

Report on the 14th Italian Information Retrieval Workshop (IIR 2024)

Eddy Maddalena

University of Udine
Italy

`eddy.maddalena@uniud.it`

Stefano Mizzaro

University of Udine
Italy

`stefano.mizzaro@uniud.it`

Kevin Roitero

University of Udine
Italy

`kevin.roitero@uniud.it`

Marco Viviani

University of Milano-Bicocca
Italy

`marco.viviani@unimib.it`

David La Barbera

University of Udine
Italy

`david.labarbera@uniud.it`

Sandip Modha

University of Milano-Bicocca
Italy

`sandip.modha@unimib.it`

Gabriella Pasi

University of Milano-Bicocca
Italy

`gabriella.pasi@unimib.it`

Francesca Da Ros

University of Udine
Italy

`francesca.daros@uniud.it`

Michael Soprano

University of Udine
Italy

`michael.soprano@uniud.it`

Abstract

IIR 2024, the 14th Italian Information Retrieval Workshop, served as the annual event for the Information Retrieval (IR) and Recommender Systems (RS) communities both in Italy and collaborating with Italian research institutions. This year's event spanned two days and featured studies on various topics within IR, RS, and Large Language Models (LLMs). Key focus areas included enhanced retrieval models, personalized information systems, conversational interfaces and user-centric systems, comparative evaluations and metrics, and practical applications in specific fields. IIR 2024 was jointly organized by the University of Udine and the University of Milano-Bicocca and was held in Udine, Italy.

Date: September 5–6, 2024.

Website: <https://iir2024.uniud.it/>.

1 Introduction

The 14th Italian Information Retrieval Workshop (IIR 2024) served as the annual event of the Italian Information Retrieval (IR) and Recommender Systems (RS) communities, as well as international research partners, for 2024. IIR 2024 was jointly organized by the University of Udine and the University of Milano-Bicocca and was held at the Sala Gusmani of Palazzo Antonini (Udine, Italy) in September 2024.

The Italian Information Retrieval Workshop (IIR), now in its 14th edition, was initiated in 2010 to provide a venue for presenting and discussing both theoretical and empirical research in IR. Since its first edition in Padua, IIR has become a key forum for researchers – particularly early-career scholars – as well as industry professionals interested in IR to connect, share ideas, and discuss their work in a collaborative and informal setting. Over the years, the Workshop has opened up to discuss various research areas in addition to search and ranking, including content representation and analysis, personalized search and recommendation, dialog systems and domain-specific applications, usability through studying human factors and interfaces, and the development of evaluation methods for both IR systems and RS. Since its first edition, IIR has been organized annually by leading research groups in the field. It has been hosted at national and international locations: Padua (2010 and 2019), Milan (2011 and 2022), Bari (2012 and 2021), Pisa (2013 and 2023), Rome (2014 and 2018), Cagliari (2015), Venice (2016), and Lugano (2017).

This event report is organized as follows. We begin by outlining the key aspects of IIR 2024 (Section 2), including the aim, topics, organization, and an overview of the received contributions. Next, we briefly describe each contribution presented at the Workshop (Section 3). Finally, we highlight emerging research trends and other relevant insights (Section 4).

2 The Workshop

IIR 2024 aimed to foster and share research on the topic of IR. It brought together researchers associated with Italian and international institutions to share ideas, present findings, and collaborate in a relaxed and engaging environment, focusing on, among other topics, the following:

- *Search and Ranking*: Research on core IR, including scalable IR. Topics such as theoretical models and foundations, retrieval and ranking models (including diversity and aggregated search), Web search (including link analysis, sponsored search, search advertising, adversarial search and spam, and vertical search), cross- and multi-lingual search, and query representation and analysis;
- *Recommendation*: Research focused on Information Filtering (IF) and RS. Topics such as collaborative filtering, content-based filtering, hybrid RS, context-aware recommendation, deep learning in RS, scalability and efficiency, evaluation metrics and methodologies, user interfaces and visualization, privacy, security, and ethics, domain-specific recommendations, social and trust-based recommendations, and explainable and interpretable recommendations;
- *Content Representation and Analysis*: Research focusing on rich content representations and analysis. Topics such as document representation (including multimodal representation), content analysis and information extraction (including readability, sentiment analysis, and opinion mining), and clustering, classification, summarization, and topic modeling;
- *Artificial Intelligence, NLP, Semantics, and Dialog*: Research bridging AI and IR, especially toward deep semantics, and dialog with intelligent agents. Topics such as question answering, conversational systems (including spoken language interfaces, dialog management systems, and intelligent chat systems), semantics and knowledge graphs, and deep learning for IR, embeddings, LLMs, and agents;

-
- *Domain-Specific Applications*: Research addressing domain-specific challenges. Topics such as social search, search in structured data (including email and entity search), multimedia search, search and recommendation for educational, legal, health (including genomics and bioinformatics), and academic domains, as well as other domains like digital libraries, enterprise, news, apps, and archival search;
 - *Human Factors and Interfaces*: Research into user-centric aspects of IR, including user interfaces, behavior modeling, privacy, and interactive systems. Topics such as mining and modeling search activity (including user and task models, click models, log analysis, behavioral analysis, and attention modeling), interactive and personalized search and recommendation, collaborative search, social tagging, crowdsourcing, and information privacy and security;
 - *Evaluation*: Research focused on the measurement and evaluation of IR and RS. Topics such as user-centered evaluation methods (including measures of user experience, performance, engagement, and search task design), test collections and evaluation metrics (including the development of new collections), evaluation of novel information access tasks and systems (such as multi-turn access), statistical methods and reproducibility in IR evaluation, and efficiency and scalability;
 - *Future Directions*: Research with theoretical or empirical contributions on new technical or social aspects of IR, especially in speculative areas or with emerging technologies. Topics such as novel approaches to IR and RS, ethics, economics, and politics, applications of search and recommendation for social good, and IR and RS with new devices (including wearable computing, neuroinformatics, sensors, Internet-of-Things, and vehicles).

In the remainder of this section, we first provide details about people involved in IIR 2024, including the *Organizing Committee* (Section 2.1), the *Program Committee* (Section 2.2), and the two *invited speakers* and their given *talks* (Section 2.3). Subsequently, we briefly summarize the 28 accepted *contributions* to the Workshop (Section 2.4).

2.1 Organizing Committee

IIR 2024 was jointly organized by the University of Udine and the University of Milano-Bicocca. The *Organizing Committee* was therefore composed of members from the two universities, under different roles, as follows:

- **General Chairs**:
 - Eddy Maddalena, *University of Udine, Italy*;
 - Stefano Mizzaro, *University of Udine, Italy*.
- **Program Chairs**:
 - Kevin Roitero, *University of Udine, Italy*;
 - Marco Viviani, *University of Milano-Bicocca, Italy*.
- **Local Organization Chair**: Francesca Da Ros, *University of Udine, Italy*.
- **Best Paper Award Chair**: Gabriella Pasi, *University of Milano-Bicocca, Italy*.¹

¹Due to a low number of long and short paper contributions, the award was not assigned this year.

-
- **Proceeding Chairs:**
 - David La Barbera, *University of Udine, Italy*;
 - Sandip Modha, *University of Milano-Bicocca, Italy*.
 - **Website Chair:** Michael Soprano, *University of Udine, Italy*.

2.2 Program Committee

The *Program Committee* was composed of the following 50 researchers, who are actively working in the areas related to the themes of the Workshop.

- Aris Anagnostopoulos, *University of Rome La Sapienza, Italy*.
- Vito Walter Anelli, *Polytechnic University of Bari, Italy*.
- Liliana Ardissono, *University of Torino, Italy*.
- Maurizio Atzori, *University of Cagliari, Italy*.
- Pierpaolo Basile, *University of Bari Aldo Moro, Italy*.
- Roberto Basili, *University of Rome Tor Vergata, Italy*.
- Giovanni Maria Biancofiore, *Polytechnic University of Bari, Italy*.
- Paolo Boldi, *University of Milano, Italy*.
- Ludovico Boratto, *University of Cagliari, Italy*.
- Davide Buscaldi, *Sorbonne Paris North University, France*.
- Giandomenico Cornacchia, *Polytechnic University of Bari, Italy*.
- Paolo Cremonesi, *Polytechnic University of Milano, Italy*.
- Danilo Croce, *University of Rome Tor Vergata, Italy*.
- Gianluca Demartini, *The University of Queensland, Australia*.
- Marco de Gemmis, *University of Bari Aldo Moro, 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*.
- Guglielmo Faggioli, *University of Padova, Italy*.
- Fabrizio Falchi, *ISTI-CNR, Italy*.
- Paolo Ferragina, *University of Pisa, Italy*.
- Maurizio Ferrari Dacrema, *Polytechnic University of Milano, Italy*.
- Claudio Lucchese, *University of Venice Ca' Foscari, Italy*.
- Alberto Carlo Maria Mancino, *Polytechnic University of Bari, Italy*.
- Stefano Marchesin, *University of Padova, Italy*.
- Massimo Melucci, *University of Padova, Italy*.
- Stefano Mizzaro, *University of Udine, Italy*.
- Alejandro Moreo, *ISTI-CNR, Italy*.
- Cristina Muntean, *ISTI-CNR, Italy*.
- Cataldo Musto, *University of Bari Aldo Moro, Italy*.
- Franco Maria Nardini, *ISTI-CNR, Italy*.

-
- Vincenzo Paparella, *Polytechnic University of Bari, Italy*.
 - Raffaele Perego, *ISTI-CNR, Italy*.
 - Marco Polignano, *University of Bari Aldo Moro, Italy*.
 - Claudio Pomo, *Polytechnic University of Bari, Italy*.
 - Beatrice Portelli, *University of Udine, Italy*.
 - Francesco Ricci, *Free University of Bozen-Bolzano, Italy*.
 - Manuela Sanguinetti, *University of Cagliari, Italy*.
 - Giovanni Semeraro, *University of Bari Aldo Moro, Italy*.
 - Gianmaria Silvello, *University of Padova, Italy*.
 - Fabrizio Silvestri, *University of Rome La Sapienza, Italy*.
 - Damiano Spina, *RMIT University, Australia*.
 - Andrea Tagarelli, *University of Calabria, Italy*.
 - Marko Tkalcic, *University of Primorska, Slovenia*.
 - Nicola Tonello, *University of Pisa, Italy*.
 - Salvatore Trani, *ISTI-CNR, Italy*.
 - Marco Viviani, *University of Milano-Bicocca, Italy*.
 - Markus Zanker, *Free University of Bozen-Bolzano, Italy*.

2.3 Invited Speakers and Talks

IIR 2024 featured two invited talks: one by *Asia J. Biega* (Max Planck Institute for Security and Privacy, Germany) on data protection (Section 2.3.1) and another by *Gabriella Pasi* (University of Milano-Bicocca, Italy) on search personalization (Section 2.3.2). The abstracts of both talks are provided below.

2.3.1 Asia J. Biega – Data Protection in Data-Driven IR Systems

Modern IR systems are characterized by extensive personal data collection, despite increasing societal costs of such practices. To prevent harm, data protection regulations specify several principles for respectfully processing user data, such as purpose limitation, data minimization, or consent. Yet, practical implementations of these principles leave much to be desired. This talk will delve into the computational and human factors that contribute to such lax implementations, and examine potential improvements.

2.3.2 Gabriella Pasi – Search Personalization: Which Information and How?

Personalization of information access has been a hot research topic for quite a long time. In particular, the task of personalized search requires two main underlying sub-tasks: modeling users and their context and leveraging users' models into proper personalization processes. In this talk, the above issues will be critically addressed, with emphasis on neural approaches, which include the issue of defining personal language models.

2.4 Overview of Received Contributions

The Workshop accepted three types of contributions: *full papers* (10 pages), *short papers* (5 pages), and *extended abstracts* (up to 4 pages). Extended abstracts provide overviews of ongoing projects or summaries of previously published work. Three members of the Program Committee reviewed each contribution. In total, IIR 2024 featured 28 contributions, comprising 5 full papers, 1 short paper, and 22 extended abstracts. Among the 28 participants, 5 chose not to include their contributions in the Proceedings, resulting in 23 contributions published in the Workshop’s CEUR-WS Proceedings volume [Maddalena et al., 2024]. The topics covered in the received contributions allowed us to regroup them into five main categories, which represented the five sessions into which the Workshop was organized over two days. In particular:

1. *Enhanced Retrieval Models* (5 contributions): Advanced techniques for improving key aspects of IR, such as fairness in ranking, learned sparse and dense retrieval, and training efficiency of Large Language Models (LLMs) within IR systems;
2. *Personalized Information Systems* (2 contributions): Customization strategies in IR systems, including Community Question Answering (CQA) and RS;
3. *Conversational AI and User-Centric Systems* (3 contributions): Implementation of conversational interfaces and user-centric systems, often powered by LLMs, across applications like IR systems, RS, and healthcare;
4. *Comparative Evaluations and Metrics* (6 contributions): Assessments and metrics in IR, focusing on Retrieval-Augmented Generation (RAG) systems in LLMs, query reformulation, sequential RS, and new evaluation metrics;
5. *Applications* (7 contributions): Practical applications of IR methodologies in e-commerce, fact-checking, the legal industry, RS, privacy evaluation, and digital humanities.

3 Contributions

In this section, we provide a brief description of each contribution within the related session: *Enhanced Retrieval Models* (Section 3.1), *Personalized Information Systems* (Section 3.2), *Conversational AI and User-Centric Systems* (Section 3.3), *Comparative Evaluations and Metrics*, (Section 3.4), and *Applications* (Section 3.5).

3.1 Session 1: Enhanced Retrieval Models

This session dealt with advanced techniques to enhance specific aspects of IR. Specifically, the contributions focused on fairness in ranking, learned sparse and dense retrieval, and the training efficiency of LLMs in IR systems.

One challenge faced by IR systems is generating fair rankings of retrieved documents. [Marcuzzi et al. \[2024\]](#) addressed this issue by incorporating fairness requirements into these algorithms. They proposed a LambdaMART-based method to train fairness-aware ranking models that simultaneously optimize fairness and effectiveness while mitigating bias propagation.

Learned sparse retrieval, which leverages LLM to enhance traditional sparse models by representing documents and queries in a sparse vector space, was another topic explored during the

Workshop. Bruch et al. [2024] proposed an inverted index organization to enable fast and effective approximate retrieval over learned sparse embeddings. In contrast, learned dense representations map queries and documents into continuous vector spaces, where similar inputs are represented by similar vectors, and retrieval is performed using an approximate k -nearest neighbors search.

Faggioli et al. [2024] operated on query representations within a projected subspace of the original dense retrieval space to estimate the importance of each dimension and truncate the representation to include only the most significant dimensions.

Personalized search aims to tailor the retrieval results returned by the model to the specific interests of individual users. Sokli et al. [2024] proposed integrating a LLM architecture based on Mixture-of-Experts (MoE) with dense retrieval models, leveraging the cognitive complexity levels defined in Bloom’s Taxonomy to enhance personalized IR based on the user’s expertise level in educational search.

Other researchers examined specific aspects of LLMs in IR tasks, emphasizing the necessity for extensive datasets and significant computational resources for training and fine-tuning. Pasin et al. [2024] proposed an Instance Selection approach that selects a relevant subset of data points to improve training efficiency leveraging QA, a specialized Quantum Computing paradigm that utilizes quantum annealers to address specific optimization problems.

3.2 Session 2: Personalized Information Systems

This session focused entirely on personalization in IR, with contributions specifically centered around personalized CQA and RS.

CQA is a form of collaborative IR in which users ask questions and receive answers from other members of an online community. Kasela et al. [2024] introduced a curated dataset for developing and evaluating models in this domain, comprising approximately one million questions and two million answers, annotated with features that reflect user interactions.

RS are central to tailoring search results and content to user preferences. Cecere et al. [2024] investigated how integrating semantic-rich embeddings of textual reviews, generated by LLMs, can enhance traditional systems based on content-based and collaborative recommendations.

3.3 Session 3: Conversational AI and User-Centric Systems

This session addressed conversational interfaces and user-centric systems, highlighting the role of conversational Artificial Intelligence, often powered by LLMs, in various applications, including IR systems, RS, and the healthcare sector.

Conversational interfaces enable users to make natural language queries, facilitating the refinement of information needs. While there are several advantages to using conversational interactions in information access systems, particularly for Conversational Search (CS) and Conversational Retrieval (CR), challenges also exist. Di Noia et al. [2024] proposed an integrated approach that combines CS and CR to enhance user satisfaction, building on findings from prior research on joint recommendation and search. The model features a dialog manager, integrated CS and CR engines, and a shared internal knowledge base, with the initial deployment focusing on conversational product search tasks.

LLMs serve as the de facto conversational agents in user interaction contexts due to their ability to understand and generate human-like text. [Di Palma et al. \[2024\]](#) investigated their application in the specific field of RS, examining the performance of two models, ChatGPT 3.5 and ChatGPT-4, under zero-shot conditions. Their research employed a role-playing prompt to re-rank recommendations across three different domains.

In the healthcare sector, conversational agents provide a variety of services, including medication guidance and emergency room assistance. [Lunardi and Coppola \[2024\]](#) introduced an LLM-powered conversational agent that utilizes a RAG framework to generate information based on a Patient Information Leaflet, a document detailing medication usage, dosage, and other relevant information. Overall, these advancements highlight the potential of conversational interfaces to enhance user experiences across various sectors

3.4 Session 4: Comparative Evaluations and Metrics

This session focused on comparative evaluations and metrics in IR, specifically addressing RAG systems in LLMs, query reformulation, sequential RS, and the development of new evaluation metrics.

Typically, RAG approaches use the “instruct” version of LLMs, which undergoes supervised fine-tuning to enhance instruction compliance and align with human preferences through advanced algorithms. [Cuconasu et al. \[2024\]](#) evaluated their performance against base versions, demonstrating that the instructed models consistently outperform their counterparts by an average of 20% in accuracy.

In IR systems, proper query formulation significantly impacts the system’s ability to retrieve relevant results. Query reformulation is often necessary to better align with the user’s information needs, frequently involving a specialized strategy called query expansion. This process adds additional terms to the original query to capture potential interpretations of the user’s intent. [Rizzo et al. \[2024\]](#) examined the performance of five LLMs in query expansion, using ten well-known datasets in few-shot and Chain-of-Thought learning settings. Their findings show that more recent LLMs consistently outperform their predecessors in zero-shot learning scenarios.

To reduce the cost of evaluating IR systems, one effective approach is to rely on simulated user search sessions. Evaluation has shifted from single-query-based metrics to more comprehensive session-based measures, yet there are still opportunities to propose additional metrics. [Zerhoubi and Granitzer \[2024\]](#) introduced Fréchet Distance as a novel metric for measuring the similarity between real and simulated user sessions, drawing on its success in other fields. They further demonstrated its effectiveness and robustness in assessing simulated user interactions across diverse scenarios.

Many state-of-the-art models for search and ranking tasks in IR are often perceived as “black boxes” due to their large number of parameters, which makes it challenging for humans to understand the relationship between input and output. [Albarelli et al. \[2024\]](#) proposed a common theoretical framework for explainability applicable to IR, outlining the general framework before focusing on specific explanation techniques within the field.

Sequential RS models user behavior over time, but they often struggle with perturbations in their training data. Traditionally, performance is assessed using Rank-Biased Overlap (RBO), which inadequately addresses certain issues, particularly those related to finite rankings. To

tackle this limitation, [Betello et al. \[2024\]](#) proposed an extension of RBO called Finite Rank-Biased Overlap (FRBO), which incorporates finite rankings. They demonstrated that removing elements from the end of a sequence can adversely affect performance by up to 60% in terms of Normalized Discounted Cumulative Gain (nDCG).

In RS, efforts have been made to propose metrics that extend beyond merely assessing the system's ability to predict user preferences and recommend relevant items, although accuracy remains a priority. To address multi-objective evaluation, [Paparella et al. \[2024\]](#) leveraged the concept of Pareto optimality from Multi-Objective Optimization and introduced quality indicators that quantitatively assess Pareto frontiers. They underscored the negative consequences of prioritizing accuracy and emphasized the need for multi-objective evaluation.

3.5 Session 5: Applications

This session focused on specific practical applications of IR approaches in the fields of e-commerce, fact-checking, the legal industry, RS, privacy evaluation, and digital humanities.

In e-commerce, the increasing demand for personalized recommendations has led to a continued focus on retrieval models. One primary objective is to leverage product reviews to enhance these recommendations. [Pisani et al. \[2024\]](#) investigated whether LLM embeddings of plain-text reviews can serve as effective inputs for improving traditional review-based recommendation algorithms by adapting their architecture to process these embeddings instead of word-level ones. In the fashion industry, online retail has driven research in fashion compatibility modeling and item retrieval. [Attimonelli et al. \[2024\]](#) proposed a model based on Conditional Generative Adversarial Networks that aims to enhance fashion image retrieval through image-to-image translation.

In fact-checking, crowdsourcing-based approaches have been employed to scale the process and address the spread of misinformation, though this requires careful consideration of several factors. [Roitero et al. \[2024\]](#) summarized two recent studies: one explores the feasibility of using crowdsourcing to build a robust foundation for large-scale fact-checking applications, while the other proposes a hybrid approach combining LLMs and human assessors to combat misinformation on a large scale.

In the legal industry, vast amounts of textual data make navigation laborious for professionals, highlighting the need for efficient summarization tools due to the complexity and domain-specific vocabulary of these texts. One such domain is food risk prevention, where [Rocchietti et al. \[2024\]](#) proposed a method for summarizing regulatory food safety documents using instruction-tuned LLMs.

In RS, the robustness of recommendation models is often measured by their ability to maintain utility under attack. Fairness, from both the consumer and provider perspectives, is a key parameter for evaluating this robustness. [Boratto et al. \[2024\]](#) provided a comprehensive analysis of robustness in fairness, specifically for graph-based RS.

In privacy evaluation, assessing the effectiveness of obfuscation mechanisms remains an open challenge. While formal frameworks like Differential Privacy are established gold standards for privacy assessment, [De Faveri et al. \[2024\]](#) argued that these mechanisms can fall short in practice. They proposed using new measures that account for the similarities between the original and obfuscated texts to better capture the actual privacy achieved.

In digital humanities, IR approaches are instrumental in managing and disseminating library and documentary heritage. [Conte et al. \[2024\]](#) described the MAGIC project, which focuses on digitizing, archiving, and making heritage content accessible and usable for both research and public engagement, with particular emphasis on managing and disseminating cultural resources.

4 Emerging Trends

The 14th Italian Information Retrieval Workshop promoted discussions on IR from multiple viewpoints, facilitating the exchange of ideas and methodologies among researchers, from which current and future research trends have emerged.

The key emerging trend is the increasing integration of LLMs across the IR landscape. Workshop contributions highlighted their application in various IR areas, including learned and sparse dense retrieval [[Bruch et al., 2024](#); [Faggioli et al., 2024](#)], fair ranking generation for retrieved documents [[Marcuzzi et al., 2024](#)], query reformulation [[Rizzo et al., 2024](#)], textual review embedding generation [[Cecere et al., 2024](#)], optimization of training efficiency [[Pasin et al., 2024](#)], and personalized search enhancement [[Sokli et al., 2024](#)]. Gabriella Pasi’s talk (Section 2.3.2), which specifically addressed the personalization of information access, emphasized user models and their role in enhancing personalization, including the integration of LLMs. LLMs play an important role also in conversational interfaces and user-centric systems, where they often act as conversational agents [[Di Palma et al., 2024](#)], frequently in conjunction with RAG-based approaches [[Lunardi and Coppola, 2024](#)], to enhance user experience and engagement [[Di Noia et al., 2024](#)]. Contributions also explored a range of practical applications of LLMs across domains such as e-commerce [[Pisani et al., 2024](#)], the fashion industry [[Attimonelli et al., 2024](#)], RS [[Boratto et al., 2024](#)], digital humanities [[Conte et al., 2024](#)], fact-checking [[Roitero et al., 2024](#)], the legal industry [[Rocchietti et al., 2024](#)], and privacy evaluation [[De Faveri et al., 2024](#)]. Asia J. Biega’s talk (Section 2.3.1) also addressed the themes of privacy and data protection, offering insights into the practical implementation of regulatory frameworks. Regarding the trend of using RAG approaches with LLMs, one contribution has focused on evaluating performance by comparing base models and their “instruct” variants [[Cuconasu et al., 2024](#)].

Another emerging trend is the development of new metrics in both IR systems and RS. Concerning IR, Researchers have introduced a general theoretical framework for the explainability of retrieval [[Albarelli et al., 2024](#)], as well as a new metric designed to reduce evaluation costs for IR systems [[Zerhoubi and Granitzer, 2024](#)]. Additionally, a new curated dataset has been introduced to facilitate the development of novel IR approaches for CQA [[Kasela et al., 2024](#)]. Considering RS, one contribution proposed a Pareto optimality-based metric for multi-objective performance evaluation [[Paparella et al., 2024](#)], while another addressed the challenges of finite rankings in RS [[Betello et al., 2024](#)].

In summary, IIR 2024 has confirmed, once again in its 14th edition, its role as a platform for promoting collaboration and innovation in research related to IR, particularly within the Italian research community. The 15th edition will be organized by the University of Cagliari, Italy, where we expect these and other emerging trends to be addressed and discussed by the scientific community with the usual interest and participation.

Acknowledgments

We acknowledge the financial contributions from the ACM SIGIR Friends Fund Program and the Department of Mathematics, Computer Science, and Physics of the University of Udine.

The 14th Italian Information Retrieval Workshop was also partially supported by the European Union’s NextGenerationEU PNRR M4.C2.1.1 – PRIN 2022 project 20227F2ZN3 MoT—The Measure of Truth: An Evaluation-Centered Machine-Human Hybrid Framework for Assessing Information Truthfulness, CUP G53D23002800006.

References

- Andrea Albarelli, Claudio Lucchese, Matteo Rizzo, and Alberto Veneri. On the Application of a Common Theoretical Explainability Framework in Information Retrieval. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 43–52. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper24.pdf>.
- Matteo Attimonelli, Claudio Pomo, Dietmar Jannach, and Tommaso Di Noia. Toward Effective Fashion Item Compatibility Modelling. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 111–114. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper4.pdf>.
- Filippo Betello, Federico Siciliano, Pushkar Mishra, and Fabrizio Silvestri. Finite Rank-Biased Overlap (FRBO): A New Measure for Stability in Sequential Recommender Systems. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 78–81. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper3.pdf>.
- Ludovico Boratto, Francesco Fabbri, Gianni Fenu, Mirko Marras, and Giacomo Medda. Comprehensive Assessment of Robustness in Fairness of GNN-based Recommender Systems against Attacks. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 62–65. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper1.pdf>.
- Sebastian Bruch, Franco Maria Nardini, Cosimo Rulli, and Rossano Venturini. Seismic: Efficient and Effective Retrieval over Learned Sparse Representation. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 103–106. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper17.pdf>.
- Nicola Cecere, Andrea Pisani, Maurizio Ferrari Dacrema, and Paolo Cremonesi. Leveraging Semantic Embeddings of User Reviews with Off-the-Shelf LLMs for Recommender Systems. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 91–94. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper25.pdf>.
- Stefania Conte, Gennaro Ferrante, Lorenza Laccetti, Andrea Mazzucchi, Yahya Momtaz, and Augusto Tortora. Content Representation and Analysis: the Magic Project and the Illuminated Dante Project Integrated Systems for Multimedia Information Retrieval. In *Proc. of the 14th*

-
- Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 66–69. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper15.pdf>.
- Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Mareek, Nicola Tonellotto, and Fabrizio Silvestri. Rethinking Relevance: How Noise and Distractors Impact Retrieval-Augmented Generation. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 95–98. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper23.pdf>.
- Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro. Content Representation and Analysis: the Magic Project and the Illuminated Dante Project Integrated Systems for Multimedia Information Retrieval. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 53–57. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper5.pdf>.
- Tommaso Di Noia, Guglielmo Faggioli, Marco Ferrante, Nicola Ferro, Fedelucio Narducci, Raffaele Perego, and Giuseppe Santucci. The CAMEO Project: A Holistic View for Conversational Agents. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 107–110. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper11.pdf>.
- Dario Di Palma, Giovanni Servedio, Vito Walter Anelli, Giovanni Maria Biancofiore, Fedelucio Narducci, Leonarda Carnimeo, and Tommaso Di Noia. Beyond Words: Can ChatGPT Support State-of-the-Art Recommender Systems? In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 13–22. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper6.pdf>.
- Guglielmo Faggioli, Nicola Ferro, Raffaele Perego, and Nicola Tonellotto. Turning on a DIME: Estimating Dimension Importance for Dense Information Retrieval. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 115–118. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper12.pdf>.
- Pranav Kasela, Marco Braga, Gabriella Pasi, and Raffaele Perego. SE-PQA: StackExchange Personalized Community Question Answering. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 103–106. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper9.pdf>.
- Riccardo Lunardi and Paolo Coppola. Conversational-Agent for Patient Information Leaflet. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 70–73. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper14.pdf>.
- Eddy Maddalena, Stefano Mizzaro, Kevin Roitero, and Marco Viviani, editors. *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, 10 2024. CEUR-WS.org. URL <https://ceur-ws.org/Vol-3802/>.

-
- Federico Marcuzzi, Claudio Lucchese, and Salvatore Orlando. LambdaFair: a Fair and Effective LambdaMART. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 87–90. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper2.pdf>.
- Vincenzo Paparella, Dario Di Palma, Vito Walter Anelli, Alessandro De Bellis, and Tommaso Di Noia. Unveiling the Potential of Recommender Systems through Multi-Objective Metrics. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 119–122. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper20.pdf>.
- Andrea Pasin, Cunha Washington, Gonçalves Marcos André, and Nicola Ferro. A Quantum Annealing-Based Instance Selection Approach for Transformer Fine-Tuning. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 58–61. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper10.pdf>.
- Andrea Pisani, Nicola Cecere, Maurizio Ferrari Dacrema, and Paolo Cremonesi. Pre-Trained LLM Embeddings of Product Reviews for Recommendation. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 91–94. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper8.pdf>.
- Daniele Rizzo, Alessandro Raganato, and Marco Viviani. Comparatively Assessing Large Language Models for Query Expansion in Information Retrieval via Zero-Shot and Chain-of-Thought Prompting. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 23–32. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper22.pdf>.
- Guido Rocchietti, Cosimo Rulli, Korbinian Randl, Cristina Ioana Muntean, Franco Maria Nardini, Raffaele Perego, Salvatore Trani, Manos Karvounis, and Jakub Janostik. LongDoc Summarization using Instruction-tuned Large Language Models for Food Safety Regulations. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 33–42. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper27.pdf>.
- Kevin Roitero, Michael Soprano, David La Barbera, Eddy Maddalena, and Stefano Mizzaro. Enhancing Fact-Checking: From Crowdsourced Validation to Integration with Large Language Models. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 74–77. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper13.pdf>.
- Effrosyni Sokli, Alessandro Raganato, and Gabriella Pasi. Incorporating Cognitive Complexity of Text in Dense Retrieval for Personalized Search. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 82–85. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper7.pdf>.
- Saber Zerhoudi and Michael Granitzer. Beyond Conventional Metrics: Assessing User Simulators in Information Retrieval. In *Proc. of the 14th Italian Information Retrieval Workshop*, volume 3802 of *CEUR Workshop Proceedings*, pages 3–12. CEUR-WS.org, 10 2024. URL <https://ceur-ws.org/Vol-3802/paper28.pdf>.