A ROBUST STACKING FRAMEWORK FOR DEEP GRAPH NETWORKS WITH MULTIFACETED NODE FEATURES

Anonymous authors
Paper under double-blind review

ABSTRACT

Graph Neural Networks (GNNs) with numerical node features and graph structure as inputs have demonstrated superior performance on various semi-supervised learning tasks with graph data. However, the numerical node features utilized by GNNs are commonly extracted from raw data which is of text or tabular (numeric/categorical) type in most real-world applications. The best models for such data types in most standard supervised learning settings with IID (non-graph) data are not simple neural network layers and thus are not easily incorporated into a GNN. Here we propose a robust stacking framework that fuses graph-aware propagation with arbitrary models intended for IID data, which are ensembled and stacked in multiple layers. Our layer-wise framework leverages bagging and stacking strategies to enjoy strong generalization, in a manner which effectively mitigates label leakage and overfitting. Across a variety of graph datasets with tabular/text node features, our method achieves comparable or superior performance relative to both tabular/text and graph neural network models, as well as existing state-of-the-art hybrid strategies that combine the two.

1 INTRODUCTION

Graph datasets comprise nodes of various data types and modalities linked by edges that encapsulate non-IID conditional dependencies between them. While it is often assumed that graph neural networks (GNN) (Kipf & Welling, 2016; Veličković et al., 2017) are preferable for handling such data relative to models originally designed for IID instances, GNNs are nonetheless subject to various limitations. In particular, the best architecture may be data-set specific and require appropriately setting many attendant structural hyperparameters, e.g., note the complex assortment of GNN architectures that populate the top of the Open Graph Benchmark (OGB) leaderboard (Hu et al., 2020). Moreover, most GNNs implicitly assume that node features are numerical, and may struggle to remain competitive with more complex text, tabular, or composite alternatives.

In fact, with richer node feature sets it has even been observed that models tailored to IID data (which in our setting simply operate on individual node features as though they were independent of the others) can at times outperform GNNs if they are combined with simple graph propagation operations to account for the graph structure (Huang et al., 2020; Chen et al., 2021). Moreover, for graph data with text features, Chien et al. (2021) has demonstrated that leveraging a BERT Transformer in addition to a GNN can greatly improve performance. And beyond these considerations, real-world applications of ML typically involve more than just a single model, GNN or otherwise. Instead they usually require an ML pipeline composed of data preprocessing and training/tuning/aggregation of many models to achieve the best results.

In this paper, we investigate how to adapt ML pipelines designed for supervised learning with IID data (e.g., Transformers for text, gradient boosted decision trees or related for tabular data) to node classification/regression tasks with graph-structured statistical dependencies between node features. We focus on using k-fold bagging (Breiman, 1996), i.e. *cross-validation*, to avoid label leakage issues, with stack ensembling methods for maximal flexibility (Wolpert, 1992; Van der Laan et al., 2007). These techniques are particularly effective for achieving high accuracy across diverse IID datasets, and are utilized in many popular AutoML frameworks (Erickson et al., 2020; LeDell & Poirier, 2020; Feurer et al., 2015), but have largely been ignored within the context of graph data.

Within this context, our goal is to design a single architecture that integrates graph propagation or message passing steps and stacked ensembles of arbitrary base models to flexibly accommodate diverse node/instance types within a unified framework. In doing so, our contributions are as follows:

- We propose a framework of stack ensembling with graph propagation called **BestowGNN** for Bagged, Ensembled, Stacked Training Of Well-balanced GNNs (see Figure 1) that can *bestow* arbitrary (non-graph) base models intended for IID data with the capability of producing highly accurate node predictions in the graph (i.e., non-IID) setting.
- Using only a single, unified architecture, our proposed methodology can match or outperform bespoke dataset-specific models that top competitive leaderboards for popular node classification/regression tasks (e.g., on OGB and elsewhere completely different network architectures typically dominate the top positions for different datasets and data types).
- Label leakage is an unavoidable issue for many layer-wise training strategies (SAGN (Sun & Wu, 2021) and GAMLP (Zhang et al., 2021)). To address this potential shortcoming, we formalize how our bagging and stacking framework can effectively mitigate the label leakage issue within the graph setting using analytical tools from differential privacy. This is the first work establishing that bagging with graph-based predictors can be useful for ameliorating label leakage.

2 Related Work

2.1 From Scalability to Layer-wise Training

Currently, GNN training suffers from high computational cost with the number of layers growing. To improve the scalability of GNNs, graph sampling scheme GraphSAGE (Hamilton et al., 2017) is adopted by uniformly sampling a fixed number of neighbours for a batch of nodes. Cluster-GCN (Chiang et al., 2019) uses graph clustering algorithms to sample a block of nodes that form a dense subgraph and runs SGD-based algorithms on these subgraphs. L^2-GCN (You et al., 2020) proposes a layer-wise training framework by disentangling feature aggregation and feature transformation to reduce time and memory complexity.

SAGN (Sun & Wu, 2021) iteratively trains models in several stages by applying graph structure-aware attention mechanisms on node features and also combines the self-training approach with label propagation to further improve performance. GAMLP (Zhang et al., 2021) proposes two attention mechanisms to explore the relation between features with different propagation steps. Both SAGN and GAMLP achieve state-of-the-art performance on two large open graph benchmarks (ogbn-products and ogbn-papers100M), demonstrating the high scalability and efficiency of layer-wise training strategies. However, SAGN and GAMLP suffer from the risk of label leakage: label information is included in the enhanced training set, and can cause performance degredation if the model extracts and relies on these labels. SAGN empirically shows that enough propagation depth can effectively alleviate label leakage, thus they only use label information at one fixed propagation step. Meanwhile, GAMLP passes label information between propagation steps using residual connections. Wang et al. (2021) further randomly masks nodes during every training epoch to mitigate label leakage issue.

2.2 Graph models with Multifaceted Node Features

Traditional GNN models are mostly studied for graphs with homogeneous sparse node features. Leading GNN models fail to achieve competitive results for heterogeneous features with tabular or text node features (Ivanov & Prokhorenkova, 2021; Huang et al., 2020; Chen et al., 2021). To remedy this, Ivanov & Prokhorenkova (2021) jointly train Gradient Boosted Decision Trees (GBDT) and GNN in an end-to-end fashion, demonstrating a significant increase in performance on graph data with tabular node features.

Chen et al. (2021) removes the need for a GNN altogether, proposing a generalized framework for iterating boosting with parameter-free graph propagation steps that share node/sample information across edges connecting related samples.

Correct and Smooth (C&S) (Huang et al., 2020) is a simple post-processing step that applies label propagation to further incorporate graph information into the outputs of a learning algorithm. Chen

et al. (2021) trains Gradient Boosted Decision Trees with label propagation incorporated into the objective function, producing competitive results for graph data with tabular node features.

Because common GNNs take numerical node features as inputs, one must establish a way to extract numerical embeddings from raw data such as text and images. For example, the embeddings of ogbn-arxiv data are computed by running the skip-gram model (Mikolov et al., 2013). Chien et al. (2021) proposes self-supervised learning to fully utilizing correlations between graph nodes, and extracts the embedding of three open graph benchmark datasets (ogbn-arxiv, ogbn-products and ogbn-papers100M). Chien et al. (2021) demonstrates the superior performance of these new embeddings for the Open Graph Benchmark datasets. Lin et al. (2021) proposes BertGCN, which combines the Bert model and transductive learning for text classification in an end-to-end fashion and achieves superior performance on a range of text classification tasks.

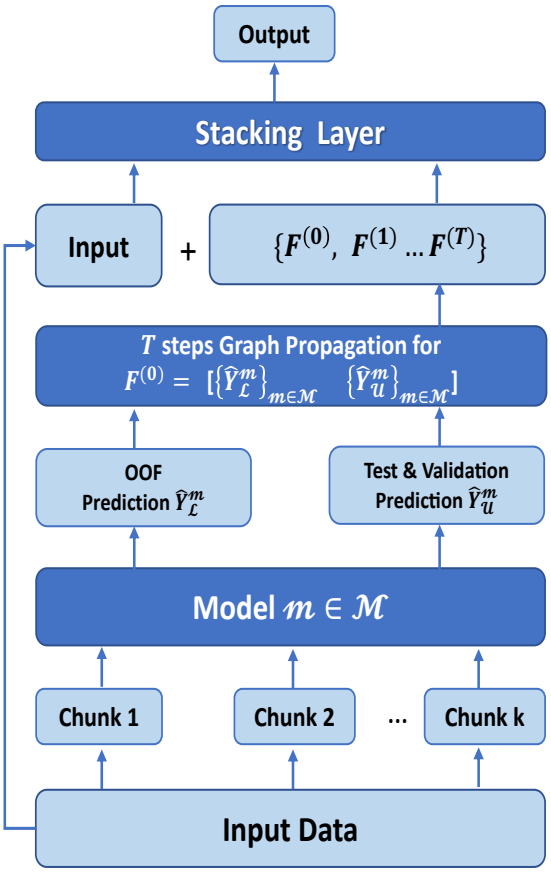

Figure 1: BestowGNN with a single base learner, 2 stacking layers, and k-fold bagging (repeated bagging not depicted here). The stacking layer repeats the operations depicted between it and the input data.

Algorithm 1 BestowGNN Training Strategy

Input: Node features and labels $(\boldsymbol{X}, \boldsymbol{Y})$ from graph $\mathcal{G}$ with labeled (training) nodes $\mathcal{L}$ and unlabeled (validation/test) nodes $\mathcal{U}$, family of models intended for IID data $\mathcal{M}$, L stacking layers, n-repeated k-fold bagging, T propagation steps.
for $l = 1$ **to** L **do** {stacking}
  for $i = 1$ **to** n **do** {repeated bagging}
    Randomly split data into k chunks $\{\boldsymbol{X}^j, \boldsymbol{Y}^j\}_{j=1}^k$
    for $j = 1$ **to** k **do**
      Train model $m \in \mathcal{M}$ on $\{\boldsymbol{X}^{-j}, \boldsymbol{Y}^{-j}\}$
      Make predictions $\hat{\boldsymbol{Y}}^j_{m,i}$ on OOF data $\boldsymbol{X}^j$
    end for
  end for
  for $m \in \mathcal{M}$ **do**
    Get OOF predictions $\hat{\boldsymbol{Y}}^m_{\mathcal{L}}$ for labeled nodes via (6)
    Get predictions $\hat{\boldsymbol{Y}}^m_{\mathcal{U}}$ for unlabeled nodes via (7)
  end for
  Concatenate all models' predictions:
  $\boldsymbol{F}^{(0)} \triangleq [\{\hat{\boldsymbol{Y}}^m_{\mathcal{L}}\}_{m\in\mathcal{M}}, \{\hat{\boldsymbol{Y}}^m_{\mathcal{U}}\}_{m\in\mathcal{M}}]$
  for $t = 0$ **to** T **do** {propagation}
    Compute $\boldsymbol{F}^{(t)}$ using (4)
  end for
  $\boldsymbol{X} \leftarrow$ concatenate $(\boldsymbol{X}, \{\boldsymbol{F}^{(0)}, ..., \boldsymbol{F}^{(T)}\})$
end for
Output: weighted prediction $\sum_{m\in\mathcal{M}} \alpha_m \hat{\boldsymbol{Y}}^m_{\mathcal{U}}$ with $\{\alpha_m\}$ fitted via Ensemble Selection

3 Background

Consider an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $n = |\mathcal{V}|$ nodes. The node feature matrix is denoted by $\boldsymbol{X} \in \mathbb{R}^{n\times d}$, and the corresponding node label matrix is $\boldsymbol{Y} \in \mathbb{R}^{n\times c}$ with d and c being the dimension of features and labels respectively. The unweighted adjacency matrix is $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. For training purposes we only have access to the labels of a subset of nodes $\{\boldsymbol{y}_i\}_{i\in\mathcal{L}}$, with $\mathcal{L} \subset \mathcal{V}$. Given feature values of all nodes $\{\boldsymbol{x}_i\}_{i\in\mathcal{V}}$, label data $\{\boldsymbol{y}_i\}_{i\in\mathcal{L}}$, and the connectivity of the graph $\mathcal{E}$, the task is to

predict the labels of the unlabeled nodes $\{\boldsymbol{y}_i\}_{i\in\mathcal{U}}$, with $\mathcal{U} = \mathcal{V} \setminus \mathcal{L}$. We denote the labeled dataset $\{\boldsymbol{x}_i, \boldsymbol{y}_i\}_{i\in\mathcal{L}}$ as $D_{\mathcal{L}}$ and the unlabeled dataset $\{\boldsymbol{x}_i\}_{i\in\mathcal{U}}$ as $D_{\mathcal{U}}$.

3.1 Bagging, Ensembling, and Stacking

For classification/regression with IID (non-graph) data, bagging, ensembling, and stacking represent practical tools that can be combined in various ways to produce more accurate predictions relative to other strategies across diverse tabular and text datasets (Shi et al., 2021; Blohm et al., 2020; Yoo et al., 2020; Fakoor et al., 2020; Bezrukavnikov & Linder, 2021; Feldman, 2021). For example, in each stacking layer of an ensemble-based architecture, bagging simply trains the same types of base models with out-of-fold predictions from the previous layer models (obtained via bagging) as extra predictive features. These base models might include various Gradient Boosted Decision Trees (Ke et al., 2017; Prokhorenkova et al., 2018), fully-connected neural networks (MLP), K Nearest Neighbors (Erickson et al., 2020), or pretrained Electra Transformer models (Clark et al., 2020). For our purposes herein, we adopt the AutoML package AutoGluon (Erickson et al., 2020), which is capable of exploiting these techniques while serving open-source code that we can readily adapt to include graph propagation.

3.2 Graph-Aware Propagation Layers

Recently there has been a surge of interest GNN architectures with layers defined in one-to-one correspondence with descent iterations that minimize a principled class of graph-regularized energy functions (Klicpera et al., 2018; Ma et al., 2020; Pan et al., 2021; Yang et al., 2021; Zhang et al., 2020; Zhu et al., 2021). IN this way GNN models can benefit from the inductive bias afforded by energy function minimizers (or close approximations thereof) whose specific form can be controlled by trainable parameters.

Following Zhou et al. (2004), one relevant energy function capable of inducing such graph-aware propagation is given by

$$\ell_Y(\boldsymbol{Y}) \triangleq (1-\lambda)\left\|\boldsymbol{Y} - m\left(\boldsymbol{X};\boldsymbol{\theta}\right)\right\|_{\mathcal{F}}^2 + \lambda \text{tr}\left[\boldsymbol{Y}^\top \boldsymbol{L}\boldsymbol{Y}\right], \tag{1}$$

where $\lambda \in (0,1)$ is a weight that determines the trade-off between the two terms. $\boldsymbol{Y} \in \mathbb{R}^{n\times d}$ is a learnable d-dimensional embedding across n nodes, and $m\left(\boldsymbol{X};\boldsymbol{\theta}\right)$ denotes a base model (parameterized by $\boldsymbol{\theta}$) that computes an initial target embedding based on the node features $\boldsymbol{X}$. $\boldsymbol{L} \in \mathbb{R}^{n\times n}$ is the graph Laplacian of $\mathcal{G}$, meaning $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A}$, where $\boldsymbol{D}$ represents the degree matrix.

Intuitively, the first term of (1) encourages $\boldsymbol{Y}$ to be close to initial target embedding, while the second term introduces the smoothness over the whole graph. On the positive side, the closed-form optimal solution of energy function (1) can be easily derived as

$$\widetilde{m}^*\left(\boldsymbol{X};\boldsymbol{\theta}\right) \triangleq \arg\min_{\boldsymbol{Y}} \ell_Y(\boldsymbol{Y}) = \boldsymbol{P}^* m\left(\boldsymbol{X};\boldsymbol{\theta}\right), \tag{2}$$

with $\boldsymbol{P}^* \triangleq (\boldsymbol{I} + \lambda\boldsymbol{L})^{-1}$. However, for large graphs the requisite inverse is impractical to compute, and alternatively iterative approximations are more practically-feasible. To this end, we may initialize as $\boldsymbol{Y}^{(0)} = m\left(\boldsymbol{X};\boldsymbol{\theta}\right)$, and it follows that $\boldsymbol{Y}$ can be approximated by iterative descent in the direction of the negative gradient. Given that

$$\frac{\partial \ell_Y(\boldsymbol{Y})}{\partial \boldsymbol{Y}} = 2\lambda\boldsymbol{L}\boldsymbol{Y} + 2\boldsymbol{Y} - 2m\left(\boldsymbol{X};\boldsymbol{\theta}\right), \tag{3}$$

the k-th iteration of gradient descent becomes

$$\boldsymbol{Y}^{(k)} = \boldsymbol{Y}^{(k-1)} - \alpha\left[\left(\lambda\boldsymbol{L} + \boldsymbol{I}\right)\boldsymbol{Y}^{(k-1)} - m\left(\boldsymbol{X};\boldsymbol{\theta}\right)\right], \tag{4}$$

where $\frac{\alpha}{2}$ serves as the effective step size. Considering that $\boldsymbol{L}$ is generally sparse, computation of (4) can leverage efficient sparse matrix multiplications, and we may also introduce modifications such as Jacobi preconditioning to speed convergence (Axelsson, 1996; Yang et al., 2021).

Furthermore, based on well-known properties of gradient descent, if k is sufficiently large and α is small enough, then

$$\widetilde{m}^*\left(\boldsymbol{X};\boldsymbol{\theta}\right) \;\approx\; \widetilde{m}^{(k)}\left(\boldsymbol{X};\boldsymbol{\theta}\right) \;\triangleq\; \boldsymbol{P}^{(k)}\left[m\left(\boldsymbol{X};\boldsymbol{\theta}\right)\right], \tag{5}$$

where the operator $\boldsymbol{P}^{(k)}(\cdot)$ computes k gradient steps via (4). The structure of these propagation steps, as well as related variants based on normalized modifications of gradient descent, equate to principled GNN layers, such as those used by GCN (Kipf & Welling, 2016), APPNP (Klicpera et al., 2018), and many others, which can be trained within a broader bilevel optimization framework as described next.

4 STACK ENSEMBLING FOR GRAPH DATA (BESTOWGNN)

For node prediction tasks (either regression or classification), each (non-graph) base model is trained within our BestowGNN framework by simply treating each node and its label as a separate IID training example and fitting the model in the usual manner. Such a model may informatively encode tabular or text features from the nodes, but its predictions will be uniformed by the additional information available in the graph structure. To enhance such models with graph information we utilize graph-aware propagation.

4.1 GRAPH-AWARE PROPAGATION

Let $\hat{\boldsymbol{Y}}_{\mathcal{L}}, \hat{\boldsymbol{Y}}_{\mathcal{U}}$ denote the predictions of labeled (i.e. training) nodes and unlabeled (i.e. validation/test) nodes, respectively. In node classification tasks, these may be predicted class probability vectors. Via iterative application of the update in (4), we can apply graph-aware propagation to predictions $\{\hat{\boldsymbol{Y}}_{\mathcal{L}}, \hat{\boldsymbol{Y}}_{\mathcal{U}}\}$ in order to ensure they reflect statistical dependencies between nodes encoded by the graph structure. We denote $\boldsymbol{F}^{(0)} \triangleq \{\hat{\boldsymbol{Y}}_{\mathcal{L}}, \hat{\boldsymbol{Y}}_{\mathcal{U}}\}$, and for each propagation step t we compute the update $\boldsymbol{F}^{(t)}$ via (4). In our method, $\hat{\boldsymbol{Y}}$ may actually be predictions from multiple models concatenated together at each node, but the propagation procedure remains identical in this case.

4.2 STACK ENSEMBLING

In stack ensembling, the predictions output by individually trained *base* models are concatenated together as features that are subsequently used to train a *stacker* model whose target is still to predict the original labels (Wolpert, 1992; Ting & Witten, 1997). A good stacker model learns how to nonlinearly combine the predictions of base models into an even more accurate prediction. This process can be iterated in multiple layers, a strategy that has been used to win high-profile prediction competitions with IID data (Koren, 2009).

In this work, we closely follow the stacking methodology of Erickson et al. (2020), but adapt it for graphs rather than IID data. We allow stacker models to access the original node features $\boldsymbol{X}$ by concatenating $\boldsymbol{X}$ with the base models' predictions when forming the features used to train each stacker model. To produce a final prediction for each node, we aggregate predictions from the topmost layer models via a simple weighted combination where weights are learned via the efficient Ensemble Selection technique of Caruana et al. (2004). Our base models before the first stacking layer are those which can effectively encode the original tabular or text features observed at the nodes (here we utilize AutoGluon which leverages models like Gradient Boosted Decision Trees for tabular features and Transformers for text features). Our stacker models are simply chosen as the same types of models as the base models.

4.3 REPEATED K-FOLD BAGGING TO MITIGATE OVER-FITTING

A problem that arises in the aforementioned stacking strategy is *label leakage*. If a base model is even slightly overfit to its training data such that its predictions memorize parts of the training labels, then subsequent stacker models will have low accuracy due to distribution shift in their features between training and inference time (their features will be highly correlated with the labels during training but not necessarily during inference). This issue is remedied by ensuring stacker models are only trained on features comprised of base model predictions on held-out nodes omitted from the base model's training set.

We achieve this while still being able to train stacker models using all labeled nodes by leveraging k-fold bagging (i.e. cross-validation) of all models (Van der Laan et al., 2007; Parmanto et al., 1996; Erickson et al., 2020). Here the training nodes are partitioned into k disjoint chunks and k copies of

each (non-graph-aware) model m are trained with a different data chunk held-out $\{\boldsymbol{X}^{-j}, \boldsymbol{Y}^{-j}\}_{j=1}^{k}$ held out from each copy. After training all k copies of model m, we can produce out-of-fold (OOF) predictions $\hat{\boldsymbol{Y}}_m^j$ for each chunk $\boldsymbol{X}^j$ by feeding it into the model copy from which it was previously held-out. Following Erickson et al. (2020), we repeat this k-fold bagging procedure over n different random partitions of the training data to further reduce variance and distribution shift that arises in stack ensembling with bagging. Thus for a labeled training node, the OOF prediction from a model of type m is averaged over n different partitions (this node is held-out from exactly one model copy in each partition):

$$\hat{\boldsymbol{Y}}_{\mathcal{L}} = \left\{ \frac{1}{n} \sum_{i=1}^{n} \hat{\boldsymbol{Y}}_{m,i}^{j} \right\}_{j=1}^{k}. \tag{6}$$

Since unlabeled (validation/test) nodes were technically held-out from every model copy, we can feed them through any copy without harming stacking performance. For a particular type of model m, we simply make predictions $\hat{\boldsymbol{Y}}_{\mathcal{U}}$ for unlabeled nodes by averaging over all n bagging repeats and all k copies of the model within each repeat:

$$\hat{\boldsymbol{Y}}_{\mathcal{U}} = \frac{1}{kn} \sum_{j=1}^{k} \sum_{i=1}^{n} \hat{\boldsymbol{Y}}_{m,i}^{j}. \tag{7}$$

For IID data, this stack ensembling procedure with bagging can produce powerful predictors, both in theory (Van der Laan et al., 2007) and in practice (Erickson et al., 2020).

4.4 Stacking with Graph-Aware Propagation

To extend this methodology to graph data, our proposed training strategy is precisely detailed in Algorithm 1. The main idea is to apply graph-aware propagation on the predictions of models at each intermediate layer of the stack. Different amounts of propagation lead to different characteristics of the data being captured in the resulting prediction (few steps of propagation means predictions are only influenced by local neighbors, whereas many propagation steps allow predictions to be influenced by more distant nodes as well). Thus we can further enrich the feature set of our stacker models by concatenating together the predictions produced after different numbers of propagation steps. With this expanded feature set, our stacker models learn to aggregate not only the predictions of different models, but differently smoothed versions of these predictions as well. This allows the stacker model to adaptively decide how to best account for dependencies induced by the graph structure.

More precisely, if we let $\boldsymbol{F}^{(t)}$ denote the predictions (concatenated across all base model types) for labeled and unlabeled nodes after t smoothing steps, then the feature input to each stacker model is given by the original node features $\boldsymbol{X}$ concatenated with $[\boldsymbol{F}^{(0)}, ..., \boldsymbol{F}^{(T)}]$. Here the predictions for labeled nodes are always OOF, obtained via bagging. Another fundamental difference between our approach and stack ensembling in the IID setting is *the use of unlabeled (test) nodes at each intermediate layer of the stack*. By including unlabeled nodes in the propagation, these nodes influence the features used to train subsequent stacker models at labeled nodes. This can even further reduce potential distribution shift in the stacker models' features between the labeled and unlabeled nodes, which ensures better generalization.

Graph machine learning models for non-IID data typically do not use bagging, seemingly because there has not been a rigorous study on the effect of bagging in relation to propagation models. Furthermore, bagging traditionally serves as a means of variance reduction which only brings limited performance benefits for large datasets (Breiman, 1996). In contrast, our stacking framework adopts bagging primarily as a means to mitigate the catastrophic effects of label leakage. While bagging can effectively mitigate label information from being directly encoded in stacker model features in the IID setting, it is not clear whether this property still holds with graph-structured dependence between nodes. A particular concern is the fact that the propagation of base model predictions across the graph implies label information is shared across the k-fold chunks used to hold-out some nodes from some models. In the next section, we theoretically study this issue and prove that bagging can still mitigate the effects of label leakage even in the non-IID graph setting. Our subsequent experiments (see Table 4) reveal that bagging produces substantial performance gains in practical applications of stack ensembling with graph propagation.

5 THEORETICAL ANALYSIS

Label utilization is a common technique in which the outputs of a model are concatenated with input features and then used to train a stacking layer. Unfortunately, layer-wise training with label utilization is susceptible to the label leakage problem. Although prior work (Sun & Wu, 2021; Zhang et al., 2021) has mentioned heuristic ways to address label leakage via graph propagation, it is unclear how generally applicable this strategy is in practice. Moreover, there is a natural trade-off between avoiding label leakage via graph propagation, and well-known oversmoothing effects in GNN models.

In this section we employ a powerful theoretical tool, Differential Privacy (Mironov, 2017), to showcase the advantage of bagging in our proposed BestowGNN. Our analysis will show that BestowGNN enjoys strong generalization under the Rényi Differential Privacy framework. In fact this is the first work that establishes that bagging in graph predictors is useful and mitigates label leakage. Specifically, BestowGNN can preserve the privacy (or information sharing) of labels between bags, that would otherwise be compromised by graph propagation.

To this end, we first introduce the definition of Rényi Differential Privacy, which is a relaxation of Differential Privacy based on the Rényi Divergence.

Definition 1. *(Rényi Differential Privacy (Mironov, 2017)). Consider a randomized algorithm $\mathcal{M}$ mapping from $\mathcal{D}$ to a real-value $\mathcal{R}$. Such an algorithm is said to have ϵ-Rényi Differential Privacy of order α if for any $D, D' \in \mathcal{D}$ with $d_H(D, D') = 1$, where d_H is the Hamming distance (D, D' are also referred to as adjacent datasets), we have that*

$$D_\alpha(\mathcal{M}(D)||\mathcal{M}(D')) \triangleq \frac{1}{\alpha - 1} \log E_{x \sim \mathcal{M}(D')} \left(\frac{\mathcal{M}(D)}{\mathcal{M}(D')} \right)^\alpha \leq \epsilon. \tag{8}$$

In plain words, this definition establishes that the output of an algorithm does not change significantly, as measured by the Rényi divergence $D_\alpha(\mathcal{M}(D)||\mathcal{M}(D'))$, when the data changes slightly. The idea behind this framework is that if each individual data sample has only a small effect on the resulting model, the model cannot be used to infer information about any single individual.

We then have the following result:

Theorem 1. *Assume base model m is a multi-layer (two-layer) perceptron and that node features $\boldsymbol{X}$ are sampled from a multivariate Gaussian as in (Jia & Benson, 2021):*

$$\boldsymbol{X} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Gamma}^{-1}), \qquad \boldsymbol{\Gamma} = c_1 \boldsymbol{I}_n + c_2 \boldsymbol{L},$$

where $\boldsymbol{I}_n$ is an identity matrix and $\boldsymbol{L}$ is the normalized graph Laplacian. Here c_1 controls a noise level and c_2 the smoothness over the whole graph. $\boldsymbol{E}(x_0; D_\mathcal{L})$ and $\boldsymbol{F}(x_0; D_\mathcal{L})$ are predictions produced by BestowGNN for a data point x_0 with and without bagging mode, respectively. If $\boldsymbol{E}$ has sensitivity 1 and lower magnitude bound L, i.e., for any two adjacent $D, D' \in D : |\boldsymbol{E}(\boldsymbol{x}_0; D) - \boldsymbol{E}(\boldsymbol{x}_0; D')| \leq 1$ and $|\boldsymbol{E}| \geq L$, then $\boldsymbol{E}$ satisfies $(\frac{1}{2}, \frac{1}{4\sigma^2 L^2} + \frac{1}{2L^2})$-Rényi Differential Privacy, where σ^2 depends on graph structure $\mathcal{G}$. Meanwhile, $\boldsymbol{F}$ has no privacy guarantee, i.e., the Rényi differential privacy loss (8) is unbounded.

The proof is deferred to the supplementary. Theorem 1 indicates that bagging with graph propagation can well preserve the privacy of $D_\mathcal{L} = \{\boldsymbol{x}_i, \boldsymbol{y}_i\}_{i \in \mathcal{L}}$ between different chunks while non-bagging would have a high risk of leaking the information of $D_\mathcal{L}$. For layer-wise training with label utilization, the output of the model $\boldsymbol{E}(x_0; D_\mathcal{L})$ is concatenated with input features and then used to train next stacking layer, and bagging can effectively mitigate the label leakage issue since the information of true label is well preserved at the first layer, while no-bagging exposes the true labels and can lead to over-fitting issue for next stacking layer.

6 EXPERIMENTS

Setup. We study the effectiveness of our approach by comparing performance against state-of-the-art baselines in node regression and classification tasks. For node regression with **tabular node features**, we consider four real-world graph datasets used for benchmarking by Ivanov & Prokhorenkova (2021): House, County, VK and Avazu. As node classification tasks, we adopt one

datasets with **numerical features**: Reddit; and two datasets with **raw text features**: OGB-Arxiv and OGB-Products. More details about the datasets are provided in the supplementary.

We compare our method against various baselines, starting with purely tabular baseline models or language models where the graph structure is ignored. Our first baseline is **Autogluon** (Erickson et al., 2020), an AutoML system for IID tabular or text data that is completely unaware of the graph structure (here we simply treat nodes as IID). Next, we consider **AutoGluon + C&S**, which performs Correct and Smooth (Huang et al., 2020) as a posthoc processing step on top of AutoGluon's predictions, in order to at least account for the graph structure during inference. For node regression tasks we also consider some popular GNN models: **GCN** (Kipf & Welling, 2016), **GAT** (Veličković et al., 2017), and a hybrid strategy **BGNN** (Ivanov & Prokhorenkova, 2021), which combines Gradient Boosted Decision Trees (also a model intended for IID data) with GNNs via end-to-end training in a manner that is graph-aware.

For node classification, we firstly consider Reddit with original numerical features. We compare with **GraphSAGE** (Hamilton et al., 2017) and **PCAPass + Tree** (Sadowski et al., 2022), which combines PCA and message passing to generate node embeddings and leverages tree-based model for node classification.

We also consider OGB-Arxiv and OGB-Products with raw text as node features (as opposed to pre-computed text embeddings as node features such as the low-dimensional homogeneous embeddings provided by OGB). We compare with **GIANT-XRT + MLP**, **GIANT-XRT + GRAPHSAGE** and **GIANT-XRT + GRAPHSAINT**, which extracts numerical embeddings from text features via a transformer trained through self-supervised learning and feed these high quality embeddings to a multi-layer perceptron or sampling based GNN model. For the smaller OGB-Arxiv dataset, we also consider standard GNN models: **GCN** (Kipf & Welling, 2016), **GAT** (Veličković et al., 2017) and **Ensemble GCN**, a natural baseline/competitor which divides all training nodes into k chunks, trains a GCN model for each chunk and then ensembles the results. Finally, we compare against SOTA model for OGB-Arxiv and OGB-Products with GIANT-XRT embedding and low-dimensional homogeneous embedding from OGB leaderboard. To our knowledge, there is not a consistent method with superior performance across each dataset. So we compare our *single* general framework with *different* SOTA models for each dataset to ensure we are competing against the best in each case; i.e., *there is no single existing model that is SOTA across them all*. We evaluate our method **BestowGNN**, which incorporates the graph information through propagation operations in each stacking layer.

Results. In Table 1 we present the results for the node regression task with tabular node features. The baseline GNN models are challenged by the tabular node features. AutoGluon is an ensemble of various base models (e.g., Gradient Boosted Decision Trees, fully-connected neural networks) intended for IID data without considering graph structure. We observe that **Autogluon + C&S** outperforms **Autogluon**, demonstrating that graph information can greatly boost the performance of models intended for IID data. Incorporating the graph structure at each stacking layer, our **BestowGNN** method performs better than **BGNN** on all datasets.

Tables 2 and 3 show the results for node classification with either raw text features or numerical embeddings. Our method BestowGNN outperforms all baselines regardless of whethor or not they leverage the raw text or OGB embeddings (or numerical Reddit embeddings). Note that OGB-Arxiv and OGB-Products have *different* SOTA models in the OGB leaderboard, for instance: **AGDN + BoT + self-KD + C&S** are architectural components from the best existing model for OGB-Arxiv, while **GAMLP + RLU + SCR + C&S** undergird the best existing model for OGB-Products. These SOTA models consisting of data-specific modules/components are manually composed to perform particularly well only for one specific dataset. In contrast, **BestowGNN** uses essentially the same architecture with minor/standard hyperparameter tuning to fit all datasets. Comparison of **BestowGNN** with AutoGluon demonstrates how incorporating graph information at each stacking layer can further improve the node classification performance of this AutoML system. More experiments details and computing cost are deferred to the supplementary.

Ablation. The key ingredients of our framework are bagging/ensembling and graph propagation. Table 4 shows an ablation study involving these components using OGB-Arxiv and OGB-Products with original OGB embeddings. From these results we observe that bagging modes can outperform

Table 1: Mean squared error results for four node regression datasets with **tabular** node features.

Data set	House	County	Vk	Avazu
GCN	0.63 ± 0.01	1.48 ± 0.08	7.25 ± 0.19	0.1141 ± 0.02
GAT	0.54 ± 0.01	1.45 ± 0.06	7.22 ± 0.19	0.1134 ± 0.01
BGNN	0.50 ± 0.01	1.26 ± 0.08	6.95 ± 0.21	0.109 ± 0.01
AutoGluon	0.618 ± 0.01	1.379 ± 0.08	7.176 ± 0.21	0.117 ± 0.018
AutoGluon + C&S	0.477 ± 0.01	1.162 ± 0.09	6.995 ± 0.21	0.107 ± 0.015
BestowGNN	**0.467 ± 0.007**	**1.145 ± 0.083**	**6.918 ± 0.220**	**0.105 ± 0.013**

Table 2: Node classification accuracy for Reddit with **numerical** node features.

Method	Reddit
PCAPass + XGBoost	96.26 ± 0.02
GraphSAGE	95.40 ± 0.22
AutoGluon	95.83 ± 0.00
BestowGNN	**96.44 ±0.00**

Table 3: Node classification accuracy for OGB-Arxiv and OGB-Products achieved by various methods. Rows labeled TEXT contain methods including SOTA models trained on the **raw text** features at each node, while those labeled OGB indicate models trained on precomputed **numerical embeddings** provided by OGB as node features. *SOTA models vary from each dataset with different embeddings/architectures, but BestowGNN has consistently superior performance for each dataset; similarly for Table 1 results above.*

OGB-Arxiv

Feature	Method	Test Acc (Validation)
OGB	GCN	73.06 ± 0.24 (74.42 ± 0.12)
	GAT + C&S	73.86 ± 0.14 (74.84 ± 0.07)
	AGDN+BoT+self-KD+C&S	74.31 ± 0.14 (75.18 ± 0.09)
	Ensemble GCN	73.22 ± 0.12 (74.64 ± 0.01)
TEXT	GIANT-XRT+MLP	73.06 ± 0.11 (74.32 ± 0.09)
	GIANT-XRT+graphSAGE	74.35 ± 0.14 (75.95 ± 0.11)
	GIANT-XRT+GCN	75.28 ± 0.17 (76.87 ± 0.04)
	GIANT-XRT+RevGAT+KD	76.15 ± 0.10 (77.16 ± 0.09)
TEXT	AutoGluon	73.05 ± 0.00 (74.33 ± 0.00)
	AutoGluon + C&S	75.34 ± 0.00 (76.67 ± 0.00)
TEXT	BestowGNN	**76.19 ± 0.02 (77.25 ± 0.05)**

OGB-Products

Feature	Method	Test Acc (Validation)
OGB	DeeperGCN + FLAG	81.93 ± 0.31 (92.21 ± 0.37)
	GAT + FLAG	81.76 ± 0.45 (92.51 ± 0.06)
	GAMLP+RLU+SCR+C&S	85.20 ± 0.08 (93.04 ± 0.05)
	Ensemble GAT	80.01 ±0.20 (93.24 ± 0.05)
TEXT	GIANT-XRT+MLP	80.49 ± 0.28 (92.10 ± 0.09)
	GIANT-XRT+graphSAGE	81.99 ± 0.45 (93.38 ± 0.05)
	GIANT-XRT+graphSAINT	84.15 ± 0.22 (93.18 ± 0.04)
	GIANT-XRT+SAGN+SLE	85.47 ± 0.29 (-)
TEXT	AutoGluon	77.10 ± 0.06 (91.78 ± 0.03)
	AutoGluon + C&S	79.03 ± 0.12 (93.62 ± 0.03)
TEXT	BestowGNN	**85.48 ± 0.03 (93.93 ± 0.02)**

Table 4: BestowGNN ablation study with (✓) and without bagging (✗). Here T is the number of graph propagation steps, thus $T = 0$ represents a baseline model that completely ignores graph structure.

STEP T	ARXIV		PRODUCTS	
	✓	✗	✓	✗
0	55.70 ± 0.33	54.14 ± 0.29	62.28 ± 0.35	62.05 ± 0.19
1	66.25 ± 0.27	64.57 ± 0.76	74.18 ± 0.21	72.61 ± 0.66
2	69.34 ± 0.16	67.37 ± 0.44	77.07 ± 0.32	74.61 ± 0.58
3	70.01 ± 0.16	68.08 ± 0.74	78.11 ± 0.19	75.79 ± 0.49
4	70.43 ± 0.21	68.72 ± 0.63	78.76 ± 0.60	76.86 ± 0.17

no-bagging modes for each number propagation step, demonstrating that bagging can effectively mitigate label leakage and over-fitting issues even in a graph-aware propagation setting.

7 DISCUSSION

While real-world graph data come with heterogeneous feature types, existing GNN models are primarily suited for (adequately preprocessed) numerical features. For IID supervised learning, it is well-known that the best models for different feature types vary based on dataset and data-type, and that a learning system aiming to output good predictions across a variety of datasets should leverage a heterogeneous collection of different types of models (Erickson et al., 2020). There is little reason the situation should be different for graph data. In this paper, we demonstrate the first working system that can utilize arbitrary heterogeneous collections of models for arbitrary graph datasets with heterogeneous feature-types (numerical, categorical, text). This is achieved by means of a novel graph-aware stack ensembling technique that takes the graph structure into account without restricting how individual models are trained. Our graph-aware propagation techniques leverage specific properties of stack ensembling that allow our proposed methodology to outperform both many complex GNNs as well as existing approaches in which propagation is only applied to the predictions output by an IID base model (e.g., AutoGluon+C&S, etc.).

REFERENCES

Owe Axelsson. *Iterative Solution Methods*. Cambridge University Press, 1996.

Oleg Bezrukavnikov and Rhema Linder. A neophyte with automl: Evaluating the promises of automatic machine learning tools. *arXiv preprint arXiv:2101.05840*, 2021.

Matthias Blohm, Marc Hanussek, and Maximilien Kintz. Leveraging automated machine learning for text classification: Evaluation of automl tools and comparison with human performance. *arXiv preprint arXiv:2012.03575*, 2020.

Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.

Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. Ensemble selection from libraries of models. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 18, 2004.

Jiuhai Chen, Jonas Mueller, Vassilis N Ioannidis, Soji Adeshina, Yangkun Wang, Tom Goldstein, and David Wipf. Convergent boosted smoothing for modeling graph data with tabular node features. *arXiv preprint arXiv:2110.13413*, 2021.

Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 257–266, 2019.

Eli Chien, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Jiong Zhang, Olgica Milenkovic, and Inderjit S Dhillon. Node feature extraction by self-supervised multi-scale neighborhood prediction. *arXiv preprint arXiv:2111.00064*, 2021.

Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*, 2020.

Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. Autogluon-tabular: Robust and accurate automl for structured data. *arXiv preprint arXiv:2003.06505*, 2020.

Rasool Fakoor, Jonas Mueller, Nick Erickson, Pratik Chaudhari, and Alexander J Smola. Fast, accurate, and simple models for tabular data via augmented distillation. In *Advances in Neural Information Processing Systems*, 2020.

Sergey Feldman. Which machine learning classifiers are best for small datasets? An empirical study. `https://www.data-cowboys.com/blog/`, 2021.

Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. Efficient and robust automated machine learning. *Advances in neural information processing systems*, 28, 2015.

William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 1025–1035, 2017.

Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.

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

Sergei Ivanov and Liudmila Prokhorenkova. Boost then convolve: Gradient boosting meets graph neural networks. *arXiv preprint arXiv:2101.08543*, 2021.

Junteng Jia and Austin R Benson. A unifying generative model for graph learning algorithms: Label propagation, graph convolutions, and combinations. *arXiv preprint arXiv:2101.07730*, 2021.

Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, 2017.

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

Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997*, 2018.

Yehuda Koren. The bellkor solution to the netflix grand prize, 2009. URL `https://www2.seas.gwu.edu/~simhaweb/champalg/cf/papers/KorenBellKor2009.pdf`.

Erin LeDell and Sebastien Poirier. H2o automl: Scalable automatic machine learning. In *Proceedings of the AutoML Workshop at ICML*, volume 2020, 2020.

Yuxiao Lin, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. Bertgcn: Transductive text classification by combining gcn and bert, 2021.

Yao Ma, Xiaorui Liu, Tong Zhao, Yozen Liu, Jiliang Tang, and Neil Shah. A unified view on graph neural networks as graph signal denoising. *arXiv preprint arXiv:2010.01777*, 2020.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.

Xuran Pan, Shiji Song, and Gao Huang. A unified framework for convolution-based graph neural networks, 2021. URL `https://openreview.net/forum?id=zUMD--Fb9Bt`.

Bambang Parmanto, Paul W Munro, and Howard R Doyle. Reducing variance of committee prediction with resampling techniques. *Connection Science*, 8(3-4):405–426, 1996.

Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, 2018.

Krzysztof Sadowski, Michał Szarmach, and Eddie Mattia. Dimensionality reduction meets message passing for graph node embeddings. *arXiv preprint arXiv:2202.00408*, 2022.

Xingjian Shi, Jonas Mueller, Nick Erickson, Mu Li, and Alexander J Smola. Benchmarking multimodal automl for tabular data with text fields. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021.

Chuxiong Sun and Guoshi Wu. Scalable and adaptive graph neural networks with self-label-enhanced training. *arXiv preprint arXiv:2104.09376*, 2021.

Kai Ming Ting and Ian H Witten. Stacking bagged and dagged models. In *International Conference on Machine Learning*, 1997.

Mark J Van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6(1), 2007.

Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

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

Yangkun Wang, Jiarui Jin, Weinan Zhang, Yong Yu, Zheng Zhang, and David Wipf. Bag of tricks for node classification with graph neural networks. *arXiv preprint arXiv:2103.13355*, 2(3), 2021.

David H Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.

Yongyi Yang, Tang Liu, Yangkun Wang, Jinjing Zhou, Quan Gan, Zhewei Wei, Zheng Zhang, Zengfeng Huang, and David Wipf. Graph neural networks inspired by classical iterative algorithms. *arXiv preprint arXiv:2103.06064*, 2021.

Jason Yoo, Tony Joseph, Dylan Yung, S Ali Nasseri, and Frank Wood. Ensemble squared: A meta automl system. *arXiv preprint arXiv:2012.05390*, 2020.

Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. L2-gcn: Layer-wise and learned efficient training of graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2127–2135, 2020.

Hongwei Zhang, Tijin Yan, Zenjun Xie, Yuanqing Xia, and Yuan Zhang. Revisiting graph convolutional network on semi-supervised node classification from an optimization perspective. *arXiv preprint arXiv:2009.11469*, 2020.

Wentao Zhang, Ziqi Yin, Zeang Sheng, Wen Ouyang, Xiaosen Li, Yangyu Tao, Zhi Yang, and Bin Cui. Graph attention multi-layer perceptron. *arXiv preprint arXiv:2108.10097*, 2021.

Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. *Advances in Neural Information Processing Systems*, 2004.

Meiqi Zhu, Xiao Wang, Chuan Shi, Houye Ji, and Peng Cui. Interpreting and unifying graph neural networks with an optimization framework. *arXiv preprint arXiv:2101.11859*, 2021.

Supplementary Materials

A PROOF OF THEOREM 1.

Preliminary 1. Firstly, we derive the format of $\boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}})$ and $\boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}})$. Suppose BestowGNN randomly splits the labeled nodes $D_{\mathcal{L}}$ into 2 disjoint chunks $D_1 = \{\boldsymbol{X}_1, \boldsymbol{Y}_1\}, D_2 = \{\boldsymbol{X}_2, \boldsymbol{Y}_2\}$. BestowGNN trains a model $m \in \mathcal{M}$ with a different data chunk held-out. Model m is defined by a set of parameters collected in $\boldsymbol{\theta}$ namely, which is defined as $m(\boldsymbol{X}; \boldsymbol{\theta})$. In the following, we will express the predicted labels from model m under the bagging and non-bagging settings. We compare the predicted labels under both settings and establish that our bagging solution is less amenable to label leakage.

The model m will learn different parameters for each chunk and those are denoted as $\boldsymbol{\theta}_1$ for the chunk I and $\boldsymbol{\theta}_2$ for the chunk II, namely $\boldsymbol{\theta}_1 = \boldsymbol{\theta}(D_1)$ and $\boldsymbol{\theta}_2 = \boldsymbol{\theta}(D_2)$. Next, BestowGNN produces prediction $\hat{\boldsymbol{Y}}_1, \hat{\boldsymbol{Y}}_2$ on out-of-fold data, i.e., $\hat{\boldsymbol{Y}}_1 = m(\boldsymbol{X}_1; \boldsymbol{\theta}_2)$ and $\hat{\boldsymbol{Y}}_2 = m(\boldsymbol{X}_2; \boldsymbol{\theta}_1)$. The prediction for unlabeled nodes is $\hat{\boldsymbol{Y}}_{\mathcal{U}} = \frac{1}{2}[m(\boldsymbol{X}_{\mathcal{U}}; \boldsymbol{\theta}_1) + m(\boldsymbol{X}_{\mathcal{U}}; \boldsymbol{\theta}_2)]$ as explained in (7). Consider one data point $\boldsymbol{x}_0$ from the unlabeled dataset $D_{\mathcal{U}}$, the prediction of $\boldsymbol{x}_0$ is given by $\hat{\boldsymbol{y}}_0 = \frac{1}{2}[m(\boldsymbol{x}_0; \boldsymbol{\theta}_1) + m(\boldsymbol{x}_0; \boldsymbol{\theta}_2)]$. Next, we perform one step graph-aware propagation on $\hat{\boldsymbol{y}}_0$.

$$\begin{aligned}\hat{\boldsymbol{y}}_0^{(1)} &= \sum_{u \in \mathcal{N}(\boldsymbol{x}_0) \cap D_{\mathcal{U}}} \hat{\boldsymbol{y}}_u + \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} \hat{\boldsymbol{y}}_v + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} \hat{\boldsymbol{y}}_w \\ &= \sum_{u \in \mathcal{N}(\boldsymbol{x}_0) \cap D_{\mathcal{U}}} \frac{1}{2}[m(\boldsymbol{x}_u; \boldsymbol{\theta}_1) + m(\boldsymbol{x}_u; \boldsymbol{\theta}_2)] + \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}_2) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}_1),\end{aligned} \tag{9}$$

where $\hat{\boldsymbol{y}}_0^{(1)}$ is the aggregated results from one-hop neighbor $\mathcal{N}(\boldsymbol{x}_0)$, which may belongs to $D_{\mathcal{U}}, D_1$ and D_2.

Next, we consider the no-bagging mode, where the predictions of $\boldsymbol{X}_1, \boldsymbol{X}_2$ are changed into $\widetilde{\boldsymbol{Y}}_1 = m(\boldsymbol{X}_1; \boldsymbol{\theta}_1)$ and $\widetilde{\boldsymbol{Y}}_2 = m(\boldsymbol{X}_2; \boldsymbol{\theta}_2)$. Notice that with bagging mode we use the parameters from a different bag, while without bagging we use the parameters from the same bag. The prediction of the test point $\boldsymbol{x}_0$ is once again $\widetilde{\boldsymbol{y}}_0 = \frac{1}{2}[m(\boldsymbol{x}_0; \boldsymbol{\theta}_1) + m(\boldsymbol{x}_0; \boldsymbol{\theta}_2)]$, which is identical to the bagging mode. We perform the same graph-aware propagation on $\widetilde{\boldsymbol{y}}_0$.

$$\begin{aligned}\widetilde{\boldsymbol{y}}_0^{(1)} &= \sum_{u \in \mathcal{N}(\boldsymbol{x}_0) \cap D_{\mathcal{U}}} \widetilde{\boldsymbol{y}}_u + \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} \widetilde{\boldsymbol{y}}_v + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} \widetilde{\boldsymbol{y}}_w \\ &= \sum_{u \in \mathcal{N}(\boldsymbol{x}_0) \cap D_{\mathcal{U}}} \frac{1}{2}[m(\boldsymbol{x}_u; \boldsymbol{\theta}_1) + m(\boldsymbol{x}_u; \boldsymbol{\theta}_2)] + \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}_1) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}_2).\end{aligned} \tag{10}$$

Next, we compare the terms among the predicted labels from the two settings, namely (9) and (10). The first term $\sum_{u \in \mathcal{N}(\boldsymbol{x}_0) \cap D_{\mathcal{U}}} \frac{1}{2}[m(\boldsymbol{x}_u; \boldsymbol{\theta}_1) + m(\boldsymbol{x}_u; \boldsymbol{\theta}_2)]$ is the same for (9) and (10) and can be cancelled. In order to facilitate the exposition of the theoretical contributions we will define functions for the different terms in (9) and (10). We define $\boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}})$, that is a function formulating the relation between training data $D_{\mathcal{L}}$ and the prediction for test data $\boldsymbol{x}_0$ under bagging mode.

$$\boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}}) := \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}(D_2)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}(D_1)). \tag{11}$$

Similarly, we define the function $\boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}})$ formulating the relation between training data $D_{\mathcal{L}}$ and the prediction for test data $\boldsymbol{x}_0$ under the no-bagging mode:

$$\boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}}) := \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}(D_1)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}(D_2)). \tag{12}$$

Notice here $\boldsymbol{\theta}(D_1)$ is the model parameters of Chunk I involving information of true label $\boldsymbol{Y}_1$. We aim to examine bagging and stacking strategies effectively preserve the information of label $\boldsymbol{Y}_1$ via

introducing randomness to the function $\boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}})$ while $\boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}})$ has high risk of leaking the information of true label $\boldsymbol{Y}_1$.

We first reiterate the definition of Rényi Differential Privacy.

Definition 1. *(Rényi Differential Privacy (Mironov, 2017)). Consider a randomized algorithm $\mathcal{M}$ mapping from $\mathcal{D}$ to real-value $\mathcal{R}$. Such an algorithm is said to have ϵ-Rényi Differential Privacy of order α ($\alpha > 1$) if any $D, D' \in \mathcal{D}$ with $d_H(D, D') = 1$, where d_H is the Hamming distance (D, D' are also referred to as adjacent datasets):*

$$D_\alpha(\mathcal{M}(D)||\mathcal{M}(D')) = \frac{1}{\alpha - 1} \log E_{x \sim \mathcal{M}(D')} \left(\frac{\mathcal{M}(D)}{\mathcal{M}(D')} \right)^\alpha \leq \epsilon. \tag{13}$$

To proceed in a quantifiable way, we rely on some preliminary results for Rényi Differential privacy and generative model for graph learning algorithms.

Proposition 1. *Rényi differential privacy is preserved by post-processing (Mironov, 2017). If $F(\cdot)$ has ϵ-Rényi Differential Privacy, then for any randomized or deterministic function g, $g(F(\cdot))$ satisfies ϵ-Rényi Differential Privacy.*

Proposition 2. *The closed-form expression of the Rényi divergence between any two Gaussian distributions is given by $D_\alpha(\mathcal{N}(\mu_0, \sigma_0^2)||\mathcal{N}(\mu_1, \sigma_1^2)) = \frac{\alpha(\mu_1 - \mu_0)^2}{2\sigma_\alpha^2} + \frac{1}{1-\alpha} \ln \frac{\sigma_\alpha}{\sigma_0^{1-\alpha}\sigma_1^\alpha}$, provided that $\sigma_\alpha^2 = (1-\alpha)\sigma_0^2 + \alpha\sigma_1^2 > 0$ (Van Erven & Harremos, 2014).*

Proposition 3. *Assume f has sensitivity 1 and lower magnitude bound L, i.e., for any pair of adjacent datasets $D, D' \in \mathcal{D}$: $|f(D) - f(D')| \leq 1$ and $|f| \geq L$, and define the Gaussian multiplicative mechanism*

$$\boldsymbol{GM}_{\mu,\sigma} f(D) = f(D)\mathcal{N}(\mu, \sigma^2).$$

Then $\boldsymbol{GM}_{\mu,\sigma} f$ satisfies $(\frac{1}{2}, \frac{1}{4\sigma^2 L^2} + \frac{1}{2L^2})$-Rényi Differential Privacy.

Proof. According to Proposition (2):

$$\begin{aligned}
&D_{1/2}\left(\mathcal{N}(f(D) + \mu, f^2(D)\sigma^2)||\mathcal{N}(f(D') + \mu, f^2(D')\sigma^2)\right) \\
&= \frac{(f(D) - f(D'))^2}{2\sigma^2\left(f^2(D) + f^2(D')\right)} + \ln[\frac{1}{2}\left(f^2(D) + f^2(D')\right)] - \ln|f(D)| - \ln|f(D')| \\
&= \frac{1}{2\sigma^2} - \frac{f(D)f(D')}{\sigma^2\left(f^2(D) + f^2(D')\right)} + \ln[\frac{1}{2}\left(f^2(D) + f^2(D')\right)] - \ln|f(D)f(D')| \\
&= \frac{1}{2\sigma^2} - \frac{f(D)f(D')}{\sigma^2\left(f^2(D) + f^2(D')\right)} + \ln|\frac{f^2(D) + f^2(D')}{2f(D)f(D')}| \\
&\leq \frac{1}{2\sigma^2}\frac{1}{f^2(D) + f^2(D')} + \ln(\frac{1}{2|f(D)f(D')|} + 1) \\
&\leq \frac{1}{4\sigma^2 L^2} + \ln(\frac{1}{2L^2} + 1) \\
&\leq \frac{1}{4\sigma^2 L^2} + \frac{1}{2L^2}.
\end{aligned}$$

The first inequality follows from $|f(D) - f(D')| \leq 1$, take square for both side $f^2(D) + f^2(D') \leq 1 + 2f(D)f(D')$. Then we have $\frac{1}{2\sigma^2} - \frac{f(D)f(D')}{\sigma^2(f^2(D)+f^2(D'))} \leq \frac{1}{2\sigma^2}\frac{1}{f^2(D)+f^2(D')}$ and $\frac{f^2(D)+f^2(D')}{2|f(D)f(D')|} \leq \frac{1}{2|f(D)f(D')|} + 1$, the first inequality holds. □

Proposition 4. *If f has sensitivity 1, i.e., for any pair of adjacent datasets $D, D' \in \mathcal{D}$: $|f(D) - f(D')| \leq 1$. Define the Gaussian additive mechanism*

$$\boldsymbol{GA}_\sigma f(D) = f(D) + \mathcal{N}(0, \sigma^2),$$

then Gaussian additive mechanism $\boldsymbol{GA}_\sigma f$ satisfies $(\alpha, \frac{\alpha}{2\sigma^2})$-Rényi Differential Privacy (Mironov, 2017).

Proposition 5. *Consider a multivariate Gaussian distribution, and the random variables are partitioned into two groups $(\boldsymbol{z}_P, \boldsymbol{z}_Q)$, the distribution is block matrix format*

$$\begin{pmatrix} \boldsymbol{z}_P \\ \boldsymbol{z}_Q \end{pmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \bar{\boldsymbol{z}}_P \\ \bar{\boldsymbol{z}}_Q \end{bmatrix}, \quad \begin{bmatrix} \boldsymbol{\Gamma}_{PP} & \boldsymbol{\Gamma}_{PQ} \\ \boldsymbol{\Gamma}_{QP} & \boldsymbol{\Gamma}_{QQ} \end{bmatrix}^{-1} \right),$$

where $\begin{bmatrix} \boldsymbol{\Gamma}_{PP} & \boldsymbol{\Gamma}_{PQ} \\ \boldsymbol{\Gamma}_{QP} & \boldsymbol{\Gamma}_{QQ} \end{bmatrix}$ is precision (inverse covariance) matrix. Then the marginal and conditional distribution can be written as

$$\boldsymbol{z}_P \sim \mathcal{N}\left(\bar{\boldsymbol{z}}_P, (\boldsymbol{\Gamma}_{PP} - \boldsymbol{\Gamma}_{PQ}\boldsymbol{\Gamma}_{QQ}^{-1}\boldsymbol{\Gamma}_{QP})^{-1} \right), \tag{14}$$

$$\boldsymbol{z}_P | \boldsymbol{z}_Q = \boldsymbol{z}_Q \sim \mathcal{N}\left(\bar{\boldsymbol{z}}_P - \boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}_{PQ}(\boldsymbol{z}_Q - \bar{\boldsymbol{z}}_Q) \right). \tag{15}$$

Before proceeding to our specific results in the main paper, we also need to describe the graph setting.

Preliminary 2. Let $\mathcal{G} = (V, E)$ be an undirected graph, where V is the set of n nodes and E is the set of edges. The adjacency matrix of $\mathcal{G}$ is $\boldsymbol{W} \in \mathcal{R}^{n \times n}$, the diagonal degree matrix is $\boldsymbol{D} \in \mathcal{R}^{n \times n}$. The normalize graph Laplacian can be written as $\boldsymbol{N} = \boldsymbol{I} - \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2} = \boldsymbol{I} - \boldsymbol{S}$. We use $\boldsymbol{X} \in \mathcal{R}^{n \times p}$ for the feature matrix, where p is the dimension of features. We assume all vertex features $\boldsymbol{X}$ are jointly sampled from a multivariate Gaussian distribution (Jia & Benson, 2021), namely

$$\boldsymbol{X} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Gamma}^{-1}), \qquad \boldsymbol{\Gamma} = c_1\boldsymbol{I}_n + c_2\boldsymbol{N}, \tag{16}$$

where $\boldsymbol{I}_n$ is identical matrix, $\boldsymbol{N}$ is normalized graph Laplacian. Here c_1 controls noise level and c_2 controls the smoothness over the whole graph.

We now proceed to our specific results in the main paper.

Theorem 1. *Assume base model m to be a multi-layer (two-layer) perceptron and node features $\boldsymbol{X}$ is sampled from a multivariate Gaussian as in Jia & Benson (2021):*

$$\boldsymbol{X} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Gamma}^{-1}), \qquad \boldsymbol{\Gamma} = c_1\boldsymbol{I}_n + c_2\boldsymbol{L},$$

where $\boldsymbol{I}_n$ is an identity matrix and $\boldsymbol{L}$ is the normalized graph Laplacian. Here c_1 controls noise level and c_2 controls the smoothness over the whole graph. $\boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}})$ and $\boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}})$ are predictions produced by BestowGNN for a data point $\boldsymbol{x}_0$ with and without bagging mode, respectively. If $\boldsymbol{E}$ has sensitivity 1 and lower magnitude bound L, i.e., for any two adjacent $D, D' \in D : |\boldsymbol{E}(\boldsymbol{x}_0; D) - \boldsymbol{E}(\boldsymbol{x}_0; D')| \leq 1$ and $|\boldsymbol{E}| \geq L$, then $\boldsymbol{E}$ satisfies $(\frac{1}{2}, \frac{1}{4\sigma^2 L^2} + \frac{1}{2L^2})$-Rényi Differential Privacy, where σ^2 depends on graph structure $\mathcal{G}$. Meanwhile, $\boldsymbol{F}$ has no privacy guarantee, i.e., the Rényi differential privacy loss (8) is unbounded.

Proof. Given the definition of function $\boldsymbol{E}$ from above, we have that

$$\begin{aligned} \boldsymbol{E}(\boldsymbol{x}_0; D_{\mathcal{L}}) &= \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}(D_2)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}(D_1)) \\ &= \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v \boldsymbol{\theta}(D_2)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w \boldsymbol{\theta}(D_1)), \end{aligned} \tag{17}$$

where the second equality follows from the MLP assumption. Similarly for $\boldsymbol{F}$ we have

$$\begin{aligned} \boldsymbol{F}(\boldsymbol{x}_0; D_{\mathcal{L}}) &= \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v; \boldsymbol{\theta}(D_1)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w; \boldsymbol{\theta}(D_2)) \\ &= \sum_{v \in \mathcal{N}(\boldsymbol{x}_0) \cap D_1} m(\boldsymbol{x}_v \boldsymbol{\theta}(D_1)) + \sum_{w \in \mathcal{N}(\boldsymbol{x}_0) \cap D_2} m(\boldsymbol{x}_w \boldsymbol{\theta}(D_2)). \end{aligned} \tag{18}$$

We now define the adjacent datasets D and D' as follows. Assume $D = D_1$; one data point $\{\boldsymbol{x}', \boldsymbol{y}'\}$ is then randomly selected from Chunk I and removed $\{\boldsymbol{x}', \boldsymbol{y}'\}$ from D_1 forming $D' = D_1 \backslash \{\boldsymbol{x}', \boldsymbol{y}'\}$. Meanwhile, the unlabeled set $D_{\mathcal{U}}$ and D_2 remain the same. Our goal is to examine the extent to which $\boldsymbol{E}$ and $\boldsymbol{F}$ may leak information pertaining to $\{\boldsymbol{x}', \boldsymbol{y}'\}$ when $\{\boldsymbol{x}', \boldsymbol{y}'\}$ is removed from D_1 as described above.

Denote $\boldsymbol{x}_v, \boldsymbol{x}_w$ as training data in chunk I and chunk II. Assume $\begin{pmatrix} \boldsymbol{x}_v \\ \boldsymbol{x}_w \end{pmatrix}$ is drawn from a multivariate Gaussian distribution:

$$\begin{pmatrix} \boldsymbol{x}_v \\ \boldsymbol{x}_w \end{pmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \quad \begin{bmatrix} \boldsymbol{\Gamma}_{vv} & \boldsymbol{\Gamma}_{vw} \\ \boldsymbol{\Gamma}_{wv} & \boldsymbol{\Gamma}_{ww} \end{bmatrix}^{-1} \right), \tag{19}$$

where $\begin{bmatrix} \boldsymbol{\Gamma}_{vv} & \boldsymbol{\Gamma}_{vw} \\ \boldsymbol{\Gamma}_{wv} & \boldsymbol{\Gamma}_{ww} \end{bmatrix} = a\boldsymbol{I} + b\boldsymbol{N}$, $\boldsymbol{I}$ is identical matrix, $\boldsymbol{N}$ is normalized graph Laplacian, a controls noise level and b controls the smoothness over the whole graph.

From Proposition 5, the condition distribution of $\boldsymbol{x}_w$ given $\boldsymbol{x}_v = \boldsymbol{x}_v$ can be written as

$$\boldsymbol{x}_w | \boldsymbol{x}_v = \boldsymbol{x}_v \sim \mathcal{N}(-\boldsymbol{\Gamma}_{ww}^{-1}\boldsymbol{\Gamma}_{wv}\boldsymbol{x}_v, \boldsymbol{\Gamma}_{ww}^{-1}).$$

Condition on the data D_1, the distribution of D_2 is a conditional multivariate Gaussian distribution with mean $-\boldsymbol{\Gamma}_{ww}^{-1}\boldsymbol{\Gamma}_{wv}\boldsymbol{x}_v$ and variance $\boldsymbol{\Gamma}_{ww}^{-1}$. Furthermore, multiplicative Gaussian distribution $\boldsymbol{x}_w\boldsymbol{\theta}(D_1)$ introduces a Gaussian random noise into (17). According to Proposition (1) and (3), $\boldsymbol{E}$ satisfies $(\frac{1}{2}, \frac{1}{4\sigma^2 L^2} + \frac{1}{2L^2})$-Rényi Differential Privacy, where σ^2 depends on Γ_{ww}^{-1} decided by graph structure.

Meanwhile, although (18) is deterministic, we can manually add Gaussian noise $\mathcal{N}(0, \sigma^2)$ such that $\boldsymbol{F}$ satisfies $\frac{\alpha}{2\sigma^2}$-Rényi Differential Privacy via Proposition (4). However, if we then let $\sigma \to 0$ to reproduce $\boldsymbol{F}$, we have that $\frac{\alpha}{2\sigma^2} \to \infty$, indicating that in fact $\boldsymbol{F}$ has no privacy guarantee. □

B Experiment Details

B.1 Data descriptions

House: node features are the property of house, edges connect the neighbors, the task is to predict the price of the house. **County**: each node is a county and edges connect two counties sharing a border, the task is to predict the unemployment rate for a county. **VK**: each node is a person and edges connect two people based on the friendships, the task is to predict the age of each person. **Avazu**: each node is a device and edges connect two devices if they appear on the same site with the same application, the target is the click-through-rate of a node. For **House**, **County**, **VK** and **Avazu** datasets, Training/validation/testing are randomly split with 6/2/2 ratio and all experiments results are averaged over 5 trails.

OGB-Arxiv, OGB-Products are standard datasets from OGB-leaderboards and all training/validation/testing splits follow the standard data splitting from OGB-leaderboards. Reddit is standard datset from Deep Graph Library (DGL).

B.2 Base models

Specifically, we consider LightGBM boosted Tress (GBM) (Ke et al., 2017), CatBoost boosted trees (CAT) (Prokhorenkova et al., 2018), fully-connected neural networks (NN), Extremely Randomized Trees (RT), Random Forests (RF), K Nearest Neighbors (KNN), Label Propagation (LP) (Huang et al., 2020) and Transformer with electra pretrained model (Text) (Training epoch is 12) (Clark et al., 2020). For the first layer, we keep the typical models, for example, Gradient Boosted Decision Trees for Tabular data, Transformer models for text data. For second stacking layer, we use all of models except extremely low-efficient models for large dataset, for example, KNN and Catboost slow down the training procedure for OGB-products dataset. All details about the base models can be found in table 5. The parameters about all models can be referred to AutoGluon (Erickson et al., 2020).

B.3 Parameters for Graph-aware propagation

We do graph-aware propagation for the prediction to incorporate the graph structure. Table 6 shows two hyperparameters considered in the propagation part: weight λ and number of propagation step T. We also present the hyperparameters for Correct and Smooth in Table 7.

Table 5: Base models

DATASET	FIRST LAYER	SECOND LAYER
HOUSE/COUNTY/VK/AVAZU	CAT, GBM, NN	KNN, GBM, RF, RT, CAT, NN
OGB-ARXIV	TEXT, GBM, NN	GBM, RF, RT, NN
OGB-PRODUCTS	TEXT, LP	GBM, RF, RT, NN
REDDIT	NN, RF, RT, GBM	NN, RF, RT, KNN, GBM

Table 6: Hyperparameters

DATASET	λ	INPUT FOR STACKING LAYER
HOUSE/COUNTY/VK/AVAZU	0.9	$(\boldsymbol{X}, \{\boldsymbol{F}_m^{(0)}, \boldsymbol{F}_m^{(1)}, \boldsymbol{F}_m^{(2)}, \boldsymbol{F}_m^{(3)}, \boldsymbol{F}_m^{(4)}, \boldsymbol{F}_m^{(5)}\})$
OGB-ARXIV	0.95	$(\boldsymbol{X}, \{\boldsymbol{F}_m^{(0)}, \boldsymbol{F}_m^{(1)}, \boldsymbol{F}_m^{(3)}, \boldsymbol{F}_m^{(5)}, \boldsymbol{F}_m^{(7)}, \boldsymbol{F}_m^{(9)}\})$
OGB-PRODUCTS	0.97	$(\boldsymbol{X}, \{\boldsymbol{F}_m^{(0)}, \boldsymbol{F}_m^{(1)}, \boldsymbol{F}_m^{(3)}, \boldsymbol{F}_m^{(5)}, \boldsymbol{F}_m^{(7)}, \boldsymbol{F}_m^{(9)}\})$
REDDIT	0.95	$(\boldsymbol{X}, \{\boldsymbol{F}_m^{(0)}, \boldsymbol{F}_m^{(1)}, \boldsymbol{F}_m^{(2)}, \boldsymbol{F}_m^{(3)}, \boldsymbol{F}_m^{(4)}, \boldsymbol{F}_m^{(5)}\})$

Table 7: Hyperparameters for C&S

DATASET	λ_1	KERNEL TYPE	λ_2	KERNEL TYPE	NUM_PROPAGATION
HOUSE/COUNTY/AVAZU	0.8	DA	0.5	DA	5
VK	0.8	DA	-	-	5
OGB-ARXIV	0.9	DA	0.1	AD	50
OGB-PRODUCTS	0.3	DAD	0.3	AD	50

Table 8: Training time tested on AWS g4dn.12xlarge machine.

DATASET	BASE MODEL	TIME(S)
HOUSE	GBM, NN	52
COUNTY	GBM, NN	18
VK	GBM, NN	119
AVAZU	GBM, NN	15
OGB-ARXIV	NN	199
OGB-PRODUCTS	NN	837

B.4 COMPUTING COST

The computing cost depends on the ensemble models we select (e.g., transformer models can take more computing resources relying on the implementation, including more emsemble models leads to more computing cost). So it's hard to consistently measure the training/inference time or memory consumption. But the computing cost is in a competitive range since the integration of the bagging and ensembling parts key to our model can be efficiently implemented, e.g., via open source packages like AutoGluon that we used. In Table 8, we present the training time of different datasets with basic ensemble models. For instance, the training time for OGB-products with OGB embeddings is around 800s, while for GraphSage it is about 1000s for 100 epochs.

ADDITIONAL REFERENCES FOR THE SUPPLEMENTARY

K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*, 2020.

N. Erickson, J. Mueller, A. Shirkov, H. Zhang, P. Larroy, M. Li, and A. Smola. Autogluon-tabular: Robust and accurate automl for structured data. *arXiv preprint arXiv:2003.06505*, 2020.

Q. Huang, H. He, A. Singh, S.-N. Lim, and A. R. Benson. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*, 2020.

J. Jia and A. R. Benson. A unifying generative model for graph learning algorithms: Label propagation, graph convolutions, and combinations. *arXiv preprint arXiv:2101.07730*, 2021.

G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, 2017.

I. Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin. CatBoost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, 2018.