







Open-Domain Conversational Search Assistant with Transformers

Rafael Ferreira^(✉) , Mariana Leite^(✉) , David Semedo^(✉) , and Joao Magalhaes^(✉) 

NOVA LINCS, School of Science and Technology,
NOVA School of Science and Technology, Almada, Portugal
{rah.ferreira,me.leite}@campus.fct.unl.pt, {df.semedo,jmag}@fct.unl.pt

Abstract. Open-domain conversational search assistants aim at answering user questions about open topics in a conversational manner. In this paper we show how the Transformer architecture [30] achieves state-of-the-art results in key IR tasks, leveraging the creation of conversational assistants that engage in open-domain conversational search with single, yet informative, answers. In particular, we propose an open-domain abstractive conversational search agent pipeline to address two major challenges: first, conversation context-aware search and second, abstractive search-answers generation. To address the first challenge, the conversation context is modeled with a query rewriting method that unfolds the context of the conversation up to a specific moment to search for the correct answers. These answers are then passed to a Transformer-based re-ranker to further improve retrieval performance. The second challenge, is tackled with recent Abstractive Transformer architectures to generate a digest of the top most relevant passages. Experiments show that Transformers deliver a solid performance across all tasks in conversational search, outperforming the best TREC CAsT 2019 baseline.

Keywords: Conversational search · Transformers · Query rewriting · Re-ranking · Answer generation

1 Introduction

Conversational search systems are an emerging research topic, and the natural evolution of the traditional search paradigm, allowing for a more natural interaction between users and search systems. Building intelligent systems able to establish and develop meaningful conversations is one of the key goals of AI and the ultimate goal of natural language research [9]. The interactions between a user and conversational systems have been studied in [32], which showed that users are willing to utilise conversational assistants as long as their needs are met with success. However, conversational search assistants still put a considerable burden on users that have to go through a list of documents, or passages, to find the information they need.

We depart from this document-based approach to conversational search, and propose an open-domain abstractive conversational assistant that is aware of the context of the conversation to generate a single and informative search-answer. We argue that by doing so, we can capture in one single and short answer the information contained on several relevant documents. Moreover, we show that Transformer architectures [30] outperform the state-of-the-art results across all the steps of the conversational system pipeline. Hence, the core contributions of this paper are twofold: first, we show that one can tightly integrate different Transformers to deliver an end-to-end conversational search pipeline with state-of-the-art results; second, abstractive answer generation can effectively compress the information of several retrieved passages into a short answer. These contributions are rooted in the groundbreaking architecture of the Transformer [30] that leverages attention mechanisms to model complex interactions between sequence data. In particular, we explore Transformer’s advantages to: (a) capture complex relations between conversation turns to rewrite a query in the middle of a conversation; (b) to look into the interactions between words in a conversation query and a candidate passage; and (c) to compress multiple retrieved passages into one single, yet informative, search-answer. The final result, is a complete conversational search assistant leveraged by the Transformer architecture.

In the following section, we discuss the related work. In Sect. 3 we detail the Transformer-based conversational search pipeline: the conversational query rewriting, the re-ranker, and abstractive answer generation. Evaluation is performed in Sect. 4 and Sect. 5 presents the key takeaway messages.

2 Related Work

Open-domain conversational search systems must account for the dialog context to provide a relevant passage. While research on interactive search systems has started long ago [1, 4, 23], the recent interest in having intelligent conversation assistants (e.g. Alexa, SIRI), has re-ignited this research field. Recent models [9, 17, 25, 31] leverage large open-domain collections (e.g. Wikipedia) to learn rich language-models using self-supervised neural networks. The applicability of these models in conversational search is twofold: grasping the dialog context and passage re-ranking. Recently, the TREC CAsT (Conversational Assistant Track) [6] task introduced a multi-turn passage retrieval dataset, enabling the development and evaluation of such models.

Conversational context-aware search models need to (a) keep track of the dialog context, and (b) select the most relevant passage. To address (a), one approach is to perform query rewriting to obtain context-independent queries. [10] observed that manually rewritten queries from QuAC [2] had enough context to be independently understandable. To automate the process, a sequence-to-sequence (seq2seq) model with attention and a copy mechanism was proposed. The model is given as input a sequence with the full conversation history and the query to be rewritten. In [31], a BERT model [7] is given as input a sequence of all terms of the current and previous queries, and is then fine-tuned on a binary

term classification task. Also using both the query and conversation history, in [17], a pre-trained T5 model [26] is fine-tuned on CANARD [10] to construct the context-independent query, and achieved state-of-the-art performance on the query-rewriting task. Task (b) is commonly addressed through re-ranking. Large pre-trained Transformer models, such as BERT [7], RoBERTa [18], and XLNet [36], have been widely adopted for re-ranking due to their generalisation capabilities. Examples of this are present in [12, 21, 22], where a Transformer-based model is fine-tuned on the question-answering relevance classification task.

Given the dialogue context, the agent must generate a natural language response. In chit-chat dialogue generation, most approaches use an encoder-decoder neural architecture that first encodes utterances and then the decoder generates a response [15, 16, 28, 29, 39]. In [15] and [16], reinforcement learning is used to overcome uninformative and general responses of standard seq2seq models. Another alternative is retrieval-based dialogue generation, in which the generator takes as input retrieved candidate documents to improve the comprehensiveness of the generated answer [28, 39]. These approaches require a large dataset with annotated dialogues, which is not feasible in our scenario. Alternatively, Transformer models have shown to be highly effective generative language models [14, 26, 38]. While both T5 [26] and BART [14] are general language models, PEGAGUS [38] focuses on abstractive summarisation, and obtained state-of-the-art results on 12 summarisation tasks.

3 Transformers-Based Conversational Search Assistant

In this section we formulate the open-domain conversational search task and describe the conversational assistant retrieval and answer generation components. The conversational search task is formally defined by a sequence of natural language conversational turns for a topic T , with queries q . For each conversation turn $T = \{q_1, \dots, q_i, \dots, q_n\}$, the conversational search task is to find relevant passages p_k for each query q_i , satisfying the user’s information need for that turn according to the conversational context. The proposed approach uses a four-stage architecture: (a) context tracking, (b) retrieval, (c) re-ranking, and (d) answer generation. An overview of the system’s architecture can be seen in Fig. 1 which we will detail in the following sections.

3.1 Conversational Query Rewriting Transformer

Due to the evolving nature of a conversational session, the current query may not include all the information needed to retrieve the answer that the user is looking for. This challenge is illustrated in the conversation presented in Table 1: in conversation turn 2, the system needs to understand that “its” refers to “Lucca’s” (explicit coreference) and in turn 3, where the important monuments should be focused in Lucca, although there is no direct evidence (implicit coreference), which makes the task even more challenging. We tackle this challenge by rewriting queries, using previous turns, making the current query context-independent.

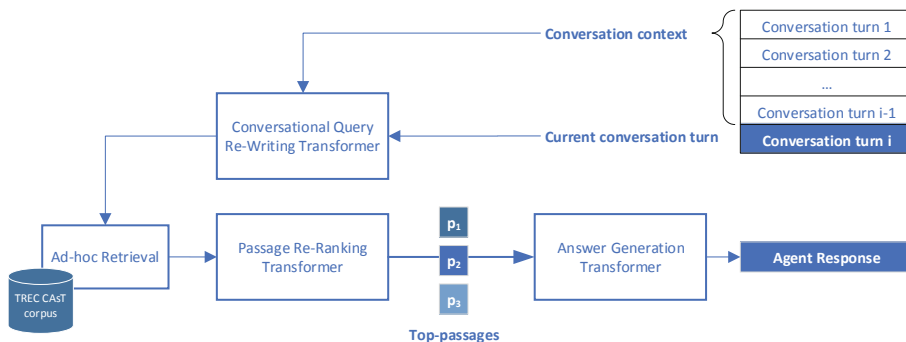


Fig. 1. The proposed Transformer-based conversational search assistant.

Table 1. Conversation example about a specific topic, in this case the city of Lucca.

Turn	Conversational query	Context-independent query
1	How is the climate in <u>Lucca</u> ?	How is the climate in <u>Lucca</u> ?
2	Tell me about <u>its</u> origins	Tell me about <u>Lucca's</u> origins
3	What monuments should I visit?	What monuments should I visit <u>in Lucca</u> ?

To perform the query rewriting task, we need a model capable of performing coreference resolution and include context from previous turns. The Text-to-text Transfer Transformer (T5) [26] can be fine-tuned to reformulate conversational queries [17] by providing as input the sequence of conversational queries and passages, and as target, the rewritten query. The training input sequence is constructed as:

$$“q_i [CTX] q_1 p_1 [TURN] q_2 p_2 [TURN] \dots [TURN] q_{i-1} p_{i-1}”, \quad (1)$$

where i is the current turn, q is a query, p_k is a passage retrieved from the index by the retrieval model, and $[CTX]$ and $[TURN]$ are special tokens. $[CTX]$ is used to separate the current query from the context (previous queries and passages) and $[TURN]$ is used to separate the historical turns (query-passage pair).

3.2 Passage Re-Ranking Transformer

With the new pre-trained neural language models, such as BERT [7] and others [18, 36], it is possible to generate contextual embeddings for a sentence and each of its tokens. These embeddings can be used as input to a model to perform passage re-ranking [21, 22]. This re-ranking step allows going beyond term matching, as the model has some understanding of both individual terms semantics as well as their interactions between queries and passages. As such, it is able to judge more thoroughly if a passage is relevant to a query.

Following this rationale, we tackle the passage re-ranking task with a BERT model [7], fine-tuned on the passage ranking task [21], through a binary relevance

classification task, where positive examples are relevant passages, and negative examples are non-relevant passages. To obtain the embedding of the query q , and passage p , a sequence with N tokens is given as input to BERT:

$$emb = BERT("[CLS] q [SEP] p"), \quad (2)$$

where $emb \in \mathbb{R}^{N \times H}$ (H is BERT embedding’s size) is the embeddings matrix of all tokens, and $[CLS]$ and $[SEP]$ are special tokens in BERT’s vocabulary, representing the classification and separation tokens, respectively. From emb we extract the embedding of the first token, which corresponds to the embedding of the $[CLS]$ token, $emb_{[CLS]} \in \mathbb{R}^H$. This embedding is then used as input to a single layer feed-forward neural network (FFNN), followed by a *softmax*, to obtain the probability of the passage being relevant to the query:

$$P(p|q) = softmax(FFNN(emb_{[CLS]})). \quad (3)$$

With $P(p|q)$ calculated for each passage p given a query q , the final rank is obtained by re-ranking according to the probability of being relevant.

3.3 Abstractive Search-Answer Generation Transformer

Having identified a set of candidate passages according to the scores given by the re-ranker model (Eq. 3), the goal is to generate a natural language response that combines the information comprised in each of the passages. To address this, we follow an abstractive summarisation approach, which unlike extractive summarisation that just selects existing sentences, it can portray both reading comprehension and writing abilities, thus allowing the generation of a concise and comprehensive digest of multiple input passages.

The Transformer [30] architecture has proved to be highly effective at modelling large dependency windows of textual sequences. Text-to-text approaches [14, 26, 38], trained over large and comprehensive collections, become effective at *understanding* different topics and retaining language regularities useful for several language tasks. Thus, to generate the agent’s response using a transformer model, we give as input the following sequence:

$$"p_1 p_2 \dots p_N", \quad (4)$$

where each p_k corresponds to one of the top- N candidate passages. With this strategy, we implicitly bias the answer generation by asking the model to summarise the passages that are deemed as more relevant according to the retrieval component.

The implicit bias of the top passages is crucial to steer the Transformer response generation. The sequence of passages of Eq. 4 is given as input to the Transformer, which will then attend to the different passages. As the multi-head attention layers look across the different passages, redundant parts will be merged, while the remaining information will be summarised, leading to a concise but comprehensive answer. The following Transformer models were considered for the task of abstractive summarisation:

- **Text-to-Text Transfer Transformer (T5)** [26] is a text-to-text model based on the encoder-decoder Transformer architecture, pre-trained on the large C4 corpus, which was derived from Common Crawl¹. A masked language modelling objective is used, where the model is trained to predict corrupted randomly sampled tokens, of varying sizes.
- **BART** [14] is a denoising autoencoder, that combines Bidirectional and Auto-Regressive Transformers. Pre-training consists of corrupting text with an arbitrary noising function and learning an autoencoder to reconstruct the original text. The best performing noise functions were text infilling (using single mask tokens to mask random sampled spans of text), and sentence shuffling (changing the order of sentences in passages).
- **PEGASUS** [38] specialises on the abstractive summarisation task. Multiple important sentences are masked and used as targets, i.e., the model is trained to generate each omitted sentence as output. As in T5, this model is not trained to reconstruct sequences.

4 Evaluation

4.1 Datasets and Protocol

CANARD Dataset [10]. This dataset was used to train and evaluate the query rewriting method. It was created by manually rewriting the queries in QuAC [2] to form non-conversational queries. The training, development, and test sets have 31.538, 3.418, and 5.571, query-rewrites respectively.

TREC CAsT Dataset [5]. This dataset was used to evaluate both the conversational search and answer generation components. There are 50 evaluation topics, each with about 10 turns. Of those in total, 20 conversational topics were labelled on average until turn depth 8 using a graded relevance that ranges from 0 (not relevant) to 4 (highly relevant). The passage collection is composed by MS MARCO [19], TREC CAR [8], and WaPo [20] datasets, which creates a complete pool of close to 47 million passages.

Experimental Protocols. To analyse query rewriting performance, we used the BLEU-4 score [24] between the model’s output and the queries rewritten by humans, on the CANARD dataset.

In the passage retrieval experiment, we used the TREC CAsT setup and the official metrics, nDCG@3 (normalised Discounted Cumulative Gain at 3), MAP (Mean Average Precision), and MRR (Mean Reciprocal Rank), along with Recall and P@3 (Precision at 3).

In the answer generation experiment, we used METEOR and the ROUGE variant ROUGE-L. For each query in TREC CAsT, we use as reference passages, all the passages with a relevance judgement of 3 and 4. Hence, the goal is to generate answers that cover, as much as possible, the information contained in all relevant passages, in one concise and summarised answer.

¹ <https://commoncrawl.org/>.

4.2 Implementation

Query Rewriting. We fine-tuned the T5 [26] model according to [17] and used the CANARD’s training set [10], providing as input the concatenation of the conversational queries and passages, and as target the rewritten query. In particular, we used the T5-BASE model and trained for 4000 steps, using a maximum input sequence length of 512 tokens, a maximum output sequence length of 64 tokens, a learning rate of 0.0001, and batches of 256 sequences.

First-Stage Retrieval. To index and search, we used the well tuned Anserini framework [35], in particular, the Python implementation Pyserini². We applied stop word removal, using Lucene’s default list, and stemming using Kstem³. We experimented with: BM25 [27], language models with Dirichlet (LMD) and Jelinek-Mercer (LMJM) smoothing [37] and from our initial analysis, LMD showed better results. This confirms previous knowledge [37] and matches the shorter queries that we observe in a conversational search scenario. Hence, LMD was the model used in all experiments.

BERT Passage Re-Ranker. To perform re-ranking, we used the BERT model implementation from Huggingface [33]. Following the state-of-the-art [21, 22], we used the LARGE version of BERT with a classification layer (feed-forward neural network) on top, that takes as input the query-passage *CLS token* embeddings vector generated by BERT, and classifies the passage as relevant or non-relevant to that query. This model was trained following [21] on the MS MARCO dataset [19]. In testing, we truncate the concatenation of the query, passage, and separator tokens to a maximum of 512 tokens (the maximum number of tokens for the BERT model).

Transformer Based Answer Generation. To generate the summarised answers, we employed the T5-BASE, BART-LARGE and PEGASUS models [33]. The T5-BASE has about 220 million parameters with 12 layers, 768 hidden-state size, 3072 feed-forward hidden-states and 12 heads. BART-LARGE holds about 406 million parameters, with a 12-layer, 1024 hidden state size and 16-head architecture. The PEGASUS model has the biggest number of parameters, 568 million, with 16 layers, 1024 hidden state size and 16-heads.

All models were fine-tuned on the summarising task with the CNN/Daily Mail dataset [13]. To generate the summary, we use 4 beams, restrict the n-grams of size 3 to only occur once, and allow for beam search early stopping when at least 4 sentences are generated. Additionally, we fix the maximum length of the summary to be of the same length of the input given to the models (which corresponds to 3 passages) and vary the minimum length from 20 to 120 words.

² <https://github.com/castorini/pyserini>.

³ <http://lexicalresearch.com/kstem-doc.txt>.

4.3 Results and Discussion

Conversation-Aware Query Rewriting. In Table 2, we show the BLEU-4 scores obtained in CANARD’s test set and in TREC CAsT’s 2019 manually rewritten queries. The rows “Human” and “Raw” are from [10], the row “T5-BASE” is from [17]. The last row corresponds to our implementation. Our results are on par with [17], being lower in the CANARD dataset but higher in TREC CAsT. We believe the minor differences in performance between our T5-Base model and the T5-BASE from [17] are due to the use of different input sequences, as the exact method of constructing the input is not specified in [17].

Table 2. BLEU-4 scores for the CANARD test set and for TREC CAsT using the manually rewritten queries of the evaluation set.

	CANARD	TREC CAsT
Human [10]	59.92	-
Raw [10]	47.44	-
T5-BASE [17]	58.08	75.07
Our T5-BASE	56.84	79.67

From the analysis of the BLEU-4 scores and outputs, we can conclude that the model is performing both coreference and context resolution, approximating the queries in a conversational format to context-independent queries. Examples of the inputs, targets, and predicted queries, are presented in Table 3. In TREC CAsT, the historical utterances do not depend on the responses of the system, so the answer is not provided as input. As we can see, T5 is capable of resolving ambiguous queries by co-reference resolution, as in example 1, but sometimes mistakes similar co-references when multiple are involved, as evidenced in example 2 and in [17], where the model predicts “throat cancer” instead of “lung cancer”. We can also note that this model is more robust than just coreference resolution, as seen in example 3, where it includes the words “Bronze Age Collapse”, even though there is no explicit mention (implicit coreference).

Transformer-Based Passage Search. Table 4 shows the results of retrieval on the TREC CAsT dataset. *Original* are the conversational queries (lower-bound), *Manual* is a baseline where the queries were manually rewritten (upper-bound), *T5* is using our query rewriting method, and the other two lines are the results of baselines retrieved from [6]. *clacBase* [3] is a method that uses AllenNLP coreference resolution [11] and a fine-tuned BM25 model with pseudo-relevance feedback, and *HistoricalQE* [34] is a method that uses a query expansion algorithm based on session and query words together with a BERT LARGE model for re-ranking. The latter was the best performing method in terms of nDCG@3 in TREC CAsT 2019 [6].

Table 3. Example of query rewriting inputs, targets and predictions.


CANARD	
Original query	What was <u>his</u> agreement with McMahon?
T5 Input query	What was his agreement with McMahon? [CTX] Superstar Billy Graham. Return to WWWF (1977–1981) [TURN] Why did he return to the WWWF? An agreement with promoter Vincent J. McMahon Senior
T5 predicted query	What was <u>Superstar Billy Graham’s</u> agreement with McMahon?
Target query	What was <u>Billy Graham’s</u> agreement with McMahon?
TREC CAsT 2019	
Original query	What are <u>its</u> symptoms?
T5 Input query	What are its symptoms? [CTX] What is throat cancer? [TURN] Is throat cancer treatable? [TURN] Tell me about lung cancer
T5 predicted query	What are <u>throat cancer’s</u> symptoms?
Target query	What are <u>lung cancer’s</u> symptoms?
Original query	What are some of the possible causes?
T5 Input query	What are some of the possible causes? [CTX] Tell me about the Bronze Age collapse? [TURN] What is the evidence for the Bronze Age collapse?
T5 predicted query	What are some of the possible causes <u>for the Bronze Age collapse?</u>
Target query	What are some of the possible causes <u>of the Bronze Age collapse?</u>

The first observation that emerges from Table 4 is the clear need for a query rewriting method to maintain the conversational context, evidenced by the low scores on all metrics using the original conversational queries. Rewriting queries (with the *T5* model) outperforms the original conversational queries by a 5–20% margin (nDCG@3), thus showing the effectiveness of this approach. The second clear observation is again the considerable improvement when Transformers are used for re-ranking. In this case, the improvement is in the 10–15% range over standard retrieval metrics. This is due to the better understanding that the fine-tuned BERT model has of the interactions between the query and passage terms.

Finally, the largest gains emerge when we combine the two Transformers to deliver state-of-the-art results. With the proposed Transformers we outperform the best TREC CAsT 2019 baseline by 3.9% in terms of nDCG@3. We consider that this improvement is mainly due to the use of a better query-rewriting method that allows the retrieval model to retrieve passages given the conversational context, providing the re-ranker with more relevant passages.

Table 4. Results of retrieval on the TREC CAsT evaluation set. The HistoricalQE [34] was the best performing model in TREC CAsT 2019.

Queries	Re-ranker	Recall	P@3	MAP	MRR	nDCG@3
Original	–	0.454	0.262	0.141	0.336	0.167
Original	BERT	0.454	0.385	0.181	0.456	0.272
T5	–	0.697	0.474	0.251	0.597	0.322
T5	BERT	0.697	0.632	0.310	0.739	0.475
TREC CAsT baselines						
clacBase [3]	–	–	–	0.246	0.640	0.360
HistoricalQE [34]	BERT	–	–	0.267	0.715	0.436
Manual baselines						
Manual	-	0.820	0.590	0.327	0.694	0.406
Manual	BERT	0.820	0.757	0.389	0.857	0.577

**Fig. 2.** Performance of the answer generation results under different metrics.

Conversational Answer Generation. Figure 2 shows the result of the answer generation step according to the ROUGE-L and METEOR metrics. The baseline is composed by the concatenation of the top 3 passages, cropped to the maximum length of the passage according to the “Summary Minimum Length” value, respecting sentence endings. In Fig. 2 all answer generation models were better than the retrieval baseline method. According to ROUGE-L the top performance is achieved around 60–90 word length answers. Since the goal is to generate short and informative answers, we were not interested in answers longer than 100 words. Actually, we believe that answers with fewer than 50 words are more natural for conversational scenarios. According to these results we observe that BART was the best answer generation method.

In Fig. 3 we analyse the retrieval and the answer generation performance over conversation turns. We see that peak performance is achieved on the first turn, which was expected given that the first turn that establishes the topic. As the conversation progresses, retrieval performance decreases, but surprisingly, answer generation performance is stable until the 6th turn. We also observed that the decreases in performance are linked to sub-topic shifts within the same conversation topic.

An interesting observation from Fig. 3 is that PEGASUS is the method that exhibits a stronger correlation with retrieval performance. We believe this is related to its generation process that has a behaviour closer to extractive summarisation, while BART and T5 demonstrate a more abstractive behaviour.

Finally, in Table 5 we illustrate the answer generation with all three Transformers. This table further confirms the abstractive versus extractive summarisation behaviours of the different Transformer-based architectures. In this example we see that T5 tries to generate new sentences by combining different sentences.




Fig. 3. Answer generation versus retrieval performance per conversation turn. The minimum length is 80 and 20 in the top and bottom graphs respectively.

Table 5. Answer generation example for the turn “*What was the first artificial satellite?*”. Summary minimum length is set to 90. Blue sentences illustrate abstractive, green sentences illustrate extractive, and red sentences illustrate wrong summaries.

Method	Answer
Retrieval passage 1	The first artificial satellite was Sputnik 1, launched by the Soviet Union on October 4, 1957, and initiating the Soviet Sputnik program, with Sergei Korolev as chief designer (there is a crater on the lunar far side which bears his name). This in turn triggered the Space Race between the Soviet Union and the United States
Retrieval passage 2	The first artificial Earth satellite was Sputnik 1. Put into orbit by the Soviet Union on October 4, 1957, it was equipped with an on-board radio-transmitter that worked on two frequencies: 20.005 and 40.002 MHz. Sputnik 1 was launched as a step in the exploration of space and rocket development. While incredibly important it was not placed in orbit for the purpose of sending data from one point on earth to another. And it was the first artificial satellite in the steps leading to today’s satellite communications
Retrieval passage 3	The first artificial satellite was Sputnik 1. It was the size of a basketball and was made by the USSR (Union of Soviet Socialist Republics) or Russia. It was launched on October 4, 1957.
T5	the first artificial satellite was Sputnik 1, launched by the u.s. or Russia. it was the size of a basketball and launched on October 4, 1957. the satellite was equipped with a radio-transmitter that worked on two frequencies. incredibly important it was not placed in orbit for sending data from one point on earth to another. in turn, it triggered the space race between the united states and the soviet union.
BART	The first artificial satellite was Sputnik 1, launched by the Soviet Union on October 4, 1957. It was equipped with an on-board radio-transmitter that worked on two frequencies: 20.005 and 40.002 MHz. This in turn triggered the Space Race between the Soviet Union and the United States. The size of a basketball, it was not placed in orbit for the purpose of sending data from one point on earth to another. And it was the first Artificial satellite in the steps leading to today’s satellite communications.
PEGASUS	The first artificial satellite was Sputnik 1, launched by the Soviet Union on October 4, 1957. Sputnik 1 was launched as a step in the exploration of space and rocket development. It was not placed in orbit for the purpose of sending data from one point on earth to another. This in turn triggered the Space Race between the USSR and the U.S. There is a crater on the lunar far side which bears his name.

5 Conclusions

In this paper we investigated how Transformer architectures can address different tasks in open-domain conversational search, with particular emphasis on the search-answer generation task. The key findings are:

- **Transformers-based Conversational Search.** Transformers can solve a number of tasks in conversational search, leading to new state-of-the-art results by outperforming the best TREC-CAsT 2019 baseline by 3.9% in terms of nDCG@3. This result is rooted on a fine-tuned bi-directional Transformer model [26] for conversational query re-writing, which attained an improvement of 5–20% (nDCG@3) over raw conversational queries. Similarly, the re-ranking task using a fine-tuned BERT LARGE model [21] improved results by 10–15% (nDCG@3) over an LMD model.
- **Search-Answer Generation.** Experiments showed that search systems can be improved with agents that abstract the information contained in multiple documents to provide a single and informative search answer. In terms of ROUGE-L we concluded that all answer generation models [14, 26, 38] performed better than the retrieval baseline.
- **Abstractive vs Extractive Answer Generation.** The examined answer generation Transformers revealed different behaviours. BART was the most effective in generating answers that were rewritten with information from different passages. This approach turned out to be better than extractive methods that copy and paste sentences from different passages.

As future research, we plan to improve conversational query rewriting methods, re-rankers with a notion of the context of the conversation, and mine possible conversation paths to steer the answer generation process towards further helping the user in exploring alternative aspects of the searched topic.

Acknowledgement. This work has been partially funded by the iFetch project, Ref. 45920 co-financed by ERDF, COMPETE 2020, NORTE 2020, the CMU Portugal project GoLocal Ref. CMUP-ERI/TIC/0046/2014 and by the project NOVA LINCS Ref. UID/CEC/04516/2013.

References

1. Belkin, N.J.: Anomalous states of knowledge as a basis for information retrieval. *Can. J. Inf. Sci.* **5**(1), 133–143 (1980)
2. Choi, E., et al.: Quac : Question answering in context. CoRR abs/1808.07036 (2018). <http://arxiv.org/abs/1808.07036>
3. Clarke, C.L.A.: Waterlooclarke at the TREC 2019 conversational assistant track. In: Voorhees, E.M., Ellis, A. (eds.) *Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC 2019, Gaithersburg, Maryland, USA, 13–15 November, 2019*. NIST Special Publication, vol. 1250. National Institute of Standards and Technology (NIST) (2019). <https://trec.nist.gov/pubs/trec28/papers/WaterlooClarke.C.pdf>

4. Croft, W.B., Thompson, R.H.: I3r: a new approach to the design of document retrieval systems. *JASIST* **38**(6), 389–404 (1987)
5. Dalton, J., Xiong, C., Callan, J.: The trec conversational assistance track (cast) (1 2020). <http://www.treccast.ai/>
6. Dalton, J., Xiong, C., Callan, J.: TREC cast 2019: The conversational assistance track overview. *CoRR abs/2003.13624* (2020). <https://arxiv.org/abs/2003.13624>
7. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR abs/1810.04805* (2018). <http://arxiv.org/abs/1810.04805>
8. Dietz, L., Gamari, B., Dalton, J.: Trec car 2.1: A data set for complex answer retrieval (7 2018). <http://trec-car.cs.unh.edu>
9. Dinan, E., Roller, S., Shuster, K., Fan, A., Auli, M., Weston, J.: Wizard of wikipedia: Knowledge-powered conversational agents. *CoRR abs/1811.01241* (2018). <http://arxiv.org/abs/1811.01241>
10. Elgohary, A., Peskov, D., Boyd-Graber, J.L.: Can you unpack that? learning to rewrite questions-in-context. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 November, 2019*. pp. 5917–5923. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/D19-1605>
11. Gardner, M., et al.: Allennlp: a deep semantic natural language processing platform. *CoRR abs/1803.07640* (2018). <http://arxiv.org/abs/1803.07640>
12. Han, S., Wang, X., Bendersky, M., Najork, M.: Learning-to-rank with BERT in tf-ranking. *CoRR abs/2004.08476* (2020). <https://arxiv.org/abs/2004.08476>
13. Hermann, K.M., et al.: Teaching machines to read and comprehend. In: *Advances in Neural Information Processing Systems*, pp. 1693–1701 (2015)
14. Lewis, M., et al.: BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR abs/1910.13461* (2019). <http://arxiv.org/abs/1910.13461>
15. Li, J., Monroe, W., Ritter, A., Jurafsky, D., Galley, M., Gao, J.: Deep reinforcement learning for dialogue generation. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1192–1202. Association for Computational Linguistics, Austin, Texas (November 2016). <https://doi.org/10.18653/v1/D16-1127>, <https://www.aclweb.org/anthology/D16-1127>
16. Li, J., Monroe, W., Shi, T., Jean, S., Ritter, A., Jurafsky, D.: Adversarial learning for neural dialogue generation. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 2157–2169. Association for Computational Linguistics, Copenhagen, Denmark (September 2017). <https://doi.org/10.18653/v1/D17-1230>, <https://www.aclweb.org/anthology/D17-1230>
17. Lin, S., Yang, J., Nogueira, R., Tsai, M., Wang, C., Lin, J.: Conversational question reformulation via sequence-to-sequence architectures and pretrained language models. *CoRR abs/2004.01909* (2020). <https://arxiv.org/abs/2004.01909>
18. Liu, Y., et al.: Roberta: A robustly optimized BERT pretraining approach. *CoRR abs/1907.11692* (2019), <http://arxiv.org/abs/1907.11692>
19. Nguyen, T., et al.: MS MARCO: A human generated machine reading comprehension dataset. *CoRR abs/1611.09268* (2016). <http://arxiv.org/abs/1611.09268>
20. NIST: Trec washington post corpus (12 2019). <https://trec.nist.gov/data/wapost/>
21. Nogueira, R., Cho, K.: Passage re-ranking with BERT. *CoRR abs/1901.04085* (2019). <http://arxiv.org/abs/1901.04085>

22. Nogueira, R., Yang, W., Cho, K., Lin, J.: Multi-stage document ranking with BERT. CoRR abs/1910.14424 (2019). <http://arxiv.org/abs/1910.14424>
23. Oddy, R.N.: Information retrieval through man-machine dialogue. *J. Documentation* **33**(1), 1–14 (1977)
24. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, pp. 311–318. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA (July 2002). <https://doi.org/10.3115/1073083.1073135>, <https://www.aclweb.org/anthology/P02-1040>
25. Qu, C., Yang, L., Chen, C., Qiu, M., Croft, W.B., Iyyer, M.: Open-retrieval conversational question answering. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 539–548. SIGIR 2020, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3397271.3401110>
26. Raffel, C., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. CoRR abs/1910.10683 (2019). <http://arxiv.org/abs/1910.10683>
27. Robertson, S., Zaragoza, H.: The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval* **3**(4), 333–389 (2009). <https://doi.org/10.1561/15000000019>
28. Song, Y., et al.: An ensemble of retrieval-based and generation-based human-computer conversation systems. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, pp. 4382–4388. International Joint Conferences on Artificial Intelligence Organization (7 2018). <https://doi.org/10.24963/ijcai.2018/609>
29. Tian, Z., Bi, W., Li, X., Zhang, N.L.: Learning to abstract for memory-augmented conversational response generation. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 3816–3825. Association for Computational Linguistics, Florence, Italy (July 2019). <https://doi.org/10.18653/v1/P19-1371>, <https://www.aclweb.org/anthology/P19-1371>
30. Vaswani, A., et al: Attention is all you need. CoRR abs/1706.03762 (2017). <http://arxiv.org/abs/1706.03762>
31. Voskarides, N., Li, D., Ren, P., Kanoulas, E., de Rijke, M.: Query resolution for conversational search with limited supervision. Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (July 2020). <https://doi.org/10.1145/3397271.3401130>
32. Vtyurina, A., Savenkov, D., Agichtein, E., Clarke, C.L.A.: Exploring conversational search with humans, assistants, and wizards. In: Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 2187–2193. CHI EA 2017, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3027063.3053175>
33. Wolf, T., et al.: Huggingface’s transformers: State-of-the-art natural language processing. CoRR abs/1910.03771 (2019). <http://arxiv.org/abs/1910.03771>
34. Yang, J., Lin, S., Wang, C., Lin, J., Tsai, M.: Query and answer expansion from conversation history. In: Voorhees, E.M., Ellis, A. (eds.) Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC 2019, Gaithersburg, Maryland, USA, 13–15 November, 2019. NIST Special Publication, vol. 1250. National Institute of Standards and Technology (NIST) (2019). https://trec.nist.gov/pubs/trec28/papers/CFDA_CLIP.C.pdf

35. Yang, P., Fang, H., Lin, J.: Anserini: enabling the use of lucene for information retrieval research. In: Kando, N., Sakai, T., Joho, H., Li, H., de Vries, A.P., White, R.W. (eds.) Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, 7–11 August, 2017, pp. 1253–1256. ACM (2017). <https://doi.org/10.1145/3077136.3080721>
36. Yang, Z., Dai, Z., Yang, Y., Carbonell, J.G., Salakhutdinov, R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding. CoRR abs/1906.08237 (2019). <http://arxiv.org/abs/1906.08237>
37. Zhai, C., Lafferty, J.: A study of smoothing methods for language models applied to ad hoc information retrieval. In: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 334–342. SIGIR 2001, Association for Computing Machinery, New York, NY, USA (2001). <https://doi.org/10.1145/383952.384019>
38. Zhang, J., Zhao, Y., Saleh, M., Liu, P.J.: PEGASUS: pre-training with extracted gap-sentences for abstractive summarization. CoRR abs/1912.08777 (2019). <http://arxiv.org/abs/1912.08777>
39. Zhuang, Y., Wang, X., Zhang, H., Xie, J., Zhu, X.: An ensemble approach to conversation generation. In: Huang, X., Jiang, J., Zhao, D., Feng, Y., Hong, Yu. (eds.) NLPC 2017. LNCS (LNAI), vol. 10619, pp. 51–62. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73618-1_5