Meta-GNN: On Few-shot Node Classification in Graph Meta-learning

Fan Zhou

fan.zhou@uestc.edu.cn

School of Information and Software Engineering, University of Electronic Science and Technology of China

Chengtai Cao

cct447131988@gmail.com

School of Information and Software Engineering, University of Electronic Science and Technology of China

Kunpeng Zhang

kzhang@rhsmith.umd.edu

University of Maryland

Goce Trajcevski

gocet25@iastate.edu

Iowa State University

Ting Zhong

zhongting@uestc.edu.cn

School of Information and Software Engineering, University of Electronic Science and Technology of China

Ji Geng

jgeng@uestc.edu.cn

School of Information and Software Engineering, University of Electronic Science and Technology of China

ABSTRACT

Meta-learning has received a tremendous recent attention as a possible approach for mimicking human intelligence, i.e., acquiring new knowledge and skills with little or even no demonstration. Most of the existing meta-learning methods are proposed to tackle few-shot learning problems such as image and text, in rather Euclidean domain. However, there are very few works applying meta-learning to non-Euclidean domains, and the recently proposed graph neural networks (GNNs) models do not perform effectively on graph few-shot learning problems. Towards this, we propose a novel graph meta-learning framework – Meta-GNN – to tackle the few-shot node classification problem in graph meta-learning settings. It obtains the prior knowledge of classifiers by training on many similar few-shot learning tasks and then classifies the nodes from new classes with only few labeled samples. Additionally, Meta-GNN is a general model that can be straightforwardly incorporated into any existing state-of-the-art GNN. Our experiments conducted on three benchmark datasets demonstrate that our proposed approach not only improves the node classification performance by a large margin on few-shot learning problems in meta-learning paradigm, but also learns a more general and flexible model for task adaption.

CCS CONCEPTS

• **Computing methodologies** → *Transfer learning; Semi-supervised learning settings.*

KEYWORDS

meta-learning; graph neural networks; node classification; few-shot learning

ACM Reference Format:

Fan Zhou, Chengtai Cao, Kunpeng Zhang, Goce Trajcevski, Ting Zhong, and Ji Geng. 2019. Meta-GNN: On Few-shot Node Classification in Graph Meta-learning. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*, November 3–7, 2019, Beijing, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3357384.3358106>

1 INTRODUCTION

Analyzing graph structure data with various deep learning methods has spurred a significant recent research interest [12]. A large number of models have been developed for solving different problems. While earlier methods mainly focused on embedding nodes in an unsupervised manner in order to capture the global/local network structure, recent efforts applied sophisticated deep learning models – which have been successfully used on Euclidean domain – to non-Euclidean graph structure data, resulting in many prominent Graph Neural Networks (GNNs)-based methods – e.g., GCN [4], GraphSAGE [3], SGC [11], etc.

Despite recent breakthroughs in GNNs, one setting that presents a challenge is that of *few-shot learning*, which aims at addressing data deficiency problem by recognizing new classes from few labeled samples. Existing GNNs models always need to re-learn their parameters to incorporate the new information when new classes are encountered, and if the number of nodes in each new class is small, models' performance will suffer a catastrophic decline. The inability to handle situations where one or very few samples exist, is one of the major challenges for the current GNNs. To address this issue, some recent studies have focused on solving the few-shot learning on graph data where GNNs are either trained via co-training and self-training [5], or extended by stacking transposed graph convolutional layers for reconstructing input features [13].

Meta-learning (i.e., *learning to learn*), has drawn a significant attention in the AI community in the recent years, due to its capability of quickly adapting to new tasks and learning transferable knowledge between tasks, with a few samples. It has been widely considered as similar to the human intelligence where humans are capable to rapidly learn new classes, with few samples and by utilizing previously learned prior knowledge. Meta-learning systems are trained by being exposed to a large number of tasks and are tested on their ability to learn new tasks. This differs from many standard

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '19, November 3–7, 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6976-3/19/11...\$15.00

<https://doi.org/10.1145/3357384.3358106>

machine learning techniques, which involve training on a single task and testing on held-out samples from that task [5, 13]. We refer to the processes of learning on a large number of tasks and testing on a new task as *meta-training* and *meta-testing*, respectively.

Adopting meta-learning for image and text learning has yielded significant progress and several models and algorithms have been recently proposed – e.g., Matching Networks [10], Prototypical Networks [9], MAML [1], etc. Despite the extensive studies and the promising results on analyzing data with Euclidean-like properties (e.g., images and text), there is surprisingly little work on applying meta-learning to graphs (non-Euclidean data). Among the main reasons are that graph data is more irregular, noisy and exhibits more complex relations among nodes, which make it difficult to directly apply existing meta-learning methods.

To bridge this gap, we present a general graph meta-learning framework called Meta-GNN, focusing on the few-shot node classification problem on graph data. To our knowledge, our work is the first to incorporate the meta-learning paradigm into GNNs, providing the capability of well adapting or generalizing to new classes that have never been encountered before, with very few samples. In summary, the main contributions of our work are as follows:

- We formulate a new graph few-shot learning paradigm for node classification. It is different from previous work in that we aim to classify nodes from new classes with only few samples each.
- We propose a generic framework for tackling few-shot node classification, which can be easily combined with any popular GNNs models and opens up a new perspective of graph structure data analysis.
- We demonstrate the superiority of our method over several state-of-the-art GNNs models on three benchmark datasets.

2 METHODOLOGY

An undirected graph, denoted as $\mathcal{G} = (V, E, \mathbf{A}, \mathbf{X})$, is a quadruple where: (1) $V = \{v_1, v_2, \dots, v_i, \dots, v_n\}$ is a node set; (2) $E = \{e_{i,j} = (v_i, v_j)\} \subseteq (V \times V)$ is an edge set; (3) $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric (typically sparse) adjacency matrix where a_{ij} denotes the weight of the edge between nodes v_i and v_j ; and (4) $\mathbf{X} \in \mathbb{R}^{n \times d}$ is a feature matrix with $\mathbf{x}_i \in \mathbb{R}^d$ representing the characteristics of a given node v_i .

Problem Definition: We consider the few-shot node classification problem within the meta-learning paradigm. Simply speaking, we aim to obtain a classifier that can be adapted to new classes that are *unseen* during the training process, given only a few samples in each new class. Formally, each node v_i in the training set $\mathcal{D}_{train}$ belongs to one of the classes in C_1 . Nodes in a given disjoint testing set $\mathcal{D}_{test}$ are associated with totally different new classes C_2 . In $\mathcal{D}_{test}$, a small number of nodes have labels/classes available. Our goal is to find a function f_θ that is able to classify the rest of the unlabeled nodes into one of the classes in C_2 , with a low misclassification rate. If the number of labeled nodes in each class is K , the task is known as $|C_2|$ -way K -shot learning problem, where K is a very small number. Thus, we address the graph meta-learning problem for node classification, in order to extract transferable knowledge from training data that will allow us to perform well on the testing data with unseen classes.

2.1 Graph Neural Networks

Modern GNNs jointly utilize the graph structure information and the node feature $\mathbf{X}$ to learn a new representation vector $\mathbf{h}_v$ of a node, usually following a neighborhood aggregation scheme. After l iterations of the aggregation, a node's representation captures the structural information from its l -hop network neighbors. An illustrative example is shown in the right part of Figure 1, whereby as the graph passes through the first layer of GNN, the red node aggregates the information from node 1, 2 and 3, and after the second layer, it gathers the information from node 5 and node 6. Formally, the l -th layer of a GNN is defined as:

$$\mathbf{a}_v^{(l)} = \mathbf{h}_v^{(l-1)} \cdot \text{AGGREGATE}^{(l)} \left(\left\{ \mathbf{h}_u^{(l-1)} : u \in \mathcal{N}(v) \right\} \right), \quad (1)$$

$$\mathbf{h}_v^{(l)} = \mathbf{h}_v^{(l-1)} \cdot \text{COMBINE}^{(l)} \left(\mathbf{a}_v^{(l)} \right), \quad (2)$$

where $\mathbf{h}_v^{(l)}$ is the feature vector of node v at the l -iteration/layer. We initialize $\mathbf{h}^{(0)} = \mathbf{X}$ and denote with $\mathcal{N}(v)$ the set of nodes adjacent to v . The final representation $\mathbf{h}_v^{(L)}$ will be used for downstream tasks such as link prediction and node classification. The choices of AGGREGATE and COMBINE operation are crucial to the task performance, and plenty of approaches have been proposed and achieved impressive results (see [12] for a recent comprehensive overview). However, as observed in the literature [7], deep neural networks generally perform poorly on few-shot learning tasks. A main cause is that the standard gradient descent method (including its many variants) requires a large number of examples to obtain a satisfactory performance, which is not upheld in the few-shot learning setting.

2.2 Meta-GNN

We now present the details of our Meta-GNN framework. Our approach is based on the methodology of meta-learning which often follows an episodic paradigm and solves a new few-shot learning problem (i.e., meta-testing task, $\mathcal{T}_{mt}$) by training on many sampled similar tasks (i.e., meta-training tasks). We refer to the corresponding training and testing sets in all tasks as *support set* and *query set*, respectively. Our method leverages MAML [1] for the gradient updates during training. After training on considerable meta-training tasks, Meta-GNN is expected to learn (as a prior knowledge) how to quickly adapt to a new task using only a few datapoints in the new task. The performance of Meta-GNN is measured by meta-testing on the new task, i.e., fine-tuning Meta-GNN on a few samples from the support set of $\mathcal{T}_{mt}$ and evaluating it on the query set of $\mathcal{T}_{mt}$.

We denote our Meta-GNN model as f_θ with parameters θ , and the training set as $\mathcal{D}_{train} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_i, y_i), \dots, (\mathbf{x}_N, y_N)\}$ where $y_i \in C_1$ and N is the number of nodes in the training set. We will fork M meta-training tasks from $\mathcal{D}_{train}$: $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_M\}$. Both the support set $\mathcal{S}_i$ and query set $\mathcal{Q}_i$ of meta-training task $\mathcal{T}_i$ are sampled from $\mathcal{D}_{train}$. In the support set, we have $\mathcal{S}_i = \{v_{i1}, v_{i2}, \dots, v_{is}\} = \{(\mathbf{x}_{i1}, y_{i1}), (\mathbf{x}_{i2}, y_{i2}), \dots, (\mathbf{x}_{is}, y_{is})\}$, where $s = |\mathcal{S}_i|$; $\mathbf{x}_{is}$ is the input vector of node v_{is} with label y_{is} .

Task Sampling: Before presenting the details of Meta-GNN, we prepare M tasks $\mathcal{T}$ by sampling from $\mathcal{D}_{train}$ according to the graph meta-learning principals. In short, we sample $|C_2|$ classes from C_1

Figure 1: Overview of Meta-GNN. The black θ , red θ and blue θ represent the randomly initialized parameters, the parameters after one meta-update and after all meta-updates, respectively. M is the number of meta-training tasks.

and then randomly sample K nodes for each class to simulate few-shot node classification. Specifically, we use the following steps to generate the graph meta-training tasks:

- (1) $C \leftarrow \text{RANDOMSAMPLE}(C_1, |C_2|)$;
- (2) $S_i \leftarrow \text{RANDOMSAMPLE}(\mathcal{D}_C, K \times |C_2|)$;
- (3) $Q_i \leftarrow \text{RANDOMSAMPLE}(\mathcal{D}_C - S_i, P)$;
- (4) $\mathcal{T}_i = S_i + Q_i$;
- (5) Repeat step (1) - (4) for M times;

Thus, we first randomly sample $|C_2|$ classes from C_1 , denoted as C . Then we can obtain $\mathcal{D}_C$, which is a subset of training set $\mathcal{D}_{train}$ containing elements (x_i, y_i) , where y_i is one of the classes in C ; Next, we randomly sample $K \times |C_2|$ nodes from $\mathcal{D}_C$ to form the support set S_i , where K is the number of nodes in each class of S_i (i.e., the number of shots). Finally, we randomly sample P nodes from the remaining nodes in $\mathcal{D}_C$ to compose the query set Q_i which, in turn, is used in constructing the meta-training task $\mathcal{T}_i = S_i + Q_i$. Repeating the above steps M times, yields M meta-training tasks. **Meta-training:** We expect to obtain a good initialization of Meta-GNN, which is generally applicable to similar tasks, and explicitly encourage the initialization parameters to perform well after a small number of gradient descent updates on a new few-shot learning task. When learning a task $\mathcal{T}_i$, we begin with feeding the support set S_i to Meta-GNN, and calculate the cross-entropy loss:

$$\mathcal{L}_{\mathcal{T}_i}(f_{\theta}) = - \left(\sum_{x_{is}, y_{is}} y_{is} \log f_{\theta}(x_{is}) + (1 - y_{is}) \log (1 - f_{\theta}(x_{is})) \right). \quad (3)$$

Then we perform parameters updates, using a simple gradient descent with one or several steps in task $\mathcal{T}_i$. For brevity, we will only describe one gradient update in the rest of this section – with a note that performing multiple gradient updates is a straightforward extension:

$$\theta'_i = \theta - \alpha_1 \frac{\partial \mathcal{L}_{\mathcal{T}_i}(f_{\theta})}{\partial \theta}. \quad (4)$$

where α_1 is the task-learning rate and the model parameters are trained to optimize the performance of $f_{\theta'_i}$ across meta-training tasks. More specifically, the meta-objective is as follows:

$$\theta = \arg \min_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta'_i}). \quad (5)$$

Table 1: Descriptive statistics and partition of three datasets.

	Cora	Citeseer	Reddit
# Nodes	2,708	3,327	232,965
# Features	1,433	3,703	602
# Labels	7	6	41

where $p(\mathcal{T})$ is the distribution of meta-training tasks. Note that the meta-optimization is performed over the model parameters θ , whereas the objective is computed using the updated model parameters θ' . This is because we need good initialization parameters θ for all the similar few-shot node classification tasks, instead of certain updated parameters θ'_i performing well on only a particular task $\mathcal{T}_i$. Essentially, Meta-GNN aims to optimize the model parameters so that it maximizes the node classification performance on a new task after one or a small number of gradient descent updates. The meta-optimization across tasks is performed via stochastic gradient descent (SGD), and the model parameters θ are updated as follows:

$$\theta \leftarrow \theta - \alpha_2 \frac{\partial \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta'_i})}{\partial \theta}. \quad (6)$$

where α_2 is the meta-learning rate – i.e., an additional hyperparameter introduced by our framework.

Meta-testing: For meta-testing, we simply need to feed the nodes of the support set of the new few-shot learning task (i.e., $\mathcal{T}_{mt}$) to the Meta-GNN, and update parameters θ'_{mt} via one or a small number of gradient descent steps using Eq.(4). Therefore, the performance of Meta-GNN can be easily evaluated on the query set of $\mathcal{T}_{mt}$.

The overall operational process of Meta-GNN is depicted in Figure 1.

3 EXPERIMENTS

We now present the empirical results of Meta-GNN performance for node classification on three commonly used datasets: Cora [8], Citeseer [8] and Reddit [3]. For reproducibility, the source code and datasets are publicly available¹. Since our work focuses on few-shot learning problem in the context of meta-learning paradigm, we make some modifications in the dataset partition, to ensure the feasibility of the few-shot learning (cf. Section 2.2).

For the Cora and Citeseer datasets, we randomly set apart two classes as meta-testing classes (i.e., the corresponding nodes forming $\mathcal{D}_{test}$ with $|C_2| = 2$), and the rest nodes (i.e., $\mathcal{D}_{train}$) are used to generate meta-training tasks following the method described in section 2.2. Similarly for the Reddit, except that $|C_2| = 5$ due to the relatively large number of unique labels. Each class in the support set of meta-training and meta-testing tasks has only one or three samples (i.e., $K = 1$ or $K = 3$) for all datasets. We observed that, when the support set is extremely small, the performance of each method is sensitive to the node selection. Therefore, we evaluate all models on 50 randomly selected nodes in the support set and report the average accuracy on Cora and Citeseer. For Reddit, as there are significant more classes and corresponding nodes in each class, we select the same nodes as support set for each run and conduct 50 different runs, and report the average Micro-F1 score. In order

¹<https://github.com/AI-DL-Conference/Meta-GNN>

to fairly and accurately assess performance of our Meta-GNN and other baselines, we conduct 10 cross-validation on Cora and Reddit and 5 on Citeseer. The statistics of datasets after pre-processing are shown in Table 1.

Baselines: For Cora and Citeseer, we compare Meta-GNN against DeepWalk [6], Node2Vec [2], GCN [4], SGC [11] and GraphSAGE [3]. For Reddit, we compare with GCN [4], SGC [11] and GraphSAGE [3]. Notably, we only modify the dataset partition to satisfy the few-shot learning setting in meta-learning paradigm, and other settings of each model are the same as its original implementation.

Implementation and Parameter Setups: To demonstrate the applicability of our framework, we implement Meta-GNN with two GNN models, i.e., SGC and GCN, forming two instances Meta-SGC and Meta-GCN. To accelerate the model convergence, we set batch size as 5 for Cora and Citeseer, but 12 for Reddit. The α_1 and α_2 are set to 0.5 and 0.003 in Meta-SGC, while the corresponding values are set to 0.1 and 0.001 in Meta-GCN respectively. The other settings of GNN module in the two models are the same as the suggestion of their original papers.

Experimental Result: Table 2 and Table 3 show the results of the performance comparison between Meta-GNN and baselines, from which we can clearly observe that our proposed model achieves the best performance across all three datasets. Generally speaking, GNN models, including ours, significantly outperform the embedding based models, e.g., DeepWalk and Node2Vec, on the few-shot learning scenarios. Among GNN models, we surprisingly find that two inductive GNN models, GraphSAGE with Mean and Pool operation/variants, have not shown competitive results, even compared with GCN and SGC. This result indicates that previous inductive graph learning models, which have shown promising results when encountering with *new nodes* (as reported in the original paper [3]), do not generalize well to *new classes*.

The superiority of Meta-GNN on Cora and Citeseer varies with the number of samples in support set, i.e., the less samples, the more improvement Meta-GNN improves over the baselines. This proves our primary motivation, i.e., adapting meta-learning into GNN models for few-shot graph learning. On the more challenging dataset Reddit, Meta-GNN can reap more improvement due to its capability of adapting to new tasks – there are more tasks for learning to learn the node representation and therefore a more general task adaption model obtained by Meta-GNN.

When comparing the two GNN modules in Meta-GNN, we did not observe any significant discrepancy – Meta-SGC performs better on Cora while Meta-GCN achieve slightly higher scores on Citeseer and Reddit, however we note that SGC and its corresponding meta-learning version Meta-SGC are ordered faster than GCN and Meta-GCN [11], respectively.

4 CONCLUSIONS

We have presented a generic graph meta-learning framework for few-shot node classification that leverages meta-learning mechanism to learn better parameter initialization of GNNs. The proposed Meta-GNN model can adapt well to new learning tasks (even new classes) with few labeled samples and significantly improves the performance in the context of few-shot node classification under meta-learning paradigm. Encouraging results have been obtained

Table 2: Accuracy comparison: Meta-GNN vs. baselines.

Datasets		Cora		Citeseer	
Model	K -shot	1-shot	3-shot	1-shot	3-shot
DeepWalk		16.06%	25.67%	14.52%	21.18%
Node2Vec		15.15%	25.66%	12.98%	20.02%
GraphSAGE		50.89%	53.12%	53.49%	55.01%
SGC		61.64%	75.67%	56.91%	65.67%
GCN		60.33%	75.15%	58.44%	67.99%
Meta-SGC		65.27%	77.19%	60.46%	68.65%
Meta-GCN		63.72%	76.78%	61.91%	69.43%

Table 3: Micro-F1 comparison: Meta-GNN vs. baselines.

Datasets		Reddit	
Model	K -shot	1-shot	3-shot
GraphSAGE		9.47%	15.89%
SGC		9.80%	16.98%
GCN		9.87%	17.17%
Meta-SGC		14.15%	20.14%
Meta-GCN		14.22%	20.71%

on three widely used datasets. In our future work, we would like to extend our framework to address more challenging problems such as few-shot graph classification and zero-shot node classification.

ACKNOWLEDGMENTS

This work was supported by National Natural Science Foundation of China (Grant No.61602097 and No.61472064), NSF grants III 1213038 and CNS 1646107 and ONR grant N00014-14-10215.

REFERENCES

- [1] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*.
- [2] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *KDD*.
- [3] Will Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *NIPS*.
- [4] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- [5] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*.
- [6] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*.
- [7] Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning. In *ICLR*.
- [8] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
- [9] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *NIPS*.
- [10] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *NIPS*.
- [11] Felix Wu, Tianyi Zhang, Amauri Holanda Souza Jr., Christopher Fifty, Tao Yu, and Kilian Q. Weinberger. 2019. Simplifying Graph Convolutional Networks. In *ICML*.
- [12] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. 2019. A Comprehensive Survey on Graph Neural Networks. *arXiv.org* (2019).
- [13] Shengzhong Zhang, Ziang Zhou, Zengfeng Huang, and Zhongyu Wei. 2018. Few-shot Classification on Graphs with Structural Regularized GCNs. *arXiv preprint* (2018).