Conversational Information Retrieval and Recommender Systems

Guglielmo Faggioli, Nicola Ferro, and Simone Merlo(✉)

University of Padua, Padua, Italy
simone.merlo@studenti.unipd.it

Abstract. Conversational systems are increasing their popularity since they allow users to interact in a simple and natural way. Information Retrieval (IR) and Recommender Systems (RS) represents two categories of systems that strongly rely on the interaction with the user. For these reasons, recently many researches increased their effort towards the development Conversational Information Retrieval (CIR) and Conversational Recommender Systems (CRS). Such systems, in fact, allow to increase the ease of use from the user perspective and also to improve the quality of the results. The aim of this tutorial is to show the best and most frequently used approaches/paradigms to build CIR and CRS systems and to understand how these can be evaluated. During the tutorial the participants will be provided with the knowledge that is needed to understand, create and evaluate CIR and CRS.

Keywords: Information Retrieval · Recommender Systems · Conversational Systems

1 Motivation

Conversational agents are rapidly growing in popularity. Such systems allow end users to seek information through natural language. This not only benefits the general public, but it might also improve the usability for specific categories of users, such as children, elderly, and visually impaired users. At the same time, the conversational interaction between the user and the systems comes with additional challenges compared to a standard information-seeking agent, which need to be properly addressed. For example, the system needs to be capable of handling complex natural language structures that include anaphoras, ellipses and co-references. It needs also to constantly maintain knowledge of the state of the conversation and past interactions to adapt its responses accordingly. The final challenge concerns the evaluation of these systems: compared to a single interaction between the user and the system, a conversation might flow in several different ways, making the evaluation of these systems a far more complete task. This problem is exasperated by Large Language Models (LLMs), a very powerful resource to dialogue with the user, but whose answers are more challenging to explain and evaluate. Traditionally, conversational agents are divided into

C. Hauff et al. (Eds.): ECIR 2025, LNCS 15576, pp. 280–290, 2025.
https://doi.org/10.1007/978-3-031-88720-8_44

Conversational Information Retrieval (CIR) systems and Conversational Recommender Systems. The former allows the user to retrieve information from a corpus that satisfies the information need they express through their utterances. The latter, on the other hand, employs an interactive process that allows the user to progressively refine its preferences naturally, through dialogue, and obtain a recommendation. While this separation is natural from the system perspective, it does not allow the exploitation of the full potential of a system to seamlessly converse with a user. At the same time, past efforts in the joint recommendation and search [50,52,63] have shown the advantages of bridging these two categories of systems.

In this tutorial, we will provide the participants with an overview of the historical development of the CIR and CRS, trying to bridge them, highlighting the similarities and differences, with the additional objective of fostering collaboration between the research communities in nformation Retrieval (IR) and Recommender Systems (RS). Additionally, we will present Conversational Agents Framework for Evaluation (CAFE), a recently developed evaluation framework that threats the conversational agent as a holistic entity, overcoming the traditional dichotomy into CIR and CRS when it comes to the evaluation,

2 Format

The workshop will be organized into three modules, *conversational search*, *conversational recommendation* and *evaluation*. In the first two modules, we will focus on the algorithmic and modelling aspects of the conversational systems, highlighting similarities and differences, and emphasizing possible synergies between the two. Finally, in the evaluation module, we will outline the challenges and opportunities the conversational context introduces in evaluating the models. We focus on the evaluation with a holistic view, to frame both search and recommendation within the same evaluation framework.

Conversational Search (1 h). Traditionally, Conversational Search (CS) agents are divided into chit-chat bots [56,59] meant to entertain the user, and task-oriented agents, devoted to completing a search task for the user [11,25]. Task-oriented agents are further divided into three main categories, pure CIR systems that retrieve the answer from a corpus [37,40,54,57,58], Question Answering (QA) systems that answer users' utterances with facts and atomic pieces of information [34], and systems that generate the answer by employing a generative model [35]. Recent advances in the LLM domain and Retrieval Augmented Generation (RAG) have blurred the borders between different categories of approaches. The tutorial will provide an overview of the development of CS systems through time and will focus on the latest advances introduced by the recent LLM-based solutions. Besides the aspects related to the matching and retrieval of the documents or the generation of the answer in response to a query, several ancillary tasks were developed in the context of the conversational search. Among them, we cite Query Performance Prediction (QPP) for CIR and the so-called mixed-initiative interaction. QPP for CIR [20,41,42,53] focuses on the

specific characteristics of the conversational setting to adjust the behaviour of the conversational agent according to the predicted quality of the response. In a similar spirit, the mixed-initiative interaction [3,10,26,46] allows the system to ask clarifying questions, in case the system detects it has not enough information to answer the user's information need. The tutorial will provide an overview of such aspects, to provide a comprehensive view of the research paths stemming from CS.

Conversational Recommendation (1 h). CRS are traditionally composed of several building blocks [28]. As for traditional RS, also for CRS personalization plays an important role [31]. Nonetheless, in CRS there are additional needs with respect to standard RS. A CRS, in fact, must be able to perform natural language understanding and to represent the user preferences in the context of a single conversation. Considering the latter, several techniques have been proposed both to keep track of the user preferences and to exploit the expressed preferences for the recommendation process. In this context Knowledge Graphs (KG), dense embeddings, and attention based mechanisms are frequently exploited [13,38,43,62]. As in many other research fields, also for CRS the advancements in the LLM domain had an important impact in the development of both the systems [38,39,44,44,49,55] and the datasets [29,36]. In this tutorial, we present the main structure and components of CRS (from traditional architectures to modern, LLM based, ones), we discuss the role of personalization in CRS and how appropriate user modelling may impact the recommendation performance and we argue the advantages and disadvantages of such systems.

Evaluating a Conversational System (1 h). Traditionally, conversational agents have been devoted to either search or recommendation, treating the two tasks separately. As a consequence, each discipline developed its evaluation framework, with limited cross-fertilization [18,19,30,45,48,51,60,61]. Similarly, shared efforts within the IR community, such as TREC-CAsT [15–17,45], TREC-iKaT [2], and datasets developed by the RS community, such as ReDial [33], its LLM-based counterpart LLM-REDIAL [36], and the ConverRSE [27] dataset, to name a few, focus exclusively on either search or recommendation. As a consequence, the evaluation of an integrated conversational agent that operates seamlessly as a searcher and recommender remains a challenging task with several under-explored aspects. In this tutorial, we present and discuss the CAFE, recently developed jointly by the IR, RS, and Natural Language Processing (NLP) communities, during a Perspective Workshop in Dagsthul [12]. The CAFE is based on six aspects that should be identified while designing the evaluation protocol for an integrated CIR and CRS system: the goals that the stakeholder aims to achieve, the relevant user aspects, the tasks that the system is expected to complete, the scope of the evaluation (i.e., whether the system should be evaluated at turn level, at a conversation level, or on the long run), the methodology (e.g., controlled lab studies, AB testing), and the measures that quantify the effectiveness of the system across all the aforementioned evaluation aspects. During the tutorial, we will present practical examples of each aspect mentioned

before and discuss how they can be identified or chosen in a real-life scenario to implement a holistic conversational agent evaluation.

3 Audience

This tutorial is expected to attract audience from a vast community including, but not limited to Information Retrieval, Recommender Systems, and Natural Language Processing at large. Furthermore, we also target the part of the communities (both IR and RS) devoted to the evaluation, who can familiarize with the conversational evaluation, a different evaluation setting that presents its own peculiarities that should be accounted for.

Target Audience. Given the broad applicability of the conversational search systems, the tutorial is tailored to target both academic and industry audience, mainly belonging to IR and RS communities. We envision a tutorial that can be useful to both research-versed audience, as well as practitioners that have more interest in the practical aspects.

Prerequisite Knowledge. This tutorial will be self-contained and has minimal prerequisite knowledge. The participants should be familiar with basic IR and RS concepts, such as ranking and filtering. Furthermore, the audience should be familiar with evaluation paradigms and shared campaigns, such as TREC. Finally, considering the recent advances introduced in the domain by contextual text representation and Large Language Models in particular, the participants to the tutorial should have basic knowledge of related concepts, such as the transformer architecture and dense encoding.

4 Tutorial History

In recent years, both tutorials related to CIR and CRS have been presented to major conferences like RecSys, SIGIR and WSDM. Previous tutorials related to CIR include: "Recent Advances in Conversational Information Retrieval" [24] and "Conversational Information Seeking: Theory and Application" [14]. Existing tutorials on CRS, instead, include: "RecSys 2021 Tutorial on Conversational Recommendation: Formulation, Methods, and Evaluation" [32] and "Tutorial on Conversational Recommendation Systems" [21–23].

Furthermore, in the past also workshops related to these research fields have been proposed [1,4–9,47].

However, the previously presented tutorials relate to 3–5 years ago but in recent years there have been many advancements in conversational systems. In our tutorial, differently form the others, we will focus on both CIR and CRS and their evaluation. We will highlight the differences and similarities between conversational systems in different contexts. Thus, this tutorial will also make it possible to better understand how to bridge the gap between CIR and CRS, going towards the novel and promising field of Joint IR and RS.

Acknowledgments. This work has received support from CAMEO, PRIN 2022 n. 2022ZLL7MW.

References

1. Acharya, P., Jones, G.J.F., Fu, X., Lipani, A., Crestani, F., Kando, N.: The 1st workshop on user modelling in conversational information retrieval (UM-CIR). In: Sakai, T., Ishita, E., Ohshima, H., Hasibi, F., Mao, J., Jose, J.M. (eds.) Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, SIGIR-AP 2024, Tokyo, Japan, 9–12 December 2024, pp. 315–317. ACM (2024). https://doi.org/10.1145/3673791.3698436
2. Aliannejadi, M., Abbasiantaeb, Z., Chatterjee, S., Dalton, J., Azzopardi, L.: TREC ikat 2023: the interactive knowledge assistance track overview. CoRR arxiv:2401.01330 (2024). https://doi.org/10.48550/ARXIV.2401.01330
3. Aliannejadi, M., Zamani, H., Crestani, F., Croft, W.B.: Asking clarifying questions in open-domain information-seeking conversations. In: Piwowarski, B., Chevalier, M., Gaussier, É., Maarek, Y., Nie, J., Scholer, F. (eds.) Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, 21–25 July 2019, pp. 475–484. ACM (2019). https://doi.org/10.1145/3331184.3331265
4. Anelli, V.W., et al.: Knowledge-aware and conversational recommender systems. In: Pera, S., Ekstrand, M.D., Amatriain, X., O'Donovan, J. (eds.) Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, 2–7 October 2018, pp. 521–522. ACM (2018). https://doi.org/10.1145/3240323.3240338
5. Anelli, V.W., et al.: Fourth knowledge-aware and conversational recommender systems workshop (kars). In: Golbeck, J., Harper, F.M., Murdock, V., Ekstrand, M.D., Shapira, B., Basilico, J., Lundgaard, K.T., Oldridge, E. (eds.) RecSys '22: Sixteenth ACM Conference on Recommender Systems, Seattle, WA, USA, 18–23 September 2022, pp. 663–666. ACM (2022). https://doi.org/10.1145/3523227.3547412
6. Anelli, V.W., et al.: Fifth knowledge-aware and conversational recommender systems workshop (kars). In: Zhang, J., Chen, L., Berkovsky, S., Zhang, M., Noia, T.D., Basilico, J., Pizzato, L., Song, Y. (eds.) Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023, Singapore, Singapore, 18–22 September 2023, pp. 1259–1262. ACM (2023). https://doi.org/10.1145/3604915.3608759
7. Anelli, V.W., Basile, P., Noia, T.D., Donini, F.M., Musto, C., Narducci, F., Zanker, M.: Third knowledge-aware and conversational recommender systems workshop (kars). In: Pampín, H.J.C., et al (eds.) RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September–1 October 2021, pp. 806–809. ACM (2021). https://doi.org/10.1145/3460231.3470933
8. Anelli, V.W., Ferrara, A., Musto, C., Narducci, F., Ragone, A., Zanker, M.: Sixth knowledge-aware and conversational recommender systems workshop (kars). In: Noia, T.D., et al. (eds.) Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, 14–18 October 2024, pp. 1245–1249. ACM (2024). https://doi.org/10.1145/3640457.3687114

9. Anelli, V.W., Noia, T.D.: 2nd workshop on knowledge-aware and conversational recommender systems - kars. In: Zhu, W., et al. (eds.) Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, 3–7 November 2019, pp. 3001–3002. ACM (2019). https://doi.org/10.1145/3357384.3358805
10. Arabzadeh, N., Seifikar, M., Clarke, C.L.A.: Unsupervised question clarity prediction through retrieved item coherency. In: Hasan, M.A., Xiong, L. (eds.) Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, 17–21 October 2022, pp. 3811–3816. ACM (2022). https://doi.org/10.1145/3511808.3557719
11. Bangalore, S., Fabbrizio, G.D., Stent, A.: Learning the structure of task-driven human-human dialogs. IEEE Trans. Speech Audio Process. **16**(7), 1249–1259 (2008). https://doi.org/10.1109/TASL.2008.2001102
12. Bauer, C., et al.: Conversational agents: a framework for evaluation (cafe) (2024)
13. Chen, X., Wang, Y., Yang, J.: UCRI: a unified conversational recommender system based on item-guided conditional generation. IEEE Intell. Syst. **39**(1), 46–55 (2024). https://doi.org/10.1109/MIS.2023.3330367
14. Dalton, J., et al.: Conversational information seeking: Theory and application. In: Amigó, E., Castells, P., Gonzalo, J., Carterette, B., Culpepper, J.S., Kazai, G. (eds.) SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, 11–15 July 2022, pp. 3455–3458. ACM (2022). https://doi.org/10.1145/3477495.3532678
15. Dalton, J., Xiong, C., Callan, J.: Cast 2020: the conversational assistance track overview. In: Voorhees, E.M., Ellis, A. (eds.) Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event, Gaithersburg, Maryland, USA, 16–20 November 2020, NIST Special Publication, vol. 1266, National Institute of Standards and Technology (NIST) (2020). https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.C.pdf
16. Dalton, J., Xiong, C., Callan, J.: TREC cast 2019: The conversational assistance track overview. CoRR arxiv:2003.13624 (2020)
17. Dalton, J., Xiong, C., Callan, J.: TREC cast 2021: the conversational assistance track overview. In: Soboroff, I., Ellis, A. (eds.) Proceedings of the Thirtieth Text REtrieval Conference, TREC 2021, online, 15–19 November 2021, NIST Special Publication, vol. 500-335. National Institute of Standards and Technology (NIST) (2021). https://trec.nist.gov/pubs/trec30/papers/Overview-CAsT.pdf
18. Faggioli, G., Ferrante, M., Ferro, N., Perego, R., Tonellotto, N.: Hierarchical dependence-aware evaluation measures for conversational search. In: Diaz, F., Shah, C., Suel, T., Castells, P., Jones, R., Sakai, T. (eds.) SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, 11–15 July 2021, pp. 1935–1939. ACM (2021). https://doi.org/10.1145/3404835.3463090
19. Faggioli, G., Ferrante, M., Ferro, N., Perego, R., Tonellotto, N.: A dependency-aware utterances permutation strategy to improve conversational evaluation. In: Hagen, M., Verberne, S., Macdonald, C., Seifert, C., Balog, K., Nørvåg, K., Setty, V. (eds.) ECIR 2022. LNCS, vol. 13185, pp. 184–198. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99736-6_13

20. Faggioli, G., Ferro, N., Muntean, C.I., Perego, R., Tonellotto, N.: A geometric framework for query performance prediction in conversational search. In: Chen, H., Duh, W.E., Huang, H., Kato, M.P., Mothe, J., Poblete, B. (eds.) Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, 23–27 July 2023, pp. 1355–1365. ACM (2023). https://doi.org/10.1145/3539618.3591625
21. Fu, Z., Xian, Y., Zhang, Y., Zhang, Y.: Tutorial on conversational recommendation systems. In: Santos, R.L.T., et al.. (eds.) RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, 22–26 September 2020, pp. 751–753. ACM (2020). https://doi.org/10.1145/3383313.3411548
22. Fu, Z., Xian, Y., Zhang, Y., Zhang, Y.: IUI 2021 tutorial on conversational recommendation systems. In: Hammond, T., Verbert, K., Parra, D. (eds.) IUI '21: 26th International Conference on Intelligent User Interfaces, College Station, TX, USA, 13–17 April 2021, Companion, pp. 1–2. ACM (2021). https://doi.org/10.1145/3397482.3450621
23. Fu, Z., Xian, Y., Zhang, Y., Zhang, Y.: WSDM 2021 tutorial on conversational recommendation systems. In: Lewin-Eytan, L., Carmel, D., Yom-Tov, E., Agichtein, E., Gabrilovich, E. (eds.) WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, 8–12 March 2021, pp. 1134–1136. ACM (2021). https://doi.org/10.1145/3437963.3441661
24. Gao, J., Xiong, C., Bennett, P.: Recent advances in conversational information retrieval. In: Huang, J.X., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., Liu, Y. (eds.) Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, 25–30 July 2020, pp. 2421–2424. ACM (2020). https://doi.org/10.1145/3397271.3401418
25. Gu, J., Ling, Z., Liu, Q.: Utterance-to-utterance interactive matching network for multi-turn response selection in retrieval-based chatbots. IEEE ACM Trans. Audio Speech Lang. Process. **28**, 369–379 (2020). https://doi.org/10.1109/TASLP.2019.2955290
26. Hashemi, H., Zamani, H., Croft, W.B.: Guided transformer: leveraging multiple external sources for representation learning in conversational search. In: Huang, J.X., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., Liu, Y. (eds.) Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, 25–30 July 2020, pp. 1131–1140. ACM (2020). https://doi.org/10.1145/3397271.3401061
27. Iovine, A., Narducci, F., de Gemmis, M.: A dataset of real dialogues for conversational recommender systems. In: Bernardi, R., Navigli, R., Semeraro, G. (eds.) Proceedings of the Sixth Italian Conference on Computational Linguistics, Bari, Italy, 13–15 November 2019, CEUR Workshop Proceedings, vol. 2481, CEUR-WS.org (2019). https://ceur-ws.org/Vol-2481/paper37.pdf
28. Jannach, D., Manzoor, A., Cai, W., Chen, L.: A survey on conversational recommender systems. ACM Comput. Surv. **54**(5), 105:1–105:36 (2022). https://doi.org/10.1145/3453154
29. Kim, M., et al.: Pearl: a review-driven persona-knowledge grounded conversational recommendation dataset. In: Ku, L., Martins, A., Srikumar, V. (eds.) Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and Virtual Meeting, 11–16 August 2024, pp. 1105–1120. Association for Computational Linguistics (2024). https://doi.org/10.18653/V1/2024.FINDINGS-ACL.65

30. Knijnenburg, B.P., Willemsen, M.C.: Evaluating recommender systems with user experiments. In: Ricci, F., Rokach, L., Shapira, B. (eds.) Recommender Systems Handbook, pp. 309–352. Springer, Heidelberg (2015). https://doi.org/10.1007/978-1-4899-7637-6_9
31. Laban, G., Araujo, T.B.: The effect of personalization techniques in users' perceptions of conversational recommender systems. In: Marsella, S., Jack, R., Vilhjálmsson, H.H., Sequeira, P., Cross, E.S. (eds.) IVA '20: ACM International Conference on Intelligent Virtual Agents, Virtual Event, Scotland, UK, 20–22 October 2020, pp. 34:1–34:3. ACM (2020), https://doi.org/10.1145/3383652.3423890
32. Lei, W., Gao, C., de Rijke, M.: Recsys 2021 tutorial on conversational recommendation: Formulation, methods, and evaluation. In: Pampín, H.J.C., Larson, M.A., Willemsen, M.C., Konstan, J.A., McAuley, J.J., Garcia-Gathright, J., Huurnink, B., Oldridge, E. (eds.) RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September–1 October 2021, pp. 842–844. ACM (2021). https://doi.org/10.1145/3460231.3473325
33. Li, R., Kahou, S.E., Schulz, H., Michalski, V., Charlin, L., Pal, C.: Towards deep conversational recommendations. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, Montréal, Canada, 3–8 December 2018, pp. 9748–9758 (2018). https://proceedings.neurips.cc/paper/2018/hash/800de15c79c8d840f4e78d3af937d4d4-Abstract.html
34. Li, Y., Li, W., Nie, L.: Dynamic graph reasoning for conversational open-domain question answering. ACM Trans. Inf. Syst. **40**(4), 82:1–82:24 (2022). https://doi.org/10.1145/3498557
35. Li, Y., Yang, N., Wang, L., Wei, F., Li, W.: Generative retrieval for conversational question answering. Inf. Process. Manag. **60**(5), 103475 (2023). https://doi.org/10.1016/J.IPM.2023.103475
36. Liang, T., et al.: LLM-REDIAL: a large-scale dataset for conversational recommender systems created from user behaviors with llms. In: Ku, L., Martins, A., Srikumar, V. (eds.) Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, 11–16 August 2024, pp. 8926–8939. Association for Computational Linguistics (2024). https://doi.org/10.18653/V1/2024.FINDINGS-ACL.529
37. Lin, S., Yang, J., Nogueira, R.F., Tsai, M., Wang, C., Lin, J.: Multi-stage conversational passage retrieval: an approach to fusing term importance estimation and neural query rewriting. ACM Trans. Inf. Syst. **39**(4), 48:1–48:29 (2021). https://doi.org/10.1145/3446426
38. Liu, S., Ao, Z., Chen, P., Kolmanic, S.: Collrec: pre-trained language models and knowledge graphs collaborate to enhance conversational recommendation system. IEEE Access **12**, 104663–104675 (2024). https://doi.org/10.1109/ACCESS.2024.3434720
39. Maes, U., Michiels, L., Smets, A.: Genui(ne) CRS: UI elements and retrieval-augmented generation in conversational recommender systems with llms. In: Noia, T.D., et al. (eds.) Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, 14–18 October 2024, pp. 1177–1179. ACM (2024). https://doi.org/10.1145/3640457.3691697
40. Mele, I., Muntean, C.I., Nardini, F.M., Perego, R., Tonellotto, N., Frieder, O.: Adaptive utterance rewriting for conversational search. Inf. Process. Manag. **58**(6), 102682 (2021). https://doi.org/10.1016/J.IPM.2021.102682

41. Meng, C., Aliannejadi, M., de Rijke, M.: Performance prediction for conversational search using perplexities of query rewrites. In: Faggioli, G., Ferro, N., Mothe, J., Raiber, F. (eds.) Proceedings of the The QPP++ 2023: Query Performance Prediction and Its Evaluation in New Tasks Workshop co-located with The 45th European Conference on Information Retrieval (ECIR), Dublin, Ireland, 6 April 2023, CEUR Workshop Proceedings, vol. 3366, pp. 25–28. CEUR-WS.org (2023). https://ceur-ws.org/Vol-3366/paper-05.pdf
42. Meng, C., Arabzadeh, N., Aliannejadi, M., de Rijke, M.: Query performance prediction: From ad-hoc to conversational search. In: Chen, H., Duh, W.E., Huang, H., Kato, M.P., Mothe, J., Poblete, B. (eds.) Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, 23–27 July 2023, pp. 2583–2593. ACM (2023). https://doi.org/10.1145/3539618.3591919
43. Ni, Y., et al.: Meta-crs: a dynamic meta-learning approach for effective conversational recommender system. ACM Trans. Inf. Syst. **42**(1), 28:1–28:27 (2024). https://doi.org/10.1145/3604804
44. Nie, G., et al.: A hybrid multi-agent conversational recommender system with LLM and search engine in e-commerce. In: Noia, T.D., Lops, P., Joachims, T., Verbert, K., Castells, P., Dong, Z., London, B. (eds.) Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, 14–18 October 2024, pp. 745–747. ACM (2024). https://doi.org/10.1145/3640457.3688061
45. Owoicho, P., Dalton, J., Aliannejadi, M., Azzopardi, L., Trippas, J.R., Vakulenko, S.: TREC cast 2022: going beyond user ask and system retrieve with initiative and response generation. In: Soboroff, I., Ellis, A. (eds.) Proceedings of the Thirty-First Text REtrieval Conference, TREC 2022, online, 15–19 November 2022, NIST Special Publication, vol. 500-338, National Institute of Standards and Technology (NIST) (2022). https://trec.nist.gov/pubs/trec31/papers/Overview_cast.pdf
46. Pal, D., Ganguly, D.: Effective query formulation in conversation contextualization: a query specificity-based approach. In: Hasibi, F., Fang, Y., Aizawa, A. (eds.) ICTIR '21: The 2021 ACM SIGIR International Conference on the Theory of Information Retrieval, Virtual Event, Canada, 11 July 2021, pp. 177–183. ACM (2021). https://doi.org/10.1145/3471158.3472237
47. Penha, G., Hauff, C.: Challenges in the evaluation of conversational search systems. In: KDD 2020 Workshop on Conversational Systems Towards Mainstream Adoption, KDD-Converse 2020, vol. 2666 (2020)
48. Pu, P., Chen, L., Hu, R.: A user-centric evaluation framework for recommender systems. In: Mobasher, B., Burke, R.D., Jannach, D., Adomavicius, G. (eds.) Proceedings of the 2011 ACM Conference on Recommender Systems, RecSys 2011, Chicago, IL, USA, 23–27 October 2011, pp. 157–164. ACM (2011). https://doi.org/10.1145/2043932.2043962
49. Salemi, A., Mysore, S., Bendersky, M., Zamani, H.: Lamp: when large language models meet personalization. In: Ku, L., Martins, A., Srikumar, V. (eds.) Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand, 11–16 August 2024, vol. 1: Long Papers, pp. 7370–7392. Association for Computational Linguistics (2024). https://doi.org/10.18653/V1/2024.ACL-LONG.399

50. Si, Z., et al.: When search meets recommendation: Learning disentangled search representation for recommendation. In: Chen, H., Duh, W.E., Huang, H., Kato, M.P., Mothe, J., Poblete, B. (eds.) Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, 23–27 July 2023, pp. 1313–1323. ACM (2023). https://doi.org/10.1145/3539618.3591786
51. Srinivasan, P., Menczer, F., Pant, G.: A general evaluation framework for topical crawlers. Inf. Retr. **8**(3), 417–447 (2005). https://doi.org/10.1007/S10791-005-6993-5
52. Thonet, T., Renders, J.-M., Choi, M., Kim, J.: Joint personalized search and recommendation with hypergraph convolutional networks. In: Hagen, M., et al. (eds.) ECIR 2022. LNCS, vol. 13185, pp. 443–456. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99736-6_30
53. Vlachou, M., Macdonald, C.: Performance predictors for conversational fashion recommendation. In: Anelli, V.W., et al. (eds.) Proceedings of the Fourth Knowledge-aware and Conversational Recommender Systems Workshop co-located with 16th ACM Conference on Recommender Systems (RecSys 2022), Seattle, WA, USA, 22 September 2022, CEUR Workshop Proceedings, vol. 3294, pp. 91–100. CEUR-WS.org (2022). https://ceur-ws.org/Vol-3294/long7.pdf
54. Voskarides, N., Li, D., Ren, P., Kanoulas, E., de Rijke, M.: Query resolution for conversational search with limited supervision. In: Huang, J.X., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., Liu, Y. (eds.) Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, 25–30 July 2020, pp. 921–930. ACM (2020). https://doi.org/10.1145/3397271.3401130
55. Wang, R., He, X., Gu, H., Wang, X.: LGCRS: llm-guided representation-enhancing for conversational recommender system. In: Wand, M., Malinovská, K., Schmidhuber, J., Tetko, I.V. (eds.) Artificial Neural Networks and Machine Learning - ICANN 2024 - 33rd International Conference on Artificial Neural Networks, Lugano, Switzerland, 17–20 September 2024, Proceedings, Part IX, Lecture Notes in Computer Science, vol. 15024, pp. 74–88. Springer, Heidelberg (2024). https://doi.org/10.1007/978-3-031-72356-8_6
56. Yan, R.: "chitty-chitty-chat bot": deep learning for conversational AI. In: Lang, J. (ed.) Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, Stockholm, Sweden, 13–19 July 2018, pp. 5520–5526. ijcai.org (2018). https://doi.org/10.24963/IJCAI.2018/778
57. Yu, S., et al.: Few-shot generative conversational query rewriting. In: Huang, J.X., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., Liu, Y. (eds.) Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, 25–30 July 2020, pp. 1933–1936. ACM (2020). https://doi.org/10.1145/3397271.3401323
58. Yu, S., Liu, Z., Xiong, C., Feng, T., Liu, Z.: Few-shot conversational dense retrieval. In: Diaz, F., Shah, C., Suel, T., Castells, P., Jones, R., Sakai, T. (eds.) SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, 11–15 July 2021, pp. 829–838. ACM (2021). https://doi.org/10.1145/3404835.3462856

59. Yu, Z., Xu, Z., Black, A.W., Rudnicky, A.I.: Strategy and policy learning for non-task-oriented conversational systems. In: Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, Los Angeles, CA, USA, 13–15 September 2016, pp. 404–412. The Association for Computer Linguistics (2016). https://doi.org/10.18653/V1/W16-3649
60. Zangerle, E., Bauer, C.: Evaluating recommender systems: survey and framework. ACM Comput. Surv. **55**(8), 170:1–170:38 (2023). https://doi.org/10.1145/3556536
61. Zhang, Y., Liu, X., Zhai, C.: Information retrieval evaluation as search simulation: A general formal framework for IR evaluation. In: Kamps, J., Kanoulas, E., de Rijke, M., Fang, H., Yilmaz, E. (eds.) Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR 2017, Amsterdam, The Netherlands, 1–4 October 2017, pp. 193–200. ACM (2017). https://doi.org/10.1145/3121050.3121070
62. Zhang, Y., et al.: Conversational recommender based on graph sparsification and multi-hop attention. Intell. Data Anal. **28**(1), 99–119 (2024). https://doi.org/10.3233/IDA-230148
63. Zhao, K., Zheng, Y., Zhuang, T., Li, X., Zeng, X.: Joint learning of e-commerce search and recommendation with a unified graph neural network. In: Candan, K.S., Liu, H., Akoglu, L., Dong, X.L., Tang, J. (eds.) WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event/Tempe, AZ, USA, 21–25 February 2022, pp. 1461–1469. ACM (2022). https://doi.org/10.1145/3488560.3498414