

Train Your Own GNN Teacher: Graph-Aware Distillation on Textual Graphs

Costas Mavromatis^{1*}, Vassilis N. Ioannidis², Shen Wang², Da Zheng², Soji Adeshina², Jun Ma², Han Zhao^{2,3}, Christos Faloutsos^{2,4}, and George Karypis^{1,2}

¹ University of Minnesota

² Amazon Web Services

³ University of Illinois at Urbana-Champaign

⁴ Carnegie Mellon University

Abstract. How can we learn effective node representations on textual graphs? Graph Neural Networks (GNNs) that use Language Models (LMs) to encode textual information of graphs achieve state-of-the-art performance in many node classification tasks. Yet, combining GNNs with LMs has not been widely explored for practical deployments due to its scalability issues. In this work, we tackle this challenge by developing a Graph-Aware Distillation framework (GRAD) to encode graph structures into an LM for graph-free, fast inference. Different from conventional knowledge distillation, GRAD jointly optimizes a GNN teacher and a graph-free student over the graph’s nodes via a shared LM. This encourages the graph-free student to exploit graph information encoded by the GNN teacher while at the same time, enables the GNN teacher to better leverage textual information from unlabeled nodes. As a result, the teacher and the student models learn from each other to improve their overall performance. Experiments in eight node classification benchmarks in both transductive and inductive settings showcase GRAD’s superiority over existing distillation approaches for textual graphs. Our code is available at: <https://github.com/cmavro/GRAD>.

Keywords: Graph Neural Networks · Language Models · Knowledge Distillation.

1 Introduction

Graph Neural Networks (GNNs) offer state-of-the-art performance on graph learning tasks in real world applications, including social networks, recommendation systems, biological networks and drug interactions. GNNs [19,10,31] learn node representations via a recursive neighborhood aggregation scheme [8], which takes as input the node features and the graph structure. In textual graphs, text-based information is associated with the graph’s nodes. Methods such as bag-of-words, word2vec [24] or pre-trained Language Models (LMs), e.g., BERT [5],

* Work done while interning at Amazon Web Services, Santa Clara. Correspondence to: mavro016@umn.edu

are used to transform raw texts of nodes into features. Transforming raw text to numerical features is usually associated with a non-negligible cost. For example, LMs that rely on transformers [30] have a very large number of parameters and their cost depends on the LM’s architecture⁵. This high computational cost results in expensive and/or slow inference during deployment.

Recent works show that combining GNNs with LMs in an end-to-end manner [16,38] leads to state-of-the-art performance for node classification and link prediction tasks. Although powerful, these models are associated with expensive inference costs. During mini-batch inference, a K -layer GNN fetches an exponential number of nodes w.r.t. K , where raw texts of the fetched nodes need to be transformed to numerical features on the fly. As a result, combining LMs with GNNs in a cascaded manner exponentially grows the LM inference cost for each target node. This cost is prohibitive for applications where fast inference is vital, such as providing instant query-product matching in e-commerce.

Aiming at a balance between effectiveness and efficiency, we seek to transfer useful knowledge from a GNN teacher to a graph-free student via distillation. As the graph-free student does not use the graph structure during inference, the inference cost depends only on the LM employed. However, existing knowledge distillation (KD) approaches for graphs either do not take full advantage of the graph structure [47] or they require powerful student models (e.g., large LMs) to achieve good performance [50].

To address these limitations, we developed a Graph-Aware Distillation approach (GRAD) that *jointly* optimizes the GNN teacher with its graph-free student via a shared LM. The shared LM serves as an interaction module that allows the two models to learn from each other. On the one hand, the GNN teacher updates the LM’s parameters with graph-aware information and distills graph knowledge as soft-labels, which are provided to the student. On the other hand, the graph-free student imitates the GNN’s predictions and leverages textual information from unlabeled nodes to improve the fine-tuning of the shared LM. This dynamic interplay between the two models stimulates the student to not only mimic the GNN predictions but to learn node features that improve its teacher’s performance.




Fig. 1. **GradBERT** is on the Pareto front: Accuracy performance w.r.t. inference time for ogbn-arxiv. L is the input sequence length and OOM (out of memory) means that the model encountered GPU failure.

⁵ For example, the inference cost of a single transformer layer is $\mathcal{O}(L^2d + Ld^2)$, where L is the sequence length and d is the number of hidden dimensions.

GRAD is formulated as a multi-task learning for the shared LM, whose goal is to achieve good performance for both the GNN teacher and the graph-free student models. We designed three different strategies for optimizing the GRAD framework. Their key differences is on how tight the teacher and student models are coupled and on how much flexibility the student model has to fit to its teacher’s predictions. As a result, GRAD can be applied in both large-scale graphs, where the student model is a powerful LM, and in small-scale graphs, where the LM is substituted by simple MLPs.

Figure 1 illustrates the superiority of GRAD when the graph-free student model is a BERT model (GRADBERT). It outperforms a fine-tuned BERT model for node classification by 3.24%, and it is as effective as combining BERT with a GNN with 2.4x-13x smaller inference time. Our contributions are summarized below:

- We analyze and identify the limitations of conventional knowledge distillation for textual graphs, which have been previously under-studied (Section 3.2).
- We present a graph-aware distillation (GRAD) framework that couples a GNN teacher and a graph-free student together to fully exploit the underlying graph structure. GRAD improves classification for both seen and unseen nodes, that are present in either large or small-scale graphs (Section 6.1 and Section 6.2).
- We developed three different strategies (Section 4) to effectively optimize GRAD framework, which we comprehensively study. This enables GRAD to scale to large graphs and achieve state-of-the-art performance in node classification tasks.

2 Background

2.1 Problem Formulation

In an input *textual graph* $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, each node $v \in \mathcal{V}$ is associated with raw text, which we denote as X_v . \mathcal{V} is the node set and \mathcal{E} is the edge set. Let N denote the total number of nodes. For node classification, the prediction targets are $\mathbf{Y} \in \mathbb{R}^{N \times m}$, where row \mathbf{y}_v is a m -dim one-hot vector for node v . The node set is divided into labeled nodes \mathcal{V}^L and unlabeled nodes \mathcal{V}^U , i.e., $\mathcal{V} = \mathcal{V}^L \cup \mathcal{V}^U$. In inductive scenarios, the input graph \mathcal{G} is divided into two subgraphs $\mathcal{G} = \mathcal{G}^{\text{tran}} \cup \mathcal{G}^{\text{ind}}$, where $\mathcal{G}^{\text{tran}}$ is used for learning (transductive part) and \mathcal{G}^{ind} is used only during inference (\mathcal{G}^{ind} is the inductive part that is not observed during learning). In transductive scenarios, we have $\mathcal{G} = \mathcal{G}^{\text{tran}}$ and $\mathcal{G}^{\text{ind}} = \emptyset$.

As introduced Section 1, utilizing the textual graph \mathcal{G} for making predictions at test time results in slow inference. Thus, we seek to learn a graph-free model τ' , e.g., a LM, that only takes node text X_v during inference. It is desired that $\tau'(X_v) \approx f(\mathcal{G}, X_v)$, where f is a model that uses the graph, e.g., a GNN, so that the graph-free model achieves as effective node classification as a graph-based model.

2.2 GNNs on Textual Graphs

To handle raw texts in textual graphs, X_v is transformed to numerical features $\mathbf{x}_v \in \mathbb{R}^d$ via a function $\tau(\cdot)$,

$$\mathbf{x}_v = \tau(X_v). \quad (1)$$

For example, LMs, such as BERT [5] for modelling $\tau(\cdot)$, transform each token of the input sequence X_v to a representation. The final \mathbf{x}_v can be obtained as the representation of a specific token, which is usually the [CLS] token.

GNNs [19,31] transform a computation graph \mathcal{G}_v that is centered around node v to a d -dimensional node representation \mathbf{h}_v . We write the GNN transformation as

$$\mathbf{h}_v = \text{GNN}(\{\tau(X_u) : u \in \mathcal{G}_v\}), \quad (2)$$

where $\tau(X_u)$ generates the input features \mathbf{x}_u of node u . For a GNN with K layers, \mathcal{G}_v includes nodes and edges (with self-loops) up to K hops away from v .

The $(k+1)$ -th GNN layer takes as input a node's representation $\mathbf{h}_v^{(k)}$ of the previous layer k as well as the representations of its 1-hop neighbors. It aggregates them to a new representation $\mathbf{h}_v^{(k+1)}$ as follows,

$$\mathbf{h}_v^{(k+1)} = \phi(\{\mathbf{h}_u^{(k)} : u \in \mathcal{N}_v\}), \quad (3)$$

where \mathcal{N}_v is the set of direct neighbors of v and $\phi(\cdot)$ is an aggregation function. For example, a common GNN update $\phi(\cdot)$, which is employed in GraphSAGE [10] and RGCN [27], can be described as follows,

$$\mathbf{h}_v^{(k+1)} = \sigma\left(\mathbf{W}_{\text{self}}^{(k)}\mathbf{h}_v^{(k)} + \sum_{u \in \mathcal{N}_v} \mathbf{W}^{(k)}\mathbf{h}_u^{(k)}\right), \quad (4)$$

where $\mathbf{W}_{\text{self}}^{(k)}, \mathbf{W}^{(k)}$ are learnable parameters and $\sigma(\cdot)$ is a nonlinearity mapping. At the first layer, we usually have $\mathbf{h}_v^{(0)} = \mathbf{x}_v$, that are text features extracted from the node. Computing \mathbf{h}_v can be as expensive as $\mathcal{O}(S^K C)$, where $\tau(\cdot)$ is a LM with inference cost C , S is the neighborhood size, and K is the number of GNN layers.

3 Towards Graph-Aware Knowledge Distillation

As discussed in Section 2.1, we seek to learn a graph-free model τ' for fast inference, that also achieves as effective node classification as a GNN model f .

3.1 Knowledge Distillation

A straightforward solution is to use the Knowledge Distillation (KD) technique [11], where a powerful teacher model transfers knowledge to a simpler student model. In our case, the teacher model corresponds to a graph-based model (GNN), while the student model is a graph-free model, such as an LM.

We follow the standard KD paradigm, in which the teacher distills knowledge via soft-labels. At a high level, the algorithmic procedure is

- **First Stage.** The GNN teacher is trained for node classification. The objective is given by

$$\mathcal{L}_{\text{nc}} = \sum_{v \in \mathcal{V}^L} l_{\text{CE}}(\hat{\mathbf{t}}_v, \mathbf{y}_v), \quad (5)$$

where l_{CE} is the standard cross-entropy and

$$\hat{\mathbf{t}}_v = \text{MLP}\left(\text{GNN}(\{\mathbf{x}_u : u \in \mathcal{G}_v\})\right) \quad (6)$$

are teacher’s label predictions (logits). Numerical features \mathbf{x}_u are learned by fine-tuning an LM $\tau(\cdot)$ in an end-to-end manner; see Eq.(1). The trained GNN teacher generates soft-labels $\hat{\mathbf{t}}_v$ for all nodes $v \in \mathcal{G}$.

- **Second Stage.** Another LM $\tau'(\cdot)$ is trained to mimic GNN’s predictions $\hat{\mathbf{t}}_v$. The LM is optimized via

$$\mathcal{L}_{\text{KD}} = \sum_{v \in \mathcal{V}^L} l_{\text{CE}}(\hat{\mathbf{s}}_v, \mathbf{y}_v) + \lambda \sum_{v \in \mathcal{V}} l_{\text{KL}}(\hat{\mathbf{s}}_v, \hat{\mathbf{t}}_v), \quad (7)$$

where

$$\hat{\mathbf{s}}_v = \text{MLP}(\tau'(X_v)) \quad (8)$$

are the student’s predictions and l_{KL} is the KL-divergence between the student’s and teacher’s logits. Hyper-parameter $\lambda \in \mathbb{R}$ controls the relative importance of the knowledge distillation term.

3.2 What Does Knowledge Distillation Learn? An Analysis

In many cases, only a subset of the nodes in a graph is labeled; it is denoted as \mathcal{V}^L . Since the GNN’s predictions $\hat{\mathbf{t}}_v$ are treated as soft-labels, the second term in Eq.(7) allows KD to fine-tune its LM τ' over both labeled and unlabeled nodes. However, as $\hat{\mathbf{t}}_v$ are pre-computed after the first stage of GNN training, the graph-free student does not use the actual graph structure during learning (second stage). In other words, nodes are treated independently from the underlying graph and the graph-free model can only infer how nodes interact via the provided soft-labels by the GNN.

Next, we quantify the importance of capturing node interactions, which benefits many applications such as community detection and label propagation in graphs [17]. As discussed in Section 2.2, the node classification objective aims at transforming a subgraph \mathcal{G}_v centered around node v to its label \mathbf{y}_v . In information theory, this is equivalent to maximizing the mutual information $I(\cdot)$,

$$\max_f \sum_{v \in \mathcal{V}^L} I_f(\mathbf{y}_v; \mathcal{G}_v), \quad (9)$$

between \mathcal{G}_v and label distribution \mathbf{y}_v , where I_f is parametrized by the GNN function f . It is always true that $I(\cdot) \geq 0$.

If we consider \mathcal{G}_v as a joint distribution of the text nodal feature set $\mathcal{X}_v = \{X_u : u \in \mathcal{G}_v\}$ and edge set \mathcal{E}_v , we have

$$I(\mathbf{y}_v; \mathcal{G}_v) = I(\mathbf{y}_v; (\mathcal{X}_v, \mathcal{E}_v)) = I(\mathbf{y}_v; \mathcal{X}_v) + I(\mathbf{y}_v; \mathcal{E}_v | \mathcal{X}_v), \quad (10)$$

which is obtained via the chain rule of mutual information. Now by setting $\tilde{\mathcal{X}}_v = \mathcal{X}_v \setminus \{X_v\}$, we can further obtain

$$I(\mathbf{y}_v; \mathcal{G}_v) = I(\mathbf{y}_v; X_v) + I(\mathbf{y}_v; \tilde{\mathcal{X}}_v | X_v) + I(\mathbf{y}_v; \mathcal{E}_v | \mathcal{X}_v). \quad (11)$$

The joint mutual information is decomposed into three terms, (i) the information that comes from the node itself, (ii) the information that comes from other nodes in \mathcal{G}_v , and (iii) the additional information that comes from the actual links between nodes.

Remark 1. If the neighbor set or the graph structure is not utilized during learning, only $I(\mathbf{y}_v; X_v)$ of Eq.(11) can be maximized.

Corollary 1. Assume that predictions $\hat{\mathbf{t}}_u$ for unlabeled nodes $u \in \mathcal{V}^U$ are obtained via Eq.(6) with a GNN f and an LM τ . The graph-free student of KD, that employs an LM τ' , solves

$$\max_{\tau'} \sum_{u \in \mathcal{V}^U} I_{\tau'}(\hat{\mathbf{t}}_u; X_u), \quad (12)$$

where

$$\begin{aligned} \hat{\mathbf{t}}_u &= f(\tau(X_{u'}) : u' \in \mathcal{G}_u), \\ \text{s.t. } f, \tau &= \arg \max_{\tilde{f}, \tilde{\tau}} \sum_{v \in \mathcal{V}^L} I_{\tilde{f}, \tilde{\tau}}(\mathbf{y}_v; \mathcal{G}_v). \end{aligned} \quad (13)$$

The student models solves Eq.(12) without directly accessing the graph structure \mathcal{G} . Since $\hat{\mathbf{t}}_u$ are obtained through a teacher GNN in Eq.(13), which maximizes all three parts of Eq.(11), it is possible that $\hat{\mathbf{t}}_u$ can still implicitly contain graph information from \mathcal{G}_u . However, due to the data-processing inequality of mutual information, it holds that

$$I_{f, \tau}(\mathbf{y}_u; \mathcal{G}_u) \geq I(\mathbf{y}_u; \hat{\mathbf{t}}_u). \quad (14)$$

This means that, despite being a function of \mathcal{G}_u , the information contained in $\hat{\mathbf{t}}_u$ w.r.t. \mathbf{y}_u is less than that of the original \mathcal{G}_u . Thus, student's performance depends on *how informative* $\hat{\mathbf{t}}_u$ are, compared to \mathcal{G}_u .

Note that the objectives of Eq.(12) and Eq.(13) are applied over different sets of nodes \mathcal{V}^U and \mathcal{V}^L , which does not ensure that GNN f is the optimal function to encode graph information for nodes $u \in \mathcal{V}^U$. For example, a GNN might overfit to nodes that appear frequently in the training subgraphs, and ignore nodes that appear infrequently. As a result, the quality of predictions $\hat{\mathbf{t}}_u$ is in question. More importantly, it is not clear how the graph-free student will perform on

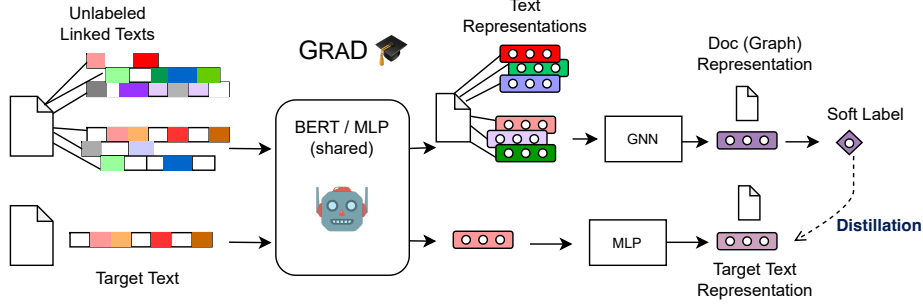


Fig. 2. GRAD framework. GRAD captures textual information among unlabeled linked texts by allowing the teacher GNN and the graph-free student to jointly update the shared text encoding function.

unseen nodes during training and, for which, soft-labels \hat{t}_u cannot be apriori computed by the GNN. However in textual graphs, nodes are associated with informative textual features (e.g., paper abstracts) and a well-trained student could potentially handle the textual information of the unseen (inductive) nodes. We provide a failing case of an under-trained student in Appendix G.

4 GRAD Framework

In this work, we take a different approach from the one shown in Eq.(12). As we cannot use the graph structure during inference, our GRaph-Aware Distillation (GRAD) improves structure utilization during training. The motivation is to use the original graph \mathcal{G} while training the graph-free student, even if GNN predictions \hat{t}_u may (or may not) implicitly contain such graph information.

Our GRAD framework does this by coupling the teacher GNN with its graph-free student with a shared function $\tau(\cdot)$. This allows the GNN to directly encode structural information into τ , that is then used by the graph-free model during inference. The overall framework is illustrated in Figure 2. GRAD is posed as a multi-task learning for the shared LM, whose goal is to collectively optimize the GNN teacher and the graph-free student models. The learning problem is given by

$$\max_{f, \tau} \sum_{v \in \mathcal{V}^L} I_{f, \tau}(\mathbf{y}_v; \mathcal{G}_v) + \sum_{u \in \mathcal{V}^U} I_{\tau}(\hat{t}_u; X_u), \quad (15)$$

where $\hat{t}_u = f(\tau(X_{u'})) : u' \in \mathcal{G}_u$,

and function τ contributes in both objectives that GRAD maximizes. This results in a *coupled* multi-objective optimization, since predictions \hat{t}_u depend on function τ and thus, tie the GNN f with the graph-free model together.

Notably, the graph-free student uses soft-labels provided by the GNN to fine-tune the LM over textual information from unlabeled nodes; see second




Fig. 3. GRAD strategies for coupling the GNN teacher and the graph-free student. SL denotes soft-label and GT denotes ground-truth label. GRAD-Joint couples the two models in a single step, GRAD-Alt couples them in an alternate fashion, while GRAD-JKD couples them in the first stage and decouples them in the second stage. For illustration purposes only, the star metric (\star) quantifies the teacher-student coupling tightness and the flexibility of the student model.

term of Eq.(15). This, consequently, ensures that the GNN leverages all the graph’s textual information (in an implicit manner) to compute $\hat{\mathbf{t}}_u$, which might have been neglected in Eq.(13). As the GNN teacher better encodes graph-aware textual information from the input (transductive) nodes, the graph-free student manages to *train its own GNN teacher* via the shared LM.

In what follows, we present three different optimization strategies (GRAD-Joint, GRAD-Alt, and GRAD-JKD) for solving Eq.(15). Figure 3 presents them at a high-level, and their key difference is on how tight they couple the teacher and the student model.

4.1 GraD-Joint

The first strategy (GRAD-Joint) optimizes GRAD framework in a single-stage, where the teacher and the student are jointly updating τ . GRAD-Joint objective is given by

$$\mathcal{L}_{\text{Joint}} = \lambda \sum_{v \in \mathcal{V}} l_{\text{KL}}(\hat{\mathbf{s}}_v, \hat{\mathbf{t}}_v) + \sum_{v \in \mathcal{V}^L} [\alpha l_{\text{CE}}(\hat{\mathbf{t}}_v, \mathbf{y}_v) + (1 - \alpha) l_{\text{CE}}(\hat{\mathbf{s}}_v, \mathbf{y}_v)], \quad (16)$$

where $\lambda \in \mathbb{R}$ and $\alpha \in [0, 1]$ are hyperparameters that control the importance of each term and ensure that the GNN models does not learn to ignore the structure. The first term is the KD loss between GNN teacher’s predictions $\hat{\mathbf{t}}_v$ and graph-free student’s predictions $\hat{\mathbf{s}}_v$, and the other two terms are for ground-truth labels \mathbf{y}_v . Predictions $\hat{\mathbf{s}}_v$ and $\hat{\mathbf{t}}_v$ are optimized via a shared encoding function $\tau(\cdot)$ (multi-task learning).

We highlight the dynamic interplay between the GNN teacher and the graph-free student in Eq.(16). The GNN teacher provides a soft-label $\hat{\mathbf{t}}_v$, which the graph-free student tries to mimic via $l_{\text{KL}}(\hat{\mathbf{s}}_v, \hat{\mathbf{t}}_v)$. However, instead of training $\tau(\cdot)$ to simply mimic the GNN predictions, term $l_{\text{CE}}(\hat{\mathbf{t}}_v, \mathbf{y}_v)$ encourages $\tau(\cdot)$ to generate text representations that are beneficial for the GNN teacher. This

co-learning between the GNN teacher and the graph-free student leads to graph-aware text representations, as the two models learn from each other.

4.2 GraD-Alt

A common challenge in multi-task learning is negative transfer [48], in which the performance improvement of a specific task leads to performance degradation of other tasks. Our second strategy (GRAD-Alt) alleviates this issue, which might be present in GRAD-Joint, by optimizing the teacher and the student model in an alternate fashion. GRAD-Alt objective is given by

$$\mathcal{L}_{\text{T-step}} = \sum_{v \in \mathcal{V}^L} l_{\text{CE}}(\hat{\mathbf{t}}_v, \mathbf{y}_v), \quad (17)$$

$$\mathcal{L}_{\text{S-step}} = \lambda \sum_{v \in \mathcal{V}} l_{\text{KL}}(\hat{\mathbf{s}}_v, \hat{\mathbf{t}}_v) + \sum_{v \in \mathcal{V}^L} l_{\text{CE}}(\hat{\mathbf{s}}_v, \mathbf{y}_v), \quad (18)$$

where the teacher objective $\mathcal{L}_{\text{T-step}}$ and the student objective $\mathcal{L}_{\text{S-step}}$ alternately update function τ . Different from GRAD-Joint, decoupling S-step from the teacher gives more flexibility to the student to fit to its teacher’s predictions.

4.3 GraD-JKD

Solving Eq.(17) and Eq.(18) consists of many alternate optimization steps that update τ . This is associated with significant computational costs when τ is a large LM and when the input graph consists of a large number of nodes.

Our third strategy (GRAD-JKD) is motivated by the benefits discussed in Section 4.1 that lead to improving the GNN teacher as well as the benefits discussed in Section 4.2 that lead to improving the flexibility of the graph-free student. GRAD-JKD consists of two distinct stages, as summarized below

- **First Stage.** The GNN teacher is jointly trained with its graph-free student via objective $\mathcal{L}_{\text{Joint}}$; see Eq.(16).
- **Second Stage.** The GNN teacher is decoupled by its student, and the student is retrained alone to mimic the teacher’s predictions via

$$\mathcal{L}_{\text{KD}} = \lambda \sum_{v \in \mathcal{V}} l_{\text{KL}}(\hat{\mathbf{s}}_v, \hat{\mathbf{t}}_v) + \sum_{v \in \mathcal{V}^L} l_{\text{CE}}(\hat{\mathbf{s}}_v, \mathbf{y}_v). \quad (19)$$

The second stage is similar to conventional KD, which is why we term this strategy GRAD-JKD (Joint +KD). In our case, however, the GNN teacher has been trained with a graph-free student, which allows it to better leverage textual information from unlabeled nodes.

4.4 Student Models

A powerful function $\tau(\cdot)$, that can encode information from both the GNN and the graph-free model, is a key factor in the co-optimization of the teacher and

student model. We use LMs, such as BERT [5], to instantiate $\tau(\cdot)$, and name our method GRADBERT. GRADBERT is fine-tuned end-to-end, which is shown to benefit node classification tasks [16].

GRADBERT requires raw texts on nodes and assumes the graphs have adequate data for training. Although not common in practical scenarios, some graphs might be of small scale and not associated with raw texts. In such cases, we model $\tau(\cdot)$ by bag-of-words or TF-IDF text vectors (if provided), followed by trainable MLPs, which we name GRADMLP.

5 Experimental Setup

5.1 Datasets

For GRADBERT, we use three widely used node classification benchmarks that provide raw text on nodes, ogbn-arxiv (Arxiv), ogbn-products (Products), and a smaller version of ogbn-papers100M (Papers1.5M) [12,3]. We also use a private heterogeneous graph (Product-Reviews) in Appendix H, and *perturbed* datasets in Appendix I. For GRADMLP, we use Cora, Citeseer, Pubmed, A-Computer, and A-Photo [35], where input features are bag-of-words or TF-IDF word encodings. The statistics of these datasets are presented in Appendix A and the evaluation metric is accuracy (%).


5.2 Implementation Details

For our GNN teacher, we use GraphSAGE [10] for both the GRADBERT and the GRADMLP students. To further reduce the training computation cost for GRADBERT, we use an 1-layer GraphSAGE and sample $S \in \{8, 12\}$ neighbors for Eq.(4), which has a training cost of $\mathcal{O}(SC)$ per node. We initialize GRADBERT parameters with SciBERT [2] for Arxiv and Papers1.5M that provides better tokenization for scientific texts. Due to computational constraints, we use a smaller LM for Products, i.e., DistilBERT [26], and we only use GRAD-Alt for small-scale graphs. Further implementation details can be found in Appendix B. Ablation studies on GRAD’s framework are conducted in Appendix D.

5.3 Compared Methods

We develop the BERT+KD baseline, which employs conventional GNN-to-BERT KD via Eq.(7). GLNN [47] is a GNN-to-MLP distillation approach that does not leverage LMs, to which we refer as MLP+KD baseline. We also compare GRADBERT with methods that fine-tune LMs on graphs. Graph-Aware BERT (GA-BERT) [16] and BERT-LP [3] fine-tune BERT models to solve the link prediction task. GIANT [3] fine-tunes BERT to solve neighborhood prediction, which is a task similar to link prediction. E2EG [6] trains GIANT for neighborhood prediction and node classification end-to-end. GLEM [50] utilizes variational inference to jointly optimize the LM and the GNN for node classification. For a fair comparison, we have the same GNN architecture (GraphSAGE) among methods that use GNNs (GRAD, KD, GA-BERT, and GLEM) during training.

Table 1. Performance comparison of graph-free methods. Bold font denotes the overall best results. We also report the number of trainable parameters of each LM model in millions (#Params).

	Arxiv		Products		Papers1.5M	
	Acc. (%)	#Params	Acc. (%)	#Params	Acc. (%)	#Params
MLP	62.91 \pm 0.60	No LM	61.06 \pm 0.08	No LM	47.24 \pm 0.39	No LM
GLNN	72.15 \pm 0.27	No LM	77.65 \pm 0.48	No LM	-	
BERT	72.81 \pm 0.12	110M	77.64 \pm 0.08	110M	61.45 \pm 0.07	110M
BERT-LP	67.33 \pm 0.54	110M	73.83 \pm 0.06	110M	-	
GIANT	73.06 \pm 0.11	110M	80.49 \pm 0.28	110M	61.10 \pm 0.19	110M
E2EG	73.62 \pm 0.14	110M	81.11 \pm 0.37	110M	-	
GLEM	74.53 \pm 0.12	138M	81.25 \pm 0.15	138M	-	
BERT+KD	74.39 \pm 0.32	110M	<u>81.91</u> \pm 0.64	66M	61.85 \pm 0.04	110M
GRADBERT 						
Joint	<u>74.92</u> \pm 0.16	110M	81.42 \pm 0.40	66M	<u>63.44</u> \pm 0.05	110M
JKD	75.05 \pm 0.11	110M	82.89 \pm 0.07	66M	63.60 \pm 0.05	110M

6 Experimental Results

In the following experiments, we assess GRAD’s performance for node classification over textual graphs. In Section 6.1, we compare GRADBERT with other methods that do not use graph structure during inference, as well as we assess the performance of GRAD’s GNN teacher. In Section 6.2, we evaluate how well different GRAD strategies generalize to inductive (unseen) nodes. Furthermore, we conduct inference time analysis to demonstrate the efficiency advantage of GRAD in Appendix C and provide qualitative examples in Appendix E.

6.1 GradBERT Results

Table 1 shows performance results for methods that do not use graph structure during inference. Clearly, GRADBERT is the method that performs the best. MLP and GLNN are methods that do not fine-tune LMs, and thus, perform poorly. GIANT and BERT-LP pretrain BERT to encode structural information, but BERT is fixed for node classification, which may be suboptimal. For example, E2EG, that adapts GIANT for node classification, improves GIANT by 0.61% points on average. GLEM relies on powerful LMs to alternately fine-tune the teacher and the student models, and thus outperforms previous methods. However, GLEM does not show a clear advantage over our baseline BERT+KD method, while it utilizes 28M or 77M more parameters. GRADBERT outperforms BERT+KD by 0.66%, 0.98%, and 1.75% points for Arxiv, Products, and Papers, respectively, which is a considerable improvement for these large-scale graphs. In Products, GRAD-JKD improves over GRAD-Joint by 1.45% points, and shows that JKD allows the student model to better fit to its teacher’s predictions.

Table 2. GNN performance comparison for different methods that jointly fine-tune LMs for node classification. The underlying GNN model is GraphSAGE. We also report the number of trainable parameters of each backbone LM model in millions (#Params).

	Arxiv		Products		Papers1.5M	
	Acc. (%)	#Params	Acc. (%)	#Params	Acc. (%)	#Params
BERT-GNN	74.78 \pm 0.52	110M	82.01 \pm 0.43	110M	65.80 \pm 0.23	110M
GA-BERT-GNN	74.97	110M	82.35	110M	-	
GLEM-GNN	75.50 \pm 0.24	138M	83.16 \pm 0.19	138M	-	
GRAD-GNN	76.42 \pm 0.21	110M	83.34 \pm 0.24	66M	66.61 \pm 0.22	110M

Table 2 evaluates the performance of GNNs combined with LMs (LM-GNN methods). GRAD’s GNN teacher performs the best, which verifies that GRAD’s student improves its teacher’s effectiveness. BERT-GNN is the baseline LM-GNN with end-to-end training, which performs the worst. GA-BERT enhances BERT with a link prediction task and leads to better results than BERT-GNN. GLEM that leverages unlabeled nodes, as GRAD does, performs better than BERT and GA-BERT. However, it does not train the LM-GNN model end-to-end, which limits its performance compared to GRAD, while it requires larger LM models. In Appendix F, we also evaluate GRADBART with state-of-the-art downstream GNNs.

In transductive settings, graph-free students could perform well by imitating their GNN teachers’ predictions without learning effective features (Section 3.2). Thus, we evaluate our GRAD approach in the inductive setting, where well-trained students are essential to generalize to features of unseen nodes, which are not connected to the existing graph during training. Table 3 shows GRADBART’s performance in inductive settings. GRADBART outperforms the baseline BERT+KD by more than 0.7% points, on average. Note that MLP+KD uses a static BERT model and cannot generalize well to unseen nodes. As Table 3 shows, methods that use conventional KD (BERT+KD and GRAD-JKD) have a higher performance drop for inductive nodes. This suggests that conventional KD favors transductive settings while it might be limited for inductive nodes, which is also analyzed in Appendix G.

6.2 GraDMLP Results

Table 4 shows the performance of different GRAD strategies for small-scale graphs, where MLPs uses as the backbone share function τ . Methods that employ KD do not couple the teacher and student models sufficiently, and thus do not learn as effective node features in these small-scale graphs. GLEM performs the worst as it requires a powerful student model that generates high-quality soft-labels, which is not the case for simple MLPs. GRAD-Alt performs the best on Citeseer and A-Comp, while GRAD-Joint performs the best on Cora and Pubmed. We suspect that, in the second case, node texts are more informative (paper keywords) and thus, coupling the teacher and student models tightly improves the performance.

Table 3. Performance comparison of different distillation strategies in the inductive setting, in which we hold 50% of validation and test nodes out of the full graph. *Full* means that methods are trained and evaluated on the full graph, while *ind* shows performance on inductive nodes. Δ_{acc} reports the relative performance degradation between the two settings.

Dataset		MLP+KD	BERT+KD	GRADBERT	
				Joint	JKD
Arxiv	full	68.26	74.43	74.45	75.15
	ind	59.34	73.45	74.22	73.81
	Δ_{acc}	-13.06%	-1.32%	-0.31%	-1.78%
Products	full	77.20	82.31	81.47	82.82
	ind	72.93	80.94	81.24	81.69
	Δ_{acc}	-5.53%	-1.66%	-0.28%	-1.36%

Table 4. Performance comparison of different GRAD strategies on transductive (tran) and inductive (ind) nodes. We sample 50% nodes from the test set as the inductive nodes to evaluate every method. We report mean accuracy over 10 runs.

Dataset		GLNN	GLEM	GRADMLP (ours)		
		(MLP+KD)	(MLP)	Joint	Alt	JKD
Cora	tran	76.72	76.78	79.39	79.20	79.04
	ind	70.74	70.44	73.00	72.86	72.51
Citeseer	tran	69.55	69.06	70.33	70.96	68.69
	ind	69.72	69.33	70.38	70.80	69.28
Pubmed	tran	74.00	76.38	78.13	77.71	77.72
	ind	73.61	75.53	77.54	76.96	76.89
A-Comp	tran	82.12	81.15	81.20	82.60	82.28
	ind	79.11	78.07	78.79	80.09	79.41
A-Photo	tran	92.21	89.83	91.68	91.70	92.36
	ind	89.96	86.49	88.83	89.03	90.09

Moreover, Figure 4 shows that GRADMLP consistently outperforms MLP+KD with different inductive rates, ranging from 10% to 75%. In these small graphs, having an inductive ration close to 90% results in training graphs with few connections among nodes. Thus, in this case, GRADMLP does not have useful graph information to leverage.

7 Related Work

Learning on Textual Graphs. A common approach for effective learning on textual graphs is to combine LMs [5,22,25] with GNNs [19,31,10]. Methods such as [16,3,41] focus on pre-training LMs over graph data to preserve node-level structure, while [42] focuses on token-level information. Methods such as [50,6,38,45] focus on fine-tuning LMs for solving node classification directly, which is more related to our work. Experiments showed that GRAD outperforms these competing approaches. Apart from node classification, leveraging GNNs




Fig. 4. Performance on inductive nodes w.r.t. inductive node rate (inductive node rate equals to $\#ind$ test nodes / $\#total$ test nodes). GRADMLP consistently outperforms MLP+KD for reasonable inductive rates.

for learning on textual graphs has been applied to question answering [23,49], sponsored search [54,20], and document classification [40,14].

Distillation Approaches. Distillation approaches mainly focus on model compression (see a survey in [9]). It is worth to mention that [46] proposed a self-distillation technique that distills knowledge from deeper layers to shallow ones of the same architecture. Inspired by self-distillation, GRAD-Joint self-distills knowledge from deeper layers (GNN) to shallow layers (LM or MLP). However, our work is motivated by capturing node interactions, and not by model compression. Moreover, [43] shows a connection between self-distillation and label smoothing, which may be interpreted as a graph-aware label smoothing in our case. Closely related to our work, GLNN [47] and ColdBrew [52] propose to distill knowledge from a graph modality to a graph-less modality. However, experiments showed that GRAD learns better graph information than GLNN for both inductive and transductive settings, while ColdBrew can only be applied to transductive settings. Graph-regularized MLPs [53,1,37,13,7] are also methods that improve MLP performance for node classification. However, their performance is inferior to the one achieved by KD approaches. Other graph-based distillation approaches focus on distilling large GNNs to smaller GNNs [39,34] or to simpler graph models [36]. Finally, the work in [4] distills knowledge from pretrained GNNs and the work in [33] applies graph-based distillation for incremental learning in recommender systems.

8 Conclusion

In this paper, we developed a graph-aware distillation approach (GRAD) that jointly trains a teacher GNN and a graph-free student via a shared LM. This allows the two models to learn from each other and improve their overall performance. We have evaluated GRAD in eight node classification benchmarks in both transductive and inductive settings, in all of which GRAD outperforms conventional knowledge distillation. GRAD is a method that achieves a balance among efficiency and effectiveness in textual graphs.

9 Limitations & Ethical Statement

GRAD relies on informative input node features to learn effective shared LMs (or MLPs) that can generalize to unseen nodes, which is the case in textual graphs. Thus, one limitation is that it is not certain how GRAD generalizes to other graphs, e.g., to featureless graphs. Moreover as a knowledge distillation approach, GRAD trades accuracy for computation efficiency and it cannot adapt to dynamic graphs with edge changes the same way as GNN could. To overcome biases encoded in the training graph, e.g., standard stereotypes in recommender graphs, GRAD needs to be retrained over the new unbiased graph.

References

1. Ando, R., Zhang, T.: Learning on graph with laplacian regularization. NIPS (2006)
2. Beltagy, I., Lo, K., Cohan, A.: SciBERT: A pretrained language model for scientific text. In: EMNLP-IJCNLP (2019)
3. Chien, E., Chang, W.C., Hsieh, C.J., Yu, H.F., Zhang, J., Milenkovic, O., Dhillon, I.S.: Node feature extraction by self-supervised multi-scale neighborhood prediction. In: ICLR (2022)
4. Deng, X., Zhang, Z.: Graph-free knowledge distillation for graph neural networks. arXiv (2021)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: ACL (2019)
6. Dinh, T.A., Boef, J.d., Cornelisse, J., Groth, P.: E2EG: End-to-end node classification using graph topology and text-based node attributes. arXiv (2022)
7. Dong, W., Wu, J., Luo, Y., Ge, Z., Wang, P.: Node representation learning in graph via node-to-neighbourhood mutual information maximization. In: IEEE/CVF CVPR (2022)
8. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: ICML (2017)
9. Gou, J., Yu, B., Maybank, S.J., Tao, D.: Knowledge distillation: A survey. IJCV (2021)
10. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. NeurIPS (2017)
11. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. arXiv (2015)
12. Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., Leskovec, J.: Open graph benchmark: Datasets for machine learning on graphs. NeurIPS (2020)
13. Hu, Y., You, H., Wang, Z., Wang, Z., Zhou, E., Gao, Y.: Graph-mlp: Node classification without message passing in graph (2021)
14. Huang, L., Ma, D., Li, S., Zhang, X., Wang, H.: Text level graph neural network for text classification. EMNLP (2019)
15. Huang, Q., He, H., Singh, A., Lim, S.N., Benson, A.: Combining label propagation and simple models out-performs graph neural networks. In: ICLR (2021)
16. Ioannidis, V.N., Song, X., Zheng, D., Zhang, H., Ma, J., Xu, Y., Zeng, B., Chilimbi, T., Karypis, G.: Efficient and effective training of language and graph neural network models. arXiv (2022)
17. Jia, J., Benson, A.R.: Residual correlation in graph neural network regression. In: KDD (2020)

18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
19. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
20. Li, C., Pang, B., Liu, Y., Sun, H., Liu, Z., Xie, X., Yang, T., Cui, Y., Zhang, L., Zhang, Q.: Adsgnn: Behavior-graph augmented relevance modeling in sponsored search. In: ACM SIGIR (2021)
21. Li, G., Müller, M., Ghanem, B., Koltun, V.: Training graph neural networks with 1000 layers. In: ICML (2021)
22. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv (2019)
23. Mavromatis, C., Karypis, G.: ReaRev: Adaptive reasoning for question answering over knowledge graphs. arXiv (2022)
24. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. NIPS (2013)
25. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. JMLR (2020)
26. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv (2019)
27. Schlichtkrull, M., Kipf, T.N., Bloem, P., Berg, R.v.d., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: ESWC (2018)
28. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. JMLR (2014)
29. Sun, C., Wu, G.: Scalable and adaptive graph neural networks with self-label-enhanced training. arXiv (2021)
30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NIPS (2017)
31. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. In: ICLR (2018)
32. Wang, M., Zheng, D., Ye, Z., Gan, Q., Li, M., Song, X., Zhou, J., Ma, C., Yu, L., Gai, Y., et al.: Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv (2019)
33. Xu, Y., Zhang, Y., Guo, W., Guo, H., Tang, R., Coates, M.: Graphsail: Graph structure aware incremental learning for recommender systems. In: CIKM (2020)
34. Yan, B., Wang, C., Guo, G., Lou, Y.: Tinygnn: Learning efficient graph neural networks. In: KDD (2020)
35. Yang, C., Liu, J., Shi, C.: Extract the knowledge of graph neural networks and go beyond it: An effective knowledge distillation framework. In: WWW (2021)
36. Yang, C., Liu, J., Shi, C.: Extract the knowledge of graph neural networks and go beyond it: An effective knowledge distillation framework. In: WWW (2021)
37. Yang, H., Ma, K., Cheng, J.: Rethinking graph regularization for graph neural networks. In: AAAI (2021)
38. Yang, J., Liu, Z., Xiao, S., Li, C., Lian, D., Agrawal, S., Singh, A., Sun, G., Xie, X.: Graphformers: Gnn-nested transformers for representation learning on textual graph. NeurIPS (2021)
39. Yang, Y., Qiu, J., Song, M., Tao, D., Wang, X.: Distilling knowledge from graph convolutional networks. In: IEEE/CVF CVPR (2020)
40. Yao, L., Mao, C., Luo, Y.: Graph convolutional networks for text classification. In: AAAI (2019)

41. Yasunaga, M., Bosselut, A., Ren, H., Zhang, X., Manning, C.D., Liang, P., Leskovec, J.: Deep bidirectional language-knowledge graph pretraining. *NeurIPS* (2022)
42. Yasunaga, M., Leskovec, J., Liang, P.: Linkbert: Pretraining language models with document links. *ACL* (2022)
43. Yuan, L., Tay, F.E., Li, G., Wang, T., Feng, J.: Revisiting knowledge distillation via label smoothing regularization. In: *IEEE/CVF CVPR* (2020)
44. Zeng, H., Zhou, H., Srivastava, A., Kannan, R., Prasanna, V.: Graphsaint: Graph sampling based inductive learning method. *arXiv* (2019)
45. Zhang, J., Zhang, H., Xia, C., Sun, L.: Graph-bert: Only attention is needed for learning graph representations. *arXiv* (2020)
46. Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: *IEEE/CVF ICCV* (2019)
47. Zhang, S., Liu, Y., Sun, Y., Shah, N.: Graph-less neural networks: Teaching old MLPs new tricks via distillation. In: *ICLR* (2022)
48. Zhang, W., Deng, L., Zhang, L., Wu, D.: A survey on negative transfer. *arXiv* (2020)
49. Zhang, X., Bosselut, A., Yasunaga, M., Ren, H., Liang, P., Manning, C.D., Leskovec, J.: GreaseLM: Graph REASoning enhanced language models. In: *ICLR* (2022)
50. Zhao, J., Qu, M., Li, C., Yan, H., Liu, Q., Li, R., Xie, X., Tang, J.: Learning on large-scale text-attributed graphs via variational inference. *arXiv* (2022)
51. Zheng, D., Ma, C., Wang, M., Zhou, J., Su, Q., Song, X., Gan, Q., Zhang, Z., Karypis, G.: Distdgl: distributed graph neural network training for billion-scale graphs. In: *2020 IEEE/ACM 10th Workshop on Irregular Applications: Architectures and Algorithms (IA3)* (2020)
52. Zheng, W., Huang, E.W., Rao, N., Katariya, S., Wang, Z., Subbian, K.: Cold brew: Distilling graph node representations with incomplete or missing neighborhoods. *ICLR* (2022)
53. Zhou, D., Bousquet, O., Lal, T., Weston, J., Schölkopf, B.: Learning with local and global consistency. *NIPS* (2003)
54. Zhu, J., Cui, Y., Liu, Y., Sun, H., Li, X., Pelger, M., Yang, T., Zhang, L., Zhang, R., Zhao, H.: Textgnn: Improving text encoder via graph neural network in sponsored search. In: *WWW* (2021)

Table 5. Statistics of the datasets.

Text Dataset	<i>raw text</i>				<i>bag-of-words / TF-IDF</i>				
	Arxiv	Products	Papers 1.5M	Product-Reviews	Cora	Citeseer	Pubmed	A-Comp	A-Photo
# Nodes	169,343	2,449,029	1,546,782	P: 2,696,062 R: 2,533,129	2,485	2,110	19,717	13,381	7,487
# Edges	1,166,243	2,449,029	1,546,782	2,096,202	5,069	3,668	44,324	245,778	119,043
Avg. Degree	13.7	50.5	17.7	0.7	2.0	1.74	2.3	18.4	15.9
# Classes	40	47	172	1,810	7	6	3	10	8
Train / Val / Test	54 / 18 / 28 (%)	8 / 2 / 90 (%)	78 / 8 / 14 (%)	Variable	#20 per class / #30 per class / rest				

A Dataset Statistics

The datasets statistics are shown in Table A. Arxiv, Papers, Cora, Citeseer, and Pubmed are citations graphs, while Products, A-Computer, and A-Photo are product co-purchase graphs. Product-Reviews is a heterogeneous graph.

B Further Implementation Details

For our teacher GNN, we use GraphSAGE [10] for both GRADBERT and GRADMLP. Our code is available at: <https://github.com/cmavro/GRAD>.

B.1 GraDBERT

For GRADBERT, training requirements are demanding, so we implement it with DGL’s [32]⁶ distributed training [51] and perform training on a p4d.24xlarge instance. To further reduce the training computation cost, we use an 1-layer GraphSAGE for Eq.(4), which has a training cost of $\mathcal{O}(SC)$ per node. We sample $S = 8$ neighbors for Arxiv and Papers, and $S = 12$ neighbors for Products. During training, we use a sequence length $L = 128$, and during inference, we use either set $L = 128$ for ablation studies or $L = 512$ for SOTA performance. We initialize GRADBERT parameters with SciBERT [2] for Arxiv and Papers1.5M and with DistilBERT [26] for Products. For Product-Reviews, we pretrain BERT over the given product descriptions.

GRAD-Joint’s hyper-parameter α is tuned amongst $\{0.2, 0.4\}$ and hyper-parameter λ amongst $\{20, 50\}$. Due to demanding computational requirements, we do not perform extensive hyper-parameter search. We set the number of epochs to 10 and evaluate on the validation set every epoch. We perform model selection based on the best validation scores. We perform mini-batch training to maximize GPU utilization, where the batch size is tuned amongst $\{16, 32, 48\}$. For inference, we use a batch size of 128. We optimize the model parameters with Adam [18] and tuning the learning rate for BERT for $\{1e^{-4}, 1e^{-5}, 1e^{-6}\}$ for the GNN amongst $\{1e^{-3}, 1e^{-4}, 1e^{-5}\}$. We set the number of hidden dimensions to $d = 128$. When combining GRADBERT with KD (GRAD-JKD), we perform 5 maximum epochs of KD with the same hyper-parameters and perform model selection based on the validation score.

⁶ <https://www.dgl.ai/>

B.2 GraDMLP

For a fair comparison to GLNN (MLP+KD), we use the official GLNN implementation and experimental setup⁷. Similar to GLNN, we use a 2-layer GraphSAGE teacher, where $S = 5$ neighbor nodes are sampled at each layer.

Hyper-parameter α of GRAD objective is tuned amongst $\{0.2, 0.4\}$ and hyper-parameter λ amongst $\{20, 50\}$. The number of epochs is tuned amongst values of $\{100, 200, 400, 1000, 2000\}$ with early stopping. We use Adam optimizer with learning rate tuned amongst $\{1e^{-4}, 1e^{-5}\}$, weight decay tuned amongst $\{5e^{-5}, 5e^{-6}\}$ and dropout [28]. We perform model selection based on the validation scores. When combining GRADMLP with KD, we tune α amongst $\{0.4, 0.6, 0.8\}$, which gives more importance to the GNN teacher. Hidden dimensions are set to $d = 128$, and we have identical MLP architectures with GLNN. We perform full-graph training. For GRAD-Alt we perform T-S steps at every optimization step of the full-graph training (1 TS cycle for every target node at each epoch).

C Inference Timing Analysis

We provide inference time analysis between GRADBERT and GNNs on two different machines (p4d.24xlarge and p3.16xlarge) for ogbn-arxiv. The p4d.24xlarge is estimated to be 2.5x faster than the p3.16xlarge, while it has 2.5x more GPU memory. As Figure 5 shows, GRADBERT is the method that that best balances efficiency and effectiveness.




Fig. 5. **GraDBERT** is on the Pareto front, being the method that best balances efficiency and effectiveness: Accuracy performance for ogbn-arxiv w.r.t. inference time. L is the input sequence length and OOM (out of memory) means that the model encountered GPU failure.

⁷ <https://github.com/snap-research/graphless-neural-networks>

D Ablation Studies

D.1 Ablation Study on GraD’s Loss

We select GRAD-Joint and perform an ablation study on how the distillation term of Eq.(16), which is controlled by hyper-parameter λ , impacts teacher-student performance. Table 6 verifies that KD loss benefits both the teacher and the student (GRAD row). The most significant benefit is observed for the student, which indicates that GNNs distill useful graph-aware information.

Table 6. KD loss effect in Eq.(16) for Arxiv dataset.

Model	Configuration	Arxiv accuracy
Student	$\alpha = 0, \lambda = 0$	72.98
	$\alpha = 0.2, \lambda = 0$	72.76
	GRAD	74.92
Teacher	$\alpha = 1, \lambda = 0$	74.78
	$\alpha = 0.8, \lambda = 0$	74.94
	GRAD	75.78

D.2 Hyper-Parameter Sensitivity Analysis

We continue with more extensive ablation studies for Products dataset (which is the dataset with the largest amount of unlabeled nodes). Table 7 show the results with $\alpha \in \{0.2, 0.8\}$ and $\lambda \in \{20, 50\}$. In general, due to the excessive amount of unlabeled nodes compared to labeled ones, higher λ might overfits to poor GNN predictions. As α controls the importance of each model, $\alpha = 0.2$ leads to a better student, and $\alpha = 0.8$ leads to a better teacher.

E Qualitative Examples

Figure 6 illustrates two examples from the Arxiv dataset in which GRADBERT and BERT+KD have different predictions. In the first example, the phrase “semidefinite programming” contained in the target text makes BERT+KD incorrectly predict *Discrete Math* as a node label, and in the second, the phrase “query service” makes BERT+KD incorrectly predict *Databases*. On the other hand, GRADBERT better leverages the graph information to decide on what label to predict. In the first example, the majority of neighbors of the target node are *Algorithm* papers and, in the second, *Security* papers. In both cases, GRADBERT predicts the correct label.

Table 7. KD loss effect in Eq.(16) for Arxiv dataset.

Model	Configuration	Products accuracy
Student	$\alpha = 0.2, \lambda = 20$	81.42
	$\alpha = 0.8, \lambda = 20$	80.91
	$\alpha = 0.2, \lambda = 50$	81.12
	$\alpha = 0.8, \lambda = 50$	81.25
Teacher	$\alpha = 0.2, \lambda = 20$	82.45
	$\alpha = 0.8, \lambda = 20$	83.34
	$\alpha = 0.2, \lambda = 50$	82.38
	$\alpha = 0.8, \lambda = 50$	83.12

F SOTA GNN Performance

We also assess the quality of GRADBERT’s learned node representations when used as input *label-aware* node features for other downstream GNNs. Table 8 shows that GRADBERT improves two successful GNNs, RevGAT and SAGN+, by 0.20%, on average. Fine-tuning GRADBERT performs better than static LMs (OGB-feat) by 3.7% points, on average. For simpler GNNs, i.e., GraphSAGE and its mini-batch version GraphSAINT [44], GRADBERT performs better than existing approaches by 1%-1.4% points. Products dataset has a very small label rate, thus we combine GRADBERT with GIANT (GIANT+GRADBERT), which is an unsupervised learning method.

G Failing KD Case

Corollary 2. *Given perfect GNN predictions $\hat{\mathbf{t}}_v$ obtained via Eq.(6), the graph-free student can achieve perfect predictions on transductive nodes by setting $X_v := \hat{\mathbf{t}}_v$ and function $\tau'(\cdot)$ to the identity matrix.*

Proof. The proof is straightforward. Perfect GNN predictions $\hat{\mathbf{t}}_v$ mean that the predictions obtained by

$$\begin{aligned} \hat{\mathbf{t}}_v &= f(\tau(X_u) : u \in \mathcal{G}_v), \\ \text{where } f, \tau &= \arg \max_{\tilde{f}, \tilde{\tau}} \sum_{v \in \mathcal{V}^L} I_{\tilde{f}, \tilde{\tau}}(\mathbf{y}_v; \mathcal{G}_v), \end{aligned} \quad (20)$$

maximize the mutual information over the unlabeled nodes

$$\hat{\mathbf{t}}_v = \arg \max_{\tilde{\mathbf{t}}_v} \sum_{v \in \mathcal{V}^U} I(\mathbf{y}_v, \mathcal{G}_v). \quad (21)$$

With $X_v := \hat{\mathbf{t}}_v$, Eq.(11) becomes

$$I(\mathbf{y}_v; \mathcal{G}_v) = I(\mathbf{y}_v; \hat{\mathbf{t}}_v) + I(\mathbf{y}_v; \tilde{\mathcal{X}}_v | \hat{\mathbf{t}}_v) + I(\mathbf{y}_v; \mathcal{E}_v | \mathcal{X}_v), \quad (22)$$

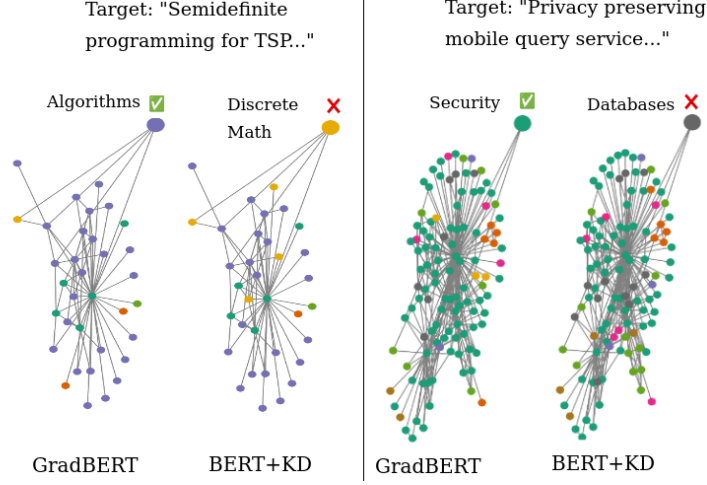


Fig. 6. Qualitative examples in which GRADBERT outperforms conventional KD. Colors denote node labels. GRADBERT leverages label-aware neighbor information to infer the correct label for the target text.

which can be further written as

$$I(\mathbf{y}_v; \mathcal{G}_v) = I(\mathbf{y}_v; \hat{\mathbf{t}}_v) + I(\mathbf{y}_v; \tilde{\mathcal{X}}_v | \hat{\mathbf{t}}_v) + I(\mathbf{y}_v; \mathcal{E}_v | (\tilde{\mathcal{X}}_v, \hat{\mathbf{t}}_v)), \quad (23)$$

Because all structural label-aware information is encoded into $\hat{\mathbf{t}}_v$, \mathcal{E}_v and $\tilde{\mathcal{X}}_v$ do not provide additional information for \mathbf{y}_v . Thus, \mathbf{y}_v is conditionally independent given $\hat{\mathbf{t}}_v$, which gives $I(\mathbf{y}_v; \tilde{\mathcal{X}}_v | \hat{\mathbf{t}}_v) = 0$ and $I(\mathbf{y}_v; \mathcal{E}_v | (\tilde{\mathcal{X}}_v, \hat{\mathbf{t}}_v)) = 0$, due to the mutual information definition. As a result, we have

$$I(\mathbf{y}_v; \mathcal{G}_v) = I(\mathbf{y}_v; \hat{\mathbf{t}}_v). \quad (24)$$

The graph-free student solves

$$\max_{\tau} I_{\tau}(\mathbf{y}_v; \tau(\hat{\mathbf{t}}_v)). \quad (25)$$

By setting τ to the identity matrix we have

$$I(\mathbf{y}_v; \hat{\mathbf{t}}_v), \quad (26)$$

which is a solution to

$$\max I(\mathbf{y}_v, \mathcal{G}_v) \quad (27)$$

as shown in Eq.(21). Thus, the graph-free student achieves perfect predictions for nodes that $\hat{\mathbf{t}}_v$ is given without the need to optimize $I(\mathbf{y}_v; \hat{\mathbf{t}}_v)$.

This is, however, a failing case as, for inductive nodes, GNN predictions are not given and the graph-free will only be able to perform random guesses (τ is an identity matrix without any learned parameters).

Table 8. Performance comparison of GNN methods in transductive settings w.r.t. the input text representations (*LM* column).

Dataset	LM	GNN	
		GraphSAGE	RevGAT
Arxiv	OGB-feat	71.49 \pm 0.27	74.26 \pm 0.17
	BERT	73.42	-
	GA-BERT	74.97	-
	GIANT	74.59 \pm 0.28	76.12 \pm 0.16
	GLEM	75.50 \pm 0.24	76.97 \pm 0.19
	GRADBERT 🏆	76.50\pm0.07	77.21\pm0.31
		GraphSAINT	SAGN+
Products	OGB-feat	80.27 \pm 0.26	84.85 \pm 0.10
	BERT	81.24	-
	GA-BERT	82.35	-
	GLEM	83.16 \pm 0.19	-
	GIANT	84.15 \pm 0.15	86.73 \pm 0.08
	+GRADBERT 🏆	85.51\pm0.51	86.92\pm0.07

RevGAT [21] is a successful GNN for ogbn-arxiv.

SAGN+, i.e., SAGN+SLE+C&S [29,15], is a successful GNN for ogbn-products.

H Product-Reviews Dataset

One major advantage of GRADBERT over simply fine-tuning BERT is that it leverages unlabeled nodes. However, in current datasets, these nodes are directly associated with the hidden labels. On the other hand, Product-Reviews dataset is a heterogeneous graph (Product-Reviews), where there are two different node types; product descriptions and product reviews. Product-nodes have review-nodes as their side information (product-nodes are not directly connected with each other). In this setting, we test whether GRADBERT can leverage this side information (review-nodes) to better predict labels on the target node type (product-nodes). We perform inductive learning, so that BERT and GRADBERT use the same product-nodes during training. As Table 9 shows, GRADBERT is able to improve BERT by classifying correctly 1,968-7,926 more products, depending on the data split. This is a considerable improvement although each product has poor graph information with an average degree of 0.7; see Table 5.

I Perturbed Node Texts

We study the impact of jointly learning the GNN teacher and graph-free student in a case where the graph structure is important while learning $\tau(\cdot)$. Following our insights from Section 3.2 that KD may simply copy GNN predictions, we develop *perturbed* datasets in which node texts cannot be easily be mapped

Table 9. Improvement of GRADBERT over BERT for Product-Reviews. GRADBERT leverages the additional side graph information to improve performance on product nodes.

Split (%)	Method	Side Info	Improvement
30/10/60	BERT	None	–
	GRADBERT	Reviews	#7,926 Products
80/10/10	BERT	None	–
	GRADBERT	Reviews	#1,968 Products

to their soft-labels. Specifically for each node of ogbn-arxiv and ogbn-products graphs, we keep half of its text the same as its original text, but replace the other half by some noisy text, that is sampled from a random node. For methods that do not use the graph during inference, it is very challenging to discriminate whether useful information comes from the original or and the noisy text parts. For methods that use structure, i.e., GNNs, the majority of the node’s neighbors will contain similar semantics with the original text, making the task much easier.

Table 10 shows GRADBERT’s performance with different α values, which control the importance of the GNN term in Eq.(16). As we increase α , GRADBERT performs better, which suggests that GRADBERT learns more useful structural information. On the other hand, BERT+KD performs poorly on this task, as it does not use any graph structure while learning $\tau(\cdot)$. We hypothesize this happens because BERT+KD cannot learn an accurate mapping $\tau(X_v) \rightarrow \hat{t}_v$ from text to labels, as analyzed in Section 3.2.

Table 10. Performance comparison of GRADBERT for perturbed datasets w.r.t. different α values (graph importance in Eq.(16)).

	Arxiv-Perturbed		Products-Perturbed	
	Val / Test		Val / Test	
<hr/>				
GRADBERT				
with $\alpha = 0.8$	49.04	47.35	87.06	63.21
with $\alpha = 0.2$	47.32	45.11	83.45	61.14
<hr/>				
BERT+KD	41.75	38.66	71.61	62.63

BERT+KD is equivalent to setting $\alpha \rightarrow 0$.