Conversational Search with Mixed-Initiative - Asking Good Clarification Questions backed-up by Passage Retrieval

Yosi Mass, Doron Cohen, Asaf Yehudai and David Konopnicki

IBM Research AI

Haifa University, Mount Carmel, Haifa, HA 31905, Israel {yosimass, doronc, davidko}@il.ibm.com Asaf.Yehudai@ibm.com

Abstract

We deal with the scenario of conversational search, where user queries are under-specified or ambiguous. This calls for a mixed-initiative setup. User-asks (queries) and system-answers, as well as system-asks (clarification questions) and user response, in order to clarify her information needs. We focus on the task of selecting the next clarification question, given the conversation context. Our method leverages passage retrieval from a background content to fine-tune two deep-learning models for ranking candidate clarification questions. We evaluated our method on two different use-cases. The first is an open domain conversational search in a large web collection. The second is a taskoriented customer-support setup. We show that our method performs well on both use-cases.

1 Introduction

A key task in information and knowledge discovery is the retrieval of relevant information given the user's information need (usually expressed by a query). With the abundance of textual knowledge sources and their diversity, it becomes more and more difficult for users, even expert ones, to query such sources and obtain valuable insights.

Thus, users need to go beyond the traditional ad-hoc (one-shot) retrieval paradigm. This requires to support the new paradigm of conversational search – a sophisticated combination of various mechanisms for exploratory search, interactive IR, and response generation. In particular, the conversational paradigm can support mixed-initiative: namely, the traditional user asks - system answers interaction in addition to system-asks (clarification questions) and user-answers, to better guide the system and reach the information needed (Krasakis et al., 2020).

Existing approaches for asking clarification questions include *selection* or *generation*. In the selection approach, the system selects clarification questions from a pool of pre-determined questions (Aliannejadi et al., 2019). In the generation approach, the system generates clarification questions using rules or using neural generative models (Zamani et al., 2020).

In this work we focus on the selection task. While the latter (i.e., generation) may represent a more realistic use-case, still there is an interest in the former (i.e., selection) as evident by the Clarifying Questions for Open-Domain Dialogue Systems (ClariQ) challenge (Aliannejadi et al., 2020). Moreover, the selection task represents a controlled and less noisy scenario, where the pool of clarifications can be mined from e.g., query logs.

In this paper we deal with content-grounded conversations. Thus, a conversation starts with an initial user query, continues with several rounds of conversation utterances (0 or more), and finally ends with one or more documents being returned to the user. Some of the agent utterances are marked as clarification questions.

The task at hand is defined as follows. Given a conversation context up to (and not including) a clarification-question utterance, predict the next clarification question. A more formal definition is given in Section 3.2 below.

Intuitively, clarification questions should be used to distinguish between several possible intents of the user. We approximate those possible intents through passages that are retrieved from a given corpus of documents. A motivating example from the (Aliannejadi et al., 2020) challenge is given in Figure 1. The user wants to get information about the topic *all men are created equal*. Through the retrieved passage, the system can ask the mentioned clarification questions.

We use two deep-learning models. The first one learns an association between conversation context and clarification questions. The second learns an association between conversation context, candidate passages and clarification questions. Evaluation was done on two different use-cases. The first one is an open domain search in a large web corpus (Aliannejadi et al., 2020). The second is an internal task-oriented customer-support setup, where users ask technical questions. We show that our method performs well on both use-cases.



Figure 1: A motivating example

2 Related work

We focus on works that deal with clarificationquestions selection. Aliannejadi et al. (2019) describes a setup very similar to ours for the aforementioned task. They apply a two-step process. In the first step, they use BERT (Devlin et al., 2019) to retrieve candidate clarification questions and, in the second step, they re-rank the candidates using multiple sources of information. Among them are the scores of retrieved documents using the clarification questions. However, they do not look at passage content as we do.

The ClariQ¹ challenge organized a competition for selecting the best clarification questions in an open-domain conversational search. The system by NTES_ALONG (Ou and Lin, 2020) was ranked first. They first retrieve candidate clarification questions and then re-rank them using a ROBERTA (Liu et al., 2019) model, that is fine-tuned on the relation between a query and a clarification question. Unlike our method, they do not exploit passage content.

In Rao and Daumé III (2018), they select clarification questions using the expected value of perfect information, namely a good question is one whose expected answer will be useful. They do not assume a background corpus of documents.

3 Clarification-questions Selection

3.1 Problem definition

A conversation C is a list of utterances, $C = \{c_0, ..., c_n\}$ where c_0 is the initial user query. Each

utterance has a speaker which is either a user or an agent.² Since we deal with content-grounded conversations, the last utterance is an agent utterance, that points to a document.

We further assume that agent utterances are tagged with a *clarification flag* where a value of 1 indicates that the utterance is a clarification question. This flag is either given as part of the dataset (e.g., in the open domain dataset, ClariQ) or is derived automatically by using a rule-based model or a classifier. We discuss such rules for the second task-oriented customer-support dataset (see Section 4.1 below).

The **Clarification-questions Selection** task is defined as follows. Given a conversation context $C^{j} = \{c_{0}, ..., c_{j-1}\}$, predict a clarification question at the next utterance of the conversation.³

3.2 Method

The proposed run-time architecture is depicted in Figure 2. It contains two indices and two fine-tuned BERT models. The *Documents index* contains the corpus of documents (recall that we deal with conversations that end with a document(s) being retrieved). This index supports passage retrieval. The *Clarification-questions index* contains the pool of clarification questions. The two BERT models are used for re-ranking of candidate clarification questions as described below.

Given a conversation context C^j , we first retrieve top-k passages from the Document index (See Section 3.3 below). We then use those passages, to retrieve candidate clarification questions from the Clarification-questions index (See Section 3.4 below). We thus have, for each passage, a list of candidate clarification questions.

The next step re-ranks those candidate clarification questions. Re-ranking is done by the fusion of ranking obtain through two BERT models. Each model re-ranks the clarification questions by their relevance to the given conversation context and the retrieved passages (see Section 3.5 below). The components of the architecture are described next in more details.

3.3 Conversation-based passage retrieval

Documents in the document index are represented using two fields. The first field contains the actual document content. The second field augments the

¹http://convai.io

²An agent can be either a human agent or a bot.

³We always return clarification questions. We leave it for future work to decide whether a clarification is required.



Figure 2: Clarification-questions selection run-time architecture

document's representation with the text of all dialogs that link to it in the train-set (Amitay et al., 2005). We refer to these two fields as text and anchor respectively. We also keep a third field anchor_and_text that contain the concatenation of the above two fields.

Given a conversation context C^{j} , Passage retrieval is performed in two steps. First, top-k documents are retrieved from the anchor_and_text field. using a disjunctive query over all words in the conversation C^{j} . Following (Ganhotra et al., 2020), we treat the dialog query as a verbose query and apply the Fixed-Point (FP) method (Paik and Oard, 2014) for weighting its words. Yet, compared to "traditional" verbose queries, dialogs are further segmented into distinct utterances. Using this observation, we implement an utterance-biased extension for enhanced word-weighting. To this end, we first score the various utterances based on the initial FP weights of words they contain. We then propagate utterance scores back to their associated words.

In the second step, candidate passages are extracted from those top-k documents using a sliding window of fixed size with some overlap. Each candidate passage p is assigned an initial score based on the coverage of terms in C^j by p. The coverage is defined as the sum over all terms in each utterance, using terms' global *idf* (inverse document frequency) and their (scaled) tf (term frequency). The final passage score is a linear combinations of its initial score and the score of the document it is extracted from. Details are given in appendix A.1

3.4 Clarification-questions retrieval

The pool of clarification questions is indexed into a Clarification index. We use the passages returned for a given conversation context C^j , to extract an initial set of candidate clarification questions as follows. For each passage P, we concatenate its content to the text of all utterances in C^j , and use it as a query to the Clarification index.

We thus have, for each passage, a list of candidate clarification questions.

3.5 Clarification-questions re-ranking

The input to this step is a conversation context C^{j} , a list of candidate passages, and a list of candidate clarification questions for each passage. We use two BERT (Devlin et al., 2019) models to rerank the candidate clarification questions. The first model, *BERT-C-cq* learns an association between conversation contexts and clarification questions. The second model, *BERT-C-P-cq* learns an association between conversation contexts, passages and clarification questions. Training and using the two models is described below.

Fine-tuning of the models. The first model, BERT-C-cq, is fine-tuned through a triplet network (Hoffer and Ailon, 2015) that is adopted for BERT fine-tuning (Mass et al., 2019). It uses triplets (C^j, cq^+, cq^-) , where cq^+ is the clarification question of conversation C at utterance c_j (as given in the conversations of the training set). Negative examples (cq^-) are randomly selected from the pool of clarification questions (not associated with C).

For fine-tuning the second model, BERT-C-P-cq, we need to retrieve relevant passages. We use a weak-supervision assumption that all passages in a relevant document (i.e., a document returned for C), are relevant as well. A triplet for the second BERT model is thus $(C^j [SEP] P, cq^+, cq^-)$, where P is a passage retrieved for C^j , [SEP] is BERT's separator token, cq^+ and cq^- are positive and negative clarification questions selected as described above for the first model.

Due to the BERT limitation on max number of tokens (512), we represent a conversation context C^{j} using the last m utterances whose total length is less than 512 characters. We also take the passage window size to be 512 characters.⁴

Re-ranking with the models. Each candidate clarification question cq_i is fed to the first model with the conversation context as (C^j, cq_i) , and to the second model as $(C^j [SEP] P, cq_i)$, where P is the passage that was used to retrieve cq_i . Final scores of the candidates is set by simple Comb-SUM (Wu, 2012) fusion of their scores from the two BERT models.

4 Experiments

4.1 Datasets

We evaluated our method on two datasets. The first, **ClariQ** (Aliannejadi et al., 2020) represents an information-seeking use-case. The second, **Support** contains conversations and technical documents of an internal customer support site. Statistics on the two datasets are given in Table 1.

The **ClariQ** dataset was built by crowd sourcing for the task of clarification-questions selection, thus it has high quality clarification questions. Each conversation has exactly three turns. Initial user query, an agent clarification question and the user response to the clarification question. The agent utterance is always a clarification question.

The **Support** dataset contains noisy logs of human-to-human conversations, that contain a lot of chit-chat utterances such as *Thanks for your help* or *Are you still there?* We thus applied the following rules to identify agent clarification questions. i) We consider only sentences in agent utterances that contain a question mark. ii) We look for question words in the text (e.g., *what, how, where, did, etc.*) and consider only the text between such a word and the question mark. iii) If no question words were found, we run the sentences with the question mark through Allennlp's constituency parser (Joshi et al., 2018), and keep sentences with a Penn-Treebank clause type of SQ or $SBARQ^5$.

The above rules can be used to detect questiontype sentences. However, we are interested in clarification questions that are related to the background collection of documents and not in chit-chat questions (such as e.g., *how are you today?*). To filter out such chit-chat question types, we apply a 4th rule as follows. iv) Recall that each conversation ends with a document answer. We send the detected question and its answer (the next user's utterance), as a passage retrieval query (see Section 3.1 above) to the Documents index and keep only those questions that returned in their top-3 results, a passage from the document of the conversation.

Table 1: Datasets statistics

	ClariQ	Support
#docs	2.7M	520
#conversations (train/dev/test)	187/50/60	500/39/43
#total clarifications	3940	704
#avg/max turns per C	3/3	8.2/80.5
#avg/max clarifications per C	14/18	1.27/5

4.2 Setup of the experiments

We use Apache Lucene⁶ for indexing the documents. We use English language analyzer and default BM25 similarity (Robertson and Zaragoza, 2009).

For the customer support dataset (**Support**) we used the anchor_and_text field for initial document retrieval, since most documents in the dataset do have training conversations.

The open-domain dataset (**ClariQ**) contains a large number of documents (2.7M), but only a small portion of them do have training conversations. Using the anchor_and_text field for retrieval will prefer that small subset of documents (since only they have anchor text). Thus for this dataset, we used the text field for retrieval.

For passage retrieval, we used a sliding window of 512 characters on retrieved documents' content. We used common values for the hyper parameters, with $\lambda = 0.5$ to combine document and passage scores, and $\mu = 2000$ for the dirichlet smoothing of the documents LM used in the FixedPoint reranking. Details of the passage retrieval are given in Apendix A.1.

The full conversations were used to retrieve passages. For feeding to the BERT models, we concatenated the last m utterances whose total length

⁴note that BERT uses tokens while for the passages and representation of conversation we use characters

⁵https://gist.github.com/nlothian/9240750 ⁶https://lucene.apache.org/

was less than 512 characters (we take full utterances that fit the above size. We do not cut utterances).

We used the pytorch huggingface implementation of BERT⁷. For the two BERT models we used bert-base-uncased (12-layers, 768-hidden, 12-heads, 110M parameters). Fine-tuning was done with the following default hyper parameters. max_seq_len of 256 tokens⁸ for the BERT-C-cq model and 384 for the BERT-C-P-cq model, learning rate of 2e-5 and 3 training epochs.

We retrieved at most 1000 initial candidate clarifications for each passage. All experiments were run on a 32GB V100 GPUs. The re-ranking times of 1000 clarification questions for each conversation took about 1 - 2 sec. For evaluation metrics we followed the ClariQ leaderboard ⁹ and used the Recall@30 as the main metrics.

4.3 Results

Table 2 reports the results on the dev sets of the two datasets.¹⁰ On both datasets, each of the BERT re-rankers showed a significant improvement over the initial retrieval from the Clarification-questions index (denoted by **IR-Base**). For example on **Support**, **BERT-C-cq** achieved R@30=0.538 compared to R@30=0.294 of **IR-Base** (an improvement of 82%).

We can further see that the two BERT models (**BERT-C-cq** and **BERT-C-P-cq**), yield quite similar results on both datasets, but, when fusing their scores (**BERT-fusion**), there is another improvement of about 2.5% over each of the rankers separately. For example on **ClariQ**, **BERT-fusion** achieved R@30=0.791, compared to R@30=0.77 of **BERT-C-cq**.

This improvement can be attributed to complementary matching that each of the two BERT models learns. The second model learns latent features that are revealed only through the retrieved passages, while the first model works better for cases where the retrieved passages are noisy. For example for query 133 in **Clariq**, *all men are created equal* (see Figure 1 above), **BERT-C-P-cq** could find nine correct clarification questions out of 14

⁷https://bit.ly/2Me0Gk1

in its top-30 (including those two in the Figure), while **BERT-C-cq** found only three of them.

Table 3 shows the official Clariq leaderboard result on the test set. We can see that our method **BERT-fusion**¹¹ was ranked forth but was the second best as a team. We note that the top performing system (NTES_ALONG) gave preferences to clarification questions from the test data, capitalizing the specific **Clariq** properties that test topics came from different domain than the train topics. This is not a valid assumption in general. In contrast, we treat all clarification questions equally in the given pool of clarification questions.

 Table 2: Retrieval quality on the dev set of the two datasets

ClariQ	R@5	R@10	R@20	R@30
IR-Base	.327	.575	.669	.706
BERT-C-cq	.352	.631	.743	.770
BERT-C-P-cq	.344	.615	.750	.774
BERT-fusion	.353	.639	.758	.791
Support				
IR-Base	.102	.153	.269	.294
BERT-C-cq	.358	.410	.487	.538
BERT-C-P-cq	.217	.294	.487	.538
BERT-fusion	.294	.410	.500	.551

Table 3: Retrieval quality on the test set of the ClariQ dataset

ClariQ	R@5	R@10	R@20	R@30
NTES_ALONG	.340	.632	.833	.874
NTES_ALONG	.341	.635	.831	.872
NTES_ALONG	.338	.624	.817	.868
BERT-fusion	.338	.631	.807	.857
TAL-ML	.339	.625	.817	.856
Karl	.335	.623	.799	.849
Soda	.327	.606	.801	.843

5 Conclusions

We presented a method for clarification-questions selection in conversational-search scenarios that end with documents as answers.

We showed that using passages, combined with deep-learning models, improves the quality of the selected clarification questions. We evaluated our method on two diversified dataset. On both datasets, the usage of passages for clarification-questions re-ranking achieved improvement of 12% - 87% over base IR retrieval.

⁸note that here we use tokens while for the passages and representation of conversation we use characters

⁹https://convai.io

¹⁰We compare our methods on the dev sets since in **Clariq** we had access only to the dev set. We note that in both datasets, the dev sets wer not used during the training, thus they can be regarded as an held-out test set

¹¹Our run was labeled CogIR in the official leaderboard

References

- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail Burtsev. 2020. Convai3: Generating clarifying questions for opendomain dialogue systems (clariq).
- Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. 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'19, page 475–484, New York, NY, USA. Association for Computing Machinery.
- Einat Amitay, Adam Darlow, David Konopnicki, and Uri Weiss. 2005. Queries as anchors: selection by association. In *Proceedings of the 16th ACM Conference on Hypertext and Hypermedia*, pages 193–201.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jatin Ganhotra, Haggai Roitman, Doron Cohen, Nathaniel Mills, R. Chulaka Gunasekara, Yosi Mass, Sachindra Joshi, Luis A. Lastras, and David Konopnicki. 2020. Conversational document prediction to assist customer care agents. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020, pages 349–356. Association for Computational Linguistics.
- Elad Hoffer and Nir Ailon. 2015. Deep metric learning using triplet network. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings.
- Vidur Joshi, Matthew Peters, and Mark Hopkins. 2018. Extending a parser to distant domains using a few dozen partially annotated examples.
- Antonios Minas Krasakis, Mohammad Aliannejadi, Nikos Voskarides, and Evangelos Kanoulas. 2020. Analysing the effect of clarifying questions on document ranking in conversational search. *Proceedings* of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Yosi Mass, Haggai Roitman, Shai Erera, Or Rivlin, Bar Weiner, and David Konopnicki. 2019. A study of bert

for non-factoid question-answering under passage length constraints. *CoRR*, abs/1908.06780.

- Wenjie Ou and Yue Lin. 2020. A clarifying question selection system from ntes_along in convai3 challenge. *CoRR*, abs/2010.14202.
- Jiaul H. Paik and Douglas W. Oard. 2014. A fixed-point method for weighting terms in verbose informational queries. In Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14, page 131–140, New York, NY, USA. Association for Computing Machinery.
- Sudha Rao and Hal Daumé III. 2018. 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)*, pages 2737–2746, Melbourne, Australia. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. The probabilistic relevance framework: Bm25 and beyond. *Found. Trends Inf. Retr.*, 3(4):333—389.
- Shengli Wu. 2012. Data Fusion in Information Retrieval, volume 13.
- Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating clarifying questions for information retrieval. In *Proceedings of The Web Conference 2020*, WWW '20, page 418–428, New York, NY, USA. Association for Computing Machinery.

A Appendix

A.1 Passage Retrieval details

We use Apache Lucene for indexing the documents, configured with English language analyzer and default BM25 similarity (Robertson and Zaragoza, 2009).

After retrieving top-k documents, candidate passages are extracted from those documents using a sliding window of fixed size with some overlap. Each retrieved passage p is assigned an initial score based on the coverage of terms in C^j by p. The coverage is defined as the sum over all terms in each utterance, using terms' global *idf* (inverse document frequency) and their (scaled) tf (term frequency). Let c be a conversation with n utterances $c = u_1, ...u_n$. Passage score is computed as a linear combination of its initial score $score_{init}(p, c)$ and the score of its enclosing document. Both scores are normalized.

$$score(p,c) = \lambda * score(d) + (1-\lambda) * score_{init}(p,c)$$
(1)

We used lambda=0.5, i.e., fixed equal weights for the document and the passage scores.

The initial passage score $score_{init}(p, c)$ is computed as a weighted sum over its utterances scores $score_{ut}(p, u_i)$. Utterance scores are discounted such that later utterances have greater effect on the passage score.

$$score_{init}(p,c) = \sum_{i=1}^{n} weight_{ut}(i) * score_{ut}(p,u_i)$$
$$weight_{ut}(i) = discount_factor^{(n-i)}$$
$$discount_factor = 0.85$$
(2)

Utterance score $score_{ut}(p, u)$ reflects utterance's terms coverage by the passage, considering terms' global *idf* (inverse document frequency) and their (scaled) *tf* (term frequency). Multiple coverage scorers are applied, which differ by their term frequency scaling schemes. Finally, the utterance score is a product of these coverage scores $score_{cov}(p, u)$.

 $score_{ut}(p, u) = \prod_{j=1}^{m} score_{cov_j}(p, u)$ $m = 2 \quad (\text{two scaling schemes are employed})$ $score_{cov_j}(p, u) = \sum_{t \in t^{pu}} idf(t) * scale_j(t, p)$ $t^{pu} = t^u \bigcap t^p \quad (\text{terms appearing in both})$ $t^p, t^u = \quad (\text{passage terms, utterance terms})$ (3)

Different scaling schemes provide different interpretations of terms' importance. We combine two tf scaling methods, one that scales by a BM25 term score, and another that scales by the minimum of tf(t) in the utterance and passage.

$$scale_1 = BM25(t, p)$$

$$scale_2 = min(tf(t, p), tf(t, c))$$
(4)

The final passage score is a linear combinations of its initial score and the score of the document it is extracted from. Candidate passage ranking exploits a cascade of scorers.