755± 0.030	0.578±0.010	–	–
Belief Propagation	BP	0.804± 0.023	0.704± 0.011	0.734± 0.024	0.587±0.015	–	–
	LinBP	0.809± 0.014	0.707± 0.004	0.744± 0.028	0.588±0.013	0.741±0.013	0.687±0.002
Our Methods	LCM-wo	0.813± 0.007	0.716± 0.006	0.754± 0.018	0.626±0.013	0.757±0.012	0.712±0.002
	LCM-L1	0.808± 0.009	0.711± 0.005	0.752± 0.013	0.609±0.008	0.749±0.018	0.699±0.004
	LCM-L2	0.811± 0.013	0.713± 0.007	0.753± 0.016	0.616±0.002	0.751±0.016	0.709±0.004
	LCM	0.833± 0.007	0.722± 0.005	0.770± 0.019	0.647±0.005	0.779±0.012	0.748±0.002

Table 1: Average accuracies and standard deviations of compared methods on the six graphs. “–” means the methods cannot be executed on our machine as they require more memory than we have. We note that all the compared graph embedding methods and graph neural networks cannot be executed on our machine for the two large-scale social graphs.

Dataset	Cora		Citeseer		Pubmed		NELL		Google+		Twitter	
	homo	heter	homo	heter	homo	heter	homo	heter	homo	heter	homo	heter
Ave. initialized weights	0.225	0.197	0.315	0.347	0.138	0.147	0.039	0.035	0.107	0.091	0.022	0.017
Ave. learnt weights	0.354	0.261	0.493	0.448	0.652	0.548	0.105	0.089	0.569	0.242	0.595	0.329

Table 2: Average initialized weights and weights learnt by our method for homogeneous edges vs. heterogeneous edges on the six graphs. “homo” and “heter” mean homogeneous edges and heterogeneous edges, respectively. Our method learns larger average weights for homogeneous edges than heterogeneous edges.

LinBP achieves classification accuracies 0.785, 0.709, and 0.787 on the three citation graphs, respectively. According to (Veličković et al. 2018), ICA consistently outperforms LP, ManiReg, and SemiEmb on the three citation graphs. However, ICA only achieves accuracies 0.751, 0.691, and 0.739 on the three citation graphs, respectively. Thus, we omit the results of these traditional methods for simplicity. The parameter settings of these methods are detailed in (Wang, Jia, and Gong 2020).

Results

Table 1 shows the classification accuracies averaged over 5 trials and the standard deviations of all compared methods on the six graphs. We note that the compared graph embedding methods and graph neural networks cannot be executed on our machine for the two large-scale social graphs due to insufficient memory. We have the following observations.

LCM outperforms the compared methods: First, we observe that graph neural networks outperform standard BP and LinBP methods. For example, GCN consistently achieves higher average accuracies than BP and LinBP on the three citation graphs and the NELL knowledge graph.

Second, LCM is consistently more accurate than BP and LinBP. For example, LCM achieves 0.018 to 0.060 higher accuracies than BP on the three citation graphs and the NELL knowledge graph; and achieves 0.015 to 0.059 higher accuracies than LinBP on the six graphs. The reason is that BP and LinBP do not learn the edge weights and coupling matrix (i.e., they use the initialized edge weights and coupling matrix), while LCM does. To further illustrate the difference between LCM and BP/LinBP, Table 2 shows the average initialized weights and average weights learnt by LCM

for homogeneous edges and heterogeneous edges on the six graphs. An edge is homogeneous if the two corresponding nodes have the same label, otherwise the edge is heterogeneous. Moreover, Figure 1 shows the learnt coupling matrices for the six graphs. Our learnt coupling matrices are diagonal dominant, which means that two linked nodes are more likely to have the same label. Moreover, homogeneous edges have larger weights than heterogeneous edges on average after learning the edge weights. In particular, LCM increases edge weights for both homogeneous edges and heterogeneous edges, but LCM increases the weights of homogeneous edges more substantially.

Third, after learning the weight matrix and the coupling matrix, our method outperforms the graph neural networks and graph embedding methods. Specifically, LCM achieves 0.009 to 0.051 higher average accuracies than the graph neural networks, and 0.071 to 0.212 higher average accuracies than the graph embedding methods on the three citation graphs and the NELL knowledge graph.

Our consistency regularization is effective: Our results also show that our consistency regularization term is effective. In particular, LCM outperforms LCM-wo, LCM-L1, and LCM-L2. Our method achieves 0.006 to 0.036 higher accuracies when using the consistency regularization term (i.e., LCM vs. LCM-wo). One possible reason is that L1/L2 regularization aims to prevent overfitting of complex models. However, pMRF is a simple model, which may not have overfitting on the datasets. Therefore, L1/L2 regularization could make LCM worse. Note that our consistency regularization does not aim to prevent traditional overfitting. Instead, it aims to improve the pMRF model itself and thus can outperform LCM-wo.

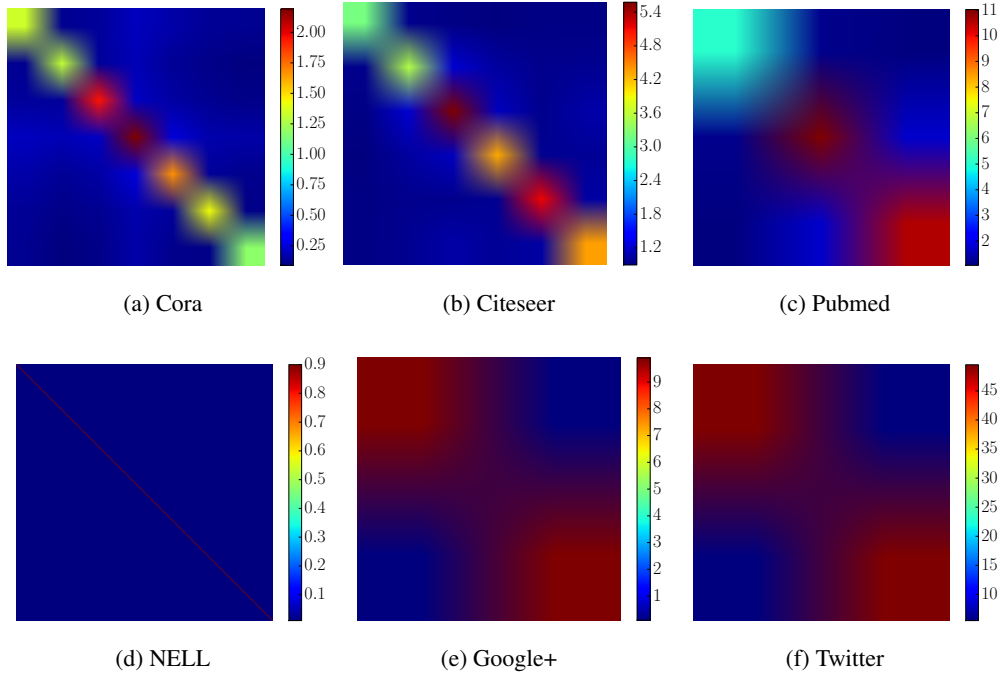


Figure 1: Learnt coupling matrices \mathbf{H} on the six graphs. The coupling matrices are diagonal dominant. For the NELL graph, it is better to zoom in the subfigure to observe the diagonal dominant property of the learnt coupling matrix.




Figure 2: Impact of the number of training nodes per class.

Impact of the number of training nodes per class: Figure 2 further shows the average accuracies of different methods on the six graph datasets when we sample a particular

number of training nodes from each class. Our method consistently outperforms the compared methods when the training dataset size is small. For instance, when each class has

Methods		Time (seconds)
Graph Embedding	DeepWalk	1,312
	Planetoid	1,033
Graph Neural Network	GCN	304
	GraphSAGE	196
	GAT	2,115
Belief Propagation	BP	443
	LinBP	8
Our Method	LCM	50

Table 3: Time of the compared methods on NELL.

only 1 training node, LCM has an around 0.08 higher average accuracy than GCN on NELL. When each class has 10 training nodes, LCM has an around 0.04 higher average accuracy than LinBP on Twitter. Note that in semi-supervised node classification, we often assume the number of labeled nodes is small, compared to the total number of nodes.

LCM is more efficient than graph embedding and graph neural network methods: For simplicity, we only show efficiency on NELL, which is the largest dataset that the compared methods can be executed on our machine. Table 3 shows the running time of the compared methods in one of our experiments. We have two observations. First, LCM is more efficient than graph embedding and graph neural network methods. Second, LCM is less efficient than LinBP. This is because LinBP does not learn the edge weight matrix and coupling matrix. However, both LCM and LinBP are much more efficient than BP. This is because BP needs to maintain messages on each edge.

Related Work

Conventional methods: Semi-supervised node classification in graph-structured data has been studied extensively and many methods have been proposed. For instance, conventional methods include label propagation (LP) (Zhu, Ghahramani, and Lafferty 2003), manifold regularization (ManiReg) (Belkin, Niyogi, and Sindhwani 2006), deep semi-supervised embedding (SemiEmb) (Weston et al. 2012), iterative classification algorithm (ICA) (Lu and Getoor 2003), and pairwise Markov Random Fields (pMRF) (Gatterbauer 2017). In particular, pMRF associates a discrete random variable with each node to model its label and defines a joint probability distribution for the random variables associated with all nodes, where the statistical correlations/independence between the random variables are captured by the graph structure. Then, either the standard Belief Propagation (BP) (Pearl 1988) or linearized BP (LinBP) (Gatterbauer et al. 2015; Gatterbauer 2017) is used for inference. Existing studies on pMRF-based semi-supervised node classification share a common limitation: they set a constant edge potential for all edges in the graph.

Graph embedding methods and graph neural networks: A recent trend is to extend neural networks to graph-structured data for semi-supervised node classification. Methods along this direction can be roughly grouped into two categories, i.e., *graph embedding* (Perozzi, Al-

Rfou, and Skiena 2014; Tang et al. 2015; Cao, Lu, and Xu 2015; Yang, Cohen, and Salakhutdinov 2016; Grover and Leskovec 2016; Ribeiro, Saverese, and Figueiredo 2017; Cui et al. 2020) and *graph neural networks* (Duvenaud et al. 2015; Atwood and Towsley 2016; Kipf and Welling 2017; Hamilton, Ying, and Leskovec 2017; Veličković et al. 2018; Xu et al. 2019; Battaglia et al. 2018; Gao, Wang, and Ji 2018; Xu et al. 2018; Ying et al. 2018; Qu, Bengio, and Tang 2019; Ma et al. 2019a; Zhang et al. 2019; Ma et al. 2019b; Wu, He, and Xu 2019; Chiang et al. 2019; Wang et al. 2020). Graph embedding methods first learn node embeddings and then learn a standard classifier (e.g., logistic regression) using the embeddings to classify nodes, where learning the embeddings and learning the classifier are performed separately. Different graph embedding methods leverage different techniques to learn node embeddings. For instance, DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) learns nodes’ embeddings via generalizing the word to vector technique (Mikolov et al. 2013) developed for natural language processing to graph data. Specifically, DeepWalk treats a node as a word in natural language, generates node sequences using truncated random walks on a graph, and leverages the skip-gram model (Mikolov et al. 2013) to learn an embedding vector for each node.

Graph neural networks learn the node embeddings and the classifier to classify nodes simultaneously. In particular, the hidden layers represent node embeddings and the last layer models a classifier to classify nodes. Different graph neural networks use different neural network architectures. For instance, Graph Convolutional Network (GCN) (Kipf and Welling 2017) uses an architecture that is motivated by spectral graph convolutions (Duvenaud et al. 2015). Specifically, the input layer is the nodes’ features. A hidden layer models nodes’ embeddings. In particular, the neural network iteratively computes a node’s embedding vector in a hidden layer via aggregating the embedding vectors of the node’s neighbors in the previous layer. State-of-the-art graph neural networks are more accurate than conventional methods such as LP and pMRF on multiple benchmark datasets. However, graph neural networks are less efficient than conventional methods such as LP and LinBP. Our method addresses the limitation of pMRF-based methods and is more accurate and efficient than graph neural networks.

Conclusion and Future Work

We propose a novel method to learn the coupling matrix in pairwise Markov Random Fields for semi-supervised node classification in graph-structured data. We formulate learning coupling matrix as an optimization problem, whose objective function is the sum of the training loss and a consistency regularization term that we propose. Moreover, we propose an iterative algorithm to solve the optimization problem. Our evaluation results on six benchmark datasets show that our optimized pMRF-based method is more accurate and efficient than state-of-the-art graph neural networks. An interesting future work is to explore the connections between our optimized pMRF-based method and graph neural networks as well as unify them in a general framework.

Acknowledgments

We would like to thank the anonymous reviewers for their insightful reviews. This work was supported by the National Science Foundation under grants No. 1937787 and 1937786. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

References

- Akoglu, L.; Chandy, R.; and Faloutsos, C. 2013. Opinion Fraud Detection in Online Reviews by Network Effects. In *ICWSM*.
- Atwood, J.; and Towsley, D. 2016. Diffusion-convolutional neural networks. In *NIPS*, 1993–2001.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Belkin, M.; Niyogi, P.; and Sindhvani, V. 2006. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *JMLR* 7: 2399–2434.
- Cao, Q.; Sirivianos, M.; Yang, X.; and Pogueiro, T. 2012. Aiding the Detection of Fake Accounts in Large Scale Social Online Services. In *NSDI*, 197–210.
- Cao, S.; Lu, W.; and Xu, Q. 2015. Grarep: Learning graph representations with global structural information. In *CIKM*, 891–900.
- Carlson, A.; Betteridge, J.; Kisiel, B.; Settles, B.; Hruschka, E. R.; and Mitchell, T. M. 2010. Toward an architecture for never-ending language learning. In *AAAI*.
- Chau, D. H. P.; Nachenberg, C.; Wilhelm, J.; Wright, A.; and Faloutsos, C. 2011. Polonium: Tera-scale graph mining and inference for malware detection. In *SIAM SDM*, 131–142.
- Chiang, W.-L.; Liu, X.; Si, S.; Li, Y.; Bengio, S.; and Hsieh, C.-J. 2019. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *KDD*, 257–266.
- Cui, G.; Zhou, J.; Yang, C.; and Liu, Z. 2020. Adaptive Graph Encoder for Attributed Graph Embedding. In *KDD*, 976–985.
- Duvenaud, D. K.; Maclaurin, D.; Iparraguirre, J.; Bombarell, R.; Hirzel, T.; Aspuru-Guzik, A.; and Adams, R. P. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *NIPS*, 2224–2232.
- Gao, H.; Wang, Z.; and Ji, S. 2018. Large-scale learnable graph convolutional networks. In *KDD*, 1416–1424.
- Gao, P.; Wang, B.; Gong, N. Z.; Kulkarni, S. R.; Thomas, K.; and Mittal, P. 2018. Sybilfuse: Combining local attributes with global structure to perform robust sybil detection. In *IEEE CNS*, 1–9.
- Gatterbauer, W. 2017. The linearization of belief propagation on pairwise markov random fields. In *AAAI*, 3747–3753.
- Gatterbauer, W.; Günnemann, S.; Koutra, D.; and Faloutsos, C. 2015. Linearized and single-pass belief propagation. *VLDB* 581–592.
- Gong, N. Z.; and Liu, B. 2018. Attribute Inference Attacks in Online Social Networks. *ACM TOPS* 21(1): 3:1–3:30.
- Gong, N. Z.; Talwalkar, A.; Mackey, L.; Huang, L.; Shin, E. C. R.; Stefanov, E.; Shi, E. R.; and Song, D. 2014. Joint link prediction and attribute inference using a social-attribute network. *ACM TIST* 5(2): 1–20.
- Gong, N. Z.; Xu, W.; Huang, L.; Mittal, P.; Stefanov, E.; Sekar, V.; and Song, D. 2012. Evolution of Social-Attribute Networks: Measurements, Modeling, and Implications using Google+. In *IMC*, 131–144.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *KDD*, 855–864.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *NIPS*, 1024–1034.
- He, D.; You, X.; Feng, Z.; Jin, D.; Yang, X.; and Zhang, W. 2018. A network-specific Markov random field approach to community detection. In *AAAI*, 306–313.
- Hooi, B.; Song, H. A.; Beutel, A.; Shah, N.; Shin, K.; and Faloutsos, C. 2016. FRAUDAR: Bounding Graph Fraud in the Face of Camouflage. In *KDD*, 895–904.
- Jia, J.; Wang, B.; Zhang, L.; and Gong, N. Z. 2017. AttrInfer: Inferring User Attributes in Online Social Networks Using Markov Random Fields. In *WWW*, 1561–1569.
- Jin, D.; Liu, Z.; Li, W.; He, D.; and Zhang, W. 2019a. Graph convolutional networks meet markov random fields: Semi-supervised community detection in attribute networks. In *AAAI*, 152–159.
- Jin, D.; You, X.; Li, W.; He, D.; Cui, P.; Fogelman-Soulié, F.; and Chakraborty, T. 2019b. Incorporating network embedding into markov random field for better community detection. In *AAAI*, 160–167.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. *ICLR*.
- Lu, Q.; and Getoor, L. 2003. Link-based classification. In *ICML*, 496–503.
- Ma, J.; Tang, W.; Zhu, J.; and Mei, Q. 2019a. A Flexible Generative Framework for Graph-based Semi-supervised Learning. In *NeurIPS*, 3276–3285.
- Ma, Y.; Wang, S.; Aggarwal, C. C.; and Tang, J. 2019b. Graph convolutional networks with eigenpooling. In *KDD*, 723–731.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, 3111–3119.
- Pandit, S.; Chau, D. H.; Wang, S.; and Faloutsos, C. 2007. Netprobe: a fast and scalable system for fraud detection in online auction networks. In *WWW*, 201–210.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems - networks of plausible inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann.

- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *KDD*, 701–710.
- Qu, M.; Bengio, Y.; and Tang, J. 2019. GMNN: Graph Markov Neural Networks. In *ICML*, 5241–5250.
- Rayana, S.; and Akoglu, L. 2015. Collective Opinion Spam Detection: Bridging Review Networks and Metadata. In *KDD*, 985–994.
- Ribeiro, L. F.; Saverese, P. H.; and Figueiredo, D. R. 2017. struc2vec: Learning node representations from structural identity. In *KDD*, 385–394.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI magazine* 29(3): 93–93.
- Tamersoy, A.; Roundy, K.; and Chau, D. H. 2014. Guilt by association: large scale malware detection by mining file-relation graphs. In *KDD*, 1524–1533.
- Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *WWW*, 1067–1077.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2018. Graph attention networks. In *ICLR*.
- Wang, B.; Gong, N. Z.; and Fu, H. 2017. GANG: Detecting Fraudulent Users in Online Social Networks via Guilt-by-Association on Directed Graphs. In *IEEE ICDM*, 465–474.
- Wang, B.; Jia, J.; and Gong, N. Z. 2019. Graph-based Security and Privacy Analytics via Collective Classification with Joint Weight Learning and Propagation. In *NDSS*.
- Wang, B.; Jia, J.; and Gong, N. Z. 2020. Semi-Supervised Node Classification on Graphs: Markov Random Fields vs. Graph Neural Networks. *arXiv preprint arXiv:2012.13085*.
- Wang, B.; Jia, J.; Zhang, L.; and Gong, N. Z. 2018. Structure-based sybil detection in social networks via local rule-based propagation. *IEEE Transactions on Network Science and Engineering* 6(3): 523–537.
- Wang, B.; Zhang, L.; and Gong, N. Z. 2017. SybilSCAR: Sybil Detection in Online Social Networks via Local Rule based Propagation. In *IEEE INFOCOM*, 1–9.
- Wang, G.; Xie, S.; Liu, B.; and Philip, S. Y. 2011. Review graph based online store review spammer detection. In *IEEE ICDM*, 1242–1247.
- Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; Liu, J.; and Hooi, B. 2020. NodeAug: Semi-Supervised Node Classification with Data Augmentation. In *KDD*, 207–217.
- Weston, J.; Ratle, F.; Mobahi, H.; and Collobert, R. 2012. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*, 639–655. Springer.
- Wu, F.; Zhang, T.; Souza Jr, A. H. d.; Fifty, C.; Yu, T.; and Weinberger, K. Q. 2019. Simplifying graph convolutional networks. In *ICML*, 6861–6871.
- Wu, J.; He, J.; and Xu, J. 2019. Demo-Net: Degree-specific graph neural networks for node and graph classification. In *KDD*, 406–415.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How powerful are graph neural networks? In *ICLR*.
- Xu, K.; Li, C.; Tian, Y.; Sonobe, T.; Kawarabayashi, K.-i.; and Jegelka, S. 2018. Representation learning on graphs with jumping knowledge networks. In *ICML*, 5449–5458.
- Yang, Z.; Cohen, W. W.; and Salakhutdinov, R. 2016. Revisiting semi-supervised learning with graph embeddings. In *ICML*, 40–48.
- Ye, Y.; Li, T.; Zhu, S.; Zhuang, W.; Tas, E.; Gupta, U.; and Abdulhayoglu, M. 2011. Combining file content and file relations for cloud based malware detection. In *KDD*, 222–230.
- Ying, Z.; You, J.; Morris, C.; Ren, X.; Hamilton, W.; and Leskovec, J. 2018. Hierarchical graph representation learning with differentiable pooling. In *NeurIPS*, 4805–4815.
- Zhang, C.; Song, D.; Huang, C.; Swami, A.; and Chawla, N. V. 2019. Heterogeneous graph neural network. In *KDD*, 793–803.
- Zhu, X.; Ghahramani, Z.; and Lafferty, J. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, 912–919.