Graph Belief Propagation Networks

Junteng Jia
Cornell University
jj585@cornell.edu

Cenk Baykal
MIT
baykal@mit.edu

Vamsi K. Potluru
J.P. Morgan AI Research
vamsi.k.potluru@jpmchase.com

Austin R. Benson
Cornell University
arb@cs.cornell.edu

Abstract

With the wide-spread availability of complex relational data, semi-supervised node classification in graphs has become a central machine learning problem. Graph neural networks are a recent class of easy-to-train and accurate methods for this problem that map the features in the neighborhood of a node to its label, but they ignore label correlation during inference and their predictions are difficult to interpret. On the other hand, collective classification is a traditional approach based on interpretable graphical models that explicitly model label correlations. Here, we introduce a model that combines the advantages of these two approaches, where we compute the marginal probabilities in a conditional random field, similar to collective classification, and the potentials in the random field are learned through end-to-end training, akin to graph neural networks. In our model, potentials on each node only depend on that node's features, and edge potentials are learned via a coupling matrix. This structure enables simple training with interpretable parameters, scales to large networks, naturally incorporates training labels at inference, and is often more accurate than related approaches. Our approach can be viewed as either an interpretable message-passing graph neural network or a collective classification method with higher capacity and modernized training.

1 Bridging Graph Neural Networks and Collective Classification

Graphs are a natural model for systems with interacting components, where the nodes represent individual objects and the edges represent their interactions [1, 2]. For example, a social network might be modeled as a graph where the nodes are users and the edges are friendships. Oftentimes, the nodes have attributes, such as a user's age, gender, or occupation in a social network, but attributes are typically incomplete due to difficulties in data collection or privacy. *Graph-based semi-supervised learning*, also called *node classification*, addresses this problem by predicting a missing attribute (i.e., label y) on some nodes given other attributes (i.e., features x), and has been used in a variety of applications, such as product category identification and protein function prediction [3, 4, 5, 6, 7].

Graph neural networks (GNNs) are a common method for semi-supervised learning on graphs [5, 8, 9, 10]. A GNN first summarizes the features and graph structure in the neighborhood of each node into a vector representation. After, each node's representation is used independently for classification. Automatic differentiation enables end-to-end training, and there are simple sub-sampling schemes to handle large graphs [6]. However, this approach implicitly assumes that node labels are conditionally independent given all features, and standard GNNs do not (directly) use correlations between training and testing labels during inference. Moreover, GNNs consist of transformation and aggregation functions parametrized by neural networks and the learned models are difficult to interpret.

On the other hand, collective classification (CC) is a class of interpretable methods based on graphical models that directly leverage label correlation for prediction [11, 12, 13, 14, 15, 16, 17]. The statistical assumptions behind CC models are also arguably more appropriate than GNNs for graph data. For instance, relation Markov networks (RMNs) [17] model the joint distribution of all node labels with a conditional random field and predict an unknown label with its marginal probabilities. Leveraging label correlation during inference then simply amounts to conditioning on the training labels. However,

Preprint. Under review.

the interpretability and convenience come with a cost. Collective classification models are learned by maximizing the joint likelihood, making end-to-end training extremely challenging. The difficulty in learning model parameters, in turn, limits the capacity of the model itself.

The fact that GNNs and CC are two solutions to a same problem inspires a series of research questions. Is there a GNN architecture that learns the joint label distribution? Can we improve the modeling capacity and training of CC without hurting its interpretability? Finally, what type of approach would bridge those two largely independently developed fields?

Here, we answer these questions in the affirmative by developing graph belief propagation networks (GBPNs). Our main idea is to learn the joint label distribution by maximizing the marginal likelihood of individual labels. Specifically, we model the node labels in a graph as a conditional random field, where the coupling-potential on each edge is the same, but the self-potential on each node only depend on its features. As such, our model is specified by a matrix for the coupling-potential that describes the affinity for nodes with certain labels to be connected, and a multi-layer perceptron (MLP) that maps a node's features to its self-potential. To compute the marginal probabilities for each node label, we use loopy belief propagation (BP). Putting everything together, we have a two-step inference algorithm that can be trained end-to-end: ① compute the self-potential on each node with an MLP, and ② estimate the marginal probabilities with a fixed number of BP iterations.

In one sense, GBPN is a message-passing GNN that projects node features to a dimension equal to the number of classes upfront, and subsequently use the projected features for propagation and aggregation. On the other hand, GBPN learns the joint distribution of node labels in a graph similar to CC, and outputs the corresponding marginal probabilities for each node as a prediction. The prediction accuracy is oftentimes higher than GNNs on benchmark datasets, and the learned coupling matrix can be used to interpret the data and predictions. We also show this approach leads to straightforward mini-batch training so that the methods can be used on large graphs.

One issue is whether the marginal probabilities estimated by BP are good approximations on graphs with loops. Empirically, BP often converges on graphs with cycles with estimated marginals that are close to the ground truth [18]. In our context, we find that BP converges in just a few iterations on several real-world graph datasets. Another possible concern is whether learning a joint distribution by maximizing the marginal likelihood produces good estimates. We show that on synthetic sampled from Markov random fields (MRFs), this approach recovers the parameters used to generate the data.

1.1 Additional related work

Collective classification. Collective classification encompasses many machine learning algorithms for relational data that use correlation in labels among connected entities in addition to features on the entities [11, 12, 13]. For example, local conditional classifiers iteratively make predictions and update node features based on these predictions [14, 15, 16, 19]. Closer to our methods are approaches that use pairwise MRFs and belief propagation [17]. The major difference is that they estimate the MRF coupling-potentials by maximizing the joint likelihood of all the nodes, and therefore require a fully labeled graph for training (see Section 2.3). In contrast, our model is learned discriminatively end-to-end on the training nodes by back-propagating through belief propagation steps.

MRFs with GNNs. A few approaches combine MRFs and graph neural networks [20, 21, 22]. They add a conditional random field layer to model the joint distribution of the node labels and maximize the joint probability of the observed labels during training. Variational methods or pseudolikelihood approaches are used to optimize the joint likelihood. Our GBPN model avoid these difficulties by optimizing for the marginal likelihood, and this leads to more accurate predictions.

Direct use of training labels. Besides MRFs, other GNN approaches have used training labels at inference in graph-based learning. These include diffusions on training labels for data augmentation [23] and post-processing techniques [24, 25, 26]. Similar in spirit, smoothing techniques model positive label correlations in neighboring nodes [27, 28]. Our approach is more statistical, as GBPN just conditions on known labels.

Learning better belief propagation algorithms. Some recent research learns a belief propagation algorithm to estimate the partition function in a graphical model [29, 30]. In comparison, we focus on the node classification problem for real-world graphs.

2 Node Classification as Statistical Inference in a Markov Random Field

Consider an attributed graph $G(V, E, \boldsymbol{X}, \boldsymbol{y})$ where V is a set n of nodes, E is a set of undirected edges, $\boldsymbol{X} \in \mathbb{R}^{n \times d}$ is a matrix whose rows correspond to node features, and $\boldsymbol{y} \in \{0, \dots, c-1\}^n$ is a vector of node labels in one of c possible classes. We partition the vertices into two disjoint sets $L \cup U = V$, where L is the training set (labeled) and U is the testing set (unlabeled). In the *node classification* task, we assume access to all node features $\{\boldsymbol{x}_i\}_{i \in V}$ and the labels on the training nodes $\{y_i\}_{i \in L}$, and the goal is to predict the testing labels $\{y_i\}_{i \in U}$. The basic idea of our approach is to define a joint distribution over node features and labels $p(\boldsymbol{X}, \boldsymbol{y})$ and then cast the node classification task as computing the marginal probability $p(y_i | \boldsymbol{X}, \boldsymbol{y}_L)$.

2.1 A Data Generation Process for Node Attributes

We treat the graph as deterministic but the node attributes randomly as follows: ① Jointly sample the node labels $p(\boldsymbol{y})$; ② Given the node labels $\boldsymbol{y}$, sample the features from $p(\boldsymbol{X}|\boldsymbol{y})$. This generation process defines a joint distribution $p(\boldsymbol{X}, \boldsymbol{y})$. Furthermore, we add the following assumptions.

- First, the joint distribution of node labels $p(\boldsymbol{y})$ is a pairwise Markov random field,

$$p(\boldsymbol{y}) = \frac{\varphi(\boldsymbol{y})}{\sum_{\boldsymbol{y}'} \varphi(\boldsymbol{y}')}, \quad \varphi(\boldsymbol{y}) = \prod_{i \in V} h_i(y_i) \prod_{(i,j) \in E} H_{ij}(y_i, y_j), \tag{1}$$

  and the coupling-potential on every edge is the same and symmetric, i.e. $H_{ij}(y_i, y_j) = H_{y_i, y_j} = H_{y_i, y_j}$, where $\boldsymbol{H} \in \mathbb{R}_+^{c \times c}$. This assumption implies each node is directly coupled only with its neighbors, a condition known as the Markov propriety in probabilistic graph models [31].

- Second, The conditional distribution $p(\boldsymbol{X}|\boldsymbol{y})$ can be factorized as

$$p(\boldsymbol{X}|\boldsymbol{y}) = \prod_{i \in V} p(\boldsymbol{x}_i | y_i) = \prod_{i \in V} f(y_i; \boldsymbol{x}_i), \tag{2}$$

  which means the features on each node is independently sampled given its label.

These two assumptions lead to simple and efficient inference algorithms, as we will show next.

2.2 Node Classification Models Derived through Probabilistic Inference

Following the generative model defined above, we derive the posterior probability as:

$$p(\boldsymbol{y}|\boldsymbol{X}) = \frac{p(\boldsymbol{X}, \boldsymbol{y})}{p(\boldsymbol{X})} = \frac{p(\boldsymbol{y})p(\boldsymbol{X}|\boldsymbol{y})}{\sum_{\boldsymbol{y}'} p(\boldsymbol{y}')p(\boldsymbol{X}|\boldsymbol{y}')} \cong \prod_{i \in V} f(y_i; \boldsymbol{x}_i) \, h_i(y_i) \prod_{(i,j) \in E} H_{ij}(y_i, y_j) \tag{3}$$

$$= \prod_{i \in V} g(y_i; \boldsymbol{x}_i) \prod_{(i,j) \in E} H_{ij}(y_i, y_j), \tag{4}$$

which is a conditional random field [32], where $\cong$ denotes "equality up to a normalization" [33]. We directly parametrize this posterior distribution using a symmetric matrix $\boldsymbol{H}$ for the coupling-potential, and a neural network function g_θ to map from a node's features to its self-potential. Similar modeling assumptions have been considered in collective classification [17, 34], where the models are trained to maximize the joint likelihood. In contrast, our model is trained on the marginal likelihood of individual labels. Next, we introduce the inference algorithm for the marginal probability $p(y_i | \boldsymbol{X})$. The algorithm only use node features for label prediction, and (once trained) can be used to predict labels in a unseen graph. Later, we consider the inference algorithm for $p(y_i | \boldsymbol{X}, \boldsymbol{y}_L)$, which gives more accurate predictions in the transductive learning setting.

Inference. Exact inference for the marginal probabilities is intractable on large-scale graphs with cycles [31]. Therefore, we resort to approximate inference using the loopy belief propagation (BP) algorithm [18]. In particular, we estimate the marginal probabilities using a two-step algorithm.

① Compute the self-potentials using a multi-layer perceptron (MLP) with parameters θ,

$$g_\theta(y_i; \boldsymbol{x}_i) = \mathtt{Cat}(y_i | \mathtt{softmax}(\mathtt{MLP}_\theta(\boldsymbol{x}_i))). \tag{5}$$

② Run T iterations of belief-propagation,

$$p_i^{(t)}(y_i) \cong p_i^{(0)}(y_i) \prod_{j \in N(i)} m_{ji}^{(t)}(y_i), \quad m_{ji}^{(t)}(y_i) \cong \sum_{y_j} H_{ji}(y_j, y_i) \frac{p_j^{(t-1)}(y_j)}{m_{ij}^{(t-1)}(y_j)}, \tag{6}$$

  where $t = 1, \dots, T$ and $N(i)$ are the neighbors of node i. The initial condition is given by,

$$p_i^{(0)}(y_i) = g(y_i; \boldsymbol{x}_i), \quad m_{ji}^{(0)}(y_j, y_i) = 1/c. \tag{7}$$

3

In other words, we set the initial belief on each node to its self-potential and update it with messages from the neighbors. The final belief $p_i(y_i) = p_i^{(T)}(y_i)$ is used to approximate $p(y_i|\boldsymbol{X})$.

In practice, for numerical stability, we perform belief-propagation updates in the log-space.

$$\log p_i^{(t)}(y_i) \cong \log p_i^{(0)}(y_i) + \sum_{j \in N(i)} \log m_{ji}^{(t)}(y_i), \tag{8}$$

$$\log m_{ji}^{(t)}(y_i) \cong \text{LSE}_{y_j} \left[\log H_{ji}(y_j, y_i) + \log p_j^{(t-1)}(y_j) - \log m_{ij}^{(t-1)}(y_j) \right], \tag{9}$$

where LSE stands for the log-sum-exp function: $\text{LSE}_{y_j} \left[f(y_i, y_j) \right] = \log \left[\sum_{y_j} \exp \left(f(y_i, y_j) \right) \right]$.

Learning. We optimize the model by maximizing the log marginal-likelihood of the training labels,

$$\theta^*, \boldsymbol{H}^* = \arg\max_{\theta, \boldsymbol{H}} \sum_{i \in L} \log p(y_i|\boldsymbol{X}) \approx \arg\max_{\theta, \boldsymbol{H}} \sum_{i \in L} \log p_i(y_i). \tag{10}$$

The gradients with respective to the parameters are computed with auto-differentiation by backpropagating through the BP iterations. Since the inference algorithm only uses node features, we can learn the parameters in one graph with abundant labels, and predict for another graph where the labels are difficult to obtain. As such, we denote this algorithm as the inductive variant of GBPN, or GBPN-I.

Direct use of training labels for inference. In the transductive learning setting, both the training and testing nodes are from the same graph. Therefore, we can directly leverage the training labels during inference. From the probabilistic inference perspective, we start with a joint distribution over testing labels by conditioning on both the features and the training labels, i.e.,

$$p(\boldsymbol{y}_U|\boldsymbol{X}, \boldsymbol{y}_L) \cong \prod_{i \in U} \left[g(y_i; \boldsymbol{x}_i) \prod_{i \in U, \ j \in L, \ (i,j) \in E} H_{ij}(y_i, y_j) \right] \prod_{i,j \in U, \ (i,j) \in E} H_{ij}(y_i, y_j). \tag{11}$$

This is nothing but another conditional random field defined on the induced subgraph $G[U]$, where the self-potentials are modified to incorporate the couplings between the training and testing nodes in the original graph. Then, similar to the inductive algorithm, we estimate the marginal probabilities by belief propagation. To better mimic how we predict the testing labels, during each training step, we randomly select half of the training nodes to be conditioned on, and optimize the log-likelihood loss on the other half. We denote this transductive variant of inference and training as GBPN.

2.3 Comparison with Other Node Classification Methods

Now that our method is fully specified, we comparing it with existing methods in terms of statistical assumption, model parametrization, training/inference objective and optimization algorithm.

GNNs [35] independently predict the label on each node using node features in the neighborhood. The implicit statistical assumption is that node labels are conditional independent given the features:

$$p(\boldsymbol{y}|\boldsymbol{X}) = \prod_{i \in V} p(y_i|\boldsymbol{X}) \tag{12}$$

Therefore, GNNs ignore label correlation and directly parametrize $p(y_i|\boldsymbol{X})$ to extract the most information from features. In contrast, GBPN models $p(\boldsymbol{y}|\boldsymbol{X})$ and uses label correlation for inference.

RMN [17] is a representative example for collective classification methods [11] that directly maximize the joint likelihood $p(\boldsymbol{y}|\boldsymbol{X})$. Those models only work for the inductive setting, where they are first trained on a fully labeled graph and then used to predict node labels in a different graph. In comparison, our model can be trained on a partially labeled graph, and it is suitable for both the inductive and transductive setting. Moreover, even on a graph that is fully labeled, methods like RMN are historically difficult to train. This is because the gradient of the joint likelihood involves a normalization factor and is therefore hard to compute. GBPN avoids this problem by differentiating through the belief-propagation iterations with backpropagation.

Pseudolikelihood EM [34] was designed to learn the correlation structure in a partially labeled graph by maximizing the evidence lower bound (ELBO) for the joint likelihood of the training labels $p(\boldsymbol{y}_L|\boldsymbol{X})$. Maximizing the ELBO boils down to directly optimizing two marginal distributions: a variational distribution q_θ and a pseudolikelihood p_ϕ used to approximate the posterior. Therefore, if p_ϕ and q_θ are neural networks (as in the GMNN model [20]), then the entire model can be trained with auto-differentiation. The learned p_ϕ can leverage training labels during inference, akin to GBPN. However, while GMNN models label correlation with neural networks, GBPN explicitly specifies the coupling-potential and is thus more interpretable. Moreover, we compare GBPM with GMNN in section 3.1 and find GBPN gives higher prediction accuracies on benchmark datasets.

Table 1: Summary of node classification methods from the probabilistic inference perspective.

method	assumption	model parametrization	training	inference	optimization
GBPN	eq. (4)	$p(\boldsymbol{y}|\boldsymbol{X})$	$p(y_i|\boldsymbol{X}, \boldsymbol{y}_L)$	$p(y_i|\boldsymbol{X}, \boldsymbol{y}_L)$	auto-diff
GNN	eq. (12)	$p(y_i|\boldsymbol{X})$	$p(y_i|\boldsymbol{X})$	$p(y_i|\boldsymbol{X})$	auto-diff
RMN	eq. (4)	$p(\boldsymbol{y}|\boldsymbol{X})$	$p(\boldsymbol{y}|\boldsymbol{X})$	$p(y_i|\boldsymbol{X}, \boldsymbol{y}_L)$	manual-diff
GMNN	eq. (4)	$p_\phi(y_i|\boldsymbol{X}, \boldsymbol{y}_{N(i)}), q_\theta(y_i|\boldsymbol{X})$	$p(\boldsymbol{y}_L|\boldsymbol{X})$	$p(y_i|\boldsymbol{X}, \boldsymbol{y}_L)$	EM (auto-diff)

2.4 Practical Considerations for Training GBPNs

One modeling choice we have to make for GBPN is the loss function for training. Although the (unweighted) negative log-likelihood loss is often used for training GNNs, the same choice for GBPN can lead to sub-optimal performance in practice, especially when the degree heterogeneity in a graph is large. The reason is that, unlike GNNs, the predictive confidence of GBPN on a node can grow exponentially with the node degree, as shown in fig. 1. Therefore, the negative log-likelihood loss associated with mis-classifying a node grows linearly with the node degree. This creates a potential misalignment between the loss function and the classification accuracy, as classification accuracy is not weighted based on node degree. We will demonstrate this misalignment empirically in Section 3.3.
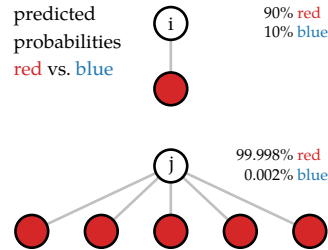
Figure 1: Predicted probabilities for low- *vs.* high-degree nodes.

We employ a straightforward way to align the loss function to the accuracy measure, which is to reweight the marginal likelihood in the loss function eq. (10) as $\log p_i'(y) \cong \alpha_i \cdot \log p_i(y)$. Here, α_i reduces the influence of high-degree nodes. In practice, we find that $\alpha_i = d(i)^{-1/2}$ gives consistently better performance than $\alpha_i = 1$ (no reweighting), where $d(i)$ is the degree of node i.

One limitation of GBPN is that it requires storing all the messages between neighbors during training, which costs $\mathcal{O}(|E|)$ per BP step. In comparison, GCN [5] and GraphSAGE [6] costs $\mathcal{O}(|V|)$ per layer, and GAT [8] costs $\mathcal{O}(|E|)$ per layer due to edge attention weights. To reduce the computational cost on large-scale network, we will introduce a mini-batch training algorithm for GBPN in Section 3.3.

3 Experiments

Now, we compare GBPN against GNNs empirically on several synthetic and real-world datasets, and also interpret the learned GBPN models, as well as discuss new sub-sampling techniques for large networks. The proposed models are summarized with pseudocode in appendix A.1 and implemented in PyTorch Geometric [36]. The source code, data, and experiments are publicly available online at `https://github.com/000Justin000/GBPN.git`. Statistics of all the datasets are in table 2 and we describe them below.

Synthetic MRF datasets. We sample data from Markov random fields, i.e., from our generative model. Specifically, we use a 51×51 grid graph, where each node belongs to one of two possible classes. The neighboring node labels are likely to be the same if the coupling-potential is positive (denoted +), and likely to be different if the coupling-potential is negative (denoted −). This setup is known as the Ising model in statistical physics and models homophily or heterophily in social networks [37], or assortativitity and disassortativity more generally [38]. We also consider a setup with three classes with either positive or negative coupling-potentials. In all settings, we use the grid coordinates as features to predict labels.

Table 2: Summary of dataset statistics.

	# nodes	# edges	# features	# classes
Ising+	2.6×10^3	1.0×10^4	2	2
Ising−	2.6×10^3	1.0×10^4	2	2
MRF+	2.6×10^3	1.0×10^4	2	3
MRF−	2.6×10^3	1.0×10^4	2	3
Cora	2.7×10^3	5.3×10^3	1433	7
CiteSeer	3.3×10^3	4.6×10^3	3703	6
PubMed	2.0×10^4	4.4×10^4	500	3
CS	1.8×10^4	8.2×10^4	6805	15
Physics	3.4×10^4	2.5×10^5	8415	5
Election	3.1×10^3	2.3×10^3	9	2
Sexual	1.9×10^3	2.1×10^3	20	2
Elliptic	2.0×10^5	2.3×10^5	165	2
Payments	2.2×10^5	2.8×10^5	573	2
arXiv	1.7×10^5	1.2×10^6	128	40
Products	2.4×10^6	6.2×10^7	100	47

Table 3: Node classification performances on synthetic data with full-batch training. The synthetic is sampled from our generative model, and, our algorithms are appropriately more accurate.

dataset	MLP	GCN	SAGE	GAT	GMNN	DeeperGNN	GBPN-I	GBPN
Ising+	$67.1_{\pm 1.6}$	$67.1_{\pm 1.8}$	$68.2_{\pm 1.5}$	$64.4_{\pm 2.2}$	$64.4_{\pm 3.1}$	$65.6_{\pm 1.6}$	$68.5_{\pm 1.8}$	$\mathbf{75.0}_{\pm 1.4}$
Ising−	$48.8_{\pm 0.8}$	$48.7_{\pm 0.7}$	$48.4_{\pm 1.0}$	$49.3_{\pm 0.9}$	$49.4_{\pm 0.9}$	$49.3_{\pm 1.1}$	$48.3_{\pm 1.1}$	$\mathbf{72.7}_{\pm 1.4}$
MRF+	$64.2_{\pm 1.9}$	$65.9_{\pm 2.0}$	$64.8_{\pm 1.9}$	$61.4_{\pm 2.3}$	$45.8_{\pm 2.9}$	$53.3_{\pm 3.8}$	$66.5_{\pm 2.0}$	$\mathbf{70.3}_{\pm 1.7}$
MRF−	$64.3_{\pm 4.1}$	$66.3_{\pm 5.0}$	$65.3_{\pm 4.6}$	$62.3_{\pm 4.0}$	$41.3_{\pm 5.3}$	$53.1_{\pm 2.9}$	$66.3_{\pm 5.0}$	$\mathbf{73.5}_{\pm 4.5}$

Citation graphs. We use the small benchmark graphs Cora, CiteSeer and PubMed [11], as well as the larger arXiv benchmark [39], where the vertices represent articles and edges represent citations. Features are derived from the article text, and the goal is to predict the research field of each article.

Co-authorship graphs. These are graphs of computer scientists and physicists [40], where nodes are researchers, edges connect researchers that have coauthored a paper, node features are paper keywords for each author, and the goal is to predict the most active field of study for each author.

U.S. election. This dataset comes from the 2016 presidential election, where nodes are U.S. counties, and edges connect counties with the strongest Facebook social ties [26]. Each node has county-level demographic features (e.g., median income) and social network user features (e.g., fraction of friends within 50 miles). The goal is to predict the party (Republican or the Democrat) that won each county.

Sexual interactions. This is a social network where connections come from sexual interactions [24, 41]. Node features include occupation and retirement status, and the goal is to predict gender.

Financial transactions. We use two financial networks, where the nodes represent transactions and the edges are payment flows. Node features are metadata associated with each transaction. The Elliptic dataset is derived from bitcoin transactions [42], and transactions are labeled as licit or illicit. The Payments dataset consists of synthetic payments that resemble real-world data, provided by J.P. Morgan Chase & Co. Transactions are labeled as fraudulent or non-fraudulent.

Co-purchasing. Here, nodes are products on Amazon, and edges connect co-purchased products. The features are the average word embeddings from the descriptions and labels are product categories.

3.1 Performance on Smaller Datasets

The datasets we consider here fits into three groups: ① small synthetic graphs; ② small real-world graphs; ③ large-scale graphs. We first consider ① and ②, where training the entire graph can fit into a single GPU. On these datasets, all models take at most 20 minutes to train on a Telsa V100 GPU.

For GBPN, we use a 2-layer MLP with 256 hidden units to map the features of each node to its self-potentials. Then we run five steps of BP iterations to compute the predictions. For standard baselines, we consider a 2-layer MLP, GCN, and GraphSAGE models with the 256 hidden units, as well as a 2-layer GAT with four attention heads × 64 hidden units per layer. We also compare GBPN with GMNN [20], which is also based on Markov random fields, as well as DeeperGNN [43], which uses several hops of neighborhood information at anode. For both of these baselines, we use the reference implementations[1] with default hyperparameters for training and inference.

We randomly split all datasets into 30% training, 20% validation, and 50% testing. All baseline methods are trained by minimizing the negative log-likelihood (NLL), while GBPN is trained with the weighted NLL loss discussed in Section 2.4. For all methods, we use an AdamW optimizer with learning rate 1.0×10^{-3} and decay rate 2.5×10^{-4} to perform full-batch training for 500 steps. The testing accuracies are averaged over 30 runs and summarized in Tables 3 and 4.

On the synthetic networks sampled from MRFs, the transductive GBPN outperforms the other methods by large margins (Table 3). The performance gap is due to two reasons. First, the underlying data assumption of our model exactly matches the data. Second, GBPN is able to directly use training labels for inference. On the empirical datasets, inductive and transductive GBPN outperform the baselines in five of seven graphs (Tables 3 and 4). GBPN is more accurate on most of the coauthorship and citation graphs, although the gains over the best baslines are modest. The differences on the Election and Sexual datasets brings additional insights. The node labels in both datasets are binary; however, election outcomes are homophilous, and the sexual interactions are heterophilous (most

---

[1]DeeperGNN is released under GNU General Public License while GMNN is under no license.

Table 4: Node classification performances on smaller empirical datasets with full-batch training.

dataset	MLP	GCN	SAGE	GAT	GMNN	DeeperGNN	GBPN-I	GBPN
Cora	72.1 ± 1.3	87.1 ± 0.7	86.9 ± 0.8	87.1 ± 0.9	86.4 ± 0.9	**87.2** ± 0.8	85.6 ± 0.7	86.4 ± 0.9
CiteSeer	71.2 ± 0.9	73.5 ± 1.0	73.2 ± 1.0	73.1 ± 1.2	72.9 ± 1.2	73.9 ± 0.8	74.7 ± 1.3	**74.8** ± 0.8
PubMed	86.5 ± 0.2	87.1 ± 0.3	87.8 ± 0.4	88.1 ± 0.3	86.7 ± 0.2	84.7 ± 0.3	88.4 ± 0.3	**88.5** ± 0.3
CS	94.2 ± 0.2	93.2 ± 0.2	93.7 ± 0.2	94.0 ± 0.3	93.3 ± 0.3	94.9 ± 0.2	**95.5** ± 0.2	**95.5** ± 0.3
Physics	95.8 ± 0.1	96.1 ± 0.1	96.3 ± 0.1	96.3 ± 0.1	96.1 ± 0.1	96.7 ± 0.1	**96.9** ± 0.1	**96.9** ± 0.1
Election	89.6 ± 0.6	88.0 ± 0.6	**90.8** ± 0.6	90.5 ± 0.7	87.3 ± 0.7	85.4 ± 0.7	90.1 ± 0.8	90.3 ± 0.9
Sexual	74.5 ± 1.4	83.9 ± 1.2	93.3 ± 0.8	93.6 ± 0.9	77.0 ± 1.7	65.0 ± 1.1	97.1 ± 0.5	**97.4** ± 0.4
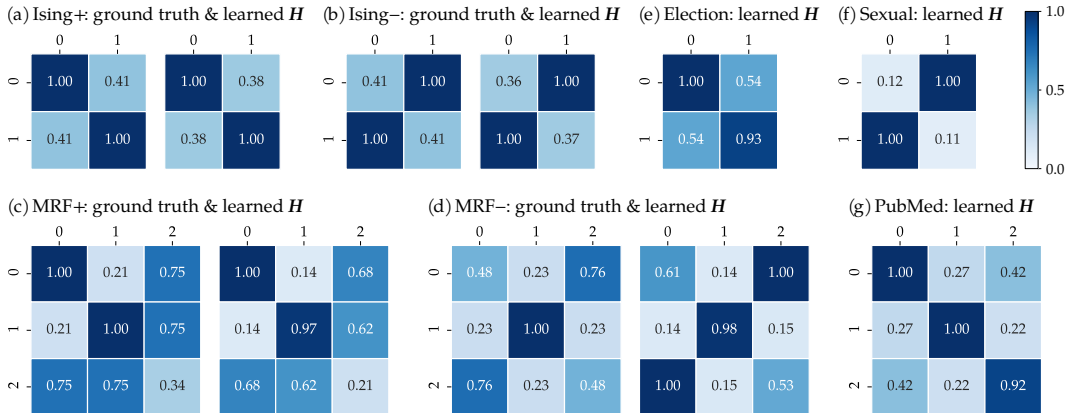
Figure 2: Visualization of the learned coupling matrices. (a–d) Synthetic datasets, with ground truth coupling matrices on the left and learned coupling matrices on the right. (e–g) Election, Sexual, and PubMed empirical datasets, where we recover homophilous, heterophilous, and mixed affinities.

relationships used to construct the dataset are heterosexual). GBPN is comparable to the best baseline on the Election network (GraphSAGE in this case), it outperforms the baselines by large margins on the Sexual network. This is because GBPN can explicitly learn heterophily with the coupling-potential (Figure 2f), while the standard GNN models are designed for homophilous graphs [24, 44]. Finally, although the transductive GBPN consistently outperforms its inductive counterpart, the gap between them is small, as the predictive probabilities on the training nodes are already quite accurate.

## 3.2 Identifiability, Interpretability and Convergence

A useful aspect of GBPN is that the learned coupling-potentials H can be interpreted as between-class affinities. We first confirm that GBPN identifies the true coupling-potentials used to generate the synthetic graphs. Indeed, Figures 2a to 2d shows that the learned H is close to the parameters used to generate the data, although the exact numbers have minor differences.

Next, we examine the learned coupling matrices for some empirical datasets (Figures 2e to 2g). The learned coupling-potentials are homophilous on the Election network and strongly heterophilous in the Sexual network, which reflect the fact that the election outcome in neighboring counties tend to be the same and that most sexual interactions in the dataset are heterosexual. On the PubMed citation network, node labels 0, 1, 2 correspond to medical publications on `experimental diabetes`, `adult-onset diabetes`, and `juvenile diabetes` respectively. The learned coupling matrix shows that papers citing each other tend to be in the same field, as shown by the large diagonal values. Beyond this, the smallest value is between `adult-onset diabetes` and `juvenile diabetes`. This aligns with the fact that those two type of diseases are largely among different age groups.

Another interesting component of the GBPN model is that we can examine intermediate predictions computed during belief propagation. Although belief propagation is not guaranteed to converge on graphs with cycles, GBPN converged quickly on all real-world graphs, oftentimes in a handful of iterations. Here, we consider the PubMed dataset in detail, using the setup in section 3.1. After training, we run belief propagation to compute the belief probability $\{p_i^{(t)}\}_{i \in V}$ for all nodes for up to $T = 20$ steps. To evaluate BP convergence, we measure the norm of the difference between the belief at step t and the converged belief, $\|p_i^{(t)} - p_i^{(\infty)}\| \approx \|p_i^{(t)} - p_i^{(20)}\|$, along with the train
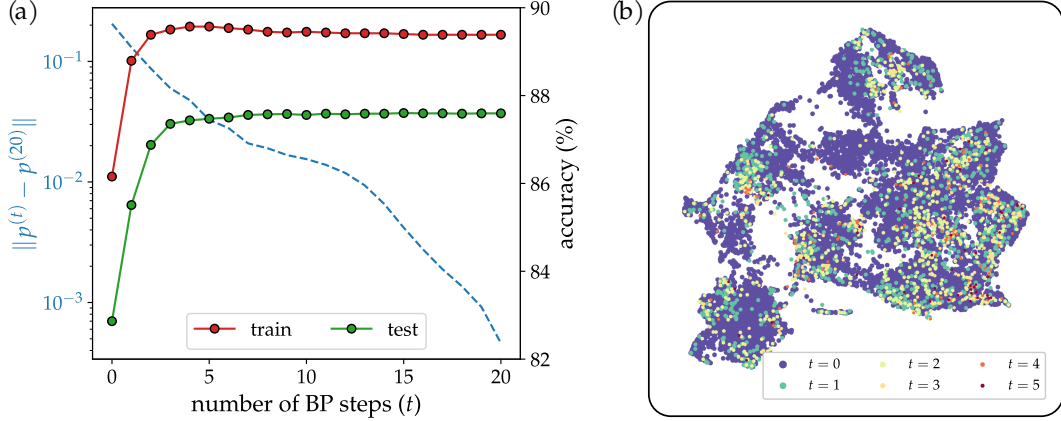
Figure 3: Convergence of GBPN on the PubMed dataset. (a) Approximate residual error at step k in the blue and train/test accuracies in red/green. (b) Nodes correctly predicted with the GBPN algorithm, where the coordinates are given by a graph spectral embedding; the color indicates the step k after which the prediction is always correct.
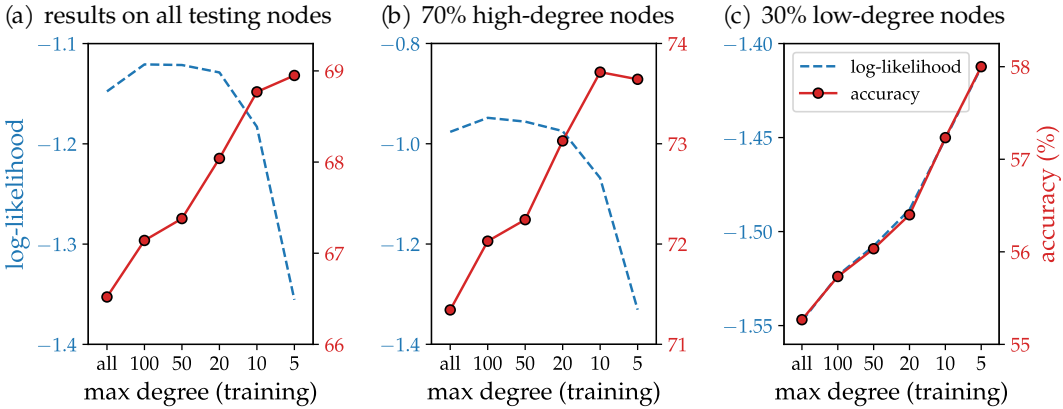
Figure 4: Visualization of the misalignment between the (unweighted) negative log-likelihood loss function and prediction accuracy on a graphs with high degree heterogeneity. (a) when training GBPN with degree regularization, the predictive log-likelihood on the testing nodes decreases while accuracy increases. (b) this misalignment is caused by 70% of nodes with the highest degrees. (c) on the rest 30% of low-degree nodes, the log-likelihood aligns perfectly with accuracy.

and test accuracies (Figure 3). The result shows that GBPN converges linearly with the number of belief-propagation steps, and the train and test accuracies change little after a handful of iterations.

To help understand BP convergence and the algorithm in general, we also visualize how each BP step refines the final prediction (Figure 3). We plot the nodes that are correctly predicted by GBPN at step 5, where the coordinate of each node is computed by compressing its spectral embeddings [45] and features together with UMAP [46], and the color indicates the step after which the prediction is always correct. Nodes corrected at later steps (e.g., $t = 5$ in dark red) are typically surrounded by others examples that are already corrected earlier. This behavior reflects the mechanism of the GBPN model: start with easy data points and iteratively correct harder examples with neighbor information.

3.3 Neighborhood Sub-sampling and Performance on Larger Datasets

In Section 2.4, we identified the misalignment between the unweighted NLL loss function and the classification accuracy on graphs with high degree heterogeneity, which prompts us to down-weight the high degree nodes. Here, we consider neighborhood sub-sampling as an alternative solution to the misalignment problem, which also reduces computation and enables efficient mini-batch training.

To demonstrate how degree heterogeneity causes the misalignment and consequently hurts accuracy, we manually reduced the degree heterogeneity by regularizing the high degree nodes during training. In particular, we train a 1-hop GBPN on the OGBN-arXiv dataset with the unweighted NLL loss,
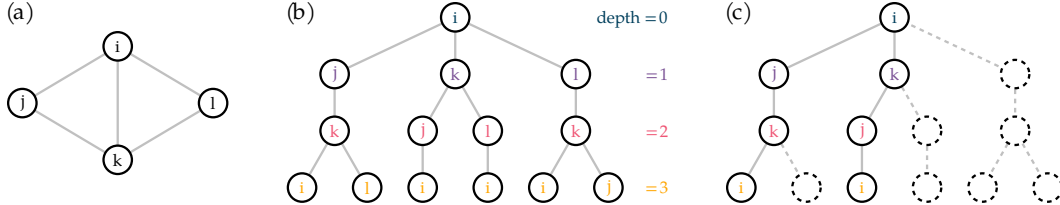
Figure 5: Demonstration of the sub-sampling algorithm on a simple graph. (a) the original graph. (b) the unrolled tree from the root node u. (c) the sub-sampled tree when the maximum degree is 2.

Table 5: Node classification performances on large-scale empirical datasets with mini-batch training.

dataset	MLP	ClusterGCN	SAGE	GAT	GBPN-I	GBPN
Elliptic	88.4 ± 0.3	81.4 ± 0.6	88.5 ± 0.5	88.3 ± 0.3	90.0 ± 0.7	**90.1** ± 0.7
Payments	6.9 ± 0.5	40.7 ± 1.5	26.6 ± 1.2	19.0 ± 1.2	47.9 ± 1.0	**53.2** ± 1.3
arXiv	54.9 ± 0.3	66.4 ± 0.5	70.2 ± 0.1	70.1 ± 0.4	70.7 ± 0.3	**71.8** ± 0.3
Products	61.3 ± 0.1	77.5 ± 0.7	78.6 ± 0.2	74.8 ± 2.2	**81.8** ± 0.3	**81.8** ± 0.2

uniformly sub-sampling a maximum number of neighbors during the aggregation step in eq. (8) (or all neighbors if the degree is below the threshold). At inference, we run GBPN with all neighbors to predict the testing labels. Figure 4a shows the average log-likelihood and accuracy on the testing nodes as a function of the maximum number of neighbors, and reducing this number improves the testing accuracy. At the same time, the log-likelihood of the testing node decreases, showing the misalignment. We further split the testing nodes into the 30% lowest degree nodes and the rest 70%, and find that the log-likelihood and accuracy aligns on the low-degree nodes (Figure 4b) but not on the high-degree nodes (Figure 4c). This justifies our heuristic of down-weighting high degree nodes.

For large-scale graphs where mini-batch training is necessary, neighborhood sub-sampling improves the classification accuracy and also reduces computational costs. In particular, computing the T-step belief of a node i requires unrolling a computation tree [47] that includes all T-hop neighbors of i, which can quickly grow to the entire graph. With neighborhood sub-sampling, we randomly select a maximum of d neighbors per node when unrolling the tree, and the sub-sampled tree has $\mathcal{O}(d^T)$ nodes (Figure 5). This is similar to sub-sampling in GNNs to address scalability [6], although our sample is a tree and also explicitly down-weights high-degree nodes.

We test the performance of GBPN with mini-batch training on four large-scale real-world datasets. The experimental setup is the same as in Section 3.1, except we only train for 100 epochs, as the model parameters are updated multiple times per epoch. We repeat each experiment only 10 times because accuracy has smaller variance. During GBPN training, we use $T = 2$ BP steps and a maximum of 5 sampled neighbors. Training on the largest dataset (Products) takes about one minute per epoch and less than two hours overall on a Telsa V100 GPU. For a fair comparison, we use a similar neighborhood expansion method for mini-batch training of GraphSAGE and GAT, with the branching factor set to 5. We replace the GCN baseline with ClusterGCN [48], where we set the average size per cluster to be 256 and the number of clusters per mini-batch to be 3; other hyperparameters are kept the same. We do not include GMNN and DeeperGNN as baselines since neither of the reference repositories implemented mini-batch training. The labels in the Elliptic and Payments datasets are highly imbalanced as most transactions are licit, so we measure the prediction performance with the F1 score ($\times 100$) on the illicit class. Performance on the arXiv and Products datasets is measured by the classification accuracy. We find that GBPN-I outperforms the baselines on all four datasets (Table 5), and GBPN provides additional gains on the Payments and arXiv datasets.

In addition to uniform sampling, we also experimented with importance sampling, which has proven useful for training certain graph neural networks [49, 50]. The results are summarized in Appendix B.

4   Conclusions and Future Work

We investigated belief propagation as a model component for node classification, developing a model that combines the advantages of graph neural networks and collective classification. Our model is easy-to-train, accurate, and scalable while maintaining interpretability and having a natural way to incorporate training labels for inference. Results on real-world datasets justify our claims. There

are several fruitful directions for future research. For instance, since BP often converges to a fix point, implicit differentiation could be used to reduce memory consumption [51]. One could also extend edge potentials to motif potentials to incorporate higher order interactions. Finally, although our algorithms are generic, improvements in node classification could lead to more personalized ad targeting or reasoning about sensitive information, which could have negative societal impact.

Disclaimer. This paper was prepared for information purposes by the AI Research Group of JPMorgan Chase & Co and its affiliates ("J.P. Morgan"), and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no explicit or implied representation and warranty and accepts no liability, for the completeness, accuracy or reliability of information, or the legal, compliance, financial, tax or accounting effects of matters contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

References

[1] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.

[2] Mark Newman. *Networks: An Introduction*. Oxford University Press, 2010.

[3] Brian Gallagher, Hanghang Tong, Tina Eliassi-Rad, and Christos Faloutsos. Using ghost edges for classification in sparsely labeled networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008.

[4] Ya Xu, Justin S Dyer, and Art B Owen. Empirical stationary correlations for semi-supervised learning on graphs. *The Annals of Applied Statistics*, pages 589–614, 2010.

[5] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

[6] 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*, page 1025–1035, 2017.

[7] Leto Peel. Graph-based semi-supervised learning for relational networks. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, pages 435–443. SIAM, 2017.

[8] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proceedings of the 6th International Conference on Learning Representations*, 2018.

[9] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6861–6871, 2019.

[10] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*, 2018.

[11] Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.

[12] David Jensen, Jennifer Neville, and Brian Gallagher. Why collective inference improves relational classification. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 593–598, 2004.

[13] Ben London and Lise Getoor. Collective classification of network data. *Data Classification: Algorithms and Applications*, 399, 2014.

[14] Jennifer Neville and David Jensen. Iterative classification in relational data. In *Proceedings of the AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, pages 13–20, 2000.

[15] Qing Lu and Lise Getoor. Link-based classification. In *Proceedings of the 20th International Conference on International Conference on Machine Learning*, pages 496–503, 2003.

[16] Sofus A Macskassy and Foster Provost. Classification in networked data: A toolkit and a univariate case study. *Journal of Machine Learning Research*, 8(May):935–983, 2007.

[17] Ben Taskar, Pieter Abbeel, and Daphne Koller. Discriminative probabilistic models for relational data. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, 2002.

[18] Kevin P. Murphy, Yair Weiss, and Michael I. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, page 467–475, 1999.

[19] Luke K McDowell, Kalyan Moy Gupta, and David W Aha. Cautious inference in collective classification. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, volume 7, pages 596–601, 2007.

[20] Meng Qu, Yoshua Bengio, and Jian Tang. GMNN: Graph Markov neural networks. In *International Conference on Machine Learning*, pages 5241–5250, 2019.

[21] Tengfei Ma, Cao Xiao, Junyuan Shang, and Jimeng Sun. CGNF: Conditional graph neural fields, 2019.

[22] Hongchang Gao, Jian Pei, and Heng Huang. Conditional random field enhanced graph convolutional neural networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.

[23] Yunsheng Shi, Zhengjie Huang, Shikun Feng, and Yu Sun. Masked label prediction: Unified massage passing model for semi-supervised classification. *arXiv:2009.03509*, 2020.

[24] Junteng Jia and Austion R Benson. Residual correlation in graph neural network regression. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 588–598, 2020.

[25] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. In *Proceedings of the 38th International Conference on Learning Representations*, 2021.

[26] Junteng Jia and Austin R. Benson. A Unifying Generative Model for Graph Learning Algorithms: Label Propagation, Graph Convolutions, and Combinations. *arXiv e-prints*, page arXiv:2101.07730, January 2021.

[27] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized PageRank. In *Proceedings of the 35th International Conference on Learning Representations*, 2018.

[28] Aleksandar Bojchevski, Johannes Klicpera, Bryan Perozzi, Amol Kapoor, Martin Blais, Benedek Rózemberczki, Michal Lukasik, and Stephan Günnemann. Scaling graph neural networks with approximate PageRank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2464–2473, 2020.

[29] KiJung Yoon, Renjie Liao, Yuwen Xiong, Lisa Zhang, Ethan Fetaya, Raquel Urtasun, Richard Zemel, and Xaq Pitkow. Inference in Probabilistic Graphical Models by Graph Neural Networks. *arXiv e-prints*, 2018.

[30] Jonathan Kuck, Shuvam Chakraborty, Hao Tang, Rachel Luo, Jiaming Song, Ashish Sabharwal, and Stefano Ermon. Belief propagation neural networks. In *Advances in Neural Information Processing Systems*, pages 667–678, 2020.

[31] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.

[32] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the 18th International Conference on Machine Learning*, page 282–289, 2001.

[33] Marc Mezard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University Press, Inc., USA, 2009.

[34] Rongjing Xiang and Jennifer Neville. Pseudolikelihood EM for Within-Network Relational Learning. In *Proceedings of the 18th IEEE International Conference on Data Mining*, page 1103–1108, 2008.

[35] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.

[36] Matthias Fey and Jan E. Lenssen. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

[37] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[38] Mark EJ Newman. Assortative mixing in networks. *Physical review letters*, 89(20):208701, 2002.

[39] 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.

[40] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of Graph Neural Network Evaluation. *arXiv e-prints*, page arXiv:1811.05868, November 2018.

[41] Martina Morris. HIV Transmission Network Metastudy Project: An Archive of Data From Eight Network Studies, 1988–2001, 2011.

[42] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao B. Schardl, and Charles E. Leiserson. EvolveGCN: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 2020.

[43] Meng Liu, Hongyang Gao, and Shuiwang Ji. Towards deeper graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020.

[44] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In *Advances in Neural Information Processing Systems*, 2020.

[45] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, page 585–591, 2001.

[46] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software*, 3(29):861, 2018.

[47] Alexander T. Ihler, John W. Fisher III, and Alan S. Willsky. Loopy belief propagation: Convergence and effects of message errors. *Journal of Machine Learning Research*, pages 905–936, 2005.

[48] 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*, page 257–266, 2019.

[49] Qingru Zhang, David Wipf, Quan Gan, and Le Song. A biased graph neural network sampler with near-optimal regret. *arXiv preprint arXiv:2103.01089*, 2021.

[50] Ziqi Liu, Zhengwei Wu, Zhiqiang Zhang, Jun Zhou, Shuang Yang, Le Song, and Yuan Qi. Bandit samplers for training graph neural networks. *arXiv preprint arXiv:2006.05806*, 2020.

[51] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep Equilibrium Models. In *Advances in Neural Information Processing Systems*, 2019.

[52] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.

[53] Taishi Uchiya, Atsuyoshi Nakamura, and Mineichi Kudo. Algorithms for adversarial bandit problems with multiple plays. In *International Conference on Algorithmic Learning Theory*, pages 375–389. Springer, 2010.

A Supplementary Materials for Reproducing the Main Results

Here we provide some implementation details of our methods to help readers reproduce and further understand the algorithms and experiments in this paper.

A.1 Additional Details on Implementations

Algorithm 1: The inference (i.e. forward propagation) algorithm for inductive GBPN.

Input : graph topology $G(V, E)$; node features $\{\boldsymbol{x}_i\}_{i \in V}$; number of BP iterations T; self-potential mapping g_θ; coupling matrix $\boldsymbol{H}$

Output : prediction for each node label $\{\hat{y}_i\}_{i \in V}$

```
/* step 0:  initialize messages */
```
1 **for** $(i, j) \in E$ **do**
2   **for** $y_j = 0 \ldots (c-1)$ **do**
3      $\log m_{ij}^{(0)}(y_j) \leftarrow -\log c$
4   **for** $y_i = 0 \ldots (c-1)$ **do**
5      $\log m_{ji}^{(0)}(y_i) \leftarrow -\log c$

```
/* step 1:  initialize beliefs */
```
6 **for** $i \in V$ **do**
7   **for** $y_i = 0 \ldots (c-1)$ **do**
8      $\log p_i^{(0)}(y_i) \leftarrow \log g_\theta(y_i; \boldsymbol{x}_i)$

```
/* step 2:  belief propagation iterations */
```
9 **for** $t = 1 \ldots T$ **do**
10   **for** $(i, j) \in E$ **do**
11   **for** $y_j = 0 \ldots (c-1)$ **do**
12        $\log m_{ij}^{(t)}(y_j) \curlyeqprec \mathrm{LSE}_{y_i} \left[\log H_{y_i, y_j} + \log p_i^{(t-1)}(y_i) - \log m_{ji}^{(t-1)}(y_i) \right]$
13   **for** $y_i = 0 \ldots (c-1)$ **do**
14        $\log m_{ji}^{(t)}(y_i) \curlyeqprec \mathrm{LSE}_{y_j} \left[\log H_{y_j, y_i} + \log p_j^{(t-1)}(y_j) - \log m_{ij}^{(t-1)}(y_j) \right]$
15   **for** $i \in V$ **do**
16   **for** $y_j = 0 \ldots (c-1)$ **do**
17        $\log p_i^{(t)}(y_i) \curlyeqprec \log p_i^{(0)}(y_i) + \sum_{j \in N(i)} \log m_{ji}^{(t)}(y_i)$

```
/* step 3:  final predictions */
```
18 **for** $i \in V$ **do**
19    $\hat{y}_i = \arg\max_{y_i} p_i^{(T)}(y_i)$

Here, we summarize the inductive GBPN variant in algorithm 1, where $\curlyeqprec$ denotes "assign after normalization." In contrast, the transductive GBPN initializes the belief $\log p_i^{(0)}(y_i)$ on each training node $i \in L$ to be $\log p_i^{(0)}(y_i) = 0$ if y_i is the ground-truth class and $\log p_i^{(0)}(y_i) = -\infty$ otherwise.

A.2 Additional Details on Experimental Setup

Architectures. The first step of GBPN is to map the features on each node into a c-dimensional initial belief (or self-potential) vector using a MLP. For this MLP, we use a 2-hidden-layer feedforward network with 256 hidden units and ReLU activation function. During training, on each dataset with a small feature dimension relative to the number of training nodes (Ising, MRF, County, Sexual, Elliptic, Payment, arXiv, Products), we set the dropout probability to 0.1; on each dataset with a large feature dimension relative to the number of training nodes (Cora, CiteSeer), we set the dropout probability to 0.6; on the rest of the datasets (PubMed, CS, Physics), we set the dropout probability to 0.3. The same activation function and dropout probabilities are used for baseline methods.

Optimization. For all methods except for GMNN and DeeperGNN, the parameters are optimized using a AdamW optimizer with with $\beta_1 = 0.9$, $\beta_2 = 0.999$, learning rate 1.0×10^{-3}, and decay rate 2.5×10^{-4}. All experiments are performed on a server with a Xeon 6254 CPU and a Telsa V100 GPU.

A.3 Additional Details on Datasets

Synthetic MRF datasets. We sample random configurations from MRFs defined on a 2-dimensional grid graph. A MRF configuration assigns a label to every node in the graph. The MRF models we use consist of two types of potentials: (i) the self-potential on each node (ii) the coupling-potential between each pairs of neighboring nodes. While the coupling-potential for each dataset is shown in fig. 2, here we describe how we set the self-potentials on each node. Let r_{i1} and r_{i2} denote the first and second grid coordinates of node i normalized between -1.0 and $+1.0$, the self-potential $h_i(y_i)$ on each node i is defined as a function of the coordinates.

- **Ising+:** $h_i(0) = \exp(-0.35 \cdot r_{i1} \cdot r_{i2})$, $h_i(1) = \exp(+0.35 \cdot r_{i1} \cdot r_{i2})$
- **Ising−:** $h_i(0) = \exp(-0.35 \cdot r_{i1} \cdot r_{i2})$, $h_i(1) = \exp(+0.35 \cdot r_{i1} \cdot r_{i2})$
- **MRF+:** $h_i(y_i) = \mathtt{sigmoid}(0.2 \cdot s_i(y_i))$; $\quad s_i(0) = 0$, $s_i(1) = r_{i1}^2 + r_{i2}^2 - 0.65$, $s_i(2) = -s_i(1)$
- **MRF−:** $h_i(y_i) = \mathtt{sigmoid}(0.6 \cdot s_i(y_i))$; $\quad s_i(0) = 0$, $s_i(1) = r_{i1}^2 + r_{i2}^2 - 0.00$, $s_i(2) = -s_i(1)$

We use the Metropolis algorithm for simulating the Ising models and Gibbs sampling [33] for simulating the 3-class MRF models. The synthetic datasets are released together with our GBPN implementation.

B Importance Sampling

In section 3.3, we observed that regularizing high degree nodes during training — by aggregating messages across a small and uniformly sampled neighborhood $N'(i)$ rather than the entire neighborhood $N(i)$ — leads to a boost in both accuracy and computational efficiency. However, one potential concern is that estimation with uniformly sampled neighbors has high variance. To this end, we explore importance sampling to reduce the variance. Surprisingly, we find the optimal sampling distribution that minimizes the variance is very close to uniform sampling on benchmark datasets.

B.1 Theoretical Optimal Distribution for Importance Sampling

We first formulate the neighbor sub-sampling procedure according to an arbitrary sampling distribution $p \in [0,1]^{|N(i)|}$ by considering a sequence of $d \leq |N(i)|$ independent and identically distributed (i.i.d.) draws of neighbor indices, $N'(i) = (j_1, \ldots, j_d)$, where each j is sampled according to p. For each of the samples $j \in N'(i)$, let

$$X_j(y_i) = \log m_{ji}^{(k)}(y_i) \tag{13}$$

be the random variable corresponding to the message from $j \in N'(i)$ for class y_i. Further denote the sum of the incoming messages over the sampled neighbors with respect to class y_i as $X(y_i)$, i.e.,

$$X(y_i) = \sum_{j \in N'(i)} X_j(y_i). \tag{14}$$

Denote the joint distribution of all d samples as p^d, we can compute the expectation of $X(y_i)$,

$$\mathbb{E}_{N'(i) \sim p^d}[X(y_i)] = |N'(i)| \, \mathbb{E}_{j \sim p}[X_j(y_i)] \qquad \text{(linearity of expectation \& i.i.d. sampling)}$$

$$= |N'(i)| \sum_{j \in N(i)} \left(\log m_{ji}^{(k)}(y_i) \right) p_j \qquad \text{(definition of } \mathbb{E}_{j \sim p}[X_j(y_i)])$$

$$= \frac{|N'(i)|}{|N(i)|} \sum_{j \in N(i)} \log m_{ji}^{(k)}(y_i). \qquad (p_j = 1/|N(i)| \text{ for uniform sampling})$$

Hence, $X(y_i)$ is an unbiased estimator for the quantity

$$\Omega(y_i) = \frac{|N'(i)|}{|N(i)|} \sum_{j \in N(i)} \log m_{ji}^{(k)}(y_i), \tag{15}$$

which represents the scaled aggregate messages from all neighbors $N(i)$. However, the estimation given by $X(y_i)$ may have high variance. Thus, we use an unbiased and low-variance estimator defined by importance sampling. In particular, we consider the following estimator:

$$Z(y_i) = \sum_{j \in N'(i)} Z_j(y_i) = \sum_{j \in N'(i)} \frac{X_j(y_i)}{|N(i)| \cdot p_j}, \tag{16}$$

which simply re-weights $X_j(y_i)$. It is easy to show $Z(y_i)$ is an unbiased estimator for $\Omega(y_i)$:

$$\mathbb{E}_{N'(i) \sim p^d}[Z(y_i)] = |N'(i)| \, \mathbb{E}_{j \sim p}[Z_j(y_i)] = \frac{|N'(i)|}{|N(i)|} \mathbb{E}_{j \sim p}[X_j/p_i] \tag{17}$$

$$= \frac{|N'(i)|}{|N(i)|} \sum_{j \in N(i)} \log m_{ji}^{(k)}(y_i). \tag{18}$$

Now, our goal is to find a probability distribution $p^* \in [0,1]^{|N(i)|}$ so that $Z(y_i)$ has smallest possible variance, i.e.,

$$p^* = \operatorname{argmin}_p \operatorname{Var}_{p^d}(Z(y_i)), \tag{19}$$

where $\operatorname{Var}_{p^d}(Z(y_i))$ is the variance of $Z(y_i)$ under the sampling distribution p^d. The variance further simplifies to

$$\begin{aligned}
\operatorname{Var}_{p^d}(Z(y_i)) &= |N'(i)| \operatorname{Var}_p(Z_j(y_i)) && \text{(i.i.d. sampling)} \\
&= \frac{|N'(i)|}{|N(i)|^2} \operatorname{Var}_p(X_j(y_i)/p_j) && \text{(definition of } Z_j(y_i)) \\
&= \frac{|N'(i)|}{|N(i)|^2} \left(\mathbb{E}_p[(X_j(y_i)/p_j)^2] - \mathbb{E}_p[X_j(y_i)/p_j]^2 \right).
\end{aligned}$$

By unbiasedness, we have that $(\mathbb{E}[X_j(y_i)/p_j])^2 = \left(\sum_{j \in N(i)} \log m_{ji}^{(k)}(y_i) \right)^2$, which implies that the distribution p that minimizes the variance is the one which minimizes the second moment $\mathbb{E}_{j \sim p}[(X_j(y_i)/p_j)^2]$. This evaluates to

$$\mathbb{E}_{j \sim p}[(X_j(y_i)/p_j)^2] = \sum_{j \in N(i)} \left(\log m_{ji}^{(k)}(y_i) \right)^2 / p_j. \tag{20}$$

The expression above is convex in p, so using the method of Lagrange multipliers,[2] we find that the minimizer is given by

$$p_j^* \propto |\log m_{ji}^{(k)}(y_i)| \quad \forall j \in |N(i)|. \tag{21}$$

Since we want a distribution that minimizes the variance over all potential classes $y_i \in \{0, \ldots, c-1\}$, we further consider minimizing the sum of variances over all classes

$$p^* = \operatorname{argmin}_p \sum_{y_i} \operatorname{Var}_{p^d}(Z(y_i)). \tag{22}$$

When the number of samples is clear from context, we will let

$$\operatorname{Var}(p) = \sum_{y_i} \operatorname{Var}_{p^d}(Z(y_i)) \tag{23}$$

denote the sum of variances over all classes under distribution p. An analogous derivation leads to the optimal distribution having the form

$$p_j^* \propto \sqrt{\sum_{y_i} (\log m_{ji}^{(k)}(y_i))^2}. \tag{24}$$

---

[2] More generally, the distribution $p_i = |x_i| / \sum_j |x_j|$ minimizes the sum $\sum_i x_i^2 / p_i$.

B.2 Online Learning

We would ideally like to sample the neighbors according to the optimal sampling distribution as defined above, but the optimal distribution requires knowledge of all of the neighbors' messages. Since we do not know the messages prior to sampling, we are faced with the dilemma where we cannot sample optimally prior to seeing the messages, but cannot see the messages until we sample. We can address via the well-studied exploration-and-exploitation trade-off in online learning with partial information, which includes the Multi-armed Bandit setting [52, 53].

We formulate the variance minimization problem as an online learning problem where the goal is to minimize the regret (sum of variances) with respect to the best sampling distribution in hindsight. In particular, we consider the linearized, partial-information setting where the losses at each epoch ℓ_τ correspond to the gradient of the variance, i.e., for each $j \in N'(i)$,

$$\ell_{\tau,j} = -\frac{\sqrt{\sum_{y_i} (\log m_{ji}^{(k)}(y_i))^2}}{p_{\tau,j}^2}, \tag{25}$$

and $\ell_{\tau,j} = 0$ for all $j \notin N'(i)$, where $p_{\tau,j}$ is the probability of sampling neighbor j at epoch τ. Note that we are in the partial information setting since we do not see the messages of all neighbors, but only those of sampled neighbors, so we will need to reweigh the losses $\ell_{\tau,j}$ by an additional $p_{\tau,j}$ term in the denominator so that we can approximate the full gradient in expectation (see [52, Algorithm 10.2] for details). Our overarching goal is to bound the expected regret with respect to the best distribution in hindsight,

$$\text{Regret}(\{p_\tau\}) = \sum_\tau \text{Var}(p_\tau) - \min_{p \in \Delta} \sum_\tau \text{Var}(p). \tag{26}$$

where $p_\tau \in \Delta$ is the sampling distribution at epoch τ and Δ is the probability simplex. By convexity of the variance with respect to the probabilities p, the regret defined above is upper bounded by the following linearized regret, i.e.,

$$\text{Regret}(\{p_\tau\}) \le \sum_\tau \langle p_\tau, \ell_\tau \rangle - \min_{p \in \Delta} \sum_\tau \langle p, \ell_\tau \rangle, \tag{27}$$

Hence, minimizing this upper bound would guarantee that the sequence of sampling distributions we used in sampling the neighbors of a node i is competitive with the best sampling distribution in hindsight. It can be shown [52, Theorem 10.2] that for an appropriate choice of the learning rate, the regret of the EXP3 algorithm with respect to sampling the neighbors of node i is bounded by $\mathcal{O}\left(\sqrt{\log(|N(i)|) \sum_\tau \mathbb{E}_{j \sim p}[\ell_{\tau,j}^2]}\right)$.

B.3 Empirical Evaluations & Discussion

We implemented the adaptive version of EXP3 with time-varying learning rates [52, Chapter 7.6]; other variants of EXP3 provide similar regret guarantees [49, 50]. Across evaluations on the data sets described in Table 2, we observed very similar performance between our regret-based, importance sampling approach and uniform sampling. Even for cases where importance sampling with EXP3 improved the accuracy, these improvements were not significant in our evaluations, i.e., they were well-within one standard deviation of uniform sampling's performance.

To dig deeper into our (unexpected) findings, we questioned whether EXP3 (and importance sampling) was providing a variance reduction at all. To this end, we plotted the variance of the distribution generated by EXP3 and compared it to the variance under the optimal sampling distribution p_τ^* in hindsight (computed as in eq. (24)) after each training epoch τ. We performed the calculation for uniform sampling. Note that unlike in the regret definition above, p_τ^* is the optimal distribution *with respect to epoch τ*, meaning that it has the lowest variance for epoch τ possible, $\text{Var}_\tau^{\text{opt}}$. We also directly compared the variance of our importance sampling distribution with that of uniform sampling, to determine whether the similarity in performance was a result of similar estimator variance.

Figure 6 depicts the results of our evaluations on the Cora, PubMed, and arXiv datasets. In Figure 6a, we plot the ratio for the variances under IMPORTANCE and UNIFORM distributions relative to the variance $\text{Var}_\tau^{\text{opt}}$ under the optimal distribution p_τ^*. We observe that after a brief burn-in period, the variance ratio of IMPORTANCE is very close to 1, implying that EXP3 in fact generates sampling
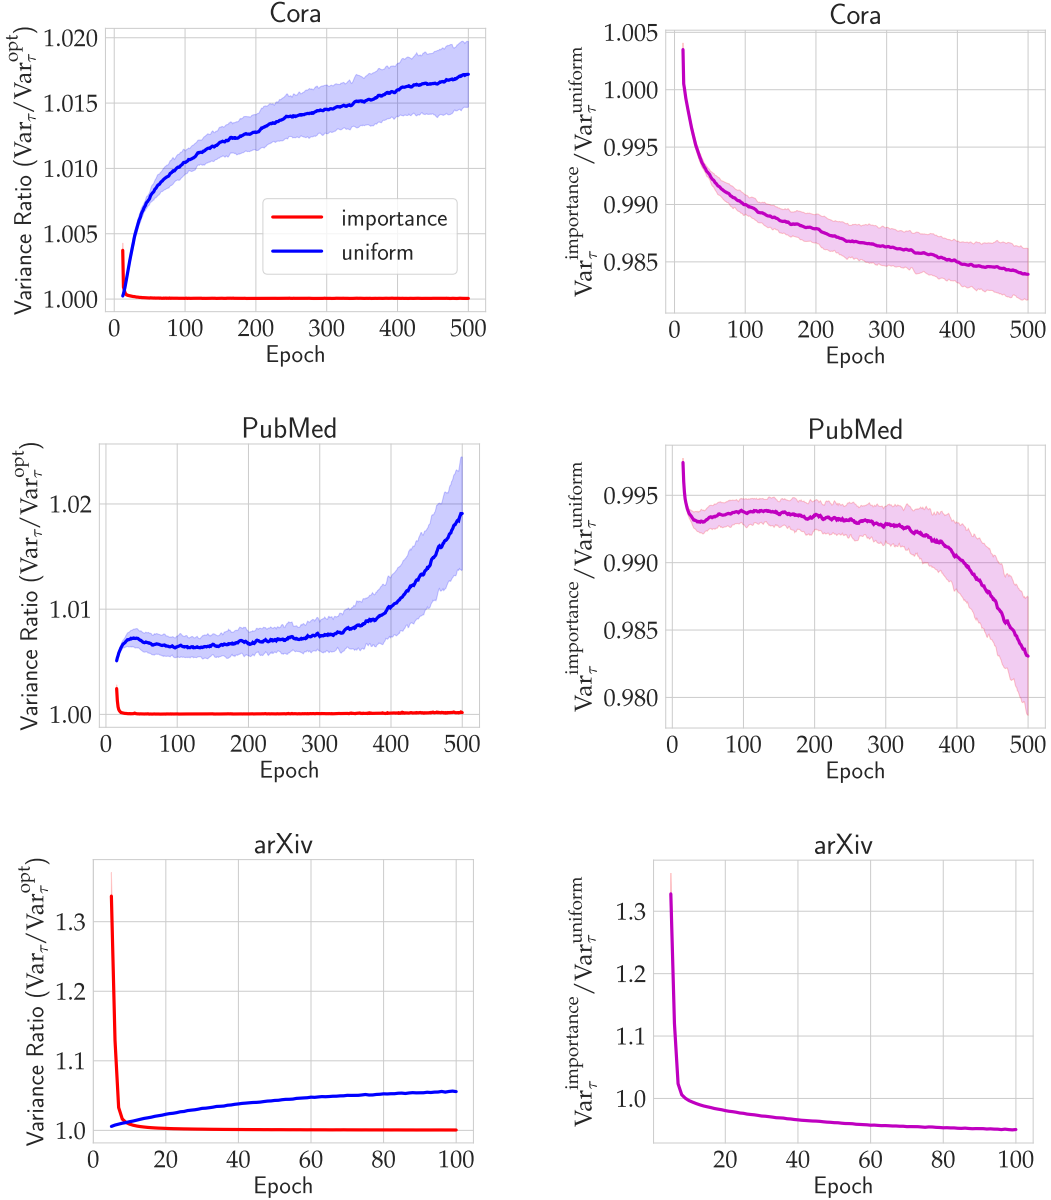
(a) Variance Relative to Optimal ($\text{Var}_\tau^{\text{opt}}$)

(b) Variance of IMPORTANCE relative to UNIFORM

Figure 6: Evaluations of the relative variance under our sampling distribution and uniform sampling on varying data sets. The ratio of the variances are averaged over all nodes per trial and shaded regions correspond to the values within one standard deviation from the mean.

distributions that are essentially optimal (p_τ^*) for each time step. In fact, we see that as time progresses, our online learning algorithm is able to learn increasingly better (closer to optimal) distributions. On the other hand, uniform sampling's performance relative to the optimal sampling distribution tends to degrade over the training process. Similarly, in fig. 6b, where we directly compute the ratio of variances between IMPORTANCE and UNIFORM, i.e.,

$$\rho_\tau = \frac{\text{Var}_\tau^{\text{importance}}}{\text{Var}_\tau^{\text{uniform}}}, \tag{28}$$

we see IMPORTANCE increasingly outperforms UNIFORM ($\rho_\tau < 1$) after the initial 10-20 epochs.

Despite the qualitatively encouraging results, the scales of the y-axis in fig. 6 help us understand why IMPORTANCE does not ultimately lead to better performance in terms of test accuracy. In particular,

Figure 6 shows that we achieve at most only 5% reduction in variance relative to uniform sampling throughout virtually the entire training process. These empirical evaluations suggest that uniform sampling is very close to the per-time-step optimal distribution p_τ^* and provide an explanation for the lack of statistically significant improvement in test accuracy we observed with importance sampling.

We envision that practitioners can bridge online learning with our framework and use, e.g., EXP3 as we described it here, on scenarios where uniform sampling of neighbors may not be close to optimal. By bridging tools with theoretical guarantees from online learning with our work, we can in effect bound the worst-case performance of our sub-sampling procedure against even adversarial inputs. In future work, we plan to investigate scenarios where importance sampling can in fact yield significant improvements when it comes to both variance reduction and increased test accuracy.