







# Estimating the Usefulness of Clarifying Questions and Answers for Conversational Search

Ivan Sekulić<sup>1</sup>(✉) , Weronika Łajewska<sup>2</sup> , Krisztian Balog<sup>2</sup> ,  
and Fabio Crestani<sup>1</sup> 

<sup>1</sup> Università della Svizzera italiana, Lugano, Switzerland

{ivan.sekulic,fabio.crestani}@usi.ch

<sup>2</sup> University of Stavanger, Stavanger, Norway

{weronika.lajewska,krisztian.balog}@uis.ch

**Abstract.** While the body of research directed towards constructing and generating clarifying questions in mixed-initiative conversational search systems is vast, research aimed at processing and comprehending users' answers to such questions is scarce. To this end, we present a simple yet effective method for processing answers to clarifying questions, moving away from previous work that simply appends answers to the original query and thus potentially degrades retrieval performance. Specifically, we propose a classifier for assessing usefulness of the prompted clarifying question and an answer given by the user. Useful questions or answers are further appended to the conversation history and passed to a transformer-based query rewriting module. Results demonstrate significant improvements over strong non-mixed-initiative baselines. Furthermore, the proposed approach mitigates the performance drops when non useful questions and answers are utilized.

**Keywords:** Conversational search · Mixed initiative · Clarifying questions

## 1 Introduction

The goal of a conversational search (CS) system is to satisfy the user's information need. To this end, several aspects of CS systems have enjoyed significant advancements, including conversational passage retrieval [22], query rewriting [21], and intent detection [13]. In recent years, the mixed-initiative paradigm of conversational search has attracted significant research attention [2, 16, 18]. Under the mixed-initiative paradigm, the CS system can at any point in a conversation offer suggestions or ask clarifying questions.

Asking clarifying questions has been identified as an invaluable component of modern-day CS systems [14]. Such questions are directed at users and aim to elucidate their underlying information needs. While there is a growing body of research on constructing and generating clarifying questions [2, 19, 23], work

aimed at processing and comprehending users' answers to such questions is scarce. Nonetheless, recent research suggests their usefulness by demonstrating improvements in passage retrieval performance after asking a clarifying question and receiving an answer [1].

To bridge the aforementioned research gap, we make a first step towards processing the answers given to clarifying questions. We hypothesize that not all information acquired through such interactions with the user would benefit the CS system, i.e., yield improvements in retrieval effectiveness. Thus, the main novelty of our approach is that we do not blindly utilize the questions and the answers, but only when they are deemed to be useful. Specifically, we focus on the task of conversational passage retrieval and design a classifier aimed at assessing usefulness of the asked clarifying question and the provided answer. We utilize the question or the answer only if they are deemed useful, by appending them to the conversational history and employing a query rewriting method to attain a more information-dense query. Results on the TREC 2022 Conversational Assistance Track (CAST'22) [11] demonstrate significant improvements in passage retrieval performance with the use of enhanced queries, as opposed to a non-mixed-initiative retrieval system (12% and 3% relative improvement in terms of Recall@1000 and nDCG, respectively). Further, when contrasting our approach to an established method that simply appends the prompted clarifying question and its answer to the original query [1], we observe differences in performance. Specifically, if neither the question nor the answer are deemed useful, but still used, there is a relative performance decrease of 13% in terms of nDCG@3, compared to non-mixed-initiative baselines. In other words, it is better not to use any information provided by such questions and answers, than to use it wrongly. Our contributions can be summarized as follows:

- We propose a simple, yet effective, method for processing answers to clarifying questions. The method is based on classifying usefulness of the prompted question and the given answer.
- We identify scenarios where asking clarifying questions resulted in improved passage retrieval, and where it decreased the retrieval performance.

## 2 Related Work

To facilitate further research in conversational search (CS), the TREC Conversational Assistance Track (CAST) [3] aims to provide a reusable benchmark composed of several pre-defined conversational trajectories over a variety of topics. The most recent edition, CAST'22, offers a subtask oriented towards mixed-initiative (MI) interactions [11]. Under the MI paradigm of CS, the system can at any point of a conversation take initiative and prompt the user with various suggestions or questions [14]. One of the most prominent usages of mixed-initiative is asking clarifying questions with a goal of elucidating users underlying information need [2]. Recent research demonstrates a positive impact of clarifying questions both on user experience [5,24] and retrieval performance [1]. Although other collections with clarifying questions and answers exist, most

notably ClariQ [1], in this work, we focus on the aforementioned CAsT’22 as we additionally have control over which kind of question is being asked. In general, two streams of approaches to constructing clarifying questions exist: (1) select an appropriate question from a pre-defined pool of questions [1, 2, 11, 16, 18]; (2) generate the question [9, 19, 23]. However, despite the abundance of research on clarifying question construction, researched aimed at processing users’ answers to such questions is scarce. To bridge this gap, Krasakis et al. [6] conduct an analysis of users’ answers and find that they vary in polarity and length, as well as that retrieval effectiveness is often hurt. Thus, in this work, we aim to automatically assess their usefulness, with a goal of mitigating this undesired effect.

### 3 Methodology

In this section, we formally define the task of conversational passage retrieval under the mixed-initiative (MI) paradigm and present our methods for each of the components of the task, i.e., query rewriting, clarifying question selection, answer processing, and passage retrieval.

#### 3.1 Task Formulation

At a current conversational turn  $t$ , given a user utterance  $u^t$  and a conversation history  $H = [(u^1, s^1), \dots, (u^{t-1}, s^{t-1})]$ , the task is to generate a system response  $s^t$ . For clarity, we omit the superscript  $t$  from the subsequent definitions. In MI CS systems, the system’s response  $s$  can either be a clarifying question  $s_{cq}$  or a ranked list of passages  $s_D$ ,  $D = [d_1, d_2, \dots, d_N]$ , where  $N$  is the number of passages retrieved and  $d_i$  is the  $i$ -th passage in the list. Similarly, user utterance  $u$  can take form of a query  $u_q$  or an answer  $u_a$  to system’s question  $s_{cq}$ . Modeling other types of user utterances, such as explicit feedback, falls out of the scope of this study. Following prior work [21], the first task, i.e., query rewriting, is aimed towards resolution of the user query  $u_q$  in the context of the conversation history, resulting in  $u'_q = \gamma(u_q|H)$ , where  $\gamma$  is a query rewriting method.

Following the MI setting introduced at TREC CAsT’22 [11], we address the problem of conversational passage retrieval through the following three components: (i) Produce system utterance  $s_{cq}$  by selecting an appropriate clarifying question  $cq$  from a given pool of questions  $PQ$ ; (ii) Process the given answer  $u_a$  and incorporate relevant information to the current query, resulting in  $u''_q = \theta(u'_q, s_{cq}, u_a)$ ; (iii) Return a ranked list of passages  $s_D$ . Next, we define our approaches to the described components. We note that a clarifying question might be needed only for ambiguous, faceted, or unclear user requests. Thus, for queries not requiring clarification, the system might opt to return a ranked list of passages without asking further questions. However, following the setup enabled by CAsT’22 track, we do not explicitly model clarification need and thus design a system that prompts the user with a clarifying question at each turn.

**Table 1.** Examples of annotated subset of ClariQ, indicating cases when clarifying question, answer, both, or neither are useful.

Query (initial request)	Clarifying Question	Answer	Useful	Prevalence
I'm looking for information on hobby stores	Do you want to know hours of operation?	No	<i>None</i>	32%
Tell me information about computer programming	Are you interested in a coding bootcamp?	No, I want to know what career options programmers have	<i>Answer</i>	53%
Find me map of USA	Do you want to see a map of US territories?	Yes	<i>Question</i>	11%
All men are created equal	Would you like to know more about the declaration of independence?	Yes, I'd like to know who wrote it	<i>Question and answer</i>	6%

### 3.2 Clarifying Question Selection

For each query  $u'_q$ , we rank the potential candidates  $cq_i$  based on their semantic similarity to  $u'_q$ . Specifically, we utilize a T5 model fine-tuned on the CANARD dataset [4], available at HuggingFace,<sup>1</sup> as our  $\gamma$  rewriting function, which yields a resolved utterance  $u'_q$ . To rank the potential candidates, we use MPNet [20] from SentenceTransformers [17], trained for semantic matching. We select  $cq_i$  with the highest score, as predicted by the MPNet:  $s_{cq} = \operatorname{argmax}_{cq_i \in PQ_f} \text{MPNet}(u'_q, cq_i)$ , where  $PQ_f$  is a pool of clarifying questions with potentially misleading, unreliable, and faulty questions automatically filtered from the pool [8].

### 3.3 Answer Processing

In this subsection, we describe our novel usefulness-based approach to processing answers given to the asked clarifying questions. To address the issue, we move away from previous approaches that simply append the question and the answer to the original query [1, 2], regardless of the actual information gain. In fact, a recent study by Krasakis et al. [6] demonstrated that such practice can cause a decrease in retrieval effectiveness. Moreover, they show that multi-word answers are informative (e.g., “yes, I’m looking for info on spiders in Europe”), thus improving retrieval performance. Similarly, short negative answers are not informative (e.g., “no”), while multi-word negative answers are (e.g., “no, I’m interested in buying aquarium cleaner”). Thus, we define four possible actions, based on the current resolved utterance  $u'_q$ , the clarifying question asked  $s_{cq}$ , and the answer  $u_a$ :

1. In case the answer is affirmative (e.g., “yes” or “Yes, that is what I’m looking for”), we expand the  $u'_q$  by appending the clarifying question asked.
2. In case the answer is deemed useful, i.e., the underlying information need is explained in greater detail, we expand  $u'_q$  by appending the answer.
3. In case the answer is affirmative and it provides additional details, we expand  $u'_q$  with both the clarifying question and the answer.

<sup>1</sup> <https://huggingface.co/castorini/t5-base-canard>.

4. If neither (1), (2), nor (3) is the case, we do not expand the utterance.

Examples of the described cases are presented in Table 1 and are all aimed at incorporating additional useful information to the current utterance. Formally:

$$u''_q = \begin{cases} u'_q, & \psi(u'_q, s_{cq}, u_a) = 0 \\ \phi(u'_q, s_{cq}), & \psi(u'_q, s_{cq}, u_a) = 1 \\ \phi(u'_q, u_a), & \psi(u'_q, s_{cq}, u_a) = 2 \\ \phi(u'_q, u_a, s_{cq}), & \psi(u'_q, s_{cq}, u_a) = 3 \end{cases} \quad (1)$$

where  $\psi(u'_q, s_{cq}, u_a)$  is a usefulness classifier, which aims to predict which of the above described actions to take. The labels 0, 1, 2, and 3 correspond to neither  $s_{cq}$  or  $u_a$  were deemed useful,  $s_{cq}$  was deemed useful,  $u_a$  was deemed useful, and both were useful, respectively. Similarly to Sect. 3.1, the function  $\phi$  rewrites the original query given the context, in this case  $s_{cq}$  or  $u_a$ .

Specifically, to model  $\psi$ , we fine-tune a large transformer-based model, namely T5 [15], for multi-class classification. To fine-tune the classifier, we manually annotate a portion of ClariQ (150 samples) for the specific aforementioned cases. The annotations were performed by two authors of the paper with an inter-annotator agreement Cohen’s kappa of 0.89. The differences in annotations were discussed and resolved consensually. Examples of annotations are presented in Table 1 and classification performance is reported in Sect. 4. We dub our novel mixed-initiative classifier-based method *MI-Clf*. Moreover, we assess the prevalence of each of the cases, and find, as presented in Table 1, that 68% of interactions contain new, useful information. In the other 32% of the cases, the answer simply negates the prompted clarifying question. Although this interaction also provides valuable insights into the user’s information need, current approaches would expand the query by appending the prompted clarifying question and the answer. However, such an expanded query contains terms the user is not interested in, which can potentially degrade retrieval performance. We compare our proposed method to such a baseline, which always extends the query as:  $u''_q = \phi(u'_q, s_{cq}, u_a)$ . This method is dubbed *MI-All*.

*Passage Retrieval.* Finally, the rewritten utterance  $u''_q$  is fed into a standard two-stage retrieve-and-rerank pipeline [7]. We utilize BM25 ( $k_1 = 0.95$ ,  $b = 0.45$ ) with RM3, where the initial query is extended with the highest-weighting terms from top- $k$  scoring passages ( $k = 10$  and number of terms  $m = 10$ ). Next, we use a point-wise monoT5 re-ranker [10] to re-rank the top 1000, followed by a pair-wise duoT5 re-ranker [12] to additionally re-rank the top 50 passages. The non-mixed-initiative baseline, dubbed *DuoT5*, uses the same retrieval pipeline.

## 4 Results

*Usefulness Classifier.* The proposed usefulness classifier, described in Sect. 3.3, achieves an average macro- $F_1$  score of 0.75 and accuracy of 89% in a stratified 5-fold evaluation on the aforementioned annotated subset of ClariQ. Next, we

**Table 2.** Performance of baselines and our mixed-initiative approaches on CAsT’22.

Approach/RunID	R@1000	MAP	MRR	NDCG	NDCG@3	NDCG@5
BM25_T5_automatic	0.3244	0.1498	0.5272	0.2987	0.3619	0.3443
BM25_T5_manual	0.4651	0.2309	0.7155	0.4228	0.5031	0.4831
our_baseline (DuoT5)	0.3846	0.1680	0.4990	0.3392	0.3593	0.3502
+MI-All	<b>0.4441</b>	0.1741	<b>0.5297</b>	0.3594	<b>0.3722</b>	0.3508
MI-Clf	0.4302	<b>0.1776</b>	0.5144	<b>0.3613</b>	0.3697	<b>0.3581</b>

employ the trained classifier to predict the usefulness of  $(u'_q, s_{cq}, u_a)$  at each turn in the CAsT’22 dataset. The question  $s_{cq}$  was classified as useful in 28% of turns, while the answer  $u_a$  in 37%. In the rest 35% of the cases, neither was predicted to be useful. While the distribution of the predictions is similar to the prevalence in human-annotated data reported in Table 1, some differences can be observed. For example, in CAsT’22 28% of the clarifying questions were deemed useful, as opposed to the 13% in ClariQ.

*Retrieval Performance.* Results of the end-to-end conversational passage retrieval task, after the applied mixed-initiative answer processing methods (*MI-All* and *MI-Clf*) are presented in Table 2. For reference, we also include the organizers’ baselines in the table. We make several observations from the presented results. First, both methods that utilize mixed-initiative show improvements over the DuoT5 method. This confirms previous findings on the positive impact of clarifications in conversational search. Second, differences between *MI-All* and *MI-Clf* are not statistically significant, across all metrics. However, we note that our classifier-based method utilizes clarifying question or the answer only when deemed useful, which is in about 70% of the cases in CAsT’22. On the contrary, *MI-All* always utilizes both the clarifying question and the answer. The equal performance of the two methods suggests that our usefulness classifier successfully includes only relevant information.

*Analysis.* In cases where the usefulness classifier predicted that neither the clarifying question  $s_{cq}$  nor the answer  $u_a$  is useful, we observe a drop of the *MI-All* method’s retrieval performance, in terms of nDCG@3 (−13%). Recall, however, is not impacted by incorporating potentially not useful information and even shows a slight increase (+3.3%). As this method always appends both  $s_{cq}$  and  $u_a$  to the query  $u'_q$ , the performance drop is expected, especially in the re-ranking stage, as the re-ranker might be confused by the additional non-relevant information. Moreover, for both MI methods, we observe higher performance gains when the answer is useful (+19.8% for *MI-All* and +23.4% for *MI-Clf* in recall), compared to cases when the question is useful (+12.5% for *MI-All* and +8.5% *MI-Clf* in recall). This could be explained by the fact that users’ answers are deemed useful when they are longer and thus provide more detail on the underlying information need [6]. On the contrary, a clarifying question can be deemed

useful even when tangibly addressing the user’s need. In other words, a good clarifying question can make a small step towards elucidating the user’s information need. However, the user’s answer can contain detailed expression of their information need, thus making further gains.

## 5 Conclusion

In this study, we proposed a classifier-based method, *MI-Clf*, for processing answers to clarifying questions in conversational search, which extends the original query only when either is deemed useful. Results on the TREC CAsT’22 dataset demonstrate clear improvements of the *MI-Clf* method over non-mixed-initiative baselines (+12% and +3% relative improvement in terms of Recall@1000 and nDCG). Moreover, we observed a performance drop for established methods that always use both the clarifying question and the answer, in cases where neither is useful (−13% in terms of nDCG@3), thus incorporating noisy information. This study makes the first steps towards improved answer processing methods.

**Acknowledgments.** This research was partially supported by the Norwegian Research Center for AI Innovation, NorwAI (Research Council of Norway, project number 309834).

## References

1. Aliannejadi, M., Kiseleva, J., Chuklin, A., Dalton, J., Burtsev, M.: Building and evaluating open-domain dialogue corpora with clarifying questions. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, pp. 4473–4484 (2021)
2. Aliannejadi, M., Zamani, H., Crestani, F., Croft, W.B.: Asking clarifying questions in open-domain information-seeking conversations. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, pp. 475–484 (2019)
3. Dalton, J., Xiong, C., Callan, J.: TREC CAsT 2019: the conversational assistance track overview. In: The Twenty-Eighth Text REtrieval Conference Proceedings, TREC 2019 (2019)
4. Elgohary, A., Peskov, D., Boyd-Graber, J.: Can you unpack that? Learning to rewrite questions-in-context. In: Empirical Methods in Natural Language Processing, EMNLP 2019, pp. 5918–5924 (2019)
5. Kiesel, J., Bahrami, A., Stein, B., Anand, A., Hagen, M.: Toward voice query clarification. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, pp. 1257–1260 (2018)
6. Krasakis, A.M., Aliannejadi, M., Voskarides, N., Kanoulas, E.: Analysing the effect of clarifying questions on document ranking in conversational search. In: Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval, ICTIR 2020, pp. 129–132 (2020)

7. Lajewska, W., Balog, K.: From baseline to top performer: a reproducibility study of approaches at the TREC 2021 conversational assistance track. In: Kamps, J., et al. (eds.) *Proceedings of the 45th European Conference on Information Retrieval, ECIR 2023*, vol. 13982, pp. 177–191. Springer, Cham (2023). [https://doi.org/10.1007/978-3-031-28241-6\\_12](https://doi.org/10.1007/978-3-031-28241-6_12)
8. Lajewska, W., Bernard, N., Kostric, I., Sekulić, I., Balog, K.: The University of Stavanger (IAI) at the TREC 2022 conversational assistance track. In: *Proceedings of the Thirty-First Text REtrieval Conference, TREC 2022* (2023)
9. Majumder, B.P., Rao, S., Galley, M., McAuley, J.: Ask what’s missing and what’s useful: improving clarification question generation using global knowledge. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2021*, pp. 4300–4312 (2021)
10. Nogueira, R., Jiang, Z., Lin, J.: Document ranking with a pretrained sequence-to-sequence model. In: *Findings of the Association for Computational Linguistics, EMNLP 2020*, pp. 708–718 (2020)
11. Owoicho, P., Dalton, J., Aliannejad, M., Azzopardi, L., Trippas, J.R., Vakulenko, S.: TREC CAsT 2022: going beyond user ask and system retrieve with initiative and response generation. In: *The Thirty-First Text REtrieval Conference Proceedings, TREC 2022* (2022)
12. Pradeep, R., Nogueira, R., Lin, J.: The expando-mono-duo design pattern for text ranking with pretrained sequence-to-sequence models. *arXiv cs.IR/2101.05667* (2021)
13. Qu, C., Yang, L., Croft, W.B., Zhang, Y., Trippas, J.R., Qiu, M.: User intent prediction in information-seeking conversations. In: *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval, CHIIR 2019*, pp. 25–33 (2019)
14. Radlinski, F., Craswell, N.: A theoretical framework for conversational search. In: *Proceedings of the 2017 Conference on Human Information Interaction and Retrieval, CHIIR 2017*, pp. 117–126 (2017)
15. Raffel, C., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* **21**(1), 5485–5551 (2020)
16. Rao, S., Daumé III, H.: Learning to ask good questions: ranking clarification questions using neural expected value of perfect information. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2018*, pp. 2737–2746 (2018)
17. Reimers, N., Gurevych, I.: Sentence-BERT: sentence embeddings using Siamese BERT-networks. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pp. 3982–3992 (2019)
18. Rosset, C., et al.: Leading conversational search by suggesting useful questions. In: *Proceedings of The Web Conference 2020, WWW 2020*, pp. 1160–1170 (2020)
19. Sekulić, I., Aliannejadi, M., Crestani, F.: Towards facet-driven generation of clarifying questions for conversational search. In: *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR 2021*, pp. 167–175 (2021)
20. Song, K., Tan, X., Qin, T., Lu, J., Liu, T.Y.: MPNet: masked and permuted pre-training for language understanding. *arXiv preprint arXiv:2004.09297* (2020)



21. Vakulenko, S., Voskarides, N., Tu, Z., Longpre, S.: A comparison of question rewriting methods for conversational passage retrieval. In: Hiemstra, D., Moens, M.-F., Mothe, J., Perego, R., Potthast, M., Sebastiani, F. (eds.) ECIR 2021. LNCS, vol. 12657, pp. 418–424. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-72240-1\\_43](https://doi.org/10.1007/978-3-030-72240-1_43)
22. Yu, S., Liu, Z., Xiong, C., Feng, T., Liu, Z.: Few-shot conversational dense retrieval. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2021, pp. 829–838 (2021)
23. Zamani, H., Dumais, S., Craswell, N., Bennett, P., Lueck, G.: Generating clarifying questions for information retrieval. In: Proceedings of The Web Conference 2020, WWW 2020, pp. 418–428 (2020)
24. Zamani, H., et al.: Analyzing and learning from user interactions for search clarification. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2020, pp. 1181–1190 (2020)