Enhancing Graph Neural Networks for Recommender Systems

Siwei Liu
s.liu.4@research.gla.ac.uk
University of Glasgow
Glasgow, Scotland

ABSTRACT

Recommender systems lie at the heart of many online services such as E-commerce, social media platforms and advertising. To keep users engaged and satisfied with the displayed items, recommender systems usually use the users' historical interactions containing their interests and purchase habits to make personalised recommendations. Recently, Graph Neural Networks (GNNs) have emerged as a technique that can effectively learn representations from structured graph data. By treating the traditional user-item interaction matrix as a bipartite graph, many existing graph-based recommender systems (GBRS) [8, 9] have been shown to achieve state-of-the-art performance when employing GNNs. However, the existing GBRS approaches still have several limitations, which prevent the GNNs from achieving their full potential. In this work, we propose to enhance the performance of the GBRS approaches along several research directions, namely leveraging additional items and users' side information, extending the existing undirected graphs to account for social influence among users, and enhancing their underlying optimisation criterion. In the following, we describe these proposed research directions.

The effectiveness of the recommender systems is commonly limited by the sparsity of user-item interactions, which leads to the well-known *cold-start* problem. A simple yet effective method to alleviate the *cold-start* problem is to incorporate side information to enrich the representations of items and users. We have previously proposed to leverage the social network and textual reviews information to augment the representations of users and items through the creation of a heterogeneous graph neural model [4]. However, there have been very few attempts to introduce other types of side information in GBRS such as the geographical information [1] and timestamps [6], which have been leveraged in traditional recommendation models. We propose to investigate the integration of multiple types of side information, of both users and items, into GBRS.

Among the different types of side information, the social relationships between users are commonly used and have been shown to be effective. In existing works [3, 5], the social influence is assumed to be bidirectional, which means that if a user u is socially connected with a friend f, then this social relation is equally associated to both users. However, we argue that a social relationship is not necessarily bidirectional. Indeed, in social networks and many online platforms such as Twitter and Weibo, there is a distinction between the 'follower' and 'followee' or the 'truster' and 'trustee'. Hence, the current simple undirected GNNs are insufficient to handle graph data with directed links between nodes. We propose to introduce this directed graph into GBRS to better leverage the social information thereby enhancing the recommendation performance by defining the direction of the social relationships.

Finally, the embedding generation process and optimisation criterion are both critical for a recommender model. However, GNNs only modify the former by augmenting the users and items' embeddings with the embeddings of their local neighbours. We also focus on the optimisation criterion because most of the GBRS approaches use Bayesian Personalised Ranking (BPR) [7], which assumes that users always prefer an interacted with item over a random non-interacted with item. However, since the users might not be aware of those not previously interacted with items, it is problematic to infer that they actually prefer the items they interacted with over the others. More precisely, one can infer that a user prefers one item over another only when this user has been exposed to this item as well as the other candidates. To correct this assumption bias, we propose to use the causal inference approach [2] to estimate the items exposure from the interaction history. By replacing the BPR optimisation criterion with a tailored causal inference approach, we aim to alleviate the over-estimation of the previously non-interacted items in GBRS.

To conclude, we have identified 3 limitations in the existing GBRS approaches. Our proposed methods and techniques aim to alleviate those limitations and help the effective deployment of GNNs in recommender systems.

KEYWORDS

Recommender Systems, Graph Neural Networks

ACM Reference Format:
Siwei Liu. 2020. Enhancing Graph Neural Networks for Recommender Systems. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20), July 25–30, 2020, Virtual Event, China.* ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3397271.3401456

REFERENCES

[1] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2013. Exploring temporal effects for location recommendation on location-based social networks. In *Proc. of Recsys.*
[2] Dawen Liang, Laurent Charlin, and David M Blei. 2016. Causal inference for recommendation. In *Proc. of UAI.*
[3] Siwei Liu, Iadh Ounis, and Craig Macdonald. 2019. Social Regularisation in a BPR-based Venue Recommendation System. In *Proc. of FDIA.*
[4] Siwei Liu, Iadh Ounis, Craig Macdonald, and Zaiqiao Meng. 2020. A Heterogeneous Graph Neural Model for Cold-Start Recommendation. In *Proc. of SIGIR.*
[5] Hao Ma, Dengyong Zhou, Chao Liu, Michael R. Lyu, and Irwin King. 2011. Recommender Systems with Social Regularization. In *Proc. of WSDM.*
[6] Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-aware recommender systems. *Comput. Surveys* (2018).
[7] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proc. of UAI.*
[8] Rianne van den Berg, Thomas N Kipf, and Max Welling. 2018. Graph Convolutional Matrix Completion. In *Proc. KDD.*
[9] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *Proc. SIGIR.*

SIGIR '20, July 25–30, 2020, Virtual Event, China

ACM ISBN 978-1-4503-8016-4/20/07.
https://doi.org/10.1145/3397271.3401456