

CAMEO: Fostering Joint Conversational Search and Recommendation

Discussion Paper

Tommaso Di Noia¹, Guglielmo Faggioli², Marco Ferrante², Nicola Ferro^{2,*},
Fedelucio Narducci¹, Raffaele Perego³ and Giuseppe Santucci⁴

¹*Polytechnic University of Bari, Italy*

²*University of Padua, Italy*

³*National Research Council, Italy*

⁴*Sapienza University of Rome, Italy*

Abstract

The rising popularity of conversational agents for accessing information stems from their natural language dialogue-based interaction, facilitating usability for a broad spectrum of users, including the elderly, children, and visually impaired individuals. Among others, two tasks that benefit the most conversational agents are search and recommendation: in the former, the user receives factual information by asking the agent; in the latter, the system refines its knowledge of the user's needs by posing them a sequence of questions. This work discusses the observations and findings of the first CAMEO (Conversational Agents: Mastering, Evaluating, Optimizing) project retreat. The retreat focused on similarities and differences of conversational search and recommendation to identify the path to construct a joint conversational search and recommendation system. Our observations highlight how all the conversational search/recommendation systems can be categorized using two axes: "exploration-disambiguation" and "search-recommendation". The first axis describes whether the question aims to gain knowledge over something unknown or allows to refine already available knowledge. The second axis describes if the user's interest is in gaining knowledge or obtaining a recommendation. Additionally, we provide insights on obtaining a dataset that can be used to train/test such a joint system. Finally, we describe how the CAMEO project will address the product search task, which we believe is the scenario where the joint conversational search and recommendation system would be the most effective.

1. Introduction

The conversational paradigm is increasingly used both for search and recommendation. For what concerns *Conversational Search (CS)*, instead of issuing a single query, the user satisfies their information need by refining it through a conversation with the agent. At the same time, concerning *Conversational Recommendation (CR)*, instead of receiving a one-shot static recommendation, users can establish a multi-turn dialogue with the CR to express their needs best. The conversational paradigm presents several advantages. It can be easily used by visually impaired people, elderly and children, who can interact with the system using their voice. At the same time, the possibility of expressing their need using natural languages alleviates the mental encumbering of formulating effective queries for generic users. Nevertheless, conversational systems also present major challenges that must be effectively addressed. For example, the

SEBD 2024: 32nd Symposium on Advanced Database Systems, June 23-26, 2024, Villasimius, Sardinia, Italy

*Corresponding author.



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

system must understand and process complex natural language structures, such as anaphoras, ellipses, and co-references. At the same time, the system should be able to keep track of the conversation's state to provide effective and satisfactory answers for the user, also considering previous interactions. Finally, the evaluation of these systems is particularly challenging for several reasons. First, being a relatively new paradigm, we still lack rich and extensive evaluation collections. Secondly, we lack proper evaluation measures and paradigms. Indeed, most of the current evaluation approaches [1, 2, 3] rely on procedures similar to those used in classic *Information Retrieval (IR)* and *Recommender Systems (RS)* and do not keep into account the structure of the conversation, nor take into consideration the fact that different users might interact in different ways with the system.

Currently, CS and CR are intended and applied as two orthogonal tasks: either the system is meant to search, or it is designed to recommend. We argue that this approach is limiting and does not exploit the full potential of a conversational system to interact seamlessly with a user. Whether to answer with a document (search) or an item (recommendation) should depend on the user's interaction and needs rather than the system's nature. Furthermore, past efforts in the joint recommendation and search [4, 5, 6] have already highlighted the positive effects that joint modeling can have on both tasks. We argue that a similar strategy can benefit both CS and CR.

The “*Conversational Agents: Mastering, Evaluating, Optimizing (CAMEO)*” project¹ addresses the abovementioned limitations while fostering the development of conversational agents designed for joint search and recommendation. CAMEO is being developed under the *Progetti di Rilevante Interesse Nazionale (PRIN)* framework and involves partners from four institutions: The National Research Council (CNR), the Polytechnic University of Bari (POLIBA), the Sapienza University of Rome and the University of Padua (UNIPD, coordinator). In this discussion paper, we present the outcomes of the first brainstorming meeting of CAMEO, organized as a retreat of three days at San Vito di Cadore, to explore the notion of joint recommendation and search for conversational and to define a common conceptual framework. Our discussion focused on formalizing the similarities and differences between CS and CR systems. We also discussed which type of architecture would fit the best. We observed that using two different engines (one for recommendation and one for search) that share some information, either in the form of a knowledge base or representation space, would be the most effective solution in our case. Finally, we discussed how a joint search and recommendation dataset could be collected.

The remainder of this paper is organized as follows: Section 2 introduces the background on CR and CS, Section 3 details the observations and analyses of the current state-of-the-art resulting from the first CAMEO's project meeting. Section 4 introduces conversational product recommendation, the first task CAMEO will focus on. Finally, Section 5 draws the conclusion and introduces our future work.

2. Background

We describe here CS and CR paradigms, highlighting their similarities and differences.

¹<https://cameo.dei.unipd.it>

2.1. Conversational Search (CS)

Our definition of CS follows the one adopted by the TREC *Conversational Assistance Track (CAST)* evaluation campaign [1, 2, 3, 7]. The main task of TREC CAST considers a scenario in which the user asks and the system responds. Each utterance of the user conveys an information need (the user wants to know something with each utterance). Thus, each utterance has a (set of) correct answer(s) – i.e., documents that satisfy the information need of the user. The system is evaluated by looking at how many correct documents have been retrieved in response to the user utterances. An example of a conversation following the CS paradigm is:

User (utterance 1): which fruits are winter fruits?

System (response):

D1: apples are a very healthy type of winter fruits.

D2: oranges are winter fruits.

D3: strawberries are not winter fruits.

D4: winter is the coldest season of the year.

User (utterance 2): tell me more about which types of apples exist

System (response):

D1: Golden Delicious is a very famous variety of apples.

D2: There are over 7500 cultivars of apples.

D3: Navel is a popular orange cultivar.

The user issues the first utterance (utterance 1), and the system retrieves documents in response. Among such documents, some will be relevant, some will not. This allows us to compute IR evaluation measures such as nDCG, AP, and Precision, and measures designed for CS [8, 7]. Then, the user issues utterance 2, and the system retrieves a new set of documents. As in the previous case, some of them are relevant, e.g., D1 and D2, some are not e.g., D3.

More recently, the *mixed-initiative* task has emerged [3, 7, 9, 10]. In this case, we assume that the system has an additional component that analyzes the query and tries to estimate whether the quality of the answer will be sufficient. Suppose the system estimates that answering the user's question is particularly challenging. In that case, it can formulate a clarifying question to ask the user to provide more information and better contextualize their utterance.

2.2. Conversational Recommendation (CR)

Regarding CR, the task presents similar challenges to CS (i.e., handling natural language and exploiting the context), but follows a quite different search paradigm. Also in this case we can have three different interaction strategies: system driven, user driven, and mixed initiative [11]. Regarding the recommendation paradigm, the CR can be categorized at least into two main approaches, based on dialogue state tracker or end-to-end models [12]. Both paradigms aim to identify the most relevant items for user needs. CR based on a dialogue-state-tracker can adopt a “slot filling” approach [13]. There are a set of characteristics and, by asking the user questions, the system obtains the value of the hidden target item for one or more of these characteristics. Once the system has sufficiently restricted the search space, it provides a recommendation. An example of conversation for CR is the following:

User (target item apples): I would like to eat something

Table 1
Differences between CS and CR

CS	CR
based on “exploration”: the user is interested in “exploring” an information need. The system “knows” the truth and provides it based on user’s questions.	based on “disambiguation”: the system wants to “disambiguate” the user’s need. The user “knows” the truth and provides it to the system, guided by its question.
The user asks, the system responds	The system asks, the user responds, or viceversa
There are no domain-dependent operations: the user conditions the domain with their questions.	The system’s questions depend on the domain.
several conversation goals: several different information needs, one per utterance.	single conversation goal: items to be recommended.
In the most typical scenario, no user profile can be used to condition the answers.	It is possible to condition the search space based on both the conversation and the user profile

System: vegetables or fruits?

User: fruits

System: winter fruits or summer fruits?

User: I prefer winter ones

System: you should try apples

When applying the slot-filling CR paradigm, it is possible to extend the task with additional context, e.g., knowing the current season, the system might avoid questions. Some information about users can also be extracted from their profiles that can store long- and/or short-term preferences (e.g., the user likes red fruits). A possible slot-filling approach could be based on “backtracking”: with each system’s question, the search space is further partitioned [13]. At a certain point, when the search space is sufficiently small, the system provides a recommendation. If the recommendation is wrong, it backtracks to the previous questions and changes them to gain additional information.

The evaluation concerns effectiveness, efficiency (or interaction cost), and linguistic aspects [14]. Regarding effectiveness, the typical aspect considered is how often the system provides the correct recommendation at the end of the conversation. Regarding efficiency, relevant aspects include the time, number of conversation turns, or tokens needed to formulate a recommendation. Finally, concerning linguistic aspects, elements that need to be evaluated include, for example, the quality of the dialog and the naturalness of the questions.

3. Discussion

3.1. Comparing CS and CR

Differences Based on the paradigms described in Section 2, we analyse here the main differences between CS and CR that should be taken into consideration when designing a joint CS and CR framework. Table 1 summarizes our observations.

The first two differences reported in Table 1 concern what the user expects to do with the system. With the archetypal CS system, the user is expected to lead the conversation. On the contrary, the typical CR system could have a mixed interaction. When a user interacts with a CS system, they aim at “exploring” the corpus to gain more information, without having any previous knowledge of it. On the other hand, when a CR system interacts with a user, the system is aware that the user has a sort of *ideal item* in mind and needs to “disambiguate” among many different possibilities to identify the right one. This leads to the third difference between the two classes of systems. The CS approach does not require specific knowledge of the domain: as the user does not have any knowledge of the relevant documents and their characteristics, it is not possible to use domain knowledge to restrict the search space. Vice-versa, when it comes to the recommendation task, the CR system needs to exploit its domain knowledge to refine the answer. For example, if the CR system works on the movie domain, the question “*with actors should have acted in it*” might be particularly effective. The same question is useless in other domains, such as music or product recommendation.

Finally, the CS system seldom employs a user profile, while it is natural for the CR system to employ characteristics of the user to provide an effective recommendation.

Similarities The first similarity between the two tasks is linked to the challenges that arise when natural language is involved. Both CS and CR systems need to deal with anaphoras, coreferences, and other complex linguistic structures.

The conversation’s structure is, in our opinion, the most relevant point of contact between CS and CR. In both cases, the conversation could be formalized as a conversation tree. When formalizing a conversation in the CS scenario, it is possible to use a tree [7, 8, 15]: some utterances serve to obtain more specific information on the previous utterances (i.e., disambiguate), while others expand the knowledge (i.e., explore). Similarly, a conversation in the CR scenario could be a tree. In the case of slot filling, it consists of choosing which branch to follow to minimize the search space and provide the best recommendation as soon as possible. We believe that the joint conversational search and recommendation model should leverage such similarity. Ideally, the search system can provide additional information to the user, so that they can be more informed and better guide the recommender system in choosing the most effective branch of the conversation.

Categorizing the conversational systems Based on our analysis, we believe that each conversational system can be characterized according to two axes:

- **exploration/disambiguation.** A pure exploration system is based on a user that asks and the system that responds. The information is available only to the system and the user has to extract from it, by “exploring”. A pure disambiguation system asks questions, and the user responds to help the system disambiguate their interest.
- **search/recommendation:** in the search scenario, the user is interested in collecting information and discovering something on a topic of interest. Vice-versa, when it comes to the recommendation, the focus is on the system providing suggestions based on a set of predefined options (i.e., the catalogue). At the same time, in search, the research is

guided by the user's query, while in recommendation there is no actual query, and the query corresponds to the user's profile.

Following this schema, we can identify five scenarios:

- **full exploration, full search.** This scenario corresponds to the current CS: the user has a piece of information and explores the corpus (guided by the system) to satisfy it.
- **full disambiguation, full recommendation.** This case is the current CR: the user has the (latent) information, and the system tries to disambiguate it by asking questions.
- **full exploration, full recommendation.** The user is not interested in receiving a recommendation but rather in exploring the catalogue of the RS. In a sense, in this scenario, the user is interested in meta-data, such as what are the most popular songs, which movies belong to a certain genre, and so on.
- **full disambiguation, full search.** The user knows a relevant document exists and needs to access it via disambiguation. An example of this scenario is represented by navigational queries used in the CS context: the user knows with certainty which page they are trying to access and guides the system in understanding it as well.
- **Mixed scenarios.** This class of systems includes for example the mixed-initiative in IR: the system is allowed to ask questions to disambiguate their information need. Other mixed situations could be a system that, besides providing recommendations, accompanies them with pieces of text explaining the item(s) in detail. For example, a user searching for a song might receive recommendations and documents describing the story of the recommended songs, details on genres, documents on the artist, and other information derived from a corpus of documents.

Due to its objective, models developed by CAMEO belong to the last category.

3.2. The unifying framework

During our brainstorming at CAMEO retreat, we identified three main strategies to design a joint conversational search and recommendation system:

- **use a unique module.** This approach would rely on using a *Large Language Model (LLM)*, an approach that proved to be effective for both CS and CR tasks. This approach would be relatively straightforward to realize, but it is hard to make it effective and evaluate it.
- **use two separate modules,** with a classifier deciding the type of module to use. In this case, some of the requests will be redirected to the IR engine while some to the RS. The major advantage is the simplicity of the approach: it is easy to define, train, and operationalize the two distinct models, as they are more or less the currently available conversational IR and RS systems. The major challenges regard i) the construction of the classifier and ii) the absence of shared knowledge as it decreases the power of the tool: it is impossible to improve a component based on the interactions of the other.
- **use two separate modules, with shared elements.** The two modules could, for example, be based on the same knowledge base or rely on the same latent representation space where documents and items are represented jointly. Then, when queried by the user,

the system uses a classifier to decide which module to use and filter the latent space based on this. Notice that this approach shares with the previous one the challenges linked to obtaining the classifier. Furthermore, several efforts highlighted greater effectiveness thanks to the usage of a joint search and recommendation space [16, 17, 18, 4, 6].

3.3. Datasets

One of the major challenges in realizing a joint conversational search and recommendation system is the fact that there are no currently available datasets.

The first possibility would be to extend an existing dataset meant for either search or recommendation with additional ground truth for the other task. For example, it could be possible to integrate the CAsT dataset with recommendation components. This requires to partially re-annotate the dataset. Nevertheless, this has the advantage that, while collecting new annotations, we could extend the existing dataset with additional meta-information such as linguistic aspects (the presence of anaphoras), links between utterances (what is the tree structure of the conversation), and external elements such as DBpedia [19] entities. An open issue concerns the construction of a “user profile”, needed to operationalize the recommendation component.

On the other hand, we could assume to extend a conversational recommendation dataset. The problem is less encumbering than in conversational search, as, by using the paradigm described in Section 2, it is possible to use any currently available RS dataset to simulate a conversation. This reduces the realism of the evaluation, but it grants access to a large amount of data. Therefore, it would be necessary to combine the items and categories available in a recommendation dataset, with documents available in a corpus. Then, we can annotate documents for relevance about questions arising on the recommendation dataset’s domain.

Considering the limitations linked to annotating the conversational search dataset for the recommendation task – namely the absence of a user’s profile, we consider the strategy based on turning a recommendation dataset into a set of simulated conversations and annotating such conversations the most effective approach.

4. The CAMEO Project

4.1. The system

Figure1 describes the conceptual architecture of the CAMEO system. It consists not only of a *conversational agent* (O1), the ultimate goal of our research, but also of a *visual analytics environment* (O2), needed to properly develop, analyze, optimize, and operate the conversational agent. Therefore, the CAMEO system targets two main categories of stakeholders: *end users* (O3) who need a conversational agent to access information and carry out their tasks; *researchers and developers* (O3) who need support to properly develop, optimize, and operate their conversational agents, but also *stakeholders* (O3) that can rely on the CAMEO visual analytics to improve their marketing strategies.

The conversational agent (O1) is the core of the system, taking as input utterances issued by the user, processing them, and producing as output the answer. It includes three main components: the *dialog manager*, responsible for the course of the dialog, the *conversational*




Figure 1: The representation of the envisioned CAMEO joint recommendation and search system.

search and *conversational recommendation* engines. The dialog manager interacts with the user, granting high quality answers from the linguistic point of view. Ideally, it could be implemented using a Large Language Model. Below the dialog manager, two separate components, the CS and RS models, are responsible of identifying the contents to be presented to the user. This is in line with our observations during the first CAMEO retreat, where we considered a system with two separate internal subsystems more effective in our specific scenario as well as more easy to evaluate. These three components share an *internal knowledge*, containing short- and long-term user preferences as well as the state of the current dialog and its dependencies. Furthermore, such shared internal knowledge will contain the joint space where represent recommendation items and search documents. The *external knowledge* includes all the external data used by the system, among which corpora of documents, catalogs of items, knowledge bases, and ontologies.

The *visual analytics environment* (O2) will allow researchers and developers to visually and interactively explain the behaviour of the conversational agent, as the dialog progresses. Moreover, it will serve to explore, analyze, and make sense of the performance of the conversational agent, easing the optimization process.

4.2. The task

After taking into account the aspects mentioned above, we consider *product search* the scenario where CAMEO might provide the biggest advantages.

Product search consists of guiding a user into finding a specific product: the user knows what they need, but does not know which specific instance of that product satisfies their needs. The user needs something possibly present in the catalogue and the system, with its knowledge,

guides them in finding the most suited item.

Assume for example a user searching for a generic product. The RS can contextualize the recommendation based on the user profile (e.g., if they prefer low-cost products, colour, material, etc). We can imagine an exploration component which provides information on the search of the user. The system can provide information to expand the user's knowledge, such as detailing the varieties of products available, the differences between brands and so on.

In the conversational scenario, we could consider a two-phase interaction. During the first (expectation) phase the user gathers knowledge about the product: the user leads the conversation and the system responds. In the second (disambiguation) phase the system guides the conversation through questions that allow to reduce the search space (the system leads the conversation and the user responds), driving the user toward a specific instance of the product.

The annotations protocol is similar to the one in IR and RS. The system provides the documents and the annotator assesses whether such documents are relevant to the user's query. For example, if the user's query is "*I would like to buy fruit*", relevant documents are those concerning fruits, such as which types of fruit exist and what are nutritional properties of different fruits. For RS, the evaluation focuses on the fact that the user was recommended an item relevant to them.

Additionally, the system could be integrated with analytics that might help in diversifying the sold products, thus increasing fairness. Assume a scenario in which a producer has certain items for which they wish to bust selling. Given two almost equivalent items – at least based on what the user considers relevant –, it would be desirable if the conversation could nudge the user towards the item that the seller wishes to sell more. Thus, the analytics integrated within CAMEO conversational systems could help in answering questions such as:

- Are there words or sets of words that better lead the user to pick one item or the other?
- Are there sequences of questions that guide the user in choosing one item over the other?
- Are documents that, if presented during the first phase, better nudge the user toward choosing a specific product?
- Is the sentiment of the questions and documents retrieved capable of influencing the "appeal" of an item to a specific user?

5. Conclusion and Future Work

In conversational search, users obtain factual information through direct queries, while in conversational recommendation systems, user preferences are refined through a series of inquiries stimulated by the agent. The CAMEO project retreat served to go in-depth into these tasks, exploring their similarities and disparities. Our findings highlighted the utility of categorizing conversational systems along axes of exploration-disambiguation and search-recommendation for the roles played by the user and the intelligent agent. Moreover, we discussed the need to acquire datasets to train and evaluate joint conversational search and recommendation systems. Looking forward, the CAMEO project aims to tackle the product search domain, envisioning a scenario where the amalgamation of conversational search and recommendation capabilities will prove most impactful.

A lot of future work for CAMEO emerged from our discussion. Firstly, the development of joint conversational search and recommendation systems necessitates a deeper understanding

of user intent and system capabilities. Additionally, creating comprehensive datasets in the product search scenario remains paramount. Finally, the methodology and the algorithmic approaches to seamlessly integrate search and recommendation functionalities will be a focal point, emphasizing efficiency, effectiveness, and appropriate evaluation metrics.

Acknowledgments

This work was partially supported by CAMEO, PRIN 2022 n. 2022ZLL7MW.

References

- [1] V. Stamatis, L. Azzopardi, A. Wilson, VES team at TREC conversational assistance track (cast) 2019, in: E. M. Voorhees, A. Ellis (Eds.), Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC 2019, Gaithersburg, Maryland, USA, November 13-15, 2019, volume 1250 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 2019. URL: <https://trec.nist.gov/pubs/trec28/papers/VES.C.pdf>.
- [2] J. Dalton, C. Xiong, J. Callan, Cast 2020: The conversational assistance track overview, in: E. M. Voorhees, A. Ellis (Eds.), Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16-20, 2020, volume 1266 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 2020. URL: <https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.C.pdf>.
- [3] J. Dalton, C. Xiong, J. Callan, TREC cast 2021: The conversational assistance track overview, in: I. Soboroff, A. Ellis (Eds.), Proceedings of the Thirtieth Text REtrieval Conference, TREC 2021, online, November 15-19, 2021, volume 500-335 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 2021. URL: <https://trec.nist.gov/pubs/trec30/papers/Overview-CAsT.pdf>.
- [4] K. Zhao, Y. Zheng, T. Zhuang, X. Li, X. Zeng, Joint learning of e-commerce search and recommendation with a unified graph neural network, in: K. S. Candan, H. Liu, L. Akoglu, X. L. Dong, J. Tang (Eds.), WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022, ACM, 2022, pp. 1461–1469. URL: <https://doi.org/10.1145/3488560.3498414>. doi:10.1145/3488560.3498414.
- [5] T. Thonet, J. Renders, M. Choi, J. Kim, Joint personalized search and recommendation with hypergraph convolutional networks, in: M. Hagen, S. Verberne, C. Macdonald, C. Seifert, K. Balog, K. Nørnvåg, V. Setty (Eds.), Advances in Information Retrieval - 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10-14, 2022, Proceedings, Part I, volume 13185 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 443–456. URL: https://doi.org/10.1007/978-3-030-99736-6_30. doi:10.1007/978-3-030-99736-6_30.
- [6] Z. Si, Z. Sun, X. Zhang, J. Xu, X. Zang, Y. Song, K. Gai, J. Wen, When search meets recommendation: Learning disentangled search representation for recommendation, in: H. Chen, W. E. Duh, H. Huang, M. P. Kato, J. Mothe, B. Poblete (Eds.), Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information

- Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023, ACM, 2023, pp. 1313–1323. URL: <https://doi.org/10.1145/3539618.3591786>. doi:10.1145/3539618.3591786.
- [7] P. Owoicho, J. Dalton, M. Aliannejadi, L. Azzopardi, J. R. Trippas, S. Vakulenko, TREC cast 2022: Going beyond user ask and system retrieve with initiative and response generation, in: I. Soboroff, A. Ellis (Eds.), Proceedings of the Thirty-First Text REtrieval Conference, TREC 2022, online, November 15-19, 2022, volume 500-338 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 2022. URL: https://trec.nist.gov/pubs/trec31/papers/Overview_cast.pdf.
- [8] G. Faggioli, M. Ferrante, N. Ferro, R. Perego, N. Tonello, Hierarchical dependence-aware evaluation measures for conversational search, in: F. Diaz, C. Shah, T. Suel, P. Castells, R. Jones, T. Sakai (Eds.), SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021, ACM, 2021, pp. 1935–1939. URL: <https://doi.org/10.1145/3404835.3463090>. doi:10.1145/3404835.3463090.
- [9] S. Vakulenko, E. Kanoulas, M. de Rijke, An analysis of mixed initiative and collaboration in information-seeking dialogues, in: J. X. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, Y. Liu (Eds.), Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020, ACM, 2020, pp. 2085–2088. URL: <https://doi.org/10.1145/3397271.3401297>. doi:10.1145/3397271.3401297.
- [10] M. Aliannejadi, L. Azzopardi, H. Zamani, E. Kanoulas, P. Thomas, N. Craswell, Analysing mixed initiatives and search strategies during conversational search, in: G. Demartini, G. Zuccon, J. S. Culpepper, Z. Huang, H. Tong (Eds.), CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, ACM, 2021, pp. 16–26. URL: <https://doi.org/10.1145/3459637.3482231>. doi:10.1145/3459637.3482231.
- [11] D. Jannach, A. Manzoor, W. Cai, L. Chen, A survey on conversational recommender systems, *ACM Computing Surveys (CSUR)* 54 (2021) 1–36.
- [12] C. Gao, W. Lei, X. He, M. de Rijke, T.-S. Chua, Advances and challenges in conversational recommender systems: A survey, *AI Open* 2 (2021) 100–126.
- [13] T. Di Noia, F. M. Donini, D. Jannach, F. Narducci, C. Pomo, Conversational recommendation: Theoretical model and complexity analysis, *Information Sciences* 614 (2022) 325–347.
- [14] A. Iovine, F. Narducci, G. Semeraro, Conversational recommender systems and natural language: A study through the converse framework, *Decision Support Systems* 131 (2020) 113250.
- [15] G. Faggioli, M. Ferrante, N. Ferro, R. Perego, N. Tonello, A dependency-aware utterances permutation strategy to improve conversational evaluation, in: M. Hagen, S. Verberne, C. Macdonald, C. Seifert, K. Balog, K. Nørsvåg, V. Setty (Eds.), Advances in Information Retrieval - 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10-14, 2022, Proceedings, Part I, volume 13185 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 184–198. URL: https://doi.org/10.1007/978-3-030-99736-6_13. doi:10.1007/978-3-030-99736-6_13.
- [16] H. Zeng, S. Kallumadi, Z. Alibadi, R. Frassetto Nogueira, H. Zamani, A personalized dense retrieval framework for unified information access, in: H. Chen, W. E. Duh, H. Huang,

- M. P. Kato, J. Mothe, B. Poblete (Eds.), Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023, ACM, 2023, pp. 121–130. URL: <https://doi.org/10.1145/3539618.3591626>. doi:10.1145/3539618.3591626.
- [17] H. Zamani, W. B. Croft, Joint modeling and optimization of search and recommendation, in: O. Alonso, G. Silvello (Eds.), Proceedings of the First Biennial Conference on Design of Experimental Search & Information Retrieval Systems, Bertinoro, Italy, August 28-31, 2018, volume 2167 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 36–41. URL: <https://ceur-ws.org/Vol-2167/paper2.pdf>.
- [18] H. Zamani, W. B. Croft, Learning a joint search and recommendation model from user-item interactions, in: J. Caverlee, X. B. Hu, M. Lalmas, W. Wang (Eds.), WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020, ACM, 2020, pp. 717–725. URL: <https://doi.org/10.1145/3336191.3371818>. doi:10.1145/3336191.3371818.
- [19] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, Z. G. Ives, Dbpedia: A nucleus for a web of open data, in: K. Aberer, K. Choi, N. F. Noy, D. Allemang, K. Lee, L. J. B. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber, P. Cudré-Mauroux (Eds.), The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007, volume 4825 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 722–735. URL: https://doi.org/10.1007/978-3-540-76298-0_52. doi:10.1007/978-3-540-76298-0_52.