Graph Neural Networks for Recommendation leveraging Multimodal Information

Daniele Malitesta
Université Paris-Saclay, CentraleSupélec, Inria
France
daniele.malitesta@centralesupelec.fr

Abstract

Recommender systems act as filtering algorithms to provide users with items that might meet their interests according to the expressed preferences and items' characteristics. As of today, the collaborative filtering paradigm, along with deep learning techniques to learn high-quality users' and items' representations, constitute the de facto standard for personalized recommendation, showing remarkable recommendation accuracy performance. Nevertheless, recommendation remains a highly-challenging task. Among the most debated open issues in the community, this thesis considers two algorithmic and conceptual ones, namely: (i) the inexplicable nature of users' preferences, especially when they come in the form of implicit feedback; (ii) the effective exploitation of the collaborative information in the designing and training of recommendation models.

In domains such as fashion, food, and media content recommendation, the shallow item's profile can be enhanced through the *multimodal* characteristics describing items [Malitesta et al., 2023]. Driven by these assumptions, in the first part of this thesis, we apply multimodal deep learning strategies for multimedia recommendation; the scope is to study and design recommendation algorithms based upon the principles of multimodality to possibly match each item's characteristic to the implicit preference expressed by the user [Deldjoo et al., 2022], thus addressing the (i) issue.

Recent collaborative filtering approaches profile users and items through embedding vectors in the latent space. However, such models disregard structural properties naturally encoded into the user-item interaction data. Indeed, recommendation datasets are easily describable under the topology of a bipartite and undirected graph, with users and items being the graph nodes connected at multiple distance hops. In this respect, the application of *graph neural networks*, recent deep learning techniques specifically tailored to learn from non-euclidean data, can provide a refined representation of users and items to mine near- and long-distance relationships in the user-item graphs [Anelli et al., 2023b]. Indeed, this is one possible solution to exploit the collaborative information, which is effectively propagated within the user-item graph, addressing the (ii) issue.

Conclusively, this thesis aims to match the two families of recommendation strategies by leveraging graph neural networks and multimodal information data [Anelli et al., 2022]. In doing so, other numerous micro-aspects within the two macro-areas (introduced above) are examined. Indeed, the thesis is a systematic compendium of careful analyses regarding, among others, reproducibility, novel evaluation dimensions [Anelli et al., 2023a], and tasks/scenarios complementary to recommendation.

Awarded by: Politecnico di Bari, Bari, Italy **on** 30 January 2024.
Supervised by: Tommaso Di Noia.
Available at: https://hdl.handle.net/11589/264941.

Selected Publications

Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta, and Claudio Pomo. Reshaping graph recommendation with edge graph collaborative filtering and customer reviews. In Wei Liu and Linsey Pang, editors, *Proceedings of the Workshop on Deep Learning for Search and Recommendation (DL4SR 2022) co-located with the 31st ACM International Conference on Information and Knowledge Management (CIKM 2022)*, volume 3317 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2022. URL https://ceur-ws.org/Vol-3317/Paper7.pdf.

Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, Vincenzo Paparella, and Claudio Pomo. Auditing consumer- and producer-fairness in graph collaborative filtering. In Jaap Kamps, Lorraine Goeuriot, Fabio Crestani, Maria Maistro, Hideo Joho, Brian Davis, Cathal Gurrin, Udo Kruschwitz, and Annalina Caputo, editors, *Advances in Information Retrieval - 45th European Conference on Information Retrieval, ECIR , Part I*, volume 13980 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2023a. doi: 10.1007/978-3-031-28244-7_3. URL https://doi.org/10.1007/978-3-031-28244-7_3.

Vito Walter Anelli, Daniele Malitesta, Claudio Pomo, Alejandro Bellogín, Eugenio Di Sciascio, and Tommaso Di Noia. Challenging the myth of graph collaborative filtering: a reasoned and reproducibility-driven analysis. In Jie Zhang, Li Chen, Shlomo Berkovsky, Min Zhang, Tommaso Di Noia, Justin Basilico, Luiz Pizzato, and Yang Song, editors, *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys*, pages 350–361. ACM, 2023b. doi: 10.1145/3604915.3609489. URL https://doi.org/10.1145/3604915.3609489.

Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. Leveraging content-style item representation for visual recommendation. In Matthias Hagen, Suzan Verberne, Craig Macdonald, Christin Seifert, Krisztian Balog, Kjetil Nørvåg, and Vinay Setty, editors, *Advances in Information Retrieval - 44th European Conference on IR Research, ECIR, Part II*, volume 13186 of *Lecture Notes in Computer Science*, pages 84–92. Springer, 2022. doi: 10.1007/978-3-030-99739-7_10. URL https://doi.org/10.1007/978-3-030-99739-7_10.

Daniele Malitesta, Giuseppe Gassi, Claudio Pomo, and Tommaso Di Noia. Ducho: A unified framework for the extraction of multimodal features in recommendation. In Abdulmotaleb El-Saddik, Tao Mei, Rita Cucchiara, Marco Bertini, Diana Patricia Tobon Vallejo, Pradeep K. Atrey, and M. Shamim Hossain, editors, *Proceedings of the 31st ACM International Conference on Multimedia, MM*, pages 9668–9671. ACM, 2023. doi: 10.1145/3581783.3613458. URL https://doi.org/10.1145/3581783.3613458.