

Grounded and Transparent Response Generation for Conversational Information-Seeking Systems

Weronika Łajewska
University of Stavanger
Stavanger, Norway
weronika.lajewska@uis.no

ABSTRACT

While previous conversational information-seeking (CIS) research has focused on passage retrieval, reranking, and query rewriting, the challenge of synthesizing retrieved information into coherent responses remains. The proposed research delves into the intricacies of response generation in CIS systems. Open-ended information-seeking dialogues introduce multiple challenges that may lead to potential pitfalls in system responses. The study focuses on generating responses grounded in the retrieved passages and being transparent about the system's limitations. Specific research questions revolve around obtaining confidence-enriched information nuggets, automatic detection of incomplete or incorrect responses, generating responses communicating the system's limitations, and evaluating enhanced responses. By addressing these research tasks the study aspires to contribute to the advancement of conversational response generation, fostering more trustworthy interactions in CIS dialogues, and paving the way for grounded and transparent systems to meet users' needs in an information-driven world.

CCS CONCEPTS

• Information systems → Presentation of retrieval results.

KEYWORDS

Conversational search; Conversational response generation; Information nugget annotation

ACM Reference Format:

Weronika Łajewska. 2024. Grounded and Transparent Response Generation for Conversational Information-Seeking Systems. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3616855.3635727>

1 INTRODUCTION

The ever-growing reliance on digital information demands the presence of transparent and trustworthy search systems in our daily interactions. A large fraction of research on conversational information-seeking (CIS) systems to date has focused on the problem of passage retrieval [16], reranking [19], and query rewriting [27]. However, identifying the top relevant passages is only

an intermediate step. Ultimately, the information from these passages would need to be synthesized into a single response in the process of *conversational response generation* [23]. The response returned by the system should encapsulate the most relevant pieces of information in an easily consumable unit [6].

Generative language models have been widely adopted for response generation [14, 28]; however, in the realm of open-ended information-seeking dialogues, the assumption that a user's query can be definitively answered by simply summarizing information from top retrieved passages falls short of reality. System responses are susceptible to various limitations, such as the failure to find a response which may result in hallucinations [10], providing a biased response only partially answering the question [9], or even presenting content with factual errors [26]. Consequently, relying solely on summarizing relevant information may lead to providing users with biased, incomplete, or, worse, incorrect responses [26].

Recognizing that users are responsible for judging the completeness, credibility, and accuracy of information provided by the system, there arises a crucial need to equip them with the necessary tools for objective assessment. Two key elements come into play for this purpose: (1) disclosing the system's confidence in the response and (2) providing transparency on system limitations. By showing how confident the system is in its response, users can assess the reliability of the information provided [15, 22, 24]. This helps differentiate between well-supported responses grounded in specific facts and responses with more uncertainty. Equally important is disclosing the limitations inherent in the response generation process. Users must be informed about potential factors that could have contributed to inaccurate responses, such as unanswerability [5, 21, 25], lack of viewpoint diversification [8, 9], or queries with impossible conditions [13]. Understanding these limitations empowers users to interpret responses more critically and aids in their decision-making process.

RQ1: How to detect factors contributing to incorrect, incomplete, or biased responses? One of the primary concerns in response generation is how to obtain information nuggets [23] (RQ1.1) defined as *minimal, atomic units of relevant information* [18] that contain key pieces of information necessary to answer the user's question. In open-ended information-seeking dialogues, the system's ability to pinpoint crucial details is essential for fact-grounded response generation. We aim to develop methods and techniques to extract these confidence-enriched information nuggets from the top retrieved passages, building foundations for more factual system responses. Given that CIS systems are susceptible to limitations, it raises a question about factors that lead to incomplete, biased, or incorrect responses (RQ1.2). We propose automated techniques to detect underlying factors leading to flawed responses by

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WSDM '24, March 4–8, 2024, Merida, Mexico

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0371-3/24/03.

<https://doi.org/10.1145/3616855.3635727>

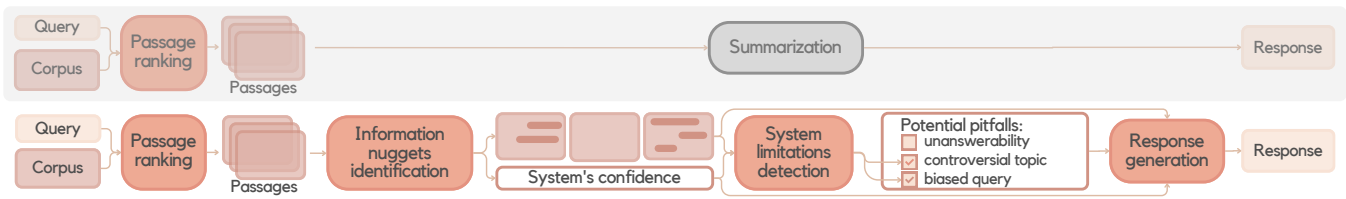


Figure 1: Overview of our approach to conversational response generation in comparison with the standard CIS pipeline.

analyzing queries and information nuggets. This approach aims to unveil patterns indicating response pitfalls, enhancing the system’s ability to identify problematic responses generated from information nuggets.

RQ2: How to generate responses transparent about the system’s confidence and limitations? Providing users with grounded and transparent responses is crucial for fostering trust and empowering users to objectively assess the presented information. We investigate how to generate responses that not only synthesize the requested information but also are grounded in the specific facts identified in the passages and articulate the system’s confidence and limitations. Designing informative responses that transparently present potential pitfalls will promote a more informed user experience and help manage user expectations [2, 20]. The accuracy of generated responses and effectiveness in communicating the system’s confidence and limitations can be evaluated via user studies. By exposing users to different variants of the responses, we aim to evaluate our approach to enhancing the transparency of CIS systems. Additionally, we aim to propose measures for automatic evaluation of the generated responses quantitatively, in terms of response completeness using the identified information nuggets.

By addressing these research tasks, our study aspires to bridge the gap between system-generated responses and user comprehension, fostering a more trustworthy interaction in CIS dialogues.

2 METHODOLOGY

Our proposed response generation module, in contrast to the summarization of top retrieved passages, identifies information nuggets and aggregates them to the response [23]. Additionally, the response (see Fig. 1) is enriched with the information about system’s confidence and the potential factors that could have contributed to flaws.

2.1 Passage Ranking

The main component of the CIS pipeline is passage ranking, retrieving the most relevant passages from the corpus. This study focuses on downstream processing following passage retrieval, and not on enhancing the ranking methods. Thus, we employ the top-performing 2021 TREC Conversational Assistance track submission reproduced in our prior work [11], serving as a strong baseline.

2.2 Information Nuggets Identification

To address RQ1.1, which concerns obtaining confidence-enriched information nuggets, we propose methods for extracting key pieces of information from passages to ensure fact-grounded response generation. In recent work, we have created the CAsT-snippets

dataset [12], enriching the TREC CAsT 2020 [7] and 2022 [17] datasets with snippet-level answer annotations. To ensure data quality, we extensively investigated effective task designs and trade-offs for snippet annotations through crowdsourcing, considering various interfaces and worker qualifications. The primary challenge in gathering snippet annotations centered around the quality control process, which couldn’t be automated due to the inherent nature of the task. In comparison to other related datasets (SaaC [23], QuaC [5]), CAsT-snippets not only contains more snippets annotated for each input text, but they are also longer on average. In future work, we aim to focus on the automatic identification of information nuggets in passages and determining the system’s confidence in the selected snippets. Information nugget identification is not a binary decision but a more granular task, as snippets can vary in relevance and complexity. They may contain exact facts fully answering the question, additional enriching details, or only address some aspects of the question. This granular nature of *answerability* presents a major challenge and plays a crucial role in incorporating the system’s confidence score in the generated response.

2.3 Detecting System Limitations

To address RQ1.2, which concerns the identification of potential system limitations while generating the response from identified information nuggets, we propose to detect factors that may lead to incomplete, or incorrect responses. As we seek to detect these limitations from the system’s perspective, we lack external ground truth for comparing different systems, particularly for evaluating the completeness of responses. In recent work, we proposed a mechanism for detecting unanswerable questions where the correct answer is not present in the corpus or cannot be retrieved. We develop a method that employs a sentence-level classifier to detect if the answer is present, then aggregates these predictions on the passage level, and eventually across the top-ranked passages to arrive at a final answerability estimate. For training and evaluation, we extend the CAsT-snippets dataset with answerability labels on the sentence, passage, and ranking levels. By assessing if a question can be at least partially answered based on the information contained in the top-ranked passages, we can mitigate the risk of generating responses from irrelevant or non-existent answers, reducing the occurrence of hallucinations [10]. In our future work, we aspire to extend our detection capabilities to encompass other limitations that influence the final step of response generation. These limitations include but are not limited to, lack of viewpoint diversification when addressing controversial topics [8], partial unanswerability, temporal considerations [3], biased queries [1], the subjectivity of the source text, and the lack of expert/background knowledge required to infer the answer.

2.4 Revealing System Confidence and Limitations

RQ2 concerns the generation of grounded responses that reveal system confidence and limitations. In recent work, we ran two crowdsourcing experiments investigating user perceptions of problems of unanswerable questions and incomplete responses in CIS systems. The results show that users find it easier to detect issues related to viewpoint diversity and response bias compared to factual errors and source validity. Our findings imply that simple source attribution is insufficient for effective system interaction, suggesting the need for explicit communication of potential inaccuracies to enhance users' assessments of the presented information. Providing users with transparent responses that acknowledge the system's limitations is paramount for fostering trust and empowering users to make informed judgments [2, 20]. We aim to generate responses that (1) synthesize the requested information, (2) ground it in specific facts identified in the passages, (3) articulate the system's confidence, and (4) reveal the system's limitations. By openly showing users the system's limitations, we aim to encourage a more critical examination of the responses, prompting them to verify the provided information. We consider diverse methods for conveying this information to users, including natural language incorporation [22], UI elements [15], and a granular confidence scale [24], drawing from prior work in recommender systems that can be adapted to the field of conversational response generation.

Evaluating generated responses in CIS systems in terms of completeness is challenging, as the system cannot be fully aware of what it does not know, even with identified information nuggets [4]. However, we recognize that snippet-level answer annotation can be a crucial first step toward automatically evaluating responses quantitatively, based on the relevant information nuggets included [18]. The recent edition of TREC CAsT proposes to evaluate summarized passages based on relevance, naturalness, and conciseness [17]. We aim to extend the evaluation of generated responses from the perspective of transparency and grounding. The question about other aspects of response generation that should be considered, in addition to transparency and grounding, remains. User feedback from our crowdsourcing experiments related to challenges in CIS response generation draws attention to the credibility of the sources, as well as the completeness, usefulness, and subjectivity of provided information that impact users' overall satisfaction.

3 RESEARCH ISSUES

In Section 2, we established a set of simplifying assumptions for our proposed approaches. Now, we highlight the open questions we have identified, seeking assistance in prioritizing them due to their expansive nature and the time constraints of the Ph.D. study. We seek for guidance on the following issues related to our research questions: (RQ1.1) How can the system articulate its confidence in the identified snippets without information about the scope of the complete response? (RQ1.2) Should we focus on identifying more limitations or delve deeper into specific ones? (RQ1.2) Can limitation detection be addressed holistically? (RQ2) Are the attempts to showcase system limitations through additional user interface elements justifiable, or does the superiority still lie in the simplicity of generating natural language responses?

REFERENCES

- [1] Leif Azzopardi. 2021. Cognitive Biases in Search: A Review and Reflection of Cognitive Biases in Information Retrieval. In *CHIIR '21*.
- [2] Leif Azzopardi, Mateusz Dubiel, M. Halvey, and Jeffrey Dalton. 2018. Conceptualizing agent-human interactions during the conversational search process. In *CAIR '18*.
- [3] Ricardo Campos, Gaël Dias, Alípio M. Jorge, and Adam Jatowt. 2015. Survey of Temporal Information Retrieval and Related Applications. *ACM Comput. Surv.* 47, 2 (2015).
- [4] Xiuyi Chen, Feilong Chen, Fandong Meng, Peng Li, and Jie Zhou. 2021. Unsupervised Knowledge Selection for Dialogue Generation. In *ACL-IJCNLP '21*.
- [5] Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question Answering in Context. In *EMNLP '18*.
- [6] J. Shane Culpepper, Fernando Diaz, and Mark D. Smucker. 2018. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *ACM SIGIR Forum* 52, 1 (2018).
- [7] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. CAsT 2020: The Conversational Assistance Track Overview. In *TREC '20*.
- [8] Tim Draws, Nava Tintarev, Ujwal Gadiraju, Alessandro Bozzon, and Benjamin Timmermans. 2021. This Is Not What We Ordered: Exploring Why Biased Search Result Rankings Affect User Attitudes on Debated Topics. In *SIGIR '21*.
- [9] Ruoyuan Gao and Chirag Shah. 2020. Toward creating a fairer ranking in search engine results. *Inf. Process. Manag.* 57, 1 (2020).
- [10] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* 55, 12 (2023).
- [11] Weronika Lajewska and Krisztian Balog. 2023. From Baseline to Top Performer: A Reproducibility Study of Approaches at the TREC 2021 Conversational Assistance Track. In *ECIR '23*.
- [12] Weronika Lajewska and Krisztian Balog. 2023. Towards Filling the Gap in Conversational Search: From Passage Retrieval to Conversational Response Generation. In *CIKM '23*.
- [13] Gyeongbok Lee, Seung-won Hwang, and Hyunsouk Cho. 2020. SQuAD2-CR: Semi-supervised Annotation for Cause and Rationales for Unanswerability in SQuAD 2.0. In *LREC '20*.
- [14] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *NIPS'20*.
- [15] Zhuoran Lu and Ming Yin. 2021. Human Reliance on Machine Learning Models When Performance Feedback is Limited: Heuristics and Risks. In *CHI '21*.
- [16] Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. Sparse, Dense, and Attentional Representations for Text Retrieval. *TACL* 9 (2021).
- [17] Paul Owoicho, Jeffrey Dalton, Mohammad Aliannejadi, Leif Azzopardi, Johanne R Trippas, and Svitlana Vakulenko. 2022. TREC CAsT 2022: Going Beyond User Ask and System Retrieve with Initiative and Response Generation. In *TREC '22*.
- [18] Virgil Pavlu, Shahzad Rajput, Peter B. Golbus, and Javed A. Aslam. 2012. IR system evaluation using nugget-based test collections. In *WSDM '12*.
- [19] Ronak Pradeep, Rodrigo Nogueira, and Jimmy Lin. 2021. The Expando-Monoduo Design Pattern for Text Ranking with Pretrained Sequence-to-Sequence Models. arXiv:2101.05667 [cs.LG].
- [20] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *CHIIR '17*.
- [21] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. In *ACL '18*.
- [22] Amy Rechkemmer and Ming Yin. 2022. When Confidence Meets Accuracy: Exploring the Effects of Multiple Performance Indicators on Trust in Machine Learning Models. In *CHI '22*.
- [23] Pengjie Ren, Zhumin Chen, Zhaochun Ren, Evangelos Kanoulas, Christof Monz, and Maarten De Rijke. 2021. Conversations with Search Engines: SERP-based Conversational Response Generation. *ACM TOIS* 39, 4 (2021).
- [24] Guy Shani, Lior Rokach, Bracha Shapira, Sarit Hadash, and Moran Tangi. 2013. Investigating confidence displays for top-N recommendations. *J. Am. Soc. Inf. Sci.* 64, 12 (2013).
- [25] Elinor Sulem, Jamaal Hay, and Dan Roth. 2022. Yes, No or IDK: The Challenge of Unanswerable Yes/No Questions. In *NAACL-HLT '22*.
- [26] Xiangru Tang, Alexander Fabbri, Haoran Li, Ziming Mao, Griffin Adams, Borui Wang, Asli Celikyilmaz, Yashar Mehdad, and Dragomir Radev. 2022. Investigating Crowdsourcing Protocols for Evaluating the Factual Consistency of Summaries. In *NAACL-HLT '22*.
- [27] Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question Rewriting for Conversational Question Answering. In *WSDM '21*.
- [28] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *ACL '20*.