

Empirical Comparison of Graph Embeddings for Trust-Based Collaborative Filtering

Tomislav Duricic^{1,2}, Hussain Hussain², Emanuel Lacic², Dominik Kowald²,
Denis Helic¹, and Elisabeth Lex¹

¹ Graz University of Technology
{tduricic,elisabeth.lex,dhelic}@tugraz.at
² Know-Center GmbH
{hhussain,elacic,dkowald}@know-center.at

Abstract. In this work, we study the utility of graph embeddings to generate latent user representations for trust-based collaborative filtering. In a cold-start setting, on three publicly available datasets, we evaluate approaches from four method families: (i) factorization-based, (ii) random walk-based, (iii) deep learning-based, and (iv) the Large-scale Information Network Embedding (LINE) approach. We find that across the four families, random-walk-based approaches consistently achieve the best accuracy. Besides, they result in highly novel and diverse recommendations. Furthermore, our results show that the use of graph embeddings in trust-based collaborative filtering significantly improves user coverage.

Keywords: Recommender Systems · Empirical Study · Graph Embeddings · Cold-Start · Trust Networks

1 Introduction

Recommender systems suffer from the well-known cold-start problem [1] that arises when users have rated no or only few items. The cold-start problem is particularly problematic in neighborhood-based recommendation approaches such as collaborative filtering (CF) [2] since the ratings of these users cannot be exploited to find similar users. Trust-based recommender systems (e.g., [3–6]) have been proposed as a potential remedy for the cold-start problem. They alleviate this problem by generating a trust network, i.e., a type of a social network in which nodes usually represent users and edges represent trust connections between users based on their explicitly expressed or implicitly derived trust relationships. Although trust is a complex and ambiguous concept from social sciences, in the context of recommender systems, we use a simple interpretation in which users trust other users in the system if they trust their opinions and ratings on different items [5]. Resulting trust network can be used to find the k -most similar users, whose items are recommended to a target user. Trust networks are, however, typically sparse [7] since only a fraction of users have trust connections, which makes finding the k -similar users challenging. In the present work, we explore the utility of *graph embeddings* to extract the k -similar users

from trust networks. To that end, we conduct experiments on three publicly available benchmark datasets often used in studies on trust-based recommender systems: Epinions [5], Ciao [8], and Filmtrust [9]. We empirically evaluate a range of state-of-the-art graph embedding approaches [10] from four distinct method families, i.e., (i) factorization-based methods, (ii) random-walk-based approaches, (iii) methods based on deep learning, and (iv) the LINE approach [11] that falls in neither of these families, with respect to their ability to deliver accurate, novel, and diverse recommendations [12] for cold-start users.

In our experimental setup, we split each dataset into a validation set (warm-start users) and a test set (cold-start users). For each graph embedding approach, we perform a hyperparameter optimization on each validation set. We then select the hyperparameters which result in highest recommendation accuracy. We generate recommendations for each target user in a CF manner by finding k -similar neighbors using the learned embeddings and ranking their items by similarity scores. Finally, we evaluate the resulting graph embeddings against a corresponding test set with respect to accuracy and beyond accuracy metrics. We compare the graph embedding approaches against five baselines from trust-based recommender systems, commonly used in cold-start settings: (i) Most Popular (MP) recommends the most frequently rated items, (ii) $Trust_{dir}$ extracts trusted users directly from a trust network, (iii) $Trust_{undir}$ ignores edge directions and extracts neighbors from the resulting undirected network, (iv) $Trust_{jac}$ applies the Jaccard coefficient on the explicit trust network, and (v) $Trust_{Katz}$ [13] computes the Katz similarity to infer transitive trust relationships between users. To quantify the algorithmic performance, we evaluate recommendation quality in terms of nDCG, novelty, diversity and user coverage.

We find that as a result of their ability to create a representation of each user in a network, graph embeddings are able to improve user coverage when compared to the baseline approaches. Our experiments also show that random-walk-based approaches, i.e., Node2vec and DeepWalk, consistently outperform other graph embedding methods on all three datasets in terms of recommendation accuracy. Finally, we find a positive correlation between novelty and accuracy in all three datasets suggesting that users in the respective platforms tend to prefer novel content. Summing up, our contributions are three-fold. Firstly, we provide a large-scale empirical study on the efficacy of graph embedding approaches in trust-based recommender systems. Secondly, unlike many previous studies, which evaluated only recommendation accuracy, we compare different approaches with respect to beyond accuracy metrics such as novelty, diversity and user coverage. Lastly, our results provide new insights into user preferences based on correlations between different recommendation quality metrics.

2 Graph Embeddings

In this study, we compare the recommendation performance of graph embedding approaches from four distinct method families [10], i.e., factorization-based methods, random-walk-based approaches, deep-learning-based approaches, and the LINE approach [11] which falls in neither of the first three families.

Factorization-based Approaches. Factorization-based approaches produce node embeddings using matrix factorization. The inner product between the resulting node embedding vectors approximates a deterministic graph proximity measure [14]. In total, we investigate five different factorization approaches:

- *Graph Factorization (GF)* [15] factorizes the adjacency matrix and determines proximity between nodes directly on the adjacency matrix.¹
- *Laplacian Eigenmaps (LE)* [16] factorizes the normalized Laplacian matrix and preserves the 1st-order proximity.¹
- *Locally Linear Embedding (LLE)* [17] minimizes the squared difference between the embedding of a node and a linear combination of its neighbors’ embeddings, weighted by the edges connecting to them. The solution of this minimization problem reduces to a factorization problem.¹
- *High-Order Proximity preserved Embedding (HOPE)* is able to preserve higher-order proximities and capture the asymmetric transitivity. [18].¹
- *Graph Representations with Global Structural Information (GraRep)* [19] can handle higher-order similarity as it considers powers of the adjacency matrix.²

Random Walk-based Approaches. RW-based approaches first identify the context of a node with a random walk and then learn the embeddings typically using a skip-gram model [10]. In total, we evaluated three different approaches:

- *DeepWalk* [20] extracts node sequences with truncated random walks and applies a skip-gram model [21] with hierarchical softmax on the node pairs.³
- *Node2vec* [22] extends DeepWalk with hyperparameters to configure the depth and breadth of the random walks. In contrast to DeepWalk, Node2vec enables to define flexible random walks, while DeepWalk only allows unbiased random walks over the graph [14].⁴
- *Role2vec* [23] uses attributed random walks to learn embeddings. As Role2vec enables to define functions that map feature vectors to types, it can learn embeddings of *types of nodes*.⁵

Deep Learning-based Approaches. Such approaches use deep neural network models to generate node embeddings. In this research paper, we studied three deep learning-based models in total:

- *Deep Neural Networks for Graph Representations (DNGR)* [24] uses random surfing to build a normalized node co-occurrence matrix and employs a stacked denoising autoencoder to learn node embeddings.⁶
- *Structural Deep Network Embedding (SDNE)* [25] finds neighbors by means of 1st and 2nd order proximity and learns node embeddings via autoencoders.⁷

¹ Implementation used: <https://github.com/palash1992/GEM-Benchmark>

² Implementation used: <https://github.com/benedekrozemberczki/role2vec>

³ Implementation used: <https://github.com/phanein/deepwalk>

⁴ Implementation used <https://github.com/aditya-grover/node2vec>

⁵ Implementation used: <https://github.com/benedekrozemberczki/role2vec>

⁶ Implementation used: <https://github.com/ShelsonCao/DNGR>

⁷ Implementation used: <https://github.com/suanrong/SDNE>

- *Graph sample and aggregate GraphSAGE* [26] is a multi-layered graph convolutional neural network, which represents nodes internally by aggregating their sampled neighborhoods and utilizes a random-walk-based cost function for unsupervised learning. GraphSAGE performs the convolution in the graph space. It uses either mean-based, GCN-based, LSTM-based, mean pooling or max pooling models for aggregation.⁸

Large-Scale Information Network Embedding. *LINE* [11] creates embeddings that preserve 1st-order and 2nd-order proximity which are represented as joint and conditional probability distributions respectively.⁹

3 Preliminaries

3.1 Datasets

We employ three open datasets commonly used when evaluating trust-based recommender systems, i.e., Epinions [5], Ciao [8], and FilmTrust [9]. For all three datasets, we create an unweighted trust network, in which each node represents a user, and each directed edge denotes a trust relationship between two users. The trust network is then an adjacency matrix \mathbf{A} where $\mathbf{A}_{u,v}$ is 1 in case of a trust link between u and v , and 0 otherwise. As a result of preliminary experiments, we found that most of the approaches achieved better accuracy results with an undirected network. One possible explanation is that removing the edge direction increases the average number of edges for each node and reduces the sparsity of the adjacency matrix. Moreover, some approaches are not able to consider link direction by design. Therefore, we convert the trust network to an undirected network in all of our experiments by removing edge direction, thus making \mathbf{A} symmetric. Furthermore, we create a ratings matrix \mathbf{R} , where each non-zero entry $\mathbf{R}_{u,i}$ represents a rating given by a user u to an item i . Table 1 shows basic statistics for all three datasets.

Table 1. Dataset statistics.

Dataset	#Users	#Items	#Edges	#Ratings	Graph density
Epinions	49,288	139,738	487,183	664,824	2×10^{-4}
Ciao	19,533	16,121	40,133	72,665	1.85×10^{-3}
Filmtrust	1,642	2,071	1,853	35,497	2.43×10^{-3}

Dataset Splits. We split each dataset into two sets: warm-start users, i.e., users with > 10 ratings and cold-start users, i.e., users with ≤ 10 item ratings. While we use the subset of warm-start users as a validation set for hyperparameter

⁸ Implementation used: <https://github.com/williamleif/GraphSAGE>

⁹ Implementation used: <https://github.com/tangjianku/LINE>

Table 2. Number of users per dataset split.

Dataset	Users with ratings		Users with ratings & trust	
	Warm-start	Cold-start	Warm-start (Validation set)	Cold-start (Test set)
Epinions	14,769	25,393	14,769	25,393
Ciao	1,020	16,591	571	2,124
Filmtrust	963	545	499	241

optimization concerning recommendation accuracy, the subset of cold-start users is used as a test set for measuring algorithm performance. Table 2 reports the number of cold-start and warm-start users in our datasets.

3.2 Experimental Setup

The initial directed trust network is converted to an undirected network by removing edge directions. The resulting undirected symmetric \mathbf{A} is then used as an input for the graph embedding methods, which, as a result, create a d -dimensional embedding for each node (i.e., user) in the graph.

Recommendation Strategy. After generating the embedding for each node in the graph, a similarity matrix \mathbf{S} is created based on the pairwise cosine similarity between nodes’ embeddings. Recommendations are generated in a kNN manner where we find the k -nearest neighbors N_k (i.e., k most similar users) for the target user u_t using the similarity matrix \mathbf{S} . We use $k = 40$ across all of our experiments as in [13]. Then, we assign a score for all items the users in N_k have interacted with:

$$score(i, u_t) = \sum_{v \in N_k} S_{u_t, v} \cdot R_v(i), \quad (1)$$

where $R_v(i)$ corresponds to the rating assigned by the user v to the item i and $S_{u_t, v}$ corresponds to the similarity score in \mathbf{S} between target user u_t and the neighbor user v from N_k . For each target user u_t with n rated items, we recommend 10 items ranked according to Eq 1 and compare them with the actual rated items.

Evaluation Metrics. Previous research has shown [27] that accuracy may not always be the only or the best criteria for measuring recommendation quality. Typically, there is a trade-off between accuracy, novelty, and diversity since users also like to explore novel and diverse content depending on the context. Therefore, in our work, we examine both novelty and diversity as well as accuracy. In particular, in our experimental setup, we use the following four accuracy and beyond-accuracy metrics for evaluation.

Normalized Discounted Cumulative Gain ($nDCG@n$) – a ranking-dependent metric measuring recommendation accuracy based on the Discounted Cumulative Gain (DCG) measure [28].

Novelty@n – corresponds to a recommender’s ability to recommend long-tail items that the target user has probably not yet seen. We compute novelty using the Expected Popularity Complement (EPC) metric [29].

Diversity@n – describes how dissimilar items are in the recommendation list. We calculate it as the average dissimilarity of all pairs of items in the recommendation list [30]. More specifically, we use cosine similarity to measure the dissimilarity of two items based on doc2vec embeddings [31] learned using the item vector from \mathbf{R} where each rating is replaced with the user id.

User Coverage – defined as the number of users for whom at least one item recommendation could have been generated divided by the total number of users in the target set [5].

Baseline Approaches. We evaluate the graph embeddings approaches against five different baselines:

- *Explicit directed trust* ($Trust_{dir}$) is a naive trust-based approach that uses the unweighted, directed trust network’s adjacency matrix for finding user’s nearest neighbors, i.e. $\mathbf{S} = \mathbf{A}$.
- *Explicit undirected trust* ($Trust_{undir}$) is similar to $Trust_{dir}$ but converts the network to an undirected one by ignoring the edge direction, thus making \mathbf{A} symmetric, i.e. $\mathbf{S} = \mathbf{A}_{undir}$.
- *Explicit trust with Jaccard* ($Trust_{jac}$) uses the Jaccard index on the undirected trust network \mathbf{A}_{undir} . \mathbf{S} is a result of calculating the pairwise Jaccard index on \mathbf{A}_{undir} . The intuition behind this algorithm is that two users are treated as similar if they have adjacent users in common, i.e., trustors and trustees.
- *Explicit trust with Katz similarity* ($Trust_{Katz}$) [13] exploits regular equivalence, a concept from network science by using Katz similarity in order to model transitive trust relationships between users.
- *Most Popular (MP)* is a non-personalized approach in recommender systems, which recommends the most frequently rated items.

4 Results

Table 3 shows our results in terms of nDCG, novelty, diversity and user coverage for $n = 10$ recommendations on cold-start users (test set). The reported results depict those hyperparameter configurations¹⁰, which achieve the highest recommendation accuracy on warm users (validation set).

4.1 Accuracy Results

To ease the interpretation of the evaluation results across all three datasets, we rank the results by nDCG and compute the average of these ranks. Correspondingly, in the Rank_{nDCG} column, we show the resulting average rank for the three datasets for recommendation accuracy.

¹⁰ Details on the hyperparameter optimization can be found at: <https://github.com/tduricic/trust-recommender/blob/master/docs/hyperparameter-optimization.md>

Table 3. Evaluation results on cold-start users for different trust-based CF approaches for $n = 10$ recommendations concerning nDCG, novelty, diversity, and user coverage comparing approaches from five different algorithm families across three different datasets. Values marked with * denote that the corresponding approach was significantly better than every other approach with respect to the appropriate metric according to a Wilcoxon signed-rank test (Bonferroni corrected, $p < 0.01$). Rank $_{nDCG}$ is calculated by summing nDCG-based ranks across the datasets and re-ranking the sums.

Cat.	Approach	Rank $_{nDCG}$	Epinions				Ciao				Filmtrust			
			nDCG	Nov.	Div.	UC	nDCG	Nov.	Div.	UC	nDCG	Nov.	Div.	UC
Baseline	Trust $_{dir}$	15	.0245	.0060	.6006	59.2%	.0140	.0028	.3700	3.9%	.2655	.0313	.2784	30.3%
	Trust $_{undir}$	15	.0260	.0063	.5960	97.0%	.0127	.0045	.3632	11.4%	.2739	.0284	.2731	42.0%
	Trust $_{jac}$	11	.0373	.0056	.6548	99.9%	.0176	.0027	.3996	12.8%	.3387	.0369	.2266	36.1%
	Trust $_{Katz}$	12	.0290	.0046	.6979		.0158	.0026	.3842	13.0%	.3681	.0322	.2185	42.9%
	MP	17	.0134	.0015	.7621*		.0135	.0012	.5666	100%	.3551	.0137	.1672	100%
Factorization	LLE	7	.0309	.0044	.6977	100 %	.0239	.0036	.4013	13.1 %	.3649	.0159	.1926	44.2 %
	LE	3	.0318	.0045	.6961		.0231	.0034	.3962		.3715	.0161	.1853	
	GF	14	.0138	.0023	.7024		.0154	.0022	.3970		.3686	.0154	.1945	
	HOPE	3	.0331	.0047	.6728		.0220	.0033	.3956		.3718	.0158	.1827	
	GraRep	7	.0298	.0042	.6704		.0209	.0030	.3974		.3694	.0147	.1859	
RW	Node2vec	1	.0413	.0064	.6581		.0228	.0036	.4042		.3904	.0151	.2235	
	DeepWalk	2	.0435*	.0067*	.6707		.0247*	.0037	.3992		.3654	.0152	.1950	
	Role2vec	6	.0363	.0054	.6910		.0149	.0024	.3933		.3695	.0151	.1919	
DL	DNGR	10	.0353	.0051	.6869		.0197	.0031	.4023		.3583	.0142	.1959	
	SDNE	12	.0184	.0022	.7412		.0175	.0028	.3921		.3687	.0152	.2003	
	GS	7	.0325	.0047	.6810		.0216	.0031	.3963		.3678	.0151	.1883	
LINE	LINE	5	.0407	.0063	.6566		.0222	.0033	.3992		.3667	.0150	.1947	

We can observe that RW-based approaches, especially Node2vec and DeepWalk, are the best performing approaches on all three datasets. In most cases, approaches based on graph embeddings outperform the baselines, except for Trust $_{jac}$ on Epinions. Contrary to a study conducted in [13], Trust $_{jac}$ achieves higher accuracy in comparison with Trust $_{Katz}$. The reason is that in the present work, we convert the trust network to an undirected network, i.e., do not consider the direction of the trust edge. HOPE and Laplacian Eigenmaps perform best among the factorization-based approaches; LINE shows a good performance on all three datasets concerning all three metrics, and GraphSAGE is the best deep learning approach. SDNE does not perform well in our experiments, which we attribute to not exploring a sufficiently broad range of hyperparameters.

4.2 Beyond-Accuracy Results

Novelty, Diversity, and User Coverage. Being superior in the case of Rank $_{nDCG}$, Node2vec also achieves high novelty and diversity scores. Plus, it performs similarly or better than other RW-based methods across all three datasets. Factorization-based approaches show average performance concerning both novelty and diversity, except for GF, which scores very low on novelty

and above average on diversity. DL approaches show average to below-average performance on novelty and average performance on diversity. Trust-based baselines achieve high novelty scores in general and, not surprisingly, MP has a high diversity score and the worst novelty score out of all approaches. Since all graph embedding approaches create a latent representation of each user in a trust network using it to generate a set of item recommendations, there are no differences among them in user coverage. Except for MP, all baselines result in lower user coverage than the graph embedding approaches. Since MP provides the same list of recommendations to all users, it always has a maximum user coverage.

Evaluation Metrics and User Preferences. Table 3 reports only mean values for each of the approaches. However, we store individual nDCG, novelty, and diversity values for each target user and each approach and dataset. By computing the Kendall rank correlation coefficient (Bonferroni corrected, $p < 0.01$) on non-zero metrics values for all approaches, we can get an insight into user preferences for each dataset. In this manner, we observe a statistically significant positive mean correlation across all three datasets between nDCG and novelty, ranging from 0.43 on Epinions to 0.36 on Filmtrust. This suggests that users on all three platforms prefer recommendations with higher novelty, especially on Epinions. We also observe a statistically significant mean negative correlation between diversity and novelty on Epinions (-0.15), which suggests that on this platform, more novel content seems to be less diverse.

5 Conclusions and Future Work

In this work, we explored the utility of graph embedding approaches from four method families to generate latent user representations for trust-based recommender systems in a cold-start setting. We found that random-walk-based approaches, (i.e., Node2vec and DeepWalk), consistently achieve the best accuracy. We additionally compared the methods concerning novelty, diversity, and user coverage. Our results showed that again, Node2vec and DeepWalk scored high on novelty and diversity. Thus, they can provide a balanced trade-off between the three evaluation metrics. Moreover, our experiments showed that we can increase the user coverage of recommendations when we utilize graph embeddings in a k -nearest neighbor manner. Finally, a correlation analysis between the nDCG, novelty, diversity scores revealed that in all three datasets, users tend to prefer novel recommendations. Hence, on these datasets, recommender systems should offer a good tradeoff between accuracy and novelty. This could also explain the superior performance of the random-walk based approaches and we plan to investigate this in more detail in follow-up work.

Limitations and Future Work. One limitation of this study is that we treated the trust networks as undirected while, in reality, they are directed. This may have resulted in loss of information, and as such, we aim to further explore how to preserve different properties of trust networks (e.g., asymmetry). Moreover, it is possible that we did not examine an ample enough space of hyperparam-

eters, which might have resulted in a more reduced performance of some of the approaches, e.g., SDNE. Furthermore, we aim to explore node properties of k -nearest neighbors for all methods to find and interpret the critical node properties preserved by the graph embeddings, which impact the recommendation accuracy, thus providing a better understanding of the complex notion of trust. Finally, we aim to incorporate user features obtained from the rating matrix into graph embeddings learned on the trust network.

References

1. A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock, "Methods and metrics for cold-start recommendations," in *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 253–260, ACM, 2002.
2. J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, "Collaborative filtering recommender systems," in *The adaptive web*, pp. 291–324, Springer, 2007.
3. S. Fazeli, B. Loni, A. Bellogin, H. Drachslar, and P. Sloep, "Implicit vs. explicit trust in social matrix factorization," in *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, pp. 317–320, ACM, 2014.
4. N. Lathia, S. Hailes, and L. Capra, "Trust-based collaborative filtering," in *IFIP international conference on trust management*, pp. 119–134, Springer, 2008.
5. P. Massa and P. Avesani, "Trust-aware collaborative filtering for recommender systems," in *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*, pp. 492–508, Springer, 2004.
6. J. O'Donovan and B. Smyth, "Trust in recommender systems," in *Proceedings of the 10th International Conference on Intelligent User Interfaces*, IUI '05, pp. 167–174, ACM, 2005.
7. Y. A. Kim, "An enhanced trust propagation approach with expertise and homophily-based trust networks," *Knowledge-Based Systems*, vol. 82, pp. 20 – 28, 2015.
8. J. Tang, H. Gao, and H. Liu, "mTrust: Discerning multi-faceted trust in a connected world," in *Proceedings of the fifth ACM international conference on Web search and data mining*, pp. 93–102, ACM, 2012.
9. G. Guo, J. Zhang, and N. Yorke-Smith, "A novel bayesian similarity measure for recommender systems," 2013.
10. P. Goyal and E. Ferrara, "Graph embedding techniques, applications, and performance: A survey," *Knowledge-Based Systems*, vol. 151, pp. 78 – 94, 2018.
11. J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "Line: Large-scale information network embedding," in *Proceedings of the 24th international conference on world wide web*, pp. 1067–1077, International World Wide Web Conferences Steering Committee, 2015.
12. Y. C. Zhang, D. Ó Séaghdha, D. Quercia, and T. Jambor, "Auralist: Introducing serendipity into music recommendation," in *Proceedings of the 5th ACM Conference on Web Search and Data Mining (WSDM-12)*, 2012.
13. T. Duricic, E. Lacic, D. Kowald, and E. Lex, "Trust-based collaborative filtering: Tackling the cold start problem using regular equivalence," in *Proceedings of the 12th ACM Conference on Recommender Systems*, RecSys '18, pp. 446–450, ACM, 2018.

14. W. L. Hamilton, R. Ying, and J. Leskovec, "Representation learning on graphs: Methods and applications," *IEEE Data Engineering Bulletin*, 2017.
15. A. Ahmed, N. Shervashidze, S. Narayanamurthy, V. Josifovski, and A. J. Smola, "Distributed large-scale natural graph factorization," in *Proceedings of the 22nd international conference on World Wide Web*, pp. 37–48, ACM, 2013.
16. M. Belkin and P. Niyogi, "Laplacian eigenmaps and spectral techniques for embedding and clustering," in *Advances in neural information processing systems*, pp. 585–591, 2002.
17. S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
18. M. Ou, P. Cui, J. Pei, Z. Zhang, and W. Zhu, "Asymmetric transitivity preserving graph embedding," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1105–1114, ACM, 2016.
19. S. Cao, W. Lu, and Q. Xu, "Grarep: Learning graph representations with global structural information," in *Proceedings of the 24th ACM international on conference on information and knowledge management*, pp. 891–900, ACM, 2015.
20. B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pp. 701–710, ACM, 2014.
21. T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in Neural Information Processing Systems 26*, pp. 3111–3119, Curran Associates, Inc., 2013.
22. A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
23. N. K. Ahmed, R. Rossi, J. Boaz Lee, T. L. Willke, R. Zhou, X. Kong, and H. Eldardiry, "Learning Role-based Graph Embeddings," *arXiv e-prints*, p. arXiv:1802.02896, Feb. 2018.
24. S. Cao, W. Lu, and Q. Xu, "Deep neural networks for learning graph representations," 2016.
25. D. Wang, P. Cui, and W. Zhu, "Structural deep network embedding," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1225–1234, ACM, 2016.
26. W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
27. M. Kaminskis and D. Bridge, "Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems," *ACM Trans. Interact. Intell. Syst.*, vol. 7, pp. 2:1–2:42, Dec. 2016.
28. K. Järvelin, S. L. Price, L. M. Delcambre, and M. L. Nielsen, "Discounted cumulated gain based evaluation of multiple-query ir sessions," in *Proceedings of ECIR'2008*, pp. 4–15, Springer, Springer, 2008.
29. S. Vargas and P. Castells, "Rank and relevance in novelty and diversity metrics for recommender systems," in *Proceedings of the fifth ACM conference on Recommender systems*, pp. 109–116, 2011.
30. B. Smyth and P. McClave, "Similarity vs. diversity," in *International conference on case-based reasoning*, pp. 347–361, Springer, 2001.
31. T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, pp. 3111–3119, 2013.