Preface to the Fifteenth Italian Information Retrieval Workshop (IIR 2025)

Ludovico Boratto[1,*], Allegra De Filippo[2,†], Stefano Marchesin[3,†] and Mirko Marras[1,*]

[1]University of Cagliari, Cagliari, Italy
[2]University of Bologna, Bologna, Italy
[3]University of Padua, Padua, Italy

About IIR 2025

The volume contains the papers presented at the 15[th] Italian Information Retrieval Workshop (IIR 2025), held in Cagliari, Italy, from 3rd to 5th September 2025. This workshop continues the tradition of serving as a forum for researchers, academics, and industry professionals interested in the field of Information Retrieval (IR). The primary objective of IIR is to provide a platform for stimulating and disseminating research in IR, especially encouraging young researchers to present and discuss their work in an informal environment.

Over the past years, IIR has established itself as a primary gathering point for the Italian IR community, while also embracing contributions and participation from international researchers. Past IIR workshops took place in Udine (2024), Pisa (2023), 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). This year's workshop featured a diverse range of topics spanning theoretical models, empirical studies, and practical applications in IR, highlighting the multidisciplinary nature and the evolving landscape of the field.

Program

The workshop's program featured 1 full paper, 2 short papers, and 24 extended abstracts. The authors of two contributions decided not to have their papers included in the proceedings. The event spanned three days and the presentations were organized into six thematic sessions:

- *Session 1 – Large Language Models and IR*
    - Marco Valentini, Antonio Ferrara, Tommaso Di Noia, Giuseppe Illuzzi, Pierangelo Colacicco. *Leveraging LLM-Powered Multi-Agent Systems to Enhance Customer Experience in Complex Product Domains.*
    - Francesco De Faveri, Guglielmo Faggioli, Nicola Ferro. *Comparing LLMs and Traditional Privacy Measures to Evaluate Query Obfuscation Approaches.*
    - Giuseppe Spillo, Cataldo Musto, Matteo Mannavola, Marco de Gemmis, Pasquale Lops, Giovanni Semeraro. *Graph Augmentation with LLMs for Knowledge-Aware Recommender System.*

- – Emiliana Murgia, Monica Landoni, Theo Huibers, Maria Soledad Pera. *Empowered or Lost? How Search Tools and Social Networks Shape Teen Learning in an AI World.*
- – Alessandro Petruzzelli, Cataldo Musto, Marco de Gemmis, Pasquale Lops, Giovanni Semeraro. *Knowledge is Power: Boosting Recommender Systems by Infusing LLMs with Domain Expertise.*

- *Session 2 – Recommender Systems: Models and Evaluation*
  - – Federico Siciliano, Antonio Purificato, Filippo Betello, Nicola Tonellotto, Fabrizio Silvestri. *Caser+ and CosRec+: Closing the Gap Between CNNs and Attention Models in SRS.*
  - – Michael Benigni, Maurizio Ferrari Dacrema, Dietmar Jannach. *Diffusion Models for Recommendation: Reproducibility and Conceptual Mismatch.*
  - – Federico Siciliano, Giulia Di Teodoro, Nicola Tonellotto, Fabrizio Silvestri. *Comparing Recommendation Losses under Negative Sampling.*
  - – Neda Afreen, Ludovico Boratto, Gianni Fenu, Mirko Marras, Alessandro Soccol. *Effective and Transparent Course Recommendation through Causal Reasoning with Language Models.*

- *Session 3 – Conversational and Personalized Systems*
  - – Alessandro Francesco Maria Martina, Alessandro Petruzzelli, Cataldo Musto, Marco de Gemmis, Pasquale Lops, Giovanni Semeraro. *Beyond Homogeneous Users: Simulating Diverse User Personas for Conversational Recommendation Datasets.*
  - – Pranav Kasela, Gabriella Pasi, Raffaele Perego. *PARK: Personalized Academic Retrieval with Knowledge-graphs.*
  - – Marco Alessio, Simone Merlo, Tommaso Di Noia, Guglielmo Faggioli, Marco Ferrante, Nicola Ferro, Cristina Ioana Muntean, Franco Maria Nardini, Fedelucio Narducci, Raffaele Perego, Giuseppe Santucci, Nicola Viterbo. *A Dataset for Joint Conversational Search and Recommendation.*
  - – Cristina Ioana Muntean, Franco Maria Nardini, Raffaele Perego, Guido Rocchietti, Cosimo Rulli. *Efficient Conversational Search via Topical Locality in Dense Retrieval.*

- *Session 4 – Neural IR, Federated Models, and Search Applications*
  - – Georgeta Bordea, Stéphane Campinas, Matteo Catena, Renaud Delbru. *Siren Federate: Bridging Document, Relational, and Graph Models for Exploratory Graph Analysis.*
  - – Francesca Pezzuti, Sean MacAvaney, Nicola Tonellotto. *Neural Web Crawling.*
  - – Marco Braga, Pranav Kasela, Alessandro Raganato, Gabriella Pasi. *Investigating Task Arithmetic for Zero-Shot Information Retrieval.*
  - – Guglielmo Faggioli, Nicola Ferro, Raffaele Perego, Nicola Tonellotto. *Counterfactual Dimension Importance Estimation (CoDIME) for Dense Information Retrieval.*
  - – Giulio D'Erasmo, Giovanni Trappolini, Nicola Tonellotto, Fabrizio Silvestri. *Eclipse: Leveraging Pseudo-Irrelevance for Dimension Importance Estimation.*

- *Session 5 – Evaluation, Metrics, and Tools*
  - – Francesco De Faveri, Guglielmo Faggioli, Nicola Ferro, Kalervo Järvelin. *MDCU: A Multi-Dimensional Cumulative Utility Metric for Information Retrieval.*
  - – Silvio Martinico, Leonardo Delfino, Domenico Erriquez, Franco Maria Nardini, Rossano Venturini, Cosimo Rulli. *kANNolo: Facilitating Research in Approximate Nearest Neighbors Search.*
  - – Andrea Pasin, Nicola Ferro. *The KIMERA Infrastructure: Shifting from Evaluation-as-a-Service to Evaluation-in-the-Cloud.*
  - – Giuseppe Scarpi, Amir Khorrami Chokami, Diego Reforgiato. *A Grading Methodology for E-commerce Reviews based on a Mixed Lexical-semantic Approach.*
  - – Ornella Irrera, Stefano Marchesin, Gianmaria Silvello. *Enhancing Ground Truth Creation for IR Through Web-Based Collaborative Annotation.*

- *Session 6 – User Modeling, Evaluation and Fairness*
  - Angelo Geninatti Cossatin. *Lean and Green: Computational Efficiency and Nudging for Sustainable Fashion Recommendation.*
  - Chen Amiraz, Florin Cuconasu, Simone Filice, Zohar Karnin. *Beyond Relevance: Quantifying Distraction of Irrelevant Passages in RAG.*
  - Salvatore Bufi, Vincenzo Paparella, Vito Walter Anelli, Tommaso Di Noia. *Less Data, More Questions: Fairness and Accuracy Under Data Minimization in Recommender Systems.*
  - Effrosyni Sokli, Pranav Kasela, Georgios Peikos, Gabriella Pasi. *Investigating Mixture of Experts in Dense Retrieval.*

The program also included two invited talks:

- *"This Will All Make Sense When I'm Older": Goals, Gaps, and the Road Ahead for Child-Aware IR*
  Maria Soledad (Sole) Pera – Delft University of Technology, Delft, The Netherlands.
- *Learned Sparse Representations of Text: A Story of Indexing and Retrieval*
  Franco Maria Nardini – Istituto di Scienza e Tecnologie dell'Informazione "Alessandro Faedo" del Consiglio Nazionale delle Ricerche (ISTI-CNR), Pisa, Italy.

Acknowledgments