Question Rewriting for Conversational Question Answering

Svitlana Vakulenko*
University of Amsterdam
s.vakulenko@uva.nl

Zhucheng Tu
Apple Inc.
zhucheng_tu@apple.com

Shayne Longpre
Apple Inc.
slongpre@apple.com

Raviteja Anantha
Apple Inc.
raviteja_anantha@apple.com

ABSTRACT

Conversational question answering (QA) requires the ability to correctly interpret a question in the context of previous conversation turns. We address the conversational QA task by decomposing it into question rewriting and question answering subtasks. The question rewriting (QR) subtask is specifically designed to reformulate ambiguous questions, which depend on the conversational context, into unambiguous questions that can be correctly interpreted outside of the conversational context. We introduce a conversational QA architecture that sets the new state of the art on the TREC CAsT 2019 passage retrieval dataset. Moreover, we show that the same QR model improves QA performance on the QuAC dataset with respect to answer span extraction, which is the next step in QA after passage retrieval. Our evaluation results indicate that the QR model we proposed achieves near human-level performance on both datasets and the gap in performance on the end-to-end conversational QA task is attributed mostly to the errors in QA.

CCS CONCEPTS

• **Information systems** Question answering; Query reformulation; Information retrieval; Information extraction.

KEYWORDS

conversational search, question answering, question rewriting

ACM Reference Format:

Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question Rewriting for Conversational Question Answering. In *preparation for ACM Conference on Web Search and Data Mining, March 8–12, 2021, Virtual Event, Israel*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3437963.3441748>

1 INTRODUCTION

Extending question answering systems to a conversational setting is an important development towards a more natural human-computer interaction [9]. In this setting a user is able to ask several

follow-up questions, which omit but reference information that was already introduced earlier in the same conversation, for example:

- (Q) - Where is Xi'an?
- (A) - Shaanxi, China.
- (Q) - What is **its** GDP?
- (A) - 95 Billion USD.
- (Q) - What is the share (*of Xi'an*) in the (*Shaanxi*) province GDP?
- (A) - 41.8% of Shaanxi's total GDP

This example highlights two linguistic phenomena characteristic of a human dialogue, which include **anaphora** (words that explicitly reference previous conversation turns) and *ellipsis* (words that can be omitted from the conversation) [35]. Therefore, conversational QA models require mechanisms capable of resolving contextual dependencies to correctly interpret such follow-up questions. While existing co-reference resolution tools are designed to handle anaphoras, they do not provide any support for ellipsis [4]. Previous research showed that QA models can be directly extended to incorporate conversation history but a considerable room for improvement still remains especially when scaling such models beyond a single input document [2, 26, 31].

In this paper we show how to extend existing state-of-the-art QA models with a QR component and demonstrate that the proposed approach improves the performance on the end-to-end conversational QA task. This QR component is specifically designed to handle ambiguity of the follow-up questions by rewriting them such that they can be processed by existing QA models as standalone questions outside of the conversation context. This setup also offers a wide range of practical advantages over the single end-to-end conversational QA model:

- **Traceability:** QR component produces the question that the QA model is actually trying to answer, which allows to separate errors that stem from an incorrect question interpretation from errors in question answering. Our error analysis makes full use of this feature by investigating different sources of errors and their correlation.
- **Reuse:** QR allows to reduce the conversational QA task to the standard QA task and leverage already existing QA models and datasets. Any new non-conversational QA model can be immediately ported into a conversational setting using QR. A single QR model can be reused across several alternative QA architectures as we show in our experiments.
- **Modularity:** In practice, information is distributed across a network of heterogeneous nodes that do not share internal

*This research was completed during the internship at Apple Inc.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

WSDM '21, March 8–12, 2021, Virtual Event, Israel.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8297-7/21/03.
<https://doi.org/10.1145/3437963.3441748>

representations. A natural language question can provide an adequate communication protocol between these distributed components [30]. Instead of sending the content of the whole conversation to a 3rd-party API, QR allows for formulating a concise question that contains only information relevant to the current information need, which also helps to determine which of the distributed systems should be used for answering the question.

Two dimensions on which QA tasks vary are: the type of data source used to retrieve the answer (e.g., a paragraph, a document collection, or a knowledge graph); and the expected answer type (a text span, a ranked list of passages, or an entity). In this paper, we experiment with two variants of the QA task: *retrieval QA*, the task of finding an answer to a given natural-language question as a ranked list of relevant passages given a document collection; and *extractive QA*, the task of finding an answer to a given natural-language question as a text span within a given passage. Though the two QA tasks are complementary to each other, in this paper we focus on the QR task and its ability to enable different types of QA models within a conversational setting. We experiment with both retrieval and extractive QA models to examine the effect of the QR component on the end-to-end QA performance. The contributions of this work are three-fold:

- (1) We introduce a novel approach for the conversational QA task based on question rewriting that sets the new state-of-the-art results on the TREC CAsT dataset for passage retrieval.
- (2) We show that the same question rewriting model (trained on the same dataset) boosts the performance of the state-of-the-art architecture for answer extraction on the QuAC dataset.
- (3) We systematically evaluate the proposed approach in both retrieval and extractive settings, and report the results of our error analysis.

2 RELATED WORK

Conversational QA is an extension of the standard QA task that introduces contextual dependencies between the input question and the previous dialogue turns. Several datasets were recently proposed extending different QA tasks to a conversational setting including extractive [2, 31], retrieval [4] and knowledge graph QA [3, 12]. One common approach to conversational QA is to extend the input of a QA model by appending previous conversation turns [3, 14, 27]. Such approach, however falls short in case of retrieval QA, which requires a concise query as input to the candidate selection step, such as BM25 [24]. Results of the recent TREC CAsT track demonstrated that co-reference models are also not sufficient to resolve the missing context in the follow-up questions [4]. A considerable gap between the performance of automated rewriting approaches and manual human annotations call for new architectures that are capable of retrieving relevant answers from large text collections using conversational context.

Query rewriting is a technique that was successfully applied in a variety of tasks, including data integration, query optimization and query expansion in sponsored search [11, 25, 33]. More recently rewriting conditioned on the conversation history was

shown to increase effectiveness of chitchat and task-oriented dialogue models [30, 34]. QR in the context of the conversational QA task was first introduced by Elgohary et al. [7], who released the CANARD dataset that contains human rewrites of questions from QuAC conversational reading comprehension dataset. Their evaluation, however, was limited to analysing rewriting quality without assessing impact from rewriting on the end-to-end conversational QA task. Our study was designed to bridge this gap and extend evaluation of question rewriting for conversational QA to the passage retrieval task as well.

The field of conversational QA is advancing rapidly due to the challenges organised by the community, such as the TREC CAsT challenge [4]. Several recent studies that are concurrent to our work proposed alternative approaches for the conversational QA task and evaluated them on the TREC CAsT dataset. Mele et al. [21] introduced an unsupervised approach that relies on a set of heuristics using part-of-speech tags, dependency parses and co-reference resolution to rewrite questions. Voskarides et al. [37] train a binary classification model on CANARD dataset that learns to pick the terms from the conversation history for query expansion. Yu et al. [41] train a sequence generation model using a set of rules applied to the MS MARCO search sessions. We demonstrate that our approach does not only outperform all the results reported on the TREC CAsT dataset previously but also boosts the performance on the answer extraction task as demonstrated on the QuAC dataset.

3 QUESTION REWRITING TASK

We assume the standard setting of a conversational (sequential) QA task, in which a user asks a sequence of questions and the system is to provide an answer after each of the questions. Every follow-up question may be ambiguous, i.e., its interpretation depends on the previous conversation turns (example: “What is its GDP?”). The task of question rewriting is to translate every ambiguous follow-up question into a semantically equivalent but unambiguous question (for this example: “What is the GDP of Xi’an?”). Then, every question is first processed by the QR component before passing it to the QA component.

More formally, given a conversation context C and a potentially implicit question Q , a question which may require the conversation context C to be fully interpretable, the task of a question rewriting (QR) model is to generate an explicit question Q' which is equivalent to Q under conversation context C and has the same correct answer A . See Figure 1 for an illustrative example that shows how QR can be combined with any non-conversational QA model.

4 APPROACH

In this section, we describe the model architectures of our QR and QA components. We use two distinct QA architectures to show that the same QR model can be successfully applied across several QA models irrespective of their implementation. One of the QA models is designed for the task of passage retrieval and the other one for the task of answer extraction from a passage (reading comprehensions). We employ state-of-the-art QA architectures that were already successfully evaluated for these tasks in non-conversational settings and show how QR allows to port existing QA models into a conversational setting.

Figure 1: Our approach for end-to-end conversational QA relies on the question rewriting component to handle conversation context and produce an explicit question that can be fed to standard, non-conversational QA components.

4.1 Question Rewriting Model

We use a model for question rewriting, which employs a unidirectional Transformer decoder [28] for both encoding the input sequence and decoding the output sequence. The input to the model is the question with previous conversation turns (we use 5 previous turns in our experiments) turned into token sequences separated with a special *[SEP]* token.

The training objective is to predict the output tokens provided in the ground truth question rewrites produced by human annotators. The model is trained via teacher forcing approach, which is a standard technique for training language generation models, to predict every next token in the output sequence given all the preceding tokens. The loss is calculated via negative log-likelihood (cross-entropy) between the output distribution $D' \in \mathbb{R}^{|V|}$ over the vocabulary V , and the one-hot vector $y_{Q'} \in \mathbb{R}^{|V|}$ for the correct token from the ground truth: $loss = -y_{Q'} \log D'$.

At training time the output sequence is shifted by one token and is used as input to predict all next tokens of the output sequence at once. At inference time, the model uses maximum likelihood estimation to select the next token from the final distribution D' (greedy decoding), as shown in Figure 2.

We further increase capacity of our generative model by learning to combine several individual distributions (D'_1 and D'_2 in Figure 2). The final distribution D' is then produced as a weighted sum of the intermediate distributions: $D' = \sum_{i=1}^m \alpha_i D'_i$ ($m = 2$ in our experiments). To produce $D'_i \in \mathbb{R}^{|V|}$ we pass the last hidden state of the Transformer Decoder $h \in \mathbb{R}^d$ through a separate linear layer for each intermediary distribution: $D'_i = W_i^H h + b^H$, where W_i^H is the weight matrix and b^H is the bias. For the weighting coefficients α_i we use the matrix of input embeddings $X \in \mathbb{R}^{n \times d}$, where n is the maximum sequence length and d is the embedding dimension, and the output of the first attention head of the Transformer Decoder $G \in \mathbb{R}^{n \times d}$ put through a layer normalization function: $\alpha_i = W_i^G \text{norm}(G) + W_i^X X + b_i^\alpha$, where all W are the weight matrices and b_i^α is the bias.

Figure 2: The question rewriting component uses the Transformer Decoder architecture, to recursively generate the tokens of an "explicit" question. At inference time, the generated output is appended to the input sequence for the next timestep in the sequence.

4.2 Retrieval QA Model

In the retrieval QA settings, the task is to produce a ranked list of text passages from a collection, ordered by their relevance to a given a natural language question [6, 22]. We employ the state-of-the-art approach to retrieval QA, which consists of two phases: candidate selection and passage re-ranking. This architecture holds

the state-of-the-art for the passage retrieval task on the MS MARCO dataset according to the recent experiments conducted by Xiong et al. [39].

In the first phase, a traditional retrieval algorithm (BM25) is used to quickly sift through the indexed collection retrieving top- k passages ranked by relevance to the input question Q' . In the second phase, a more computationally-expensive model is used to re-rank all question-answer candidate pairs formed using the previously retrieved set of k passages.

For re-ranking, we use a binary classification model that predicts whether the passage answers a question, i.e., the output of the model is the relevance score in the interval $[0, 1]$. The input to the re-ranking model is the concatenated question and passage with a separation token in between (see Figure 3 for the model overview). The model is initialized with weights learned from unsupervised pre-training on the language modeling (masked token prediction) task (BERT) [5]. During fine-tuning, the training objective is to reduce cross-entropy loss, using relevant passages and non-relevant passages from the top- k candidate passages.

Figure 3: Retrieval QA component includes two sequential phases: candidate selection (BM25) followed by passage re-ranking (Transformer Encoder).

4.3 Extractive QA Model

The task of extractive QA is given a natural language question and a single passage find an answer as a contiguous text span within the given passage [29]. Our model for extractive QA consists of a Transformer-based bidirectional encoder (BERT) [5] and an output layer predicting the answer span. This type of model architecture corresponds to the current state-of-the-art setup for several reading comprehension benchmarks [15, 19].

The input to the model is the sequence of tokens formed by concatenating a question and a passage separated with a special $[SEP]$ token. The encoder layers are initialized with the weights of a Transformer model pre-trained on an unsupervised task (masked token prediction). The output of the Transformer encoder is a hidden vector T_i for each token i of the input sequence.

For fine-tuning the model on the extractive QA task, we add weight matrices W^s, W^e and biases b^s, b^e that produce two probability distributions over all the tokens of the given passage separately for the start (S) and end position (E) of the answer span. For each token i the output of the Transformer encoder T_i is passed through a linear layer, followed by a softmax normalizing the output logits over all the tokens into probabilities:

$$S_i = \frac{e^{W^s \cdot T_i + b^s}}{\sum_{j=1}^n e^{W^s \cdot T_j + b^s}} \quad E_i = \frac{e^{W^e \cdot T_i + b^e}}{\sum_{j=1}^n e^{W^e \cdot T_j + b^e}} \quad (1)$$

The model is then trained to minimize cross-entropy between the predicted start/end positions (S_i and E_i) and the correct ones from the ground truth (y_S and y_E are one-hot vectors indicating the correct start and end tokens of the answer span):

$$loss = - \sum_{i=1}^n y_S \log S_i - \sum_{i=1}^n y_E \log E_i \quad (2)$$

At inference time all possible answer spans from position i to position j , where $j \geq i$, are scored by the sum of end and start positions' probabilities: $S_i + E_j$. The output of the model is the maximum scoring span (see Figure 4 for the model overview).

21% of the CANARD (QuAC) examples are Not Answerable (NA) by the provided passage. To enable our model to make No Answer predictions we prepend a special $[CLS]$ token to the beginning of the input sequence. For all No Answer samples we set both the ground truth start and end span positions to this token's position (0). Likewise, at inference time, predicting this special token is equivalent to a No Answer prediction for the given example.

Figure 4: Extractive QA component predicts a span of text in the paragraph P' , given an input sequence with the question Q' and passage P' .

5 EXPERIMENTAL SETUP

To evaluate the question rewriting approach we perform a range of experiments with several existing state-of-the-art QA models. In

Table 1: Datasets used for training and evaluation (with the number of questions).

Tasks	Question Rewriting	Retrieval QA	Extractive QA
Train	CANARD (35k)	MS MARCO (399k)	+MultiQA (75k) CANARD (35k)
Test	TREC CAsT (173) CANARD (5.5k)	TREC CAsT (173)	CANARD (5.5k)

the following subsections we describe the datasets used for training and evaluation, the set of metrics for each of the components, our baselines and details of the implementation. Our experimental setup is designed to evaluate gains in performance, reuse and traceability from introducing the QR component:

RQ1: How does the proposed approach perform against competitive systems (performance)?

RQ2: How does non-conversational pre-training benefit the models with and without QR (reuse)?

RQ3: What is the proportion of errors contributed by each of the components (traceability)?

5.1 Datasets

We chose two conversational QA datasets for the evaluation of our approach: (1) CANARD, derived from Question Answering in Context (QuAC) for extractive conversational QA [2], and (2) TREC CAsT for retrieval conversational QA [4]. See Table 1 for the overview of the datasets. Since TREC CAsT is relatively small we used only CANARD for training QR. The same QR model trained on CANARD is evaluated on both CANARD and TREC CAsT.

CANARD [7] contains sequences of questions with answers as text spans in a given Wikipedia article. It was built upon the QuAC dataset [2] by employing human annotators to rewrite original questions from QuAC dialogues into explicit questions. CANARD consists of 40.5k pairs of question rewrites that can be matched to the original answers in QuAC. We use CANARD splits for training and evaluation. We use the question rewrites provided in CANARD and articles with correct answer spans from QuAC. In our experiments, we refer to this joint dataset as CANARD for brevity.

To further boost performance of the extractive QA model we reuse the approach from Fisch et al. [8] and pre-train the model on MultiQA dataset, which contains 75k QA pairs from six standard QA benchmarks, and then fine-tune it on CANARD to adapt for the domain shift. Note that MultiQA is a non-conversational QA dataset.

TREC CAsT contains sequences of questions with answers to be retrieved as text passages from a large collection of passages: MS MARCO [22] with 8.6M passages + TREC CAR [6] with 29.8M passages. The official test set that we used for evaluation contains relevance judgements for 173 questions across 20 dialogues (topics). We use the model from Nogueira and Cho [23] for retrieval QA, which was tuned on a sample from MS MARCO with 12.8M query-passage pairs and 399k unique queries.

5.2 Metrics

Our evaluation setup corresponds to the one used in TREC CAsT, which makes our experiments directly comparable to the official TREC CAsT results. Mean average precision (*MAP*), mean reciprocal rank (*MRR*), normalized discounted cumulative gain (*NDCG@3*) and precision on the top-passage (*P@1*) evaluate quality of passage retrieval. Top-1000 documents are considered per query with a relevance judgement value cut-off level of 2 (the range of relevance grades in TREC CAsT is from 0 to 4).

We use *F1* and Exact Match (*EM*) for extractive QA, which measure word token overlap between the predicted answer span and the ground truth. We also report accuracy for questions without answers in the given passage (*NA Acc*).

Our analysis showed that *ROUGE* recall calculated for unigrams (*ROUGE-1* recall) correlates with the human judgement of the question rewriting performance (Pearson 0.69), which we adopt for our experiments as well. *ROUGE* [18] is a standard metric of lexical overlap, which is often used in text summarization and other text generation tasks. We also calculate question similarity scores with the Universal Sentence Encoder (*USE*) model [1] (Pearson 0.71).

5.3 QR Baselines

The baselines were designed to challenge the need for a separate QR component by incorporating previous turns as direct input to custom QA components. Manual rewrites by human annotators provide the upper-bound performance for a QR approach and allows for an ablation study of the down-stream QA components.

Original. Original questions from the conversational QA datasets without any question rewriting.

Original + *k*-DT. Our baseline approach for extractive QA prepends the previous *k* questions to the original question to compensate for the missing context. The questions are separated with a special token and used as input to the Transformer model. We report the results for $k = \{1, 2, 3\}$.

Original + *k*-DT*. Since in the first candidate selection phase we use BM25 retrieval function which operates on a bag-of-words representation, we modify the baseline approach for retrieval QA as follows. We select keywords from *k* prior conversation turns (not including current turn) based on their inverse document frequency (IDF) scores and append them to the original question of the current turn. We use the keyword-augmented query as the search query for *Anserini* [40], a Lucene toolkit for replicable information retrieval research, and if we use BERT re-ranking we concatenate the keyword-augmented query with the passages retrieved from the keyword-augmented query. We use the keywords with IDF scores above the threshold of 0.0001, which was selected based on a 1 million document sample of the MS MARCO corpus.

Human. To provide an upper bound, we evaluate all our models on the question rewrites manually produced by human annotators.

5.4 QR Models

In addition to the baselines described above, we chose several alternative models for question rewriting of the conversational context: (1) co-reference resolution as in the TREC CAsT challenge;

(2) PointerGenerator used for question rewriting on CANARD by Elgohary et al. [7] but not previously evaluated on the end-to-end conversational QA task; (3) CopyTransformer extension of the PointerGenerator model that replaces the bi-LSTM encoder-decoder architecture with a Transformer Decoder model. All models, except for co-reference, were trained on the train split of the CANARD dataset. Question rewrites are generated turn by turn for each dialogue recursively using already generated rewrites as previous turns. This is the same setup as in the TREC CAsT evaluation.

Co-reference. Anaphoric expressions in original questions are replaced with their antecedents from the previous dialogue turns. Co-reference dependencies are detected using a publicly available neural co-reference resolution model that was trained on OntoNotes [16].¹

PointerGenerator. A sequence-to-sequence model for text generation with bi-LSTM encoder and a pointer-generator decoder [7].

CopyTransformer. The Transformer decoder, which, similar to pointer-generator model, uses one of the attention heads as a pointer [10]. The model is initialized with the weights of a pre-trained GPT2 model [28, 38] (Medium-sized GPT-2 English model: 24-layer, 1024-hidden, 16-heads, 345M parameters) and then fine-tuned on the question rewriting task.

Transformer++. The Transformer-based model described in Section 4.1. Transformer++ is initialized with the weights of the pre-trained GPT2 model, same as in CopyTransformer.

5.5 QA Models

Our retrieval QA approach is implemented as proposed in [23] using Anserini for the candidate selection phase with BM25 (top-1000 passages) and *BERT_{LARGE}* for the passage re-ranking phase (Anserini + BERT). Both components were fine-tuned only on the MS MARCO dataset ($k_1 = 0.82$, $b = 0.68$).²

We train several models for extractive QA on different variants of the training set based on the CANARD training set [7]. All models are first initialized with the weights of the *BERT_{LARGE}* model pre-trained using the whole word masking [5].

CANARD-O. The baseline models were trained using original (implicit) questions of the CANARD training set with a dialogue context of varying length (Original and Original + k -DT). The models are trained separately for each $k = \{0, 1, 2, 3\}$, where $k = 0$ corresponds to the model trained only on the original questions without any previous dialogue turns.

CANARD-H. To accommodate input of the question rewriting models, we train a QA model that takes human rewritten question from the CANARD dataset as input without any additional conversation context, i.e., as in the standard QA task.

MultiQA $\rightarrow$ CANARD-H. Since the setup with rewritten questions does not differ from the standard QA task, we experiment with pretraining the extractive QA model on the MultiQA dataset with explicit questions [8], using parameter choices introduced by Longpre et al. [20]. We further fine-tune this model on the target

Table 2: Evaluation results of the QR models. *Human performance is measured as the difference between two independent annotators' rewritten questions, averaged over 100 examples. This provides an estimate of the upper bound.

Test Set	Question	ROUGE	USE	EM
CANARD	Original	0.51	0.73	0.12
	Co-reference	0.68	0.83	0.48
	PointerGenerator	0.75	0.83	0.22
	CopyTransformer	0.78	0.87	0.56
	Transformer++	0.81	0.89	0.63
	Human*	0.84	0.90	0.33
TREC CAsT	Original	0.67	0.80	0.28
	Co-reference	0.71	0.80	0.13
	PointerGenerator	0.71	0.82	0.17
	CopyTransformer	0.82	0.90	0.49
	Transformer++	0.90	0.94	0.58
	Human*	1.00	1.00	1.00

CANARD dataset to better adopt it for the domain shift in CANARD QA samples (see Figure 5).

6 RESULTS

RQ1: How does the proposed approach perform against competitive systems (performance)? Our approach, using question rewriting for conversational QA, consistently outperforms the baselines that use previous dialogue turns, in both retrieval and extractive QA tasks (see Tables 3-5). Moreover, it shows considerable improvement over the latest results reported on the TREC CAsT dataset (see Table 4).

The precision-recall trade-off curve in Figure 6 shows that question rewriting performance is close to the performance achieved by manually rewriting implicit questions. Our results also indicate that the QR performance metric is able to correctly predict the model that performs consistently better across both QA tasks (see Table 2).

Passage re-ranking with BERT always improves ranking results (almost a two-fold increase in MAP, see Table 3). Keyword-based baselines (Original + k -DT*) prove to be very strong outperforming both Co-reference and PointerGenerator models on all three performance metrics. Both MRR and NDCG@3 are increasing with the number of turns used for sampling keywords, while MAP is slightly decreasing, which indicates that it brings more relevant results at the very top of the rank but non-relevant results also receive higher scores. In contrast, the baseline results for Anserini + BERT model indicate that the re-ranking performance for all metrics decreases if the keywords from more than 2 previous turns are added to the original question.

Similarly for extractive QA, the model incorporating previous turns proved to be a very strong baseline (see Table 5). The performance results also suggest that all models do not discriminate well the passages that do not have an answer to the question (71% accuracy on the human rewrites). We notice that the baseline models, which were trained with the previous conversation turns, tend to adopt a conservative strategy by answering more questions as

¹<https://github.com/kentonl/e2e-coref>

²<https://github.com/nyu-dl/dl4marco-bert>

Table 3: Retrieval QA results on the TREC CAsT test set.

QA Input	QA Model	MAP	MRR	NDCG@3
Original	Anserini	0.089	0.245	0.131
Original + 1-DT*		0.133	0.343	0.199
Original + 2-DT*		0.130	0.374	0.213
Original + 3-DT*		0.127	0.396	0.223
Co-reference		0.109	0.298	0.172
PointerGenerator		0.100	0.273	0.159
CopyTransformer		0.148	0.375	0.213
Transformer++		0.190	0.441	0.265
Human		0.218	0.500	0.315
Original	Anserini	0.172	0.403	0.265
Original + 1-DT*	+BERT	0.230	0.535	0.378
Original + 2-DT*		0.245	0.576	0.404
Original + 3-DT*		0.238	0.575	0.401
Co-reference		0.201	0.473	0.316
PointerGenerator		0.183	0.451	0.298
CopyTransformer		0.284	0.628	0.440
Transformer++		0.341	0.716	0.529
Human		0.405	0.879	0.589

Table 4: Comparison with the state-of-the-art results reported on the TREC CAsT test set.

Approach	NDCG@3
Mele et al. [21]	0.397
Voskarides et al. [37]	0.476
Yu et al. [41]	0.492
Ours	0.529

“Unanswerable” (NA). Controlling for this effect, we show that our QR model gains a higher performance level by actually answering questions that have answers.

RQ2: How does non-conversational pre-training benefit the models with and without QR (reuse)? We observe that pre-training on MultiQA improves performance of all extractive QA model. However, it is much more prominent for the systems using question rewrites (7% increase in EM and 6% in F1 when using human-rewritten questions). The models do not require any additional fine-tuning, when using QR, since the type of input to the QA model remains the same (non-conversational). While for Original + k -DT models fine-tuning is required also to adopt the model for the new type of input data (conversational). Note that, in this case, we had to fine-tune all models but for another reason. The style of questions in CANARD is rather different from other QA datasets in MultiQA. Figure 5 demonstrates that a small portion of the training data is sufficient to adopt a QR-based model trained with non-conversational samples to work well on CANARD.

RQ3: What is the proportion of errors contributed by each of the components (traceability)? We measure the effect from question rewriting for each of the questions by comparing the answers produced for the original, the model-rewritten (Transformer++) and

Table 5: Extractive QA results on the CANARD test set. F1 and EM is calculated for both answerable and unanswerable questions, while NA Acc only for unanswerable questions.

QA Input	Training Set	EM	F1	NA Acc
Original	CANARD-O	38.68	53.65	66.55
Original + 1-DT		42.04	56.40	66.72
Original + 2-DT		41.29	56.68	68.11
Original + 3-DT		42.16	56.20	68.72
Original	CANARD-H	39.44	54.02	65.42
Human		42.36	57.12	68.20
Original	MultiQA →	41.32	54.97	65.84
Original + 1-DT	CANARD-H	43.15	57.03	68.64
Original + 2-DT		42.20	57.33	69.42
Original + 3-DT		43.29	57.87	71.50
Co-reference		42.70	57.59	66.20
PointerGenerator		41.93	57.37	63.16
CopyTransformer		42.67	57.62	68.02
Transformer++		43.39	58.16	68.29
Human		45.40	60.48	70.55

Figure 5: Effect from fine-tuning the MultiQA model on a portion of the target CANARD-H dataset due to the domain shift between the datasets.

the human-rewritten question (see Tables 6-7). This approach allows us to pinpoint the cases, in which QR contributes to the QA performance, and distinguish them from cases in which the answer can be found using the original question as well.

Assuming that humans always produce correct question rewrites, we can attribute all cases in which these rewrites did not result in a correct answer as errors of the QA component (rows 1-4 in Tables 6-7). The next two rows 5-6 show the cases, where human rewrites succeeded but the model rewrites failed, which we consider to be a likely error of the QR component. The last two rows are true positives for our model, where the last row combines cases where the original question was just copied without rewriting (numbers in brackets) and other cases when rewriting was not required. Since

Figure 6: Precision-recall curve illustrating model performance on the TREC CAsT test set for Anserini + BERT.

there is no single binary measure for the answer correctness, we select different cut-off thresholds for our QA metrics.

The majority of errors stem from the QA model: 29% of the test samples for retrieval and 55% for extractive estimated for P@1 and F1, comparing to 11% and 5% for QR respectively. Note that it is a rough estimate since we cannot automatically distinguish the cases that failed both QA and QR.

Overall, we observe that the majority of questions in extractive QA setup can be correctly answered without rewriting or accessing the conversation history. In other words, the extractive QA model tends to return an answer even when given an incomplete ambiguous question. This finding also explains the low NA Acc results reported in Table 5. Our results provide evidence of the deficiency of the reading comprehension setup, which was also reported in the previous studies on non-conversational datasets [13, 17, 32].

In contrast, in the retrieval QA setup, only 10% of the questions in TREC CAsT were rewritten by human annotators that did not need rewriting to retrieve the correct answer. These results highlight the difference between the retrieval and extractive QA training/evaluation setup. More details on our error analysis approach and results can be found in Vakulenko et al. [36].

7 CONCLUSION

We showed in an end-to-end evaluation that question rewriting is effective in extending standard QA approaches to a conversational setting. Our results set the new state-of-the-art on the TREC CAsT 2019 dataset. The same QR model also shows superior performance on the answer span extraction task evaluated on the CANARD/QuAC dataset. Based on the results of our analysis, we conclude that QR is a challenging but also very promising task that can be effectively implemented into conversational QA approaches.

We also confirmed that the QR metric provides a good indicator for the end-to-end QA performance. Thereby, it is reliable to be used for the QR model selection, which would help to avoid more costly end-to-end evaluation in the future.

The experimental evaluation and the detailed analysis we provide increases our understanding of the main error sources. QR performance is sufficiently high on both CANARD and TREC CAsT datasets, while the QA performance even given human rewritten

Table 6: Break-down analysis of all retrieval QA results for the TREC CAsT dataset. Each row represents a group of QA samples that exhibit similar behaviour. ✓ indicates that the answer produced by the QA model was correct or ✗ – incorrect, according to the thresholds provided in the right columns. We consider three types of input for every QA sample: the question from the test set (Original), generated by the best QR model (Transformer++) or rewritten manually (Human). The numbers correspond to the count of QA samples for each of the groups. The numbers in parenthesis indicate how many questions do not require rewriting, i.e., should be copied from the original.

Original	QR	Human	P@1		NDCG@3	
			= 1	> 0	≥ 0.5	= 1
✗	✗	✗	49 (14)	10 (1)	55 (20)	154 (49)
✓	✗	✗	0	0	0	0
✗	✓	✗	2	0	1	0
✓	✓	✗	0	1	1	0
✗	✗	✓	19	10	25	4
✓	✗	✓	0	1	0	0
✗	✓	✓	48	63	47	11
✓	✓	✓	55 (37)	88 (52)	44 (33)	4 (4)
Total			173 (53)			

Table 7: Break-down analysis of all extractive QA results for the CANARD dataset, similar to Table 6.

Original	QR	Human	F1 > 0	F1 ≥ 0.5	F1 = 1
✗	✗	✗	847 (136)	1855 (235)	2701 (332)
✓	✗	✗	174	193	181
✗	✓	✗	19	35 (2)	40 (1)
✓	✓	✗	135	153	120
✗	✗	✓	141	288	232
✓	✗	✓	65 (1)	57 (1)	40
✗	✓	✓	226	324	269
✓	✓	✓	3964 (529)	2666 (428)	1988 (333)
Total			5571 (666)		

questions for both tasks lags behind. This result suggest that the major improvement in conversational QA will come from improving the standard QA models.

In future work we would like to evaluate QR performance with the joint model integrating passage retrieval and answer span extraction. Recent results reported by Qu et al. [26] indicate that there is sufficient room for improvement in the history modeling phase.

On the other hand, the QR-QA type of architecture is generic enough to incorporate other types of context, such as a user model or an environmental context obtained from multi-modal data (deictic reference). Experimental evaluation of QR-QA performance augmented with such auxiliary inputs is a promising direction for future work.

ACKNOWLEDGEMENTS

We would like to thank our colleagues Srinivas Chappidi, Bjorn Hoffmeister, Stephan Peitz, Russ Webb, Drew Frank, and Chris DuBois for their insightful comments.

REFERENCES

- [1] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. Universal Sentence Encoder for English. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 169–174.
- [2] Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question Answering in Context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2174–2184.
- [3] Philipp Christmann, Rishiraj Saha Roy, Abdalghani Abujabal, Jyotsna Singh, and Gerhard Weikum. 2019. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 729–738.
- [4] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2019. CAsT 2019: The Conversational Assistance Track Overview. In *Proceedings of the 28th Text REtrieval Conference*. 13–15.
- [5] 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*. 4171–4186.
- [6] Laura Dietz, Ben Gamari, Jeff Dalton, and Nick Craswell. 2018. TREC Complex Answer Retrieval Overview. TREC.
- [7] Ahmed Elgohary, Denis Peskov, and Jordan Boyd-Graber. 2019. Can You Unpack That? Learning to Rewrite Questions-in-Context. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*. 5920–5926.
- [8] Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. MRQA 2019 Shared Task: Evaluating Generalization in Reading Comprehension. *arXiv preprint arXiv:1910.09753* (2019).
- [9] Jianfeng Gao, Michel Galley, and Lihong Li. 2019. Neural Approaches to Conversational AI. *Foundations and Trends in Information Retrieval* 13, 2-3 (2019), 127–298.
- [10] Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-Up Abstractive Summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4098–4109.
- [11] Mihajlo Grbovic, Nemanja Djuric, Vladan Radosavljevic, Fabrizio Silvestri, and Narayan Bhamidipati. 2015. Context- and Content-aware Embeddings for Query Rewriting in Sponsored Search. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*. 383–392.
- [12] Daya Guo, Duyu Tang, Nan Duan, Ming Zhou, and Jian Yin. 2018. Dialog-to-Action: Conversational Question Answering Over a Large-Scale Knowledge Base. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 2946–2955*.
- [13] Robin Jia and Percy Liang. 2017. Adversarial Examples for Evaluating Reading Comprehension Systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 2021–2031.
- [14] Ying Ju, Fubang Zhao, Shijie Chen, Bowen Zheng, Xuefeng Yang, and Yunfeng Liu. 2019. Technical report on Conversational Question Answering. *CoRR abs/1909.10772* (2019).
- [15] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite BERT for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942* (2019).
- [16] Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-Order Coreference Resolution with Coarse-to-Fine Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 687–692.
- [17] Mike Lewis and Angela Fan. 2019. Generative Question Answering: Learning to Answer the Whole Question. In *7th International Conference on Learning Representations*.
- [18] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.
- [19] 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* (2019).
- [20] Shayne Longpre, Yi Lu, Zhucheng Tu, and Chris DuBois. 2019. An Exploration of Data Augmentation and Sampling Techniques for Domain-Agnostic Question Answering. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*. 220–227.
- [21] Ida Mele, Cristina Ioana Muntean, Franco Maria Nardini, Raffaele Perego, Nicola Tonello, and Ophir Frieder. 2020. Topic Propagation in Conversational Search. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 2057–2060.
- [22] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches*.
- [23] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085* (2019).
- [24] Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019. Document Expansion by Query Prediction. *CoRR abs/1904.08375* (2019).
- [25] Yannis Papakonstantinou and Vasilis Vassalos. 1999. Query Rewriting for Semistructured Data. In *SIGMOD 1999, Proceedings ACM SIGMOD International Conference on Management of Data*. 455–466.
- [26] Chen Qu, Liu Yang, Cen Chen, Minghui Qiu, W. Bruce Croft, and Mohit Iyyer. 2020. Open-Retrieval Conversational Question Answering. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 539–548.
- [27] Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W. Bruce Croft, and Mohit Iyyer. 2019. Attentive History Selection for Conversational Question Answering. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1391–1400.
- [28] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019).
- [29] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 2383–2392.
- [30] Pushpendre Rastogi, Arpit Gupta, Tongfei Chen, and Lambert Mathias. 2019. Scaling Multi-Domain Dialogue State Tracking via Query Reformulation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 97–105.
- [31] Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics* 7 (2019), 249–266.
- [32] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Semantically Equivalent Adversarial Rules for Debugging NLP models. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. 856–865.
- [33] Shariq Rizvi, Alberto O. Mendelzon, S. Sudarshan, and Prasan Roy. 2004. Extending Query Rewriting Techniques for Fine-Grained Access Control. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, Paris, France, June 13-18, 2004*. 551–562.
- [34] Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. Improving Multi-turn Dialogue Modelling with Utterance ReWriter. *arXiv preprint arXiv:1906.07004* (2019).
- [35] Andrew L Thomas. 1979. Ellipsis: The Interplay of Sentence Structure and Context. *Lingua Amsterdam* 47, 1 (1979), 43–68.
- [36] Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2020. A Wrong Answer or a Wrong Question? An Intricate Relationship between Question Reformulation and Answer Selection in Conversational Question Answering. *Proceedings of the 2020 EMNLP Workshop SCAI: The 5th International Workshop on Search-Oriented Conversational AI* (2020).
- [37] Nikos Voskarides, Dan Li, Pengjie Ren, Evangelos Kanoulas, and Maarten de Rijke. 2020. Query Resolution for Conversational Search with Limited Supervision. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 921–930.
- [38] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Transformers: State-of-the-art Natural Language Processing. *arXiv preprint arXiv:1910.03771* (2019).
- [39] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. *arXiv preprint arXiv:2007.00808* (2020).
- [40] Peilin Yang, Hui Fang, and Jimmy Lin. 2017. Anserini: Enabling the Use of Lucene for Information Retrieval Research. In *Proceedings of the 40th International Conference on Research and Development in Information Retrieval*. 1253–1256.
- [41] Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. Few-Shot Generative Conversational Query Rewriting. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 1933–1936.