

Report on the 13th Italian Information Retrieval Workshop (IIR 2023)

Guglielmo Faggioli

University of Padova
Italy

`guglielmo.faggioli@unipd.it`

Antonio Ferrara

Politecnico di Bari
Italy

`antonio.ferrara@poliba.it`

Franco Maria Nardini

ISTI-CNR
Italy

`francomaria.nardini@isti.cnr.it`

Nicola Tonellotto

University of Pisa
Italy

`nicola.tonellotto@unipi.it`

Abstract

The 13th Italian Information Retrieval Workshop is the thirteenth edition of the annual conference of the Italian information retrieval and recommender systems communities. The two days of the conference gathered interesting studies and research work on a wide range of topics on information retrieval, recommender systems, and natural language processing, such as Search and Ranking, Recommendation, Content Analysis, and Classification, Artificial Intelligence, NLP, Semantics, and Dialog, Domain-Specific Applications, Human Factors and Interfaces, and Evaluation. It was hosted by the National Research Council (CNR) of Italy and the University of Pisa in a conference facility in Pisa, Italy.

Date: 8–9 June 2023.

Website: <http://iir2023.isti.cnr.it/>.

1 Introduction

The 13th Italian Information Retrieval Workshop (IIR 2023) is the thirteenth annual conference of the Italian information retrieval and recommender systems communities.

IIR 2023 took place on June 2023 in Pisa. The Italian Information Retrieval (IIR) Workshop was created in 2009 at the European Summer School in Information Retrieval. Its purpose is to provide a meeting forum for stimulating and disseminating research in Information Retrieval, where early-stage researchers (including MSc and PhD students) can network and discuss their research results in an informal way, and to give them the the opportunity to meet and discuss with senior researchers. IIR is organized every year by one of the main research groups in Italy in different locations. Previous IIR workshops took place in Milan (2022), Bari (2021), Padua (2019) Rome (2018), Lugano (2017), Venice (2016), Cagliari (2015), Rome (2014), Pisa (2013), Bari (2012), Milan (2011), and Padua (2010).

In 2023, the conference received 33 contributions from Italian and International institutions, with 24 papers published on its proceedings. The papers presented at the workshop covered several different research topics, including core information retrieval and recommender systems algorithms, privacy, fairness, domain-specific applications, NLP tasks, and evaluation.

IIR has been *free of charge* and relied on the hosting institutions for organizational and financial support. IIR 2023 has been partially supported by ACM SIGIR and ISTI-CNR.

2 The Workshop

2.1 Keynotes

IIR 2023 featured two invited keynote speakers. We present the title and summary of each talk below along with a short biography of each speaker.

Sean Macavaney — Re-Thinking Re-Ranking

Abstract Re-ranking systems take a “cascading” approach, wherein an initial candidate pool of documents is ranked and filtered to produce a final result list. This approach exhibits a fundamental relevance misalignment problem: the most relevant documents may be filtered out by a prior stage as insufficiently relevant, ultimately reducing recall and limiting the potential effectiveness. In this talk, I challenge the cascading paradigm by proposing methods that efficiently pull in additional potentially-relevant documents during the re-ranking process, using the long-standing Cluster Hypothesis. I demonstrate that these methods can improve the efficiency and effectiveness of both bi-encoder and cross-encoder retrieval models at various operational points. Cascading is dead, long live re-ranking!

Speaker’s Biography Sean Macavaney is a Lecturer in Machine Learning at the University of Glasgow and a member of the Terrier Team. His research primarily focuses on effective and efficient neural retrieval. He completed his PhD at Georgetown University in 2021, where he was a member of the IR Lab and an ARCS Endowed Scholar. He was a co-recipient of the ECIR 2023 Best Short Paper award.

Christine Bauer — Three views of a secret: Embracing multiple perspectives in the evaluation of information retrieval and recommender systems

Abstract As users are increasingly confronted with information and choice overload, we need the ‘right’ information, at the ‘right’ time, in the ‘right’ place, in the ‘right’ way, to the ‘right’ person. Information retrieval and recommender systems are effective means to address this goal. When optimizing and evaluating such systems, we often disregard that a ‘typical’ user is not the only stakeholder interested in a well-functioning system. Beyond ignoring the needs of specific stakeholders, this eventually leads to a malfunctioning system for anyone. In this talk, I will demonstrate that we need to consider the demands of the various stakeholders and provide insights into how we can embrace those needs when evaluating our systems.

Speaker’s Biography Christine Bauer is EXDIGIT Professor for Interactive Intelligent Systems at the Department of Artificial Intelligence and Human Interfaces (AIHI) at the Paris Lodron University Salzburg, Austria. Her research centers on interactive intelligent systems. In recent years, she worked on context-aware recommender systems in the music and media domains. The core interests in her research activities are fairness and multi-method evaluations. She has authored more than 100 papers and holds several best paper awards and several awards for her reviewing activities. She is on the Editorial Board of ACM Transactions on Recommender Systems (TORS) and co-organizes the Workshop series “Perspectives on the Evaluation of Recommender Systems (PERSPECTIVES)”.

2.2 Papers

The purpose of the Italian Information Retrieval (IIR) workshop series is to provide a forum for stimulating and disseminating research in information retrieval, where Italian researchers and researchers affiliated with Italian institutions can network and discuss their research results in an informal way.

The contributions to IIR 2023 mainly address the following relevant topics:

- *Search and Ranking*: Research on core Information Retrieval (IR) algorithmic topics, including IR at scale
- *Recommendation, Content Analysis, and Classification*: Research focusing on recommender systems (RS), rich content representations and content analysis.
- *Artificial Intelligence, NLP, Semantics, and Dialog*: Research bridging AI and IR –, especially toward deep semantics — and dialog with intelligent agents.
- *Domain-Specific Applications*: Research focusing on domain-specific challenges.
- *Human Factors and Interfaces*: Research into user-centric aspects of IR, including user interfaces, behavior modeling, privacy, and interactive systems.
- *Evaluation*: Research that focuses on the measurement and evaluation of IR systems.
- *Future Directions*: Research with theoretical or empirical contributions on new technical or social aspects of IR, especially in more speculative directions or with emerging technologies.

We received submissions of original papers presenting new research results and extended abstracts containing descriptions of ongoing projects or summarizing already published results. The program committee involved 38 researchers, equally divided among IR and RS. At least two program committee members reviewed each submission. The reviewers looked at originality, technical depth, style of presentation, and impact. Finally, the committee decided to accept 5 original papers and 28 extended abstracts for presentation at the workshop. Out of these, 24 papers were further published on the proceedings through the CEUR-WS publisher¹ [Nardini et al., 2023].

The program also included two invited talks. The first talk, titled “Re-Thinking Re-Ranking”, was given by Sean Macavaney from the University of Glasgow. The second was by Christine Bauer from the Paris Lodron University Salzburg and was entitled “Three views of a secret: Embracing multiple perspectives in the evaluation of information retrieval and recommender systems”.

¹<http://ceur-ws.org/Vol-3448/>

2.3 Organization

2.3.1 General Chairs

Franco Maria Nardini, *ISTI-CNR, Italy*

Nicola Tonello, *University of Pisa, Italy*

2.3.2 Program Chairs

Guglielmo Faggioli, *University of Padova, Italy*

Antonio Ferrara, *Polytechnic University of Bari, Italy*

2.3.3 Webmaster & Graphics Design

Beatrice Rapisarda *ISTI-CNR, Italy*

2.3.4 Steering Committee

Aris Anagnostopoulos, *University of Rome “La Sapienza”, Italy*

Luca Becchetti, *University of Rome “La Sapienza”, Italy*

Fabio Crestani, *Università della Svizzera Italiana, Switzerland*

Emanuele Di Buccio, *University of Padova, Italy*

Massimo Melucci, *University of Padova, Italy*

Stefano Mizzaro, *University of Udine, Italy*

Tommaso Di Noia, *Polytechnic University of Bari, Italy*

Gabriella Pasi, *University Milano Bicocca, Italy*

Raffaele Perego, *ISTI-CNR, Italy*

Fabrizio Silvestri, *University of Rome “La Sapienza”, Italy*

Marko Tkalcić, *Free University of Bolzano, Italy*

Nicola Tonello, *University of Pisa, Italy*

Franco Maria Nardini, *ISTI-CNR, Italy*

Giorgio Maria Di Nunzio, *University of Padova, Italy*

Salvatore Orlando, *University Ca’ Foscari Venezia, Italy*

2.4 Program Committee

Maristella Agosti, *University of Padova, Italy*

Aris Anagnostopoulos, *Sapienza University of Rome, Italy*

Marco Angelini, *University of Rome “La Sapienza”, Italy*

Vito Walter Anelli, *Polytechnic University of Bari, Italy*

Liliana Ardissono, *University of Torino, Italy*

Pierpaolo Basile, *University of Bari Aldo Moro, Italy*

Giovanni Maria Biancofiore, *Polytechnic University of Bari, Italy*

Paolo Boldi, *University of Milan, Italy*

Ludovico Boratto, *University of Cagliari, Italy*

Giandomenico Cornacchia, *Polytechnic University of Bari, Italy*

Paolo Cremonesi, *Polytechnic University of Milano, Italy*
Danilo Croce, *University of Rome “Tor Vergata”, Italy*
Gianluca Demartini, *University of Queensland, Italy*
Emanuele Di Buccio, *University of Padova, Italy*
Tommaso Di Noia, *Polytechnic University of Bari, Italy*
Giorgio Maria Di Nunzio, *University of Padova, Italy*
Francesco Donini, *Università della Tuscia, Italy*
Andrea Esuli, *ISTI-CNR, Italy*
Maurizio Ferrari Dacrema, *Polytechnic University of Milano, Italy*
Nicola Ferro, *University of Padova, Italy*
Claudio Lucchese, *University Ca’ Foscari Venezia, Italy*
Daniele Malitesta, *Polytechnic University of Bari, Italy*
Stefano Marchesin, *University of Padova, Italy*
David Massimo, *Free University of Bolzano, Italy*
Alberto Carlo Maria Mancino, *University Polytechnic of Bari, Italy*
Ida Mele, *IASI-CNR, Italy*
Alejandro Moreo, *ISTI-CNR, Italy*
Cristina Muntean, *ISTI-CNR, Italy*
Stefano Mizzaro, *University of Udine, Italy*
Raffaele Perego, *ISTI-CNR, Italy*
Marco Polignano, *University of Bari, Italy*
Claudio Pomo, *Polytechnic University of Bari, Italy*
Francesco Ricci, *Free University of Bozen-Bolzano, Italy*
Giovanni Semeraro, *University of Bari “Aldo Moro”, Italy*
Gianmaria Silvello, *University of Padova, Italy*
Marko Tkalcić, *Free University of Bozen Bolzano, Italy*
Salvatore Trani, *ISTI-CNR, Italy*
Markus Zanker, *Free University of Bozen-Bolzano, Italy*

3 Summary

3.1 Session 1: User Satisfaction and Personalization

The initial session centered on strategies aimed at enhancing user satisfaction through the delivery of personalized content, while also assessing potential user dissatisfaction. In this context, [Fabris et al. \[2023\]](#) addressed the limitations of the family of pairwise fairness measures and introduced a new metric, called “Dissatisfaction Induced by Pairwise Swaps,” to overcome these limitations. [Maurera et al. \[2023\]](#) delved into a specific type of user-system interaction, namely impressions, and provided a comprehensive overview of their role in recommender systems from three distinct perspectives: models, datasets, and evaluation. [Massimo and Ricci \[2023\]](#) detailed their reinforcement learning-based RS approach, known as QEXP, which leverages logs of Point of Interest (POI) visits to address the cold-start problem while suggesting new POIs based on tourist movement patterns. [Braga et al. \[2023\]](#) initiated their discussion by highlighting the intricacies of fine-tuning a personalized BERT model. To tackle this challenge, they introduced an approach that person-

alizes dialogue systems using adapter layers and topic modeling, incorporating a user embedding vector as an external input to encode personal information into the model. Finally, [Manzo et al. \[2023\]](#) addressed the context of public streaming platforms and proposed an innovative approach to personalized recommendations for streaming media content. Their model prioritizes serendipity and multicultural diversity while minimizing the need for sharing personal data.

3.2 Session 2: Privacy, Fairness, and Explainability

The focus of the second session centered primarily on exploring fairness and explainability within various information access methodologies. Several approaches investigated the fairness aspects of recommender systems, while others provided explainable solutions to information access. In particular, [Atzori et al. \[2023\]](#) examined the bias in the exposure of demographic groups in forensic face rankings by introducing a framework comprising six face models and a public dataset. Their research revealed that biases in forensic face rankings are significant, necessitating the development of specialized fairness concepts and mitigation strategies. [Balloccu et al. \[2023\]](#) summarized their work on replicating and evaluating three state-of-the-art path reasoning recommendation approaches through a standardized protocol utilizing publicly available datasets. Their findings indicated that these reasoning methods, compared to knowledge-aware baselines, exhibited reduced utility and coverage but offered increased serendipity and diversity. [Di Palma et al. \[2023\]](#) aimed to address concerns related to fairness in graph-based Collaborative Filtering (CF) recommendations by proposing a taxonomy of graph CF solutions. They subsequently used this taxonomy in their analysis, revealing that while graph CF models demonstrated superior accuracy performance based on Consumer and Producer (CP)-fairness metrics, this came at the expense of user fairness. [Boratto et al. \[2023\]](#) summarized their previous work on defining desirable properties for a fair recommendation system in the context of recommender systems. They identified eight properties and assessed how existing mitigation procedures aligned with these properties using two public datasets. [Afreen et al. \[2023\]](#) delved into the performance of existing path reasoning methods, specifically for educational recommendation. Afreen et al. showed promising results when these methods were applied to two large-scale online course datasets. [Giachelle et al. \[2023\]](#) extended the capabilities of the Semantic Knowledge Extractor Tool (SKET) by introducing SKET eXplained (SKET X). This web-based system provides visual explanations that detail the models, rules, and parameters associated with each prediction, catering to pathologists and experts to facilitate the comprehension of SKET predictions.

3.3 Session 3: NLP

The focus of the third session was on NLP at large. Among the NLP tasks addressed this year, we recognize authorship verification, fact-checking, the development of Italian LLMs, and topic modeling. [Corbara and Moreo \[2023\]](#) proposed their results to tackle the adversarial Authorship Verification task by augmenting the training set with synthetic textual examples. Corbara and Moreo experimented with SVM and Neural Networks, and generation strategies based on language modeling and GAN training, empirically showing that data augmentation allows for improved performance in an adversarial setup. To obtain annotations for the fact-checking task, [La Barbera et al. \[2023\]](#) investigated the possibility of employing crowd workers. Their experiments were meant

to identify the best setting for misinformation assessment using crowdsourcing. Their results indicated that the crowd can effectively address misinformation, at a lower cost than experts. [Bacciu et al. \[2023\]](#) presented Fauno, the first and largest open-source Italian conversational Large Language Model (LLM). Through Fauno, Bacciu et al. demonstrated the possibility of obtaining a fine-tuned conversational bot with a single GPU is possible. Additionally, Bacciu et al. released and described the collection of datasets for conversational AI in Italian that they used for fine-tuning Fauno. [Ravenda et al. \[2023\]](#) addressed the topic-modeling task, by proposing an approach based on Mixtures of Normalizing Flows, which extracts topics from a collection of time-varying documents in a dynamic way. Their experiments highlighted that the approach proposed, which dynamically updates the topic representations is on par or better than state-of-the-art solutions.

3.4 Session 4: Efficient and Domain Specific Approaches

The fourth session was divided into two parts. The first concerned how to improve the efficiency of the models, while the latter was focused on domain-specific approaches.

Two papers investigated how to improve the performance in the IR domain. The former focused on accelerating the query efficiency via index decomposition, while the latter investigated the complexity of solving IR problems using Quantum Annealers (QA). [Mackenzie et al. \[2023\]](#) introduced a novel technique to improve the query efficiency, named “postings clipping”. The approach proposed by Mackenzie et al. differed from current efforts because every posting remains in the original posting list, while the high-impact postings are sliced and moved to a different list. [Pellini et al. \[2023\]](#) started by observing that theoretical studies indicate that specific optimization problems are more difficult to solve on QA. Therefore, they proposed a preliminary empirical investigation to recognize such hard-to-be-solved problems. Pellini et al. observed a correlation between the features of a problem and its complexity to be solved on a QA.

Regarding domain-specific approaches, the investigated scenarios involved investigative intelligence, retrieval of clinical trials, and the generation of profiles for tourists. [Campinas et al. \[2023\]](#) described how the Siren Federate plugin integrates relational search capabilities into Elasticsearch and how it allows to tackle some of the challenges in Investigative Intelligence. [Di Nunzio et al. \[2023\]](#) presented their previous and current work concerning the Retrieval of Clinical Trials. Their solutions are based on three techniques: query reformulation, the merging of the ranking lists obtained using summarized queries, and query expansion based on pseudo-relevance feedback. To provide personalized tourist profiles, [Merinov et al. \[2023\]](#) proposed an approach to synthesize individual-level user profiles from population-level data in the tourism domain. Merinov et al. utilized a discrete choice behavioral model to merge choices made by the population at large with the choices in the real data.

3.5 Session 5: Evaluation

The final session concerned the evaluation. The submitted works regarded mostly IR and included the theoretical analysis of evaluation models, the empirical assessment of LambdaMART, the definition of a runs repository to ease experimentation, and the query performance prediction. [Faggioli et al. \[2023a\]](#) expanded the current evaluation strategy in IR, which is based on General Linear Models (GLiMs), to Generalized Linear Models (GLMs), observing an improved

power of the statistical analyses while maintaining comparable stability. Bassani [2023] introduced “ranxhub,” an online repository designed for the sharing of pre-computed runs, which are ranked lists of documents retrieved for specific query sets by retrieval models. Ranxhub aimed to expedite research, enhance transparency, and reduce the environmental footprint associated with experimentation. Lucchese et al. [2023] analyzed LambdaMART gradients, presenting strategies to mitigate gradient inconsistencies. Their approach involved computing the full gradient only for documents that should be ranked within the top-k positions, ensuring more consistent results. Faggioli et al. [2023b] described an approach based on the geometric properties of dense spaces to improve Query Performance Prediction (QPP) in the conversational search domain. The model proposed by Faggioli et al. is based on measuring the volume of the hypercuboid that encompasses the retrieved documents.

4 Lessons Learned

The 13th Italian Information Retrieval Workshop encouraged the discussion on IR and RS from several perspectives, fostering the exchange of ideas and methodologies between people working on various research areas.

A key lesson was the significance of addressing the needs of different stakeholders when evaluating information retrieval and recommender systems. Christine Bauer’s keynote stressed the importance of considering multiple perspectives to ensure fairness and effectiveness of the system. Moreover, the presentation by Sean Macavaney showcased novel approaches for improving re-ranking in information retrieval challenging conventional methods, thus underscoring the need for flexible and innovative approaches in our community.

The papers presented at the workshop covered a wide spectrum of research topics, ranging from core information retrieval algorithms to privacy, fairness, domain-specific applications, and NLP tasks. The presentations and the discussions offered a platform for new and experienced researchers from the Italian IR and RS communities to learn from one another, share findings, exchange ideas, and stay current with the latest advancements in the field.

The major takeaways are strictly connected with the importance of user-centric design and adaptability in information retrieval and recommender systems design, as well as the significance of topics related to the trustworthiness of such systems such as privacy, fairness, explainability.

Finally, the workshop emphasized not only the need to develop new models but also the urgency for rigorous evaluation models and resources for reproducible and easy experimentation in the field of Information Retrieval and Recommender Systems.

In summary, the 13th Italian Information Retrieval Workshop has once again proven itself as an important platform for promoting collaboration, innovation, and the continuous advancement of Information Retrieval and Recommender Systems research, underscoring its enduring importance in our Italian community.

Acknowledgments

We thank CEUR-WS for hosting the proceedings. We thank the High-Performance Computing Laboratory (HPC Lab) at the Information Science and Technologies Institute (ISTI) of the Na-

tional Research Council of Italy (CNR) and the Special Interest Group on Information Retrieval (SIGIR) for their financial support.

References

- Neda Afreen, Giacomo Balloccu, Ludovico Boratto, Gianni Fenu, and Mirko Marras. Towards explainable educational recommendation through path reasoning methods. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 131–136. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-29.pdf>.
- Andrea Atzori, Gianni Fenu, and Mirko Marras. Fairness of exposure in forensic face rankings. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 91–96. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-12.pdf>.
- Andrea Bacciu, Giovanni Trappolini, Andrea Santilli, Emanuele Rodolà, and Fabrizio Silvestri. Fauno: The italian large language model that will leave you senza parole! In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 9–17. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-24.pdf>.
- Giacomo Balloccu, Ludovico Boratto, Christian Cancedda, Gianni Fenu, and Mirko Marras. Knowledge-aware recommendations: Exploring the interplay between utility, explanation quality, and fairness in path reasoning methods. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 111–116. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-26.pdf>.
- Elias Bassani. Towards a repository for information retrieval runs. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 125–130. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-04.pdf>.
- Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. Consumer fairness benchmark in recommendation. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 60–65. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-27.pdf>.
- Marco Braga, Alessandro Raganato, Gabriella Pasi, et al. Personalization in bert with adapter modules and topic modelling. In Franco Maria Nardini, Nicola Tonello, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop*

-
- (IIR 2023). Pisa, Italy, June 8-9, 2023, pages 24–29. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-13.pdf>.
- Stéphane Campinas, Matteo Catena, and Renaud Delbru. Relational search and its application to investigative intelligence scenarios. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 117–118. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-28.pdf>.
- Silvia Corbara and Alejandro Moreo. Enhancing adversarial authorship verification with data augmentation. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 73–78. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-11.pdf>.
- Giorgio Maria Di Nunzio, Guglielmo Faggioli, and Stefano Marchesin. An analysis of a methodology and experimental results for the retrieval of clinical trials. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 53–57. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-21.pdf>.
- Dario Di Palma, Vito Walter Anelli, Daniele Malitesta, Vincenzo Paparella, Claudio Pomo, Yashar Deldjoo, and Tommaso Di Noia. Examining fairness in graph-based collaborative filtering: A consumer and producer perspective. 3448:79–84, 2023. URL <https://ceur-ws.org/Vol-3448/paper-30.pdf>.
- Alessandro Fabris, Gianmaria Silvello, Gian Antonio Susto, and Asia J Biega. Dissatisfaction induced by pairwise swaps. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 66–71. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-10.pdf>.
- Guglielmo Faggioli, Nicola Ferro, Norbert Fuhr, et al. Ir systems evaluation via generalized linear models. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 105–110. CEUR-WS, 2023a. URL <https://ceur-ws.org/Vol-3448/paper-02.pdf>.
- Guglielmo Faggioli, Nicola Ferro, Cristina Muntean, Raffaele Perego, Nicola Tonellotto, et al. A spatial approach to predict performance of conversational search systems. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, volume 3448, pages 41–46. CEUR-WS, 2023b. URL <https://ceur-ws.org/Vol-3448/paper-07.pdf>.
- Fabio Giachelle, Stefano Marchesin, and Gianmaria Silvello. Sket x: A visual analytics tool for explaining knowledge extraction results. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval*

-
- Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 119–124. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-06.pdf>.
- David La Barbera, Michael Soprano, Kevin Roitero, Eddy Maddalena, Stefano Mizzaro, et al. Fact-checking at scale with crowdsourcing: Experiments and lessons learned. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 85–90. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-18.pdf>.
- Claudio Lucchese, Federico Marcuzzi, Salvatore Orlando, et al. Does lambdamart do what you expect? In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, volume 3448, page 72. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-16.pdf>.
- Joel Mackenzie, Antonio Mallia, Alistair Moffat, and Matthias Petri. Accelerating learned sparse indexes via term impact decomposition. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 47–52. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-01.pdf>.
- Gaetano Manzo, Yvan Pannatier, Gabriel Autès, Michaël De Lucia, Jean-Gabriel Piguet, and Jean-Paul Calbimonte. Serendipity and diversity boosting for personalized streaming media recommendation. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 30–40. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-03.pdf>.
- David Massimo and Francesco Ricci. Cold-start management in poi recommendation via reinforcement learning and spatial proximity exploration. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 58–59. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-09.pdf>.
- Fernando B Pérez Maurera, Maurizio Ferrari Dacrema, Pablo Castells, and Paolo Cremonesi. Impressions in recommender systems: Present and future. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, volume 3448, pages 97–104. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-23.pdf>.
- Pavel Merinov, David Massimo, and Francesco Ricci. Behaviour-aware tourist profiles data generation. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 3–8. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-08.pdf>.

Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors. *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, volume 3448, 2023. CEUR-WS. URL <http://ceur-ws.org/Vol-3448/>.

Riccardo Pellini, M Ferrari Dacrema, Paolo Cremonesi, et al. Towards improved qubo formulations of ir tasks for quantum annealers. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 137–142. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-05.pdf>.

Federico Ravenda, Andrea Raballo, Antonietta Mira, and Fabio Crestani. Incremental mixture of normalizing flows for dynamic topic modelling. In Franco Maria Nardini, Nicola Tonellotto, Guglielmo Faggioli, and Antonio Ferrara, editors, *Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023)*. Pisa, Italy, June 8-9, 2023, pages 18–23. CEUR-WS, 2023. URL <https://ceur-ws.org/Vol-3448/paper-19.pdf>.