

Report on the Hands-On PhD Course on Responsible AI from the Lens of an Information Access Researcher

Damiano Spina

RMIT University
Australia

`damiano.spina@rmit.edu.au`

Kevin Roitero

University of Udine
Italy

`kevin.roitero@uniud.it`

Stefano Mizzaro

University of Udine
Italy

`stefano.mizzaro@uniud.it`

Vincenzo Della Mea

University of Udine
Italy

`vincenzo.dellamea@uniud.it`

Francesca Da Ros

University of Udine
Italy

`francesca.daros@uniud.it`

Michael Soprano

University of Udine
Italy

`michael.soprano@uniud.it`

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, Alessandro Tremamunno*

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. This paper presents a four-day, PhD-level course on Responsible Artificial Intelligence conducted at the University of Udine by Dr. Damiano Spina. Using a hands-on approach, the course aimed to illustrate the application of responsible AI concepts in research. Using case studies based on existing IR research, the course explored responsible AI concepts such as positionality, participatory research, fairness, diversity, and ethics. The course engaged 23 participants, both online and in person, including PhD students at various stages, postdoctoral researchers, professors, and academic staff. It featured four sessions and five interactive group activities. Of the 23 attendees, 20 (87%) actively participated in the activities, and 14 (61%) completed the final survey. We believe the hands-on activities discussed in this paper can assist practitioners and educators in the design of responsible AI content for information retrieval curriculum.

Date: May 21–24, 2024.

Website: <https://www.damianospina.com/teaching/responsible-ai>.

*Affiliation not shown for all authors due to space limitations (see Appendix F for details).

1 Introduction

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. Given the multidisciplinary nature of these challenges, it is often not straightforward to advance knowledge in this area.

The first author has designed the course described in this report to share with PhD students enrolled in the Computer Science and AI PhD program at the University of Udine (Italy).¹

The material of the course is based the experience of working in multi-disciplinary teams within the ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S).² The collaborations included working with, 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 and retrieval systems – in particular, search engines and conversational assistants.

The course comprised four in-person sessions featuring presentations, along with five interactive group activities. The enthusiastic participation of students and the positive feedback received indicate that students appreciated this more interactive approach to knowledge sharing. The course learning outcomes included an enhanced understanding of the terminology and methods employed by multidisciplinary teams dedicated to advancing knowledge in responsible AI, as well as the application of these concepts in the context of information access and retrieval research.

All supplementary materials – including data and code used to analyze the post-course feedback survey – are available to the research community in an OSF repository.³ A shorter version of this course has been accepted as a half-day tutorial at the 2025 ACM SIGIR CHIIR conference [Spina, 2025].⁴

The purpose of this report is to detail the activities conducted during the course and to share the outcomes with the Information Retrieval (IR) community. We believe – including the students who attended the course – that: (i) the activities designed for this course may help other educators designing curricula for information access and retrieval [Bauer et al., 2023], and (ii) the material delivered in this course can assist others in integrating responsible AI concepts into their research.

Section 2 outlines the topics and activities provided within the course. Section 3 describes the outcomes of the first group activity (research topics). Section 4 describes the outcomes of the second group activity (keywords); Section 5 discusses the case studies crafted by the groups (Activities 3–5). Section 6.4 summarizes the results of a post-course feedback survey and we conclude the report in Section 7.

¹<https://dmif.uniud.it/it/didattica/dottorato/iai/phd-course-in-computer-science-and-artificial-intelligence>

²<https://admsctentre.org.au>

³https://bit.ly/responsibleIR_course

⁴<https://www.damianospina.com/tutorial/responsible-ai>

2 Course Outline

The course took place at the University of Udine (Italy) over four days, from May 21st to May 24th, 2024. It included four sessions (Section 2.1) and a series of interactive group activities (Section 2.2).

2.1 Sessions

The course consisted of four sessions, one session per day:

- **Day 1 – Responsible AI: A Multidisciplinary Problem. How Can We Contribute?:** The first day provided an introduction to two phenomena that are crucial – but not necessarily obvious – to consider when carrying out multidisciplinary research: understanding who is involved in the discussions and the differences in understanding (common) terminology. The session introduced the concept of *positionality* [Olteanu et al., 2023], by inviting participants to reflect on factors (socioeconomic, discipline background, level of experience, etc.) that may have shaped the identity of the researchers who were going to discuss Responsible AI (Activity 1). After the group discussion, participants were invited to identify and embrace differences in the understanding of key terminology related to responsible AI (Activity 2). This activity was inspired by a workshop developed by Prof. Anthony McCosker [Burgess et al., 2022]. It draws on Raymond Williams’ Keywords project [Williams, 2014], which has helped explore and socialize the language associated with new technologies.
- **Day 2: Mixed Methods for Designing and Evaluating Presentation Strategies for Fact-checked Content:** The second session introduced the concept of mixed methods and participatory research. These concepts were illustrated with the work carried out in the ADM+S project titled “Quantifying and Measuring Bias and Engagement” [Spina et al., 2024]. In this project, a multidisciplinary team collaborated with fact-checkers to design and evaluate presentation strategies of verified information in screens and voice-only channels [Hettiachchi et al., 2023; Spina et al., 2023]. In this session, participants were asked to work in groups to identify research projects of their interest (case studies) on which they think it would be beneficial to apply mixed methods (Activity 3).
- **Day 3 – Fairness and Diversity: Two Sides of the Same Coin?:** The third session introduced the notion of fairness and diversity. In particular, it gave an overview of the Cranfield paradigm for offline evaluation of information retrieval systems [Mizzaro, 1997; Sanderson, 2010], and introduced axiomatic analysis of effectiveness measures [Amigó et al., 2018a, 2022], as well as diversity-aware evaluation [Amigó et al., 2018b] and fairness-aware rankings [Ekstrand et al., 2022; Hiemstra, 2023; Kiesel et al., 2021; Pathiyan Cherumanal et al., 2021, 2023; Pathiyan Cherumnal et al., 2023]. Finally, groups were asked to discuss fairness and diversity in the context of their case studies (Activity 4).
- **Day 4 – Characterizing Information Processing Activities with Physiological Signals from Multiple Wearable Devices:** The last session aimed to introduce the concept of ethical considerations by describing recent work in conducting lab user studies to characterize information processing activities with neurophysiological signals obtained via EEG (electroencephalogram), wristbands, and eye-tracking [Ji, 2023; Ji et al., 2023a,b, 2024a]. Then, the use of wearable devices to characterize cognitive biases in spoken conversational

search was also discussed [Kiesel et al., 2021; Ji et al., 2024b; Pathiyan Cherumanal et al., 2024]. Finally, groups were asked to revisit their case studies to discuss the ethical considerations that should be taken into account in their proposed projects (Activity 5).

2.2 Activities

Throughout the four days, participants were asked to reflect and apply the topics discussed in the sessions to their research through a series of interactive activities. To facilitate communication and ideas exchange, participants were divided into four groups, each symbolized by a typical Australian animal (Marram/Kangaroo 🦘, Warin/Wombat 🦡, Gurrborra/Koala 🦙, and Dulai Wurrung/Platypus 🦇⁵).

A total of five activities were proposed (their handouts are reported in Appendix A):

- **Activity 1 – Research Topics:** Each participant was asked to individually reflect upon their research activity, first identifying its key elements in the form of an abstract, and then considering how their identity as a scholar influences their research (positionality statement). After these reflections, each participant shared and discussed their experiences with the other group members, eventually drafting a positionality statement. The activity was conducted within Day 1. The results of this activity are reported in Section 3.
- **Activity 2 – Keywords:** Each group was asked to collect and cluster keywords and terms that might be highly important for understanding their research matters. Among these keywords, the group should identify those that might be divisive or cause misunderstandings (e.g., differences between technical and ordinary use of a word). Finally, the group should draft activities, tools, or materials for facilitating dialogue on these terms. The activity was conducted within Day 1. The results of this activity are reported in Section 4.
- **Activity 3 – Mixed Methods:** Each group was asked to discuss how mixed methods could be helpful in the research activities conducted by the participants. The group should draft possible activities and guess the outcomes (i.e., prepare a series of case studies). The activity was conducted within Day 2. The results of this activity are reported in Section 5 along with those of the next activities.
- **Activity 4 – Fairness and Diversity:** Considering the case studies outlined in Activity 3, each group was called to reflect upon the concepts of fairness and diversity. The activity was conducted within Day 3. The results of this activity are reported in Section 5 along with those of the previous and the next activity.
- **Activity 5 – Ethical Considerations:** Considering the case studies outlined in Activity 3, each group was called to reflect upon the ethical consequences of the research. Specifically, each group was asked to identify the beneficiaries of the research and conduct a risk assessment. The activity was conducted within Day 4. The results of this activity are reported in Section 5 along with those of the previous two activities.

⁵Note that the first name is in Woi Wurrung and Boon Wurrung languages: https://deadlystory.com/page/aboriginal-country-map/Aboriginal_Country_Completed/Wurundjeri/Wurundjeri_Language

3 Activity 1: Research Topics

This section summarizes the results of Activity 1. A handout for the activity is provided in Appendix A.1, and the content created by each group is included in Appendix B.

3.1 Objective

The activity aimed to guide participants through a reflective and collaborative process, starting with the individual development of a short research abstract, followed by an assessment of their scholarly positionality and its impact on their research. The activity concluded with a group discussion to share and exchange perspectives.

3.2 Procedure

The procedure involved having participants individually outline their research as a short abstract. Each participant then assessed their scholarly position and reflected on how it could influence their research. Finally, participants shared and discussed their perspectives within their groups.

The activity was conducted simultaneously by all course participants. Each participant had 10-15 minutes to report on their research and positionality statement using a shared template. They were also able to view contributions from others. Following this, group discussions were held, with the group leader summarizing the discussion in the shared document. Despite the limited time, all participants successfully outlined their topics and positions. After the course, we reorganized and proofread the content to improve its clarity and accuracy.

3.3 Outcomes

Figure 1 shows a word cloud containing all the research topics of the course participants. Many of these topics involve artificial intelligence to varying degrees. Some participants focus on applying artificial intelligence techniques—primarily machine learning—in healthcare, such as in oncology, (Appendix B.1.1), digital pathology (Appendix B.4.5 and Appendix B.1.3), and immunohistochemistry (Appendix B.4.4 and Appendix B.3.4) while others focuses on forestry and agriculture (Appendix B.2.4) or the environment in general (Appendix B.3.2). Regarding healthcare in particular, one research topic focuses on the development of data structures for the efficient storage of genetic and biological data (Appendix B.4.2).

Other researchers address topics that impact artificial intelligence more broadly. An important aspect is the interpretability or explainability of the outputs of these techniques (Appendix B.2.1). Another area of research involves approaches to integrating visual and textual data (Appendix B.4.3), while another focuses on virtual reality (Appendix B.3.5).

Another key research topic is misinformation. One participant (Appendix B.1.5) proposes hybrid human-in-the-loop approaches to detect misinformation at scale, while another (Appendix B.3.1) aims to leverage human intelligence for this purpose.

The formulation of real-world problems into mathematical representations—defined as models consisting of decision variables, constraints, and objective functions—is another area of study.

Applications include waste collection (Appendix B.1.2) and problems solvable using combinatorial optimization (Appendix B.2.2). One participant focuses specifically on performing efficient inference (Appendix B.3.3).

Digital humanities is another research field that has emerged among course participants, specifically focusing on intertextuality among literary works (Appendix B.1.4) and the use of the meta-verse for educational purposes (Appendix B.4.1). Finally, another topic involves cellular-based geolocation (Appendix B.2.3).

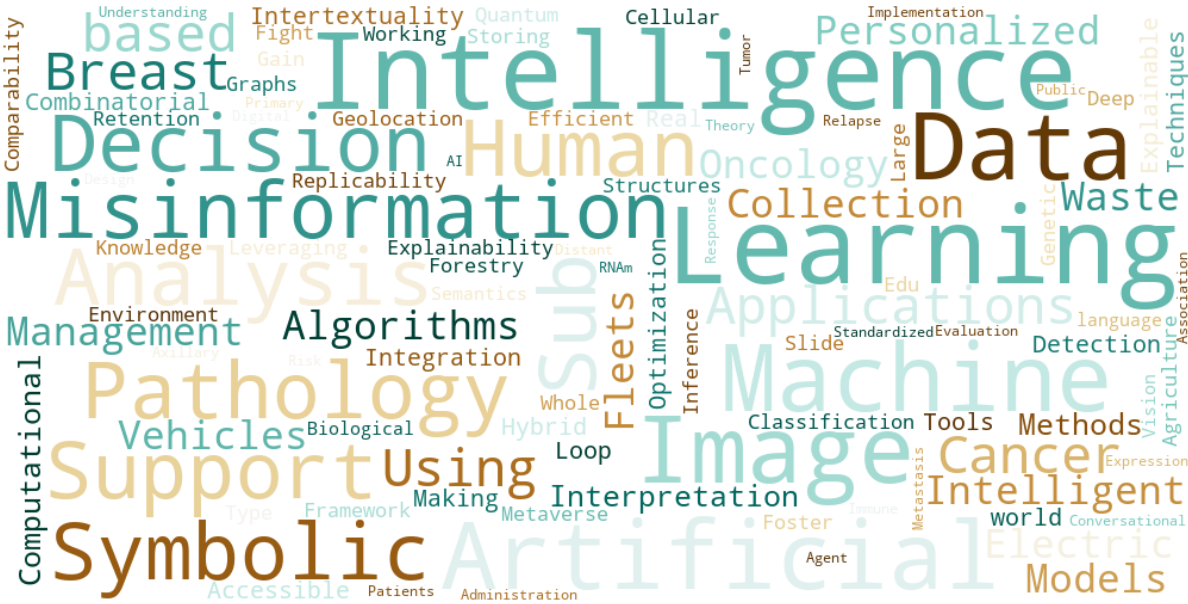


Figure 1. Word cloud of the research topics emerged during Activity 1.

Turning to each group discussion, the Marram/Kangaroo 🦘 group members acknowledged their shared background in computer science but agreed that additional domain-specific knowledge is necessary to further develop their research topics (Appendix B.1.6).

The Warin/Wombat 🦣 group members, similarly, acknowledged that their background relates only to computer science and information engineering, despite their research spanning additional topics, and that they lack domain-specific knowledge. They further highlighted that such a knowledge gap becomes evident when collaborating with other scholars (Appendix B.2.5).

The Gurrborra/Koala 🦘 group members, on the other hand, recognize differences in their educational backgrounds and acknowledge that their research topics often deal with real-world applications involving ethical concerns. Furthermore, their data often originate from people who participate in their experiments, thus particular attention must be paid when processing them. They also believe that a multidisciplinary team is needed to increase the quality of research (Appendix B.3.6).

The Dulai Wurrung/Platypus 🦨 group members, in their discussion, identified three factors that impact their research: their own internal biases, personal values, and external expectations. They believe that dealing with internal biases is the main challenge and propose being open-minded and embracing constructive ideas to improve the quality of their research (Appendix B.4.7).

In summary, most of the groups focused on discussing their educational backgrounds, the importance of having domain-specific knowledge, and proposed ideas to improve their research quality.

4 Activity 2: Keywords

This section summarizes the results of Activity 2. A handout for the activity is provided in Appendix A.2, and the content created by each group is included in Appendix C.

4.1 Objective

The activity aimed to cluster keywords and terms that are highly important for understanding participants' research topics. In their groups, participants were asked to brainstorm ideas along with key terms and refine them across three distinct steps (*layers*).


Initially, they were asked to propose terms and words significant to their research and computer science-related topics and group them according to specific criteria, such as by topic or shared issues (*Layer 1: Gather and Curate*). Next, they were tasked with identifying words that might be misunderstood or have different meanings in technical and real-world contexts (*Layer 2: Scrutinize*). Finally, the participants were asked to suggest tools and strategies to overcome such misunderstandings (*Layer 3: Translate*).

4.2 Procedure

The participants brainstormed terms and words in an online collaborative workspace using Miro.⁶ Each group was given a digital board divided into three columns, one for each layer, where they added their keywords using sticky notes. The activity took place in real-time, allowing each group to see the sticky notes added by the others.

We reproduced each board in Appendix C. Sticky notes that belong to the same cluster share a unique color across each group board. For the third layer, where participants often wrote short texts instead of single words, we have consolidated these into keywords for clarity. We also report the original texts, as generated by the participants, in the appendix. Below, we comment on the keywords that emerged.

4.3 Outcomes

The Marram/Kangaroo  group (Appendix B.1) identified six clusters of keywords. Among these, five refer to issues shared across their research topics, while one relates with a specific research field. Specifically, they point out that aspects such as credibility, ethics, equality, and privacy are related to trust, and that reliability and fairness favor explainability. Openness is important to support transparency, while quality control enforces the overall perception of safety. Additionally, biases should be considered. A particular focus is on healthcare, intended as a human-centered field of research. The group members argue that data regulations could support the notion of

⁶<https://miro.com/>

safety. Additionally, educational content should generally be massively available, and models should not be evaluated using results that have not been verified for quality.

Also the Warin/Wombat 🦘 group (Appendix B.2) identified six clusters of keywords. The group members focused solely on issues shared among their research topics, specifically emphasizing the need for fairness, explainability, impact, and anonymity. They also considered ethics and trust essential for ensuring privacy and highlighted the need for control over algorithms. Additionally, they explored the relationship between diversity, culture, and sustainability in relation to transparency, as well as the connection between inclusion, equality, and accountability in relation to safety. The group further elaborated on the theme of transparency, recommending the avoidance of metaphors in describing tools and processes. They suggested that stakeholders begin by discussing the project to reach a consensus, using examples to clarify their points and incorporating a syllabus that provides simple, clear definitions.

The Gurrborra/Koala 🦘 group (Appendix B.3), on the other hand, identified four clusters of keywords focused on issues and topics shared among the members' research areas. They emphasized the need to consider data completeness, as well as aggregation and generalization approaches for computing metrics. Additionally, they identified four elements that influence the social impact of their research: society, rewards, data availability, and general ethical concerns. They noted that researcher behavior and expertise affect trust and consent. The group also highlighted the importance of people and their biases, though they did not elaborate further. They discussed tools and topics related to three of the four clusters identified. Specifically, researchers should conduct meta-studies on data collection processes to assess their quality. Metrics should also be defined to measure the social impact of new technologies and track changes in society over time. Different levels and notions of trust and consent should be established, and tools enabling people to access their data directly should be developed.

Finally, the five clusters of keywords identified by the Dulai Wurrung/Platypus 🦘 group (Appendix B.4) all focused aspects shared by their research fields. They proposed the idea of family and related values, noting that people's backgrounds influence their biases. According to them, friends and responsibility shape personal experiences, while teamwork and goals help set expectations. They also mentioned the concept of culture without further elaboration and emphasized the importance of being open-minded when discussing words that might be misunderstood. The group further discussed four out of the five clusters of keywords. Most of their discussion centered around the idea of building a translator so that even the most technical content can be understood by everyone using unambiguous vocabulary. They argued that such a tool requires a semantic summarizer capable of standardizing and simplifying complex and technical texts, as well as a form of keyword segmentation.

The brainstorming of terms relevant to each group member's research, along with an examination of potentially misunderstood or ambiguous words, followed by a discussion on tools and strategies to address these issues, led to the emergence of several shared themes and ideas, as shown in Figure 2. The word cloud includes many concepts associated with responsible AI, such as Bias, Ethics, Fairness, Privacy, Explainability, Accountability, Credibility, and Trust. We can also see concepts related to governance, evaluation, and impact, such as Impact, Social, Quality, Control, Open, Metrics, and Data. Finally, there are keywords that are more domain-specific, such as Healthcare, Education, and Truthfulness.



5 Activities 3–5: Case Studies

This section summarizes the results of Activities 3-5. Their handouts are provided in Appendix A.3, Appendix A.4, and Appendix A.5, and the content created by each group is included in Appendix D.

5.1 Objective

The participants were asked to design case studies on topics related to their research background using mixed methods. They were then asked to reflect on fairness and diversity and outline ethical considerations. This included examining the reasons behind their research, identifying potential beneficiaries, and conducting risk management to address potential misuse of the research.

5.2 Procedure

Each group was allocated 10-15 minutes per session for the activities. During this time, group members discussed the topics orally, with all groups working simultaneously. The group leader was responsible for recording each consideration in a shared document using a predefined template.

More specifically, during Activity 3 (see Appendix A.3), each group developed at least one case study. The participants were asked to define the overall problem, list the research questions, formulate a positionality statement, and discuss the expected outcomes. After formulating the case study, the participants engaged in a group discussion.

For Activity 4 (see Appendix A.4) and Activity 5 (see Appendix A.5), the groups were required to further develop their case study. They resumed their oral discussions and expanded on the case study template by addressing fairness, diversity, and ethical considerations. To stimulate the

discussion on ethical considerations, participants were provided with a set of four questions to explore further.

After the course, we sorted and reorganized the user-generated content based on the predefined templates used to document each activity, with the goal of improving clarity and presentation. Below, we summarize the topics of each case study and the outcomes of the group discussions.

5.3 Outcomes

The Marram/Kangaroo 🦘 group addressed cancer diagnosis using AI approaches and their validation through crowdsourcing (Appendix D.1). The discussion focused on the ethical considerations of using AI in cancer diagnosis and radiotherapy planning, emphasizing the protection of patient privacy and the responsible use of medical data. It was noted that the combination of human expertise and AI has the potential to significantly enhance cancer diagnostics and treatment if properly managed.

The Warin/Wombat 🦘 group addressed solutions for performing geolocation using radio signals (Appendix D.2). In their discussion, they acknowledged the usefulness of user-based studies, as well as their scarcity, in validating the alignment of models and algorithms with real-world needs. They also believe that integrating user feedback and expert validation would lead to more robust and applicable outcomes.

The Gurrborra/Koala 🦘 group discussed understanding the perception of and commitment to longitudinal studies conducted on crowdsourcing platforms (Appendix D.3). They emphasized the need to understand the barriers to conducting longitudinal studies on crowdsourcing platforms, noting the absence of well-established guidelines for them. Fairer and better-designed studies will improve outcomes for both participants and requesters of crowdsourced work.

The Dulai Wurrung/Platypus 🦘 group focused on building a robust annotation tool for digital pathology and an age-based conversational agent for healthcare applications (Appendix D.4). They discussed using a conversational agent to manage communication across age groups, similar to how physicians adjust their vocabulary for patients. The agent would adapt its interaction style based on user preferences, using AI to enhance engagement, particularly for younger and older users.

6 Course Survey

The final survey, comprising a total of 30 questions, was divided into four main categories (see Appendix E for the entire questionnaire): demographics (4 questions), education, research background, and familiarity with course topics (9 questions), session and activities evaluation (11 questions), and course feedback (6 questions). Collecting such a variety of information is essential to understanding the diversity of our participants pool and gauging pre-existing knowledge and expertise.

Out of the 23 participants (Figure 3), 20 (86.96%) actively participated in the course activities and 14 (60.87%) completed the survey. From now on we refer to *participant* as someone who took the survey. Each participant took approximately 12 minutes to complete the questionnaire (an average time of 12 minutes and 17 seconds, with a standard deviation of 7 minutes and 11 seconds).



Figure 3. Group photo of the course participants (Dagstuhl style, but in Udine).

In the following sections, we report the summarized results of the survey.

6.1 Demographics

The 4 questions regarding demographic composition were designed to assess the diversity of participants (e.g., gender, age groups, etc.).

Information regarding the gender of the participants was collected through close-ended questions, with options including “female”, “male”, “non-binary”, and “prefer not to say”. Out of the 14 participants, 11 (78.57%) identified as male, and 3 (21.43%) identified as female.

Regarding age groups, data were collected through close-ended questions, comprising different age intervals from 18 to 64 years old. The options included 18–24, 25–34, 35–44, 45–54, 55–64, and prefer not to say. Among the 14 participants, 5 (35.71%) were aged 18–24, and 9 (64.29%) were aged 25–34.

Regarding country of origin, which was investigated through an open-ended question, 12 participants (85.72%) were from Italy, 1 (7.14%) was from Morocco, and 1 (7.14%) was from Pakistan.

All participants knew English. Among other languages spoken, Italian was spoken by 13 participants (92.86%), and Friulan by 4 participants (28.57%). Further details on the languages can be found in Figure 4.

6.2 Education, Research Background, and Familiarity with Course Topics

The 6 questions regarding educational, research, and work background were designed to gauge participants’ prior knowledge and expertise related to the course topics. The 3 questions, focused on familiarity with course content, aimed to evaluate whether participants had prior experience with conducting user studies (covered in Session 2) and with ethical and institutional review boards (covered in Session 3).

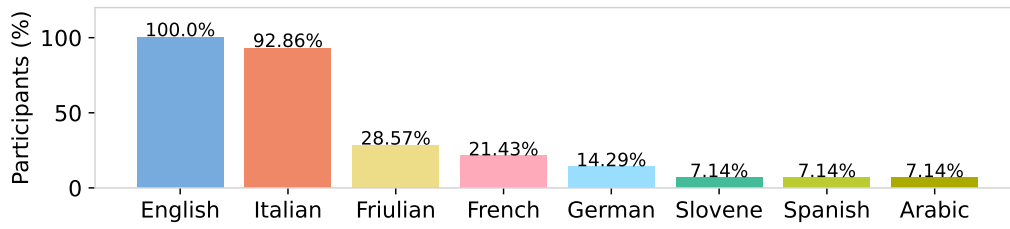


Figure 4. Languages spoken by course participants.

Additionally, the participants' backgrounds in various computer science and AI-related topics suggest that the methods proposed in the course (e.g., user studies, bias and fairness considerations, etc.) address a wide range of subjects, extending beyond just information access and fact-checking.

At the time the course was delivered (May 2024), 11 participants (78.57%) were PhD candidates, 2 participants (14.29%) were postdoctoral researchers, and 1 participant (7.14%) was a research assistant. Among the PhD candidates, 1 participant (7.14%) was a visiting PhD student.

Out of the 11 PhD candidates, 7 (63.64%) were in their first year of the doctoral program, 3 (27.27%) were in their second year, and 1 (9.09%) was in their third year. Additionally, 2 students (18.18% of the PhD candidates) were conducting their PhD in collaboration with industrial partners.

Regarding the assessment of research fields, some participants indicated more than one field. Out of the 14 participants, 8 (57.14%) indicated Machine Learning, 2 (14.29%) indicated Information Retrieval, 2 (14.29%) indicated Human-Computer Interaction, 2 (14.29%) indicated Optimization, 2 (14.29%) indicated Medical Imaging, and 1 (7.14%) indicated Formal Methods. A visual representation is shown in Figure 5.

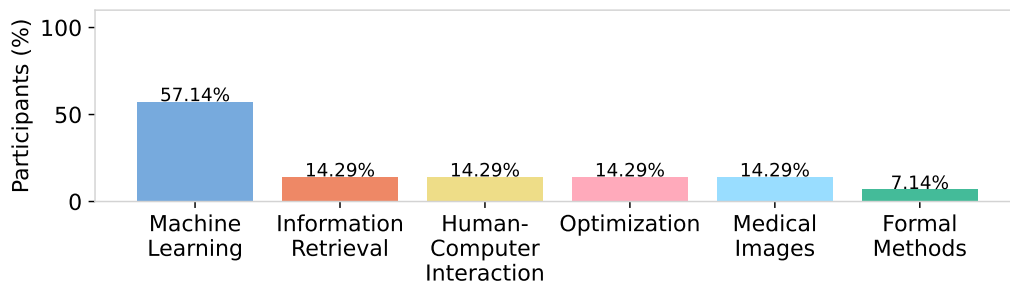


Figure 5. Research fields in which course participants are involved.

Regarding the assessment of Bachelor's degrees, all 14 participants held a Bachelor's degree, and none reported holding more than one. Among them, 10 participants (71.43%) had a Bachelor's degree in Computer Science-related fields (e.g., Computer Science, Internet of Things, Big Data & Web, and Web Technologies & Multimedia), 3 participants (21.43%) had a Bachelor's degree in Engineering-related fields (e.g., Electrical and Industrial Engineering), and 1 participant (7.14%) had a Bachelor's degree in Mathematics.

Regarding the assessment of Master's degrees, all 14 participants held a Master's degree, and none reported holding more than one. Among them, 11 participants (78.57%) had a Master's degree in Computer Science-related fields (e.g., Computer Science and Artificial Intelligence & Cybersecurity), while the remaining 3 participants (21.43%) had a Master's degree in Engineering-related fields (e.g., Electrical and Industrial Engineering). All in all, the majority of participants had a Computer Science-related background. Further details on the participants' educational backgrounds are shown in Figure 6.

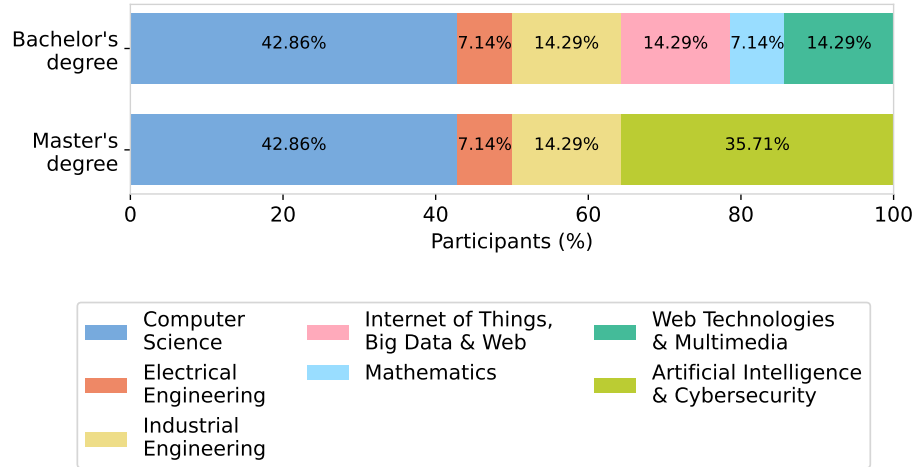


Figure 6. Distribution of course participants by degree and field of study.

Investigating work experience (i.e., any job conducted outside academia), 4 participants (28.57%) reported having such experience. Among them, 2 participants (50% of those with work experience) reported having worked for 2 years or more, 1 participant (25%) reported 1 year of work experience, and 1 participant (25%) did not specify the duration.

To assess participants' experience with user studies, we asked if they had conducted any in the past five years. Each participant selected one of the following options: 0 studies, 1–2 studies, 3–5 studies, or more than 5 studies. The majority of participants (9, or 64.29%) had never conducted a user study, while 2 participants (14.29%) had conducted more than 5 user studies (Figure 7).

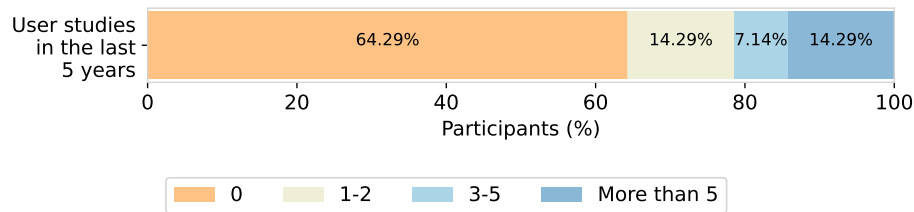


Figure 7. Number of user studies conducted by course participants over the past five years.

To assess participants' experience with ethical and institutional review boards, we asked about their familiarity with these subjects. Each participant had to choose one of the following options to indicate their level of familiarity: “not familiar at all”, “slightly familiar”, “very familiar”, or “extremely familiar”. Half of the participants revealed not having familiarity at all with human

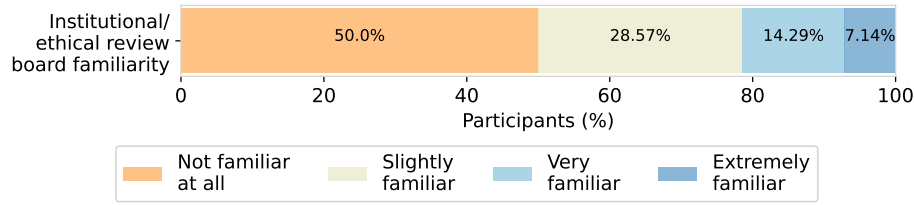


Figure 8. Familiarity level of course participants with ethical and institutional review board processes.

ethics approval processes (7 participants, 50%); considering the remaining participants, only a few had consistent familiarity: 2 of them (14.29%) were very familiar and one of them (7.14%) was extremely familiar (Figure 8).

To assess participants' experience with ethical and institutional review boards, we asked about their familiarity with these subjects. Each participant was asked to select one of the following options to indicate their level of familiarity: "not familiar at all", "slightly familiar", "very familiar", or "extremely familiar".

Half of the participants (7 participants, 50%) reported having no familiarity with human ethics approval processes. Among the remaining participants, only a few had substantial familiarity: 2 participants (14.29%) were very familiar, and 1 participant (7.14%) was extremely familiar (Figure 8).

6.3 Sessions and Activities

The 11 questions about the evaluation of sessions and activities were designed to understand the participants' perceptions of each component of the course.

The participants' perception was evaluated using Magnitude Estimation [Stevens, 1951, 1975]. Each participant was instructed to assign numbers greater than zero to reflect their perceived scores. This approach establishes a ratio scale, enabling direct comparison of normalized values.

Figure 9 presents the normalized scores evaluating the sessions, distinguishing between overall assessments and individual session evaluations. The overall sessions received a mean score of 0.76 (± 0.32) with a median of 0.95, based on the normalized values. In contrast, individual sessions were generally rated lower. Specifically, Session 1 had a mean score of 0.69 (± 0.36) and a median of 0.72; Session 2, 0.68 (± 0.43) with a median of 1.00; Session 3, 0.48 (± 0.42) with a median of 0.56; and Session 4, 0.57 (± 0.44) with a median of 0.67.

Figure 10 presents the normalized scores evaluating the activities, distinguishing between the overall assessment and individual activities. Overall, activities were appreciated less than sessions. The overall activities received a mean score of 0.51 (± 0.37) with a median of 0.52. Activity 1 had a mean score of 0.65 (± 0.35) and a median of 0.67; Activity 2, 0.49 (± 0.38) with a median of 0.33; Activity 3, 0.42 (± 0.43) with a median of 0.33; and Activity 4, 0.35 (± 0.34) with a median of 0.33.

6.4 Overall Course Feedback

The 6 questions regarding course feedback were designed to understand to what extent the participants appreciated the course and to identify potential areas for improvement.

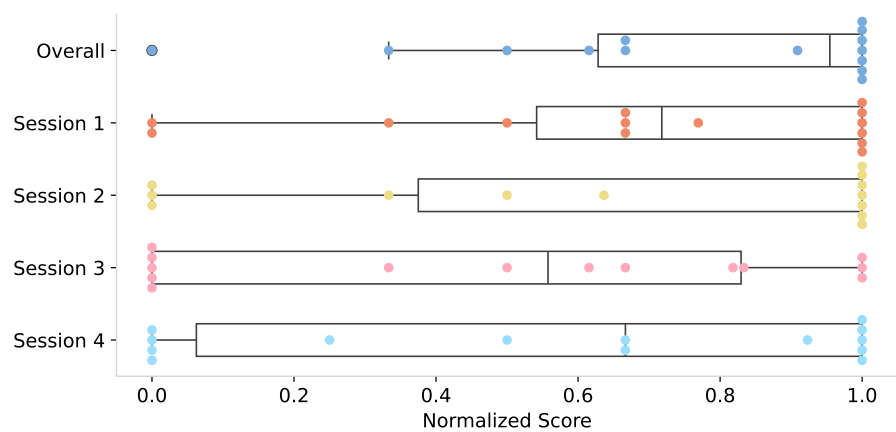


Figure 9. Evaluation of each course session and the overall experience.

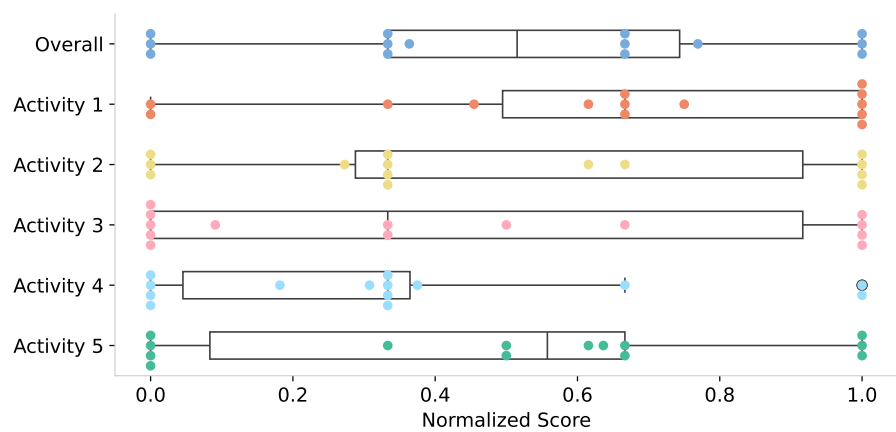


Figure 10. Evaluation of each course activity and the overall experience.

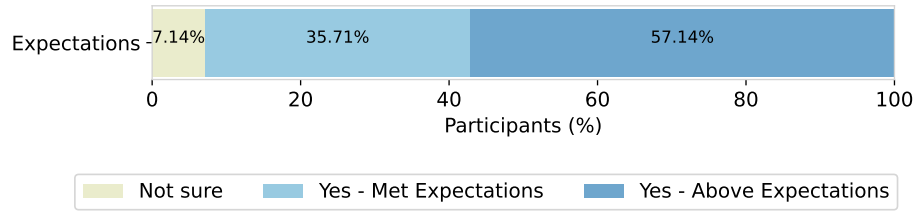


Figure 11. Opinion of course participants about whether the course met their expectations.

The expectations toward the course were assessed with a close-ended question using five levels, ranging from “above” to “far below” expectations. Only 1 participant (7.14%) declared being neutral about their expectations, and no participant reported negative expectations (see Figure 11 for details).

The likelihood of recommending the course to peers was assessed with a close-ended question using a discrete scale from 1 to 10, where 10 represented the highest likelihood. As shown in Figure 12, the majority of participants indicated they would recommend the course.

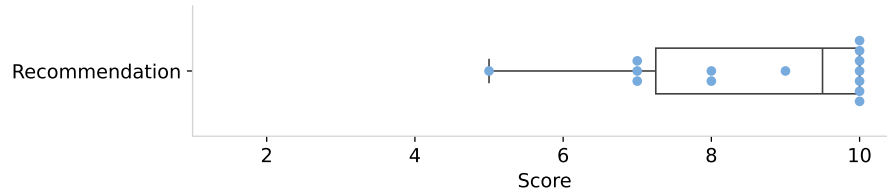


Figure 12. Likelihood of course participants recommending the course to their peers.

The effectiveness of the course methods was assessed based on the following dimensions: quality of instructional materials, quality of learning activities, use of technologies, and group activities. These dimensions were evaluated using close-ended questions with responses ranging from “extremely ineffective” to “extremely effective”. For all dimensions, the majority of participants agreed that they were effective. The only exceptions are as follows: In assessing the use of technologies, 1 participant (7.14%) reported it as somewhat ineffective. In the group activities, 1 participant (7.14%) reported it as extremely ineffective.

7 Conclusion

This paper describes the results of a teaching and learning activity to apply responsible AI concepts in the context of information retrieval research.

The positive feedback received by the students highlights the benefits of the interactive approach in introducing concepts such as positionality, participatory research, fairness, diversity, and ethics.

We believe that – complementing existing resources [Fröbe et al., 2024; Markov and de Rijke, 2019; Bauer et al., 2023]– the hands-on activities discussed in this paper can assist practitioners and educators in the design of responsible AI content for information retrieval curriculum.

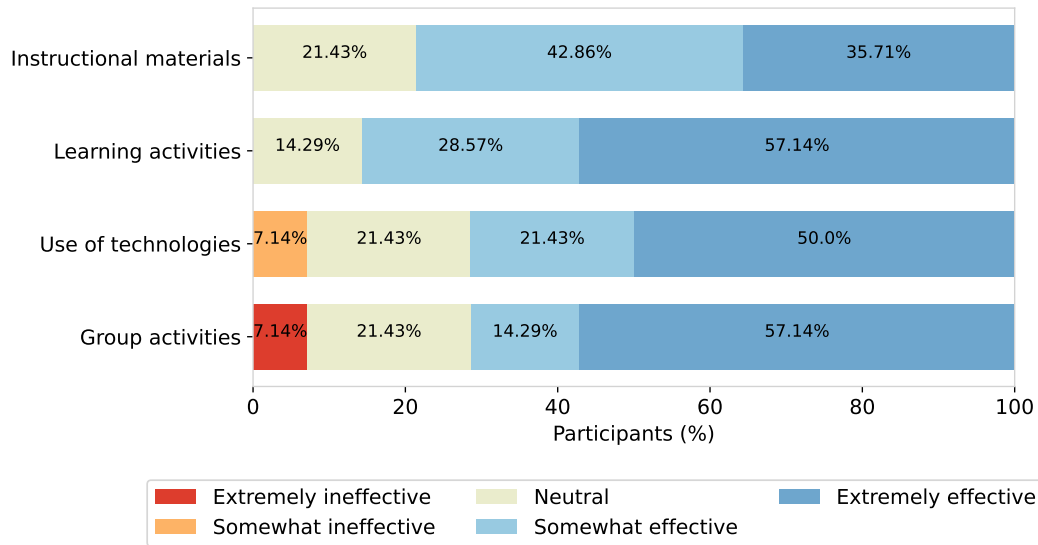


Figure 13. Opinion of course participants about the effectiveness of the course methods.

Acknowledgments

The course described in this report was designed and developed in the unceded lands of the Wurundjeri and Boon Wurrung peoples of the eastern Kulin Nation. We pay our respects to their Ancestors and Elders, past, and present.

Most of the research included in the course was supported by the Australian Research Council under the DECRA scheme (DE200100064) and the ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S, CE200100005).

The course was organized as a training activity for students enrolled in the PhD Course on Computer Science and Artificial Intelligence at the University of Udine.

References

- Raghad Alfaisal, Haslinda Hashim, and Ummu Husna Azizan. Metaverse system adoption in education: a systematic literature review. *Journal of Computers in Education*, 11(1):259–303, Mar 2024. ISSN 2197-9995. URL <https://doi.org/10.1007/s40692-022-00256-6>.
- Enrique Amigó, Hui Fang, Stefano Mizzaro, and ChengXiang Zhai. Are we on the Right Track? An Examination of Information Retrieval Methodologies. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18, page 997–1000, New York, NY, USA, 2018a. Association for Computing Machinery. ISBN 9781450356572. URL <https://doi.org/10.1145/3209978.3210131>.
- Enrique Amigó, Damiano Spina, and Jorge Carrillo-de Albornoz. 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*, SIGIR '18, page 625–634, New York, NY, USA, 2018b. Association for Computing Machinery. ISBN 9781450356572. URL <https://doi.org/10.1145/3209978.3210024>.
- Enrique Amigó, Stefano Mizzaro, and Damiano Spina. Ranking Interruptus: When Truncated Rankings Are Better and How to Measure That. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, page 588–598, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450387323. URL <https://doi.org/10.1145/3477495.3532051>.
- Giulia Lucrezia Baroni, Laura Rasotto, Kevin Roitero, Angelica Tulisso, Carla Di Loreto, and Vincenzo Della Mea. Optimizing Vision Transformers for Histopathology: Pretraining and Normalization in Breast Cancer Classification. *Journal of Imaging*, 10(5):108, 2024. URL <https://doi.org/10.3390/jimaging10050108>.
- Christine Bauer, Ben Carterette, Nicola Ferro, Norbert Fuhr, Joeran Beel, Timo Breuer, Charles L. A. Clarke, Anita Crescenzi, Gianluca Demartini, Giorgio Maria Di Nunzio, Laura Dietz, Guglielmo Faggioli, Bruce Ferwerda, Maik Fröbe, Matthias Hagen, Allan Hanbury, Claudia Hauff, Dietmar Jannach, Noriko Kando, Evangelos Kanoulas, Bart P. Knijnenburg, Udo Kruschwitz, Meijie Li, Maria Maistro, Lien Michiels, Andrea Papenmeier, Martin Potthast, Paolo Rosso, Alan Said, Philipp Schaer, Christin Seifert, Damiano Spina, Benno Stein, Nava Tintarev, Julián Urbano, Henning Wachsmuth, Martijn C. Willemsen, and Justin Zobel. Report on the Dagstuhl Seminar on Frontiers of Information Access Experimentation for Research and Education. *SIGIR Forum*, 57(1), dec 2023. ISSN 0163-5840. URL <https://doi.org/10.1145/3636341.3636351>.
- Rodrigo Bavaresco, Diórgenes Silveira, Eduardo Reis, Jorge Barbosa, Rodrigo Righi, Cristiano Costa, Rodolfo Antunes, Marcio Gomes, Clauter Gatti, Mariangela Vanzin, Saint Clair Junior, Elton Silva, and Carlos Moreira. Conversational agents in business: A systematic literature review and future research directions. *Computer Science Review*, 36:100239, 2020. ISSN 1574-0137. URL <https://doi.org/10.1016/j.cosrev.2020.100239>.

-
- Christian Blum and Andrea Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*, 35(3):268–308, September 2003. URL <https://doi.org/10.1145/937503.937505>.
- Marília D. V. Braga, Leonie R. Brockmann, Katharina Klerx, and Jens Stoye. A Linear Time Algorithm for an Extended Version of the Breakpoint Double Distance. In *22nd International Workshop on Algorithms in Bioinformatics (WABI 2022)*, volume 242 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:16, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-243-3. URL <https://doi.org/10.4230/LIPIcs.WABI.2022.13>.
- Erik Brand, Kevin Roitero, Michael Soprano, Afshin Rahimi, and Gianluca Demartini. A Neural Model to Jointly Predict and Explain Truthfulness of Statements. *Journal of Data and Information Quality*, 15(1), 12 2022. ISSN 1936-1955. URL <https://doi.org/10.1145/3546917>.
- Jean Burgess, Kath Albury, Anthony McCosker, and Rowan Wilken. *Everyday Data Cultures*. Polity Press, Cambridge, 2022. ISBN 9781509547562. URL <https://www.wiley.com/en-us/Everyday+Data+Cultures-p-9781509547562>.
- Davide Ceolin, Giuseppe Primiero, Michael Soprano, and Jan Wielemaker. Transparent Assessment of Information Quality of Online Reviews Using Formal Argumentation Theory. *Information Systems*, 110:102107, 7 2022. ISSN 0306-4379. URL <https://doi.org/10.1016/j.is.2022.102107>.
- Sabyasachi Dash, Sushil Kumar Shakyawar, Mohit Sharma, and Sandeep Kaushik. Big data in healthcare: management, analysis and future prospects. *Journal of Big Data*, 6(1):54, Jun 2019. ISSN 2196-1115. doi: 10.1186/s40537-019-0217-0. URL <https://doi.org/10.1186/s40537-019-0217-0>.
- Vincenzo Della Mea, Mihai Horia Popescu, and Kevin Roitero. Underlying Cause of Death Identification from Death Certificates via Categorical Embeddings and Convolutional Neural Networks. In *2020 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 1–6, 2020. URL <https://doi.org/10.1109/ICHI48887.2020.9374316>.
- Vincenzo Della Mea, Mihai Horia Popescu, and Kevin Roitero. Underlying cause of death identification from death certificates using reverse coding to text and a NLP based deep learning approach. *Informatics in Medicine Unlocked*, 21:100456, 2020. ISSN 2352-9148. URL <https://doi.org/10.1016/j.imu.2020.100456>.
- Radosvet Desislavov, Fernando Martínez-Plumed, and José Hernández-Orallo. Trends in AI inference energy consumption: Beyond the performance-vs-parameter laws of deep learning. *Sustainable Computing: Informatics and Systems*, 38:100857, 2023. ISSN 2210-5379. URL <https://doi.org/10.1016/j.suscom.2023.100857>.
- Jianfeng Dong, Xirong Li, Chaoxi Xu, Xun Yang, Gang Yang, Xun Wang, and Meng Wang. Dual encoding for video retrieval by text. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4065–4080, 2021. URL <https://doi.org/10.1109/TPAMI.2021.3059295>.

-
- Tim Draws, David La Barbera, Michael Soprano, Kevin Roitero, Davide Ceolin, Alessandro Checco, and Stefano Mizzaro. The Effects of Crowd Worker Biases in Fact-Checking Tasks. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 2114–2124. Association for Computing Machinery, 2022. ISBN 9781450393522. URL <https://doi.org/10.1145/3531146.3534629>.
- Michael D. Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. Fairness in Information Access Systems. *Foundations and Trends in Information Retrieval*, 16(1–2):1–177, 7 2022. ISSN 1554-0669. URL <https://doi.org/10.1561/15000000079>.
- Alex Falcon, Giuseppe Serra, and Oswald Lanz. Improving semantic video retrieval models by training with a relevance-aware online mining strategy. *Computer Vision and Image Understanding*, 245:104035, 2024. ISSN 1077-3142. URL <https://doi.org/10.1016/j.cviu.2024.104035>.
- Romanos Fasoulis, Georgios Paliouras, and Lydia E. Kavraki. Graph representation learning for structural proteomics. *Emerging Topics in Life Sciences*, 5(6):789–802, 10 2021. ISSN 2397-8554. URL <https://doi.org/10.1042/ETLS20210225>.
- Jacques Ferlay, Morten Ervik, Frederic Lam, Mathieu Laversanne, Murielle Colombet, Les Mery, Marion Piñeros, Ariana Znaor, Isabelle Soerjomataram, and Freddie Bray. Global Cancer Observatory: Cancer Today. <https://gco.iarc.who.int/today>, 2024. Lyon, France: International Agency for Research on Cancer.
- Sebastian Försch, Frederick Klauschen, Peter Hufnagl, and Wilfried Roth. Artificial intelligence in pathology. *Deutsches Ärzteblatt International*, 118(12):199, 2021. URL <https://doi.org/10.3238/arztebl.m2021.0011>.
- Maik Fröbe, Harris Scells, Theresa Elstner, Christopher Akiki, Lukas Gienapp, Jan Heinrich Reimer, Sean MacAvaney, Benno Stein, Matthias Hagen, and Martin Potthast. Resources for Combining Teaching and Research in Information Retrieval Coursework. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, page 1115–1125, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704314. URL <https://doi.org/10.1145/3626772.3657886>.
- Bala Subrahmanyam Garimella, Hari Sharan Garlapati, Sriharini Choul, Rajesh Cherukuri, and Pallavi Lanke. Advancing Healthcare Accessibility: Development of an AI-Driven Multimodal Chatbot. In *CONIT 2024*, pages 1–10, 2024. URL <https://doi.org/10.1016/10.1109/CONIT61985.2024.10626795>.
- Erik Garrison, Andrea Guarracino, Simon Heumos, Flavia Villani, Zhigui Bao, Lorenzo Tattini, Jörg Hagmann, Sebastian Vorbrugg, Santiago Marco-Sola, Christian Kubica, David G. Ashbrook, Kaisa Thorell, Rachel L. Rusholme-Pilcher, Gianni Liti, Emilio Rudbeck, Sven Nahnsen, Zuyu Yang, Mwaniki N. Moses, Franklin L. Nobrega, Yi Wu, Hao Chen, Joep de Lig, Peter H. Sudmant, Nicole Soranzo, Vincenza Colonna, Robert W. Williams, and Pjotr Prins. Building pangenome graphs. *bioRxiv*, 2023. URL <https://doi.org/10.1101/2023.04.05.535718>.

-
- Jonathan Herington, Melissa D. McCradden, Kathleen Creel, Ronald Boellaard, Elizabeth C. Jones, Abhinav K. Jha, Arman Rahmim, Peter J.H. Scott, John J. Sunderland, Richard L. Wahl, Sven Zuehlsdorff, and Babak Saboury. Ethical Considerations for Artificial Intelligence in Medical Imaging: Deployment and Governance. *Journal of Nuclear Medicine*, 2023. ISSN 0161-5505. doi: 10.2967/jnumed.123.266110. URL <https://jnm.snmjournals.org/content/early/2023/08/23/jnumed.123.266110>.
- Danula Hettiachchi, Kaixin Ji, Jenny Kennedy, Anthony McCosker, Flora D. Salim, Mark Sander-son, Falk Scholer, and Damiano Spina. Designing and Evaluating Presentation Strategies for Fact-Checked Content. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, New York, NY, USA, 2023. Association for Computing Machinery. URL <https://doi.org/10.1145/3583780.3614841>.
- Djoerd Hiemstra. Was Fairness in IR Discussed by Cooper and Robertson in the 1970's? *SIGIR Forum*, 56(2), 1 2023. ISSN 0163-5840. URL <https://doi.org/10.1145/3582900.3582924>.
- Mahdi S Hosseini, Babak Ehteshami Bejnordi, Vincent Quoc-Huy Trinh, Lyndon Chan, Danial Hasan, Xingwen Li, Stephen Yang, Taehyo Kim, Haochen Zhang, Theodore Wu, et al. Computational pathology: a survey review and the way forward. *Journal of Pathology Informatics*, page 100357, 2024. URL <https://doi.org/10.1016/j.jpi.2023.100357>.
- Ekta Jain, Ankush Patel, Anil V Parwani, Saba Shafi, Zoya Brar, Shivani Sharma, and Sam-bit K Mohanty. Whole slide imaging technology and its applications: Current and emerging perspectives. *International Journal of Surgical Pathology*, 32(3):433–448, 2024.
- Kaixin Ji. Quantifying and Measuring Confirmation Bias in Information Retrieval Using Sen-sors. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*, Ubi-Comp/ISWC '23 Adjunct, page 236–240, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400702006. URL <https://doi.org/10.1145/3594739.3610765>.
- Kaixin Ji, Damiano Spina, Danula Hettiachchi, Flora D. Salim, and Falk Scholer. Examining the Impact of Uncontrolled Variables on Physiological Signals in User Studies for Information Processing Activities. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, New York, NY, USA, 2023a. ACM. URL <https://doi.org/10.1145/3539618.3591981>.
- Kaixin Ji, Damiano Spina, Danula Hettiachchi, Falk Scholer, and Flora D. Salim. Towards Detect-ing Tonic Information Processing Activities with Physiological Data. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*, UbiComp/ISWC '23 Ad-junct, page 1–5, New York, NY, USA, 2023b. Association for Computing Machinery. ISBN 9798400702006. URL <https://doi.org/10.1145/3594739.3610679>.
- Kaixin Ji, Danula Hettiachchi, Flora D. Salim, Falk Scholer, and Damiano Spina. Characterizing Information Seeking Processes with Multiple Physiological Signals. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*,

-
- SIGIR '24, New York, NY, USA, 2024a. ACM. URL <https://doi.org/10.1145/3626772.365779>.
- Kaixin Ji, Sachin Pathiyan Cherumanal, Johanne R. Trippas, Danula Hettiachchi, Flora D. Salim, Falk Scholer, and Damiano Spina. Towards Detecting and Mitigating Cognitive Bias in Spoken Conversational Search. In *Adjunct Proceedings of the 26th International Conference on Mobile Human-Computer Interaction*, MobileHCI '24 Adjunct, New York, NY, USA, 2024b. Association for Computing Machinery. ISBN 9798400705069. URL <https://doi.org/10.1145/3640471.3680245>.
- Bijan Khosrawi-Rad, Heidi Rinn, Ricarda Schlimbach, Pia Gebbing, Xingyue Yang, Christoph Lattemann, Daniel Markgraf, and Susanne Robra-Bissantz. Conversational Agents in Education – A Systematic Literature Review. In *ECIS 2022 Research Papers*, 2022. URL https://aise1.aisnet.org/ecis2022_rp/18.
- Johannes Kiesel, Damiano Spina, Henning Wachsmuth, and Benno Stein. The Meant, the Said, and the Understood: Conversational Argument Search and Cognitive Biases. In *CUI 2021 - 3rd Conference on Conversational User Interfaces*, CUI '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450389983. URL <https://doi.org/10.1145/3469595.3469615>.
- Scott Kirkpatrick, Daniel Gelatt, and Mario P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, 1983. URL <https://doi.org/10.1126/science.220.4598.671>.
- David La Barbera, Stefano Mizzaro, and Kevin Roitero. A Hybrid Human-In-The-Loop Framework for Fact Checking. In Debora Nozza, Lucia C. Passaro, and Marco Polignano, editors, *Proceedings of the Sixth Workshop on Natural Language for Artificial Intelligence (NL4AI 2022) co-located with 21th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2022), Udine, November 30th, 2022*, volume 3287 of *CEUR Workshop Proceedings*, pages 13–23. CEUR-WS.org, 2022. URL <https://ceur-ws.org/Vol-3287/paper4.pdf>.
- Marie-Louise Lackner, Christoph Mrkvicka, Nysret Musliu, Daniel Walkiewicz, and Felix Winter. Exact methods for the Oven Scheduling Problem. *Constraints*, 28(2):320–361, 2023. URL <https://doi.org/10.1007/s10601-023-09347-2>.
- Latrice G. Landry, Nadya Ali, David R. Williams, Heidi L. Rehm, and Vence L. Bonham. Lack Of Diversity In Genomic Databases Is A Barrier To Translating Precision Medicine Research Into Practice. *Health Affairs*, 37(5):780–785, 2018. URL <https://doi.org/10.1377/hlthaff.f.2017.1595>.
- Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models, 2023. URL <https://doi.org/10.48550/arxiv.2304.03271>.
- C. K. Lim, K. L. Tan, A. A. Zaidan, and B. B. Zaidan. A proposed methodology of bringing past life in digital cultural heritage through crowd simulation: a case study in George Town, Malaysia. *Multimedia Tools and Applications*, 79(5):3387–3423, Feb 2020. ISSN 1573-7721. URL <https://doi.org/10.1007/s11042-019-07925-2>.

-
- Fangfang Liu, Thomas Hardiman, Kailiang Wu, Jelmar Quist, Patrycja Gazinska, Tony Ng, Arnie Purushotham, Roberto Salgado, Xiaojing Guo, Sarah E. Pinder, and Anita Grigoriadis. Systemic immune reaction in axillary lymph nodes adds to tumor-infiltrating lymphocytes in triple-negative breast cancer prognostication. *npj Breast Cancer*, 7(1):86, Jul 2021. ISSN 2374-4677. doi: 10.1038/s41523-021-00292-y. URL <https://doi.org/10.1038/s41523-021-00292-y>.
- Junyu Liu, Minzhao Liu, Jin-Peng Liu, Ziyu Ye, Yunfei Wang, Yuri Alexeev, Jens Eisert, and Liang Jiang. Towards provably efficient quantum algorithms for large-scale machine-learning models. *Nature Communications*, 15(1):434, Jan 2024. ISSN 2041-1723. URL <https://doi.org/10.1038/s41467-023-43957-x>.
- Riccardo Lunardi and Paolo Coppola. Conversational-Agent for Patient Information Leaflet. In *Proceedings of the 14th Italian Information Retrieval Workshop*, 2024. URL <https://hdl.handle.net/11390/1292426>.
- Riccardo Lunardi, David La Barbera, and Kevin Roitero. The Elusiveness of Detecting Political Bias in Language Models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, CIKM '24, page 3922–3926, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704369. URL <https://doi.org/10.1145/3627673.3680002>.
- Carlos López, Ramon Bosch, Guifre Orero, Anna Korzynska, Marcial García-Rojo, Gloria Bueno, María del Milagro Fernández-Carrobles, Albert Gibert-Ramos, Lukasz Roszkowiak, Cristina Callau, Laia Fontoura, Maria-Teresa Salvadó, Tomás Álvaro, Joaquín Jaén, Albert Roso-Llorach, Montserrat Llobera, Julia Gil, Montserrat Onyos, Benoît Plancoulaine, Jordi Baucells, and Marylène Lejeune. The Immune Response in Nonmetastatic Axillary Lymph Nodes Is Associated with the Presence of Axillary Metastasis and Breast Cancer Patient Outcome. *The American Journal of Pathology*, 190(3):660–673, 2020. ISSN 0002-9440. URL <https://doi.org/10.1016/j.ajpath.2019.11.002>.
- Carlos López, Albert Gibert-Ramos, Ramón Bosch, Anna Korzynska, Marcial García-Rojo, Gloria Bueno, Joan Francesc García-Fontgivell, Salomé Martínez González, Laia Fontoura, Andrea Gras Navarro, Esther Sauras Colón, Júlia Casanova Ribes, Lukasz Roszkowiak, Albert Roso, Marta Berenguer, Montserrat Llobera, Jordi Baucells, and Marylène Lejeune. Differences in the Immune Response of the Nonmetastatic Axillary Lymph Nodes between Triple-Negative and Luminal A Breast Cancer Surrogate Subtypes. *The American Journal of Pathology*, 191(3): 545–554, 2021. ISSN 0002-9440. URL <https://doi.org/10.1016/j.ajpath.2020.11.008>.
- Ilya Markov and Maarten de Rijke. What Should We Teach in Information Retrieval? *SIGIR Forum*, 52(2):19–39, January 2019. ISSN 0163-5840. URL <https://doi.org/10.1145/3308774.3308780>.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2630–2640, 2019. URL <https://doi.org/10.1109/ICCV.2019.00272>.

-
- Stefano Mizzaro. Relevance: the whole history. *Journal of the American Society for Information Science*, 48(9):810–832, sep 1997. ISSN 0002-8231. URL <https://dl.acm.org/doi/10.5555/262192.262203>.
- Benjamin Moxley-Wyles, Richard Colling, and Clare Verrill. Artificial intelligence in pathology: an overview. *Diagnostic Histopathology*, 26(11):513–520, 2020. URL <https://doi.org/10.1016/j.mpdhp.2020.08.004>.
- Soraia Raupp Musse, Vinicius Jurinic Cassol, and Daniel Thalmann. A history of crowd simulation: the past, evolution, and new perspectives. *Vis. Comput.*, 37(12):3077–3092, December 2021. ISSN 0178-2789. URL <https://doi.org/10.1007/s00371-021-02252-w>.
- Alexandra Olteanu, Michael Ekstrand, Carlos Castillo, and Jina Suh. Responsible AI Research Needs Impact Statements Too. arXiv, 2023. URL <https://doi.org/10.48550/arXiv.2311.11776>.
- Mohamed Omar, Mohammad K. Alexanderani, Itzel Valencia, Massimo Loda, and Luigi Marchionni. Applications of Digital Pathology in Cancer: A Comprehensive Review. *Annual Review of Cancer Biology*, 8:245–268, 2024. URL <https://doi.org/10.1146/annurev-cancerbio-062822-010523>.
- Sachin Pathiyan Cherumanal, Damiano Spina, Falk Scholer, and W. Bruce Croft. Evaluating Fairness in Argument Retrieval. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management, CIKM '21*, page 3363–3367, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384469. URL <https://doi.org/10.1145/3459637.3482099>.
- Sachin Pathiyan Cherumanal, Kaixin Ji, Danula Hettiachchi, Johanne R. Trippas, Falk Scholer, and Damiano Spina. RMIT_IR at the NTCIR-17 FairWeb-1 Task. In *Proceedings of 17th Conference on Evaluation of Information Access Technologies, NTCIR-17*, 2023. URL <https://doi.org/10.20736/0002001315>.
- Sachin Pathiyan Cherumanal, Falk Scholer, Johanne R. Trippas, and Damiano Spina. Towards Investigating Biases in Spoken Conversational Search. In *Companion Publication of the 26th International Conference on Multimodal Interaction, ICMI '24 Companion*, 2024. URL <https://doi.org/10.1145/3686215.3690156>.
- Sachin Pathiyan Cherumnal, Marwah Alaofi, Reham Abdullah Altalhi, Elham Naghizade, Falk Scholer, and Damiano Spina. RMIT CIDDA IR at the TREC 2022 Fair Ranking Track. In *Proceedings of TREC 2022*, 2023. URL https://trec.nist.gov/pubs/trec31/papers/rmit_cidda_ir.F.pdf.
- Gordon Pennycook, Ziv Epstein, Mohsen Mosleh, Antonio A. Arechar, Dean Eckles, and David G. Rand. Shifting Attention To Accuracy Can Reduce Misinformation Online. *Nature*, 592(7855): 590–595, 4 2021. ISSN 1476-4687. URL <https://doi.org/10.1038/s41586-021-03344-2>.

-
- Raquel Perez-Lopez, Narmin Ghaffari Laleh, Faisal Mahmood, and Jakob Nikolas Kather. A guide to artificial intelligence for cancer researchers. *Nature Reviews Cancer*, 24(6):427–441, 6 2024. ISSN 1474-1768. URL <https://doi.org/10.1038/s41568-024-00694-7>.
- Mihai Horia Popescu, Kevin Roitero, Stefano Travasci, and Vincenzo Della Mea. Automatic Assignment of ICD-10 Codes to Diagnostic Texts using Transformers Based Techniques. In *2021 IEEE 9th International Conference on Healthcare Informatics (ICHI)*, pages 188–192, 2021. URL <https://doi.org/10.1109/ICHI52183.2021.00037>.
- Mihai Horia Popescu, Can Celik, Vincenzo Della Mea, and Robert Jakob. Preliminary Validation of a Rule-Based System for Mortality Coding Using ICD-11. *Studies in Health Technology and Informatics*, 294:679–683, 2022a. URL <https://doi.org/10.3233/SHTI220555>.
- Mihai Horia Popescu, Kevin Roitero, and Vincenzo Della Mea. Explainable Classification of Medical Documents Through a Text-to-Text Transformer. In Francesco Calimeri, Mauro Dragoni, and Fabio Stella, editors, *1st AIXIA Workshop on Artificial Intelligence For Healthcare (HC@AIXIA 2022) co-located with the 21st International Conference of the Italian Association for Artificial Intelligence (AIXIA 2022), Udine, Italy, November 30 2022*, volume 3307 of *CEUR Workshop Proceedings*, pages 57–66. CEUR-WS.org, 2022b. URL <https://ceur-ws.org/Vol-3307/paper6.pdf>.
- Wullianallur Raghupathi and Viju Raghupathi. Big data analytics in healthcare: promise and potential. *Health Information Science and Systems*, 2(1):3, Feb 2014. ISSN 2047-2501. doi: 10.1186/2047-2501-2-3. URL <https://doi.org/10.1186/2047-2501-2-3>.
- Kevin Roitero, David La Barbera, Michael Soprano, Gianluca Demartini, Stefano Mizzaro, and Tetsuya Sakai. How Many Crowd Workers Do I Need? On Statistical Power When Crowdsourcing Relevance Judgments. *ACM Transactions on Information Systems*, 42(1), aug 2023a. ISSN 1046-8188. URL <https://doi.org/10.1145/3597201>.
- Kevin Roitero, Michael Soprano, Beatrice Portelli, Massimiliano De Luise, Damiano Spina, Vincenzo Della Mea, Giuseppe Serra, Stefano Mizzaro, and Gianluca Demartini. Can the crowd judge truthfulness? a longitudinal study on recent misinformation about covid-19. *Personal and Ubiquitous Computing*, 27(1):59–89, 2 2023b. ISSN 1617-4917. URL <https://doi.org/10.1007/s00779-021-01604-6>.
- Francesc Rosselló and Gabriel Valiente. Analysis of Metabolic Pathways by Graph Transformation. In Hartmut Ehrig, Gregor Engels, Francesco Parisi-Presicce, and Grzegorz Rozenberg, editors, *Graph Transformations*, pages 70–82, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. ISBN 978-3-540-30203-2. URL https://doi.org/10.1007/978-3-540-30203-2_7.
- Leena Salmela, Riku Walve, Eric Rivals, and Esko Ukkonen. Accurate self-correction of errors in long reads using de Bruijn graphs. *Bioinformatics*, 33(6):799–806, 06 2016. ISSN 1367-4803. URL <https://doi.org/10.1093/bioinformatics/btw321>.
- Mark Sanderson. Test Collection Based Evaluation of Information Retrieval Systems. *Foundations and Trends in Information Retrieval*, 4(4):247–375, 2010. ISSN 1554-0669. URL <http://dx.doi.org/10.1561/1500000009>.

-
- Paul Shaw. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *Proceedings of the 4th International Conference on Principles and Practice of Constraint Programming*, CP '98, pages 417–431, Berlin, Heidelberg, 1998. Springer-Verlag. URL https://doi.org/10.1007/3-540-49481-2_30.
- Kate Smith-Miles and Mario Andrés Muñoz. Instance Space Analysis for Algorithm Testing: Methodology and Software Tools. *ACM Comput. Surv.*, 55(12), 2023. doi: 10.1145/3572895.
- Michael Soprano, Kevin Roitero, David La Barbera, Davide Ceolin, Damiano Spina, Stefano Mizzaro, and Gianluca Demartini. The Many Dimensions of Truthfulness: Crowdsourcing Misinformation Assessments on a Multidimensional Scale. *Information Processing & Management*, 58(6):102710, 2021. ISSN 0306-4573. URL <https://doi.org/10.1016/j.ipm.2021.102710>.
- Michael Soprano, Kevin Roitero, Ujwal Gadiraju, Eddy Maddalena, and Gianluca Demartini. Longitudinal Loyalty: Understanding The Barriers To Running Longitudinal Studies On Crowdsourcing Platforms. *ACM Transactions on Social Computing*, 7(1–4), 9 2024a. URL <https://doi.org/10.1145/3674884>.
- Michael Soprano, Kevin Roitero, David La Barbera, Davide Ceolin, Damiano Spina, Gianluca Demartini, and Stefano Mizzaro. Cognitive Biases in Fact-Checking and Their Countermeasures: A Review. *Information Processing & Management*, 61(3):103672, 2024b. ISSN 0306-4573. URL <https://doi.org/10.1016/j.ipm.2024.103672>.
- Kenneth Sörensen. Metaheuristics—the metaphor exposed. *International Transactions in Operational Research*, 22(1):3–18, 2015. URL <https://doi.org/10.1111/itor.12001>.
- Damiano Spina. Responsible AI From the Lens of an Information Retrieval Researcher: A Hands-On Tutorial. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval*, CHIIR '25, New York, NY, USA, 2025. Association for Computing Machinery.
- Damiano Spina, Mark Sanderson, Daniel Angus, Gianluca Demartini, Dana McKay, Lauren L. Saling, and Ryen W. White. Human-AI Cooperation to Tackle Misinformation and Polarization. *Commun. ACM*, 66(7):40–45, 6 2023. ISSN 0001-0782. URL <https://doi.org/10.1145/3588431>.
- Damiano Spina, Danula Hettiachchi, and Anthony McCosker. 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, 2024. URL <https://doi.org/10.60836/k8xh-en92>.
- S. S. Stevens. Mathematics, Measurement, and Psychophysics. In *Stevens' Handbook of Experimental Psychology*, pages 1–49. Wiley, Oxford, England, 1951. URL <https://doi.org/10.1002/9781119170174>.
- S. S. Stevens. *Psychophysics: Introduction to its perceptual, neural, and social prospects*. Psychophysics: Introduction to its perceptual, neural, and social prospects. John Wiley & Sons, Oxford, England, 1975. URL <https://awspntest.apa.org/record/1975-20087-000>.

-
- Jerry Swan, Steven Adriaensen, Alexander E.I. Brownlee, Kevin Hammond, Colin G. Johnson, Ahmed Kheiri, Faustyna Krawiec, J.J. Merelo, Leandro L. Minku, Ender Özcan, Gisele L. Pappa, Pablo García-Sánchez, Kenneth Sörensen, Stefan Voß, Markus Wagner, and David R. White. Metaheuristics “In the Large”. *European Journal of Operational Research*, 297(2): 393–406, 2022. URL <https://doi.org/10.1016/j.ejor.2021.05.042>.
- Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer Nature, Cham, Switzerland, 2022. ISBN 9783030343712. URL <https://doi.org/10.1007/978-3-030-34372-9>.
- Paul J van Diest, Rachel N Flach, Carmen van Dooijeweert, Seher Makineli, Gerben E Breimer, Nikolas Stathonikos, Paul Pham, Tri Q Nguyen, and Mitko Veta. Pros and cons of artificial intelligence implementation in diagnostic pathology. *Histopathology*, 84(6):924–934, 2024. URL <https://doi.org/10.1111/his.15153>.
- Jacky Visser, John Lawrence, and Chris Reed. Reason-Checking Fake News. *Communications of the ACM*, 63(11):38–40, 10 2020. ISSN 0001-0782. URL <https://doi.org/10.1145/3397189>.
- Lulu Wang. Early Diagnosis of Breast Cancer. *Sensors*, 17(7), 2017. ISSN 1424-8220. doi: 10.3390/s17071572. URL <https://www.mdpi.com/1424-8220/17/7/1572>.
- Minjuan Wang, Haiyang Yu, Zerla Bell, and Xiaoyan Chu. Constructing an Edu-Metaverse Ecosystem: A New and Innovative Framework. *IEEE Transactions on Learning Technologies*, 15(6): 685–696, 2022. URL <https://doi.org/10.1109/TLT.2022.3210828>.
- Raymond Williams. *Keywords: A Vocabulary of Culture and Society*. Oxford University Press, Oxford, UK, 2014. ISBN 9780199393213. URL <https://www.vitalsource.com/products/keywords-raymond-williams-v9780199393237>.
- Michael Wray, Hazel Doughty, and Dima Damen. On Semantic Similarity in Video Retrieval. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3649–3659, 2021. URL <https://doi.org/10.1109/CVPR46437.2021.00365>.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5288–5296, 2016. URL <https://doi.org/10.1109/CVPR.2016.571>.
- Kelvin Xu, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, attend and tell: neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, page 2048–2057. JMLR.org, 2015. URL <https://doi.org/10.5555/3045118.3045336>.
- Shanwen Yang, Tianrui Li, Xun Gong, Bo Peng, and Jie Hu. A review on crowd simulation and modeling. *Graphical Models*, 111:101081, 2020. ISSN 1524-0703. URL <https://doi.org/10.1016/j.gmod.2020.101081>.

A Activity Handouts

This appendix presents the outlines of the five activities given to the course participants.

A.1 Activity 1 – Research Topic

1. Write a couple of sentences about a research topic you're currently working on (e.g., your PhD Thesis).
2. Reflect on how your identity as a scholar may influence the way you're conducting the research.
3. Discuss your reflection with your peers.
4. Draft a positionality statement relevant to your research.

A.2 Activity 2 – Keywords

1. Layer 1 – Gather and curate:
 - Gather *keywords*, cluster and roughly organize them around themes, issues ... semantically;
 - Identify terms that carry the most tension, polysemy, are the most iconic or feel like *stars of the show*... those with large affective valence and complex feels...
2. Layer 2 – Scrutinize:
 - Nominate a set of landmine terms that generated tensions and uncertainty.
 - Scrutinize: Why those words? Who speaks on their behalf? What are their interests? What are the discontinuities between their technical or disciplinary and ordinary use? Who is included or excluded through their use?
3. Layer 3 – Translate:
 - How to we build inclusive dialogues around all this in contexts that will count?
 - What kinds of tools and materials do we need to produce?
 - What next steps do we want to take?

A.3 Activity 3 – Mixed Methods

1. **Case Studies:** Discuss and characterize with peers how your research could benefit from participatory methods.
 - What type of activities would you do?
 - What would you expect to find out?
2. Write a summary of the discussion.

A.4 Activity 4 – Fairness and Diversity

1. Discuss with your peers:
 - The concept of fairness and diversity in our case studies.
 - Evaluation in relation to the problem and the dimensions of effectiveness, diversity, and fairness.
2. Write a summary of the discussion.

A.5 Activity 5 – Ethical Considerations

1. Discuss with your peers what are the benefits and risks of this type of research.
 - Why is this research needed?
 - Who benefits from it?
 - What can go wrong?
 - How do we manage risk?
2. Pick one of the case studies/research topics and discuss the ethical considerations.
3. Write a summary of the discussion.

B Research Topics

This appendix presents the content provided by the participants in each group for Activity 1 (see also Appendix [A.1](#) and Section [3](#)).

B.1 Group: Marram/Kangaroo

B.1.1 Artificial Intelligence and Data Analysis for Personalized Oncology (Mubashara Rehman)

Abstract. For planning of radiotherapy treatments of head and neck cancers, a Computed Tomography (CT) of the patients is usually employed. However, for head and neck cancer patients, the quality of standard CT based on X-rays generated with kilo-Voltage tube potentials is severely degraded by streak artifacts occurring in the presence of metallic implants such as dental fillings. Some radiotherapy devices offer the possibility of acquiring Mega-Voltage CT (MVCT) for daily patient setup verification which, because of the higher energy of X-rays used, is almost completely free from artifacts and thus is more suitable for radiotherapy treatment planning. We leverage the advantages of kVCT scans with those of MVCT scans (artifact-free). We propose developing a deep learning-based approach capable of generating artifact-free MVCT images from acquired kVCT images.

Positionality Statement. I did not have any previous academic and professional experience in medical imaging, I studied Software Engineering and then worked in corporate as a Project Manager and previously as a Quality Assurance officer. When I started my Ph.D. on this topic, things were difficult initially, but eventually, with the help of our clinicians, all ambiguities were clarified. I do not have any personal biases affecting my research at all. Everything is going very well so far. However, the ethical concerns of medical data availability and patients' data privacy laws influence the research progress in the medical domain.

B.1.2 Models and Algorithms for the Intelligent Management of Waste Collection Through Fleets of Electric Vehicles (Francesco Taverna)

Abstract. Due to the increase in waste production and its impact from both an economic and ecological point of view, the efficiency in waste collection has become a key aspect in the transition toward the so-called smart green cities. The use of electric fleets can reduce the ecological impact, but it also poses new challenges. The vehicle routing problem is a well-known issue in the optimization framework, but its variant with electric vehicles brings two new challenges: the driving range and long recharge times. This research project focuses on optimizing the routes of electric vehicles using both exact and meta-heuristic approaches.

Positionality Statement. My background is more closely aligned with the theoretical aspects of operations research; therefore, the results may lack insights from an operational context. Specifically, I am not an expert on workers' rights in waste collection.

B.1.3 Machine Learning for Decision Support in Image Interpretation in Pathology (Hafsa Akebli)

Abstract. Whole slide images (WSIs) provide detailed gigapixel data crucial for disease diagnosis in digital pathology. However, their analysis is time-consuming for pathologists. This research applies deep learning, including convolutional neural networks (CNNs), vision transformers, and graph neural networks, to assist in WSI analysis with a focus on disease pattern detection. The study also aims to develop a generalized model capable of handling variations in staining, magnification, and other differences in WSIs. By incorporating these techniques, the research seeks to reduce the workload of pathologists, improve diagnostic accuracy, and enhance patient outcomes.

Positionality Statement. As a PhD student with a generalist industrial engineering degree, my training has equipped me with strong problem-solving and technical skills. My previous professional experience includes working as a supply chain controller and an operations project manager, roles that enhanced my analytical and organizational capabilities. While I do not have a background in the medical field and am not a specialist in deep learning, I recognize the transformative potential of AI in healthcare, especially after witnessing the challenges faced during the COVID-19 pandemic. I am motivated by the opportunity to reduce false diagnoses, improve pathology workflows, and contribute to better patient outcomes. Through my research, I aim to support advancements in digital pathology and foster innovation in healthcare technology.

B.1.4 Computational Methods for Intertextuality (Alessandro Tremamunno)

Abstract. Intertextuality is a concept that involves exploring connections between literary works, spanning from direct quotations to highly allusive references across varying scopes (from quoting single sentences to entire works that reference previous ones, such as Joyce's *Ulysses* and Homer's *Odyssey*). My research focuses on developing computational methods to create tools for humanities researchers, enabling them to more effectively discover and analyze intertextual connections.

Positionality Statement. My educational background is primarily scientific; I graduated from a scientific high school and focused on AI during my Master's degree. However, I also studied humanistic subjects, which I have always been passionate about. Working in Digital Humanities benefits from a predisposition towards humanistic subjects, yet a gap remains between my background as a computer scientist and that of a humanist. Furthermore, my specialization in Artificial Intelligence may sometimes lead me to exclusively consider AI methods as viable options, potentially dismissing other approaches prematurely.

B.1.5 Hybrid-Human-In-The-Loop for Misinformation Detection (David La Barbera)

Abstract. Developing techniques to combine NLP and crowdsourcing to deploy hybrid human-in-the-loop systems capable of identifying misinformation at scale [La Barbera et al., 2022].

Positionality Statement. As a computer scientist working in the field of misinformation detection using both NLP and crowdsourcing, I'm aware that my mental model can affect my research. Specifically, I may lack awareness of how human biases subtly influence AI models and the collected data, which can then be propagated and affect system performance and results.

B.1.6 Summary of Group Discussion

The group discussion shows that we all have a computer science background and we lack domain-specific knowledge of the data/ information that we are considering for our research. This can become an obstacle when planning computationally-driven solutions for problems in fields of research that extend beyond the realm of informatics.

B.2 Group: Warin/Wombat

B.2.1 Automatic Coding of Clinical Documents: Leveraging Symbolic and Sub-Symbolic Approaches for Enhanced Interpretability and Explainability (Mihai Horia Popescu)

Abstract. The healthcare industry is experiencing a rapid increase in medical data, driven by the widespread use of electronic health records and applications generating continuous patient data [Dash et al., 2019]. A significant portion of this data is unstructured, posing challenges for effective utilization. Standardizing the meaning of such textual data is essential for automated processing, decision-making, and the implementation of public health policies [Raghupathi and Raghupathi, 2014]. This is typically achieved through the coding and classification of medical text using standardized terminologies and classifications such as the International Statistical Classification of Diseases and Related Health Problems (ICD).

Traditionally, medical coding has been performed manually by trained professionals, but it is labor-intensive, time-consuming, and error-prone. To improve this process, two main approaches have been used: symbolic techniques, which rely on logic and human-defined rules for deterministic outcomes, and subsymbolic techniques, which use machine learning and neural networks to handle large, noisy datasets and adaptively learn from data patterns [Popescu et al., 2022a].

My research focuses on advancing the automation of clinical coding through both symbolic and sub-symbolic techniques, with particular attention to enhancing the interpretability of sub-symbolic methods. Specifically, I am working on the following topics:

- Developing rule-based systems for the automated selection of the Underlying Cause of Death (UCOD) using classification-independent rules [Popescu et al., 2021].
- Exploring the coding of death certificates using sub-symbolic approaches by investigating various machine learning techniques to enhance UCOD selection [Della Mea et al., 2020; Della Mea et al., 2020].
- Enhancing the interpretability and explainability of deep learning models to produce more transparent models without sacrificing prediction accuracy, thereby enabling hu-

man users to understand, trust, and effectively manage AI systems in clinical contexts [Popescu et al., 2022b].

Positionality Statement. Working with experts in the statistical and medical domains presents unique challenges in aligning on priorities for system implementation and improvement, particularly regarding the interpretability and explainability of results. As I focus on symbolic and sub-symbolic techniques, I face the added complexity of developing a rule-based system without extensive domain knowledge. Additionally, the explanations of the rules are often tailored for medical practitioners and are incomplete, which complicates the translation into terms comprehensible and actionable for a computer scientist.

B.2.2 Making Combinatorial Optimization Accessible: Real-world Applications, Tools, Explainability, Replicability, and Comparability (Francesca Da Ros)

Abstract. Metaheuristics are high-level, problem-independent strategies designed to efficiently explore search spaces and find near-optimal solutions [Blum and Roli, 2003]. Despite numerous successful applications, metaheuristics face criticisms related to explainability, replicability, and comparability [Swan et al., 2022; Sörensen, 2015]. My research aims to explore and evaluate methods in the field of metaheuristics to enhance and ensure the transparency of the search processes. My research is multifaceted and involves the following steps:

1. Analysis of real-world combinatorial optimization problems; my focus is on the Oven Scheduling Problem, a recently introduced parallel batch scheduling problem [Lackner et al., 2023].
2. Modeling the problem using a variety of algorithms, including Simulated Annealing [Kirkpatrick et al., 1983] and Large Neighborhood Search [Shaw, 1998].
3. Component-based analysis to examine problem-specific components and existing metaheuristics, with a focus on general algorithmic components.
4. Instance Space Analysis [Smith-Miles and Muñoz, 2023] and algorithm selection to determine the most suitable algorithms for different instances.
5. Development of tools to enhance replicability and comparability of metaheuristics, using C++ and Julia.
6. Application of Large Language Models in combinatorial optimization

Positionality Statement. Since my background is in information engineering and management and my primary interest is in the algorithmic aspects of combinatorial optimization, I may overlook application-specific details of the combinatorial optimization problem. For instance, when defining the weights of the components of a cost function, I might tend to miss important details. Therefore, collaborating with experts in the application field is essential.

B.2.3 Cellular-based Geolocation (Alberto Marturano)

Abstract. Accurate and reliable positioning of a device is essential for numerous applications like communications, emergency services, asset tracking, and vehicular navigation. While Global Navigation Satellite Systems (GNSS) provide high accuracy in open-sky environments, their

performance degrades significantly in challenging settings such as urban canyons or indoor locations due to obstacles like buildings and walls that block or reflect GNSS signals. Consequently, alternative positioning solutions are increasingly necessary. Geolocation of a device using cellular signals is a valid alternative and can be achieved opportunistically through various techniques such as fingerprinting and trilateration, enhanced by machine learning. My research focuses on these techniques, aiming to improve the accuracy and reliability of cellular-based geolocation methods in scenarios where GNSS is not available or reliable.

Positionality Statement. My research project is in collaboration with u-blox, a company specializing in semiconductor and module creation. Additionally, I am currently employed by u-blox, which means there is both a theoretical academic interest and the practical objective of creating something beneficial for the company. My background is in Computer Science, and I have limited knowledge of the physical properties of digital signals and cellular protocols.

B.2.4 Machine Learning and Deep Learning for Forestry and Agriculture (Mehdi Fasihi)

Abstract. In my research, I apply AI to solve problems in forestry and agriculture. For instance, my current project focuses on estimating carbon storage using remote sensing data.

Positionality Statement. I have extensive experience in software engineering, and currently, there are no limitations to my work.

B.2.5 Summary of Group Discussion

The group has recognized a common challenge: its members hold degrees in computer science and information engineering, and their research spans various domains such as formal methods, natural language processing, and optimization. However, their projects often extend into specific application domains (e.g., medical field, ancient languages) where specialized knowledge is lacking. This knowledge gap becomes apparent when addressing field-specific issues and collaborating with specialists from these sectors. Clear communication and explanation of specific terminology are essential for effective collaboration, facilitating the bridging of expertise between computer science and the requirements of other fields.

B.3 Group: Gurrborra/Koala 🐼

B.3.1 Leveraging Human Intelligence to Fight Misinformation (Michael Soprano)

Abstract. The spread of online misinformation has important effects on the stability of democracy since the information that is consumed every day influences human decision-making processes [Visser et al., 2020]. Understanding which information to trust is pivotal for the whole society. Judging the truthfulness of the published information is a complex process which involves several activities that have been traditionally performed by expert fact-checkers, such as journalists. However, the sheer size of digital content on the web and social media and the ability to immediately access and share it, among other factors, have

made it difficult to perform timely fact-checking at scale. The rate at which such a problem propagates continues to increase, largely aided by the increasing popularity of social media platforms [Pennycook et al., 2021].

A different approach to deal with the misinformation spreading problem can be to rely on the vast amount of people who consume information. A crowd of non-expert judges could perform the fact-checking activity instead of the expert ones. Specifically, I am exploring three distinct lines of research related to the topic of misinformation spreading:

1. Are human assessors able to detect and objectively categorize online (mis)information? Can their judgment be compared and related to those of experts? What is the optimal environment for obtaining the best results when judging information truthfulness? Can a multidimensional notion of truthfulness be defined? [Soprano et al., 2024a; Roitero et al., 2023a; Draws et al., 2022; Ceolin et al., 2022; Roitero et al., 2023b; Soprano et al., 2021]
2. What is the impact of cognitive biases on human assessors when judging information truthfulness? Is it possible to detect this kind of bias? Are there countermeasures to mitigate their effects? Can a bias-aware judgment pipeline for fact-checking be defined? [Soprano et al., 2024b]
3. Can the truthfulness judgments collected be leveraged using machine learning-based approaches? Is it possible to develop an approach capable of predicting information truthfulness while generating a natural language explanation to support the prediction itself? Are machine-generated explanations useful for human assessors in enhancing their judgment of the truthfulness of information items? [Brand et al., 2022]

Positionality Statement. My background has been rooted in computer science since secondary school. Consequently, I feel that I may lack a deep understanding of human behavior, beliefs, thoughts, and actions. I tend to focus more on systems and technical aspects rather than broader processes. However, much of my work involves humans in processes, often with significant financial implications, and my research can impact their lives. I believe I currently lack an adequate “ethical toolkit” to fully address these challenges.

B.3.2 Quantum Machine Learning applications for Environment (Daniele Lizzio Bosco)

Abstract. Deep Learning has revolutionized numerous fields, driving breakthroughs in image recognition, natural language processing, and more. However, these advancements come at a high cost, with training large models requiring vast computational resources and significant resource consumption. A clear example is given by Large Language Models, which have a tremendous cost in terms of energy [Desislavov et al., 2023], but also of water consumption [Li et al., 2023].

Quantum Computing leverages the principles of quantum mechanics, such as superposition and entanglement, to perform computations beyond the reach of classical systems. This has given rise to Quantum Machine Learning, which integrates quantum algorithms with

classical methods, aiming to reduce the computational cost of training and deploying AI models by exploiting quantum advantages [Liu et al., 2024].

However, due to the current limitations of quantum hardware, practical applications remain largely unexplored. My research focuses on using hybrid quantum-classical models compatible with current and near-future quantum devices, investigating whether and how quantum machine learning models can effectively address real-world challenges and provide tangible advantages, both as in increased performances or in reduced computational costs.

Positionality Statement. Due to my technical background, I primarily focus on methods and theoretical foundations rather than specific applications. Consequently, when trying to apply a model to a specific task, I often lack domain-relevant information crucial for understanding real-world problems and their implications.

B.3.3 Enhancing Inference Efficiency in Backbone Models: Optimizing Architecture and Modules (Moritz Nottebaum)

Abstract. The goal is to increase efficiency in model inference through architectural or module adaptations. Results are assessed based on the actual inference speed on common hardware, relative to quantitative outcomes. Peak memory consumption serves as a secondary evaluation criterion. In this regard, there is a special focus on backbone models and downstream tasks utilizing these models, as efficiency insights in this domain have the greatest impact.

Positionality Statement. My work is directly related to improving tracking systems. Living in a democratic environment with strong personal rights, I am not overly concerned about being monitored by cameras, as this information cannot realistically be used to control or harm me. However, in countries like China, where even crossing a red traffic light as a pedestrian is tracked, advancements in tracking technology may leave a negative impression and may not align with the goals researchers aim to achieve.

B.3.4 Breast Cancer Sub-Type Classification from Whole Slide Images (WSI) (Muhammad Zaka-Ud-Din)

Abstract. Breast Cancer (BC) remains one of the leading causes of mortality among women worldwide, posing a significant public health challenge [Ferlay et al., 2024]. However, this concerning trend can be substantially mitigated through timely and accurate diagnosis. Early diagnosis and precise delineation of affected areas are critical not only for understanding the stage and subtype of the cancer but also for informing effective treatment strategies [Wang, 2017]. In this context, the tasks of identifying affected areas, classifying the subtypes, and segmenting the tumor area plays a pivotal role in providing physicians with the necessary insights to devise tailored and effective treatment plans, ultimately improving patient outcomes and survival rates.

Considering these facts, my PhD work is particularly focused on designing models to deal with the issues in early diagnosis and relapse detection. The designing process of such model considers the availability of computational resources and data diversity due to the diverse data sources in the medical domain. In addition to these considerations in this research

direction I always need to consider different ethical implications regarding data availability and its use [Herington et al., 2023].

Positionality Statement. As a multi-disciplinary PhD researcher focusing on Computer Vision applications in different domains, I have significant advantages over many application-focused computer scientists. This uniqueness of multi-disciplinary research allows me to apply my computer skills across various domains of everyday life, including medical, industry, remote sensing, biodiversity, and other industrial applications [Szeliski, 2022]. However, as my current research is focused on computer vision applications in the medical domain, ethical considerations regarding the availability of medical data impose potential limitations on the application of computer vision in the medical field. As in the medical field, data annotation and the privacy of individuals contributing to data collection are paramount. Due to privacy concerns, much of the data I work with cannot be shared or made publicly available, which limits peer validation and reuse. These practices are indispensable for fostering trust in AI systems and advancing responsible innovation in sensitive fields like healthcare [Herington et al., 2023].

B.3.5 Crowd simulation in Immersive Virtual Reality (iVR) (Massimiliano Pascoli)

Abstract. My research focuses on crowd simulation [Musse et al., 2021; Yang et al., 2020], a field with applications in training and safety simulations, videogames, cultural heritage [Lim et al., 2020], and more. Specifically, I explore the use of immersive Virtual Reality (iVR) for realistic crowd simulations, particularly in emergency and safety training scenarios. iVR allows us to recreate dangerous situations safely and present realistic stimuli to users, leveraging its potential to study perception, user affect, and behavior in high-fidelity environments. My PhD work is structured around three key areas that have to be considered combined together for a robust crowd simulation framework: i) performance and model optimization, ii) perception and behaviors, and iii) graphics and appearance. The first area focuses on efficiently simulating large numbers of autonomous agents, while the second investigates the psychological and perceptual aspects of how users interact with and are influenced by virtual crowds. The third area aims to identify effective visual representations for diverse humanoid agents. Additionally, I aim to apply these findings to real-world scenarios, bridging research with practical implementations that have intrinsic ethical implications.

Positionality Statement. As a computer scientist working in the field of crowd simulation, I recognize the ethical responsibilities inherent in designing systems that aim to emulate human behavior, including psychological and social dynamics, in virtual environments. My work, which utilizes immersive Virtual Reality (iVR) to simulate crowd scenarios, particularly for emergency and safety training, has significant implications for safety, human behavior research, and the broader societal applications of AI technologies.

A key consideration in my research is the representation of human data. I often rely on real-world data to create realistic simulations, which may inadvertently carry biases or fail to represent the diverse populations my models aim to simulate comprehensively. For example, much of the evaluation data I collect comes from university students aged 20-30, a demographic not fully representative of the broader population. This limitation could lead

to less generalizable models or inadvertently reinforce stereotypes or inaccuracies when applied to different cultural, demographic, or social contexts. I am committed to addressing these challenges by being transparent about the limitations of my datasets and exploring ways to include broader, more diverse data sources whenever feasible.

Furthermore, the privacy of individuals contributing to data collection is paramount. Due to privacy concerns, much of the data I work with cannot be shared or made publicly available, which limits peer validation and reuse. I strive to ensure that data is collected, stored, and analyzed in a manner that respects privacy and adheres to ethical standards, including anonymization and compliance with legal regulations such as GDPR. These practices are essential not only for protecting individuals but also for maintaining trust in AI systems.

I also acknowledge that cultural and contextual factors significantly influence crowd behaviors, and my current simulations may reflect biases inherent to the social environments I am most familiar with. For example, the social norms, movement patterns, and decision-making processes of crowds in Western contexts may differ markedly from those in other regions or cultures. This underscores the importance of adapting models to diverse populations and validating their applicability in different contexts to avoid reinforcing cultural biases or misrepresentations.

In designing iVR simulations, I am mindful of their potential impact on users. These simulations aim to recreate realistic and sometimes high-stress scenarios to study user behavior. While this can provide valuable insights, it also raises ethical concerns about the psychological impact on participants, especially in scenarios involving emergencies or threats. I strive to ensure that all studies are conducted with informed consent, clear debriefings, and support for participants as needed.

Overall, I am committed to conducting my research in a manner that prioritizes inclusivity, transparency, and respect for the individuals and communities it seeks to model and serve. By explicitly acknowledging these ethical considerations and actively working to mitigate them, I aim to contribute to the development of responsible AI systems that not only advance technological capabilities but also respect and reflect the diversity and complexity of human behaviors and societies.

B.3.6 Summary of Group Discussion

In our discussion, we realized that we come from different backgrounds, and that some of our research topics are more connected to real-world applications with clear ethical implications (e.g., biases in data, privacy concerns, etc.), while others are more technical. Nonetheless, all of us base our research on data that must be treated responsibly. This data originates from real-world individuals and, in some cases, may also be sensitive or personal. Finally, we believe that a multidisciplinary team is essential for improving the quality of research (due to the multiple perspectives brought into the research process) and for better assessing ethical issues.

B.4 Group: Dulai Wurrung/Platypus

B.4.1 An Edu-Metaverse Framework to Foster Knowledge Gain and Retention (Biagio Tomasetig)

Abstract. In recent years, the metaverse has gained interest as a potential digital domain for various applications, including education [Alfaisal et al., 2024]. Within this scope, Edu-Metaverses are specifically designed for educational purposes [Wang et al., 2022]. These systems offer an innovative learning process, allowing students to acquire knowledge at their own pace and tailored to their needs. Effectively designing online shared virtual environments requires identifying critical factors influencing and balancing user experience, engagement, and learning outcomes regarding knowledge gain and retention.

Positionality Statement. As educational institutions adopt new technologies, Edu-Metaverses will contribute to formal and informal learning contexts. Therefore, I aim to investigate how Edu-Metaverses can facilitate knowledge acquisition, improve learner engagement, and enhance long-term retention. These technologies can significantly enhance learning quality in various environments, and Edu-Metaverses specifically address diverse learning needs while promoting collaboration, creativity, and critical thinking [Wang et al., 2022]. Furthermore, integrating Artificial Intelligence techniques into Edu-Metaverses can optimize content delivery, foster interactivity, and enhance the overall learning experience, making education more impactful and accessible. They could contribute to a more inclusive and valuable educational experience as essential tools for educators and learners.

From my side, it is crucial to research Edu-Metaverse systems through comprehensive evaluations and user testing that examine their impact on learners and educators. My research evaluates various factors within these virtual environments, considering both the learner's and the teacher's perspectives. These include usability, engagement, user experience, accessibility, and educational outcomes for students, as well as teaching effectiveness, ease of use, adaptability, and support for instructional strategies for educators. By testing my systems with diverse user groups - including individuals of varying ages, cultural backgrounds, and educational levels - I aim to check that the Edu-Metaverse is inclusive and effective for all users. Also, I am committed to addressing these technologies' ethical and social concerns, such as privacy, equity of access, and the potential for over-reliance on digital environments. These considerations are vital to fostering a responsible and equitable implementation of Edu-Metaverses in education.

I envision the Edu-Metaverse as a transformative tool that, when designed thoughtfully and implemented ethically, has the potential to revolutionize education. This resource could empower teachers and learners, enrich education, foster inclusivity, and meet everyone's unique needs. As we look to the future, the Edu-Metaverse could become a necessary tool, helping to create a more accessible, engaging, and effective educational landscape for future generations.

B.4.2 Data Structures for Storing and Working on Large Graphs for Genetic/Biological Data (Francesco Nascimben)

Abstract. Graphs can be used to represent and manipulate knowledge stemming from biology (e.g. protein structures [Fasoulis et al., 2021], metabolic pathways [Rosselló and Valiente, 2004]) and genetics (e.g. genomic distance [Braga et al., 2022], pangenomes [Garrison et al., 2023]): the relevance of this specific sub-field of graph theory keeps growing stronger, as the amount of data and computational power available to computer scientists and biologists increases. My research project aims to develop new techniques for efficiently dealing with the above-mentioned data, either by developing new algorithms and data structures for already established graph encodings of biological/genetic data or by defining entirely new representations.

Positionality Statement. Being a computer scientist with somewhat of a mathematical background, I'm fascinated by the theoretical problems my research topic poses, encompassing both graph theory and string compression (nucleotide strings being the standard for genome representation). However, a number of practical issues related to biology and genomics must also be taken into account to design efficient and robust techniques, such as errors occurring during the sampling process (due to limitations of state-of-the-art tools, e.g. long read sequencing [Salmela et al., 2016]) or bias in the data itself (for instance, it is known that most Genome-Wide Association Studies are currently ethnically biased towards people of European ancestry [Landry et al., 2018], thus not being adequately representative of mankind as a whole). Constant dialogue with biologists and other experts in the field is thus required, in order for my work to properly account for these issues.

B.4.3 Using Semantics for Vision-and-language Understanding (Alex Falcon)

Abstract. The main question in my research involves understanding how to link visual concepts with textual ones. This affects how an automated tool interacts and understands the contents of visual data (e.g. to generate descriptions of the visual contents [Xu et al., 2015], or rank the data based on their relevance to a textual query [Dong et al., 2021]). While this problem has been studied for several years now, there are still many shortcomings in existing solutions [Wray et al., 2021]. Specifically, I'm working on how to use semantic nuances of text (e.g. synonyms) to reduce the domain gap existing between the multiple modalities [Falcon et al., 2024]. An expected outcome is to improve the understanding capabilities of multimodal systems since, for instance, existing systems often neglect that two different yet equivalent descriptions can both be important for the same video.

Positionality Statement. I have a CS background with a focus on AI and machine learning, so it is common to analyze the data and based on that decide how to perform the modeling and subsequent learning. However, data itself is typically the problem. The annotations used in my field are typically created by non-expert humans (e.g., textual annotations are either represented by descriptions of the visual contents provided by crowdsource workers [Xu et al., 2016], by human-provided alt-text, or by automatically generated subtitles [Miech et al., 2019]), resulting in descriptions that are often unrelated (think of philosophical quotes in Instagram posts), incomplete, vague, and imprecise. While these represent biases and

limitations out of my control (to some extent), there are other biases and limitations that are due to the techniques developed in my research. First, by working on linguistic nuances, etc, I'm assuming that the main problem in cross-modal understanding is strictly related to it. Second, I'm by no means a linguist, so more often than not I'm relying on publicly available tools which have their own biases to identify and deal with synonyms or hypo/hypernyms, etc in the data. Third, while doing so, many other aspects of the data are neglected, e.g. even when rephrasing a caption in multiple different ways, there are always aspects of the video not initially present in the caption, and so neglected from there on.

B.4.4 Risk Association of Relapse for Distant Metastasis in Breast Cancer Patients with the Immune Response of Primary and Axillary Tumor Using Digitized Histopathological Images Analysis (Alessio Fiorin)

Abstract. The work aims to identify specific immune response patterns in primary tumor and axillary lymph nodes that correlate with survival and relapse outcomes in breast cancer patients [López et al., 2021, 2020; Liu et al., 2021]. The research leverages artificial intelligence techniques to analyze immune biomarkers' quantification, localization, morphology, and functions extracted from digitized histopathological images.

Positionality Statement. I am a computer scientist specializing in artificial intelligence and cybersecurity. While I do not have a background in histopathology and immunology, I can bring a technical perspective. I enjoy collaborating with teams from diverse fields since different points of view can lead to innovative and challenging projects. I am aware of the difficulties regarding the communication issues that can arise in multidisciplinary teams due to differences in terminology and expertise. To address this, I aim to develop advanced skills in histopathological image analysis and facilitate and make more fluent work in multidisciplinary teams to carry out projects helpful for people's healthcare.

B.4.5 Artificial Intelligence for Decision Support in Pathology (Laura Rasotto)

Abstract. The main objective of the research project deals with the study of methodologies and techniques for decision support in pathology [Perez-Lopez et al., 2024]. The development of digital pathology [Omar et al., 2024], thanks to microscopy images and the most recent gigabyte-sized Whole-Slide Images (WSI) [Jain et al., 2024], offers an opportunity to extract richer insights from histological samples. This is made possible through advanced image analysis techniques and the recent developments of deep learning [Försch et al., 2021; Moxley-Wyles et al., 2020]. Thus, starting from an analysis of histopathological images and the understanding of the various phases of the pathologists' work [Hosseini et al., 2024], the goal is not only to develop artificial intelligence systems and techniques that can carry out the tasks of tissue segmentation and classification [Baroni et al., 2024], but also to develop a methodology that can integrate image analysis models within pathology information systems.

Positionality Statement. My position as a computer scientist places me between new technologies and healthcare. Until now, my understanding of the real daily work of pathologists and the challenges they encounter has been mainly obtained through literature reviews [van

[Diest et al., 2024](#)]. In the future, I will collaborate with pathologists and colleagues who are already familiar with this field. Hopefully, they will arrive from different backgrounds, making it possible to share different points of view. Additionally, people and situations that surround me, such as close friends diagnosed with cancer, give me the opportunity to strongly believe in the importance of connecting my interests in computer science to the medical field. Another important aspect is my commitment to raising awareness about the importance of sharing medical data for research purposes. When I explain my work to others outside the field, I emphasize how the availability of various and huge datasets is fundamental for the development of effective models in digital pathology. For now, I make an effort to decrease the lack of knowledge of the people, highlighting the potential advantages of responsibly sharing their medical data.

B.4.6 Conversational Agent AI-Based for Public Administration: Theory, Design, Implementation, and Evaluation (Riccardo Lunardi)

Abstract. Conversational Agent AI-Based for Public Administration: Theory, Design, Implementation, and Evaluation The research explores the deployment of conversational agents in online public administration services, with a focus on reducing the digital divide among vulnerable populations, particularly older adults and individuals with disabilities. The study primarily aims to enhance accessibility and usability in the healthcare sector [[Garimella et al., 2024](#)], addressing barriers that limit equal access to essential services. There are existing studies focused on similar goals within slightly different contexts [[Lunardi and Coppola, 2024](#)], such as improving accessibility in education [[Khosrawi-Rad et al., 2022](#)] and customer services [[Bavaresco et al., 2020](#)]. Recognizing the potential biases in both human design choices and the underlying machine learning models [[Lunardi et al., 2024](#)], special attention will be given to identifying and mitigating these biases. Evaluation metrics, including usability, task success rates, and user satisfaction, will ensure the agents provide meaningful support to all users.

Positionality Statement. My research focuses on improving the accessibility of healthcare online services through conversational agents, targeting vulnerable populations such as older adults and individuals with disabilities. My personal experiences have shaped this focus: I have observed challenges faced by older adults in navigating digital systems, particularly within my family and my personal social contexts. These experiences have heightened my sensitivity to their needs, potentially overlooking the challenges faced by others. By integrating feedback from diverse perspectives and critically examining both human and model biases, I aim to balance personal experiences with a rigorous, inclusive research methodology that prioritizes equitable outcomes for all users.

B.4.7 Summary of Group Discussion

The group discussion highlighted three main factors that influence our research, i.e., internal bias, external expectations, and personal values. We follow our values and beliefs, e.g., honesty, creativity, integrity, respect, and ambition, in everyday life scenarios, as we do in our research. The main challenge is dealing with our biases, which depend on our background, culture, and

experiences. Sometimes, we feel conditioned by people close to us, like family, relatives, friends, and colleagues. We should be open-minded and actively question our biases by listening to others' perspectives. At the same time, we should embrace constructive ideas consistent with our ideals, especially those promoting virtuous behavior.

C Keywords

This appendix presents the content provided by the participants in each group for Activity 2 (see also Appendix A.2 and Section 4).

C.1 Group: Marram/Kangaroo

Table 1. Keywords brainstormed by the members of the Marram (Kangaroo) group.

Layer 1: Gather and Curate	Layer 2: Scrutinise	Layer 3: Translate
<div>Credibility</div> <div>Ethics</div> <div>Equality</div> <div>Privacy</div> <div>Bias</div> <div>Reliability</div> <div>Fairness</div> <div>Open</div> <div>Healthcare</div> <div>Quality Control</div>	<div>Trust</div> <div>Explainability</div> <div>Transparency</div> <div>Human-centered</div> <div>Safety</div>	<div>Model Evaluation</div> <div>Educational Content</div> <div>Data Regulations</div>

The texts provided for Layer 3 are summarized into keywords as follows:

- **Model Evaluation:** “Create limit cases and build instances that are evaluated with completely uncertain results”.
- **Educational Content:** “Educational content designed for widespread use”.

C.2 Group: Warin/Wombat 🐨

Table 2. Keywords brainstormed by the members of the Warin (Wombat) group.

Layer 1: Gather and Curate	Layer 2: Scrutinise	Layer 3: Translate
<div>Diversity</div> <div>Culture</div> <div>Sustainability</div> <div>Inclusion</div> <div>Equality</div> <div>Accountability</div> <div>Fairness</div> <div>Explainability</div> <div>Impact</div> <div>Anonymity</div> <div>Ethics</div> <div>Trust</div> <div>Control</div>	<div>Transparency</div> <div>Safety</div> <div>Privacy</div> <div>Algorithm</div>	<div>Avoid Metaphors</div> <div>Explain With Examples</div> <div>Syllabus</div> <div>Shareholder Discussion</div>

The texts provided for Layer 3 are summarized into keywords as follows:

- **Avoid Methaphors:** “Avoid using metaphors, as they could be used to justify the mathematical or computational processes”.
- **Explain with Examples:** “Explain with examples and images, as they are easily understandable”.
- **Syllabus:** “A syllabus with simple, clear definitions to establish a shared point of view”.
- **Shareholder Discussion:** “As a first step, the project shareholders should discuss and reach an agreement, sharing their points of view”.

C.3 Group: Gurrborra/Koala 🐼

Table 3. Keywords brainstormed by the members of the Gurrborra (Koala) group.

Layer 1: Gather and Curate	Layer 2: Scrutinise	Layer 3: Translate
<div>Aggregation</div> <div>Data Completeness</div> <div>Generalization</div> <div>Society</div> <div>Reward</div> <div>Ethical Concerns</div> <div>Data Availability</div> <div>Behavior</div> <div>Expertise</div> <div>People</div> <div>Bias</div>	<div>Metrics</div> <div>Social Impact</div> <div>Trust</div> <div>Consent</div>	<div>Meta-studies Quality</div> <div>Metrics Social Impact</div> <div>Trust Consent Levels</div> <div>Tools Data Access</div>

The texts provided for Layer 3 are summarized into keywords as follows:

- **Meta-studies Quality:** “Meta-studies on how data is collected to assess its quality.”
- **Metrics Social Impact:** “We need actual metrics to measure the social impacts of technology in both the short and long terms, as social changes can occur rapidly and persist for extended periods.”
- **Trust Consent Levels:** “Different levels of trust and consent must be defined: legal, professional, ethical, etc.”
- **Tools Data Access:** “Tools that allow people to have direct access to the data they provide or that is being used by researchers.”

C.4 Group: Dulai Wurrung/Platypus 🦋

Table 4. Keywords brainstormed by the members of the Duali Wurrung (Platypus) group.

Layer 1: Gather and Curate	Layer 2: Scrutinise	Layer 3: Translate
Background	Bias	Universal Translator
Family	Values	
Friends Responsibility	Personal Experience	Semantic Summarized
Teamwork Goals	Expectations	
	Open-minded	Translator Tool
Culture		Keyword Segmentation

The texts provided for Layer 3 are summarized into keywords as follows:

- **Universal Translator:** “The implementation of a universal background translator (a practical application of the Feynman method), where technical or specialized information is transformed using vocabulary from the universal background language, so that even the most technical concepts can be understood by everyone through unambiguous terminology”.
- **Semantic Summarizer:** “A semantic summarizer that can standardize and simplify a complex and technical text”.
- **Translator Tool:** “Background translation tool – a universal background tool”.
- **Keyword Segmentation:** “An application that, given an input (text or speech-to-text), segments keywords and, on hover, provides explanations (definitions, related keywords, etc.), not necessarily using terminology from other fields”.

D Case Studies

This appendix presents the content provided by the participants in each group for Activities 3–5 (see also Appendix A.3, Appendix A.4, Appendix A.5, and Section 5).

D.1 Diagnosis of Cancer and Enhancing Radiotherapy Treatment: A Deep Learning Approach (Group: Marram/Kangaroo)

Aim. Using AI tools for detecting instances of cancer from CT scans and the subsequent use of crowdsourcing as a means to check the reliability of these results.

Problem Statement. For planning radiotherapy treatments for head and neck cancers, we have datasets from the XYZ Cancer Research Institute, Kilo-Voltage CT (KVCT), and Mega-Voltage CT (MVCT). In head and neck cancer patients, the quality of standard CT images, based on X-rays generated with kilo-voltage tube potentials, is significantly degraded by streak artifacts caused by metallic implants such as dental fillings. Some radiotherapy devices provide the capability to acquire MVCT scans for daily patient setup verification. Due to the higher energy of the X-rays used in MVCT, these scans are almost entirely free from artifacts, making them more suitable for radiotherapy treatment planning.

The CT dataset is processed using advanced computer vision AI tools. During this stage, parameters such as color, resolution, and size are normalized and optimized. The processed CT images are then divided into slices or patches, which serve as inputs for deep learning models, including Convolutional Neural Networks (CNNs) and Transformers. These models are trained on ground truth labels representing cancer tumor grades to predict the labels of the test data.

After labeling the CT dataset, expert pathologists review the AI predictions within a highly trained crowdsourcing framework. This step ensures accuracy and reliability by providing a double-check mechanism conducted by experienced professionals.

Research Questions.

- RQ1 To what extent are metal artifacts effectively removed from the dataset, and how does their removal impact the performance of subsequent models?
- RQ2 What methodologies does AI employ to process CT images, and how efficient is it in detecting and grading cancer tumors?
- RQ3 How does AI's performance in detecting and grading cancer tumors compare to that of human experts, and how is this performance perceived by medical professionals?
- RQ4 What limitations and challenges does AI face in processing CT images and detecting cancer tumors?
- RQ5 How do variations in preprocessing steps (e.g., normalization, resolution optimization) affect the accuracy and reliability of AI predictions?
- RQ6 What are the implications of using crowdsourced expert validation for enhancing the reliability of AI predictions in cancer diagnosis?

Methodologies. To comprehensively address our problem statement, we employed three distinct methodologies, each tailored to a specific stage of the research:

1. Leveraging clinicians' expertise and Artificial Intelligence (AI) for denoising and removing metal artifacts from the dataset.
2. Utilizing deep learning tools for cancer cell detection.
3. Applying crowdsourcing techniques to verify the obtained results.

Positionality Statement. Our background is in computer science, and we lack direct medical expertise. Thus, our evaluation metrics for each of the three steps involved may need input from additional perspectives to help us define more qualitative measures to evaluate the outcomes. Our constraints include medical data quality and availability, computational resources, and model generalization. In addition to these, patient privacy and consent are crucial, as safeguarding patient privacy and obtaining informed consent for the use of medical data in research and AI model development are essential but challenging tasks.

Expected Outcomes. These multi-faceted methodologies combine the strengths of human expertise and advanced AI techniques, ensuring robust and reliable outcomes throughout the research process.

Fairness. For this case study, imbalanced datasets present the biggest challenge. This imbalance could introduce bias into the results. Therefore, the proportion of the artifact dataset in the training set must be equal to that of the clean dataset. If this is not possible, loss functions such as the Focal Loss Function can be used to mitigate the imbalance in the training dataset.

Diversity. In the mentioned case study, diversity includes variations in modality and in the region of interest. For our research to be unbiased, the AI model needs to be generalized.

Ethical Considerations.

- *Why is this research needed?* To enable early and accurate diagnosis of cancer cells.
- *Who benefits from it?* Patients through faster prognosis and better radiotherapy planning.
- *What can go wrong?* Potential misdiagnosis or reliance on biased data.
- *How do we manage risk?* Establishing robust validation and monitoring protocols.

Summary of Group Discussion

The core of the discussion revolved around the ethical considerations of using AI in cancer diagnosis and radiotherapy planning, mainly the safeguard of patient privacy, and ensuring responsible use of medical data. Validation of AI predictions by expert pathologists through crowdsourcing was recognized as essential for reliability. The discussion also addressed challenges in fairness and diversity, such as the presence of biases in imbalanced datasets, and the importance of providing patient data from different social as well as ethnic backgrounds. Nevertheless, the combination of human expertise and AI innovation can significantly enhance cancer diagnostics and treatment, if the challenges previously outlined can be properly addressed.

D.2 Geolocation Based on Radio Signals (Group: Warin/Wombat 🦘)

Aim. The goal is to clearly and effectively communicate the accuracy of the estimated position provided by a geolocation system.

Problem Statement. Many optimal solutions in the literature do not adequately address the issue of communicating geolocation accuracy. Displaying accuracy as a percentage can be confusing for users, and the concept of confidence in precision is often misunderstood. For instance, stating that “you are within 100 meters (estimated error) with 85% confidence (probability)” can be unclear.

Research Question.

RQ1 What are the best methods to communicate accuracy in geolocation effectively?

Methodologies. Provide participants with multiple representations of the same accuracy and ask them to evaluate and rank these representations based on their clarity and effectiveness in conveying the information. Additionally, provide participants with a specific metric, visualization, KPI, or other relevant measure, and ask them to interpret the information.

Positionality Statement.

- All authors have a background in computer science or information engineering.
- One author works in the geolocation services field.

Expected Outcomes.

- Users may prefer highly precise information on either confidence or range, but they may struggle to balance the trade-off between the two aspects.
- We aim to gain a better understanding of how users interpret geolocation accuracy data and identify the most effective and usable ways to represent this information. This should lead to improvements in how such data is communicated, making it more accessible and understandable for users.

Fairness.

- Ensure that the accuracy and confidence data are reliable, unbiased, and free from outliers that could introduce errors in both position and error estimation.
- Continuously use user feedback to enhance the reliability of error estimation procedures.

Diversity.

- Offer users multiple confidence interval options, allowing them to select the most appropriate one for their specific use case, rather than relying on a single, fixed confidence level.
- Enable users to set a threshold for the maximum acceptable uncertainty value for a given confidence level.

Ethical Considerations.

- *Why is this research needed?* Current solutions for communicating geolocation accuracy are suboptimal. As technology advances, research should keep pace to ensure usability and effectiveness.
- *Who benefits from it?* Anyone involved in creating or improving geolocation systems, as well as researchers working on these systems and end-users who rely on accurate and understandable geolocation information.
- *What can go wrong?* Misuse of positional information (e.g., stalking) or the perpetration of industrial espionage by selling relevant data (e.g., information about product tracking).
- *How do we manage risk?* Implement controlled access to geolocation systems and ensure user anonymity to protect against misuse of positional data.

Summary of Group Discussion

In our field, user-based studies are quite rare. However, we recognize their potential value. Engaging with users can facilitate the validation of our results, bridging the gap between technical research and practical application. This process provides insights that ensure our models and algorithms align with real-world needs. Moreover, collaboration helps present the final results in a more comprehensible and relevant manner.

Overall, by integrating user feedback and expert validation, we can achieve more robust and applicable outcomes in our research.

D.3 Longitudinal Studies Conducted on Crowdsourcing Platforms (Group: Gurrborra/Koala 🐼)

Aim. Understanding the barriers to running longitudinal studies on crowdsourcing platforms.

Problem Statement. Crowdsourcing tasks are commonly used to gather a large volume of human labels efficiently. Some tasks are single-use, performed only once by workers, while others involve repeated participation, known as longitudinal studies. Despite their frequent use in crowdsourcing, there is limited understanding of the factors that affect worker participation in such studies across various crowdsourcing platforms.

Research Questions.

- RQ1 What is the current perception of workers regarding longitudinal studies on commercial micro-task crowdsourcing platforms? How have their past experiences been, and what are their views on participating and committing to future studies? What characteristics do they prefer in a longitudinal study?
- RQ2 What recommendations should researchers and practitioners follow when designing and conducting longitudinal studies on commercial micro-task crowdsourcing platforms?
- RQ3 What best practices should commercial micro-task crowdsourcing platforms adopt to effectively support and improve the conduct of longitudinal studies?

Methodologies. The study will involve a mixed methods approach, combining both qualitative and quantitative techniques to provide a comprehensive analysis of participants' experiences. The methodology will include a survey, consisting of both open-ended and closed-ended questions. The open-ended questions aim to gather qualitative data, allowing participants to share their thoughts, feelings, and insights in their own words. This data will be analyzed qualitatively to identify recurring themes and patterns. The closed-ended questions, on the other hand, will collect quantitative data, offering measurable insights into participants' views and experiences. These responses will be analyzed using statistical methods to uncover trends and relationships across the sample.

Positionality Statement. This research aims to enhance participants' experiences in these studies by providing useful guidelines, considerations, and evidence for those who design and publish them.

Expected Outcomes. Recommendations and best practices for enhancing the design and execution of longitudinal studies on crowdsourcing platforms. These insights will focus on improving worker participation, increasing engagement over time, and fostering long-term commitment, with the goal of optimizing the overall effectiveness and sustainability of such studies.

Fairness.

- Double-check the desirable properties of the metrics used in the study, including an analysis of how the study's metrics align with these properties and comparing them against user satisfaction (e.g., through a pilot test).
- Ensure fair compensation for contributions, such as consistent hourly payments or rewarding rates, maintained throughout the duration of the entire study.

Diversity.

- Filters, statistical constraints, and sampling methods embedded in the crowdsourcing platform to ensure diversity.

Ethical Considerations.

- *Why is this research needed?* This research is necessary because there are currently no well-established guidelines, recommendations, or best practices for conducting longitudinal studies on crowdsourcing platforms.
- *Who will benefit from it?* People from around the world who are recruited on platforms will benefit from the formulation of fairer, well-designed longitudinal studies, as these studies require long-term commitment and effort. Additionally, those who design these studies will gain a broader understanding of workers' needs, fears, and expectations.
- *What can go wrong?* The risk is that we may develop an inaccurate perception of longitudinal studies on crowdsourcing platforms, especially if we recruit a sample that is too narrow or lacks sufficient diversity to extract statistically significant insights.

-
- *How do we manage risk?* We will manage this risk by replicating the study at a later time, possibly multiple times. If we find that the responses from different groups of individuals remain consistent, we will be able to confidently argue that we have accurately framed the perception of longitudinal studies.

Summary of Group Discussion

Research on the perception of longitudinal studies is necessary, as no standard guidelines currently exist for conducting such studies on crowdsourcing platforms. Fairer and better-designed studies will benefit global participants and provide researchers and businesses with valuable insights into workers' needs and user confidence. However, there is a risk of recruiting a non-diverse sample, which could lead to inaccurate perceptions. To mitigate this risk, the study will be repeated multiple times to ensure consistency and validity in the findings.

D.4 Age-based Conversational Agent For Healthcare Applications (Group: Dulai Wurrung/Platypus 🦘)

Aim. To present and interact with health applications using a conversational agent in an effective way for all ages.

Problem Statement. Digital healthcare applications are becoming indispensable in delivering medical services, yet their usability often falls short of engaging diverse age groups. Younger users may favor interactive and fast-paced interfaces, while older adults may face barriers due to unfamiliarity with technology or accessibility issues. This disparity in engagement can reduce the overall effectiveness of healthcare tools. To address this challenge, an age-based conversational agent is proposed to provide adaptive, personalized interactions tailored to the needs and preferences of users of different age groups, integrating a gamified user experience. The aim is to make healthcare information accessible, reliable, and engaging for all, while fostering trust and usability.

Research Questions.

- RQ1 How to present content to different age groups?
- RQ2 How to adapt the content to the age group e.g. terms used, pacing, mental load?
- RQ3 How to balance adaptability with healthcare communication reliability?
- RQ4 How to control the veracity of responses generated by the tool?
- RQ5 How to analyze the risk in answering a specific question (e.g. related to life-or-death situations) and decide whether the tool should answer or not?
- RQ6 How to evaluate the system to determine its real-world utility?
- RQ7 Which AI models to use? Is it better to use a LLM or to use a domain-specific model?

Methodologies. Our methodology will be based on a user-centric design, conducting in-person questionnaires and performing different user testing in real-world scenarios. We will use ML/DL models both to analyze the input question and other eventual input data, and to

predict user preferences and adjust interaction styles dynamically, e.g., varying the terminology used, adding voice modulation, using different models or LLMs. We will evaluate the system via qualitative and quantitative metrics, i.e., engagement, user satisfaction, usability and user experience (qualitative evaluation through user questionnaires), robustness, faithfulness, relevance of the answer to the input question (quantitative models metrics).

Positionality Statement. We are researchers with varied backgrounds coming from computer science, and this background leads us to develop and evaluate the system by solely focusing on objective metrics. We acknowledge that there are age-related differences in technology adoption and preferences, and we would like to bridge-the-gap via a technological solution. Our technical background coupled with a lack of knowledge in diverse professional fields including medical and psychological is severely limiting, possibly leading to a lack of empathy when communicating with the user. Therefore, this could greatly affect how we develop and evaluate both the conversational agent and the questionnaires.

Expected Outcomes. A conversational agent that can adapt its way to present content based on the kind of input it receives to different age groups.

Fairness. The conversational agent should be able to adapt to different age ranges. Thus, fairness is given by providing each user with the same information but with a different presentation strategy to improve the experience of the user and make the experience more engaging. To build a framework like that, we should collect user preferences about content presentations (e.g. vocal types, tone intensity, slang). So, we are able to build types of presentation modalities in order to optimize the conversation in an age-oriented way.

Diversity. We would like to create an inclusive conversational agent that meets the needs of users from various cultural, linguistic, and physical backgrounds. This includes incorporating i) multiple language support, ii) cultural sensitivity to local idioms, expressions and culturally relevant contents, and iii) accessibility features (e.g., text-to-speech for visually impaired users, visual cues, caption, sign language options for hearing impaired users, or simplified language for non-native speakers). The system will use age-sensitive presentation strategies, e.g., playful tones, formal, empathetic.

Ethical Considerations.

- *Why is this research needed?* Because the type of communication between people of different ages varies, as communication can be more or less expressive and with age-dependent terms.
- *Who Benefit from it?* People with illnesses who do not have access to the expertise of specialists in a certain period of time. Also, the general public can access its personal health portfolio.
- *What can go wrong?* It may happen that the conversational model misinterprets the user's queries due to the different types of dialogue. In addition, there may be biases due to the data used to train the conversational model. Subsequently, errors might occur in the responses of the conversational model, e.g. due to hallucination issues. For this reason, responses could be provided with references from research articles.
- *How do we manage risk?* The conversational agent bases its replies from the knowledge on the data used to train it, so it will be important to hide personal information via

data anonymization. Companies might also use the built models/solutions so we need a way to certify the AI models through data certification and a standardized protocol validation. Regarding the data, we need the patients' consent; they have to know that their data is used only for specific purposes explicitly agreed by the patient (e.g. only within the conversation and not used to increase the knowledge of the conversational model). The data should be stored and shared using secure databases with the best cybersecurity standards in order to reduce the probability of data leakage or data theft. Also, there should be a data protection regulation and agreement (like the GDPR) among stakeholders in order to assure data quality and privacy protection and avoid data misuse of the exchanged data.

D.4.1 Summary of Group Discussion

The group had an interesting discussion about using a conversational agent to manage communication between people of different ages, noting that this issue mirrors the challenges faced by physicians who adjust their vocabulary based on the patient. The agent would adapt its interaction style to user preferences, using AI models to personalize experiences and improve engagement, particularly for younger and older users, and adapting to user preferences. Key challenges include balancing adaptability with communication reliability, ensuring accuracy in critical situations, and evaluating system effectiveness. The agent will prioritize fairness by personalizing content without altering core information and incorporating multilingual support and accessibility features. Ethical considerations include data privacy, user consent, and ensuring accurate, reliable responses.

E Course Survey

This appendix provides the survey used to gather demographic information about the participants (questions 1–4), insights into their educational background and previous experiences (questions 5–13), evaluation of sessions and activities (questions 14–24), and general feedback about the overall course (questions 25–30).

The questions are shown in the order they were presented to the participants. The text of each question is in *italics*, the expected answer type is in normal font, and additional details are in monospaced font.

Demographic Information

1: *Age group* (closed-ended, radio button)

- 18–25
- 26–34
- 35–44
- 45–54
- 55–64
- Prefer not to say

2: *Gender* (closed-ended, radio button)

- Female
- Male
- Non-binary
- Prefer not to say

3: *Country of origin* (textual field)

- Non-empty text

4: *Languages you speak* (closed-ended, checkbox)

- Italian
- English
- Arabic
- Spanish
- Friulan
- German
- Portuguese
- Japanese
- French
- Russian
- Chinese

-
- Romanian
 - Other (please specify)

Studies and Experiences

5: *Research field* (closed-ended, checkbox)

- Human-computer Interaction
- Information Retrieval
- Machine Learning
- Optimization
- Formal Methods
- Other (please, specify)

6: *Which Year of your PhD you are currently in:* (closed-ended, radio button)

- Year 1
- Year 2
- Year 3
- Other (please, specify)

7: *Are you a visiting student at University of Udine?* (closed-ended, radio button)

- Yes
- No

8: *Are you doing a PhD with industry?* (closed-ended, radio button)

- Yes
- No

9: *Master's degree in:* (textual field)

- Non-empty text

10: *Bachelor's degree in:* (textual field)

- Non-empty text

11: *Do you have any work experience? If so, indicate number of years, role, etc.:* (textual field)

- Non-empty text

12: *How many user studies have you conducted in the past 5 years?* (closed-ended, radio button)

-
- 0
 - 1–2
 - 3–5
 - More than 5

13: *How familiar are you with the institutional/ethical review board (IRB) review process?* (closed-ended, radio button)

- Not familiar at all
- Slightly familiar
- Moderately familiar
- Very familiar
- Extremely familiar

Sessions and Activities

In this section, we present the instructions given to participants on how to answer questions 14–24 and their formulations.

Instructions

For each of the following questions, please provide a numerical rating to indicate your evaluation. You may use any numbers that seem appropriate to you – whole numbers, fractions, or decimals. However, you may not use negative numbers or zero.

Don't worry about running out of numbers – there will always be a larger number than the largest you use, and a smaller number than the smallest you use.

Try to judge each aspect in relation to the previous ones. For example, if you feel as half as satisfied with the current aspect of the course as with the previous one, then assign a score that is half of your previously assigned score.

Questions

14: *How would you rate the overall course experience?*

15: *How would your rate Session 1 – Responsible AI: A Multidisciplinary Problem – How Can We Contribute?*

16: *How would your rate Session 2 – Mixed Methods for Designing and Evaluating Presentation Strategies for Fact-checked Content?*

17: *How would your rate Session 3 – Fairness and Diversity: Two Sides of the Same Coin?*

18: *How would your rate Session 4 – Characterizing Information Processing Activities with Physiological Signals from Multiple Wearable Devices?*

19: *How would you rate the group activities overall?*

20: *How would you rate Activity 1 – Positionality?*

21: *How would you rate Activity 2 – Keywords?*

22: *How would you rate Activity 3 – Case Studies for Mixed Methods?*

23: *How would you rate Activity 4 – Fairness and Diversity?*

24: *How would you rate Activity 5 – Ethical Considerations?*

Course Feedback

25: *How likely are you to recommend this course to other PhD students or colleagues?* (closed-ended, buttons)

- 1 (Not at all likely)
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10 (Extremely likely)

26: *How effective were the following aspects regards to your experience in this course?* (closed-ended, buttons)

Instructional materials used in this course

- Extremely ineffective
- Somewhat ineffective
- Neutral
- Somewhat effective
- Extremely effective

Learning activities used in this course

- Extremely ineffective
- Somewhat ineffective
- Neutral
- Somewhat effective
- Extremely effective

Use of technologies in the class

-
- Extremely ineffective
 - Somewhat ineffective
 - Neutral
 - Somewhat effective
 - Extremely effective

Group activities organized after the class

- Extremely ineffective
- Somewhat ineffective
- Neutral
- Somewhat effective
- Extremely effective

27: *Did the course meet your expectation?* (closed-ended, radio button)

- No – Far Below Expectations
- No – Below Expectations
- Not Sure
- Yes – Met Expectations
- Yes – Above Expectations

28: *Describe what you liked the most from this course* (textual field)

- Non-empty text

29: *Describe what you liked the least from this course* (textual field)

- Non-empty text

30: *Any additional comments?* (textual field)

- Non-empty text

F Authors and Affiliations

This is the full list of the organizers and participants of the PhD course delivered by the first author, who contributed to this report:

- Damiano Spina (RMIT University, Australia)
- Kevin Roitero (University of Udine, Italy)
- Stefano Mizzaro (University of Udine, Italy)
- Vincenzo Della Mea (University of Udine, Italy)
- Farnesca Da Ros (University of Udine, Italy)
- Michael Soprano (University of Udine, Italy)
- Hafsa Akebli (University of Udine, Italy)
- Alex Falcon (University of Udine, Italy)
- Mehdi Fasihi (University of Udine, Italy)
- Alessio Fiorin (Universitat Rovira i Virgili, Spain)
- David La Barbera (University of Udine, Italy)
- Daniele Lizzio Bosco (University of Udine, Italy)
- Riccardo Lunardi (University of Udine, Italy)
- Alberto Marturano (University of Udine, Italy)
- Zaka-Ud-Din Muhammad (University of Udine, Italy)
- Francesco Nascimben (University of Udine, Italy)
- Moritz Nottebaum (University of Udine, Italy)
- Massimiliano Pascoli (University of Udine, Italy)
- Mihai Horia Popescu (University of Udine, Italy)
- Laura Rasotto (University of Udine, Italy)
- Mubashara Rehman (University of Udine, Italy)
- Francesco Taverna (University of Udine, Italy)
- Biagio Tomasetig (University of Udine, Italy)
- Alessandro Tremamunno (University of Udine, Italy)