

GRAIN: Improving Data Efficiency of Graph Neural Networks via Diversified Influence Maximization

Wentao Zhang^{†‡}, Zhi Yang^{†§}, Yexin Wang[†], Yu Shen[†], Yang Li[†], Liang Wang[†], Bin Cui^{†§}

[†]School of EECS & Key Laboratory of High Confidence Software Technologies, Peking University [‡]Tencent Inc.

[§]Center for Data Science, Peking University & National Engineering Laboratory for Big Data Analysis and Applications

[†]{wentao.zhang, yangzhi, yexinwang, shenyu, liyang.cs, bin.cui}@pku.edu.cn [†] liang_wang_now@yahoo.com

ABSTRACT

Data selection methods, such as active learning and core-set selection, are useful tools for improving the data efficiency of deep learning models on large-scale datasets. However, recent deep learning models have moved forward from independent and identically distributed data to graph-structured data, such as social networks, e-commerce user-item graphs, and knowledge graphs. This evolution has led to the emergence of Graph Neural Networks (GNNs) that go beyond the models existing data selection methods are designed for. Therefore, we present GRAIN, an efficient framework that opens up a new perspective through connecting data selection in GNNs with *social influence maximization*. By exploiting the common patterns of GNNs, GRAIN introduces a novel feature propagation concept, a diversified influence maximization objective with novel influence and diversity functions, and a greedy algorithm with an approximation guarantee into a unified framework. Empirical studies on public datasets demonstrate that GRAIN significantly improves both the performance and efficiency of data selection (including active learning and core-set selection) for GNNs. To the best of our knowledge, this is the first attempt to bridge two largely parallel threads of research, data selection, and social influence maximization, in the setting of GNNs, paving new ways for improving data efficiency.

PVLDB Reference Format:

Wentao Zhang, Zhi Yang, Yexin Wang, Yu Shen, Yang Li, Liang Wang, Bin Cui. Grain: Improving Data Efficiency of Graph Neural Networks via Diversified Influence Maximization. PVLDB, 14(11): 2473 - 2482, 2021. doi:10.14778/3476249.3476295

PVLDB Availability Tag:

The source code of this research paper has been made publicly available at <https://github.com/zwt233/Grain>.

1 INTRODUCTION

Data selection methods, such as active learning and core-set selection, improve the data efficiency of deep learning on large datasets by identifying the most informative training examples. In particular, active learning assists the learning procedure by prioritizing the selection of valuable unlabeled samples for human labeling, under the goal of maximizing the model performance with minimal

labeling cost [1, 2, 11, 16, 33, 45]. Core-set selection techniques aim to find a small subset that accurately approximates the full dataset by selecting representative examples [4, 18, 22, 41, 51, 52].

Graph Neural Networks (GNNs) have achieved state-of-the-art performance across various graph-based tasks such as node classification [5, 9, 49, 61] and link prediction [10, 20, 55, 56]. However, GNNs require plenty of labeled data to achieve satisfactory performance, and substantial training times. Therefore, introducing data selection methods for GNNs is crucial. Unfortunately, existing data selection methods fail to address the following new challenges posed by GNNs in terms of both performance and efficiency:

First, GNNs are graph-based semi-supervised learning models, which aggregate the feature of a node with its nearby neighbors. However, most existing data selection methods are designed to learn models on independent and identically distributed (i.i.d) data, which fail to model the interactions imposed by the graph structure. Several attempts [38, 42] have been made for data selection on graphs, but they are ineffective to GNNs since they cannot capture both graph structure and node features. As a result, it is necessary to design a new data selection criterion that is in coherence with the characteristics of GNNs to select the most valuable samples.

Second, GNNs incur substantial training costs and are hard to scale to large graphs because they need to perform a recursive neighborhood expansion to compute the hidden representations of a given node. Most existing active learning techniques select a batch of samples with guidance from the previously trained model and have to retrain a computationally-expensive model once a new labeled example comes. So it is expensive to directly apply them to GNNs, which hampers their applicability in real-world applications.

In this paper, we propose GRAIN, a novel data selection framework towards efficient data selection for GNNs. The working philosophy of GRAIN is to (1) take the essential component of GNNs – feature propagation – as a type of influence propagation, and (2) maximize the efficiency of feature propagation with analogy to *social influence maximization* [30]. Concretely, we measure the sensitivity of node v to node u (i.e., the “influence” of u on v) by computing how much the input feature of u affects the aggregated feature of v through feature propagation. Then we *maximize the influence* of the selected nodes (i.e., seed nodes) to get more unlabeled nodes influenced and involved in the downstream model training. In contrast to the traditional methods that consider informative scores of individual nodes, GRAIN is more effective in exploiting the interaction among nodes, by *explicitly* maximizing the number of unlabeled nodes influenced by labeled ones.

Although the definition of feature influence is inspired by [50, 58], the non-obvious observation that data selection for GNNs can be

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 11 ISSN 2150-8097. doi:10.14778/3476249.3476295

modeled as influence maximization is a major contribution. Moreover, we are making contributions in bridging these two largely parallel threads of research, active learning and social influence maximization, in new GNN settings: First, unlike the classic linear threshold (LT) model and independent cascade (IC) model [24] which use social influence maximization, we define a novel feature influence-based propagation model. Besides considering the direct influence via feature propagation, we further recognize the *indirect influence* of GNNs over feature space: nodes that are close in feature space are likely to have the same label. Thus, when a node is influenced via direct feature propagation, the propagated influence/label signal can also be viewed as indirectly interpolating a smooth signal [32] to nodes with close feature distance from the influenced nodes in feature space. To further improve the effectiveness of feature influence, we incorporate the diversity of influenced crowds to enforce the indirect influence, which encourages the influenced nodes to cover more semantic categories.

Based on the above perspective, we propose a novel data selection criterion for GNNs by unifying both the magnitude and the diversity of the influenced crowd into a *Diversified Influence Maximization* (DIM) criterion. For effective influence maximization, we construct new influence and diversity functions in GRAIN framework. We also prove that the proposed diversity functions and the objective functions have good properties of submodularity and monotonicity, which enable a simple greedy algorithm to find a near-optimal result. Meanwhile, GRAIN has the advantages of high efficiency and scalability over the existing methods, as it reduces the training cost by separating feature propagation from the training of GNNs.

In summary, the core contributions of this paper are: (1) **Problem Connection.** We open up a novel perspective for improving the data efficiency of GNNs by connecting GNN active learning with social influence maximization. (2) **New Criterion.** We propose a fundamentally new data selection criterion—*diversified influence maximization*—for GNNs by considering both the magnitude and diversity of feature influence. Moreover, we propose a greedy algorithm to solve this problem with an approximation guarantee. (3) **Novel Functions.** We propose novel influence and diversity functions based on both direct and indirect influence of GNNs over the graph and feature space, respectively, which have good properties of submodularity and monotonicity. (4) **High Performance and Efficiency.** Through experiments on real-world graphs with typical GNNs, we demonstrate that GRAIN significantly outperforms the state-of-the-art baselines on both performance and efficiency.

2 PRELIMINARY

In this section, we first describe the notations and define two data selection problems: Active Learning and Core-set Selection. Then we introduce GNNs and social influence maximization.

2.1 Data Selection Problems

We are given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{E}| = N$ nodes, and its adjacent matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. Each node $v_i \in \mathcal{V}$ is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^d$, which forms the feature matrix $\mathbf{X} \in \mathbb{R}^{N \times d}$. The ground-truth label for node v_i is a one-hot vector $\mathbf{y}_i \in \mathbb{R}^C$ where C is the number of classes and the c -th element is 1 only if node v_i belongs to class c . The entire node set \mathcal{V} is partitioned into

the training set \mathcal{V}_{train} , validation set \mathcal{V}_{val} and test set \mathcal{V}_{test} . The training algorithm is denoted as M , which is GNN in this paper.

Active Learning. Given the unlabeled node set \mathcal{V}_{train} and a loss function ℓ (e.g., generalization error), the goal of active learning is to select a subset of \mathcal{B} nodes S from \mathcal{V}_{train} to label so that the lowest loss on the test set \mathcal{V}_{test} can be achieved after applying S to algorithm M . The formulation of active learning is as follows:

$$\arg \min_{S:|S|=\mathcal{B}} \mathbb{E}_{v_i \in \mathcal{V}_{test}} [\ell(\mathbf{y}_i, P(\hat{\mathbf{y}}_i | \mathbf{x}_i; M_S))], \quad (1)$$

where $P(\hat{\mathbf{y}}_i | \mathbf{x}_i; M_S)$ is the label distribution of node v_i predicted by M_S , and M_S is trained under the supervision of the labeled set S .

Recent researches [14, 27, 44, 62] focus on using deep learning models as the learning algorithm M and solve the problem of active learning in the batch setting. During b rounds, they select $\frac{\mathcal{B}}{b}$ data points to label in each round. We denote the selected points during the i -th round as s_i . At the k -th round, these methods refit the learning algorithm M on all the labeled data $\bigcup_{i \in [k]} s_i$ to avoid any correlation between selections [27, 60]. Specifically, Wolf et al. [53], and Sener et al. [44] propose the greedy k-centers method to select data points that maximize the minimal distance between the current point and the labeled ones. Settles [46], Shen et al. [47], and Gal et al. [14] select points based on predictive confidence, i.e., the highest probability belonging to a certain class predicted by M .

While the above methods concentrate on supervised learning, AGE [3] and ANRMAB [15] are proposed for active learning on semi-supervised scenarios (e.g., graph) by using both the node features and graph structure. Recently, clustering-based active learning methods have been proposed for GNNs, such as Featprop [57] and LSCALE [37]. Featprop uses propagated node features as representations to cluster unlabeled nodes and labels the cluster centers. LSCALE proposes a learned latent space for clustering, which combines the unsupervised learning features and supervised hidden representations to select nodes for labeling. Compared with the related work, GRAIN fully utilizes the hidden information of unlabeled nodes, by *explicitly* maximizing the number of unlabeled nodes influenced by labeled ones over the graph and feature space (i.e., getting more unlabeled nodes involved in GNN training).

Core-set Selection. Suppose that \mathcal{V}_{train} is full labeled, core-set selection can be defined as methods that search for a subset of data points that maintain a similar level of quality (e.g., classification error) with the entire training set \mathcal{V}_{train} . Concretely, the goal of core-set selection is to select a subset of \mathcal{B} nodes $S \subset \mathcal{V}_{train}$ that achieves comparable result as the whole training set \mathcal{V}_{train} , that is:

$$\arg \min_{S:|S|=\mathcal{B}} \mathbb{E}_{v_i \in \mathcal{V}_{test}} [\ell(\mathbf{y}_i, P(\hat{\mathbf{y}}_i | \mathbf{x}_i; M_S)) - \ell(\mathbf{y}_i, P(\hat{\mathbf{y}}_i | \mathbf{x}_i; M_{\mathcal{V}_{train}}))], \quad (2)$$

where M_S and $M_{\mathcal{V}_{train}}$ are algorithms trained under the supervision of the labeled set S and the full training set \mathcal{V}_{train} respectively.

Besides the greedy k-centers method [44, 53] mentioned above, we introduce another two core-set selection techniques in previous work: forgetting events, and max entropy. Toneva et al. [48] define forget events as the number of times an example is incorrectly classified after having been correctly classified earlier during training and select the points with the highest number of forgetting events. Lewis et al. [31] and Settles [46] rank the entropy of the predictions from algorithm M and keep the points with the highest entropy.

2.2 Graph Neural Networks

Graph Neural Networks (GNNs) [17, 26, 49] define a multi-layer message passing process, through which the feature representation of a node in the next layer could be the aggregation of its neighborhood in the current layer. They differ in the way of defining the recursive function f for message passing:

$$\mathbf{X}^{(k)} \leftarrow f(\mathbf{X}^{(k-1)}, \mathbf{A}, \Theta^{(k)}), \quad (3)$$

where $\mathbf{X}^{(k)}$ and $\Theta^{(k)}$ are the output embedding and trainable parameters of layer k . Naturally, the input $\mathbf{X}^{(0)}$ satisfies $\mathbf{X}^{(0)} = \mathbf{X}$.

For example, Graph Convolution Network (GCN) [26] has a specific form of the function f as:

$$\mathbf{X}^{(k)} = \delta\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}^{(k-1)} \Theta^{(k)}\right), \quad (4)$$

where $\tilde{\mathbf{D}}$ is the diagonal degree matrix of $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$, and \mathbf{I}_N is the identity matrix. $\Theta^{(k)}$ is the layer-specific trainable weight matrix and $\delta(\cdot)$ is the non-linear activation function. For classification tasks, a K -layer GCN applies a softmax function on the aggregated representation in the final layer to obtain the prediction score for each class. That is, $\hat{\mathbf{Y}} = \text{softmax}\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}^{(K-1)} \Theta^{(K)}\right)$.

2.3 Social Influence Maximization

The influence maximization (IM) problem in social networks aims to select \mathcal{B} nodes so that the number of nodes activated (or influenced) in the social networks is maximized [24]. Namely, given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the formulation is as follows:

$$\max_S |\sigma(S)|, \text{ s.t. } S \subseteq \mathcal{V}, |S| = \mathcal{B}, \quad (5)$$

where $\sigma(S)$ is the set of nodes activated by the seed set S under certain influence propagation models, such as Linear Threshold (LT) and Independent Cascade (IC) models [24]. The maximization of $\sigma(S)$ is NP-hard. However, if $\sigma(S)$ is nondecreasing and submodular with respect to S , a greedy algorithm can provide an approximation guarantee of $(1 - \frac{1}{e})$ [40]. Great efforts have been devoted to the IM problem [8, 23, 59, 63], and some researches also explore the diversity over the activated nodes in social influence maximization [12, 30, 36]. But they are designed for specific models or tasks and cannot be applied to GNNs directly.

This paper focuses on *bridging* data selection and diversified influence maximization (DIM) in new GNN settings, rather than addressing classic social influence maximization problems in previous work. We propose a novel propagation model and diversity functions by exploiting the characteristics of GNNs, including both direct and indirect feature influence. We further propose a fundamentally new selection criterion by connecting GNN data selection with DIM, and demonstrate that it outperforms traditional learning-based AL approaches (i.e., achieves state-of-the-art performance for GNN data selection problem). GRAIN leverages the literature dedicated to (diversified) social influence maximization and exhibits both the feasibility and potential of such connection, which opens up a promising future direction in GNN-based active learning.

3 GRAIN FRAMEWORK

In this section, we present GRAIN, a new data selection framework for GNNs. As illustrated in Figure 1, at each round of data selection, GRAIN takes as input a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, feature matrix \mathbf{X} , and

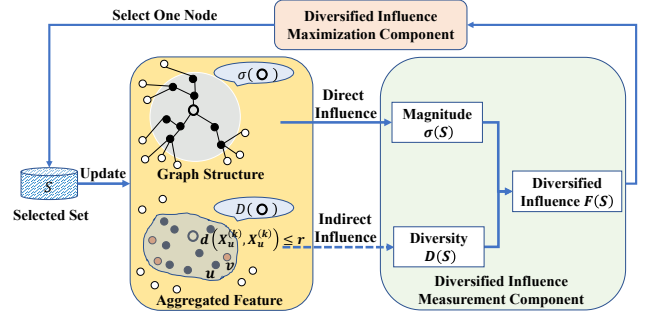


Figure 1: The workflow of GRAIN.

computes the proposed influence and diversity measurements based on the direct and indirect influence over graph structure and k -steps aggregated embedding $\mathbf{X}^{(k)}$, respectively. Next, GRAIN combines the influence and diversity into a unified criterion – maximizing the diversified influence and selects a node based on this new criterion. The procedure is repeated until the labeling budget \mathcal{B} exhausts. Below, we introduce each component of GRAIN in detail.

3.1 Feature Influence Model

Decoupled Feature Propagation. Graph neural networks define a multi-layer feature propagation process. While feature propagation and transformation might be intertwined in many GNNs, recent studies have observed that GNNs primarily derive their benefits from performing feature smoothing over graph neighborhoods rather than learning non-linear hierarchies of features as implied by the analogy to CNNs [13, 20, 54]. Thus, we separate the essential operation of GNNs – feature propagation – inherited from GNNs by removing the neural network Θ and non-linear activation for feature transformation. Specifically, we construct a parameter-free K -step feature propagation process for a target K -layer GNN:

$$\mathbf{X}^{(k)} \leftarrow f(\mathbf{X}^{(k-1)}, \mathbf{T}, \mathbf{X}^{(0)}), \quad \forall k = 1, \dots, K. \quad (6)$$

where \mathbf{T} is the generalized transition matrix, e.g., the random walk transition matrix $\mathbf{T}_{rw} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, the symmetric transition matrix $\mathbf{T}_{sym} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$, the triangle-IA matrix $\mathbf{T}_{tr} = \mathbf{D}_T^{-1} \mathbf{A}_T$, etc.

Feature Influence Viewpoint. Under the given decoupled propagation mechanism, $\mathbf{X}^{(k)}$ is the aggregated feature obtained by propagating features from nodes within k -hop neighborhood. Hence, $\mathbf{X}^{(k)}$ captures the information from the subtree of height k rooted at individual nodes. By taking feature propagation as a type of influence propagation, we open up a new perspective for the data selection problems in GNNs from the viewpoint of influence maximization (IM). Inspired by [50, 58], we measure the feature influence of a node u on v by how much a change in the input feature of u affects the *aggregated* feature of v after k iterations.

Definition 3.1 (Feature Influence). The feature influence score of node u on node v after k -step propagation is the L1-norm of the expected Jacobian matrix:

$$I(v, u, k) = \left\| \mathbb{E}[\partial \mathbf{X}_v^{(k)} / \partial \mathbf{X}_u^{(0)}] \right\|_1. \quad (7)$$

The normalized influence score is defined as:

$$I_v(u, k) = \frac{I(v, u, k)}{\sum_{w \in \mathcal{V}} I(v, w, k)}. \quad (8)$$

Table 1: Propagation mechanisms adopted in various GNNs. Note that each edge weight in A_T means the number of different triangles it belongs to and Triangle. IA is the abbreviation of Triangle-induced Adjacency.

Prorogation Mechanism	Formula
Normalized Adjacency [26]	$X^{(k)} = T_{sym}X^{(k-1)}$
Random Walk [54]	$X^{(k)} = T_{rw}X^{(k-1)}$
PPR [28]	$X^{(k)} = (1 - \alpha)T_{rw}X^{(k-1)} + \alpha X^{(0)}$
Triangle. IA [13]	$X^{(k)} = T_{tr}X^{(k-1)}$
S ² GC [64]	$X^{(k)} = \frac{1}{k}((1 - \alpha)T^k X^{(0)} + \alpha X^{(0)} + (k - 1)X^{(k-1)})$
GBP [7]	$X^{(k)} = \theta_k T^k X^{(0)} + X^{(k-1)}$

Given the k -step feature propagation mechanism (6), the feature influence score $I_v(u, k)$ captures the sum over probabilities of all possible influential paths from v to u . For example, let the generalized transition matrix be $T = T_{rw}$, the $I_v(u, k)$ is the probability that a random walk starting at v ends at u after taking k steps:

$$I_v(u, k) = \sum_{\mathcal{P}_k^{v \rightarrow u}} \prod_{i=k}^1 \tilde{a}_{v_{(i-1)}, v_{(i)}}, \quad (9)$$

where $\mathcal{P}_k^{v \rightarrow u}$ be a path $[v_{(k)}, v_{(k-1)}, \dots, v_{(0)}]$ of length k from node v to u and $\tilde{a}_{v_{(i-1)}, v_{(i)}}$ is the normalized weight of edge $(v_{(i)}, v_{(i-1)})$.

Discussions. GNN model family adopts a variety of propagation mechanisms, such as Random Walk and Personalized PageRank (PPR), as summarized in Table 1. The performance of different mechanisms depends on the task, graph structure, and features jointly. For example, in social networks, triangle-based propagation might help distinguishing edges representing weak or strong ties. On graphs with noisy connectivity, PPR may work well. One could also add other propagation kernels following Eq. (6). GRAIN can be applied to a large variety of GNNs by i) adopting similar propagation kernels used by these GNNs and ii) computing node influence with Eq. (8). In this way, we offer a general approach for GNN active learning from the novel perspective of feature influence.

The decoupled feature propagation in Eq. (6) enables GRAIN to perform node selection in a model-free manner, i.e., it selects the nodes to label once and for all before the GNN model starts. It is worth pointing that such decoupling in GRAIN *does not strictly* require the decoupled feature propagation in downstream GNN models. Besides, GRAIN also supports GNN with intertwined feature propagation and DNN transformation (e.g., GCN) and self-supervised GNN model (e.g., MVGRL [19]). The main reason for this generality is that: for both coupled and decoupled GNNs, (1) they primarily benefit from performing feature propagation over graph rather than learning non-linear hierarchies of features [54], and (2) their node influence distributions are consistently connected to probabilities of possible influential paths, e.g., [58] proves that the influence distribution in GCN is equivalent to the random walk distribution given by Eq. (9) when choosing $T = T_{rw}$ in Eq. (6).

3.2 Diversified Influence Maximization

Influence Function. Intuitively, the weak influence of a label node u on v with small probability $I_v(u, k)$ would have limited impact on v due to few influence paths to propagate labels.

Therefore, we define the activation of a node v by requiring the maximum influence $I_v(S, k) = \max_{u \in S} I_v(u, k)$ of S on the node v is larger than a certain threshold value.




Figure 2: The relationship between seed nodes S ($|S| = 20$) and the test accuracy of GCN model trained on S of Cora.

Definition 3.2 (Activated Node Set). Given a threshold θ , k -step feature propagation (6), and a set of seeds S , the activated node set $\sigma(S)$ is a subset of nodes in \mathcal{V} that can be activated by S :

$$\sigma(S) = \bigcup_{v \in \mathcal{V}, I_v(S, k) > \theta} \{v\}, \quad (10)$$

Note that the threshold θ controls the value at which a node v is activated (or significantly influenced) by the labeled nodes set S . An inactive node v becomes active if $I_v(S, k) > \theta$. Therefore, a larger θ requires the labeled node set S to impose a stronger influence on the node v in order to activate it. The budget \mathcal{B} here is the number of nodes to label, i.e., $|S| \leq \mathcal{B}$. Given a small budget, the overall influence of the labeled node set S on other nodes is relatively weak. In this case, we set a small threshold $\theta \rightarrow 0$ to make nodes more easily to be activated, so that we can choose nodes (to label) that can activate more unlabeled nodes for model training. Otherwise, few nodes can be activated by S . On the contrary, given enough labeling budget, the overall influence of the label node set S is relatively larger, and we should set a larger threshold $\theta > 0$ to ensure sufficient influence of S on each activated node v .

Since increasing the influence of nodes that are already activated cannot benefit $\sigma(S)$, it is easy to derive the following theorem:

THEOREM 3.3. $|\sigma(S)|$ is nondecreasing and submodular with respect to S , i.e., $\forall S \subseteq T, v \notin T, |\sigma(T)| \geq |\sigma(S)|$ and $|\sigma(S \cup \{v\})| - |\sigma(S)| \geq |\sigma(T \cup \{v\})| - |\sigma(T)|$.

Influence Maximization. To increase the feature/label influence (smoothness) effect on graphs, we should select nodes that can influence more unlabeled nodes. Due to the impact of graph structure, the speed of expansion, i.e., the growth of the influence, can change dramatically given different sets of label nodes. This observation motivates us to address the graph data selection problem in the viewpoint of influence maximization defined in Eq. (5). To illustrate this insight, we randomly select different sets of $\mathcal{B} = 20$ labeled nodes and train a GCN model with the different labeled set on Cora. We sort these label sets in terms of $\sigma(S)$, and Figure 2(a) plots the relationships between influence magnitude $\sigma(S)$ and the GCN model accuracy trained with the label supervision of S . Even with the same number of 20 labeled nodes, Figure 2(a) shows that the accuracy tends to increase along with their influence magnitude, implying the potential gain of exploring the influence viewpoint.

Introducing Diversity. From Figure 2(a), we also observe that the node set $\sigma(S)$ with a similar influence still leads to different accuracy (i.e., high variance). This is because simply maximizing $\sigma(S)$

fails to model the interactions between nodes, i.e., activating one node could often affect the utility of activating another one. Since the aggregated features $\mathbf{X}^{(k)}$ contains both the feature and graph structure information, it is reasonable to assume that nodes that are close together in the aggregated feature space will be likely to have the same label. So when a node $v \in \sigma(S)$ is activated over the graph structure, the propagated influence/label signal can also be viewed as *indirectly* interpolating a smooth signal to nodes with close feature distance from v at the feature space, which we refer to as *indirect influence*. So besides the direct influence by feature propagation over the graph structure, we also consider enforcing this indirect influence over feature space to further improve the overall effectiveness of influence. To this end, the diversity of activated nodes $\sigma(S)$ needs to be explored in the feature space. Specifically, we expect more nodes in $\sigma(S)$ can scatter over different regions of \mathcal{V} in the feature space, aim to allow each node in \mathcal{V} can find a close node belonging to $\sigma(S)$. We shall introduce our diversity functions in the later Section 3.3.

Proposed Criterion. In this paper, we propose a new GNN data selection criterion to consider the magnitude of influence and the diversity of the influence simultaneously, from a novel viewpoint of diversified influence maximization. Specifically, our GRAIN framework adopts a diversified influence maximization objective:

$$\max_S F(S) = \frac{|\sigma(S)|}{|\hat{\sigma}|} + \gamma \frac{D(S)}{\hat{D}}, \text{ s.t. } S \subseteq \mathcal{V}, |S| = \mathcal{B}. \quad (11)$$

where $\sigma(S)$ is the influence function, $D(S)$ represents the diversity of the influenced crowd; and $\gamma \in [0, 1]$ is a trade-off parameter. $\hat{\sigma}$ and \hat{D} are normalization factors, commonly chosen as the maximum possible values of $|\sigma(S)|$ and $D(S)$, respectively.

To demonstrate our diversified influence principle, we measure the diversity of two nodes $u, v \in \sigma(S)$ with their Euclidean distance over k -hop aggregated feature space: $D(u, v) = \frac{1}{2} \left\| \frac{\mathbf{X}_u^{(k)}}{\|\mathbf{X}_u^{(k)}\|} - \frac{\mathbf{X}_v^{(k)}}{\|\mathbf{X}_v^{(k)}\|} \right\|$. Larger distance indicates higher dissimilarity of two representations. We compute the diversity of nodes in $\sigma(S)$ by their average of pairwise-distance. Figure 2(b) plots the relationship between influence diversity and the model accuracy when the influence magnitude $|\sigma(S)|$ is fixed to 400, demonstrating that the diversity significantly benefits the model performance.

3.3 Submodular Diversity Functions

However, pair-wise diversity is not monotone and submodular. To guarantee diversity and efficient greedy maximization, we next discuss a general recipe for constructing monotone and submodular diversity functions $D(S)$. Our scheme relies on enforcing the indirect influence by enabling each node in \mathcal{V} can find a close node belonging to $\sigma(S)$. The diversity is enforced by avoiding repeatedly adding similar nodes to $\sigma(S)$ (i.e., those have already been affected by the indirect influence.) Based on this scheme, we define two submodular diversity functions, i.e., the Nearest Neighbor(NN) based diversity and the coverage based diversity.

Nearest Neighbor (NN)-based Diversity. We measure the indirect influence for each node $v \in \mathcal{V}$ by v 's distance to the activated node $u \in \sigma(S)$ closest to it in the feature space. Let $d(\cdot)$ measures the euclidean distance of two aggregated feature vectors, and $d_{max} =$




Figure 3: The coverage-based diversity of a seed node.

$\max_{u, v \in \mathcal{V}} d(\mathbf{X}_u^{(k)}, \mathbf{X}_v^{(k)})$ is the maximum pairwise distance. We define the diversity function as:

Definition 3.4 (NN-diversity Function). Given the k -step feature propagation mechanism (6), the diversity function of seed set S is:

$$D_{NN}(S) = \sum_{u \in \mathcal{V}} \left(d_{max} - \min_{v \in \sigma(S)} d(\mathbf{X}_u^{(k)}, \mathbf{X}_v^{(k)}) \right). \quad (12)$$

THEOREM 3.5. *The function $D_{NN}(S)$ is nondecreasing and submodular with respect to S .*

Coverage-based Diversity. Notice that $D_{NN}(S)$ is to minimize the total distance, it might incur relatively high performance variation without considering the variance of distance (and thus the indirect influence). So we further propose a coverage-based strategy emphasizing variance-reduction for more stable performance. Specifically, we assume indirect influence is only valid within each r -radius ball G_u centered at each activated node $u \in \sigma(S)$: i.e., $G_u = \{v \mid \forall v \in \mathcal{V}, d(\mathbf{X}_u^{(k)}, \mathbf{X}_v^{(k)}) \leq r\}$. As illustrated in Figure 3, our diversity function is the aggregated indirect influence of $\sigma(S)$:

Definition 3.6 (Ball-diversity Function). Given a set of groups $\{G_u\}$ and $\sigma(S)$, the diversity function of seed set S is:

$$D_{ball}(S) = \left| \bigcup_{u \in \sigma(S)} G_u \right|. \quad (13)$$

The influence function $|\sigma(S)|$ can be treated as the special case of $D_{ball}(S)$ with ball radius $r = 0$, i.e., ignoring the indirect influence. Similar to $|\sigma(S)|$, $D_{ball}(S)$ is apparently increasing with larger $|S|$ since more nodes can be covered. Since more nodes have not been covered with larger $|S|$, it is harder to cover more nodes, and thus the marginal gain decreases accordingly. So we have:

THEOREM 3.7. *The function $D_{ball}(S)$ is nondecreasing and submodular with respect to S .*

Note that the diversity function enforcing the *indirect influence* of GNNs is completely new. The fundamental difference between our diversity functions and the classic approach is that the node coverage is determined by both the propagation over the graph and the distance over the feature space. For example, in our ball-diversity given by Eq. (13), the coverage of labeled node set S in the feature space is the ball-covered regions centered on node set $\sigma(S)$ influenced by S over the graph. By contrast, the classic coverage approach [43] only considers the covered region centered on S , which is designed for independent and identically distributed (i.i.d) data and fails to model the influence imposed by the graph structure.




Figure 4: The test accuracy across different labeling budgets for model training.

Algorithm 1: Greedy node selection

Input: Graph \mathcal{G} , feature X , propagation mechanism f and layer number k for target GNN, budget \mathcal{B} .

Output: Seed set S

```

1 for  $i = 1, 2, \dots, k$  do
2    $\mathbf{X}^{(k)} \leftarrow f(\mathbf{X}^{(k-1)}, \mathbf{T}, \mathbf{X}^{(0)})$ ;
3    $S = \emptyset$ ;
4   for  $t = 1, 2, \dots, \mathcal{B}$  do
5     for  $v \in \mathcal{V}_{train} \setminus S$  do
6       update  $\sigma(S \cup \{v\})$  and  $D(S \cup \{v\})$  based on (10) and
          (12) (or (13)), respectively;
7        $v^* = \arg \max_{v \in \mathcal{V}_{train} \setminus S} F(S \cup \{v\}) - F(S)$ ;
8        $S = S \cup \{v^*\}$ ;
9   return  $S$ 

```

3.4 Selection Algorithm

Greedy Algorithm. Algorithm 1 provides a sketch of our greedy node selection method for GNNs. Without losing generality, we consider a batch setting with \mathcal{B} rounds where one node is selected in each iteration. Given the propagation mechanism f and layer k inherited from a target GNN, we first perform the propagation based on equation (6) (line 2). Notice the marginal gain $F(S \cup \{v\}) - F(S)$ of each node $v \in \mathcal{V}_{train} \setminus S$ is closely correlated to the current label set S , which is decreasing as the S grows. Once a vertex is selected and added to the label set S , we update the marginal gain of each node $v \in \mathcal{V}_{train} \setminus S$ based on the new set S . Specifically, the influence and diversity score of $\sigma(S \cup \{v\})$ and $D(S \cup \{v\})$ would be updated according to Equations (10) and (12) (or (13)), respectively (lines 5-6). Next, we select the node v^* generating the maximum marginal gain, and the selected nodes set S are updated (line 7-8). For monotone and submodular F , the final selected node set S is within a factor of $(1 - \frac{1}{e})$ of the optimal set S^* : $F(S) \geq (1 - \frac{1}{e})F(S^*)$. *Efficiency Optimization.* Compared to existing learning-based methods, GRIN provides high efficiency and scalability advantages, as it avoids the training cost by separating the propagation mechanism from the neural networks. By leveraging the existing works on scalable and parallelizable social influence maximization, we could enable GRIN to effectively deal with large-scale graphs. The key idea is to identify and dismiss uninfluential nodes in order to dramatically reduce the amount of computation for evaluating

influence spread. For example, we can use the degree of nodes or the distribution of random walkers throughout the nodes [25] to filter out a vast number of uninfluential nodes.

4 EXPERIMENTS

4.1 Experimental Settings

Datasets. We evaluate GRIN in both inductive and transductive settings [17] on three citation networks (i.e., Citeseer, Cora, and PubMed) [26], one large social network (Reddit), and the largest benchmark dataset ogbn-papers100M [21].

GNN Models. We conduct experiments using the widely used GCN model and demonstrate the generalization of GRIN on other GNNs such as SGC [54], APPNP [29] and MVGRL [19] in Section 4.5.

Baselines. We compare GRIN (ball-D) and GRIN (NN-D) with the following baselines: Random, Degree, AGE [3], ANRMAB [15], K-Center-Greedy (KCG) [44].

Settings. For each method, we use the hyper-parameter tuning toolkit [34, 35] or follow the original papers to find the optimal hyperparameters. To eliminate randomness, we repeat each method ten times and report the mean performance. Specifically, each learning-based active learning strategy chooses a small set of labeled nodes as an initial pool. Like AGE, we consider the label balance, and two nodes are randomly selected for each class. Note that GRIN can guide the node selection from scratch under the new diversified influence maximization criterion.

4.2 Performance Comparison

Active Learning. Let C be the number of classes for each dataset (e.g., 7 for Cora and 3 for PubMed). We choose the budget $\mathcal{B} = |S|$ from a range of $2C$ to $20C$ labeled nodes, and report the test accuracy of the GCN model trained on the selected labeled node set S along with the number of labeled nodes for training in Figure 4. As the labeling cost is proportional to the labeling budget, Figure 4 equivalently shows the improvement in terms of labeling cost.

Compared to the other baselines, both GRIN (ball-D) and GRIN (NN-D) quickly boost the accuracy at the beginning and consistently outperform the baselines as the number of labeled nodes grows. Concretely, the competitive baseline AGE has to label 120 nodes to achieve the accuracy of 71.4% on Citeseer, while GRIN




Figure 5: The label rate (the percentage of the whole available node set) to reach an accuracy gap from the best accuracy.

Table 2: Test accuracy using 20C labeled nodes in active learning scenario. OOT means “out of time”.

Method	Cora	Citeseer	PubMed	Reddit	ogbn-papers100M
Random	78.8	70.8	78.9	91.1	51.2
Degree	81.8	70.9	78.3	91.4	51.5
AGE	82.5	71.4	79.4	91.6	OOT
ANRMAB	82.4	70.6	78.2	91.5	OOT
KCG	82.6	71	79.3	91.3	51.6
GRAIN (NN-D)	83.3	73.7	80.8	92.5	52.6
GRAIN (ball-D)	84.2	74.2	81.8	92.3	52.9

(NN-D) only needs 60 labeled nodes to achieve similar results, indicating that GRAIN could cut the cost (e.g., money) by half for users. This improvement demonstrates the effectiveness of our DIM selection criterion. Moreover, GRAIN can avoid the sensitivity to model accuracy inherited from the learning-based methods like AGE and ANRMAB, especially when the model is under-fitted given a small labeling budget. Concretely, GRAIN (ball-D) outperforms AGE by a margin of 8.8% on Citeseer when 18 labeled nodes are used.

To demonstrate the improvement of GRAIN on the final performance, we also provide the test accuracy using all the 20C labeled nodes. Table 2 shows that AGE and ANRMAB outperform the Random and Degree method on most graph datasets, and GRAIN (ball-D) and GRAIN (NN-D) further boost the performance by a significant margin. GRAIN (ball-D) improves the test accuracy of the best baseline AGE by 1.7-2.8% on the three citation networks while GRAIN (NN-D) also outperforms AGE by a margin of 0.9% on Reddit. Note that GRAIN (ball-D) outperforms GRAIN (NN-D) on three citation networks, while GRAIN (NN-D) performs better than GRAIN (NN-D) on Reddit. It is because the citation networks have low degrees (i.e., are sparse). Thus the model performance is more vulnerable to the variance of node feature distance (indirect influence). GRAIN (NN-D) performs better on dense graphs (Reddit) since it can minimize the total distance, while GRAIN (ball-D) outperforms GRAIN (NN-D) on sparse graphs by emphasizing variance-reduction.

We also conduct an experiment on the largest benchmark dataset ogbn-papers100M. For such large-scale graphs, the learning-based methods such as AGE and ANRMAB require extremely long training time. In our experiment, both AGE and ANRMAB fail to finish the training within two weeks, and this indicates that GRAIN (26.1 hours) could achieve at least an order of magnitude speedup. Besides, we observe that GRAIN (ball-D) outperforms the second-best method KCG by a large margin of 1.3%.




Figure 6: End-to-end runtime (at log scale) using GPU.

Core-set Selection. Core-set selection starts with a large labeled or unlabeled dataset and aims to find a small subset that accurately approximates the entire dataset. We first train the exact model with all training labels (i.e., 18217) and get the test accuracy (86.5%) on PubMed. As shown in Figure 5, we evaluate the number of labeled nodes for each core-set selection method to achieve the corresponding accuracy gap. It is evident that both GRAIN (ball-D) and GRAIN (NN-D) significantly outperform the other baselines and achieve the same accuracy gap using much fewer labeled nodes. Concretely, to achieve an accuracy gap of 2% on PubMed, AGE needs 3.2% of all labeled training nodes while GRAIN (ball-D) needs only 1% of these nodes, which means GRAIN (ball-D) outperform AGE by 3.2× in terms of data efficiency.

4.3 Efficiency and Scalability Analysis

Another advantage of GRAIN is its high efficiency in node selection. We report the end-to-end runtime of active learning methods in Figure 6. Note that the end-to-end runtime results include the overhead from both the node importance measurement and node selection, and do not consider the human-intensive Oracle labeling time which depends on the complexity of a specific task and monetary budget. Moreover, GRAIN is oracle-free and model-free, that is, the labeling process is not required in the node selection. For comparison, the learning-based process of AGE and ANRMAB has to wait for the oracle to provide labels in each iteration.

Figure 6(a) shows that GRAIN (ball-D) obtains a speedup of 37×, 44×, and 231× over ANRMAB on Cora, PubMed, and Reddit respectively on GPU. To test the scalability on large-scale graphs, we sample different scales on ogbn-papers100M. Since training GCN on such a large graph will lead to the out-of-memory exception, instead, we use SGC as the training model for both AGE and GRAIN. Figure 6(b) demonstrates GRAIN achieves a linear scaling on the ogbn-papers dataset, and at least one order of magnitudes faster than AGE. To select the same number of labeled nodes on the 100M graph scale, GRAIN (ball-D) and GRAIN (NN-D) take 26.1 hours and 42.5 hours respectively, while it takes AGE more than one year to achieve this by estimating its runtime trend.

4.4 Ablation Study

Our method combines both the influence magnitude and the diversity measures. To verify the necessity of each component, we evaluate GRAIN (ball-D) while disabling one measure at a time. We evaluate GRAIN (ball-D): (i) without the diversity and only maximize $|\sigma(S)|$ (called "No Diversity"); (ii) without the influence magnitude

Table 3: Influence of different components.

Method	Cora	Δ	Citeseer	Δ	PubMed	Δ
No Magnitude	81.1	-3.1	70.8	-3.4	76.7	-5.1
No Diversity	82.2	-2.0	71.2	-3.0	79.9	-1.9
Classic Coverage	82.3	-1.9	71.5	-2.7	80.2	-1.6
GRAIN (ball-D)	84.2	-	74.2	-	81.8	-

Table 4: Test accuracy of different models on PubMed.

Method	SGC	APPNP	GCN	MVGRL
Random	77.6	79.2	78.9	79.3
Degree	77.3	78.6	78.3	78.7
AGE	78.8	79.9	79.4	79.9
ANRMAB	77.8	78.7	78.2	78.9
KCG	78.2	79.7	79.3	79.8
GRAIN (NN-D)	80.2	81.6	80.8	81.8
GRAIN (ball-D)	81.1	82.0	81.8	82.1

and the goal is to cover the maximum number of nodes with the balls generated from selected nodes S (called "No Magnitude"); (iii) replace $|\sigma(S)|$ with S when computing diversity (called "Classic Coverage"). Table 3 displays the results of these methods.

Influence Magnitude. The influence magnitude has a significant impact on model performance on all datasets, and it is more important than diversity since removing it will lead to a significant performance gap. For example, the gap on PubMed is 5.1%, which is much higher than the other gap (1.9%). The higher the influence magnitude is, the more labeled nodes we can use to train a GNN.

Influence Diversity. The test accuracy decreases in all three datasets if we ignore the influence diversity. For example, the performance gap is as large as 3.0% if the influence diversity is removed on Citeseer. The higher the influence diversity is, the more nodes are influenced in the aggregated feature space.

To further demonstrate the novelty of our diversity function, we add a baseline that adopts a classic coverage approach for diversity measurement. The result shows the test accuracy of GRAIN has decreased by a large margin if we replace $\sigma(S)$ with S in the ball-diversity function, which verifies the necessity of considering propagation in the diversity measurement.

4.5 Generalization

In addition to GCN, GRAIN can also be applied to a large variety of GNN variants. GCN, SGC, MVGRL and APPNP are four representative GNNs [6] which adopt different message passings. Unlike the coupled GCN, both SGC and APPNP are decoupled, while their orderings when doing feature propagation and transformation are different. Besides, MVGRL is a classic self-supervised GNN. We test the generalization ability of GRAIN by evaluating the aforementioned four types of GNNs on 20C nodes selected by GRAIN and other baselines in the AL scenario, and the corresponding results are shown in Table 4. The results suggest that both GRAIN (ball-D) and GRAIN (NN-D) consistently outperform the other baselines, regardless of the coupled and decoupled GNNs. Moreover, the result also shows our proposed method GRAIN can significantly outperform the compared baselines on the top of more sophisticated self-supervised GNN such as MVGRL [19]. As shown in the table, the test accuracy of GRAIN (ball-D) could outperform KCG




Figure 7: Distribution of the seed nodes and activated nodes selected by GRAIN (ball-D) and AGE. The red triangle refers to the seed nodes, the blue circle refers to the activated nodes, and the gray x refers to the non-activated nodes.

and AGE by more than 2% on PubMed. Therefore, we conclude that GRAIN can generalize to different types of GNNs well.

4.6 Model Interpretability

In this section, we illustrate the insight of our method and conduct experiments to explain why GRAIN obtains better performance from the perspective of influence magnitude and influence diversity. To depict the data distribution, we randomly choose 60 nodes from Citeseer and select 12 of them using GRAIN (ball-D) and AGE as the labeled set. Then we mark the seed nodes, activated nodes, and non-activated nodes in Figure 7 and use t-SNE [39] to visualize them in the aggregated feature space.

For influence magnitude $|\sigma(S)|$, we observe that the number of non-activated nodes of AGE is larger than GRAIN, which means more unlabeled nodes are activated and get involved in the model training using GRAIN (ball-D). For influence diversity $D(S)$, it is evident that the activated nodes of GRAIN (ball-D) scatter over different regions of the whole dataset, while the nodes gather in some specific area of AGE. In this way, more unlabeled nodes in GRAIN (ball-D) can be affected by the indirect influence from $\sigma(S)$, and the performance is boosted accordingly.

5 CONCLUSION

GNNs are emerging deep learning models that arise naturally from the requirements of applying neural network models on graphs. Efficient and scalable data selection for GNN training is demanding but still challenging due to its inherent complexity. This paper advocates a novel perspective for GNN data selection by connecting it with social influence maximization. GRAIN represents a critical step in this direction by showing both the feasibility and potential of this connection. To this end, we define a new feature influence model to exploit the common patterns of GNNs and propose novel submodular influence and diversity functions. Experiments show that GRAIN outperforms competitive baselines by a large margin in terms of both model performance and efficiency.

ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China (No. 2018YFB1004403), NSFC (No. 61832001, 6197200), Beijing Academy of Artificial Intelligence (BAAI), and PKU-Tencent joint research Lab. Wentao Zhang and Zhi Yang contributed equally, and Bin Cui is the corresponding author.

REFERENCES

- [1] Xiangnan Aggarwal, Charu C. and Kong, Qanquan Gu, Jiawei Han, and Philip S. Yu. 2014. Active Learning: A Survey. 571–605.
- [2] Mustafa Bilgic, Lilyana Mihalkova, and Lise Getoor. 2010. Active Learning for Networked Data. In *In Proceedings of the 27th International Conference on Machine Learning*.
- [3] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. 2017. Active learning for graph embedding. *arXiv preprint arXiv:1705.05085* (2017).
- [4] Trevor Campbell and Tamara Broderick. 2018. Bayesian coresets construction via greedy iterative geodesic ascent. In *International Conference on Machine Learning*, PMLR, 698–706.
- [5] Jie Chen, Tengfei Ma, and Cao Xiao. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. In *International Conference on Learning Representations*.
- [6] Lei Chen, Zhengdao Chen, and Joan Bruna. 2020. On Graph Neural Networks versus Graph-Augmented MLPs. *arXiv preprint arXiv:2010.15116* (2020).
- [7] Ming Chen, Zhewei Wei, Bolin Ding, Yaliang Li, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. 2020. Scalable Graph Neural Networks via Bidirectional Propagation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- [8] Wei Chen, Laks VS Lakshmanan, and Carlos Castillo. 2013. Information and influence propagation in social networks. *Synthesis Lectures on Data Management* 5, 4 (2013), 1–177.
- [9] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-GCN: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. KDD.
- [10] Ganqu Cui, Jie Zhou, Cheng Yang, and Zhiyuan Liu. 2020. Adaptive Graph Encoder for Attributed Graph Embedding. In *Proceedings of SIGKDD 2020*.
- [11] Sanjoy Dasgupta. 2005. Analysis of a greedy active learning strategy. *Advances in neural information processing systems* 17 (2005), 337–344.
- [12] Debadeepta Dey, Tian Yu Liu, Martial Hebert, and J Andrew Bagnell. 2013. Contextual sequence prediction with application to control library optimization. *Proceedings of robotics: Science and systems VIII* (2013).
- [13] Fabrizio Frasca, Emanuele Rossi, Davide Eynard, Benjamin Chamberlain, Michael Bronstein, and Federico Monti. 2020. SIGN: Scalable Inception Graph Neural Networks. In *ICML 2020 Workshop on Graph Representation Learning and Beyond*.
- [14] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, PMLR, 1183–1192.
- [15] Li Gao, Hong Yang, Chuan Zhou, Jia Wu, Shirui Pan, and Yue Hu. 2018. Active discriminative network representation learning. In *IJCAI International Joint Conference on Artificial Intelligence*.
- [16] Daniel Golovin and Andreas Krause. 2011. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. In *Journal of Artificial Intelligence Research*, 42, 427–486.
- [17] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 1025–1035.
- [18] Sarel Har-Peled and Akash Kushal. 2007. Smaller coresets for k-median and k-means clustering. In *Discrete & Computational Geometry*, 37(1), 3–19.
- [19] Kaveh Hassani and Amir Hosein Khasahmadi. 2020. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*, PMLR, 4116–4126.
- [20] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 639–648.
- [21] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. 2021. OGB-LSC: A Large-Scale Challenge for Machine Learning on Graphs. *arXiv preprint arXiv:2103.09430* (2021).
- [22] Jonathan Huggins, Trevor Campbell, and Tamara Broderick. 2016. Coresets for scalable bayesian logistic regression. In *In Advances in Neural Information Processing Systems*. 4080–4088.
- [23] Kyomin Jung, Wooram Heo, and Wei Chen. 2012. Irie: Scalable and robust influence maximization in social networks. In *2012 IEEE 12th International Conference on Data Mining*. IEEE, 918–923.
- [24] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. 137–146.
- [25] Seungkeol Kim, Dongeun Kim, Jinoh Oh, Jeong-Hyon Hwang, Wook-Shin Han, Wei Chen, and Hwanjo Yu. 2017. Scalable and parallelizable influence maximization with random walk ranking and rank merge pruning. *Information Sciences* 415 (2017), 171–189.
- [26] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [27] Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. 2019. BatchBALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning. In *NeurIPS*.
- [28] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Personalized Embedding Propagation: Combining Neural Networks on Graphs with Personalized PageRank. *CoRR abs/1810.05997* (2018).
- [29] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997* (2018).
- [30] Andreas Krause, Ajit Singh, and Carlos Guestrin. 2008. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research* 9, 2 (2008).
- [31] David D Lewis and William A Gale. 1994. A sequential algorithm for training text classifiers. In *SIGIR'94*. Springer, 3–12.
- [32] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [33] Xin Li and Yuhong Guo. 2013. Adaptive active learning for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 859–866.
- [34] Yang Li, Yu Shen, Wentao Zhang, Yuanwei Chen, Huajun Jiang, Mingchao Liu, Jiawei Jiang, Jinyang Gao, Wentao Wu, Zhi Yang, et al. 2021. OpenBox: A Generalized Black-box Optimization Service. *arXiv preprint arXiv:2106.00421* (2021).
- [35] Yang Li, Yu Shen, Wentao Zhang, Jiawei Jiang, Bolin Ding, Yaliang Li, Jingren Zhou, Zhi Yang, Wentao Wu, Ce Zhang, et al. 2021. VolcanoML: Speeding up End-to-End AutoML via Scalable Search Space Decomposition. *arXiv preprint arXiv:2107.08861* (2021).
- [36] Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. 510–520.
- [37] Juncheng Liu, Yiwei Wang, Bryan Hooi, Renchi Yang, and Xiaokui Xiao. 2020. Active Learning for Node Classification: The Additional Learning Ability from Unlabelled Nodes. *arXiv preprint arXiv:2012.07065* (2020).
- [38] Jun Long, Jianping Yin, Wentao Zhao, and En Zhu. 2008. Graph-based active learning based on label propagation. In *International Conference on Modeling Decisions for Artificial Intelligence*. Springer, 179–190.
- [39] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. In *Journal of Machine Learning Research*, 9(Nov). 2579–2605.
- [40] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical programming* 14, 1 (1978), 265–294.
- [41] Chongjia Ni, Cheung-Chi Leung, Lei Wang, Nancy F Chen, and Bin Ma. 2015. Unsupervised data selection and word-morph mixed language model for tamil low-resource keyword search. In *In Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference*. 4714–4718.
- [42] Adarsh Prasad, Stefanie Jegelka, and Dhruv Batra. 2014. Submodular meets structured: finding diverse subsets in exponentially-large structured item sets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2*. 2645–2653.
- [43] Adarsh Prasad, Stefanie Jegelka, and Dhruv Batra. 2014. Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger (Eds.). 2645–2653.
- [44] Ozan Sener and Silvio Savarese. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *International Conference on Learning Representations*.
- [45] Burr Settles. 2009. *Active learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
- [46] Burr Settles. 2012. Active learning. *Synthesis lectures on artificial intelligence and machine learning* 6, 1 (2012), 1–114.
- [47] Yanyao Shen, Hyokun Yun, Zachary C Lipton, Yakov Kronrod, and Animashree Anandkumar. 2017. Deep Active Learning for Named Entity Recognition. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*. 252–256.
- [48] Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. 2018. An Empirical Study of Example Forgetting during Deep Neural Network Learning. In *International Conference on Learning Representations*.
- [49] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [50] Hongwei Wang and Jure Leskovec. 2020. Unifying graph convolutional neural networks and label propagation. *arXiv preprint arXiv:2002.06755* (2020).
- [51] Kai Wei, Yuzong Liu, Katrin Kirchhoff, Chris Bartels, and Jeff Bilmes. 2014. Submodular subset selection for large-scale speech training data. In *In Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference*. 3311–3315.
- [52] Kai Wei, Yuzong Liu, Katrin Kirchhoff, and Jeff A Bilmes. 2013. Using document summarization techniques for speech data subset selection. In *In HLT-NAACL*.

- [53] Gert W Wolf. 2011. Facility location: concepts, models, algorithms and case studies.
- [54] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*. PMLR, 6861–6871.
- [55] Shiwen Wu, Fei Sun, Wentao Zhang, and Bin Cui. 2020. Graph neural networks in recommender systems: a survey. *arXiv preprint arXiv:2011.02260* (2020).
- [56] Shiwen Wu, Yuanxing Zhang, Chengliang Gao, Kaigui Bian, and Bin Cui. 2020. GARG: Anonymous Recommendation of Point-of-Interest in Mobile Networks by Graph Convolution Network. *Data Science and Engineering* 5, 4 (2020), 433–447.
- [57] Yuexin Wu, Yichong Xu, Aarti Singh, Yiming Yang, and Artur Dubrawski. 2019. Active learning for graph neural networks via node feature propagation. *arXiv preprint arXiv:1910.07567* (2019).
- [58] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In *ICML*. 5449–5458.
- [59] Yu Yang, Enhong Chen, Qi Liu, Biao Xiang, Tong Xu, and Shafqat Ali Shad. 2012. On approximation of real-world influence spread. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 548–564.
- [60] Wentao Zhang, Jiawei Jiang, Yingxia Shao, and Bin Cui. 2020. Snapshot boosting: a fast ensemble framework for deep neural networks. *Sci. China Inf. Sci.* 63, 1 (2020), 112102.
- [61] Wentao Zhang, Xupeng Miao, Yingxia Shao, Jiawei Jiang, Lei Chen, Olivier Ruas, and Bin Cui. 2020. Reliable data distillation on graph convolutional network. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1399–1414.
- [62] Wentao Zhang, Yu Shen, Yang Li, Lei Chen, Zhi Yang, and Bin Cui. 2021. ALG: Fast and Accurate Active Learning Framework for Graph Convolutional Networks. In *Proceedings of the 2021 International Conference on Management of Data*. 2366–2374.
- [63] Yuying Zhao, Yunfei Hu, Pingpeng Yuan, and Hai Jin. 2021. Maximizing Influence Over Streaming Graphs with Query Sequence. *Data Science and Engineering* (2021), 1–19.
- [64] Hao Zhu and Piotr Koniusz. 2021. Simple spectral graph convolution. In *International Conference on Learning Representations*.