



HAL
open science

Link Prediction in Bipartite Networks

Şükrü Demir İnan Özer, Günce Keziban Orman, Vincent Labatut

► **To cite this version:**

Şükrü Demir İnan Özer, Günce Keziban Orman, Vincent Labatut. Link Prediction in Bipartite Networks. 28th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems (KES), Sep 2024, Sevilla, Spain. hal-04605758

HAL Id: hal-04605758

<https://hal.science/hal-04605758>

Submitted on 7 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Link Prediction in Bipartite Networks

Şükrü Demir İnan Özer¹, Günce Keziban Orman¹, and Vincent Labatut²

¹Galatasaray University, Computer Engineering Department, Istanbul, Turkey

²Laboratoire Informatique d’Avignon – UPR 4128, Avignon Université, Avignon, France

28th International Conference on Knowledge-Based and
Intelligent Information & Engineering Systems – KES 2024

Abstract

Bipartite networks serve as highly suitable models to represent systems involving interactions between two distinct types of entities, such as online dating platforms, job search services, or e-commerce websites. These models can be leveraged to tackle a number of tasks, including link prediction among the most useful ones, especially to design recommendation systems. However, if this task has garnered much interest when conducted on unipartite (i.e. standard) networks, it is far from being the case for bipartite ones. In this study, we address this gap by performing an experimental comparison of 19 link prediction methods able to handle bipartite graphs. Some come directly from the literature, and some are adapted by us from techniques originally designed for unipartite networks. We also propose to repurpose recommendation systems based on graph convolutional networks (GCN) as a novel link prediction solution for bipartite networks. To conduct our experiments, we constitute a benchmark of 3 real-world bipartite network datasets with various topologies. Our results indicate that GCN-based personalized recommendation systems, which have received significant attention in recent years, can produce successful results for link prediction in bipartite networks. Furthermore, purely heuristic metrics that do not rely on any learning process, like the Structural Perturbation Method (SPM), can also achieve success.

Keywords— Bipartite graphs; Link prediction; Recommendation systems; Graph representation learning.

1 Introduction

Networks are a very common paradigm to model real-world systems, as they allow explicitly representing the interactions between the objects that constitute these systems [6]. A whole subclass of such systems involves two different types of interacting objects, such as in user–item purchase interactions or job seeker–job advertisement applications. Their graph-based representation takes the form of *bipartite* networks, which contain two distinct types of nodes, and in which links always connect nodes of different types [26]. Due to their more constrained structure, bipartite networks require specific processing compared to *unipartite* (i.e. standard, a.k.a. homogeneous) networks.

This is notably the case when performing link prediction, a task that consists in determining which links miss from a network known to be incomplete, or which links are likely to form in the forthcoming state of an evolving network. This is a major research topic of network science, which has been studied widely on unipartite networks [21]. It is useful in a number of applications, including recommender systems, spam detection, privacy control, network routing, and bibliometrics [15]. One can broadly distinguish three main approaches: (i) ranking candidate pairs of nodes using a heuristic metric that does not rely on any learning process [21], (ii) leveraging the network topological features to train a predictor into predicting such score [1], and (iii) perform representation learning to produce graph embeddings and use them during classification instead of features [10].

Although extensively studied in unipartite networks, link prediction in bipartite networks remains relatively underexplored [17, 2]. In this work, our goal is to address this gap, by proposing a comprehensive experimental study that compares a selection of link prediction methods on a bipartite network benchmark. We apply the most widespread methods from each one of the three families of approaches identified above for unipartite networks. In the process, we adapt them to bipartite networks when needed. In addition, we propose a fourth approach by leveraging state-of-the-art GCN-based personalized recommendation techniques [8] to perform link prediction in bipartite networks. Our contributions are as follows:

- Proposing a taxonomy of link prediction methods specifically designed for bipartite networks. We supplement it with visual representation to aid readers in understanding how these methods differ.
- Repurposing personalized recommendation models for accurate link prediction in bipartite networks.
- Investigating the effectiveness of using node centrality measurements as features to train supervised learning models for link prediction.
- Assessing the performance of 19 link prediction methods on three real-world datasets with varying sizes and densities.

The rest of this paper is organized as follows. In Section 2, we define the relevant concepts and describe in further detail the methods that we use. Then, in Section 3, we present our experimental setup, interpret our results, and discuss their implications. Lastly, we summarize our study and findings to conclude in Section 4, and discuss the future directions that could be pursued.

2 Methods

Bipartite networks are characterized by their nodes being distributed over two mutually exclusive sets U and V , and their edges $L \subseteq U \times V$ connecting only nodes from two different sets. Let us denote $L' = (U \times V) \setminus L$ as the set of links missing from the network. The problem of link prediction can then be defined as finding the proper links in L' that are more likely to occur in the next evolution of the network. This prediction can be performed through many different approaches, many of which rely on the correct scoring of candidate links of L' by learning-based methods. They all use the current network structure, which is encoded in L . In this section, we describe the different approaches that we later assess for the link prediction task in bipartite networks. For the remainder of the paper, we will refer to the nodes in U as “left nodes” and the others in V as “right nodes”.

2.1 Traditional Link Scores

Traditional link scores are the metrics that quantify each link $\in L'$ by using the structure of L . These scores are then ranked to choose the most likely links for predictions. To evaluate the performances of different link scores and to compare them with other approaches, we adopt a standard supervised method [31]. We split L into two sets: L_{train} and L_{test} . We use the graph structure of L_{train} to calculate these scores for all possible links, i.e. the set of $L_{test} \cup L'$. The links in this set are then ranked in decreasing order according to the scores they received. The first $|L_{test}|$ links having the highest scores constitute our prediction set, L_{pred} . The success of predictions is evaluated by comparing the links from L_{pred} with the links from L_{test} .

Most of the existing link scores rely on the concept of triadic closure observed in social networks [22], which proposes that the more neighbors two nodes share, the more likely they are to be connected. In bipartite networks, however, two nodes of the same type cannot be directly connected, and there is consequently no triangle. Hence, we employ link prediction methods that do not rely on triadic closure. In the following scores, $x, v \in U$, $y, u \in V$, k , A and l are the nodes from two sets, the degree of a node, the adjacency matrix and the length of the path, respectively.

Kovács et al. [14] introduced the *L3 Index* after observing that the triadic closure principle was ineffective for certain networks. L3 is the degree-normalized count of paths of length three connecting two nodes.

$$L3(x, y) = \sum_{u, v} \frac{A_{xu}A_{uv}A_{vy}}{\sqrt{k_u \times k_v}}. \quad (1)$$

To compare how the length of the paths affect the link prediction performance, we also propose to use two variations of L3, that we refer to as L5 and L7. Those indices are computed the same way, except they count the paths of length five and seven, respectively.

Similarly, the *Katz Index* measures the connectivity between nodes by considering the number of paths penalized with an attenuation factor α [13]. Longer paths are penalized in an exponentially increasing manner.

$$Katz(x, y) = \sum_{l=1}^{\infty} \alpha^l (A^l)_{xy}, \quad (2)$$

Local Paths (LP) [19] could be viewed as a combination of Katz and L3. It is computed like Katz, but only for paths with length of two and three. The paths of length three have an attenuation factor ϵ :

$$LP(x, y) = (A^2)_{x,y} + \epsilon(A^3)_{x,y}. \quad (3)$$

In bipartite networks, even powers of A allow counting paths that connect nodes of the same type. However, in our case, x and y belong to different node sets, therefore the term $(A^2)_{x,y}$ from Eq. 3 takes 0 for all (x, y) .

The *Preferential Attachment Index* (PA) assumes that nodes with higher degrees are more likely to be linked together [3].

$$PA(x, y) = k_x \times k_y. \quad (4)$$

The *Geodesic Distance Index* (Dist) is the reciprocal of the distance between two nodes, i.e. of the length of the shortest path between these nodes, without any normalization. By taking the reciprocal, we penalize links connecting nodes that have longer paths between them.

The *Structural Perturbation Method* (SPM) studies the change in Eigenvalues when the adjacency matrix is perturbed, to measure the structural consistency of the network [20]. The links to be predicted are ranked according to their scores in the perturbed matrix \tilde{A} , which is obtained through first-order perturbation.

$$\tilde{A} = \sum_{k=1}^N (\lambda_k + \Delta\lambda_k) x_k x_k^T, \quad (5)$$

where λ_k and x_k are respectively the Eigenvalue and the corresponding Eigenvector of matrix A^R , after the links in the perturbation set are removed from the network. Expressions $\lambda_k + \Delta\lambda_k$ and $x_k + \Delta x_k$ are the corrected Eigenvalue and its corresponding Eigenvector.

2.2 Link Prediction via Topological Features

One of the common methods of link prediction relies on supervised learning techniques. In this work, we propose to represent each node pair with network-related topological features and label them with the positive or negative classes depending on whether there is a link in L_{train} or not, respectively. Then, we apply classification methods. A detailed schema of the framework we propose can be seen in Fig. 1. At first, we split the network set into training and testing subsets. Then, all node pairs from each network are represented by the PageRank [23] scores of the nodes involved, as well as the PA score of each link. We also employ degree, closeness, betweenness, Eigenvector, and Katz centrality measures as reviewed by Landherr et al. [18] with detailed formulas provided in their study.

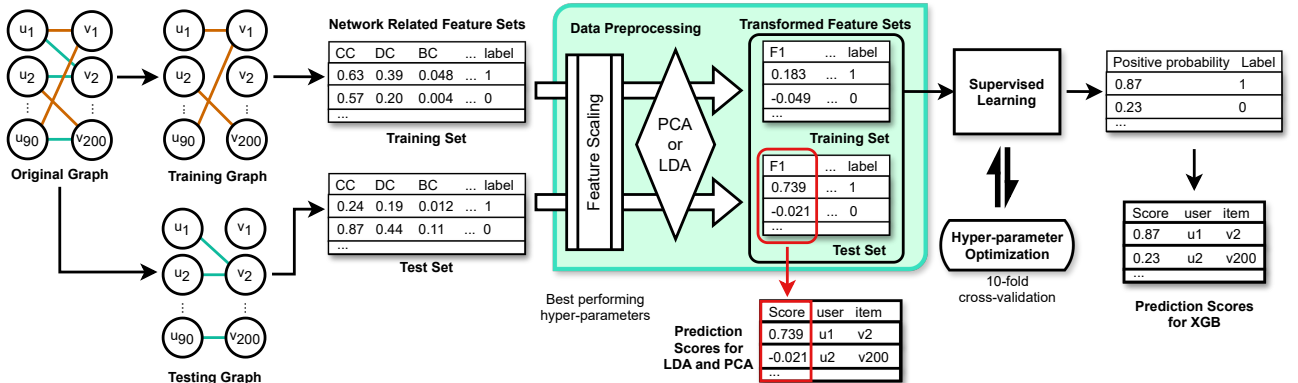


Figure 1: Representation of the link prediction method that relies on network topological features (Section 2.2).

For an unbiased training process, we use uniformly random negative sampling to include as many negative instances as positive ones, since the class numbers are highly imbalanced. We employ the XGBoost [5] classifier model for supervised learning. It is an ensemble learning model that trains multiple gradient-boosted decision trees as estimators. We refer to this method as XGB in the rest of the paper. We also perform dimension reduction into one single feature by applying Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA). We evaluate XGB, PCA, and LDA results separately in the experiments.

2.3 Embedding-based Link Prediction

In recent years, network representation learning, a.k.a. graph embedding, has gained high attention because it allows for the representation of network objects, i.e., nodes, links, or whole networks, as lower-dimensional real-valued vectors. The aim of these embedding techniques is to reflect the topological differences between the objects in the network into the novel embedded features. Some popular methods developed for this goal include DeepWalk [24], node2vec [10], GraphSAGE [11], and LINE [27]. These generic embedding methods are all unsupervised. They employ techniques such as random walks, skip-gram models, and neural networks to generate vector representations. However, they are not developed to handle bipartite networks and may

overlook the differences in node types. Embedding techniques developed for bipartite networks are rare: in this article, we use metapath2vec, BiNE, and BiGI.

Dong et al. [7] introduced metapath2vec (MP2V) as an unsupervised embedding method for multipartite networks. Building on node2vec [10], it applies skip-gram maximization along with a novel approach that employs meta-path guided random walk. These strategies allow capturing both structural and semantic correlations across different node types. BiNE (Bipartite Network Embedding) attempts to capture the structural properties of bipartite networks by executing biased random walks [9]. This bias allows BiNE to model the implicit and explicit relationships between the nodes of the network to preserve said structural properties such as the distribution of node degrees. The major limitation of BiNE is that it may not be effective for nodes with limited or no links, as it solely relies on the information obtained from observed links [9]. To overcome the drawback of random walk methods in capturing global structure, Cao et al. [4] proposed BiGI (Bipartite Graph Embedding via Mutual Information Maximization). This method uses two-layer multilayer perceptrons for its loss function.

As the methods explained above take inspiration from previous node embedding algorithms such as node2vec, they also inherit their generic nature. This certainly allows for a more versatile algorithm that can cater to a wide array of tasks. At the same time, it hinders the accuracy of the algorithm, as generic approaches are unable to model the fine details that are often crucial.

These embedding methods are not directly developed for link prediction tasks. However, the openly available versions of BiNE and BiGI allow for link prediction by feeding embedding vectors to a logistic regressor while metapath2vec only produces vector embeddings of the nodes. Using the same approach used in BiNE, we utilize embeddings produced by metapath2vec to train a logistic regression model. Each node pair is represented by combining left node and right node embedding vectors. This binary classifier not only labels the node pairs but also generates a prediction score for each of them. This score estimates how likely that node pair is to belong to the positive class, which is to say, how likely it is for a link to exist between them. We treat these scores as link prediction scores and rank them accordingly, as illustrated in Fig. 2.

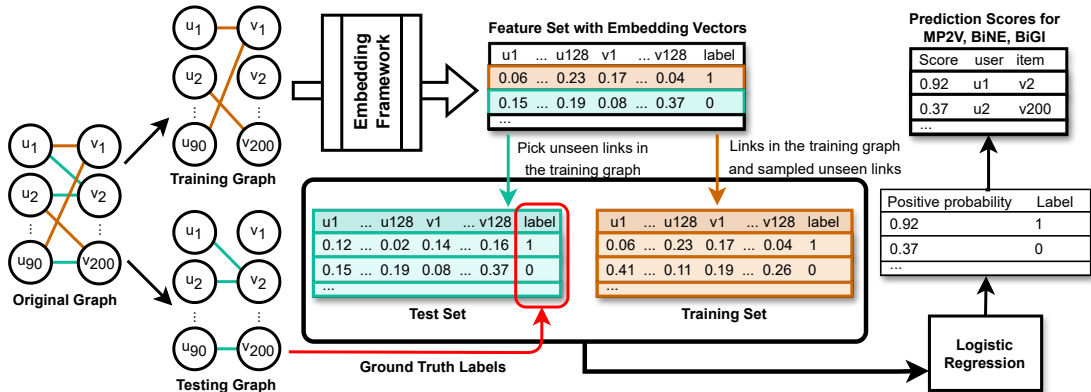


Figure 2: Representation of the link prediction method that relies on graph embeddings (Section 2.3).

2.4 Personalized Recommendation Conversion

Graph neural networks (GNN) and graph convolution networks (GCN) have been frequently used to handle bipartite networks in recent years. These techniques are usually employed for recommendation tasks. Basically, the user-item purchase relation is modeled under the bipartite network structure. Then, the main objective of those methods is to rank the k most likely items for each user.

These approaches also generate separate embedding vectors for the nodes in U and V , corresponding to user and item sets, respectively. Unlike the previously mentioned embedding-based methods however, those embedding vectors are generated specifically for the personalized recommendation task and trained with supervised learning. The embedding vectors are extracted through a GCN (or GNN), where each node’s embedding vector affects those of its neighbors. On top of that, the training and prediction phases of these models are designed with this task in mind, allowing them to optimize the model for a more accurate ranking of items.

The output of personalized recommendation models is a ranked list of items for each user. These items are sorted based on ranking scores, which are derived from the inner product of user and item embedding vectors. In this work, we reformulate the output sets of these approaches to employ them in link prediction tasks. In order to capitalize on the effective item recommendations made from these methods, we consider the ranking score of an item v for a user u as the predictive score assigned to the connection between nodes u and v , as

both user and item representations influence the ranking score of the item. Our proposed approach is shown in Fig. 3.

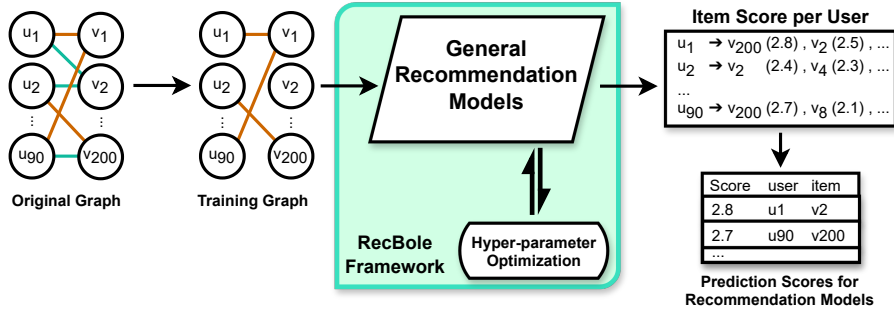


Figure 3: Representation of the link prediction method based on the repurposing of item recommendation techniques (Section 2.4).

In our experiments, we use the following personalized recommendation models:

- Bayesian Personalized Ranking (BPR) uses batches that consist of one user and two items as pairwise training data and optimizes a personalized ranking-focused metric to correctly rank item pairs for each user instead of scoring a single item [25].
- Neural Graph Collaborative Filtering (NGCF) utilizes high-order connectivity in user-item interactions to improve the user and item embeddings [29]. The initial embedding is enhanced through first-order and high-order propagation. This allows for n -hop neighbors to contribute to the embedding of users and items because the propagation layers can capture the collaborative signal when multiple of them are stacked.
- Light Graph Convolution Network (LightGCN) attempts to create node embeddings by iteratively performing graph convolution [12]. Its authors suggest that feature transformation and nonlinear activation of GCNs are not necessary for the collaborative filtering task. Instead, LightGCN continuously performs weighted sum neighbor aggregation.
- Self-supervised Graph Learning (SGL) aims to improve the recommendation performance of the previous GCN models like LightGCN by generating multiple representation of nodes [30]. Different representations are obtained through modifying the network structure via node dropout, link dropout, and random walk.
- Diffusion Recommender Model (DiffRec) utilizes neural networks for de-noising purposes and obtain interaction probabilities [28].

3 Experiments and Results

In our experiments, we use three bipartite networks coming from different domains. *MovieLens* contains users’ ratings of movies on a scale of 1 to 5. We do not consider the ratings and only work with the interaction network. *Setur*, supplied by the Turkish travel agency Setur, contains user-hotel interactions. *LastFM* dataset contains the interactions of users with songs. As weighted links are not supported by some of the methods, the repeating interactions between the same user and song are ignored. We ignore all node and link attributes, besides node types. We apply a minimum degree filter to *Setur* and *LastFM* that removes left nodes (user nodes) with fewer than 8 interactions, for the sake of supervised learning experiments and prediction tasks. This filter is not applied to *MovieLens*, as the lowest degree observed for the left nodes is already 20. The topological properties of the networks used in the study are given in Table 1.

Table 1: Properties of the datasets used in this study. $\langle k \rangle$ denotes the average degree.

| Dataset | #Left nodes | #Right nodes | #Links | $\langle k \rangle$ | Density |
|------------------|-------------|--------------|---------|---------------------|---------|
| <i>Setur</i> | 324 | 1,305 | 3,798 | 4.18 | 0.00898 |
| <i>MovieLens</i> | 943 | 1,682 | 100,000 | 68.88 | 0.06305 |
| <i>LastFM</i> | 512 | 4,014 | 18,492 | 5.91 | 0.00900 |

3.1 Experimental Setup, Performance Evaluation and Hyper-parameter Tuning

We split the network into training and test networks with 90% and 10% of the existing links, respectively. Instead of uniformly randomly selecting 10% of all the links, we draw 10% of the links connected to each left node. Thus, in both the training set and the test set, none of the left nodes are isolates. As a result, during prediction, we are able to rule out the well-defined cold-start problem in recommender systems, which could create additional difficulties for any of our methods.

We assess prediction quality using two key metrics: *Area Under the Precision-Recall Curve* (AUPR) and *Area Under the Receiver Operating Characteristic Curve* (AUROC). The AUPR metric gauges the area under the curve plotted with precision on the y -axis and recall on the x -axis. Precision denotes the ratio of true positives to the items predicted positively, while recall is the proportion of true positives among all positive items. This metric essentially evaluates how effectively the predictor maintains high precision while lowering the positive class threshold to include more positively predicted items. Similarly, AUROC measures the area under the curve, where the y -axis represents recall and the x -axis signifies the false positive rate. The false positive rate is determined by dividing the numbers of false positives by all negative items. Both AUPR and AUROC range from 0 to 1. They do not depend on thresholds, offering general assessments that can cover different prediction scenarios. These two metrics are commonly used for link prediction performance evaluation. In [31], it is suggested to favor AUPR, since AUROC can be deceptive in case of class imbalance. It should also be noted that, in practice, AUPR takes much lower values than AUROC when assessing link prediction [31, 16].

For the implementation of personalized recommendation models, we use the RecBolelibrary, with the suggested configuration settings. They are tuned via random search with early stopping, using NDCG@10 as the validation metric and a training-to-validation data ratio of 8:1. We have obtained the best performances of the learning with these parameters. The NetworkX Python package is used to calculate the network topological features described in Section 2.2. Link prediction indices do not require hyperparameter tuning. The same goes for PCA and LDA methods. But a random search with 10-fold cross-validation is applied to the XGB method to obtain the best performance. During the validation step, the F1 score of the positive class is used to evaluate the success of the parameters, as the prediction of negative links did not matter. Embedding-based methods are executed with the hyperparameters suggested in their respective papers to generate vector embeddings with a length of 128.

3.2 Results

The AUPR and AUROC results of our experiments on three different datasets are presented in Table 2. It is split in four-part, according to the different categories of methods introduced in Section 2. In addition to predictive performance, the table shows the execution time of every method on all three data sets. These do not include the time spent tuning hyper-parameters.

For *MovieLens*, SPM (category *Traditional Link Scores*) gives the best results, while for *Setur* and *LastFM*, NGCF and DiffRec (*Personalized Recommendation Conversion*) are the best, respectively. We cannot generalize this success to their whole respective categories, though, because methods from the same category exhibit a high variability in terms of performance. Still, categories *Link Prediction via Topological Features* and *Embedding-based Link Prediction* seem to be one step behind the others, overall. Independently of the method category, SPM and DiffRec seem to be the best performing methods. Especially, DiffRec obtains significantly higher AUPR scores for *LastFM*. It achieves an AUPR of 0.4061, while the highest score achieved by methods from other categories is 0.1315, representing a 208.8% increase. Dist, BiGI, and MP2V seem to get the lowest scores in general.

Evaluating each method individually, in the *Personalized Recommendation Conversion* category, LightGCN performs best for *Setur*, but receives relatively low AUPR scores for both other datasets. On top of obtaining the best AUROC scores for *LastFM* and *Setur*, NGCF also manages to consistently be in the top 3 when the methods are ranked according to AUPR scores. BPR, the only recommendation model that does not employ GCNs, is ranked the lowest for *MovieLens* and *Setur*, as expected. However, it manages to obtain the third-highest AUPR and second-highest AUROC scores for *LastFM*. In category *Traditional Link Scores*, SPM outperforms the others, most notably for *LastFM*, where SPM receives a 32.43% higher AUPR score compared to the second best performing method. Following SPM, LP is consistently one of the best-performing methods. In all three datasets, there is a small difference in the AUPR scores of Katz and LP indices (Δ AUPR= 0.2%). L5 and L7 score worse than L3 for every dataset on every metric except once, where L5 achieves a negligibly higher AUPR score than L3 for *LastFM*. This shows us the ineffectiveness of using paths longer than three, since we proposed paths longer than five and seven specifically for this work. Moreover, L3 performed better than LP for *MovieLens*, but worse for both other datasets. We attribute this to degree normalization in L3, as both methods are identical otherwise. Finally, Dist performed the worst.

In category *Link Prediction via Topological Features*, the highest AUPR scores are achieved by PCA. Only for *MovieLens* does XGB get better results. XGB also consistently ranks highest in terms of AUROC. Finally, in

Table 2: Performances obtained with all considered methods, for the three selected datasets, in terms of AUPR and AUROC, and runtime.

| Method | <i>MovieLens</i> | | | <i>Setur</i> | | | <i>LastFM</i> | | |
|----------|------------------|---------------|-------------|---------------|---------------|-------------|---------------|---------------|-------------|
| | AUPR | AUROC | Time (s) | AUPR | AUROC | Time (s) | AUPR | AUROC | Time (s) |
| Katz | 0.0560 | 0.8648 | 16.51 | 0.1394 | 0.6286 | 4.00 | 0.0992 | 0.7847 | 25.27 |
| SPM | 0.0778 | 0.8879 | 1189.00 | 0.1412 | 0.6472 | 200.00 | 0.1315 | 0.8153 | 4024.00 |
| L3 | 0.0621 | 0.8962 | 4.92 | 0.0096 | 0.6950 | 1.35 | 0.0444 | 0.8154 | 10.60 |
| L5 | 0.0532 | 0.8759 | 6.10 | 0.0077 | 0.6515 | 1.40 | 0.0447 | 0.7921 | 13.92 |
| L7 | 0.0490 | 0.8648 | 7.48 | 0.0074 | 0.6379 | 1.61 | 0.0421 | 0.7921 | 12.84 |
| PA | 0.0464 | 0.8552 | 2.20 | 0.0867 | 0.6348 | 1.24 | 0.0480 | 0.7803 | 6.24 |
| Dist | 0.0130 | 0.7181 | 5.00 | 0.0026 | 0.6168 | 2.09 | 0.0073 | 0.7641 | 15.55 |
| LP | 0.0561 | 0.8652 | 5.95 | 0.1398 | 0.6976 | 3.60 | 0.0993 | 0.8132 | 14.50 |
| XGB | 0.0179 | 0.7788 | 183.00 | 0.0039 | 0.7207 | 40.41 | 0.0337 | 0.7983 | 133.00 |
| PCA | 0.0119 | 0.5560 | 1.92 | 0.0553 | 0.6191 | 0.32 | 0.0376 | 0.7576 | 1.47 |
| LDA | 0.0151 | 0.7417 | 1.43 | 0.0061 | 0.6453 | 0.25 | 0.0139 | 0.7416 | 1.35 |
| DiffRec | 0.0613 | 0.8462 | 534.00 | 0.0541 | 0.6368 | 239.00 | 0.4061 | 0.8309 | 791.00 |
| LightGCN | 0.0368 | 0.8266 | 184.00 | 0.1316 | 0.7126 | 13.52 | 0.0730 | 0.8342 | 9.64 |
| BPR | 0.0270 | 0.7944 | 36.17 | 0.0473 | 0.7009 | 15.47 | 0.2152 | 0.8678 | 79.00 |
| SGL | 0.0443 | 0.8143 | 763.00 | 0.0508 | 0.7066 | 92.00 | 0.0652 | 0.8674 | 443.00 |
| NGCF | 0.0324 | 0.8048 | 477.00 | 0.0724 | 0.7250 | 36.31 | 0.2174 | 0.8732 | 413.00 |
| BiNE | 0.0275 | 0.8091 | 3966.00 | 0.0080 | 0.5334 | 650.00 | 0.0250 | 0.7632 | 2002.00 |
| BiGI | 0.0368 | 0.3423 | 4834.00 | 0.0065 | 0.4434 | 204.00 | 0.0082 | 0.5379 | 1066.00 |
| MP2V | 0.0078 | 0.5469 | 4011.00 | 0.0012 | 0.5692 | 643.00 | 0.0220 | 0.6298 | 1912.00 |

category *Embedding-based Link Prediction*, for *MovieLens*, BiGI has a better AUPR score than others, although it has the lowest AUROC score in the entire table. For the other datasets, it cannot replicate this success and is surpassed by BiNE.

Examining the experimental results, it is evident that there is a considerable variation in the scores obtained from the performance metrics depending on the dataset. To make an interpretation of the results, we take into account the properties of the datasets seen in Table 1. Indeed, as the density of the data declines, so does the accuracy of the predictions made by the L3 and Dist link prediction indices. This decrease can be explained by the difference in L3 index scores between low-probability and high-probability links becoming less pronounced. Such a decrease, however, is not observed for the SPM, Katz, LP, and PA indices. Nevertheless, they seem to be affected by the size of the network. *Setur* and *LastFM* have similar densities, but the former is smaller. This could be the reason for their success with *Setur*. Another group of methods that does not see a decline in evaluation metrics in sparse networks is *Personalized Recommendation Conversion*. On the contrary, their performance seems to increase as the data becomes sparser. This may be attributed to their unique architecture. The *Personalized Recommendation* models used in our study, except for BPR, employ GCNs to generate representations of the nodes. While *metapath2vec* and BiNE methods also generate node embeddings, they rely on random walks to do so. BiGI utilizes GNNs, whose effectiveness is apparent in the AUPR score achieved on the *MovieLens* dataset, matching LightGCN’s AUPR score.

Moreover, here, *Personalized Recommendation Conversion* methods are optimized for scoring items to make accurate recommendations to users. In contrast, *Embedding-based Methods* treat the link prediction task as a binary classification problem by fitting a logistic regression model to make predictions. They are similar in this way to methods that use topological network features. It is also possible that the success of these methods are limited by the simple classifier used during the prediction step.

The methods that rely on topological features do not apply graph representational learning at all. It could be argued that this puts them at a disadvantage compared to methods that can exploit the network structure through path-based algorithms or graph representational learning. While topological features may accurately describe a node’s role in the network, they do not offer any insight on how two specific nodes interact with each other. The same line of thought could be applied to *Embedding-based* methods. They indeed generate embedding vectors, but no matter how successful this process is, those embeddings are generalized representations of nodes. By contrast, *Personalized Recommendation Conversion* models generate embeddings for the explicit purpose of assigning a score to a pair of nodes. As a result, although we do not use exact-ranked items but convert them to link prediction, it seems to be effective in terms of AUPR and AUROC accuracy.

Evaluation metrics used to score predictions should also be considered while interpreting our results. The link prediction task primarily focuses on predicting future links, in other words, positive instances. High AUPR

scores indicates that the predictor can maintain a high ratio of true positives to positively predicted items, as the positive class g threshold gets lower to include more positives. This is observed as high precision values while recall increases. In our case, the AUPR metric is sensitive to small changes in the number of true positives. In the link prediction task, where there is an abundance of negative instances, AUROC should be handled with care. Due to the high number of negatives, the false positive rate may not vary significantly even after substantial changes in the number of false positives. For these reasons, we primarily compare AUPR scores while paying attention to AUROC scores.

Regarding the execution time of the methods, generally, it can be seen that the more successful methods, such as SPM and GCN-based, take longer to make predictions. Although it is not the best-performing method on any data set, LP shows considerable success in prediction accuracy while maintaining low runtime. The quickest category is *Link Prediction via Topological Features*, due to optimized algorithms of PCA and LDA. Except for SPM, which is an outlier, *Traditional Link Scores* are the second quickest. Nevertheless, their scalability falls behind that of other methods. Indeed, on average, they take 18 times longer to finish processing *LastFM* compared to *Setur*. Most of the indices predict link scores through the adjacency matrix, which represents every link possible in the network. As a consequence, such methods do not scale well in large but sparse real-world networks. Other methods seem to scale off the number of links present in the network in addition to the number of nodes. This is evidenced by their lengthy execution times for *MovieLens*, which is smaller than for *LastFM* in terms of the number of nodes and all possible links. *Embedding-based Link Prediction* methods take longer to work than other methods, despite their low prediction accuracy. This can be attributed to the fact that they use logistic regression to obtain prediction scores, which may not be able to handle the embedding dimensions.

To summarize our findings, methods of category *Personalized Recommendation Conversion*, that use a GCN architecture, are more effective in link prediction compared to unsupervised embedding methods based on random walks, traditional link scores, and classifiers trained with network topological features. We attribute the success of these models to the fact that they specifically generate node embeddings with a supervised technique using the already purchased/preferred items and go through a training phase to make recommendations, which can be formulated as link prediction problems on a per-user basis. *Traditional Link Scores* suffer in sparse networks due to a lack of preexisting information. But surprisingly, SPM, a simple link score, highly performs in several cases. It outperforms all other methods we employed in our experiments except DiffRec and NGCF. In cases where it is vital to make predictions as fast as possible, the LP index can be used to get results in the shortest possible time without too much compromising predictive performance.

4 Conclusion

In this paper, we compared different link prediction methods that could be applied to bipartite networks. We evaluated 19 different methods that we classified into four different categories, with three different datasets of varying sizes and density. Most of the methods that we used were not previously applied to link prediction in bipartite networks. Hence, our work carries especially experimental novelties for link prediction methodologies in bipartite networks. The key result of our experiments can be listed as it follows.

- Our major contribution as a novel link prediction approach, repurposing recommendation models that generate personalized ranked lists of items for link prediction tasks in bipartite networks, outperforms almost all the other methods.
- Specifically, the DiffRec model achieved an AUPR score of 0.4061 on the *LastFM* network, which has a low density (0.009) and a high number of nodes (4, 526).
- SPM, which is a traditional link score, is one of the best performing methods. Although this method was not developed for bipartite networks, our experiments show that it can be a very good link predictor for bipartite networks as well.
- Methods that use bipartite network embeddings and network topological features for link prediction tasks produce the worst results.

Personalized Recommendation methods generate node embeddings within the GCN architecture using supervised learning from interacted states, unlike unsupervised embedding methods. Our experimental results indicate that for link prediction, supervised learning with GCNs yields superior performance in graph representation learning. This study could be further expanded by including more datasets with varying network sizes, link counts, and densities. An ensemble model could be used for personalized recommendation models, since our results imply that different recommendation models were able to better capture the structural properties of different networks.

Acknowledgements

This work is supported by the Galatasaray University Research Fund (BAP) within the scope of project number FBA-2023-1204, titled *Öneri Sistemlerinin Performans Değerlendirmesi*, and the bilateral project of the Scientific and Technological Research Council of Türkiye (TÜBİTAK), under grant number 122N701, with CampusFrance, within the scope of Hubert Curien Partnerships (PHC) project number 49032VB.

References

- [1] Al Hasan, M., Chaoji, V., Salem, S., Zaki, M., 2006. Link prediction using supervised learning, in: SDM06: workshop on link analysis, counter-terrorism and security, pp. 798–805.
- [2] Aslan, S., Kaya, M., 2018. Topic recommendation for authors as a link prediction problem. *Future Generation Computer Systems* 89, 249–264.
- [3] Barabási, A., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., Vicsek, T., 2002. Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications* 311, 590–614.
- [4] Cao, J., Lin, X., Guo, S., Liu, L., Liu, T., Wang, B., 2021. Bipartite graph embedding via mutual information maximization, in: *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, New York, NY, USA. p. 635–643.
- [5] Chen, T., Guestrin, C., 2016. Xgboost: A scalable tree boosting system, in: *Proceedings of the 22nd ACM SIGKDD*, p. 785–794.
- [6] Costa, L.d.F., Oliveira Jr, O.N., Travieso, G., Rodrigues, F.A., Villas Boas, P.R., Antiqueira, L., Viana, M.P., Correa Rocha, L.E., 2011. Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Advances in Physics* 60, 329–412.
- [7] Dong, Y., Chawla, N.V., Swami, A., 2017. metapath2vec: Scalable representation learning for heterogeneous networks, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA. p. 135–144.
- [8] Gao, C., Zheng, Y., Li, N., Li, Y., Qin, Y., Piao, J., Quan, Y., Chang, J., Jin, D., He, X., Li, Y., 2023. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ACM Trans. Recomm. Syst.* 1.
- [9] Gao, M., Chen, L., He, X., Zhou, A., 2018. Bine: Bipartite network embedding, in: *The 41st International ACM SIGIR*, Association for Computing Machinery, New York, NY, USA. p. 715–724.
- [10] Grover, A., Leskovec, J., 2016. Node2vec: Scalable feature learning for networks, in: *Proceedings of the 22nd ACM SIGKDD*, p. 855–864.
- [11] Hamilton, W.L., Ying, R., Leskovec, J., 2017. Inductive representation learning on large graphs, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA. p. 1025–1035.
- [12] He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M., 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation, in: *Proceedings of the 43rd International ACM SIGIR*, p. 639–648.
- [13] Katz, L., 1953. A new status index derived from sociometric analysis. *Psychometrika* 18, 39–43.
- [14] Kovács, I.A., Luck, K., Spirohn, K., Wang, Y., Pollis, C., Schlabach, S., Bian, W., Kim, D.K., Kishore, N., Hao, T., Calderwood, M.A., Vidal, M., Barabási, A.L., 2019. Network-based prediction of protein interactions. *Nature Communications* 10, 1240.
- [15] Kumar, A., Singh, S.S., Singh, K., Biswas, B., 2020a. Link prediction techniques, applications, and performance: A survey. *Physica A: Statistical Mechanics and its Applications* 553, 124289.
- [16] Kumar, A., Singh, S.S., Singh, K., Biswas, B., 2020b. Link prediction techniques, applications, and performance: A survey. *Physica A: Statistical Mechanics and its Applications* 553, 124289.

- [17] Kunegis, J., De Luca, E.W., Albayrak, S., 2010. The link prediction problem in bipartite networks, in: Hüllermeier, E., Kruse, R., Hoffmann, F. (Eds.), *Computational Intelligence for Knowledge-Based Systems Design*, pp. 380–389.
- [18] Landherr, A., Friedl, B., Heidemann, J., 2010. A critical review of centrality measures in social networks. *Business & Information Systems Engineering* 2, 371–385.
- [19] Lü, L., Jin, C.H., Zhou, T., 2009. Similarity index based on local paths for link prediction of complex networks. *Phys. Rev. E* 80, 046122.
- [20] Lü, L., Pan, L., Zhou, T., Zhang, Y.C., Stanley, H.E., 2015. Toward link predictability of complex networks. *Proceedings of the National Academy of Sciences* 112, 2325–2330.
- [21] Lü, L., Zhou, T., 2011. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 390, 1150–1170.
- [22] Opsahl, T., 2013. Triadic closure in two-mode networks: Redefining the global and local clustering coefficients. *Social Networks* 35, 159–167. Special Issue on Advances in Two-mode Social Networks.
- [23] Page, L., Brin, S., Motwani, R., Winograd, T., et al., 1999. The pagerank citation ranking: Bringing order to the web .
- [24] Perozzi, B., Al-Rfou, R., Skiena, S., 2014. Deepwalk: Online learning of social representations, in: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, USA. p. 701–710.
- [25] Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L., 2009. Bpr: Bayesian personalized ranking from implicit feedback, in: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, AUAI Press, Arlington, Virginia, USA. p. 452–461.
- [26] Saoub, K.R., 2021. *Graph Theory: an introduction to proofs, algorithms, and applications*. Chapman and Hall/CRC.
- [27] Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q., 2015. Line: Large-scale information network embedding, in: *Proceedings of the 24th International Conference on World Wide Web*, Republic and Canton of Geneva, CHE. p. 1067–1077.
- [28] Wang, W., Xu, Y., Feng, F., Lin, X., He, X., Chua, T.S., 2023. Diffusion recommender model, in: *Proceedings of the 46th International ACM SIGIR*, Association for Computing Machinery, New York, NY, USA. p. 832–841.
- [29] Wang, X., He, X., Wang, M., Feng, F., Chua, T.S., 2019. Neural graph collaborative filtering, in: *Proceedings of the 42nd International ACM SIGIR*, p. 165–174.
- [30] Wu, J., Wang, X., Feng, F., He, X., Chen, L., Lian, J., Xie, X., 2021. Self-supervised graph learning for recommendation, in: *Proceedings of the 44th International ACM SIGIR*, Association for Computing Machinery, New York, NY, USA. p. 726–735.
- [31] Yang, Y., Lichtenwalter, R., Chawla, N., 2015. Evaluating link prediction methods. *Knowledge and Information Systems* 45.