



First International Workshop on Graph-Based Approaches in Information Retrieval (IRonGraphs 2024)

Ludovico Boratto¹ , Daniele Malitesta² , Mirko Marras¹  ,
Giacomo Medda¹ , Cataldo Musto³ , and Erasmo Purificato⁴ 

¹ University of Cagliari, Cagliari, Italy

{ludovico.boratto,mirko.marras}@acm.org, giacomo.medda@unica.it

² Polytechnic University of Bari, Bari, Italy

daniele.malitesta@poliba.it

³ University of Bari, Bari, Italy

cataldo.musto@uniba.it

⁴ Otto von Guericke University Magdeburg, Magdeburg, Germany

erasmo.purificato@ovgu.de

Abstract. In the dynamic field of information retrieval, the adoption of graph-based approaches has become a notable research trend. Fueled by the growing research on Knowledge Graphs and Graph Neural Networks, these approaches rooted in graph theory have shown significant promise in enhancing the effectiveness and relevance of information retrieval results. With this motivation in mind, this workshop serves as a platform, bringing together researchers and practitioners from diverse backgrounds, to delve into and discuss the integration of modern graph-based methodologies into information retrieval methods. The workshop website is available at <https://irongraphs.github.io/ecir2024/>.

Keywords: Graphs · Graph Neural Networks · Knowledge Graphs · Information Retrieval · Algorithms · Search · Recommendation

1 Motivation and Relevance

In recent years, Information Retrieval (IR) has witnessed a shift in focus, driven by the increasing complexity of data structures, user demands for personalised experiences, and the growth of interconnected information sources. Traditional linear models, while foundational, often struggle to leverage rich interconnections that define the digital age. In this scenario, graph-based approaches have proven their ability to model complex relationships and semantic connections [11, 13, 19].

Graph-based approaches, such as those driven by Graph Neural Networks (GNNs) [14, 16, 24, 28, 29], allow us to extract knowledge from networks of connections, changing the way we approach IR [4, 12, 15, 17, 18, 22, 23, 25–27]. For

instance, by traversing the graph through its edges, we can better understand the context and meaning embedded in the data. This semantic understanding enables personalised and context-aware IR systems. As another example, by analysing user interactions within the graph, we can better investigate beyond-accuracy aspects, such as fairness and explainability [1–3, 5–10, 20, 21].

Given the importance of graph-based approaches and the rapidly-changing techniques driving search and recommendation, the *First International Workshop on Graph-based Approaches in Information Retrieval* (IRonGraphs 2024) represents the first ECIR’s workshop that aims to collect novel contributions in this continuously growing field and provide a common ground for interested researchers and practitioners within the European IR community.

2 Workshop Vision, Objectives, and Outcomes

The workshop vision is centred around the idea that graphs are a crucial data structure for modeling various domains. The success of graph representation learning has prompted the proposal for a dedicated event focused on IR topics. The resulting objectives include raising awareness of graph-based approaches in the IR community, identifying areas within IR that can benefit from such approaches, soliciting contributions addressing graph-based approaches, familiarising the IR community with graph-based research and practices, and uncovering gaps in IR research based on real-world needs. As outcomes, the workshop includes compiling and publishing workshop proceedings in the *CCIS* series by *Springer* as well as sharing slides and recordings associated with accepted papers. Additional outcomes will be represented by a strengthened community of researchers in graph-based approaches within IR, fostering future collaboration.

3 Workshop Format and Content Outline

The workshop involves a series of 40-minute spot thematic and interactive sessions (possibly 4 in total), wherein accepted contributions are grouped into 3–4 papers based on their topic or applicative domain. Each session comprises paper presentations lasting 7 to 10 min each (30 min in total), followed by a 10-minute discussion on the papers. Workshop organisers will facilitate these discussions, extracting lessons learned and providing brainstorming points on the session’s theme. Thematic sessions will be complemented by two 60-minute keynote talks, one from academia and another from industry. Insights from thematic sessions and keynote talks will inform a final discussion lasting 60 min, possibly with renowned experts, to establish a roadmap for shared initiatives.

4 Intended Audience

This workshop targets a diverse audience, encompassing individuals interested in exploring innovative contributions to graph-based approaches. This audience

includes researchers specialising in information retrieval, machine learning, and deep learning, as well as practitioners from both academia and industry. The workshop’s emphasis on graphs as ubiquitous data structures across various domains highlights its interdisciplinary nature, extending beyond algorithmic considerations, fostering engagement among participants with varied expertise.

5 Workshop Organisers’ Biography

Ludovico Boratto (webpage: <https://www.ludovicoboratto.com/>) is an Assistant Professor at the Dept. of Mathematics and Computer Science of the University of Cagliari (Italy). His research interests focus on recommender systems and their impact on stakeholders. He has co-authored over 60 papers published in top-tier conference proceedings and journals. His research brought him to give talks and tutorials at top-tier research centres and conferences, including CIKM, UMAP, RecSys, ICDE, ECIR, WSDM, ICDM, DSAA, and ECAI. 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 part of the program committees of the main Web conferences, where he received four outstanding contribution awards. In 2012, he got his Ph.D. at the University of Cagliari. From May 2016 to April 2021, he joined Eurecat as a Senior Research Scientist in the Data Science and Big Data Analytics research group. In 2010 and 2014, he visited Yahoo! Research in Barcelona.

Daniele Malitesta (webpage: <https://danielemalitesta.github.io>) is a PhD candidate at the Polytechnic University of Bari (Italy). His current research focuses on recommendation algorithms leveraging side information, with a focus on graph- and multimodal-based recommender systems. He has published papers at top-tier conferences, such as SIGIR, ECIR, RecSys, and MM, and served as a reviewer at venues such as ICLR, NeurIPS, SIGIR, RecSys, ECIR, and LoG. He is a co-developer of Elliot, a framework for the rigorous evaluation and reproducibility of recommender systems. Last summer, he visited Pasquale Minervini at the University of Edinburgh during a PhD internship. Recently, he presented a tutorial entitled “Graph Neural Networks for Recommendation: Reproducibility, Graph Topology, and Node Representation” at LoG 2023.

Mirko Marras (webpage: <https://www.mirkomarras.com/>) is an Assistant Professor at the Dept. of Mathematics and Computer Science of the University of Cagliari (Italy). Before that, he was a postdoctoral researcher at EPFL (Switzerland) and a visiting scholar at Eurecat (Spain) and New York University (USA). His research ranges across various domains impacted by user modelling and personalization, including business, education, entertainment, and healthcare. He has co-authored more than 80 papers in top-tier conferences and journals and has given tutorials at ECML-PKDD, RecSys, ICDE, ECIR, WSDM, ICDM, and UMAP. He has also co-chaired several workshops on related themes at ECIR,

WSDM, ICCV, EDM, and ECML-PKDD. He is part of the program committees of top-tier conferences in the field, where he received three outstanding reviewer awards. He is an associate editor for Springer’s Journal of Ambient Intelligence and Humanized Computing and Neural Processing Letters.

Giacomo Medda (webpage: <https://jackmedda.github.io/>) is a PhD candidate at the Dept. of Mathematics and Computer Science of the University of Cagliari (Italy). In 2022, he has been a visiting scholar at Eurecat (Spain), where he worked on graph-based recommender systems with Dr. Francesco Fabbri. His research has focused on analysing and developing beyond-accuracy algorithms to improve the fairness and explainability of recommender and speaker recognition systems. He has published papers in top-tier conferences, such as CIKM and ECIR, and journals, such as IPM and PRLetters. He has served as a reviewer for flagship workshops, such as BIAS and KaRS, and journals, such as Neural Processing Letters and TOIS.

Cataldo Musto (webpage: <http://www.di.uniba.it/~swap/musto>) is an Assistant Professor at the Dept. of Informatics, University of Bari. He completed his PhD in 2012 under the supervision of Prof. Giovanni Semeraro. His research focuses on the adoption of natural language processing techniques for fine-grained semantic content representation in recommender systems and user modelling platforms. He acts as a program committee member for RecSys and UMAP, and he organized several events on user modelling and recommender systems. In 2016 and 2017, he gave a tutorial at UMAP conference on semantics-aware representation in content-based personalized systems. Recently, he organized the workshop on Explainable User Modeling (ExUM), jointly held with UMAP 2019–2023.

Erasmus Purificato (webpage: <https://erasmopurif.com/>) has been a Research Assistant in the “Human-Centred AI” group at the Otto von Guericke University Magdeburg (OVGU) and in the “Human-Centred Technologies for Educational Media” department at the Leibniz Institute for Educational Media | Georg Eckert Institute (GEI), since February 2020. He is doing his PhD in Computer Science at OVGU with a project entitled “Human-Centred Fairness Analysis on Graph Neural Network-Based Models for Behavioural User Profiling”. From July 2021, he has been appointed by the Guglielmo Marconi University of Rome, Italy, as an Adjunct Professor in “Software Engineering”. He co-organised the HCAI4U Workshop at CHIItaly, the APEX-UI Workshop at IUI, and the IEEE Autumn School ISACT at ICHMS. He has given tutorials at CIKM and UMAP. 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 and TITS.

References

1. Abdelrazek, M., Purificato, E., Boratto, L., De Luca, E.W.: FairUP: a framework for fairness analysis of graph neural network-based user profiling models. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR, p. 3165-3169. SIGIR 2023, ACM (2023)
2. Anelli, V.W., Deldjoo, Y., Noia, T.D., Malitesta, D., Paparella, V., Pomo, C.: Auditing consumer- and producer-fairness in graph collaborative filtering. In: Proceedings of the 45th European Conference on Information Retrieval, ECIR. LNCS, vol. 13980, pp. 33–48. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-28244-7_3
3. Anelli, V.W., et al.: How neighborhood exploration influences novelty and diversity in graph collaborative filtering. In: Proceedings of the 2nd Workshop on Multi-Objective Recommender Systems co-located with 16th ACM Conference on Recommender Systems, RecSys. CEUR Workshop Proceedings, vol. 3268. CEUR-WS.org (2022)
4. Anelli, V.W., Malitesta, D., Pomo, C., Bellogín, A., Sciascio, E.D., Noia, T.D.: Challenging the myth of graph collaborative filtering: a reasoned and reproducibility-driven analysis. In: Proceedings of the 17th ACM Conference on Recommender Systems, RecSys, pp. 350–361. ACM (2023)
5. Balloccu, G., Boratto, L., Cancedda, C., Fenu, G., Marras, M.: Faithful path language modelling for explainable recommendation over knowledge graph. CoRR [abs/2310.16452](https://arxiv.org/abs/2310.16452) [arXiv:2310.16452](https://arxiv.org/abs/2310.16452) (2023)
6. Balloccu, G., Boratto, L., Cancedda, C., Fenu, G., Marras, M.: Knowledge is power, understanding is impact: Utility and beyond goals, explanation quality, and fairness in path reasoning recommendation. In: Proceedings of the 45th European Conference on Information Retrieval, ECIR. LNCS, vol. 13982, pp. 3–19. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-28241-6_1
7. Balloccu, G., Boratto, L., Fenu, G., Marras, M.: Post processing recommender systems with knowledge graphs for recency, popularity, and diversity of explanations. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR, pp. 646–656. ACM (2022)
8. Balloccu, G., Boratto, L., Fenu, G., Marras, M.: Reinforcement recommendation reasoning through knowledge graphs for explanation path quality. *Knowl. Based Syst.* **260**, 110098 (2023)
9. Boratto, L., Fabbri, F., Fenu, G., Marras, M., Medda, G.: Counterfactual graph augmentation for consumer unfairness mitigation in recommender systems. In: Proceeding of the 32nd ACM International Conference on Information and Knowledge Management, CIKM, pp. 3753–3757. ACM (2023)
10. Boratto, L., Fenu, G., Marras, M., Medda, G.: Practical perspectives of consumer fairness in recommendation. *Inf. Process. Manag.* **60**(2), 103208 (2023)
11. Cambria, E., Mao, R., Han, S., Liu, Q.: Sentic parser: a graph-based approach to concept extraction for sentiment analysis. In: Proceeding of the IEEE International Conference on Data Mining Workshops, ICDM - Workshops, pp. 1–8. IEEE (2022)
12. Chen, J., Zhu, G., Hou, H., Yuan, C., Huang, Y.: AutoGSR: neural architecture search for graph-based session recommendation. In: Proceeding of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR, pp. 1694–1704. ACM (2022)

13. Halilaj, L., Dindorkar, I., Lüttin, J., Rothermel, S.: A knowledge graph-based approach for situation comprehension in driving scenarios. In: Verborgh, R., et al. (eds.) *ESWC 2021*. LNCS, vol. 12731, pp. 699–716. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77385-4_42
14. Hamilton, W.L., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: *Proceedings of the Annual Conference on Neural Information Processing Systems, NIPS*, pp. 1024–1034 (2017)
15. Kamphuis, C.: Graph databases for information retrieval. In: Jose, J.M., et al. (eds.) *ECIR 2020*. LNCS, vol. 12036, pp. 608–612. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45442-5_79
16. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: *Proceeding of the 5th International Conference on Learning Representations, ICLR*. OpenReview.net (2017)
17. Liu, S., Ounis, I., Macdonald, C.: An MLP-based algorithm for efficient contrastive graph recommendations. In: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR*, pp. 2431–2436. ACM (2022)
18. Malitesta, D., Pomo, C., Anelli, V.W., Mancino, A.C.M., Sciascio, E.D., Noia, T.D.: A topology-aware analysis of graph collaborative filtering. *CoRR* **abs/2308.10778** [arXiv:2308.10778](https://arxiv.org/abs/2308.10778) (2023)
19. Mayank, M., Sharma, S., Sharma, R.: DEAP-FAKED: knowledge graph based approach for fake news detection. In: *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM*, pp. 47–51. IEEE (2022)
20. Medda, G., Fabbri, F., Marras, M., Boratto, L., Fenu, G.: GNNUERS: fairness explanation in GNNs for recommendation via counterfactual reasoning. *CoRR* **abs/2304.06182** [arXiv:2304.06182](https://arxiv.org/abs/2304.06182) (2023)
21. Purificato, E., Boratto, L., De Luca, E.W.: 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, CIKM*, p. 4399–4403. CIKM 2022, ACM (2022)
22. Spillo, G., Musto, C., Polignano, M., Lops, P., de Gemmis, M., Semeraro, G.: Combining graph neural networks and sentence encoders for knowledge-aware recommendations. In: *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization, UMAP*, p. 1–12. UMAP 2023, ACM (2023)
23. Thonet, T., Renders, J.-M., Choi, M., Kim, J.: Joint personalized search and recommendation with hypergraph convolutional networks. In: Hagen, M., et al. (eds.) *ECIR 2022*. LNCS, vol. 13185, pp. 443–456. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99736-6_30
24. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: *Proceedings of the 6th International Conference on Learning Representations, ICLR*. OpenReview.net (2018)
25. Witschel, H.F., Riesen, K., Grether, L.: KvGR: a graph-based interface for explorative sequential question answering on heterogeneous information sources. In: Jose, J.M., et al. (eds.) *ECIR 2020*. LNCS, vol. 12035, pp. 760–773. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45439-5_50
26. Yi, Z., Ounis, I., Macdonald, C.: Graph contrastive learning with positional representation for recommendation. In: *Proceedings of the 45th European Conference on Information Retrieval, ECIR*. *Lecture Notes in Computer Science*, vol. 13981, pp. 288–303. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-28238-6_19

27. Yu, H.C., Dai, Z., Callan, J.: PGT: pseudo relevance feedback using a graph-based transformer. In: Hiemstra, D., Moens, M.-F., Mothe, J., Perego, R., Potthast, M., Sebastiani, F. (eds.) ECIR 2021. LNCS, vol. 12657, pp. 440–447. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72240-1_46
28. Zhang, C., Song, D., Huang, C., Swami, A., Chawla, N.V.: Heterogeneous graph neural network. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, SIGKDD, pp. 793–803 (2019)
29. Zhang, Z., Cui, P., Zhu, W.: Deep learning on graphs: a survey. *IEEE Trans. Knowl. Data Eng.* **34**(1), 249–270 (2022)