Report on the 1st Workshop on User Modelling in Conversational Information Retrieval (UM-CIR 2024) at SIGIR-AP 2024 and Workshop on Emerging Trends in Conversational Information Retrieval (ET-CIR 2024) at National Institute of Informatics (NII)

Praveen Acharya
Dublin City University
Ireland
praveen.acharya2@mail.dcu.ie

Gareth J. F. Jones
Dublin City University
Ireland
gareth.jones@dcu.ie

Xiao Fu
University College London
U.K.
xiao.fu.20@ucl.ac.uk

Fabio Crestani
Universitá della Svizzera Italiana
Switzerland
fabio.crestani@usi.ch

Noriko Kando
National Institute of Informatics
Japan
kando@nii.ac.jp

Abstract

Conversational Information Retrieval (CIR) has seen growing interest for a number of years, gaining momentum in recent years, driven in large part by the emergence of generative AI and the increasing capabilities of conversational agents. To date, much of the CIR research within the information retrieval community has concentrated on modelling the interactive processes between users and systems, with comparatively less attention given to explicitly modelling the user within these interactions. The UM-CIR 2024 sought to provide a forum for examination of the role of the user in CIR. Recognising that CIR continues to be an evolving area of investigation, the follow on ET-CIR 2024 workshop provided a platform for discussion of emerging trends in the field. Each workshop featured a combination of keynote and invited talks, panel discussions, and paper presentations that collectively emphasized the importance of incorporating user models and perspectives in the design and evaluation of CIR systems.

Date: 12–13 December 2024.

Website: https://um-cir.github.io.

1 Introduction

Conversational Information Retrieval (CIR) has been the subject of a number of previous workshops and meetings. These have explored definitions and the potential of CIR. Recent advances, especially in generative AI, have led to growing interest in CIR within the information retrieval research community. Much of the current research on CIR has concentrated on system interactions, with limited attention given to modelling the user involved in these processes. The *First Workshop on User Modelling in Conversational Information Retrieval (UM-CIR 2024)* held on December 12, 2024, as a half-day event at SIGIR-AP 2024, hosted by Waseda University in Tokyo, Japan [Acharya et al., 2024b] sought address this gap by bringing together researchers to share perspectives and proposals on user modelling in CIR. Topics included algorithm and system design, personalization strategies, user simulation, and evaluation methods. The half-day event featured a keynote talk and paper presentations, summarised in Section 2. The workshop concluded with a panel discussion, offering an integrated view of the presented work and exploring future directions for user modelling in CIR.

While there has been much previous discussion of CIR, it remains an evolving area of investigation. Recognising that further exploration is needed of the potential directions for the continuing development of CIR, the *Workshop on Emerging Trends in Conversational Information Retrieval (ET-CIR)* was held the day after UM-CIR, on December 13, 2024 at the National Institute of Informatics (NII) in Tokyo, Japan. The workshop provided a forum for discussion on current and emerging trends in CIR, in particular noting the impact on CIS of advances in generative AI and conversational agents. The workshop aimed to gather diverse perspectives, insights and research directions relevant to CIR system design and interaction. The full-day event featured an invited keynote talk, paper presentations, and thematic discussions exploring various aspects of CIR. These included in-depth analyses and broader overviews of the field's development, and details of the most fruitful of these discussions are summarised in Section 3. ET-CIR was strongly related in some of its topics examined at UM-CIR, and therefore involved a few of the same speakers and many common participants, offering further opportunities for community engagement and collaboration.

2 UM-CIR 2024

2.1 Keynote: Dr Johanne Trippas, RMIT University, Australia

Dr Johanne Trippas, Vice-Chancellor's Research Fellow at RMIT University, Australia, delivered a thought-provoking invited keynote on the future of interactive information retrieval through immersive personalisation. Her talk, titled *Immersive Personalisation for Conversational Information-seeking: The Future of Interactive Information Retrieval*, outlined a vision in which user interactions with information systems evolve from static interfaces to dynamic, conversational experiences that are deeply tailored to individual users. [Trippas, 2024]

Dr. Trippas challenged the traditional model of search, which is largely based on document retrieval and fixed query-result patterns. Drawing on recent research and advances in large language models (LLMs), she emphasised that users typically do not conceptualise their information needs in terms of documents. Instead, they pursue goals shaped by context, preferences, and evolving

needs. Immersive personalisation aims to meet these needs by leveraging adaptive conversational interfaces that anticipate and respond to users in real time.

A key theme of the keynote was the redefinition of information spaces, moving beyond the familiar "ten blue links" to systems that support interactive exploration and visualisation of complex multidimensional knowledge domains. This future, Dr. Trippas argued, promises a more intuitive and powerful form of engagement with information, designed around how users think, learn, and solve problems.

However, the keynote also raised important ethical and practical questions. Dr. Trippas discussed the need to balance personalisation with privacy, ensure fairness and inclusivity, and design systems that remain accessible to all users. She called for a user-centred approach that does not simply respond to queries but anticipates needs, adapts to context, and supports critical engagement with information. Her presentation provided both inspiration and direction for the CIR community, highlighting the transformative potential of immersive personalisation while underscoring the responsibilities that come with designing such systems.

2.2 Paper Presentations

The workshop also featured talks by presenters for the papers accepted for inclusion in the workshop proceedings.

Xiao Fu: The first presentation was given by Xiao Fu from University College London (UCL), UK, on *An Evaluation Framework for Conversational Information Retrieval Using User Simulation* [Fu et al., 2024]. Xiao introduced a novel evaluation framework for CIR. The proposed framework utilises a two-stage user simulation approach, comprising an *Action Predictor*, informed by psychological principles of priming, and a *Response Generator* leveraging LLMs to generate realistic user responses. The Action Predictor categorises user actions into *stopping*, *following-up*, or *switching topics*, guided by statistical signals and textual patterns. Meanwhile, the Response Generator produces personalised and diversified responses based on textual user profiles.

The presentation emphasised the importance of robust, automated evaluation paradigms that transcend traditional IR metrics, which often fail to capture contextual satisfaction in conversational interactions. The proposed user simulation aims to align closely with genuine user behaviour by integrating indirect assessment methods involving real-user annotations, thus providing realistic insights into user satisfaction and effort. Additionally, the work anticipates future directions in CIR, particularly the integration and evaluation of RAG systems, addressing potential challenges such as data leakage associated with the broader deployment of LLMs.

Fabio Crestani: The second talk was given by Fabio Crestani, from the Università della Svizzera Italiana (USI) in Lugano, Switzerland. It was titled: *Responses to Conversational Information Retrieval Clarifying Questions with User Emotions*. Fabio's presentation explored the integration of emotional aspects into user simulation for CIR, a growing area aimed at improving the evaluation of the CIR system. Current data sets used to train user simulators often lack emotional content, limiting their ability to mimic realistic user behaviour. The paper, resulting from the good work of a Master's student, who is the paper's first author, describes an initial approach for injecting emotions into existing CIR data, focusing on user responses to clarifying questions (CQs) [Ecevit et al., 2024]. Using the Qulac [Aliannejadi et al., 2019] dataset as a base, the authors

generate emotionally rich responses via LLMs, employing three different prompt strategies: CO-STAR, TIDD-EC, and a custom-designed format. The paper evaluates 13 combinations of LLMs and prompts, assessing the quality of the responses generated in terms of emotional expression, naturalness, and perceived usefulness. The results offer valuable insights into the effectiveness of different prompting techniques for simulating emotional variability in user behaviour, pointing toward more realistic and diverse user simulation in CIR, ultimately supporting more robust and human-aligned system evaluation frameworks.

Praveen Acharya: The final talk was given by Praveen Acharya from Dublin City University (DCU), Ireland, on *A Framework for Knowledge Integration in Conversational Information Retrieval* [Acharya et al., 2024c]. This proposes a novel framework aimed at improving CIR by incorporating the user's prior knowledge about the search topic into the dialogue process designed to enhance the system's ability to accurately interpret the user's information needs, maintain contextual relevance, and take appropriate subsequent actions throughout the conversation.

The framework is structured around three core components: *Knowledge Extractor, Knowledge Updater and Current Knowledge State*. The *Knowledge Extractor* is responsible for identifying and extracting relevant information from various sources encountered during the conversational search process. The *Knowledge Updater* integrates the extracted information into the CIR system's existing knowledge representation, ensuring that the system continuously maintains an accurate and evolving understanding of the conversation. Finally, the *Current Knowledge State* captures the system's current understanding of the user's knowledge and information needs, as well as the context of the ongoing conversation. This dynamic representation enables the system to generate contextually appropriate responses and make informed decisions throughout the interaction.

The paper primarily focuses on the theoretical formulation of the framework, including its conceptual structure and potential operational benefits. However, the authors leave the full technical implementation and empirical evaluation of the framework for future work, noting the need for further research to explore practical deployment strategies, algorithmic solutions, and performance benchmarking.

3 ET-CIR 2024

The ET-CIR Workshop was a follow-up to the UM-CIR workshop that was specifically organized to discuss emerging trends in CIR. The workshop featured invited talks from experts and researchers in the field of CIR introducing their personal perspective on trends that are emerging in personalized information retrieval and areas which they see as important in the future of CIR research, and an open discussion exploring topics felt to be important by the participants.

3.1 Invited Presentations

Qingyao Ai: The first invited talk was by Qingyao Ai, Tsinghua University, China, titled *Asking Clarifying Questions with Large Language Models*. Qingyao discussed the ambiguity often observed to be present in user queries during conversational search and explored how LLMs can be leveraged to generate clarifying questions to improve the conversational search process. He addressed key challenges related to how, when, and what to ask in order to effectively reduce

ambiguity. While highlighting the potential of LLMs as a foundational component in the future of conversational search, he also emphasized the limitations of using LLMs in isolation, particularly their lack of deep user understanding and contextual awareness, which are critical for effective search interactions.

Mohammad Aliannejadi: The second invited talk was by Mohammad Aliannejadi, University of Amsterdam, The Netherlands, titled *LLMs and IR Evaluation: Towards Building Reusable Conversational Test Collections*. The talk centred on recent advances in LLM-based evaluation methods, exploring both the opportunities and potential risks of using large language models as evaluators. In addition, nugget-based evaluation was introduced as a promising approach for assessing generated content. Drawing on insights from organizing the TREC Interactive Knowledge Assistance Track (iKAT) [Aliannejadi et al., 2024], this method was presented as a way to enable more reusable and consistent evaluation of generated text.

Johanne Trippas: The third invited talk was by Johanne Trippas, RMIT University, Australia, titled *Re-evaluating the Command-and-Control Paradigm in Conversational Search Interactions*. The talk was based on research analysing real-world usage of conversational voice assistants. The findings revealed a strong prevalence of command-and-control interactions, with limited evidence of conversational search. Most usage was focused on simple tasks, highlighting that current interactions with voice assistants remain narrow in scope. This suggests that users may be constrained not by the technology itself, but by a lack of awareness about its broader capabilities – they simply do not know that these systems can support more complex or conversational interactions.

Fabio Crestani: The fourth invited talk was by Fabio Crestani, USI, Switzerland, titled *Conversational Information Retrieval Requires Personalization*. The talk was grounded in the idea that CIR is fundamentally a dialogue between the user and the system, and that conversations are inherently personal, shaped by the individuals involved. As such, effective CIR systems must incorporate a degree of personalization tailored to the specific user, which has implications for both system design and evaluation. The talk explored core principles of personalization in information retrieval, discussed potential approaches for implementing personalization in CIR, and introduced the concept of simulating users as a means to facilitate the design and evaluation of personalized CIR systems. It also outlined the key challenges associated with this approach.

Gareth Jones: Finally, the last invited talk was by Gareth Jones, Dublin City University, Ireland, titled *Conversational Search Futures: Proactivity, Personality, and Presentation*. The talk focused on how most current CIR systems are reactive, primarily operating through simple query-response cycles. It explored the potential for CIR systems to become more proactive – anticipating user needs and guiding the interaction more intelligently. The speaker also raised the idea of incorporating personality into agents within CIR systems, posing the question of whether agent personality could enhance the user experience and support more effective conversational search. Additionally, the talk highlighted a key challenge in CIR: understanding how users engage with content during the search process and how that interaction shapes the overall search experience.

3.2 Discussions

Whilst briefly raising a number of topics, the lively discussion session during the workshop primarily focused on the role of *nuggets* in evaluating CIR. Nuggets were described as discrete pieces

of valuable information that contribute to satisfying a user's information need. The discussion highlighted several key points:

- **Diverse User Trajectories**: Even with the same dialogue system, users from different backgrounds may pursue distinct conversational paths. This divergence influences how and when nuggets are encountered or required.
- **Nature of Nuggets**: Nuggets can appear in various forms—triples, sentences, points, or even entire documents. There is currently no standard definition, adding complexity to their detection and evaluation.
- **Gold Nuggets vs. Fool's Gold**: A *gold nugget* directly addresses a user's core information need, while hallucinated content, although seemingly relevant, was likened to "fool's gold": deceptive and ultimately unhelpful.
- **Information Overload**: Individual documents may contain many nuggets, and across large document sets, the sheer volume can overwhelm users. Efficient selection and presentation are therefore essential.
- **Nuggets in Interaction**: At the interaction level, the sequence and combination of nuggets influence conversational flow. Ordering can guide the dialogue toward or away from the successful resolution of the user's goal.
- **Compositionality**: Combinations of nuggets can lead to emergent insights, effectively creating new nuggets. This reflects the compositional and cumulative nature of information seeking.
- **Historical Context**: Nuggets have been explored in the context of complex question answering, reinforcing their potential utility in CIR. Their evaluation requires nuanced metrics that account for both content and context.

The discussion underscored the need for clearer definitions, structured representations, and context-aware evaluation strategies for nuggets to advance CIR system development and benchmarking.

4 Concluding Remarks

The UM-CIR 2024 and ET-CIR 2024 workshops served as venues for discussions on user-centric approaches and emerging trends in Conversational Information Retrieval (CIR). A key theme was the reimagining of information spaces — moving beyond traditional search paradigms toward systems that support more intuitive, interactive exploration. These discussions emphasized the need for CIR systems to be fundamentally user-centred, while also recognizing the significant challenges involved in designing such systems. In particular, the difficulty of effectively evaluating user-centric CIR systems was identified as a critical area for future research, essential to the development and validation of these next-generation interfaces.

5 Acknowledgements

This work was conducted with the financial support of Taighde Éireann – Research Ireland under Grant No. 18/CRT/6223 and partially as part of the SFI ADAPT Centre at DCU (Grant No.

13/RC/2106_P2) (www.adaptcentre.ie), Swiss Federal Institute of Technology (ETH) Leading House Asia Project on "Evaluation of Empathic Conversational Search for Depression Assessment and Support", Swiss National Science Foundation (SNSF) Grant "Personality and Conversational Information Access" and National Institute of Informatics Internship programme.

References

Praveen Acharya, Charles L. A. Clarke, Fabio Crestani, Xiao Fu, Gareth J. F. Jones, Noriko Kando, Makoto P. Kato, Aldo Lipani, and Yiqun Liu, editors. *Joint Proceedings of the 1st Workshop on Evaluation Methodologies,Testbeds and Community for Information Access Research (EMTCIR 2024) and the 1st Workshop on User Modelling in Conversational Information Retrieval (UM-CIR 2024) co-located with the 2nd International ACM SIGIR Conference on Information Retrieval in the Asia Pacific (SIGIR-AP 2024), Tokyo, Japan, December 12, 2024*, volume 3854 of *CEUR Workshop Proceedings*, 2024a. CEUR-WS.org. URL https://ceur-ws.org/Vol-3854.

Praveen Acharya, Gareth JF Jones, Xiao Fu, Aldo Lipani, Fabio Crestani, and Noriko Kando. The 1st workshop on user modelling in conversational information retrieval (um-cir). In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 315–317, 2024b.

Praveen Acharya, Noriko Kando, and Gareth J. F. Jones. A framework for knowledge integration in conversational information retrieval. In Acharya et al. [2024a]. URL https://ceur-ws.org/Vol-3854/um-cir-3.pdf.

Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. Asking clarifying questions in open-domain information-seeking conversations. In *International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, SIGIR '19, 2019.

Mohammad Aliannejadi, Zahra Abbasiantaeb, Shubham Chatterjee, Jeffery Dalton, and Leif Azzopardi. Trec ikat 2023: The interactive knowledge assistance track overview. *arXiv preprint arXiv:2401.01330*, 2024.

Isin Su Ecevit, Lili Lu, and Fabio Crestani. Responses to conversational information retrieval clarifying questions with user emotions. In Acharya et al. [2024a]. URL https://ceur-ws.org/Vol-3854/um-cir-2.pdf.

Xiao Fu, Aldo Lipani, and Noriko Kando. An evaluation framework for conversational information retrieval using user simulation. In Acharya et al. [2024a]. URL https://ceur-ws.org/Vol-3854/um-cir-1.pdf.

Johanne R. Trippas. Immersive personalisation for conversational information-seeking: The future of interactive information retrieval. In Acharya et al. [2024a]. URL https://ceur-ws.org/Vol-3854/um-cir-keynote-1.pdf.