

Building and Evaluating Open-Domain Dialogue Corpora with Clarifying Questions

Mohammad Aliannejadi¹, Julia Kiseleva², Aleksandr Chuklin³,
Jeffrey Dalton⁴, Mikhail Burtsev⁵

¹University of Amsterdam, ²Microsoft Research, ³Google Research

⁴University of Glasgow, ⁵MIPT & AIRI

m.aliannejadi@uva.nl, julia.kiseleva@microsoft.com
chuklin@google.com, jeff.dalton@glasgow.ac.uk, burtcev.ms@mipt.ru

Abstract

Enabling open-domain dialogue systems to ask clarifying questions when appropriate is an important direction for improving the quality of the system response. Namely, for cases when a user request is not specific enough for a conversation system to provide an answer right away, it is desirable to ask a clarifying question to increase the chances of retrieving a satisfying answer. To address the problem of ‘asking clarifying questions in open-domain dialogues’: (1) we collect and release a new dataset focused on open-domain single- and multi-turn conversations, (2) we benchmark several state-of-the-art neural baselines, and (3) we propose a pipeline consisting of offline and online steps for evaluating the quality of clarifying questions in various dialogues. These contributions are suitable as a foundation for further research.

1 Introduction

The ultimate goal of a conversational system is to assist users by returning an appropriate answer in response to their requests (Kiseleva et al., 2016a; Li et al., 2021). Recent progress on neural approaches to natural language processing (Devlin et al., 2019; Liu et al., 2019; Clark et al., 2020), and the availability of large amounts of conversational data have triggered a renaissance in end-to-end neural open-domain chatbots (Adiwardana et al., 2020; Roller et al., 2021; Zhang et al., 2020; Burtsev et al., 2017; Dalton et al., 2020). There has been great progress on suggesting measures to evaluate what makes a conversation satisfying for users using various human evaluation techniques (Li et al., 2019a; See et al., 2019b). Those efforts showed that suggested large pre-trained models do not always perform seamlessly (See et al., 2019a), and there are still several challenges needed to be solved for open-domain conversational systems (Huang et al., 2020).

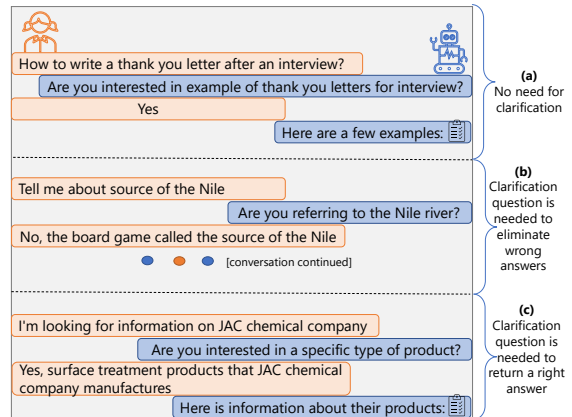


Figure 1: Examples of clarification questions embedded into the open domain conversations: (a) represents a clear user request, so clarification was unnecessary; (b) and (c) demonstrate a situation when the request is ambiguous and a system needs to act.

Nass and Moon (2000) conclude that people have similar expectations from talking to bots and humans. This similarity is a possible explanation for why sometimes user requests might be ambiguous and incomplete, as shown in Fig. 1 (b) and (c). This ambiguity is especially challenging to handle in a dialogue setting, where a system is limited by returning only one answer in response to each request, unlike in web search setup where diversification of results is possible and acceptable (Vallet and Castells, 2012). Previous research has shown that users are much more forgiving about system mistakes if they can act on them with minimal efforts spent (Kocielnik et al., 2019; Kiseleva et al., 2016b). Therefore it is more appropriate to ask a clarifying question in user request ambiguity rather than generating incorrect answers. There are separate attempts to explore the following related tasks: (1) identifying a moment when the question should be asked in the course of conversation (Hancock et al., 2019); and (2) retrieving a clarification question (Rao and III, 2018; Wang et al., 2018). In this paper, we aim to combine these related aspects and

study the following problem of *generating clarifying questions for open-domain conversations*: the system must identify wherever the question is ambiguous, and, if so then instead of trying to answer it directly, it should ask a good clarifying question (Fig. 1). One possible stumbling block preventing the community from studying the problem of open-domain clarifying question generation to enhance user experience while interacting with a conversational bot (Huang et al., 2020) is the lack of suitable datasets, which we address in this work. To summarise, the main contributions of this work are:

- C1** releasing a dataset dedicated to the problem of asking a clarifying question in open-domain dialogue systems. The dataset includes single (~15K) and multi-turn (~1.5M) conversations and covers ~300 various topics and it is suited to study: (1) *when* a clarifying question should be asked given the current context of the conversation; and (2) *which* question should be asked;
- C2** benchmarking several state-of-the-art (SoTA) neural models; and
- C3** building an evaluation pipeline that provides fast iteration and involves two stages: (1) offline (automatic evaluation); and (2) online (human-in-a-loop to converse with the system).¹

We release the collected dataset, offline evaluation pipeline, and the code for running explored neural SoTA models. These models can be employed as baselines for the task.²

2 Related work

Our work is broadly relevant to two strands of research: learning to ask clarifying questions in open-domain conversational settings (Section 2.1) and evaluating dialogue systems (Section 2.2).

2.1 Learning to ask clarifying questions

Information retrieval community has paid close attention to the problem of ambiguity in user search queries. Previously this problem was addressed through the diversification of search result pages (Radlinski and Dumais, 2006; Kong and Allan, 2016, 2014), including via usage of personal and contextual data (Jiang et al., 2015; Kato

and Tanaka, 2016). Recently, (Rosset et al., 2020; Aliannejadi et al., 2019; Zamani et al., 2020a) suggest techniques to address ambiguity by generating clarifying questions.

Where the general settings are: (1) a user is issuing an *ambiguous* keyword query; (2) a search engine’s goal is to suggest conversational clarifying questions to help to find the required information (Krasakis et al., 2020; Lotze et al., 2021; Sekulic et al., 2021; Aliannejadi et al., 2021). These works also resulted in a number of datasets, e.g. *Qulac* (Aliannejadi et al., 2019) and *MIM-ICS* (Zamani et al., 2020b), which consists of queries, issued by real users, and behavioral signals such as clicks. Braslavski et al. (2017) focus on characteristics, forms, and general patterns of clarifying questions.

Suggesting a clarifying questions is closely related to question answering (Q&A) (Kwiatkowski et al., 2019; Soleimani et al., 2021) and question generation (QG) domains (Gao et al., 2019; Chai and Wan, 2020). Trienes and Balog (2019) made an attempt to understand unclear questions, Li et al. (2017) suggesting an RL-based method for deciding *when* to ask for user feedback in Q&A setup.

Recently, proactive bot behavior has started to attract researchers’ attention in dialogue settings yet remains rather untouched (Huang et al., 2020). Rao and III (2018) designed a model to rank a candidate set of clarification questions by their usefulness to the given post at Stack Exchange, which targeted the problem *which* question to ask. The resulted dataset was released, but it covers specific narrow topics. In contrast, Hancock et al. (2019) focused on *when* to ask a question in order to self-retrain a bot, which has been resulted in releasing a dataset. Wang et al. (2018) studied QG techniques in application to open-domain conversations.

2.2 Evaluating Dialogue Systems

Dialogue systems are generally separated into two types: task-oriented and open-domain. The task-oriented ones usually have clear criteria for evaluation, e.g. turn correction ratio, inappropriate utterance ratio, proxies for accuracy, and success rate (Takanobu et al., 2019; Li et al., 2016; Su et al., 2018; Li et al., 2020). Despite significant efforts to introduce automatic metrics to evaluate open-domain conversations (Reiter, 2018; Novikova et al., 2017; Lowe et al., 2017), it remains area for exploration (Li et al., 2019a,b, 2021). To the best

¹The pipeline was designed as part of the ConvAI3 (Aliannejadi et al., 2020) data challenge (<https://convai.io>)

²Available at <https://github.com/aliannejadi/ClariQ>

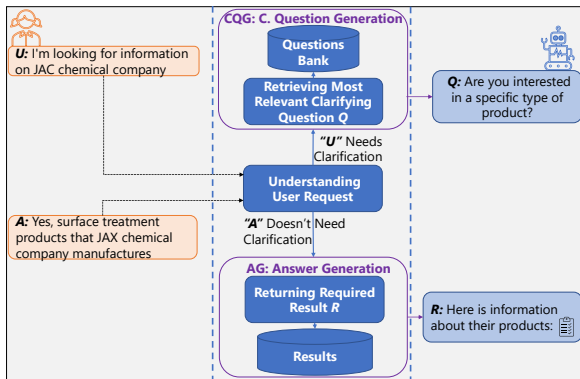


Figure 2: A pipeline for asking clarifying questions for open-domain conversations using example Fig 1 (C).

of our knowledge the current standard approach for evaluating open-domain dialogues requires employing human assessments via crowdsourcing platforms (Zhang et al., 2018; Li et al., 2019a) or engaging volunteers to participate in research competitions (Burtsev et al., 2018; Dinan et al., 2020; Burtsev and Logacheva, 2020; Aliannejadi et al., 2020).

Therefore, we can conclude that understanding and generating open-domain clarification questions is a major component in conversational information-seeking systems, which is still under exploration. Hence, our efforts on collecting datasets and investigating the performance of the neural SoTAs are timely and useful for future research in this area.

3 Problem Setting

Our main goal is to collect a dataset to enable studying clarifying questions generation whenever appropriate, as depicted in examples in Fig. 1. Fig. 2 demonstrates a pipeline that makes it possible to process user requests in the open domain as follows: ‘User Request Understanding’ (URU) decides which module to call either ‘Clarifying Question Generation’ (CQG) or ‘Answer Generation’ (AG). In this work, we focus on the first two. We aim to collect the following data:

- **User Request (U):** an initial user request in the conversational form, e.g., ‘What is Fickle Creek Farm?’ with a label reflecting whether clarification is needed;
- **Set of clarification questions ($\{Q\}$):** a set of possible reasonable clarifying questions that address multiple aspects/facets of U , e.g., Q_1 : ‘Do you want to know the location of fickle creek farm?’, Q_2 : ‘Would you like to know the history




Figure 3: The pipeline highlights steps for two main goals: to collect required datasets and to perform the reproducible evaluation.

of fickle creek farm?’³;

- **User Answers (A):** each question is supplied with a user answer, e.g., the answer to Q_1 is A_1 : ‘No, I want to find out where can I purchase fickle creek farm products’, the answer to Q_2 is A_2 : ‘I just need general information about fickle creek’

The collected dataset of pairs $(U, \{Q, A\})$ can be easily transformed to a set of single-turn conversations consisting of the coherent and consistent triples (U, Q, A) as shown in the example in Fig. 2. We ask items in a triple, U , Q and A , satisfy the following requirements:

- R1** user requests (U) must cover various conversational topics to represent open-domain dialogues;
- R2** the final collection of U should contain both types: ambiguous and unambiguous;
- R3** each inquiry U to the system should be in the conversational form;
- R4** the need for clarification should be predetermined as a label for each U in the collection;
- R5** each clarifying question (Q) should be reasonable, coherent with U and address multiple facets of every ambiguous request U ; and
- R6** each user answer A should be consistent with the clarifying question from the system.

After collecting of single-turn conversations, they are used to train various conversational agents. To collect multi-turn conversations, the two best-performing agents are utilized to converse with crowdsourced workers, who evaluate a system quality and reply to suggested clarifying questions. Finally, the two agents are evaluated using Acute-eval framework (Li et al., 2019a), which is best available practice for online evaluation of open-domain dialogue systems. Overall, our pipeline for *data collection* and *evaluation* is summarized in Fig. 3.

³Candidate clarifying questions should also address out-of-collection facets.

Table 1: An example of facets for the incomplete query.

Topic	Facet
Neil Young	Find albums by Neil Young to buy
	Find biographical information about Neil Young
	Find lyrics or sheet music for Neil Young’s songs
	Find a list of Neil Young tour dates.

4 Data Collection

Following the suggested pipeline in Fig. 3 this section describes our results regarding the collected datasets to initiate more follow-up studies of the problem of ‘asking clarifying questions in open-domain dialogues’, namely, P_1 : single-turn dialogues (Sec. 4.1) and P_3 multi-turn ones (Sec. 4.2).

4.1 P_1 : Crowdsourcing Single-Turn Dialogues

Collecting Conversational User Requests Our collection of single-turn conversations is built on top of the TREC Web track 2009-2014⁴ data, which was originally designed to evaluate search result diversification. The TREC collection contains 300 search topics. The presence of varied search topics, which express different user information needs, helps us imitate open domain user-system interactions, which are required in **R1**. Each topic in the collection is specific, ambiguous, or faceted (Clarke et al., 2009). In this work, we use the term ‘facet’ to refer to the subtopics of both faceted and ambiguous topics. For clarity, the example of mapping from the search topic to set of facets is provided in Tab. 1. Faceted and ambiguous topics make an ideal case to study the effect of clarifying questions as they can be interpreted in various ways, which is required by **R2**. User information needs are expressed in the form of short search topic description⁵ (Tab. 1) because the TREC Web track collection was designed for web search needs. Therefore, to satisfy requirement **R3** and to make requests lexical diverse, we ask expert annotators to convert those short keyword queries to the proper conversational request. The examples of such conversion are presented in Tab. 2. To satisfy the requirement **R4**, annotators are asked to provide a score for each request to reflect if clarification is needed ranging from 1 to 4, where ‘1’

stands for very low or no need for clarification and ‘4’ indicates a highly ambiguous request (examples are provided in Tab. 2). Two annotators assess clarification need of a query. In case of disagreement, we assign an additional annotator to make the final assessment. We achieve a high inter-annotator agreement on this task (Cohen’s $\kappa = 0.78$).

Collecting Clarifying Questions $\{Q\}$ To collect $\{Q\}$ for every U that satisfies **R5**: (1) we utilized Qulac dataset⁶ by converting topics into conversational requests; (2) we significantly extended it by crowdsourcing more data through Human Intelligence Task (HIT) on Amazon Mechanical Turk⁷, which design follows a general strategy proposed in (Aliannejadi et al., 2019). Namely, we asked the workers to imagine themselves acting as a conversational agent⁸ where an imaginary user had asked them about a topic. Then, we described the concept of facet to them, supporting it with multiple examples. Finally, we ask Turkers to do the following:

- discover the facets of each U using a preferred search engine and scan the results in the first three pages; and
- generate six questions related to U , aiming to address the facets they had figured out.

We assigned two workers per HIT, resulting in 12 questions per U in the first round. To preserve the questions’ language diversity, we limited each worker to a maximum of two HITs. HITs were available to workers residing in the U.S. who had an approval rate of over 97%.

Controlling Quality of Clarifying Questions

To estimate the quality of the collected questions, we aim to address two main concerns: (1) *how good are the collected clarifying questions?*; and (2) *Is the set of clarifying questions diverse (in other words, addressing different facets associated with the topic)?* Given the high complexity of this task, we appointed two expert annotators. They

⁴<https://trec.nist.gov/data/webmain.html>

⁵sometimes also can be referred as ‘keyword query’

⁶<https://github.com/aliannejadi/qulac>

⁷<http://www.mturk.com>

⁸such as Microsoft Cortana, Alexa, or Google Assistant.

Table 2: Examples of keyword queries converted to conversational requests with assigned clarification score (C. score).

Search Topic	Conversational Request	C. Score
average charitable donation	What is average charitable donation?	1
gmat prep classes	How to prepare for the GMAT?	2
von willebrand disease	What is von Willebrand Disease?	3
land surveyor	I'm interested to know about land surveyor	3
alexian brothers hospital	Give me information about Alexian Brothers hospitals	3
worm	I'm looking for information on worm	4

were instructed to read all the collected questions on each topic, marking invalid and duplicate questions. Annotators were asked to match a question to a facet if its answer would address the facet. Finally, to ensure that all facets were covered by at least one question, we asked the annotators to generate an additional question for each facet that needed more specific questions.

Collecting Answers To satisfy **R6**, we designed another HIT to collect coherent and consistent answers to the clarifying questions. The task started with detailed instructions followed by several examples. The workers were given U and a facet description. Then we instruct them to assume that they had submitted the initial user request U with their actual information need being the given facet. Then workers were required to write the answer to the one clarifying question that was presented to them. If a question required information other than what workers were provided with, they were instructed to use a ‘No answer’ tag. Each worker was allowed to complete a maximum of 100 HITs to ensure language diversity. Workers were based in the U.S. with an approval rate of 95% or greater.

Controlling quality of Collected Answers During the course of data collection, we performed regular quality checks on the collected answers. The checks were done manually on 10% of submissions per worker. In case we observed any invalid submissions among the sampled answers of one worker, we then studied all the submissions from that workers. Invalid submissions were then removed from the collection, and the worker was banned. Finally, we assigned all invalid answers to other workers to complete. Moreover, we employed basic behavioral check techniques in the design of the HIT. For example, we disabled copy/paste features of text inputs and tracked workers’ keystrokes. This enabled us to detect and reject low-quality submissions.

As an outcome, we have a high-quality collec-

tion of single-turn conversations in the form of required triples: (U, Q, A) , which is marked as P_1 in our pipeline in Fig. 3. Tab. 3 provides a statistics on collected dataset of single-turn conversations.

4.2 P_3 : Crowdsourcing Multi-Turn Dialogues

The collected dataset of single-turn is sufficient to train and evaluate several conversational agents. More technical details on training and evaluation are provided in Sec. 5. For now, we assume that as a result of P_2 : DA is one of the best-performing trained dialogue agents. We assume that the trained DA can have a conversation with users. Namely, it should either ask a clarification question or give a factual answer to the user’s request at each dialog step. Therefore, the trained DA is capable of:

- providing clarification question whenever appropriate in the course of the conversation;
- interpreting user’s answer to the clarifying question.

To collect multi-turn conversations, we utilize best-performing dialogue agents that can accommodate an arbitrary number of turns, having two goals in mind:

G1 evaluating the quality of the agents in multi-turn settings where they converse with real humans; and

G2 collecting a new dataset of multi-turn conversations with respect to clarification questions.

To reach **G1**, the idea is to run the agents multiple times with different turns and evaluate them accordingly. For that purpose, we design a HIT similar to the ones described in Sec. 4.1 with the difference in the context of a conversation. We instructed crowd workers to understand the user’s actual information need and imagine they are looking for the same information. Then, follow a conversation on that as presented in Fig. 4. The workers were instructed to answer the last question in the conversation while considering the conversation’s




Figure 4: An example of the task provided at the HIT for multi-turn conversations, which asks to submit the answer given dialogue history/context.

context, which consists of previous questions and answers. The context of conversation could include 1-2 rounds of question-answer interactions. To check the quality of clarifying questions returned by the trained dialogues, we instructed workers to indicate if the question was not understandable or was not in a proper language. As a result, such questions were removed from the collection. We use the same quality check procedure for the collected answers as described in the previous section. Tab. 3 provides a statistics on collected multi-turn conversations, to achieve **G2**.

Synthetic Multi-Turn Conversations We also generate synthetic multi-turn conversations for training purposes. To do so, for each topic, we create a set of all possible combinations of questions (2 or 3 questions) together with their corresponding answers.

5 Models and Evaluation

Following the suggested pipeline in Fig. 3, we explain our contributions regarding the *evaluation* of ‘asking clarifying questions in open-domain dialogues’ problem, namely:

- P_2 : how the dialogue agents are trained and automatically evaluated based on single-turn conversations (Sec. 5.1);
- P_4 : how evaluation of multi-turn conversations

Table 3: Statistics over collected data.

# search topics	298
# faceted & ambiguous topics	250
# single topics	48
# facets	1070
# questions	3,304
Average terms per question	9.74 ± 2.62
Average terms per answer	8.48 ± 4.40
# synthetic conversations (all)	1,596,757
# synthetic conversations (2 turns)	203,050
# synthetic conversations (3 turns)	1,393,707
# human-machine dialogues (1 turn)	15,226
# human-machine dialogues (3 turns)	499

is performed from both perspectives: offline automatic manner and having human-in-the-loop using Acute-eval (Li et al., 2019a) (Sec. 5.2).

We design our experiments to collect answers to the following research questions:

RQ1 When to ask clarifying questions during open-domain dialogues? (Sec. 5.1.1)

RQ2 Which clarifying question to ask for a given context of a conversation? (**a.** the single-turn conversations case is described in Sec. 5.1.2; **b.** multi-turn one – Sec. 5.2)

5.1 P_2 : Evaluating Single-Turn Agents

The collected dataset, described in Sec. 4.1, is split into training (70%), validation (dev) (10%), and test (20%) sets. We split the data based on the search topic and maintained the same split for all single-turn and multi-turn experiments. During the evaluation procedure, the following is used: (1) a set of conversational user requests, and (2) a set of questions (i.e., question bank), which contains all collected questions on all the topics.

5.1.1 Predicting Clarification Need

Task The task is, given a user request, return a score from 1 (no need for clarifying questions) to 4 (cannot provide any answers without user clarification) indicating the necessity of asking clarifying questions (as depicted in module ‘Understanding User Request’ in Fig. 2).

Automatic Evaluation To evaluate the performance of the suggested classifier, we use Precision, Recall, F1-Measure, and Mean Squared Error (MSE). Tab. 4 presents the collected results of various classification methods, which includes Roberta-based classifier (Liu et al., 2019), BART (Chipman et al., 2010), and BERT-based classifier (Devlin et al., 2019). Based on the supplied results, we can

Table 4: Performance for predictors returning classification need score based on *dev* and *test* sets. The best-performing model is marked in bold.

Model		Precision	Recall	F1-Measure	MSE
RoBERTa-based	dev	0.6039	0.5600	0.5551	0.6200
	test	0.5981	0.6557	0.6070	0.5409
BART	dev	0.7008	0.7000	0.6976	0.5200
	test	0.4813	0.4754	0.4756	0.7705
BERT-based	dev	0.5218	0.4800	0.5000	0.8200
	test	0.3931	0.4918	0.4253	0.6557

answer **RQ1**: the task is rather difficult and potentially can benefit from more exploration despite the reasonable performance of the proposed baselines.

5.1.2 Returning Clarifying Question

Task The task is, given a user request which needs clarification, return the most suitable clarifying question from the supplied question bank (as shown in module CQG in Fig. 2).

Automatic Evaluation We introduce two main strategies for evaluation: (1) document relevance and (2) question relevance.

Document Relevance To estimate the relevance of the retrieved documents we use the following standard metrics: Mean Reciprocal Rank (MRR) (Voorhees, 1999; Radev et al., 2002), Precision (P)@[1,3,5,10,20], Normalized Discounted Cumulative Gain (nDCG)@[1,3,5,20] (Wang et al., 2013). These metrics are computed as follows: a selected clarifying question, together with its corresponding answer, is added to the original user request. The updated query is then used to retrieve (or re-rank) documents from the collection. The quality of the question is then evaluated by measuring how much the question and its answer affect document retrieval performance when added to the initial request. We evaluate document relevance based on the relevance assessments provided by the TREC Web Track.

Question Relevance Models are also evaluated in how well they can rank relevant questions higher than other questions in the question bank. For this task, which we call ‘question relevance,’ the models are evaluated in terms of Recall@[10,20,30]. Since the precision of models is evaluated in the document relevance task, here we focus only on recall.

The suggested evaluation metrics are collected for a number of baselines (**B**) and fine-tuned state-of-the-art NLP models (**M**):

Table 5: A set of document relevance related metrics reported on *dev* and *test* sets. NDCG@3 (in bold) reported on the test set is used as the main metric to rank the quality of the models.

Model		MRR	P@1	NDCG@3	NDCG@5
B ₁ : Worst Q.	dev	0.0841	0.0125	0.0252	0.0313
	test	0.0541	0.0000	0.0097	0.0154
B ₂ : No Q.	dev	0.3000	0.2063	0.1475	0.1530
	test	0.3223	0.2268	0.1134	0.1059
B ₃ : Best Q.	dev	0.4882	0.4187	0.3337	0.3064
	test	0.4881	0.4275	0.2107	0.1759
M ₁ : Roberta	dev	0.3640	0.2813	0.2002	0.1954
	test	0.3190	0.2342	0.1265	0.1130
M ₂ : ELECTRA	dev	0.3761	0.3000	0.2113	0.1955
	test	0.3140	0.2379	0.1229	0.1097
M ₃ : BERT	dev	0.3596	0.2750	0.1879	0.1882
	test	0.3044	0.2119	0.1131	0.1021
M ₄ : BM25	dev	0.3096	0.2313	0.1608	0.1530
	test	0.3134	0.2193	0.1151	0.1061
M ₅ : BERT +BM25	dev	0.3180	0.2437	0.1625	0.1550
	test	0.3216	0.2453	0.1196	0.1097
M ₆ : Roberta +BM25	dev	0.3606	0.2813	0.1942	0.1891
	test	0.3045	0.2156	0.1108	0.1025

- **B₁: Worst Question**, when the dialogue system returns the least relevant question;
- **B₂: No Question**, when the system never returns any clarifying question;
- **B₃: Best Question**, which show oracle performance as it always returns the most relevant question from the bank.
- **M₁**: fine-tuned version of the pre-trained Roberta (Liu et al., 2019);
- **M₂**: fine-tuned sequence classification model based on pre-trained ELECTRA (Clark et al., 2020; Ou and Lin, 2020);
- **M₃**: fine-tuned version of pre-trained BERT (Devlin et al., 2019)
- **M₄**: BM25 (Robertson et al., 1995); and
- ‘+BM25’ for **M₅** and **M₆**: means that BM25 on top on neural baseline for the final re-ranking of the questions from the bank.

Based on results of automatic evaluations of all the

Table 6: A set of question relevance related metrics reported on *dev* and *test* sets. Recall@30 (in bold) reported on the test set is used as the main metric to rank the quality of the models.

Model		R@5	R@10	R@20	R@30
M ₁ : Roberta	dev	0.3649	0.6694	0.8265	0.8587
	test	0.3395	0.6251	0.8176	0.8568
M ₂ : ELECTRA	dev	0.3604	0.6749	0.8478	0.8761
	test	0.3404	0.6329	0.8335	0.8744
M ₃ : BERT	dev	0.3492	0.6196	0.7337	0.7632
	test	0.3438	0.6228	0.7987	0.8409
M ₄ : BM25	dev	0.3245	0.5638	0.6675	0.6913
	test	0.3170	0.5705	0.7292	0.7682
M ₅ : BERT+BM25	dev	0.3454	0.6166	0.7354	0.7621
	test	0.3272	0.6061	0.8013	0.8433
M ₆ : Roberta+BM25	dev	0.3637	0.6409	0.7484	0.7793
	test	0.3361	0.6219	0.7960	0.8360

methods suggested above are reported in Tab. 5 and Tab 6 for single-turn conversations, to answer **RQ2.a**, we can conclude the performance of the best-performing fine-tuned neural SoTAs is reasonable and we can use them for multi-turn conversation.

5.2 P₄: Evaluating Multi-Turn Conversations

Task The task is, given an ongoing conversation with multiple turns, select or generate the next question that would clarify the user’s intent best. The main goal is to learn from previous user feedback and ask a question that would lead to the highest information gain.

Automatic Evaluation Similar to the single-turn task, we evaluate the effectiveness of the baseline models based on document relevance. Therefore, we utilize the whole conversation context, clarifying questions, and human responses to retrieve documents from the collection and assess the quality of a question based on its impact on ranking performance. Note that we do not evaluate multi-turn models in terms of question relevance, since the question relevance is intended to evaluate recall of questions related to the search topic. Due to complexity and costs of the evaluation, we pick two best-performing models from Sec. 5.1.2 for this task. To do so, we use the synthetic training data to fine-tune ELECTRA and Roberta similarly to our single-turn setup, but in the multi-turn case the whole history context is considered as a user request. The module for deciding whenever the request needs clarification is preserved. We see in Tab. 7 that ELECTRA outperforms Roberta in

Table 7: A set of document relevance related metrics reported on test set for multi-turn conversations. NDCG@3 (in bold) is used as the main metric to rank the quality of the models.

Model	MRR	P@1	NDCG@3	NDCG@5
ELECTRA	0.1798	0.1161	0.0553	0.0536
Roberta	0.1669	0.1067	0.0522	0.0494

Table 8: Pairwise comparison of ELECTRA and Roberta on the multi-turn conversations. The values report the percentage of judgements in which ELECTRA wins Roberta in terms of HU, EG, IN, KL, and CL.⁹

Comparison	HU	EG	IN	KL	CL
ELECTRA vs. Roberta	0.57	0.59	0.56	0.57	0.56

terms of all evaluation metrics by a margin. One promising future research line might be exploring what properties of these two models lead to that difference in their effectiveness for this task.

Human Evaluation To ensure that our automatic evaluation reflects the dialogues’ quality, we conduct a pairwise human evaluation on two of the baselines. We use and extend the Acute-eval human annotation framework (Li et al., 2019a) to evaluate 120 randomly sampled dialogue pairs. For consistency, we use the four questions that the authors suggest measuring Humanness (HU), Engagingness (EG), Interestingness (IN), Knowledgeable (KL), and add a fifth one, specific to our task, on Clarification (CL). We modify the crowdsourcing task to inform the annotators about the conversation’s main goal (i.e., information seeking). Furthermore, it is crucial to ensure that the annotators consider the model’s ability to understand the user’s feedback and incorporate the additional knowledge when asking its next question. Therefore, we added another question to examine this aspect of the conversation. As shown in Fig. 5, after showing two full conversations to the annotators, they evaluated the model’s ability of clarification by answering the following question: ‘Which one asks better (or more reasonable) clarifying questions?’. Tab. 8 reports the results of our human evaluation of 120 dialogue pairs in terms of percentage of cases that ELECTRA beats Roberta based on the human annotation. We see that ELECTRA is judged to be the best model in most cases for all five aspects. It is interesting to see that the human annotation is in line with the proposed automatic annotation, suggesting that our approach approximates the true quality of the models. Moreover, we see that our new evaluation dimension, Clarification, achieves




Figure 5: Sample on-boarding dialogue for clarification annotation using Acute-eval framework.

a similar result to other dimensions, which suggests that it should be included in the evaluation framework for the open-domain dialogues.

Based on reported results for automatic evaluation (Tab. 7), which is aligned with human-in-a-loop one (Tab. 8), we can conclude that suggested methods can be solid baselines for the follow-up research as an answer **RQ2.b**.

6 Conclusions

Asking clarifying questions when a user request is ambiguous is essential for developing human-like open-domain dialogue systems. In this work, we introduce a large-scale dataset that covers almost 300 different topics and is suitable as a foundation for further research. The collected dataset includes both single- and multi-turn conversations. We benchmark several state-of-the-art neural models that were fine-tuned for asking clarifying questions. Based on how these models performed, we conclude that they are solid baselines for future research that more fully explores the problem space. In this paper, we also suggest an offline automatic evaluation pipeline, which agrees with human-in-loop evaluation.

We publicly release the collected datasets, the code for training baselines, and our evaluation procedure in order to push forward the state-of-the-art.

Acknowledgement

This work was supported in part by the NWO Innovational Research Incentives Scheme Vidi (016.Vidi.189.039) and an Engineering and Physical Sciences Research Council grant EP/V025708/1. We thank Microsoft Research, Google, and Amazon Science for their support to collect the data and organize the ConvAI3 data challenge. We also thank all the participants of the data challenge for taking part in the competition and releasing their models. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

References

- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Mohammad Aliannejadi, Leif Azzopardi, Hamed Zamani, Paul Thomas, and Nick Craswell. 2021. Analysing mixed initiatives and search strategies during conversational search. In *International Conference on Information and Knowledge Management (CIKM)*.
- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail S. Burtsev. 2020. ConvAI3: Generating clarifying questions for open-domain dialogue systems (ClariQ). *arXiv preprint arXiv:2009.11352*.
- Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking clarifying questions in open-domain information-seeking conversations. In *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 475–484.
- Pavel Braslavski, Denis Savenkov, Eugene Agichtein, and Alina Dubatovka. 2017. What do you mean exactly? analyzing clarification questions in CQA. In *Conference on Conference Human Information Interaction and Retrieval (CHIIR)*, pages 345–348.
- Mikhail Burtsev, Aleksandr Chuklin, Julia Kiseleva, and Alexey Borisov. 2017. Search-oriented conversational AI (SCAI). In *ACM SIGIR International Conference on Theory of Information Retrieval (ICTIR)*, pages 333–334.
- Mikhail Burtsev and Varvara Logacheva. 2020. Conversational intelligence challenge: Accelerating research with crowd science and open source. *AI Magazine*, 41(3):18–27.

- Mikhail Burtsev, Varvara Logacheva, Valentin Malykh, Iulian Vlad Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W Black, Alexander Rudnicky, and Yoshua Bengio. 2018. The first conversational intelligence challenge. In *The NIPS'17 Competition: Building Intelligent Systems*, pages 25–46.
- Zi Chai and Xiaojun Wan. 2020. Learning to ask more: Semi-autoregressive sequential question generation under dual-graph interaction. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 225–237.
- Hugh A Chipman, Edward I George, Robert E McCulloch, et al. 2010. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations (ICLR)*.
- Charles L. A. Clarke, Nick Craswell, and Ian Soboroff. 2009. Overview of the TREC 2009 web track. In *Text REtrieval Conference (TREC)*.
- Jeff Dalton, Aleksandr Chuklin, Julia Kiseleva, and Mikhail Burtsev. 2020. Proceedings of the 5th international workshop on search-oriented conversational AI (SCAI). In *International Workshop on Search-Oriented Conversational AI (SCAI)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2020. The second conversational intelligence challenge (ConvAI2). In *The NeurIPS'18 Competition*, pages 187–208.
- Yifan Gao, Piji Li, Irwin King, and Michael R. Lyu. 2019. Interconnected question generation with coreference alignment and conversation flow modeling. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4853–4862.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazaré, and Jason Weston. 2019. Learning from dialogue after deployment: Feed yourself, chatbot! In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 3667–3684.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems (TOIS)*, 38(3):1–32.
- Di Jiang, Kenneth Wai-Ting Leung, Lingxiao Yang, and Wilfred Ng. 2015. Query suggestion with diversification and personalization. *Knowledge-Based Systems*, 89:553–568.
- Makoto P Kato and Katsumi Tanaka. 2016. To suggest, or not to suggest for queries with diverse intents: Optimizing search result presentation. In *ACM International Conference on Web Search and Data Mining (WSDM)*, pages 133–142.
- Julia Kiseleva, Kyle Williams, Ahmed Hassan Awadallah, Aidan C Crook, Imed Zitouni, and Tasos Anastasakos. 2016a. Predicting user satisfaction with intelligent assistants. In *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 45–54.
- Julia Kiseleva, Kyle Williams, Jiepu Jiang, Ahmed Hassan Awadallah, Aidan C Crook, Imed Zitouni, and Tasos Anastasakos. 2016b. Understanding user satisfaction with intelligent assistants. In *Conference on Human Information Interaction and Retrieval (CHIIR)*, pages 121–130.
- Rafal Kocielnik, Saleema Amershi, and Paul N Bennett. 2019. Will you accept an imperfect AI? exploring designs for adjusting end-user expectations of AI systems. In *CHI Conference on Human Factors in Computing Systems (CHI)*, pages 1–14.
- Weize Kong and James Allan. 2014. Extending faceted search to the general web. In *ACM International Conference on Conference on Information and Knowledge Management (CIKM)*, pages 839–848.
- Weize Kong and James Allan. 2016. Precision-oriented query facet extraction. In *ACM International on Conference on Information and Knowledge Management (CIKM)*, pages 1433–1442.
- Antonios Minas Krasakis, Mohammad Aliannejadi, Nikos Voskarides, and Evangelos Kanoulas. 2020. Analysing the effect of clarifying questions on document ranking in conversational search. In *ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*, pages 129–132.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Jiwei Li, Alexander H. Miller, Sumit Chopra, Marc Aurelio Ranzato, and Jason Weston. 2017. Dialogue learning with human-in-the-loop. In *International Conference on Learning Representations (ICLR)*.
- Margaret Li, Jason Weston, and Stephen Roller. 2019a. Acute-eval: Improved dialogue evaluation with optimized questions and multi-turn comparisons. *arXiv preprint arXiv:1909.03087*.

- Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- Ziming Li, Julia Kiseleva, and Maarten de Rijke. 2019b. Dialogue generation: From imitation learning to inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 6722–6729.
- Ziming Li, Sungjin Lee, Baolin Peng, Jinchao Li, Julia Kiseleva, Maarten de Rijke, Shahin Shayandeh, and Jianfeng Gao. 2020. Guided dialogue policy learning without adversarial learning in the loop. In *Findings of the Association for Computational Linguistics (EMNLP)*, pages 2308–2317.
- Ziming Li, Dookun Park, Julia Kiseleva, Young-Bum Kim, and Sungjin Lee. 2021. DEUS: A data-driven approach to estimate user satisfaction in multi-turn dialogues. *arXiv preprint arXiv:2103.01287*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Tom Lotze, Stefan Klut, Mohammad Aliannejadi, and Evangelos Kanoulas. 2021. Ranking clarifying questions based on predicted user engagement. *arXiv preprint arXiv:2103.06192*.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1116–1126.
- Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues*, 56(1):81–103.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for NLG. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2241–2252.
- Wenjia Ou and Yue Lin. 2020. A clarifying question selection system from ntes.along in ConvAI3 challenge. *arXiv preprint arXiv:2010.14202*.
- Dragomir R. Radev, Hong Qi, Harris Wu, and Weiguo Fan. 2002. Evaluating web-based question answering systems. In *International Conference on Language Resources and Evaluation (LREC)*, pages 1153–1156.
- Filip Radlinski and Susan Dumais. 2006. Improving personalized web search using result diversification. In *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 691–692.
- Sudha Rao and Hal Daumé III. 2018. Learning to ask good questions: Ranking clarification questions using neural expected value of perfect information. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2737–2746.
- Ehud Reiter. 2018. A structured review of the validity of bleu. *Computational Linguistics*, 44(3):393–401.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at TREC-3. *Nist Special Publication Sp*, 109:109.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building an open-domain chatbot. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 300–325.
- Corby Rosset, Chenyan Xiong, Xia Song, Daniel Campos, Nick Craswell, Saurabh Tiwary, and Paul Bennett. 2020. Leading conversational search by suggesting useful questions. In *The Web Conference*, pages 1160–1170.
- Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D. Manning. 2019a. Do massively pretrained language models make better storytellers? In *Conference on Computational Natural Language Learning (CoNLL)*, pages 843–861.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019b. What makes a good conversation? how controllable attributes affect human judgments. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1702–1723.
- Ivan Sekulic, Mohammad Aliannejadi, and Fabio Crestani. 2021. User engagement prediction for clarification in search. In *European Conference on Information Retrieval (ECIR)*, pages 619–633.
- Amir Soleimani, Christof Monz, and Marcel Worring. 2021. NLQuAD: A non-factoid long question answering data set. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 1245–1255.
- Shang-Yu Su, Xiujun Li, Jianfeng Gao, Jingjing Liu, and Yun-Nung Chen. 2018. Discriminative deep dyna-q: Robust planning for dialogue policy learning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3813–3823.
- Ryuichi Takanobu, Hanlin Zhu, and Minlie Huang. 2019. Guided dialog policy learning: Reward estimation for multi-domain task-oriented dialog. In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 100–110.

- Jan Trienes and Krisztian Balog. 2019. Identifying unclear questions in community question answering websites. In *European Conference on Information Retrieval (ECIR)*, pages 276–289.
- David Vallet and Pablo Castells. 2012. Personalized diversification of search results. In *ACM SIGIR conference on Research and Development in Information Retrieval (SIGIR)*, pages 841–850.
- EM Voorhees. 1999. Proceedings of the 8th text retrieval conference. *TREC-8 Question Answering Track Report*, pages 77–82.
- Yansen Wang, Chenyi Liu, Minlie Huang, and Liqiang Nie. 2018. Learning to ask questions in open-domain conversational systems with typed decoders. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2193–2203.
- Yining Wang, Liwei Wang, Yuanzhi Li, Di He, and Tie-Yan Liu. 2013. A theoretical analysis of NDCG type ranking measures. In *Conference on Learning Theory (COLT)*, pages 25–54.
- Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020a. Generating clarifying questions for information retrieval. In *The Web Conference*, pages 418–428.
- Hamed Zamani, Gord Lueck, Everest Chen, Rodolfo Quispe, Flint Luu, and Nick Craswell. 2020b. MIMICS: A large-scale data collection for search clarification. In *ACM International Conference on Information and Knowledge Management (CIKM)*, pages 3189–3196.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2204–2213.
- 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 *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 270–278.