

Node Attribute Prediction with Weighted and Directed Edges on Single and Multilayer Networks

Yiguang Zhang¹, Kristen M. Altenburger², Poppy Zhang², Tsutomu Okano², Shawndra Hill²

¹Columbia University*

²Meta

Abstract

With the rapid development of digital platforms, users can now interact in endless ways from writing business reviews and comments to sharing information with their friends and followers. As a result, organizations have numerous digital social networks available for graph learning problems with little guidance on how to select the right graph or how to combine multiple edge types. For example, while user-to-user interactions are directed in nature, many graph learning approaches use the undirected version of the network. In this paper, we introduce edge direction, edge weight, and multi-relational data for node prediction tasks. We first adapt an existing node attribute prediction method for binary prediction, LINK-Naive Bayes, to account for both edge direction and weights on single-layer networks. We compare predictive performance metrics across various node attribute prediction tasks for an ads click prediction task on Facebook and for a publicly available dataset from the Open Graph Benchmark (OGB). We observe meaningful predictive performance improvements when incorporating edge direction and weight, and performance that's competitive with the OGB Leaderboard. We then introduce an approach called MultiLayerLINK-NaiveBayes that can combine multiple network layers during training and observe superior performance over the single-layer results. Ultimately, whether edge direction, edge weights, and multi-layers are practically useful will depend on the particular setting. Our approach enables practitioners to quickly combine multiple layers and edge types.

Introduction

Inferring user attributes has a rich history in the graph mining literature including predicting demographic attributes (Dong et al. 2017; Brea et al. 2014), marketing behavior (Hill, Provost, and Volinsky 2006; Goel and Goldstein 2014; Zubcsek, Katona, and Sarvary 2017), integrity outcomes (Chakrabarti et al. 2014; Fakhræi et al. 2015; Wang, Gong, and Fu 2017) and more. The typical within-network setting includes known ground-truth labels for a subset of users and a single network on how all users are connected. The goal is then to infer the missing labels or attributes for the remaining nodes (Neville and Jensen 2000; Welling and Kipf 2016; Perozzi, Al-Rfou, and Skiena 2014;

Sen et al. 2008; Altenburger and Ugander 2021). The related across-layer task trains on a set of labeled nodes and edge type and predicts for the same node set for a different edge type. Examples include making attribute predictions over time (Henderson et al. 2011). A recent review paper has also unified the various types of graph representation learning methods (Chami et al. 2020). We focus on the within-network task in this paper.

Feature representations for node attribute prediction problems can generally be classified by whether the features are identity-independent (dependent) and label-independent (dependent) (Altenburger and Ugander 2021). Identity-dependence captures whether a feature depends on the identity labels of the node $i=1,2,\dots,n$. Label-dependence (Gallagher and Eliassi-Rad 2008) captures whether the feature depends on the attribute label. Node embedding feature representations are a popular class of methods for representing users via a low-dimensional feature. Approaches like node2vec (Grover and Leskovec 2016), DeepWalk (Perozzi, Al-Rfou, and Skiena 2014), and LINE (Tang et al. 2015) are all identity-dependent and label-independent. LINK (Zheleva and Getoor 2009) is another promising identity-dependent and label-independent feature representation where the feature is the row in the adjacency matrix.

While many within-network tasks train on a single network or single edge relationship (Henderson et al. 2012, 2011; Altenburger and Ugander 2021; Lim et al. 2021), real-world node prediction settings have numerous user-to-user networks to consider for a given prediction task. It is not clear which is the optimal network to select for a given problem. For example, Facebook (FB) users can be connected by a friendship edge or interaction edges including comments, reactions, and more. These same users can also follow one another on Instagram (IG) and interact. On other social network platforms like Twitter, users can be represented by follower relationships or interaction edges (Huberman, Romero, and Wu 2008). These types of networks are considered multi-layer networks (Kivelä et al. 2014) where the two aspects are the platform (FB/IG) and interaction or friendship/follower edge type.

With the rise in multi-relational edge types in networks, recent methods propose new ways to incorporate multi-layer edges. Embedding approaches can additionally incorporate multi-relational edges into the embedding framework (Feng

*all work was done while intern at Meta

et al. 2019; Bordes et al. 2013; Lerer et al. 2019; Zitnik and Leskovec 2017). Graph Neural Network (GNN) approaches include approaches like heterogeneous graph transformers (Hu et al. 2020b) which can incorporate different node and edge types. One major limitation of GNNs is that they make a homophily assumption (McPherson, Smith-Lovin, and Cook 2001), whereby connected units are assumed to be more similar to one another. In real-world networks, the homophily assumption can be unreasonable (Zhu et al. 2020; Loveland et al. 2023). Some recent work has focused on improving the performance of GNN in low homophily or heterophily settings (Zhu et al. 2020; Rossi et al. 2023; Ma et al. 2021). Meanwhile, other promising feature representations like LINK (Zheleva and Getoor 2009), which create a large sparse feature vector from the row of an adjacency matrix, do not require the network to be homophilous or heterophilous (Altenburger and Ugander 2018). Another advantage for LINK is that it can scale easily to networks with millions of nodes.

In this paper, we extend the LINK approach (Zheleva and Getoor 2009) to weighted and directed edges on single-layer networks and then extend the method to multi-relational edges. This paper follows other recent approaches for leveraging multilayers such as modeling multilayer interdependence (Aguiar, Taylor, and Ugander 2022) or extracting k -cores from multilayers (Hashemi, Behrouz, and Lakshmanan 2022). We refer the interested reader to reviews (Boccaletti et al. 2014; Kivela et al. 2014) on multilayer networks. Our work is also related to the “graph sanitation” task (Xu, Du, and Tong 2022), which learns the optimal graph representation from a single graph.

In our setting, our goal is to learn the relative value of different types of edges based on weight, direction, or multilayers for improving the Node Attribute Prediction Task. In the “Node Prediction on Single-Layer Networks with Edge Direction and Weight” section, we first introduce how to incorporate edge weights and direction information into LINK by utilizing an “edge binning” strategy that enables us to distinguish between different types of edges. We observe improved predictive performance when incorporating edge direction and/or edge weight for a business prediction task and to a protein-protein network that’s publicly available. For the undirected protein-protein network, we observe a neutral ROC-AUC performance on the unweighted network but observe improved performance for the different versions incorporating edge weight information. This illustrates how an undirected/unweighted network can seemingly appear to provide no information for node prediction, but can contain meaningful information once edge direction and weight are included. Then, in the “Node Attribute Prediction with Multi-layer Networks” section, we introduce a method called MultiLayerLINK-NaiveBayes which combines multiple layers for network prediction and observe additional predictive performance gains beyond the single-layer setting. We conclude with a discussion of our findings and future research directions in the field of node attribute prediction.

We summarize our contributions below:

- While many GNN methods for node attribute prediction tasks rely on similarity (homophily) or dissimilarity (het-

erophily) in the network, LINK is a network prediction methodology that is homophily/heterophily-agnostic. We extend the methodology for edge direction and weights on the single-layer network. We then extend to the multi-layer setting. An advantage of our approach is that it is domain-agnostic.

- For single-layer networks, we propose the concept of edge binning as a way to categorize edge types by direction/weight. We propose a unified methodology to incorporate this edge binning information for node attribute prediction tasks. We demonstrate this methodology on an ad click prediction task from Facebook and on a publicly available dataset from the Open Graph Benchmark, showing that accounting for edge direction/weight can increase the predictive performance of node attribute prediction tasks. While recent work includes edge direction into Graph Neural Networks (Rossi et al. 2023), this paper is the first to our knowledge to incorporate edge direction/weight into LINK.
- For multi-layer networks, we introduce a new approach to leverage edges across different types of network layers and demonstrate further predictive performance gains relative to single-layer networks. We include the time complexity of this algorithm as well.

Broader Perspectives and Ethical Considerations: This research was reviewed and approved by an internal review board-like organization in industry. We limit all analysis to users who have explicitly linked their Facebook and Instagram accounts. All data was de-identified and analyzed in aggregate. Node attribute prediction challenges in particular and machine learning on graphs in general have become commonplace across organizations with rich social network data. Multilayer networks provide a natural language for analyzing different user-level networks and for framing various prediction tasks. New methods to improve how to optimally combine or select networks for a prediction task have tremendous potential for helping organizations improve the user experience. Beyond what is presented here, we believe there is room for improved methods that evaluate the representativeness of node labels for prediction tasks and adequately accounts for any label bias.

Node Prediction on Single-layer Networks with Edge Direction and Weight

While existing node embedding approaches can incorporate multi-relational information (Feng et al. 2019; Bordes et al. 2013; Lerer et al. 2019), other existing feature representations do not. LINK-logistic regression (Zheleva and Getoor 2009) is an interpretable node prediction approach that represents a node by the row of the graph adjacency matrix. This large, sparse feature vector can then be utilized by a regularized logistic regression model for binary node prediction tasks. This method can also perform optimally with or without homophily (Zheleva and Getoor 2009; Altenburger and Ugander 2018). For this paper, given the flexibility of LINK to work in a variety of settings and given its prior high predictive performance, we adapt the LINK-Naive Bayes approach (Altenburger and Ugander 2018) to account for both




Figure 1: (left) We illustrate a single-layer network prediction problem across different edge types. Each layer represents a particular edge type and, we compare predictive performance of each layer. (right) We illustrate how multi-layer prediction is done with respect to different edge types.

edge direction and weights. We note that LINKX is a similar approach (Lim et al. 2021) that embeds the node features and adjacency matrix separately via an MLP. The focus of this work is on LINK-Naive Bayes where we have the adjacency matrix as an input signal. While LINK refers to using the rows of the adjacency matrix as features, the scope of this paper is on node attribute prediction and not link prediction.

Undirected and Unweighted Edges

We first review the derivation of LINK-Naive Bayes for a single layer in an undirected, unweighted network $G = (V, E)$, where each $v \in V$ represents a user, and $(u, v) \in E$ if and only if u and v interact with each other. We show in Figure 1 (left) a schematic of the within-layer task. Each rectangle represents a network for a particular edge type or layer. The prediction task, as will be explained, will be done within each network such as only within the friendship network.

Suppose we have a binary label prediction task, where each node has a label being either $+1$ or -1 . For a target node u , the task is to predict its label using known labels from neighbors of u in the network layer. Particularly, for each node $i \in G$, we define $x_i := \mathbb{1}_{\{(i,u) \in E(G)\}}$ as the indicator of i being a neighbor of u , and a random variable $X \in \{0, 1\}^N$, where N is the total number of nodes in the network. We review the Naive Bayes classification rule $LR(x)$ via the likelihood ratio and by making a conditional independence assumption:

$$\begin{aligned}
&= \frac{P(y_{train} = +1) \cdot P(X|y_{train} = +1)}{P(y_{train} = -1) \cdot P(X|y_{train} = -1)} \\
&= \frac{P(y_{train} = +1) \cdot \prod_{i=1}^N P(X_i = x_i|y_{train} = +1)}{P(y_{train} = -1) \cdot \prod_{i=1}^N P(X_i = x_i|y_{train} = -1)} \\
&= \frac{P(y_{train} = +1)}{P(y_{train} = -1)} \cdot \prod_{i:x_i=0} \frac{P(X_i = x_i|y_{train} = +1)}{P(X_i = x_i|y_{train} = -1)} \\
&\quad \cdot \prod_{i:x_i=1} \frac{P(X_i = x_i|y_{train} = +1)}{P(X_i = x_i|y_{train} = -1)}
\end{aligned}$$

Note that by setting n_{+1} (n_{-1}) as the number of $+1$ -labeled (-1 -labeled) nodes in the training sample and $d_{i,+1}$ ($d_{i,-1}$) as node i 's degree with $+1$ (-1) nodes in the training sample, we have the following empirical estimations:

$$\begin{aligned}
\hat{P}(y_{train} = +1) &= \frac{n_{+1}}{n_{+1} + n_{-1}}, \\
\hat{P}(y_{train} = -1) &= \frac{n_{-1}}{n_{+1} + n_{-1}}, \\
\hat{P}(X_i = 1|y_{train} = +1) &= \frac{d_{i,+1} + 1}{n_{+1} + 2}, \\
\hat{P}(X_i = 1|y_{train} = -1) &= \frac{d_{i,-1} + 1}{n_{-1} + 2}, \\
\hat{P}(X_i = 0|y_{train} = +1) &= \frac{n_{+1} - d_{i,+1} + 1}{n_{+1} + 2}, \\
\hat{P}(X_i = 0|y_{train} = -1) &= \frac{n_{-1} - d_{i,-1} + 1}{n_{-1} + 2}.
\end{aligned}$$

Assuming that for sparse networks that $n_{+1}, n_{-1} \gg d_{i,+1}, d_{i,-1}$, the final derivation for the log of the likelihood ratio can be written as follows (Altenburger and Ugander 2018):

$$\log(LR(x)) \approx \underbrace{\sum_{i=1}^N x_i \cdot \log \left[\frac{d_{i,+1} + 1}{d_{i,-1} + 1} \cdot \frac{n_{-1} + 2}{n_{+1} + 2} \right]}_{\text{one-hop aggregation}} \cdot \underbrace{\quad}_{\text{two-hop aggregation}}. \quad (1)$$

Directed and Weighted Edges

As we observed when comparing user-to-user networks in the ‘‘Analysis of Multilayer Networks’’ section, edge direction and weights can contain different information across different interaction networks. To better use network signals from directed and weighted networks, we propose two extensions of LINK-Naive Bayes to incorporate edge direction and weight.

Create Edge Bins: One central idea for our approach is to create ‘‘bins’’ for edges. We will describe the different ways to create edge bins and will reference this set-up throughout the rest of the paper. An edge bin is a way to place edges into mutually exclusive categories based on edge weight and/or edge direction. As shown in Figure 2, we illustrate edge bins based on direction (left) and edge bins based on weight (right). This means that instead of treating all edges equivalently as in the current LINK-Naive Bayes approach, edge bins will be a way for us to distinguish between different types of edges. We can also combine different edge bins to create finer bins such as combining direction \times weight bins.

(Version I) Learning Weights for Edge Bins: We change the previously introduced observed features $x_i \in \{0, 1\}$ and the random variable $X \in \{0, 1\}^N$ to reflect edge weights and/or directions. Specifically, we first divide the edges into $|W|$ bins by pre-selected metrics. For example, as introduced in the last section, edges with different weights are




Figure 2: We illustrate binning edges. (left) We show 4 categories for binning edges based on whether a source node connects with a target node, a target node connects with a source node, both source and target connect, or neither. (right) We show 4 categories for binning edges based on the strength of the connection between the source and target node.

divided into bins by their values, and directed edges are divided into bins by whether the edge is an outgoing one from the target node, an incoming one, or both. We denote the set of bins as $\{Bin_w, w \in W\}$, where W is the set of bins. We denote $\mathbb{1}_{x_i=w}$ if the edge between the target node j and source node i belongs to Bin_w .

An easy extension to test is optimizing the weights on different edge bins for two-hop aggregations when aggregating scores among a source node’s connections (right-hand side of Equation 1). Specifically, let λ_w be the weight of Bin_w and denote $\vec{\lambda}$ the vector of λ_w ’s: $\vec{\lambda} = (\lambda_w)_{w=1, \dots, |W|}$. Denoting α as the intercept term in Logistic regression, we can write the classifier’s score function $f(\cdot)$ as the following:

$$f(\alpha, \vec{\lambda}, x) = \alpha + \underbrace{\sum_{w=1}^{|W|} \lambda_w \sum_{i=1}^N \mathbb{1}_{x_i=w} \cdot \log \left[\frac{d_{i,+1} + 1}{d_{i,-1} + 1} \cdot \frac{n_{-1} + 2}{n_{+1} + 2} \right]}_{\text{one-hop aggregation}}}_{\text{two-hop aggregation by edge bin } w} \quad (2)$$

We can then estimate α and $\vec{\lambda}$ with logistic regression utilizing a validation split, which is minimizing the log loss for:

$$l := \sum_{(X,y) \in \text{validation data}} - \left(\frac{y+1}{2} \right) \log \left(\frac{e^{f(\alpha, \vec{\lambda}, X)}}{1 + e^{f(\alpha, \vec{\lambda}, X)}} \right) - \left(\frac{y-1}{2} \right) \log \left(\frac{1}{1 + e^{f(\alpha, \vec{\lambda}, X)}} \right).$$

Note that we can regard Equation 2 as a more general form of Equation 1 as it recovers Equation 1 in the case that, there are only two bins with $Bin_1 = \{x_i : x_i = 1\}$, $Bin_2 = \{x_i : x_i = 0\}$, and $\lambda_1 = 1, \lambda_2 = 0, \alpha = C_1$. The terms $d_{i,+1}$ ($d_{i,-1}$) still refer to the undirected/unweighted degree counts a source node has with other nodes labeled +1 (-1).

(Version II/II*) Classifier within each Edge Bin: An alternative approach is to apply a Naive Bayes classifier within

each edge bin. Here, we further denote $d_{i,+1,w}$ ($d_{i,-1,w}$) as the number of +1-labeled (-1-labeled) neighbors s of node i with edge type of (i, s) in Bin_w . Note that $d_{i,+1} = \sum_{w=1}^{|W|} d_{i,+1,w}$. We then have

$$\begin{aligned} \hat{P}(y_{train} = +1) &= \frac{n_{+1}}{n_{+1} + n_{-1}}, \\ \hat{P}(X_i = w | y_{train} = +1) &= \frac{d_{i,+1,w} + 1}{n_{+1} + |W|}; \\ \hat{P}(y_{train} = -1) &= \frac{n_{-1}}{n_{+1} + n_{-1}}, \\ \hat{P}(X_i = w | y_{train} = -1) &= \frac{d_{i,-1,w} + 1}{n_{-1} + |W|}. \end{aligned}$$

For the log-likelihood ratio $\log(LR(x))$, we have the following formula (see the full derivation in the “Derivation of Version II of Naive Bayes Classifier” appendix section):

$$\approx \sum_{i=1}^N \sum_{w=1}^{|W|} \mathbb{1}_{x_i=w} \cdot \log \left[\frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]$$

Just like in Version I, we can make the model more flexible, by learning different weights λ_w for the edge bins via a logistic regression model:

$$f(\alpha, \vec{\lambda}, x) = \alpha + \underbrace{\sum_{w=1}^{|W|} \lambda_w \sum_{i=1}^N \mathbb{1}_{x_i=w} \cdot \log \left[\frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]}_{\text{one-hop aggregation by edge bin}}}_{\text{two-hop aggregation by edge bin}}$$

Intuitively, Version I only uses the information of the direction and weights in turning λ_w ’s, while Version II also uses this edge information to estimate the posterior distribution of y . An alternative way for creating this classifier is to separate out one-hop and two-hop edge bins via Version II*. For example, for a given node i , we could construct the two-hop aggregations based on the out-degree edge bin but allow different weights between a node’s one-hop in-degree, out-degree, and reciprocal edges. Otherwise, Version II, would only allow out-degree aggregations among out-degree connections. We then learn different weights $\lambda_{w,w'}$ via the following:

$$f(\alpha, \vec{\lambda}, x) = \alpha + \underbrace{\sum_{w=1}^{|W|} \sum_{w'=1}^{|W|} \lambda_{w,w'} \sum_{i=1}^N \mathbb{1}_{x_i=w} \log \left[\frac{d_{i,+1,w'} + 1}{d_{i,-1,w'} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]}_{\text{one-hop aggregation by edge bin}}}_{\text{two-hop aggregation by edge bin}}$$

Empirical Results

Application on Facebook: There are numerous user-to-user interaction networks available across the Facebook and Instagram platforms. Users can be friends on Facebook and

follow one another on Instagram. Users can also interact with one another via more dynamic interactions like commenting on other user’s content or reacting. In the Appendix Section “Analysis of Real-World Multilayer Networks”, we compare user networks across Facebook and Instagram illustrating how users behave differently across different interactions. In this section, we’ll focus on one such network created by user-to-user comments. We first compare the predictive performance of the extended LINK-Naive Bayes classifier (Version I and Version II¹) using the largest user-to-user interaction network – the *comment* network from Facebook. Specifically, we construct a directed, weighted user-to-user network of over 100M users based on users’ commenting activities in the last week of July 2022. A directed edge exists if a source user commented on a target user’s content in the last week of July 2022, and we create edge weights based on the frequency of interactions between users. For binary prediction task, we consider the users who were exposed to Retail ads on 2022-07-30 and code $y_{train} = +1$ if a user clicked on an ad and $y_{train} = -1$ otherwise following previous ad click prediction problems (He et al. 2014). We restrict this sample to active U.S. users who are older than 18 years-old with at least 20 Facebook friends.

We set-up a within-network task where we randomly select 80% nodes as training nodes, 10% as validation, and the remaining 10% as test nodes. Specifically, we first randomly select 80% nodes where we observe their labels and merge these training labels into the edgelist. Among the rest 20% nodes, we use 1/2 to train the weights in the logistic regressions and use the remaining 1/2 as the test nodes. Note that we don’t train the weights using the 80% nodes, since the NB estimators utilize the label information of the 80% nodes in training, which makes them different from the remaining 20% nodes.

We report predictive performance results with normalized entropy (NE) or normalized cross-entropy (He et al. 2014), which is standard for ad click prediction tasks. NE computes the predictive log loss normalized by the entropy of the baseline click-through-rate, and lower values are better. (See the “Formula for Normalized Entropy” appendix section for more details.) We report the relative percent change in test normalized entropy (NE) when comparing the different versions to the undirected, unweighted version in Table 1. We do not report raw NE due to the sensitivity of this prediction task but all results are above a random classifier. Also note that class imbalance is not an issue for the NE metric since it uses the log-likelihood ratio, which doesn’t directly measure the performance of a predictor under an imbalanced dataset. (In the next section, we’ll report raw predictive results on a publicly available dataset). To test the usefulness of edge direction and weight, we consider the following three ways of dividing the edges into mutually exclusive bins:

- **undirected, weighted:** we ignore the edge direction, and divide the edges by considering if the weight is of 0, 1, 2 – 5, > 5. We select these bins in a way to distinguish between infrequent vs. more frequent interac-

tions. In practice, a researcher will have to select edge bins based on the empirical distribution of edge weights in their dataset or by other heuristics.

- **directed, unweighted:** we treat all the edge weights the same, and divide the edges into four bins by their directions. That is, we distinguish whether the edge is only an outgoing edge from the target node, only an incoming edge from the target node, both an outgoing and incoming edge, or no edge between them.
- **directed, weighted:** we divide all the edges by their weight values and directions. Specifically, we consider if the out-going edge weight is of 0, 1, 2 – 5, > 5 and if the in-coming edge weight is of 0, 1, 2 – 5, > 5. There are 16 bins in total.

Across one train/test split, we observe in Table 1 that the performance of Version I is better than Version II. Considering the edge direction and weights using the method introduced in Version I both seem to help improve the NE compared to the undirected/unweighted version. (Note lower NE means better performance.) We observe that considering the edge direction reduces NE by 0.09% while considering the edge weight reduces NE by 0.01%.

undirected, unweighted	100% (reference)	
Version I	undirected, weighted	-0.01%
	directed, unweighted	-0.09%
	directed, weighted	-0.11%
Version II	undirected, weighted	+0.02%
	directed, unweighted	-0.02%
	directed, weighted	-0.06%
friend’s conversion rate	+0.03%	
user embeddings	+0.02%	

Table 1: Undirected/unweighted comment networks have variability in predictive performance. We observe overall higher performance with Version I compared to Version II. Both Versions exhibit higher performance than friend conversion and user embeddings from PBG.

To compare the LINK-Naive Bayes approach with other node embedding methods, we also consider the following two baseline models:

- **friend’s conversion rate:** for each user u , we take the ratio of conversion rate among u ’s undirected neighborhoods as the u ’s conversion probability. We can view this as a homophily-based feature.
- **user embeddings (PBG):** we implement PyTorch-BigGraph (PBG) (Lerer et al. 2019) to train a 64-dimensional embedding from the user-to-user comment network. PBG scales to large graphs and can support multi-relation graphs. PBG employs a block decomposition of the adjacency matrix and trains on edges from one bucket at a time. We implement the embeddings with a 0.01 learning rate and a negative batch size of 50.

Application on obgn-proteins: We next demonstrate the LINK-Naive Bayes approach on the Node Property Prediction task from the publicly available Open Graph Benchmark (OGB). OGB includes several large-scale datasets for

¹We only include Version II* results for the obgn-proteins.

Method	Avg. ROC-AUC
undirected, unweighted (reference)	0.51
Version I (undirected, weighted)	0.65
Version II (undirected, weighted)	0.57
Version II* (undirected, weighted)	0.70

Table 2: Compare the predictive performance on the `ogbn-proteins` dataset with the LINK-Naive Bayes method using Version I/II as well as an undirected, unweighted version as a reference. We observe meaningful performance improvement in both Versions I and II.

node prediction among other tasks and includes data splits for training models (Hu et al. 2020a). Another advantage of using a dataset from OGB is that we can compare the predictive performance of our approach to an exhaustive set of baselines on the public leaderboard.

We include results on the `ogbn-proteins` dataset, which is an undirected and weighted network consisting of 132,534 nodes, 39,561,252 edges, and 112 different labels to predict. Nodes represent proteins and edges indicate a biological association between proteins. Each edge is associated with an 8-dimensional feature representing the strength of a relationship between two proteins. This is the only binary task dataset in OGB and the only one that allows us to set-up multi-layer networks to compare with in the next section. For this section, we will treat the network as a single layer by averaging edge weights across the 8-dimensional feature vector. For both Version I/II, we construct 4 edge bins based on intervals constructed from the 25th, 50th, 75th percentiles of the edge weight. Here we only construct edge bins based on edges that exist between nodes in order to accelerate calculations. Consistent with the data splitting recommendation (Hu et al. 2020a) and with reporting ROC-AUC metrics, we split the nodes into train/validate/test splits according to what’s provided and report average ROC-AUC across the 112 tasks in Table 2. We do not incorporate any additional benchmarks as the public leaderboard already tracks top submissions.

For all performance metrics for this dataset, we consider ROC-AUC metrics to be consistent with results presented on the public leaderboard². We see that Version I/II/II* outperform the undirected, unweighted version. We use a paired Wilcoxon signed rank test (with a Bonferroni correction) to test whether the ROC-AUC metrics from Version I/II/II* are greater than the undirected, unweighted version and observe p -values < 0.01 . We also test for pairwise differences between the different versions and observe a statistically significant difference. When comparing this performance to the public leaderboard as of September 2023, we can see that both Version I/II* are comparable to `node2vec`’s test performance of 0.6881 at #22 on the leaderboard. We see other baseline methods like MLP have 0.7204 and GCN have 0.7251 ROC-AUC, respectively. The protein network has higher density than social networks and has a small graph diameter (Hu et al. 2020a), which may explain in part the

²https://ogb.stanford.edu/docs/leader_nodeprop/

lower performance of a method like `node2vec`.

While we observe increased average ROC-AUC performance across the different versions relative to an undirected and unweighted version of the `ogbn-proteins` dataset, it is possible that LINK may still have more limited performance relative to the leaderboard due the train/test splits being based on species. There are more within-species edges relative to between-species edges. In the hypothetical case where there are no edges between species, the across-layer task would look more similar to an across-network task and LINK would not be applicable.

Summary of Empirical Results: When comparing predictive performance for both datasets, our takeaways are that:

- Adding the edge direction and/or weight can help improve predictive performance compared to an undirected/unweighted network. We observe edge weight is helpful for `ogbn-proteins` but not ad click prediction.
- Version I has better predictive performance than Version II for the ad click prediction task but the reverse is true for the `ogbn-proteins`.
- Both sets of empirical results should be interpreted as a case study of this methodology. Whether one method is better than another will also depend in part on the underlying data generating mechanism and the network itself.

Node Prediction on Multi-layer Networks

Users can exhibit different interaction patterns across platforms. We provide simple descriptive statistics around interaction network differences across the Facebook and Instagram platforms in the Appendix. Such differences could potentially result in varying predictive information when using the same approach but different network layers. As a quick illustration, in Table 3, we compute the percent change in test NE using Version I (directed/weighted) for the ad click prediction task for five Facebook user-to-user interaction networks – *comment*, *tag*, *reshare*, *mention*, and *direct post*. These edges represent different types of interactions that can occur between users. A comment edge indicates a user commented on another user’s content. A tag edge indicates a user tags another user on a post. This is similar to a mention but mentions can happen in comments or replies where a user writes another user’s name and the user is notified. Finally, a reshare means an edge exists if a user shares the same content from another user. We observe heterogeneity in the predictive performance using these different interaction networks.

<i>comment</i>	<i>mention</i>	<i>reshare</i>	<i>tag</i>	<i>post</i>
100% (ref)	+0.10%	+0.13%	+0.20%	+0.15%

Table 3: We compare the relative NE performance across different interaction types using the comment network as the reference network. We observe variability in predictive performance.

While the comment network has the best predictive performance, there is still some predictive signal in other network layers that’s above a random classifier, albeit lower

than just the comment network. Therefore, the question is whether we can achieve higher performance when *combining* networks than any single network layer in isolation. All the above differences motivate us to further extend the LINK-Naive Bayes approach for multi-layer networks, in which we can consider multiple user-to-user networks simultaneously during training. We call this new approach MultiLayerLINK-NaiveBayes. We show in Figure 1 (right) a schematic of combining multi-layers for a within-layer task.

Model for Multi-layer Networks

We assume a K -layer network $\mathcal{G} = \left(\bigcup_{k=1}^K V^{(k)}, \bigcup_{k=1}^K E^{(k)} \right)$ contains a set of k single-layer networks $G^{(1)} = (V^{(1)}, E^{(1)}), \dots, G^{(K)} = (V^{(K)}, E^{(K)})$, where $V^{(k)}$ and $E^{(k)}$ represent the node set and edges representing different types of edges in network i , respectively. For the most comprehensive case, we consider $E^{(k)}$ to encompass both directed and weighted edges. Now, suppose we want to have the same prediction task as before, where each node in $\mathcal{V} := \bigcup_{k=1}^K V^{(k)}$ has a label being either $+1$ or -1 , and for a target node u , we want to predict its label using known labels from neighbors of u at different layers.

(Version 0) Linear Combination of Single layers: The simplest approach is to treat the multi-layer score as a linear combination of the classifiers built on each layer separately. Specifically, let $n_{+1}^{(k)}$ ($n_{-1}^{(k)}$) denote the number of $+1$ -labeled (-1 -labeled) nodes in the k th layer among training samples, and $d_{i,+1}^{(k)}$ ($d_{i,-1}^{(k)}$) as node i 's degree with $+1$ (-1) nodes in the k th layer among training samples. We similarly represent the set of bins as $\{Bin_{w^{(k)}}, w^{(k)} \in W^{(k)}\}$. Suppose layer k has $N^{(k)}$ nodes, and they are divided into $|W^{(k)}|$ bins in the corresponding single-layer prediction task. We denote $\mathbb{1}_{x_i^{(k)}=w^{(k)}}$ if the edge between the target node and the source node i is of $Bin_{w^{(k)}}$ at layer k . For each bin in layer k , we set a weight $\lambda_{w^{(k)}}$, and define $\vec{\lambda} = \bigcup_{k=1}^K \bigcup_{w^{(k)} \in W^{(k)}} \lambda_{w^{(k)}}$. Following Equation (2), we can write the score function as:

$$f(\alpha, \vec{\lambda}, x) = \alpha + \underbrace{\sum_{k=1}^K \sum_{w^{(k)} \in W^{(k)}} \lambda_{w^{(k)}} \sum_{i=1}^{N^{(k)}} \mathbb{1}_{x_i^{(k)}=w^{(k)}} \log \left[\frac{d_{i,+1}^{(k)} + 1}{d_{i,-1}^{(k)} + 1} \cdot \frac{n_{-1}^{(k)} + 2}{n_{+1}^{(k)} + 2} \right]}_{\text{one-hop aggregation at layer } k} \underbrace{\quad}_{\text{two-hop aggregation by edge bin at layer } k} \quad (3)$$

This score function is a linear combination of the scores from different layers. As before, we can then estimate α and $\vec{\lambda}$ with Logistic regression to minimize the log loss.

(Version I) Learning Weights for Edge Bins: Alternatively, we can couple edges between layers and place them into edge bins based on the product category of bins at all single layers. That is, we first construct edge bins based on edge weight or direction as introduced in the ‘‘Create Edge

Bins’’ subsection. For a given edge that exists across multiple network layers, we categorize it into one of the bins based on edge weight/direction for that layer. Then to couple an edge across multiple layers, we consider the joint categorization of edge bins across multiple layers. The final bins in the multi-layer network are the product category of bins at all single layers. That is, if at each layer k , the edges are divided into $|W^{(k)}|$ bins, then the total number of bins for the K -layer network is $\prod_{k=1}^K |W^{(k)}|$. As before, we denote the set of bins as $\{Bin_{\tilde{w}}, \tilde{w} \in \tilde{W}\}$, where \tilde{W} is the set of bins’ labels, and we denote $\mathbb{1}_{x_i=\tilde{w}}$ if the relationship between the target node and the source node i is of $Bin_{\tilde{w}}$.

We then assign different weights for different bins generated at different layers. Let $\lambda_{\tilde{w}}^{(k)}$ be the weight of $Bin_{\tilde{w}}$ for layer k , and denote by $\vec{\lambda}$ the vector of $\lambda_{\tilde{w}}^{(k)}$'s: $\vec{\lambda} = \bigcup_{k=1}^K \bigcup_{\tilde{w} \in \tilde{W}} \lambda_{\tilde{w}}^{(k)}$. We can write the classifier’s score function as the following:

$$f(\alpha, \vec{\lambda}, x) = \alpha + \underbrace{\sum_{k=1}^K \sum_{\tilde{w} \in \tilde{W}} \lambda_{\tilde{w}}^{(k)} \sum_{i=1}^N \mathbb{1}_{x_i=\tilde{w}} \cdot \log \left[\frac{d_{i,+1}^{(k)} + 1}{d_{i,-1}^{(k)} + 1} \cdot \frac{n_{-1}^{(k)} + 2}{n_{+1}^{(k)} + 2} \right]}_{\text{one-hop aggregation at layer } k} \underbrace{\quad}_{\text{two-hop aggregation by coupling edge bins across layers}}$$

As explained, in this method we construct the finer bins $\{Bin_{\tilde{w}}, \tilde{w} \in \tilde{W}\}$ by the Cartesian product of each layer’ bins. We then weight the scores in each layer using the finer bins, and the score function is a linear combination of every layers’ scores. The advantage of using the finer bins is that it contains multi-layer information of a node, which might capture some interaction effect across different layers, that cannot be captured by Version 0 above.

(Version II/II*) Classifier within each Edge Bin: Similar to the Version II in the single-layer network section, we can apply a Naive Bayes classifier within each bin for multi-layer networks. We denote $d_{i,+1,\tilde{w}}$ ($d_{i,-1,\tilde{w}}$) as the number of $+1$ -labeled (-1 -labeled) neighbors s of node i with edge type of (i, s) in $Bin_{\tilde{w}}$. We also denote by n_{-1} (n_{+1}) the number of nodes in \mathcal{V} with the label -1 ($+1$).

With exactly the same derivation as Version II in the previous single-layer section, we have the log-likelihood ratio $\log(LR(x))$ as follows:

$$\approx \sum_{i=1}^N \sum_{\tilde{w} \in \tilde{W}} \mathbb{1}_{x_i=\tilde{w}} \cdot \log \left[\frac{d_{i,+1,\tilde{w}} + 1}{d_{i,-1,\tilde{w}} + 1} \cdot \frac{n_{-1} + |\tilde{W}|}{n_{+1} + |\tilde{W}|} \right]$$

Again, we can set different weights $\lambda_{\tilde{w}}$ on different bins and train the weights with Logistic regression to minimize the log-loss.

$$f(\alpha, \vec{\lambda}, x) = \alpha + \sum_{\tilde{w} \in \tilde{W}} \lambda_{\tilde{w}} \sum_{i=1}^N \mathbb{1}_{x_i=\tilde{w}} \cdot \log \left[\frac{d_{i,+1,\tilde{w}} + 1}{d_{i,-1,\tilde{w}} + 1} \cdot \frac{n_{-1} + |\tilde{W}|}{n_{+1} + |\tilde{W}|} \right]$$

Intuitively, in this version we still construct the finer bins $\{Bin_{\tilde{w}}, \tilde{w} \in \tilde{W}\}$ by the Cartesian product of each layer’s bins. In Version I we only use the bins division to weight the one-hop scores in different layers, while in Version II we calculate the one-hop score directly using the finer bins. To create Version II*, we similarly separate out the one-hop and two-hop edge bins as in the single-layer set-up.

Empirical Results

Application on Facebook: To test the performance of the multi-layer LINK-Naive Bayes classifier, we use the two largest Facebook interaction networks – *comment* and *mention* – to construct a two-layer network. As before, we select active U.S. users with at least 20 Facebook friends who were exposed to Retail ads on 2022-07-30. The number of edge bins in the two-layer network is the product of the number of bins for each network layer. Since our single-layer results comparing edge direction and weight show that only edge direction improves predictive performance across both versions in Table 1, for simplicity, we ignore the edge weights and we build directed, unweighted user-to-user networks based on uses’ commenting or mentioning activities. Within each layer, for a directed pair of nodes (u, v) , there are four possible relationships between them: there is no edge between u and v , there is a directed edge going from u to v , there is a directed edge going from v to u , and there are reciprocal edges.

We then set-up a within-network task similar to the within-network section where we randomly select 80% nodes as the training sample, 10% as validation sample, and the remaining 10% as test sample. We report the percent change in test NE in Table 4 relative to the baseline of using the single-layer comment network.

comment (Version I in Sec 3)	100% (reference)
mention (Version I in Sec 3)	+0.09%
comment + mention (Version 0)	−0.03%
comment + mention (Version I)	−0.07%
comment + mention (Version II)	+0.01%
comment + mention (PBG)	+0.06%

Table 4: We compare MultiLayerLINK-NaiveBayes across the different Versions 0/I/II and compare to the single-layer network prediction performance. All NE metrics are compared relative to the single-layer network (comment for Version I).

We observe that:

- For Version 0: adding the multi-layer information (from the mention network) is helpful compared to just the comment network in isolation.
- For the multi-layer Version I approach: it outperforms Version 0. We expect this since Version I uses the finer bins (Cartesian product of each layer’s bins) to weight the scores, which can capture the interaction effect across different layers that is missing in Version 0.
- For Version II: it performs even worse than the single layer baseline. One explanation could be that, our data set

is imbalanced (only about 6 percent of nodes have label 1), so when we calculate the Naive Bayes scores directly for each bin (which might only contain a small number of data points), the estimation of the posterior probability might be highly volatile.

Application on obgn-proteins: We next compare predictive performance on the *obgn-proteins* described in the earlier single-layer section. We create multi-layer networks by utilizing the 8-dimensional edge weight vector to construct 8 separate network layers corresponding to different edge weights. Within each layer, for an undirected pair of nodes (u, v) , we bin edges based on a 75th percentile cut-off for edge weights and only consider bins where an edge exists. We again use the same train/validation/test splits described in the earlier section and that’s available in the public data. We report the Average ROC-AUC in Table 5. We observe that:

- For Version 0, we observe that Version I/II/II* have higher average predictive performance.
- For the multi-layer set-up, Version I/II/II* both have higher average ROC-AUC compared to their single-layer counterparts with p -values < 0.01 .
- We do not observe a meaningful performance difference in the multi-layer setting for Version II vs. II*.

Method	Avg. ROC-AUC
Version 0 (undirected, weighted)	0.68
Version I (undirected, weighted)	0.72
Version II (undirected, weighted)	0.75
Version II* (undirected, weighted)	0.76

Table 5: Compare the predictive performance on the *obgn-proteins* dataset with the MultilayerLINK-Naive Bayes method using Version 0/I/II/II* for the multi-layer network set-up.

Time Complexity of MultiLayerLINK-NaiveBayes

We finally compute the time complexity of the multilayer approach. For a k -layer network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \cup_{k=1}^K V^{(k)}$, and $\mathcal{E} = \cup_{k=1}^K E^{(k)}$, assume that we have l bins; that is, $\vec{\lambda}$ is of dimension l in the logistic regression step. The time complexities for different steps are:

- for getting the NB estimators, the time complexity is $O(|\mathcal{E}|)$,
- for calculating the score functions for all nodes, the time complexity is $O(|\mathcal{E}|)$, and
- for training the logistic regression, the time complexity is $O(l \cdot |\mathcal{V}|)$.

In total, the complexity for our model is $O(|\mathcal{E}| + l \cdot |\mathcal{V}|)$.

Discussion and Future Work

Social network platforms enable users to connect with one another in numerous ways. As a result, organizations have

an opportunity to combine and learn from different user-to-user network representations. In this paper, we first demonstrate that the various IG/FB user-to-user networks exhibit different properties and minimal edge overlap. Given these empirical differences between user networks, we first extend a single-layer network prediction approach called LINK-Naive Bayes to account for edge direction and weights. We introduce a business-relevant prediction task, ad click prediction, as well as utilize publicly available data from the Open Graph Benchmark to compare our different approaches. We observe the most performance gain for ad click prediction in incorporating edge direction. For the Open Graph Benchmark dataset, ROC-AUC results improve from 0.51 in the undirected and unweighted case to 0.57-0.70 in the different versions for incorporating weight information. Next, we introduce MultiLayerLINK-NaiveBayes, which incorporates multi-layer network signals for the within-network prediction task. We observe higher predictive performance gain for Version I for ad click prediction and Version I/II/II* relative to their single-layer counterparts. The results in our paper should be viewed as a case study on how to deploy LINK to multi-layered networks. We may see in other settings that other versions do better than what was presented in this work.

Limitations: While we observe how edge direction and weight along with multilayer networks can improve LINK, these improvements are only in the top 20 of the leaderboard of the Open Graph Benchmark. One hypothesis is that the protein-protein dataset could have a stronger homophily or heterophily signal in which case the GNN variations are performing better. Many of the other baseline methods in the leaderboard incorporate hyperparameter tuning, so it may be possible to further tune LINK by parameterizing it more. Next, while we observe predictive performance gains with the extensions introduced in this paper, it is difficult to assess the practical value of these gains. Future work can focus on evaluating how offline performance gains may translate to real-world results.

Future work could consider comparing approaches on synthetic graphs (Tsitsulin et al. 2022) for generalizing when this method does better to alternatives. While the results in this paper were shown for binary classification and multiplex networks, the methodology can generalize to multi-class prediction and more general multilayer networks. For multi-class prediction, we could extend by treating it as a one-vs-rest binary task across each class. For more general multilayer networks, we could extend the methodology by creating edge bins based on how nodes are connected across layers. We leave these extensions for future work. Beyond node prediction, LINK-NB is applicable to link prediction as discussed in the “Link Prediction” appendix section.

We also consider additional possible directions for extending our work:

- *Temporal networks:* Although our study focuses on predictions on network representations at a particular time point, user networks are dynamic over time. If networks are correlated over time, then we’d expect our approach to generalize to prediction over time. However, if net-

works are very different across time, then future work would need to account for the changes in edges over time. One option is to encode temporal motifs (Paranjape, Benson, and Leskovec 2017) or to account for time intervals between new edges appearing.

- *Missingness mechanisms and noisy observations:* In our current study, we assume node labels are missing completely at random (MCAR) (Little and Rubin 2019). Across various business prediction tasks, the MCAR assumption is likely unreasonable. For example, for demographic user prediction³ we may expect training data labels to vary by other characteristics. In these cases, further work is needed to understand how robust our approach is to violations of MCAR. We also assume we observe accurate ground-truth of node labels and edges. GNN can under perform when there are noisy edges (Liu et al. 2022) and future work may consider how robust LINK and its variants are to noise.
- *Cold start for new users:* Finally, we assume the set of users in the sample is known and fixed. However, in practice, new users are joining interactions over time. We consider new users a cold start problem because there aren’t existing network signals to leverage. One thought is trying to identify similar nodes in the network and using those as a proxy.

In total, combining multi-layers network signals and accounting for both edge direction and weights for within-network prediction is a critical step forward in enabling organizations to optimally learn from network representations.

Acknowledgements

We thank Alex Dow, Johan Ugander, and Alyssa Wang for helpful comments.

Code Availability: The code to replicate the analysis on the `ogbn-proteins` is available at <https://github.com/facebookresearch/LINKExtension>.

References

- Aguiar, I.; Taylor, D.; and Ugander, J. 2022. A tensor factorization model of multilayer network interdependence. *arXiv preprint arXiv:2206.01804*.
- Altenburger, K. M.; and Ugander, J. 2018. Monophily in social networks introduces similarity among friends-of-friends. *Nature Human Behaviour*, 2(4): 284–290.
- Altenburger, K. M.; and Ugander, J. 2021. Which Node Attribute Prediction Task Are We Solving? Within-Network, Across-Network, or Across-Layer Tasks. In *ICWSM*, 38–48.
- Boccaletti, S.; Bianconi, G.; Criado, R.; Del Genio, C. I.; Gómez-Gardenes, J.; Romance, M.; Sendina-Nadal, I.; Wang, Z.; and Zanin, M. 2014. The structure and dynamics of multilayer networks. *Physics Reports*, 544(1): 1–122.

³<https://tech.facebook.com/artificial-intelligence/2022/06/adult-classifier/>

- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems*, 26.
- Brea, J.; Burrone, J.; Minnoni, M.; and Sarraute, C. 2014. Harnessing mobile phone social network topology to infer users demographic attributes. In *Proceedings of the 8th Workshop on Social Network Mining and Analysis*, 1–9.
- Chakrabarti, D.; Funiak, S.; Chang, J.; and Macskassy, S. 2014. Joint inference of multiple label types in large networks. In *International Conference on Machine Learning*, 874–882. PMLR.
- Chami, I.; Abu-El-Haija, S.; Perozzi, B.; Ré, C.; and Murphy, K. 2020. Machine learning on graphs: A model and comprehensive taxonomy. *arXiv preprint arXiv:2005.03675*, 1.
- De Domenico, M.; Solé-Ribalta, A.; Cozzo, E.; Kivelä, M.; Moreno, Y.; Porter, M. A.; Gómez, S.; and Arenas, A. 2013. Mathematical formulation of multilayer networks. *Physical Review X*, 3(4): 041022.
- Dong, Y.; Chawla, N. V.; Tang, J.; Yang, Y.; and Yang, Y. 2017. User modeling on demographic attributes in big mobile social networks. *ACM Transactions on Information Systems (TOIS)*, 35(4): 1–33.
- Fakhraei, S.; Foulds, J.; Shashanka, M.; and Getoor, L. 2015. Collective spammer detection in evolving multi-relational social networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1769–1778.
- Feng, M.-H.; Hsu, C.-C.; Li, C.-T.; Yeh, M.-Y.; and Lin, S.-D. 2019. Marine: Multi-relational network embeddings with relational proximity and node attributes. In *The World Wide Web Conference*, 470–479.
- Gallagher, B.; and Eliassi-Rad, T. 2008. Leveraging label-independent features for classification in sparsely labeled networks: An empirical study. In *International Workshop on Social Network Mining and Analysis*, 1–19. Springer.
- Goel, S.; and Goldstein, D. G. 2014. Predicting individual behavior with social networks. *Marketing Science*, 33(1): 82–93.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 855–864.
- Hashemi, F.; Behrouz, A.; and Lakshmanan, L. V. 2022. FirmCore Decomposition of Multilayer Networks. In *Proceedings of the ACM Web Conference 2022*, 1589–1600.
- He, X.; Pan, J.; Jin, O.; Xu, T.; Liu, B.; Xu, T.; Shi, Y.; Atallah, A.; Herbrich, R.; Bowers, S.; et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, 1–9.
- Henderson, K.; Gallagher, B.; Eliassi-Rad, T.; Tong, H.; Basu, S.; Akoglu, L.; Koutra, D.; Faloutsos, C.; and Li, L. 2012. Rolx: structural role extraction & mining in large graphs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1231–1239.
- Henderson, K.; Gallagher, B.; Li, L.; Akoglu, L.; Eliassi-Rad, T.; Tong, H.; and Faloutsos, C. 2011. It’s who you know: graph mining using recursive structural features. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 663–671.
- Hill, S.; Provost, F.; and Volinsky, C. 2006. Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science*, 21(2): 256–276.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020a. Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33: 22118–22133.
- Hu, Z.; Dong, Y.; Wang, K.; and Sun, Y. 2020b. Heterogeneous graph transformer. In *Proceedings of the Web Conference 2020*, 2704–2710.
- Huberman, B. A.; Romero, D. M.; and Wu, F. 2008. Social networks that matter: Twitter under the microscope. *arXiv preprint arXiv:0812.1045*.
- Kivelä, M.; Arenas, A.; Barthelemy, M.; Gleeson, J. P.; Moreno, Y.; and Porter, M. A. 2014. Multilayer networks. *Journal of Complex Networks*, 2(3): 203–271.
- Lerer, A.; Wu, L.; Shen, J.; Lacroix, T.; Wehrstedt, L.; Bose, A.; and Peysakhovich, A. 2019. Pytorch-biggraph: A large scale graph embedding system. *Proceedings of Machine Learning and Systems*, 1: 120–131.
- Lim, D.; Li, X.; Hohne, F.; and Lim, S.-N. 2021. New benchmarks for learning on non-homophilous graphs. *arXiv preprint arXiv:2104.01404*.
- Little, R. J.; and Rubin, D. B. 2019. *Statistical Analysis with Missing Data*, volume 793. John Wiley & Sons.
- Liu, Y.; Zheng, Y.; Zhang, D.; Chen, H.; Peng, H.; and Pan, S. 2022. Towards unsupervised deep graph structure learning. In *Proceedings of the ACM Web Conference 2022*, 1392–1403.
- Loveland, D.; Zhu, J.; Heimann, M.; Fish, B.; Shaub, M. T.; and Koutra, D. 2023. On Performance Discrepancies Across Local Homophily Levels in Graph Neural Networks. *arXiv preprint arXiv:2306.05557*.
- Ma, Y.; Liu, X.; Shah, N.; and Tang, J. 2021. Is Homophily a Necessity for Graph Neural Networks? In *International Conference on Learning Representations*.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1): 415–444.
- Neville, J.; and Jensen, D. 2000. Iterative classification in relational data. In *Proc. AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, 13–20.
- Paranjape, A.; Benson, A. R.; and Leskovec, J. 2017. Motifs in temporal networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 601–610.

Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 701–710.

Rossi, E.; Charpentier, B.; Di Giovanni, F.; Frasca, F.; Günnemann, S.; and Bronstein, M. 2023. Edge Directionality Improves Learning on Heterophilic Graphs. *arXiv preprint arXiv:2305.10498*.

Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI Magazine*, 29(3): 93–93.

Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, 1067–1077.

Tsitsulin, A.; Rozemberczki, B.; Palowitch, J.; and Perozzi, B. 2022. Synthetic Graph Generation to Benchmark Graph Learning. *arXiv preprint arXiv:2204.01376*.

Ugander, J.; Karrer, B.; Backstrom, L.; and Marlow, C. 2011. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*.

Wang, B.; Gong, N. Z.; and Fu, H. 2017. GANG: Detecting fraudulent users in online social networks via guilt-by-association on directed graphs. In *2017 IEEE International Conference on Data Mining (ICDM)*, 465–474. IEEE.

Welling, M.; and Kipf, T. N. 2016. Semi-supervised classification with graph convolutional networks. In *J. International Conference on Learning Representations (ICLR 2017)*.

Xu, Z.; Du, B.; and Tong, H. 2022. Graph sanitation with application to node classification. In *Proceedings of the ACM Web Conference 2022*, 1136–1147.

Zheleva, E.; and Getoor, L. 2009. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. In *Proceedings of the 18th International Conference on World Wide Web*, 531–540.

Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33: 7793–7804.

Zitnik, M.; and Leskovec, J. 2017. Predicting multicellular function through multi-layer tissue networks. *Bioinformatics*, 33(14): i190–i198.

Zubcsek, P. P.; Katona, Z.; and Sarvary, M. 2017. Predicting mobile advertising response using consumer colocation networks. *Journal of Marketing*, 81(4): 109–126.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **yes**

- (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **yes**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **yes**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **yes**
- (e) Did you describe the limitations of your work? **yes**
- (f) Did you discuss any potential negative societal impacts of your work? **yes**
- (g) Did you discuss any potential misuse of your work? **yes**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **yes**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **yes**

2. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **yes**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **yes**
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **yes**
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **n/a**
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **yes**
- (f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? **n/a**

3. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- (a) If your work uses existing assets, did you cite the creators? **yes**
- (b) Did you mention the license of the assets? **n/a**
- (c) Did you include any new assets in the supplemental material or as a URL? **yes**
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **n/a**
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **yes**
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see ?)? **n/a**
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see ?)? **n/a**

Appendix

Analysis of Real-World Multilayer Networks

A multilayer network (Kivelä et al. 2014; Boccaletti et al. 2014; De Domenico et al. 2013) consists of a node set V of N nodes that can belong to different layers L such as a Facebook friendship edge or a comment edge between users. A multiplex network is a type of multilayer network that is limited to the same set of nodes across different network layers. While the approaches presented in this paper are applicable to the general multilayer set-up, we illustrate the methodology on multiplex networks.

While prior work analyzed the social structure on the FB friendship graph (Ugander et al. 2011), no work to our knowledge has done an empirical comparison of different network representations of users across different interactions. We can view these different representations as a multilayer network (Kivelä et al. 2014) and compare users (nodes) between layers. For each platform (FB and IG), we pull data for 2-weeks in June 2022 for active U.S users over 18-years old with at least 20 Facebook friends, and construct user-to-user interaction networks (FB/IG), friend and follow networks (FB/IG). Interaction networks are directed and weighted and can include numerous interaction types including comments, reshares, likes, comment likes, and mentions. We direct the interaction from the source user creating the interaction to the target user receiving the interaction. We limit all analysis to users who have explicitly linked their FB and IG accounts. All data was de-identified and analyzed in aggregate.

On IG, the out-degree distribution across all interactions has more weight on small degrees, while in-degree distribution is more heavy-tailed. That is, there is a larger proportion of users receiving many interactions than the proportion of users sending many interactions. We also observe that different interactions have different patterns with respect to who people interact with: direct reshares and mentions happen mainly between users who follow each other (“reciprocal follows”). Comment, overall likes, and likes on comments also happens more among reciprocal follows. On FB, in terms of degree distributions, unlike IG, for different interactions, the out-degrees do not always have more weight on small degrees except for comment and reshares. Users primarily interact with those with reciprocal follows. When comparing networks across IG and FB, the edge overlap is small indicating that a user’s connections on IG are different than those on FB. We find small correlations between structural properties like degree. In total, these empirical observations indicate that different user network layers may provide different values for node prediction tasks.

Derivation of Version II of Naive Bayes Classifier

Recall that we have the estimate

$$\begin{aligned}\hat{P}(y_{train} = +1) &= \frac{n_{+1}}{n_{+1} + n_{-1}}, \\ \hat{P}(X_i = w | y_{train} = +1) &= \frac{d_{i,+1,w} + 1}{n_{+1} + |W|}; \\ \hat{P}(y_{train} = -1) &= \frac{n_{-1}}{n_{+1} + n_{-1}}, \\ \hat{P}(X_i = w | y_{train} = -1) &= \frac{d_{i,-1,w} + 1}{n_{-1} + |W|}.\end{aligned}$$

By the formula of the Naive Bayes classifier (conditional independence), we have that $LR(x)$ is

$$\begin{aligned}&= \frac{P(y_{train} = +1)}{P(y_{train} = -1)} \cdot \prod_{i=1}^N \frac{P(X_i = x_i | y_{train} = +1)}{P(X_i = x_i | y_{train} = -1)} \\ &= \frac{P(y_{train} = +1)}{P(y_{train} = -1)} \cdot \prod_{w=1}^{|W|} \prod_{i:x_i=w} \frac{P(X_i = w | y_{train} = +1)}{P(X_i = w | y_{train} = -1)} \\ &= \frac{n_{+1}}{n_{-1}} \prod_{w=1}^{|W|} \prod_{i:x_i=w} \frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|}.\end{aligned}$$

Plugging in our estimates into the above equation, and taking log on both sides, we get

$$\begin{aligned}\log(LR(x)) &\approx C_2 + \\ &\sum_{i=1}^N \sum_{w=1}^{|W|} \mathbb{1}_{x_i=w} \log \left[\frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]\end{aligned}$$

which gives the expression for the Version II approach.

Generalizations of LINK-NB to Link Prediction

In this section we discuss how to apply LINK-NB for edge attributes prediction. Suppose that each pair of nodes (z_1, z_2) in the training data has a binary label $e_{(z_1, z_2)} \in \{+1, -1\}$. For example, $e_{(z_1, z_2)}$ can be the indicator of whether z_1 and z_2 are friends.

Our features are the rows of the adjacency matrix that corresponds to z_1, z_2 . That is, when ignoring edge directions and edge weights, for all node $i \in G$, we define $x_i := (\mathbb{1}_{i \text{ is a neighbor of } z_1}, \mathbb{1}_{i \text{ is a neighbor of } z_2})$ as the feature vector of i . Note that there are four possible values for x_i , and we can group the nodes by their values into four bins, which we call them as Bin_w for $w = 1, \dots, 4$. In the more complicated cases where we consider edge weights or directed edges or multiple layers, similar to node LINK-NB, we can also group nodes by values of x_i into generalized bins $\{Bin_w | w = 1, \dots, |W|\}$. We denote $\mathbb{1}_{x_i=w}$ as $x_i \in Bin_w$.

We have the following expression for the likelihood ratio for the Naive Bayes estimator

$$\begin{aligned}LR(x) &= \frac{P(e_{train} = +1) \cdot P(x | e_{train} = +1)}{P(e_{train} = -1) \cdot P(x | y_{train} = +1)} \quad (4) \\ &= \frac{P(e_{train} = +1) \cdot \prod_{i=1}^N P(x_i | e_{train} = +1)}{P(e_{train} = -1) \cdot \prod_{i=1}^N P(x_i | e_{train} = -1)} \quad (5)\end{aligned}$$

Let n_{+1} and n_{-1} be the count of (z_1, z_2) pairs with $e_{(z_1, z_2)} = +1$ and $e_{(z_1, z_2)} = -1$ respectively. Denote by $d_{i,+1,w}$ (and $d_{i,-1,w}$) the number of (z_1, z_2) pairs with label $+1$ (and -1), such that its corresponding $x_i = w$. Similar to previous sections, we have the following estimates,

$$\begin{aligned}\hat{P}(e_{train} = +1) &= \frac{n_{+1}}{n_{+1} + n_{-1}}, \\ \hat{P}(e_{train} = -1) &= \frac{n_{-1}}{n_{+1} + n_{-1}}, \\ \hat{P}(x_i = w | e_{train} = +1) &= \frac{d_{i,+1,w} + 1}{n_{+1} + |W|}, \\ \hat{P}(x_i = w | e_{train} = -1) &= \frac{d_{i,-1,w} + 1}{n_{-1} + |W|}.\end{aligned}$$

Substituting the above estimates into (5), we get the following expression for the log likelihood ratio $\log(LR(x))$:

$$\approx C_4 + \sum_{i=1}^N \sum_{w=1}^{|W|} \mathbb{1}_{x_i=w} \cdot \log \left[\frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]$$

Similar to the node attributes prediction case, we can then set different weights, λ_j 's, on different categories of x_i , and train α and the weights with Logistic regression with the score function

$$f(\alpha, \vec{\lambda}, x) = \alpha + \sum_{w=1}^{|W|} \lambda_j \sum_{i=1}^N \mathbb{1}_{x_i=w} \cdot \log \left[\frac{d_{i,+1,w} + 1}{d_{i,-1,w} + 1} \cdot \frac{n_{-1} + |W|}{n_{+1} + |W|} \right]$$

Formula for Normalized Entropy

Normalized entropy is the predictive log loss normalized by the entropy of the average empirical clicking rate. The lower the value is, the better predictive power the model has. Suppose there are N samples in the dataset, where each sample has a label $y_i \in \{+1, -1\}$, then the normalized entropy is defined as:

$$NE = \frac{-\frac{1}{N} \sum_{i=1}^n \left(\frac{1+y_i}{2} \log(p_i) + \frac{1-y_i}{2} \log(1-p_i) \right)}{-(p * \log(p) + (1-p) * \log(1-p))}, \quad (6)$$

where p_i is the estimated probability of user i 's clicking rate, and p is the average clicking rate in the dataset.