Responsible AI From the Lens of an Information Retrieval Researcher: A Hands-On Tutorial

Damiano Spina
RMIT University
Melbourne, Australia
damiano.spina@rmit.edu.au

Abstract

While the concept of responsible AI is becoming more and more popular, practitioners and researchers may often struggle to characterize responsible practices in their own work. Using a hands-on approach – where participants are invited to discuss terminology from their own perspectives – this tutorial aims to illustrate the application of responsible practices in Information Retrieval (IR). Using case studies based on existing IR research, the tutorial will explore responsible AI concepts such as positionality, participatory research, fairness, diversity, and ethics.

CCS Concepts

• **Information systems → Users and interactive retrieval**.

Keywords

positionality, fairness, diversity, ethical considerations, multidisciplinary research

ACM Reference Format:
Damiano Spina. 2025. Responsible AI From the Lens of an Information Retrieval Researcher: A Hands-On Tutorial. In *2025 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '25), March 24–28, 2025, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3698204.3716479

1 Proposed Title and Format

The title of the proposed tutorial is:

> Responsible AI from the Lens of an Information Retrieval Researcher: A Hands-On Tutorial

I am proposing a half-day tutorial.

2 Motivation

The rapid advancement of Artificial Intelligence (AI) has brought immense benefits across numerous domains, but it is also well-known that it raises important challenges in terms of ethical and responsible practices in the development, application, and use of these technologies.

The concept of responsible practices in conducting Information Retrieval research (IR) is not new [2, 5, 9]. However, given the multidisciplinary nature of the challenges faced in responsible AI and IR (e.g., defining and measuring fairness in a ranking), it is often not straightforward to advance knowledge in this area. In the last few years, I had the opportunity to work in different multidisciplinary teams – including, in addition to computer scientists, experts with a background in media and communication, fact-checking, psychology, and law – to address problems related to responsible AI in the context of information access systems – in particular, information retrieval and conversational information-seeking systems. In this tutorial, I will share my experience on these specific areas, discussing the problems faced and the solutions adopted.

I propose to start the tutorial by facilitating a reflective/positioning activity where participants can discuss different keywords related to responsible AI and information access (e.g., "fairness", "bias", "diversity", "search engine", "recommender system", "fact-checking", "fake news", "echo chambers", "filter bubbles"). Then, through a blend of existing findings and emerging avenues, we will characterize the benefits (the "good"), challenges ("the bad"), and the opportunities (the "unknown") that we faced while advancing knowledge in specific information access tasks including presentation strategies for fact-checked content, fairness-aware rankings via search results diversification, characterizing information processing activities via physiological signals, among others.

A longer version of this proposed tutorial has been delivered as a 10-hour course to more than 20 students enrolled in the AI PhD program at the University of Udine (Italy) [17]. Anonymous student feedback from a survey with 14 responses reflects high satisfaction (8.64/10 average rating) and positive comments, e.g., "I totally enjoyed the group activities. They were a good way to interact with other students coming from different fields/sub-fields, so the chats we had were sometimes strange yet overall engaging".

3 Syllabus

The tutorial is organized in two parts (Table 1). In the first part of the tutorial, we will discuss the concept of *positionality*. Participants will be invited to conduct a group activity to discuss the terminology associated with responsible IR. This activity was inspired by a workshop developed by Prof. Anthony McCosker [3]. It draws on Raymond Williams' Keywords project [20], which has helped to explore and socialize the language associated with new technologies.

In the second part, rather than providing definitions of concepts associated with responsible IR, the tutorial will follow a hands-on approach, where concepts such as fairness, diversity, participatory research, and ethics are instantiated in case studies based on existing research, including:

- Evaluating fairness and diversity in rankings [1, 15]
- Designing and evaluating presentation strategies for fact-checking reports [6, 13, 16, 18]

CHIIR '25, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-1290-6/25/03
https://doi.org/10.1145/3698204.3716479

Table 1: Tentative schedule of the half-day tutorial.

Time	Event
09:00–09:15	Acknowledgment of Country & Opening
09:15–10:45	**Part I**: Positionality and Group Activity (*keywords*)
10:45–11:00	Morning tea break
11:15–12:00	**Part II**: Case Studies (fairness & diversity, participatory research, and ethical considerations)
12:15–12:30	Wrap up & Closing

- Quantifying and measuring bias and engagement [10–12, 16]
- Characterizing information seeking processes using neurophysiological signals [7, 8, 14]
- ethical considerations of search engines in the era of generative AI [4, 19]

The report by Spina et al. [17] discusses in detail the content and activities that will be covered in the tutorial. The webpage of the tutorial contains additional information: https://www.damianospina.com/tutorial/responsible-ai/.

4 Tutorial Outcomes

The tutorial is designed for practitioners and researchers interested in knowing more about how to apply responsible practices in their work. It is potentially relevant to all conference attendees, including students, early-career, and experienced researchers.

The learning outcomes of this half-day tutorial are two-fold:

(1) a better understanding of the concept of positionality, and the acknowledgment of differences in terminology and methods when working in multidisciplinary teams; and
(2) the application of responsible AI concepts in the context of interactive information retrieval research.

5 Presenter

Dr. Damiano Spina is a Senior Lecturer at the School of Computing Technologies, RMIT University, Australia, an Associate Investigator at the ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S), and RMIT Research Lead at the Australian Internet Observatory (AIO).

His research focuses on interactive information retrieval and evaluation of information access systems. Dr. Spina completed his PhD in Computer Science in 2014 (UNED, Spain). He has published more than 70 peer-reviewed scientific publications in the field of information retrieval. He serves as editorial board member for IP&M and ACM TOIS, and he is an active Program Committee member of various IR conferences. He is the recipient of an ARC Discovery Early Career Researcher Award (DECRA) and the 2021 RMIT Award for Research Impact (Technology).

Dr. Spina is an active mentor and PhD supervisor, including co-supervision in multidisciplinary teams with scholars and students from media and communication, psychology, and design. He has co-organized workshops in international conferences and shared tasks for evaluation campaigns at CLEF and IberLEF.
E-mail: damiano.spina@rmit.edu.au
Website: https://www.damianospina.com

Acknowledgments

The author acknowledges that this tutorial proposal was designed on the unceded lands of the Woi Wurrung and Boon Wurrung language groups of the eastern Kulin Nation. The author acknowledges their Ancestors and Elders, past, and present, and their connection to Country.

This work is supported by the Australian Research Council by means of a DECRA (DE200100064) and the ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S, CE200100005).

References

[1] Enrique Amigó, Damiano Spina, and Jorge Carrillo-de Albornoz. 2018. An Axiomatic Analysis of Diversity Evaluation Metrics: Introducing the Rank-Biased Utility Metric. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) *(SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 625–634. https://doi.org/10.1145/3209978.3210024

[2] Nolwenn Bernard and Krisztian Balog. 2023. A Systematic Review of Fairness, Accountability, Transparency and Ethics in Information Retrieval. *ACM Comput. Surv.* (Dec. 2023). doi:10.1145/3637211

[3] Jean Burgess, Anthony McCosker, Kath Albury, and Rowan Wilken. 2022. *Everyday Data Cultures*. Wiley.

[4] Simon Coghlan, Hui Xian Chia, Falk Scholer, and Damiano Spina. 2025. Control Search Rankings, Control the World: What is a Good Search Engine? *AI and Ethics* (2025), 31 pages. https://www.damianospina.com/publication/coghlan-2025-control/coghlan-2025-control

[5] Michael D. Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2022. Fairness in Information Access Systems. *Found. Trends Inf. Retr.* 16, 1–2 (July 2022), 1–177. doi:10.1561/1500000079

[6] Danula Hettiachchi, Kaixin Ji, Jenny Kennedy, Anthony McCosker, Flora D. Salim, Mark Sanderson, Falk Scholer, and Damiano Spina. 2023. Designing and Evaluating Presentation Strategies for Fact-Checked Content. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management* (Birmingham, United Kingdom) *(CIKM '23)*. Association for Computing Machinery, New York, NY, USA, 751–761. https://doi.org/10.1145/3583780.3614841

[7] Kaixin Ji, Danula Hettiachchi, Flora D. Salim, Falk Scholer, and Damiano Spina. 2024. Characterizing Information Seeking Processes with Multiple Physiological Signals. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) *(SIGIR '24)*. Association for Computing Machinery, New York, NY, USA, 1006–1017. doi:10.1145/3626772.3657793

[8] Kaixin Ji, Sachin Pathiyan Cherumanal, Johanne R. Trippas, Danula Hettiachchi, Flora D. Salim, Falk Scholer, and Damiano Spina. 2024. Towards Detecting and Mitigating Cognitive Bias in Spoken Conversational Search. In *Adjunct Proceedings of the 26th International Conference on Mobile Human-Computer Interaction* (Melbourne, VIC, Australia) *(MobileHCI '24 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 10, 10 pages. doi:10.1145/3640471.3680245

[9] Diane Kelly. 2009. Methods for Evaluating Interactive Information Retrieval Systems with Users. *Found. Trends Inf. Retr.* 3, 1—2 (Jan. 2009), 1–224. doi:10.1561/1500000012

[10] Johannes Kiesel, Damiano Spina, Henning Wachsmuth, and Benno Stein. 2021. The Meant, the Said, and the Understood: Conversational Argument Search and Cognitive Biases. In *CUI 2021 - 3rd Conference on Conversational User Interfaces* (Bilbao (online), Spain) *(CUI '21)*. Association for Computing Machinery, New York, NY, USA, Article 20, 5 pages. https://doi.org/10.1145/3469595.3469615

[11] Weronika Łajewska, Krisztian Balog, Damiano Spina, and Johanne Trippas. 2024. Can Users Detect Biases or Factual Errors in Generated Responses in Conversational Information-Seeking?. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region* (Tokyo, Japan) *(SIGIR-AP 2024)*. Association for Computing Machinery, New York, NY, USA, 92–102. doi:10.1145/3673791.3698409

[12] Weronika Łajewska, Damiano Spina, Johanne Trippas, and Krisztian Balog. 2024. Explainability for Transparent Conversational Information-Seeking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) *(SIGIR '24)*. Association for Computing Machinery, New York, NY, USA, 1040–1050. doi:10.1145/3626772.3657768

[13] Sachin Pathiyan Cherumanal, Ujwal Gadiraju, and Damiano Spina. 2024. Everything We Hear: Towards Tackling Misinformation in Podcasts. In *Proceedings of the 26th International Conference on Multimodal Interaction* (San Jose, Costa Rica) *(ICMI '24)*. Association for Computing Machinery, New York, NY, USA, 596–601. doi:10.1145/3678957.3678959

[14] Sachin Pathiyan Cherumanal, Falk Scholer, Johanne R Trippas, and Damiano Spina. 2024. Towards Investigating Biases in Spoken Conversational Search. In *Companion Proceedings of the 26th International Conference on Multimodal Interaction* (San Jose, Costa Rica) *(ICMI Companion '24)*. Association for Computing Machinery, New York, NY, USA, 61–66. doi:10.1145/3686215.3690156

[15] Sachin Pathiyan Cherumanal, Damiano Spina, Falk Scholer, and W. Bruce Croft. 2021. Evaluating Fairness in Argument Retrieval. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (Virtual Event, Queensland, Australia) *(CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 3363–3367. https://doi.org/10.1145/3459637.3482099

[16] Damiano Spina, Danula Hettiachchi, and Anthony McCosker. 2024. *Quantifying and Measuring Bias and Engagement in Automated Decision-Making*. Technical Report. ARC Centre of Excellence for Automated Decision-Making and Society, RMIT University, Melbourne, Australia. doi:10.60836/k8xh-en92

[17] Damiano Spina, Kevin Roitero, Stefano Mizzaro, Vincenzo Della Mea, Francesca Da Ros, Michael Soprano, Hafsa Akebli, Alex Falcon, Mehdi Fasihi, Alessio Fiorin, David La Barbera, Daniele Lizzio Bosco, Riccardo Lunardi, Alberto Marturano, Zaka-Ud-Din Muhammad, Francesco Nascimben, Moritz Nottebaum, Massimiliano Pascoli, Mihai Horia Popescu, Laura Rasotto, Mubashara Rehman, Francesco Taverna, Biagio Tomasetig, and Alessandro Tremamunno. 2024. Report on the Hands-On PhD Course on Responsible AI from the Lens of an Information Access Researcher. *SIGIR Forum* 58, 2 (2024), 61 pages. https://sigir.org/wp-content/uploads/2025/01/p07.pdf

[18] Damiano Spina, Mark Sanderson, Daniel Angus, Gianluca Demartini, Dana Mckay, Lauren L. Saling, and Ryen W. White. 2023. Human-AI Cooperation to Tackle Misinformation and Polarization. *Commun. ACM* 66, 7 (June 2023), 40–45. doi:10.1145/3588431

[19] Johanne R. Trippas, Damiano Spina, and Falk Scholer. 2025. Adapting Generative Information Retrieval Systems to Users, Tasks, and Scenarios. In *Information Access in the Era of Generative AI*, Ryen W. White and Chirag Shah (Eds.). The Information Retrieval Series, Vol. 51. Springer Nature Switzerland AG, Cham, Switzerland. doi:10.1007/978-3-031-73147-1_4

[20] Raymond Williams. 2014. *Keywords: A vocabulary of culture and society*. Oxford University Press.