Tutorial on User Profiling with Graph Neural Networks and Related Beyond-Accuracy Perspectives

Erasmo Purificato
Otto von Guericke University Magdeburg
Magdeburg, Germany
Leibniz Institute for Educational Media | Georg Eckert Institute
Brunswick, Germany
erasmo.purificato@ovgu.de

Ludovico Boratto
Department of Mathematics and Computer Science, University of Cagliari
Cagliari, Italy
ludovico.boratto@acm.org

Ernesto William De Luca
Otto von Guericke University Magdeburg
Magdeburg, Germany
Leibniz Institute for Educational Media | Georg Eckert Institute
Brunswick, Germany
ernesto.deluca@ovgu.de

ABSTRACT

The proposed tutorial aims to introduce the UMAP community to modern user profiling approaches leveraging graph neural networks (GNNs). We will begin by discussing the conceptual foundations of user profiling and GNNs and providing a literature review of the two topics. We will then present a systematic overview of the state-of-the-art GNN architectures designed for user profiling, including the types of data that are typically used for this purpose. We will also discuss ethical considerations and beyond-accuracy perspectives (i.e. fairness and explainability), which can arise within the potential applications of adopting GNNs for user profiling. In the practical session of the tutorial, attendees will have the opportunity to understand concretely how recent GNN models for user profiling are built and trained with open-source tools and publicly available datasets. The audience will also be engaged in investigating the impact of the presented models on case studies involving bias detection and mitigation, as well as user profiles explanations. The tutorial will end with an analysis of existing and emerging open challenges in the field and their future research directions.

CCS CONCEPTS

• **Human-centered computing → User models**; • **Social and professional topics → User characteristics**; • **Applied computing → Law, social and behavioral sciences**.

KEYWORDS

User Profiling, Graph Neural Networks, Fairness, Explainability

ACM Reference Format:
Erasmo Purificato, Ludovico Boratto, and Ernesto William De Luca. 2023. Tutorial on User Profiling with Graph Neural Networks and Related Beyond-Accuracy Perspectives. In *UMAP '23: Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (UMAP '23), June 26–29, 2023, Limassol, Cyprus.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3565472.3595616

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
UMAP '23, June 26–29, 2023, Limassol, Cyprus
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9932-6/23/06.
https://doi.org/10.1145/3565472.3595616

1 INTRODUCTION

The principal goal of **user profiling** is to infer an individual's interests, personality traits, or behaviours from generated data to create an efficient user representation, i.e. a *user model*, whose construction is particularly important in the context of adaptive and personalised systems [8], as it is usually considered one of the main practices adopted to develop recommender systems [17]. The importance of this particular field of research has been sharpened in the last years due to the massive amount of data shared daily by millions of online users, making *user profiling* a key topic in many scenarios, such as social networks [16], e-commerce [28] and recommenders [11].

Initially, early user profiling techniques focused solely on examining static characteristics (*explicit user profiling*) obtained from online forms and surveys [18]. However, these methods proved ineffective as users are unwilling to share their personal information directly. As a result, modern systems have turned to profile users implicitly based on their actions and interactions (*implicit user profiling*, also known as *behavioural user profiling*) [14]. One common way to represent these behaviours is through graph structures, where edges represent interactions between users, represented by nodes. **Graph Neural Networks** (GNNs) [12, 15, 26, 34, 35] have proven to be successful in modelling graph data across various fields, such as recommender systems [13, 32], natural language processing [30] and text mining [24], and several contributions have been already emerged for user profiling [4, 5, 21, 29].

The effectiveness of existing user profiling techniques is commonly evaluated by assessing the ability of the models to classify different types of profiles [5]. In recent years, the growing concerns about the *ethical principles* that should guide the development and use of artificial intelligence (AI) systems [9] led, in every research field, to an increased focus on issues such as transparency [27], privacy [20], and social equity [10].

One of the main goal of the proposed tutorial is to illustrate and discuss two specific *beyond-accuracy perspectives* related to GNN-based models for user profiling, namely **fairness** and **explainability**. About *fairness*, a particular issue with using GNNs to classify user profiles is that, like any machine learning system trained on historical data, these neural networks can learn biases present in the data and reflect those biases in their output. This is often due to the structure of the graph and the message-passing process used by GNNs, which can amplify discrimination because nodes belonging to the same group are more likely to be connected to each other than nodes belonging to different groups [22]. The topic of *algorithmic fairness* is becoming more prominent in the field of decision-making systems. There have been many studies published on techniques for detecting and reducing biases in machine learning models [2, 3], but only a few of these focus specifically on fairness in graph neural networks (GNNs) for user profiling models, such as FairGNN [7] and our analysis on two state-of-the-art models for behavioural user profiling [19]. Another rising research topic for automated systems is *explainability*, aiming to expose AI models to humans in an interpretable manner [23] by dealing with specific regulations, like the EU GDPR, which explicitly requires users to be able to understand why and how a particular result from a system is obtained. Except for some sporadic work, such as [1, 6], the explainable user models are still an almost open research field to address, which can be tackled by applying general explainability approaches for "general" GNNs (e.g. [31, 33] to architectures designed for user profiling. The latter is the direction we are exploring in our research and whose results will be presented within the tutorial sessions.

2 OUTLINE

In the first session of our tutorial, we are going to present the existing literature and the most recent advances of user profiling approaches in machine learning, focusing on GNN-based models designed for that purpose. We will also cover two beyond-accuracy perspectives of this field, fairness and explainability, to illustrate how the state-of-the-art approaches and techniques are able to deal with those timely aspects. Moreover, we will prepare some practical exercises for the second session to keep the audience engaged and make the tutorial interactive. We plan to publish all the slides, Jupyter notebooks, reference lists, pre-processed datasets, and any source code used during the tutorial on our webpage. The expected outline, including the description of the practical session, is briefly illustrated below[1].

- **User profiling session** - *Lecture slides and Jupyter Notebooks in Python* - 90 mins
  - **Opening and instructors' presentation** (5 mins)
  - **Introduction to user profiling** (15 mins): a historical overview of the research on user profiling is given to set the basis for understanding the recent advances. In particular, starting from the definition of the key terms in the field (e.g. the difference between *explicit* and *implicit user profiling* [14, 18]), we will shortly illustrate several user profiling contributions in different domains [11, 16, 17, 28].
  - **Introduction to GNNs** (15 mins): as for the user profiling part, in this portion of the tutorial, we will cover the most important notions about *graph neural networks*, such as basic terminology and most popular architectures (e.g. GCN [15] and GAT [26]), with the aim to create a common background to make the audience able to follow the core tutorial sections.
  - **GNN-based models for user profiling** (25 mins): we will present in detail the current state-of-the-art GNN-based model for user profiling, such as CatGCN [4], RHGN [29] and others, describing their architectures, their training procedures, and discussing their strengths and weaknesses.
  - **Hands on state-of-the-art GNN-based models for user profiling** (25 mins): to show in practice how the described GNN models are designed and implemented, we will execute and explain some of the models, such as CatGCN [4] and RHGN [29], in their original configuration, also illustrating the used datasets for every contribution (e.g. Alibaba[2] and JD[3]).
  - **Q&A** (5 mins)
- **Beyond-accuracy perspectives session** - *Lecture slides and Jupyter Notebooks in Python* - 90 mins
  - **Overview on beyond-accuracy perspectives for user profiling with GNNs** (25 mins): the *fairness* and *explainability* aspects encountered in user profiling research will be discussed in the last part of the first session. Several points of view will be considered: (1) analysis of the beyond-accuracy perspectives in existing state-of-the-art models [19]; (2) description of GNN-based models designed for the precise purpose, such as FairGNN [7]; (3) application of general GNN approaches for fairness or explainability to a specific user profiling case study, such as GNNExplainer [31].
  - **Hands on FairGNN** (10 mins): we will illustrate FairGNN [7], a state-of-the-art GNN-based framework for debiasing, and run it on the datasets used in the original publication (i.e. NBA[4] and Pokec [25]).
  - **Use cases on related beyond-accuracy perspectives** (35 mins): we will apply the three approaches described in the last section of the theoretical session by making use of a standardised framework developed for analysing fairness and explainability in GNN-based models for user profiling. In particular, the framework can get a graph structure in *NetworkX*[5] or *Neo4J*[6] format and generate the specific input for each of the GNNs. The participants will be able to compute different types of biases (i.e. binary and multiclass) and execute several debiasing strategies and explainability methods.
  - **Discussion and open challenges** (15 mins)
  - **Closing** (5 mins)

3 TARGET AUDIENCE

The tutorial is intended for a *beginner* or *intermediate* audience and is open to researchers, industry technologists and practitioners. It covers all the necessary background on user profiling and graph neural networks to make it accessible to people without prior knowledge about these topics, which is not assumed. Basic knowledge of Python programming and familiarity with common data science libraries, such as Pandas and NumPy, are preferred. It is

[1] Due to the page limit for this proposal, the literature referenced in the outline is not meant to be complete. It only provides leads to the most relevant papers. The proposed tutorial covers more than 50 references, with most of the challenges and solutions presented in detail.

[2] https://tianchi.aliyun.com/dataset/dataDetail?dataId=56
[3] https://github.com/guyulongcs/IJCAI2019_HGAT
[4] https://www.kaggle.com/datasets/noahgift/social-power-nba
[5] https://networkx.org/
[6] https://neo4j.com/

worth noticing that, despite the technical aspects, the key concepts illustrated during the tutorial and their applications touch a range of interdisciplinary fields, making also the proposed outline of interest to an interdisciplinary audience. Upon completing this tutorial, the participants will gain an understanding of key concepts related to graph neural network models for user profiling. They will be able to analyse the fairness and consider the impact on relevant parties. Additionally, they will be able to identify challenges and opportunities in this area.

4 INSTRUCTORS' BIOGRAPHY

Erasmo Purificato (webpage: https://erasmopurif.com/; email: erasmo.purificato@ovgu.de) has been a Research Assistant in the "Human-Centred AI" group at the Otto von Guericke University Magdeburg (Germany) and in the "Human-Centred Technologies for Educational Media" department at the Leibniz Institute for Educational Media | Georg Eckert Institute (Germany), since February 2020. He is doing his Ph.D. in Computer Science with a project entitled "Beyond-Accuracy Perspectives on Graph Neural Network-Based Models for Behavioural User Profiling," focusing on Human-Centred AI and Responsible AI techniques. From July 2021, he has been appointed by the Guglielmo Marconi University of Rome (Italy) as an Adjunct Professor in "Software Engineering". Between September 2017 and December 2019, he worked at Blue Reply company in Turin as a Machine Learning Engineer and Technical Leader in the Cognitive&Data business unit, developing projects in content-based image retrieval fields. He co-organised the APEx-UI 2022 workshop co-located with IUI 2022 and the IEEE Autumn School ISACT 2022 co-located with ICHMS 2022, where he gave a talk about "Beyond-Accuracy Perspectives: Explainability and Fairness". He has been part of the program committee of several conferences and workshops, such as SIGIR, RecSys, UMAP, HT and ExUM, and served as a reviewer for journals, such as TORS.

Ludovico Boratto (webpage: https://www.ludovicoboratto.com/; email: ludovico.boratto@acm.org) is Assistant Professor at the Department of Mathematics and Computer Science of the University of Cagliari (Italy). His research interests focus on recommender systems and their impact on stakeholders, considering accuracy and beyond-accuracy evaluation metrics. He has authored more than 60 papers and published his research in top-tier conferences and journals. His research activity also brought him to give talks and tutorials at top-tier conferences and research centres (Yahoo! Research). He is the editor of the book "Group Recommender Systems: An Introduction" by Springer. He is an editorial board member of the "Information Processing & Management" journal (Elsevier) and "Journal of Intelligent Information Systems" (Springer) and guest editor of several journals' special issues. He is regularly part of the program committees of the main Web conferences. In 2012, he got his Ph.D. at the University of Cagliari (Italy), where he was a research assistant until May 2016. From May 2016 to April 2021, he joined Eurecat as Senior Research Scientist in the Data Science and Big Data Analytics research group.

Ernesto William De Luca (webpage: https://ernestodeluca.eu/; email: deluca@gei.de) is head of the "Human-Centred Technologies for Educational Media" department at the Leibniz Institute for Educational Media | Georg Eckert Institute (Germany) and from October 2019 has been appointed as a Full Professor in "Research Infrastructures for Digital Humanities" at the Otto von Guericke University Magdeburg (Germany), where he leads the "Human-Centred AI" group. In addition, in May 2015, the Guglielmo Marconi University of Rome (Italy) appointed him an associate professor in "Computational Engineering". He studied computational linguistics and then gained his Ph.D. in computer science. His research includes machine learning, human-machine interaction, Natural Language Processing, user and data modelling, the Semantic Web and Information Retrieval. He has written over 100 papers for national and international conferences and journals and organized and chaired numerous workshops and conferences. He is a regular reviewer and programme committee member of different high-profile journals and conferences.

REFERENCES

[1] Krisztian Balog, Filip Radlinski, and Shushan Arakelyan. 2019. Transparent, scrutable and explainable user models for personalized recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 265–274.

[2] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. http://www.fairmlbook.org.

[3] Simon Caton and Christian Haas. 2020. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053* (2020).

[4] Weijian Chen, Fuli Feng, Qifan Wang, Xiangnan He, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. CatGCN: Graph Convolutional Networks with Categorical Node Features. *IEEE Transactions on Knowledge and Data Engineering* (2021).

[5] Weijian Chen, Yulong Gu, Zhaochun Ren, Xiangnan He, Hongtao Xie, Tong Guo, Dawei Yin, and Yongdong Zhang. 2019. Semi-supervised user profiling with heterogeneous graph attention networks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 2116–2122.

[6] David Jaime Tena Cucala, Bernardo Cuenca Grau, Egor V Kostylev, and Boris Motik. 2021. Explainable GNN-Based Models over Knowledge Graphs. In *International Conference on Learning Representations*.

[7] Enyan Dai and Suhang Wang. 2021. Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 680–688.

[8] Christopher Ifeanyi Eke, Azah Anir Norman, Liyana Shuib, and Henry Friday Nweke. 2019. A survey of user profiling: State-of-the-art, challenges, and solutions. *IEEE Access* 7 (2019), 144907–144924.

[9] European-Commission. 2019. *Ethics guidelines for trustworthy AI*. Publications Office. https://doi.org/doi/10.2759/346720

[10] Elizabeth Gómez, Carlos Shui Zhang, Ludovico Boratto, Maria Salamó, and Mirko Marras. 2021. The winner takes it all: geographic imbalance and provider (un) fairness in educational recommender systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1808–1812.

[11] Yulong Gu, Zhuoye Ding, Shuaiqiang Wang, and Dawei Yin. 2020. Hierarchical user profiling for e-commerce recommender systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 223–231.

[12] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).

[13] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.

[14] Sumitkumar Kanoje, Sheetal Girase, and Debajyoti Mukhopadhyay. 2015. User profiling trends, techniques and applications. *arXiv preprint arXiv:1503.07474* (2015).

[15] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Conference Track Proceedings*.

[16] Lizi Liao, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2018. Attributed social network embedding. *IEEE Transactions on Knowledge and Data Engineering* 30, 12 (2018), 2257–2270.

[17] Mohammad Naiseh, Nan Jiang, Jianbing Ma, and Raian Ali. 2020. Personalising explainable recommendations: Literature and conceptualisation. In *World Conference on Information Systems and Technologies*. Springer, 518–533.

[18] Danny Poo, Brian Chng, and Jie-Mein Goh. 2003. A hybrid approach for user profiling. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*. IEEE, 9–13.

[19] Erasmo Purificato, Ludovico Boratto, and Ernesto William De Luca. 2022. Do Graph Neural Networks Build Fair User Models? Assessing Disparate Impact and Mistreatment in Behavioural User Profiling. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 4399–4403.

[20] Erasmo Purificato, Sabine Wehnert, and Ernesto William De Luca. 2021. Dynamic Privacy-Preserving Recommendations on Academic Graph Data. *Computers* 10, 9 (2021), 107.

[21] Afshin Rahimi, Trevor Cohn, and Timothy Baldwin. 2018. Semi-supervised user geolocation via graph convolutional networks. *arXiv preprint arXiv:1804.08049* (2018).

[22] Tahleen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. 2019. Fairwalk: towards fair graph embedding. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 3289–3295.

[23] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. 2019. *Explainable AI: interpreting, explaining and visualizing deep learning*. Vol. 11700. Springer Nature.

[24] Zhiqing Sun, Jian Tang, Pan Du, Zhi-Hong Deng, and Jian-Yun Nie. 2019. Divgraphpointer: A graph pointer network for extracting diverse keyphrases. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 755–764.

[25] Lubos Takac and Michal Zabovsky. 2012. Data analysis in public social networks. In *International scientific conference and international workshop present day trends of innovations*, Vol. 1. Present Day Trends of Innovations Lamza Poland.

[26] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).

[27] Qianwen Wang, Yao Ming, Zhihua Jin, Qiaomu Shen, Dongyu Liu, Micah J Smith, Kalyan Veeramachaneni, and Huamin Qu. 2019. Atmseer: Increasing transparency and controllability in automated machine learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[28] Chuhan Wu, Fangzhao Wu, Junxin Liu, Shaojian He, Yongfeng Huang, and Xing Xie. 2019. Neural demographic prediction using search query. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining*. 654–662.

[29] Qilong Yan, Yufeng Zhang, Qiang Liu, Shu Wu, and Liang Wang. 2021. Relation-aware Heterogeneous Graph for User Profiling. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 3573–3577.

[30] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7370–7377.

[31] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. Gnn explainer: A tool for post-hoc explanation of graph neural networks. *arXiv preprint arXiv:1903.03894* (2019).

[32] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 974–983.

[33] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. 2022. Explainability in graph neural networks: A taxonomic survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).

[34] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla. 2019. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 793–803.

[35] Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2022. Deep Learning on Graphs: A Survey. *IEEE Transactions on Knowledge and Data Engineering* 34, 1 (2022), 249–270.